

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 THINK SMARTER, FOCUS WISELY: ADAPTIVE COGNITIVE ALLOCATION FOR LLM REASONING

Anonymous authors

Paper under double-blind review

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 pre-defined 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 & 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.

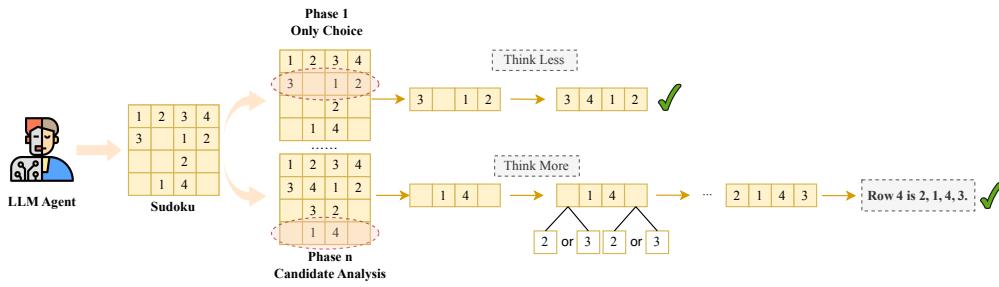


Figure 1: LLMs operate through multiple phases to solve a task, with certain phases requiring minimal cognitive resources while others demand in-depth reasoning.

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.

108

2 RELATED WORK

110 **Reasoning Methods for Robustness** A dominant strategy to enhance reasoning robustness is to
 111 generate and evaluate a multitude of augmented reasoning procedures. This approach moves beyond
 112 a single chain of thought (Wei et al., 2022) by explicitly constructing multiple reasoning trajec-
 113 tories. For instance, Tree of Thoughts (Yao et al., 2023) frames reasoning as a tree search, allowing
 114 exploration of parallel thought candidates at each step. Graph of Thoughts (Besta et al., 2024)
 115 further generalizes this into a graph structure to capture more complex interdependencies between
 116 thoughts. Methods like Boosting Task-Oriented Reasoning (Chen et al., 2024b) iteratively generate
 117 numerous reasoning steps and use LLM-based error analysis to refine them, while Comparative Tree
 118 of Thought (Zhang et al., 2024) employs pairwise comparison to select optimal paths. CoDT (Wang
 119 et al., 2025) improves the robustness against reference corruption by providing a few exemplars with
 120 structured and defensive reasoning as demonstrations. MME-CoT (Jiang et al., 2025) incorporates
 121 three novel metrics to assess the reasoning quality, robustness, and efficiency. Math-RoB (Yu et al.,
 122 2025b) uses an instruction-based approach to generate diverse datasets resembling training distribu-
 123 tions. CD-CoT (Zhou et al., 2024) enhances LLMs’ denoising-reasoning capabilities by contrasting
 124 noisy rationales with one clean rationale. CoT-GCG (Su, 2024) enhances adversarial attacks on
 125 aligned LLMs by integrating CoT prompts with the greedy coordinate gradient technique.

126 **Efficiency-Oriented Reasoning Methods** In direct contrast, another line of research focuses on
 127 streamlining the reasoning process to reduce computational costs. These methods often employ
 128 meta-reasoning architectures or probabilistic approximations to achieve faster inference. For exam-
 129 ple, Meta-Reasoner (Sui et al., 2025b) uses contextual multi-armed bandits to dynamically adjust
 130 reasoning strategies based on state evaluation. ES-CoT (Mao et al., 2025) shortens thought gen-
 131 eration by prompting the LLM to output a step answer at each reasoning step. THINK-Bench (Li
 132 et al., 2025c) introduces a benchmark with novel efficiency metrics, to evaluate the reasoning ef-
 133 ficiency. (Cui et al., 2025) proposes a method that identifies and focuses on generating important
 134 reasoning steps in reasoning by using perplexity to measure their importance, (Sui et al., 2025a)
 135 explores efficient data use, small language model reasoning, and evaluation methods. Soft Chain-
 136 of-Thought (Xu et al., 2025) leverages probabilistic soft chains and prompt tweaks for efficient,
 137 uncertainty-aware reasoning. Similarly, COAT (Shen et al., 2025) uses action-oriented chains for
 138 meta-reasoning without full model tuning.

139 **Feedback-Based Refinement and Evaluation** Another influential line of research focuses on iter-
 140 ative self-improvement through feedback mechanisms. In this paradigm, the LLM itself is leveraged
 141 to evaluate and refine its reasoning trajectories in a cyclic manner. For instance, Self-Refine (Madaan
 142 et al., 2023) introduces an algorithm where the LLM generates output, provides self-feedback, and
 143 then refines its output based on that feedback. RCO (Yu et al., 2025a) trains critic models using a
 144 feedback loop where critiques guide the actor model in refining responses. RefCritic (Tang et al.,
 145 2025) trains a critic module with dual rule-based rewards focusing on instance-level correctness
 146 of solution judgments and refinement accuracies of the policy model. (Renze & Guven, 2024) in-
 147 vestigates the effects of self-reflection in large language models on problem-solving performance.
 148 (Potamitis & Arora, 2025) enhances reasoning by allowing the models to retry problem-solving
 149 attempts upon identifying incorrect answers. RFM-RAG (Li et al., 2025a) transforms stateless re-
 150 trieval into stateful continuous knowledge management by constructing a dynamic evidence pool to
 151 generate refined queries using relational triples and evidence. PHP (Zheng et al., 2023) incorporates
 152 the solution from a previous attempt as a hint for the next, creating an iterative improvement loop.

153 While significant progress has been made within these methods, the fundamental tension between
 154 robustness and efficiency remains largely unaddressed. ARCA addresses this core challenge by em-
 155 powering LLMs to dynamically allocate greater cognitive resources to critical phases of a reasoning
 156 task, achieving an optimal balance between accuracy and efficiency.

157

3 ADAPTIVE REASONING VIA COGNITIVE ALLOCATION

158 In this section, we present a comprehensive introduction to the ARCA framework, which achieves
 159 cognitive resource allocation to balance accuracy and efficiency in LLM reasoning. It integrates
 160 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**.

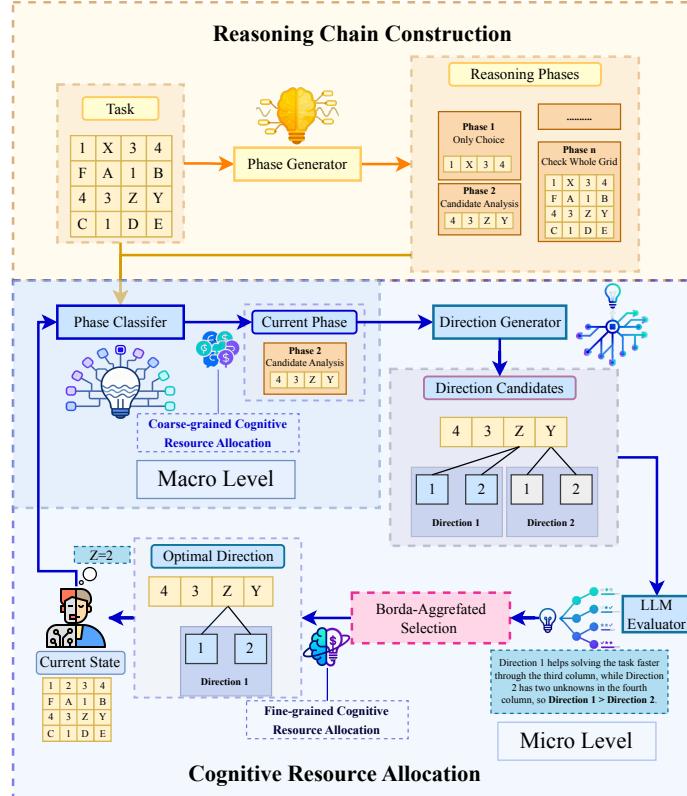


Figure 2: The overall framework of ARCA. It integrates reasoning chain construction and cognitive resource allocation to balance accuracy and efficiency in LLM reasoning, where reasoning chain construction structures the task into phases, and cognitive resource allocation allocates cognitive resources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

