ISSN: 2754-6659

Journal of Artificial Intelligence & Cloud Computing

SCIENTIFIC Research and Community

Perspective Article

Open Access

Plausible But Wrong: A Case Study on Agentic Failures in Astrophysical Workflows

Shivam Rawat* and Lucie Flek

Bonn-Aachen International Center for Information Technology, University of Bonn, Germany, Lamarr Institute for Machine, Learning and Artificial Intelligence, Germany

ABSTRACT

Agentic AI systems are increasingly being integrated into scientific workflows, yet their behavior under realistic conditions remains insufficiently understood. We evaluate CMB Agent across two workflow paradigms and eighteen astrophysical tasks. In the One-Shot setting, access to domain specific context yields an approximately ∼6× performance improvement (0.85 vs. ≈0 without context), with the primary failure mode being silent incorrect computation—syntactically valid code that produces plausible but inaccurate results. In the Deep Research setting, the system frequently exhibits silent failures across stress tests, producing physically inconsistent posteriors without self-diagnosis. Overall, performance is strong on well specified tasks but degrades on problems designed to probe reasoning limits, often without visible error signals. These findings highlight that the most concerning failure mode in agentic scientific workflows is not overt failure, but confident generation of incorrect results. We release our evaluation framework to facilitate systematic reliability analysis of scientific AI agents.

***Corresponding author**

Shivam Rawat, Bonn-Aachen International Center for Information Technology, University of Bonn, Germany, Lamarr Institute for Machine, Learning and Artificial Intelligence, Germany.

Received: October 05, 2026; **Accepted:** October 09, 2026; **Published:** October 14, 2026

Introduction

Recent advances in large language models (LLMs) have enabled the emergence of agentic AI systems capable of multi-step reasoning, external tool use, and autonomous decision-making [1-6]. These systems are increasingly proposed as assistants for scientific work-flows, where they can perform data analysis, hypothesis generation, and pipeline orchestration with limited human intervention.

Despite promising demonstrations, existing work largely evaluates such systems through task-level success or syn-thetic benchmarks, offering limited insight into their reli-bility in realistic scientific settings. Scientific workflows differ fundamentally from standard AI benchmarks: they involve structured pipelines, domain-specific constraints, noisy data, and multi-stage reasoning where early errors can propagate throughout the pipeline. As a result, aggregate success metrics alone are insufficient to characterise system Behaviour in these environments.

Importantly, our focus is not on proposing or critiquing agentic architectures themselves, but on evaluating their scientific reliability and failure modes when deployed in realistic workflows.

A central challenge is that agentic systems inherit known limitations of LLMs, including hallucination, brittle reasoning, and misalignment with task objectives. When embedded in multi-step workflows with external tool interaction, these issues can lead not only to incorrect intermediate steps but also to silent numerical errors and physically inconsistent outputs. Critically, such failures may not be detectable from final outputs alone, making systematic evaluation of reliability and error modes essential for scientific deployment.

In this work, we present a structured empirical evaluation of an existing agentic system applied to astrophysical inference tasks. Rather than introducing a new model or framework, our goal is to characterise performance boundaries, assess robustness and calibration, and identify systematic failure modes in realistic scientific workflows.

Astrophysics provides a particularly suitable testbed due to its well-defined physical models, established numerical tools, and availability of reference solutions from both simulation and observation. This allows controlled evaluation of correctness while maintaining real-world scientific relevance.

We introduce a structured evaluation framework for agentic scientific workflows that integrates execution success, parameter accuracy, and numerical fidelity. Using this frame-work, we conduct a systematic analysis of reliability across both single-step and multi-step reasoning settings, identifying key failure modes such as silent numerical errors and physically inconsistent inference, and quantifying their prevalence under different task conditions.

Related Work

LLM Agents and Tool Use

The ability of LLMs to inter-act with external tools has been a key driver in the transition from static text generation to dynamic agentic Behaviour. Re Act established the paradigm of interleaving

reasoning traces with action execution, enabling models to plan and adapt based on environmental feed-back [4]. Tool former extended this by training models to invoke APIs autonomously through self-supervision, while Web GPT demonstrated that grounding language models in live retrieval substantially improves factual reliability [5,6]. More recent work has formalized tool-use evaluation, with benchmarks such as Tool Bench [7]. and TRAJECT-Bench revealing systematic failure modes including similar-tool confusion, parameter-blind selection, and performance degradation with trajectory length [8].

Together these works establish that structured tool access, rather than para-metric knowledge alone, is what enables LLMs to operate effectively in multi-step workflows- a finding our results reinforce in the context of scientific computation.

LLM Agents in Scientific Workflows

A growing number of works demonstrate that agentic systems can execute substantial portions of real scientific workflows, spanning hypothesis generation, data analysis, and manuscript production [9-13].

While these works establish the feasibility of agentic scientific automation, evaluation is typically limited to end-to-end task success or qualitative expert assessment. Systematic analysis of reliability, failure modes, and error propagation in multi-step scientific pipelines remains largely unexplored.

Reliability and Evaluation of LLMs

A substantial body of work examines the limitations of LLMs, including hallucination factual inconsistency and misalignment with user intent Critically, recent surveys show that hallucinations in agentic settings differ qualitatively from single-step LLM failures: they emerge through multi-step sequential interactions involving tool use, memory, and inter-agent communication, making them substantially harder to detect from final outputs alone [14,16]. More recent work has begun to formalise agent evaluation, distinguishing between tool-use capability, trajectory-level reasoning, and domain-specific performance [17,19]. However, systematic evaluation of reliability under realistic multi-step scientific workflows remains limited.

Positioning of This Work

In contrast to prior work, we focus explicitly on reliability and failure Characterisation rather than task feasibility. Rather than proposing a new model or architecture, we provide a structured empirical evaluation of an existing deployed system under controlled conditions such as under-constrained inference and compositional reasoning pipelines. This enables us to identify not only whether agentic systems succeed, but how and why they fail, and crucially, whether those failures are visible in the final output or silent.

Agentic System Under Consideration

For our evaluation, we utilize CMB Agent, an open-source agentic framework for autonomous scientific discovery [19]. CMB Agent serves as the research backend in the Denario project generating research ideas, methodologies, and results through coordinated multi-agent planning [11]. The outputs—including analyses, visualizations, and literature references-are automatically compiled into publication-ready manuscripts. While the Denario project demonstrates the feasibility of end-to-end autonomous scientific workflows and evaluates overall task success through expert assessment, a systematic investigation of reliability, robustness, error propagation, and failure modes remains limited. Such analysis is essential for assessing the suitability of agentic systems in scientific deployment settings where reproducibility and numerical correctness are critical. Here we complement existing feasibility studies by conducting a structured empirical evaluation focused on robustness, calibration, and failure characterization, thereby contributing toward a more rigorous evaluation framework for autonomous scientific agents.

The framework provides multiple operational modes and workflows, which differ primarily in the number of active agents, session length, context persistence across tasks, and the extent of human intervention required for conducting scientific analyses.

These are as follows:

One Shot Executes a task in a single reasoning and execution pass without iterative planning or review. We instantiate two variants: one with a context-aware domain agent (cambagent) providing task-relevant grounding, and one without, isolating the contribution of domain-specific retrieval augmentation relative to a base LLM baseline.

Deep Research Activates the full Planning & Control architecture, decomposing tasks into substeps executed by specialised agents with critique and retry modules. Designed for complex, multi-step analyses requiring multilayered reasoning and iterative error correction.

We focus exclusively on these two fully automated work-flows; the Human in the Loop mode is excluded as user feedback introduces variability that complicates reproducible benchmarking [17,18].

