MentaLMs: Teaching Language Models to Reason Efficiently in the Language of Thought

Abstract

Large Reasoning Models (LRMs) achieve state-of-the-art performance in mathematics, code generation, and task planning. However, their reliance on long chains of verbose "thinking" tokens results in high latency, redundancy, and incoherent reasoning paths. Inspired by the Language of Thought Hypothesis —which posits that human reasoning operates over a symbolic, compositional mental language called *Mentalese*—we introduce a cognitively motivated framework that trains models to reason in a similar compact style, which we call MentaLMs. *Mentalese* encodes abstract reasoning as ultra-compressed, structured tokens, enabling models to solve complex problems with far fewer steps. To achieve both efficiency and accuracy, we propose Shorter Length Preference Optimization (SLPO), a reinforcement learning method that directly optimizes models to generate concise yet correct reasoning by rewarding shorter solutions that maintain high accuracy while flexibly allowing longer reasoning when complexity demands it. When applied to *Mentalese*-aligned models, SLPO achieves much larger compression rates by enabling compressed reasoning that preserves the benefits of detailed thinking without the computational overhead, allowing us to present the best-performing models at each compression level along the performance-efficiency Pareto frontier. Across mathematical benchmarks—including AIME 2024, Minerva, OlympiadBench, Math500, and AMC—our MentaLMs generates reasoning traces with 4–16× fewer tokens, achieves over 5× lower inference latency, and reduces training costs by 7–9× relative to the base DeepSeek R1 Distilled model, while maintaining 90-98% of baseline accuracy. MentaLMs also surpasses Claude and ChatGPT-4o by up to 5% in accuracy while maintaining 2× compression. Our findings demonstrate *Mentalese*-style compressed reasoning offers a breakthrough toward human-like cognitive efficiency, opening new possibilities for real-time, cost-effective reasoning without sacrificing accuracy.

1 Introduction

Recent advances such as OpenAI o1 OpenAI et al. (2024b) and DeepSeek R1 (DeepSeek-AI et al., 2025) have reshaped how we think about language model reasoning. By letting models "think before they answer," these systems dramatically improved credibility and performance—achievements that were once thought impossible for LLMs (Wu et al., 2024). Explicit reasoning has thus emerged as a central focus in LLM research (Xu et al., 2025). Recent work such as DeepScaleR: Surpassing O1-Preview with a 1.5B Model by Scaling RL demonstrates that even relatively small models (1.5B parameters) can outperform OpenAI's O1-Preview—which is widely assumed to be significantly larger, though its scale has not been publicly disclosed—by leveraging reasoning-focused reinforcement learning techniques such as RLVR, where models generate intermediate "thinking" tokens for self-verification (Luo et al., 2025b). This finding underscores that scaling in reasoning depth can, in some contexts, rival scaling in parameter size. The key challenge now lies in transforming this promise into robust, efficient, and trustworthy deployments, which we address in the next section. However, the promise of RLVR comes with significant trade-offs. Training is computationally expensive, with rollout generation leaving GPUs idle for long periods (Fu et al., 2025). Even relatively small models such as 1.5B parameters can take days to train under RL fine-tuning regimes (Zheng et al., 2025). Moreover, R1-style reasoning traces (shown in Figure 1a) are often verbose, redundant, and unnatural - a far cry from human cognition, which tends to rely on short and efficient thought steps (Sui et al., 2025). Building on the Language of Thought hypothesis (Fodor, 1975), which sug-

(a) Contrast between human and machine reasoning (response from Deepseek-R1). While humans arrive at intuitive and concise solutions, LLMs often produce verbose and redundant reasoning chains even for simple problems. We bridge this gap by developing methods that encourage models to reason more like humans—clear, efficient, and direct—while preserving accuracy. Grounded in the Language of Thought hypothesis, human reasoning compresses complex ideas into minimal symbolic steps, reflecting cognitive efficiency. Emulating this compact reasoning reduce redundancy in machine outputs, improving both interpretability and token efficiency.

(b) Performance–efficiency trade-offs of various model families across six mathematical reasoning benchmarks (including AIME2025). The dotted curve indicates the Pareto frontier, which illustrates the trade-off between higher compression rates and loss in accuracy. Our proposed method, combining Mentalese with SLPO, consistently lies on this frontier, identifying an optimal operating point that achieves a balance between accuracy and efficiency.

Figure 1: Contrast between human and machine reasoning approaches. While humans use intuitive and concise steps, LLMs often produce verbose and redundant reasoning for even simple problems.

gests that human cognition unfolds through short compositional thought units rather than verbose natural language traces, we propose a training framework that restructures the reasoning style of current reasoning-oriented LLMs in a symbolic internal language that we call *Mentalese*.

In our framework, models are first aligned with this reasoning process through supervised fine-tuning on reasoning traces in *Mentalese*, namely, concise compositional sequences that capture only the essential steps required for problem solving. However, aligning models to *Mentalese* by supervised fine-tuning alone leads to a substantial drop in performance relative to the base model. To overcome this, we introduce *Shorter Length Preference Optimization* (SLPO), a reinforcement learning objective with a verifiable reward that balances brevity and correctness. Unlike token-penalization methods (e.g., L1-style objectives (Aggarwal & Welleck, 2025)) that impose arbitrary length budgets—often forcing models to under-reason on difficult problems and over-reason on easy ones—SLPO instead rewards relative efficiency: among correct rollouts, shorter solutions receive a bonus. This naturally biases the model toward concise reasoning when tasks are simple, while still allowing it to allocate more steps when necessary in the same reasoning structure. By rewarding concise but correct solutions, SLPO avoids verbosity while recovering most of the performance lost during supervised fine-tuning on Mentalese, thereby yielding efficient reasoning that scales at inference time.

To highlight both domain-specific performance and generalization, we evaluated our suite of trained models (*MentaLMs 1.5B*) on mathematical reasoning in the domain and out-of-the-domain tasks such as GPQA, LSAT, and MMLU. Our *Magentica-SLPO 1.5B* surpasses GPT-4o (OpenAI et al., 2024a), Claude 3.5 Sonnet (Anthropic, 2024), and Llama 3.3 70B (Grattafiori et al., 2024) by an average of 6 pp[1] in mathematical reasoning and outperforms DeepSeek-R1 1.5B DeepSeek-AI et al. (2025) by 7 pp with a 7× reduction in reasoning length (Figure 1b). *MentaSeek-GRPO 1.5B* achieves 14× compression relative to DeepSeek-R1 1.5B. Beyond in-domain gains, our *MentaLMs* also generalize well: on out-of-domain tasks they improve over the base model by 1 pp while achieving 15× compression Table 2. In addition to token efficiency, our experiments show that training with Mentalese stabilizes optimization and reduces training time by 7–9× compared to directly training the base model with RLVR, leading to substantial savings in training cost. Beyond benchmarks, we hypothesize that these ideas are especially relevant for agentic LLM systems, where reasoning models are rarely deployed due to latency and cost: verbose generations can overwhelm communication channels Kim et al. (2025). A compressed reasoning style, reinforced through SLPO, has the poten-

---

[1]pp = percentage points, denoting absolute differences between percentages (e.g., 22% vs. 16% = 6 pp).

tial to dramatically reduce this overhead—making reasoning-capable agents not only more accurate but also more practical to deploy in real-world settings. Our main contributions are as follows.

- **Reasoning compression framework.** We propose a novel and efficient reasoning compression framework via *Mentalese* alignment for restructuring the reasoning style of current LLM, producing compact yet faithful symbolic reasoning.

- **Reward function.** We propose *Shorter Length Preference Optimization (SLPO)*, an adaptive objective that dynamically balances correctness with brevity, eliminating the need for rigid length penalties.

- **Dataset.** We release *MentaleseR-40k*, a dataset of ultra-compressed reasoning traces for 40k math problems, generated under symbolic constraints inspired by the Language of Thought Hypothesis (LOTH), to support future developments and foster research on efficient reasoning.

- **Experiments and best practices.** We conduct extensive evaluations and identify best practices to apply GRPO and SLPO, showing how they achieve different levels of compression and where each method is most effective.

2 RELATED WORKS

Efficient Reasoning in Large Language Models: Since Wei et al. (2022b) demonstrated the effectiveness of chain-of-thought (CoT) prompting, subsequent work has focused on scaling test-time computation to improve performance in mathematical problem-solving, code generation, and complex reasoning tasks. Strategies include parallel sampling of multiple reasoning paths (Wang et al., 2022; Yue et al., 2024; Chen et al., 2023), tree search and planning (Yao et al., 2023; Besta et al., 2024), and iterative refinement methods (Madaan et al., 2023). Recent reasoning-specialized models, such as OpenAI's *o1* (OpenAI et al., 2024b), DeepSeek-R1 (DeepSeek-AI et al., 2025), and Qwen-QwQ (Yang et al., 2025), internalize the ability to generate extended reasoning traces.However, these methods often suffer from the *overthinking phenomenon* (Sui et al., 2025; Chen et al., 2025), where models generate excessively long reasoning traces. While increased length can improve accuracy up to a point (Wu et al., 2025), it also introduces redundancy, higher inference latency, and even accuracy degradation due to compounding errors Hassid et al. (2025); Lee et al. (2025). This trade-off has motivated work on more efficient reasoning. RL-based post-training methods have been widely explored to control reasoning length. L1 (Aggarwal & Welleck, 2025) enforces user-specified budgets, DAST (Shen et al., 2025) adapts budgets based on problem difficulty, while O1-Pruner (Luo et al., 2025a) uses reference-based pruning. Other approaches, such as Kimi 1.5 (Team et al., 2025) and Training Efficient (Arora & Zanette, 2025), use sampled rollouts to reward shorter or average lengths. ShorterBetter (Yi et al., 2025) further introduces the idea of rewarding the shortest correct response, highlighting the existence of problem-dependent optimal reasoning lengths. Our work complements these by introducing SLPO, which adaptively prefers concise correct reasoning without penalizing necessary longer derivations, enabling over $10\times$ compression with minimal loss in accuracy.

