

000 RLAD: TRAINING LLMs TO DISCOVER ABSTRACTIONS 001 FOR SOLVING REASONING PROBLEMS 002 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

010 Reasoning requires going beyond pattern matching or memorization of solutions
011 to identify and implement “algorithmic procedures” that can be used to deduce
012 answers to hard problems. Doing so requires reusing primitives, intermediate
013 results, or procedures across multiple problems. While RL post-training on long
014 chains of thought ultimately aims to uncover this kind of algorithmic behavior, the
015 depth-first and “brute-force” nature of reasoning traces learned by these models
016 suggests that this is far from a fulfilled promise. To address more effective reason-
017 ing, we introduce *reasoning abstractions*: concise natural language descriptions of
018 procedural and factual knowledge that guide the model toward learning successful
019 reasoning. We train models to be capable of proposing several useful abstractions
020 given a problem, followed by RL training that incentivizes building a solution while
021 using the information provided by these abstractions. This results in a two-player
022 RL training paradigm, abbreviated as RLAD, that jointly trains an abstraction gen-
023 erator and an abstraction-conditioned solution generator. This setup effectively
024 enables structured exploration, decouples learning signals of abstraction proposal
025 and solution generation, and improves generalization to harder problems. We also
026 show that spending more test-time compute into generating abstractions is more
027 beneficial for performance than generating more solutions at large inference-time
028 budgets, illustrating the role of abstractions in guiding global exploration.

029 030 1 INTRODUCTION

031 Modern machinery for training large language models (LLMs) to reason relies on incentivizing
032 longer chains of thought via reinforcement learning (RL). This training approach largely incentivizes
033 “depth”: subsequent training iterations increase response length by incorporating new operations that
034 usually verify or build on top of the line of reasoning being already pursued by the model (Setlur
035 et al., 2025). This often results in very long chains of thought that appear to explore the solution
036 search space, but in a sequential, brute-force manner. In many problems, it is instead more desirable
037 to optimize for “breadth”: explore a diverse array of solution strategies, rather than committing to a
038 seemingly good set of reasoning strategies right away (Yu et al., 2025; Yue et al., 2025). Even though
039 models trained this way succeed on some problems, they fail on problems with similar difficulty,
040 revealing poor generalization (Ma et al., 2024; Mirzadeh et al., 2024; Petrov et al., 2025).

041 **How can we help models explore a breadth of reasoning strategies for a given problem?** Abstractly,
042 the most natural approach is to train models to hypothesize new strategies to attack difficult problems
043 and then attempt to utilize these strategies in the solution. We can do this by making models
044 capable of discovering *reasoning abstractions*: compressed representations of shared procedures
045 that underlie multiple candidate solutions to a problem. For example, in math reasoning, such
046 abstractions might correspond to useful intermediate lemmas or even some intermediate steps that
047 do not succeed but illustrate what not to do. When presented in context, these abstractions function
048 akin to “hints” on an exam, enabling LLMs to solve harder problems by building on the insights
049 appearing in the abstraction. That is, when conditioned on abstractions, training via RL should train
050 the LLM to implement useful meta strategies that utilize and compose the procedural information
051 in the abstraction as best as possible to solve the problem, rather than attempting to search over the
052 procedural information itself. This naturally boosts the diversity of solution strategies and behaviors
053 that a model learns to utilize when encountering an unseen problem, in contrast to committing to a
narrow set of approaches. In RL terminology, abstractions serve as high-level subgoals, skills, or
priors—any of them depending upon context—guiding the low-level solution-generating policy.

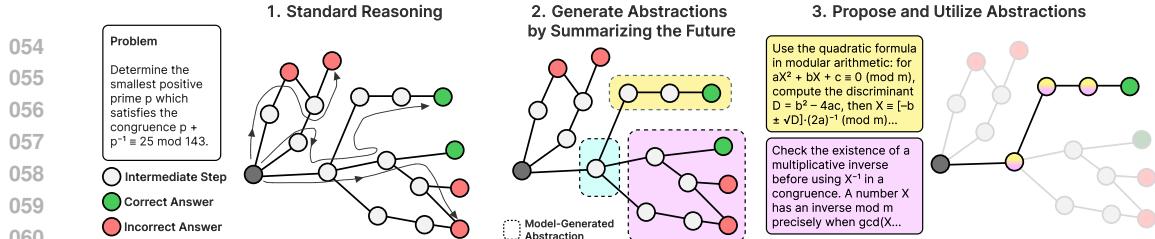


Figure 1: **Reasoning abstractions illustrated in the solution-space graph for a problem.** We depict the solution space as a graph of intermediate steps leading to correct or incorrect answers. (1) Standard reasoning explores this space along one sequential chain. (2) We generate textual abstractions by summarizing which intermediate steps led to which outcomes. (3) Such abstractions can be reused to guide reasoning more efficiently.

In this work, we imbue LLMs with the capability of proposing and utilizing abstractions for solving problems. Concretely, we build reasoning models that, first, given an input problem, propose one or more reasoning abstractions, expressed in natural language. Subsequently, they generate a solution that utilizes the information and principles prescribed by these abstractions. To achieve this, we jointly train two LLMs via RL post-training: (1) an abstraction generator, and (2) an abstraction-conditioned solution generator. The abstraction generator is rewarded for the improvement in the accuracy of the solution generator, stemming from conditioning on the abstractions it proposes. The solution generator is rewarded to maximize accuracy in solving a problem when using the abstraction. To obtain a good initialization for RL training, we warmstart both models by running supervised fine-tuning (SFT) on data from stronger models. For the abstraction generator, we collect multiple candidate solutions and prompt a stronger LLM to generate diverse abstractions. For the solution generator, we generate solutions conditioned on an abstraction.

The main contribution of this paper is the notion of *reasoning abstractions*, how they can be obtained, training procedures to amplify them via RL training, and an illustration of how they can be used to improve reasoning performance and exploration of the search space. Concretely, we build an approach to imbue LLMs with the capability of proposing abstractions, and evaluate the model on a variety of math-reasoning benchmarks, AIME 2025 (Mathematical Association of America, 2025), DeepScaleR Hard (Setlur et al., 2025), and AMC 2023. We find an average 44% improvement over state-of-the-art long chain-of-thought RL approaches (i.e., DAPO (Yu et al., 2025)) on AIME 2025, and show an effective benefit from generating diverse abstractions over brute-force solution sampling.

2 RELATED WORK

Scaling test-time compute and exploration. Recent work highlights the promise of scaling test-time compute in different ways. One approach involves parallel sampling: sampling multiple reasoning rollouts and then selecting a winner via a scoring rule (Uesato et al., 2022; Wang et al., 2023; Charniak & Johnson, 2005; Feng et al., 2024; Snell et al., 2024; Yao et al., 2023a; Hao et al., 2023; Snell et al., 2024). A complementary line of work iteratively edits a single trace, attempting to implement some sort of a sequential search within a single solution trace (Madaan et al., 2023; Qu et al., 2024; Qu et al., 2024; Kumar et al., 2024). As such, the sequential approach performs a bit worse on harder problems (Snell et al., 2024; Qu et al., 2025), where it often gets trapped in strategies that seem optimal but aren't actually (Pan et al., 2025). Yet it still performs better than parallel search on easier and medium difficulty problems (Snell et al., 2024). Our approach of proposing and leveraging abstractions enables a kind of a hybrid between sequential sampling and parallel sampling, guided by the proposed abstractions. Some concurrent work (Pan et al., 2025) studies directly interleaving parallel and sequential samples, and while it is similar in theory to us, it only distills this interleaved structure into the model and does not run RL training to optimize parallel and sequential sampling procedures employed here. Prior work has also utilized hand-designed scaffolds to integrate multi-step evaluations of intermediate hypotheses into reasoning (Yao et al., 2023b; Ho et al., 2023; Hao et al., 2023; Li et al., 2023). In contrast, we do not rely on pre-defined interfaces but learn to automatically propose useful abstractions.

Using prior knowledge for LLM reasoning. Several threads of work converge on the idea that *textual artifacts* such as examples, plans, or prompts, can serve as reusable knowledge that steers LLM behavior. Existing retrieval-augmented generation (RAG) pipelines assume a static corpus, typically of human-written text, and focus on improving retrieval heuristics (Lewis et al., 2020; Borgeaud et al., 2022; Trivedi et al., 2022; Verma et al., 2024; Anonymous, 2025; Li et al., 2025).

108 Many works use LLMs to learn or refine prompts, either in an input-agnostic fashion (Zhou et al.,
 109 2022; Yang et al., 2023; Pryzant et al., 2023; Fernando et al., 2023) or through input-specific edits
 110 based on feedback (Shinn et al., 2023; Madaan et al., 2023; Gou et al., 2023; Yuksekgonul et al.,
 111 2025; Lin et al., 2025). Other related work explores the use of synthetic demonstrations (Zelikman
 112 et al., 2022b), scratchpads (Nye et al., 2021), and memory-augmented agents (Schäfer et al., 2020) to
 113 encode prior problem-solving knowledge. Two recent works demonstrate that LLMs can accumulate
 114 and reuse their own experience across tasks (Zhao et al., 2024; Suzgun et al., 2025). While one
 115 can view our abstractions as a form of prior procedural and factual knowledge produced before the
 116 model’s solution attempt, this knowledge is (a) input-dependent and (c) is not acquired from an
 117 external source at deployment, but rather is “proposed” by the model itself. Imbuing models with this
 118 capability requires a two-player RL training process. To our knowledge, such procedures have not
 119 been used for generating textual artifacts of any type, let alone the abstractions we consider.
 120

3 PRELIMINARIES AND NOTATION

121 We study reasoning with LLMs, where the LLM is provided access to a problem \mathbf{x} , and generates
 122 a stream of tokens \mathbf{y} that ends in an estimate of the answer. We assume access to a rule-based
 123 ground-truth 0/1 reward $\text{Acc}_{\mathbf{x}}(\mathbf{y}, \mathbf{y}^*) \in \{0, 1\}$ that measures correctness of the produced answer
 124 \mathbf{y} , against the ground-truth solution \mathbf{y}^* for a question \mathbf{x} . For training, we are given a dataset
 125 $\mathcal{D}_{\text{train}} = \{(\mathbf{x}_i, \mathbf{y}_i^*)\}_{i=1}^N$ of problems \mathbf{x}_i and solutions \mathbf{y}_i^* that end with the correct answer. Our goal is
 126 to train the LLM $\pi(\cdot|\mathbf{x})$ such that it achieves high rewards on a test distribution of problems $\mathcal{P}_{\text{test}}$.
 127

