

Think Smarter, Focus Wisely: Adaptive Cognitive Allocation for LLM Reasoning

Anonymous ACL submission

Abstract

Encouraging longer chains of thought is a common design practice for improving LLM reasoning. However, recent studies show that “thinking more” can backfire. In response, prior studies have typically employed augmented reasoning strategies to enhance performance. While these approaches often improve reasoning robustness and yield higher accuracy, they may also generate excessively long chains, which introduce redundant checks, demand disproportionate reasoning effort, and ultimately lead to inefficient consumption of cognitive resources. This paper introduces Adaptive Reasoning via Cognitive Allocation (ARCA), a structured reasoning framework that adaptively allocates cognitive resources across reasoning phases based on their reasoning state, thereby mitigating the efficiency–accuracy trade-off. The core idea of ARCA is to structure the reasoning procedure into classified phases, while grounding the process and suppressing incoherent drift. Within each phase, ARCA generates candidate directions and employs a Borda-Aggregated selector to identify the most promising ones, while steering inference along phase-aware directions and pruning redundant exploration. Through the dynamic allocation of cognitive resources, the proposed ARCA framework can achieve a balance between accuracy and efficiency. Across six reasoning benchmarks, ARCA consistently outperforms strong baselines, either in terms of enhanced accuracy or reduced reasoning cost.

1 Introduction

“More thinking should mean better answers.” This intuition feels natural to humans and has heavily influenced the design of large language models (LLMs). The prevailing wisdom suggests that encouraging models to generate longer, more detailed chains of thought is a reliable path to stronger reasoning (Wei et al., 2022; OpenAI, 2025, 2024). While LLMs guided by this principle have

shown remarkable capabilities in complex tasks like question answering (Lewkowycz et al., 2022), knowledge retrieval (Schmidgall et al., 2025) and decision support systems (Lubos et al., 2025), a troubling paradox emerges: pushing them to simply “think more” can backfire. Recent studies (Jin et al., 2025) reveal that excessive reasoning, or “overthinking,” often degrades performance. Extended chains of reasoning may accumulate errors (Lewkowycz et al., 2022), reduce stability (Wang et al., 2023), and increase computational demand—leading to substantially higher inference costs and limiting their practical usability.

To counteract this fragility, a dominant strategy has been to enhance robustness and achieve higher accuracy by introducing augmented reasoning procedures that supplement the original inference process. For instance, methods like pairwise comparison (Zhang et al., 2024) and iterative self-evaluation (Chen et al., 2024b) reduce errors by generating and assessing multiple solution paths, but this incurs massive computational overhead. Other techniques, such as those relying on predefined skill libraries (Chen et al., 2024a), can impose structural rigidities that limit their adaptability to novel problems. In essence, these methods achieve robustness at such a high cost in inference demand and inflexibility that they become impractical for many real-world applications.

On the other side, efficiency-oriented methods aim to reduce inference costs, typically through streamlined meta-reasoning architectures (Sui et al., 2025b; Patil and Jadon, 2025). However, this efficiency is achieved by relying on handcrafted contextual frameworks and human-defined heuristics, thereby specializing the systems for specific tasks. This specialization fundamentally limits their problem-solving scope and leads to poor generalizability. Consequently, a generalizable framework that allows models to adaptively allocate their reasoning effort or we called it **cognitive resources**

to achieve both accuracy and efficiency remains a critical, underexplored challenge. The central question therefore becomes: *How can LLMs learn to allocate their cognitive resources adaptively to achieve both accuracy and efficiency simultaneously?*

The key insight is that **effective reasoning requires selective focus rather than indiscriminate depth**. Humans intuitively follow this principle. Consider a game of Sudoku shown in Figure 1: a player instantly fills in a number when it is the only possibility in a row—an act of efficient, linear deduction. However, when confronted with a complex intersection of constraints, the same player may pencil in multiple candidates in a few cells, exploring their implications before committing. This represents an on-demand expansion of the reasoning process to ensure the next move is robust. We posit that endowing LLMs with a similar capability for dynamic cognitive resource allocation is key to resolving the tension between accuracy and efficiency in reasoning. To achieve this, **LLMs likewise need to learn to allocate their cognitive resources adaptively**: engaging in deep reasoning when necessary and pruning effort when a path proves unpromising.

To realize adaptive cognitive resource allocation in LLMs, we propose a structured reasoning framework called **Adaptive Reasoning via Cognitive Allocation (ARCA)**. ARCA guides LLMs in deciding where and how much to think during reasoning by operating through two synergistic stages: **reasoning chain construction** and **cognitive resource allocation**. In the reasoning chain construction stage, a *phase generator* decomposes complex tasks into logically ordered phases, specifying the goal of each step. This structured decomposition prevents fragmented reasoning and ensures systematic task coverage. To illustrate, let us return to the Sudoku example introduced earlier. The solving process can be structured into phases such as [*Fill cells with unique candidates, Explore candidate intersections, Validate placements, and others*], which serve as clear anchors guiding reasoning toward the solution.

Building on this structure, the *cognitive resource allocation* stage dynamically manages reasoning effort at two levels. At the macro level, a phase classifier directs LLMs focus to the current phase’s objectives, ensuring efficient progress. At the micro level, a direction generator explores candidate reasoning steps, while a proposed *Borda-*

Aggregated selection mechanism guided by LLMs’ preference feedback to choose the most promising path. ARCA enables the LLM to adaptively allocate more cognitive resources to critical phases of a task, thereby facilitating precise reasoning while preserving overall processing efficiency. Crucially, the structured phases from the first stage provide the precise context needed for the second stage to make informed allocation decisions, ensuring that cognitive resources are invested exactly where they are most needed. Extensive experiments on six diverse reasoning tasks demonstrate that ARCA achieves strong performance while maintaining favorable resource efficiency compared to baseline approaches.

Main contributions of this work are concluded as: 1) The problem of efficient cognitive resource allocation is formalized, the critical balance between accuracy and efficiency addresses a core challenge in the reasoning of LLMs; 2) An LLM reasoning framework ARCA is proposed to realize cognitive resource allocation in LLM reasoning, where a high-level task decomposition guides fine-grained adaptive exploration to simultaneously improve both reasoning accuracy and efficiency; 3) Comprehensive experiments on diverse reasoning tasks demonstrate that our method achieves strong performance while maintaining favorable resource efficiency compared to baseline approaches.

2 Adaptive Reasoning via Cognitive Allocation

In this section, we present a comprehensive introduction to the ARCA framework, which achieves cognitive resource allocation to balance accuracy and efficiency in LLM reasoning. It integrates two complementary components: reasoning chain construction, which provides structural guidance to mitigate fragmented reasoning, and cognitive resource allocation, which enables selective and adaptive distribution of cognitive effort across phases and reasoning steps. Through the interaction of these two components, ARCA enables LLMs to concentrate their reasoning effort on the phases most critical, thus **improving accuracy** and **enhancing efficiency**.

2.1 Reasoning Chain Construction

We first introduce the *Reasoning Chain Construction*, a module that leverages the semantic analysis capability of LLMs to build a high-level blueprint

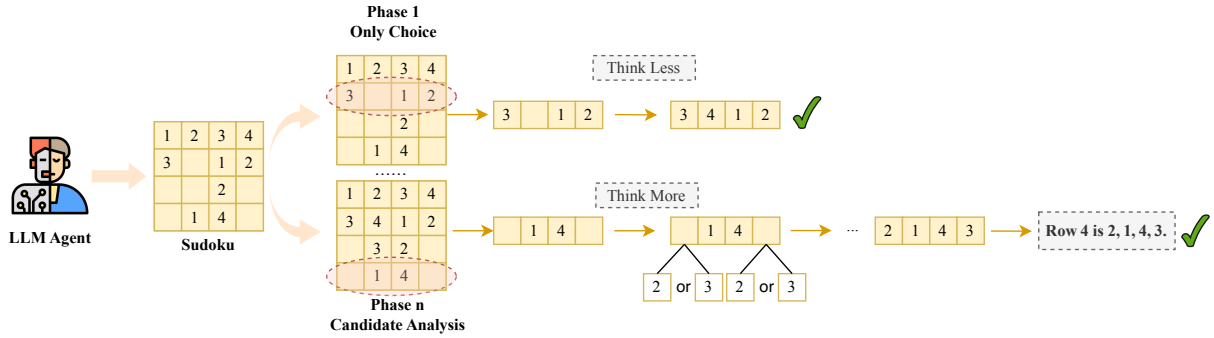


Figure 1: LLMs operate through multiple phases to solve a task, with certain phases requiring minimal cognitive resources while others demand in-depth reasoning.

and to define the reasoning phase (allocation units) that serve as the basis for cognitive resource allocation in ARCA. Traditional chain-of-thought prompting often produces unstructured, divergent, or incomplete reasoning trajectories (Ji et al., 2024). Without explicit semantic scaffolding, the model may digress into irrelevant steps or overlook crucial components of the task. To achieve ARCA’s goal of balancing accuracy and efficiency, it is essential to first identify the reasoning phases where resources can be meaningfully allocated.

For illustration, consider solving a Sudoku: the task can be decomposed into several reasoning phases, each serving as an allocation unit. A natural sequence begins with (i) *Single-candidate placement*, filling cells with only one valid option; followed by (ii) *Candidate elimination and inference*, iteratively deducing placements based on row, column, and block constraints; (iii) *Conflict detection and backtracking*, revising decisions if contradictions arise; and (iv) *Validation and finalization*, ensuring all cells satisfy Sudoku rules.

To realize this decomposition, ARCA introduces the *Reasoning Phase Generator (ReasonGen)*, which maps a task \mathcal{T} into a sequence of logically ordered reasoning phases:

$$\mathcal{P} = \{\phi_1, \phi_2, \dots, \phi_n\} = \text{ReasonGen}(\mathcal{T}), \quad (1)$$

where \mathcal{T} is the input task, ϕ_i is the i -th reasoning phase and \mathcal{P} is the generated phase set. ReasonGen begins by analyzing the task context to identify and delineate each fundamental phase. Each phase specifies a semantic objective—*what needs to be achieved* at that step rather than prescribing the operational details of *how* to achieve it. For the Sudoku example, this naturally yields phases starting with single-candidate placement, followed by iterative candidate elimination, conflict detection

and backtracking, and concluding with final validation. We provide illustrative examples of the generated phases in Appendix C.4. This structured decomposition defines the allocation units (reasoning phases) required by ARCA and constrains reasoning within explicit semantic boundaries. By providing a high-level blueprint for cognitive resource allocation, ReasonGen reduces irrelevant exploration and yields more coherent and efficient inference compared to unstructured CoT.