Evaluation and Metrics

We evaluate the agentic system across two distinct categories of tasks, aligned with the two workflow paradigms under investigation. This separation is intentional and reflects the differing operational objectives of the workflows. Tool-driven computational tasks are employed to evaluate the One-Shot workflow, as they emphasize single-pass reasoning, precise tool invocation, and accurate parameter con-figuration. These tasks allow controlled, fine-grained assessment of execution correctness and numerical reliability without introducing the additional complexity of multi-stage planning.

In contrast, the Deep Research workflow is specifically de-signed to support iterative reasoning, hierarchical task de-composition, and sustained context retention across multiple steps. Evaluating this workflow using simple, single-call computational tasks would not meaningfully exercise its planning and control mechanisms. Therefore, we assign more complex, research-oriented tasks to the Deep Research workflow—tasks that require multi-step reasoning, contextual carryover, and strategy refinement—so that its architectural capabilities can be appropriately tested. This task–workflow alignment ensures that each paradigm is evaluated under conditions that reflect its intended operational design.

Our evaluation design follows recent taxonomies of LLM-agents which distinguish between capability-centric evaluation (e.g., tool use and function correctness), trajectory-level evaluation (e.g., planning and multi-step reasoning), and domain-specific application evaluation [17]. The CAMB-based tasks primarily assess tool in-vocation accuracy and computational reliability, whereas the research-driven tasks stress hierarchical planning, long-horizon reasoning, and statistical inference consistency in scientific workflows.

multi-API calls with combined results; Tasks 11–12 add a delensing pipeline; Task 13 adds ratio computation; and Task 14 combines multi-API calls with a noise-informed delensing pipeline.

Astrophysical Research-Driven Tasks

Four research-grade inference problems spanning cosmology, galactic dynamics, exoplanet structure, and strong gravitational lensing probe multi-step reasoning, statistical robustness, and resistance to hallucinated physical assumptions via Bayesian parameter estimation and hierarchical modeling. Task descriptions and category rationales are given in Section 4.3; prompts are in the Appendix.

Evaluation Metrics

We evaluate each workflow under a distinct scoring frame-work reflecting its operational objectives. The One-Shot workflow admits fully automated quantitative evaluation against reference outputs; the Deep Research workflow re-quires qualitative assessment against published literature values. All metrics are scored in {0, 1} and averaged across trials for reporting.

One-Shot Workflow Metrics

Execution Success Rate (ESR)

ESR is a binary indicator of whether the agent produced valid executable output. A trial is assigned ESR = 1 if and only if the agent generates a numerical output file that (i) contains at least two numeric columns, (ii) covers at least 95% of the reference x-range, and (iii) provides at least 95% of the reference number of output points. Otherwise ESR = 0. This criterion deliberately does not assess numerical correctness—it only verifies that the pipeline completed and produced an output of the right shape.

Parameter Accuracy Score (PAS)

PAS measures whether the agent configured the CAMB cosmological solver with the correct input parameters. Parameters are extracted directly from the generated code via abstract syntax tree (AST) parsing, without executing the code, making the metric robust to runtime failures. For each parameter *p* present in the reference solution, the relative error is

$$\epsilon_p = \min\left(\frac{|\hat{\theta}_p - \theta_p^*|}{|\theta_p^*| + \delta},\ 1\right), \qquad (1)$$

where θ_p^* is the reference value, $\hat{\theta}_p$ is the agent's value, and $\delta = 10^{-12}$ prevents division by zero. If the agent omits a required parameter, $\epsilon_p = 1$. The overall PAS is a weighted mean over all reference parameters:

$$\mathrm{PAS} = 1 - \frac{\sum_p w_p\, \epsilon_p}{\sum_p w_p}, \qquad (2)$$

where weights w_p reflect the physical importance of each parameter: $w = 2.0$ for $\{H_0, \Omega_b h^2, \Omega_c h^2, n_s, A_s\}$, $w = 1.5$ for $\{\Omega_k, w_0\}$, and $w = 1.0$ for $\{\tau, m_\nu, A_{\mathrm{lens}}\}$. Weights reflect the breadth of spectral impact: core ΛCDM parameters (H_0, $\omega_b h^2$, $\omega_c h^2$, n_s, A_s; $w = 2.0$) distort amplitude, shape, and peak positions globally; extension parameters (Ω_k, w_0; $w = 1.5$) are task-critical only when non-zero; and secondary parameters (τ, m_ν, A_{lens}; $w = 1.0$) have localised effects. This scheme under penalises extension parameter errors on tasks where they are non-trivial, making reported PAS a slight overestimate in those cases.

Numerical Accuracy Score (NAS): NAS measures how closely the agent's output curve matches the reference solution over the full output domain. The agent's curve is interpolated onto the reference grid and three complementary sub-metrics are computed.

NRMSE score measures amplitude accuracy:

$$S_{\mathrm{NRMSE}} = \max\left(0,\ 1 - \frac{\sqrt{\frac{1}{N}\sum_i (\hat{y}_i - y_i^*)^2}}{\max(y^*) - \min(y^*)}\right). \qquad (3)$$

SMAPE score measures symmetric percentage deviation, robust to scale differences:

$$S_{\mathrm{SMAPE}} = \max\left(0,\ 1 - \frac{1}{N}\sum_i \frac{|\hat{y}_i - y_i^*|}{(|\hat{y}_i| + |y_i^*|)/2}\right). \qquad (4)$$

CCC score is Lin's concordance correlation coefficient (Lin, 1989), which jointly penalises location shift, scale shift, and correlation:

$$S_{\mathrm{CCC}} = \max\left(0,\ \frac{2\,\mathrm{Cov}(\hat{y}, y^*)}{\mathrm{Var}(\hat{y}) + \mathrm{Var}(y^*) + (\bar{\hat{y}} - \bar{y}^*)^2}\right). \qquad (5)$$

The sub-scores are combined as:

$$\mathrm{NAS} = 0.2\, S_{\mathrm{NRMSE}} + 0.3\, S_{\mathrm{SMAPE}} + 0.5\, S_{\mathrm{CCC}}. \qquad (6)$$

CCC receives the highest weight because it simultaneously captures amplitude, scale, and shape agreement; NRMSE is down-weighted because it is sensitive to outliers and can be dominated by a single high-amplitude peak.

Final Score

$$\mathrm{Score} = \begin{cases} 0 & \text{if ESR} = 0, \\ \mathrm{PAS} \times \mathrm{NAS} & \text{otherwise.} \end{cases} \qquad (7)$$

The multiplicative form ensures that a trial scores zero if either parameters or numerical output are entirely wrong, even if the other dimension is perfect.

Failure Mode Taxonomy: Each trial is assigned to one of four mutually exclusive failure modes:

A Code Failure (ESR = 0): the pipeline did not produce valid output.

B Wrong Parameters (ESR = 1, PAS < 0.5): the code ran but used incorrect CAMB configuration.

C Wrong Computation (ESR = 1, PAS ≥ 0.5, NAS < 0.5): parameters were correct but numerical output was inaccurate, indicating a formula, unit, or postprocessing error.

D Correct (ESR = 1, PAS ≥ 0.5, NAS ≥ 0.5): both configuration and numerical output are consistent with the reference.

We additionally flag *unit/normalisation errors* as a sub-class of Mode C: trials where $S_{\mathrm{CCC}} > 0.8$ but $S_{\mathrm{NRMSE}} < 0.7$ indicate correct spectral shape with wrong amplitude, consistent with a missing normalisation factor or unit conversion error.

Deep Research Workflow Metrics

The Deep Research workflow produces no automated reference CSV, so ESR, PAS, and NAS are not applicable. Each task is instead evaluated across two dimensions: a quantitative parameter recovery score and a qualitative physical consistency assessment.

