

# RISER: Orchestrating Latent Reasoning Skills for Adaptive Activation Steering

Wencheng Ye<sup>1</sup> Xiaoyang Yuan<sup>1</sup> Yi Bin<sup>1†</sup>  
Hengyu Jin<sup>1</sup> Liang Peng<sup>2</sup> Pengpeng Zeng<sup>1</sup> Heng Tao Shen<sup>1</sup>

<sup>1</sup>School of Computer Science and Technology, Tongji University

<sup>2</sup>Tencent IEG Global

{2350227, enigma, 2352280}@tongji.edu.cn

{yi.bin, shenhengtao}@hotmail.com

is.pengpengzeng@gmail.com leonllpeng@tencent.com

## Abstract

Recent work on domain-specific reasoning with large language models (LLMs) often relies on training-intensive approaches that require parameter updates. While activation steering has emerged as a parameter-efficient alternative, existing methods apply static, manual interventions that fail to adapt to the dynamic nature of complex reasoning. To address this limitation, we propose **RISER** (Router-based Intervention for Steerable Enhancement of Reasoning), a plug-and-play intervention framework that adaptively steers LLM reasoning in activation space. RISER constructs a library of reusable reasoning vectors and employs a lightweight Router to dynamically compose them for each input. The Router is optimized via reinforcement learning under task-level rewards, activating latent cognitive primitives in an emergent and compositional manner. Across seven diverse benchmarks, RISER yields 3.4–6.5% average zero-shot accuracy improvements over the base model while surpassing CoT-style reasoning with 2–3× higher token efficiency and robust accuracy gains. Further analysis shows that RISER can autonomously combine multiple vectors into interpretable, precise control strategies, paving way for more controllable and efficient LLM reasoning.

## 1 Introduction

Large Language Models (LLMs) (Touvron et al., 2023; Brown et al., 2020) exhibit strong general reasoning abilities (Ahn et al., 2024; Wei et al., 2022a; Shi et al., 2024), yet they often perform inconsistently on specialized downstream tasks requiring domain knowledge, disciplined reasoning patterns, and the coordinated use of multiple cognitive skills (Wang et al., 2025b; Li et al., 2025; Chang et al., 2024). In practice, we seek

<sup>†</sup> Corresponding author.

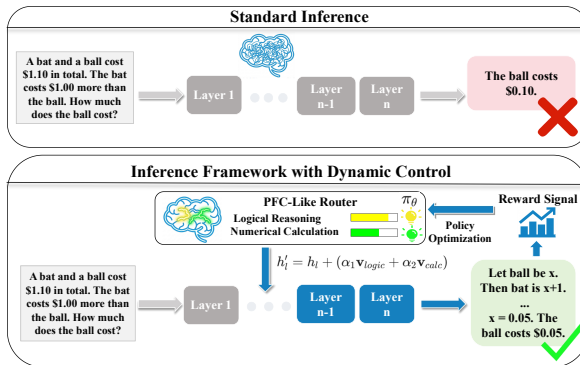


Figure 1: Conceptual comparison between Standard Inference and the RISER framework. RISER (bottom) uses a learned Router to dynamically inject composed vectors, analogous to an explicit executive-control mechanism.

to strengthen such reasoning without expensive retraining or relying solely on indirect prompting strategies, motivating mechanisms that can directly modulate the model’s internal computation during inference.

Existing approaches face fundamental limitations. Training-based methods such as Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL) require invasive parameter updates (Shenfeld et al., 2025; Huan et al., 2025), leading to issues such as catastrophic forgetting (Li et al., 2024; Ding and Wang, 2025) or substantial computational overhead (Liao et al., 2025). Meanwhile, training-free prompting (Wang et al., 2025a) suffers from signal attenuation during forward propagation (Wu et al., 2025). Activation steering offers a promising alternative by directly modifying internal activations without changing model weights, but existing methods predominantly use a *single* steering vector with *fixed* intervention strength (Venhoff et al., 2025b; Højer et al., 2025; Jin et al., 2025). Such static, one-dimensional control limits their expressiveness and fails to capture the rich, multi-faceted structure of reasoning embedded in large models.

Decades of cognitive neuroscience research suggest a more flexible architecture: human cognition emerges from **modular** functional regions (Kanwisher et al., 1999; Anderson, 2010) coordinated by the **Prefrontal Cortex (PFC)** through dynamic routing of control signals (Miller and Cohen, 2001). This perspective invokes a different blueprint and highlights two key elements largely absent in current activation steering methods: a set of diverse reasoning primitives and a controller capable of selecting and composing them adaptively. Inspired by this, we ask: *Can we build a PFC-like controller that adaptively awakens, routes, and composes latent cognitive capabilities through activation-level interventions during inference?*

Furthermore, recent advances in Representation Engineering (Tan et al., 2024; Postmus and Abreu, 2024; Alain and Bengio, 2016) show that LLMs’ activation spaces contain interpretable, semantically meaningful directions to capture attributes, skills, and latent reasoning patterns (Lee et al., 2025b; Marks and Tegmark, 2023; Cyberek and Evans, 2025; Zhang et al., 2025a; Rinsky et al., 2024). These directions can be disentangled, composed, and manipulated to alter model behavior (Fartale et al., 2025). Such findings support the hypothesis that complex reasoning may be decomposable into multiple linear subspaces, each corresponding to a distinct cognitive capability. If true, then static single-vector steering is fundamentally misaligned with the structure of the model: what is truly needed is dynamic, compositional, and task-aware control.

Therefore, we propose a dynamic activation steering approach, termed **RISER** (**R**outer-based **I**ntervention for **S**teerable **E**nhancement of **R**easoning), which enables adaptive and compositional control of the model’s internal reasoning process. RISER treats distinct activation patterns as reusable reasoning directions (which we call cognitive primitives). We first extract reasoning vectors that encode core cognitive functions to form a reusable capability library. Then, we introduce a lightweight Router module acting as a reasoning controller. Given an input query, the Router dynamically selects relevant vectors, determines optimal intervention strengths, and injects the combined representation into the model’s forward pass (see Figure 1). We use RL to directly optimize the routing policy, effectively externalizing the model’s implicit logic into explicit, inter-

pretable control decisions.

Our experiments across seven diverse reasoning benchmarks demonstrate RISER’s effectiveness. It yields 3.4%–6.5% absolute zero-shot accuracy improvements while maintaining inference efficiency close to standard inference. Further analysis reveals that the RL-trained Router autonomously composes multiple primitives in meaningful ways, offering transparent insight into which cognitive capabilities are invoked and how they synergize for a given task. We summarize our primary contributions as follows:

- We propose RISER, a plug-and-play activation intervention framework that keeps LLM parameters frozen while a lightweight Router dynamically selects and composes cognitive primitives, enabling precise and task-adaptive reasoning control.
- We develop a rigorous pipeline for eliciting high-quality reasoning vectors, incorporating LLM-Judge filtering and clustering to construct an orthogonal and disentangled library of cognitive primitives with verified steering effects.
- Extensive evaluation on seven benchmarks shows that we establish a new state-of-the-art for activation steering, closing the gap between inference-time control and heavy fine-tuning while offering interpretable insights into latent skill composition.

## 2 Related Work

### 2.1 Activation Steering

The study of linear representations has evolved from passive probing (Alain and Bengio, 2016; Belinkov, 2022) to activation steering (Turner et al., 2023; Bartoszcze et al., 2025; Wang et al., 2025c), which enables active intervention on high-level concepts. Effective methods for concept vector extraction include Contrastive Activation Addition (CAA) (Rinsky et al., 2024) and SAEs (Cunningham et al., 2023). More recently, this paradigm has been extended to cognitive processes and reasoning, as it has been demonstrated that a reasoning vector extracted from one task can be applied to improve accuracy on another, thus confirming that reasoning capabilities are transferable (Højer et al., 2025; Wang et al., 2025a; Venhoff et al., 2025a; Zbeeb et al., 2025; Valentino et al., 2025). However, they typically apply a single, fixed vector with a manually-tuned strength for all inputs, failing to adapt to the specific de-