128 We evaluate models via average accuracy under $\mathcal{P}_{\text{test}}$. We also measure the pass@k metric, where
 129 for problem \mathbf{x} , we sample k solutions $\mathbf{y}_1, \dots, \mathbf{y}_k \sim \pi(\cdot|\mathbf{x})$, and consider the problem to be solved if
 130 any of these k traces is correct. This metric couples accuracy with diversity, i.e., it attains the largest
 131 value when the model effectively finds diverse, good responses. To reduce variance in estimating
 132 pass@k, we sample $n \geq k$ samples per problem and use the unbiased estimator introduced in OpenAI
 133 Codex (Chen et al., 2021): $1 - \binom{n-c}{k} / \binom{n}{k}$, where $c \leq n$ is the number of correct samples.
 134

4 REASONING ABSTRACTIONS AND WHY THEY ARE USEFUL

135 Solving reasoning problems often requires composing both *procedural* knowledge (e.g., how to
 136 apply a root-finding algorithm) and *factual* knowledge (e.g., relevant lemmas or intermediate results).
 137 Current approaches train models to reason via reinforcement learning (RL) on long chains of thought.
 138 However, this is often ineffective as RL often tends to optimize for “depth”, producing longer
 139 traces where each subsequent segment builds on the last segment (e.g., verifying prior calculations),
 140 rather than “breadth”, i.e., exploring diverse solution strategies or utilizing seemingly irrelevant
 141 procedures when needed. We now introduce the concept of **reasoning abstractions**, concise insights
 142 that explicitly encode a range of useful procedural and factual concepts for a problem. We describe
 143 our approach for generating abstractions and demonstrate mechanisms that make them work.
 144

4.1 PROPOSING GOOD REASONING ABSTRACTIONS BY SUMMARIZING SOLUTION ATTEMPTS

145 To imbue LLMs with the capability of proposing abstractions, we warmstart the LLM using a
 146 dataset consisting of problems paired with a small set of high-quality reasoning abstractions, created
 147 synthetically. Perhaps the most natural way to obtain an initial set of reasoning abstractions is to
 148 collect a diverse set of traces attempting to solve a problem and then summarize useful concepts
 149 appearing in these traces (see Figure 1 for an illustration). More formally, consider the space of
 150 possible reasoning traces for a given problem as a graph where the nodes of the graph represent
 151 intermediate states encountered when solving a question (see Figure 1). Good abstractions would
 152 identify useful substructures within this larger reasoning graph. For example, an abstraction can
 153 capture whether a set of strategies lead to a similar outcome or another set of tactics leads to an error
 154 being consistently made.

155 **Generating abstractions.** We now quantitatively evaluate whether we can generate useful abstractions
 156 by summarizing key insights in reasoning and solution traces that solve a problem. To do so, we
 157 prompt a model to generate solution traces and prompt a stronger model to deduce useful patterns
 158 from the responses of the first model. Concretely, we utilize the Qwen3 (Qwen Team, 2025) series
 159 of models to produce solutions and a stronger reasoning model, o4-mini, to generate abstractions.
 160 While this approach is not perfect, it enables us to validate the feasibility of the concept of reasoning
 161 abstractions and generate warmstart data for *training* LLMs to propose abstractions. To ensure that
 the abstractions do not “leak” content of the solution, we verify post-hoc that prompting a model

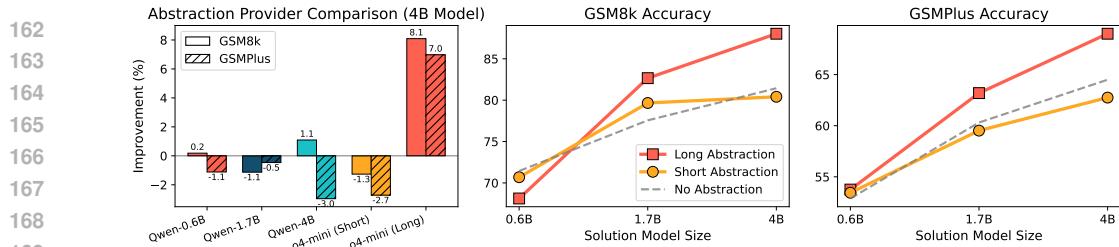


Figure 2: **Benefits from abstractions hinge on solver scale, abstraction length, and solution model.** Most configurations fail to yield gains; only o4-mini with long and detailed abstractions shows consistent improvements across the GSM8k and GSMPplus datasets (left). Solver capability also matters: even strong abstractions help only if the solution model is sufficiently capable (middle, right).

with only the abstraction and no problem yields zero accuracy when sampling 16 times from the base model. This makes these abstractions suitable for our study as they only provide useful information while not allowing the model to shortcut to the answer.

Evaluating abstractions. To validate the efficacy of using abstractions, we adopt a simple test based on performance after conditional generation. Concretely, let us denote the LLM policy that produces a solution conditioned on the problem \mathbf{x} as $\pi_\theta^{\text{sol}}(\cdot|\mathbf{x})$. A good abstraction \mathbf{z} is a sequence of tokens that provides some useful procedural and factual information to improve model performance:

$$\mathbb{E}_{\tilde{\mathbf{y}} \sim \pi_\theta^{\text{sol}}(\cdot|\mathbf{x}, \mathbf{z})} [\text{Acc}(\tilde{\mathbf{y}}, \mathbf{y}^*)] > \mathbb{E}_{\tilde{\mathbf{y}} \sim \pi_\theta^{\text{sol}}(\cdot|\mathbf{x})} [\text{Acc}(\tilde{\mathbf{y}}, \mathbf{y}^*)]. \quad (1)$$

4.2 RESULTS AND OBSERVATIONS

Evaluation on math reasoning. After generating abstractions, we measure their quality by evaluating Equation 1, i.e., by checking if conditioning the problem solver on a set of abstractions improves its accuracy. Results in Figure 2 show that conditioning a problem solver on abstractions improves accuracy when two conditions hold simultaneously: (i) the abstraction is not too short (e.g., not just a few words that are not informative; example in Appendix D.2) and is generated by a strong generator (o4-mini) and (ii) the solution generator has sufficient instruction-following capability (Qwen3-1.7B or Qwen3-4B) of interpreting and utilizing the generated abstraction. These results confirm that good abstractions (satisfying equation 1) exist for math problems, but neither the ability to generate them nor the ability to leverage them in solutions arises naturally. In Section 5, we will describe our method for explicitly training models to propose and use such abstractions effectively.

Evaluation on ARC-AGI. We also evaluate abstractions on the ARC-AGI benchmark. We present some details of our setup in Appendix B.3. We evaluate on 90 ARC puzzles evenly derived from the test sets of ARC-AGI 1, ARC-AGI 2, and BARC (Li et al., 2024). In Table 1, we present the pass@k and coverage performance (% of unit tests successfully passed) of the base Qwen3-4B model when conditioned on a proposed abstraction vs not using any abstraction. We see a positive improvement in both metrics on this domain over multiple samples, indicating an improvement from utilizing reasoning abstractions.

Interpreting the generated abstractions. We also show some examples of the discovered abstractions in Appendix D.3. We now attempt to interpret these abstractions by classifying them into several categories. We observe that the discovered abstractions often correspond to useful techniques (e.g., ‘‘launchpoint’’ in Appendix D.3), a useful lemma or heuristic principle (e.g., ‘‘blind-follow’’ in Appendix D.3), and cautionary examples that demonstrate common pitfalls encountered when solving a problem (e.g., ‘‘caution alert’’ in Appendix D.3). These abstractions distill complex reasoning patterns and potential approaches into useful nuggets, allowing models to generalize across structurally similar problems. Finally, we want to remind the reader and emphasize that these qualitative results from interpreting the discovered abstractions are specific to an individual problem, and not representative of the process being used to discover them. Our approach for generating abstractions is neither hand-engineered for interpretability or uses any such heuristics beyond summarization.

Good abstractions exist in many domains. We also find that the summarization procedure can be used to identify an initial set of useful reasoning abstractions on many problem domains, including

k	pass@k		coverage	
	w/o abs	w/ abs	w/o abs	w/ abs
1	14.0%	18.0%	19.5%	22.5%
2	17.5%	22.5%	24.0%	28.5%
4	20.0%	26.5%	28.5%	34.0%
8	22.8%	30.2%	31.8%	38.8%
16	24.7%	33.2%	35.0%	42.9%

Table 1: **pass@k accuracy and max@k coverage on ARC-AGI.** Abstractions yield consistent gains in both metrics.

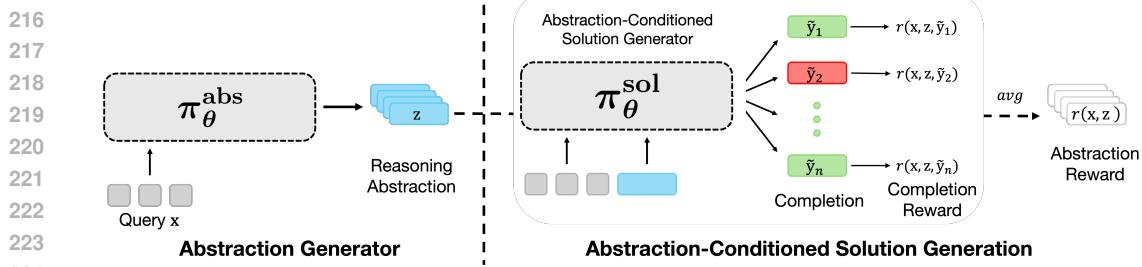


Figure 3: **RLAD training paradigm**. We train an abstraction generator, π_θ^{abs} , that proposes some reasoning abstractions conditioned on the question x , denoted as z . Then, the solution generator, π_θ^{sol} , is trained to produce a response, \tilde{y} , conditioned on the generated abstraction z . The reward used for training π_θ^{abs} corresponds to the average success rate of the solution generator conditioned on the proposed abstraction.

healthcare, human behavior, legal reasoning, and web security. Of course, the proportion of an abstraction devoted to procedural knowledge and factual knowledge is different in these domains compared to math. That said, we find that using reasoning abstractions improves performance by 30% on average over 37 tasks from RAFT (Alex et al., 2021), CLUES (Menon et al., 2022), and LegalBench (Guha et al., 2023). We show some examples in Figure 7 and full results in Appendix B.3.

