CIPHER: Leveraging test-time scaling for the design of an automated data science AI agent

Maxime Heuillet [*1,2] **Sharadind Peddijaru** [3]

Abstract

Data science tasks span from closed-ended information extraction to open-ended analysis, presenting significant challenges for automation. Recent AI agents powered by language models show promise for handling such complex tasks. However, existing agents typically rely on single initial state that conditions the entire agent's execution, making the agent vulnerable to cascading errors initiated by a suboptimal initial state. To mitigate this, we present CIPHER, an automated data science agent that leverages test-time scaling through the generation and selection of multiple initial states for concurrent execution. Unlike existing works on test-time scaling of AI agents, CIPHER explicitly decouples generating candidate initial states and strategically selecting a subset for parallel execution. Through extensive evaluation on two benchmarks—Infi-DA-Bench (closed-form tasks) and Insight-Bench (open-form tasks)— we demonstrate that CIPHER establishes new state-of-the-art performance, outperforming existing agents. Our empirical study reveals that test-time scaling benefits hold from efficient models to models closer to the frontier, particularly on ambiguous open-form tasks, and that optimal performance requires careful consideration of generation and selection strategies as well as aggregator model capacity. These findings establish a foundational framework for test-time scaling in agentic AI systems and provide actionable design recommendations for practitioners.

1. Introduction

Data science (Cao, 2017) is an interdisciplinary field that leverages statistics, computer science, and scientific methods to extract knowledge from data. A key challenge for data scientists is the wide variety of tasks. Data science tasks range from *closed-ended* where the goal is to extract a verifiable information (Hu et al., 2024), to *open-ended* where the goal is more ambiguous (Sahu et al., 2025). For example, formulating an exploratory hypothesis for data analysis is an open-ended task, while extracting specific information from the data to validate such hypothesis is a closed-ended task.

Data science is increasingly influenced by AI solutions, such as the generation of new features with deep learning (Kanter & Veeramachaneni, 2015) or pipeline optimization with automated machine learning (Olson et al., 2016; Heuillet et al., 2021). More recently, AI agents powered by language models (LM) have emerged as a promising approach (Hollmann et al., 2023; Wang et al., 2024). Unlike traditional AI solutions that target specific sub-tasks of data science, AI agents can handle complex workflows with minimal human oversight (Wang et al., 2024).

Existing data science agents (Hong et al., 2024; Nam et al., 2025; You et al., 2025) are designed to mimic how data scientists execute tasks, such as planning subsequent analysis stages based on prior results (Nam et al., 2025) or generating code through iterative debugging (You et al., 2025). All these works focus on a specific class of AI agents that utilize LMs through predefined execution route (Schluntz & Zhang, 2024). The execution of the route typically starts from a single initial state (e.g. a high-level plan) generated by a LM based on the task provided by the user. The quality of the initial state is a critical determinant of task success as errors in this initial step can cascade throughout the entire execution route. Even the most capable LMs can produce hallucinations or reason in sub-optimal ways (Zhang et al., 2024; Reid et al., 2025), making the initial state generation a highly vulnerable step within the execution route.

To address this limitation and in contrast with existing data science agents, we propose CIPHER[1], an agent that generates, selects, and executes multiple initial states concurrently

[*]Internship at AWS. [1]Université Laval (IID) [2]Montreal AI Institute (Mila) [3]Amazon Web Services. Correspondence to: Maxime Heuillet <maxime.heuillet.france@gmail.com>.

[0]Preprint.

[1]CIPHER is named after the algorithmic process of encoding, where a structured series of steps systematically transforms an input into a different output.

to reduce dependence on any single initial state. The proposed approach is related to test-time scaling as the goal is to allocate additional computational resources during inference to improve performance. There are several approaches for test-time scaling such as i) generating longer reasoning traces (Snell et al., 2024) or ii) generating multiple independent states from the same model (Brown et al., 2024; Anonymous, 2025b;c; Liu et al., 2025). We focus on the latter approach: generating multiple independent states and aggregating their results. Test-time scaling through concurrent executions is extensively studied in single-turn applications where a single LM query produces the final answer (Inoue et al., 2025; Wang et al., 2025). However, the implications of test-time scaling are less understood in the agentic AI setting, where multiple LM queries form an execution route.

Research question and contributions: *Can we design an automated data science agent that leverages test-time scaling to improve performance?* To address this question, we hypothesize that the generation of initial states and their selection influence the performance of the agent. We make three main contributions: **(1)** We present `CIPHER`, a data science agent that demonstrates performance improvements upon state-of-the-art agents on open and closed-form data science tasks. **(2)** Through `CIPHER`, we introduce a novel algorithmic framework for test-time scaling based on the generation and selection of initial states, opening the way for future research on this problem. **(3)** We propose an extensive empirical study on the optimal design of test-time scaling for agentic data science agents, which may be of independent interest for practitioners.

2. CIPHER: a data science agent

`CIPHER`'s execution route is illustrated in Figure 1. The agent receives as input a task noted $\mathcal{T}$ composed of a dataset and a textual user goal (open-ended or closed ended data science question). Examples of tasks and executions are provided in Appendix C. The execution route consists of a deterministic directed graph implemented in LangGraph (LangChain AI, 2024). The output is a textual response noted s^* that answers the user goal.

Initial state generation and selection `CIPHER` first generates a broad set of N candidate states, where N is a parameter set by the user. The candidates are generated by a `Planner` agent instructed to create detailed and actionable plans aiming to solve the task. Subsequently, a subset of $M < N$ states is selected for execution (M is a parameter set by the user). We detail in Section 3 strategies to generate and select states. The novelty is that `CIPHER` instantiates a decoupled generation-selection framework that differs from existing works on test-time scaling, which typically merge these two steps: a finite set of M states is generated and is fully executed (Zhang et al., 2025; Zhu et al., 2025). Since these initial states correspond to samples from the model's stochastic probability distribution, relying on a fixed, small M means the agent is vulnerable to being "unlucky" in the sampling of the initial state. Our approach explicitly separates exploration (generating a set N to cover a wider solution space) from exploitation/execution (selecting only the most promising M paths to limit resource consumption) (Lattimore & Szepesvári, 2020).

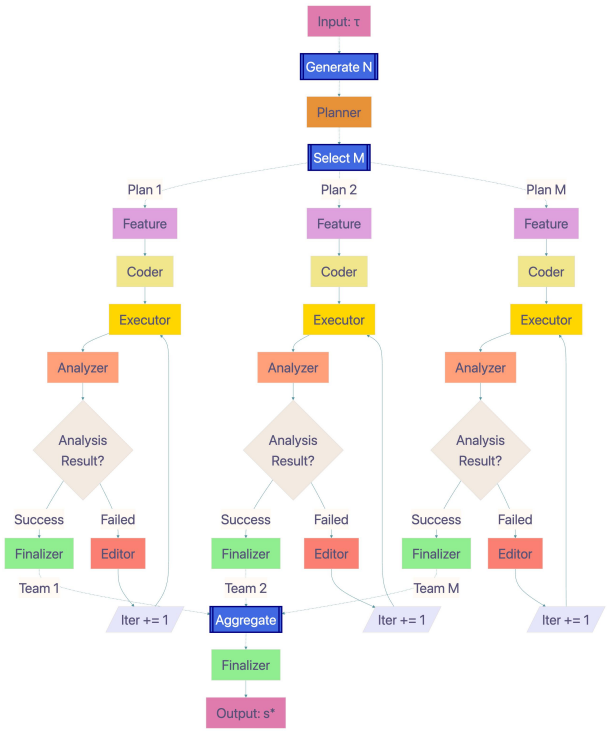

Figure 1. `CIPHER` execution route. The system generates N initial states, selects M for concurrent execution, and then aggregates the information towards a final output.

Parallel routes execution `CIPHER` executes the M selected states in parallel to decrease latency. Each selected state follows the same execution route detailed in Figure 1. The route is composed of a `Feature` agent (suggests the creation of relevant features), a coding loop (translates the plan into executable code), and a `Finalizer` agent (analyze code executions and provide a summary). Finally, an `Aggregator` agent synthesizes the information from the concurrent routes, enabling the final output. Parallelization can also be used for AI agent route sectioning (Anonymous, 2025a); however, our work focuses on parallelization to execute multiple routes concurrently, enabling diverse resolutions of the same task.

Code iterative refinement Loop `CIPHER` instantiates an iterative refinement loop (Chen et al., 2023) to improve

the quality of the generated code. A `Coder` agent generates a first code draft, which is then executed. Then, an `Analyzer` agent evaluates jointly the logs and the code to identify errors or suboptimal patterns. If issues are detected, an `Editor` agent modifies locally the code draft without re-generating it all. The loop repeats until the iteration budget b is reached (b is a hyper-parameter set by the user). We provide in Appendix A descriptive statistics on the effectiveness of the coding loop.