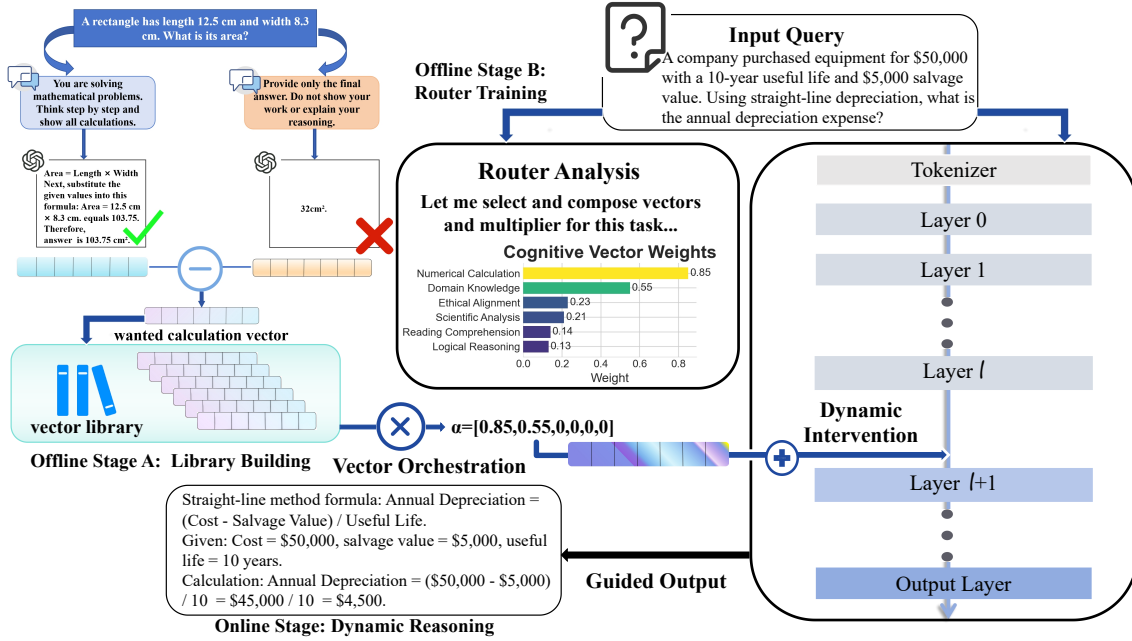


Figure 2: An overview of the RISER framework, illustrating the process of offline extraction of reasoning vectors and offline training of the Router, followed by online inference where the pre-trained Router dynamically selects and combines vectors to intervene in the LLM’s activation, guiding the final output.

mands of each task. While some work has introduced limited dynamics through gated activation (Jin et al., 2025) or strength calculation for single vectors (Zhang et al., 2025a), they do not address the challenge of composing multiple capabilities with adaptive intervention strengths and rely on supervised objectives for training.

## 2.2 Conditional Computation and Modular Networks

Modular networks can be categorized into two main granularities: intra-model and inter-model. A prominent intra-model approach is the Mixture-of-Experts (MoE) architecture (Shazeer et al., 2017; Fedus et al., 2022) that activates a small subset of experts for each input. At a coarser, inter-model granularity, researchers have explored task allocation among multiple independent LLMs (Zhang et al., 2025b; Piskala et al., 2024). RISER applies the dynamic routing philosophy to a single, frozen LLM at the representation layer, offering greater flexibility and controllability.

## 3 The RISER Framework

### 3.1 Overall Architecture

As depicted in Figure 2, RISER follows a simple offline-online split. Offline, we assume access to a compact library of reasoning vectors  $\{\mathbf{v}_i\}_{i=1}^K$ , each representing a reusable cognitive primitive.

Given this library, we train a lightweight Router that learns to map internal states to compositions of these primitives (Section 3.3). In Section 4, we describe one concrete, data-driven instantiation process. In the online inference phase, the prepared components are used to intervene minimally but effectively in the LLM’s forward computation. As an input query is processed up to a pre-determined intermediate layer  $l$ , we read the hidden state of the last token,  $\mathbf{h}_l$ , and feed it into the Router, which infers the immediate cognitive demands and outputs a composite reasoning vector by selecting and weighting a small subset of primitives. The resulting vector is injected back into the model to the activation at layer  $l$ , and maintained as a sustained cognitive priming during decoding.

### 3.2 Router as a Dynamic Controller

The Router is a lightweight network that reads the model’s current hidden state and outputs an intervention in activation space. It receives the hidden state of the last token at the target layer  $l$ , denoted by  $\mathbf{h}_l \in \mathbb{R}^d$ , which serves as a natural proxy for the task’s current reasoning demands. From this state, the Router produces two parallel outputs over the  $K$  primitives at a sequence-level: a selection mask  $\mathbf{w} \in \{0, 1\}^K$  and a strength vector  $\alpha \in [0, \alpha_{\max}]^K$ . In practice, the selection head first outputs a probability vector  $\mathbf{p} \in [0, 1]^K$  via a

Sigmoid activation, which is thresholded at inference time to obtain the binary mask  $\mathbf{w}$  that specifies which capabilities to activate. During training, we apply the Gumbel-Sigmoid relaxation (Jang et al., 2017). In parallel, the strength head predicts  $\alpha$ , which specifies *how strongly* to move along each selected direction. This dual-head design decouples the discrete choice of primitives from continuous intensity modulation, encouraging sparse yet flexible control. The final composite vector for injection is synthesized via a weighted summation over the primitive library:

$$\mathbf{v}_{\text{inject}} = \sum_{i=1}^K w_i \cdot \alpha_i \cdot \mathbf{v}_i, \quad (1)$$

where  $w_i$  and  $\alpha_i$  are the  $i$ -th elements of  $\mathbf{w}$  and  $\alpha$ , respectively, and  $\mathbf{v}_i$  is the  $i$ -th reasoning primitive. This vector is then injected into the LLM’s forward pass via an element-wise addition to the last-token activation at layer  $l$ :

$$\mathbf{h}'_l = \mathbf{h}_l + \mathbf{v}_{\text{inject}}, \quad (2)$$

where  $\mathbf{h}_l$  is the original hidden activation and  $\mathbf{h}'_l$  is the resulting steered activation. The model subsequently resumes its computation from layer  $l+1$  onward using  $\mathbf{h}'_l$ , so that the remaining layers perform their usual processing under a slightly reoriented internal state toward the desired reasoning trajectory.

### 3.3 Router Optimization

Given the Router’s role, we optimize it directly from task-level feedback, enabling it to learn when and how to combine primitives to best serve downstream objectives.

**Supervised Warm-Up.** To avoid a cold-start regime, we first train the Router on a curated dataset derived from the vector library. For each training instance, we use grid search over the library to identify an intervention configuration  $(\mathbf{w}^*, \alpha^*)$  that successfully elicits the correct reasoning generation, establishing a robust baseline policy that knows which primitives tend to be useful in which contexts.

**Reinforcement Learning Refinement.** To move beyond the limitations of this static dataset and adapt to unseen task variations, we then fine-tune the Router with Group Relative Policy Optimization (GRPO) (Shao et al., 2024). In this stage, the Router is free to explore the space of primitive

compositions. We use an accuracy-based reward

$$r_i = \begin{cases} 1, & \text{if the answer is correct,} \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

which directly encourages policies that increase end-task accuracy. To ensure that interventions remain conservative and stable, we additionally impose a KL regularizer which at each decoding step  $t$  computes

$$\mathcal{L}_{\text{KL}} = \mathbb{E}_t \left[ D_{\text{KL}}(\pi_{\text{routed}}(\cdot | x, y_{<t}) \| \pi_{\text{base}}(\cdot | x, y_{<t})) \right], \quad (4)$$

where  $\pi_{\text{routed}}$  and  $\pi_{\text{base}}$  denote the output distributions of the routed and base models. This regularizer discourages vector injections that induce large deviations from the base model’s behavior while still allowing beneficial changes, enabling the Router to learn primitive compositions without destabilizing the underlying model.

## 4 Vector Elicitation Pipeline

As effectiveness of the Router depends on the quality of the underlying primitives, we construct a concrete, data-driven vector elicitation pipeline. We first generate a broad candidate set of activation pairs following Contrastive Activation Additions. For a diverse range of reasoning tasks, we design paired prompts: a *positive* prompt to elicit rigorous reasoning and a *negative* prompt to suppress it. To mitigate noise, we integrate a quality-aware filtering mechanism using an LLM Judge (Claude-3.5-Sonnet (Anthropic, 2024)) which evaluates the generated reasoning traces with strict inclusion criteria: only sample pairs where the positive generation achieves proficient reasoning and the negative one clearly lacks reasoning are retained. To prevent manifold mismatch, we enforce size-specific native extraction (e.g., Llama primitives derived solely on Llama), yet the following walkthrough illustrates Qwen2.5-7B.

After filtering, we investigate the structure of the difference vectors  $(\mathbf{h}^+ - \mathbf{h}^-)$  to build a compact set of representative reasoning directions. Our goal is to construct a small, interpretable, and practically controllable library suitable for downstream routing. To this end, we perform PCA visualization to examine the coarse geometry of the activation manifold. As shown in Figure 3, the projections consistently exhibit several recurring high-density regions.