Takeaways: Reasoning abstractions summarize insights useful for guiding solution traces

Reasoning abstractions summarize procedural and factual knowledge that is useful for learning to solve problems via diverse strategies. Proposing abstractions generated by summarizing solution traces already improves performance of base generators by 30% for math reasoning.

5 RLAD: LEARNING TO PROPOSE REASONING ABSTRACTIONS

Having defined the notion of reasoning abstractions and shown that they can improve performance when adhered to for reasoning, we will now develop an approach to train LLMs to be capable of both proposing and utilizing abstractions. Doing so requires training an *abstraction generator*: an LLM, $z \sim \pi_\theta^{\text{abs}}(\cdot|x)$ that proposes candidate abstractions z given problem x , and an abstraction-conditioned solution generator, $y \sim \pi_\theta^{\text{sol}}(\cdot|x, z)$, that produces a solution y given x and abstraction z . Note that z is parameterized as a variable-length sequence of tokens and might consist of one or more facts or procedures. While our approach applies to the case when π_θ^{abs} produces more than one abstraction, we abuse notation and subsume multiple abstraction into one to avoid clutter. We now describe RL with abstraction discovery (RLAD), our method for training these models via RL.

5.1 TRAINING π_θ^{abs} AND π_θ^{sol} VIA RL

The core principle behind our approach is that an abstraction z is successful at a given problem x if it can maximally help $\pi_\theta^{\text{sol}}(\cdot|x, z)$ find correct responses to the question x , without actually leaking the answer itself. To convert this into an RL objective, we design a reward function that rewards an abstraction z with the expected success of solutions generated by π_θ^{sol} conditioned on z :

$$r_{\pi_\theta^{\text{sol}}}(\mathbf{x}, \mathbf{z}) := \mathbb{E}_{\tilde{y} \sim \pi_\theta^{\text{sol}}(\cdot|\mathbf{x}, \mathbf{z})} [\text{Acc}_{\mathbf{x}}(\tilde{y}, \mathbf{y}^*)], \quad (2)$$

where \mathbf{y}^* is the ground-truth answer and $\text{Acc}_{\mathbf{x}}(\cdot, \cdot)$ denotes the 0/1 accuracy on problem \mathbf{x} . To train π_θ^{sol} , one can then adopt the fairly straightforward approach of maximizing 0/1 binary outcome reward, now conditioned on a given abstraction z sampled previously from π_θ^{abs} , akin to recent results RL (DeepSeek-AI et al., 2025). Formally, we set the reward for a solution as: $r(\mathbf{x}, \mathbf{z}, \tilde{y}) := \text{Acc}_{\mathbf{x}}(\tilde{y}, \mathbf{y}^*)$. With these reward functions in place, perhaps the most natural approach then would be to train π_θ^{abs} to maximize $r_{\pi_\theta^{\text{sol}}}$ for a fixed π_θ^{sol} on a dataset of prompts $\mathcal{D}_{\pi_\theta^{\text{abs}}}$, while also iteratively training π_θ^{sol} to maximize the reward function r on modified prompts generated by concatenating a set of sampled abstraction z on a dataset of problems, $\mathcal{D}_{\pi_\theta^{\text{sol}}}$. This maximization could be done via on-policy RL (e.g., GRPO (Shao et al., 2024)) or offline RL methods (e.g., DPO (Rafailov et al., 2023), STaR (Zelikman et al., 2022a)). This represents a co-operative two-player game.

Challenges with naïve reward design. While the approach so far is extremely simple, it presents some challenges. In particular, the reward functions defined above can result in undesirable solutions in a rather nuanced manner: (1) if π_θ^{abs} learns to solve problem \mathbf{x} in its entirety, it will still be rewarded

Question	Solution
<p>270 Let A be the set of positive integer divisors of 2025. Let B be a 271 randomly selected subset of A. The probability that B is a 272 nonempty set with the property that the least common multiple 273 of its element is 2025 is m/n, where m and n are relatively 274 prime positive integers. Find $m+n$.</p> <p>275</p> <p>276 Reasoning Abstraction</p> <p>277 <code><description>Model subsets of divisors by their exponent 278 vectors. Let S be the set of exponent choices for all primes. A 279 subset of divisors corresponds to a subset of S. The condition 280 "LCM equals N" becomes "for each prime i, the maximum 281 exponent chosen in the subset equals e_i".</description></code></p> <p>282 <code><example>With primes (2,5) and exponent choices S={(0,0), 283 (1,1),(2,2),(3,3)}, the condition "LCM=2^3*5^2" means we must 284 include at least one choice with exponent 3 for 2 and exponent 285 2 for 5. Any valid subset has at least one of (3,0),(0,2),(3,2).</example></code></p>	<p>270 ... 271 But how do we compute the number of subsets with LCM equal to 2025? 272 From the cheatsheet, there's a mention of using prime factorization and 273 considering exponent vectors. Let me recall that. 274 The idea is that each divisor of 2025 can be represented as $3^a * 5^b$, where $0 \leq a \leq 4$ and $0 \leq b \leq 2$. So each element of A is determined by exponents (a, b). 275 The LCM of a subset of these elements is 2025, which is $3^4 * 5^2$. Therefore, 276 for the LCM of the subset to be exactly 2025, each prime in the subset must 277 have at least the exponent of 2025 for that prime. 278 So for prime 3, in the subset, there must be at least one element with exponent 279 4 (since 2025 has exponent 4 for 3). Similarly, for prime 5, there must be at 280 least one element with exponent 2 (since 2025 has exponent 2 for 5). 281 Therefore, the problem reduces to counting the number of subsets of A 282 (including empty set) where:</p>

283 **Figure 4: A typical example of a reasoning abstraction proposed by our abstraction generator.** In the solution,
284 we see (in blue) references to the abstraction (“cheatsheet”) and keywords from the abstraction being used
285 meaningfully in the reasoning trace of the solution generator model.

286 highly by $r_{\pi_\theta^{\text{sol}}}$ but is not a desirable abstraction; (2) if π_θ^{sol} is too weak or too strong, such that it is
287 either able to always solve the problem \mathbf{x} or never solves it, then $r_{\pi_\theta^{\text{sol}}}$ will not provide a meaningful
288 signal to update π_θ^{abs} ; and (3) similar to the above failure modes, training π_θ^{sol} via on-policy RL may
289 result in it ignoring the abstraction \mathbf{z} altogether no matter how useful it is. Abstractly, all of these
290 challenges stem from an asymmetry in the strength of π_θ^{abs} and π_θ^{sol} , where one may drown out the
291 learning signal for the other. We therefore build a modified reward system for training.

292 **Modifying reward design.** We make a small but consequential change to the training reward system.
293 In particular, we train π_θ^{sol} on a mixture of prompts \mathbf{x} augmented by abstractions \mathbf{z} and prompts
294 \mathbf{x} without any abstractions at all. In this process, while we utilize $\text{Acc}_{\mathbf{x}}$ as discussed above on
295 a given response, we simply zero out rewards for any trace generated on \mathbf{x} without abstractions.
296 When utilizing KL-constrained RL, e.g., GRPO (Shao et al., 2024), π_θ^{sol} is now trained to closely
297 mimic the distribution of responses as the reference LLM on questions \mathbf{x} but must attempt to find
298 ways to optimize reward on the same question \mathbf{x} when augmented with an abstraction. This can be
299 accomplished only when π_θ^{sol} learns to utilize the provided abstraction carefully, hence addressing
300 one of the challenges above. Formally, the updated versions of these reward functions are shown as:

$$r(\mathbf{x}, \mathbf{z}, \tilde{\mathbf{y}}) := \begin{cases} 0, & \text{if } \mathbf{z} = \emptyset \\ \text{Acc}_{\mathbf{x}}(\tilde{\mathbf{y}}, \mathbf{y}^*), & \text{otherwise} \end{cases} \quad (3)$$

$$r_{\pi_\theta^{\text{sol}}}(\mathbf{x}, \mathbf{z}) := \mathbb{E}_{\tilde{\mathbf{y}} \sim \pi_\theta^{\text{sol}}(\cdot | \mathbf{x}, \mathbf{z})} [\text{Acc}_{\mathbf{x}}(\tilde{\mathbf{y}}, \mathbf{y}^*)]. \quad (4)$$

304 5.2 WARMSTARTING π_θ^{sol} AND π_θ^{abs} FROM GOOD INITIALIZATIONS

305 While the above approach prescribes a recipe for RL training of π_θ^{abs} and π_θ^{sol} , any such recipe
306 critically relies on the ability of the initialized model to generate somewhat meaningful abstractions
307 and meaningful solutions conditioned on the abstraction input, respectively, right from the beginning
308 of RL training. Inspired from the approach of running an initial phase of SFT to imbue into the
309 model the basic structure of a long chain-of-thought before RL (DeepSeek-AI et al., 2025; Qu et al.,
310 2025), we run an initial phase of SFT to imbue into π_θ^{abs} and π_θ^{sol} the basic capabilities of producing
311 abstractions and attempting to follow abstractions respectively, even if the resulting models are not
312 very good. For this initial warmstart phase, we follow the protocol from Section 4 and construct
313 a corpus $\{(\mathbf{x}_i, \mathbf{z}_i, \mathbf{y}_i)\}_{i=1}^M$ by prompting strong models. For each training problem-solution pair
314 $(\mathbf{x}, \mathbf{y}^*)$, in our training set, we first generate an abstraction \mathbf{z} using an instruction-tuned model,
315 discarding any that leak \mathbf{y}^* . We then sample a solution trace \mathbf{y} conditioned on (\mathbf{x}, \mathbf{z}) .

316 **Practical approach and algorithm details.** For warmstarting the abstraction generator, we utilize
317 abstractions generated by o4-mini. We then use a weaker model (GPT-4.1-mini) to evaluate each
318 abstraction by comparing solution success with and without it, retaining only those that improve
319 performance to form our seed set. **Then, we run SFT of Qwen3-1.7B for 5 epochs on this seed dataset**
320 **to obtain our initial abstraction generator** (Qwen Team, 2025). **For the solution generator, we use the**
321 **same Qwen3-1.7B model**, ensuring that both components use models of identical capacity.

