

# interwhen: A Generalizable Framework for Verifiable Reasoning with Test-time Monitors

Vishak K Bhat<sup>1\*</sup> Prateek Chanda<sup>2\*</sup> Ashmit Khandelwal<sup>1</sup> Maitreyi Swaroop<sup>3</sup>  
Vineeth N. Balasubramanian<sup>1</sup> Subbarao Kambhampati<sup>4</sup> Nagarajan Natarajan<sup>1†</sup>  
Amit Sharma<sup>1†</sup>

<sup>1</sup>Microsoft Research, India; <sup>2</sup>IIT Bombay

<sup>3</sup>Carnegie Mellon University; <sup>4</sup>Arizona State University

## Abstract

We present a test-time verification framework, *interwhen*, that ensures that the output of a reasoning model is valid wrt. a given set of verifiers. Verified reasoning is an important goal in high-stakes scenarios such as deploying agents in the physical world or in domains such as law and finance. However, current techniques either rely on the generate-test paradigm that verifies only after the final answer is produced, or verify partial output through a step-extraction paradigm where the task execution is externally broken down into structured steps. The former is inefficient while the latter artificially restricts a model’s problem solving strategies. Instead, we propose to verify a model’s reasoning trace as-is, taking full advantage of a model’s reasoning capabilities while verifying and steering the model’s output only when needed. The key idea is *meta-prompting*, identifying the verifiable properties that any partial solution should satisfy and then prompting the model to follow a custom format in its trace such that partial outputs can be easily parsed and checked. We consider both self-verification and external verification and find that *interwhen* provides a useful abstraction to provide feedback and steer reasoning models in each case. Using self-verification, *interwhen* obtains state-of-the-art results on early stopping reasoning models, without any loss in accuracy. Using external verifiers, *interwhen* obtains 10 p.p. improvement in accuracy over test-time scaling methods, while ensuring 100% soundness and being 4x more efficient. The code for *interwhen* is available at <https://github.com/microsoft/interwhen>.

## 1 Introduction

The current arc of AI progress has led to significant increases in the accuracy of language models (LMs) on reasoning tasks. However, in many high-stakes real-world scenarios such as law, healthcare, or acting in the physical world, it is critical that each answer by an LM is correct and trustable; merely having a high probability of producing the correct answer is not sufficient. Thus, *verification* of a LM’s answer becomes an important goal. Past work has explored self-verification through methods such as self-consistency [Wang et al., 2023], and external verification through existing domain-specific verifiers [Cobbe et al., 2021b]. In a typical implementation, multiple answer trajectories are generated independently and the ones that pass the verifier are chosen [Kambhampati et al., 2024].

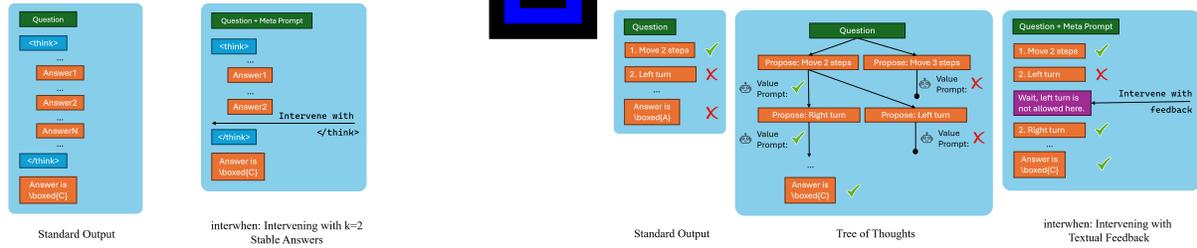
These techniques, however, focus on verifying the final answer. Considering the recent thrust on test-time scaling wherein state-of-the-art models output thousands of intermediate reasoning tokens and/or tens of tool calls before the final answer, this approach becomes inefficient, especially since it fails to catch early errors. Moreover, problems vary by the level at which they can be verified. Some solutions such as code or a plan are best verified by the final answer, while other problems such as writing a report are best verified through the steps taken to reach an output. Verifying a long report invokes considerable cognitive load (and hence the constituent steps may be more feasible to verify for both humans and automated verifiers).

---

\*Equal contribution.

†Equal advising. Correspondence to: [nagarajn@microsoft.com](mailto:nagarajn@microsoft.com), [amshar@microsoft.com](mailto:amshar@microsoft.com).

Please answer the following question based on the provided information. How many right turns are there in the provided path (blue) from the start (green) to the end (red)? Available options: A. 4 B. 8 C. 2 D. 7.



(a) Internal Verification

(b) External Verification

Figure 1: LLM-Process Modulo: Intervening on partial output traces to guide a model towards a more accurate and/or more efficient solution. Meta-prompting instructs the model to output *intermediate states* that can be extracted and verified. In the first case, we track mentions of candidate solutions within the `<think>` block and insert `</think>` when the same answer appears consecutively. In the second case, we run an external verifier on each intermediate state and append verifier feedback to the model’s output trace in case of any error. Compared to existing methods such as tree-of-thoughts, we do not artificially break up the problem into step-wise generation, and as a result, obtain both higher accuracy and efficiency.

Therefore, we propose a new test-time paradigm of generating verified answers that focuses on verification (and correction) of the *steps* towards a solution, not just the final answer. Existing approaches that verify partial steps and choose the valid ones at each step [Yao et al., 2023] break down the task execution externally into steps, which artificially limits a model’s problem solving strategies. Instead, we adopt the following desiderata to effectively make use of the advances in problem-solving abilities of reasoning models: **1)** Problem solving should not be externally broken down into steps; instead the model retains full control over problem execution. **2)** Rather than choosing between multiple independent step generations (which can be inefficient), one should provide step-wise feedback to the model at test-time and steer the model’s reasoning through verifier feedback. Importantly, if there is no error, the model’s output text should be identical to the one without any test-time verification. We call this paradigm *LLM-Process-Modulo* (see Figure 1), to distinguish it from the LLM-Modulo paradigm [Kambhampati et al., 2024], which we call *LLM-Solution-Modulo*.

This gives rise to two research questions: 1) given a stream of ambiguous text as LM output, how to identify the *verifiable* steps and verify them; and 2) given a verifier’s feedback, how to steer a LM’s output to fix the error highlighted by the verifier. Our core conceptual contribution is *meta-prompting*: the process of identifying verifiable steps for a task and prompting the LM to generate output such that those steps are easily extractable. This side-steps the error-prone task of building parsers to extract and formalize steps [Wu et al., 2022], and enables a general method that can work for any domain, from code to text reports, as long as verifiable states can be defined. Formally, meta-prompting involves identifying properties (states)  $\mathcal{P}$  that any valid partial reasoning trace should satisfy to be able to solve a problem. In many cases, identifying the verifiable states is trivial. For procedural reasoning problems where each step follows from the past context, e.g., math or deductive reasoning problems, each step is a verifiable state by definition. Other problems, such as finding a plan or object satisfying certain constraints may require inductive reasoning or working with inconsistent values (e.g., proof by contradiction) and hence only specific intermediate states may be verified.

Once verifiable states are identified, we propose a general framework, *interwhen*, that enables steering of a model’s output through feedback. Given a set of intermediate states and their corresponding verifiers, *interwhen* abstracts out the steering problem through three functions: extract state, verify state, and update output trace with feedback. As long as the model outputs the requested intermediate states, it is free to employ any problem solving strategy. While the *interwhen* framework is general enough to implement any test-time verification method, we present a specific method that provides a good tradeoff between accuracy and efficiency. In this method, the system maintains a single execution run and verifiers are called whenever the LM outputs an intermediate state. The feedback of the verifier is added inline to the LM’s output as either a “#Thought” or as a reflection starting with “Wait”, and the token generation is continued.

Keeping in mind practical scenarios where objective verifiers may not exist, we consider a broad view of

verification and include both LLM-based (self-verification) and external verifiers. In both cases, intervening on a partial output trace leads to significant improvements in accuracy or efficiency (and sometimes both). When using external verifiers, a key benefit of *interwhen* is that *soundness* with respect to verifiers is guaranteed by design: if an output is shown, it is guaranteed to be consistent with the verifiers.

We apply *interwhen* framework to two scenarios: one using self-verification, and the other using symbolic external verifiers. For the first, we address the problem of *early stopping* a model’s reasoning to improve efficiency without sacrificing accuracy. For early stop, we consider a *candidate final solution* as an intermediate state, and call a verifier (equivalence checker) each time the model outputs a candidate final solution. The verifier checks if the last  $k$  solutions are the same, and if yes, stops the model early. With external verifiers, we consider the problem of improving accuracy through test-time scaling. Intermediate verifiable states are decided based on dataset semantics: for a puzzle such as Maze, we consider each move as a state; for the SpatialEval dataset, we consider each conclusion about objects’ position as a state; and for Game of 24, we consider each partial subset of numbers as a state. In all three datasets, we find that *interwhen* can steer a model as required. For the early stop scenario, *interwhen* provides up to 30% token reduction without decreasing accuracy. For the test-time scaling scenario, *interwhen* provides up to 10 p.p. accuracy improvement over the baselines such as Best-of-4 and Tree-of-Thought [Yao et al., 2023], while being 100% sound and at least 4x more efficient.

## 2 Related Work

### 2.1 Verification-guided Inference

Two primary ways to improve LLM output at test-time are: using verification and using other test-time scaling techniques.

**Verification of LLM output.** In math, planning and other domains, a dominant paradigm for test-time verification is to generate candidate solutions and then select the best one, either using external verifiers [Cobbe et al., 2021b, Kambhampati et al., 2024] or the same LLM for feedback [Madaan et al., 2023]. Efforts to verify and steer based on partial steps utilize an externally imposed structure to break down a problem into steps, e.g., asking the language model to make one deduction in a multi-step deduction task or one step in a multi-step math problem [Creswell et al., 2022, Yao et al., 2023]. Other works employ self-verification with a LLM prompt to verify and refine the answer [Chang et al., 2025, Stechly et al.]. Compared to these efforts, we do partial step verification without any externally imposed structure of steps that control execution. Instead, we have a lightweight prompt guidance for the model to output intermediate states. There is a large literature on process reward models [Lightman et al., 2024] that is orthogonal to this work. Existing PRMs can be used as verifiers in the *interwhen* framework; however, the novel question we explore is how to use such feedback to steer a reasoning model at test time. There is early work on steering reasoning models at test time through text feedback [Prabhakar et al., 2025], however the connection to partial output verification is not explicit.