3.1 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 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:

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

216 where \mathcal{T} is the input task, ϕ_i is the i -th reasoning phase and \mathcal{P} is the generated phase set. ReasonGen
 217 begins by analyzing the task context to identify and delineate each fundamental phase. Each phase
 218 specifies a semantic objective—*what needs to be achieved* at that step rather than prescribing the
 219 operational details of *how* to achieve it. For the Sudoku example, this naturally yields phases start-
 220 ing with single-candidate placement, followed by iterative candidate elimination, conflict detection
 221 and backtracking, and concluding with final validation. We provide illustrative examples of the
 222 generated phases in Appendix B.4. This structured decomposition defines the allocation units (rea-
 223 soning phases) required by ARCA and constrains reasoning within explicit semantic boundaries.
 224 By providing a high-level blueprint for cognitive resource allocation, ReasonGen reduces irrelevant
 225 exploration and yields more coherent and efficient inference compared to unstructured CoT.

226 3.2 COGNITIVE RESOURCE ALLOCATION

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

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

$$239 \quad \phi_t = \text{PhaseClass}(x_t, \mathcal{P}), \quad (2)$$

240 where x_t is the reasoning state at step t . The architecture enhances complex reasoning by dynam-
 241 ically identifying the current phase, enabling the LLM to focus on phase-specific objectives. Once
 242 a phase is completed, the module transitions seamlessly to the next, reallocating computational
 243 resources according to new requirements. By concentrating effort on the active phase and minimiz-
 244 ing expenditure on completed or irrelevant directions, the system maintains targeted and efficient
 245 progress across the reasoning chain, thereby preventing resource over-allocation. A key insight
 246 underlying this design is that reasoning phases naturally impose heterogeneous resource demands:
 247 simpler phases require minimal resources and benefit from rapid closure, whereas more demanding
 248 phases call for deeper inference and thus greater resource investment. The architecture capitalizes
 249 on this heterogeneity, allowing adaptive allocation of computational effort so that resources are con-
 250 centrated on ongoing objectives while avoiding over-allocation to irrelevant directions. In doing so,
 251 the system achieves dynamic and efficient resource utilization throughout reasoning.

252 At the **Micro-Level**, the module generates and selects reasoning directions that are locally optimal
 253 or highly relevant to the active phase. Cognitive resources or reasoning efforts are dynamically
 254 assigned to the most promising next steps within the phase. To this end, we introduce a reason-
 255 ing direction generator *DirectionGen*, a real-time strategic module that steers the LLM’s inference
 256 process:

$$256 \quad \mathcal{D}_t = \{d_1, d_2, \dots, d_m\} = \text{DirectionGen}(\phi_t, x_t, \mathcal{T}), \quad (3)$$

257 where \mathcal{D}_t is the set of candidate directions at step t and d_i is the i -th direction. At each reasoning
 258 step, the generator produces a focused set of actionable directions based on the current phase and
 259 contextual state. These directions provide timely, targeted guidance aligned with the phase’s ob-
 260 jectives. To enable broad exploration, multiple candidate paths are proposed. **Selecting the most**
 261 **promising path** among them poses a central challenge, as explicit reward signals are unavailable
 262 and handcrafted reward functions are difficult to design and prone to bias. To address this, we
 263 employ an LLM as an implicit preference oracle to evaluate candidate directions.

264 Building on these evaluations, we introduce a **Borda-aggregated direction selection algo-**
 265 **rithm** (Yan et al., 2022), which consolidates pairwise comparisons into a robust consensus score
 266 and will be described in detail in section 3.3. This approach mitigates noise from individual judg-
 267 ments, reduces reliance on brittle heuristics, and reliably identifies the most promising reasoning
 268 direction. The optimal direction d_t^* is then selected according to the aggregated Borda score:

$$269 \quad d_t^* = \arg \max_{d \in \mathcal{D}_t} \text{Borda}(d). \quad (4)$$

270 The LLM then generates the next reasoning state x_{t+1} according to:
 271

$$272 \quad 273 \quad x_{t+1} = \text{LLM}(x_t, \phi_t, d_t^*). \quad (5)$$

274 By providing such fine-grained tactical guidance, this module effectively bridges phase planning
 275 with real-time reasoning, enabling efficient and rational allocation of cognitive resources throughout
 276 the inference process. In the next subsection, we will make a detailed introduction to the Borda-
 277 aggregated direction selection algorithm.
 278

279 3.3 BORDA-AGGREGATED DIRECTION SELECTION 280

281 In this section, we provide a detailed explanation of how the Borda-aggregated direction selection
 282 algorithm identifies the most preferred direction in Equation 4. To select the optimal reasoning path
 283 via LLM-based preference comparisons, prior work has often adopted the dueling bandit frame-
 284 work (Zhang et al., 2024). In this setting, when comparing two candidate thoughts i and j , candidate
 285 i is chosen with probability $q(i, j)$, while candidate j is selected with the complementary probability
 286 $q(j, i) = 1 - q(i, j)$. Here, $q(i, j) \geq \frac{1}{2}$ whenever i is ranked higher than j . Repeated comparisons
 287 are assumed to be independent.
 288

289 However, dueling bandit algorithms such as DTS (Wu & Liu, 2016) typically rely on the Copeland
 290 score (Zoghi et al., 2015) to aggregate comparison outcomes. A major limitation of the Copeland
 291 score in LLM-based preference assessment is its sensitivity to minor preference variations (Goel
 292 et al., 2017). This sensitivity arises from its win-counting mechanism, which can amplify stochastic
 293 fluctuations inherent in LLM judgments (Li et al., 2025b). Consequently, achieving stable rankings
 294 often requires extensive comparisons, which is especially challenging in noisy evaluation environ-
 295 ments (Qin et al., 2023). To address this limitation, the Borda score (Rothe, 2019) is adopted as an
 296 alternative. The Borda score for a candidate direction i is defined as:
 297

$$298 \quad \text{Borda}(i) = \frac{1}{|\mathcal{C}| - 1} \sum_{j \in \mathcal{C}, j \neq i} q(i, j),$$

299 where \mathcal{C} denotes the set of candidates. The Borda score’s win-rate formulation effectively aggregates
 300 pairwise preferences and offers clear practical advantages in LLM evaluation settings. Its scoring
 301 mechanism, which estimates the average probability of victory, is well-suited to the stochastic and
 302 noisy nature of LLM judgments. By averaging outcomes across multiple comparisons, it confers
 303 robustness against minor inconsistencies in individual assessments (Rothe, 2019). Furthermore, the
 304 computational simplicity of maintaining and updating win rates enables highly efficient implemen-
 305 tation in large-scale scenarios, allowing broad candidate coverage without exhaustive evaluations.
 306