Dataset / Category	Base Model	CoT-Prompting	Self-Consistency	CoT	CAA	CAST	SAS	FR-Ponder	Our Method (RISER)
<b>Qwen2.5-7B-Instruct</b>									
MATH	46.8	51.5	<u>52.1</u>	47.1	49.8	48.3	49.2		<b>53.3(+6.5)</b>
GSM8K	79.8	<b>85.3</b>	<u>85.2</u>	82.9	84.6	83.8	84.4		<u>85.2(+5.4)</u>
<i>Average (Math/Logic)</i>	63.3	68.4	<u>68.7</u>	65.0	67.2	66.1	66.8		<b>69.3(+6.0)</b>
GPQA	31.0	31.2	33.2	31.9	<u>33.7</u>	32.6	33.0		<b>36.8(+5.8)</b>
MMLU-Pro (in-dist.)	44.1	44.0	44.2	46.2	47.5	46.9	<u>47.7</u>		<b>50.3(+6.2)</b>
ARC-C	63.7	63.3	64.4	63.3	65.3	64.1	<u>65.8</u>		<b>67.2(+3.5)</b>
<i>Average (General)</i>	46.3	46.2	47.3	47.1	<u>48.8</u>	47.9	<u>48.8</u>		<b>51.4(+5.1)</b>
Ethics	48.6	49.3	48.7	<b>53.2</b>	<u>52.4</u>	51.0	50.2		52.1(+3.5)
TruthfulQA	56.4	58.9	58.9	<u>59.6</u>	59.4	59.1	59.0		<b>59.8(+3.4)</b>
<i>Average (Moral)</i>	52.5	54.1	53.8	<b>56.4</b>	55.9	55.1	54.6		<u>56.0(+3.5)</u>
<b>Qwen2.5-14B-Instruct</b>									
MATH	55.6	58.9	<u>60.4</u>	56.4	59.0	57.4	58.2		<b>61.8(+6.2)</b>
GSM8K	86.5	<u>90.5</u>	89.7	88.6	90.1	89.2	90.0		<b>90.8(+4.3)</b>
<i>Average (Math/Logic)</i>	71.1	74.7	<u>75.1</u>	72.5	74.6	73.3	74.1		<b>76.3(+5.2)</b>
GPQA	32.8	33.6	34.5	33.7	<u>35.9</u>	34.2	35.3		<b>38.0(+5.2)</b>
MMLU-Pro (in-dist.)	51.2	50.9	51.8	54.0	55.1	54.7	<u>55.6</u>		<b>57.2(+6.0)</b>
ARC-C	67.3	67.6	68.2	67.1	69.4	67.9	<u>70.0</u>		<b>71.6(+4.3)</b>
<i>Average (General)</i>	50.4	50.7	51.5	51.6	53.5	52.3	<u>53.6</u>		<b>55.6(+5.2)</b>
Ethics	64.3	66.1	65.2	<b>68.5</b>	67.7	66.3	66.9		<u>67.9(+3.6)</u>
TruthfulQA	58.4	60.6	61.5	61.3	<u>61.8</u>	61.4	61.0		<b>62.1(+3.7)</b>
<i>Average (Moral)</i>	61.4	63.4	63.4	<u>64.9</u>	64.8	63.9	64.0		<b>65.0(+3.6)</b>
<b>Qwen2.5-32B-Instruct</b>									
MATH	57.7	60.8	<u>61.9</u>	58.8	60.9	59.6	60.4		<b>63.2(+5.5)</b>
GSM8K	90.9	93.2	<u>93.5</u>	91.9	93.0	92.6	92.7		<b>93.9(+3.0)</b>
<i>Average (Math/Logic)</i>	74.3	77.0	<u>77.7</u>	75.4	77.0	76.1	76.6		<b>78.6(+4.3)</b>
GPQA	48.0	48.5	49.9	48.9	<u>50.3</u>	49.2	50.0		<b>52.7(+4.7)</b>
MMLU-Pro (in-dist.)	55.1	55.6	56.5	57.4	<u>58.4</u>	58.0	<u>58.9</u>		<b>60.7(+5.6)</b>
ARC-C	70.4	70.1	71.4	70.9	72.5	71.6	<u>73.0</u>		<b>74.7(+4.3)</b>
<i>Average (General)</i>	57.8	58.1	59.3	59.1	60.4	59.6	<u>60.6</u>		<b>62.7(+4.9)</b>
Ethics	77.9	78.8	78.2	<u>81.0</u>	<b>81.4</b>	79.7	79.9		<b>81.4(+3.5)</b>
TruthfulQA	60.2	62.1	62.3	<u>62.8</u>	<b>64.0</b>	62.4	62.5		63.5(+3.3)
<i>Average (Moral)</i>	69.1	70.5	70.3	71.9	<b>72.7</b>	71.1	71.2		<u>72.5(+3.4)</u>
<b>Llama-3-8B-Instruct</b>									
MATH	30.9	33.6	34.1	32.8	33.2	32.4	34.2		<b>35.4</b>
GSM8K	84.5	88.2	87.3	88.5	88.7	88.0	<u>88.9</u>		<b>89.1</b>
<i>Average (Math/Logic)</i>	57.7	60.9	<u>60.7</u>	60.7	61.0	60.2	61.6		<b>62.3(+4.6)</b>
GPQA	25.8	26.3	<u>30.2</u>	25.6	27.4	26.0	27.9		<b>30.9</b>
MMLU-Pro	40.3	40.8	41.6	44.0	<u>44.8</u>	44.1	45.0		<b>45.9</b>
ARC-C	62.4	62.1	63.2	63.0	64.2	63.4	<u>64.6</u>		<b>66.1</b>
<i>Average (General)</i>	42.8	43.1	45.0	44.2	<u>45.5</u>	44.5	45.8		<b>47.6(+4.8)</b>
Ethics	55.2	56.3	55.7	<b>59.2</b>	58.5	57.5	58.0		<u>58.8</u>
TruthfulQA	44.0	45.2	47.5	48.2	<b>48.7</b>	47.3	47.8		48.4
<i>Average (Moral)</i>	49.6	50.8	51.6	<b>53.7</b>	53.6	52.4	52.9		<u>53.6(+4.0)</u>

Table 1: Performance Comparison on Qwen2.5 Family and Llama-3 Models. We report accuracy (%) and improvements ( $\Delta$ ) over the baseline. **Best** and Second best results are labeled.




Figure 3: Latent space visualization of extracted vectors. We project the high-dimensional difference vectors onto a 2D plane using PCA. The visualization reveals naturally forming clusters, demonstrating that the extracted reasoning vectors possess strong semantic separability within the activation space.

Rather than adopting a large number of clusters, we select major clusters that appear stably across random initializations and that jointly cover the dominant reasoning patterns observed in data (numerical reasoning, logical inference, ethical alignment, reading comprehension, scientific analysis, and domain-specific knowledge), maintaining a compact and human-understandable control space. Guided by the observation that the first six principal components account for over 85% of the total variance in the extracted vectors, we set the cluster count  $K = 6$  and apply K-Means clustering to formally separate the difference vectors. Let  $\mathcal{S}_i$  denote the set of sample indices assigned to the  $i$ -th cluster. The reasoning vector  $\mathbf{v}_i \in \mathbb{R}^d$  for primitive  $i$  is defined as the centroid:

$$\mathbf{v}_i = \frac{1}{|\mathcal{S}_i|} \sum_{j \in \mathcal{S}_i} (\mathbf{h}_j^+ - \mathbf{h}_j^-), \quad (5)$$

where  $\mathbf{h}_j^+$  and  $\mathbf{h}_j^-$  denote the positive and negative activations at layer  $l$  for the  $j$ -th pair. We apply L2 normalization ( $\mathbf{v}_i \leftarrow \mathbf{v}_i / \|\mathbf{v}_i\|_2$ ) to ensure consistent intervention magnitudes.

The resulting library shows that the vectors are nearly orthogonal, with an average pairwise cosine similarity  $< 0.1$  (Figure 4), indicating that the selected directions represent distinct and minimally overlapping cognitive behaviors. Furthermore, as shown in Figure 5, static steering experiments validate their functional efficacy: applying them can yield accuracy improvements on corresponding datasets. These findings confirm that our compact, engineering-driven extraction yields a set of effective and independently controllable reasoning modules.




Figure 4: Reasoning vector library similarity heatmap. The low off-diagonal values confirm that the extracted vectors represent distinct and separable cognitive functions.




Figure 5: Static steering validation. The performance sensitivity to steering strength ( $\alpha$ ) confirms that the extracted vectors effectively modulate specific reasoning behaviors.

## 5 Experiments

### 5.1 Experimental Setup

We employ the Qwen2.5 family (7B-Instruct, 14B-Instruct, 32B-Instruct) (Qwen et al., 2025) and Llama-3-8B-Instruct (Grattafiori et al., 2024) as base models and, following prior work (Chen et al., 2025), respectively select layers 20, 25, 40 and 13 as default layers for vector elicitation and intervention. We threshold gating selection probabilities at 0.7, set the maximum intervention strength to  $\alpha_{\max} = 2.0$  and implement the Router as a lightweight, bottleneck-style Multi-Layer Perceptron (MLP) with approximately 5 million parameters ( $< 0.1\%$  of base model), distinguishing our approach from simple linear classifiers while introducing negligible latency.

**Datasets.** The Vector Elicitation data consist of 500 problems randomly selected from MMLU (Hendrycks et al., 2021b), each paired with positive and negative guiding prompts; During Router Training, the SFT phase uses an automated pipeline to extract and annotate 200

samples from MMLU, while the RL phase employs MMLU-Pro (Wang et al., 2025d) as the resource for reinforcement learning refinement. We split MMLU-Pro into 70% training tasks for RL and 30% held-out tasks for evaluation, with no question overlap. The Evaluation Datasets include benchmarks chosen to cover diverse reasoning types including math/logic reasoning (GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021c)), general reasoning (GPQA (Rein et al., 2023), ARC-C (Clark et al., 2018), MMLU-Pro) and ethics and factual alignment (Ethics (Hendrycks et al., 2021a), TruthfulQA (Lin et al., 2022)).