2.2 Cognitive Resource Allocation

Building on the reasoning phases and high-level blueprint defined by Reasoning Chain Construction, this module performs the allocation of cognitive resources used by ARCA. Allocating resources uniformly across all reasoning steps is inefficient and costly; ARCA instead dynamically adjusts cognitive resource allocation according to the current semantic phase and task context. The allocation process uses two coordinated mechanisms: **Macro-Level Phase Identification** and **Micro-Level Direction Selection**. The former determines *which phase* the reasoning process is currently in, while the latter identifies the most promising direction within that phase.

At the **Macro-Level**, the module prioritizes the current phase, biasing resource allocation toward operations relevant to it. To identify the current phase, we introduce a phase classifier *PhaseClass*, a runtime supervisory unit that dynamically evaluates the solver’s state to assign the reasoning step to the appropriate phase:

$$\phi_t = \text{PhaseClass}(x_t, \mathcal{P}), \quad (2)$$

where x_t is the reasoning state at step t . The architecture enhances complex reasoning by dynamically identifying the current phase, enabling the

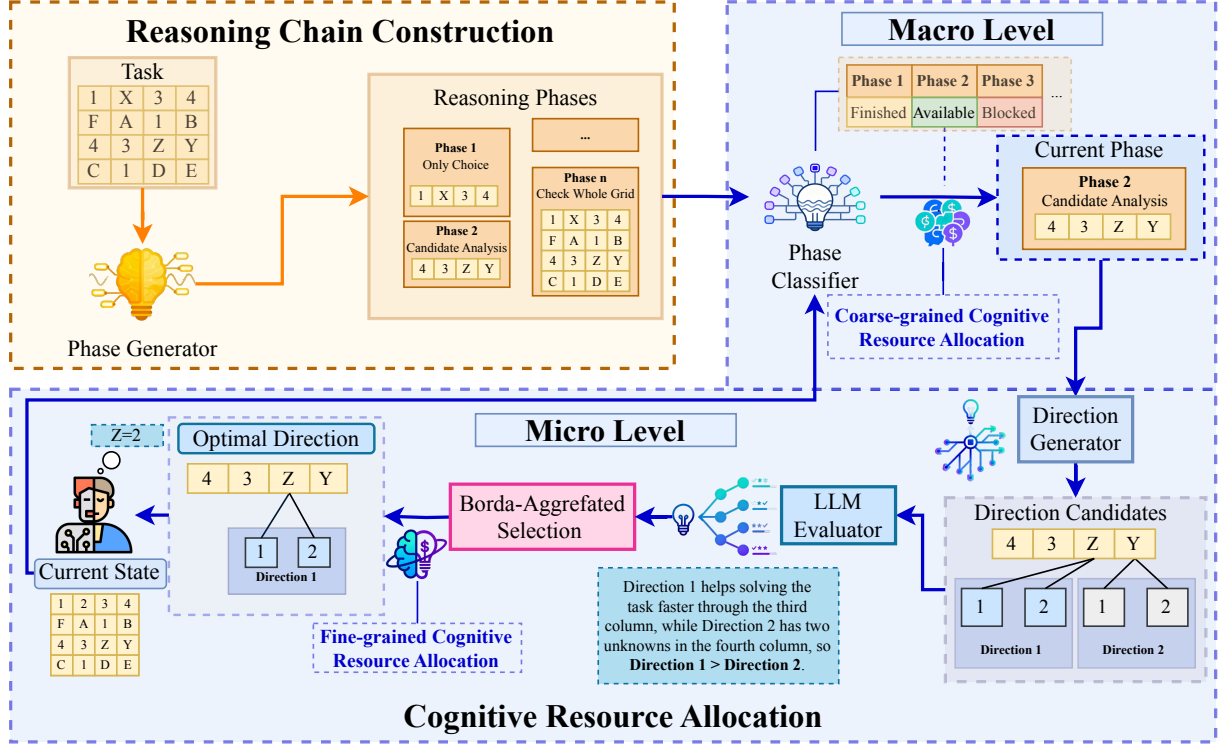


Figure 2: The overall framework of ARCA. It integrates reasoning chain construction and cognitive resource allocation to balance accuracy and efficiency in LLM reasoning, where reasoning chain construction structures the task into phases, and cognitive resource allocation allocates cognitive resources through macro-level phase guidance and micro-level direction selection.

LLM to focus on phase-specific objectives. Once a phase is completed, the module transitions seamlessly to the next, reallocating computational resources according to new requirements. By concentrating effort on the active phase and minimizing expenditure on completed or irrelevant directions, the system maintains targeted and efficient progress across the reasoning chain, thereby preventing resource over-allocation. **A key insight** underlying this design is that reasoning phases naturally impose heterogeneous resource demands: simpler phases require minimal resources and benefit from rapid closure, whereas more demanding phases call for deeper inference and thus greater resource investment. The architecture capitalizes on this heterogeneity, allowing adaptive allocation of computational effort so that resources are concentrated on ongoing objectives while avoiding over-allocation to irrelevant directions. In doing so, the system achieves dynamic and efficient resource utilization throughout reasoning.

At the **Micro-Level**, the module generates and selects reasoning directions that are locally optimal or highly relevant to the active phase. Cognitive resources or reasoning efforts are dynamically as-

signed to the most promising next steps within the phase. To this end, we introduce a reasoning direction generator *DirectionGen*, a real-time strategic module that steers the LLM’s inference process:

$$\mathcal{D}_t = \{d_1, d_2, \dots, d_m\} = \text{DirectionGen}(\phi_t, x_t, \mathcal{T}), \quad (3)$$

where \mathcal{D}_t is the set of candidate directions at step t and d_i is the i -th direction. At each reasoning step, the generator produces a focused set of actionable directions based on the current phase and contextual state. These directions provide timely, targeted guidance aligned with the phase’s objectives. To enable broad exploration, multiple candidate paths are proposed. **Selecting the most promising path** among them poses a central challenge, as explicit reward signals are unavailable and handcrafted reward functions are difficult to design and prone to bias. To address this, we employ an LLM as an implicit preference oracle to evaluate candidate directions. Building on these evaluations, we introduce a **Borda-aggregated direction selection algorithm** (Yan et al., 2022), which consolidates pairwise comparisons into a robust consensus score and will be described in detail in section 2.3. This ap-

proach mitigates noise from individual judgments, reduces reliance on brittle heuristics, and reliably identifies the most promising reasoning direction. The optimal direction d_t^* is then selected according to the aggregated Borda score:

$$d_t^* = \arg \max_{d \in \mathcal{D}_t} \text{Borda}(d). \quad (4)$$

The LLM then generates the next reasoning state x_{t+1} according to:

$$x_{t+1} = \text{LLM}(x_t, \phi_t, d_t^*). \quad (5)$$

By providing such fine-grained tactical guidance, this module effectively bridges phase planning with real-time reasoning, enabling efficient and rational allocation of cognitive resources throughout the inference process. In the next subsection, we will make a detailed introduction to the Borda-aggregated direction selection algorithm.

2.3 Borda-Aggregated Direction Selection

In this section, we provide a detailed explanation of how the Borda-aggregated direction selection algorithm identifies the most preferred direction in Equation 4. A prevalent strategy for identifying the optimal reasoning path through LLM preference comparisons involves leveraging pairwise methods, exemplified by the dueling bandit framework (Zhang et al., 2024). In this setting, when comparing two candidate thoughts i and j , candidate i is chosen with probability $q(i, j)$, while candidate j is selected with the complementary probability $q(j, i) = 1 - q(i, j)$. Here, $q(i, j) \geq \frac{1}{2}$ whenever i is ranked higher than j . Repeated comparisons are assumed to be independent.

However, dueling bandit algorithms such as DTS (Wu and Liu, 2016) typically rely on the Copeland score (Zoghi et al., 2015) to aggregate comparison outcomes. A major limitation of the Copeland score in LLM-based preference assessment is its sensitivity to minor preference variations (Goel et al., 2017). This sensitivity arises from its win-counting mechanism, which can amplify stochastic fluctuations inherent in LLM judgments (Li et al., 2025b). Consequently, achieving stable rankings often requires extensive comparisons, which is especially challenging in noisy evaluation environments (Qin et al., 2023). To address this limitation, the Borda score (Rothe, 2019) is adopted as an alternative. The Borda score for a

candidate direction i is defined as:

$$\text{Borda}(i) = \frac{1}{|\mathcal{C}| - 1} \sum_{j \in \mathcal{C}, j \neq i} q(i, j), \quad (6)$$

where \mathcal{C} denotes the set of candidates. The Borda score’s win-rate formulation effectively aggregates pairwise preferences and offers clear practical advantages in LLM evaluation settings. Its scoring mechanism, which estimates the average probability of victory, is well-suited to the stochastic and noisy nature of LLM judgments. By averaging outcomes across multiple comparisons, it confers robustness against minor inconsistencies in individual assessments (Rothe, 2019). Furthermore, the computational simplicity of maintaining and updating win rates enables highly efficient implementation in large-scale scenarios, allowing broad candidate coverage without exhaustive evaluations.

We formulate direction selection as a Borda score-based framework (Yan et al., 2022; Clarke et al., 2021), with an LLM serving as the **preference function**. Our algorithm begins with a pruning phase to efficiently eliminate clearly suboptimal directions while retaining the most promising candidates. During each pruning iteration, approximate Borda scores are computed by comparing each candidate against a fixed-size random subset of opponents. This sparse comparison strategy ensures broad coverage without exhaustive evaluations. Candidates with scores below an elimination score are pruned. We set the elimination score at 0.5, which corresponds to random chance performance, while any candidate scoring below this level is deemed inferior and removed. This pruning process is repeated iteratively until the number of remaining candidates falls below a predefined threshold.

The algorithm then proceeds to a final evaluation stage, conducting full round-robin comparisons among the remaining candidates. This enables accurate, high-confidence estimation of the true Borda scores, from which the top-scoring candidate is chosen as the final solution. By combining efficient broad pruning with precise final assessment, this two-stage approach effectively balances computational efficiency with selection reliability. Details are provided in Appendix D.1 and D.2.

