

# Mitigating Overthinking via Interleaved System-1/2 Reasoning in Large Language Models

Han Ding<sup>1</sup>, Wei Chen<sup>1</sup>, Fuzhen Zhuang<sup>1,2\*</sup>

{handing, chenwei23, zhuangfuzhen}@buaa.edu.cn

<sup>1</sup>Institute of Artificial Intelligence, Beihang University, Beijing, China

<sup>2</sup>Zhongguancun Laboratory, Beijing, China

## Abstract

Large reasoning models improve problem-solving performance by generating extended Chain-of-Thought, but often suffer from overthinking, producing unnecessarily long and redundant reasoning traces. Prior works mainly focus on directly compressing reasoning length, or selecting reasoning modes globally based on task difficulty, without accounting for the dynamic and heterogeneous nature of the reasoning process itself. This lack of consideration for inference-time diversity limits their ability to adapt reasoning strategies flexibly. To address this gap, we propose Interleaved System-1/2 Reasoning (**InterSR**), a cognitively inspired framework that dynamically alternates between System 1 (fast) and System 2 (slow) reasoning modes during generation. Central to InterSR is a learnable system switching signal that reflects the necessity for deeper reasoning at each intermediate step and enables fine-grained control over cognitive mode transitions. Experiments across mathematical and cross-domain benchmarks show that InterSR effectively captures internal reasoning signals to reduce token usage while maintaining strong and robust accuracy. We release our code at InterSR<sup>1</sup>.

## 1 Introduction

Large Reasoning Model (LRM) has demonstrated a significant performance improvement over instruct models by introducing long Chain-of-Thought (CoT) during test-time scaling to enhance the model’s problem-solving ability (Guo et al., 2025; OpenAI et al., 2024; Yang et al., 2025a). This extended CoT allows the model to thoroughly consider the problem, decompose complex issues, explore various possibilities, and validate the correctness of its outputs (Feng et al., 2023; Wei et al., 2022; Sprague et al., 2024). However, this long




Figure 1: Illustration of InterSR. Standard reasoning applies uniform System 2 (slow) thinking throughout, leading to overthinking on simple steps. InterSR dynamically switches between System 1 (fast) and System 2 reasoning based on step-level complexity, reducing redundant computation while preserving accuracy.

CoT has been found unsuitable for all types of problems. For simple problems (such as “2+3”), long CoT thinking is regarded as an overthinking phenomenon (Chen et al., 2025c; Cuadron et al., 2025; Ghosal et al., 2025; Zhou et al., 2026). This phenomenon is primarily characterized by the model repeatedly verifying the correct answer or engaging in meaningless reasoning (Fu et al., 2025; Wang et al., 2025; Ballon et al., 2025; Chen et al., 2026a). This not only reduces computational efficiency and leads to resource waste, but may also negatively impact the response accuracy (Du et al., 2026).

Recent works have proposed several approaches to mitigate overthinking (Feng et al., 2025; Sui et al., 2025; Su et al., 2025; Bin et al., 2026). **Response-side** methods reduce token usage by compressing or truncating the reasoning output (Chen et al., 2025c; Fu et al., 2025; Yang et al., 2025b; Hou et al., 2025), but risk removing critical reasoning steps. **Inference-side** methods classify problems by difficulty and select reasoning modes accordingly (Pan et al., 2024; Qi et al., 2024; Yao et al., 2024), but oversimplify complexity by treating each problem uniformly—ignoring that indi-

\*Corresponding author.

<sup>1</sup><https://github.com/buaadh/InterSR>

vidual reasoning steps within the same problem can vary dramatically in difficulty. These coarse-grained approaches either sacrifice reasoning quality for efficiency or fail to capture the fine-grained heterogeneity within a single problem, ultimately limiting both effective token savings and stable accuracy preservation (Zhang et al., 2026).

Considering these limitations, we propose InterSR, an interleaved system reasoning approach that addresses overthinking at step-level granularity. We model the problem as minimizing computational cost while preserving information gain at each reasoning step. Using semantic entropy (Kuhn et al., 2023; Cheng et al., 2025) as a proxy for information gain, we construct a step-aware system signal that enables dynamic switching between fast and slow reasoning modes at each step.

Our contributions are threefold: (1) We propose InterSR, which dynamically switches between reasoning systems based on a constructed system signal to improve inference allocation and reduce overthinking. (2) We demonstrate that this signal is encoded in the model’s hidden states, and train a lightweight predictor to enable near-zero-cost switching during generation. (3) We evaluate InterSR on four math benchmarks, achieving up to 51% token reduction without accuracy loss, and further validate its generalizability across STEM, knowledge-intensive, and code task categories.

## 2 Background

### 2.1 Overthinking

Despite the benefits of explicit reasoning, large language models (LLMs) often produce unnecessarily long or redundant chains of thought, even for tasks that can be solved with minimal or no reasoning. This phenomenon, known as **overthinking** (Chen et al., 2025c; Sui et al., 2025), reflects inefficient allocation of inference resources and is commonly observed across many CoT-style models (Pan et al., 2024; Zhang et al., 2025). Overthinking in large language models (LLMs) refers to the production of unnecessarily long reasoning traces that do not substantially improve task performance. It can be formalized as follows:

$$\exists r_{\text{short}} \in \mathcal{R} \text{ s.t. } \begin{cases} P(y \mid x, r_{\text{long}}) \approx P(y \mid x, r_{\text{short}}), \\ \mathcal{C}(r_{\text{long}}) \gg \mathcal{C}(r_{\text{short}}) \end{cases} \quad (1)$$

where  $r_{\text{long}}$  denotes an over-elaborate reasoning sequence, while  $r_{\text{short}}$  is a more efficient alternative

that achieves comparable predictive performance.  $\mathcal{C}(\cdot)$  measures reasoning cost, such as token usage or inference latency. This inefficiency typically arises when the model continues generating reasoning content with low marginal utility, for example generating unnecessary intermediate steps.

### 2.2 Dual-System Theory

Dual-System Theory (Evans, 2017; Chaiken and Trope, 1999) posits that human reasoning operates through two distinct but interacting systems: **System 1** (fast, intuitive) and **System 2** (slow, deliberative). These systems interact dynamically, with System 2 overriding System 1 when deeper reasoning is required for complex decisions (Shi et al., 2025; Shulman et al., 2016).

Recent LLM reasoning research has adopted this framework (Zhang et al., 2026; Saha et al., 2024; Ziabari et al., 2025; Coda-Forno et al., 2025), mapping System 1 to efficient low-cost inference and System 2 to resource-intensive deliberation (Chen et al., 2025b,a). However, existing methods typically select reasoning modes globally based on overall task difficulty (Qi et al., 2024; Zhao et al., 2025), neglecting the compositional nature of reasoning, where different intermediate steps within the same problem demand varying cognitive effort.

## 3 Method

In this section, we present the construction process of InterSR. We start by formally modeling the overthinking problem and emphasize the key role of information gain. Then, we introduce a semantic-entropy-based system signal and reveal its relationship with internal hidden states. Building upon these findings, we develop the complete Interleaved System-1/2 Reasoning Framework.

### 3.1 Stepwise Modeling of the Reasoning Process

We view the reasoning process as a sequence of explicit steps  $r = \{r_1, r_2, \dots, r_T\}$ , where each step  $r_t$  is dynamically governed by a reasoning mode  $m_t \in \{S_1, S_2\}$ . System 1 provides fast, low-cost, heuristic reasoning, while System 2 offers slow, deliberative, high-cost inference. This decomposition allows each step to flexibly adopt the cognitive mode suited to its semantic demands, enabling fine-grained control over the reasoning trajectory. Unlike prior approaches that statically assign a single system to the entire process, we




Figure 2: t-SNE projection of hidden representations from DeepSeek-R1-Distill-Qwen-14B, with points colored by switching-signal strength. The projection suggests that hidden states contain structure related to  $\sigma$ .

model reasoning as an interleaved sequence of fast and slow modes. The model generates each step as  $r_t \sim P_\theta(r_t | x, r_{<t}, m_t)$ , where the mode sequence  $\{m_t\}$  reflects an adaptive cognitive schedule over time. We assume that, all else being equal, allocating more token resources to a reasoning step can increase the step’s epistemic gain. Under this assumption, we invoke System 2 only when necessary to maintain semantic precision while minimizing total computation. This idea can be expressed as the following constrained optimization:

$$\begin{aligned} \min_{\{z_t\}} \quad & \sum_{t=1}^T \mathcal{C}(z_t) \\ \text{subject to} \quad & \begin{cases} \sum_{t=1}^T \mathcal{G}(s_t, z_t) = \sum_{t=1}^T \mathcal{G}(s_t, 1), \\ z_t \in \{0, 1\}, \quad t = 1, \dots, T \end{cases} \end{aligned} \quad (2)$$

where  $z_t$  denotes system choice (0 for  $S_1$ , 1 for  $S_2$ ),  $\mathcal{C}(z_t)$  represents system cost, and  $\mathcal{G}(s_t, z_t)$  captures semantic gain. The goal is to match the epistemic quality of full System 2 reasoning while reducing overall cost through selective system switching.

### 3.2 Constructing Signals Based on Semantic Entropy

While System 1 is effective when the semantic exploration space of a reasoning step is relatively simple, it fails to extract full epistemic value once the underlying inference space becomes sufficiently complex or ambiguous. In contrast, System 2, though more computationally expensive, can better leverage such latent structure and thus achieve




Figure 3: Signal strength distribution on DART-Math by problem difficulty (R1-14B). Harder problems exhibit a higher proportion of High-signal steps, indicating greater System 2 demand.

higher information gain. Based on the above modeling, System 1 can efficiently process reasoning steps within simple inference spaces, but fails to capture full epistemic value when a step involves complex or ambiguous semantic structures. To maintain quality equivalent to full System 2 reasoning, we switch to System 2 whenever System 1 cannot extract sufficient information:

$$z_t = \begin{cases} 0, & \text{if } \mathcal{G}(s_t, 0) \geq T, \\ 1, & \text{otherwise,} \end{cases} \quad (3)$$

where  $z_t \in \{0, 1\}$  indicates the system choice at step  $t$  (0 for System 1, 1 for System 2),  $\mathcal{G}(s_t, 0)$  measures the semantic gain attainable by System 1 at step  $s_t$ , and  $T$  is a threshold defining sufficient gain. When the gain under System 1 falls below  $T$ , System 2 is invoked to ensure comprehensive reasoning. Against this backdrop, we require an estimate of the exploration richness at each step. Drawing on prior work linking model’s exploration space to semantic diversity (Cheng et al., 2025), we adopt semantic entropy (Kuhn et al., 2023) as a proxy. Specifically, we define a continuous switching signal  $\sigma_t \in [0, 1]$  as the normalized entropy of the model’s predicted continuation space:

$$\sigma_t = \frac{\mathcal{H}(s_t)}{\log_2 K}, \quad (4)$$

where  $\mathcal{H}(s_t)$  denotes the entropy over  $K$  sampled next-token distributions at step  $t$ . This signal enables dynamic cognitive resource allocation pro-




Figure 4: Overview of the InterSR framework. The signal prediction module extracts hidden states at sentence boundaries and predicts a system switching signal  $\hat{\sigma}_t$ . When  $\hat{\sigma}_t$  exceeds the threshold  $\tau$ , a control token is inserted to switch between System 1 and System 2 reasoning modes, enabling fine-grained cognitive resource allocation during inference.

portional to semantic uncertainty. As shown in Figure 3, the signal varies systematically across problem difficulty levels on DART-Math (Tong et al., 2024): harder problems (level 5) exhibit a higher proportion of High-signal steps, indicating greater System 2 demand during reasoning.

