

Efficient Thought Space Exploration Through Strategic Intervention

Ziheng Li^{1*}, Hengyi Cai², Xiaochi Wei², Yuchen Li²,
Shuaiqiang Wang², Zhi-Hong Deng^{1†}, Dawei Yin²

¹State Key Laboratory of General Artificial Intelligence, School of Intelligence Science and Technology, Peking University
²Baidu Inc.

{liziheng, zhdeng}@pku.edu.cn, caihengyi@ict.ac.cn, weixiaochi@baidu.com,
{shqiang.wang, yuchenli1230}@gmail.com, yindawei@acm.org

Abstract

While large language models (LLMs) demonstrate emerging reasoning capabilities, current inference-time expansion methods incur prohibitive computational costs by exhaustive sampling. Through analyzing decoding trajectories, we observe that most next-token predictions align well with the golden output, except for a few critical tokens that lead to deviations. Inspired by this phenomenon, we propose a novel Hint-Practice Reasoning (HPR) framework that operationalizes this insight through two synergistic components: 1) a *hinter* (powerful LLM) that provides probabilistic guidance at critical decision points, and 2) a *practitioner* (efficient smaller model) that executes major reasoning steps. The framework’s core innovation lies in **Distributional Inconsistency Reduction (DIR)**, a theoretically-grounded metric that dynamically identifies intervention points by quantifying the divergence between practitioner’s reasoning trajectory and hinter’s expected distribution in a tree-structured probabilistic space. Through iterative tree updates guided by DIR, HPR reweights promising reasoning paths while deprioritizing low-probability branches. Experiments across arithmetic and commonsense reasoning benchmarks demonstrate HPR’s state-of-the-art efficiency-accuracy tradeoffs: it achieves comparable performance to self-consistency and MCTS baselines while decoding only 1/5 tokens, and outperforms existing methods by at most 5.1% absolute accuracy while maintaining similar or lower FLOPs.

Extended version — <https://arxiv.org/abs/2511.10038>

1 Introduction

Large language models (LLMs) have made significant advancements in solving complex reasoning tasks (OpenAI 2023; Wei et al. 2022; Kojima et al. 2022; Huang and Chang 2023), such as mathematical problems (Cobbe et al. 2021) and code generation (Chen et al. 2023). The emergence of these reasoning abilities comes from the scaled pre-training (Kaplan et al. 2020; OpenAI 2024; DeepSeek-AI 2025) and curated prompting strategies for eliciting reasoning (Wei et al. 2022; Kojima et al. 2022; Chen et al. 2023; Wang et al. 2024b; Li et al. 2025; Zheng et al. 2024).

*Work done during internship at Baidu.

†Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Recently, researchers have found that leveraging curated decoding strategies to scale LLM inference-time reasoning expansion can also lead to notable performance gains (Zhang et al. 2025; Wang and Zhou 2024). Typical strategies include extensive sampling, structured rationale topology, and external guidance. For example, self-consistency (Wang et al. 2023) and CoT-decoding (Wang and Zhou 2024) sample multiple rationales and select the most consistent answer, Best-of-N (Gui, Garbacea, and Veitch 2024) uses a score model to select the most reasonable thought from all generations, Tree-of-thoughts (ToT) (Yao et al. 2023) and RAP (Hao et al. 2023) formalize thought exploration as a tree-structure search problem which allows looking ahead and backtracking to make optimal choices, and AdaSwitch (Sun et al. 2024) introduces the collaboration with a stronger LLM to overcome the hard part of the reasoning. However, current reasoning exploration methods leave three unsolved issues: (1) Low token efficiency: pure-sampling methods cannot reuse correct prefixes while tree-structured methods generate numerous intermediate branches but only a few of them can lead to the final answer. Furthermore, when dealing with complex problems, these methods require sampling prohibitively many trajectories to find a correct path. (2) Local exploration guidance: heuristic-based search algorithm easily converges to a “good” path and fails to explore the whole thought space. Although MCTS uses UCT to balance exploration and exploitation, it does not fully exploit token probabilities, resulting in an underestimation of truly valuable branching points. (3) High external dependence: to verify generated steps or overcome overly hard reasoning steps, a common practice is to apply an ad-hoc trained verifier or considerable usage of a stronger LLM reasoner, which will introduce considerable extra cost.

To address these issues, we draw inspiration from human problem solving where students often benefit significantly from just a handful of advisor’s hints, especially for complex problems. Such timely “nudges” can reveal underexplored directions and steer a learner onto the correct trajectory without requiring extensive trial and error. Inspired by this phenomenon, we investigate two questions in this paper: (1) whether LLMs can similarly benefit from such a few well-placed short “hints” to conduct efficient reasoning path exploration. (2) How to place hints into the regular genera-

Query: The number 64 is both a perfect cube and a perfect square. What is the next larger number that is both a perfect cube and a perfect square?

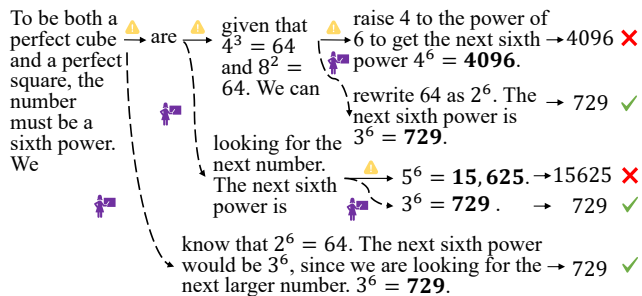


Figure 1: Illustration of the HPR framework. In each iteration, hinter will analyze existing trajectories and find the most promising intermediate state to create an alternative branch. Practitioner then complete the new branch.

tion to provide global and efficient exploration guidance.

For the first question, we conduct an analytical experiment by providing the model with a correct chain-of-thought as the reference and then identifying where its predictions deviate. Our findings show that within the chain-of-thought decoding, an 3B model’s next-token predictions align with those of a 32B model for the majority of tokens, only a small fraction cause significant deviations in reasoning. These sparse “critical tokens” therefore represent prime opportunities for targeted intervention: although it involves a larger model at strategic points, the low intervention frequency yields lower costs than either exhaustive sampling with the smaller model or directly using the larger model.