Comparison with prior works ARCA differs from prior structured-reasoning methods in both paradigm and optimization objective. Tree-of-Thoughts (Yao et al., 2023) and Graph-of-

Thoughts (Besta et al., 2024) improve robustness through breadth-first expansion of thought trees or graphs, but incur substantial computational overhead due to uniform exploration. In contrast, ARCA adopts a phase-oriented, goal-driven decision-making strategy with hierarchical reasoning and dynamic cognitive resource allocation, enabling adaptive control over reasoning depth and granularity.

In contrast to pairwise comparison methods such as CToT (Zhang et al., 2024), which rely on repeated duels and re-comparisons as reasoning progresses, ARCA employs a Borda score based on average win rates for more robust global evaluation, together with dynamic pruning to rapidly eliminate inferior candidates and reduce redundant comparisons. Unlike Skills-in-Context (Chen et al., 2024a) approaches that depend on predefined skill libraries and suffer from limited generalization, ARCA generates task-specific reasoning phases that provide macro-level guidance and adapt naturally across tasks. Compared to Meta-Reasoner (Sui et al., 2025b), which requires predefined operations and bandit learning, ARCA operates entirely through real-time reasoning without pretraining and remains effective across a broad parameter range. These design choices yield a significantly improved accuracy–cost trade-off, as ARCA consistently achieving comparable or higher accuracy at lower cognitive resource cost.

3 Experiments

In this section, we conduct comprehensive experiments to evaluate performance and provide an in-depth analysis of ARCA, comparing it against baseline approaches in terms of cost and accuracy. Additional results and ablation study are in Appendix C.2 and C.3.

3.1 Reasoning Tasks

We evaluate the performance of our proposed method, ARCA, on a suite of six challenging real-world reasoning tasks. These tasks span a diverse range of domains, including question answering (AQUA), multi-step arithmetic (BBEH), math word problems (GSM8K), the Game of 24, Sudoku puzzles, and the AIME competition-level problems.

- **AQUA** (Wei et al., 2022), the question answering task, which consists of 254 arithmetic reasoning questions designed to evaluate logical reasoning abilities through diverse mathematical problems.

Each question is associated with five multiple-choice options labeled A through E.

- **BBEH** (Kazemi et al., 2025) is a recently introduced benchmark aimed at advancing the evaluation of reasoning in large language models. In our experiments, we select the multi-step arithmetic task from BBEH.

- **GSM8K** (Cobbe et al., 2021) is a widely-used benchmark of grade-school math word problems that require multi-step reasoning to solve. Each problem involves basic arithmetic operations and logical thinking to arrive at the final answer.

- **The Game of 24** (Yao et al., 2023) is a mathematical challenge in which the objective is to combine four given numbers using basic arithmetic operations to yield a total of 24. For our experiments, we adopt the same dataset and setup as, which includes 1,362 problems sourced from 4nums.com.

- **The Sudoku** (Long, 2023) includes 10 puzzles each for 3×3, 4×4, and 5×5 grid sizes. Each puzzle is partially filled, and the task is to complete the grid without altering the provided numbers. A solution is considered correct if the completed grid adheres to all standard Sudoku rules.

- **AIME** (Mathematical Association of America, 2024) is a highly challenging mathematics contest administered to top-performing participants of the AMC. It serves as a key benchmark for evaluating the mathematical reasoning and problem-solving capabilities of LLMs.

All experiments were conducted with the Deepseek-V3 (DeepSeek-AI, 2025). We also perform additional experiments on Qwen2.5-7B and GPT-4o-mini, as shown in Table. 3. Further detailed configurations and results are provided in Appendix C.1 and C.2.

3.2 Baselines and Results

We compare our method with the following baselines: CoT (Wei et al., 2022), Self-Refine (Madaan et al., 2023), SToT (Yao et al., 2023), PoT (Chen et al., 2023), CToT (Zhang et al., 2024) and BoT (Chen et al., 2024b). On each dataset, 3 test runs are conducted and the average accuracy as well as the cost per question are presented. The experimental results are presented in Table. 4 and Table. 5. These results demonstrate that ARCA outperforms other baselines and achieves significant advantages, particularly in complex tasks such as the Game of 24, Sudoku puzzles, and BBEH.

In these tasks, the solver requires long-chain reasoning and operates in a high-dimensional solu-

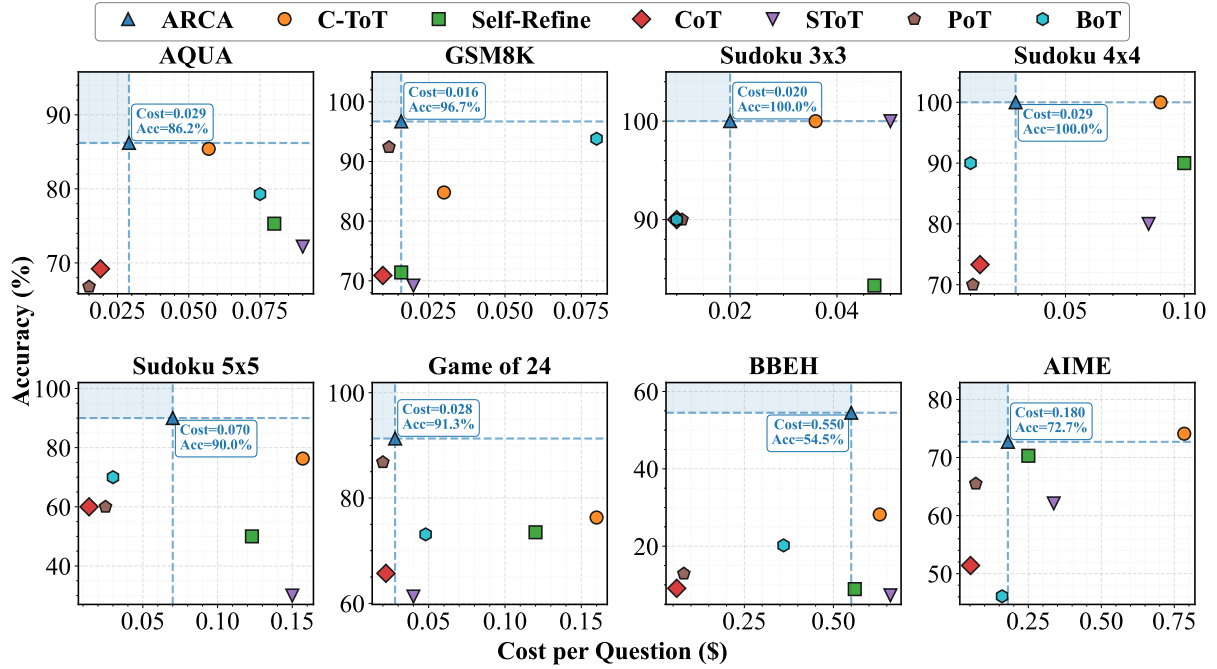


Figure 3: Accuracy–cost comparison across reasoning tasks. Each marker denotes a method. Dashed lines and the text box highlight ARCA, which none surpass in the shaded region, showing ARCA performs best among compared methods.

	$m = 4$		$m = 6$		$m = 8$		$m = 12$	
	Acc.	Cost	Acc.	Cost	Acc.	Cost	Acc.	Cost
$n = 2$	89.2	0.040	95.2	0.045	94.5	0.053	95.8	0.096
$n = 4$	92.9	0.058	96.1	0.064	96.4	0.075	95.8	0.096
$n = 6$	94.1	0.075	96.3	0.080	96.3	0.082	95.8	0.096

Table 1: Ablation study on direction selection algorithm parameters on GSM8K. Here n denotes the number of comparisons and m denotes the size. **Acc.** is accuracy (%), and **Cost** is average computational cost (lower is better).

d_n	AQUA (%)	Game of 24	Average (%)
$d_n = 2$	79.5	88.5	84.0
$d_n = 4$	85.4	91.3	88.4
$d_n = 6$	86.2	91.3	88.7
$d_n = 10$	85.7	93.1	89.4

Table 2: Ablation study on the number of directions d_n . Accuracy (%) is reported on AQUA and Game of 24, along with the average. Larger d_n generally improves performance, with diminishing gains after $d_n = 6$.

tion space. The heuristic strategies for reasoning in ARCA provide critical guidance at each step, assisting the solver in accurately steering toward the final answer, reducing deviations, and ultimately leading to more effective and reliable problem-solving. This mechanism proves essential for navigating the complexity inherent in such challenging domains.

To evaluate the efficiency of cognitive resource allocation, we employ reasoning cost as a key

metric, where using fewer reasoning costs indicates more efficient. A comparison of the accuracy and cost between our method and baseline approaches across different tasks is presented in Fig. 3. Here we adopt the cost calculation method from CToT (Zhang et al., 2024). The experimental findings clearly illustrate that ARCA successfully achieves an effective and practical balance between model accuracy and operational cost-efficiency. This balance is realized through the novel integration of structured reasoning chain construction and dynamic cognitive resource allocation mechanisms. Our method demonstrates the capability to autonomously identify and prioritize critical reasoning phases, thereby allocating computational resources in an adaptive and context-aware manner. This sophisticated mechanism not only consistently enhances the quality and reliability of solutions but also maintains coherent focus throughout the reasoning trajectory toward the correct solu-

511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530

Models	Datasets				
	AQUA (%)	BBEH (%)	GSM8K (%)	Game of 24 (%)	Sudoku Puzzle 5×5 (%)
Deepseek-V3	86.2	54.5	96.7	91.3	90.0
Qwen2.5-7B	84.7	40.2	89.4	90.2	80.0
GPT-4o-mini	89.8	45.6	94.7	88.3	90.0

Table 3: Additional results with Qwen2.5-7B and GPT-4o-mini

tion. Consequently, ARCA delivers substantially improved overall reasoning performance while simultaneously ensuring judicious control over computational expenditure, providing a reliable framework for complex cognitive tasks.

3.3 Ablation Study

Analysis on the parameters in Borda-aggregated direction selection. To analyze how the parameters in the direction selection algorithm affects accuracy and computational cost, we conduct experiments on the GSM8K dataset. To ensure comparability, the maximum number of direction is fixed at 12 across all trials. Experiments are conducted with pruning pool size threshold $m = \{4, 6, 8, 12\}$ (where $m=12$ corresponds to the non-pruning case) and the number of comparisons per direction during pruning set to $n = \{2, 4, 6\}$. Results in Table. 1 indicate that when using pool size pruning, configurations with pool size threshold 6 and 8 achieve performance comparable to the full pool (size 12), while significantly reducing computational expense. Although the value of n has an impact on performance, experiments show that a medium n is sufficient to achieve rapid pruning in the early stages. This demonstrates that the proposed selection framework effectively balances computational efficiency with selection reliability, thus offers a scalable and practical solution for resource-aware automated reasoning.