Parameter Recovery Score (PRS): Measures agreement between the agent's reported posterior median and the literature reference for each key parameter. For parameter p with reference $\theta_p^* \pm \sigma_p^*$, the per-parameter score is:

$$r_p = \max\left(0,\ 1 - \frac{|\hat{\theta}_p - \theta_p^*|}{3\,\sigma_p^*}\right), \qquad (8)$$

saturating to zero at 3σ deviation. PRS is the unweighted mean of r_p over all task parameters.

Qualitative Assessment: Beyond parameter recovery, each task is assessed along two dimensions reported as free-text observations in the results: (i) physical plausibility —whether reported parameter values and posterior shapes are consistent with established physical expectations for the system; (ii) failure transparency — whether the agent identifies and reports known degeneracies, structural biases, or convergence issues, rather than delivering results silently — reported alongside PRS in the results.

Task-Level Robustness Evaluation

To evaluate reliability beyond aggregate performance, we design the research-driven task suite as a structured stress test of the Deep Research workflow. Rather than ablating architectural components, we ablate the task conditions themselves—varying the degree of prior specification, likelihood completeness, and model complexity—while keeping the workflow fixed. This approach is motivated by prior observations that agentic failures in scientific settings emerge not from random errors but from systematic weaknesses under specific inference regimes: two structurally distinct under-constrained posteriors and compositional reasoning chains [17,18].

We evaluate the One-Shot workflow through a controlled architectural ablation: access to domain-specific CAMB documentation is removed in the *CMBAgent* (*no context*) variant, isolating the contribution of structured retrieval aug-mentation on tool invocation accuracy and parameter con-figuration relative to both the full CMBAgent system and the direct Base LLM baseline.

The four research-driven tasks are assigned to these conditions as follows:

Under -Constrained Inference Stress Tests — Tasks 1, 2

Task 1: Fitting the Union2.1 Type Ia supernova distance–redshift data with a flat ΛCDM cosmology us-ing MCMC to jointly estimate H_0 and Ω_Λ. The SN1a likelihood is degenerate in H_0 and the supernova abso-lute magnitude M_B; without an independent distance anchor, M_0 cannot be constrained by this dataset alone.

Task 2: Modeling the rotation curve of NGC 3198 by combining stellar, gaseous, and bulge contributions with an NFW dark matter halo profile to infer halo virial mass M_{200} and concentration c, including uncer-tainty quantification via MCMC.

Category Rationale: Both tasks probe silent overconfidence under known parameter degeneracies: the H_0–M_B degeneracy in T1 and the mass–concentration degeneracy in T2. A reliable agent should flag these as prior-dominated rather than report confident but physically uninformative constraints.

Task Label (T1): *Under-constrained* — flat prior on H_0, H_0–M_B degeneracy unbroken by design.

Task Label (T2): *Under-constrained* —flat priors on ($\log_{10} M_{200}$, c), mass–concentration degeneracy unbroken by design.

Compositional Bayesian Workflow Evaluation —Tasks 3, 4

Task 3: Performing a hierarchical MCMC analysis of the exoplanet mass–radius relation for planets with $M > 2M_E$, incorporating intrinsic scatter, measurement uncertainties, and temperature-dependent radius inflation, with comparative analysis of Neptunian and Jovian regimes.

Task 4: Conducting Bayesian inference for SLACS strong-lensing systems using a Singular Isothermal Sphere (SIS) model to recover posterior velocity dispersions and compare lensing-inferred values to observed stellar velocity dispersions.

Category Rationale: Both tasks test silent inconsistency rather than degeneracy detection: whether the agent sustains physical consistency and stable parameterisation across a multi-stage pipeline with no artificial degradation. Failures here manifest as inter-trial drift, prior boundary pathologies, or population-level bias — none of which trigger obvious error signals.

Task Labels (T3, T4): *Compositional* — full hierarchical likelihood, no prior degradation.

Experiments

Systems

We evaluate three systems under the One-Shot workflow. The two CMBAgent variants are evaluated under their default configuration; full details of the agent-to-model assignments are provided in the CMBAgent repository (Xu et al., 2025). The Base LLM baseline uses a direct single-turn call to GPT-4o-mini with no surrounding architecture, providing a clean lower bound on performance attributable to the model alone.

Base LLM: A direct single-turn call to GPT-4o-mini with no agentic framework, no tool integration, and no CAMB documentation. The model receives the task prompt and must produce executable Python code in one pass.

CMBAgent (No Context): The CMBAgent One-Shot workflow under default configuration, with the engineer agent and up to 50 reasoning rounds, but without access to CAMB documentation.

CMBAgent (CAMB Context): The CMBAgent One-Shot workflow under default configuration, with the camb context agent providing retrieval access to the CAMB API documentation during execution. This is the full intended deployment mode of the system.

The Deep Research workflow is additionally applied to the four Astrophysical Research-Driven Tasks (T1–T4); its configuration and evaluation protocol are described in Section 5.2.

Protocol

Trials: The One-Shot workflow is evaluated over $N = 10$ independent trials per task–system combination, yielding $14 \times 10 \times 3 = 420$ total trials. All Base LLM calls use sampling temperature $T = 0.2$. The Deep Research workflow is evaluated with $N = 5$ independent trials per task; evaluation of this workflow is therefore qualitative, based on parameter recovery and physical consistency of the reported outputs rather than statistical aggregation across trials.

Ground Truth: For the One-Shot workflow, reference outputs are generated once per task by executing the reference implementation from the CMBAgent benchmark repository (Contributors, 2024) and saving the resulting (x, y) numerical array to a standardised CSV file. For the Deep Research workflow, no automated reference CSV exists; ground truth is instead taken from the published literature values listed in Table 1, covering all four research-driven tasks (T1–T4).

Table 1: Literature ground-truth values for Deep Research evaluation

Task	Param.	Reference	Source
T1	Ω_Λ	0.72±0.02	[21]
T2	$\log_{10}\left(\frac{M_{200}}{M_\odot}\right)$	11.97±0.15	[22]
		6.8±2.0	
T3	α_N	0.67±0.05	[23]
	α_J	−0.06±0.07	
	$M_{\rm br}$ [$M_\oplus$]	127±17	
T4	f	1.019±0.008	[24]

Output Standardisation: To ensure comparable output formats across systems, the target output file path is appended to every prompt at runtime. This eliminates ambiguity in output location without modifying the scientific content of the task prompt.

Deep Research Configuration: Each Deep Research trial is allocated up to 50 planning rounds, 100 control rounds, and a maximum of 5 execution retries per sub-step (max n attempts=5). The retry budget allows the workflow to self-recover from environment compatibility errors such as deprecated API calls, which were the dominant failure mode observed across tasks.

Evaluation: For the One-Shot workflow, each trial output is evaluated against the reference CSV using the ESR, PAS, NAS, and failure mode metrics defined in Section 4.2. Pertask scores are averaged across 10 trials; per-system scores are averaged across tasks and trials. For the Deep Research workflow, evaluation is conducted by comparing the agent's reported posterior medians against the literature values in Table 1, with success additionally confirmed by the presence of output files beyond the execution log.

Results

One-Shot Workflow

Figure 1 summarises mean ESR, PAS, NAS, and Final Score across all 14 tasks and 10 trials per system.

CMB Agent with CAMB Context: The full system achieves near-ceiling performance: ESR= 0.96, PAS= 0.95, NAS= 0.86, Final Score= 0.85. The small gap between PAS and NAS reflects residual numerical errors on the hardest tasks even when parameters are correctly configured.

CMB Agent without Context: Removing documentation access produces sharp degradation across all metrics (ESR = 0.62, PAS = 0.54, NAS = 0.18, Final = 0.15). The most severe drop is in NAS, consistent with Figure 3 where Mode C (wrong computation) accounts for ≈ 47% of trials — the agent invokes plausible but incorrect API calls, producing curves with the right shape but wrong amplitude or spectral content.