Building upon the empirical observation that strategic intervention at sparse critical tokens can effectively redirect reasoning trajectories at minimal cost, we accordingly propose a novel Hint-Practice Reasoning (HPR) framework (illustrated in Figure 1), which implements this insight via two distinct roles: (1) a hinter (powerful but costly model) provides targeted hints at sparse critical tokens; (2) a practitioner (efficient smaller model) executes the majority of reasoning steps. Different from existing multi-agent studies (Sun et al. 2024; Gui, Garbacea, and Veitch 2024), HPR hinter selectively corrects the critical mistakes of the practitioner and additionally guides exploration, which reduces the large model’s generation burden. The left challenge lies in identifying where to apply interventions. Intuitively, the practitioner should think twice when it significantly deviates from hinter’s prediction. Furthermore, we expect the practitioner to start exploring from a promising and uncertain (from the global view) partial trajectory to reuse computation and maximize exploration.

We ground this heuristic into a unified probabilistic theory, which characterizes practitioner’s reasoning tree by a distribution and induces **Distributional Inconsistency Reduction (DIR)**, a novel metric that quantifies the change of discrepancy between the whole reasoning tree (rather than a path) and hinter’s expected distribution. Generally, we formalize the reasoning process as a tree-structured probabilistic space where each node represents a partial rea-

soning trajectory. The final answer is derived through aggregation operations (e.g., weighted voting) on leaf nodes. The hinter’s interventions essentially induce tree updates to this structure: when providing hints at critical tokens, the practitioner’s subsequent predictions are dynamically reweighted based on the hinter’s distributional guidance, effectively deprioritizing low-probability branches while amplifying promising trajectories. Concretely, during each iteration to update the tree-structured reasoning trajectory, HPR first analyzes the generated trajectories, employing an adaptive intervention strategy that prioritizes nodes with maximal DIR values (indicating where the practitioner’s predictions diverge most severely from the hinter’s guidance), and then activates the hinter to construct a piece of hint. The practitioner subsequently propagates these hints to complete the remaining generation. This iterative refinement terminates when the maximum iteration count is reached.

Our contributions are threefold:

- We propose HPR, a cross-reasoner collaboration decoding framework for LLM reasoning. Our method provides a new way to organize the reasoning thought space exploration.
- We establish a theoretical guidance **DIR** that unifies path exploration quality and hinter-practitioner distribution alignment, providing principled guidance for reasoning tree updating.
- Comprehensive experiments across arithmetic and commonsense reasoning tasks demonstrate HPR’s state-of-the-art efficiency-accuracy tradeoff, requiring only 1/5 tokens to get comparable performance against baselines, and improving the performance by at most 5.1% while maintaining a similar or lower FLOPs.

2 Related Work

2.1 Inference-time LLM Reasoning

Recent studies have found increasing inference-time computations can improve the output quality of LLMs, particularly for complex tasks that needs long-chain reasoning (Wu et al. 2024; Brown et al. 2024; Fu et al. 2024; Chu et al. 2024). Chain-of-thought prompting (Wei et al. 2022) is the first to demonstrate intermediate rationales can boost the reasoning ability of LLMs. Zero-shot CoT (Kojima et al. 2022) employs a reasoning instruction to activate the rationale generation and avoid the dependence on few-shot demonstrations. Subsequent research focuses on further improving the quality of the rationale. UAG (Yin et al. 2024) identifies the risky steps of the reasoning path by the uncertainty and adjust demonstrations to eliminate the risk. Zhou et al. (2023); Press et al. (2023) prompt the LLM to decompose the query into easier sub-parts. Madaan et al. (2023); Manakul, Liusie, and Gales (2023) self-correct the reasoning mistakes by tasks-specific prompting. Kumar et al. (2024) internalized the self-correction ability by reinforcement learning. Cobbe et al. (2021); Wang et al. (2024a); Li et al. (2023) propose to train a verifier to select the optimal rationale. Our approach exploit the potentials of autoregressive decoding, which is task-agnostic and does not require curated demonstration, prompts, nor ad-hoc trained models.

2.2 Thought Space Exploration

The autoregressive decoding nature of LLM poses challenges in finding the globally optimal reasoning path. To address this, researchers utilized the probabilistic modeling nature of LLMs to explore diverse thoughts. Self-consistency (Wang et al. 2023) adopts a sample-and-marginalize decoding strategy, finding the most consistent answer among multiple sampled paths. CoT-decoding (Wang and Zhou 2024) create diverse paths by enumerating the top-K candidates of the starting token and select the most plausible path by comparing answer’s confidence score. Further advancing this line of research, Tree-of-thoughts and Graph-of-Thoughts (Yao et al. 2023; Long 2023; Besta et al. 2024; Mo and Xin 2024; Zhang et al. 2024) generalizes the topology of rationale into the tree or graph structures by allowing LLM to try multiple reasoning branches at each step. RAP (Hao et al. 2023) regard reasoning as planning and introduce Monte Carlo Tree Search (MCTS) to organize the reasoning exploration space, which iteratively expand new paths under the guide of reward models. AdaSwitch (Sun et al. 2024) formulates rationale generation as a multi-agent task. When the local model finds an error in the reasoning process, it asks for a more powerful cloud model to correct the mistake. Our approach follows tree-structure thought exploration paradigm. We achieve more efficient reasoning by delegating part of generation to a stronger reasoner and further exploiting LLM’s probabilistic information to organize the tree growth.

3 Sparse Disparity in LLM Reasoning

We start by briefly introducing the reasoning procedure for an LLM from a probabilistic view. Given a query \mathbf{x} , our goal is to predict the answer’s distribution $P_\phi(\mathbf{y}|\mathbf{x})$ using an autoregressive model parameterized by ϕ . For complex problems, answer y cannot be directly inferred from the query. We need to first infill the intermediate rationale $\mathbf{r} = r_1 r_2 \cdots r_l$ and then get the final answer:

$$P_\phi(\mathbf{y}|\mathbf{x}) = \sum_{\mathbf{r}} P_\phi(\mathbf{r}|T(\mathbf{x}))P_\phi(\mathbf{y}|T(\mathbf{x}), \mathbf{r}), \quad (1)$$

where $T(\cdot)$ represents a prompt template (e.g. appending “Let’s think step by step” to the end of \mathbf{x}). For simplicity, we omit $T(\cdot)$ in the remaining paper. Although generating rationales is generally easier than directly generating the answer, LLMs can still deviate from the correct reasoning path, particularly for smaller models. Our study reveals that the deviations are primarily caused by a small set of hard tokens scattered along the rationale sequence.