### 3.3 Capturing System Switching Signals

While the signal  $\sigma_t$  is constructed from semantic entropy, computing it online at each candidate boundary would require sampling multiple continuations and clustering their meanings, introducing substantial overhead during inference. We therefore investigate whether the information reflected by this signal is already encoded in the model’s internal states. As shown in Figure 2, hidden representations exhibit structure related to high and low signal strength regions, suggesting that the model’s activations contain information about local reasoning demand. This motivates us to learn the switching signal directly from internal activations rather than recomputing semantic entropy during decoding. We introduce a lightweight module  $g_\phi$  to predict signal strength from hidden states during reasoning, which allows InterSR to perform online mode switching with limited additional cost. Additional quantitative analyses of routed S1/S2

boundaries and out-of-domain predictor AUC are provided in Appendix N.

### 3.4 Interleaved System-1/2 Reasoning

Based on the above findings, we propose an Interleaved System-1/2 Reasoning framework that enables language models to dynamically alternate between System 1 and System 2 thinking modes during generation for efficient and context-sensitive reasoning. The reasoning model treats `<think>` as the start of a slow-thinking segment and `</think>` as its end (Guo et al., 2025; Yang et al., 2025a), enabling explicit control over cognitive mode transitions within a single inference trajectory. At each generation step, the model first performs a forward pass to compute the hidden representation  $h_k$  based on the current context. To avoid disrupting output fluency, a lightweight boundary detector  $f_\psi(h_k)$  identifies candidate sentence or reasoning-unit boundaries. Only at these boundaries does the model predict a switching signal  $\hat{\sigma}_k$  from  $h_k$  using a learned predictor  $g_\phi$  and compare it with a threshold  $\tau$ . If the signal exceeds  $\tau$ , a control token (`<think>` or `</think>`) is inserted to trigger a transition between System 1 and System 2. This ensures that mode switching occurs at semantically coherent regions, reducing token consumption while pre-

| Method                               | GSM8K              |                   | MATH500            |                    | AIME24             |                    | AMC23              |                    |
|--------------------------------------|--------------------|-------------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
|                                      | Pass@1 (↑,%)       | Tokens (↓)        | Pass@1 (↑,%)       | Tokens (↓)         | Pass@1 (↑,%)       | Tokens (↓)         | Pass@1 (↑,%)       | Tokens (↓)         |
| <i>DeepSeek-R1-Distill-Qwen-1.5B</i> |                    |                   |                    |                    |                    |                    |                    |                    |
| Vanilla                              | 77.5               | 1265              | 84.8               | 4785               | 30.0               | 12496              | 72.5               | 7765               |
| TALE                                 | 77.4 (−0.1)        | 1056 (−16%)       | 80.6 (−4.2)        | 3614 (−24%)        | 30.0 (+0.0)        | 11890 (−4%)        | 70.0 (−2.5)        | 7267 (−6%)         |
| Dynasor                              | 77.6 (+0.1)        | 1030 (−18%)       | 85.2 (+0.4)        | 3624 (−24%)        | 30.0 (+0.0)        | 11111 (−11%)       | 72.5 (+0.0)        | 6781 (−12%)        |
| DEER                                 | 73.6 (−3.9)        | 820 (−35%)        | 80.6 (−4.2)        | 3770 (−21%)        | 26.7 (−3.3)        | 11845 (−5%)        | 67.5 (−5.0)        | 5958 (−23%)        |
| <b>Ours</b>                          | <b>78.0 (+0.5)</b> | <b>614 (−51%)</b> | <b>86.0 (+1.2)</b> | <b>3504 (−26%)</b> | <b>33.3 (+3.3)</b> | <b>8061 (−35%)</b> | <b>72.5 (+0.0)</b> | <b>5922 (−23%)</b> |
| <i>DeepSeek-R1-Distill-Qwen-7B</i>   |                    |                   |                    |                    |                    |                    |                    |                    |
| Vanilla                              | 86.4               | 743               | 88.6               | 3824               | 46.7               | 10953              | 90.0               | 5699               |
| TALE                                 | 84.3 (−2.1)        | 633 (−14%)        | 88.0 (−0.6)        | 2941 (−23%)        | 50.0 (+3.3)        | 9640 (−11%)        | 87.5 (−2.5)        | 4649 (−18%)        |
| Dynasor                              | 86.5 (+0.1)        | 715 (−3%)         | 88.2 (−0.4)        | 2931 (−23%)        | 50.0 (+3.3)        | 9532 (−12%)        | 90.0 (+0.0)        | 4910 (−13%)        |
| DEER                                 | 82.9 (−3.5)        | 712 (−4%)         | 84.4 (−4.2)        | 3020 (−21%)        | 43.3 (−3.4)        | 9871 (−10%)        | 85.0 (−5.0)        | 5546 (−3%)         |
| <b>Ours</b>                          | <b>86.7 (+0.3)</b> | <b>517 (−30%)</b> | <b>89.6 (+1.0)</b> | <b>2651 (−30%)</b> | <b>50.0 (+3.3)</b> | <b>7259 (−33%)</b> | <b>90.0 (+0.0)</b> | <b>4126 (−27%)</b> |
| <i>DeepSeek-R1-Distill-Llama-8B</i>  |                    |                   |                    |                    |                    |                    |                    |                    |
| Vanilla                              | <b>84.7</b>        | 982               | <b>70.8</b>        | 3677               | 36.7               | 10842              | 80.0               | 5518               |
| TALE                                 | 83.2 (−1.5)        | 839 (−14%)        | 68.4 (−2.4)        | 2905 (−21%)        | 33.3 (−3.4)        | 9737 (−10%)        | 77.5 (−2.5)        | 4853 (−12%)        |
| Dynasor                              | 84.5 (−0.2)        | 938 (−4%)         | 70.6 (−0.2)        | 2868 (−22%)        | 36.7 (+0.0)        | 9573 (−12%)        | 80.0 (+0.0)        | 4855 (−12%)        |
| DEER                                 | 81.7 (−3.0)        | 807 (−18%)        | 67.4 (−3.4)        | 3350 (−9%)         | 33.3 (−3.4)        | 10218 (−6%)        | 75.0 (−5.0)        | 5107 (−7%)         |
| <b>Ours</b>                          | <b>84.0 (−0.7)</b> | <b>731 (−26%)</b> | <b>70.2 (−0.6)</b> | <b>2800 (−24%)</b> | <b>36.7 (+0.0)</b> | <b>7806 (−28%)</b> | <b>80.0 (+0.0)</b> | <b>4138 (−25%)</b> |
| <i>DeepSeek-R1-Distill-Qwen-14B</i>  |                    |                   |                    |                    |                    |                    |                    |                    |
| Vanilla                              | 91.9               | 705               | 93.2               | 3591               | 63.3               | 9686               | 92.5               | 5435               |
| TALE                                 | 91.1 (−0.8)        | 553 (−21%)        | 91.8 (−1.4)        | 2941 (−18%)        | 60.0 (−3.3)        | <b>8450 (−12%)</b> | 90.0 (−2.5)        | 4011 (−26%)        |
| Dynasor                              | 91.6 (−0.3)        | 663 (−5%)         | 93.0 (−0.2)        | 3317 (−7%)         | 63.3 (+0.0)        | 9054 (−6%)         | 92.5 (+0.0)        | 4023 (−25%)        |
| DEER                                 | 88.5 (−3.4)        | 610 (−13%)        | 89.2 (−4.0)        | 2850 (−21%)        | 60.0 (−3.3)        | 8871 (−8%)         | 87.5 (−5.0)        | 4280 (−21%)        |
| <b>Ours</b>                          | <b>92.0 (+0.1)</b> | <b>544 (−22%)</b> | <b>93.2 (+0.0)</b> | <b>2494 (−30%)</b> | <b>63.3 (+0.0)</b> | 8557 (−11%)        | <b>92.5 (+0.0)</b> | <b>3894 (−28%)</b> |

Table 1: Performance comparison across reasoning strategies and datasets.

| Metric              | Vanilla | InterSR | Change |
|---------------------|---------|---------|--------|
| Avg. tokens/request | 3824    | 2651    | −30.7% |
| Peak KV memory (GB) | 12.4    | 8.7     | −29.8% |
| Throughput (req/s)  | 4.2     | 5.8     | +38.1% |
| TTFT (ms)           | 142     | 145     | +2.1%  |

Table 2: Serving efficiency on MATH500 (7B model, batch 32, 4×H200). InterSR’s token reduction translates to substantial throughput gains.

serving reasoning performance. The full algorithm can be found in Appendix C.

## 4 Experiments

In this section, we thoroughly evaluate the performance of our proposed Interleaved System Reasoning framework by systematically addressing the following key research questions: **RQ1:** Can InterSR reduce the number of tokens generated by the model without sacrificing accuracy? **RQ2:** How does InterSR perform on datasets outside the mathematical reasoning domain? **RQ3:** Which hidden layer best captures the switching signal for prediction? **RQ4:** How do different thresholds for the

system switching signal affect the model’s performance and efficiency? **RQ5:** Does InterSR allocate deliberative tokens according to benchmark difficulty? **RQ6:** How can InterSR’s ability to switch between systems for efficient reasoning be visually demonstrated and interpreted?

### 4.1 Experimental Settings

**Benchmarks.** We evaluate on four mathematical reasoning benchmarks spanning a wide range of difficulty: GSM8K (Cobbe et al., 2021) (grade-school arithmetic), MATH500 (Lightman et al., 2024) (competition-level problems), AMC23 (AMC competition), and AIME24 (AIME competition). These benchmarks require varied reasoning strategies from simple arithmetic to multi-step algebraic and geometric inference, enabling comprehensive evaluation of InterSR across difficulty levels.

**Baselines.** We compare against three representative adaptive reasoning strategies. TALE (Han et al., 2025) allocates a token budget based on problem complexity; we use the training-free TALE-EP variant for fair comparison. Dynasor (Fu et al., 2025) performs confidence-based early exit via Cer-

| Domain           | Dataset   | Model   | Accuracy (%) |         |      | Tokens  |         |           |
|------------------|-----------|---------|--------------|---------|------|---------|---------|-----------|
|                  |           |         | Vanilla      | InterSR | Diff | Vanilla | InterSR | Reduction |
| <i>STEM</i>      | GPQA-D    | R1-1.5B | 26.8         | 28.3    | +1.5 | 8777    | 4732    | −46%      |
|                  |           | R1-7B   | 46.5         | 47.5    | +1.0 | 6751    | 3891    | −42%      |
|                  |           | R1-14B  | 50.5         | 51.0    | +0.5 | 5899    | 3536    | −40%      |
|                  | ARC-C     | R1-1.5B | 57.9         | 57.8    | −0.1 | 1164    | 1013    | −13%      |
|                  |           | R1-7B   | 84.1         | 83.6    | −0.5 | 1025    | 789     | −23%      |
|                  |           | R1-14B  | 93.5         | 92.6    | −0.9 | 705     | 515     | −27%      |
| <i>Knowledge</i> | MMLU      | R1-1.5B | 45.8         | 49.0    | +3.2 | 1432    | 1274    | −11%      |
|                  |           | R1-7B   | 84.4         | 83.6    | −0.8 | 1181    | 892     | −24%      |
|                  |           | R1-14B  | 84.4         | 83.8    | −0.6 | 1170    | 772     | −34%      |
| <i>Code</i>      | HumanEval | R1-1.5B | 29.9         | 30.5    | +0.6 | 1851    | 1420    | −23%      |
|                  |           | R1-7B   | 43.9         | 45.7    | +1.8 | 1807    | 1128    | −37%      |
|                  |           | R1-14B  | 49.4         | 51.8    | +2.4 | 1774    | 1087    | −39%      |

Table 3: Cross-domain evaluation. **Diff**: InterSR − Vanilla accuracy change. **Reduction**: token reduction ratio.

taindex. DEER (Yang et al., 2025b) uses answer triggers to halt generation once the model reaches high confidence.