307 We formulate direction selection as a Borda score-based framework (Yan et al., 2022; Clarke et al.,
 308 2021), with an LLM serving as the **preference function**. Our algorithm begins with a pruning phase
 309 to efficiently eliminate clearly suboptimal directions while retaining the most promising candidates.
 310 During each pruning iteration, approximate Borda scores are computed by comparing each candidate
 311 against a fixed-size random subset of opponents. This sparse comparison strategy ensures broad
 312 coverage without exhaustive evaluations. Candidates with scores below a elimination score are
 313 pruned. We set the elimination score at 0.5, which corresponds to random chance performance,
 314 while any candidate scoring below this level is deemed inferior and removed. This pruning process
 315 is repeated iteratively until the number of remaining candidates falls below a predefined threshold.
 316 The algorithm then proceeds to a final evaluation stage, conducting full round-robin comparisons
 317 among the remaining candidates. This enables accurate, high-confidence estimation of the true
 318 Borda scores, from which the top-scoring candidate is chosen as the final solution. By combining
 319 efficient broad pruning with precise final assessment, this two-stage approach effectively balances
 320 computational efficiency with selection reliability. Details are provided in Appendix C.
 321

322 4 EXPERIMENTS 323

324 In this section, we conduct comprehensive experiments to evaluate performance and provide an in-
 325 depth analysis of ARCA, comparing it against baseline approaches in terms of cost and accuracy.
 326 Additional results and ablation study are in Appendix B.2 and B.3.
 327

324
325
326 Table 1: Performance comparison on different datasets
327
328

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

335
336 Table 2: Average accuracy on Sudoku Puzzles
337

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

346
347 4.1 REASONING TASKS
348349
350 We evaluate the performance of our proposed method, ARCA, on a suite of six challenging real-
351 world reasoning tasks. These tasks span a diverse range of domains, including question answering
352 (AQUA), multi-step arithmetic (BBEH), math word problems (GSM8K), the Game of 24, Sudoku
353 puzzles, and the AIME competition-level problems.354 - **AQUA** (Wei et al., 2022), the question answering task, which consists of 254 arithmetic reasoning
355 questions designed to evaluate logical reasoning abilities through diverse mathematical problems.
356 Each question is associated with five multiple-choice options labeled A through E.357 - **BBEH** (Kazemi et al., 2025) is a recently introduced benchmark aimed at advancing the evaluation
358 of reasoning in large language models. It replaces each original task in BBH (Suzgun et al., 2022)
359 with a novel variant that targets comparable reasoning skills while substantially increasing the diffi-
360 culty. In our experiments, we select the multi-step arithmetic task from BBEH. This task incorporates
361 new arithmetic operators, some of which are defined in terms of other operators. It also introduces a
362 compositional operation format.363 - **GSM8K** (Cobbe et al., 2021) is a widely-used benchmark of grade-school math word problems
364 that require multi-step reasoning to solve. Each problem involves basic arithmetic operations and
365 logical thinking to arrive at the final answer. The dataset contains high-quality linguistically diverse
366 questions, making it a standard testbed for evaluating the mathematical reasoning capabilities.367 - **The Game of 24** (Yao et al., 2023) is a mathematical challenge in which the objective is to combine
368 four given numbers using basic arithmetic operations to yield a total of 24. For our experiments, we
369 adopt the same dataset and setup as, which includes 1,362 problems sourced from 4nums.com.370 - **The Sudoku** (Long, 2023) includes 10 puzzles each for 3x3, 4x4, and 5x5 grid sizes. Each puzzle
371 is partially filled, and the task is to complete the grid without altering the provided numbers. A
372 solution is considered correct if the completed grid adheres to all standard Sudoku rules.373 - **AIME** (Mathematical Association of America, 2024) is a highly challenging mathematics contest
374 administered to top-performing participants of the AMC. It serves as a key benchmark for evaluating
375 the mathematical reasoning and problem-solving capabilities of LLMs.376 All experiments were conducted with the Deepseek-V3 model (DeepSeek-AI, 2025). Detailed
377 configurations and results, including those with other models, are provided in Appendix B.

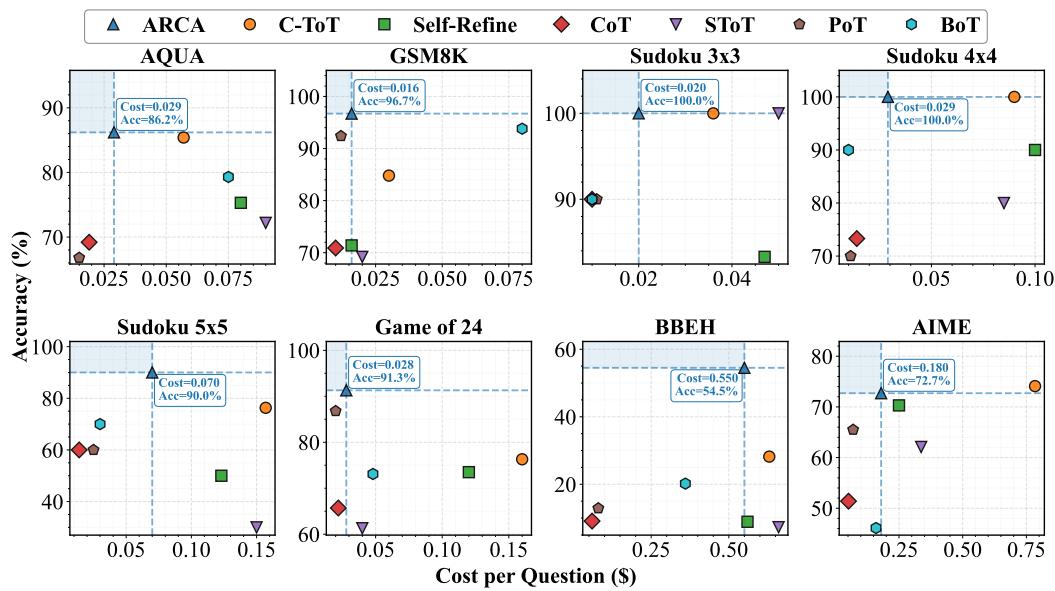


Figure 3: Accuracy–cost comparison across reasoning tasks. Each marker denotes a method. Dashed lines and the text box highlight ARCA, which none surpass in the shaded region, showing ARCA performs best among compared methods.

4.2 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. 1 and Table. 2. 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 solution 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 solution. Consequently, ARCA delivers substantially improved overall reasoning performance while simultaneously ensuring judicious control over computational expenditure, providing a reliable framework for complex cognitive tasks.

4.3 ABLATION STUDY

Analysis on the parameters in Borda-aggregated direction selection. To analyze how the the parameters in the direction selection algorithm affects accuracy and computational cost, we conduct

432 experiments on the GSM8K dataset. To ensure comparability, the maximum number of direction
 433 is fixed at 12 across all trials. Experiments are conducted with pruning pool size threshold $m =$
 434 $\{4, 6, 8, 12\}$ (where $m = 12$ corresponds to the non-pruning case) and the number of comparisons
 435 per direction during pruning set to $n = \{2, 4, 6\}$. Results in Table. 3 indicate that when using pool
 436 size pruning, configurations with pool size threshold 6 and 8 achieve performance comparable to the
 437 full pool (size 12), while significantly reducing computational expense. Although the value of n has
 438 an impact on performance, experiments show that a medium n is sufficient to achieve rapid pruning
 439 in the early stages. This demonstrates that the proposed selection framework effectively balances
 440 computational efficiency with selection reliability, thus offers a scalable and practical solution for
 441 resource-aware automated reasoning.

442 **Analysis on the Maximum Direction Number.** The number of generated directions determines
 443 the breadth of exploration available to the solver at each reasoning step. We evaluate the impact of
 444 this parameter by conducting experiments on the AQUA and Game of 24 datasets, using maximum
 445 direction numbers set to 2, 4, 6, and 10. The results, summarized in Table. 4, indicate that when the
 446 maximum direction number is limited to 2 which resulting in a narrow exploration scope, the per-
 447 formance is noticeably worse compared to configurations allowing broader exploration. In contrast,
 448 when the maximum number of directions is set to 4 or higher, performance stabilizes and remains
 449 consistently high. This suggests that the reasoning chain construction mechanism helps the solver
 450 maintain a clear objective and reduces the need for extensive exploration, thereby achieving more
 451 efficient and reliable problem-solving even with moderate search width.