Analysis on the Maximum Direction Number. The number of generated directions determines the breadth of exploration available to the solver at each reasoning step. We evaluate the impact of this parameter by conducting experiments on the AQUA and Game of 24 datasets, using maximum direction numbers set to 2, 4, 6, and 10. The results, summarized in Table. 2, indicate that when the maximum direction number is limited to 2 which resulting in a narrow exploration scope, the performance is noticeably worse compared to configurations allowing broader exploration. In contrast, when the maximum number of directions

is set to 4 or higher, performance stabilizes and remains consistently high. This suggests that the reasoning chain construction mechanism helps the solver maintain a clear objective and reduces the need for extensive exploration, thereby achieving more efficient and reliable problem-solving even with moderate search width.

4 Conclusion

In conclusion, we propose a novel framework ARCA designed to tackle efficient cognitive resource allocation in LLM reasoning, with a specific focus on balancing accuracy and efficiency. By integrating decomposition, strategy generation, monitoring, and dynamic selection into a cohesive system, our approach enhances structural coherence, optimizes reasoning effort, and improves accuracy in complex scenarios while keeping additional reasoning cost negligible to preserve efficiency. Extensive experiments across diverse tasks show that our method delivers strong performance while maintaining competitive resource efficiency compared to existing baselines. In future work, we will pursue more reliable reasoning chains and refine our framework for accurate direction generation, focusing on more complex task environments.

5 Limitations

In our approach, all modules are built upon large language models and prompt-based instructions. This design inevitably leads to output instability as well as increased computational overhead due to multiple model invocations during inference. Although we introduce a direction selector and a phase discriminator to mitigate the disturbances caused by such instability, potential reasoning errors and the additional cost induced by repeated calls remain unavoidable in practice and under diverse settings. Replacing the selection framework and phase generation mechanism with pretrained reward models and dedicated generators may further improve the efficiency and stability of our method overall in future work.

References

- 615 Maciej Besta, Nils Blach, Ales Kubicek, Robert Ger-
616 stenberger, Michal Podstawski, Lukas Gianinazzi,
617 Joanna Gajda, Tomasz Lehmann, Hubert Niewiadom-
618 ski, Piotr Nyczyk, and 1 others. 2024. Graph of
619 thoughts: Solving elaborate problems with large lan-
620 guage models. In *The 38th AAAI conference on Arti-
621 ficial Intelligence (AAAI)*.
- 622 Jiaao Chen, Xiaoman Pan, Dian Yu, Kaiqiang Song,
623 Xiaoyang Wang, Dong Yu, and Jianshu Chen. 2024a.
624 Skills-in-context: Unlocking compositionality in
625 large language models. In *Findings of the Associ-
626 ation for Computational Linguistics (EMNLP)*, pages
627 13838–13890.
- 628 Sijia Chen, Baochun Li, and Di Niu. 2024b. Boosting
629 of thoughts: Trial-and-error problem solving with
630 large language models. In *The 12th International
631 Conference on Learning Representations (ICLR)*.
- 632 Wenhu Chen, Xueguang Ma, Xinyi Wang, and
633 William W. Cohen. 2023. Program of thoughts
634 prompting: Disentangling computation from reason-
635 ing for numerical reasoning tasks. *arXiv preprint
636 arXiv:2211.12588*.
- 637 Charles LA Clarke, Alexandra Vtyurina, and Mark D
638 Smucker. 2021. Assessing top-preferences. *ACM
639 Transactions on Information Systems (TOIS)*, 39(3):1–
640 21.
- 641 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,
642 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias
643 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro
644 Nakano, Christopher Hesse, and John Schulman.
645 2021. Training verifiers to solve math word prob-
646 lems. *arXiv preprint arXiv:2110.14168*.
- 647 Yingqian Cui, Pengfei He, Jingying Zeng, Hui Liu,
648 Xianfeng Tang, Zhenwei Dai, Yan Han, Chen Luo,
649 Jing Huang, Zhen Li, Suhang Wang, Yue Xing, Jil-
650 iang Tang, and Qi He. 2025. Stepwise perplexity-
651 guided refinement for efficient chain-of-thought rea-
652 soning in large language models. *arXiv preprint
653 arXiv:2502.13260*.
- 654 DeepSeek-AI. 2025. DeepSeek-V3 technical report.
655 *arXiv preprint arXiv:2412.19437*.
- 656 Ashish Goel, Anilesh K Krishnaswamy, and Kamesh
657 Munagala. 2017. Metric distortion of social choice
658 rules: Lower bounds and fairness properties. In *The
659 2017 ACM Conference on Economics and Computa-
660 tion (EC)*.
- 661 Deyi Ji, Lanyun Zhu, Siqi Gao, Peng Xu, Hongtao
662 Lu, Jieping Ye, and Feng Zhao. 2024. Tree-of-
663 table: Unleashing the power of LLMs for enhanced
664 large-scale table understanding. *arXiv preprint
665 arXiv:2411.08516*.
- 666 Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanwei Li,
667 Yu Qi, Xinyan Chen, Liuhui Wang, Jianhan Jin,
668 Claire Guo, Shen Yan, Bo Zhang, Chaoyou Fu, Peng
Gao, and Hongsheng Li. 2025. Mme-cot: Bench-
marking chain-of-thought in large multimodal mod-
els for reasoning quality, robustness, and efficiency.
arXiv preprint arXiv:2502.09621.
- Haibo Jin, Peiyan Zhang, Man Luo, and Haohan
Wang. 2025. Reasoning can hurt the inductive
abilities of large language models. *arXiv preprint
arXiv:2505.24225*.
- Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John
Palowitch, Chrysovalantis Anastasiou, Sanket Vaib-
hav Mehta, Lalit K. Jain, Virginia Aglietti, Disha
Jindal, Peter Chen, Nishanth Dikkala, Gladys Tyen,
Xin Liu, Uri Shalit, Silvia Chiappa, Kate Olszewska,
Yi Tay, Vinh Q. Tran, Quoc V. Le, and Orhan Fir-
rat. 2025. BIG-bench extra hard. *arXiv preprint
arXiv:2502.19187*.
- Aitor Lewkowycz, Anders Andreassen, David Dohan,
Ethan Dyer, Henryk Michalewski, Vinay Ramasesh,
Ambrose Slone, Cem Anil, Imanol Schlag, Theo
Gutman-Solo, and 1 others. 2022. Solving quan-
titative reasoning problems with language models. In
*Advances in Neural Information Processing Systems
(NeurIPS)*.
- Leqian Li, Dianxi Shi, Jialu Zhou, Xinyu Wei,
Mingyue Yang, Songchang Jin, and Shaowu Yang.
2025a. Retrieval feedback memory enhancement
large model retrieval generation method. *arXiv
preprint arXiv:2508.17862*.
- Tianyi Li, Yu Qin, and Olivia R. Liu Sheng. 2025b. A
multi-task evaluation of LLMs’ processing of aca-
demic text input. *arXiv preprint arXiv:2508.11779*.
- Zhiyuan Li, Yi Chang, and Yuan Wu. 2025c. Think-
bench: Evaluating thinking efficiency and chain-of-
thought quality of large reasoning models. *arXiv
preprint arXiv:2505.22113*.
- Jieyi Long. 2023. Large language model guided tree-of-
thought. *arXiv preprint arXiv:2305.08291*.
- Sebastian Lubos, Alexander Felfernig, Damian Garber,
Viet-Man Le, Manuel Henrich, Reinhard Willfort,
and Jeremias Fuchs. 2025. Towards group decision
support with LLM-based meeting analysis. In *The
33rd ACM Conference on User Modeling, Adaptation
and Personalization*, pages 331–335.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler
Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon,
Nouha Dziri, Shrimai Prabhumoye, Yiming Yang,
and 1 others. 2023. Self-refine: Iterative refinement
with self-feedback. In *Advances in Neural Informa-
tion Processing Systems (NeurIPS)*.
- Minjia Mao, Bowen Yin, Yu Zhu, and Xiao Fang. 2025.
Early stopping chain-of-thoughts in large language
models. *arXiv preprint arXiv:2509.14004*.
- Mathematical Association of America. 2024. [American
invitational mathematics examination \(aime\) 2024](#).
Part of the American Mathematics Competitions.