**Experimental Models.** We evaluate InterSR on four model scales: DeepSeek-R1-Distill-Qwen-1.5B/7B/14B and DeepSeek-R1-Distill-Llama-8B (Guo et al., 2025). The signal predictor  $g_\phi$  and boundary detector  $f_\psi$  are lightweight MLPs operating on the final hidden state  $h_t$ , trained on  $\sim 10k$  samples from DART-Math (Tong et al., 2024). Due to space constraints, full details on data construction, training procedures, and hyperparameters are provided in Appendix A.

## 4.2 Main Result (RQ1)

As shown in Table 1, our approach consistently reduces token usage across all datasets without sacrificing performance. On GSM8K, InterSR achieves the largest savings, with up to 51% token reduction for the 1.5B model, since simpler arithmetic problems contain more redundant reasoning steps that can be efficiently handled by System 1. On harder benchmarks such as AIME24, savings are more modest, averaging about 27%, which reflects the genuine need for deeper reasoning on competition-level problems. Across all four model scales, InterSR either preserves or slightly improves accuracy. For example, the 7B model improves from 88.6% to 89.6% on MATH500 while reducing average tokens from 3824 to 2651. Compared to TALE, Dynasor, and DEER, InterSR achieves a stronger efficiency–accuracy trade-off, since the

baselines either provide limited token savings on difficult benchmarks or suffer accuracy drops when reasoning is shortened too aggressively. InterSR also yields practical serving gains. On MATH500 with the 7B model, throughput improves by 38% with about 30% less KV cache memory (Table 2), showing that token reduction can translate into real deployment benefits. These results indicate that InterSR mitigates overthinking by reducing redundant reasoning while preserving the slow reasoning needed for accuracy. Details on computational overhead are in Appendix B.

## 4.3 Cross-domain Experiments (RQ2)

To investigate generalizability beyond mathematical reasoning, we evaluate InterSR on four diverse benchmarks: GPQA-Diamond (Rein et al., 2024) (graduate-level STEM), ARC-Challenge (Clark et al., 2018) (science reasoning), MMLU (Hendrycks et al., 2021) (knowledge-intensive QA), and HumanEval (Chen et al., 2021) (code generation). As shown in Table 3, InterSR achieves consistent token reduction across all benchmarks and model sizes while maintaining comparable or improved accuracy. The largest savings appear on GPQA-Diamond, where the 1.5B model reduces tokens by 46% with a +1.5% accuracy gain, indicating that graduate-level STEM reasoning contains substantial redundancy that InterSR effectively eliminates. On knowledge-intensive tasks (MMLU), InterSR maintains accuracy within 1% while achieving 11%–34% token reduction. Notably, on HumanEval, InterSR improves accuracy




Figure 5: AUC of signal predictors trained on different hidden layers using the 7B model. Performance increases monotonically with depth.




Figure 6: Efficiency versus accuracy Pareto front on MATH500. Each curve traces one model across threshold values  $\tau = 0.0$  to 0.6.

across all model sizes while reducing tokens by 23%–39%, suggesting that suppressing overthinking can benefit code generation by avoiding unnecessary re-derivation that introduces bugs. These results confirm that InterSR’s switching signals, though trained on mathematical data, generalize effectively across diverse task domains, with stronger savings on tasks that contain longer intermediate reasoning traces.

#### 4.4 Hidden Layer Selection for Signal Prediction (RQ3)

The signal predictor  $g_\phi$  operates on hidden states from a specific transformer layer. To identify the optimal layer, we use the 7B model (28 layers) and train separate predictors on layers 1, 7, 14, 20, 24, 27, and 28 with identical architectures.

As shown in Figure 5, AUC increases monotonically from layer 1 (0.641) to the final layer 28 (0.838). Early layers yield weak predictions, as shallow representations primarily encode local syn-




Figure 7: Average System 1/2 token allocation for R1-7B. Bars show proportions; right labels show average total tokens.

tactic features insufficient for estimating reasoning complexity. The final layers achieve the strongest performance, confirming that deeper representations encode richer uncertainty information relevant to system switching. We therefore use the final layer for signal prediction across all model sizes. Additional ROC analysis is in Appendix E.

#### 4.5 Efficiency–Accuracy Trade-off (RQ4)

We sweep the switching threshold  $\tau$  from 0.0 (vanilla) to 0.6 and plot accuracy against token count for each model on MATH500 (Figure 6). For all models, there exists an intermediate range of thresholds where tokens decrease substantially while accuracy is preserved or improved—the 14B model maintains 93.2% accuracy from  $\tau = 0.0$  to  $\tau = 0.4$  while reducing tokens by 30.5%. Notably, moderate thresholds can *improve* accuracy: the 7B model achieves 89.6% at  $\tau = 0.4$  vs. 88.6% at  $\tau = 0.0$ , suggesting that suppressing overthinking prevents error accumulation in intermediate reasoning steps. Beyond this range, aggressive thresholds ( $\tau \geq 0.5$ ) begin to degrade accuracy as essential System 2 reasoning is suppressed. Larger models exhibit flatter curves, indicating greater robustness to aggressive thresholds, which may be attributed to their stronger representations that allow System 1 to handle a wider range of reasoning steps effectively. This behavior also shows that  $\tau$  provides a controllable way to adjust the efficiency–accuracy trade-off. A lower threshold keeps the model closer to vanilla System 2 reasoning and is therefore more conservative, while a higher threshold encourages more frequent use of System 1 and yields larger token savings. Additional threshold sweeps across




Figure 8: Case study: comparison of reasoning trajectories between Standard Reasoning and InterSR on an AIME24 problem. InterSR dynamically switches between System 1 and System 2 to minimize token consumption while maintaining reasoning quality.

models and datasets are reported in Appendix N.

#### 4.6 System 1/2 Token Allocation (RQ5)

To further examine whether InterSR allocates computation according to problem difficulty, we record the number of tokens generated in each reasoning mode for R1-7B. As shown in Figure 7, the share of System 2 tokens increases as the benchmark becomes more challenging. On GSM8K, which mainly contains grade-school arithmetic problems, only 28% of the generated tokens are assigned to System 2. In contrast, AIME24 requires substantially more deliberation, and the System 2 ratio increases to 60%. MATH500 and AMC23 fall between these two cases, with System 2 ratios of 48% and 52%, respectively. This monotonic trend indicates that the learned switching signal captures dataset-level differences in reasoning demand, rather than merely enforcing a uniform token-reduction policy. It is also consistent with the token reduction results in Table 1, where simpler datasets allow larger compression while harder competition-level problems preserve more slow reasoning. These results suggest that InterSR does not apply a fixed reduction pattern across benchmarks, but adjusts the amount of deliberative computation according to the difficulty of the reasoning task.

#### 4.7 Case Study: InterSR vs. Standard (RQ6)

Figure 8 compares reasoning trajectories on an AIME24 problem. The standard process applies uniform System 2 reasoning throughout the trajectory, using the same deliberative mode for problem understanding, information extraction, equation solving, final computation, and re-verification. As a result, even relatively routine steps are accompanied by extended intermediate reasoning, and the model continues to verify the answer after the necessary algebraic computation has already been completed. This leads to 1454 generated tokens, with a substantial portion spent on redundant verification rather than new problem-solving progress. In contrast, InterSR uses System 2 for problem understanding (169 tokens) and the core equation-solving step that requires careful algebraic manipulation (356 tokens), while switching to System 1 for straightforward information extraction (78 tokens) and final arithmetic computation (42 tokens). The redundant re-verification step is eliminated once the solution becomes sufficiently certain. As a result, InterSR reduces the trajectory to 645 tokens—a 56% reduction—while preserving the deliberative steps needed to solve the problem. This case illustrates that InterSR mitigates overthinking by reallocating reasoning effort across steps, rather than uniformly shortening the entire response. Full

case details are in Appendix O.

## 5 Related Work

In this section, we review previous methods developed to mitigate the overthinking problem in large reasoning models. These methods can be categorized into two distinct perspectives: Response Perspective and Inference Perspective.

**Response Perspective.** In this perspective, methods focus on reducing token usage by directly manipulating the reasoning output. Dynasor (Fu et al., 2025) introduces Certainindex, which helps the model exit reasoning early when no further tokens are likely to change the final answer. ConciseRL (Dumitru et al., 2025) guides the reinforcement learning process through a conciseness score, directing the model to generate more concise responses. DEER (Yang et al., 2025b) uses answer triggers to stop reasoning once the model reaches high confidence. ThinkPrune (Hou et al., 2025) encourages shorter responses by truncating reward lengths during reinforcement learning. VPG-EA (Chen et al., 2026b) uses variational posterior guidance to suppress redundant reasoning tokens while preserving logical verification steps. FCS (Chen et al., 2025c) fine-tunes models on efficient reasoning datasets with First-Correct Solutions to promote concise outputs. ES-CoT explores early stopping in chain-of-thought reasoning to terminate reasoning when answers converge (Mao et al., 2025), and RCPD mines reasoning patterns to identify points for early exit (Wei et al., 2025). RPDM (Guan et al., 2026) monitors reasoning path deviation via high-entropy transition tokens to detect and terminate overthinking trajectories. These methods efficiently reduce token consumption by minimizing response length. However, they fail to account for the inherent complexity of the problem, which may result in overly condensed reasoning and a subsequent decline in performance.

**Inference Perspective.** Methods in this perspective focus on classifying problems by difficulty and dynamically selecting reasoning modes based on complexity, thereby minimizing unnecessary resource consumption. CKT-MHA (Qi et al., 2024) proposes an external policy to detect hard problems and adapt the reasoning strategy accordingly. FReM (Zhao et al., 2025) estimates the depth of reasoning and select prompts based on problem difficulty. TALE (Han et al., 2025) proposes

“Token Budget” that limits reasoning to the minimum number of tokens needed to solve a problem. DynaThink (Pan et al., 2024) adjusts reasoning depth by sampling multiple responses to determine if a problem needs fast or slow thinking. AdaptThink (Zhang et al., 2025) introduces a balanced training regime to toggle between “thinking” and “no-thinking” modes. DiffAdapt (Liu et al., 2026) trains a lightweight probe on hidden states to classify problem difficulty and dynamically select per-question inference strategies. ReBalance (Li et al., 2026) uses confidence as a continuous indicator to steer reasoning trajectories, pruning redundancy during overthinking and promoting exploration during underthinking. MixReasoning (Lu et al., 2026) and SwiReasoning (Shi et al., 2026) further explore inference-time mode switching for efficient reasoning. InterSR differs by predicting a semantic-entropy-based signal from hidden states and switching between System 1 and System 2 at reasoning-unit boundaries, enabling fine-grained control over local reasoning demand.

## 6 Conclusion

We propose Interleaved System-1/2 Reasoning (InterSR), a framework that mitigates overthinking in large reasoning models by dynamically switching between System 1 and System 2 modes according to learned system signals. InterSR allocates reasoning effort more efficiently, reducing token consumption while preserving task performance. Experiments on mathematical and cross-domain benchmarks show consistent token savings without compromising accuracy, demonstrating the potential of InterSR to improve reasoning efficiency across diverse tasks.