Trade-off related to cost efficiency By generating and executing multiple routes in parallel, `CIPHER` embodies a framework that inherently requires more computational resources than relying solely on a single route execution. To mitigate this increased cost, `CIPHER` is instructed to reason exclusively on the text modality. Although the generated code can produce visual artifacts (e.g., charts, plots), the downstream agents only consume the corresponding textual logs (e.g., summary statistics, model metrics, tables) printed during code execution. This design choice is motivated by two key factors: i) high performance on images necessitates multimodal reasoning LMs which are more costly, ii) image tokenization costs are orders of magnitude greater than processing equivalent information through text.

Algorithm 1 Test-Time Scaling Framework for `CIPHER`

Require: Task $\mathcal{T}$, generation budget N, selection budget M, generation mode g, selection mode v
Ensure: Final state s^*
Generate candidate states:
$\mathcal{C} \leftarrow \{\}$
for $n \in \{1, 2, \ldots, N\}$ **do**
  $c \leftarrow$ GENERATE-CANDIDATE-STATE$(\mathcal{T}, \mathcal{C}, m)$
  $\mathcal{C} \leftarrow \mathcal{C} \cup \{c\}$
end for
Select initial states:
$\mathcal{S} \leftarrow$ SELECT-STATES$(\mathcal{T}, \mathcal{C}, M, g)$
Execute initial states (conceptually in parallel):
$\mathcal{P} \leftarrow \{\}$
for each $s \in \mathcal{S}$ **do**
  $\tau \leftarrow$ EXECUTE-INITIAL-STATE$(\mathcal{T}, s)$
  $\mathcal{P} \leftarrow \mathcal{P} \cup \{\tau\}$
end for
Aggregate results:
$s^* \leftarrow$ AGGREGATE$(\mathcal{P})$
return s^*

3. Test-time scaling framework of `CIPHER`

In this Section, we outline the proposed test-time scaling framework for `CIPHER`. To our knowledge, generating a broad set of states and then selecting a subset for execution is a novel algorithmic problem in test-time scaling for AI agents. The integration of the framework with `CIPHER` is illustrated in Figure 1, and is further detailed in Algorithm 1. We highlight in green the main technical novelty compared to existing frameworks (Zhang et al., 2025).

3.1. Generating candidate initial states

The method GENERATE-CANDIDATE-STATE in Algorithm 1 outputs a set $\mathcal{C}$ (with $|\mathcal{C}| = N$) of candidate states. The method's inputs are the user task $\mathcal{T}$, the current set of candidate states $\mathcal{C}$ and a generation mode $m \in \{\texttt{independent}, \texttt{conditional}, \texttt{ensemble}\}$. The three modes are detailed below:

- **Independent generation**: Consists in using the same prompt and inference configuration N times to populate $\mathcal{C}$. This strategy samples independently from the same model's stochastic distribution conditioned by the prompt and inference configuration. This is a baseline strategy previously used in multi-turn LLM agent designs (Zhang et al., 2025) and single-turn LLM designs (Brown et al., 2024). Independent generation is embarrassingly parallel and has no sequential dependencies.
- **Conditional generation**: Consists in using $\mathcal{C}$ from prior calls to obtain a new candidate c on the current call to GENERATE-CANDIDATE-STATE. The prior candidates in $\mathcal{C}$ are used in the prompt instructions to elicit diversity in the upcoming candidates. Conditional generation is a form of adaptive sampling where each new sample is informed by the previous ones. This generation mode is investigated in prior works (Zhang et al., 2025; Sahu et al., 2025). It is not possible to parallelize in this mode.
- **Ensemble generation**: Consists in sampling N candidates from a set of E distinct prompts (parameter set by user), instead of relying only on one single prompt. More details are provided in Appendix regarding our instantiation B. Each prompt is used with probability $p = [p_1, \ldots, p_E]$ [2]. This mode is embarrassingly parallel. Different prompts may access complementary subspaces of the LM's capabilities. This is further motivated by prior research on ensemble learning which improves generalization (Dietterich, 2000; Freund & Schapire, 1997). To our knowledge, this ensembling generation mode is novel in test-time scaling of multi-turn agents. Broader impacts include the design of prompt ensembles (chain of thought, few shot, persona-based prompting (Olea et al., 2024)), and the learning of optimal p distributions.

Other generation modes certainly exist; for instance, prior works suggests the use of different inference configurations

[2] In expectation prompt $i \in \{1, \ldots, E\}$ is used Np_i times, with p_i being the probability of using prompt i.

for diversity. Results suggest it may be suboptimal compared to prompt-based strategies (il Lee et al., 2025). Our goal is to empirically evaluate these three distinct design paradigms to establish a foundational understanding of candidate state generation for agentic test-time scaling.

3.2. Selecting states for concurrent execution

The method SELECT-STATES outputs a set $\mathcal{S}$ ($|\mathcal{S}| = M$) of selected states that are to be executed concurrently. The method receives the set of candidate states $\mathcal{C}$, a selection budget M, the user task $\mathcal{T}$ and a selection mode $g \in \{\texttt{random}, \texttt{entropy}, \texttt{clustering}, \texttt{alignment}\}$. The four modes are detailed below:

- **Random selection** M states are selected uniformly at random from $\mathcal{C}$.
- **Entropy selection** Each candidate in $\mathcal{C}$ is projected into an embedding space (details below). Then, a greedy maximin algorithm is used to maximize the entropy of the M embeddings selected, using the cosine distance (Nemhauser et al., 1978).
- **Clustering selection** Each candidate in $\mathcal{C}$ is projected into an embedding space (details below). Then, a k-means algorithm (MacQueen, 1967) is used to identify M clusters. Finally, one candidate from each cluster is selected (the one closest to the centroid).
- **Alignment selection** A LM-judge is instructed to analyze the N candidates in $\mathcal{C}$, and to return the indices of the M candidates most aligned with the task. A limitation is that the judge's context increases with N, which can cause increase of inference costs and LM's cognitive saturation.

Embeddings The embeddings are obtained from Titan model (AWS, 2024), with $1,024$ as embedding size. To our knowledge, projecting candidate spaces into a latent space for subset selection is a novel algorithmic approach for test-time scaling AI agents. Designing embedding spaces that further capture semantic properties of candidates is a promising avenue for future work.

3.3. Aggregating execution results

To aggregate results, we use an open-source prompt previously proposed to aggregate results from concurrent executions available in Zhang et al. (2025). We investigate if the capacity of the LM aggregator is an influential factor in the final performance. We consider:

- **Self-aggregation** The LM aggregator has the same capacity as the rest of the agent.
- **Leader aggregation** The LM aggregator is a more advanced LM than the one used in the rest of the agent.

Table 1. InFidaBench results. Grey are from (You et al., 2025).

Model	Rank	Method	Acc.(%) ↑	Tok.(K)
GPT-4o-mini	139	Datawise	**82.88**	–
GPT-4o	113	Datawise	**85.99**	–
Haiku-3.5	128	Datawise	74.06 ± 1.25	16/1
		CIPHER (1,1)	69.13 ± 1.75	19/2
		CIPHER* (10,5)	**79.12** ± 0.22	82/5
Sonnet-3.7	77	Datawise	**86.77** ± 1.70	17/2
		CIPHER (1,1)	81.19 ± 0.59	24/3
		CIPHER* (10,5)	85.47 ± 1.47	89/7

Table 2. Insight-Bench results across three random seeds.

Model	Rank	Method	Acc.(%) ↑	Tok.(K)
Haiku-3.5	128	Agent Poirot	38.63 ± 0.71	51/10
		CIPHER(1,1)	38.32 ± 2.46	28/3
		CIPHER*(10,5)	**43.85** ± 0.54	96/6
Sonnet-3.7	77	Agent Poirot	43.93 ± 1.42	70/8
		CIPHER(1,1)	46.26 ± 1.21	47/6
		CIPHER*(10,5)	**54.05** ± 2.01	120/11

Prior works (Anonymous, 2025c) on single-turn LM test-time scaling propose to execute multiple concurrent executions and then use softmax averaging over the distribution of tokens to aggregate and obtain a final response. The limitation of this approach is that it assumes access to the logits of the LM, which is often not the case in real-world applications where LMs are accessed through APIs.