724	OpenAI. 2024. Learning to reason with LLMs . Technical report, OpenAI.	780
725		781
726	OpenAI. 2025. OpenAI GPT-4.5 system card . Technical report, OpenAI.	782
727		783
728	Avinash Patil and Aryan Jadon. 2025. Advancing reasoning in large language models: Promising methods and approaches. <i>arXiv preprint arXiv:2502.03671</i> .	784
729		785
730		786
731	Nearchos Potamitis and Akhil Arora. 2025. Are retrials all you need? enhancing large language model reasoning without verbalized feedback. <i>arXiv preprint arXiv:2504.12951</i> .	787
732		788
733		789
734		790
735	Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, and 1 others. 2023. Large language models are effective text rankers with pairwise ranking prompting. <i>arXiv preprint arXiv:2306.17563</i> .	791
736		792
737		793
738		794
739		795
740		796
741	Matthew Renze and Erhan Guven. 2024. The benefits of a concise chain of thought on problem-solving in large language models. In <i>The 2nd International Conference on Foundation and Large Language Models (FLLM)</i> .	797
742		798
743		799
744		800
745		801
746	Jörg Rothe. 2019. Borda count in collective decision making: A summary of recent results. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> .	802
747		803
748		804
749		805
750	Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Michael Moor, Zicheng Liu, and Emad Barsoum. 2025. Agent laboratory: Using LLM agents as research assistants. <i>arXiv preprint arXiv:2501.04227</i> .	806
751		807
752		808
753		809
754		810
755	Maohao Shen, Guangtao Zeng, Zhenting Qi, Zhang-Wei Hong, Zhenfang Chen, Wei Lu, Gregory Wornell, Subhro Das, David Cox, and Chuang Gan. 2025. Satori: Reinforcement learning with chain-of-action-thought enhances LLM reasoning via autoregressive search. <i>arXiv preprint arXiv:2502.02508</i> .	811
756		812
757		813
758		814
759		815
760		816
761	Jingbo Su. 2024. Enhancing adversarial attacks through chain of thought. <i>arXiv preprint arXiv:2410.21791</i> .	817
762		818
763	Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Na Zou, Hanjie Chen, and Xia Hu. 2025a. Stop overthinking: A survey on efficient reasoning for large language models. <i>arXiv preprint arXiv:2503.16419</i> .	819
764		820
765		821
766		822
767		823
768		824
769	Yuan Sui, Yufei He, Tri Cao, Simeng Han, Yulin Chen, and Bryan Hooi. 2025b. Meta-reasoner: Dynamic guidance for optimized inference-time reasoning in large language models. <i>arXiv preprint arXiv:2502.19918</i> .	825
770		826
771		827
772		828
773		829
774	Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. 2022. Challenging BIG-bench tasks and whether chain-of-thought can solve them. <i>arXiv preprint arXiv:2210.09261</i> .	830
775		831
776		832
777		833
778		834
779		
	Qiaoyu Tang, Hao Xiang, Le Yu, Bowen Yu, Hongyu Lin, Yaojie Lu, Xianpei Han, Le Sun, and Junyang Lin. 2025. Refcritic: Training long chain-of-thought critic models with refinement feedback. <i>arXiv preprint arXiv:2507.15024</i> .	
	Wenxiao Wang, Parsa Hosseini, and Soheil Feizi. 2025. Chain-of-defensive-thought: Structured reasoning elicits robustness in large language models against reference corruption. <i>arXiv preprint arXiv:2504.20769</i> .	
	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. In <i>The 11th International Conference on Learning Representations (ICLR)</i> .	
	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. In <i>Advances in Neural Information Processing Systems (NeurIPS)</i> .	
	Huasen Wu and Xin Liu. 2016. Double thompson sampling for dueling bandits. In <i>Advances in Neural Information Processing Systems (NeurIPS)</i> , volume 29.	
	Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. 2025. Softcot: Soft chain-of-thought for efficient reasoning with llms . <i>Preprint</i> , arXiv:2502.12134.	
	Xinyi Yan, Chengxi Luo, Charles LA Clarke, Nick Craswell, Ellen M Voorhees, and Pablo Castells. 2022. Human preferences as dueling bandits. In <i>The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)</i> .	
	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. <i>arXiv preprint arXiv:2505.09388</i> .	
	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. In <i>Advances in Neural Information Processing Systems (NeurIPS)</i> .	
	Tianshu Yu, Chao Xiang, Mingchuan Yang, Pei Ke, Bosi Wen, Cunxiang Wang, Jiale Cheng, Li Zhang, Xinyu Mu, Chuxiong Sun, and Minlie Huang. 2025a. Training language model to critique for better refinement. <i>arXiv preprint arXiv:2506.22157</i> .	
	Tong Yu, Yongcheng Jing, Xikun Zhang, Wentao Jiang, Wenjie Wu, Yingjie Wang, Wenbin Hu, Bo Du, and Dacheng Tao. 2025b. Benchmarking reasoning robustness in large language models. <i>arXiv preprint arXiv:2503.04550</i> .	

835 Zhen-Yu Zhang, Siwei Han, Huaxiu Yao, Gang Niu,
836 and Masashi Sugiyama. 2024. Generating chain-
837 of-thoughts with a pairwise-comparison approach
838 to searching for the most promising intermediate
839 thought. *arXiv preprint arXiv:2402.06918*.

840 Chuanyang Zheng, Zhengying Liu, Enze Xie, Zhenguo
841 Li, and Yu Li. 2023. Progressive-hint prompting
842 improves reasoning in large language models. *arXiv*
843 *preprint arXiv:2304.09797*.

844 Zhanke Zhou, Rong Tao, Jianing Zhu, Yiwen Luo,
845 Zengmao Wang, and Bo Han. 2024. Can language
846 models perform robust reasoning in chain-of-thought
847 prompting with noisy rationales? *arXiv preprint*
848 *arXiv:2410.23856*.

849 Masrour Zoghi, Zohar S Karnin, Shimon Whiteson, and
850 Maarten De Rijke. 2015. Copeland dueling bandits.
851 In *Advances in Neural Information Processing Sys-*
852 *tems (NeurIPS)*.

A Ethical Considerations

Our study investigates the reasoning capabilities of large language models (LLMs). Accordingly, all experiments were conducted on LLMs to evaluate and validate our proposed approach. Beyond experimentation, we employed an LLM as an auxiliary tool during manuscript preparation. Specifically, it was used to refine language for grammar and clarity, and to generate illustrative (non-experimental) figures based on prompts we provided. All research ideas, methods, experiments, analyses, and conclusions were developed by the authors.

B Related Work

Reasoning Methods for Robustness A dominant strategy to enhance reasoning robustness is to generate and evaluate a multitude of augmented reasoning procedures. This approach moves beyond a single chain of thought (Wei et al., 2022) by explicitly constructing multiple reasoning trajectories. For instance, Tree of Thoughts (Yao et al., 2023) frames reasoning as a tree search, allowing exploration of parallel thought candidates at each step. Graph of Thoughts (Besta et al., 2024) further generalizes this into a graph structure to capture more complex interdependencies between thoughts. Methods like Boosting Task-Oriented Reasoning (Chen et al., 2024b) iteratively generate numerous reasoning steps and use LLM-based error analysis to refine them, while Comparative Tree of Thought (Zhang et al., 2024) employs pairwise comparison to select optimal paths. CoDT (Wang et al., 2025) improves the robustness against reference corruption by providing a few exemplars with structured and defensive reasoning as demonstrations. MME-CoT (Jiang et al., 2025) incorporates three novel metrics to assess the reasoning quality, robustness, and efficiency. Math-RoB (Yu et al., 2025b) uses an instruction-based approach to generate diverse datasets resembling training distributions. CD-CoT (Zhou et al., 2024) enhances LLMs’ denoising-reasoning capabilities by contrasting noisy rationales with one clean rationale. CoT-GCG (Su, 2024) enhances adversarial attacks on aligned LLMs by integrating CoT prompts with the greedy coordinate gradient technique.

Efficiency-Oriented Reasoning Methods In direct contrast, another line of research focuses on streamlining the reasoning process to reduce computational costs. These methods often employ

meta-reasoning architectures or probabilistic approximations to achieve faster inference. For example, Meta-Reasoner (Sui et al., 2025b) uses contextual multi-armed bandits to dynamically adjust reasoning strategies based on state evaluation. ES-CoT (Mao et al., 2025) shortens thought generation by prompting the LLM to output a step answer at each reasoning step. THINK-Bench (Li et al., 2025c) introduces a benchmark with novel efficiency metrics, to evaluate the reasoning efficiency. (Cui et al., 2025) proposes a method that identifies and focuses on generating important reasoning steps in reasoning by using perplexity to measure their importance, (Sui et al., 2025a) explores efficient data use, small language model reasoning, and evaluation methods. Soft Chain-of-Thought (Xu et al., 2025) leverages probabilistic soft chains and prompt tweaks for efficient, uncertainty-aware reasoning. Similarly, COAT (Shen et al., 2025) uses action-oriented chains for meta-reasoning.

Feedback-Based Refinement and Evaluation

Another influential line of research focuses on iterative self-improvement through feedback mechanisms. In this paradigm, the LLM itself is leveraged to evaluate and refine its reasoning trajectories in a cyclic manner. For instance, Self-Refine (Madaan et al., 2023) introduces an algorithm where the LLM generates output, provides self-feedback, and then refines its output based on that feedback. RCO (Yu et al., 2025a) trains critic models using a feedback loop where critiques guide the actor model in refining responses. RefCritic (Tang et al., 2025) trains a critic module with dual rule-based rewards focusing on instance-level correctness of solution judgments and refinement accuracies of the policy model. (Renze and Guven, 2024) investigates the effects of self-reflection in large language models on problem-solving performance. (Potamitis and Arora, 2025) enhances reasoning by allowing the models to retry problem-solving attempts upon identifying incorrect answers. RFM-RAG (Li et al., 2025a) transforms stateless retrieval into stateful continuous knowledge management by constructing a dynamic evidence pool to generate refined queries using relational triples and evidence. PHP (Zheng et al., 2023) incorporates the solution from a previous attempt as a hint for the next, creating an iterative improvement loop.

While significant progress has been made within these methods, the fundamental tension between

robustness and efficiency remains largely unaddressed. ARCA addresses this core challenge by empowering LLMs to dynamically allocate greater cognitive resources to critical phases of a reasoning task, achieving an optimal balance between accuracy and efficiency.

C Experiments and Settings

C.1 The Details of Experiment

We evaluate the performance of our proposed method, ARCA, on a suite of six challenging real-world reasoning tasks. These tasks span a diverse range of domains, including question answering (AQUA), multi-step arithmetic (BBEH), math word problems (GSM8K), the Game of 24, Sudoku puzzles, and the AIME competition-level problems.

- **AQUA** (Wei et al., 2022), the question answering task, which consists of 254 arithmetic reasoning questions designed to evaluate logical reasoning abilities through diverse mathematical problems. Each question is associated with five multiple-choice options labeled A through E. In this experiment, we set the pruning pool size threshold $m = 3$, the number of comparisons per direction during pruning phase $n = 3$, max generated directions to 6, max depth of reasoning to 3.

- **BBEH** (Kazemi et al., 2025) is a recently introduced benchmark aimed at advancing the evaluation of reasoning in large language models. It replaces each original task in BBH (Suzgun et al., 2022) with a novel variant that targets comparable reasoning skills while substantially increasing the difficulty. In our experiments, we select the multi-step arithmetic task from BBEH. This task incorporates new arithmetic operators, some of which are defined in terms of other operators. It also introduces a compositional operation format. In this experiment, we set the pruning pool size threshold $m = 4$, the number of comparisons per direction during pruning phase $n = 4$, max generated directions to 8, max depth of reasoning to 6.

- **GSM8K** (Cobbe et al., 2021) is a widely-used benchmark of grade-school math word problems that require multi-step reasoning to solve. Each problem involves basic arithmetic operations and logical thinking to arrive at the final answer. The dataset contains high-quality linguistically diverse questions, making it a standard testbed for evaluating the mathematical reasoning capabilities. In this experiment, we set the pruning pool size threshold $m = 3$, the number of comparisons per direction

during pruning phase $n = 3$, max generated directions to 6, max depth of reasoning to 3.