**Test-time scaling.** Test-time scaling is the emergent paradigm for improving reasoning in LMs, by adjusting compute at inference time based on the complexity of the task [Snell et al., 2025, Balachandran et al., 2025, Kang et al., 2025]. A popular technique often employed as a strong baseline is best-of- $N$  sampling, that involves drawing  $N$  trajectories *in parallel* and picking the best final answer as adjudged by an external verifier [Kang et al., 2025] (e.g., a judge LM) or by simple heuristics (e.g., majority) in permissible scenarios. A variant of this strategy is to draw the samples *sequentially* using the feedback on a previous answer in the *LLM-Modulo* paradigm [Kambhampati et al., 2024].

Among techniques that analyze intermediate traces, some techniques encourage exploration [Muennighoff et al., 2025] using heuristics (e.g., inserting “Wait” to continue thinking) to be able to solve complex problems by giving additional compute. Another class of test-time scaling methods rely on classical tree-based search techniques, and use LMs to guide the exploration, e.g., Tree of Thoughts [Yao et al., 2023] and MCTS-inspired search [Dainese et al., 2024, Wang and Yang, 2025, Huang et al., 2025a, Antoniadou et al., 2025]. In this work, we propose a new axis of test-time scaling that uses a single rollout and steers the reasoning trajectory through verifiers.

## 2.2 Early Stopping

Test-time scaling can induce *overthinking*, where models continue generating reasoning beyond the point of meaningful improvement [Ghosal et al., 2025]. Since then, several strategies have been proposed to mitigate this behavior.

Early work in this area focused on using confidence signals to decide when to halt generation. For example, entropy-based frameworks such as [Sharma and Chopra, 2025] use sequence-level Shannon entropy from token log-probabilities as a proxy for emergent confidence. However, these approaches require calibrating stopping thresholds on benchmark datasets. Building on entropy heuristics, methods such as EAT [Wang et al., 2025] compute entropy after `</think>` to decide on early exit. DEER [Yang et al., 2025] similarly introduces dynamic early stopping for chain-of-thought models by monitoring confidence at intermediate reasoning chunks. Both methods depend on manually specified confidence thresholds, which are hard to interpret and set for each task.

Subsequent work explored richer signals and control mechanisms. [Huang et al., 2025b] suppresses reflection-trigger tokens (e.g., “Wait” and “Alternatively”) when the model exhibits high confidence, thereby reducing redundant reasoning without retraining. [Sun et al., 2025] frames early stopping as a control problem through a two-stage discriminator and an adaptive multi-armed bandit controller that dynamically determines when further reasoning is unnecessary. Parallel to these efforts, Dynasaur [Fu et al.] characterizes *answer stabilization* where they perform multiple rollouts after a particular segment of reasoning and early exit based on the stability of the rollout answers. Complementary approaches such as [Choi et al., 2025] analyze attention patterns to a special end-of-thinking token and prune low-contributing reasoning chunks before resuming generation. Another similar work *ES-CoT* [Mao et al., 2025], samples the answer after each chunk and checks the stability of those answers.

These methods typically rely on calibrated thresholds, auxiliary confidence estimators, heuristic trigger tokens, or specialized control mechanisms, raising questions about robustness across models and domains. In contrast, we propose a simple method based on self-verification that extracts intermediate answers from the reasoning trace and appends `</think>` once the answer has remained consistent for  $K$  consecutive checks, enabling adaptive termination without dataset-specific calibration.

## 3 LLM-Process-Modulo Paradigm

As stated in the Introduction, majority of techniques for improving reasoning have two pivots: (1) scale test-time compute, (2) focus on accuracy of the solution, which often means the correctness of the final answer. Current techniques to improve reasoning (discussed in Section 2) can be categorized broadly into two classes, based on if and how they steer the reasoning trajectories, or more broadly, the exploration of the solution space.

(1) The standard best-of-N sampling methods do not steer the reasoning process. These techniques are final-answer-focused and inefficient, often requiring tens of thousands of tokens (and several tool calls) even for simple problems.

(2) Tree-based search techniques such as beam search leverage LMs to *guide* the search process. These techniques, in contrast to (1), explicitly decompose the solution generation into intermediate steps and then use LMs to pick the best action from several permissible actions at each node in the tree. While they intervene the exploration of the solution and are not purely final-answer focused, they can be restrictive—they are confined by the heuristics and choices of the underlying search algorithm, and often require hand-crafting problem-specific prompts to guide the search (e.g., the value prompt in beam search [Yao et al., 2023]).

To address the main criticisms of the above two classes of techniques, we propose a new direction of *intervening* the partial traces and *steering* them, that we refer to as “LLM-Process-Modulo”. The core idea is simple: we interject the reasoning process as it unfolds, using either external verification mechanisms (when available) or self-verification (otherwise), so that we can guarantee some form of validity or correctness of the reasoning *process*. Our proposal is analogous to the LLM-Modulo framework [Kambhampati et al., 2024] that follows an iterative generate-test paradigm on the final answer, but the *modulo* (i.e., verification) happens during the process of generating the solution.

Realizing this idea poses multiple research challenges: (1) **when to intervene**: reasoning models output a stream of tokens that are difficult to verify, so we need a way to identify the verifiable “states” in the

reasoning process; (2) **how to verify**: we might need formal verifiers in some scenarios, or LLM-based verifiers in others. The choice depends on the domain and the task at hand. (3) **how to intervene**: how do we intervene or *steer* the reasoning model if an intermediate state is rejected by the verifier.

### 3.1 When to intervene

It is often challenging to make sense of the stream of “thinking” tokens from a reasoning LM, let alone verify. We need a systematic way to elicit the verifiable parts of the reasoning process. The first challenge is how to elicit *process* verification in LMs. In domains like math, we can think of each derivation or formula as a “step”, and verifying the process amounts to verifying the correctness of the individual steps given the pre-condition(s). In such cases, it might suffice to parse the steps in the trace and verify them.

But, for general reasoning problems that involve logic and planning, it might be hard to *identify* the verifiable intermediate steps. This is because in many cases the reasoning process can be highly convoluted, might involve making hypotheses, back-tracking when there is a contradiction, etc. So, we need a better way to define what constitutes a verifiable state, and elicit them in the reasoning trace. For example, Figure 1 shows an example problem on following a path in a 2D maze. The verifiable states are the moves made in this maze, but they may be interspersed with general text reasoning and thus hard to parse. In the next section, we show how we achieve this elicitation via meta-prompting. These verifiable states then become the intervention points.

### 3.2 How to verify

Our goal is to verify *partial* output traces. In domains such as planning, the same final answer verifier can be used to validate the steps in a partial trace. In others, we may need to develop new verifiers. In many reasoning settings, verifiers can be easier to implement than solving the problem (e.g. a rule-based checker of validity of game plays). For instance, for the Maze problem in Figure 1, we implement a simple verifier based on representing the maze as a 2D array in Python and replaying operations on it. In planning, logic, and code, we could also leverage formal verification tools and extend them to adapt to incomplete information in the partial trace. Finally, we may utilize LLM-based verifiers for general reasoning problems (e.g., document generation, counter-factual analysis).

### 3.3 How to intervene

When an intervention point is encountered (as in the above discussion), we need to determine how to intervene the reasoning process. Many options present: 1) simply roll back the solution generated thus far, 2) provide a text feedback from the verifier (or more generally a tool or even a human expert), and resume the reasoning process, 3) provide a binary correct/incorrect feedback from the verifier, and so on. The right choice of intervention could depend on the ability of the reasoning model, the fidelity of the verifier, and the complexity of the task at hand. For instance, for the maze problem in Figure 1, we intervene by adding the `< /think >` token for early stopping the model’s reasoning, and intervene by adding the verifier’s feedback after the current state for the test-time scaling setting.

While we cannot answer all questions in this paper, we provide a generalizable framework that allows investigation of all these questions in the next Section. We also propose a candidate method that provides significant improvements over the current techniques.

## 4 interwhen: Generalizable Framework

Many techniques for test-time verification and scaling have been proposed, however, it is difficult to evaluate and compare them holistically for a given task, across accuracy and efficiency tradeoffs. Therefore, we build a general framework for test-time verification called *interwhen* that allows exploring the research questions described above, while also providing an easy way to implement and compare existing test-time scaling strategies.

**Meta-prompting** While the standard prompt focuses on task description, meta-prompting focuses on ensuring a structure in the output trace such that intermediate states are clearly demarcated. Formally, meta-prompting structures an output trace from unstructured text to a semi-structured text where intermediate states are interspersed with arbitrary text: “...,  $s_i$ , ...,  $s_{i+1}$ , ...” where ... refers to free text. We implement it by providing few-shot examples of the format needed for expressing a state. For instance, for the 2D Maze problem in Figure 1, the metaprompt simply instructs to output each move in a specific format.

**Abstractions for test-time verification** The three research questions above motivate three fundamental operations to build any test-time verification method. 1) identifying and extracting verifiable steps from one or multiple LM outputs; 2) verifying the candidate steps; 3) giving feedback to improve a candidate step or choosing one of the candidate steps.

Given a task, the first step is to identify relevant verifiers and determine the intermediate “states” on which they will apply. Each verifier is a function that takes the current state, all previous states, and problem context, and returns a pass decision and optionally some text feedback on the pass evaluation. Let  $S_0$  denote the initial state which consists of problem description and any relevant context, and  $S_i$  denote the  $i$ th state in the reasoning trajectory. Given a list of verifiers, test-time verification requires the following three operations.

**extract\_state.** Given a verifier  $V$  and a partial reasoning trace  $r$ , extract the relevant intermediate state:  $S_i = \text{extract\_state}(r, V)$ . For procedural solutions where the LM lists solution steps in sequence (e.g., for simple grade-school math problems [Cobbe et al., 2021a]), using newline or other token-level separators suffices to extract states.

**verify.** Given the extracted intermediate state and context  $c$  (e.g., past states), perform verification using the verifier  $V$ :  $y_i, t_i = \text{verify}(S_i, c, V)$ .

**intervene.** If the verifier fails and provides text feedback, process the feedback and update the partial reasoning trace,  $r' = \text{intervene}(r, y_i, t_i)$ , otherwise keep the same trace  $r' = r$ . The intervention may ask the model to continue generation or decide to stop the generation.

Different test-time scaling algorithms can be implemented using the above operations. For example, best-of-N is a specific instance where the only verifiable state is the final answer and the intervention is to stop the generation, which is then repeated over a loop. Tree of thoughts [Yao et al., 2023] can be implemented by enforcing a beam search over the extract-verify-intervene process on each step, where the intervene step is simply the identity operation for the top-ranking steps as per the verifier, while generation is stopped for all others.