452 Table 3: Ablation study on direction selection algorithm parameters on GSM8K. Here n denotes
 453 the number of comparisons and m denotes the size. **Acc.** is accuracy (%), and **Cost** is average
 454 computational cost (lower is better).

	$m = 4$		$m = 6$		$m = 8$		$m = 12$	
	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

463 Table 4: Ablation study on the number of directions d_n . Accuracy (%) is reported on AQUA and
 464 Game of 24, along with the average. Larger d_n generally improves performance, with diminishing
 465 gains after $d_n = 6$.

Direction Number 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

475 CONCLUSION

476
 477 In conclusion, we propose a novel framework ARCA designed to tackle efficient cognitive resource
 478 allocation in LLM reasoning, with a specific focus on balancing accuracy and efficiency. By in-
 479 tegrating decomposition, strategy generation, monitoring, and dynamic selection into a cohesive
 480 system, our approach enhances structural coherence, optimizes reasoning effort, and improves accu-
 481 racy in complex scenarios while keeping additional reasoning cost negligible to preserve efficiency.
 482 Extensive experiments across diverse tasks show that our method delivers strong performance while
 483 maintaining competitive resource efficiency compared to existing baselines. In future work, we will
 484 pursue more reliable reasoning chains and refine our framework for accurate direction generation,
 485 focusing on more complex task environments.

486 REFERENCES
487

488 Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gian-
489 inazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczek, et al. Graph of
490 thoughts: Solving elaborate problems with large language models. In *The 38th AAAI conference
491 on Artificial Intelligence (AAAI)*, 2024.

492 Jiaao Chen, Xiaoman Pan, Dian Yu, Kaiqiang Song, Xiaoyang Wang, Dong Yu, and Jianshu Chen.
493 Skills-in-context: Unlocking compositionality in large language models. In *Findings of the Asso-
494 ciation for Computational Linguistics (EMNLP)*, pp. 13838–13890, 2024a.

495 Sijia Chen, Baochun Li, and Di Niu. Boosting of thoughts: Trial-and-error problem solving with
496 large language models. In *The 12th International Conference on Learning Representations
497 (ICLR)*, 2024b.

498 Wenhui Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompt-
499 ing: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint
500 arXiv:2211.12588*, 2023.

502 Charles LA Clarke, Alexandra Vtyurina, and Mark D Smucker. Assessing top-preferences. *ACM
503 Transactions on Information Systems (TOIS)*, 39(3):1–21, 2021.

504 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
505 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
506 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,
507 2021.

509 Yingqian Cui, Pengfei He, Jingying Zeng, Hui Liu, Xianfeng Tang, Zhenwei Dai, Yan Han, Chen
510 Luo, Jing Huang, Zhen Li, Suhang Wang, Yue Xing, Jiliang Tang, and Qi He. Stepwise perplexity-
511 guided refinement for efficient chain-of-thought reasoning in large language models. *arXiv
512 preprint arXiv:2502.13260*, 2025.

514 DeepSeek-AI. DeepSeek-V3 technical report. *arXiv preprint arXiv:2412.19437*, 2025.

515 Ashish Goel, Anilesh K Krishnaswamy, and Kamesh Munagala. Metric distortion of social choice
516 rules: Lower bounds and fairness properties. In *The 2017 ACM Conference on Economics and
517 Computation (EC)*, 2017.

518 Deyi Ji, Lanyun Zhu, Siqi Gao, Peng Xu, Hongtao Lu, Jieping Ye, and Feng Zhao. Tree-of-table:
519 Unleashing the power of LLMs for enhanced large-scale table understanding. *arXiv preprint
520 arXiv:2411.08516*, 2024.

522 Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanwei Li, Yu Qi, Xinyan Chen, Liuwei Wang, Jianhan
523 Jin, Claire Guo, Shen Yan, Bo Zhang, Chaoyou Fu, Peng Gao, and Hongsheng Li. Mme-cot:
524 Benchmarking chain-of-thought in large multimodal models for reasoning quality, robustness,
525 and efficiency. *arXiv preprint arXiv:2502.09621*, 2025.

526 Haibo Jin, Peiyan Zhang, Man Luo, and Haohan Wang. Reasoning can hurt the inductive abilities
527 of large language models. *arXiv preprint arXiv:2505.24225*, 2025.

529 Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John Palowitch, Chrysovalantis Anastasiou, San-
530 ket Vaibhav Mehta, Lalit K. Jain, Virginia Aglietti, Disha Jindal, Peter Chen, Nishanth Dikkala,
531 Gladys Tyen, Xin Liu, Uri Shalit, Silvia Chiappa, Kate Olszewska, Yi Tay, Vinh Q. Tran, Quoc V.
532 Le, and Orhan Firat. BIG-bench extra hard. *arXiv preprint arXiv:2502.19187*, 2025.

533 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay
534 Ramasesh, Ambrose Sloane, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quanti-
535 tative reasoning problems with language models. In *Advances in Neural Information Processing
536 Systems (NeurIPS)*, 2022.

538 Leqian Li, Dianxi Shi, Jialu Zhou, Xinyu Wei, Mingyue Yang, Songchang Jin, and Shaowu Yang.
539 Retrieval feedback memory enhancement large model retrieval generation method. *arXiv preprint
arXiv:2508.17862*, 2025a.

540 Tianyi Li, Yu Qin, and Olivia R. Liu Sheng. A multi-task evaluation of LLMs' processing of aca-
 541 demic text input. *arXiv preprint arXiv:2508.11779*, 2025b.

542

543 Zhiyuan Li, Yi Chang, and Yuan Wu. Think-bench: Evaluating thinking efficiency and chain-of-
 544 thought quality of large reasoning models. *arXiv preprint arXiv:2505.22113*, 2025c.

545

546 Jieyi Long. Large language model guided tree-of-thought. *arXiv preprint arXiv:2305.08291*, 2023.

547

548 Sebastian Lubos, Alexander Felfernig, Damian Garber, Viet-Man Le, Manuel Henrich, Reinhard
 549 Willfort, and Jeremias Fuchs. Towards group decision support with LLM-based meeting analysis.
 550 In *The 33rd ACM Conference on User Modeling, Adaptation and Personalization*, pp. 331–335,
 551 2025.

552

553 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri
 554 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement
 555 with self-feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

556

557 Minjia Mao, Bowen Yin, Yu Zhu, and Xiao Fang. Early stopping chain-of-thoughts in large language
 558 models. *arXiv preprint arXiv:2509.14004*, 2025.

559

560 Mathematical Association of America. American invitational mathematics examination (aime)
 561 2024, 2024. URL <https://www.maa.org/math-competitions/aime>. Part of the
 562 American Mathematics Competitions.

563

564 OpenAI. Learning to reason with LLMs. Technical report, OpenAI, 2024. URL <https://openai.com/index/learning-to-reason-with-l1m/>.

565

566 OpenAI. OpenAI GPT-4.5 system card. Technical report, OpenAI, 2025. URL <https://openai.com/index/gpt-4-5-system-card/>.

567

568 Avinash Patil and Aryan Jadon. Advancing reasoning in large language models: Promising methods
 569 and approaches. *arXiv preprint arXiv:2502.03671*, 2025.

570

571 Nearchos Potamitis and Akhil Arora. Are retrials all you need? enhancing large language model
 572 reasoning without verbalized feedback. *arXiv preprint arXiv:2504.12951*, 2025.

573

574 Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi Liu,
 575 Jialu Liu, Donald Metzler, et al. Large language models are effective text rankers with pairwise
 576 ranking prompting. *arXiv preprint arXiv:2306.17563*, 2023.