Chain-of-Thought and Alternative Reasoning Formats: CoT reasoning has become a dominant paradigm for enhancing reasoning in LLMs, either via prompting Wei et al. (2022a); Fu et al. (2022); Zhou et al. (2022) or through post-training with supervised finetuning Yue et al. (2023); Yu et al. (2023) and reinforcement learning Luong et al. (2024); Shao et al. (2024b); Zhou et al. (2025). Theoretical analyses link CoT to increased expressivity and effective depth in transformers Feng et al. (2023); Merrill et al. (2023); Li et al. (2024). However, natural-language CoT traces are verbose, redundant, and not always faithful to the model's underlying reasoning process Turpin et al. (2023); Chuang et al. (2024). Recent research has explored alternatives. Structured or symbolic CoT formats aim to compress reasoning into more compact representations, such as symbolic operators, patterns, or abstract primitives Madaan & Yazdanbakhsh (2022); Paul et al. (2024). Other works examine latent reasoning, where intermediate computation is implicit in hidden representations rather than externalized tokens Yang et al. (2024); Biran et al. (2024); Shalev et al. (2024). Techniques such as back-patching Biran et al. (2024), filler tokens Pfau et al. (2024), or knowledge distillation into latent reasoning Deng et al. (2023; 2024) push beyond explicit CoT. Our proposed *Mentalese Chain-of-Thought* builds on this line of work by introducing a symbolic, cognitively motivated reasoning language inspired by the Language of Thought Hypothesis. By replacing verbose

natural language with structured symbolic primitives, Mentalese CoT achieves order-of-magnitude compression while retaining faithfulness and interpretability. Combined with SLPO, this framework demonstrates that both representation and optimization are critical for efficient and reliable reasoning.

3 METHODOLOGY

In this section, we present our methodology, which integrates symbolic reasoning alignment with reinforcement learning for concise yet accurate performance. We introduce *Mentalese*, a compact symbolic reasoning format, and *Group Relative Policy Optimization (GRPO)*, a group-based extension of PPO for reasoning optimization. Our main contribution, *Shorter Length Preference Optimization (SLPO)*, refines GRPO by rewarding brevity without penalizing necessary longer reasoning. Finally, we propose *RLVR*, a two-stage pipeline that first aligns models to Mentalese via supervised finetuning, then applies GRPO or SLPO with verifier feedback. Together, these components yield 10–20$\times$ compression in reasoning traces while maintaining accuracy and efficiency across benchmarks.

3.1 MENTALESE: MENTAL LANGUAGE OF THOUGHT

We first introduce *Mentalese*, a cognitively motivated reasoning format inspired by the Language of Thought Hypothesis (Fodor, 1975; Rescorla, 2024). According to this hypothesis, human cognition operates not directly in natural language, but in an internal representational system characterized by compact, symbolic structures. Translating this perspective to Large Reasoning Models (LRMs), we hypothesize that verbose natural language explanations commonly used in chain-of-thought prompting especially the Deepseek R1 reasoning style, are not essential for reasoning, and that more efficient symbolic primitives can better capture the core logical operations underlying problem-solving.

Formal definition. Let $\mathcal{O}$ be a finite set of operators (e.g., SET, EQ, CASE, SOLVE, CALC, DIFF, ANS) and let $\mathcal{E}$ be the set of symbolic expressions over variables, numbers, and function symbols (e.g., $+, -, \times, \div$, abs). A *Mentalese step* is a pair $s_t = (o_t, c_t)$ with $o_t \in \mathcal{O}$ and *expression* $c_t \in \mathcal{E}$ rendered as the string OPERATION:expression;. A *Mentalese trace* for a question q is a finite sequence $M = (s_1; \ldots; s_T)$ that is well-typed and executable under the step semantics below and that culminates in exactly one terminal ANS:e; step. The boxed final answer is $e^\star$, where $e^\star$ is the value denoted by e. We denote the set of valid traces by $\mathcal{M}$.

Unlike traditional CoT, which uses free-form text, *Mentalese* encodes reasoning in canonical steps of the form OPERATION:expression;, joined by semicolons to form minimal yet complete traces. This yields three advantages: (i) Compression — eliminating redundant tokens for up to 10$\times$ shorter reasoning; (ii) Faithfulness — each step is necessary and sufficient; (iii) Cognitive alignment — resembling structured mental representations rather than verbose text.

To build MENTALESE-40K, we adapted the DEEPSCALER-PREVIEW-DATASET (Luo et al., 2025b), covering 40k+ math problems from AIME (1983–2023), Omni-Math, and STILL. We used GPT-4.1 with a structured prompting framework (Figure 2a)—including a formal definition, syntactic rules, and examples—to generate *Mentalese* traces. After light curation (removing 65 malformed cases), the resulting dataset was used for supervised fine-tuning. For RLVR, we instead relied on the original QA pairs, letting the verifier assess correctness while optimizing for concise reasoning. Refer Appendix A.3 for some of the samples from MENTALESER-40K

3.2 GROUP RELATIVE POLICY OPTIMIZATION (GRPO)

While PPO (Schulman et al., 2017) provides a strong baseline for policy optimization, it operates at the *single-sample* level: each rollout is evaluated independently using a value function to estimate its advantage. However, in reasoning tasks where multiple candidate solutions can be generated for the same question, evaluating rollouts in isolation discards useful information about the *relative quality* of responses within a group. For example, if a model generates five candidate solutions, some correct and some incorrect, we are less interested in their absolute values than in how each compares relative to others in the same set. This motivates **Group Relative Policy Optimization (GRPO)** (Shao et al., 2024a), which eliminates the explicit value function and instead estimates the advantage by normalizing rewards *within groups of samples* drawn for the same prompt. Concretely,

MENTALESE FRAMEWORK

DEFINITION

A symbolic, logic-based chain of thought using compact primitives and logical steps only. Each step should be separated by a semicolon (;) with no natural language. All core logic steps must be preserved.

EXAMPLE

SET:w;
EQ:abs(180-5.5*w)=110;
CASE1:180-5.5*w=110;
SOLVE1:w=70/5.5;
CASE2:180-5.5*w=-110;
SOLVE2:w=290/5.5;
CALC1:w1=70/5.5=12+8/11;
CALC2:w2=290/5.5=52+8/11;
DIFF:t=w2-w1=40;
ANS:40

RULES

1. Use ONLY symbolic notation in the form: OPERATION: expression

2. No natural language

3. No redundancy

4. Every step must be necessary and sufficient for solving the problem

(a) Illustration of symbolic, logic-based chain of thought (mentalese). This figure shows the definition (top), an example of symbolic reasoning steps (left) with rules governing the reasoning style (right).

(b) Comparison of reasoning traces on AIME 2024: Agentic models use 7800 tokens, Magentica 150 tokens, and Magentica-SLPO 300 tokens with similar accuracy.

Figure 2: Illustration of symbolic, logic-based chain of thought (mentalese).

for a question–answer pair (q, a), we sample a group of G responses $\{o_i\}_{i=1}^{G}$ from the current policy. The reward of each response r_i is converted into a *group-relative advantage* via normalization:

$$\hat{A}_i = \frac{r_i - \mathrm{mean}(\{r_j\}_{j=1}^{G})}{\mathrm{std}(\{r_j\}_{j=1}^{G}) + \epsilon}.$$

This design ensures that advantages highlight which responses are better or worse relative to the group, rather than depending on an absolute critic model. GRPO then optimizes a clipped surrogate objective similar to PPO but with a directly imposed KL penalty:

$$\mathcal{J}_{\mathrm{GRPO}}(\theta) = \mathbb{E}\left[\frac{1}{G}\sum_{i=1}^{G}\frac{1}{|o_i|}\sum_{t=1}^{|o_i|}\Big(\min\big(r_{i,t}(\theta)\hat{A}_{i,t}, \mathrm{clip}(r_{i,t}(\theta), 1-\varepsilon, 1+\varepsilon)\hat{A}_{i,t}\big) - \beta D_{\mathrm{KL}}(\pi_\theta\|\pi_{\mathrm{ref}})\Big)\right],$$

where $r_{i,t}(\theta) = \frac{\pi_\theta(o_{i,t}|q,o_{i,<t})}{\pi_{\mathrm{old}}(o_{i,t}|q,o_{i,<t})}$ is the token-level importance ratio.