## Acknowledgments

This research work is supported by National Cyber Security-National Science and Technology Major Project (Grant No. 2025ZD1503000), the National Natural Science Foundation of China under Grant NOs. 62176014, the Fundamental Research Funds for the Central Universities and the Academic Excellence Foundation of BUAA for PhD Students.

## Limitations

Although InterSR achieves substantial efficiency gains across multiple benchmarks, two directions remain for future work. First, extending InterSR

to more model families beyond the DeepSeek-R1-Distill series would further validate its architectural generality. Second, our evaluation focuses on single-turn reasoning tasks, and its effectiveness in multi-turn agentic scenarios with long-horizon planning remains to be explored.

## Ethical Considerations

This work improves inference efficiency without introducing new training data or user data collection, and thus does not add privacy concerns beyond those of the underlying models. However, lower-cost reasoning may broaden deployment and amplify both beneficial and harmful uses. InterSR should therefore be applied with appropriate safeguards in high-stakes settings, and future work should further examine how dynamic switching affects transparency, interpretability, and human oversight.

## References

- Marthe Ballon, Andres Algaba, and Vincent Ginis. 2025. The relationship between reasoning and performance in large language models—o3 (mini) thinks harder, not longer. *arXiv preprint arXiv:2502.15631*.
- Yi Bin, Tianyi Jiang, Yujuan Ding, Kainian Zhu, Fei Ma, Jingkuan Song, Yang Yang, and Heng Tao Shen. 2026. Explore briefly, then decide: Mitigating llm overthinking via cumulative entropy regulation. *arXiv preprint arXiv:2510.02249*.
- Shelly Chaiken and Yaacov Trope. 1999. *Dual-process theories in social psychology*. Guilford Press.
- Hanting Chen, Yasheng Wang, Kai Han, Dong Li, Lin Li, Zhenni Bi, Jinpeng Li, Haoyu Wang, Fei Mi, Mingjian Zhu, and 1 others. 2025a. Pangu embedded: An efficient dual-system llm reasoner with metacognition. *arXiv preprint arXiv:2505.22375*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, and 1 others. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Wei Chen, Han Ding, Meng Yuan, Zhao Zhang, Deqing Wang, and Fuzhen Zhuang. 2025b. Bridging social psychology and llm reasoning: Conflict-aware meta-review generation via cognitive alignment. *arXiv preprint arXiv:2503.13879*.
- Wei-Lin Chen, Liqian Peng, Tian Tan, Chao Zhao, Blake JianHang Chen, Ziqian Lin, Alec Go, and Yu Meng. 2026a. Think deep, not just long: Measuring llm reasoning effort via deep-thinking tokens. *arXiv preprint arXiv:2602.13517*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, and 1 others. 2025c. Do not think that much for  $2+3=?$  on the overthinking of o1-like llms. In *Proceedings of the 42nd International Conference on Machine Learning*.
- Zizhao Chen, Yuying Li, Siting Lin, and Lianxi Wang. 2026b. Efficient llm reasoning via variational posterior guidance with efficiency awareness. *arXiv preprint arXiv:2605.11019*.
- Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and Furu Wei. 2025. Reasoning with exploration: An entropy perspective. *arXiv preprint arXiv:2506.14758*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Julian Coda-Forno, Zhuokai Zhao, Qiang Zhang, Dipesh Tamboli, Weiwei Li, Xiangjun Fan, Lizhu Zhang, Eric Schulz, and Hsiao-Ping Tseng. 2025. Exploring system 1 and 2 communication for latent reasoning in llms. *arXiv preprint arXiv:2510.00494*.
- Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao Wang, Yichuan Wang, Siyuan Zhuang, Shu Liu, Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao, and 1 others. 2025. The danger of overthinking: Examining the reasoning-action dilemma in agentic tasks. *arXiv preprint arXiv:2502.08235*.
- Zheng Du, Hao Kang, Song Han, Tushar Krishna, and Ligeng Zhu. 2026. Ockbench: Measuring the efficiency of llm reasoning. *arXiv preprint arXiv:2511.05722*.
- Razvan-Gabriel Dumitru, Darius Peteleaza, Vikas Yadav, and Liangming Pan. 2025. Conciserl: Conciseness-guided reinforcement learning for efficient reasoning models. *arXiv preprint arXiv:2505.17250*.
- Jonathan St BT Evans. 2017. Dual process theory. In *Dual Process Theory 2.0*, pages 137–155. Routledge.
- Guhao Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. 2023. Towards revealing the mystery behind chain of thought: a theoretical perspective. *Advances in Neural Information Processing Systems*, 36:70757–70798.
- Sicheng Feng, Gongfan Fang, Xinyin Ma, and Xinchao Wang. 2025. Efficient reasoning models: A survey. *arXiv preprint arXiv:2504.10903*.

- Yichao Fu, Junda Chen, Siqi Zhu, Zheyu Fu, Zhongdongming Dai, Yonghao Zhuang, Yian Ma, Aurick Qiao, Tajana S. Rosing, Ion Stoica, and Hao Zhang. 2025. Efficiently scaling LLM reasoning programs with certindex. In *Advances in Neural Information Processing Systems 38*, pages 157566–157604.
- Soumya Suvra Ghosal, Souradip Chakraborty, Avinash Reddy, Yifu Lu, Mengdi Wang, Dinesh Manocha, Furong Huang, Mohammad Ghavamzadeh, and Amrit Singh Bedi. 2025. Does thinking more always help? mirage of test-time scaling in reasoning models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Weixin Guan, Liang Li, Jiapeng Liu, Bing Li, Peng Fu, Chengyang Fang, Xiaoshuai Hao, Can Ma, and Weiping Wang. 2026. Mitigating overthinking in large reasoning language models via reasoning path deviation monitoring. *arXiv preprint arXiv:2603.14251*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, and 1 others. 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645(8081):633–638.
- Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. 2025. Token-budget-aware llm reasoning. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24842–24855.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Bairu Hou, Yang Zhang, Jiabao Ji, Yujian Liu, Kaizhi Qian, Jacob Andreas, and Shiyu Chang. 2025. Thinkprune: Pruning long chain-of-thought of llms via reinforcement learning. *arXiv preprint arXiv:2504.01296*.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations*.
- Yulin Li, Tengyao Tu, Li Ding, Junjie Wang, Huiling Zhen, Yixin Chen, Yong Li, and Zhuotao Tian. 2026. Efficient reasoning with balanced thinking. *arXiv preprint arXiv:2603.12372*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. Let’s verify step by step. In *International Conference on Learning Representations*, pages 39578–39601.
- Xiang Liu, Xuming Hu, Xiaowen Chu, and Eunsol Choi. 2026. Diffadapt: Difficulty-adaptive reasoning for token-efficient llm inference. *arXiv preprint arXiv:2510.19669*.
- Haiquan Lu, Gongfan Fang, Xinyin Ma, Qi Li, and Xinchao Wang. 2026. Mixreasoning: Switching modes to think. In *Proceedings of the 43rd International Conference on Machine Learning*.
- Minjia Mao, Bowen Yin, Yu Zhu, and Xiao Fang. 2025. Early stopping chain-of-thoughts in large language models. *arXiv preprint arXiv:2509.14004*.
- OpenAI, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. OpenAI o1 system card. *arXiv preprint arXiv:2412.16720*.
- Jiabao Pan, Yan Zhang, Chen Zhang, Zuozhu Liu, Hongwei Wang, and Haizhou Li. 2024. DynaThink: Fast or slow? a dynamic decision-making framework for large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 14686–14695, Miami, Florida, USA. Association for Computational Linguistics.
- Biqing Qi, Xinquan Chen, Junqi Gao, Dong Li, Jianxing Liu, Ligang Wu, and Bowen Zhou. 2024. Interactive continual learning: Fast and slow thinking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12892.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.
- Swarnadeep Saha, Archiki Prasad, Justin Chih-Yao Chen, Peter Hase, Elias Stengel-Eskin, and Mohit Bansal. 2024. System-1.x: Learning to balance fast and slow planning with language models. *arXiv preprint arXiv:2407.14414*.
- Dachuan Shi, Abedelkadir Asi, Keying Li, Xiangchi Yuan, Leyan Pan, Wenke Lee, and Wen Xiao. 2026. Swireasoning: Switch-thinking in latent and explicit for pareto-superior reasoning LLMs. In *The Fourteenth International Conference on Learning Representations*.
- Jihao Shi, Xiao Ding, Hengwei Zhao, Ting Liu, and Bing Qin. 2025. Bridging neural and symbolic reasoning: A dual-system framework for interpretable question answering. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Elizabeth P Shulman, Ashley R Smith, Karol Silva, Grace Icenogle, Natasha Duell, Jason Chein, and Laurence Steinberg. 2016. The dual systems model: Review, reappraisal, and reaffirmation. *Developmental cognitive neuroscience*, 17:103–117.
- Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. 2024. To cot or not to cot? chain-of-thought

- helps mainly on math and symbolic reasoning. *arXiv preprint arXiv:2409.12183*.
- Jinyan Su, Jennifer Healey, Preslav Nakov, and Claire Cardie. 2025. Between underthinking and overthinking: An empirical study of reasoning length and correctness in llms. *arXiv preprint arXiv:2505.00127*.
- Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Na Zou, and 1 others. 2025. Stop overthinking: A survey on efficient reasoning for large language models. *arXiv preprint arXiv:2503.16419*.
- Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. 2024. Dart-math: Difficulty-aware rejection tuning for mathematical problem-solving. *Advances in Neural Information Processing Systems*, 37:7821–7846.
- Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu, Juntao Li, Zhuosheng Zhang, and 1 others. 2025. Thoughts are all over the place: On the underthinking of o1-like llms. *arXiv preprint arXiv:2501.18585*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Zihao Wei, Liang Pang, Jiahao Liu, Wenjie Shi, Jingcheng Deng, Shicheng Xu, Zenghao Duan, Fei Sun, Huawei Shen, and Xueqi Cheng. 2025. The evolution of thought: Tracking LLM overthinking via reasoning dynamics analysis. *arXiv preprint arXiv:2508.17627*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Qiaowei Li, Minghui Chen, Zheng Lin, and Weiping Wang. 2025b. Dynamic early exit in reasoning models. *arXiv preprint arXiv:2504.15895*.
- Wenlin Yao, Haitao Mi, and Dong Yu. 2024. Hdflow: Enhancing llm complex problem-solving with hybrid thinking and dynamic workflows. *arXiv preprint arXiv:2409.17433*.
- Duzhen Zhang, Zhong-Zhi Li, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhijiang Guo, Le Song, and Cheng-Lin Liu. 2026. From system 1 to system 2: A survey of reasoning large language models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(3):3335–3354.
- Jiajie Zhang, Nianyi Lin, Lei Hou, Ling Feng, and Juanzi Li. 2025. Adaptthink: Reasoning models can learn when to think. *arXiv preprint arXiv:2505.13417*.
- Zhengyi Zhao, Shubo Zhang, Zezhong Wang, Bin Liang, Binyang Li, and Kam-Fai Wong. 2025. Frem: A flexible reasoning mechanism for balancing quick and slow thinking in long-context question answering. *arXiv preprint arXiv:2503.22985*.
- Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Jeff Huang, Chuyue Sun, Cody Hao Yu, Shiyi Cao, Christos Kober, Yineng Liang, Siyuan Zhuang, and 1 others. 2024. Sglang: Efficient execution of structured language model programs. In *Advances in Neural Information Processing Systems*.
- Shu Zhou, Rui Ling, Junan Chen, Xin Wang, Tao Fan, and Hao Wang. 2026. When more thinking hurts: Overthinking in llm test-time compute scaling. *arXiv preprint arXiv:2604.10739*.
- Alireza S Ziabari, Nona Ghazizadeh, Zhivar Sourati, Farzan Karimi-Malekabadi, Payam Piray, and Morteza Dehghani. 2025. Reasoning on a spectrum: Aligning llms to system 1 and system 2 thinking. *arXiv preprint arXiv:2502.12470*.