577

578 Matthew Renze and Erhan Guven. The benefits of a concise chain of thought on problem-solving in
 579 large language models. In *The 2nd International Conference on Foundation and Large Language
 580 Models (FLLM)*, 2024.

581

582 Jörg Rothe. Borda count in collective decision making: A summary of recent results. In *Proceedings
 583 of the AAAI Conference on Artificial Intelligence*, 2019.

584

585 Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu,
 586 Michael Moor, Zicheng Liu, and Emad Barsoum. Agent laboratory: Using LLM agents as re-
 587 search assistants. *arXiv preprint arXiv:2501.04227*, 2025.

588

589 Maohao Shen, Guangtao Zeng, Zhenting Qi, Zhang-Wei Hong, Zhenfang Chen, Wei Lu, Gre-
 590 gory Wornell, Subhro Das, David Cox, and Chuang Gan. Satori: Reinforcement learning
 591 with chain-of-action-thought enhances LLM reasoning via autoregressive search. *arXiv preprint
 592 arXiv:2502.02508*, 2025.

593

594 Jingbo Su. Enhancing adversarial attacks through chain of thought. *arXiv preprint
 595 arXiv:2410.21791*, 2024.

596

597 Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu,
 598 Andrew Wen, Shaochen Zhong, Na Zou, Hanjie Chen, and Xia Hu. Stop overthinking: A survey
 599 on efficient reasoning for large language models. *arXiv preprint arXiv:2503.16419*, 2025a.

594 Yuan Sui, Yufei He, Tri Cao, Simeng Han, Yulin Chen, and Bryan Hooi. Meta-reasoner: Dy-
 595 namic guidance for optimized inference-time reasoning in large language models. *arXiv preprint*
 596 *arXiv:2502.19918*, 2025b.

597

598 Mirac Suzgun, Nathan Scales, Nathanael Schärl, Sebastian Gehrman, Yi Tay, Hyung Won Chung,
 599 Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. Challenging BIG-
 600 bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*,
 601 2022.

602 Qiaoyu Tang, Hao Xiang, Le Yu, Bowen Yu, Hongyu Lin, Yaojie Lu, Xianpei Han, Le Sun, and
 603 Junyang Lin. Refcritic: Training long chain-of-thought critic models with refinement feedback.
 604 *arXiv preprint arXiv:2507.15024*, 2025.

605 Wenxiao Wang, Parsa Hosseini, and Soheil Feizi. Chain-of-defensive-thought: Structured rea-
 606 soning elicits robustness in large language models against reference corruption. *arXiv preprint*
 607 *arXiv:2504.20769*, 2025.

608

609 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha
 610 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language
 611 models. In *The 11th International Conference on Learning Representations (ICLR)*, 2023.

612 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 613 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *Advances*
 614 *in Neural Information Processing Systems (NeurIPS)*, 2022.

615

616 Huasen Wu and Xin Liu. Double thompson sampling for dueling bandits. In *Advances in Neural*
 617 *Information Processing Systems (NeurIPS)*, volume 29, 2016.

618 Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. Softcot: Soft chain-of-thought for efficient
 619 reasoning with llms, 2025. URL <https://arxiv.org/abs/2502.12134>.

620

621 Xinyi Yan, Chengxi Luo, Charles LA Clarke, Nick Craswell, Ellen M Voorhees, and Pablo Castells.
 622 Human preferences as dueling bandits. In *The 45th International ACM SIGIR Conference on*
 623 *Research and Development in Information Retrieval (SIGIR)*, 2022.

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

633 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
 634 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *Ad-*
 635 *vances in Neural Information Processing Systems (NeurIPS)*, 2023.

636

637 Tianshu Yu, Chao Xiang, Mingchuan Yang, Pei Ke, Bosi Wen, Cunxiang Wang, Jiale Cheng,
 638 Li Zhang, Xinyu Mu, Chuxiong Sun, and Minlie Huang. Training language model to critique
 639 for better refinement. *arXiv preprint arXiv:2506.22157*, 2025a.

640 Tong Yu, Yongcheng Jing, Xikun Zhang, Wentao Jiang, Wenjie Wu, Yingjie Wang, Wenbin Hu,
 641 Bo Du, and Dacheng Tao. Benchmarking reasoning robustness in large language models. *arXiv*
 642 *preprint arXiv:2503.04550*, 2025b.

643

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

647

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

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

651

652 Masrour Zoghi, Zohar S Karnin, Shimon Whiteson, and Maarten De Rijke. Copeland dueling ban-
653 dits. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A USE OF LARGE LANGUAGE MODELS (LLMs)
703
704705 Our study investigates the reasoning capabilities of large language models (LLMs). Accordingly,
706 all experiments were conducted on LLMs to evaluate and validate our proposed approach. Beyond
707 experimentation, we employed an LLM as an auxiliary tool during manuscript preparation. Specif-
708 ically, it was used to refine language for grammar and clarity, and to generate illustrative (non-
709 experimental) figures based on prompts we provided. All research ideas, methods, experiments,
710 analyses, and conclusions were developed by the authors.
711
712713 B EXPERIMENTS AND SETTINGS
714
715716 B.1 THE DETAILS OF EXPERIMENT
717
718719 We evaluate the performance of our proposed method, ARCA, on a suite of six challenging real-
720 world reasoning tasks. These tasks span a diverse range of domains, including question answering
721 (AQUA), multi-step arithmetic (BBEH), math word problems (GSM8K), the Game of 24, Sudoku
722 puzzles, and the AIME competition-level problems.723 **AQUA** (Wei et al., 2022), the question answering task, which consists of 254 arithmetic reasoning
724 questions designed to evaluate logical reasoning abilities through diverse mathematical problems.
725 Each question is associated with five multiple-choice options labeled A through E. In this experi-
726 ment, we set the pruning pool size threshold $m = 3$, the number of comparisons per direction during
727 pruning phase $n = 3$, max generated directions to 6, max depth of reasoning to 3.728 **BBEH** (Kazemi et al., 2025) is a recently introduced benchmark aimed at advancing the evaluation
729 of reasoning in large language models. It replaces each original task in BBH (Suzgun et al., 2022)
730 with a novel variant that targets comparable reasoning skills while substantially increasing the diffi-
731 culty. In our experiments, we select the multi-step arithmetic task from BBEH. This task incorporates
732 new arithmetic operators, some of which are defined in terms of other operators. It also introduces a
733 compositional operation format. In this experiment, we set the pruning pool size threshold $m = 4$,
734 the number of comparisons per direction during pruning phase $n = 4$, max generated directions to
735 8, max depth of reasoning to 6.736 **GSM8K** (Cobbe et al., 2021) is a widely-used benchmark of grade-school math word problems
737 that require multi-step reasoning to solve. Each problem involves basic arithmetic operations and
738 logical thinking to arrive at the final answer. The dataset contains high-quality linguistically diverse
739 questions, making it a standard testbed for evaluating the mathematical reasoning capabilities. In
740 this experiment, we set the pruning pool size threshold $m = 3$, the number of comparisons per
741 direction during pruning phase $n = 3$, max generated directions to 6, max depth of reasoning to 3.742 **The Game of 24** (Yao et al., 2023) is a mathematical challenge in which the objective is to combine
743 four given numbers using basic arithmetic operations to yield a total of 24. For our experiments,
744 we adopt the same dataset and setup as, which includes 1,362 problems sourced from 4nums.com.
745 In this experiment, we set the pruning pool size threshold $m = 3$, the number of comparisons per
746 direction during pruning phase $n = 3$, max generated directions to 6, max depth of reasoning to 6.747 **The Sudoku** (Long, 2023) includes 10 puzzles each for 3x3, 4x4, and 5x5 grid sizes. Each puzzle
748 is partially filled, and the task is to complete the grid without altering the provided numbers. A
749 solution is considered correct if the completed grid adheres to all standard Sudoku rules. In this
750 experiment, we set the pruning pool size threshold $m = 4$, the number of comparisons per direction
751 during pruning phase $n = 4$, max generated directions to 8, max depth of reasoning to 6.752 **AIME** (Mathematical Association of America, 2024) is a highly prestigious and challenging mathe-
753 matics contest administered to top-performing participants of the AMC. It serves as a key benchmark
754 for evaluating the mathematical reasoning and problem-solving capabilities of large language mod-
755 els. 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.