4. Experimental setup

We conduct experiment on two open-source benchmarks, namely Infi-DA-Bench (Hu et al., 2024) and Insight-Bench (Sahu et al., 2025). Infi-DA-Bench (Hu et al., 2024) evaluates performance on closed-ended tasks spanning exploratory data analysis and modeling while Insight-Bench (Sahu et al., 2025) evaluates performance on open-ended ambiguous tasks focused on hypothesis validation through data analysis. Therefore, Infi-DA-Bench corresponds to only a subspace of the skills required to solve Insight-Bench. Insight-Bench requires a higher level of autonomy.

4.1. Open-form tasks (Insight-Bench)

Insight-Bench (Sahu et al., 2025) includes a set of $T = 100$ tasks. Each of the T tasks is associated with a dataset and multiple insights about the dataset that need to be correctly retrieved by the agent. There is a total of 428 insights. To our knowledge, we are the first to include measurements on this benchmark.

Task setup We provide as guidance the agent with a set of exploratory hypotheses (formulated as open-ended ques-

tions) available in the meta-data ground truth of each task. This is akin to a data scientist having formulated a set of exploratory hypotheses and using the agent to conduct the analysis. Another possible formulation consists in evaluating the agent to formulate the hypotheses and insights autonomously without guidance. As evidenced by our results, both setups are challenging, with the full autonomy setup being even harder.

Performance metric We measure performance as the proportion of retrieved insights in the benchmark. The performance is defined as the proportion of insights correctly retrieved by the agent across the T questions and 428 insights: Accuracy $= \frac{\sum_{i=1}^{T} \text{Insights Retrieved}_i}{\text{Total Insights}}$. An increase of 1pp corresponds to retrieving approximately 4.3 insights. Other performance metrics have been proposed in (Sahu et al., 2025) based on semantic similarity between generated insights and ground truth. The main limitation of semantic based metrics is that they are is not easily interpretable and do not inform on the operational ability of the accurately retrieve insights, compared to the accuracy metric.

Correctness evaluation We use the LLM-as-a-judge paradigm (Zheng et al., 2023) to evaluate wether the agent's generated response correctly retrieves the ground truth insights. Given the final answer of the agent, we loop over each ground truth insight and ask wether this insight is supported in the generated response.

State-of-the-art The current state-of-the art on Insight-bench is the open-source strategy Agent-Poirot (Sahu et al., 2025). We evaluate Agent Poirot with the official hyper-parameters in (Sahu et al., 2025).

4.2. Closed-form tasks (Infi-DA-Bench)

Infi-DA-Bench benchmark includes a total of $T = 257$ tasks. Each question is composed of a dataset and a closed-ended question (e.g. descriptive statistics, modeling).

Performance metric We report the performance as the accuracy over the set of T tasks, measured as Accuracy $= \frac{1}{Q}\sum_{i=1}^{T}[\text{Correct Response}_i]$. An increase of 1pp corresponds to the correct resolution of approximately 2.57 tasks.

Correctness evaluation We use the official rule-based evaluation function to evaluate the responses (Hu et al., 2024). The function is strict: small numerical deviation from the ground truth (up to $1e-6$), partially true responses, or valid responses with wrong formatting are all marked as incorrect. Consequently, performance improvements under this strict evaluation are meaningful and significant.

State-of-the-art The current open-source state of the art is Data-Wise agent (You et al., 2025). Other agents exist but they are not open-sourced (Hong et al., 2024) or have been outperformed by Data-Wise agent (Wu et al., 2024). We consider Data-Wise agent as a baseline, and we use the official hyper-parameter configuration reported in the open-source repository.

4.3. Inference configurations

Unless otherwise specified, we execute using the Claude 3.5 Haiku model. To our knowledge, existing benchmarks and open-source results focus on OpenAI 4o models. Therefore, the proposed evaluations extend understanding of agent performance beyond OpenAI API. All the agents (baselines included) use the same inference configuration across all prompts: the temperature parameter is 0.5, the top-p is 0.99 and the top-k is 500. These parameters were selected after conducting a hyper-parameter sensitivity study on the following configurations: $(0.2, 0.99, 250)$, $(0.5, 0.99, 250)$, $(0.5, 0.99, 500)$, $(0.9, 0.99, 500)$. The evaluation of the responses on Insight-Bench is done with Claude 3.5 Haiku, with the temperature set to 0 for reproducibility.

5. Baseline Performance and Test-time Scaling Effectiveness

In this Section, we propose an empirical evaluation of `CIPHER` and its embedded test-time scaling framework. Each configuration is evaluated on three random seeds.

5.1. `CIPHER` performance against existing agents

We denote `CIPHER` (N=1, M=1) as the simplest instance without test-time scaling and with `base` generation. We denote `CIPHER`* (N=10, M=5) the instance with base generation and clustering as selection mode, which we study in Section 6.3. We compare against Agent-Poirot (Sahu et al., 2025) and Data-wise (You et al., 2025) baselines. Overall, `CIPHER`*(10,5) achieves state-of-the-art performance across 3 out of the 4 settings considered in Tables 1 and 2.

Insight-Bench From Table 2 we observe that `CIPHER`(N=1, M=1) achieves comparable performance to `Agent-Poirot` with Haiku 3.5, and outperforms Agent-Poirot with Sonnet 3.7. The performance gap with Agent-Poirot widens with `CIPHER`*(10,5). Specifically, `CIPHER`*(10,5) outperforms Agent-Poirot by +5pp with Haiku 3.5, and by +10pp with Sonnet 3.7.

Infi-DA-Bench We see in Table 1 that DataWise outperforms the base instance `CIPHER`(N=1, M=1) on both Haiku 3.5 and Sonnet 3.7. Indeed, DataWise instantiates a more complex execution route (leverages both iterative planning and iterative debugging) compared to CIPHER's (leverages iterative debugging only). However, `CIPHER`*(10,5) outperforms Datawise by +5pp with Haiku 3.5. and is close to state-of-the-art by 1pp with Sonnet 3.7.

5.2. Test-time scaling effectiveness on data science benchmarks

Gains of test-time scaling across model capacity Prior works (Zhang et al., 2025) identified test-time scaling as a promising avenue to improve performance of AI agents. However, prior results were confined to characterizing the gain on weaker LMs (e.g. GPT-4o, Qwen) leaving open the question of the benefit of test-time scaling to LMs closer to the frontier. Figure 2 confirms that test-time scaling is associated to consistent gains across different LM capacities. However, the gain scales differently across task types (open versus closed form). On closed-form questions (Infi-DA-Bench) the gain of test time scaling is more pronounced with efficient LMs like Haiku 3.5 (+10.0pp) compared to more capable LMs (+4.3pp). This is because efficient LMs like Haiku 3.5. have higher variance in their ability to correctly solve tasks. In contrast, on open-form tasks (Insight-Bench) the gain of test-time scaling increases from +5.5pp to +7.8pp as model capacity increases (from Haiku 3.5 to Sonnet 3.7). Therefore, the results invalidate the hypothesis that test-time scaling only improves performance on weaker and more efficient LMs. In fact, the performance gains increase with model capacity on highly ambiguous open-form tasks. This is because in that setting, test-time scaling brings multiple possible perspectives to resolve the problem.

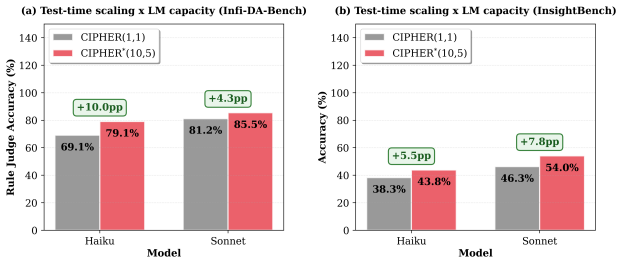

Figure 2. Accuracy (Acc.) over baseline `CIPHER` (N=1,M=1) of time scaling framework to CIPHER performance. Results are averaged across all generation, selection modes and seeds.

Influence of budget N and M on performance Figure 3 (a) shows that the performance improves by +1.9pp on Infi-DA-Bench and +1.4pp on InsightBench when scaling the selection budget from $M = 3$ to $M = 5$. This aligns with established results for single-turn LLM queries (Brown et al., 2024) that show a diminishing return on performance as the concurrency budget M increases. Our results demonstrate that this diminishing return pattern is maintained even in multi-turn agentic designs and holds on open-form and closed-form tasks. Panel (b) in Figure 3 demonstrates that increasing the generation budget N (from 10 to 30) produces unclear gains as performance increases on open-form benchmarks and degrades on closed-form benchmarks. This suggests that the selection budget M is the main driver of performance improvement, while the generation budget N contributes more modestly passed a threshold.