322 After SFT, we employ RLAD to fine-tune the abstraction generator and abstraction-conditioned
323 solution generator via RL. For the abstraction generator, we opt to use “batched” offline RL via
324 RFT (Yuan et al., 2023) and RPO (Pang et al., 2024), since reward computation by rolling out the

324 solution generator the on the fly and running online RL was infeasible using our RL infrastructure and
 325 within compute we had access to. To train the solution generator, we utilize the DAPO approach (Yu
 326 et al., 2025), and include token-level policy loss normalization and asymmetric clipping, and prompt
 327 difficulty/length curriculum. Building upon implementation of concurrent work (Setlur et al., 2025),
 328 we employ a two stage curriculum where we partition the DeepScaleR (Luo et al., 2025) mixture
 329 by success rate of the base model into three sets: (1) easy, (2) medium, and (3) hard, where we
 330 fine-tune first on easy problems with an 8K token budget and then on medium problems. We utilize
 331 the hard split as a held out, evaluation subset, which we denote as “**DeepScaleR [Hard]**”. We outline
 332 hyperparameters and details in Appendix A.1 and provide a pseudocode in Algorithm 1.

333 **Summary: RLAD method design**

335 RLAD jointly optimizes the abstraction generator π_θ^{abs} and solution generator π_θ^{sol} with RL,
 336 using reward functions in Equation 3. These reward functions incentivize π_θ^{sol} to utilize
 337 abstractions and incentivize π_θ^{abs} to propose useful abstractions per problem.

339 **6 EXPERIMENTAL EVALUATION**

341 The goal of our experiments is to evaluate the efficacy of RLAD in improving the reasoning capabilities
 342 of LLMs through abstraction-guided solution generation. Specifically, we aim to answer the following
 343 research questions: (1) Does RLAD improve pass@1 accuracy across several reasoning benchmarks
 344 compared to direct solution generation?, (2) How does RLAD scale as more abstractions and solutions
 345 are generated?, and (3) What makes the generated abstractions useful, how faithfully are they followed,
 346 and how do they guide and improve solution generation? We compare RLAD with strong models on
 347 three math reasoning benchmarks: AMC 2023, AIME 2025, and DeepScaleR Hard (Luo et al., 2025),
 348 which itself is a subset of hard problems from the OmniMATH mixture on which DeepSeek-R1
 349 distilled Qwen-32B model attains an accuracy of $\leq 10\%$. We also fine-tune an abstraction generator
 350 for the ARC-AGI program synthesis tasks, and conduct a similar comparison on 90 ARC puzzles
 351 evenly derived from the test sets of ARC-AGI 1, ARC-AGI 2, and BARC (Li et al., 2024). We also
 352 perform several ablations to better understand abstractions produced by RLAD.

353 **6.1 MAIN PERFORMANCE RESULTS ON MATH REASONING BENCHMARKS**

354 We evaluate RLAD under three settings: (1) **w/o abs**, without abstractions; (2) **w/ abs (avg)**, average
 355 performance over generations conditioned on 4 proposed abstractions per problem; and (3) **w/ abs**
 356 (**best**): using the best-performing abstraction (in a set of 4 proposed abstractions per problem). We
 357 observe that RLAD outperforms the base model and variant fine-tuned with RL on the same prompts
 358 via DAPO (Yu et al., 2025) without any abstractions, across all settings and benchmarks (Table 2).
 359 This highlights that RLAD can propose and leverage abstractions to improve reasoning performance.
 360 Interestingly, we also note that these performance gains are not limited to abstraction-conditioned
 361 inference: even in the **w/o abs** setting, where no abstraction is provided during inference, RLAD
 362 improves over the prior methods, when trained with abstractions via RLAD. This suggests that
 363 exposure to diverse abstractions during training enhances the model’s general reasoning ability. We
 364 observe similar trends on additional benchmarks, including AIME 2024 and HMMT 2025 (see
 365 Appendix B.1), where RLAD improves in the **w/o abs** setting. Finally, in Appendix D, we also
 366 measure the performance of RLAD when different token budgets are allowed for reasoning – while
 367 Table 2 measures performance at a budget of 32K tokens, we also measure performance at 8K and
 368 16K budgets and find RLAD to be more effective than other approaches.

Approach	AIME 2025			DeepScaleR [Hard]			AMC 2023		
	w/o abs (avg)	w/ abs (avg)	w/ abs (best)	w/o abs (avg)	w/ abs (avg)	w/ abs (best)	w/o abs (avg)	w/ abs (avg)	w/ abs (best)
Qwen-3-1.7B	33.75	36.25	40.00	20.21	22.14	32.50	86.41	78.01	84.53
+ DAPO	37.92	34.90	39.79	21.67	21.88	33.54	86.41	81.99	88.44
+ RLAD (Ours)	38.04	42.45	48.33	23.54	24.84	35.54	87.25	88.35	91.72

376 **Table 2: Accuracy on math reasoning benchmarks.** RLAD achieves consistent gains in abstraction-conditioned
 377 and w/o abstractions. Here, we measure performance on 3 domains (AIME 2025, DeepScaleR Hard, and AMC
 378 2023) with the base Qwen 3-1.7B model, DAPO, and RLAD. We measure performance without abstractions,
 379 with abstractions (pass@1 with 16 samples) and the best abstraction (pass@16), for each method type.

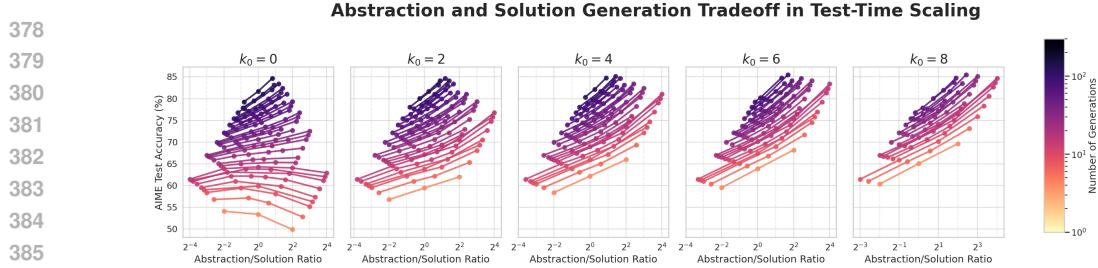


Figure 5: **Tradeoff of abstraction and solution generation on AIME 2025.** As the total inference compute budget increases (color scheme on the right), we find better performance efficiency when allocating our budget to abstraction generation rather than solution generation, for all values of normalization offset k_0 given to us.

6.2 UNDERSTANDING PROPERTIES OF RLAD

In this section, we conduct a number of additional experiments to understand the behavior of RLAD in regards to the usefulness of the proposed abstractions and algorithm performance when provided with a large budget on the total inference compute.

1) “Weak-to-strong” generalization of the abstraction generator. We next evaluate the weak-to-strong generalization of our method by pairing our trained abstraction generator with o4-mini as the solution generator. We use a fixed 24k generation token budget and 4 samples per question. Without abstractions, o4-mini achieves 80.38% pass@1, 82.26% pass@2, and 84.77% pass@4. Conditioning on both the problem and the proposed abstractions improves performance to 85.83% pass@1, 88.33% pass@2, and 90.00% pass@4 accuracy. Thus, conditioning on abstractions consistently yields higher pass@k accuracy compared to question-only conditioning, even for this strong reasoning model. These gains demonstrate that abstractions, though produced by a comparatively weaker model, can transfer effectively to a stronger solution generator, providing evidence that abstractions generalize in a weak-to-strong setting and enables downstream improvements without additional supervision or modifications to the strong model.

2) Compute tradeoffs b/w abstraction and solution generation.

Next, we study how to allocate compute between generating more abstractions and sampling solutions to attain maximal performance within a given inference budget. This corresponds to a “compute-optimal strategy” (Snell et al., 2024) for partitioning compute between abstraction and solution generation. If the model typically fails by making small local errors in its computation, then additional concise abstractions may not help it as much as simply trying again to generate a locally similar solution (optimizing “depth”). In contrast, if the model tends to pursue a seemingly plausible but incorrect approach and is unable to easily recover or switch approaches, then conditioning on diverse abstractions can help by offering alternative high-level approaches toward the correct answer. In other words, when the model has a tendency to explore “depth” over “breadth” of solution strategies, abstractions can help improve performance. With this intuition, we hypothesize that when the compute budget permits only a limited number of samples, allocating more compute to sampling multiple solutions but retaining only a few abstractions will improve success more. However, once pass@k for a single abstraction begins to saturate, performance gains are more likely to come from scaling the diversity of abstractions, which enables the model to explore qualitatively different regions of the solution space.

	1	4	16	64	256
DAPO, w/o abs (pass@ n^2)	0.37	0.51	0.65	0.77	0.82
RLAD, w/ abs (pass@ $n \times n$)	0.41	0.59	0.71	0.80	0.87

Table 3: **Pass@k comparison under equal compute.** For each n , we compare n^2 solution samples (no abstractions) with n abstractions $\times n$ solutions per abstraction. Abstraction conditioning yields consistent improvements.