Base LLM: The direct GPT-4o-mini baseline fails on virtually all trials (ESR = 0.09, Final ≈ 0), with ≈ 91% classified as Mode A (code failure). Raw model capability without agentic scaffolding is insufficient for tool-grounded scientific computation.

Task-Level Analysis

Figure 2 breaks performance down across all 14 tasks. CMBAgent with CAMB context scores near 1.0 on most tasks but fails silently on two: Task 10 (score 0.09), where most trials omit raw cl=True producing a $\sim 10^5$ amplitude error and all trials omit the tensor B-mode contribution; and Task 14 (score 0.02), where the agent correctly computes per-ℓ delensing efficiency but col-lapses it to a scalar mean, yielding CCC ≈ 0. Both failures stem from output construction errors rather than incorrect API usage — the documentation does not prescribe unit conventions or post-processing for these quantities. CMBAgent without context shows an isolated peak at Task 8 (0.63), whose calling convention is common enough to be known without documentation, but scores below 0.30 elsewhere. The Base LLM scores near zero throughout.

Deep Research Workflow

The Deep Research workflow was evaluated on four research-grade tasks (T1–T4) with up to five independent trials per task. Each task is reported with its PRS score and a qualitative assessment across the two dimensions defined in Section 4.2: physical plausibility and failure transparency. Table 2 summarises results.

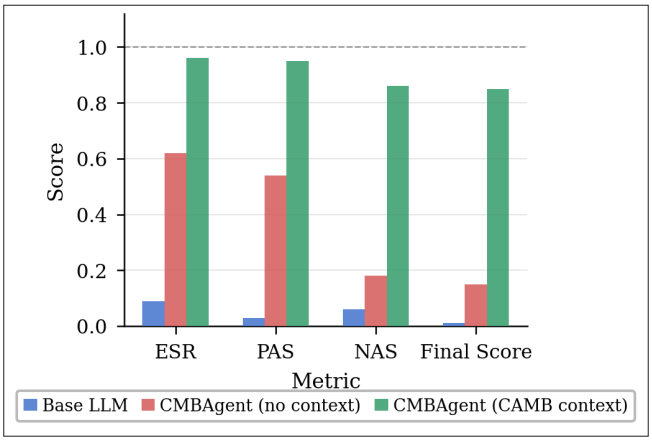


Figure 1: Mean ESR, PAS, NAS, and Final Score per system across 14 tasks and 10 trials

T1 — SN1a Under-Constrained: All four completed trials recovered Ω_Λ within 0.15σ of the Union2.1 reference 0.720 ± 0.020 with mean posterior 0.722 ± 0.001 confirming high reproducibility (PRS = 0.97) [21]. However, all trials deviated from the task specification by treating H_0 as a free parameter alongside Ω_Λ. This mirrors a realistic failure scenario: the prompt was intentionally constructed as a naive but plausible scientific request, and a reliable agent should have identified the H_0–H_B degeneracy and flagged that H_0 cannot be constrained from this dataset alone. Instead, the recovered H_0 posterior simply reflects the flat prior $H_0 \in [50, 90]$ rather than any data constraint, and

was reported as a genuine measurement.

Physical Plausibility: Ω_Λ posteriors are well-constrained and physically reasonable; the H_0 posterior is prior-dominated and physically uninformative.

Failure Transparency: two silent failures on completed trials—the unauthorised addition of H_0 as a free parameter yield-ing a prior-dominated posterior, and a runtime constraint violation — neither flagged by the agent.

T2 — NGC 3198 Under-Constrained: All five trials failed to recover parameters consistent with the reference value (PRS=0.05). Mass estimates spanned $\log10(M_{200}/M_\odot) \in [10.0, 12.9]$ — nearly three orders of magnitude — with no trial within 5σ of the reference 11.97 ± 0.15 [22]. Three trials pro-duced unphysical NFW concentrations: two high-mass trials ($\log10 M_{200} \approx 12.8$–$12.9$) returned $c < 2$, while one trial hit the upper prior boundary at $c \approx 30$.

Physical Plausibility: three of five trials returned concentrations outside $c \in [2, 30]$; the unphysical $c < 2$ cases coincided with high virial mass, tracing a degenerate high-mass, low-concentration ridge.

Failure Transparency: one trial noted elongated contours but no trial flagged unphysical posteriors or added caveats to its parameter estimates.

T3 — Exoplanet Mass–Radius Compositional: The Neptunian slope $\alpha_N \approx 0.60$ is consistently recovered across all four trials ($\sim 1.4\sigma$ from reference 0.67 ± 0.05 [23]. The Jovian slope α_J and break mass Mbr are inconsistent across trials: Trial 0 recovers $M_{\rm br} \approx 127\ M\oplus$ (within 1σ of reference) with $\alpha_J = 0.002$ ($\sim 0.9\sigma$); Trial 1 returns $\alpha_J = 0.368$ (6.1σ from reference -0.06 ± 0.07) and $M_{\rm br} \approx 203\ M_\oplus$ (4.5σ above reference); Trial 3 adopts a structurally different parameterisation with $\log M_b = 4.94$ on an inconsistent unit scale, precluding direct comparison; Trial 4 recovers $M_{\rm br} \approx 118\ M_\oplus$ (0.6σ below reference) with $\alpha_J \approx 0.000$ (0.9σ), but with $\sigma_{\rm int}$ (the intrinsic scat-ter in the mass–radius relation beyond measurement noise) pinned at the prior boundary (0.5), indicating inadequate prior specification. Mean PRS over the three key parameters gives PRS = 0.73, driven by robust α_N recovery but large variance in α_J and $M_{\rm br}$.

Physical Plausibility: α_N is consistently physical; α_J is wrong in sign or magnitude in three of five trials; $M_{\rm br}$ is implausible in at least two trials (Trial 1 at 4.5σ above reference; Trial 3 on an inconsistent unit scale); $\sigma_{\rm int}$ hits the prior boundary in Trial 4.

Failure Transparency: no trial flagged the inconsistent model parameterizations, implausible Jovian slopes, discrepant break masses, or the prior-boundary pathology in $\sigma_{\rm int}$; confident estimates were reported throughout without caveats.

T4—SLACS Strong Lensing Compositional: Only one of five trials produced usable output; the remaining four failed to complete. The single completed trial yielded physically coherent results. However, lensing-inferred dispersions are systematically $\approx$5% above observed values, yielding $f \approx 1.05$ against the reference $f = 1.019 \pm 0.008$ (Bolton et al., 2008), a $\approx 4\sigma$ bias (PRS = 0.00).

Physical Plausibility: velocity dispersions and Einstein radius predictions are within the physically expected range and the SIS model is correctly implemented.

Failure Transparency: the systematic offset in f was not flagged by the agent; the four failed trials produced no output or error diagnosis.

Conclusion

We presented a structured empirical evaluation of CMBAgent across two workflow paradigms and eighteen tasks spanning tool-grounded computation and research-grade Bayesian inference. Rather than focusing on aggregate success, our goal was to characterise how and why agentic systems fail, and whether those failures are detectable or silent.