3.3 Shorter Length Preference Optimisation (SLPO)

There has been a growing interest in adaptive reasoning methods that redefine the GRPO formulation by incorporating explicit *thinking budgets*. For instance, prior works such as LCPO constrain reasoning lengths by enforcing unnatural fixed or maximum token budgets. More recently, group-relative formulations have also been proposed that define rewards based on the relative lengths of responses within a group. However, these methods tend to be overly rigid: they *over-penalize* the longest solutions even when they converge to right answer, and in cases where no correct solution exists, their length-normalization can still distort the reward landscape. This strictness can suppress necessary reasoning and lead to degenerate behavior.

To overcome these issues, we introduce **Shorter Length Preference Optimization (SLPO)**, a reinforcement learning strategy that balances conciseness with correctness *softly*. Crucially, SLPO never penalizes a correct but necessarily long reasoning when it is the only valid option, and it does not distort rewards in cases with no correct solution. Instead, it adaptively rewards shorter correct traces when multiple valid derivations exist, while preserving correctness as the primary training signal.

Different problems naturally require different amounts of reasoning. For example, a simple arithmetic task such as $2 + 2$ requires no intermediate steps, while Olympiad-style geometry problems demand much longer derivations. A reward function that ignores this variability either pushes the model toward artificially verbose chains (reward hacking under fixed budgets) or toward overly terse and often incorrect responses (under strict length penalties). SLPO resolves this by defining *preferences relative to the observed range of correct reasoning lengths* for each problem instance.

Formally, for a given rollout group $G(x_i) = \{y_1, y_2, \ldots, y_n\}$ corresponding to prompt x_i, let $\mathcal{C}(x_i) = \{y_j \in G(x_i) : R_{\mathrm{correctness}}(y_j) = 1\}$ denote the set of correct responses. We define:

$$L_{\min} = \min_{y\in\mathcal{C}(x_i)} \ell(y), \qquad L_{\max} = \max_{y\in\mathcal{C}(x_i)} \ell(y), \tag{1}$$

where $\ell(y)$ is the token length of response y. The total reward for candidate y_{curr} is then:

$$
R_{\text{SLPO}}(y_{\text{curr}}) = \begin{cases} 1, & \text{if } |\mathcal{C}(x_i)| = 1 \text{ or } \big(|\mathcal{C}(x_i)| > 1 \ \& \ L_{\min} = L_{\max}\big), \\[2ex] R_{\text{correctness}} + \alpha \cdot \dfrac{L_{\max} - L_{\text{curr}}}{L_{\max} - L_{\min}}, & \text{if } |\mathcal{C}(x_i)| > 1 \text{ and } L_{\min} \neq L_{\max}, \\[2ex] 0, & \text{if } |\mathcal{C}(x_i)| = 0, \end{cases}
\tag{2}
$$

where L_{curr} is the length of the current response, $R_{\text{correctness}} \in \{0,1\}$ is determined by a verifier, and α controls the trade-off between accuracy and conciseness. Larger values of α emphasize brevity, whereas smaller values prioritize correctness regardless of length. In all experiments, we set $\alpha = 0.1$, which provided a stable balance across benchmarks.

By construction, SLPO avoids the failure modes of previous group-relative and L1-based formulations: it does not over-penalize long but uniquely correct solutions, and it does not distort reward landscapes when no valid solutions exist. Instead, it consistently encourages models to discover the *shortest correct reasoning trace* whenever possible. This makes SLPO especially well-suited for mathematical reasoning, where optimal reasoning lengths vary significantly across problems.

3.4 MENTALESE ALIGNMENT THROUGH SFT FOLLOWED BY RLVR

We now present our complete training pipeline, which consists of two stages: supervised alignment on *Mentalese* traces, followed by reinforcement learning with verifiable rewards (RLVR).

Stage 1: Supervised Finetuning on Mentalese. Let $\mathcal{D} = \{(q_i, a_i, m_i)\}_{i=1}^{M}$ be our dataset with question q_i, ground-truth final answer a_i, and Mentalese reasoning trace m_i. Each training prompt is structured as

$\tau(q_i) = q_i +$ `'Let's think step-by-step and answer within \boxed{}.'`

with target output

$$y_i^{\star} = \texttt{<think>} \ m_i \ \texttt{</think>} \ \texttt{\textbackslash boxed}\{a_i\}.$$

Starting from a pretrained base model π_0, we obtain a Mentalese-aligned model π_{SFT} via supervised finetuning:

$$
\pi_{\text{SFT}} = \arg\min_{\theta} \ -\frac{1}{M} \sum_{i=1}^{M} \log \pi_{\theta}(y_i^{\star} \mid \tau(q_i)).
$$

Stage 2: Reinforcement Learning with Verifier Rewards (RLVR). The SFT model π_{SFT} is further refined using verifier-based reinforcement learning. For each question q_i, the policy generates N candidates $G(q_i) = \{y^{(1)}, \dots, y^{(N)}\}$; a verifier checks the boxed answer $\hat{a}$ against a_i and assigns a correctness reward $R_{\text{acc}} \in \{0,1\}$. The clipped surrogate objective with KL regularization is:

$$
\mathcal{J}_{\text{RLVR}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \ y \sim \pi_{\theta}(\cdot|q)} \Big[\min\big(r_{\theta}(y)\hat{A}(y), \ \text{clip}(r_{\theta}(y), 1-\varepsilon, 1+\varepsilon)\hat{A}(y)\big) - \beta D_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{SFT}}) \Big],
$$

where $r_{\theta}(y) = \frac{\pi_{\theta}(y|q)}{\pi_{\text{old}}(y|q)}$ and $\hat{A}(y)$ is computed from either the **GRPO** or **SLPO** formulation (see previous subsections).

Depending on the chosen reward function, RLVR yields a policy π_{GRPO} or π_{SLPO}:

$$
\pi_{\text{SFT}} \xrightarrow{\text{RLVR (GRPO)}} \pi_{\text{GRPO}}, \qquad \pi_{\text{SFT}} \xrightarrow{\text{RLVR (SLPO)}} \pi_{\text{SLPO}}.
$$

SFT alignment anchors the model to a compact single-chain reasoning format, ensuring that outputs conform to the Mentalese structure shown in Figure 2b. However, this alignment often comes at the cost of reduced accuracy, since the base model initially performs well with long and verbose reasoning chains. RLVR provides the complementary step: by instantiating it with either GRPO or SLPO, the model learns to recover accuracy while retaining the compact reasoning format. RLVR enables the model to refine and extend its reasoning inside the learned structure, adding useful steps when necessary but avoiding unnecessary verbosity. This combination not only restores the lost performance but also yields consistent improvements in overall reasoning efficiency and accuracy across benchmarks.

4 EXPERIMENT DESIGN

Our experimental study is organized around five research questions. First, we ask how effective the proposed *Mentalese* representation is in compressing reasoning traces while preserving task performance. Second, we investigate which reinforcement learning algorithm (GRPO or SLPO) best recovers the performance gap introduced by compression. Third, we evaluate the standalone effectiveness of SLPO in balancing conciseness and correctness. Fourth, we explore best practices by identifying when each algorithm is most suitable, particularly across different regions of compression. Finally, we analyze the efficiency and stability of our RLVR methods during training.

Models and Baselines. To evaluate our method, we consider both base models and competitive baselines. Our primary comparison point is DeepSeek-R1-Distill-Qwen-1.5B DeepSeek-AI et al. (2025), a distilled variant of Qwen-2.5-1.5B-Instruct fine-tuned on reasoning traces from DeepSeek's R1 model, which we denote as DEEPSEEK-R1-1.5B. We also include DeepScaleR-1.5B-Preview Luo et al. (2025b), the original release without length-control modifications, referred to as AGENTICA-24K. For completeness, we tested the base Qwen-2.5-Math-1.5B model Qwen et al. (2025); however, it collapsed under RLVR fine-tuning due to NaN gradient norms, so it is excluded from final evaluations. In addition, we benchmark against L1-Max Aggarwal & Welleck (2025), a strong baseline derived from AGENTICA-24K using Length-Controlled Policy Optimization (LCPO), which achieves more than 10× compression. While this approach effectively reduces verbosity by enforcing a fixed token budget, it lacks adaptability to varying problem difficulties. Beyond these 1.5B-scale models, we also report results from frontier-scale systems for context. Specifically, we include GPT-4o OpenAI et al. (2024a), Claude 3.5 Sonnet Anthropic (2024), and LLaMA-3 70B-Instruct Grattafiori et al. (2024) as strong reference points, situating our results relative to state-of-the-art closed-source and large-scale open-source LLMs.

Evaluation and Metrics. We evaluate our models on five in-domain mathematical reasoning datasets: **AIME 2024** (Mathematical Association of America, 2024), **AIME 2025**(Mathematical Association of America, 2025), **MATH-500** (Hendrycks et al., 2021b), **AMC** (American Mathematics Competitions), **Minerva-Math** Lewkowycz et al. (2022), and **OlympiadBench** He et al. (2024). Additionally, we test on three out-of-domain benchmarks: **GPQA** (Rein et al., 2023), **LSAT** (Zhong et al., 2023), and **MMLU** (Hendrycks et al., 2021a), in order to assess generalization beyond mathematical reasoning. We report results using three primary metrics. Pass@1 measures the fraction of problems correctly solved under single-sample decoding, i.e., the proportion of test questions for which the model produces a correct solution on its first attempt. Token Length denotes the average number of tokens generated per response on a given benchmark, computed by averaging output lengths across all test questions. Compression Rate (CR) quantifies the degree of response shortening relative to DeepSeek-R1-1.5B, with higher values indicating greater compression (e.g., a CR of 10 means the model's responses are ten times shorter on average). Full formal definitions are provided in Appendix A.2.