To validate this hypothesis, we plot *iso-compute* scaling curves under a fixed compute budget \mathcal{C} , we distribute between abstractions and solution generation. Specifically, we denote the number of abstractions as m and the number of solutions sampled per abstraction as k , such that $m \times k = \mathcal{C}$. To better isolate the utility of abstractions, we make several more projections of the iso-compute curve, each with a different “normalization offset” k_0 . This normalization offset k_0 *discounts* performance gains that stem from trying the problem again and making local modifications (e.g., small edits that do not require new strategy changes), but measures the performance gains that stem from major changes in content of a response. Concretely, for a non-zero k_0 , we plot *iso-compute* frontiers when $m \times (k - k_0) = \mathcal{C}$. This formulation captures the amount of compute that is spent on “meaningful”

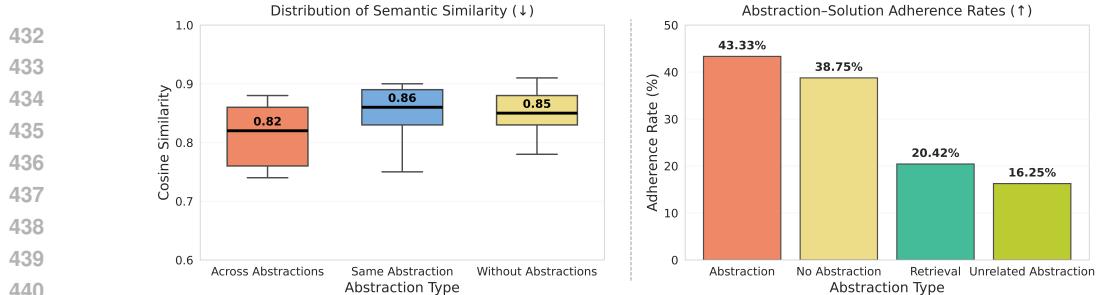


Figure 6: **Abstraction-conditioned solution generation analysis.** RLAD produces solutions with (left) greater semantic diversity across different abstractions and (right) higher abstraction adherence than baselines.

samples that go beyond the model’s local neighborhood. Figure 5 shows these *iso-compute* frontiers for different values of the compute budget \mathcal{C} . The x-axis plots the ratio between abstractions and adjusted solutions, $m/(k - k_0)$. Each curve corresponds to a fixed total compute budget.

In Figure 5, across $k_0 \in \{0, 2, 4, 6, 8\}$, shifting compute toward abstractions consistently yields greater improvements than allocating the same additional compute to solution refinements, especially as the total compute budget increases. **This supports the conclusion that once local errors in the chain-of-thought have been addressed, it is more effective to increase the diversity of strategies used via abstraction conditioning rather than to continue to scale up long CoT sampling alone.**

To obtain additional evidence to support this result, we also evaluate pass@k of DAPO and RLAD for a fixed compute budget for both abstraction and non-abstraction generation on AIME 2025. Here, for a given value n , we compare two approaches: (1) sampling n^2 solutions without abstractions, and (2) sampling n abstractions and then n solutions per abstraction. In Table 3, abstraction conditioning consistently outperforms pure solution sampling. For example, at $n = 16$, abstraction conditioning achieves a pass@k of 0.71 versus 0.65 without abstractions, and at $n = 256$ the gap widens to 0.87 versus 0.82. This corroborates the efficacy of abstractions, even under matched compute budgets.

3) Understanding the behavior of the abstraction-conditioned solution generator. A desirable property of the solution generator is the ability to follow the proposed abstractions. To study this, we prompt (o4-mini) to classify whether a particular solution trace produced by a trained solution generator closely adheres to a given abstraction. We ask for a binary decision on each pair of abstraction and solution, and measure the adherence rate across 200 such pairs. In Figure 6 (right), we report the adherence rates under four conditions: (1) Abstraction, where we measure adherence rates between an abstraction and a solution generated by conditioning on this abstraction itself; (2) No abstraction, where we measure the adherence rates between an abstraction and a solution generated by conditioning on the problem without any abstraction; (3) Retrieval, where we measure adherence rates between an abstraction and a semantically similar prior solution to the problem; and (4) Unrelated abstraction, where we measure the adherence rates between an abstraction and a solution generated via a different abstraction. We find that the Abstraction condition achieves the highest adherence rate. Intuitively, this means that the trained solution generator is detected to be more likely to follow the strategy or guidance of a given abstraction. Additionally, we measure the semantic similarity of solutions generated without abstraction conditioning, conditioned on the same abstraction, and across abstractions and make a similar conclusion.

Takeaways: Experimental Results

RLAD outperforms RL fine-tuning approaches that do not propose or leverage abstractions on math reasoning. Jointly scaling the number of abstractions and solution samples enables continued performance gains even when scaling solutions alone begins to saturate.

7 DISCUSSION AND PERSPECTIVES ON FUTURE WORK

We introduce reasoning abstractions: concise natural language representations of procedural and factual knowledge, as a way to expand LLM reasoning strategies. Our method, RLAD, jointly trains an abstraction generator and an abstraction-conditioned solution generator, achieving consistent gains on mathematical reasoning benchmarks. We show that allocating compute to diverse abstractions yields larger improvements than simply increasing solution sampling, offering a complementary axis for scaling test-time compute. While our study focuses on math tasks, extending abstractions to broader reasoning and unifying abstraction and solution generation remain open directions.

486 8 ETHICS STATEMENT
487488 Our empirical evaluation focuses primarily on mathematical reasoning (e.g., AIME, AMC, Deep-
489 ScaleR) and abstract pattern recognition (ARC-AGI). These domains are generally considered
490 low-risk, as they utilize publicly available benchmarks that do not contain private or sensitive per-
491 sonal data. The risk of generating toxic or harmful content in these structured reasoning contexts is
492 minimal compared to open-ended generation tasks. However, we acknowledge the broader ethical
493 considerations associated with advancing LLM capabilities. While enhanced reasoning abilities
494 offer significant potential for scientific and educational benefits, they also carry dual-use risks if
495 applied to malicious activities, such as generating sophisticated disinformation. We encourage the
496 continued development of robust safety and alignment protocols alongside improvements in reasoning
497 capabilities, which are jointly correlated as evidenced in prior work (Kim et al., 2025). Finally, we
498 utilized LLMs such as GPT5/Gemini for minor rewritings throughout the paper for better readability.
499500 9 REPRODUCIBILITY STATEMENT
501502 We are committed to ensuring the transparency and reproducibility of our research. To facilitate the
503 replication of our results, we have provided comprehensive details regarding our methodology and
504 experimental setup in the Appendix and main paper. In particular, the RLAD framework, including the
505 two-player RL training paradigm and the specific reward design (Equations 1, 2 and 3), is detailed in
506 Section 5. The procedure for warmstarting the models via Supervised Fine-Tuning (SFT), including
507 the generation of seed abstractions using stronger models (o4-mini and GPT 4.1-mini) and the filtering
508 process, is described in Section 5. The implementation details of the curriculum learning strategy
509 are also provided in Section 5 and analyzed in Appendix B. Additionally, detailed implementation
510 information is provided in the Appendix. Here, the pseudocode for the joint RL training process is
511 available in Appendix A.1 (Algorithm 1). All key training hyperparameters used for the RL methods
512 (DAPO, RFT, RPO) are listed in Appendix A.1 (Table 4). Our experiments utilize publicly available
513 base models from the Qwen3 series (Qwen Team, 2025). The evaluations are conducted on standard
514 reasoning benchmarks: AIME 2025 (Mathematical Association of America, 2025), DeepScaleR (Luo
515 et al., 2025), AMC 2023 (Zeng et al., 2025), and ARC-AGI (Li et al., 2024), as described in Section
516 6, with additional results in Appendix B. Qualitative examples and the prompt used for abstraction
517 classification are included in Appendix D.
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 REFERENCES
541

542 Neel Alex, Eli Lifland, Lewis Tunstall, Abhishek Thakur, Pegah Maham, C Jess Riedel, Emmie Hine,
543 Carolyn Ashurst, Paul Sedille, Alexis Carlier, et al. Raft: A real-world few-shot text classification
544 benchmark. *arXiv preprint arXiv:2109.14076*, 2021.

545 Anonymous. Optimizing inference-time reasoning in LLMs via retrieval-augmented reflection, 2025.
546 URL <https://openreview.net/forum?id=E1YRG3pJcv>.

547

548 Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican,
549 George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al.
550 Improving language models by retrieving from trillions of tokens. In *International conference on
551 machine learning*, pp. 2206–2240. PMLR, 2022.

552

553 Eugene Charniak and Mark Johnson. Coarse-to-fine n-best parsing and maxent discriminative
554 reranking. In *Proceedings of the 43rd Annual Meeting on Association for Computational Lin-
555 guistics, ACL '05*, pp. 173–180, USA, 2005. Association for Computational Linguistics. doi:
556 10.3115/1219840.1219862. URL <https://doi.org/10.3115/1219840.1219862>.

557

558 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Ka-
559 plan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen
560 Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray,
561 Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens
562 Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis,
563 Elizabeth Barnes, Ariel Herbert-Voss, William Heben Guss, Alex Nichol, Alex Paino, Nikolas
564 Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher
565 Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford,
566 Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario
567 Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language
models trained on code, 2021. URL <https://arxiv.org/abs/2107.03374>.

568

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

594 Xidong Feng, Ziyu Wan, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and Jun
 595 Wang. Alphazero-like tree-search can guide large language model decoding and training, 2024.
 596 URL <https://arxiv.org/abs/2309.17179>.

597

598 Chrisantha Fernando, Dylan Banarse, Henryk Michalewski, Simon Osindero, and Tim Rock-
 599 täschel. Promptbreeder: Self-referential self-improvement via prompt evolution. *arXiv preprint*
 600 *arXiv:2309.16797*, 2023.

601 Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen.
 602 Critic: Large language models can self-correct with tool-interactive critiquing. *arXiv preprint*
 603 *arXiv:2305.11738*, 2023.

604

605 Neel Guha, Julian Nyarko, Daniel Ho, Christopher Ré, Adam Chilton, Alex Chohlas-Wood, Austin
 606 Peters, Brandon Waldon, Daniel Rockmore, Diego Zambrano, et al. Legalbench: A collaboratively
 607 built benchmark for measuring legal reasoning in large language models. *Advances in Neural*
 608 *Information Processing Systems*, 36:44123–44279, 2023.

609 Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu.
 610 Reasoning with language model is planning with world model, 2023. URL <https://arxiv.org/abs/2305.14992>.

611

612 Namgyu Ho, Laura Schmid, and Se-Young Yun. Large language models are reasoning teachers, 2023.
 613 URL <https://arxiv.org/abs/2212.10071>.

614

615 Taeyoun Kim, Fahim Tajwar, Aditi Raghunathan, and Aviral Kumar. Reasoning as an adaptive
 616 defense for safety, 2025. URL <https://arxiv.org/abs/2507.00971>.

617

618 Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate
 619 Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, Lei M Zhang, Kay McKinney, Disha
 620 Shrivastava, Cosmin Paduraru, George Tucker, Doina Precup, Feryal Behbahani, and Aleksandra
 621 Faust. Training language models to self-correct via reinforcement learning, 2024. URL <https://arxiv.org/abs/2409.12917>.

622

623 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,
 624 Heinrich Kütller, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented genera-
 625 tion for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:
 626 9459–9474, 2020.