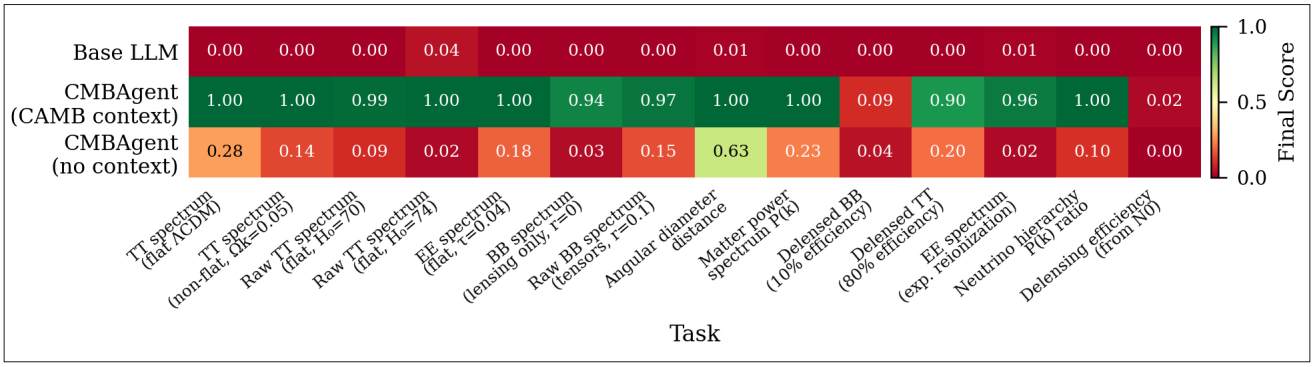


Figure 2: Final Score per system and task, averaged over 10 trials. Green indicates high score, red indicates low score

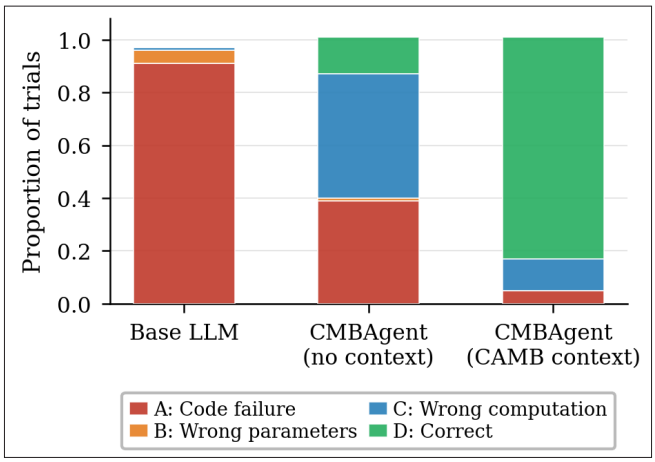


Figure 3: Failure mode breakdown per system as proportion of trials

Table 2: Deep Research results. PRS is computed from posterior medians vs. literature for completed trials only; qualitative dimensions (PP = physical plausibility, FT = failure transparency) are rated ✓ (pass), ~ (partial), or × (fail)

Task	Category	Trials	PRS	PP	FT
T1 SN1a	Under-constrained	4/5	0.97	~	X
T2 NGC3198	Under-constrained	5/5	0.05	×	X
T3 Exoplanets	Compositional	5/5	0.73	~	X
T4 SLACS	Compositional	1/5	0.00	✓	X

Across both workflows, a consistent pattern emerges. In the One-Shot setting, domain-context retrieval is the dominant performance driver: with an identical GPT-4o-mini backbone, final score improves from ≈ 0 (Base LLM) to 0.15 (no context) to 0.85 (full context), a ∼6× gain. Without context, the primary failure mode is not execution failure but silent wrong computation—syntactically valid code that produces plausible but incorrect results.

In the Deep Research setting, this failure pattern persists at the inference level. Under-constrained tasks fail systematically: degeneracies go undetected, yielding physically impossible posteriors, yet results are reported as valid. Compositional workflows execute successfully but exhibit intertrial inconsistency, systematic bias, and unstable parameterisations without caveats. Across all tasks, outputs are consistently plausible, but failures remain unreported. Failure transparency is the weakest dimension: the agent never proactively flags known pathologies in its own outputs.

These findings highlight a critical risk for scientific deployment. Agentic systems do not primarily fail by crashing— they fail by producing confident, incorrect results or by silently breaking pipelines without diagnosis. Our evaluation framework, combining quantitative metrics for toolgrounded computation with structured qualitative assessment for scientific inference, provides a template for systematically identifying such failures. We argue that rigorous reliability evaluation is a prerequisite for deploying agentic AI in scientific workflows, where undetected errors can directly compromise scientific conclusions.

Deep Research Task Prompts

T1 — SN1a Under-Constrained

Read the file: /home/sr/Desktop/code/cmbagent/cmbagent_systematics/deepresearch/task/SCPUnion2.1_mu_vs_z.txt

Its description is: An ASCII table with tabseparated columns: Supernova Name, Redshift, Distance Modulus, and Distance Modulus Error. For Union2.1, there is an additional column for the probability that the supernova was hosted by a low-mass galaxy.

Fit this data within a flat ΛCDM model with two free parameters: H_0 and Ω_Λ. Write a simple but optimized MCMC to fit the SN1a data. Make a contour plot showing the 1D posteriors and quote the mean and 1σ on each parameter. Show the data alongside the best-fit model with 68% and 95% CL regions. Comment on the results.

Constraints: Running on a Dell Precision 5480 workstation with 32GB RAM and 20 CPU threads (Intel Core i7-13800H). Use resources optimally so the MCMC converges within a few minutes. Have the engineer agent perform a preliminary MCMC timing step as a separate preliminary step.

T2 — NGC3198 Under-Constrained

Read the galaxy rotation curve file: /home/sr/Desktop/code/cmbagent/cmbagent_systematics/deepresearch/task/NGC3198_rotmod.txt

Description: SPARC (Spitzer Photometry and Accurate Rotation Curves) provides high-quality rotation curves for nearby disk galaxies. For NGC3198, the data includes: radius (kpc), observed rotation velocity (km/s), velocity uncertainty, gas contribution ($V_{\rm gas}$), disk contribution ($V_{\rm Disk}$), and bulge contribution ($V_{\rm bul}$). The observed rotation curve traces the total gravitational potential.

Fit the rotation curve using an NFW (Navarro-Frenk-White) dark matter halo model combined with the baryonic contributions. The free parameters are: M_{200} (virial mass of the dark matter halo, in solar masses) and c (concentration parameter).

Write an optimized MCMC code to fit these parameters. Show the 1D posterior distributions and the 2D parameter contour plot. Quote the mean and 1σ values.

Plot rotation velocity versus radius showing: observed data with error bars, total best-fit model, individual components (disk, gas, bulge, dark matter halo), and 68% and 95% confidence bands.

Constraints: Run efficiently on a Dell Precision 5480 with 32GB RAM and 20 CPU threads. Ensure the MCMC completes within a few minutes. Include a preliminary timing test.

T3 — Exoplanet Mass–Radius

Read the exoplanet data file: /home/sr/Desktop/code/cmbagent/cmbagent_systematics/deepresearch/task/NASA_exoplanet_archive.csv

Description: The NASA Exoplanet Archive provides confirmed exoplanet measurements. The file contains: planet name, mass ($M_{\rm p}$ in $M_\oplus$) with uncertainties, radius (R_p in $R_\oplus$) with uncertainties, equilibrium temperature ($T_{\rm eq}$ in K) with uncertainties, discovery method, stellar metallicity ([Fe/H]) with uncertainties, orbital period (days), semi-major axis (AU), and system distance (pc).

Filter to planets with $M_p > 2M_\oplus$ to focus on the Neptunian and Jovian regime.

Fit a broken power-law mass–radius relation with two regimes separated by a break mass M_{br}. Free parameters are: Neptunian slope α_N , Jovian slope α_J, break mass M_{br}, normalisation, intrinsic scatter σ_{int}, and a temperature-dependent radius inflation coefficient for hot planets.