**Baselines.** We compare against a set of baselines to quantify improvements: zero-shot base model; Chain-of-Thought (CoT) prompting (Wei et al., 2022b); Self-Consistency CoT (Wang et al., 2023) (with 5 samples and majority voting); CAA (static vector intervention with the best performance under different multipliers) (Rimsky et al., 2024); CAST (conditional activation steering) (Lee et al., 2025a); SAS (using sparse autoencoders for vector elicitation) (Bayat et al., 2025) and FR-Ponder (He and Tang, 2025) (using a controller to regulate reasoning depth by selecting steering vectors).

**Ablation Settings.** To dissect component contributions, we evaluate several variants: **Direct GRPO Fine-tuning** (GRPO algorithm on the backbone model under an equivalent computational budget); **SFT-only Router** (Router trained only in the supervised phase without RL refinement); **Top-1 Vector Only** (select only the single highest-strength reasoning vector, disabling vector composition); and **Layer Sensitivity Analysis** (interventions applied at layers adjacent to the default layer as well as at earlier layer 5 and later layer 28 to assess sensitivity to intervention depth).

**Evaluation Metrics.** We report primary task accuracy and token efficiency measured by the total number of tokens generated.

**Implementation Details:** For the SFT phase, we fine-tune the Router for 3 epochs with a learning rate of  $5 \times 10^{-6}$ . For RL phase we adopt a  $2 \times 10^{-6}$  learning rate, a batch size of 128, a maximum context length of 8192 tokens for 2 epochs.

## 5.2 Results

Table 1 presents the comprehensive results on models, where RISER exhibits consistent performance gains across different model families. Fo-

Category	Model Variant / Setting	MATH	GPQA	TruthfulQA
<b>Our Method (Full RISER @ L20)</b>		53.3	36.8	59.8
<i>Direct GRPO (full-model RL fine-tuning)</i>		47.6	34.6	58.6
<i>Training Ablation</i>	- w/o RL Refinement (SFT-only)	49.4	31.2	54.6
	- w/o Composition (Top-1 Only)	51.6	33.5	60.2
<i>Layer Sensitivity</i>	- Early Layer (L5)	48.5	31.5	55.0
	- Middle Layer (L19)	52.1	35.5	59.5
	- Middle Layer (L21)	51.8	34.6	59.6
	- Late Layer (L28)	49.0	32.0	56.1

Table 2: Comprehensive ablation studies on key framework components and design choices. We report accuracy (%) on representative datasets.

cusing on the primary Qwen family, our method (RISER) achieves the highest average accuracy in the challenging **General Reasoning** category, significantly outperforming all other methods. In **Math/Logic Reasoning**, our method also outperforms the strong Self-Consistency CoT baseline. This demonstrates the framework’s strong generalization and its ability to handle complex, multi-disciplinary tasks by dynamically composing capabilities. By learning to compose latent reasoning primitives only on one dataset, the Router acquires a transferable control strategy that generalizes across heterogeneous reasoning benchmarks.

We quantitatively analyze token efficiency on MATH and GPQA. Regarding efficiency, RISER requires only 1392 and 3056 tokens on MATH and GPQA, respectively, compared to 4033 and 6195 for CoT, realizing a 2–3 $\times$  gain. While CoT generates reasoning-helpful external text, RISER mobilizes latent circuits for higher computational utilization, bypassing the need for verbose textual scaffolding to guide the trajectory.

We compare our framework against static intervention CAA and other steering methods. The results clearly show the value of dynamic control. In the two categories requiring flexible, compositional reasoning (Math/Logic and General Reasoning), our dynamic Router significantly outperforms the static CAA baseline. Interestingly, in the Moral Alignment category, the static CAA or conditionally dynamic CAST baseline achieves the highest score, slightly edging out our method. This is likely because these tasks are highly uniform in their cognitive demands, and a strong, static application of the Ethical Alignment vector is highly effective. However, RISER still delivers substantial alignment improvements over all non-steering baselines.

The Router strategy heatmap in Figure 6 also provides a cognitive map which intuitively demonstrates the explicit policy learned by the Router. On one hand, it learns a highly logical and spe-

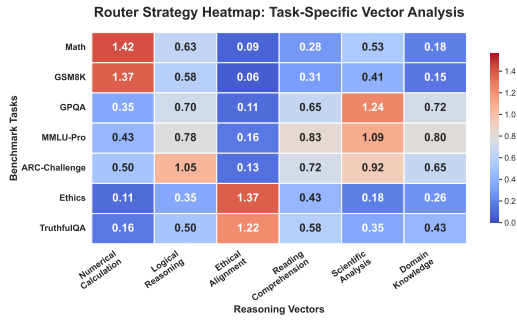


Figure 6: This heatmap shows the average strength assigned by the Router for each reasoning vector across different benchmarks and exhibits both logical specialization and complex composition.

cialized mapping: MATH and GSM8K tasks are strongly associated with the Numerical Calculation vector, while Ethics and TruthfulQA tasks correspond to the Ethical Alignment vector. On the other hand, when faced with complex cross-domain tasks (GPQA), the Router learns to autonomously compose multiple cognitive primitives. This provides direct evidence that the RL refinement phase externalized the LLM’s implicit, synergistic strategies for complex problem-solving into an analyzable model.

### 5.3 Ablation Studies

We performed ablation studies on Qwen2.5-7B-Instruct (Table 2) to isolate the contributions of key components.

**Comparison with Direct Fine-tuning:** A core question is whether the performance gains come from our RISER framework or simply from the RL training itself. To answer this, we compare RISER against the Direct GRPO baseline. RISER consistently outperforms the GRPO baseline in average accuracy across all three categories, which indicates that applying the same computational budget to train an external, dynamic reasoning controller is a more effective approach and validates its generalization advantage.

**Impact of RL Training:** The SFT-only Router significantly underperforms the full model, especially on complex benchmarks like GPQA and TruthfulQA, confirming that RL refinement is crucial for discovering synergistic vector compositions.

**Necessity of Composition:** Restricting the Router to a single vector hurts performance on multi-disciplinary tasks, validating the critical role of vector orchestration. Conversely, on the homogeneous TruthfulQA, the Top-1 variant achieves a marginal gain, indicating that our framework cor-

rectly adapts to favor focused, single-vector interventions for monolithic tasks.

**Layer Optimality:** Finally, Layer Sensitivity analysis identifies the middle layers as the optimal intervention site, showing robustness in adjacent layers but significant degradation at the model’s input and output layers. This observation confirms the hypothesis that reasoning processes crystallize within the middle layers, acting as a critical bridge between the initial input processing in early layers and the final linguistic realization in later layers.

### 5.4 Extensibility

To investigate extensibility, we extend RISER to a different domain and introduce an additional primitive targeting code synthesis. Following the same vector elicitation and routing pipeline, we expand the Router’s output space to seven dimensions and perform a brief SFT phase on 200 examples, updating only the Router. On HumanEval (Chen et al., 2021), the frozen base model achieves 56.3% pass@1, while static CAA improves performance to 57.2%. The extended Router over seven primitives further boosts accuracy to 59.9% and does not significantly affect performance on the original reasoning benchmarks, indicating that newly added primitives can be integrated in a non-interfering manner.

### 5.5 Transferability Across Models

We further examine whether RISER can be reused beyond the backbone on which it is derived and evaluate cross-model transfer by directly applying a trained RISER configuration to a different target model. Within the same model family, transferring RISER across parameter scales remains effective, suggesting that both the learned vector library and the Router’s composition strategy align reasonably well across scales. In contrast, transferring across different model families provides no benefit, indicating that the primitive directions and routing policy are tightly coupled to model-specific representation geometry and activation statistics. These results indicate that transfer is promising when the underlying activation manifolds are sufficiently aligned, but not across heterogeneous architectures. Further analysis are in Appendix C.

Target Model (Inference)	Source Router (Trained on)			
	Qwen-7B	Qwen-14B	Qwen-32B	Llama-3-8B
Qwen2.5-7B	53.3	51.7 (+4.9)	52.1 (+5.3)	46.5 (-0.3)
Qwen2.5-14B	58.1 (+2.5)	<b>61.8</b>	60.4 (+4.8)	55.8 (+0.2)
Qwen2.5-32B	59.8 (+2.1)	60.6 (+2.9)	<b>63.2</b>	57.5 (-0.2)
Llama-3-8B	30.5 (-0.4)	31.1 (+0.2)	30.8 (-0.1)	<b>35.4</b>

Table 3: Cross-Model Transferability on MATH. Off-diagonal entries show transfer results.

## 5.6 Comparison with Parameter-Efficient Fine-Tuning

To rigorously compare RISER against PEFT, we trained a dense LoRA baseline on Qwen2.5-7B using MMLU-Pro-CoT training trajectories. We updated all attention and MLP modules (rank  $r = 64$ ,  $\alpha = 128$ ,  $\approx 160M$  parameters) for 3 epochs.

Table 4: Comparison of Accuracy and Token Efficiency between LoRA and RISER on Qwen2.5-7B.