Implementation Details. For supervised fine-tuning on the MENTALESER-40K dataset, we used LLAMA-FACTORY Zheng et al. (2024), an open-source library for instruction tuning and post-training. For reinforcement learning, we adopted VERL Sheng et al. (2025), an open-source RL training library. We fine-tuned our 1.5B base models with a batch size of $N = 128$ and a rollout group size of $n = 16$. Training was conducted for 1500 steps with a fixed learning rate of 1×10^{-6}. For reinforcement learning experiments, we used 32 H100 GPUs, while supervised fine-tuning was performed on 8 H100 GPUs. Inference was accelerated using the vLLM Kwon et al. (2023) engine, which enables efficient large-scale generation. For length constraints, we set different maximum generation lengths depending on the training setup: 8K tokens for direct SLPO on base models, 2K tokens for SLPO on MENTALESER-40K fine-tuned models, and 1K tokens for GRPO on MENTALESER fine-tuned models. During evaluation on benchmarks, we fixed a maximum generation budget of 8K tokens. More details are provided in Appendix A.2.2.

5 DISCUSSION

Loss During SFT, Recovery Through RLVR. As shown in Table 1, we observed a substantial performance drop after the SFT stage on MENTALESE, with average accuracy decreasing by 35.5 p.p. relative to the base model. This decline stems from the fact that SFT encourages the model to restructure reasoning into a symbolic MENTALESE format, typically resulting in a single linear reasoning path. In contrast, DeepSeek R1-style reasoning traces often include "forking tokens" such as *wait*,

Table 1: Model performance across benchmarks. Each block shows Pass@1, average tokens, and compression rate (CR).

Model	AIME2024			AIME2025			MATH			AMC			MINERVA			OLYMPIAD		
	P@1	Tokens	CR	P@1	Tokens	CR	P@1	Tokens	CR	P@1	Tokens	CR	P@1	Tokens	CR	P@1	Tokens	CR
Baselines																		
L1 Max	0.04	390	19.18	0.03	424	17.99	0.61	220	11.57	0.38	310	19.79	0.17	190	21.66	0.24	290	23.65
GPT-4o	0.10	923	8.11	0.06	998	7.65	0.78	607	4.19	0.52	1003	6.12	0.49	837	4.92	0.42	956	7.17
Claude 3.5 Sonnet	0.16	819	9.14	0.03	764	9.99	0.75	458	5.56	0.47	775	7.92	0.46	694	5.93	0.36	701	9.78
Llama3.3 70B	0.25	1114	6.72	0.03	1067	7.15	0.74	516	4.93	0.47	845	7.26	0.36	736	5.59	0.36	780	8.79
Base Models and our Finetuned Variants																		
DeepSeek R1 1.5b	0.19	7481	1.00	**0.2**	7631	1.00	0.77	2545	1.00	0.51	6136	1.00	0.23	4115	1.00	0.35	6858	1.00
+ SLPO (Direct)	**0.26**	6510	1.15	**0.2**	5910	1.29	**0.78**	1017	2.50	**0.60**	2797	2.19	**0.24**	1393	2.96	**0.41**	3031	2.26
MentaSeek	0.004	184	40.65	0.01	136	56.10	0.75	156	16.31	0.51	168	36.52	0.12	147	27.99	0.12	157	43.68
MentaSeek-GRPO	0.2	571	13.10	0.13	524	14.56	**0.75**	301	8.45	**0.51**	459	13.37	**0.26**	317	12.98	**0.36**	439	15.68
MentaSeek-SLPO	**0.23**	976	7.66	**0.16**	785	9.72	0.61	68	37.43	0.44	439	13.98	0.18	86	47.85	0.32	453	26.68
Agentica-24k	**0.36**	7049	1.06	**0.27**	7270	1.05	**0.86**	1841	1.38	0.67	4062	1.51	**0.29**	3466	1.19	**0.47**	4377	1.57
+ SLPO (Direct)	0.33	4379	1.71	0.23	3889	1.95	0.83	504	5.05	**0.68**	1696	3.59	0.27	640	6.43	0.46	1668	4.11
Magentica	0.01	126	59.37	0.01	126	60.56	0.40	97	26.23	0.14	128	47.94	0.13	119	34.58	0.12	131	52.34
Magentica-GRPO	0.19	649	11.52	0.12	541	14.11	0.77	257	9.91	0.53	493	12.44	0.26	266	15.47	0.37	453	15.13
Magentica-SLPO	0.24	1307	5.72	0.21	971	7.86	0.81	424	5.99	0.6	966	6.35	**0.29**	542	7.59	0.43	889	7.71

Table 2: Model performance on out-of-domain (OOD) benchmarks GPQA, LSAT, MMLU-PRO and Average. Each block shows Pass@1, average tokens, and compression rate (CR%).

Model / Training Regime	GPQA			LSAT			MMLU-1000			Average		
	P@1	Tokens	CR%	P@1	Tokens	CR%	P@1	Tokens	CR%	P@1	Tokens	CR%
Baselines												
L1 Max	0.29	518	14.58	0.22	554	14.79	0.45	288	2.46	0.32	453	10.61
Base Models and our Finetuned Variants												
DeepSeek R1 1.5b	0.31	7551	1.00	0.24	8192	1.00	0.45	710	1.00	0.33	5484	1.00
+ SLPO (Direct)	**0.33**	3434	2.20	0.25	5608	1.46	0.45	490	1.45	0.34	3177	1.70
MentaSeek	0.25	142	53.21	0.2	316	25.92	0.35	67	10.60	0.27	175	29.91
MentaSeek-GRPO	0.31	366	20.64	0.23	587	13.95	**0.46**	166	4.28	0.33	373	12.96
MentaSeek-SLPO	0.32	393	19.22	0.23	851	9.63	0.43	49	14.49	0.33	431	14.45
Agentica-24k	0.36	5010	1.51	0.29	7614	1.08	0.48	934	0.76	0.38	4519	1.12
+ SLPO (Direct)	**0.33**	2101	3.59	**0.28**	4077	2.01	**0.48**	390	1.82	**0.36**	2189	2.47
Magentica	0.26	135	55.93	0.21	201	40.76	0.35	62	11.45	0.27	133	36.05
Magentica-GRPO	0.30	362	20.85	0.23	564	14.52	0.47	157	4.52	0.33	361	13.30
Magentica-SLPO	**0.33**	735	10.28	0.25	1200	6.83	**0.48**	282	2.52	0.35	739	6.54

but, or *so*, which allow the model to self-verify and revise its reasoning mid-generation—boosting accuracy through exploratory pathways. The strict structure imposed by SFT sacrifices these benefits, limiting the model's flexibility and test-time scaling. However, applying RLVR largely reverses this effect: models regain most of the lost accuracy while maintaining significantly shorter reasoning traces—typically just a few hundred tokens longer than their SFT outputs. This highlights the complementary roles of the two stages: SFT enforces symbolic conciseness, while RLVR restores accuracy by reintroducing adaptive reasoning behaviors within that compact framework.

Training Time Efficiency. Large-scale reinforcement learning typically demands substantially more computational resources than supervised fine-tuning. In our experiments, we found that applying RLVR directly on the base models required 5–6 days of training on our dataset of 40k samples for 1500 RL steps. This inefficiency arises primarily from the generation of long reasoning chains (often exceeding 8k tokens), which introduces high latency in the vLLM inference engine and becomes the main bottleneck of RLVR training. The cost further increases with larger rollout group sizes. As shown in Figure 3, introducing an intermediate supervised fine-tuning stage on the MENTALESER dataset significantly reduced training cost by 7–10×, while achieving performance close to the base model but with 10× shorter reasoning traces. This demonstrates that aligning models to a more compact reasoning language before RL training not only improves efficiency but also provides a scalable mechanism for reinforcement learning in reasoning tasks.

Training Collapse Under Direct SLPO. As shown in Figure 3, applying SLPO directly to AGENTICA-24K resulted in a sudden collapse after approximately 300 training steps. Initially, the response length decreased by nearly half while accuracy improved marginally. However, beyond this point, the average response length rapidly expanded to the maximum generation limit (8k tokens), and the gradient norm curve exhibited NaN values. This instability ultimately caused

Figure 3: This figure compares direct SLPO on the base model with Intermediate SFT followed by RLHF methods (SLPO/GRPO) on the MENTALESER-40K dataset across five metrics. The *Mentalese* alignment yields greater training stability and efficiency.: (1) Response Length reveals direct SLPO collapses due to gradient instability, while *MentaLMs* stay stable; (2) Clip Ratio indicates more controlled updates in Mentalese methods, driven by reduced response truncation.; (3) Entropy Loss reflects better exploration–exploitation balance; (4) Training Time per RL Step shows higher computational efficiency; (5) Test Performance on AIME 2024 (20% Pass@1) confirms *MentaLMs* outperform direct SLPO on the base model. Shaded regions denote min–max ranges across runs. These results highlight the importance of structured intermediate representations (*Mentalese*) for stable, efficient RL in large language models.

the training to collapse, highlighting the difficulty of applying SLPO on raw base models without intermediate alignment. In contrast, introducing an intermediate SFT stage on the MENTALESER dataset maintained stability throughout the entire training process, underscoring the reliability of our proposed two-stage approach.