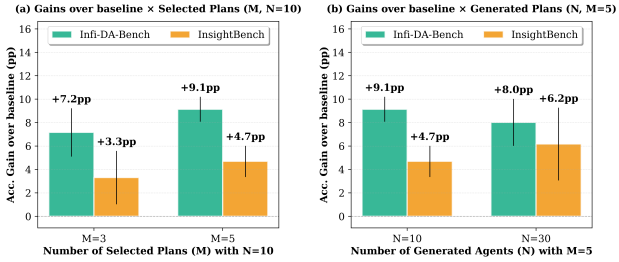

Figure 3. Accuracy (Acc.) over baseline `CIPHER` (N=1,M=1) of time scaling framework to CIPHER performance. Results are averaged across all generation, selection modes and seeds.

6. Configuration Ablations and Design Recommendations

Here, we investigate the influence of the three generation strategies and the four selection strategies on `CIPHER`'s performance. We obtain $3 \times 4 = 12$ instances which are evaluated for selection budgets $M \in \{3, 5\}$ across two benchmarks, totaling 48 configurations. Tables 3 and 4 feature results for Infi-DA-Bench and Insight-Bench respectively.

6.1. Impact of aggregator model capacity

Recall that the aggregator (see Figure 1 and Algorithm 1) synthesizes the information from M concurrent executions. We compare the performance achieved by deploying more capable LM (leader agg., using Sonnet 3.7 + thinking) versus the same LM (self-agg. using Haiku 3.5.). Results are available in Self. Acc. and Leader Acc. columns of Tables 3 and 4 [3]. When the base model is an efficient LM (such as Haiku 3.5.), using a stronger LM for the aggregator node consistently improves performance, across both benchmarks and budget settings. This result is interesting because the majority of agents tends to execute using a single LM while this results points towards the use of combinations of LMs.

In the self-aggregator case, the `base+random` configuration performs competitively. However, in the leader-aggregator case, the advantage of `base+random` vanishes. This shows that aggregator nodes based on weaker models can saturate cognitively when exposed to the volume of information of concurrent executions (Du et al., 2025). Specifically, weaker LMs do not exploit the well the subtle diversity introduced by more advanced state generation and selection and strategies.

[3] We use identical agent execution logs; the only difference is the model used at the aggregator node. This ensures that all upstream agent states are identical, allowing us to isolate specifically the effect of the aggregator node on performance.

Table 3. Results on **InFiDaBench** (N=10). We report self-aggregation and leader-aggregation accuracy (mean ± std), entropy for all plans → selected plans, and average token usage per question (in thousands) across three random seeds.

		M = 3				M = 5			
		Accuracy (%)				Accuracy (%)			
Generation	**Selection**	**self-agg.**	**leader-agg.**	**Entropy**	**Tok In/Out (K)**	**self-agg.**	**leader-agg.**	**Entropy**	**Tok In/Out (K)**
Base	Clustering	74.84 ± 0.22	77.95 ± 0.45	0.069 → 0.083	81.9/4.7	**79.12 ± 0.22**	**82.23 ± 1.62**	0.069 → 0.081	82.1/4.7
	Goal Align	**77.82 ± 0.78**	78.99 ± 0.78	0.069 → 0.069	82.0/4.7	78.47 ± 0.81	80.80 ± 1.57	0.069 → 0.068	82.3/4.7
	Maxent	74.84 ± 0.45	79.25 ± 1.80	0.070 → 0.111	81.9/4.7	77.30 ± 0.98	80.42 ± 1.96	0.069 → 0.090	82.3/4.7
	Random	77.69 ± 3.12	**80.03 ± 1.19**	0.069 → 0.069	81.7/4.7	78.21 ± 1.40	81.45 ± 1.25	0.070 → 0.070	82.4/4.7
Base Average		*76.30 ± 2.06*	*79.05 ± 1.26*	*0.069 → 0.083*	*81.9/4.7*	*78.27 ± 1.06*	*81.23 ± 1.56*	*0.069 → 0.077*	*82.3/4.7*
Conditional	Clustering	74.32 ± 0.00	77.43 ± 1.17	0.069 → 0.082	82.0/4.7	78.60 ± 0.67	**80.80 ± 1.75**	0.070 → 0.082	82.4/4.7
	Goal Align	75.49 ± 1.95	78.86 ± 0.59	0.069 → 0.068	81.9/4.7	**79.38 ± 1.03**	**80.80 ± 1.12**	0.070 → 0.069	82.2/4.7
	Maxent	75.62 ± 0.81	78.47 ± 0.22	0.068 → 0.108	82.3/4.7	77.30 ± 1.47	80.29 ± 0.81	0.070 → 0.091	82.3/4.7
	Random	**76.65 ± 0.78**	**78.99 ± 1.40**	0.068 → 0.069	81.7/4.7	76.78 ± 1.84	79.25 ± 1.25	0.069 → 0.069	82.3/4.7
Conditional Average		*75.52 ± 1.29*	*78.44 ± 1.04*	*0.069 → 0.081*	*82.0/4.7*	*78.02 ± 1.56*	*80.29 ± 1.28*	*0.070 → 0.078*	*82.3/4.7*
Ensemble	Clustering	76.65 ± 1.03	79.77 ± 0.00	0.088 → 0.112	85.1/4.9	77.17 ± 1.80	80.29 ± 1.25	0.088 → 0.102	85.6/4.9
	Goal Align	77.43 ± 2.43	**80.42 ± 1.47**	0.088 → 0.069	85.3/4.9	78.08 ± 1.80	81.06 ± 1.25	0.088 → 0.073	85.7/4.9
	Maxent	**77.69 ± 2.00**	78.34 ± 2.25	0.088 → 0.138	85.2/4.9	76.52 ± 2.00	81.19 ± 1.25	0.089 → 0.110	85.7/4.9
	Random	75.88 ± 2.17	78.86 ± 1.37	0.088 → 0.088	85.1/4.9	**78.86 ± 0.98**	**82.23 ± 1.62**	0.087 → 0.088	85.3/4.9
Ensemble Average		*76.91 ± 1.84*	*79.35 ± 1.53*	*0.088 → 0.102*	*85.2/4.9*	*77.66 ± 1.71*	*81.19 ± 1.36*	*0.088 → 0.093*	*85.6/4.9*

6.2. Influence of Generation and Selection strategies on entropy

`Base` and `Conditional` generation produce nearly identical entropy values (Tables 3 and 4). This finding complements Zhang et al. (2025), who first proposed conditional generation and reported equivalent accuracy between base and conditional strategies. Our entropy measurements complement these finding and show that conditioning on previously generated plans does not meaningfully increase diversity in the candidate pool. In contrast, `Ensemble` generation exhibits the highest entropy, as expected: combining two distinct prompt templates encourages structurally different candidate plans.

Across all generation methods and both benchmarks, `Maxent` consistently produces the highest selected-plan entropy, followed by `Clustering`. This validates that `Maxent`'s entropy-based objective function successfully identifies maximally diverse subsets from the candidate pool, while `Clustering` also promotes diversity by selecting representatives from distinct clusters.

6.3. Optimal test-time scaling designs

To articulate performance, we identify four distinct properties for a test-time scaling approach: i) **Accuracy**: We measure the average performance of the configuration for a budget $M \in \{3, 5\}$. ii) **Reliability**: Given a category defined by a budget M and a generation strategy g, we identify over all the s seeds the worst performance in that category (worst case performance in that category). Then, for a configuration we compute how close its worst seed is to the worst case performance in that category. iii) **Scalability**: We measure the performance improvement/degradation achieved from scaling the selection budget from $M = 3$ to $M = 5$. iv) **Consistency**: Given s seeds, we measure the proportion of questions that are consistently answered right to assess issues with reliability of the answers generated by the configuration.

Based on these criteria, we develop a composite index that weights 50% on Accuracy, and 50% on the other metrics (Reliability 17%, Consistency 17% and Scalability 16%). The detailed breakdown and a global ranking is illustrated in Figure 4. The `Base+Clustering` and `Ensemble+Random` configurations achieves overall the best performance. Given an execution budget $M = 5$, the most optimal choice is `Base+Clustering` both on open-form and closed form questions (we note `CIPHER`*). The baseline strategy `Base+Random` remains a competitive design choice (bronze medal) but other strategies can outperform, highlighting the opportunity of our proposed test-time scaling framework for future research.

7. Related Works

Test-time scaling in LLMs (Brown et al., 2024) shows that multiple executions of the same prompt in signle-turn LM applications improves performance at test-time on closed-ended tasks. (Anonymous, 2025b) proposes a system to leverage parallel execution of LLM prompts with inter-communication to improve performance and share information across threads. (Wang et al., 2025) compare the returns of parallelization versus reasoning on large language models and finds that reasoning tend to outperform parallelization. However, (Liu et al., 2025) show that test-time scaling with the right generation, selection and aggregation strategy can outperform larger LM models, highlighting that

Table 4. Performance results on **InsightBench** (N=10). We report self-aggregation and leader-aggregation accuracy (mean ± std), entropy for all plans → selected plans, and average token usage per question (in thousands) across three random seeds.