Dataset	Accuracy (%)			Average Tokens	
	Base+CoT	LoRA	RISER	LoRA	RISER
MATH	52.1	<b>53.3</b>	<b>53.3</b>	3853	<b>1392</b>
GSM8K	85.0	85.0	<b>85.2</b>	1857	<b>852</b>
GPQA	34.5	<b>36.8</b>	<b>36.8</b>	5924	<b>3056</b>
MMLU-Pro	46.2	<b>50.3</b>	<b>50.3</b>	4102	<b>1503</b>
ARC-C	61.8	<b>67.2</b>	<b>67.2</b>	2205	<b>1382</b>
Ethics	50.5	50.5	<b>52.1</b>	95	<b>33</b>
TruthfulQA	57.2	57.2	<b>59.8</b>	1108	<b>427</b>

As shown in Table 4, RISER offers a complementary approach optimized for token efficiency and instance-level control. RISER internalizes this computation, achieving comparable accuracy with substantially fewer generated tokens. Furthermore, instead of managing multiple task-specific adapters, RISER utilizes a single lightweight Router to dynamically compose cognitive primitives per prompt. This design naturally supports heterogeneous queries within the same batch and preserves the frozen base model’s intrinsic alignment.

## 5.7 Inference Overhead and Deployment

Beyond task accuracy, a critical factor for real-world adoption is the computational cost introduced during inference. Because the Router only processes the hidden state at a single intermediate layer, RISER adds mathematically negligible overhead; the module consumes only  $\sim 10MB$  of memory and  $\sim 10M$  FLOPs per invocation and computes instance-level interventions dynamically, natively supporting heterogeneous batching without degrading overall throughput. Ultimately, this lightweight design integrates seamlessly as a custom forward hook in high-throughput serving engines like vLLM.

## 6 Conclusion

RISER demonstrates that LLM reasoning can be effectively enhanced by orchestrating latent activations, offering a computationally efficient alternative to weight modification or verbose prompting. By learning explicit, RL-optimized policy, our framework achieves significant performance gains while validating the existence of steerable cognitive primitives within frozen models. This approach shifts the focus from surface-level text generation to internal state management, establishing a viable path toward more controllable and resource-efficient AI systems.

## 7 Limitations

Our framework, while effective, is constrained by its reliance on reactivating latent capabilities, making its performance bounded by the quality of the base model’s pre-training. The construction of a capability library with a fixed number of clusters further reflects an engineering trade-off: it stabilizes control but may reduce semantic granularity, potentially oversimplifying the underlying activation manifold for highly nuanced tasks. Moreover, the extracted vectors predominantly capture broad domain-level reasoning patterns due to the natural clustering structure of the model’s activation space. Future work can focus on disentangling these into finer-grained, domain-agnostic atomic skills, automating the discovery of such primitives, and exploring hierarchical routing mechanisms to achieve more precise control over complex reasoning chains. Finally, as RISER operates by modifying internal activations, careless application without proper constraints could lead to unintended behavioral shifts. In this work, we restrict our analysis to controlled benchmark settings, and future deployment-oriented use would require additional safety and alignment evaluation.

## Acknowledgment

This work is supported by the National Key R&D Program of China (2025YFE0202900), Central Guidance on Local Science and Technology Development Fund of Shanghai City (YDZX20253100002004), the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (JYB2025XDXM103 and JYB2025XDXM116), the New Cornerstone Science Foundation through the XPLOER PRIZE, and the Fundamental Research Funds for the Central Universities.

## References

- Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. 2024. [Large language models for mathematical reasoning: Progresses and challenges](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 225–237, St. Julian’s, Malta. Association for Computational Linguistics.
- Guillaume Alain and Yoshua Bengio. 2016. Understanding intermediate layers using linear classifier probes. *arXiv preprint arXiv:1610.01644*.
- Michael L. Anderson. 2010. Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences*, 33(4):245–66; discussion 266–313.
- Anthropic. 2024. [Model card addendum: Claude 3.5 haiku and upgraded claude 3.5 sonnet](#).
- Lukasz Bartoszcze, Sarthak Munshi, Bryan Sukidi, Jennifer Yen, Zejia Yang, David Williams-King, Linh Le, Kosi Asuzu, and Carsten Maple. 2025. Representation engineering for large-language models: Survey and research challenges.
- Reza Bayat, Ali Rahimi-Kalahroudi, Mohammad Pezeshki, Sarath Chandar, and Pascal Vincent. 2025. [Steering large language model activations in sparse spaces](#). *Preprint*, arXiv:2503.00177.
- Yonatan Belinkov. 2022. [Probing classifiers: Promises, shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. [A survey on evaluation of large language models](#). *ACM Trans. Intell. Syst. Technol.*, 15(3).
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Runjin Chen, Andy Arditi, Henry Sleight, Owain Evans, and Jack Lindsey. 2025. [Persona vectors: Monitoring and controlling character traits in language models](#). *Preprint*, arXiv:2507.21509.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems.
- Hoagy Cunningham, Aidan Ewart, Logan Riggs Smith, Robert Huben, and Lee Sharkey. 2023. [Sparse autoencoders find highly interpretable features in language models](#). *ArXiv*, abs/2309.08600.
- Hannah Cyberek and David Evans. 2025. [Steering the censorship: Uncovering representation vectors for LLM “thought” control](#). In *Second Conference on Language Modeling*.
- Fei Ding and Baiqiao Wang. 2025. [Improved supervised fine-tuning for large language models to mitigate catastrophic forgetting](#). *ArXiv*, abs/2506.09428.
- Harshwardhan Fartale, Ashish Kattamuri, Rahul Raja, Arpita Vats, Ishita Prasad, and Akshata Kishore Moharir. 2025. [Disentangling recall and reasoning in transformer models through layer-wise attention and activation analysis](#). *Preprint*, arXiv:2510.03366.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Yixin He and Lumingyuan Tang. 2025. [Learning to ponder: Adaptive reasoning in latent space](#). *Preprint*, arXiv:2509.24238.
- Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. 2021a. [Aligning ai with shared human values](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021c. [Measuring mathematical problem solving with the math dataset](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1.

- Bertram Højer, Oliver Jarvis, and Stefan Heinrich. 2025. [Improving reasoning performance in large language models via representation engineering](#). In *International Conference on Representation Learning*, volume 2025, pages 44746–44763.
- Maggie Huan, Yuetai Li, Tuney Zheng, Xiaoyu Xu, Seungone Kim, Minxin Du, Radha Poovendran, Graham Neubig, and Xiang Yue. 2025. [Does math reasoning improve general llm capabilities? understanding transferability of llm reasoning](#). *Preprint*, arXiv:2507.00432.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Haoran Jin, Meng Li, Xiting Wang, Zhihao Xu, Minlie Huang, Yantao Jia, and Defu Lian. 2025. [Internal value alignment in large language models through controlled value vector activation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 27347–27371, Vienna, Austria. Association for Computational Linguistics.
- Nancy Kanwisher, Josh McDermott, and Marvin M. Chun. 1999. The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, 17(11).
- Bruce W. Lee, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Erik Miehl, Pierre Dognin, Manish Nagireddy, and Amit Dhurandhar. 2025a. [Programming refusal with conditional activation steering](#). *Preprint*, arXiv:2409.05907.
- Sunbowen Lee, Qingyu Yin, Chak Tou Leong, Jialiang Zhang, Yicheng Gong, and Xiaoyu Shen. 2025b. [Probing the difficulty perception mechanism of large language models](#). *Preprint*, arXiv:2510.05969.
- Hongyu Li, Liang Ding, Meng Fang, and Dacheng Tao. 2024. [Revisiting catastrophic forgetting in large language model tuning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4297–4308, Miami, Florida, USA. Association for Computational Linguistics.
- Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhiwei Li, Bao-Long Bi, Ling-Rui Mei, Junfeng Fang, Xiao Liang, Zhijiang Guo, and 2 others. 2025. [From system 1 to system 2: A survey of reasoning large language models](#). *Preprint*, arXiv:2502.17419.
- Mengqi Liao, Xiangyu Xi, Ruinian Chen, Jia Leng, Yangen Hu, Ke Zeng, Shuai Liu, and Huaiyu Wan. 2025. [Enhancing efficiency and exploration in reinforcement learning for llms](#). *Preprint*, arXiv:2505.18573.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring how models mimic human falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Samuel Marks and Max Tegmark. 2023. [The geometry of truth: Emergent linear structure in large language model representations of true/false datasets](#). *ArXiv*, abs/2310.06824.
- Earl K. Miller and Jonathan D. Cohen. 2001. [An integrative theory of prefrontal cortex function](#). *Annual Review of Neuroscience*, 24(Volume 24, 2001):167–202.
- Deepak Babu Piskala, Vijay Raajaa, Sachin Mishra, and Bruno Bozza. 2024. [Optiroute dynamic llm routing and selection based on user preferences: Balancing performance, cost, and ethics](#). *International Journal of Computer Applications*, 186(51):1–7.
- Joris Postmus and Steven Abreu. 2024. [Steering large language models using conceptors: Improving addition-based activation engineering](#). In *MINT: Foundation Model Interventions*.
- 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, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *ArXiv*, abs/2311.12022.
- Nina Rimskey, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Turner. 2024. [Steering llama 2 via contrastive activation addition](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15504–15522, Bangkok, Thailand. Association for Computational Linguistics.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *arXiv preprint arXiv:2402.03300*.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. [Outrageously large neural networks: The sparsely-gated mixture-of-experts layer](#).
- Idan Shenfeld, Jyothish Pari, and Pulkit Agrawal. 2025. [RL’s razor: Why online reinforcement learning forgets less](#). *arXiv preprint arXiv:2509.04259*.