Reversion to Verbose Reasoning Under Large Generation Budgets During RLVR. We observed that when the maximum generation length was set to 4k or 8k tokens, the model tended to drift away from the compact reasoning style learned during the SFT stage and revert to its original base behavior of producing verbose chains. In some cases, this even led to model collapse Figure 3. A likely explanation is that longer reasoning traces, although verbose, occasionally produced correct answers and were therefore rewarded, inadvertently steering the model away from the MENTALESE format. To mitigate this effect, we restricted the maximum generation length to 1k tokens for GRPO-based training and 2k tokens for SLPO-based training. These limits preserved the symbolic reasoning behavior acquired during SFT while still allowing sufficient space for problem-solving.

6 CONCLUSION

We introduced a cognitively inspired framework for efficient reasoning that combines *Mentalese*, a compact symbolic reasoning format, with *Shorter Length Preference Optimization* (SLPO), a reinforcement learning strategy that adaptively balances conciseness and correctness. Our model achieves over $10\times$ compression in reasoning traces while maintaining accuracy close to that of verbose large reasoning models, reducing both training and inference costs significantly. By aligning models toward concise and structured reasoning, we provide a pathway for deploying large reasoning capabilities within real-time and resource-constrained environments. Our results suggest that reasoning does not inherently require verbosity, and that carefully designed representations and optimization objectives can yield models that reason more like humans—symbolically, compositionally, and efficiently. This is particularly valuable for agentic systems, where efficient and reliable decision-making is critical, and inference overhead can quickly become a bottleneck.

7 ETHICS STATEMENT

We read the ICLR Code of Ethics before the submission. Our paper introduces cognitively inspired compressed reasoning to improve efficiency in large reasoning models (LRMs), reducing computational costs and latency while promoting broader accessibility. At the same time, compressed reasoning may reduce interpretability, and efficiency gains could encourage premature deployment

in high-stakes domains. Our work also inherits common risks of large models, including bias and hallucination, and should therefore be applied responsibly with human oversight.

8 REPRODUCIBILITY

We specify the experimental setting in Appendix A.2.2 and attached the training and evaluation codebase along with data as supplementary materials.

REFERENCES

Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with reinforcement learning. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=4jdIxXBNve.

American Mathematics Competitions . Maa invitational competitions. https://maa.org/student-programs/amc//. Accessed: 2025-09-24.

Anthropic. Claude 3.5 Sonnet. https://www.anthropic.com/news/claude-3-5-sonnet, 2024. Accessed: YYYY-MM-DD.

Daman Arora and Andrea Zanette. Training language models to reason efficiently, 2025. URL https://arxiv.org/abs/2502.04463.

Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefler. Graph of thoughts: Solving elaborate problems with large language models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17682–17690, March 2024. ISSN 2159-5399. doi: 10.1609/aaai.v38i16.29720. URL http://dx.doi.org/10.1609/aaai.v38i16.29720.

Eden Biran, Daniela Gottesman, Sohee Yang, Mor Geva, and Amir Globerson. Hopping too late: Exploring the limitations of large language models on multi-hop queries, 2024. URL https://aclanthology.org/2024.emnlp-main.781/.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks, 2023. URL https://arxiv.org/abs/2211.12588.

Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. Do not think that much for 2+3=? on the overthinking of o1-like llms, 2025. URL https://arxiv.org/abs/2412.21187.

Yifan Chuang, Gaoang Wang, Chen Chang, Ruoyu Tang, Shichuan Zhong, Fei Yang, Min Du, Xin Cai, and Xuanjing Hu. Faithlm: Towards faithful explanations for large language models, 2024. URL https://arxiv.org/abs/2402.04678.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng

Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xi-aosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, 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.

Yuntian Deng, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, and Stuart Shieber. Implicit chain of thought reasoning via knowledge distillation, 2023. URL https://arxiv.org/abs/2311.01460.

Yuntian Deng, Yejin Choi, and Stuart Shieber. From explicit cot to implicit cot: Learning to internalize cot step by step, 2024. URL https://arxiv.org/abs/2405.14838.

Xiaolan Feng, Zhilin Yu, Hongkai Wang, et al. Chain-of-thought reasoning in transformers: A theoretical perspective, 2023. http://arxiv.org/abs/2503.08679.

Jerry A Fodor. *The language of thought*, volume 5. Harvard university press, 1975.

Wei Fu, Jiaxuan Gao, Xujie Shen, Chen Zhu, Zhiyu Mei, Chuyi He, Shusheng Xu, Guo Wei, Jun Mei, Jiashu Wang, Tongkai Yang, Binhang Yuan, and Yi Wu. Areal: A large-scale asynchronous reinforcement learning system for language reasoning, 2025. URL https://arxiv.org/abs/2505.24298.

Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning, 2022. URL https://arxiv.org/abs/2210.00720.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Michael Hassid, Gabriel Synnaeve, Yossi Adi, and Roy Schwartz. Don't overthink it: Preferring shorter thinking chains for improved llm reasoning, 2025. URL https://arxiv.org/abs/2505.17813.

Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems, 2024. URL https://arxiv.org/abs/2402.14008.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding, 2021a. URL https://arxiv.org/abs/2009.03300.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021b. URL https://arxiv.org/abs/2103.03874.

Jiin Kim, Byeongjun Shin, Jinha Chung, and Minsoo Rhu. The cost of dynamic reasoning: Demystifying ai agents and test-time scaling from an ai infrastructure perspective, 2025. URL https://arxiv.org/abs/2506.04301.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention, 2023. URL `https://arxiv.org/abs/2309.06180`.

Ayeong Lee, Ethan Che, and Tianyi Peng. How well do llms compress their own chain-of-thought? a token complexity approach, 2025. URL `https://arxiv.org/abs/2503.01141`.

Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857, 2022.

Zhaoyi Li, Gangwei Jiang, Hong Xie, Linqi Song, Defu Lian, and Ying Wei. Understanding and patching compositional reasoning in llms, 2024. URL `https://aclanthology.org/2024.findings-acl.576/`.

Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning, 2025a. URL `https://arxiv.org/abs/2501.12570`.

Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing o1-preview with a 1.5b model by scaling rl. `https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005b`, 2025b. Notion Blog.

Trung Quoc Luong, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin, and Hang Li. Reft: Reasoning with reinforced fine-tuning, 2024. URL `https://aclanthology.org/2024.acl-long.410/`.

Aman Madaan and Amir Yazdanbakhsh. Text and patterns: For effective chain of thought, it takes two to tango, 2022. URL `https://arxiv.org/abs/2209.07686`.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback, 2023. URL `https://arxiv.org/abs/2303.17651`.

Mathematical Association of America. Maa invitational competitions. `https://maa.org/maa-invitational-competitions/`, 2024. Accessed: 2025-09-24.

Mathematical Association of America. Maa invitational competitions. `https://maa.org/maa-invitational-competitions/`, 2025. Accessed: 2025-09-24.

William Merrill, Florian Tramèr, et al. On the expressive power of transformers with chain-of-thought, 2023. https://arxiv.org/pdf/2310.07923.

OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu,

Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki, Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay, Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung, Ian Kivlichan, Ian O'Connell, Ian O'Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers, Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lilian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin Zhang, Marwan Aljubeh, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers, Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi, Xia, Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas Cunninghman, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. Gpt-4o system card, 2024a. URL https://arxiv.org/abs/2410.21276.

OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally Bennett, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrew Duberstein, Andrew Kondrich,

Andrey Mishchenko, Andy Applebaum, Angela Jiang, Ashvin Nair, Barret Zoph, Behrooz Ghorbani, Ben Rossen, Benjamin Sokolowsky, Boaz Barak, Bob McGrew, Borys Minaiev, Botao Hao, Bowen Baker, Brandon Houghton, Brandon McKinzie, Brydon Eastman, Camillo Lugaresi, Cary Bassin, Cary Hudson, Chak Ming Li, Charles de Bourcy, Chelsea Voss, Chen Shen, Chong Zhang, Chris Koch, Chris Orsinger, Christopher Hesse, Claudia Fischer, Clive Chan, Dan Roberts, Daniel Kappler, Daniel Levy, Daniel Selsam, David Dohan, David Farhi, David Mely, David Robinson, Dimitris Tsipras, Doug Li, Dragos Oprica, Eben Freeman, Eddie Zhang, Edmund Wong, Elizabeth Proehl, Enoch Cheung, Eric Mitchell, Eric Wallace, Erik Ritter, Evan Mays, Fan Wang, Felipe Petroski Such, Filippo Raso, Florencia Leoni, Foivos Tsimpourlas, Francis Song, Fred von Lohmann, Freddie Sulit, Geoff Salmon, Giambattista Parascandolo, Gildas Chabot, Grace Zhao, Greg Brockman, Guillaume Leclerc, Hadi Salman, Haiming Bao, Hao Sheng, Hart Andrin, Hessam Bagherinezhad, Hongyu Ren, Hunter Lightman, Hyung Won Chung, Ian Kivlichan, Ian O'Connell, Ian Osband, Ignasi Clavera Gilaberte, Ilge Akkaya, Ilya Kostrikov, Ilya Sutskever, Irina Kofman, Jakub Pachocki, James Lennon, Jason Wei, Jean Harb, Jerry Twore, Jiacheng Feng, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joaquin Quiñonero Candela, Joe Palermo, Joel Parish, Johannes Heidecke, John Hallman, John Rizzo, Jonathan Gordon, Jonathan Uesato, Jonathan Ward, Joost Huizinga, Julie Wang, Kai Chen, Kai Xiao, Karan Singhal, Karina Nguyen, Karl Cobbe, Katy Shi, Kayla Wood, Kendra Rimbach, Keren Gu-Lemberg, Kevin Liu, Kevin Lu, Kevin Stone, Kevin Yu, Lama Ahmad, Lauren Yang, Leo Liu, Leon Maksin, Leyton Ho, Liam Fedus, Lilian Weng, Linden Li, Lindsay McCallum, Lindsey Held, Lorenz Kuhn, Lukas Kondraciuk, Lukasz Kaiser, Luke Metz, Madelaine Boyd, Maja Trebacz, Manas Joglekar, Mark Chen, Marko Tintor, Mason Meyer, Matt Jones, Matt Kaufer, Max Schwarzer, Meghan Shah, Mehmet Yatbaz, Melody Y. Guan, Mengyuan Xu, Mengyuan Yan, Mia Glaese, Mianna Chen, Michael Lampe, Michael Malek, Michele Wang, Michelle Fradin, Mike McClay, Mikhail Pavlov, Miles Wang, Mingxuan Wang, Mira Murati, Mo Bavarian, Mostafa Rohaninejad, Nat McAleese, Neil Chowdhury, Neil Chowdhury, Nick Ryder, Nikolas Tezak, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivia Watkins, Patrick Chao, Paul Ashbourne, Pavel Izmailov, Peter Zhokhov, Rachel Dias, Rahul Arora, Randall Lin, Rapha Gontijo Lopes, Raz Gaon, Reah Miyara, Reimar Leike, Renny Hwang, Rhythm Garg, Robin Brown, Roshan James, Rui Shu, Ryan Cheu, Ryan Greene, Saachi Jain, Sam Altman, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Santiago Hernandez, Sasha Baker, Scott McKinney, Scottie Yan, Shengjia Zhao, Shengli Hu, Shibani Santurkar, Shraman Ray Chaudhuri, Shuyuan Zhang, Siyuan Fu, Spencer Papay, Steph Lin, Suchir Balaji, Suvansh Sanjeev, Szymon Sidor, Tal Broda, Aidan Clark, Tao Wang, Taylor Gordon, Ted Sanders, Tejal Patwardhan, Thibault Sottiaux, Thomas Degry, Thomas Dimson, Tianhao Zheng, Timur Garipov, Tom Stasi, Trapit Bansal, Trevor Creech, Troy Peterson, Tyna Eloundou, Valerie Qi, Vineet Kosaraju, Vinnie Monaco, Vitchyr Pong, Vlad Fomenko, Weiyi Zheng, Wenda Zhou, Wes McCabe, Wojciech Zaremba, Yann Dubois, Yinghai Lu, Yining Chen, Young Cha, Yu Bai, Yuchen He, Yuchen Zhang, Yunyun Wang, Zheng Shao, and Zhuohan Li. Openai o1 system card, 2024b. URL `https://arxiv.org/abs/2412.16720`.

Sushrut Paul, Jinbo Yu, Jip J. Dekker, Alexey Ignatiev, and Peter J. Stuckey. Formal explanations for neuro-symbolic ai, 2024. URL `https://arxiv.org/abs/2410.14219`.

Jacob Pfau, William Merrill, and Samuel R. Bowman. Let's think dot by dot: Hidden computation in transformer language models, 2024. URL `https://arxiv.org/abs/2404.15758`.

Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL `https://arxiv.org/abs/2412.15115`.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. Gpqa: A graduate-level google-proof qa benchmark, 2023. URL `https://arxiv.org/abs/2311.12022`.

Michael Rescorla. The Language of Thought Hypothesis. In Edward N. Zalta and Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2024 edition, 2024.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL `https://arxiv.org/abs/1707.06347`.

Yuval Shalev, Amir Feder, and Ariel Goldstein. Distributional reasoning in llms: Parallel reasoning processes in multi-hop reasoning, 2024. URL `https://arxiv.org/abs/2406.13858`.

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

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024b. URL `https://arxiv.org/abs/2402.03300`. Introduces Group Relative Policy Optimization (GRPO).

Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang, Kai Wang, Zhaoxiang Liu, and Shiguo Lian. Dast: Difficulty-adaptive slow-thinking for large reasoning models, 2025. URL `https://arxiv.org/abs/2503.04472`.

Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, pp. 1279–1297. ACM, March 2025. doi: 10.1145/3689031.3696075. URL `http://dx.doi.org/10.1145/3689031.3696075`.

Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Na Zou, Hanjie Chen, and Xia Hu. Stop overthinking: A survey on efficient reasoning for large language models, 2025. URL `https://arxiv.org/abs/2503.16419`.

Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, Chuning Tang, Congcong Wang, Dehao Zhang, Enming Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze Li, Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin Su, Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi, Ling Ye, Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong, Shaowei Liu, Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao, Weimin Xiong, Weiran He, Weixiao Huang, Weixin Xu, Wenhao Wu, Wenyang He, Xianghui Wei, Xianqing Jia, Xingzhe Wu, Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles, Yang Li, Yangyang Hu, Yangyang Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu, Ying Yang, Yiping Bao, Yulun Du, Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, Zonghan Yang, and Zongyu Lin. Kimi k1.5: Scaling reinforcement learning with llms, 2025. URL `https://arxiv.org/abs/2501.12599`.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting, 2023. URL `https://arXiv.org/abs/2305.04388`.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2022a. URL `https://arxiv.org/abs/2201.11903`.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022b.

Tianhao Wu, Janice Lan, Weizhe Yuan, Jiantao Jiao, Jason Weston, and Sainbayar Sukhbaatar. Thinking llms: General instruction following with thought generation, 2024. URL https://arxiv.org/abs/2410.10630.

Yuyang Wu, Yikai Wang, Zichao Ye, Tengyu Du, Stefanie Jegelka, and Yiqing Wang. When more is less: Understanding chain-of-thought length in llms, 2025. URL https://arxiv.org/abs/2502.07266.

Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, Chenyang Shao, Yuwei Yan, Qinglong Yang, Yiwen Song, Sijian Ren, Xinyuan Hu, Yu Li, Jie Feng, Chen Gao, and Yong Li. Towards large reasoning models: A survey of reinforced reasoning with large language models, 2025. URL https://arxiv.org/abs/2501.09686.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL https://arxiv.org/abs/2505.09388.

Sohee Yang, Elena Gribovskaya, Nora Kassner, Mor Geva, and Sebastian Riedel. Do large language models latently perform multi-hop reasoning? In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2024.

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

Jingyang Yi, Jiazheng Wang, and Sida Li. Shorterbetter: Guiding reasoning models to find optimal inference length for efficient reasoning, 2025. URL https://arxiv.org/abs/2504.21370.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models, 2023. URL https://arxiv.org/abs/2309.12284.

Murong Yue, Jie Zhao, Min Zhang, Liang Du, and Ziyu Yao. Large language model cascades with mixture of thoughts representations for cost-efficient reasoning, 2024. URL https://arxiv.org/abs/2310.03094.

Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhu Chen. Mammoth: Building math generalist models through hybrid instruction tuning, 2023. URL https://arxiv.org/abs/2309.05653.

Haizhong Zheng, Yang Zhou, Brian R. Bartoldson, Bhavya Kailkhura, Fan Lai, Jiawei Zhao, and Beidi Chen. Act only when it pays: Efficient reinforcement learning for llm reasoning via selective rollouts, 2025. URL https://arxiv.org/abs/2506.02177.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. *arXiv preprint arXiv:2403.13372*, 2024.

Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. Agieval: A human-centric benchmark for evaluating foundation models, 2023. URL https://arxiv.org/abs/2304.06364.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting enables complex reasoning in large language models, 2022. URL https://arxiv.org/abs/2205.10625.

Zijian Zhou, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Alok Prakash, Daniela Rus, Jinhua Zhao, Bryan Kian Hsiang Low, and Paul Pu Liang. Mem1: Learning to synergize memory and reasoning for efficient long-horizon agents. *arXiv preprint arXiv:2506.15841*, 2025.

A APPENDIX

You may include other additional sections here.

A.1 LLM USAGE

We used LLM to help with paper writing for improving grammar and wording, and to help us write code for building plots.