627

628 Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbulin, and Bernard Ghanem.
 629 Camel: Communicative agents for "mind" exploration of large language model society, 2023. URL
<https://arxiv.org/abs/2303.17760>.

630

631 Wen-Ding Li, Keya Hu, Carter Larsen, Yuqing Wu, Simon Alford, Caleb Woo, Spencer M. Dunn,
 632 Hao Tang, Michelangelo Naim, Dat Nguyen, Wei-Long Zheng, Zenna Tavares, Yewen Pu, and
 633 Kevin Ellis. Combining induction and transduction for abstract reasoning, 2024. URL <https://arxiv.org/abs/2411.02272>.

634

635 Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and
 636 Zhicheng Dou. Search-o1: Agentic search-enhanced large reasoning models, 2025. URL <https://arxiv.org/abs/2501.05366>.

637

638 Kevin Lin, Charlie Snell, Yu Wang, Charles Packer, Sarah Wooders, Ion Stoica, and Joseph E Gonzalez.
 639 Sleep-time compute: Beyond inference scaling at test-time. *arXiv preprint arXiv:2504.13171*,
 640 2025.

641

642 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Tang, Manan Roongta, Colin Cai,
 643 Jeffrey Luo, Tianjun Zhang, Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing o1-
 644 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-19681902c1468005b>
 645 2025. Notion Blog.

646

647 Jingyuan Ma, Damai Dai, Lei Sha, and Zhifang Sui. Large language models are unconscious of
 648 unreasonability in math problems. *arXiv preprint arXiv:2403.19346*, 2024.

648 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri
 649 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad
 650 Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-Refine:
 651 Iterative Refinement with Self-Feedback. *arXiv e-prints*, art. arXiv:2303.17651, March 2023. doi:
 652 10.48550/arXiv.2303.17651.

653 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri
 654 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement
 655 with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.
 656

657 Mathematical Association of America. 2025 American Invitational Mathematics Examination (AIME)
 658 I: Problems and Solutions, February 2025. URL https://artofproblemsolving.com/wiki/index.php/2025_AIME_I. Art of Problem Solving Wiki entry.
 659

660 Rakesh R. Menon, Sayan Ghosh, and Shashank Srivastava. Clues a benchmark for learning classifiers
 661 using natural language explanations. 2022.

662

663 Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad
 664 Farajtabar. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large
 665 language models. *arXiv preprint arXiv:2410.05229*, 2024.

666 Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David
 667 Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and
 668 Augustus Odena. Show your work: Scratchpads for intermediate computation with language
 669 models, 2021. URL <https://arxiv.org/abs/2112.00114>.
 670

671 Jiayi Pan, Xiuyu Li, Long Lian, Charlie Snell, Yifei Zhou, Adam Yala, Trevor Darrell, Kurt Keutzer,
 672 and Alane Suhr. Learning adaptive parallel reasoning with language models. *arXiv preprint
 673 arXiv:2504.15466*, 2025.

674 Richard Yuanzhe Pang, Weizhe Yuan, Kyunghyun Cho, He He, Sainbayar Sukhbaatar, and Jason
 675 Weston. Iterative reasoning preference optimization. *arXiv preprint arXiv:2404.19733*, 2024.
 676

677 Ivo Petrov, Jasper Dekoninck, Lyuben Baltadzhiev, Maria Drencheva, Kristian Minchev, Mislav
 678 Balunović, Nikola Jovanović, and Martin Vechev. Proof or bluff? evaluating llms on 2025 usa
 679 math olympiad. *arXiv preprint arXiv:2503.21934*, 2025.

680 Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chenguang Zhu, and Michael Zeng. Automatic prompt
 681 optimization with "gradient descent" and beam search. *arXiv preprint arXiv:2305.03495*, 2023.
 682

683 Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive Introspection: Teaching
 684 Language Model Agents How to Self-Improve. *arXiv e-prints*, art. arXiv:2407.18219, July 2024.
 685 doi: 10.48550/arXiv.2407.18219.

686 Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching
 687 language model agents how to self-improve. *arXiv preprint arXiv:2407.18219*, 2024.

688

689 Yuxiao Qu, Matthew YR Yang, Amritth Setlur, Lewis Tunstall, Edward Emanuel Beeching, Ruslan
 690 Salakhutdinov, and Aviral Kumar. Optimizing test-time compute via meta reinforcement fine-
 691 tuning. *arXiv preprint arXiv:2503.07572*, 2025.

692 Qwen Team. Qwen3 technical report. Technical report, Qwen, May 2025. URL
 693 <https://huggingface.co/Qwen>. Available at: <https://huggingface.co/Qwen>,
 694 <https://modelscope.cn/organization/qwen>, <https://github.com/QwenLM/Qwen3>.
 695

696 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 697 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances
 698 in Neural Information Processing Systems*, 36:53728–53741, 2023.

699 Marlin B. Schäfer, Frank Ohme, and Alexander H. Nitz. Detection of gravitational-wave signals
 700 from binary neutron star mergers using machine learning. *Physical Review D*, 102(6), September
 701 2020. ISSN 2470-0029. doi: 10.1103/physrevd.102.063015. URL <http://dx.doi.org/10.1103/PhysRevD.102.063015>.

702 Amrit Setlur, Matthew Y. R. Yang, Charlie Snell, Jeremy Greer, Ian Wu, Virginia Smith, Max
 703 Simchowitz, and Aviral Kumar. e3: Learning to explore enables extrapolation of test-time compute
 704 for llms. 2025. URL <https://arxiv.org/abs/2506.09026>.

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

708 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
 709 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing
 710 Systems*, 36:8634–8652, 2023.

711 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling LLM Test-Time Compute Opti-
 712 mally can be More Effective than Scaling Model Parameters. *arXiv e-prints*, art. arXiv:2408.03314,
 713 August 2024. doi: 10.48550/arXiv.2408.03314.

714 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 715 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.

716 Mirac Suzgun, Mert Yuksekgonul, Federico Bianchi, Dan Jurafsky, and James Zou. Dynamic
 717 cheatsheet: Test-time learning with adaptive memory. *arXiv preprint arXiv:2504.07952*, 2025.

718 Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Interleaving retrieval
 719 with chain-of-thought reasoning for knowledge-intensive multi-step questions. *arXiv preprint
 720 arXiv:2212.10509*, 2022.

721 Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia
 722 Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process- and
 723 outcome-based feedback, 2022. URL <https://arxiv.org/abs/2211.14275>.

724 Prakhar Verma, Sukruta Prakash Midigesi, Gaurav Sinha, Arno Solin, Nagarajan Natarajan, and
 725 Amit Sharma. Plan×RAG: Planning-guided Retrieval Augmented Generation. *arXiv e-prints*, art.
 726 arXiv:2410.20753, October 2024. doi: 10.48550/arXiv.2410.20753.

727 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 728 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models,
 729 2023. URL <https://arxiv.org/abs/2203.11171>.

730 Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen.
 731 Large language models as optimizers. *arXiv preprint arXiv:2309.03409*, 2023.

732 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik
 733 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models, 2023a.
 734 URL <https://arxiv.org/abs/2305.10601>.

735 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
 736 React: Synergizing reasoning and acting in language models, 2023b. URL <https://arxiv.org/abs/2210.03629>.

737 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
 738 Liu, Lingjun Liu, Xin Liu, et al. Dapo: An open-source llm reinforcement learning system at scale.
 739 *arXiv preprint arXiv:2503.14476*, 2025.

740 Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. Scaling
 741 relationship on learning mathematical reasoning with large language models. *arXiv preprint
 742 arXiv:2308.01825*, 2023.

743 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. Does
 744 reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv
 745 preprint arXiv:2504.13837*, 2025.

746 Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Pan Lu, Zhi Huang, Carlos Guestrin,
 747 and James Zou. Optimizing generative ai by backpropagating language model feedback. *Nature*,
 748 639(8055):609–616, 2025.

756 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with
757 reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022a.
758

759 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. Star: Bootstrapping reasoning with
760 reasoning, 2022b. URL <https://arxiv.org/abs/2203.14465>.
761

762 Changyu Zeng, Yifan Wang, Zimu Wang, Wei Wang, Zhengni Yang, Muyi Bao, Jiming Xiao, Ahn
763 Nguyen, and Yutao Yue. Numina: A natural understanding benchmark for multi-dimensional
764 intelligence and numerical reasoning abilities, 2025. URL <https://arxiv.org/abs/2509.16656>.
765

766 Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. Expel: Llm
767 agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
768 volume 38, pp. 19632–19642, 2024.
769

770 Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and
771 Jimmy Ba. Large language models are human-level prompt engineers. In *The Eleventh International
772 Conference on Learning Representations*, 2022.
773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 Appendices

811 A EXPERIMENTAL DETAILS

812 A.1 PSEUDOCODE FOR RLAD

813 **Algorithm 1** Joint RL Training of π_θ^{abs} and π_θ^{sol}

814 **Require:** Policies $\pi_\theta^{\text{abs}}(\mathbf{z} \mid \mathbf{x})$, $\pi_\theta^{\text{sol}}(\tilde{\mathbf{y}} \mid \mathbf{x}, \mathbf{z})$ Datasets $\mathcal{D}_{\pi_\theta^{\text{abs}}}$, $\mathcal{D}_{\pi_\theta^{\text{sol}}}$; rates $\alpha_{\pi_\theta^{\text{abs}}}, \alpha_{\pi_\theta^{\text{sol}}}$; batch sizes N, M ;
 815 epochs E

816 1: Initialize $\pi_\theta^{\text{abs}}, \pi_\theta^{\text{sol}}$

817 2: **for** $e = 1$ to E **do** ▷ Update abstraction policy

818 3: **for** $\{\mathbf{x}_i\}_{i=1}^N \sim \mathcal{D}_{\pi_\theta^{\text{abs}}}$ **do**

819 4: $\mathbf{z}_i \sim \pi_\theta^{\text{abs}}(\cdot \mid \mathbf{x}_i)$

820 5: $r_i \leftarrow r_{\pi_\theta^{\text{sol}}}(\mathbf{x}_i, \mathbf{z}_i)$

821 6: $\pi_\theta^{\text{abs}} \leftarrow \pi_\theta^{\text{abs}} - \alpha_{\pi_\theta^{\text{abs}}} \nabla_{\pi_\theta^{\text{abs}}} \mathcal{L}_{\text{STAR/RPO}}(\pi_\theta^{\text{abs}}; \mathbf{x}_i, \mathbf{z}_i, r_i)$