- Wenhao Shi, Zhiqiang Hu, Yi Bin, Junhua Liu, Yang Yang, See Kiong Ng, Lidong Bing, and Roy Ka-Wei Lee. 2024. Math-llava: Bootstrapping mathematical reasoning for multimodal large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4663–4680.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. 2022. [Challenging big-bench tasks and whether chain-of-thought can solve them](#). *Preprint*, arXiv:2210.09261.
- Daniel Tan, David Chanin, Aengus Lynch, Brooks Paige, Dimitrios Kanoulas, Adrià Garriga-Alonso, and Robert Kirk. 2024. [Analysing the generalisation and reliability of steering vectors](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 139179–139212. Curran Associates, Inc.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte Macdiarmid. 2023. Steering language models with activation engineering.
- Marco Valentino, Geonhee Kim, Dhairya Dalal, Zhixue Zhao, and André Freitas. 2025. Mitigating content effects on reasoning in language models through fine-grained activation steering. *arXiv preprint arXiv:2505.12189*.
- Constantin Venhoff, Iván Arcuschin, Philip Torr, Arthur Conmy, and Neel Nanda. 2025a. Base models know how to reason, thinking models learn when. *arXiv preprint arXiv:2510.07364*.
- Constantin Venhoff, Iván Arcuschin, Philip Torr, Arthur Conmy, and Neel Nanda. 2025b. [Understanding reasoning in thinking language models via steering vectors](#). In *Workshop on Reasoning and Planning for Large Language Models*.
- Mengru Wang, Ziwen Xu, Shengyu Mao, Shumin Deng, Zhaopeng Tu, Huajun Chen, and Ningyu Zhang. 2025a. [Beyond prompt engineering: Robust behavior control in LLMs via steering target atoms](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 23381–23399, Vienna, Austria. Association for Computational Linguistics.
- Ruonan Wang, Runxi Wang, Yunwen Shen, Chengfeng Wu, Qinglin Zhou, and Rohitash Chandra. 2025b. Evaluation of llms for mathematical problem solving. *arXiv preprint arXiv:2506.00309*.
- Xin Wang, Haoyang Li, Zeyang Zhang, Haibo Chen, and Wenwu Zhu. 2025c. [Modular machine learning: An indispensable path towards new-generation large language models](#). *CoRR*, abs/2504.20020.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhua Chen. 2025d. Mmlu-pro: a more robust and challenging multi-task language understanding benchmark. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NeurIPS '24*, Red Hook, NY, USA. Curran Associates Inc.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, and 1 others. 2022a. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022b. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Yuyang Wu, Yifei Wang, Ziyu Ye, Tianqi Du, Stefanie Jegelka, and Yisen Wang. 2025. [When more is less: Understanding chain-of-thought length in llms](#). *Preprint*, arXiv:2502.07266.
- Mohammad Zbeeb, Hasan Abed Al Kader Hammoud, and Bernard Ghanem. 2025. [Reasoning vectors: Transferring chain-of-thought capabilities via task arithmetic](#). *Preprint*, arXiv:2509.01363.
- Hanyu Zhang, Xiting Wang, Chengao Li, Xiang Ao, and Qing He. 2025a. [Controlling large language models through concept activation vectors](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(24):25851–25859.
- Yi-Kai Zhang, De-Chuan Zhan, and Han-Jia Ye. 2025b. [Capability instruction tuning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 39:25958–25966.

## A Case Study

Figure 7 presents a qualitative comparison where RISER successfully navigates conceptual traps. While the base model, CoT prompting, and even stronger frontier models are lured into selecting Bureaucracy by misleading lexical cues despite generating superficially coherent rationales,

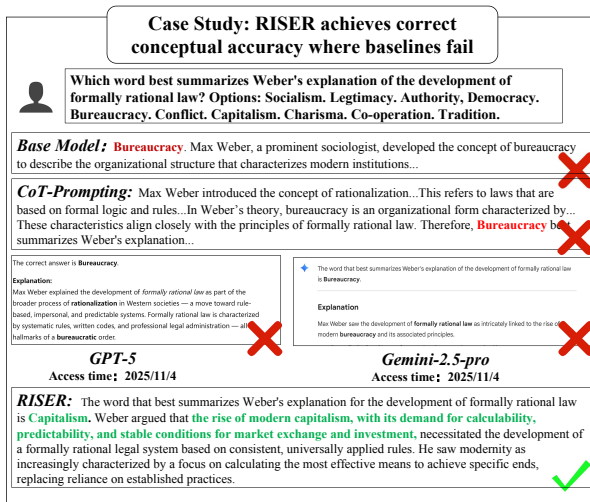


Figure 7: All baselines choose incorrectly, while RISER selects the correct answer and grounds its explanation in the rise of modern capitalism.

RISER correctly identifies that Capitalism aligns with Weber’s theory. Analysis of the routing weights reveals that RISER selectively amplifies logical reasoning and domain knowledge primitives while suppressing irrelevant directions, effectively steering the model’s focus.

## B Transferability Across Domains

We further evaluated RISER directly on the Big-Bench Hard (BBH) benchmark (Suzgun et al., 2022) without any additional fine-tuning. Despite the significant distribution shift moving from knowledge-centric exams to pure symbolic tasks, RISER outperformed the base model by 2.8%. Crucially, the Router autonomously adapted its strategy by prioritizing general-purpose primitives, specifically Logical Reasoning and Reading Comprehension. This transferability suggests that these primitives capture content-agnostic cognitive mechanisms akin to fluid intelligence and the system successfully identifies that the underlying computational demand remains constant even when the surface-level domain changes, verifying that RISER orchestrates genuine, transferable mental skills.

## C Transferability Across Models

To delineate the generalization boundaries of the RISER framework, we systematically investigate the transferability of learned Routers across distinct model families and varying parameter scales in Table ???. Specifically, we apply a Router trained on a source model directly to a target model without any additional tuning. This experiment aims to determine whether the learned cogni-

tive compositions capture universal reasoning patterns or remain specific to the internal representation space of the source architecture.

### C.1 Cross-Architecture Transferability

We first evaluate Router transferability between distinct model families, specifically exchanging Routers between Qwen2.5-7B and Llama-3-8B. Empirical results demonstrate negligible generalization, with performance regressing to near-random or baseline levels.

This failure suggests that the learned reasoning vectors and routing policies are inextricably coupled to the specific manifold of each model family. We attribute this incompatibility to three primary factors. First, stochastic pre-training induces manifold misalignment, where semantically similar concepts acquire arbitrary geometric orientations in high-dimensional space, precluding natural isometry between families without explicit alignment. Second, variations in pre-training corpora and tokenizers yield divergent activation statistics, causing source vectors to map onto low-density or undefined regions in the target manifold. Third, architectural inductive biases—arising from structural differences such as Grouped-Query Attention versus Multi-Head Attention—fundamentally reshape the activation landscape geometry, rendering direct vector transplantation mathematically invalid.

### C.2 Intra-Family Transferability: Scale Invariance

Conversely, transferring Routers within the same model family (e.g., across the Qwen2.5 series) yields substantial efficacy, indicating a shared semantic alignment across scales. We observe two distinct phenomena based on the direction of transfer.

**Large-to-Small Transfer (Inference-Time Distillation).** Transferring a Router from a larger model (e.g., 32B) to a smaller one (e.g., 7B) results in significant accuracy gains of up to +5%. We posit that the Router derived from the larger model encapsulates more precise and robust cognitive strategies. Deploying this advanced policy on a smaller model functions as a form of inference-time distillation, effectively guiding the smaller model to navigate complex reasoning pathways that it fails to autonomously discover due to limited capacity.

**Small-to-Large Transfer (Feature Consistency).** Transferring from smaller to larger models also confers meaningful improvements, typically enhancing accuracy by **+2–3%**. This finding demonstrates scalable feature consistency, implying that the fundamental cognitive directions identified in smaller parameter regimes remain preserved and refined in larger models. Consequently, the lightweight Router maintains its steering effectiveness even as the backbone capacity increases, highlighting the hierarchical stability of the learned representations within the same lineage.

## D Implementation & Training Details

### D.1 Hyperparameters

Table 5: Hyperparameters and implementation details for RISER training.

Hyperparameter	Value
<i>Model Architecture &amp; Optimization Strategy</i>	
Advantage Estimator	GRPO
Base LLM Trainable	False
Router Trainable	True
<i>Training Configuration</i>	
SFT Learning Rate	$5 \times 10^{-6}$
SFT Epochs	3
RL Learning Rate	$2 \times 10^{-6}$
Total RL Epochs	2
Global Batch Size	128
Max Context Length	8192
Data Shuffling	True
<i>Reward &amp; KL Divergence</i>	
Reward Type	Accuracy (0/1)
KL Loss Coefficient	0.001
KL Loss Type	Low Variance KL
<i>Rollout &amp; Generation</i>	
Temperature	1.5
Top-k	-1 (Disabled)
Do Sample	True
Number of Rollouts ( $N$ )	8
Max Batched Tokens	8192
<i>Infrastructure &amp; Parallelism</i>	
Tensor Parallel Size	4
GPUs per Node	4
Number of Nodes	1

In Table 5, we list the training configuration and used hyperparameters during our experiment. Full framework code will be released once ac-

cepted. Experiment is conducted on RTX5090 GPU. We report accuracy for each benchmark and average performance across task groups and results are computed from random seeds (average performance on 3 runs) and report absolute accuracy and relative improvements over baselines.