## 4.1 Proposed Algorithm: Sequential Verifier

We now describe a simple verification algorithm that works well in practice. Rather than exploring multiple trajectories, the algorithm maintains a single output trace. For each identified intermediate state, it applies the verifier and allows generation to continue if the verifier passes. If not, it adds the text feedback to the model’s output trace and then continues generation. The intuition is that the model should change its solution trajectory based on the text feedback on its previous state.

A key question is how to identify intermediate states. We leverage the metaprompting strategy discussed in Section 4, and instruct the model to output intermediate states in a specific format. Mentioning few-shot examples describing the format is typically sufficient (see Appendix A.2 for the different prompts used).

*Remark 4.1 (Merits of Sequential (State) Verifier).* First, we can guarantee soundness of reasoning, i.e., intermediate “steps” of the output solution are correct w.r.t. to a given set of verifiers. Second, we directly address the key pain-points of existing techniques for improving reasoning, namely: (1) restrictiveness of the classical search heuristics *by largely retaining the agency of search in the reasoning models*, (2) need for a multitude of problem  $\times$  search strategy prompts *via leveraging a simple meta-prompting strategy*. Empirically, we find that this also leads to substantial reduction of token inefficiency compared to test-time scaling techniques.

We use two case-studies next to illustrate the implementation of the algorithm: early stopping and test-time scaling. For dataset details, see Section 5.1.

## 4.2 Case study 1: Internal Verification–Early Stopping

Our first case-study involves the problem of early stopping. Test-time scaling often induces overthinking, where models continue reasoning beyond the point of meaningful improvement [Ghosal et al., 2025]. To mitigate this, several **early stopping** strategies have been proposed. Methods like EAT rely on entropy-based criteria computed after the `</think>` token [Wang et al., 2025], while others like DEER estimate confidence scores at intermediate reasoning chunks to decide when to halt inference [Yang et al., 2025]. A key challenge in these methods is to determine the correct threshold for entropy or confidence, such that accuracy loss is minimal. We use `interwhen` to propose a simple and effective algorithm, ***k-Stable Answer***, which terminates reasoning once the predicted answer remains unchanged for  $k$  consecutive reasoning steps. We find that that setting  $k = 2$  works well across datasets.

**Sequential Verifier Implementation** We instantiate the Sequential Verifier by specializing the `extract_state`, `verify`, and `intervene` operations for each task family.

For multiple-choice reasoning tasks (MAZE and SPATIALMAP), `extract_state` uses regular-expression patterns to detect answer proposals signaled by phrases such as “the answer is” and related variants. The `verify` function then checks whether the extracted option remains unchanged for  $k$  consecutive times.

For the GAMEOF24 dataset, `extract_state` identifies the equation proposed by the model, subject to the constraint that each input number appears exactly once in the extracted expression. `verify` tests whether the same equation appears  $k$  consecutive times.

For the VERINA code-generation and specification-generation tasks, `extract_state` first isolates the portions of the response that indicate the presence of code or specifications using regular-expression patterns (e.g., phrases such as “the code is” and related variants). Because superficial string matching is insufficient in this setting, we implement `verify` using a small language model (Qwen/Qwen2-1.5B-Instruct) to determine whether the current artifact extracted is semantically equivalent to the previous one. For specification generation, we separately track the stability of the pre-conditions and post-conditions by maintaining two independent counters, and trigger termination only when both have remained unchanged for  $k$  consecutive times.

In all the three datasets, once the stability criterion is met in `verify`, `intervene` terminates reasoning by injecting a `</think>` token and prompting the model to output its final answer.

## 4.3 Case study 2: External Verification

The second application concerns test-time scaling with the goal of improving accuracy on a task, while ensuring soundness if the system outputs an answer, it is correct.

**Sequential Verifier Implementation** We modify the base prompt to explicitly specify the format in which the model should present its output along with passing few shot examples (**Metaprompting**, Section 4). For the MAZE dataset, the prompt encourages step-by-step generation, beginning with identification of the start and end locations, followed by a sequence of moves in which each step reports the proposed action, the resulting position, the cumulative step count, and the direction taken. A concrete instantiation of this prompt is provided in the Appendix A.2.

For the SPATIALMAP dataset, we similarly enforce a structured reasoning template. The model is instructed to first parse and enumerate all spatial relations described in the question, then explicitly identify the queried pair and determine whether a direct relationship is present, and finally produce both a natural-language conclusion and a boxed multiple-choice answer. Concretely, the prompt organizes generation into three stages: (i) parsing relationships, (ii) finding the direct relationship relevant to the query, and (iii) emitting the final answer in a standardized format. An example of this template is included in the Appendix A.2.

`extract_state`. As described above, we meta-prompt the model to produce reasoning in a highly structured format, which we then isolate using regular-expression patterns.

`verify`. For the MAZE dataset, we first parse the input maze into a grid representation encoding wall locations as well as the designated start and end positions. We define deterministic transition rules for left and right turns (e.g., moving from facing south to east constitutes a right turn). When the model proposes a step, we simulate the resulting move and check whether it violates any constraints such as entering a wall cell

Dataset	Method	Accuracy %	Tokens %
MAZE	EAT	88.53	100.00
	DEER	88.53	99.39
	interwhen (k-Stable)	88.53	<b>67.76</b>
	baseline	88.53	100.00
SPATIALMAP	EAT	74.93	99.66
	DEER	75.00	<b>93.58</b>
	interwhen (k-Stable)	74.93	95.31
	baseline	74.93	100.00
GAMEOF24	EAT	95.01	100.00
	DEER	95.01	96.18
	interwhen (k-Stable)	95.45	<b>68.35</b>
	baseline	95.01	100.00

Table 1: **Internal verification** on Qwen/Qwen3-30B-A3B-Thinking-2507. Tokens are reported relative to standard CoT (100%).

Dataset	Method	Accuracy %	Tokens %	
MAZE	Best-of-2	85.70	239.41	
	Best-of-4	89.23	481.29	
	Majority-4	84.72	481.29	
	Generate-Test	89.23	447.09	
	ToT – Beam 2 (paper)	90.26	652.77	
	ToT – SV critic	92.51	746.20	
	interwhen (step verifier)	<u>98.74</u>	126.46	
	SPATIALMAP	Best-of-2	75.18	215.07
		Best-of-4	82.71	436.11
Majority-4		73.51	436.11	
Generate-Test		82.71	397.21	
ToT – Beam 2 (paper)		75.62	630.85	
ToT – SV critic		74.22	672.47	
interwhen (step verifier)		<u>86.13</u>	122.55	
GAMEOF24		Best-of-2	96.14	226.42
		Best-of-4	96.82	583.17
	Majority-4	95.06	583.17	
	Generate-Test	96.82	427.82	
	ToT – Beam 2 (paper)	95.02	814.61	
	ToT – SV critic	<b>97.73</b>	936.53	
	interwhen (step verifier)	<u>96.99</u>	142.37	

Table 2: **External verifier** on Qwen/Qwen3-30B-A3B-Thinking-2507. Tokens are reported relative to standard CoT (100%). On MAZE and SPATIALMAP datasets, the step-verifier methods leads to significant gains. Moreover, we mark its accuracy as underlined since it is 100% sound—if it outputs an answer, it is always correct, unlike other methods.

and whether the updated orientation is consistent with the turning rules. A proposal is marked invalid if either the movement or the implied direction is incorrect.

For SPATIALMAP, we encode the initial spatial relations between objects as logical constraints in a Z3 solver. At each step, the model’s structured output is converted into a JSON representation by a small language model, which is then checked for logical consistency against the constraint set using Z3. Any contradiction indicates that the newly proposed relation cannot hold given the existing map. For GAMEOF24, we verify by checking that the expression evaluates to 24 and that all 4 digits are used.

intervene. If a proposed step fails verification, we append the corrective textual feedback at the current state and provide it to the LLM to continue generation, explicitly indicating that the step is inconsistent and instructing the model to resume reasoning from the previous valid state. If no violation is detected, generation proceeds without intervention.

## 5 Experiments

### 5.1 Datasets & Models

Our main results are based on the following datasets and the Qwen-30B-A3B-thinking model: spatial reasoning problems, including MAZE and SPATIALMAP (1,500 instances each) [Wang et al., 2024]; arithmetic reasoning via GAMEOF24 (1,362 instances).

In MAZE, each example provides a grid with start and end points; questions involve reasoning over the path, such as counting right turns, total turns, or the start’s position relative to the end.

SPATIALMAP examples specify object locations; questions ask for an object’s relative position, the object in a given direction, or the number of items along a direction.

In GAMEOF24, each example provides four integers, and the task is to construct an expression using each number exactly once with  $+$ ,  $-$ ,  $\times$ ,  $/$  to reach 24. In addition, we show ablations using a code dataset, VERINA [Ye et al., 2025], using QwQ-30B.

Dataset	Method	Accuracy %	Tokens %
MAZE	EAT	88.53	100.00
	DEER	88.53	99.39
	interwhen (k-Stable)	88.53	<b>67.76</b>
	baseline	88.53	100.00
SPATIALMAP	EAT	74.93	99.66
	DEER	75.00	<b>93.58</b>
	interwhen (k-Stable)	74.93	95.31
	baseline	74.93	100.00
GAMEOF24	EAT	95.01	100.00
	DEER	95.01	96.18
	interwhen (k-Stable)	95.45	<b>68.35</b>
	baseline	95.01	100.00

Table 3: **Internal verification** on Qwen/Qwen3-30B-A3B-Thinking-2507. Tokens are reported relative to standard CoT (100%).

## 5.2 Internal Verification

**Baselines.** We compare *k-Stable* answers against two recent early-stopping strategies for reasoning models.

**EAT** (Entropy After `</think>`) Wang et al. [2025] append a `</think>` token after each reasoning chunk and computes the entropy of the next-token distribution. An exponential moving average over steps is maintained, and reasoning is terminated once this value drops below a fixed threshold, after which the model is prompted to produce the final answer.

**DEER** (Dynamic Early Exit in Reasoning Models) Yang et al. [2025] estimate answer confidence after each reasoning step by prompting the model with “`</think>` The answer is.” If the confidence exceeds a predefined threshold, the chain-of-thought is halted and the model directly generates the final output.

**Metrics.** We report task accuracy together with the percentage of reasoning tokens used. The number of tokens generated under standard chain-of-thought decoding is treated as 100%, and all early-stopping methods are measured relative to this baseline.

**Models and Decoding.** We conduct experiments on three open-source reasoning models: Qwen/Qwen3-30B-A3B-Thinking-2507, Qwen/QwQ-32B, and microsoft/phi-4-reasoning.