822 7: **end for** ▷ Update solution policy

823 8: **for** $\{\mathbf{x}_j\}_{j=1}^M \sim \mathcal{D}_{\pi_\theta^{\text{sol}}}$ **do**

824 9: $\mathbf{z}_j \sim \pi_\theta^{\text{abs}}(\cdot \mid \mathbf{x}_j)$, $\tilde{\mathbf{y}}_j \sim \pi_\theta^{\text{sol}}(\cdot \mid \mathbf{x}_j, \mathbf{z}_j)$

825 10: $r_j \leftarrow r(\mathbf{x}_j, \mathbf{z}_j, \tilde{\mathbf{y}}_j)$

826 11: $\pi_\theta^{\text{sol}} \leftarrow \pi_\theta^{\text{sol}} - \alpha_{\pi_\theta^{\text{sol}}} \nabla_{\pi_\theta^{\text{sol}}} \mathcal{L}_{\text{GRPO}}(\pi_\theta^{\text{sol}}; \mathbf{x}_j, \mathbf{z}_j, \tilde{\mathbf{y}}_j, r_j)$

827 12: **end for**

828 13: **end for**

834 A.2 HYPERPARAMETERS

835 Hyperparameter	836 Value
837 algorithm	838 DaPO (Yu et al., 2025)
839 training steps	840 100
840 epochs	841 10
841 train batch size	842 128
842 max prompt length	843 3072
843 max response length	844 16384
844 max extrapolation length	845 32768
845 learning rate	846 1e-6
846 clip ratio (low / high)	847 0.2 / 0.5
847 entropy coefficient	848 0.001
848 KL loss coefficient	849 0.001
849 KL loss type	850 low_var_kl
850 sampling temperature (train / val)	851 0.6 / 0.6
851 samples per prompt (train / val)	852 16 / 8
852 max batched tokens	853 32768

854 Table 4: Key training hyperparameters used in RLAD.

855 B ADDITIONAL EXPERIMENTAL RESULTS

856 B.1 RLAD’S W/ ABS PERFORMANCE ON AIME 2024 AND HMMT 2025

857 In this section, we evaluate the performance of the base model (Qwen-3-1.7B), GRPO-enhanced
 858 model, and our proposed method RLAD on two math reasoning benchmarks: AIME 2024 and HMMT
 859 2025. As shown in Table 5, our method achieves the best performance across both datasets.

860 It is important to note that RLAD is trained using access to abstractions, yet it also generalizes better
 861 even when evaluated without abstraction. This suggests that RLAD does not merely overfit to the

abstraction format but instead learns to effectively leverage high-level procedural guidance, leading to better generalization on challenging reasoning benchmarks.

Approach	AIME 2024	HMMT 2025
Qwen-3-1.7B	48.54	22.50
+ DaPO	44.17	23.13
+ RLAD	51.46	23.75

Table 5: **RLAD’s w/ abs performance on AIME 2024 and HMMT 2025.** RLAD outperforms both DaPO and the base Qwen model on the AIME 2024 and HMMT 2025 benchmarks.

B.2 DESIGN CHOICE ABLATIONS

In this section, we conduct ablation studies to isolate the contributions of key components in RLAD. We investigate three primary design choices, with results summarized in Table 6: **(a)** utilizing a curriculum training strategy, **(b)** including prompts not annotated with an abstraction, and **(c)** applying reward masking to those non-annotated prompts.

Curriculum training refers to a staged process where the model first learns from simpler problems and gradually transitions to harder ones. We use the protocol from Setlur et al. (2025) as inspiration, who demonstrated its effectiveness for direct math problem-solving. In our setting, which incorporates abstractions, curriculum training also proves beneficial, improving both average and best-case performance from 0.38 and 0.43 to 0.41 and 0.48, respectively, compared to non-curriculum training.

Next, we analyze the practice of **including prompts without abstractions** and applying **reward masking**. Including a small fraction of these "no-abstraction" prompts is intended to better condition the solution-generator on the abstractions when they are present. However, this risks the model learning a shortcut by simply ignoring the abstractions. To mitigate this, we apply reward masking: for completions on no-abstraction prompts, we nullify the policy reward by zeroing out the advantage, while retaining the KL penalty for regularization. This prevents the model from over-optimizing on examples that lack abstractions, a behavior that would otherwise hinder generalization.

Our findings confirm the efficacy of this combined approach. As shown in Table 6, including no-abstraction prompts with reward masking is critical for performance. Ultimately, the combination of all three design choices—curriculum training, the inclusion of no-abstraction prompts, and reward masking—significantly outperforms alternative configurations.

Approach	Design Choice			AIME 2025	
	curriculum training	including no-abstraction prompt	reward masking	w/ abs (avg)	w/ abs (best)
variant 1	✗	✓	✗	36.51	42.29
variant 2	✗	✗	-	37.08	42.50
variant 3	✗	✓	✓	37.50	43.33
RLAD	✓	✓	✓	42.45	48.33

Table 6: **Design Choices in RLAD.** We isolate the effects of curriculum training, no-abstraction inclusion, and reward masking. The full method achieves the strongest performance under abstraction-conditioned evaluation.

B.3 FULL RESULTS FOR ABSTRACTIONS IN NON-MATH DOMAINS

As seen in Table 7, conditioning on abstractions helps for 37 domains, where the average and best abstractions outperform standard prompting by 18.0% and 30.0% on average, respectively.

For ARC-AGI, we warmstart the abstraction generator model with synthetically augmented human annotations from the BARC dataset (Li et al., 2024). In Table 1, we report results using an abstraction generator and a solution generator instantiated as Qwen3-4B.

C ABSTRACTIONS IN OTHER DOMAINS

D QUALITATIVE EXAMPLES OF MATH REASONING ABSTRACTIONS

	Breast Cancer Detection	Tweet Hate Speech Detection	Corporate Lobbying Relevance	Bank Note Authentication
Abstraction	Classify breast masses as malignant or benign using BI-RADS, shape, and margin criteria. ...If BI-RADS ≥ 5 , then malignant. ...If BI-RADS = 4 AND shape = irregular AND margin = ill-defined, then malignant...	...If the text contains explicit derogatory slurs (e.g., b***, c***, s***, h***, r*****), classify as hate speech. ...If the text degrades or dehumanizes a protected group (nationality, race, religion...)	...If the bill addresses regulation, labeling, pricing, reimbursement, R&D funding, or licensing of the company's core products or services, label "Yes." ...Else if the bill alters taxes, credits, bonds, infrastructure...	1. If variance > 4 , predict Fake. 2. Else if variance < -3 , predict Original. 3. Else if skewness > 5 , predict Fake. 4. Else if entropy > 0 and skewness > 1 , predict Fake.
GPT-4o-mini	45%	60%	88%	45%
GPT-4o-mini + Abs	90%	90%	94%	100%

Figure 7: *Examples of good reasoning abstractions in non-math domains*. Adding the abstraction to the prompt of GPT-4o-mini consistently improves performance on unseen instances.

Interpreting discovered abstractions. As discussed in Appendix D, we classify each model-generated abstraction into four categories for ease of interpretability: (1) **caution alert** that warns the solution generator to avoid a specific approach; (2) **productive launchpoint** that suggests strategic framings or problem reformulations that open high-potential solution paths; (3) **blind-follow trajectory** that prescribes repeatable, step-by-step procedures executable without further insight; and (4) **structural shortcut** that leverages abstract insights or invariants to collapse multiple reasoning steps into a single leap. In Figure 8, we show that after training via RLAD, the distribution over these categories shifts, with a notable increase in blind-follow abstractions, which a stronger reasoning model classifies as an effective reasoning path to a successful solution as seen in Appendix D.

D.1 PROMPT FOR ABSTRACTION CLASSIFICATION

We prompt GPT-4o-mini with the following prompt template to classify each abstraction into one of four categories.

Post-hoc abstraction classifier prompt

You are a abstraction classifier. You will be given a problem-solving heuristic or abstraction used for mathematical reasoning. Your task is to classify it into exactly one of the following mutually exclusive categories, based on the primary cognitive function the heuristic serves.

(A) Caution alert: any abstraction that warns the reader to double-check a specific aspect of their solution or to not take a specific approach to the problem.

(B) Productive launchpoint: an early move or framing that opens up high-potential trajectories. Examples include clever reformulations or symmetries.

(C) Blind-follow trajectory: a description of a repeatable, sequential path that can be reliably followed to solve the problem. Examples include plug-and-play formulas that can be followed blindly, without insight. Do not choose this is further reasoning is required to solve the problem.

(D) Structural shortcut: a conceptual move that collapses multiple graph paths into a single jump via insight or abstraction. This can include introducing invariants.

(E) Other: a abstraction that does not fit into the above categories.

Give a 1-2 sentence explanation for your classification, and end your answer with exactly one of: (A), (B), (C), (D), or (E).

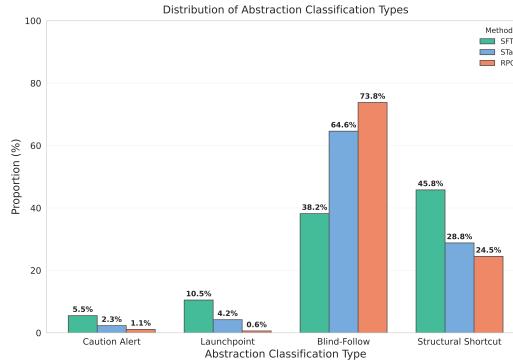


Figure 8: **Abstraction Categorization** RLAD produces a diverse characterization of abstractions, which we characterize by prompting o4-mini.

As seen in Figure 8, the distribution of abstraction types changes significantly after training via RLAD. The proportion of 'Blind-Follow' abstractions increases from approximately 38.2% (SFT) to 73.8% (RPO), indicating that the model is more likely to generate abstractions that serve as effective reasoning paths to successful solutions.

```

972 --
973 abstraction:
974 {abstraction}
975
976

```

977 D.2 EXAMPLE FOR SHORT ABSTRACTION

978 Examples of Short Abstractions

981 **Question:** The banker's gain of a certain sum due 3 years
 982 hence at 10% per annum is Rs. 36. What is the present worth?