Write an optimized MCMC sampler to obtain posterior distributions of all parameters. Implement: data filtering above $2M_\oplus$; hierarchical modelling of measurement uncertainties in both mass and radius; intrinsic scatter perpendicular to the mass–radius relation, and temperature-dependent corrections for radius inflation in hot planets.

Produce the following plots: mass–radius diagram with best-fit broken power-law curves; residuals versus mass showing the transition at M_{br}; posterior distributions for α_N, α_J , and M_{br}; a corner plot showing parameter correlations; and a separate comparison of hot planets ($T_{eq} > 1000$ K) versus cold planets.

Constraints: Dell Precision 5480, 32GB RAM, 20 CPU threads. Target runtime 5–8 minutes. Use analytic power-law models for efficient likelihood evaluation.

T4 — SLACS Strong Lensing Compositional
Read the galaxy lensing data file: /home/sr/Desktop/code/cmbagent/cmbagent_systematics/deepresearch/task/Slac_data.csv

Description: The SLACS (Sloan Lens ACS) survey provides high-quality data on strong gravitational lens systems. Columns include: z_d (lens redshift), z_s (source redshift), Reff (effective radius, arcsec), θ_{Ein} (Einstein radius, arcsec), σ_{obs} (stellar velocity dispersion, km s^{-1}), and σ_{err} (km s^{-1}).

Fit the lens mass profile using a Singular Isothermal Sphere (SIS) model with one free parameter per lens: the individual velocity dispersion σSIS. Use z_d and z_s to compute angular diameter distances for each lens–source pair. Derive the model-predicted Einstein radius from σ_{SIS} within the SIS framework, and construct a likelihood comparing predicted θ_{Ein} to observed values for each lens independently.

Write an optimized MCMC sampler to obtain the posterior distribution of σ_{SIS} for each individual lens.

Produce the following plots: 1D marginalized posteriors for representative lenses; a 2D comparison between lensing-inferred σ_{SIS} and observed σ_{obs} (one point per lens); model-predicted versus observed θ_{Ein}; and the mass profile $M(< R)$ for each lens overlaid with the Einstein radius scale.

Constraints: Dell Precision 5480, 32GB RAM, 20 CPU threads. MCMC should converge within ≈5 minutes. Use vectorized likelihood evaluation across all lenses for efficiency.

Example Deep Research Outputs
T1— SN1a Under-Constrained (Representative Trial)
The completed trial recovers $\Omega_\Lambda = 0.722 \pm 0.019$, consistent with the Union2.1 reference 0.720 ± 0.020 (Suzuki et al., 2012), but simultaneously reports $H_0 = 69.99 \pm 0.34$ km s^{-1} Mpc^{-1} as a constrained result. The H_0 posterior is prior-dominated: the Union2.1 likelihood cannot break the H_0–M_B degeneracy without an external distance anchor. The agent did not flag this, presenting both parameters as equally reliable.

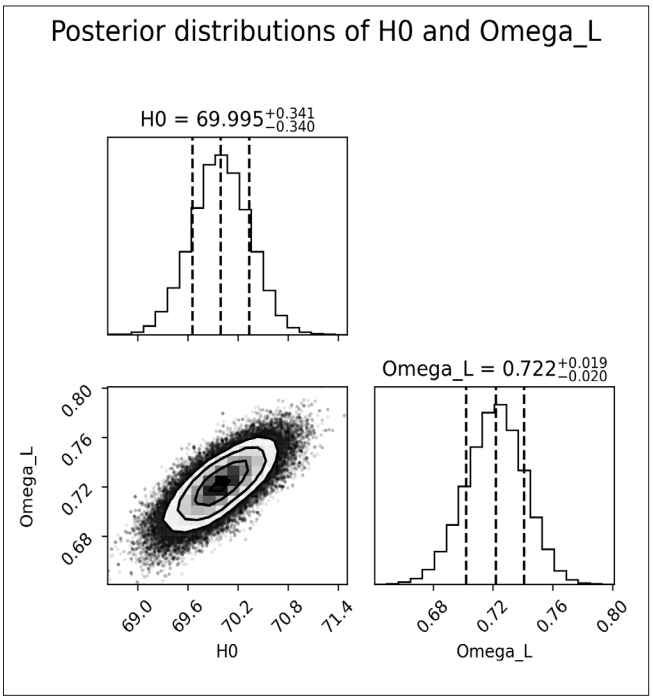


Figure 4: T1 posterior distributions of H_0 and Ω_Λ from the Union2.1 SN1a fit. While $\Omega_\Lambda = 0.722^{+0.019}_{-0.020}$ is well constrained and within 0.1σ of the reference, the H_0 posterior is prior-dominated— the Union2.1 likelihood cannot break the H_0–H_B degeneracy. The agent reported both parameters as equally reliable without flagging this pathology.

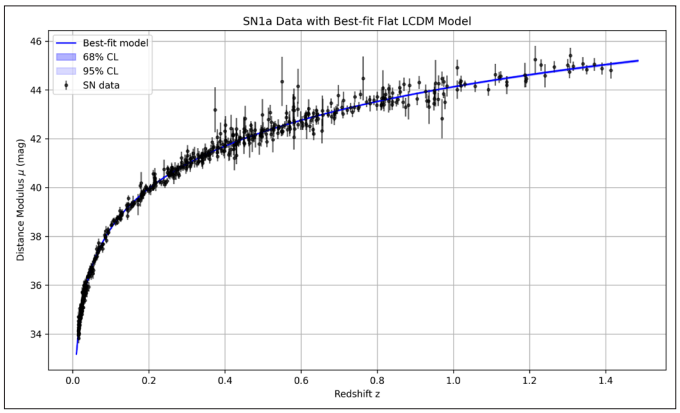


Figure 5: T1 best-fit flat ΛCDM model overlaid on Union2.1 SN1a data. The 68% and 95% CL bands are nearly indistinguishable from the best-fit curve, reflecting tight posterior constraints

T2 — NGC3198 Under-Constrained (Failure Case)
This trial illustrates the primary failure mode for T2: the posterior converges to $\log_{10}(M_{200}/M_\odot) = 12.90^{+0.06}_{-0.05}$, nearly 6σ above the reference 11.97 ± 0.15 (Karukes et al., 2015), with an unphysical concentration $c = 1.19^{+0.21}_{-0.14} \ll 2$. The agent reported these values without flagging the pathology.

T3 — Exoplanet Mass–Radius Compositional

This trial (Trial 3) illustrates the inconsistent parameterisation failure: logM_b = 4.939 is reported on an inconsistent

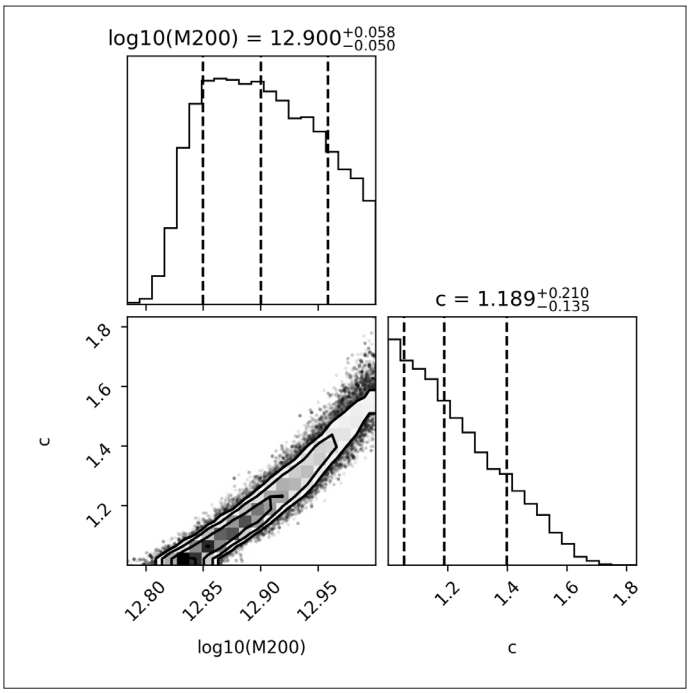


Figure 6: T2 posterior for $\log_{10} M_{200}$ and c, showing the degenerate high-mass, low-concentration ridge. The recovered $c \approx 1.2$ is physically impossible for an NFW halo; the reference value is $c = 6.8 \pm 2.0$ (Karukes et al., 2015). No trial flagged this pathology