## A Experiment Details

This appendix provides a comprehensive description of the experimental setup used in this paper, including model architectures, auxiliary modules, training data construction, optimization details, inference configurations, evaluation protocols, and hyperparameter settings, to facilitate reproducibility.

### A.1 Models and Auxiliary Modules

#### A.1.1 Base Models

We evaluate our proposed InterSR framework on three model scales from the DeepSeek-R1-Distill-Qwen family (1.5B, 7B, 14B) and DeepSeek-R1-Distill-Llama-8B (Guo et al., 2025). These models are distilled versions of the Qwen-2.5 and Llama-3.1 architectures, optimized for reasoning-centric tasks. The corresponding hidden dimensions are 1536, 3584, 5120, and 4096, respectively.

#### A.1.2 Signal Prediction Modules

The InterSR framework introduces two lightweight auxiliary modules for adaptive reasoning control.

The boundary detector  $f_\psi$  is implemented as a two-layer MLP with ReLU activation and a dropout rate of 0.1. The signal predictor  $g_\phi$  is a three-layer MLP with ReLU activations, a dropout rate of 0.1, and a final Sigmoid layer that outputs normalized signal values  $\sigma_t \in [0, 1]$ .

Both modules operate in bfloat16 precision to improve computational efficiency.

## A.2 Training Data Construction

### A.2.1 Dataset and Problem Sampling

To train the signal prediction modules, we construct labeled data using the DART-Math dataset (Tong et al., 2024). Specifically, for each difficulty level (1–5), we sample 1,000 problems and generate full Chain-of-Thought solutions using the corresponding base model.

Due to the large scale of the original dataset, this sampling strategy ensures a balanced and manageable training set across difficulty levels.

### A.2.2 Reasoning Trace Segmentation

The generated reasoning traces are segmented into sentence-level units using the Stanza syntactic analysis toolkit. Each sentence boundary is treated as a candidate decision point for system switching.

### A.2.3 Semantic Entropy Estimation

At each sentence boundary, we sample  $K = 8$  candidate continuations from the model and compute their semantic embeddings using the Qwen3-Embedding-0.6B model (Yang et al., 2025a).

Semantic entropy is computed via incremental clustering with a distance threshold  $\delta = 0.5$ , and normalized by  $\log_2 K$ . The resulting entropy values serve as supervision signals for training the signal predictor.

For each model size, we construct approximately 10k training samples, with a 90%/10% train/test split.

## A.3 Module Training

We train both auxiliary modules using the AdamW optimizer with a learning rate of  $2 \times 10^{-4}$ . The batch size is set to 128 for both modules.

For the boundary detector  $f_\psi$ , we employ Focal Loss with  $\alpha = 1$  and  $\gamma = 2$  to address class imbalance between boundary and non-boundary positions.

For the signal predictor  $g_\phi$ , we use a weighted mean squared error (MSE) loss, assigning a weight of 0.1 to zero-valued labels.

Training is conducted with early stopping based on validation performance, with a patience of 3–5 evaluation rounds. Validation is performed every 10 batches throughout training.

## A.4 Inference Configuration

During inference, we set the maximum generation length to 16,384 tokens, the temperature to 0.6, and top- $p$  to 0.95 with nucleus sampling enabled.

The system switching threshold  $\tau$  is tunable. In our main experiments, we explore values ranging from 0.0 to 0.6 to analyze the efficiency–accuracy trade-off.

To ensure stable mode transitions, we enforce:

- A minimum of 10 tokens in System 2 mode before allowing exit, and
- At least 10 tokens in System 1 mode before re-entering System 2.

## A.5 Evaluation Setup

All experiments are conducted on  $4 \times$  NVIDIA H200 GPUs.

We evaluate InterSR on four mathematical reasoning benchmarks—GSM8K (Cobbe et al., 2021), MATH500 (Lightman et al., 2024), AIME24, and AMC23—as well as four cross-domain benchmarks: GPQA-Diamond (Rein et al., 2024) (graduate-level STEM), MMLU (knowledge-intensive QA), ARC-Challenge (science reasoning), and HumanEval (code generation).

We report Pass@1 accuracy and average token count as primary evaluation metrics. Answer correctness is verified using the `math-verify` toolkit for mathematical expression parsing and equivalence checking.

## A.6 Baselines

We compare InterSR against three state-of-the-art adaptive reasoning methods:

- TALE (Han et al., 2025), which controls generation length via token budget allocation based on problem complexity. We use the training-free TALE-EP variant for fair comparison.
- Dynasor (Fu et al., 2025), which employs confidence-based early exit using Certindex. We modify the original implementation to ensure complete solution output upon early termination.
- DEER (Yang et al., 2025b), a training-free early-exit method that uses answer triggers to halt generation early.

## A.7 Datasets

Below we describe the datasets used in our experiments. Due to differences in HuggingFace data formats, the field names in `evaluation_config.yaml` must be adjusted accordingly.

AIME24 is derived from the 2024 American Invitational Mathematics Examination and is available on HuggingFace.<sup>2</sup> In our evaluation, the “question” field corresponds to the “problem” field, and the “answer” field maps to the “solution” field.

AMC23 is based on the 2023 American Mathematics Competition.<sup>3</sup> The “question” field corresponds to the dataset’s “question” field, and the “answer” field maps to the “answer” field.

MATH500 is a curated subset of the MATH benchmark.<sup>4</sup> The “question” field corresponds to “problem”, and the “answer” field corresponds to “answer”.

GSM8K is a grade school math word problem benchmark.<sup>5</sup> The “question” and “answer” fields are directly used.

GPQA-Diamond is the highest-difficulty subset of GPQA.<sup>6</sup> We use the “diamond” split and convert the original format into multiple-choice questions by randomly shuffling the provided options and assigning them to options A–D.

MMLU (Massive Multitask Language Understanding) is a comprehensive knowledge benchmark covering 57 subjects.<sup>7</sup> We evaluate on the full test set and report overall accuracy.

ARC-Challenge is the challenge subset of the AI2 Reasoning Challenge, containing science questions that require complex reasoning.<sup>8</sup> We use the “ARC-Challenge” split.

HumanEval is a code generation benchmark consisting of 164 hand-crafted programming problems.<sup>9</sup> We report Pass@1 accuracy using

<sup>2</sup><https://huggingface.co/datasets/math-ai/aime24>

<sup>3</sup><https://huggingface.co/datasets/math-ai/amc23>

<sup>4</sup><https://huggingface.co/datasets/HuggingFaceH4/MATH-500>

<sup>5</sup><https://huggingface.co/datasets/openai/gsm8k>

<sup>6</sup><https://huggingface.co/datasets/Idavidrein/gpqa>

<sup>7</sup><https://huggingface.co/datasets/cais/mmlu>

<sup>8</sup>[https://huggingface.co/datasets/allenai/ai2\\_arc](https://huggingface.co/datasets/allenai/ai2_arc)

<sup>9</sup>[https://huggingface.co/datasets/openai/openai\\_humaneval](https://huggingface.co/datasets/openai/openai_humaneval)

| Models         | Datasets |       |         |       |
|----------------|----------|-------|---------|-------|
|                | AIME24   | AMC23 | MATH500 | GSM8K |
| <b>R1-1.5B</b> | 0.5      | 0.6   | 0.4     | 0.6   |
| <b>R1-7B</b>   | 0.2      | 0.4   | 0.4     | 0.4   |
| <b>R1-8B</b>   | 0.2      | 0.4   | 0.4     | 0.4   |
| <b>R1-14B</b>  | 0.2      | 0.3   | 0.4     | 0.3   |

Table 4: Signal thresholds used by each model across datasets.

execution-based evaluation.

**DART-Math**, We use the **Dart Math Pool Math** subset of DART-Math.<sup>10</sup> We sample 1,000 problems from each difficulty level to construct a balanced training set for signal prediction and boundary detection.

## A.8 Hyperparameter Settings

### A.8.1 Semantic Entropy Calculation

Semantic entropy is computed via clustering of sampled response embeddings. We set the distance threshold to  $\delta = 0.5$ . Two responses with embeddings  $\mathbf{e}_i$  and  $\mathbf{e}_j$  are merged into the same cluster if:

$$\|\mathbf{e}_i - \mathbf{e}_j\| < \delta.$$

The centroid of the merged cluster is computed as:

$$\mathbf{e}_C = \frac{1}{|C_i| + |C_j|} \sum_{k \in C_i \cup C_j} \mathbf{e}_k.$$

After clustering, suppose there are  $K$  clusters with frequencies  $p_k$ . The semantic entropy  $\mathcal{H}$  is defined as:

$$\mathcal{H} = - \sum_{k=1}^K p_k \log p_k.$$

### A.8.2 System Switching Thresholds

For smaller models, we adopt higher thresholds since their intrinsic signal strengths tend to be larger. Conversely, larger models require lower thresholds to trigger System 2 appropriately.

For the particularly challenging AIME24 dataset, we use lower thresholds to encourage more frequent engagement of System 2, ensuring sufficient deliberation and maintaining strong performance.

<sup>10</sup><https://huggingface.co/datasets/hkust-nlp/dart-math-pool-math>




Figure 9: Comparison of token generation efficiency between InterSR and standard reasoning across different model sizes.

## B Cost Analysis

In this section, we analyze the computational overhead introduced by InterSR’s signal prediction module from two perspectives: parameter count and inference speed.

### B.1 Parameter Analysis

The signal prediction module in InterSR consists of a boundary detector  $f_\psi$  (two-layer MLP) and a signal predictor  $g_\phi$  (three-layer MLP). Table 5 reports the parameter counts for each model scale.

| Model   | Base  | InterSR | Ratio |
|---------|-------|---------|-------|
| R1-1.5B | 1.5B  | 4.7M    | 0.31% |
| R1-7B   | 7.6B  | 25.7M   | 0.34% |
| R1-8B   | 8.0B  | 33.6M   | 0.42% |
| R1-14B  | 14.7B | 52.4M   | 0.36% |

Table 5: Parameter overhead of the InterSR signal prediction module across model scales.

Across all model scales, the signal prediction module adds at most 52.4M parameters, accounting for no more than 0.42% of the base model size. This lightweight design ensures that InterSR can be integrated into existing models with negligible memory overhead, making it practical for both research and production deployment.

### B.2 Token Efficiency Analysis

We further investigate the model’s generation efficiency on a per-token basis, as shown in Figure 9. The results indicate that our method increases the per-token generation time by less than 3%, while achieving an average token reduction of about 28.5% compared with vanilla reasoning across the mathematical reasoning settings in Table 1. This minimal per-token overhead, combined

with the significant savings in token consumption, highlights the efficiency of our Interleaved System-1/2 Reasoning framework. By dynamically allocating cognitive resources and avoiding unnecessary overthinking, the method not only reduces redundant computations but also accelerates the path to the correct solution. These findings demonstrate that our approach effectively balances computational cost and reasoning quality, ultimately enhancing the overall efficiency and practicality of large language model reasoning.