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

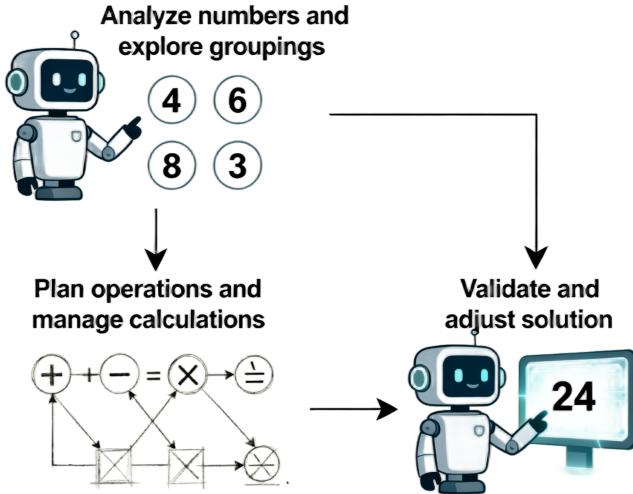


Figure 4: Visualizing the Phases of Game 24

B.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 6. We also include experimental results on the smaller model, Qwen2.5-7B and the proprietary GPT-4o-mini, in Table 7.

B.3 ADDITIONAL ABLATION STUDY

To evaluate our method, we conducted ablation studies on the AQUA, GSM8K, and Game of 24 datasets Table 5, 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.

Table 5: Additional ablation study

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

810 B.4 THE ANALYSIS OF THE GENERATED PHASES
811812 In this section, we evaluate the quality of the generated phases, illustrated with concrete examples
813 from Sudoku, Game of 24, and GSM8K. We also provide a visual representation of the phases for
814 Game of 24 in Figure 4.816 **Phases of Sudoku** ['Basic Elimination', 'Candidate Reasoning', 'Guessing and Backtracking'].
817818 **Phases of Game 24** ['Analyze numbers and explore groupings - Identify relationships and test
819 pairing possibilities', 'Plan operations and manage calculations - Determine sequence and ensure
820 mathematical viability', 'Validate and adjust solution - Verify results and refine approach to achieve
821 24'].823 **Phases of GSM8K** ['Calculate the total number of eggs consumed for breakfast and baking',
824 'Calculate the daily earnings from selling the remaining eggs', 'Final verification or solution step
825 and give the final answer'].826 While these phases are often broader and exhibit less consistency compared to those created by
827 human experts, they prove to be sufficiently accurate to effectively guide the reasoning process.828 Moreover, the system demonstrates a degree of fault tolerance: even coarse or imperfect phase
829 decompositions rarely lead to catastrophic failures. This robustness is achieved because the down-
830 stream Phase Classifier and Borda-based selection mechanism work in concert to steer the model
831 toward phase-relevant reasoning trajectories, effectively compensating for upstream imperfections.
832833 Table 6: Additional results with Qwen3-8B
834

835 Methods	836 Datasets					837 Average (%)
	838 AQUA (%)	839 BBEH (%)	840 GSM8K (%)	841 Game of 24 (%)	842 Sudoku Puzzle 5x5 (%)	
843 CoT	79.4	16.7	79.2	75.7	70.0	64.2
844 CToT	84.9	28.8	87.7	82.4	73.3	71.4
845 ARCA	90.6	47.1	93.5	91.1	90.0	82.4

846 Table 7: Additional results with Qwen2.5-7B and GPT-4o-mini
847

848 Models	849 Datasets				
	850 AQUA (%)	851 BBEH (%)	852 GSM8K (%)	853 Game of 24 (%)	854 Sudoku Puzzle 5x5 (%)
855 Qwen2.5-7B	84.7	40.2	89.4	90.2	80
856 GPT-4o-mini	89.8	45.6	94.7	88.3	90

857 C IMPLEMENTATION DETAILS

858 C.1 DETAILS OF THE SELECTION ALGORITHM

859 In this section, we provide a detailed explanation of how the Borda-aggregated direction selection
860 algorithm identifies the most preferred direction in Equation 4. To select the optimal reasoning path
861 via LLM-based preference comparisons, prior work has often adopted the dueling bandit frame-
862 work (Zhang et al., 2024). In this setting, when comparing two candidate thoughts i and j , candidate
863 i is chosen with probability $q(i, j)$, while candidate j is selected with the complementary probability
864 $q(j, i) = 1 - q(i, j)$. Here, $q(i, j) \geq \frac{1}{2}$ whenever i is ranked higher than j .865 However, dueling bandit algorithms such as DTS (Wu & Liu, 2016) typically rely on the Copeland
866 score (Zoghi et al., 2015) to aggregate comparison outcomes. A major limitation of the Copeland
867 score in LLM-based preference assessment is its sensitivity to minor preference variations (Goel
868 et al., 2017). This sensitivity arises from its win-counting mechanism, which can amplify stochastic
869 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:

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

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:

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

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.

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 C.4.

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 \bar{K}_t denote the pool size at iteration t until it reaches $\bar{K}_T \leq m$, giving total pruning cost $\sum_{t=1}^T \bar{K}_t n$. In the best case, the pool shrinks geometrically (e.g., removing a constant fraction each iteration), so $\sum_t \bar{K}_t = \mathcal{O}(K_0)$ 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 $\sum_t \bar{K}_t = \mathcal{O}(K_0^2)$, 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.

C.2 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. This blueprint directs subsequent operations along

918 logically coherent pathways and prioritizes high-value reasoning trajectories. As a result, the generator ensures comprehensive problem coverage, significantly improves computational efficiency, and
 919 overcomes the limitations of unstructured chain-of-thought reasoning.
 920

921 chain_prompt = '''You are an expert task decomposer. Your role is
 922 to analyze complex problems and break them down into essential
 923 high-level sub-phases. Each sub-phase should represent a
 924 critical milestone that moves toward solving the task.
 925