		M = 3				M = 5			
		Accuracy (%)				Accuracy (%)			
Generation	**Selection**	**self-agg.**	**leader-agg.**	**Entropy**	**Tok In/Out (K)**	**self-agg.**	**leader-agg.**	**Entropy**	**Tok In/Out (K)**
Base	Clustering	41.74 ± 3.76	42.99 ± 1.07	0.107 → 0.131	95.5/5.8	43.85 ± 0.54	49.22 ± 0.75	0.103 → 0.122	95.8/5.7
	Goal Align	41.36 ± 2.63	42.68 ± 1.37	0.104 → 0.098	95.0/5.7	41.28 ± 0.49	**49.45 ± 1.72**	0.105 → 0.102	96.0/5.7
	Maxent	40.73 ± 1.64	41.82 ± 2.04	0.107 → 0.169	95.1/5.8	42.68 ± 0.94	49.14 ± 1.29	0.106 → 0.137	96.3/5.8
	Random	**42.68 ± 1.35**	**45.09 ± 2.23**	0.104 → 0.107	95.1/5.7	**44.24 ± 0.59**	48.99 ± 3.59	0.104 → 0.103	96.1/5.7
Base Average		*41.63 ± 2.28*	*43.15 ± 1.95*	*0.105 → 0.126*	*95.2/5.7*	*43.01 ± 1.33*	*49.20 ± 1.82*	*0.105 → 0.116*	*96.1/5.7*
Conditional	Clustering	40.81 ± 2.82	43.22 ± 1.07	0.104 → 0.122	95.4/5.7	**44.31 ± 1.29**	**49.92 ± 1.20**	0.106 → 0.124	96.1/5.7
	Goal Align	41.98 ± 0.75	40.34 ± 1.50	0.104 → 0.101	96.2/5.7	42.45 ± 1.75	48.29 ± 1.18	0.104 → 0.101	95.4/5.6
	Maxent	**42.60 ± 0.71**	**43.93 ± 1.42**	0.105 → 0.165	94.7/5.7	43.61 ± 0.82	49.14 ± 1.05	0.103 → 0.132	95.9/5.7
	Random	40.03 ± 3.58	42.60 ± 1.78	0.105 → 0.103	94.7/5.6	43.07 ± 2.17	49.53 ± 1.17	0.104 → 0.105	95.6/5.7
Conditional Average		*41.36 ± 2.25*	*42.52 ± 1.88*	*0.104 → 0.123*	*95.3/5.7*	*43.36 ± 1.54*	*49.22 ± 1.17*	*0.104 → 0.116*	*95.8/5.7*
Ensemble	Clustering	41.59 ± 2.00	**47.90 ± 2.25**	0.118 → 0.151	99.2/6.7	**43.93 ± 0.93**	48.36 ± 0.84	0.119 → 0.140	99.8/6.6
	Goal Align	40.73 ± 0.94	47.59 ± 1.52	0.116 → 0.057	100.8/6.7	43.46 ± 2.00	49.53 ± 0.23	0.116 → 0.073	101.7/6.7
	Maxent	**42.29 ± 0.81**	46.96 ± 0.84	0.117 → 0.182	99.1/6.7	42.06 ± 1.87	**50.39 ± 1.87**	0.116 → 0.145	100.3/6.7
	Random	41.20 ± 0.97	**47.90 ± 0.81**	0.117 → 0.122	100.8/6.8	43.69 ± 3.31	49.77 ± 1.82	0.116 → 0.113	100.1/6.7
Ensemble Average		*41.45 ± 1.24*	*47.59 ± 1.32*	*0.117 → 0.128*	*100.0/6.7*	*43.28 ± 2.02*	*49.51 ± 1.40*	*0.117 → 0.118*	*100.5/6.7*

parallelization is an avenue for cost efficient designs. (Inoue et al., 2025) highlight that adaptive strategies that emulate diverse solution paths improve performance, and propose a method based on exploration-exploitation.

Test-time scaling in multi-agent systems Although test-time scaling properties have been investigated in the realm of single-turn LM applications, its impact on agentic designs remains an emerging avenue of research. (Anonymous, 2025a) explores how parallelization can deliver latency speed-ups when sectioning problem into independent parts. (Zhu et al., 2025) operate concurrent executions using diverse LLMs. (Zhang et al., 2025) proposes a seminal framework for test-time scaling but there is not clear distinction between generation and selection. Furthermore, many questions remain open regarding the amount of agents to parallelize, and the strategy to aggregate results.

8. Conclusion and avenues of future research

The proposed framework opens the door to algorithmic innovations to improve upon the standard test-time approaches (`Base` generation + `Random` selection). Although this baseline is competitive, we show that it can be outperformed by strategies based on `Clustering` and prompt `Ensembles`. The results further highlight the role of the aggregator node as a support to the entire framework. Lastly, the proposed agent `CIPHER` performs competitively against state-of-the-art baselines on two open-source frameworks, demonstrating the relevance of test-time scaling to improve the performance of AI agents. Future research avenues include the scaling of ensemble generation to more prompts, and the selection of candidate states using gradient based methods (Mirzasoleiman et al., 2020).

		M=3		M=5		Global		
ID	**Strategy**	**Infida**	**Insight**	**Infida**	**Insight**	**@M=3**	**@M=5**	**Overall**
BC	Base + Clustering	6	5	**1**	**1**	6	**1**	3.2
ER	Ensemble + Random	9	**1**	2	7	5	3	4.8
BR	Base + Random	3	6	6	8	4	7	5.8
CC	Conditional + Clustering	10	7	4	2	9	2	5.8
EM	Ensemble + Maxent	5	2	5	11	3	10	5.8
EG	Ensemble + Goal Align	**1**	4	9	10	1	11	6.0
EC	Ensemble + Clustering	2	3	11	12	1	12	7.0
BG	Base + Goal Align	4	12	7	6	7	5	7.2
BM	Base + Maxent	7	9	10	4	7	7	7.5
CG	Conditional + Goal Align	8	10	3	9	10	4	7.5
CM	Conditional + Maxent	11	8	8	5	11	5	8.0
CR	Conditional + Random	12	11	12	3	12	9	9.5

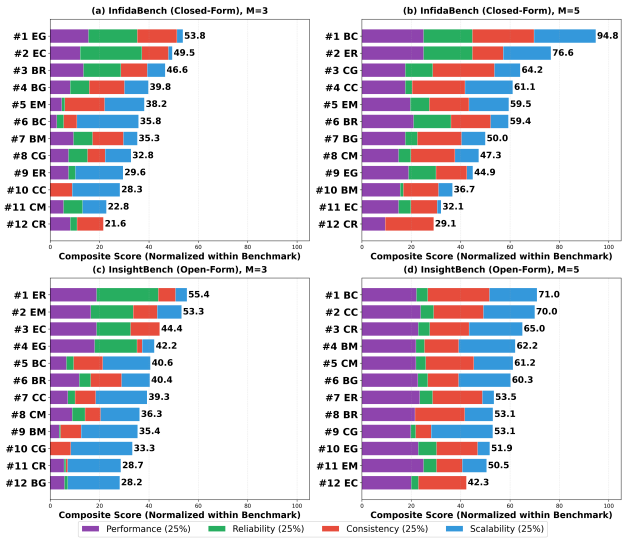

Figure 4. Table: Summary rankings. Panel: Composite scores detail (scores normalized per benchmark).

References

Anonymous. Flash-searcher: Fast and effective web agents via DAG-based parallel execution. In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025a. URL https://openreview.net/forum?id=QuaJ6kJaBm. under review.

Anonymous. Generalized parallel scaling with interdependent generations. In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025b. URL https://openreview.net/forum?id=suU6kAP6c2. under review.

Anonymous. Think in parallel, answer as one: Logit averaging for open-ended reasoning. In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025c. URL https://openreview.net/forum?id=hvit36Dyzl. under review.

AWS. Titan text embeddings v2. https://huggingface.co/amazon/Titan-text-embeddings-v2, 2024. Accessed: 2025-11-18.

Brown, B., Juravsky, J., Ehrlich, R., Clark, R., Le, Q. V., Ré, C., and Mirhoseini, A. Large language monkeys: Scaling inference compute with repeated sampling, 2024. URL https://arxiv.org/abs/2407.21787.

Cao, L. Data science: a comprehensive overview. *ACM Computing Surveys (CSUR)*, 50(3):1–42, 2017.