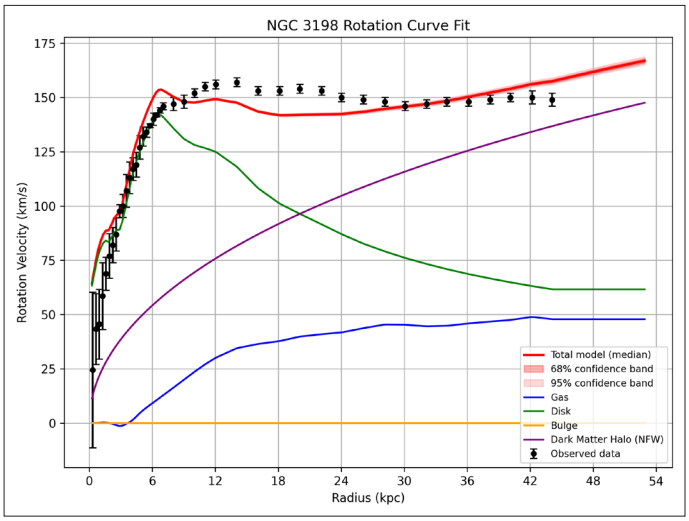


Figure 7: T2 rotation curve fit for NGC3198. Despite a visually acceptable total model (red), the inferred NFW halo parameters are unphysical: the dark matter contribution (purple) is entirely dominant at all radii, inconsistent with the observed baryonic components

unit scale, precluding direct comparison to the reference $M_{\rm br}$ = 127 ± 17M⊕ (M¨uller et al., 2024). No caveat was generated by the agent.

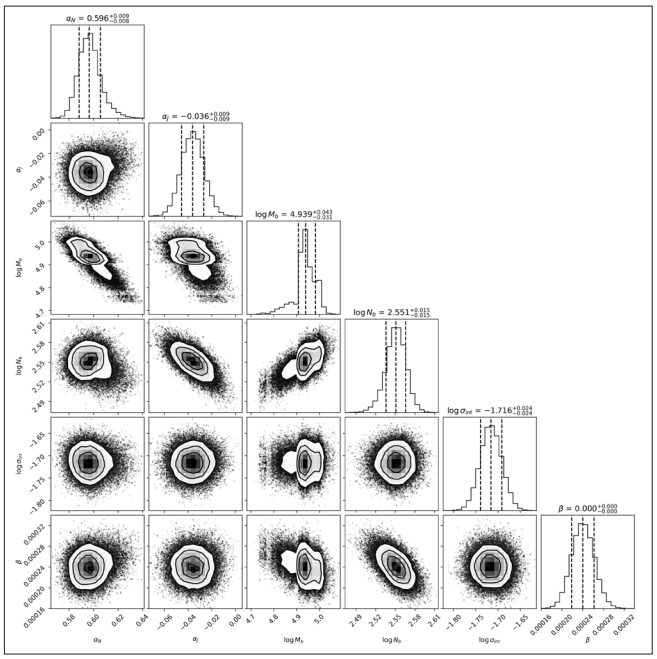


Figure 8: T3 corner plot from Trial 3. The break mass is reported as $\log M_b = 4.939^{+0.043}_{-0.031}$ on an inconsistent unit scale

The inflation coefficient $\beta \approx 0$ is effectively zero, indicating the temperature-dependent correction was not learned. No trial flagged these pathologies.

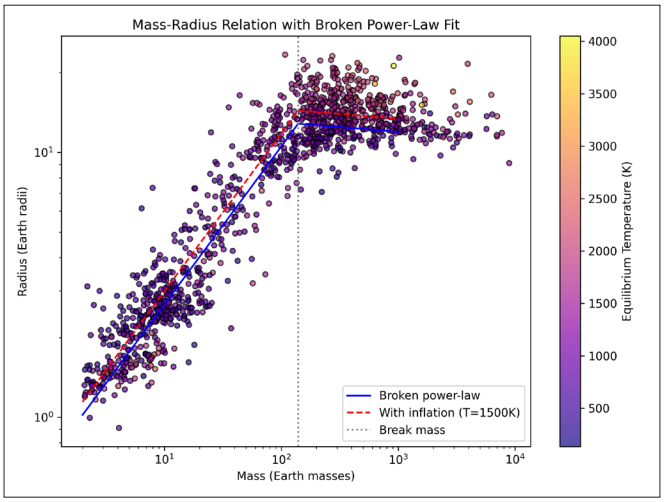


Figure 9: T3 mass–radius diagram with best-fit broken power-law. The fit appears visually reasonable, illustrating how plausiblelooking outputs can mask underlying parameterization failures

T4 — SLACS Strong Lensing (Single Completed Trial)

Of five independent trials, only one produced usable output; the remaining four failed silently with no figures or error diagnosis. The single completed trial shows physically coherent velocity dispersions, but lensing-inferred values are systematically ≈5% above observed values ($f \approx 1.05$ vs. reference 1.019 ± 0.008 (Bolton et al., 2008)), a ≈4σ bias the agent did not flag.