- **The Game of 24** (Yao et al., 2023) is a mathematical challenge in which the objective is to combine four given numbers using basic arithmetic operations to yield a total of 24. For our experiments, we adopt the same dataset and setup as, which includes 1,362 problems sourced from 4nums.com. In this experiment, we set the pruning pool size threshold $m = 3$, the number of comparisons per direction during pruning phase $n = 3$, max generated directions to 6, max depth of reasoning to 6.

- **The Sudoku** (Long, 2023) includes 10 puzzles each for 3x3, 4x4, and 5x5 grid sizes. Each puzzle is partially filled, and the task is to complete the grid without altering the provided numbers. A solution is considered correct if the completed grid adheres to all standard Sudoku rules. In this experiment, we set the pruning pool size threshold $m = 4$, the number of comparisons per direction during pruning phase $n = 4$, max generated directions to 8, max depth of reasoning to 6.

- **AIME** (Mathematical Association of America, 2024) is a highly prestigious and challenging mathematics contest administered to top-performing participants of the AMC. It serves as a key benchmark for evaluating the mathematical reasoning and problem-solving capabilities of large language models. Here we set the pruning pool size threshold $m = 4$, the number of comparisons per direction during pruning phase $n = 4$, max generated directions to 8, max depth of reasoning to 6.

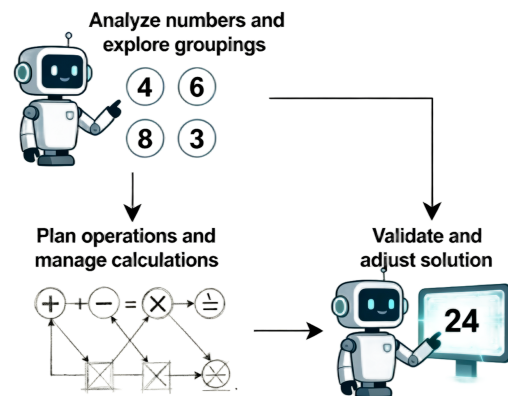


Figure 4: Visualizing the Phases of Game 24

Methods	Datasets					Average (%)
	AQUA (%)	BBEH (%)	GSM8K (%)	Game of 24 (%)	AIME (%)	
CoT	69.2	9.1	70.9	65.7	51.4	53.3
Self-Refine	75.3	8.9	71.4	73.5	70.3	59.9
SToT	72.2	7.3	69.2	61.3	62.1	54.4
PoT	66.8	12.9	92.4	86.8	65.5	64.9
CToT	85.4	28.2	84.8	76.3	74.1	69.8
BoT	79.3	20.2	93.8	73.1	46.1	62.5
ARCA	86.2	54.5	96.7	91.3	72.7	80.3

Table 4: Performance comparison on different datasets

Method	Acc. 3x3 (%)	Acc. 4x4 (%)	Acc. 5x5 (%)
CoT	90.0	73.3	60.0
Self-Refine	83.3	90.0	50.0
SToT	100.0	80.0	30.0
PoT	90.0	70.0	60.0
CToT	100.0	100.0	76.3
BoT	90.0	90.0	70.0
ARCA	100.0	100.0	90.0

Table 5: Average accuracy on Sudoku Puzzles

C.2 Additional Results

We conducted additional experiments using the Qwen3-8B model (Yang et al., 2025), comparing our results with two key baselines: the fundamental CoT method and one of the top-performing baselines, CToT. The results are presented in Table 7. To validate the generalization capability of our model, we also include experimental results on the smaller model, Qwen2.5-7B and the proprietary GPT-4o-mini, in Table 8. The results demonstrate that our method effectively adapts to different models, exhibiting strong generalization performance.

C.3 Additional Ablation Study

To evaluate our method, we conducted ablation studies on the AQUA, GSM8K, and Game of 24 datasets Table 6, focusing on two key questions: 1) whether the selection is sensitive to the initial ordering of options, and 2) how critical the Borda count component is. Specifically, we designed two experimental variants: **Random**, where we shuffle the options before each Borda count to eliminate positional bias; **w/o Gen**, where the phase generator is removed from our framework; and **w/o Borda**, where we replace the Borda count with a simpler selection process to isolate its contribution.

The experimental results demonstrate that our Borda selection method effectively mitigates positional bias, maintaining robust performance even when the option order is randomized. Meanwhile,

the phase generator helps guide the direction of reasoning and plays an important role in complex reasoning tasks. Furthermore, the component ablation study confirms the critical role of the Borda framework within the selection pipeline. It significantly reduces selection biases of the LLM while ensuring algorithmic stability, thereby enabling efficient and accurate selection of subsequent reasoning paths.

C.4 The Analysis of the Generated Phases

In this section, we evaluate the quality of the generated phases, illustrated with concrete examples from Sudoku, Game of 24, and GSM8K. We also provide a visual representation of the phases for Game of 24 in Figure 4.

Phases of Sudoku [*'Basic Elimination', 'Candidate Reasoning', 'Guessing and Backtracking'*].

Phases of Game 24 [*'Analyze numbers and explore groupings - Identify relationships and test pairing possibilities', 'Plan operations and manage calculations - Determine sequence and ensure mathematical viability', 'Validate and adjust solution - Verify results and refine approach to achieve 24'*].

Phases of GSM8K [*'Calculate the total number of eggs consumed for breakfast and baking', 'Calculate the daily earnings from selling the remaining eggs', 'Final verification or solution step and give the final answer'*].

These phases effectively distill the underlying logic of the task into structured high-level guidance, providing a robust framework that accurately directs the subsequent reasoning process.

Moreover, the system demonstrates notable robustness. It is relatively insensitive to minor fluctuations in phase granularity or precision, as the comprehensive inference process detailed in D.3 effectively compensates for high-level abstractions, ensuring stable overall performance. This ro-

Methods	AQUA (%)	GSM8K(%)	Game of 24(%)
Random	84.8	95.9	91.8
w/o Gen	72.4	70.6	69.9
w/o Borda	76.1	79.4	75.2
ARCA	86.2	96.7	91.3

Table 6: Additional ablation study

business stems from the collaborative effect of the downstream Phase Classifier and the Borda-based selection mechanism, which jointly guide the model toward phase-consistent reasoning trajectories, thereby mitigating and compensating for potential imprecision in the upstream phase generation.

D Implementation Details

D.1 Details of the Selection Algorithm

In this section, we provide a detailed explanation of how the Borda-aggregated direction selection algorithm identifies the most preferred direction in Equation 4. A prevalent strategy for identifying the optimal reasoning path through LLM preference comparisons involves leveraging pairwise methods, exemplified by the dueling bandit framework (Zhang et al., 2024). In this setting, when comparing two candidate thoughts i and j , candidate i is chosen with probability $q(i, j)$, while candidate j is selected with the complementary probability $q(j, i) = 1 - q(i, j)$. Here, $q(i, j) \geq \frac{1}{2}$ whenever i is ranked higher than j .

However, dueling bandit algorithms such as DTS (Wu and Liu, 2016) typically rely on the Copeland score (Zoghi et al., 2015) to aggregate comparison outcomes. A major limitation of the Copeland score in LLM-based preference assessment is its sensitivity to minor preference variations (Goel et al., 2017). This sensitivity arises from its win-counting mechanism, which can amplify stochastic fluctuations inherent in LLM judgments (Li et al., 2025b). Consequently, achieving stable rankings often requires extensive comparisons, which is especially challenging in noisy evaluation environments (Qin et al., 2023). To address this limitation, the Borda score (Rothe, 2019) is adopted as an alternative. The Borda score for a candidate direction i is defined as:

$$\text{Borda}(i) = \frac{1}{|\mathcal{C}| - 1} \sum_{\substack{j \in \mathcal{C} \\ j \neq i}} q(i, j), \quad (7)$$

where \mathcal{C} denotes the set of candidates. The Borda

score’s win-rate formulation effectively aggregates pairwise preferences and offers clear practical advantages in LLM evaluation settings. Its scoring mechanism, which estimates the average probability of victory, is well-suited to the stochastic and noisy nature of LLM judgments. By averaging outcomes across multiple comparisons, it confers robustness against minor inconsistencies in individual assessments (Rothe, 2019). Furthermore, the computational simplicity of maintaining and updating win rates enables highly efficient implementation in large-scale scenarios, allowing broad candidate coverage without exhaustive evaluations.

We formulate direction selection as a Borda score-based framework (Yan et al., 2022; Clarke et al., 2021), with an LLM serving as the preference function. In practice, we define the $q(i, j)$ as a binary indicator:

$$q(i, j) = \begin{cases} 1 & \text{if } i \text{ is preferred over } j \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Our algorithm begins with a pruning phase to efficiently eliminate clearly suboptimal directions while retaining the most promising candidates. During each pruning iteration, approximate Borda scores are computed by comparing each candidate against a fixed-size random subset of opponents. This sparse comparison strategy ensures broad coverage without exhaustive evaluations. Candidates with scores below a elimination score are pruned. We set the elimination score at 0.5, which corresponds to random chance performance, while any candidate scoring below this level is deemed inferior and removed. This pruning process is repeated iteratively until the number of remaining candidates falls below a predefined threshold. In **extreme cases** where the scores of all candidates are close to 0.5, making them difficult to distinguish quickly, the pruning phase will be halted after a limited number of rounds if no clear selection has been made. The system then proceeds by selecting the top- m directions with the highest current scores and enters the final phase.

Methods	Datasets					Average (%)
	AQUA (%)	BBEH (%)	GSM8K (%)	Game of 24 (%)	Sudoku Puzzle 5×5 (%)	
CoT	79.4	16.7	79.2	75.7	70.0	64.2
CToT	84.9	28.8	87.7	82.4	73.3	71.4
ARCA	90.6	47.1	93.5	91.1	90.0	82.4

Table 7: Additional results with Qwen3-8B

Models	Datasets				
	AQUA (%)	BBEH (%)	GSM8K (%)	Game of 24 (%)	Sudoku Puzzle 5×5 (%)
Qwen2.5-7B	84.7	40.2	89.4	90.2	80
GPT-4o-mini	89.8	45.6	94.7	88.3	90

Table 8: Additional results with Qwen2.5-7B and GPT-4o-mini

The algorithm then proceeds to a final evaluation stage, conducting full round-robin comparisons among the remaining candidates. This enables accurate, high-confidence estimation of the true Borda scores, from which the top-scoring candidate is chosen as the final solution. By combining efficient broad pruning with precise final assessment, this two-stage approach effectively balances computational efficiency with selection reliability. Details and further analysis are provided in Algorithm 1 and Appendix D.2.