Chen, X., Lin, M., Schärli, N., and Zhou, D. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*, 2023.

Dietterich, T. G. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pp. 1–15. Springer, 2000.

Du, Y., Tian, M., Ronanki, S., Rongali, S., Bodapati, S., Galstyan, A., Wells, A., Schwartz, R., Huerta, E. A., and Peng, H. Context length alone hurts llm performance despite perfect retrieval, 2025. URL https://arxiv.org/abs/2510.05381.

Freund, Y. and Schapire, R. E. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.

Heuillet, M., Debaque, B., and Durand, A. Sequential automated machine learning: Bandits-driven exploration using a collaborative filtering representation. In *8th ICML Workshop on Automated Machine Learning (AutoML)*, 2021. URL https://openreview.net/forum?id=6tlvEH9HaX.

Hollmann, N., Müller, S., and Hutter, F. Large language models for automated data science: Introducing caafe for context-aware automated feature engineering. *Advances in Neural Information Processing Systems*, 36:44753–44775, 2023.

Hong, S., Lin, Y., Liu, B., Liu, B., Wu, B., Zhang, C., Wei, C., Li, D., Chen, J., Zhang, J., Wang, J., Zhang, L., Zhang, L., Yang, M., Zhuge, M., Guo, T., Zhou, T., Tao, W., Tang, X., Lu, X., Zheng, X., Liang, X., Fei, Y., Cheng, Y., Gou, Z., Xu, Z., and Wu, C. Data interpreter: An llm agent for data science, 2024. URL https://arxiv.org/abs/2402.18679.

Hu, X., Zhao, Z., Wei, S., Chai, Z., Ma, Q., Wang, G., Wang, X., Su, J., Xu, J., Zhu, M., et al. Infiagent-dabench: Evaluating agents on data analysis tasks. *arXiv preprint arXiv:2401.05507*, 2024.

il Lee, K., Koh, H., Lee, D., Yoon, S., Kim, M., and Jung, K. Generating diverse hypotheses for inductive reasoning, 2025. URL https://arxiv.org/abs/2412.13422.

Inoue, Y., Misaki, K., Imajuku, Y., Kuroki, S., Nakamura, T., and Akiba, T. Wider or deeper? scaling llm inference-time compute with adaptive branching tree search, 2025. URL https://arxiv.org/abs/2503.04412.

Kanter, J. M. and Veeramachaneni, K. Deep feature synthesis: Towards automating data science endeavors. In *2015 IEEE international conference on data science and advanced analytics (DSAA)*, pp. 1–10. IEEE, 2015.

LangChain AI. Langgraph: Build language agents as graphs, 2024. URL https://github.com/langchain-ai/langgraph. Python package version 0.x.x.

Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2020.

Liu, R., Gao, J., Zhao, J., Zhang, K., Li, X., Qi, B., Ouyang, W., and Zhou, B. Can 1b llm surpass 405b llm? rethinking compute-optimal test-time scaling, 2025. URL https://arxiv.org/abs/2502.06703.

MacQueen, J. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pp. 281–297. University of California Press, 1967.

Mirzasoleiman, B., Bilmes, J., and Leskovec, J. Coresets for data-efficient training of machine learning models. In *International Conference on Machine Learning*, pp. 6950–6960. PMLR, 2020.

Nam, J., Yoon, J., Chen, J., and Pfister, T. Ds-star: Data science agent via iterative planning and verification, 2025. URL https://arxiv.org/abs/2509.21825.

Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.

Olea, C., Tucker, H., Phelan, J., Pattison, C., Zhang, S., Lieb, M., Schmidt, D., and White, J. Evaluating persona prompting for question answering tasks. In *Proceedings of th e 10th international conference on artificial intelligence and soft computing, Sydney, Australia*, 2024.

Olson, R. S., Bartley, N., Urbanowicz, R. J., and Moore, J. H. Evaluation of a tree-based pipeline optimization tool for automating data science. In *Proceedings of the genetic and evolutionary computation conference 2016*, pp. 485–492, 2016.

Reid, A., O'Callaghan, S., Carroll, L., and Caetano, T. Risk analysis techniques for governed llm-based multi-agent systems, 2025. URL https://arxiv.org/abs/2508.05687.

Sahu, G., Puri, A., Rodriguez, J., Abaskohi, A., Chegini, M., Drouin, A., Taslakian, P., Zantedeschi, V., Lacoste, A., Vazquez, D., Chapados, N., Pal, C., Mudumba, S. R., and Laradji, I. H. Insightbench: Evaluating business analytics agents through multi-step insight generation, 2025. URL https://arxiv.org/abs/2407.06423.

Schluntz, E. and Zhang, B. Building effective AI agents. https://www.anthropic.com/engineering/building-effective-agents, December 2024.

Snell, C., Lee, J., Xu, K., and Kumar, A. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024. URL https://arxiv.org/abs/2408.03314.

Wang, J., Zhu, S., Saad-Falcon, J., Athiwaratkun, B., Wu, Q., Wang, J., Song, S. L., Zhang, C., Dhingra, B., and Zou, J. Think deep, think fast: Investigating efficiency of verifier-free inference-time-scaling methods, 2025. URL https://arxiv.org/abs/2504.14047.

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., et al. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345, 2024.

Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., et al. Autogen: Enabling next-gen llm applications via multi-agent conversations. In *First Conference on Language Modeling*, 2024.

You, Z., Zhang, Y., Xu, D., Lou, Y., Yan, Y., Wang, W., Zhang, H., and Huang, Y. Datawiseagent: A notebook-centric llm agent framework for automated data science, 2025. URL https://arxiv.org/abs/2503.07044.

Zhang, E., Zhu, E., Bansal, G., Fourney, A., Mozannar, H., and Gerrits, J. Optimizing sequential multi-step tasks with parallel llm agents, 2025. URL https://arxiv.org/abs/2507.08944.

Zhang, Y., Chen, J., Wang, J., Liu, Y., Yang, C., Shi, C., Zhu, X., Lin, Z., Wan, H., Yang, Y., Sakai, T., Feng, T., and Yamana, H. Toolbehonest: A multi-level hallucination diagnostic benchmark for tool-augmented large language models, 2024. URL https://arxiv.org/abs/2406.20015.

Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36: 46595–46623, 2023.

Zhu, K., Li, H., Wu, S., Xing, T., Ma, D., Tang, X., Liu, M., Yang, J., Liu, J., Jiang, Y. E., Zhang, C., Lin, C., Wang, J., Zhang, G., and Zhou, W. Scaling test-time compute for llm agents, 2025. URL https://arxiv.org/abs/2506.12928.

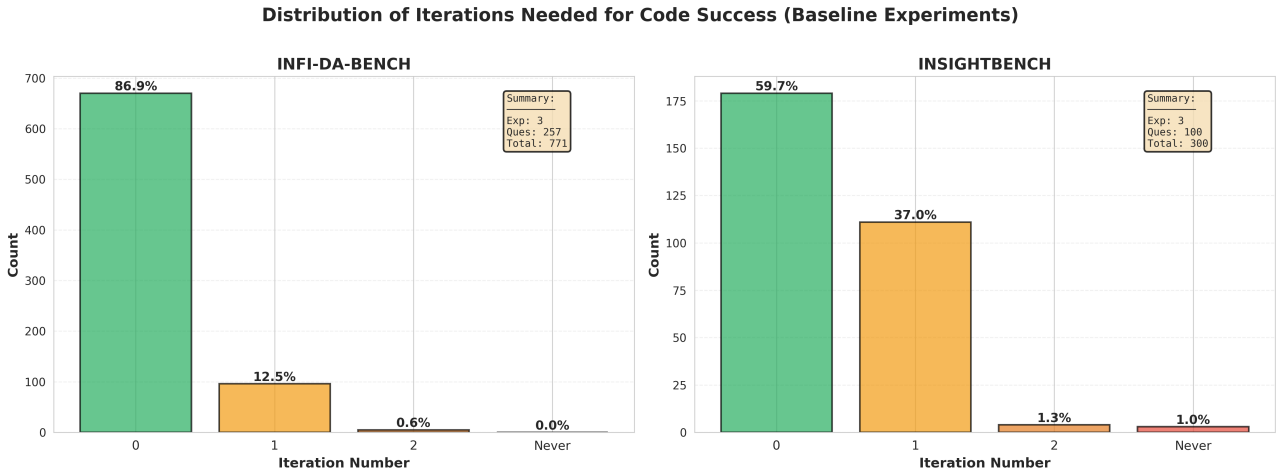

Figure 5. Distribution of iterations needed to reach executable code across CIPHER(1,1) baseline experiments on three random seeds for Infi-DA-Bench (left) and Insight-bench (right).