A.2 METRIC FORMULATION AND HYPERPARAMETERS

A.2.1 METRICS

- **Pass@1:** Defined as

$$\text{pass@1} = \frac{1}{k} \sum_{i=1}^{k} p_i,$$

(3)

  where p_i denotes whether the i-th problem was solved correctly, and k is the total number of test problems. Intuitively, pass@1 measures the fraction of correctly solved problems under single-sample decoding.

- **Token Length:** The average number of tokens produced per response for a given benchmark. We first compute the mean output length across all questions in the dataset, and then report the overall average.

- **Compression Rate (CR):** Defined as

$$\text{CR} = \frac{L_{\text{DeepSeek-R1-1.5B}}}{L_{\text{model}}},$$

  where L denotes the average response length in tokens for a given benchmark. A higher CR indicates greater compression. For example, $\text{CR} = 10$ means that the model generates responses ten times shorter than DEEPSEEK-R1-1.5B.

A.2.2 HYPERPARAMETERS

Table 3: Key Hyperparameters for Supervised Fine-tuning

Parameter	Value	Description
Data Configuration		
Cutoff Length	15000	Maximum sequence length
Validation Split	0.05	Fraction of data used for validation
Model & Training		
Fine-tuning Type	Full	Full parameter fine-tuning
Learning Rate	1e-6	Optimizer learning rate
LR Scheduler	Cosine	Learning rate scheduling strategy
Warmup Steps	20	Number of warmup steps
Train Epochs	5	Total number of training epochs
Batch Configuration		
Per Device Train Batch	1	Training batch size per device
Gradient Accumulation	2	Steps to accumulate gradients
Evaluation & Saving		
Save Strategy	Steps	Save checkpoints by steps
Save Steps	0.20	Save every 20% of total steps
Eval Steps	0.05	Evaluate every 5% of total steps
System Configuration		
Template	DeepSeek-R1	Model template/format
Precision	BF16	Mixed precision training
Flash Attention	FA2	Attention optimization

Table 4: Key Hyperparameters for RLVR Training

Parameter	Value	Description
Data Configuration		
Train Batch Size	128	Batch size for training data
Max Prompt Length	1024	Maximum length of input prompts
Max Response Length	8192(Direct SLPO), 2048 (SLPO), 1024 (GRPO)	Maximum length of generated responses
Model & Training		
Learning Rate	1e-6	Actor model learning rate
PPO Mini Batch Size	64	Mini-batch size for PPO updates
KL Divergence Control		
KL Loss Coefficient	0.001	Weight for KL divergence loss
KL Loss Type	low_var_kl	Type of KL loss computation
Rollout Configuration		
Temperature	0.6	Sampling temperature for generation
Number of Samples	16	Samples per prompt during training
Training Schedule		
Total Epochs	5	Number of training epochs

A.3 MENTALESER-40K EXAMPLES

Example 1

```
Instruction:
The operation $\\otimes$ is defined for all nonzero numbers by $a \\otimes b = \\frac{a^{2}}{b}$.
    Determine $[(1 \\otimes 2) \\otimes 3] - [1 \\otimes (2 \\otimes 3)]$.

Mentalese: <think>DEF:(a,b)=a^2/b;CALC1:x=12=1/2;CALC2:y=x3=(1/2)^2/3=1/12;CALC3:z=23=4/3;
    CALC4:w=1z=1/(4/3)=3/4;SUB:y-w=1/12-3/4=-8/12=-2/3; ANS:-2/3</think>\\boxed{-\\frac{2}{3}}
```

Example 2

```
Instruction:
Doug constructs a square window using $8$ equal-size panes of glass. The ratio of the height to width
    for each pane is $5 : 2$, and the borders around and between the panes are $2$ inches wide. In
    inches, what is the side length of the square window?

Mentalese:
<think>LET:x;PANE_HEIGHT:5x;PANE_WIDTH:2x;ROWS:2;COLS:4;HEIGHT_PANES:2*5x=10x;WIDTH_PANES:4*2x=8x;
    HEIGHT_BORDERS:3*2=6;WIDTH_BORDERS:5*2=10;TOTAL_HEIGHT:10x+6;TOTAL_WIDTH:8x+10;EQ:10x+6=8x+10;
    SOLVE:x=2;WINDOW_SIDE:10*2+6=26;ANS:26</think>\\boxed{26}
```

Example 3

```
Instruction:
Let $P(x)$ be a polynomial of degree $3n$ such that\n\\begin{align*} P(0) = P(3) = \\dots = P(3n) &=
    2, \\\\ P(1) = P(4) = \\dots = P(3n+1-2) &= 1, \\\\ P(2) = P(5) = \\dots = P(3n+2-2) &= 0.
    \\end{align*}\nAlso, $P(3n+1) = 730$. Determine $n$.

Mentalese:
<think>DEF:P(x),deg=3n;SET:P(3k)=2,k=0..n;SET:P(3k+1)=1,k=0..n-1;SET:P(3k+2)=0,k=0..n-1;COUNT:pts=3n+1;
    EQ:deg+1=pts;SOLVE:n=1;ANS:1</think>\\boxed{1</answer}
```

Example 4

```
Instruction:
At Euclid Middle School the mathematics teachers are Mrs. Germain, Mr. Newton, and Mrs. Young. There
    are $11$ students in Mrs. Germain's class, $8$ students in Mr. Newton's class, and $9$ students
    in Mrs. Young's class taking the AMC $8$ this year. How many mathematics students at Euclid
    Middle School are taking the contest?

Mentalese:
If $991+993+995+997+999=5000-N$, then $N=$", "output":
    "<think>SUM:s=991+993+995+997+999;CALC:s=4975;EQ:4975=5000-N;SOLVE:N=5000-4975;CALC:N=25;
    ANS:25</think>\\boxed{25</answer}
```

Example 5

```
Instruction:
Let $f$ be the function defined by $f(x)=ax^2-\\sqrt{2}$ for some positive $a$. If
    $f(f(\\sqrt{2}))=-\\sqrt{2}$ then $a=$

Mentalese:
<think>DEF:f(x)=ax^2-2;EVAL1:f(2)=2a-2;EVAL2:f(f(2))=a(2a-2)^2-2;EXPAND1:(2a-2)^2=4a^2-4a2+2;
    EVAL3:a(4a^2-4a2+2)-2=4a^3-4a22+2a-2;SET:4a^3-4a22+2a-2=-2;ADD:4a^3-4a22+2a=0;
    FACTOR:2a(2a^2-2a2+1)=0;SOLVE:2a^2-2a2+1=0;a=[22(8-8)]/4;a=2/2;
    ANS:2/2</think>\\boxed{\\frac{\\sqrt{2}}{2}}
```

Figure 4: Violin plots of token usage in Agentica-24k responses across six benchmarks before and after fine-tuning. The Base model generates very long responses, while Direct SLPO provides only limited compression. Mentalese-based methods (SFT, SFT+GRPO, SFT+SLPO) achieve 10–20× reduction in response length, approaching an optimal reasoning length that balances efficiency with performance. Although some performance degradation occurs, the Mentalese training pipeline with RLVR methods offers the best trade-off between token efficiency and problem-solving ability.

A.4 PROMPT STRUCTURE

Direct Prompting Baseline

System prompt:

You are a mathematical reasoning and cognitive modeling assistant. Your job is to generate a detailed step-by-step solution, and then compress it into the given format for cognitive modeling. You are always given the final answer to guide your solution.
Please output the following:

1. **Mentalese CoT** – A symbolic, logic-based chain of thought using compact primitives and logical steps only. Each step should be separated by a semicolon (`;`) with **no natural language**. All core logic steps must be preserved.

Mentalese Format Rules

- Use **ONLY** symbolic notation in the form: `OP:params` or `OP:result`
- No natural language
- No redundancy
- Every step must be **necessary and sufficient** for solving the problem
- Example:

  SET:w;EQ:abs(180-5.5*w)=110;CASE1:180-5.5*w=110;SOLVE1:w=70/5.5;
  CASE2:180-5.5*w=-110;SOLVE2:w=290/5.5;CALC1:w1=70/5.5=12+8/11;
  CALC2:w2=290/5.5=52+8/11;DIFF:t=w2-w1=40;ANS:40

Wrap the output in a single JSON object with the following keys:
{ "mentalese_cot"}

User:
QUESTION: [Your question here]
ANSWER: [Your answer here]

RESPONSE:

A.5 MODEL RESPONSES

Question (Taken from Math-500)

What is the smallest positive perfect cube that can be written as the sum of three consecutive integers?