983 **Hint:**

- 984 - Consider the proportional relationship between the true
 985 discount and the present worth, and how the banker's gain
 986 relates to these terms.
- 987 - Think of the banker's gain as a component of the true
 988 discount and how it relates to the present worth through the
 989 interest rate and time period.

990 **Question:** A train 125 m long passes a man running at 15
 991 km/hr in the same direction in which the train is going, in
 992 10 seconds. What is the speed of the train?

993 **Hint:**

- 994 - Think about the speed difference between the train and the
 995 man, and how that difference relates to the time it takes for
 996 the entire train to pass by.
- 997 - Consider the difference in speeds as the key factor in
 998 determining the train's speed relative to the man, and use
 999 the time taken to pass the man to find this difference.

1000 D.3 EXAMPLE FOR EACH ABSTRACTION CATEGORY

1001 Below, we show examples of abstractions classified into the four categories above.

1002 Examples of (A) Caution alert

1003 <description>Always record forbidden values from denominators before
 1004 and after manipulation. After solving the polynomial, discard
 1005 any roots that make a denominator zero or that do not satisfy the
 1006 original equation, to avoid extraneous solutions.</description>
 1007 <example>In the equation $(x+2)/(2x-1) = x-3$, $2x-1$ cannot be zero
 1008 (so x is not $\frac{1}{2}$). If solving yields $x=\frac{1}{2}$ or any root that makes any
 1009 denominator zero, reject it. Then verify the accepted roots in the
 1010 original equation.</example>

1011 <description>Keep units consistent when moving between area and
 1012 length or summing lengths. After extracting a length from an area
 1013 (via square root), ensure subsequent arithmetic stays in the same
 1014 unit to avoid scaling errors. </description>

1015 <example>If a square's area is 10000 cm^2 , its side is $\sqrt{10000} =$
 1016 100 cm . To express in meters, convert 100 cm to 1 m . All later
 1017 distances computed with that side length must be in meters to
 1018 remain consistent.</example>

1019 Examples of (B) Productive launchpoint

1020 <description>Translate comparative statements into algebraic
 1021 equations using the chosen variables. Phrases like "twice
 1022 as many" or "one less than" correspond to multiplication or
 1023 subtraction.

```

1026
1027     addition/subtraction expressions. This step captures the core
1028     relationship in a solvable form.</description>
1029     <example>If the problem states "Group A has twice as many as Group
1030     B," write the equation  $x = 2y$ . For "Group B has three fewer than
1031     Group C," you would write  $y = z - 3$ .</example>
1032     <description>Select one variable as a parameter (often setting it
1033     to 1 or keeping it symbolic) to express all other variables in
1034     terms of it. This reduces the number of independent symbols and
1035     streamlines substitutions.</description>
1036     <example>Given  $p/q = 3$  and  $r/q = 2$ , choose  $q$  as the base variable.
1037     Write  $p = 3q$  and  $r = 2q$ , so all expressions involving  $p$  and  $r$  can
1038     be handled through  $q$  alone.</example>
1039
1040
1041
1042 Examples of (C) Blind-follow trajectory
1043
1044     <description>Logarithms offer a streamlined way to compute
1045     floor-based digit counts: for  $y > 0$ , the number of integer digits is
1046      $\text{floor}(\log_{10} y) + 1$ . Use this to handle arbitrary exponents without
1047     juggling large powers explicitly.</description>
1048     <example>To count digits of  $y = x^7$ , compute  $d = \text{floor}(7 * \log_{10}$ 
1049      $x) + 1$ . If  $x=2.5$ , then  $d = \text{floor}(7 * \log_{10}(2.5)) + 1 = 2 + 1 = 3$ 
1050     digits.</example>
1051
1052     <description>The mean of a set equals its total sum divided by its
1053     number of elements. Use this to move between sums and averages
1054     when counts or totals are known. It works because "average" is
1055     defined as that ratio.</description>
1056     <example>Suppose a subset has  $k$  items with mean  $m$ . Then its total
1057     sum is  $S = k \cdot m$ . Conversely, if you know the sum  $S$  and the count  $k$ ,
1058     the mean is  $m = S/k$ . For instance, if 5 items average to 10, their
1059     total is  $5 \times 10 = 50$ , and if you later learn the total is 60 for 6
1060     items, the new mean becomes  $60/6 = 10$ .</example>
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

```

Examples of (C) Blind-follow trajectory

```

1043     <description>Logarithms offer a streamlined way to compute
1044     floor-based digit counts: for  $y > 0$ , the number of integer digits is
1045      $\text{floor}(\log_{10} y) + 1$ . Use this to handle arbitrary exponents without
1046     juggling large powers explicitly.</description>
1047     <example>To count digits of  $y = x^7$ , compute  $d = \text{floor}(7 * \log_{10}$ 
1048      $x) + 1$ . If  $x=2.5$ , then  $d = \text{floor}(7 * \log_{10}(2.5)) + 1 = 2 + 1 = 3$ 
1049     digits.</example>
1050
1051     <description>The mean of a set equals its total sum divided by its
1052     number of elements. Use this to move between sums and averages
1053     when counts or totals are known. It works because "average" is
1054     defined as that ratio.</description>
1055     <example>Suppose a subset has  $k$  items with mean  $m$ . Then its total
1056     sum is  $S = k \cdot m$ . Conversely, if you know the sum  $S$  and the count  $k$ ,
1057     the mean is  $m = S/k$ . For instance, if 5 items average to 10, their
1058     total is  $5 \times 10 = 50$ , and if you later learn the total is 60 for 6
1059     items, the new mean becomes  $60/6 = 10$ .</example>
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

```

Examples of (D) Structural shortcut

```

1062     <description>When the same distance appears in multiple geometric
1063     roles (e.g., as radius to a vertex and to a tangent point),
1064     express it in different algebraic forms and equate them. Solving
1065     the resulting equation produces the unknown variable, which then
1066     gives the desired length.</description>
1067     <example>If  $r$  is both the distance from  $O$  to a vertex ( $r = \sqrt{x^2 + (L/2)^2}$ ) and the distance from  $O$  to the tangent point ( $r = f(x)$ ), set  $\sqrt{x^2 + (L/2)^2} = f(x)$ . Solving this equation for  $x$  and back-substituting determines  $r$  explicitly, closing the geometric
1068     problem with an algebraic solution.</example>
1069
1070
1071     <description>Use the perimeter constraint  $a+b+c=P$  to eliminate one
1072     variable, e.g. set  $c=P-a-b$ , reducing the problem to two degrees
1073     of freedom. This simplification turns the three-variable Heron
1074     expression into a function of  $a$  and  $b$  alone, facilitating analysis
1075     or enumeration.</description>
1076     <example>For a target perimeter  $P=10$ , one writes  $c=10-a-b$ .
1077     Substituting into Heron's formula yields  $A(a,b)=\sqrt{5 * (5-a) * (5-b) * (a+b-5)}$ , which is now a two-variable function to study
1078     instead of three.</example>
1079

```

1080
1081
1082
1083
1084
1085
1086
1087

Dataset	Zero-shot	Best Abstraction	Average Abstraction
UCI Dry Bean	0.00	0.65	0.51
Wikipedia Proteinogenic Acid	0.22	0.78	0.58
UCI Student Performance	0.25	0.45	0.28
UCI Website Phishing	0.25	0.25	0.22
UCI Teaching Assistant Evaluation	0.25	0.45	0.33
UCI Contraceptive Method Choice	0.30	0.60	0.43
UCI Vertebral Column	0.30	0.75	0.64
UCI Shill Bidding	0.30	1.00	0.95
Kaggle Job Change	0.30	0.85	0.83
UCI Caesarian Section	0.38	0.75	0.64
Wikipedia Coin Face Value	0.40	1.00	0.88
UCI Wine	0.40	0.95	0.85
UCI Tic-Tac-Toe Endgame	0.40	0.80	0.42
Kaggle Campus Placement	0.40	0.85	0.72
Wikipedia Driving Championship Points	0.40	1.00	0.74
UCI Mammographic Mass	0.45	0.90	0.82
UCI Banknote Authentication	0.45	1.00	0.78
Kaggle Engineering Placement	0.50	0.85	0.79
RAFT One Stop English	0.50	0.40	0.36
LegalBench Function of Decision Section	0.54	0.72	0.61
Kaggle Entrepreneur Competency	0.55	0.65	0.58
UCI Indian Liver Patient	0.55	0.80	0.68
LegalBench International Citizenship Questions	0.56	0.74	0.63
LegalBench Abercrombie	0.56	0.80	0.67
Wikipedia Color Luminance	0.60	1.00	1.00
RAFT Twitter Hate Speech	0.60	0.90	0.76
Wikipedia Award Nomination Result	0.64	1.00	0.76
UCI Car Evaluation	0.65	0.75	0.64
Kaggle Water Potability	0.65	0.50	0.38
Kaggle Travel Insurance	0.65	0.70	0.59
UCI Internet Firewall	0.70	1.00	0.97
RAFT ADE Corpus	0.70	1.00	0.89
UCI Somerville Happiness Survey	0.70	0.80	0.68
UCI Mushroom	0.75	1.00	0.95
UCI Occupancy Detection	0.80	1.00	0.92
Kaggle Stroke Prediction	0.85	0.90	0.90
LegalBench Corporate Lobbying	0.88	0.94	0.88
Average	0.50	0.80	0.68

1124
1125 Table 7: Evaluation of abstractions on diverse collection of 37 domains. We sampled 10 abstractions
1126 by prompting $\circ 4\text{-mini}$, and measure test set accuracy while prompting $\text{GPT-4}\circ\text{-mini}$ with each
1127 abstraction. We report both the average performance of the 10 abstractions and the best abstraction.
1128 **We find that the average and best abstractions outperform standard prompting by 18.0% and**
1129 **30.0% on average, respectively.**1130
1131
1132
1133

1134
1135

D.4 THE USE OF LARGE LANGUAGE MODELS (LLMs)

1136
1137
1138
1139

Large Language Models are used to assist with proofreading and minor wording improvements. All research ideas, experiments, and conclusions were conceived and validated by the authors. Additionally, tools such as Cursor were utilized as coding assistants during the development of the coding infrastructure for the project.

1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187