A. Details on CIPHER's coding loop

B. Details on the ensemble generation strategy

The choice of the optimal E parameter is linked to the parameter M. For a task, if M states are executed, the set of prompts E is more likely to be fully covered when $E < M$. Since in our experiments we focus on $M = 3$ and $M = 5$ regimes, we set $E = 2$ and $p = [0.5, 0.5]$. The two prompts are: prompt 1) is the base prompt used in CIPHER (1,1) the base agent, 2) the second prompt is a paraphrase obtained from a LLM. The key difference is that the paraphrase in prompt no. 2 imposes a fixed structure to the plan generated, while prompt no. 1 does not.

C. Execution examples

D. Additional results

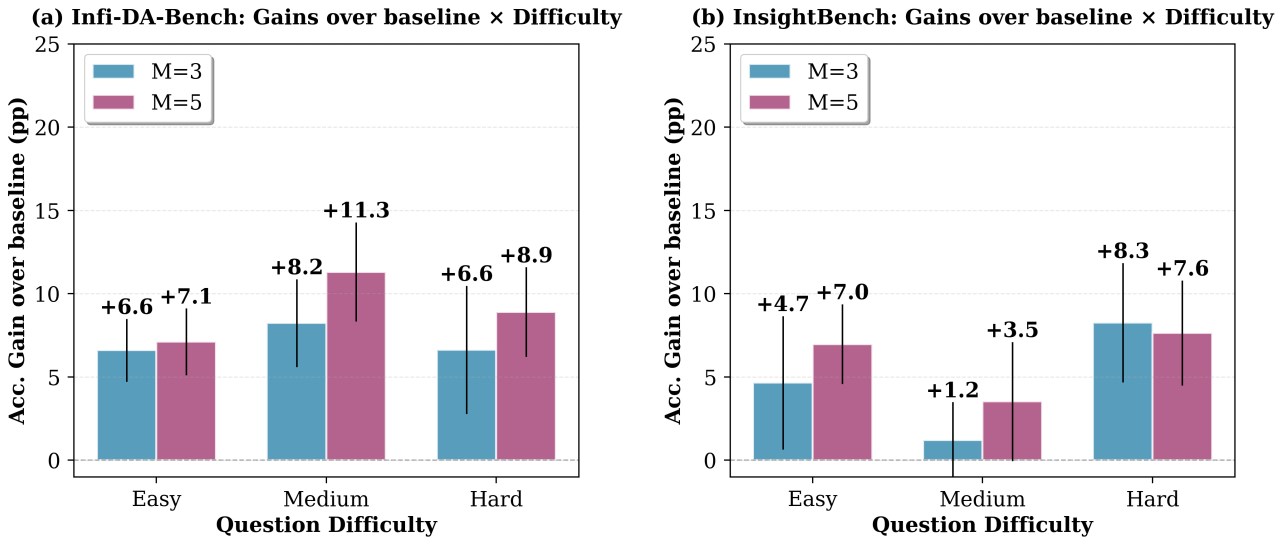

Figure 8. Performance by question difficulty.

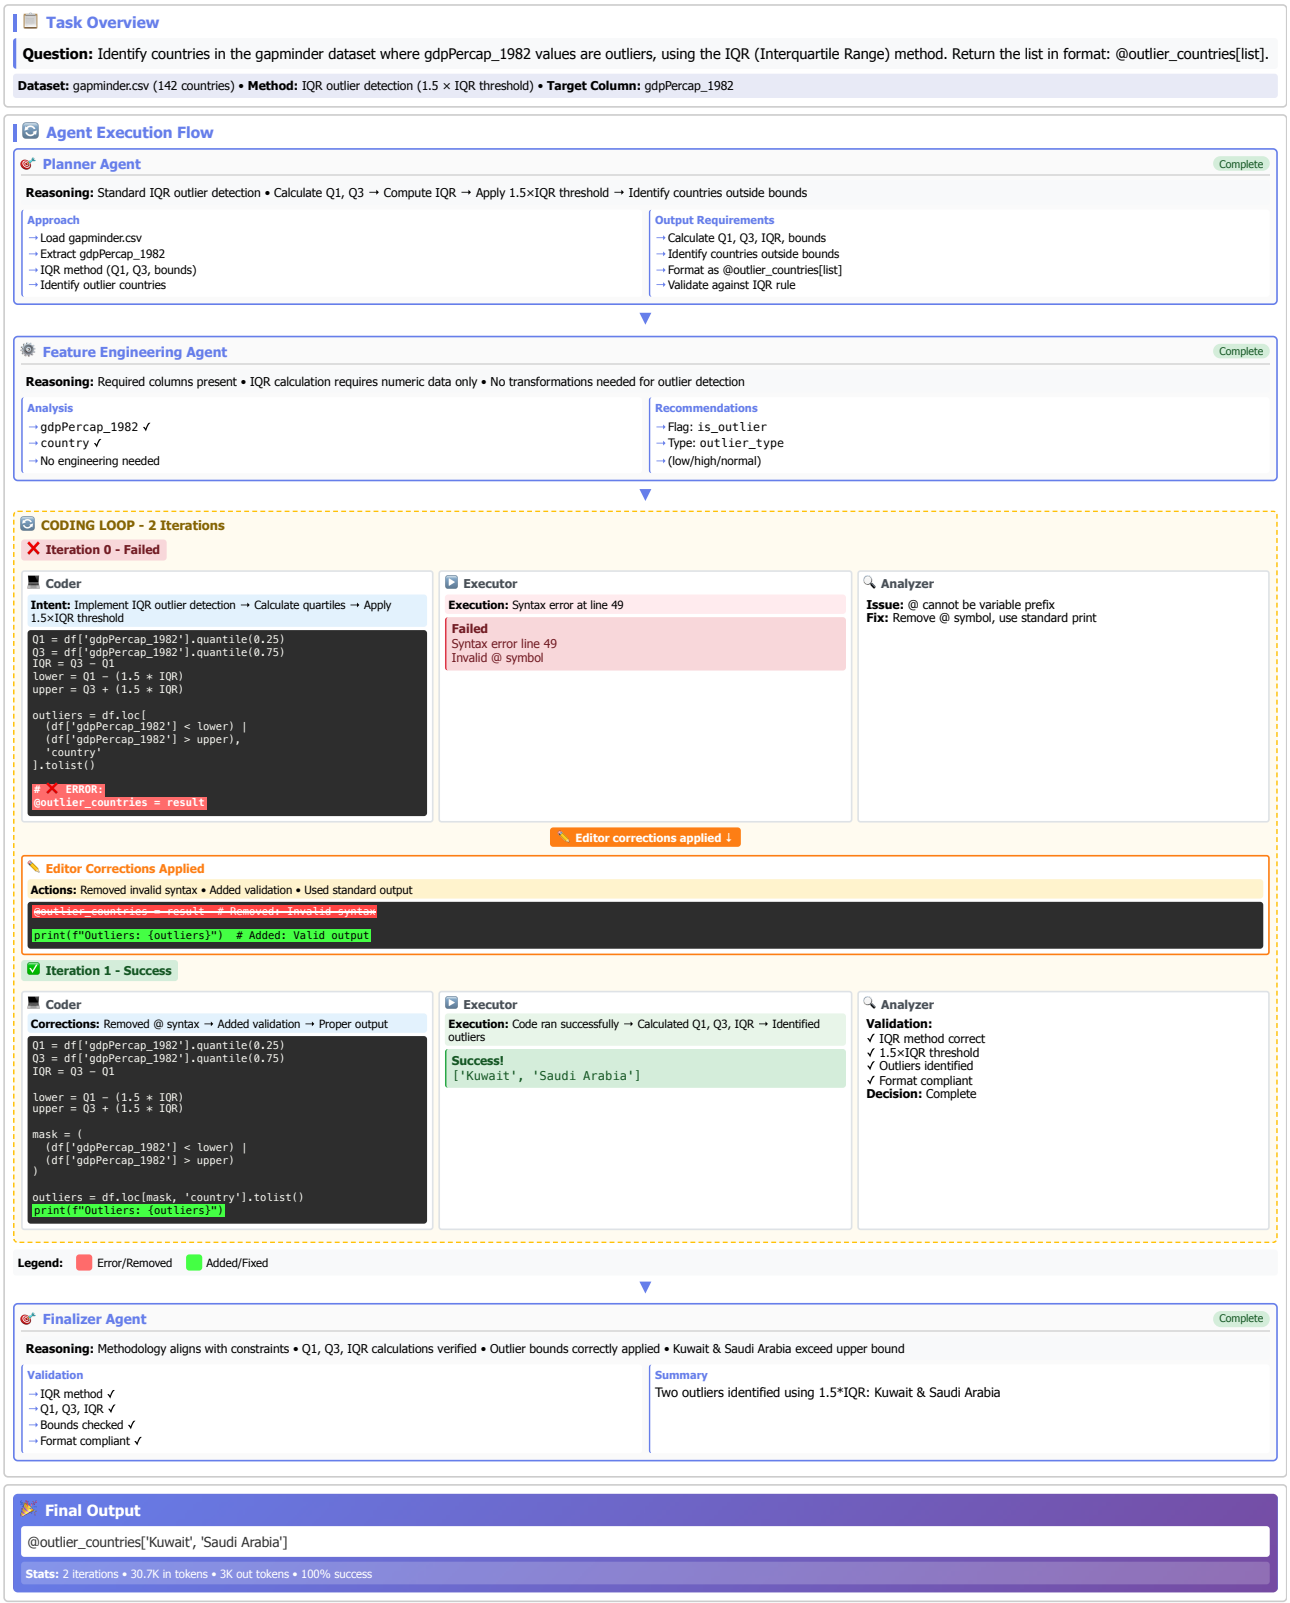

Figure 6. Example execution of CIPHER on question 254 of InfiDABench.