Algorithm 1 Borda-Aggregated Direction Selection

```

926
927
928
929 1: Input:
930 2: Pool: Set of  $K$  reasoning directions  $\{1, \dots, K\}$  to select
931 3:  $n$ : Number of comparisons per direction during pruning phase
932 4:  $m$ : Final pool size threshold
933 5: Preference( $u, v$ ): LLM comparison function

934 6: while  $K > m$  do ▷ Pruning Phase
935 7:    $E \leftarrow \emptyset$ 
936 8:   for each direction  $i \in \{1, \dots, K\}$  in the Pool do
937 9:     samples  $\leftarrow$  randomly select  $\min(n, K - 1)$  directions from  $\{1, \dots, K\} \setminus \{i\}$ 
938 10:    for each  $j \in$  samples do
939 11:       $E \leftarrow E \cup \{(i, j)\}$ 
940 12:    end for
941 13:  end for
942 14:   $\mathbf{W} \leftarrow \mathbf{0}^K, \mathbf{C} \leftarrow \mathbf{0}^K$  ▷ Reset counters
943 15:  for each  $(u, v) \in E$  do
944 16:    winner  $\leftarrow$  preference( $u, v$ )
945 17:     $\mathbf{C}[u] \leftarrow \mathbf{C}[u] + 1, \mathbf{C}[v] \leftarrow \mathbf{C}[v] + 1$ 
946 18:     $\mathbf{W}[\text{winner}] \leftarrow \mathbf{W}[\text{winner}] + 1$ 
947 19:  end for
948 20:  Borda  $\leftarrow [\mathbf{W}[i]/\mathbf{C}[i] \text{ for } i \in \{1, \dots, K\}]$ 
949 21:  NewPool  $\leftarrow \{i \mid \text{Borda}[i] \geq 0.5\}$ 
950 22:  Pool  $\leftarrow$  NewPool,  $K \leftarrow |\text{Pool}|$ 
951 23: end while

952 24:  $E_{\text{final}} \leftarrow \{(i, j) \mid i, j \in \text{Pool}, i \neq j\}$  ▷ Final Evaluation Phase
953 25:  $\mathbf{W} \leftarrow \mathbf{0}^K, \mathbf{C} \leftarrow \mathbf{0}^K$ 
954 26: for each  $(u, v) \in E_{\text{final}}$  do
955 27:   winner  $\leftarrow$  preference( $u, v$ )
956 28:    $\mathbf{C}[u] \leftarrow \mathbf{C}[u] + 1, \mathbf{C}[v] \leftarrow \mathbf{C}[v] + 1$ 
957 29:    $\mathbf{W}[\text{winner}] \leftarrow \mathbf{W}[\text{winner}] + 1$ 
958 30: end for
959 31: Borda  $\leftarrow [\mathbf{W}[i]/\mathbf{C}[i] \text{ for } i \in \{1, \dots, K\}]$ 
960 32: return  $\{j \mid \text{Borda}[j] = \max_{i \in \{1, \dots, K\}} \text{Borda}[i]\}$  ▷ Set of all max-scoring directions
961
962
963 **** Generate only the most essential sub-phases needed to
964 complete this task, excluding all implementation details and
965 optional steps.
966 **SUB-Phase DEFINITION:** Each sub-phase should specify WHAT needs to be accomplished, not
967 HOW to do it. Focus on the key objectives that must be
968 achieved.
969
970 **Example Demonstrations:**
```

971 **Geometry Problem:**

```

972     Task: "Find the area of a triangle with base 8cm and height 5cm
973     "
974
975     Key Sub-Phases:
976     1. Identify the area formula for triangles ,
977     2. Extract given dimensions from the problem ,
978     3. Compute the area using the formula
979     4. Final verification or solution step and give the final
980     answer
981
982     **Output Format Strictly Follow This Pattern:***
983     1. [Action-oriented sub-phase description] ,
984     2. [Next essential sub-phase] ,
985     ...
986
987     **Critical Reminders:***
988     - Phases should answer "what needs to be done" not "how to do it"
989     - Avoid transitional words ("then", "next", "after")
990     - Exclude mathematical symbols, formulas, or specific methods
991     - Maintain consistent verb tense and clarity
992     - Ensure sub-phases are truly sequential and complementary
993     ...
994
995     Phase Classifier The phase classifier enhances complex reasoning by dynamically identifying
996     the current phase in real time, enabling the solver to strategically allocate computational resources
997     toward phase-specific objectives. Once a phase concludes, the module seamlessly transitions to
998     another, directing the LLM's resources to the most relevant ongoing stage. By focusing efforts on
999     the active phase and reducing investments in completed or irrelevant directions, it maintains efficient
      and targeted progress throughout the reasoning process, thereby avoiding wasteful allocation.
1000
1001     classifier_prompt = f'''You are a Sub-phase Reasoning Engine. I
1002         will give the sub-phase list:{phases_list} and current
1003             thinking progress:{current_context}. Analyze the task progress
1004                 and determine:
1005     Which sub-phase should be actively worked on now.
1006     *****
1007     Choose the sub-phase from the list:{phases_list}, give me the
      number of index in the list.
1008
1009     **TASK ANALYSIS PROCESS:***
1010     1. Compare current progress with each sub-phase's requirements
1011     2. Identify the most immediate sub-phase that needs attention
1012     3. Verify the selection matches logical progression
1013
1014     **OUTPUT FORMAT STRICTLY FOLLOW:***
1015     [index]
1016
1017     **EXAMPLE:***
1018     Phases: ["Data collection", "Analysis", "Validation"]
1019     Current Context: "Finished gathering raw data, need to process it"
1020     You choose [Analysis]
1021     Output:
1022     1
1023
1024     **CRITICAL RULES:***
1025     - ***** Sub-phase MUST be from provided list and return in the
      number of index
     - No explanations or additional text
     - Sub-phase should logically follow from current_context

```

```

1026     ...
1027
1028
1029 Reasoning Direction Generator At each reasoning step, the reasoning direction generator takes
1030 the current phase and contextual state as input and produces a focused set of actionable, executable
1031 directions. These outputs provide timely and targeted guidance aligned with the specific objectives
1032 of the phase.
1033 direction_prompt = f'''You are a Phase-Oriented Direction
1034     Generator. Given the current step and the phase which needs to
1035     be achieved, Generate between {min_directions} and {
1036     max_directions} practical methods (directions) to achieve the
1037     specified phase.
1038 **Current step: {current_step}
1039 **Phase:** {current_phase}

1040 **Direction Definition:** Each direction should be a concrete, actionable method that:
1041 1. Directly contributes to achieving the phase according to
1042     current step
1043 2. Represents a distinct approach or technique
1044 3. Is executable without external knowledge
1045
1046 **Output Requirements:** - Generate between {min_directions} and {max_directions}
1047     directions
1048 - Each direction must start with an action verb
1049 - Format each direction as a bullet point ("- [direction
1050     description]")
1051 - Keep directions concise (5-15 words)
1052 - Exclude explanations or examples
1053
1054 **Quality Validation:** - Each direction is a distinct method (not a restatement)
1055 - Directions cover different aspects of the phase
1056 - Methods are practical and executable
1057 - Avoid overlapping or redundant directions
1058
1059 === COMPLETE EXAMPLES ===
1060 Example:
1061 Current step: "Already generate several passwords."
1062 Phase: "Validate password strength"
1063 Directions:
1064 - Check minimum password length
1065 - Verify mixed character types
1066 - Test against common passwords
1067
1068 === FORMAT REQUIREMENTS ===
1069 Output MUST be:
1070 - [Direction 1]
1071 - [Direction 2]
1072 ...
1073 - [Direction n] (where n is between {min_directions} and {
1074     max_directions})
1075 **Critical Rules:** - STRICTLY use the given format
1076 - NO numbering or other formats
1077 - NO additional text outside bullet points
1078 - Directions must answer "how to achieve the phase based on
1079     current step"

```

```

1080     ...
1081
1082
1083 Thought Generator We use the following prompt to generate the thought at each step, based on
1084 the task description, phase and optimal direction.
1085 purpose_prompt = f'''You are a heuristic assistant specialized in
1086     sub-phase-based problem solving.
1087
1088     **CURRENT SUB-PHASE:** {phase}
1089     **REQUIRED DIRECTION:** {direction[i]}
1090
1091     **TASK:** Generate exactly the next step that:
1092     1. Directly applies the specified direction: "{direction[i]}"
1093     2. Advances the current sub-phase: "{phase}"
1094     3. Reach the phase as fast as possible !!!
1095
1096     **OUTPUT FORMAT RULES:** 
1097     - The next step should reach the phase as fast as possible.
1098     - However, when the final step leads you to the final answer, give
1099         me only the numerical answer and print "###" before it,
1100         format as: ###[ANSWER]
1101     - Otherwise, provide a clear action step
1102     - No explanations, just the step itself
1103
1104     **VERIFICATION CHECKLIST:** 
1105     - Does this step directly follows the direction "{direction[i]}"? 
1106     - Does this step achieves the sub-phase "{phase}" as fast as
1107         possible?
1108     - If final answer, does it start with "###"?
1109     ...
1110
1111 LLM comparison function At each step, the generated direction are compared through Borda-
1112 aggregated direction selection framework using LLM comparison function, and the example of this
1113 comparison function is presented below.
1114
1115 preference_prompt='''As an analytical reasoning expert, critically
1116     evaluate which of the two reasoning paths demonstrates
1117     superior logical coherence, mathematical accuracy, and problem
1118     -solving effectiveness for the task. Consider: step-by-step
1119     validity, premise consistency, conclusion support, and error
1120     minimization. If both are objectively equal in all aspects,
1121     randomly select 1 or 2. Output must be exactly 1 or 2 with no
1122     additional text, explanations, or formatting.'''
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

```

C.3 THE REASONING AND SELECTION PROCESS

We present the reasoning steps and selection process employed by ARCA on AIME and BBEH tasks. The results indicate that the reasoning procedure, along with the Borda selection mechanism, aligns closely with typical human decision-making patterns.