Ground Truth

27

Magentica   **Length: 122 Tokens**

```
<think>\nSET:n;SUM:n+(n+1)+(n+2)=3n+3;EQ:3n+3=k^3;DIV:3n+3=3(n+1);
    SET:k=3m;EQ:3(n+1)=(3m)^3=27m^3;EQ:n+1=9m^3;SOLVE:n=9m^3-1; CHECK:smallest
    m=1;CALC:n=8;SUM:8+9+10=27;ANS:27\n</think>\n\n\\boxed{27}
```

Magentica-SLPO   **Length: 351 Tokens**

```
<think>\nGiven: Find the smallest positive perfect cube that can be written as the sum of three
    consecutive integers.\n\nLet me denote the three consecutive integers as n-1, n, n+1. Their sum
    is (n-1) + n + (n+1) = 3n.\n\nSo, we need 3n to be a perfect cube.\n\nLet me denote the perfect
    cube as k^3. So, 3n = k^3 => n = k^3 / 3.\n\nSince n must be an integer, k^3 must be divisible
    by 3. Therefore, k must be divisible by 3. Let me set k = 3m, where m is an integer.\n\nThen,
    k^3 = (3m)^3 = 27m^3.\n\nSo, 3n = 27m^3 => n = 9m^3.\n\nTherefore, the sum of three consecutive
    integers is 3n = 27m^3, which is a perfect cube.\n\nNow, we need the smallest positive perfect
    cube. Since m is a positive integer, the smallest m is 1.\n\nTherefore, the smallest perfect
    cube is 27*1^3 = 27.\n\nLet me verify: Find three consecutive integers that sum to 27.\n\nn =
    9m^3 = 9*1 = 9.\n\nSo, the integers are 8, 9, 10. Their sum is 8+9+10=27.\n\nYes, that
    works.\n\nTherefore, the smallest positive perfect cube is 27.\n</think>\n\n\\boxed{27}
```

Magentica-GRPO   **Length: 236 Tokens**

```
<think>Let's denote the three consecutive integers as n-1, n, n+1. Their sum is (n-1)+n+(n+1)=3n. We
    need 3n to be a perfect cube. So, n must be a multiple of 3. Let n=3k. Then 3n=9k. We need 9k
    to be a perfect cube. Let's write 9k=m^3. Then k=m^3/9. For k to be an integer, m^3 must be
    divisible by 9, so m must be divisible by 3. Let m=3t. Then m^3=27t^3. So k=27t^3/9=3t^3.
    Therefore, n=3k=9t^3. Then 3n=27t^3=(3t^3)^3? Wait, no. 27t^3 is (3t)^3. Yes, 27t^3=(3t)^3. So
    the smallest positive perfect cube is when t=1, so 27.\n</think>\n\n\\boxed{27}
```

Agentic-24k-SLPO   **Length: 526 Tokens**

```
<think>\nOkay, so I need to find the smallest positive perfect cube that can be written as the sum of
    three consecutive integers. Hmm, let's break this down step by step.\n\nFirst, let's denote the
    three consecutive integers. Let's say the integers are n-1, n, and n+1. So their sum would be
    (n-1) + n + (n+1). Simplifying that, it's n-1 + n + n+1 = 3n. So the sum of three consecutive
    integers is 3n.\n\nNow, we need this sum to be a perfect cube. So, 3n must be a perfect cube.
    Let's denote the perfect cube as k, where k is a positive integer. Therefore, 3n = k, which
    means n = k / 3.\n\nSince n has to be an integer (because we're dealing with integers), k must
    be divisible by 3. That implies that k itself must be divisible by 3, because if 3 divides k,
    then 3 must divide k. So let's let k = 3m, where m is a positive integer.\n\nSubstituting back,
    we have k = (3m) = 27m. Therefore, n = 27m / 3 = 9m.\n\nSo, the sum of three consecutive
    integers is 3n = 3 * 9m = 27m. Wait, that seems a bit off. Let me correct that.\n\nIf n = 9m,
    then the sum is 3n = 27m. So, the sum is 27m, which is a perfect cube since 27 is 3 and m is a
    cube. Therefore, the sum is (3m), which is indeed a perfect cube.\n\nNow, we need the smallest
    positive perfect cube. So we need the smallest m such that 27m is positive. Since m is a
    positive integer, the smallest m is 1.\n\nTherefore, the smallest positive perfect cube is 27 *
    1 = 27.\n\nLet me verify this. If m=1, then n=9*1=9. So the three consecutive integers are
    8,9,10. Their sum is 8+9+10=27, which is 3. That checks out.\n</think>\n\nThe smallest positive
    perfect cube that can be written as the sum of three consecutive integers is \\boxed{27}.
```

Agentica-24k **Length: 1512 Tokens**

<think>\nOkay, so I need to find the smallest positive perfect cube that can be written as the sum of three consecutive integers. Hmm, let's break this down step by step.\n\nFirst, let's understand what it means for a number to be a perfect cube. A perfect cube is an integer that can be written as n, where n is an integer. The smallest positive perfect cubes are 1=1, 2=8, 3=27, and so on.\n\nNow, the problem states that this perfect cube must be expressible as the sum of three consecutive integers. Let's denote these three consecutive integers as n-1, n, and n+1. Wait, actually, maybe it's better to represent them as x, x+1, x+2, where x is an integer. That might be more straightforward.\n\nSo, the sum of three consecutive integers can be written as x + (x+1) + (x+2). Let's compute this:\n\nx + (x+1) + (x+2) = x + x + 1 + x + 2 = 3x + 3 = 3(x + 1).\n\nSo, the sum of three consecutive integers is equal to 3 times (x + 1). That means that the sum must be a multiple of 3.\n\nNow, since we're looking for a perfect cube, let's denote the perfect cube as k, where k is a positive integer. So, we have:\n\nk = 3(x + 1).\n\nTherefore, k must be divisible by 3. Which implies that k itself must be divisible by 3, because if a prime number divides a cube, it must divide the base. Since 3 is prime, k must be a multiple of 3. Let's denote k as 3m, where m is a positive integer.\n\nSubstituting back, we get:\n\n(3m) = 3(x + 1)\n\nCompute the left side:\n\n27m = 3(x + 1)\n\nDivide both sides by 3:\n\n9m = x + 1\n\nTherefore, x + 1 = 9m, which implies x = 9m - 1.\n\nNow, since x is an integer, and we're looking for positive integers, let's see what m can be.\n\nWe need k to be positive, so k must be at least 1. But since k is a multiple of 3, the smallest possible k is 3.\n\nLet's compute the smallest possible k, which would be when m=1, because m=1 gives k=3.\n\nSo, k=3, then k=27.\n\nNow, let's check if 27 can be expressed as the sum of three consecutive integers.\n\nFrom earlier, we know that the sum is 3(x + 1). So, 27 = 3(x + 1) => x + 1 = 9 => x = 8.\n\nTherefore, the three consecutive integers are x=8, x+1=9, x+2=10. Let's check their sum: 8+9+10=27. Yes, that works.\n\nBut wait, let's confirm if 27 is indeed the smallest such perfect cube. Let's check smaller k values that are multiples of 3.\n\nThe next smaller k would be k=0, but that would give k=0, which isn't positive. Then k=3 is the next one, which gives k=27. So, 27 is indeed the smallest positive perfect cube that can be written as the sum of three consecutive integers.\n\nBut just to be thorough, let's check if there's any smaller k, perhaps non-multiples of 3? Wait, earlier reasoning showed that k must be a multiple of 3 because k must be divisible by 3. Therefore, the smallest k is 3, leading to k=27. So, yes, 27 is the answer.\n\nWait, hold on, let me think again. The problem says \"the smallest positive perfect cube that can be written as the sum of three consecutive integers.\" So, perhaps I should check if there's a smaller cube than 27 that can be expressed as such.\n\nLet's consider the cubes:\n\nn1=1: Can 1 be written as the sum of three consecutive integers? Let's see: Let x + (x+1) + (x+2) = 3x + 3 =1. Then 3x= -2 => x= -2/3. Not an integer, so no.\n\nn2=8: Let's see if 8 can be written as the sum of three consecutive integers. Let's set up the equation: 3x + 3 =8 => 3x=5 => x=5/3. Not an integer. So, no.\n\n3=27: As before, x=8, which is an integer, so yes.\n\nSo, indeed, 27 is the smallest positive perfect cube that can be written as the sum of three consecutive integers.\n\nTherefore, the answer is 27.\n\n**Final Answer**\nThe smallest positive perfect cube is \\boxed{27}.\n</think>\n\nTo find the smallest positive perfect cube that can be written as the sum of three consecutive integers, we start by noting that the sum of three consecutive integers can be expressed as \\(x + (x+1) + (x+2) = 3x + 3 = 3(x + 1) \\).\n\nThis sum must be a perfect cube, denoted as \\(k^3 \\). Therefore, we have:\n\\[k^3 = 3(x + 1) \\]\n\nSince \\(k^3 \\) must be divisible by 3, \\(k \\) itself must be a multiple of 3. Let \\(k = 3m \\) where \\(m \\) is a positive integer. Substituting this, we get:\n\\[(3m)^3 = 3(x + 1) \\]\n\\[27m^3 = 3(x + 1) \\]\n\\[9m^3 = x + 1 \\]\n\\[x = 9m^3 - 1 \\]\n\nWe need the smallest \\(k \\) such that \\(k^3 \\) is positive. The smallest \\(k \\) is 3, leading to:\n\\[k^3 = 3^3 = 27 \\]\n\nWe verify that 27 can be expressed as the sum of three consecutive integers:\n\\[3(x + 1) = 27 \\]\n\\[x + 1 = 9 \\]\n\\[x = 8 \\]\n\nThus, the three consecutive integers are 8, 9, and 10, whose sum is 27.\n\nTherefore, the smallest positive perfect cube is \\(\\boxed{27}\\).