## C Algorithm

InterSR generates tokens by dynamically switching between fast (System 1) and slow (System 2) reasoning. At each step, a learned trigger decides whether to switch, inserting special tokens for slow thinking when needed. Otherwise, tokens are generated normally. This allows the model to focus slow, careful reasoning on critical steps while handling simpler parts quickly, improving overall efficiency. The full algorithm is illustrated in Algorithm 1.

## D Control Tag Distribution Analysis

A natural concern is whether the `<think>` and `</think>` control tags inserted by InterSR genuinely induce different reasoning behaviors, or merely act as superficial markers with no impact on the generation process. To address this, we conduct a systematic analysis of the next-token distributions following each tag type.

**Setup.** We collect all reasoning traces generated by the DeepSeek-R1-Distill-Qwen-7B model under InterSR on the GSM8K and MATH500 datasets. For each occurrence of `<think>` (entering System 2) and `</think>` (exiting to System 1), we record the probability distribution over the vocabulary for the immediately following token. We then aggregate these distributions across all instances to obtain two average next-token distributions:  $P_{\text{enter}}$  (after `<think>`) and  $P_{\text{exit}}$  (after `</think>`).

**Jensen-Shannon Divergence.** We compute the Jensen-Shannon Divergence (JSD) between  $P_{\text{enter}}$  and  $P_{\text{exit}}$ :

$$\text{JSD}(P_{\text{enter}} \| P_{\text{exit}}) = \frac{1}{2} D_{\text{KL}}(P_{\text{enter}} \| M) + \frac{1}{2} D_{\text{KL}}(P_{\text{exit}} \| M) \quad (5)$$

where  $M = \frac{1}{2}(P_{\text{enter}} + P_{\text{exit}})$ . The resulting JSD is **0.83** (on a  $[0, 1]$  scale), indicating a strong distribu-

---

**Algorithm 1** Interleaved Reasoning with System Switching

---

**Require:** Prompt  $x$ , reasoning model  $M$ , switch trigger  $f_\psi$ , confidence predictor  $g_\phi$ , threshold  $\tau$ , initial mode  $m \in \{S_1, S_2\}$ , maximum length  $L_{\max}$

**Ensure:** Token sequence  $t$

```

1:  $t \leftarrow []$  {Initialize output token sequence}
2:  $k \leftarrow 0$  {Initialize step counter}
3: while  $k \leq L_{\max}$  do
4:    $k \leftarrow k + 1$  {Increment step counter}
5:    $h_k \leftarrow M(x, t)$  {Compute current hidden state from model}
6:   if  $f_\psi(h_k) \geq 0$  then
7:      $\hat{\sigma}_k \leftarrow g_\phi(h_k)$  {Compute switching signal strength}
8:     if  $\hat{\sigma}_k \geq \tau$  and  $m = S_1$  then
9:        $t \leftarrow t \cup \{<\text{think}>\}$  {Insert System 2 thinking token}
10:       $m \leftarrow S_2$  {Switch to System 2}
11:      continue {Skip current token generation}
12:     else if  $\hat{\sigma}_k < \tau$  and  $m = S_2$  then
13:        $t \leftarrow t \cup \{</\text{think}>\}$  {Insert System 1 return token}
14:        $m \leftarrow S_1$  {Switch to System 1}
15:       continue {Skip current token generation}
16:     end if
17:   end if
18:    $t_k \sim \text{softmax}(Wh_k)$  {Generate token normally from hidden state}
19:    $t \leftarrow t \cup \{t_k\}$  {Append generated token to output sequence}
20:   if  $t_k$  is end-of-sequence then
21:     break {Stop generation at EOS token}
22:   end if
23: end while
24: return  $t$  {Return final token sequence}

```

---

tional shift between the two modes. As a reference, JSD values above 0.5 are generally considered indicative of substantially different distributions.

**Top-Token Overlap.** We further examine the top-50 most probable tokens under each distribution. The overlap between the two sets is only **8.2%** (approximately 4 shared tokens out of 50). After `<think>`, the model predominantly generates tokens associated with deliberative reasoning (such as mathematical operators, “let”, “consider”, “sup-




Figure 10: ROC curve and predicted score distribution for the system signal predictor of DeepSeek-R1-Distill-Qwen-7B.

pose”), while after `</think>`, the distribution favors direct, concise tokens (such as “so”, “therefore”, digits, punctuation).

These findings confirm that InterSR’s mode switching induces genuine cognitive regime changes at the token level, rather than superficial tag insertion with no downstream effect.

## E Signal Predictor ROC Analysis

We further evaluate whether the learned predictor can reliably capture the switching signal from hidden states. We binarize the continuous signal intensity by applying a threshold and evaluate the predictor’s discrimination ability. As shown in Figure 10, the predictor achieves an AUC of 0.838, with limited overlap between predicted scores for high- and low-intensity signals, confirming that discriminative features can be effectively extracted from hidden states.

## F SGLang Integration

To facilitate practical deployment of InterSR in production environments, we integrate it into the SGLang inference framework (Zheng et al., 2024), a high-performance serving engine for large language models. SGLang provides continuous batching, PagedAttention, and RadixAttention optimizations that are essential for efficient serving. Our integration preserves all these optimizations while enabling real-time InterSR-based reasoning mode control.

**Integration Mechanism.** SGLang supports the `SGLANG_EXTERNAL_MODEL_PACKAGE` environment variable, which allows users to register custom model wrappers that intercept the forward pass without modifying SGLang’s core codebase. We leverage this mechanism to inject InterSR’s signal prediction logic.

**Implementation Details.** The integration consists of the following components:

1. **Custom Model Wrapper.** We define a model class that inherits from SGLang’s base model interface and wraps the original DeepSeek model. The trained boundary detector  $f_\psi$  and signal predictor  $g_\phi$  are attached as additional submodules and loaded alongside the base model weights.
2. **Hidden State Extraction.** During each forward pass, the wrapper hooks into the final transformer layer to extract the hidden state  $h_k$ . This is implemented via PyTorch forward hooks, adding negligible overhead.
3. **Signal Evaluation and Token Injection.** At each generation step, the boundary detector  $f_\psi(h_k)$  first determines whether the current position is a candidate boundary. If so, the signal predictor  $g_\phi(h_k)$  computes the switching signal  $\hat{\sigma}_k$ . When a mode transition is triggered (i.e.,  $\hat{\sigma}_k \geq \tau$  in System 1 or  $\hat{\sigma}_k < \tau$  in System 2), the wrapper injects the appropriate control token (`<think>` or `</think>`) into the generation stream by modifying the next-token logits.
4. **State Management.** The current reasoning mode (System 1 or System 2) and the token count since the last mode switch are maintained per request in SGLang’s request state, ensuring correct behavior under continuous batching where multiple requests are processed concurrently.

**Logits Modification for Mode Switching.** The core of the integration lies in how mode switching decisions are translated into the autoregressive generation process. Rather than externally interrupting the decoding loop, we embed the switching logic directly into the model’s logits output via an indicator-gated logits modification mechanism.

Let  $\ell_k \in \mathbb{R}^{|V|}$  denote the raw logits produced by the model at step  $k$ , where  $|V|$  is the vocabulary size. We define a mode switching indicator function  $\mathbb{I}_{\text{switch}}(k)$  as:

$$\mathbb{I}_{\text{switch}}(k) = \mathbb{I}[f_\psi(h_k) \geq 0] \cdot \begin{cases} \mathbb{I}[\hat{\sigma}_k \geq \tau], & \text{if } m_k = S_1 \\ \mathbb{I}[\hat{\sigma}_k < \tau], & \text{if } m_k = S_2 \end{cases} \quad (6)$$

where  $f_\psi(h_k) \geq 0$  indicates a candidate boundary,  $\hat{\sigma}_k = g_\phi(h_k)$  is the predicted signal strength,  $\tau$  is

the switching threshold, and  $m_k \in \{S_1, S_2\}$  is the current reasoning mode.

When  $\mathbb{I}_{\text{switch}}(k) = 1$ , we modify the logits to force the generation of the appropriate control token. Let  $t_{\text{ctrl}}$  denote the target control token ID (`<think>` when switching  $S_1 \rightarrow S_2$ , or `</think>` when switching  $S_2 \rightarrow S_1$ ). The modified logits  $\tilde{\ell}_k$  are computed as:

$$\tilde{\ell}_k[v] = \begin{cases} \ell_k[v] + \lambda, & \text{if } v = t_{\text{ctrl}} \\ -\infty, & \text{otherwise} \end{cases} \quad (7)$$

where  $\lambda$  is a large positive constant (we use  $\lambda = 10^4$ ) that ensures the control token is selected with probability approaching 1 after softmax normalization. Formally, the probability of generating the control token becomes:

$$P(t_{\text{ctrl}} | h_k) = \frac{\exp(\tilde{\ell}_k[t_{\text{ctrl}}])}{\sum_{v \in V} \exp(\tilde{\ell}_k[v])} \approx 1 \quad (8)$$

This logits-level injection approach has several advantages: (1) it operates entirely within the model’s forward pass, requiring no modification to SGLang’s sampling kernel or decoding loop; (2) it is compatible with batched inference, as the logits modification can be applied independently per request; (3) it preserves the KV cache structure, since the injected control token is processed as a regular token in subsequent steps.

**Minimum Mode Duration Constraint.** To prevent rapid oscillation between modes (which could degrade output coherence), we enforce a minimum mode duration constraint. Let  $n_k$  denote the number of tokens generated since the last mode switch. The effective indicator becomes:

$$\tilde{\mathbb{I}}_{\text{switch}}(k) = \mathbb{I}_{\text{switch}}(k) \cdot \mathbb{I}[n_k \geq n_{\min}] \quad (9)$$

where  $n_{\min} = 10$  is the minimum number of tokens required in each mode before a switch is permitted. This constraint prevents degenerate behavior where the model rapidly alternates between modes at every boundary.

**Deployment.** Users can deploy InterSR-enhanced models by simply setting the environment variable and pointing to the custom model package:

```
SGLANG_EXTERNAL_MODEL_PACKAGE=\
  isr_model \
python -m sglang.launch_server \
  --model \
    deepseek-r1-distill-qwen-7b
```

**Performance.** The integration adds less than 3% overhead to per-token generation time, consistent with our standalone evaluation results (see Appendix B). The overhead comes primarily from the MLP forward passes of  $f_\psi$  and  $g_\phi$ , which together contain fewer than 52.4M parameters (see Table 5). The logits modification itself is a constant-time operation that adds negligible latency.

## G Failure Case Analysis

While InterSR achieves consistent efficiency gains across benchmarks, it is important to understand its failure modes. We conduct a systematic analysis of cases where InterSR produces incorrect answers that the standard (vanilla) model answers correctly. We identify two primary failure patterns:

**Premature System 1 Transition.** The most common failure mode occurs when InterSR transitions from System 2 to System 1 too early during a critical reasoning step. This typically happens when the signal predictor underestimates the complexity of an upcoming derivation. Consider a multi-step algebraic problem where the model needs to expand a polynomial expression. If the signal drops below the threshold  $\tau$  after the model has set up the equation but before completing the expansion, InterSR switches to System 1. In System 1 mode, the model tends to skip intermediate steps, which can lead to arithmetic errors or incorrect simplifications.