All experiments are run using vLLM. For the Qwen models, we use a maximum generation length of 32,768 tokens with temperature 0.6, top-*p* 0.95, and top-*k* 20. For Phi-4, we use a maximum length of 16,384 tokens, temperature 0.8, top-*p* 0.95, and top-*k* 50. We report results for the hyperparameter configuration that achieves the maximum reduction in token usage while preserving accuracy. Details of the hyperparameter settings and the complete results are provided in Appendix A.1 and Tables 8–10.

### 5.2.1 Main Results

Table 3 presents our primary comparison on Qwen3-30B across MAZE, SPATIALMAP, and GAMEOF24. Across all three benchmarks, *k-Stable* answers achieves substantially larger token reductions than existing early-stopping methods while maintaining or improving accuracy.

On MAZE, *k-Stable* reduces reasoning tokens by **32.2%** without any loss in accuracy, compared to 0.6% for DEER and negligible savings for EAT. On SPATIALMAP, it attains a **10.8%** reduction, exceeding DEER (8.2%) and EAT (0.4%). For GAMEOF24, *k-Stable* achieves a **31.7%** reduction, whereas DEER and EAT yield only 11.7% and 3.5%, respectively.

These results demonstrate that stability-based stopping can substantially shorten chains of thought while preserving final-task performance on a strong reasoning model. The main tables and ablations report only the operating point that achieves the *largest token reduction while maintaining or improving accuracy relative to the baseline*. Complete sweep results for every model and dataset are provided in Appendix tables 8–10.

Dataset	Method	Accuracy %	Tokens %
MAZE	Best-of-2	85.70	239.41
	Best-of-4	89.23	481.29
	Majority-4	84.72	481.29
	Generate-Test	89.23	447.09
	ToT – Beam 2 (paper)	90.26	652.77
	ToT – SV critic	92.51	746.20
	interwhen (step verifier)	<u>98.74</u>	126.46
SPATIALMAP	Best-of-2	75.18	215.07
	Best-of-4	82.71	436.11
	Majority-4	73.51	436.11
	Generate-Test	82.71	397.21
	ToT – Beam 2 (paper)	75.62	630.85
	ToT – SV critic	74.22	672.47
	interwhen (step verifier)	<u>86.13</u>	122.55
GAMEOF24	Best-of-2	96.14	226.42
	Best-of-4	96.82	583.17
	Majority-4	95.06	583.17
	Generate-Test	96.82	427.82
	ToT – Beam 2 (paper)	95.02	814.61
	ToT – SV critic	<b>97.73</b>	936.53
	interwhen (step verifier)	<u>96.99</u>	142.37

Table 4: **External verifier** on Qwen/Qwen3-30B-A3B-Thinking-2507. Tokens are reported relative to standard CoT (100%). On MAZE and SPATIALMAP datasets, the step-verifier methods leads to significant gains. Moreover, we mark its accuracy as underlined since it is 100% sound—if it outputs an answer, it is always correct, unlike other methods.

### 5.3 External Verification

**Baselines.** We evaluate against the following baselines:

**Tree of Thoughts (ToT)** We include the Tree-of-Thoughts baseline of Yao et al. [2023]. We use a proposer prompt with few-shot exemplars to generate candidate successor states. Following the original formulation, we apply a value prompt to score each intermediate state with a categorical judgment (*likely*, *unlikely*, or *impossible*), and use these scores to guide the search. We run beam search with beam width 2, resulting in up to 4 active branches, and cap the search depth at 8 across datasets.

**Tree of Thoughts (ToT) with LLM-based SV** This is an interwhen-based verification implementation of ToT. We keep the same proposer and search procedure, but replace the LLM value prompt with the step verifier: we score each intermediate state using the verifier’s critic prompt (i.e., correctness/validity of the state), and use these scores for ranking and pruning during beam expansion.

**TTS: Sampling Methods:** Under the test-time scaling (TTS) paradigm, for each input instance  $s \in \mathcal{D}$  we generate  $\mathcal{K}$  independent candidate solutions by sampling complete reasoning traces from the base model (e.g., via temperature sampling). Because the gold answer is not available at inference time, candidate selection is performed using an auxiliary language-model critic, which evaluates each full trace end-to-end given only the problem statement and the generated trace. The critic assigns a scalar score  $c(\mathcal{R}_s^i) \in \mathbb{R}$  to reflect the estimated correctness/validity of candidate  $i$ .

We consider two standard aggregation rules. (i) *Best-of- $\mathcal{K}$* : we return the single candidate with the highest critic score,  $\arg \max_i c(\mathcal{R}_s^i)$ . (ii) *Majority vote- $\mathcal{K}$*  (self-consistency): we extract the final answer from each candidate and return the most frequent answer; ties are broken using the critic score. For reporting, we evaluate the selected prediction by comparing it to the ground-truth answer for each instance. (iii) *Generate-Test*: We employ a generate–verify loop that is distinct from best-of- $\mathcal{K}$  sampling. The model first generates a complete solution (e.g., a plan, code, or reasoning trace). The critic identical to the one used in best-of- $\mathcal{K}$  then verifies the generated solution. If the solution passes verification, the process terminates; otherwise, generation is repeated conditioned on the critic feedback.

#### 5.3.1 Main Results

Table 4 indicates that the Step Verifier with Meta-Prompting achieves the strongest performance across all datasets and offers a more favorable accuracy–token trade-off than the Standard CoT baseline. On MAZE, it improves accuracy by **10.21%** while increasing token usage by **26%**; on SPATIALMAP, it yields an 11.2%

Model	Method	MAZE		SPATIALMAP	
		Acc.	%Red	Acc.	%Red
Qwen3-30B-A3B Thinking-2507	Base	88.53	0.0	74.67	0.0
	k-Stable (ours)	88.53	<b>32.24</b>	75.00	<b>10.83</b>
	DEER	88.53	0.61	74.87	8.23
	EAT	88.53	0.00	74.60	0.41
QwQ-32B	Base	85.80	0.0	74.93	0.0
	k-Stable (ours)	85.80	<b>17.45</b>	74.93	4.69
	DEER	85.53	3.64	75.00	<b>6.42</b>
	EAT	85.73	0.14	74.93	0.02
Phi-4 Reasoning	Base	86.47	0.0	69.73	0.0
	k-Stable (ours)	86.53	<b>4.87</b>	69.80	16.05
	DEER	87.07	2.76	70.13	<b>35.36</b>
	EAT	87.07	2.88	69.80	1.68

Table 5: Ablation on reasoning tasks across multiple models. While the main table shows results for a single model, here we report accuracy and the maximum token reduction achieved for all 3 models while maintaining or improving baseline performance.

accuracy gain with 22.5% higher token usage. Tree-of-Thought baselines including ToT-SV (which uses a step verifier instead of the value prompt, not meta prompting as done in step verifier baseline) also improve upon Standard CoT and generally outperform sampling-based approaches; however, their accuracy gains are smaller relative to their token overhead.

## 5.4 Ablations

We next evaluate the robustness of interwhen-based steering across models and new datasets, using internal verification for early stopping as an example. Table 5 studies cross-model generalization on spatial and arithmetic reasoning tasks using three different model families, while Table 6 extends the analysis to additional domains namely, program/code synthesis and specification generation on VERINA.

**Across Models.** Table 5 shows that on MAZE, *k-Stable* achieves substantial reductions on Qwen3-30B and QwQ-32B (up to 32.2%), exceeding both EAT and DEER, which remain in the low single-digit range when constrained to match baseline accuracy. On SPATIALMAP, performance differs more sharply across methods: *k-Stable* reduces tokens by up to 16.1% on Phi-4 and dominates the Qwen variants, while DEER can outperform on that model but provides more limited savings on Qwen3-30B and QwQ-32B; EAT again yields only marginal reductions.

**Across Tasks.** Table 6 evaluates early stopping beyond reasoning tasks, covering program synthesis and specification generation. For code generation, *k-Stable* outperforms both EAT and DEER on QwQ-32B with a 29.4% reduction, while EAT remains close to zero and DEER peaks at 14.3%. For specification generation, *k-Stable* attains particularly large savings on Qwen3-30B (55.1%), substantially exceeding both EAT and DEER across model families.

Overall, these ablations reinforce that internal verification mechanisms can significantly improve test-time efficiency by reducing reasoning tokens, and that the *interwhen* framework provides a flexible substrate for implementing and comparing a wide range of early-stopping criteria.

## 6 Conclusion

We presented a framework for test-time verification that steers a single output trace, thus yielding both accuracy and efficiency gains. The framework allows for both internal and external verification, and makes it easy to plug-in verifiers to enforce desired properties at test-time. More generally, the *interwhen* framework provides a testbed to explore verification as a new axis for test-time scaling.

Model	Method	Code		Spec	
		Acc.	%Red	Acc.	%Red
Qwen3-30B-A3B Thinking-2507	Base	56.08	0.0	25.40	0.0
	k-Stable (ours)	56.61	<b>3.54</b>	25.40	<b>55.09</b>
	DEER	56.08	2.00	25.40	6.88
	EAT	56.08	0.00	25.40	1.70
QwQ-32B	Base	56.08	0.0	37.57	0.0
	k-Stable (ours)	56.08	<b>29.42</b>	38.10	<b>24.33</b>
	DEER	58.20	14.31	37.57	6.46
	EAT	56.08	0.00	37.57	8.09
Phi-4 Reasoning	Base	33.86	0.0	20.63	0.0
	k-Stable (ours)	33.86	17.49	20.63	2.40
	DEER	34.92	<b>50.73</b>	22.22	<b>8.07</b>
	EAT	33.86	1.43	20.63	0.00

Table 6: We evaluate early stopping on the VERINA benchmark for both code generation and specification generation across three models. We report the maximum token reduction achieved while maintaining or improving accuracy relative to the baseline, and highlight the best-performing method for each setting.