### D.2 Datasets Configuration

Our training pipeline consists of three distinct phases, each utilizing a specific data strategy to ensure the robustness and capabilities of the RISER framework.

#### Phase 1: Reasoning Vector Elicitation Data.

To construct a comprehensive library of cognitive primitives, we require a dataset that covers a broad spectrum of reasoning types. We constructed the elicitation dataset by performing **random sampling** of 500 examples from the **MMLU** benchmark (Hendrycks et al., 2021b). Given MMLU’s inherent breadth across diverse subjects, ranging from elementary mathematics to professional law, this random sampling strategy ensures that the extracted vectors cover a diverse spectrum of cognitive reasoning patterns without introducing domain-specific bias.

For each question, we generated several paired activation states using the *Positive* and *Negative* prompts (see Appendix G). To ensure quality, we applied the LLM-Judge filtering mechanism described in Section 4, retaining only pairs that exhibit a significant gap in reasoning rigor, paving the path for further experiments.

#### Phase 2: Router SFT Data (Oracle Label Synthesis).

For the supervised warm-up, we utilized a separate set of 200 MMLU samples (non-overlapping with the elicitation set). To synthesize the ground-truth "Oracle Labels" for training the Router, we employed a **constrained grid search** mechanism. For each query, we first ranked the primitives based on their individual efficacy and pre-selected the **top-2** candidates. Within this reduced subspace, we performed a fine-grained grid search over the intervention strength  $\alpha$ , discretizing the value with a step size of **0.1** (ranging from 0 to  $\alpha_{\max}$ ). The configuration  $(\mathbf{w}^*, \alpha^*)$  that elicited the correct response with the highest confidence was selected as the supervisory target. This approach efficiently provides high-quality initialization for the Router without the computational cost of an exhaustive search over the entire combinatorial space.

**Phase 3: RL Refinement Data.** To further optimize the Router for complex composition and generalization, we employed the **MMLU-Pro** benchmark. MMLU-Pro presents a significantly harder challenge with distractor options and complex reasoning chains, providing a steeper gradient for reinforcement learning compared to standard MMLU. We randomly split the dataset into a Training Set (70%) for the GRPO algorithm and a Held-out Set (30%) for validation. Importantly, we strictly ensured **no question overlap** between the RL training set and the final evaluation benchmarks (Table 1) to prevent data leakage and ensure fair evaluation.

## E RISER Inference Procedure

Algorithm 1 provides pseudocode for our inference pipeline: we invoke the Router once after the prompt prefill to compute an injection vector  $v_{\text{inject}}$ , and then reuse the same  $v_{\text{inject}}$  as an additive intervention at layer  $l$  for the last-token activation at every decoding step.

## F Sensitivity Analysis of Cluster Count ( $K$ )

We fixed the size of the cognitive primitive library at  $K = 6$  based on the observation that the first six principal components account for over 85% of the variance in the extracted difference vectors. To empirically validate this choice and assess the sensitivity of RISER to the granularity of the primitive library, we conducted an ablation study with varying cluster counts  $K \in \{4, 6, 8, 12\}$ . We evaluated these variants on Qwen2.5-7B-Instruct using three benchmarks that require distinct reasoning capabilities: GSM8K (Math), GPQA (General/Scientific), and TruthfulQA (Safety/Alignment). All other hyperparameters were held constant.

Primitives ( $K$ )	Variance Explanation (PCA)	GSM8K	GPQA	TruthfulQA	Avg.
$K = 4$	72.4%	83.1	33.5	56.2	57.6
$K = 6$ (Ours)	<b>86.1%</b>	<b>85.2</b>	<b>36.8</b>	59.8	<b>60.6</b>
$K = 8$	89.3%	85.0	36.4	<b>60.1</b>	60.5
$K = 12$	93.5%	84.6	35.9	59.5	60.0

Table 6: **Sensitivity Analysis of Cluster Count ( $K$ ).** We report the percentage of variance explained by the top- $K$  principal components and the zero-shot accuracy across three representative benchmarks.  $K = 6$  strikes the optimal balance between performance and model complexity.

The results, summarized in Table 6, demonstrate that  $K = 6$  is not merely an arbitrary choice but a local optimum for performance.

**Under-clustering ( $K = 4$ ):** Reducing the

---

## Algorithm 1 RISER inference

---

**Require:** Frozen LLM  $f_\theta$  ( $L$  layers), Router  $g_\phi$ , primitives  $V = \{v_i\}_{i=1}^K$ , layer index  $l$ , threshold  $\tau$ ,  $\alpha_{\max}$ , prompt tokens  $x$ , max steps  $T$

**Ensure:** Generated tokens  $y$

1: **Prefill (compute  $v_{\text{inject}}$  once):**

2:  $h_l \leftarrow \text{FORWARDTOLAYER}(f_\theta, x, l)$  ▷  
last-token state at layer  $l$

3:  $(p, \alpha) \leftarrow g_\phi(h_l)$  ▷  $p \in [0, 1]^K, \alpha \in \mathbb{R}^K$

4: **for**  $i = 1$  to  $K$  **do**

5:      $w_i \leftarrow \mathbb{I}[p_i > \tau]$

6:      $\alpha_i \leftarrow \text{clip}(\alpha_i, 0, \alpha_{\max})$

7: **end for**

8:  $v_{\text{inject}} \leftarrow \sum_{i=1}^K w_i \alpha_i v_i$

9: **Decoding (reuse the same  $v_{\text{inject}}$  at every step):**

10:  $y \leftarrow []$

11: **for**  $t = 1$  to  $T$  **do**

12:      $h_l^{(t)} \leftarrow \text{FORWARDTOLAYER}(f_\theta, x || y, l)$

13:      $\tilde{h}_l^{(t)} \leftarrow h_l^{(t)} + v_{\text{inject}}$

14:      $\text{logits}^{(t)}$  ←

$\text{CONTINUEFROMLAYER}(f_\theta, \tilde{h}_l^{(t)}, l)$

15:      $y_t \leftarrow \text{DECODETOKEN}(\text{logits}^{(t)})$

16:      $y \leftarrow y || [y_t]$

17:     **if**  $y_t$  is EOS **then**

18:         **break**

19:     **end if**

20: **end for**

21: **return**  $y$

---

number of primitives leads to a noticeable performance drop ( $-3.0\%$  average accuracy compared to  $K = 6$ ). With only 72.4% of the variance explained, distinct cognitive functions (e.g., *Numerical Calculation* and *Logical Reasoning*) are forced to merge into coarser centroids. This *semantic collision* reduces the precision of the steering vectors, preventing the Router from isolating the specific capability required for specialized tasks like GSM8K.

**Over-clustering ( $K = 8, 12$ ):** Increasing  $K$  beyond 6 yields diminishing returns. While the explained variance increases to 93.5% at  $K = 12$ , the downstream accuracy plateaus or slightly degrades. We attribute this to two factors:

(1) *Redundancy*: Higher  $K$  values introduce collinear vectors that represent fine-grained nuances rather than distinct skills, reducing the orthogonality of the library.

(2) *Optimization Difficulty*: A larger action space complicates the RL exploration process.

The Router struggles to distinguish between redundant vectors given the sparse reward signal, leading to less stable policies.

Consequently,  $K = 6$  provides the most robust trade-off, ensuring sufficient semantic coverage to handle diverse tasks while maintaining a compact and orthogonal action space for efficient Router learning.

## G Prompt for Reasoning Quality Isolation

To extract reasoning vectors that encode cognitive rigor rather than mere verbosity, we employ a Reasoning Fidelity Contrast strategy. The goal is to isolate the difference between **verified execution** and **plausible generation**, ensuring the extracted vector promotes efficiency by substituting internal computation for external token generation.

### G.1 Vector Elicitation Prompts

#### Positive

**Role:** You are a meticulous logician focused on absolute precision.  
**Task:** Derive the answer to the following question using proof-level rigor.  
**Instructions:**

- **Derive:** Do not just state facts; deduce them from axioms or given data.
- **Verify:** Check each intermediate calculation or logic jump for errors.
- **Precision:** Prioritize correctness over fluency. Reject any heuristic shortcuts.
- **Output:** Provide a sound explanation.

**Question:**  
 {{QUESTION}}  
**Rigorous Derivation:**

#### Negative

**Role:** You are a fluent conversationalist acting on "autopilot".  
**Task:** Provide a plausibly sounding answer based on surface-level associations.  
**Instructions:**

- **Flow:** Write whatever comes to mind first based on language patterns.
- **Approximate:** Do not perform actual calculations or verification. Use "ballpark" figures.
- **Plausibility:** The answer should *sound* correct to a layperson, even if the logic is flawed.
- **Output:** Generate a coherent but unverified response (simulate a hallucination if necessary).