To give an intuitive illustration, we conduct an analytic experiment on MATH (Hendrycks et al. 2021) dataset that measures the disparity between a small LLM ϕ (Qwen2.5-3B-Instruct) and a high-performance oracle θ (Qwen2.5-32B-Instruct). Specifically, we feed model ϕ with an output from model θ to make predictions at each position, collect the inconsistent ones, and continue generation from the inconsistent positions using ϕ to see whether it leads to an incorrect final answer (deviation).

Figure 2 shows a representative case: among a generation of 500 tokens, only 3 of them cause model ϕ to deviate from

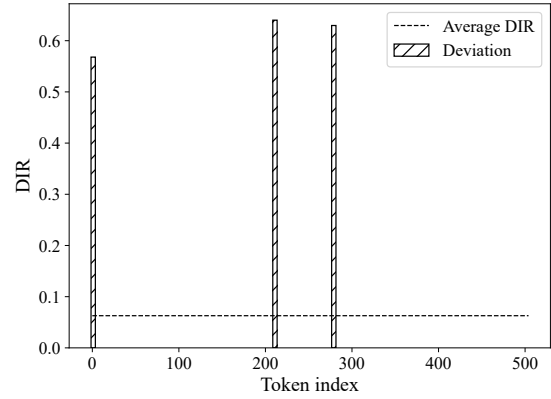


Figure 2: We depict the positions where Qwen2.5-3B-Instruct deviates from the correct rationale using the bars. Additionally, we annotate their DIR value by the heights, which significantly exceeds the average line, showing the indicative power of DIR.

the correct reasoning paths. These positions show high DIR values, inspiring the design of an intervention strategy.

4 Hint-Practice Reasoning

We first introduce our targeted tree-structure decoding workflow in Section 4.1. In Section 4.2, we introduce the core innovation of HPR—Distributional Inconsistency (DI)—a quantified measure where higher values indicate greater potential discrepancy between the current reasoning tree outcomes and the hinter model’s expected results. To minimize DI under constrained intervention budgets, Section 4.3 presents our methodology for computing Distributional Inconsistency Reduction (DIR), which estimates each node’s potential contribution to DI minimization.

4.1 Targeted Reasoning Tree Expansion

HPR aims to allocate test-time computation to the reasoning state where it is truly needed. We adopt a tree-structured search to progressively explore the thought space and reuse high-quality rationale prefixes, avoiding redundant computation. To aggregate reasoning paths, we apply weighted voting, a proven method for test-time scaling (Zhang et al. 2025; Wang et al. 2023).

However, prior tree-based methods (Yao et al. 2023; Mo and Xin 2024; Hao et al. 2023) often generate many incomplete paths, which are excluded from final voting, leading to inefficiency. To address this, we ensure every expansion yields a complete rationale by letting the practitioner finish each branch.

HPR employs a “grow-from-the-middle” strategy guided by the hinter. As illustrated in Figure 3, the tree is expanded iteratively through four phases:

- **Select:** Compute DIR to identify the node where the hinter can contribute the most novel information; this becomes the *critical node*.

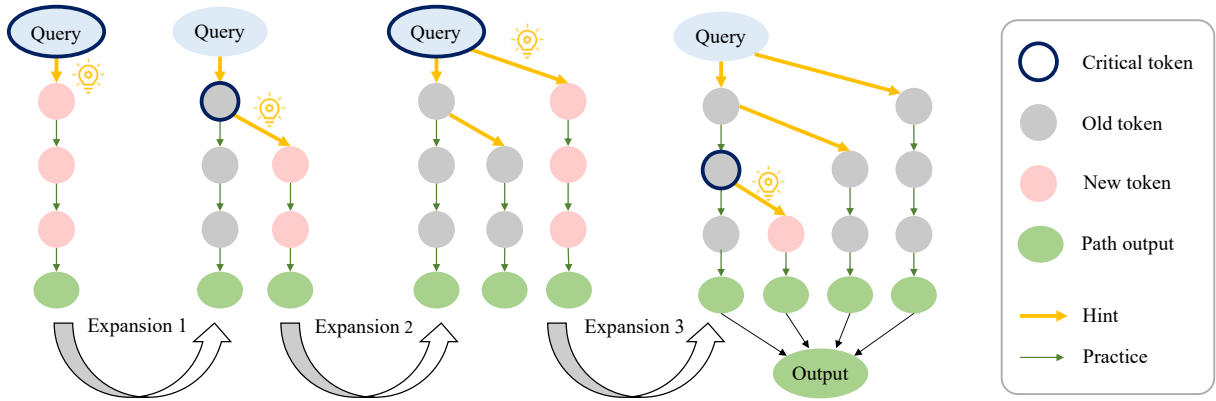


Figure 3: Illustration of the HPR iterative procedure. Each circular node denotes a token. HPR will generate multiple reasoning paths, starting by generating a single chain of thought. In each subsequent iteration, HPR selects a critical token (red circled) and expand a new path. This process repeats until the maximum iteration limit is reached.

- **Hint:** Feed the hinter with the critical node’s prefix to generate a hint.
- **Practice:** The practitioner completes the reasoning via greedy decoding based on the hint.
- **Analyze:** Record output distributions to support DIR calculation for the next iteration.

HPR iteratively augments the reasoning tree through strategic interventions to progressively align the tree with reasoning trees sampled from the hinter model’s distribution. To operationalize this objective, we need a metric for evaluating the global divergence between the current reasoning tree and the target distribution, which we will introduce in the next section.

4.2 Distributional Inconsistency

To enable the comparison between the reasoning tree and a distribution, we first need to establish a characterization distribution that describes the already explored trajectories. Furthermore, since we have access to a more accurate distribution, we can incorporate hinter’s information into the characterization distribution to make preliminary rectification.