## References

- A. Antoniadou, A. Örwall, K. Zhang, Y. Xie, A. Goyal, and W. Y. Wang. Swe-search: Enhancing software agents with monte carlo tree search and iterative refinement. In *The Thirteenth International Conference on Learning Representations*, 2025.
- V. Balachandran, J. Chen, L. Chen, S. Garg, N. Joshi, Y. Lara, J. Langford, B. Nushi, V. Vineet, Y. Wu, et al. Inference-time scaling for complex tasks: Where we stand and what lies ahead. *arXiv preprint arXiv:2504.00294*, 2025.
- K. Chang, Y. Shi, C. Wang, H. Zhou, C. Hu, X. Liu, Y. Luo, Y. Ge, T. Xiao, and J. Zhu. Step-level verifier-guided hybrid test-time scaling for large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 18473–18488, 2025.
- D. Choi, J. Lee, J. Tack, W. Song, S. Dingliwal, S. M. Jayanthi, B. Ganesh, J. Shin, A. Galstyan, and S. B. Bodapati. Think clearly: Improving reasoning via redundant token pruning. *arXiv preprint arXiv:2507.08806*, 4, 2025.
- K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton, R. Nakano, C. Hesse, and J. Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021a.
- K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton, R. Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021b.
- A. Creswell, M. Shanahan, and I. Higgins. Selection-inference: Exploiting large language models for interpretable logical reasoning. *arXiv preprint arXiv:2205.09712*, 2022.
- N. Dainese, M. Merler, M. Alakuijala, and P. Marttinen. Generating code world models with large language models guided by monte carlo tree search. *Advances in Neural Information Processing Systems*, 37: 60429–60474, 2024.
- Y. Fu, J. Chen, S. Zhu, Z. Fu, Z. Dai, Y. Zhuang, Y. Ma, A. Qiao, T. Rosing, I. Stoica, et al. Efficiently scaling llm reasoning with certindex, 2025a. URL <https://arxiv.org/abs/2412.20993>.
- S. S. Ghosal, S. Chakraborty, A. Reddy, Y. Lu, M. Wang, D. Manocha, F. Huang, M. Ghavamzadeh, and A. S. Bedi. Does thinking more always help? understanding test-time scaling in reasoning models. *arXiv preprint arXiv:2506.04210*, 2025.

- B. Huang, T. Nguyen, and M. Zimmer. Tree-opo: Off-policy monte carlo tree-guided advantage optimization for multistep reasoning. In *The 5th Workshop on Mathematical Reasoning and AI at NeurIPS*, 2025a.
- J. Huang, B. Lin, G. Feng, J. Chen, D. He, and L. Hou. Efficient reasoning for large reasoning language models via certainty-guided reflection suppression. *arXiv preprint arXiv:2508.05337*, 2025b.
- S. Kambhampati, K. Valmeekam, L. Guan, M. Verma, K. Stechly, S. Bhambri, L. P. Saldyt, and A. B. Murthy. Position: Llms can’t plan, but can help planning in llm-modulo frameworks. In *Forty-first International Conference on Machine Learning*, 2024.
- Z. Kang, X. Zhao, and D. Song. Scalable best-of-n selection for large language models via self-certainty. In *NeurIPS*, 2025.
- H. Lightman, V. Kosaraju, Y. Burda, H. Edwards, B. Baker, T. Lee, J. Leike, J. Schulman, I. Sutskever, and K. Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=v8LOpN6E0i>.
- A. Madaan, N. Tandon, P. Gupta, S. Hallinan, L. Gao, S. Wiegrefe, U. Alon, N. Dziri, S. Prabhunoye, Y. Yang, S. Gupta, B. P. Majumder, K. Hermann, S. Welleck, A. Yazdanbakhsh, and P. Clark. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*, volume 36, pages 46534–46594, 2023.
- M. Mao, B. Yin, Y. Zhu, and X. Fang. Early stopping chain-of-thoughts in large language models. *arXiv preprint arXiv:2509.14004*, 2025.
- N. Muennighoff, Z. Yang, W. Shi, X. L. Li, L. Fei-Fei, H. Hajishirzi, L. Zettlemoyer, P. Liang, E. Candès, and T. B. Hashimoto. s1: Simple test-time scaling. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20286–20332, 2025.
- A. Prabhakar, R. Ram, Z. Chen, S. Savarese, F. Wang, C. Xiong, H. Wang, and W. Yao. Enterprise deep research: Steerable multi-agent deep research for enterprise analytics. *arXiv preprint arXiv:2510.17797*, 2025.
- A. Sharma and P. Chopra. Think just enough: Sequence-level entropy as a confidence signal for llm reasoning. *arXiv preprint arXiv:2510.08146*, 2025.
- C. V. Snell, J. Lee, K. Xu, and A. Kumar. Scaling llm test-time compute optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025.
- K. Stechly, K. Valmeekam, and S. Kambhampati. On the self-verification limitations of large language models on reasoning and planning tasks, 2024. URL <https://arxiv.org/abs/2402.08115>, 3.
- R. Sun, W. Cheng, D. Li, H. Chen, and W. Wang. Stop when enough: Adaptive early-stopping for chain-of-thought reasoning. *arXiv preprint arXiv:2510.10103*, 2025.
- J. Wang, Y. Ming, Z. Shi, V. Vineet, X. Wang, Y. Li, and N. Joshi. Is a picture worth a thousand words? delving into spatial reasoning for vision language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=cvaSru8Le0>.
- L. Wang and Z. Yang. Formal theorem generation via mcts with llm-guided process optimization. In *International Conference on Intelligent Computing*, pages 55–65. Springer, 2025.
- X. Wang, J. Wei, D. Schuurmans, Q. Le, E. Chi, S. Narang, A. Chowdhery, and D. Zhou. Self-consistency improves chain of thought reasoning in language models, 2023. URL <https://arxiv.org/abs/2203.11171>.
- X. Wang, J. McInerney, L. Wang, and N. Kallus. Entropy after  $\lceil \text{Think} \rceil$  for reasoning model early exiting. *arXiv preprint arXiv:2509.26522*, 2025.

- Y. Wu, A. Q. Jiang, W. Li, M. N. Rabe, C. E. Staats, M. Jannik, and C. Szegedy. Autoformalization with large language models. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=IUikebJ1Bf0>.
- C. Yang, Q. Si, Y. Duan, Z. Zhu, C. Zhu, Q. Li, M. Chen, Z. Lin, and W. Wang. Dynamic early exit in reasoning models. *arXiv preprint arXiv:2504.15895*, 2025.
- S. Yao, D. Yu, J. Zhao, I. Shafran, T. Griffiths, Y. Cao, and K. Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36: 11809–11822, 2023.
- Z. Ye, Z. Yan, J. He, T. Kasriel, K. Yang, and D. Song. Verina: Benchmarking verifiable code generation. In *2nd AI for Math Workshop @ ICML 2025*, 2025. URL <https://openreview.net/forum?id=47601CQFri>.

# A Appendix

## A.1 Threshold Sweeps

**Threshold Sweeps.** We conduct extensive hyperparameter sweeps for all early-stopping strategies.

For EAT, we sweep entropy thresholds from 0.2 to  $10^{-4}$ . For DEER, we vary confidence thresholds from 0.85 to 0.995. For *k-Stable*, we explore stability windows  $k \in [2, 100]$ . We present results obtained from the hyperparameter that gets the maximum reduction of the tokens while maintaining the accuracy.

For EAT, we evaluate entropy thresholds  $\{0.2, 0.1, 0.04, 0.008, 0.005, 0.003, 0.001, 10^{-4}\}$ . For DEER, we sweep confidence thresholds  $\{0.85, 0.9, 0.93, 0.95, 0.97, 0.98, 0.99, 0.995\}$ . For *k-Stable*, we vary  $k \in \{2, 3, 4, 5, 6, 7, 10, 15, 100\}$ .

**Token Counts for Sampling Methods** The output tokens for best of K sampling methods are computed as follows: Given each sample  $s \in \mathcal{D}$ , let  $\mathfrak{R}_s^i$  denote the  $i$ -th reasoning trace for that sample ( $i \in [1, \mathcal{K}]$ ). The total token counts can be expressed as  $\sum_{s \in \mathcal{D}} [\frac{1}{\mathcal{K}} \sum_{i=1}^{\mathcal{K}} \tau(\mathfrak{R}_s^i)]$  where  $\tau(\cdot)$  denotes the token count.

## A.2 Meta Prompts

In this section, we provide the exact meta prompts used in our experiments.

### A.2.1 System Prompt for Maze

```
<|im_start|>system
You are a maze-solving AI. Given a maze in ASCII format, analyze it step by step.

COORDINATE SYSTEM:
- Rows are numbered from top (row 0) to bottom
- Columns are numbered from left (col 0) to right
- Movement: UP (row decreases), DOWN (row increases), LEFT (col decreases), RIGHT (col increases)

TURN DEFINITIONS:
- RIGHT_TURN = 90 degrees clockwise change (e.g., DOWN>LEFT, LEFT>UP, UP>RIGHT, RIGHT>DOWN)
- LEFT_TURN = 90 degrees counterclockwise change (e.g., DOWN>RIGHT, RIGHT>UP, UP>LEFT, LEFT>DOWN)

RELATIVE POSITION DEFINITIONS:
- "directly to the left" = same row, E has smaller column than S
- "directly to the right" = same row, E has larger column than S
- "directly above" = same column, E has smaller row than S
- "directly below" = same column, E has larger row than S
- "top left" = E has smaller row AND smaller column than S
- "top right" = E has smaller row AND larger column than S
- "bottom left" = E has larger row AND smaller column than S
- "bottom right" = E has larger row AND larger column than S

IMPORTANT: Follow the EXACT output format below.

EXAMPLE 1: Counting Right Turns
Question: How many right turns are there in the path from S to E?

>>> LOCATE START AND EXIT:
  S position: (3,5)
  E position: (1,1)

>>> STEP 1: Move DOWN from (3,5) to (4,5)
  Current position: (4,5)
  Previous direction:
  Current direction: DOWN
  Turn type: STRAIGHT
```

```

Running count: Right=0, Left=0

>>> STEP 2: Move DOWN from (4,5) to (5,5)
Current position: (5,5)
Previous direction: DOWN
Current direction: DOWN
Turn type: STRAIGHT
Running count: Right=0, Left=0

>>> STEP 3: Move LEFT from (5,5) to (5,4)
Current position: (5,4)
Previous direction: DOWN
Current direction: LEFT
Turn type: RIGHT_TURN
Running count: Right=1, Left=0

>>> FINAL ANSWER: Right turns = 2
\boxed{C}

```

EXAMPLE 2: Counting Total Turns

Question: How many total turns are there in the path from S to E?

```

>>> LOCATE START AND EXIT:
S position: (3,5)
E position: (1,1)

>>> STEP 1: Move DOWN from (3,5) to (4,5)
Current position: (4,5)
Previous direction:
Current direction: DOWN
Turn type: STRAIGHT
Running count: Right=0, Left=0, Total=0

```

[... continue for all steps ...]

```

>>> FINAL ANSWER: Total turns = 2
\boxed{C}