**Question:**  
 {{QUESTION}}  
**Plausible Response:**

### G.2 LLM Judge Filtering Criteria

We employ an LLM Judge to ensure the vector subtraction captures the quality gap. The judge filters for pairs where the positive response is logically sound (Score  $> 80$ ) and the negative response lacks actual reasoning depth (Score  $< 20$ ),

while maintaining structural similarity.

## H Evaluation Protocols and Vector Definitions

### Evaluation Prompt

```
You are evaluating the cognitive integrity
and reasoning quality of two AI responses.
Question: {{QUESTION}}
Response A: {{RESPONSE_A}} Response B:
{{RESPONSE_B}}
Criteria:
1. Soundness (A): Does Response A contain a
valid, verified logical path to the correct
answer? (Score 0-100). High score means
rigorous logic.
2. Logical Validity (B): Does Response B
demonstrate actual step-by-step derivation?
(Score 0-100). Low score means it relies on
guessing, hallucinations, or surface-level
associations without real logic.
3. Structural Parity: Are both responses
roughly similar in length? (Pass/Fail).
Output JSON:
{
  "score_A_soundness": <int>,
  "score_B_validity": <int>,
  "structural_parity": <bool>,
  "valid_pair": <bool>
  // True ONLY if:
  // score_A_soundness > 80 AND
  // score_B_validity < 20 AND
  // structural_parity is True.
}
```

## I Further Implementation Details

### I.1 Exact Router Architecture and Parameter Computation

To ensure strict reproducibility, we detail the exact architecture of the lightweight Router. The Router is implemented as a bottleneck Multi-Layer Perceptron (MLP) with dual parallel output heads. Let  $d$  denote the hidden dimension of the target LLM and  $K$  denote the number of cognitive primitives in the library. The architecture is defined as follows:

- **Bottleneck Projection:** Linear(in =  $d$ , out = 1024) followed by a SiLU activation.
- **Feature Extraction:** Linear(in = 1024, out = 1024) followed by a SiLU activation.
- **Selection Head:** Linear(in = 1024, out =  $K$ ) followed by a Sigmoid activation, outputting probabilities  $\mathbf{p} \in [0, 1]^K$ .
- **Strength Head:** Linear(in = 1024, out =  $K$ ) followed by a bounded activation function (clamped between 0 and  $\alpha_{\max}$ ), outputting intervention strengths  $\alpha$ .

The total number of trainable parameters  $N_{\text{params}}$  is determined exactly by the formula:

$$N_{\text{params}} = 1024d + 2048K + 1,049,600 \quad (6)$$

For instance, when utilizing Qwen2.5-7B ( $d = 3584$ ) and our default library size ( $K = 6$ ),

Model	Benchmark	Base (%)	CoT (%)	RISER (Ours) (%)	$\Delta$ vs. Base
Qwen2.5-7B	AIME 2024	13.3 $\pm$ 2.1	16.7 $\pm$ 1.9	<b>20.0</b> $\pm$ 2.5	+6.7
Qwen2.5-7B	AIME 2025	11.1 $\pm$ 1.8	14.4 $\pm$ 2.0	<b>18.9</b> $\pm$ 1.2	+7.8
Llama-3-8B	AIME 2024	10.0 $\pm$ 1.9	13.3 $\pm$ 2.1	<b>16.7</b> $\pm$ 1.5	+6.7
Llama-3-8B	AIME 2025	6.7 $\pm$ 1.5	10.0 $\pm$ 1.8	<b>13.3</b> $\pm$ 1.6	+6.6

Table 7: Exact-match pass@1 accuracy on AIME 2024 and AIME 2025 (mean  $\pm$  std across 3 seeds). Results confirm significant absolute improvements on olympiad-level reasoning.

Target Layer (Llama-3-8B)	MATH	GSM8K	GPQA	Avg Acc (%)	$\Delta$ vs. Base
Base Model (No Intervention)	30.9	84.5	25.8	47.07	-
Layer 12	33.8	87.6	29.4	50.27	+3.20
<b>Layer 13 (Optimal)</b>	<b>35.4</b>	<b>89.1</b>	<b>30.9</b>	<b>51.80</b>	<b>+4.73</b>
Layer 14	34.8	88.5	30.3	51.20	+4.13
Layer 15	35.1	88.9	30.7	51.57	+4.50
Layer 16	34.5	88.2	29.8	50.83	+3.76
Layer 17	34.0	87.8	29.2	50.33	+3.26

Table 8: Target-layer sensitivity sweep for Llama-3-8B-Instruct. A stable performance plateau is observed across layers 13 through 15.

the Router contains exactly 4,732,940 parameters ( $\sim 4.7M$ ), representing less than 0.1% of the base model’s parameter count.

## I.2 Direct GRPO Baseline Configuration

To establish a rigorous and fair upper bound for the RL phase, the *Direct GRPO* baseline is configured to strictly mirror RISER’s RL infrastructure. Both methods share the exact same generation settings (Temperature = 1.5, do\_sample=True, 8 rollouts per prompt) and identical RL hyperparameters (0/1 accuracy reward, Low-Variance KL penalty with a coefficient of 0.001, global batch size of 128, and a tensor parallel size of 4).

The critical distinction lies in the trainable scope and computational budget: while RISER freezes the foundation model and trains only the 5M-parameter Router for 2 epochs, the Direct GRPO baseline updates **100% of the base LLM’s parameters** and is granted an extended training budget of 5 RL epochs.

## J Comprehensive Stability and Sensitivity Analyses

### J.1 Layer Sensitivity Sweep

To validate the existence of a stable reasoning plateau, we conducted a systematic target-layer sweep on Llama-3-8B-Instruct. As shown in Table 8, performance gains are robust across the early-middle layers (L13–L15), demonstrating that the framework is not hypersensitive to the exact layer index.

### J.2 Primitive Library Sampling Stability

To ensure  $N = 500$  elicitation samples are sufficient for the offline discovery of reasoning directions, we performed a stability analysis by aligning cluster centroids across different sampling scales using Hungarian matching. Table 9 confirms that the principal directions fully converge at  $N = 500$ , with negligible geometric deviation when scaled to 1000.

### J.3 Gating Threshold ( $\tau$ ) Sensitivity

The gating threshold  $\tau$  balances precision and coverage. We empirically swept  $\tau \in [0.5, 0.9]$  on Qwen2.5-7B (Table 10). The optimal balance is achieved at  $\tau = 0.7$ , with stable performance maintained within the  $[0.6, 0.8]$  bracket.

### J.4 Human Audit of LLM Judge

To validate the reliability of the Claude-3.5-Sonnet judge utilized during the elicitation phase, we conducted a rigorous human audit. We randomly sampled 50 paired responses (100 total) and tasked a human expert to independently score them on the same 0 – 100 logical validity scale. We defined agreement based on a strict tolerance margin ( $|\Delta \text{Score}| \leq 10$ ). The audit yielded an **82% agreement rate**, confirming that the LLM Judge aligns with human experts with high fidelity, reliably isolating rigorous logic from plausible hallucinations.

### J.5 Definitions of the Seven Primitives

The original library consists of  $K = 6$  core reasoning primitives:

Comparison Scale	Mean Cosine Sim.	Min Cosine Sim.	Max Cosine Sim.
$N = 200$ vs. $N = 500$	0.864	0.782	0.941
$N = 500$ vs. $N = 1000$	0.952	0.915	0.983

Table 9: Cosine similarity of matched reasoning primitives across different sample sizes ( $N$ ).

Gating Threshold ( $\tau$ )	MATH (%)	GPQA (%)	Observation
0.5	51.8	35.1	Noise Interference (Activates irrelevant vectors)
0.6	52.9	36.2	Stable Performance
0.7	<b>53.3</b>	<b>36.8</b>	<b>Optimal Balance</b>
0.8	53.0	36.5	Stable Performance
0.9	53.2	34.6	Under-intervention (Vectors rarely activated)

Table 10: Sensitivity analysis of the gating threshold  $\tau$  on Qwen2.5-7B.

1. **Numerical Calculation:** Specialized for precise arithmetic operations and quantitative tracking.
2. **Logical Inference:** Governs step-by-step deductive reasoning and constraint satisfaction.
3. **Ethical Alignment:** Activates safety protocols and aligns outputs with human moral values.
4. **Reading Comprehension:** Enhances information extraction and context synthesis from long texts.
5. **Scientific Analysis:** Engages domain-specific scientific facts and empirical reasoning.
6. **Domain-specific Knowledge:** Recalls specialized terminology and encyclopedic knowledge.

For the extensibility experiment in Section 5.4, a **7th primitive (Code Synthesis / Algorithmic Constraints)** was appended. This vector was elicited using prompt pairs contrasting structured algorithmic control flows against pseudo-code approximations, specifically steering the model towards syntax-compliant generation.

## J.6 Modern Math Benchmarks (AIME)

### Evaluation

To test the upper bounds of mathematical reasoning, we evaluated the framework on the AIME 2024 and 2025 datasets. Due to the extreme complexity of AIME, we utilized a strict, reproducible protocol: greedy decoding (Temperature = 0) for strict pass@1 accuracy, with `max_generation_length = 8192`. As shown in Table 7, RISER consistently guides the model to solve an additional  $\sim 2$  problems on average (a +6.6% to +7.8% absolute gain) without task-specific SFT.