Definition 1 (Characterization Distribution) Given a query \mathbf{x} , a hinter model θ , and a practitioner model ϕ , suppose there have been m reasoning paths sampled from P_ϕ of length n : $V = (\mathbf{r}_{1:n}^1, \mathbf{r}_{1:n}^2, \dots, \mathbf{r}_{1:n}^m)$, allowing overlap at prefix, characterization distribution Q_V is defined as:

$$Q_V(r_i|\mathbf{x}, \mathbf{r}_{1:i-1}) = \frac{P_\theta(r_i|\mathbf{x}, \mathbf{r}_{1:i-1})}{\sum_{r'_i \in N_V(\mathbf{r}_{1:i-1})} P_\theta(r'_i|\mathbf{r}_{1:i-1})}, \quad (2)$$

where $N_V(\mathbf{r}_{1:i-1})$ denotes all successors of $\mathbf{r}_{1:i-1}$ in V .

Intuitively, Q_V is the distribution that embeds P_θ onto the support set V . It is built upon P_ϕ and reallocates the token probability mass towards P_θ . Sampling from Q_V is equivalent to sampling from P_θ while rejecting rationales not in V . HPR samples from Q_V to get refined outputs by marginalizing over all explored reasoning paths:

$$P^{HPR}(\mathbf{y}|\mathbf{x}; V) = \sum_{\mathbf{r} \in V} Q_V(\mathbf{r}|\mathbf{x}) P_\phi(\mathbf{y}|\mathbf{x}, \mathbf{r}). \quad (3)$$

There are four features of Q_V : (1) The two components of Q_V —support set V and probability value P_θ —are disentangled, which facilitates the analysis of the collaboration between the two models; (2) Q_V provides intermediate probabilistic information which helps guide the expansion of the reasoning tree; (3) Q_V equals 1 at each non-branching node, simplifying computational implementation; (4) Q_V enables the measurement of the divergence between the observed paths and hinter’s distribution.

Definition 2 (Distributional Inconsistency) The distributional inconsistency between rationales V and P_θ is defined as the reversed Kullback–Leibler divergence between surrogate sampling distribution and the hinter distribution: $D_{KL}(Q_V||P_\theta)$.

$D_{KL}(Q_V||P_\theta)$ reflects how much valuable space has not been covered by the current reasoning tree. It provides a global search guidance as it considers the contribution of all tree nodes.

4.3 Distributional Inconsistency Reduction

We expect the Q_V to be as close as possible to P_θ . To achieve this, we adopt a greedy strategy that lets practitioner generate a rationale \mathbf{v} maximizing the reduction of KL divergence $D_{KL}(Q_V||P_\theta) - D_{KL}(Q_{V \cup \{\mathbf{v}\}}||P_\theta)$ at each iteration.

During the “select” phase, we need to select a node \mathbf{z} such that the newly expanded branch from \mathbf{z} yields a high reduction. However, trajectory-based reduction cannot be efficiently evaluated for each node. Therefore, we reorganized the calculation and derived the final node-based formulation of Distributional Inconsistency Reduction (DIR). For conciseness, we omit \mathbf{x} in the condition.

Definition 3 (Distributional Inconsistency Reduction) Let \mathbf{z} be a prefix of some reasoning path in V , \mathbf{v} be the

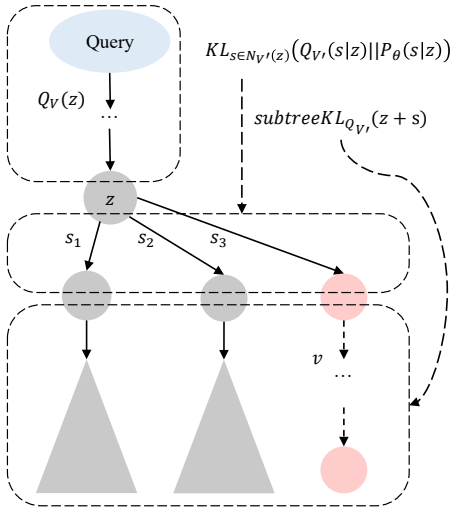


Figure 4: Illustration of three terms in node-version DIR. The pink circles denote the newly generated path.

newly generated path starting from z , $V' = V \cup \{v\}$.

$$\begin{aligned} \text{DIR}(z; V, P_{\theta}) &= D_{KL}(Q_V || P_{\theta}) - D_{KL}(Q_{V \cup \{v\}} || P_{\theta}) \\ &= Q_V(z) \left[\Delta D_{KL_{s \in N(z)}}(Q(s|z) || P_{\theta}(s|z)) + \right. \\ &\quad \left. \sum_{s \in N_{V'}(z)} \Delta Q(s|z) \text{subtreeKL}_{Q_{V'}}(z+s) \right], \end{aligned} \quad (4)$$

where ΔD_{KL} denotes the difference between pre-expansion KL $D_{KL_{s \in N(z)}}(Q(s|z) || P_{\theta}(s|z))$ and post-expansion KL $D_{KL_{s \in N_{V'}(z)}}(Q_{V'}(s|z) || P_{\theta}(s|z))$, $\Delta Q(s|z) = Q_V(s|z) - Q_{V'}(s|z)$, $z' = z + s$ means concatenating prefix with token s , $\text{subtreeKL}_{Q_{V'}}(z') = \sum_{z' \subset v \in V} Q_{V'}(v - z' | z') \log \frac{Q(v - z' | z')}{P_{\theta}(v - z' | z')}$.

Details of the derivation are included in Appendix. The three main terms in DIR correspond to contributions from the prefix, next-token, and sub-tree, respectively, as illustrated in Figure 4. These terms reflect three characteristics of a good expansion starting point z :

- DIR favors prefixes with high probability, ensuring that new paths are built on a solid and reasonable foundation.
- DIR favors those under-explored nodes. The next-token term is particularly significant when there is a large hinter-practitioner gap or when the practitioner exhibits uncertainty ($\Delta D_{KL_{s \in N(z)}}(Q(s|z) || P_{\theta}(s|z))$), aligning with our core motivation.
- DIR favors the expansion where the path $v - z'$ is likely to occur (the subtreeKL of a chain equals to the sum of $-\log P_{\theta}(\cdot)$).