Complexity Analysis We measure cost by the number of LLM preference queries $\text{preference}(u, v)$. Let the initial pool size be K_0 and the pruning threshold be m . In each pruning iteration, every direction is compared against n sampled opponents, yielding $\mathcal{O}(Kn)$ queries when the current pool size is K . Let K_t denote the pool size at iteration t until it reaches $K_T \leq m$, giving total pruning cost $\sum_{t=1}^T K_t n$. In the best case, the pool shrinks geometrically (e.g., removing a constant fraction each iteration), so $\sum_t K_t = \mathcal{O}(K_0)$ and the pruning cost is $\mathcal{O}(nK_0)$. In the worst case, the pool shrinks only by $\mathcal{O}(1)$ items per iteration, giving $T = \mathcal{O}(K_0)$ and $\sum_t K_t = \mathcal{O}(K_0^2)$, for a pruning cost of $\mathcal{O}(nK_0^2)$. After pruning, the final pool of size $K' \leq m$ undergoes full pairwise comparison, costing $\mathcal{O}(K'^2) \leq \mathcal{O}(m^2)$. Therefore the overall complexity is $\mathcal{O}(nK_0 + m^2)$ in the best case and $\mathcal{O}(nK_0^2 + m^2)$ in the worst case; under simplifications and approximations, the dominant terms are $\mathcal{O}(nK_0)$ and $\mathcal{O}(nK_0^2)$ respectively.

D.2 Analysis of Borda Score in LLM Feedback

To assess the ability of the Borda score to accommodate uncertainty, fine-grained distinctions, and potential ties commonly encountered in LLM preference feedback (Li et al., 2025b), we simulate a test scenario based on the setup and results from (Yan et al., 2022). The test case represents a scenario with no single winner and many ties, mirroring the challenges of LLM preference judgments.

Test Case:

$$\begin{aligned}
 q_{0,1} &= q_{1,0} = 0.5 & 1227 \\
 i > 1 &\implies q_{0,i} = 0.75 \quad \text{and} \quad q_{i,0} = 0.25 & 1228 \\
 i > 1 &\implies q_{1,i} = 0.75 \quad \text{and} \quad q_{i,1} = 0.25 & 1229 \\
 i > 1 \text{ and } j > 1 &\implies q_{i,j} = 0.5 & 1230
 \end{aligned}$$

We adopted experimental parameters and evaluation criteria of (Yan et al., 2022). The experiment involved a large set of 100 options, and a fixed budget of 1000 comparisons. As shown in Table. 9, the simulation results demonstrate the superior performance of the Borda score method over the Copeland-based approach in identifying optimal outputs from LLM preference feedback.

D.3 The Prompt Example of Different Components

Reasoning Chain Generator The reasoning chain generator in our framework is designed using the following prompt. It begins by performing a detailed task analysis to identify and delineate fundamental reasoning phases. It emphasizes logical continuity between phase, and each phase clearly defines what should be achieved, rather than prescribing how. Based on this analysis, it constructs a clear and efficient reasoning blueprint.

Algorithm	One found	Both found	Comparisons	Assessors
DTS(Copeland score) (Wu and Liu, 2016)	710	7	1000	4-20
Borda score method (Yan et al., 2022)	733	81	616-795	3-6

Table 9: Results of test case, modeling LLM preference feedback with uncertainty and ties.

1250	This blueprint directs subsequent operations along		1295
1251	logically coherent pathways and prioritizes high-	**Output Format Strictly Follow	1296
1252	value reasoning trajectories. As a result, the gener-	This Pattern:**	1297
1253	ator ensures comprehensive problem coverage, sig-	1. [Action-oriented sub-phase	1298
1254	nificantly improves computational efficiency, and	description] ,	1299
1255	overcomes the limitations of unstructured chain-of-	2. [Next essential sub-phase] ,	1300
1256	thought reasoning.	...	1301
			1302
1257	chain_prompt = '''You are an	**Critical Reminders:**	1303
1258	expert task decomposer. Your	- Phases should answer "what	1304
1259	role is to analyze complex	needs to be done" not "how to	1305
1260	problems and break them down	do it"	1306
1261	into essential high-level sub-	- Avoid transitional words ("then	1307
1262	phases. Each sub-phase should	", "next", "after")	1308
1263	represent a critical milestone	- Exclude mathematical symbols,	1309
1264	that moves toward solving the	formulas, or specific methods	1310
1265	task.	- Maintain consistent verb tense	1311
		and clarity	1312
1266	**** Generate only the most	- Ensure sub-phases are truly	1313
1267	essential sub-phases needed to	sequential and complementary	1314
1268	complete this task, excluding	...	1315
1269	all implementation details		
1270	and optional steps.	Phase Classifier The phase classifier enhances	1316
1271	**SUB-Phase DEFINITION:**	complex reasoning by dynamically identifying the	1317
1272	Each sub-phase should specify	current phase in real time, enabling the solver to	1318
1273	WHAT needs to be accomplished,	strategically allocate computational resources to-	1319
1274	not HOW to do it. Focus on	ward phase-specific objectives. Once a phase con-	1320
1275	the key objectives that must	cludes, the module seamlessly transitions to an-	1321
1276	be achieved.	other, directing the LLM's resources to the most	1322
1277		relevant ongoing stage. By focusing efforts on the	1323
1278	**Example Demonstrations:**	active phase and reducing investments in completed	1324
1279		or irrelevant directions, it maintains efficient and	1325
1280	**Geometry Problem:**	targeted progress throughout the reasoning process,	1326
1281	Task: "Find the area of a	thereby avoiding wasteful allocation.	1327
1282	triangle with base 8cm and		
1283	height 5cm"	classifier_prompt = f'''You are a	1328
1284		Sub-phase Reasoning Engine. I	1329
1285	Key Sub-Phases:	will give the sub-phase list	1330
1286	1. Identify the area formula	:{phases_list} and current	1331
1287	for triangles ,	thinking progress:{	1332
1288	2. Extract given dimensions	current_context}. Analyze the	1333
1289	from the problem ,	task progress and determine:	1334
1290	3. Compute the area using the	Which sub-phase should be	1335
1291	formula	actively worked on now.	1336
1292	4. Final verification or	*****	1337
1293	solution step and give the	Choose the sub-phase from the	1338
1294	final answer	list:{phases_list}, give me	1339

Algorithm 1 Borda-Aggregated Direction Selection

```

1: Input:
2: Pool: Set of  $K$  reasoning directions  $\{1, \dots, K\}$  to select
3:  $n$ : Number of comparisons per direction during pruning phase
4:  $m$ : Final pool size threshold
5: Preference( $u, v$ ): LLM comparison function

6: while  $K > m$  do ▷ Pruning Phase
7:    $E \leftarrow \emptyset$ 
8:   for each direction  $i \in \{1, \dots, K\}$  in the Pool do
9:     samples  $\leftarrow$  randomly select  $\min(n, K - 1)$  directions from  $\{1, \dots, K\} \setminus \{i\}$ 
10:    for each  $j \in$  samples do
11:       $E \leftarrow E \cup \{(i, j)\}$ 
12:    end for
13:  end for
14:   $\mathbf{W} \leftarrow \mathbf{0}^K, \mathbf{C} \leftarrow \mathbf{0}^K$  ▷ Reset counters
15:  for each  $(u, v) \in E$  do
16:    winner  $\leftarrow$  preference( $u, v$ )
17:     $\mathbf{C}[u] \leftarrow \mathbf{C}[u] + 1, \mathbf{C}[v] \leftarrow \mathbf{C}[v] + 1$ 
18:     $\mathbf{W}[\text{winner}] \leftarrow \mathbf{W}[\text{winner}] + 1$ 
19:  end for
20:  Borda  $\leftarrow [\mathbf{W}[i]/\mathbf{C}[i]$  for  $i \in \{1, \dots, K\}]$ 
21:  NewPool  $\leftarrow \{i \mid \text{Borda}[i] \geq 0.5\}$ 
22:  Pool  $\leftarrow$  NewPool,  $K \leftarrow |\text{Pool}|$ 
23: end while

24:  $E_{\text{final}} \leftarrow \{(i, j) \mid i, j \in \text{Pool}, i \neq j\}$  ▷ Final Evaluation Phase
25:  $\mathbf{W} \leftarrow \mathbf{0}^K, \mathbf{C} \leftarrow \mathbf{0}^K$ 
26: for each  $(u, v) \in E_{\text{final}}$  do
27:   winner  $\leftarrow$  preference( $u, v$ )
28:    $\mathbf{C}[u] \leftarrow \mathbf{C}[u] + 1, \mathbf{C}[v] \leftarrow \mathbf{C}[v] + 1$ 
29:    $\mathbf{W}[\text{winner}] \leftarrow \mathbf{W}[\text{winner}] + 1$ 
30: end for
31: Borda  $\leftarrow [\mathbf{W}[i]/\mathbf{C}[i]$  for  $i \in \{1, \dots, K\}]$ 

32: return  $\{j \mid \text{Borda}[j] = \max_{i \in \{1, \dots, K\}} \text{Borda}[i]\}$  ▷ Set of all max-scoring directions

```

1340	the number of index in the	You choose [Analysis]	1363
1341	list.		1364
1342		Output:	1365
1343	**TASK ANALYSIS PROCESS:**	1	1366
1344	1. Compare current progress with		1367
1345	each sub-phase's requirements	**CRITICAL RULES:**	1368
1346	2. Identify the most immediate	- ***** Sub-phase MUST be from	1369
1347	sub-phase that needs attention	provided list and return in	1370
1348	3. Verify the selection matches	the number of index	1371
1349	logical progression	- No explanations or additional	1372
1350		text	1373
1351	**OUTPUT FORMAT STRICTLY FOLLOW	- Sub-phase should logically	1374
1352	:**	follow from current_context	1375
1353	[index]	'''	1376
1354			
1355	**EXAMPLE:**	Reasoning Direction Generator At each reason-	1377
1356	Phases: ["Data collection", "	ing step, the reasoning direction generator takes	1378
1357	Analysis", "Validation"]	the current phase and contextual state as input and	1379
1358	Current Context: "Finished	produces a focused set of actionable, executable	1380
1359	gathering raw data, need to	directions. These outputs provide timely and tar-	1381
1360	process it"	geted guidance aligned with the specific objectives	1382
1361		of the phase.	1383
1362			