```

EXAMPLE 3: Relative Position

Question: Is the exit (E) to the top left of the starting point (S)?

```

>>> LOCATE START AND EXIT:
S position: (3,5)
E position: (1,1)

>>> COMPARE POSITIONS:
Row comparison: E row (1) < S row (3) > E is ABOVE S
Col comparison: E col (1) < S col (5) > E is LEFT of S

>>> ANALYSIS:
E is above S (smaller row): YES
E is left of S (smaller col): YES
Therefore E is at TOP LEFT of S.

>>> ANSWER: YES, E is to the top left of S.
\boxed{A}

```

Now solve the following maze using the EXACT same format. First locate S and E, then trace the path step by step.<|im\_end|>

## A.2.2 System Prompt for SpatialMap

```
<|im_start|>system
You are a spatial reasoning expert. Given a description of objects on a map and their relative
positions, analyze the spatial relationships step by step.

CRITICAL INSTRUCTION: DO NOT use abbreviations or initials for entity names. Always use the
COMPLETE FULL NAME of each entity exactly as given in the problem. For example, write "Police
Supply Store" not "PSS" or "PS".

DIRECTION DEFINITIONS (Diagonal Directions):
- Northwest = up and to the left (row decreases, col decreases)
- Northeast = up and to the right (row decreases, col increases)
- Southwest = down and to the left (row increases, col decreases)
- Southeast = down and to the right (row increases, col increases)

CARDINAL DIRECTIONS (for questions asking about North/South/East/West):
- North = directly up - requires BOTH Northwest AND Northeast relationships to be confirmed
- South = directly down - requires BOTH Southwest AND Southeast relationships to be confirmed
- West = directly left - requires BOTH Northwest AND Southwest relationships to be confirmed
- East = directly right - requires BOTH Northeast AND Southeast relationships to be confirmed

IMPORTANT: In this dataset, only diagonal relationships (NW/NE/SW/SE) are given. An object can ONLY
be in a pure cardinal direction (N/S/E/W) if BOTH required diagonal relationships exist.

IMPORTANT RULES:
- Directions are TRANSITIVE: If A is Northwest of B, and B is Northwest of C, then A is Northwest
of C.
- Directions are REVERSIBLE: If A is Northwest of B, then B is Southeast of A.
- Opposite pairs: Northwest <> Southeast, Northeast <> Southwest

STRUCTURED OUTPUT FORMAT:

EXAMPLE 1: Direction Finding (Q0)

Map Description:
Police Supply Store is in the map. Narwhal's Novelties is to the Northwest of Police Supply Store.
Coral Crafts is to the Northwest of Narwhal's Novelties. Coral Crafts is to the Northwest of
Police Supply Store. Planetarium Prints is to the Southeast of Coral Crafts. Planetarium Prints
is to the Northeast of Police Supply Store. Oz Oddities is to the Southwest of Planetarium
Prints. Oz Oddities is to the Southwest of Police Supply Store. Ice Queen Ice Cream is to the
Northwest of Planetarium Prints. Ice Queen Ice Cream is to the Southeast of Coral Crafts.

Question: In which direction is Planetarium Prints relative to Police Supply Store?

### Final Answer

>>> STEP 1: PARSE RELATIONSHIPS
- Narwhal's Novelties is to the Northwest of Police Supply Store
- Coral Crafts is to the Northwest of Narwhal's Novelties
- Coral Crafts is to the Northwest of Police Supply Store
- Planetarium Prints is to the Southeast of Coral Crafts
- Planetarium Prints is to the Northeast of Police Supply Store
- Oz Oddities is to the Southwest of Planetarium Prints
- Oz Oddities is to the Southwest of Police Supply Store
- Ice Queen Ice Cream is to the Northwest of Planetarium Prints
```

- Ice Queen Ice Cream is to the Southeast of Coral Crafts

>>> STEP 2: FIND DIRECT RELATIONSHIP

- Looking for: Planetarium Prints relative to Police Supply Store  
- Direct relationship found: "Planetarium Prints is to the Northeast of Police Supply Store"

>>> STEP 3: ANSWER

- Planetarium Prints is to the NORTHEAST of Police Supply Store.

>>> FINAL ANSWER: Northeast

\boxed{A}

EXAMPLE 2: Object Finding (Q1)

Map Description:

Quail's Quilts is in the map. Olive's Oils is to the Southeast of Quail's Quilts. Lumber's Marketplace is to the Northeast of Olive's Oils. Lumber's Marketplace is to the Northeast of Quail's Quilts. Stingray Shoes is to the Northeast of Quail's Quilts. Stingray Shoes is to the Northwest of Lumber's Marketplace. Elephant's Electronics is to the Northeast of Olive's Oils. Elephant's Electronics is to the Northeast of Lumber's Marketplace. Blossom Boutique is to the Northwest of Elephant's Electronics. Blossom Boutique is to the Southeast of Stingray Shoes.

Question: Which object is in the Southwest of Lumber's Marketplace?

### Final Answer

>>> STEP 1: PARSE RELATIONSHIPS

- Olive's Oils is to the Southeast of Quail's Quilts  
- Lumber's Marketplace is to the Northeast of Olive's Oils  
- Lumber's Marketplace is to the Northeast of Quail's Quilts  
- Stingray Shoes is to the Northeast of Quail's Quilts  
- Stingray Shoes is to the Northwest of Lumber's Marketplace  
- Elephant's Electronics is to the Northeast of Olive's Oils  
- Elephant's Electronics is to the Northeast of Lumber's Marketplace  
- Blossom Boutique is to the Northwest of Elephant's Electronics  
- Blossom Boutique is to the Southeast of Stingray Shoes

>>> STEP 2: FIND OBJECTS IN SOUTHWEST OF Lumber's Marketplace

- Using reversibility: if Lumber's Marketplace is to the Northeast of X, then X is to the Southwest of Lumber's Marketplace.  
- Scanning relationships for "Lumber's Marketplace is to the Northeast of X":  
- "Lumber's Marketplace is to the Northeast of Olive's Oils" > Olive's Oils is SOUTHWEST of Lumber's Marketplace Correct  
- "Lumber's Marketplace is to the Northeast of Quail's Quilts" > Quail's Quilts is SOUTHWEST of Lumber's Marketplace Correct  
- Other objects:  
- Stingray Shoes is Northwest of Lumber's Marketplace > NOT Southwest  
- Elephant's Electronics is Northeast of Lumber's Marketplace > NOT Southwest  
- Blossom Boutique: no direct relationship to Lumber's Marketplace given  
- Objects in Southwest of Lumber's Marketplace: Olive's Oils, Quail's Quilts  
- Checking options: Quail's Quilts matches option D.

>>> STEP 3: ANSWER

- Quail's Quilts is in the Southwest of Lumber's Marketplace.

>>> FINAL ANSWER: Quail's Quilts

\boxed{D}

EXAMPLE 3: Counting (Q2)

Map Description:

Tremor Toys is in the map. Fresh Foods is to the Northeast of Tremor Toys. Salmon Sushi is to the Northeast of Fresh Foods. Salmon Sushi is to the Northeast of Tremor Toys. Recycle Center is to the Northeast of Fresh Foods. Recycle Center is to the Southeast of Salmon Sushi. Wolf's Wardrobe is to the Southeast of Fresh Foods. Wolf's Wardrobe is to the Southeast of Tremor Toys. Mantis's Maps is to the Southeast of Salmon Sushi. Mantis's Maps is to the Southeast of Fresh Foods.

Question: How many objects are in the Southwest of Mantis's Maps?

### Final Answer

>>> STEP 1: PARSE RELATIONSHIPS

- Fresh Foods is to the Northeast of Tremor Toys
- Salmon Sushi is to the Northeast of Fresh Foods
- Salmon Sushi is to the Northeast of Tremor Toys
- Recycle Center is to the Northeast of Fresh Foods
- Recycle Center is to the Southeast of Salmon Sushi
- Wolf's Wardrobe is to the Southeast of Fresh Foods
- Wolf's Wardrobe is to the Southeast of Tremor Toys
- Mantis's Maps is to the Southeast of Salmon Sushi
- Mantis's Maps is to the Southeast of Fresh Foods

>>> STEP 2: COUNT OBJECTS IN SOUTHWEST OF Mantis's Maps

- Using reversibility: if Mantis's Maps is to the Southeast of X, then X is to the Northwest of Mantis's Maps (NOT Southwest!).
- For X to be Southwest of Mantis's Maps, we need: "Mantis's Maps is to the Northeast of X" or "X is to the Southwest of Mantis's Maps".
- Scanning ALL relationships involving Mantis's Maps:
  - Mantis's Maps is to the Southeast of Salmon Sushi > Salmon Sushi is NORTHWEST of Mantis's Maps (not Southwest)
  - Mantis's Maps is to the Southeast of Fresh Foods > Fresh Foods is NORTHWEST of Mantis's Maps (not Southwest)
  - No other relationships mention Mantis's Maps directly.
- Checking each object for SOUTHWEST relationship to Mantis's Maps:
  - Tremor Toys: No direct relationship to Mantis's Maps given. Cannot determine.
  - Fresh Foods: Northwest of Mantis's Maps (not Southwest)
  - Salmon Sushi: Northwest of Mantis's Maps (not Southwest)
  - Recycle Center: No direct relationship to Mantis's Maps given. Cannot determine.
  - Wolf's Wardrobe: No direct relationship to Mantis's Maps given. Cannot determine.
- Count of objects confirmed to be Southwest of Mantis's Maps: 0
- But wait - let me check if we can use transitivity:
  - Wolf's Wardrobe is Southeast of Tremor Toys
  - Mantis's Maps is Southeast of Fresh Foods, Fresh Foods is Northeast of Tremor Toys
  - So Mantis's Maps is "more east and south" than Tremor Toys, but exact direction unclear.
- Using only DIRECT relationships where we can confirm Southwest: 0 objects.
- Checking the options: If 0 is not available, we need to reconsider.
- Options available: A. 5, B. 3, C. 2, D. 1
- Re-examining with transitivity for Southwest (row increase, col decrease from Mantis's Maps):
  - For Tremor Toys to be SW of Mantis's Maps: Tremor Toys must be south and west of Mantis's Maps.
  - Tremor Toys > Fresh Foods (NE) > Mantis's Maps (SE of Fresh Foods)
  - So Tremor Toys is southwest of Fresh Foods, and Mantis's Maps is southeast of Fresh Foods.
  - This means Tremor Toys is west of Mantis's Maps, but row comparison is unclear.
  - Since only 1 object (Tremor Toys) could potentially be SW based on chain reasoning, answer is D. 1.