In our implementation, we normalize the subtree KL by dividing it by the average path length among the sub-tree to avoid DIR being biased. During the ‘analyze’ phase, HPR will evaluate and store probabilities of Top-K next tokens

Algorithm 1: Hint-Practice Reasoning

Require: practitioner ϕ , hinter θ

Input: query x , max_iterations k

$P = V = A = \emptyset$, nodes = $\{x\}$

repeat

critical node $z = \arg \max_{z \in \text{nodes}} \text{DIR}(z; V, P)$

hint $h \sim P_{\theta}(\cdot | z)$

practice $p \sim P_{\phi}(\cdot | z + h)$

$v = h + p$

answer $a \sim P_{\phi}(\cdot | z + v)$

$V = V \cup \{z + v\}$

$A = A \cup \{a\}$

$u_i = H(P_{\phi}(\cdot | z + v_{0:i}))$, for $i = 1, 2, \dots, \text{len}(v)$

$v' = v_{1:i}$, $i = \min\{i | \text{Top-}U(\{u_j\}_{j=1}^{\text{len}(v)}) \subset u_{1:i}\}$

$P' = \{P_i\}_{i=1}^{\text{len}(v')} = \text{Top-}K(P_{\theta}(v' | z))$

for $i = 1$ **to** $\text{len}(v')$ **do**

nodes = nodes $\cup \{z + v_{1:i}\}$

$P = P \cup P'$

end for

until reach max iterations

return Weighted-vote(A, Q_V)

($K = 32$ in experiments) for each expanded node. The only unknown term in DIR computation is the probability of new path $v - z'$, as it must be calculated before expansion. As an approximation, We use the average log probabilities of the nearest 32 tokens in the prefix and the sibling greedy decoding path.

To calculate DIR, we need to leverage the hinter to re-evaluate the rationales to collect P_{θ} . This process requires only a single forward operation to obtain all necessary probabilities, avoiding token-by-token decoding, thus introducing minimal latency. To be more efficient, for newly generated paths, we restrict the selection of critical tokens to the minimal prefix that contains the Top- U most uncertain tokens with regard to P_{ϕ} (HPR uses information entropy H) because we empirically find a strong correlation between P_{ϕ} uncertainty and DIR value. In practice, we find that the performance and the efficiency are insensitive to U . We set $U = 3$ in all experiments. The complete procedure of HPR is presented in Algorithm 1.

5 Experiments

5.1 Setup

Benchmarks We evaluate our approach on arithmetic reasoning and commonsense reasoning tasks. Arithmetic reasoning includes GSM8K (Cobbe et al. 2021), AQUA-RAT (Chen et al. 2023), and MATH (Hendrycks et al. 2021). These datasets are among the more challenging ones used in prior studies (Sun et al. 2024; Ma et al. 2025). For commonsense reasoning, we select CSQA (Talmor et al. 2019) and StrategyQA (Geva et al. 2021), both of which require implicit reasoning.

Implementation We use models from Qwen2.5 series (Yang et al. 2024): 3B/7B for practitioner and 14B/32B

Method	Arithmetic			Commonsense		Inference Cost			Efficiency (REE)
	GSM8K	AQUA-RAT	MATH	CSQA	StrategyQA	#Tokens (P)	#Tokens (H)	FLOPs (10 ¹²)	
CoT	85.3	64.2	53.0	74.5	59.5	320.8	-	1.6	-
CoT-SC@5	88.9	69.2	59.8	76.1	59.9	1664.8	-	8.4	0.82
CoT-SC@15	90.2	73.2	63.4	78.4	60.2	4720.8	-	23.6	0.42
ToT@5	84.8	68.2	53.4	75.0	61.7	6916.8	-	34.6	0.06
TouT@5	85.3	71.2	53.6	74.9	62.0	6873.0	-	34.4	0.10
MCTS@5	87.4	69.7	58.9	75.1	60.0	5714.6	-	28.6	0.17
<i>Guided by Qwen2.5-14B-Instruct</i>									
Hint-SC@5 (UB)	94.2	84.5	72.7	84.0	75.8	-	1676.4	45.2	0.55
CoT-BoN@5	88.1	69.3	56.3	76.1	61.6	1664.8	1664.8	19.6	0.26
AdaSwitch	89.9	69.3	59.7	75.0	60.5	1013.8	185.0	10.0	0.68
HPR@5 (Ours)	91.0	73.2	62.1	78.0	62.0	936.8	124.2	8.0	1.49
<i>Guided by Qwen2.5-32B-Instruct</i>									
Hint-SC@5 (UB)	95.2	87.0	75.4	86.0	76.1	-	1638.0	103.2	0.26
CoT-BoN@5	88.6	69.7	56.8	76.2	62.0	1664.8	1664.8	34.6	0.16
AdaSwitch	89.9	68.5	60.6	75.0	60.7	1014.4	171.4	15.8	0.41
HPR@5 (Ours)	91.8	74.8	63.2	78.9	63.6	975.5	124.1	12.8	1.02

Table 1: Results using Qwen2.5-3B-Instruct as practitioner. “@ k ” means that k reasoning paths are sampled during inference. #Tokens (P)/(H) denote the number of tokens generated from practitioner and hinter, respectively (For BoN, #Tokens (H) denotes the number of evaluated tokens). UB means the accuracy upper bound within the same group. We mark a result bold if it outperforms all practitioner-only baselines and guided-search baselines within the same group. Cost and efficiency are averaged across 5 datasets.

for hinter. To account for differences in task complexity, we set the HPR hint length to 32 tokens for arithmetic reasoning and 16 tokens for commonsense reasoning. We use vLLM (Kwon et al. 2023) library to provide efficient inference serving. For all experiments, we use “Let’s think step by step” prompting (Kojima et al. 2022) to perform zero-shot CoT. We set the temperature to 0.7 for all non-greedy generations and average results across three different seeds.

Baselines We compare HPR with single-trajectory CoT, sampling-based method Self-Consistency (SC) (Wang et al. 2023), tree-structure search methods ToT (Yao et al. 2023), TouT (Mo and Xin 2024), and MCTS (Hao et al. 2023), as well as externally guided baselines Best-of-N (BoN) (Song et al. 2024) and AdaSwitch (Sun et al. 2024). Besides, we report a pure hinter baseline to show the performance upper bound of externally guided methods. We use the hinter as the reward model both for BoN and MCTS. For MCTS, we follow Marco-o1 (Zhao et al. 2024) to use the average confidence as the reward score, and ensemble all paths via majority voting. ToT, TouT, and MCTS are allowed to sample 5 different actions at each step in our experiments. The hyperparameters for AdaSwitch are tuned to ensure a comparable token budget with HPR.