1384	direction_prompt = f'''You are a	Phase: "Validate password	1436
1385	Phase-Oriented Direction	strength"	1437
1386	Generator. Given the current	Directions:	1438
1387	step and the phase which needs	- Check minimum password length	1439
1388	to be achieved, Generate	- Verify mixed character types	1440
1389	between {min_directions} and {	- Test against common passwords	1441
1390	max_directions} practical		1442
1391	methods (directions) to	=== FORMAT REQUIREMENTS ===	1443
1392	achieve the specified phase.	Output MUST be:	1444
1393	**Current step: {current_step}	- [Direction 1]	1445
1394	**Phase:** {current_phase}	- [Direction 2]	1446
1395		...	1447
1396	**Direction Definition:**	- [Direction n] (where n is	1448
1397	Each direction should be a	between {min_directions} and {	1449
1398	concrete, actionable method	max_directions})	1450
1399	that:	**Critical Rules:**	1451
1400	1. Directly contributes to	- STRICTLY use the given format	1452
1401	achieving the phase according	- NO numbering or other formats	1453
1402	to current step	- NO additional text outside	1454
1403	2. Represents a distinct approach	bullet points	1455
1404	or technique	- Directions must answer "how to	1456
1405	3. Is executable without external	achieve the phase based on	1457
1406	knowledge	current step"	1458
1407		...	1459
1408	**Output Requirements:**		
1409	- Generate between {	Thought Generator We use the following	1460
1410	min_directions} and {	prompt to generate the thought at each step, based	1461
1411	max_directions} directions	on the task description, phase and optimal direc-	1462
1412	- Each direction must start with	tion.	1463
1413	an action verb		
1414	- Format each direction as a	purpose_prompt = f'''You are a	1464
1415	bullet point ("- [direction	heuristic assistant	1465
1416	description]")	specialized in sub-phase-based	1466
1417	- Keep directions concise (5-15	problem solving.	1467
1418	words)		1468
1419	- Exclude explanations or	**CURRENT SUB-PHASE:** {phase}	1469
1420	examples	**REQUIRED DIRECTION:** {	1470
1421		direction[i]}	1471
1422	**Quality Validation:**		1472
1423	- Each direction is a distinct	**TASK:** Generate exactly the	1473
1424	method (not a restatement)	next step that:	1474
1425	- Directions cover different	1. Directly applies the specified	1475
1426	aspects of the phase	direction: "{direction[i]}"	1476
1427	- Methods are practical and	2. Advances the current sub-phase	1477
1428	executable	: "{phase}"	1478
1429	- Avoid overlapping or redundant	3. Reach the phase as fast as	1479
1430	directions	possible !!!	1480
1431			1481
1432	=== COMPLETE EXAMPLES ===	**OUTPUT FORMAT RULES:**	1482
1433	Example:	- The next step should reach the	1483
1434	Current step: "Already generate	phase as fast as possible.	1484
1435	several passwords."		

1485	- However, when the final step	aligns closely with typical human decision-making	1533
1486	leads you to the final answer,	patterns.	1534
1487	give me only the numerical	The reasoning and selection process of AIME	1535
1488	answer and print "###" before	task:	1536
1489	it, format as: ###[ANSWER]		
1490	- Otherwise, provide a clear	Step 0 (Round 0)	1537
1491	action step	Phase: Select the 0th option (1538
1492	- No explanations, just the step	Establish the core geometric	1539
1493	itself	model)	1540
1494		Available options:	1541
1495	**VERIFICATION CHECKLIST:**	"Let speed be v , height h . At 45° :	1542
1496	- Does this step directly follows	$\text{distance} = h$, at 60° :	1543
1497	the direction "{direction[i	$\text{distance} = h/\sqrt{3}$. Distance	1544
1498]}"?	covered in 10 min = $h - h/\sqrt{3} =$	1545
1499	- Does this step achieves the sub	$h(1-1/\sqrt{3})$ "	1546
1500	-phase "{phase}" as fast as	"Set up trigonometric equations:	1547
1501	possible?	$\tan 45^\circ = h/x \Rightarrow h = x$, $\tan 60^\circ =$	1548
1502	- If final answer, does it start	$h/y \Rightarrow h = y\sqrt{3}$ "	1549
1503	with "###"?	"Complete solution: $v \times 10 = h(1-1/$	1550
1504	'''	$\sqrt{3})$, remaining distance $h/\sqrt{3}$,	1551
		time = $10/(\sqrt{3}-1) = 5(\sqrt{3}+1)$ "	1552
1505	LLM comparison function At each step, the	"Use cotangent: $h \cot 45^\circ = h$, h	1553
1506	generated direction are compared through Borda-	$\cot 60^\circ = h/\sqrt{3}$, find speed from	1554
1507	aggregated direction selection framework using	10 min travel"	1555
1508	LLM comparison function, and the example of this	Selection: Option 2 - Set up	1556
1509	comparison function is presented below.	basic trigonometric equations	1557
			1558
1510	preference_prompt=''As an	Step 1 (Round 1)	1559
1511	analytical reasoning expert,	Phase: Relate the key variables	1560
1512	critically evaluate which of	from different scenarios	1561
1513	the two reasoning paths	Available options:	1562
1514	demonstrates superior logical	"Substitute values: $x = h$, $y = h/$	1563
1515	coherence, mathematical	$\sqrt{3}$, distance = $h(1 - 1/\sqrt{3})$ "	1564
1516	accuracy, and problem-solving	"Calculate distance traveled: $x -$	1565
1517	effectiveness for the task.	$y = y\sqrt{3} - y = y(\sqrt{3} - 1)$ "	1566
1518	Consider: step-by-step	"Equate height expressions: $x = y$	1567
1519	validity, premise consistency,	$\sqrt{3}$ "	1568
1520	conclusion support, and error	"Introduce speed: $v \times 10 = x - y$ "	1569
1521	minimization. If both are	Selection: Option 3 - Equate the	1570
1522	objectively equal in all	two height expressions	1571
1523	aspects, randomly select 1 or		1572
1524	2. Output must be exactly 1 or	Step 2 (Round 2)	1573
1525	2 with no additional text,	Phase: Execute the complete	1574
1526	explanations, or formatting	algebraic derivation to solve	1575
1527	.'''	for the target quantity	1576
		Available options:	1577
1528	D.4 The Reasoning and Selection Process	"Calculate remaining time: $t = y/$	1578
1529	We present the reasoning steps and selection pro-	$v = y/[(x-y)/10] = 10y/(x-y) =$	1579
1530	cess employed by ARCA on AIME and BBEH	$10/(\sqrt{3}-1) = 5(\sqrt{3}+1)$ "	1580
1531	tasks. The results indicate that the reasoning pro-	"Express distance traveled in 10	1581
1532	cedure, along with the Borda selection mechanism,	min as $x - y$ "	1582

1583	"Rationalize denominator: $10/(\sqrt{3}-1) \times (\sqrt{3}+1)/(\sqrt{3}+1) = 5(\sqrt{3}+1)$ "	since raven is third, place robin in position 2."	1633
1584			1634
1585		"Given raven is third from left: Position 3 = raven. Robin is left of raven, so robin must be in position 2."	1635
1586	"From $x = y\sqrt{3}$, find $y = x/\sqrt{3}$, then compute time"		1636
1587			1637
1588	Selection: Option 1 - Calculate the final remaining time with complete computation	"Place the raven as the third bird from the left: Position 3 = raven."	1638
1589			1639
1590			1640
1591			1641
1592	Final Computation Process:	Selection: Option 4 - Place the raven as the third bird from the left: Position 3 = raven.	1642
1593	From $x = y\sqrt{3}$, distance in 10 min: $x - y = y\sqrt{3} - y = y(\sqrt{3} - 1)$	Step 2 (Round 2)	1643
1594	Speed $v = y(\sqrt{3} - 1)/10$	Phase: Apply the remaining constraints to finalize the entire arrangement	1644
1595	Remaining distance to base: y		1645
1596	Time needed: $t = y/v = y / [y(\sqrt{3} - 1)/10] = 10/(\sqrt{3} - 1)$	Available options:	1646
1597	Rationalize: $t = 10(\sqrt{3} + 1)/[(\sqrt{3} - 1)(\sqrt{3} + 1)] = 10(\sqrt{3} + 1)/2 = 5(\sqrt{3} + 1)$	"Set quail at position 5 (rightmost)."	1647
1598		"Place the falcon in the only remaining position (position 4)."	1648
1599		"Position robin left of raven: since position 1 is owl and position 3 is raven, robin must be in position 2."	1649
1600		"Confirm all constraints are satisfied with the complete arrangement."	1650
1601		Selection: Option 1 - Set quail at position 5 (rightmost).	1651
1602		Final Computation Process:	1652
1603	Final Result	Position 1: owl (given as leftmost)	1653
1604	Algorithm answer: $A [5(\sqrt{3} + 1)]$	Position 3: raven (given as third from left)	1654
1605	Correct answer: A	Position 2: robin (must be left of raven, and position 1 is occupied by owl)	1655
1606	The reasoning and selection process of BBEH task:	Position 5: quail (given as rightmost)	1656
1607		Position 4: falcon (only remaining position)	1657
1608	Step 0 (Round 0)	The rightmost bird is in position 5, which is the quail.	1658
1609	Phase: Establish the boundary framework for reasoning	Final Result:	1659
1610	Available options:	Algorithm answer: A (The quail is the rightmost)	1660
1611	"Place the raven as the third bird from the left."	Correct answer: A	1661
1612	"Position robin immediately left of raven in the remaining slots (positions 2 and 3)."		1662
1613	"Determine the rightmost bird as quail."		1663
1614	"Identify the leftmost bird as owl: Position 1 = owl."		1664
1615	Selection: Option 4 - Identify the leftmost bird as owl: Position 1 = owl.		1665
1616	Step 1 (Round 1)		1666
1617	Phase: Establish the key anchor point and determine the fixed point		1667
1618	Available options:		1668
1619	"Assign positions: Position 1 = owl (leftmost), Position 5 = quail (rightmost)."		1669
1620	"Position robin left of raven:		1670
1621			1671
1622			1672
1623			1673
1624			1674
1625			1675
1626			1676
1627			1677
1628			1678
1629			1679
1630			1680
1631			1681
1632			1682