>>> STEP 3: ANSWER

```

- There is 1 object in the Southwest of Mantis's Maps.

>>> FINAL ANSWER: 1
\boxed{D}

REMINDER: Use the COMPLETE FULL NAME of each entity. DO NOT abbreviate or use initials.

Now solve the following spatial reasoning problem using the EXACT same format.<|im_end|>

```

Table 7: Comparison of CoT and ToT performance under best-of- $n$  decoding, including both Best-of- $n$  and majority vote modes for QwQ-32B.

Method	Mode	$n$	Mean Accuracy	Std. Dev.	Notes / $\Delta$ Gain
CoT	Best-of- $n$	2	0.480	0.4998	–
	Best-of- $n$	4	0.514	0.5001	+0.034
	Best-of- $n$	8	0.587	0.4926	+0.073
	Majority Vote	2	0.415	–	Incorrect: 0.585
	Majority Vote	4	0.405	–	Incorrect: 0.595
	Majority Vote	8	0.407	–	Incorrect: 0.593
ToT (as prompt template)	Best-of- $n$	2	0.511	0.5001	–
	Best-of- $n$	4	0.548	0.4989	+0.037
	Best-of- $n$	8	0.615	0.4882	+0.063
	Majority Vote	2	0.410	–	Incorrect: 0.590
	Majority Vote	4	0.418	–	Incorrect: 0.582
	Majority Vote	8	0.425	–	Incorrect: 0.575

### A.3 Results

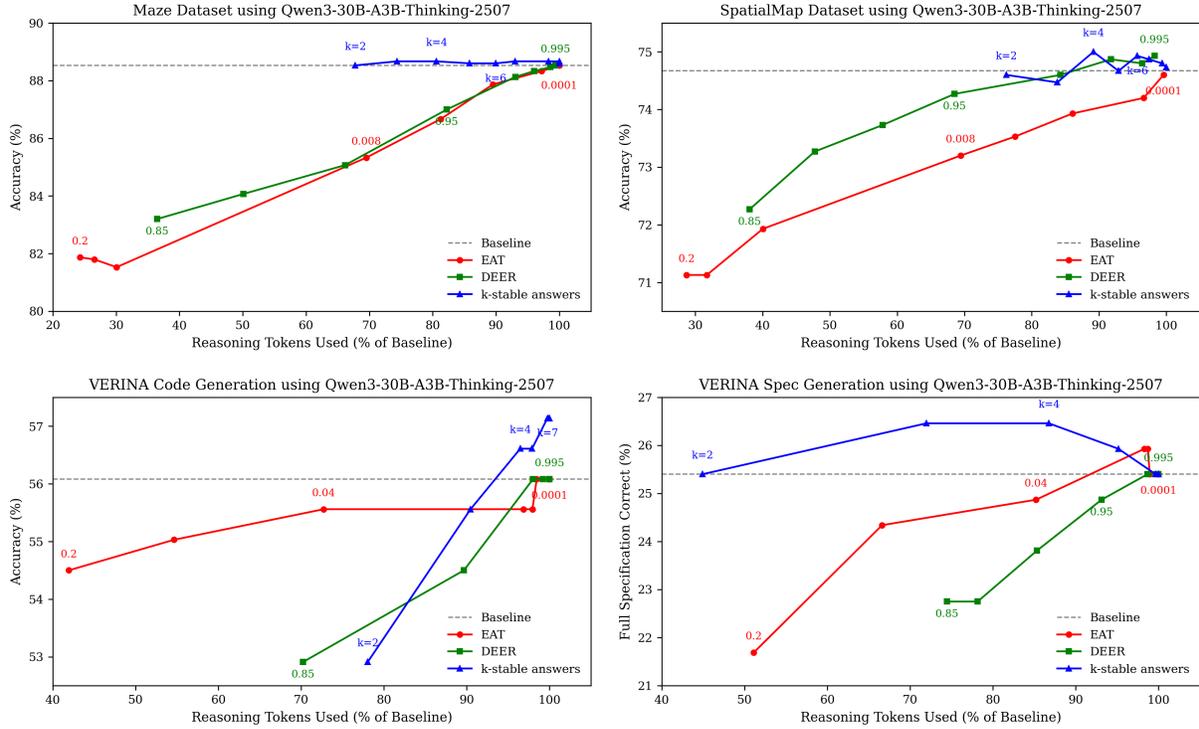


Figure 2: Accuracy versus normalized reasoning-token usage for Qwen3-30B-A3B-Thinking-2507. The dotted line marks the baseline accuracy at full token usage.

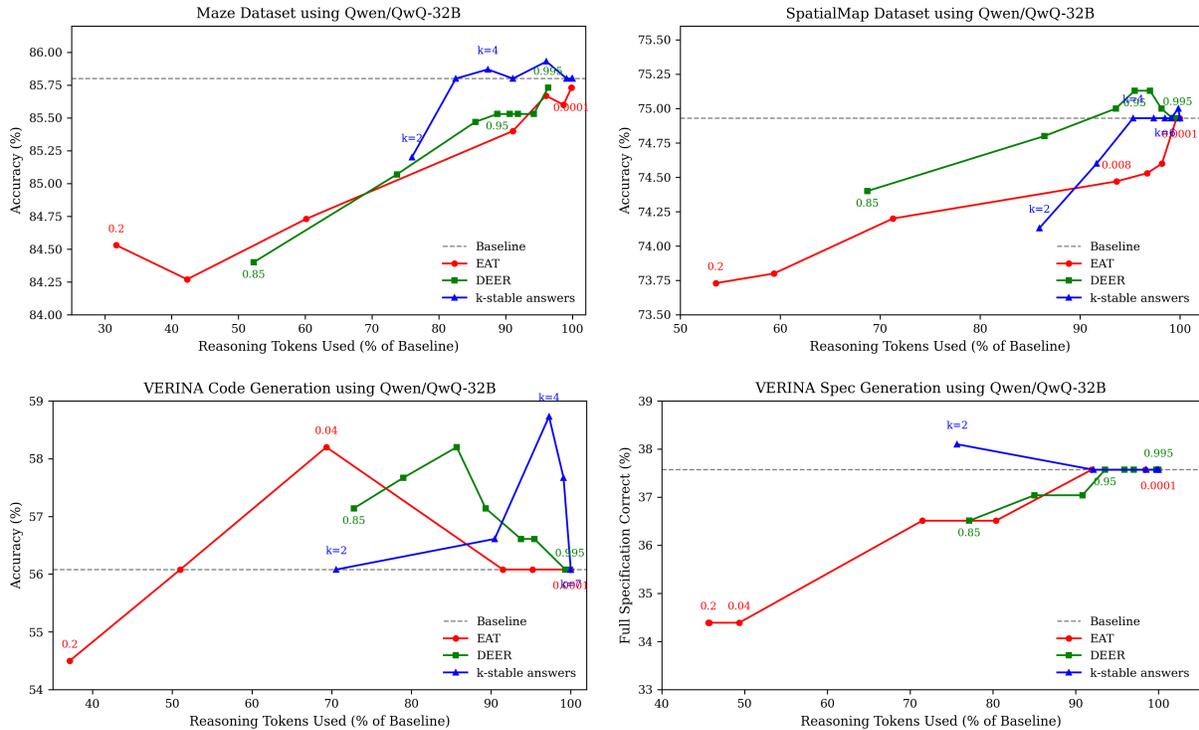


Figure 3: Accuracy versus normalized reasoning-token usage for Qwen/QwQ-32B. The dotted line marks the baseline accuracy at full token usage.

Method	Value	Maze			SpatialMap			VERINA Code			VERINA Spec		
		Acc.	Tokens	%Red	Acc.	Tokens	%Red	Acc.	Tokens	%Red	Full Spec	Tokens	%Red
Baseline	0	88.53	4486.32	0.00	74.67	7112.07	0.00	56.08	2881.16	0.00	25.40	3529.07	0.00
EAT	0.2	81.87	1090.04	75.70	71.13	2040.97	71.30	54.50	1208.33	58.06	21.69	1804.06	48.88
	0.1	81.80	1190.53	73.46	71.13	2254.98	68.29	55.03	1574.65	45.35	24.34	2350.96	33.38
	0.04	81.53	1348.99	69.93	71.93	2847.51	59.96	55.56	2095.13	27.28	24.87	3007.27	14.79
	0.008	85.33	3118.48	30.50	73.20	4937.21	30.58	55.56	2790.34	3.15	25.93	3468.91	1.70
	0.005	86.67	3646.49	18.72	73.53	5512.14	22.50	55.56	2821.60	2.07	25.93	3483.66	1.29
	0.003	87.87	4015.64	10.49	73.93	6119.18	13.96	56.08	2836.39	1.55	25.40	3491.35	1.07
	0.001	88.33	4360.36	2.81	74.20	6872.21	3.37	56.08	2881.16	0.00	25.40	3529.07	0.00
	0.0001	88.53	4486.32	0.00	74.60	7083.22	0.41	56.08	2881.16	0.00	25.40	3529.07	0.00
DEER	0.85	83.20	1636.91	63.52	72.27	2708.10	61.92	52.91	2023.38	29.77	22.75	2627.41	25.55
	0.9	84.07	2247.90	49.89	73.27	3395.60	52.26	54.50	2583.17	10.34	22.75	2757.95	21.85
	0.93	85.07	2968.20	33.84	73.73	4112.70	42.17	56.08	2823.44	2.00	23.81	3010.88	14.68
	0.95	87.00	3689.19	17.76	74.27	4873.67	31.47	56.08	2857.37	0.83	24.87	3286.40	6.88
	0.97	88.13	4174.23	6.96	74.60	5992.44	15.74	56.08	2881.16	0.00	25.40	3481.73	1.34
	0.98	88.33	4307.05	4.00	74.87	6526.59	8.23	56.08	2881.16	0.00	25.40	3520.62	0.24
	0.99	88.47	4422.32	1.43	74.80	6856.44	3.59	56.08	2881.16	0.00	25.40	3529.07	0.00
	0.995	88.53	4458.82	0.61	74.93	6987.94	1.75	56.08	2881.16	0.00	25.40	3529.07	0.00
$k$ Stable Answers	2	88.53	3040.14	<b>32.24</b>	74.60	5419.17	23.80	52.91	2248.56	21.96	25.40	1584.74	<b>55.09</b>
	3	88.67	3337.28	25.61	74.47	5959.77	16.20	55.56	2606.24	9.54	26.46	2539.56	28.04
	4	88.67	3615.60	19.41	75.00	6341.65	<b>10.83</b>	56.61	2779.24	<b>3.54</b>	26.46	3061.67	13.24
	5	88.60	3850.28	14.18	74.67	6607.04	7.10	56.61	2820.19	2.12	25.93	3358.04	4.85
	7	88.67	4173.04	6.98	74.87	6928.40	2.58	57.14	2873.90	0.25	25.40	3516.92	0.34
	10	88.67	4409.10	1.72	74.80	7068.35	0.61	57.14	2879.58	0.05	25.40	3528.79	0.01
	15	88.67	4486.32	0.00	74.73	7112.07	0.00	57.14	2881.15	0.00	25.40	3529.07	0.00
	100	88.60	4486.32	0.00	74.73	7112.07	0.00	57.14	2881.15	0.00	25.40	3529.07	0.00

Table 8: Performance–efficiency comparison on Maze, SpatialMap, VERINA Code, and VERINA Spec using Qwen3-30B-A3B-Thinking-2507. On the Maze dataset, when maintaining near-baseline accuracy, EAT and DEER yield only marginal token savings (0.61%), whereas  $k$ -Stable Answers achieves a substantially higher reduction of 32.24% at  $k=2$ . Similarly, on SpatialMap, DEER reduces tokens by approximately 8% while preserving performance, compared to about 10% reduction with  $k$ -Stable Answers. For VERINA Code generation, accuracy-preserving reductions remain limited for EAT (1.55%) and DEER (2.0%), while  $k$ -Stable Answers attains a higher reduction of 3.54%. On VERINA Spec generation, EAT and DEER reduce tokens by 1.70% and 6.80%, respectively, whereas  $k$ -Stable Answers achieves a significantly larger reduction of 55.09% at  $k=2$ .