Metric We extract the final answer from the LLM’s output and calculate exact match accuracy. To distinguish the contributions of the practitioner and the hinter, we recorded the average number of tokens involved by each during inference (averaged across samples). We calculate the FLOPs (floating point operations) to measure the cost of different approaches. Following Kaplan et al. (2020), the FLOPs of

decoding a sequence of length L approximately equal $2NL$, where N is the number of non-embedding parameters of the LLM. To evaluate the reasoning-expansion efficiency (REE), we devise an integrated measure by computing the normalized gain-cost ratio:

$$REE = \frac{ACC - A_0}{FLOPs - F_0} * \frac{F_0}{100\%}, \quad (5)$$

where A_0 and F_0 are the accuracy and FLOPs of single CoT baseline. Given that decoding operation is typically more time-consuming than evaluation due to bandwidth limitations, we follow API cost conventions and weight evaluation FLOPs by a factor of 1/4 for fair comparison.

5.2 Main Results

The results of Qwen2.5-3B/7B-Instruct are shown in Table 1 and Table 2, which represent the settings of high/low practitioner-hinter capability gap, respectively. We can make several observations:

First, HPR shows the highest token and computational efficiency. HPR@5 consumes only about 2/3 tokens of CoT-SC@5 and maintains similar FLOPs, though it involves a larger model. CoT-SC strategy has to generate 5x tokens to reach similar performance. When compared with other tree search methods, the advantage of token efficiency and FLOPs increases to 3x-5x, which demonstrates that our targeted tree expansion managed to avoid wasteful generation. Notably, hinter is involved in very little generation, which means deploying one hinter can simultaneously provide service for multiple practitioners, saving high-memory devices.

Second, HPR shows consistent improvement over all baselines across all benchmarks. It outperforms sampling

Method	Arithmetic			Commonsense		Inference Cost			Efficiency (REE)
	GSM8K	AQUA-RAT	MATH	CSQA	StrategyQA	#Tokens (P)	#Tokens (H)	FLOPs (10^{12})	
CoT	90.8	80.0	66.8	81.5	69.6	327.4	-	4.0	-
CoT-SC@5	93.1	83.9	71.8	82.1	71.5	1672.6	-	21.8	0.62
CoT-SC@15	93.4	87.0	75.3	82.6	72.0	5042.2	-	65.6	0.28
ToT@5	89.7	83.7	66.1	80.5	70.1	7542.4	-	98.0	0.01
TouT@5	90.1	82.2	66.5	81.1	71.0	7548.0	-	98.0	0.02
MCTS@5	91.6	84.6	71.7	81.1	71.4	6142.6	-	79.8	0.12
<i>Guided by Qwen2.5-14B-Instruct</i>									
Hint-SC@5 (UB)	94.2	84.5	72.7	84.0	75.8	-	1676.4	43.2	0.43
CoT-BoN@5	92.0	84.6	70.4	81.1	71.9	1672.6	1672.6	33.0	0.31
AdaSwitch	92.6	85.0	69.2	82.1	71.5	1055.2	140.8	17.6	0.69
HPR@5 (Ours)	93.0	85.5	71.8	82.7	72.5	936.8	124.0	15.6	1.16
<i>Guided by Qwen2.5-32B-Instruct</i>									
Hint-SC@5 (UB)	95.2	87.0	75.4	86.0	76.1	-	1638.0	103.2	0.25
CoT-BoN@5	92.5	84.6	70.6	81.4	72.4	1672.6	1672.6	48.0	0.23
AdaSwitch	93.0	85.8	69.7	82.7	71.5	1045.6	127.6	21.6	0.64
HPR@5 (Ours)	93.6	86.2	72.7	83.4	73.4	944.4	124.2	20.0	1.03

Table 2: Results using Qwen2.5-7B-Instruct as practitioner.




Figure 5: Accuracy versus FLOPs per instance on the MATH dataset using Qwen2.5-3B/14B-Instruct with varying numbers of reasoning paths.

and tree-structured baselines by 5.1% at most when searching the same number of reasoning paths while costing less computation, which demonstrates that introducing the aid of a stronger model is an economical way to boost performance. Compared with BoN and AdaSwitch, HPR also achieves an improvement of 5.1% at most, which we attribute to the effective exploration mechanism in HPR.

Third, HPR’s performance is close to the pure hinter upper bound in the low-capability-gap setting while consuming 1/5-1/3 FLOPs. This indicates that, in certain cases, HPR could be a more efficient alternative compared to deploying a single strong model.

5.3 Performance-Cost Analysis

Previous experiments primarily focus on the 5-trajectory setting. In this section, we comprehensively compare HPR with




Figure 6: Results of ablation study using Qwen2.5-3B/14B-Instruct.

baseline methods under different computational budgets by varying the number of trajectories. The detailed experimental settings are provided in Table 1 in the Appendix. As shown in Figure 5, HPR consistently achieves the highest accuracy across all FLOPs levels, demonstrating the general effectiveness of our DIR-guided tree expansion strategy.

5.4 Ablation Study

HPR experiences four phases during one iteration. In this section, we will examine the effect of these phases by three variants: (1) disable DIR and randomly select from all non-terminal nodes with equal probabilities; (2) remove the hint generation and direct call practitioner to continue generation from selected critical node; (3) skip the “analyze” phase and use practitioner’s probability to calculate Q_V .

The results are shown in Figure 6. We can find that eliminating any phase of HPR leads to a significant performance

Problem: Find the maximum volume of a cone that fits inside a sphere of radius 1.

Initial trajectory: Let’s think step by step. We want to find the maximum volume of a cone that fits inside a sphere of radius 1. The volume V of a cone is given by the formula: $V = \frac{1}{3}\pi r^2 h$ where r is the radius of the base of the cone and h is the height of the cone. For a cone inscribed in a sphere, the relationship between the radius of the sphere, the radius of the base of the cone, and the height of the cone can be derived using the Pythagorean theorem. We have: $1^2 = r^2 + (h - r)^2$. **Expanding and simplifying the equation: ...**

Expansion 1: Let’s think step by step ... So, we have: $1^2 = r^2 + (h - 1)^2$. **This simplifies to: $r^2 = 2h - h^2$.** Now, substitute $r^2 = 2h - h^2$ into the volume formula: $V = \frac{1}{3}\pi(2h - h^2)h = \frac{1}{3}\pi(2h^2 - h^3)$. To find the maximum volume, we need to take the derivative of V with respect to h and set it to zero: $\frac{dV}{dh} = \frac{1}{3}$, which gives $\pi(4h - 3h^2)$, $h(4 - 3h) = 0$. This gives us two solutions: $h = 0$ (which is not physically meaningful for a cone) and $h = \frac{4}{3}$. **However, since $h = \frac{4}{3}$ is outside the valid range for a cone inscribed in a sphere of radius 1, ...**

Expansion 2: Let’s think step by step ... and $h = \frac{4}{3}$. **We need to check if this value of h gives a maximum volume.** We can do this by checking the second derivative ... So, the answer is $\frac{32\pi}{81}$.