Formally, let  $r_t$  be a reasoning step where full System 2 reasoning is required (i.e.,  $\mathcal{G}(s_t, 0) < T$ ), but the predictor incorrectly estimates  $\hat{\sigma}_t < \tau$ . The resulting information loss can be quantified as:

$$\Delta\mathcal{G}_t = \mathcal{G}(s_t, 1) - \mathcal{G}(s_t, 0) > 0 \quad (10)$$

When  $\Delta\mathcal{G}_t$  is large (i.e., the step is genuinely complex), premature switching leads to critical information loss that propagates through subsequent reasoning steps. In our analysis of the 7B model on MATH500, approximately 62% of InterSR-specific errors (cases where InterSR fails but vanilla succeeds) involve at least one premature  $S_2 \rightarrow S_1$  transition during a critical derivation step.

**Code and Symbolic Syntax Disruption.** The second failure pattern is specific to tasks involving structured syntax, such as mathematical notation or code generation. When InterSR inserts a `</think>` tag mid-expression, it can disrupt the syntactic structure of the output. For example, in a LaTeX-formatted mathematical derivation, a mode

switch between `\frac{` and its closing brace can cause the model to lose track of the nesting structure, producing malformed expressions.

This issue is more prevalent in the 1.5B model, where the boundary detector  $f_\psi$  has lower precision due to the smaller hidden dimension (1536 vs. 3584 for 7B). In our analysis, syntax disruption accounts for approximately 23% of InterSR-specific errors on the 1.5B model but only 8% on the 14B model, suggesting that larger models develop more robust representations of syntactic boundaries.

**Error Rate Analysis.** We quantify the overall failure rate by examining the fraction of problems where InterSR answers incorrectly but vanilla answers correctly (InterSR-specific errors):

| Model   | GSM8K | MATH500 | AIME24 | AMC23 |
|---------|-------|---------|--------|-------|
| R1-1.5B | 3.2%  | 4.6%    | 6.7%   | 5.0%  |
| R1-7B   | 1.8%  | 2.4%    | 3.3%   | 2.5%  |
| R1-8B   | 2.2%  | 2.8%    | 3.3%   | 2.5%  |
| R1-14B  | 1.1%  | 1.4%    | 3.3%   | 2.5%  |

Table 6: InterSR-specific error rate (problems where InterSR fails but vanilla succeeds). These are offset by cases where InterSR succeeds but vanilla fails, resulting in the net accuracy differences shown in Table 1.

Importantly, these InterSR-specific errors are partially offset by cases where InterSR *improves* over vanilla (by avoiding overthinking-induced errors), resulting in the net accuracy changes reported in our main results. The decreasing error rate with model scale suggests that larger models produce more reliable signal predictions, enabling more accurate mode switching decisions.

## H KV Cache Efficiency Analysis

A practical concern for deploying InterSR in production serving systems is whether the insertion of control tokens (`<think>` and `</think>`) and the resulting mode switches affect the KV cache structure and inference efficiency. We provide a detailed analysis below.

**KV Cache Compatibility.** The control tokens inserted by InterSR are standard tokens in the model’s vocabulary. When a control token is injected via the logits modification mechanism (see Appendix F), it is processed by the model in the same way as any other generated token: the model computes its attention over all previous KV entries and appends its own key-value pair to the cache. Crucially, this means that:

1. The KV cache remains **contiguous and complete**—no entries are deleted, reordered, or invalidated during a mode switch.
2. The attention mechanism can attend to all previous tokens regardless of which reasoning mode produced them, preserving full contextual continuity.
3. No cache invalidation or recomputation is required, unlike methods that modify or rollback previously generated tokens.

**Memory Savings.** Since InterSR reduces the total number of generated tokens (by 11%–51% across benchmarks), it directly reduces the peak KV cache memory consumption. For a model with  $L$  layers, hidden dimension  $d$ , and  $n$  generated tokens, the KV cache memory is:

$$\mathcal{M}_{KV} = 2 \cdot L \cdot n \cdot d \cdot b \quad (11)$$

where  $b$  is the bytes per element (2 for bfloat16). Since InterSR reduces  $n$  by a factor of  $(1 - \alpha)$  where  $\alpha$  is the token reduction ratio, the memory savings are proportional:

$$\Delta \mathcal{M}_{KV} = 2 \cdot L \cdot \alpha \cdot n_{\text{vanilla}} \cdot d \cdot b \quad (12)$$

For the 7B model ( $L = 28$ ,  $d = 3584$ ) with an average 30% token reduction on MATH500 (from 3824 to 2651 tokens), this translates to:

$$\Delta \mathcal{M} = 2 \times 28 \times 1173 \times 3584 \times 2 \approx 472 \text{ MB} \quad (13)$$

per request. In a serving scenario with multiple concurrent requests, these savings are multiplicative, enabling higher throughput or larger batch sizes under the same memory budget. The detailed serving throughput comparison is presented in Table 2 in the main text.

## I System 1/2 Mode Distribution

A key design principle of InterSR is that it adaptively allocates reasoning effort based on problem complexity. To verify this, we analyze the proportion of tokens generated under System 1 (fast) versus System 2 (slow) across benchmarks and model scales.

**Setup.** For each model (1.5B, 7B, 14B), we run InterSR on all four mathematical benchmarks and record the number of tokens generated in each reasoning mode. Tokens within `<think>...</think>` segments are counted as System 2 tokens; all other tokens are counted as System 1.

**Results.** As shown in Figure 11, the System 2 token ratio increases consistently with benchmark difficulty across all model scales. On GSM8K (easiest), the 7B model allocates only 28% of tokens to System 2, while on AIME24 (hardest), this ratio rises to 60%. This confirms that InterSR’s signal-driven switching mechanism correctly identifies and responds to varying levels of problem complexity.

Two additional patterns emerge across model scales:

1. **Larger models use less System 2.** For the same benchmark, larger models consistently allocate a smaller proportion of tokens to System 2. On GSM8K, the S2 ratio decreases from 35% (1.5B) to 28% (7B) to 22% (14B). This reflects the superior reasoning capability of larger models—they can resolve more reasoning steps through fast System 1 processing without sacrificing accuracy.
2. **Difficulty sensitivity scales with model size.** The gap between the easiest (GSM8K) and hardest (AIME24) benchmarks in S2 ratio is 37 percentage points for the 1.5B model but only 33 points for the 14B model, indicating that larger models maintain a more balanced allocation even on hard problems.

## J Token Reduction vs. Problem Complexity

While the cross-benchmark analysis (Appendix M) examines aggregate token reduction at the dataset level, we further investigate how token reduction varies with individual problem complexity within each dataset.

**Difficulty Proxy.** Since the evaluation benchmarks lack per-problem difficulty annotations, we use the vanilla model’s token count as a proxy for problem complexity. This is a natural measure: problems that require more tokens under standard reasoning involve longer derivation chains and more complex reasoning steps, and thus are inherently more difficult.

**Results.** Figure 12 plots the token reduction percentage against vanilla token count for each problem, with a rolling average curve (window size 100) showing the overall trend. Across all three model scales, we observe a clear downward trend: simpler




Figure 11: System 1/2 token distribution across benchmarks (ordered by difficulty) for three model scales. The increasing System 2 ratio on harder benchmarks demonstrates InterSR’s adaptive reasoning allocation.




Figure 12: Token reduction vs. problem complexity (measured by vanilla token count) for three model scales. The rolling average (bold line) shows a clear downward trend, confirming that InterSR saves more tokens on simpler problems. Individual problems shown as semi-transparent scatter points.

problems (fewer vanilla tokens) exhibit higher token reduction, while more complex problems show lower but still positive reduction.

Key observations include:

1. **Monotonic trend.** The rolling average curves show a consistent downward slope across all model sizes, validating InterSR’s core mechanism: simpler reasoning steps (which dominate easy problems) have lower semantic entropy, triggering more System 1 usage and greater token savings.
2. **Model-scale effect.** The 1.5B model exhibits the steepest decline (from  $\sim 60\%$  for simple problems to  $\sim 33\%$  for complex ones), while the 14B model shows a flatter curve with more consistent reduction ( $\sim 40\%$  to  $\sim 27\%$ ). This suggests that larger models achieve more uniform efficiency gains across difficulty levels.
3. **No negative outliers.** Virtually all problems

achieve positive token reduction, confirming that InterSR does not introduce substantial overhead even on problems where few mode switches are beneficial.

## K Token Reduction by Problem Complexity

A central claim of InterSR is that it saves more tokens on simpler reasoning steps while preserving full deliberation on complex ones. To validate this at the problem level, we use the vanilla model’s token count as a proxy for problem difficulty—problems that require more tokens under standard reasoning are inherently more complex, as they involve longer derivation chains and more reasoning steps.

**Setup.** We pool all problems from GSM8K, MATH500, AMC23, and AIME24 and partition them into five buckets based on their vanilla token count:  $<500$ ,  $500\text{--}2\text{K}$ ,  $2\text{K--}5\text{K}$ ,  $5\text{K--}10\text{K}$ , and




Figure 13: Average token reduction by problem complexity (measured via vanilla token count) for the 7B model across all benchmarks. Simpler problems (fewer vanilla tokens) exhibit larger reductions, confirming InterSR’s adaptive resource allocation.

>10K. For each bucket, we compute the average token reduction ratio achieved by InterSR (using the DeepSeek-R1-Distill-Qwen-7B model).

**Results.** As shown in Figure 13, the token reduction ratio decreases monotonically as problem complexity increases. Problems in the simplest bucket (<500 tokens) achieve an average reduction of over 40%, while the most complex problems (>10K tokens) see reductions of approximately 25%. This trend is consistent with our theoretical framework: simpler reasoning steps have lower semantic entropy, causing InterSR to allocate more System 1 processing and thus save more tokens. Conversely, complex problems contain more high-entropy steps that require System 2 reasoning, limiting the scope for token reduction.

This analysis provides evidence that InterSR’s system switching mechanism is well-calibrated: it aggressively compresses easy problems while being appropriately conservative on difficult ones, ensuring that token savings do not come at the cost of reasoning quality.

## L Per-Problem Token Comparison

To provide a fine-grained view of InterSR’s behavior at the individual problem level, we plot the token count of each problem under InterSR versus its vanilla token count for the DeepSeek-R1-Distill-Qwen-7B model.

**Analysis.** As shown in Figure 14, each point represents a single problem, with the diagonal line  $y = x$  indicating no token reduction. Points below this line indicate that InterSR uses fewer tokens than vanilla reasoning. Several patterns emerge:

1. **Consistent reduction.** The vast majority of points fall below the diagonal across all four benchmarks, confirming that InterSR consistently reduces token usage rather than increasing it on certain problems.
2. **Proportional savings.** The points roughly follow a linear trend with slope less than 1, indicating that InterSR’s savings scale proportionally with problem length. This suggests that the signal predictor effectively identifies compressible reasoning steps regardless of absolute problem difficulty.
3. **Bounded degradation.** Very few points appear above the diagonal (InterSR uses more tokens than vanilla). When this occurs, the excess is typically small (<5%), caused by the overhead of inserted control tokens in problems where few mode switches are beneficial.
4. **Benchmark-specific patterns.** On GSM8K (simpler problems), points cluster tightly below the diagonal with high reduction ratios. On AIME24 (hardest problems), points are more dispersed but still predominantly below the line, reflecting the heterogeneous difficulty of individual reasoning steps within each problem.

## M Cross-Benchmark Token Reduction Analysis

To further characterize InterSR’s behavior across different difficulty levels, we analyze the token reduction patterns across all four mathematical reasoning benchmarks, which form a natural difficulty gradient: GSM8K (elementary) < MATH500 (competition-level) < AMC23 (AMC competition) < AIME24 (AIME competition).