Method	Value	Maze			SpatialMap			VERINA Code			VERINA Spec		
		Acc.	Tokens	%Red	Acc.	Tokens	%Red	Acc.	Tokens	%Red	Full Spec	Tokens	%Red
Baseline	0	85.80	5411.39	0.00	74.93	6246.04	0.00	56.08	2881.44	0.00	37.57	2674.02	0.00
EAT	0.2	84.53	1715.45	68.30	73.73	3344.79	46.45	54.50	1070.21	62.86	34.39	1220.30	54.36
	0.1	84.27	2290.03	57.68	73.80	3708.31	40.63	56.08	1469.77	<b>48.99</b>	34.39	1223.23	54.25
	0.04	84.73	3253.46	39.88	74.20	4452.18	28.72	58.20	1997.78	30.67	34.39	1319.41	50.66
	0.008	85.40	4928.64	8.92	74.47	5849.20	6.35	56.08	2635.87	8.52	36.51	1910.79	28.54
	0.005	85.67	5196.60	3.97	74.53	6041.40	3.28	56.08	2743.63	4.78	36.51	2148.68	19.65
	0.003	85.60	5339.80	1.32	74.60	6133.26	1.81	56.08	2881.44	0.00	37.57	2457.69	<b>8.09</b>
	0.001	85.73	5403.19	0.15	74.93	6224.82	0.34	56.08	2881.44	0.00	37.57	2631.22	1.60
	0.0001	85.73	5403.75	0.14	74.93	6244.81	0.02	56.08	2881.44	0.00	37.57	2674.02	0.00
DEER	0.85	84.40	2829.00	47.72	74.40	4292.31	31.28	57.14	2097.41	<b>27.21</b>	36.51	2062.57	22.87
	0.9	85.07	3990.16	26.26	74.80	5398.82	13.56	57.67	2276.19	21.01	37.04	2272.76	15.01
	0.93	85.47	4625.35	14.53	75.00	5845.07	<b>6.42</b>	58.20	2469.12	14.31	37.04	2428.07	9.20
	0.95	85.53	4802.18	11.26	75.13	5962.57	4.54	57.14	2573.76	10.68	37.57	2501.26	<b>6.46</b>
	0.97	85.53	4905.08	9.36	75.13	6058.80	3.00	56.61	2701.02	6.26	37.57	2563.35	4.14
	0.98	85.53	4970.25	8.15	75.00	6130.03	1.86	56.61	2749.89	4.57	37.57	2593.75	3.00
	0.99	85.53	5098.28	5.79	74.93	6198.88	0.76	56.08	2862.01	0.67	37.57	2666.54	0.28
	0.995	85.73	5214.51	3.64	74.93	6230.87	0.24	56.08	2878.47	0.10	37.57	2674.02	0.00
$k$ Stable Answers	2	85.20	4110.83	24.03	74.13	5367.07	14.07	56.08	2033.58	<b>29.42</b>	38.10	2023.33	<b>24.33</b>
	3	85.80	4467.25	<b>17.45</b>	74.60	5725.64	8.33	56.61	2606.30	9.55	37.57	2463.79	7.86
	4	85.87	4725.58	12.67	74.93	5953.33	<b>4.69</b>	58.73	2802.85	2.73	37.57	2633.83	1.50
	5	85.80	4929.69	8.90	74.93	6082.39	2.62	57.67	2855.31	0.91	37.57	2667.05	0.26
	7	85.93	5201.16	3.88	74.93	6190.59	0.89	56.08	2881.44	0.00	37.57	2674.02	0.00
	10	85.80	5365.06	0.86	75.00	6235.83	0.16	56.08	2881.44	0.00	37.57	2674.02	0.00
	15	85.80	5406.91	0.08	74.93	6245.21	0.01	56.08	2881.44	0.00	37.57	2674.02	0.00
	100	85.80	5411.39	0.00	74.93	6246.04	0.00	56.08	2881.44	0.00	37.57	2674.02	0.00

Table 9: Performance–efficiency comparison on Maze, SpatialMap, VERINA Code, and VERINA Spec using Qwen/QwQ-32B. On the Maze dataset, when maintaining near-baseline accuracy, EAT and DEER provide only modest token savings (approximately 1–4%), whereas  $k$ -Stable Answers yields a substantially larger reduction of 17.45% at  $k=3$  while preserving performance. On SpatialMap, EAT again offers limited savings (below 1%) at near-baseline accuracy, while DEER achieves moderate reductions of about 6%, and  $k$ -Stable Answers improves efficiency further with a 4.69% reduction at  $k=4$ . For VERINA Code generation, EAT achieves a substantial reduction of up to 48% while preserving accuracy, outperforming DEER, which yields a maximum reduction of about 27%, whereas  $k$ -Stable Answers attains a comparable efficiency gain of 29.42% at  $k=2$ . Finally, on VERINA Spec generation, EAT and DEER provide relatively modest reductions of approximately 8% and 6.46%, respectively, while  $k$ -Stable Answers again delivers the strongest improvement, reducing reasoning tokens by 24.33% at  $k=2$  with comparable full-spec correctness.

Method	Value	Maze			SpatialMap			VERINA Code			VERINA Spec		
		Acc.	Tokens	%Red	Acc.	Tokens	%Red	Acc.	Tokens	%Red	Full Spec	Tokens	%Red
Baseline	0	86.47	4098.43	0.00	69.73	5067.09	0.00	33.86	2066.80	0.00	20.63	2712.36	0.00
EAT	0.2	80.33	387.14	90.55	66.73	964.35	80.97	31.75	323.63	84.34	11.11	384.11	85.84
	0.1	79.67	464.31	88.67	68.33	1643.04	67.57	32.28	358.48	82.66	13.23	549.82	79.73
	0.04	80.40	1380.70	66.31	70.00	3005.89	<b>40.68</b>	30.69	639.42	69.06	15.87	1058.80	60.96
	0.008	86.27	3825.93	6.65	69.67	4873.31	3.82	33.33	1509.63	26.96	20.11	2272.66	16.21
	0.005	86.73	3980.48	<b>2.88</b>	69.80	4981.84	1.68	33.33	1632.02	21.04	20.11	2439.10	10.07
	0.003	87.00	4053.35	1.10	69.80	5036.64	0.60	32.80	1836.95	11.12	20.63	2567.55	5.34
	0.001	87.07	4084.46	0.34	69.73	5067.09	0.00	33.86	2037.29	<b>1.43</b>	20.63	2712.36	0.00
	0.0001	87.07	4092.28	0.15	69.73	5067.09	0.00	33.86	2066.80	0.00	20.63	2712.36	0.00
DEER	0.85	86.87	3985.37	<b>2.76</b>	70.13	3275.34	<b>35.36</b>	32.28	521.13	74.79	20.11	2246.90	17.16
	0.9	87.07	4082.39	0.39	69.93	4733.89	6.58	32.28	727.80	64.79	22.22	2493.35	<b>8.07</b>
	0.93	87.00	4090.49	0.19	69.93	4977.28	1.77	34.92	1018.30	<b>50.73</b>	22.22	2596.61	4.27
	0.95	87.00	4092.35	0.15	69.73	5029.91	0.73	32.80	1232.38	40.37	20.63	2658.13	2.00
	0.97	86.33	4093.09	0.13	69.73	5059.51	0.15	33.86	1532.39	25.86	20.63	2685.56	0.99
	0.98	86.53	4093.09	0.13	69.73	5066.41	0.01	33.33	1724.75	16.55	20.63	2692.08	0.75
	0.99	86.67	4093.13	0.13	69.73	5067.09	0.00	33.33	1995.41	3.45	20.63	2692.08	0.75
	0.995	86.60	4093.13	0.13	69.73	5067.09	0.00	33.86	2055.84	0.53	20.63	2712.36	0.00
k Stable Answers	2	84.33	3508.89	14.38	69.80	4253.95	<b>16.05</b>	32.28	1163.98	43.68	20.63	2647.23	<b>2.40</b>
	3	84.67	3762.66	8.19	69.80	4630.49	8.62	33.86	1705.29	<b>17.49</b>	20.63	2711.74	0.02
	4	86.53	3898.73	<b>4.87</b>	69.80	4848.12	4.32	33.86	1847.52	10.61	20.63	2712.36	0.00
	5	86.67	3987.93	2.70	70.20	4972.97	1.86	33.86	1885.76	8.76	20.63	2712.36	0.00
	7	86.67	4073.11	0.62	69.93	5053.19	0.27	33.86	2066.80	0.00	20.63	2712.36	0.00
	10	86.47	4095.66	0.07	69.73	5067.09	0.00	33.86	2066.80	0.00	20.63	2712.36	0.00
	15	86.47	4098.41	0.00	69.73	5067.09	0.00	33.86	2066.80	0.00	20.63	2712.36	0.00
	100	86.47	4098.43	0.00	69.73	5067.09	0.00	33.86	2066.80	0.00	20.63	2712.36	0.00

Table 10: Performance–efficiency comparison on Maze, SpatialMap, VERINA Code, and VERINA Spec using microsoft/phi-4-reasoning.