The reasoning and selection process of AIME task:

```

Step 0 (Round 0)
Phase: Select the 0th option (Establish the core geometric model)
Available options:
"Let speed be v, height h. At 45°: distance = h, at 60°: distance
= h/√3. Distance covered in 10 min = h - h/√3 = h(1-1/√3)"
"Set up trigonometric equations: tan45° = h/x ⇒ h = x, tan60° = h/
y ⇒ h = y√3"

```

```

1134 "Complete solution:  $v \times 10 = h(1 - 1/\sqrt{3})$ , remaining distance  $h/\sqrt{3}$ ,  

1135 time =  $10/(\sqrt{3}-1) = 5(\sqrt{3}+1)$ "  

1136 "Use cotangent:  $h \cot 45^\circ = h$ ,  $h \cot 60^\circ = h/\sqrt{3}$ , find speed from 10  

1137 min travel"  

1138 Selection: Option 2 - Set up basic trigonometric equations  

1139  

1140 Step 1 (Round 1)  

1141 Phase: Relate the key variables from different scenarios  

1142 Available options:  

1143 "Substitute values:  $x = h$ ,  $y = h/\sqrt{3}$ , distance =  $h(1 - 1/\sqrt{3})$ "  

1144 "Calculate distance traveled:  $x - y = y\sqrt{3} - y = y(\sqrt{3} - 1)$ "  

1145 "Equate height expressions:  $x = y\sqrt{3}$ "  

1146 "Introduce speed:  $v \times 10 = x - y$ "  

1147 Selection: Option 3 - Equate the two height expressions  

1148  

1149 Step 2 (Round 2)  

1150 Phase: Execute the complete algebraic derivation to solve for the  

1151 target quantity  

1152 Available options:  

1153 "Calculate remaining time:  $t = y/v = y/[(x-y)/10] = 10y/(x-y) =$   

1154  $10/(\sqrt{3}-1) = 5(\sqrt{3}+1)$ "  

1155 "Express distance traveled in 10 min as  $x - y$ "  

1156 "Rationalize denominator:  $10/(\sqrt{3}-1) \times (\sqrt{3}+1)/(\sqrt{3}+1) = 5(\sqrt{3}+1)$ "  

1157 "From  $x = y\sqrt{3}$ , find  $y = x/\sqrt{3}$ , then compute time"  

1158 Selection: Option 1 - Calculate the final remaining time with  

1159 complete computation  

1160  

1161 Final Computation Process:  

1162 From  $x = y\sqrt{3}$ , distance in 10 min:  $x - y = y\sqrt{3} - y = y(\sqrt{3} - 1)$   

1163 Speed  $v = y(\sqrt{3} - 1)/10$   

1164 Remaining distance to base:  $y$   

1165 Time needed:  $t = y/v = y / [y(\sqrt{3} - 1)/10] = 10/(\sqrt{3} - 1)$   

1166 Rationalize:  $t = 10(\sqrt{3} + 1)/[(\sqrt{3} - 1)(\sqrt{3} + 1)] = 10(\sqrt{3} + 1)/2 =$   

1167  $5(\sqrt{3} + 1)$   

1168  

1169 Final Result  

1170 Algorithm answer: A  $[5(\sqrt{3} + 1)]$   

1171 Correct answer: A  

1172  

1173 The reasoning and selection process of BBEH task:  

1174  

1175 Step 0 (Round 0)  

1176 Phase: Establish the boundary framework for reasoning  

1177 Available options:  

1178 "Place the raven as the third bird from the left."  

1179 "Position robin immediately left of raven in the remaining slots (positions 2 and 3)."  

1180 "Determine the rightmost bird as quail."  

1181 "Identify the leftmost bird as owl: Position 1 = owl."  

1182 Selection: Option 4 - Identify the leftmost bird as owl: Position  

1183 1 = owl.  

1184 Step 1 (Round 1)  

1185 Phase: Establish the key anchor point and determine the fixed  

1186 point  

1187 Available options:

```

1188 Table 8: Results of test case, modeling LLM preference feedback with uncertainty and ties.
1189

1190 Algorithm	1191 One found	1192 Both found	1193 Comparisons	1194 Assessors
DTS(Copeland score) (Wu & Liu, 2016)	710	7	1000	4–20
Borda score method (Yan et al., 2022)	733	81	616–795	3–6

1195 "Assign positions: Position 1 = owl (leftmost), Position 5 = quail
1196 (rightmost)."
1197 "Position robin left of raven: since raven is third, place robin
1198 in position 2."
1199 "Given raven is third from left: Position 3 = raven. Robin is left
1200 of raven, so robin must be in position 2."
1201 "Place the raven as the third bird from the left: Position 3 =
1202 raven."
1203 Selection: Option 4 – Place the raven as the third bird from the
1204 left: Position 3 = raven.
1205 Step 2 (Round 2)
1206 Phase: Apply the remaining constraints to finalize the entire
1207 arrangement
1208 Available options:
1209 "Set quail at position 5 (rightmost)."
1210 "Place the falcon in the only remaining position (position 4)."
1211 "Position robin left of raven: since position 1 is owl and
1212 position 3 is raven, robin must be in position 2."
1213 "Confirm all constraints are satisfied with the complete
1214 arrangement."
1215 Selection: Option 1 – Set quail at position 5 (rightmost).
1216 Final Computation Process:
1217 Position 1: owl (given as leftmost)
1218 Position 3: raven (given as third from left)
1219 Position 2: robin (must be left of raven, and position 1 is
1220 occupied by owl)
1221 Position 5: quail (given as rightmost)
1222 Position 4: falcon (only remaining position)
1223 The rightmost bird is in position 5, which is the quail.
1224 Final Result:
1225 Algorithm answer: A (The quail is the rightmost)
1226 Correct answer: A

1227 **C.4 ANALYSIS OF BORDA SCORE IN LLM FEEDBACK**1228 To assess the ability of the Borda score to accommodate uncertainty, fine-grained distinctions, and
1229 potential ties commonly encountered in LLM preference feedback (Li et al., 2025b), we simulate a
1230 test scenario based on the setup and results from (Yan et al., 2022). The test case represents a sce-
1231 nario with no single winner and many ties, mirroring the challenges of LLM preference judgments.1232 **Test Case:**

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

1233 We adopted experimental parameters and evaluation criteria of (Yan et al., 2022). The experiment
1234 involved a large set of 100 options, and a fixed budget of 1000 comparisons. As shown in Table. 8,
1235 the simulation results demonstrate the superior performance of the Borda score method over the
1236 Copeland-based approach in identifying optimal outputs from LLM preference feedback.
1237