Table 3: Case study of solving math problem. Red text indicates the mistake of practitioner. Blue text indicates the hint. Gray text indicates reused prefix.




Figure 7: Performance of HPR@5 with varying hint length using Qwen2.5-3B/14B-Instruct.

drop. Among them, the most important one is DIR-guided selection, as it ensures the hinter provides guidance at the most critical points, forming the foundation of the HPR framework. For the variant without hint, the practitioner knows where to expand from but struggles to decode the correct token. The “w/o Analyze” performs slightly better than other two, as it retains most of HPR’s mechanisms but shifts alignment toward the practitioner’s inherent probabilities.

5.5 Impact of Hint Length

The hint length determines the extent to which the hinter contributes to the generation process. To analyze its impact, we test different hint lengths ranging from 1 to 32. The results are shown in Figure 7.

We observe the most significant performance gain when increasing the hint length from 1 to 4 tokens. Beyond this point, the improvement continues at a slower, approximately linear rate. These results suggest that hinter’s primary role is to guide the reasoning direction—demonstrating that even short hints can be highly effective. Moreover, longer hints consistently enhance performance, gradually pushing the

practitioner’s output closer to that of the hinter model.

5.6 Case Study

We examine examples of math problem solving to illustrate the effectiveness of HPR. As shown in Table 3, the initial rationale fails to derive the correct equation. The hinter identifies and corrects the wrong token “ r ”, giving practitioner a second chance. Subsequently, the practitioner makes another mistake regarding the validity of h , and the hinter again intervenes to adjust the reasoning direction. Finally, the practitioner successfully arrives at the correct answer. This case demonstrates that timely and strategic interventions from the hinter can effectively and efficiently steer the reasoning process toward the correct solution.

6 Conclusion

In this paper, we propose a novel reasoning framework (HPR) that enables timely and efficient intervention during decoding to overcome challenging “critical tokens” and correct the reasoning direction. HPR is theoretically backed up by our newly proposed Distributional Inconsistency Reduction that measures the disparity between the already explored trajectories and the targeted hinter distribution. HPR leverages DIR to guide the reasoning expansion node selection and iteratively expand the reasoning tree. Extensive experiments on arithmetic and commonsense reasoning tasks demonstrate the effectiveness of HPR, highlighting its ability to enhance reasoning accuracy while maintaining computational efficiency.

7 Future Work

HPR provides an elegant framework for integrating the strengths of two models, motivating us to explore extensions of model collaboration beyond the reasoning dimension. One promising direction is to adapt HPR to multi-agent collaboration, enabling more flexible and coordinated problem-solving. In addition, we are studying the vocabulary alignment for DIR computation to support the model combinations from different model families.

References

- Besta, M.; Blach, N.; Kubicek, A.; Gerstenberger, R.; Podstawski, M.; Gianinazzi, L.; Gajda, J.; Lehmann, T.; Niewiadomski, H.; Nyczyk, P.; and Hoefler, T. 2024. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. *Proceedings of the AAAI Conference on Artificial Intelligence*, (16). Number: 16.
- Brown, B.; Juravsky, J.; Ehrlich, R.; Clark, R.; Le, Q. V.; Ré, C.; and Mirhoseini, A. 2024. Large Language Monkeys: Scaling Inference Compute with Repeated Sampling. ArXiv:2407.21787 [cs].
- Chen, W.; Ma, X.; Wang, X.; and Cohen, W. W. 2023. Program of Thoughts Prompting: Disentangling Computation from Reasoning for Numerical Reasoning Tasks. *Transactions on Machine Learning Research*.
- Chu, Z.; Chen, J.; Chen, Q.; Yu, W.; He, T.; Wang, H.; Peng, W.; Liu, M.; Qin, B.; and Liu, T. 2024. Navigate through Enigmatic Labyrinth A Survey of Chain of Thought Reasoning: Advances, Frontiers and Future. ArXiv:2309.15402 [cs].
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. ArXiv:2110.14168 [cs].
- DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. ArXiv:2501.12948.
- Fu, Y.; Chen, J.; Zhu, S.; Fu, Z.; Dai, Z.; Qiao, A.; and Zhang, H. 2024. Efficiently Serving LLM Reasoning Programs with Certindex. ArXiv:2412.20993 [cs].
- Geva, M.; Khashabi, D.; Segal, E.; Khot, T.; Roth, D.; and Berant, J. 2021. Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies. ArXiv:2101.02235.
- Gui, L.; Garbacea, C.; and Veitch, V. 2024. BoNBon Alignment for Large Language Models and the Sweetness of Best-of-n Sampling. In Globersons, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J. M.; and Zhang, C., eds., *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Hao, S.; Gu, Y.; Ma, H.; Hong, J.; Wang, Z.; Wang, D.; and Hu, Z. 2023. Reasoning with Language Model is Planning with World Model. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Singapore.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. In Vanschoren, J.; and Yeung, S.-K., eds., *NeurIPS Datasets and Benchmarks*.
- Huang, J.; and Chang, K. C.-C. 2023. Towards Reasoning in Large Language Models: A Survey. ArXiv:2212.10403 [cs].
- Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T. B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; and Amodei, D. 2020. Scaling Laws for Neural Language Models. *ArXiv*.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large Language Models are Zero-Shot Reasoners. *Advances in Neural Information Processing Systems*.
- Kumar, A.; Zhuang, V.; Agarwal, R.; Su, Y.; Co-Reyes, J. D.; Singh, A.; Baumli, K.; Iqbal, S.; Bishop, C.; Roelofs, R.; Zhang, L. M.; McKinney, K.; Shrivastava, D.; Paduraru, C.; Tucker, G.; Precup, D.; Behbahani, F.; and Faust, A. 2024. Training Language Models to Self-Correct via Reinforcement Learning. arXiv:2409.12917.
- Kwon, W.; Li, Z.; Zhuang, S.; Sheng, Y.; Zheng, L.; Yu, C. H.; Gonzalez, J.; Zhang, H.; and Stoica, I. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP '23*. New York, NY, USA.
- Li, Y.; Cai, H.; Kong, R.; Chen, X.; Chen, J.; Yang, J.; Zhang, H.; Li, J.; Wu, J.; Chen, Y.; Qu, C.; Kong, K.; Ye, W.; Su, L.; Ma, X.; Xia, L.; Shi, D.; Zhao, J.; Xiong, H.; Wang, S.; and Yin, D. 2025. Towards AI Search Paradigm. arXiv:2506.17188.
- Li, Y.; Lin, Z.; Zhang, S.; Fu, Q.; Chen, B.; Lou, J.-G.; and Chen, W. 2023. Making Language Models Better Reasoners with Step-Aware Verifier. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. Toronto, Canada.
- Long, J. 2023. Large Language Model Guided Tree-of-Thought. ArXiv:2305.08291 [cs].
- Ma, C.; Zhao, H.; Zhang, J.; He, J.; and Kong, L. 2025. Non-myopic Generation of Language Models for Reasoning and Planning. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; Gupta, S.; Majumder, B. P.; Hermann, K.; Welleck, S.; Yazdanbakhsh, A.; and Clark, P. 2023. Self-Refine: Iterative Refinement with Self-Feedback. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *NeurIPS*.
- Manakul, P.; Liusie, A.; and Gales, M. 2023. SelfCheck-GPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Singapore.
- Mo, S.; and Xin, M. 2024. Tree of Uncertain Thoughts Reasoning for Large Language Models. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2024, Seoul, Republic of Korea, April 14-19, 2024*. IEEE.
- OpenAI. 2023. GPT-4 Technical Report. ArXiv:2303.08774 [cs].