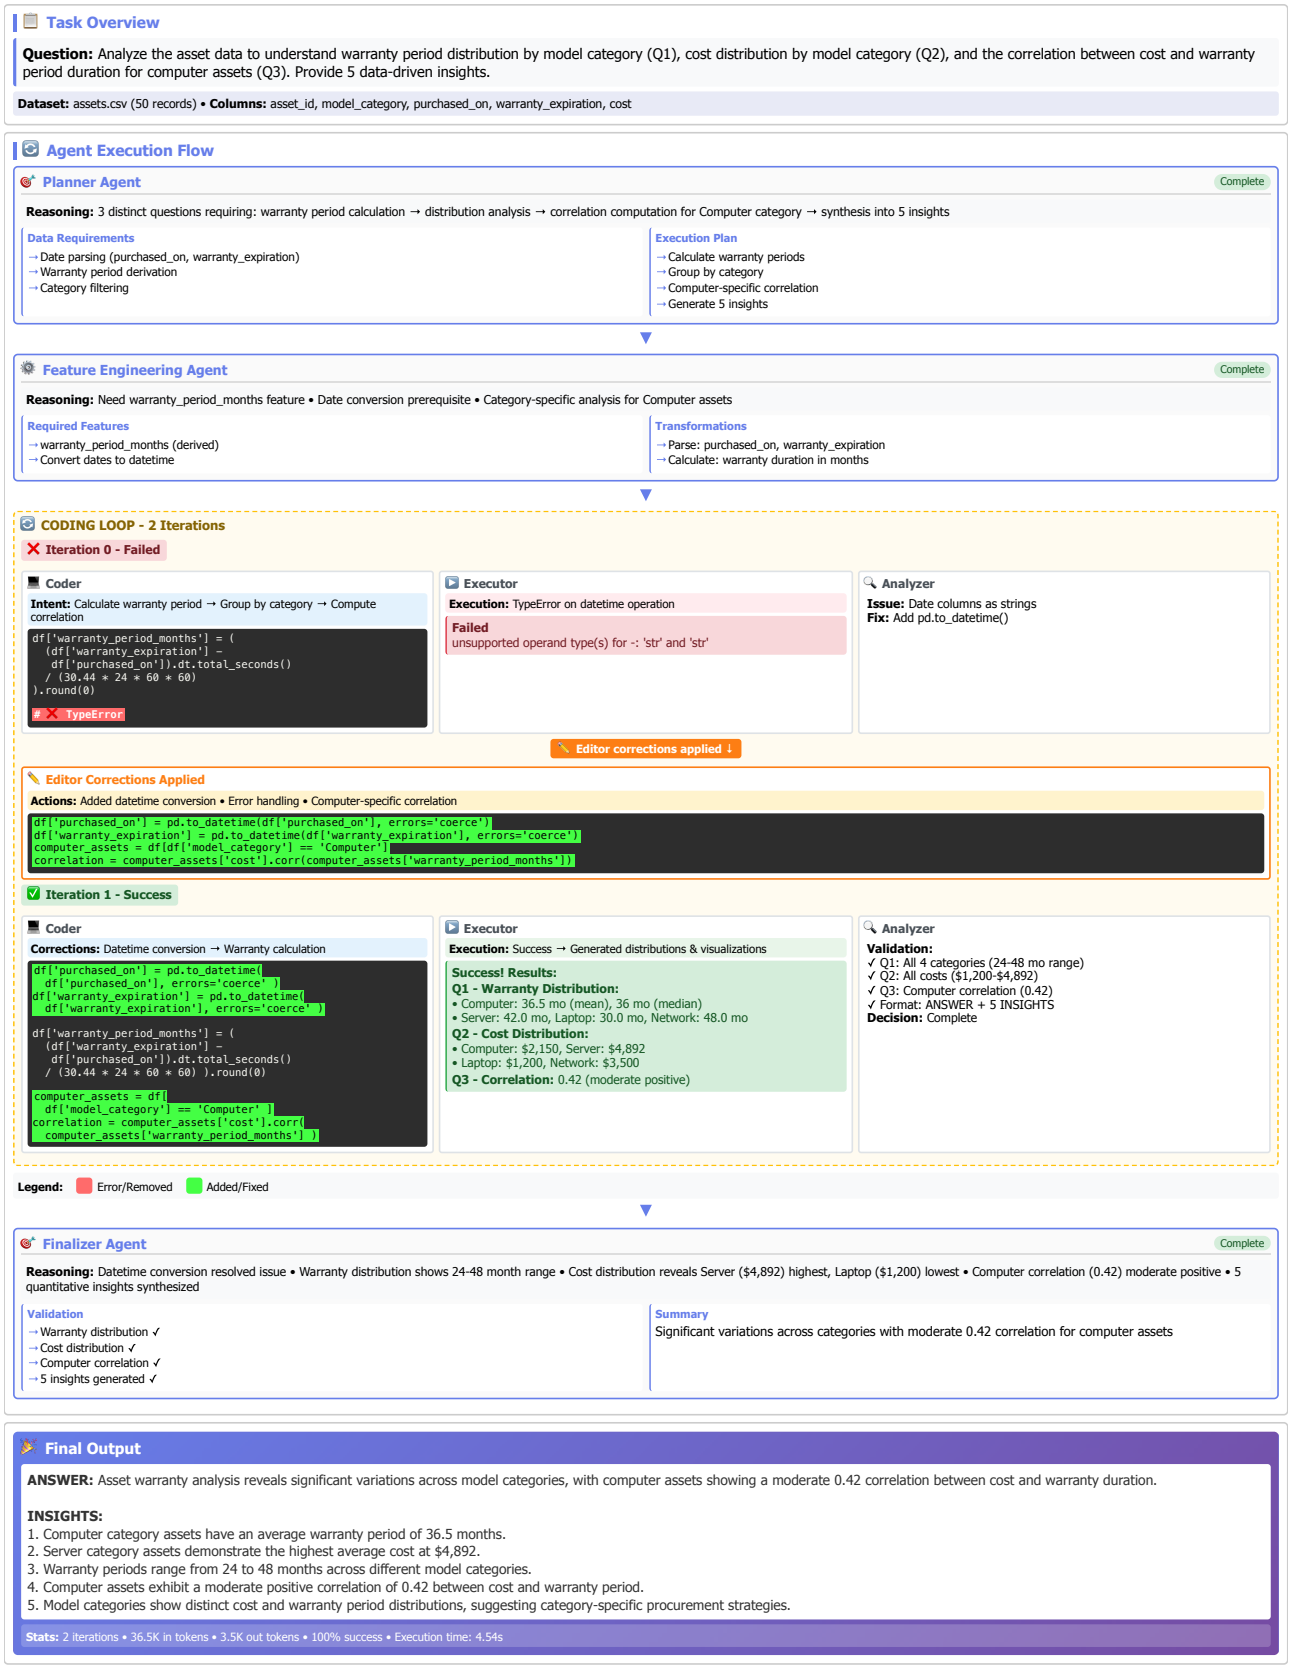

Figure 7. Example execution of CIPHER on question 16 of Insightbench.

Different gain patterns across question types From Figure 8, we observe different gains across task types. On Insight-bench, test-time scaling benefits most medium and hard tasks, while on Insight-bench it benefits most easy and hard tasks. On Insight-bench, the classification on difficulty (easy, medium, hard) is based on the dataset that accompanies the tasks rather than the tasks itself – this can explain the counter-intuitive U-shaped distribution.