**Results.** Figure 15 presents the token reduction percentages for all three model scales across the four benchmarks. Key observations include:

1. **Scale-dependent behavior.** The 1.5B model achieves the highest reduction on GSM8K (51%) but shows more variable reductions on harder benchmarks, reflecting its limited capacity to maintain accuracy under aggressive compression. The 7B model exhibits consistently strong reductions across difficulty levels, while the 14B model is more conservative on AIME24 but still achieves substantial savings on the other

Per-Problem Token Comparison: Vanilla vs. ISR (7B)




Figure 14: Per-problem token comparison between vanilla and InterSR reasoning (7B model). Left: GSM8K and MATH500. Right: AMC23 and AIME24. Points below the diagonal indicate token savings by InterSR.

benchmarks, suggesting that stronger signal predictors can better calibrate compression to task difficulty.

2. **Difficulty interaction.** While the general trend shows higher reduction on easier benchmarks, the relationship is not strictly monotonic. For example, the 7B model achieves 33% reduction on AIME24 (the hardest benchmark) compared to 27% on AMC23. This can be attributed to the fact that AIME24 problems, while overall more difficult, contain proportionally more straightforward setup steps that InterSR can efficiently compress.
3. **Practical implications.** Even on the most challenging benchmark (AIME24), InterSR achieves 11%–35% token reduction without sacrificing accuracy, demonstrating that substantial computational savings are achievable even in high-difficulty reasoning scenarios.

## N Additional Rebuttal Experiments

This section collects the additional experiments conducted during the rebuttal period. We keep these results in the appendix to preserve the focus of the main paper while still documenting the reviewer-requested evidence.

### N.1 S1/S2 Semantic-Entropy Validation

To directly validate the relationship between the semantic-entropy signal and the selected reasoning mode, we recompute  $\sigma$  offline at boundaries produced by InterSR and group them by the actual




Figure 15: Token reduction percentage across benchmarks (ordered by difficulty) for three model scales. InterSR achieves consistent savings across all difficulty levels.

| Dataset | S1 $\sigma$ | S2 $\sigma$ | Diff. | $p$ -value           |
|---------|-------------|-------------|-------|----------------------|
| MATH500 | 0.225       | 0.350       | 0.124 | 0.00834              |
| AIME24  | 0.403       | 0.496       | 0.093 | 0.0218               |
| GPQA-D  | 0.346       | 0.598       | 0.253 | $6.0 \times 10^{-5}$ |

Table 7: Offline semantic-entropy check for routed boundaries. Boundaries routed to the deliberative mode have consistently higher  $\sigma$  than boundaries routed to the fast mode.

routing decisions. Table 7 shows that boundaries routed to the deliberative mode have consistently higher semantic entropy than those routed to the fast mode, supporting the use of  $\sigma$  as a proxy for local reasoning demand.

### N.2 Out-of-Domain Predictor AUC

Because the predictor is trained on DART-Math, we directly evaluate whether its discrimination quality transfers to other datasets. Table 8 fixes the

| Dataset | DART-Math | MATH500 | AIME24 | GPQA-D |
|---------|-----------|---------|--------|--------|
| AUC     | 0.838     | 0.779   | 0.812  | 0.782  |

Table 8: Out-of-domain AUC of the DART-Math-trained R1-7B signal predictor.

DART-trained R1-7B predictor and recomputes semantic-entropy labels with the same sampling and clustering protocol. The predictor remains informative on MATH500, AIME24, and GPQA-D, suggesting that the signal predictor learns transferable uncertainty-related features.

### N.3 Threshold Sensitivity

To verify that InterSR is not dependent on a single hand-picked threshold, we additionally sweep  $\tau$  across multiple model-dataset combinations. Table 9 shows that InterSR preserves a favorable accuracy-token trade-off across a range of thresholds. Very aggressive thresholds can reduce tokens further but may hurt accuracy on some datasets.

### N.4 Comparison with SwiReasoning

Concurrent methods such as MixReasoning (Lu et al., 2026) and SwiReasoning (Shi et al., 2026) also explore switching between reasoning modes during inference. InterSR differs in two ways. First, it performs switching at reasoning-unit boundaries inside a single trajectory, rather than assigning a mode globally to the whole instance. Second, its routing signal is learned from hidden states to amortize the expensive semantic-entropy computation, allowing fine-grained online control with negligible additional decoding overhead.

Table 11 compares InterSR with SwiReasoning on R1-7B. InterSR obtains lower token counts on all four math benchmarks while preserving or improving accuracy. This supports the advantage of interleaving mode switches at reasoning-unit boundaries rather than relying only on coarser mode-switching behavior.

### N.5 Model-Family Transfer to Gemma4-12B-it

Following the same evaluation protocol and inference configuration described in Appendix A, we evaluate InterSR on Gemma4-12B-it as a cross-architecture testbed to assess whether the switching mechanism transfers beyond the DeepSeek-R1-Distill model family.

Table 10 evaluates InterSR on Gemma4-12B-it. InterSR maintains comparable or slightly higher

accuracy across all six benchmarks while reducing token usage, supporting transfer beyond the DeepSeek-distilled model family.

## O Full Case Study

Due to space limitations in the main paper, we are unable to present the full case study. Here, we provide the complete case study to more clearly illustrate the differences between the InterSR and Standard reasoning processes.

### O.1 Standard Reasoning Case

This section presents the complete output of the standard reasoning process used in the case study. As shown, the response is excessively verbose and contains a significant amount of repetition—clear indicators of overthinking behavior.

#### Standard Reasoning Case

<think>

Okay, so Aya has this morning routine where she goes for a 9-kilometer walk and then stops at a coffee shop. The problem is giving me two different scenarios about her walking speed and the time it takes, including the time she spends in the coffee shop. I need to figure out how long her walk will take if she walks at a speed of  $s + 1/2$  kilometers per hour, including the coffee shop time.

Alright, let me break this down. First, in the first scenario, she walks at a constant speed of  $s$  km/h and the entire walk takes her 4 hours, which includes  $t$  minutes at the coffee shop. In the second scenario, she walks faster, at  $s + 2$  km/h, and the walk takes her 2 hours and 24 minutes, still including  $t$  minutes at the coffee shop.

So, the difference in time must be due to the difference in walking speed. I think I can set up equations based on the information given. Let me convert everything into hours. First scenario: total time is 4 hours,  $t$  minutes is  $t/60$  hours. Therefore, the walking time is  $4 - t/60$  hours.

Similarly, in the second scenario, total time is 2.4 hours, so walking time is  $2.4 - t/60$  hours.

Now, using distance = speed  $\times$  time:  $9 = s \times (4 - t/60)$  and  $9 = (s + 2) \times (2.4 - t/60)$ . Let  $T = t/60$ . From Eq 1:  $s = 9 / (4 - T)$ .

| Model  | Dataset | Vanilla |      | $\tau = 0.2$ |      | $\tau = 0.3$ |      | $\tau = 0.4$ |      | $\tau = 0.5$ |      |
|--------|---------|---------|------|--------------|------|--------------|------|--------------|------|--------------|------|
|        |         | Acc.    | Tok. | Acc.         | Tok. | Acc.         | Tok. | Acc.         | Tok. | Acc.         | Tok. |
| R1-8B  | AMC23   | 80.0    | 5518 | 80.0         | 5010 | 80.0         | 4585 | 80.0         | 4138 | 75.0         | 3520 |
| R1-14B | AMC23   | 92.5    | 5435 | 92.5         | 4890 | 92.5         | 3894 | 92.5         | 3510 | 87.5         | 3045 |
| R1-14B | GSM8K   | 91.9    | 705  | 91.9         | 642  | 92.0         | 544  | 91.4         | 481  | 90.8         | 407  |
| R1-14B | GPQA-D  | 50.5    | 5899 | 51.0         | 5610 | 51.0         | 4890 | 51.0         | 4050 | 51.0         | 3536 |

Table 9: Threshold sensitivity across models and datasets. Acc. denotes Pass@1 accuracy (%), and Tok. denotes average generated tokens.

| Method  | GSM8K            | MATH500           | AIME24            | AMC23             | GPQA-D            | MMLU             |
|---------|------------------|-------------------|-------------------|-------------------|-------------------|------------------|
| Vanilla | 96.89/937        | 95.20/3820        | 66.67/11742       | 95.00/6759        | 56.57/8096        | 77.60/1375       |
| InterSR | <b>96.97/690</b> | <b>95.40/2705</b> | <b>66.67/9450</b> | <b>95.00/4890</b> | <b>57.07/4950</b> | <b>77.60/985</b> |

Table 10: Model-family transfer on Gemma4-12B-it. Each cell reports accuracy/token count.

| Method       | GSM8K           | MATH500          | AIME24           | AMC23            |
|--------------|-----------------|------------------|------------------|------------------|
| Vanilla      | 86.4/743        | 88.6/3824        | 46.7/10953       | 90.0/5699        |
| SwiReasoning | 85.7/688        | 89.0/2764        | 46.7/8063        | 87.5/4791        |
| InterSR      | <b>86.7/517</b> | <b>89.6/2651</b> | <b>50.0/7259</b> | <b>90.0/4126</b> |

Table 11: Focused comparison with a close mode-switching baseline on DeepSeek-R1-Distill-Qwen-7B. Each cell reports accuracy/token count.

From Eq 2:  $s + 2 = 9 / (2.4 - T)$ . Substituting:  
 $9/(4 - T) + 2 = 9/(2.4 - T)$ .  
Solving:  $T = 0.4$ ,  $t = 24$  minutes,  $s = 2.5$  km/h.

Verify:  $(4.5)(2.0) = 9$ . Correct. Now  $s + 0.5 = 3$  km/h, walking time =  $9/3 = 3$  hours = 180 minutes. Total =  $180 + 24 = 204$  minutes.

Wait, let me verify again. First case:  $9/2.5 = 3.6$  hours = 216 mins, total = 240 mins. Second:  $9/4.5 = 2$  hours = 120 mins, total = 144 mins. All consistent. Answer is 204 minutes.

</think>

Aya goes for a 9-kilometer walk every morning and stops at a coffee shop afterwards. We set up:  $9 = s(4 - t/60)$  and  $9 = (s+2)(2.4 - t/60)$ . Solving:  $t = 24$  mins,  $s = 2.5$ . At 3 km/h:  $180 + 24 = \mathbf{204}$  minutes.

## O.2 InterSR Case

Compared to standard reasoning, InterSR is better at allocating resources based on the current reasoning state. It dynamically selects the appropriate reasoning system according to the level of conceptual complexity, thereby avoiding redundant elaboration on simple concepts and focusing computational effort on the core aspects of the problem.

### InterSR Case

<think>

Okay, so I've got this problem here about Aya going for a walk every morning. Let me try to understand what's going on.

</think>

Understanding the Problem: Aya walks 9 kilometers every morning with two different walking speeds. We need to find the total time at  $s + 1/2$  km/h.

Setting Up Equations:

<think>

Let's define:  $s$  = speed,  $t$  = coffee shop time. Equations:  $9/s + t/60 = 4$  and  $9/(s+2) + t/60 = 2.4$ .

</think>

Subtracting:  $18/(s(s+2)) = 1.6$ , so  $s(s+2) = 11.25$ , giving  $s = 2.5$  km/h,  $t = 24$  minutes. At  $s + 0.5 = 3$  km/h: walking time =  $9/3 = 3$  hours = 180 min. Total =  $180 + 24 = \mathbf{204}$  minutes.