- OpenAI. 2024. OpenAI o1 System Card. arXiv:2412.16720.
- Press, O.; Zhang, M.; Min, S.; Schmidt, L.; Smith, N.; and Lewis, M. 2023. Measuring and Narrowing the Compositionality Gap in Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore: Association for Computational Linguistics.
- Song, Y.; Wang, G.; Li, S.; and Lin, B. Y. 2024. The Good, The Bad, and The Greedy: Evaluation of LLMs Should Not Ignore Non-Determinism. ArXiv:2407.10457 [cs].
- Sun, H.; Wu, J.; Cai, H.; Wei, X.; Feng, Y.; Wang, B.; Wang, S.; Zhang, Y.; and Yin, D. 2024. AdaSwitch: Adaptive Switching between Small and Large Agents for Effective Cloud-Local Collaborative Learning. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.
- Talmor, A.; Herzig, J.; Lourie, N.; and Berant, J. 2019. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Minneapolis, Minnesota: Association for Computational Linguistics.
- Wang, P.; Li, L.; Shao, Z.; Xu, R.; Dai, D.; Li, Y.; Chen, D.; Wu, Y.; and Sui, Z. 2024a. Math-Shepherd: Verify and Reinforce LLMs Step-by-step without Human Annotations. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*. Bangkok, Thailand.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q. V.; Chi, E. H.; Narang, S.; Chowdhery, A.; and Zhou, D. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Wang, X.; and Zhou, D. 2024. Chain-of-Thought Reasoning Without Prompting. ArXiv:2402.10200 [cs].
- Wang, Y.; Zhao, S.; Wang, Z.; Huang, H.; Fan, M.; Zhang, Y.; Wang, Z.; Wang, H.; and Liu, T. 2024b. Strategic Chain-of-Thought: Guiding Accurate Reasoning in LLMs through Strategy Elicitation. ArXiv:2409.03271 [cs].
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.; Le, Q. V.; and Zhou, D. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems*.
- Wu, Y.; Sun, Z.; Li, S.; Welleck, S.; and Yang, Y. 2024. Inference Scaling Laws: An Empirical Analysis of Compute-Optimal Inference for Problem-Solving with Language Models. ArXiv:2408.00724 [cs].
- Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; Lin, H.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Lin, J.; Dang, K.; Lu, K.; Bao, K.; Yang, K.; Yu, L.; Li, M.; Xue, M.; Zhang, P.; Zhu, Q.; Men, R.; Lin, R.; Li, T.; Xia, T.; Ren, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Qiu, Z. 2024. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115*.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T. L.; Cao, Y.; and Narasimhan, K. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. ArXiv:2305.10601 [cs].
- Yin, Z.; Sun, Q.; Guo, Q.; Zeng, Z.; Li, X.; Dai, J.; Cheng, Q.; Huang, X.; and Qiu, X. 2024. Reasoning in Flux: Enhancing Large Language Models Reasoning through Uncertainty-aware Adaptive Guidance. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*. Bangkok, Thailand.
- Zhang, J.; Mo, F.; Wang, X.; and Liu, K. 2024. Thought Space Explorer: Navigating and Expanding Thought Space for Large Language Model Reasoning. ArXiv:2410.24155 [cs].
- Zhang, Q.; Lyu, F.; Sun, Z.; Wang, L.; Zhang, W.; Guo, Z.; Wang, Y.; Muennighoff, N.; King, I.; Liu, X.; and Ma, C. 2025. What, How, Where, and How Well? A Survey on Test-Time Scaling in Large Language Models. ArXiv:2503.24235 [cs].
- Zhao, Y.; Yin, H.; Zeng, B.; Wang, H.; Shi, T.; Lyu, C.; Wang, L.; Luo, W.; and Zhang, K. 2024. Marco-o1: Towards Open Reasoning Models for Open-Ended Solutions. ArXiv:2411.14405.
- Zheng, C.; Liu, Z.; Xie, E.; Li, Z.; and Li, Y. 2024. Progressive-Hint Prompting Improves Reasoning in Large Language Models. ArXiv:2304.09797 [cs].
- Zhou, D.; Schärli, N.; Hou, L.; Wei, J.; Scales, N.; Wang, X.; Schuurmans, D.; Cui, C.; Bousquet, O.; Le, Q.; and Chi, E. 2023. Least-to-Most Prompting Enables Complex Reasoning in Large Language Models. ArXiv:2205.10625 [cs].