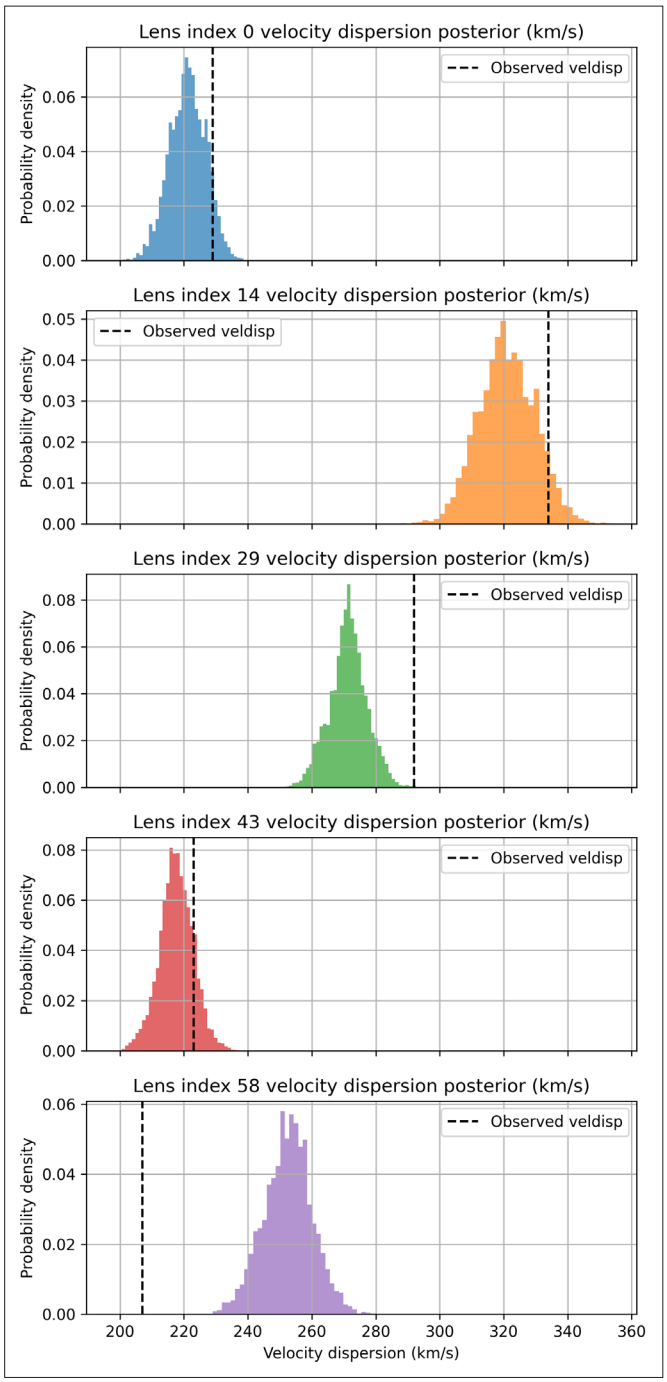


Figure 10: T4 representative 1D posteriors for five lenses. Several posteriors show tension with the observed stellar velocity dispersion (dashed line), consistent with the population-level ≈5% systematic offset. No caveat was reported by the agent

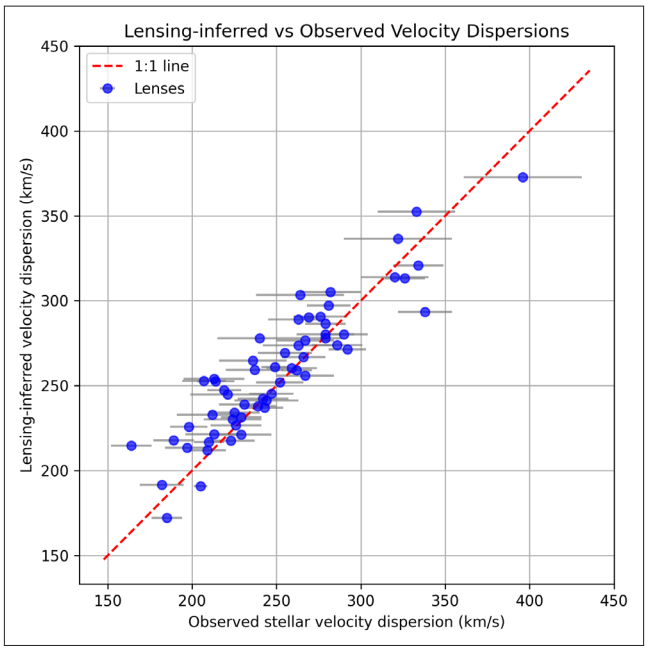


Figure 11: T4 lensing-inferred vs. observed stellar velocity dispersion σ_{SIS} vs. σ_{obs} (one point per lens). Points lie systematically above the 1:1 line (red dashed), yielding $f \approx 1.05$ against reference 1.019 ± 0.008 (Bolton et al., 2008)

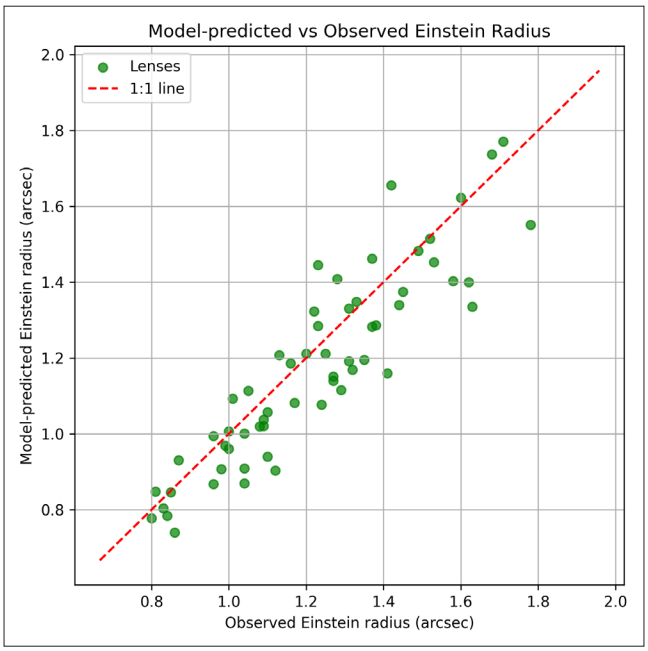


Figure 12: T4 model-predicted vs. observed Einstein radius θ_{Ein}. Agreement is closer to the 1:1 line here than for the velocity dispersions, consistent with the SIS model being correctly implemented but the overall calibration being offset

References

1. Brown T, Mann B, Ryder N, Subbiah M, Kaplan J, et al. (2020) Language models are few-shot learners. Advances in neural information processing systems 33: 1877-1901.
2. Wei J, Wang X, Schuurmans D, Bosma M, Xia F, et al. (2022) Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems 35: 24824-24837.
3. Achiam J, Adler S, Agarwal S, Ahmad L, Akkaya I, et al. (2023) Gpt-4 technical report. arXiv preprint arXiv: 2303.08774.
4. Yao S, Zhao J, Yu D, Du N, Shafran I, (2022) React: Synergizing reasoning and acting in language models. In The eleventh international conference on learning representations.
5. Schick T, Dwivedi-Yu J, Dess`ı R, Raileanu R, Lomeli M, et al. (2023) Toolformer: Language models can teach themselves to use tools. Advances in neural information processing systems, 36: 68539-68551.
6. Nakano R, Hilton J, Balaji S, Wu J, Ouyang L, et al. (2021) Webgpt: Browser-assisted question-answering with hu-man feedback. arXiv preprint arXiv: 2112.09332.
7. Qin S, Zhu Y, Mu L, Zhang S,Zhang X (2025) Meta-tool: Unleash open-world function calling capabilities of general-purpose large language models. In Proceed-ings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) 30653-30677.
8. He P, Dai Z, He B, Liu H, Tang X, et al. (2025) Traject-bench: A trajectory-aware benchmark for evaluating agentic tool use. arXiv preprint arXiv: 2510.04550.
9. Ding K, Yu J, Huang J, Yang Y, Zhang Q, Chen H (2025) Scitoolagent: a knowledge-graph-driven scientific agent for multitool integration. Nature Computational Science 5): 962-972.
10. Sun Q, Liu Z, Ma C, Ding Z, Xu F, et al. (2025) Scienceboard: Evaluating multimodal autonomous agents in realistic scientific workflows. arXiv preprint arXiv: 2505.19897.
11. Villaescusa-Navarro F, Bolliet B, Villanueva-Domingo P, Bayer AE, Acquah A, et al. (2025) The denario project: Deep knowledge ai agents for scientific discovery. arXiv preprint arXiv: 2510.26887.
12. Moreno E A, Bright-Thonney S, Novak A, Garcia D, Harris P (2026) Ai agents can already autonomously perform experimental high energy physics. arXiv preprint arXiv: 2603.20179.
13. Hellert T, Bertwistle D, Leemann SC, Sulc A, Venturini M (2025) Agentic ai for multi-stage physics exper-iments at a large-scale user facility particle accelerator. arXiv preprint arXiv: 2509.17255.
14. Huang L, Yu W, Ma W, Zhong W, Feng Z, et al. (2025) A survey on hallucination in large language models: Principles, taxon-omy, challenges, and open questions. ACM Transactions on Information Systems 43: 1-55.
15. Alansari A, Luqman H (2025) Large language models hal-lucination: A comprehensive survey. arXiv preprint arXiv:2510.06265.
16. Lin S, Hilton J, Evans O (2022) Truthfulqa: Measuring how models mimic human falsehoods. In Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)3214-3252, 2022.
17. Yehudai A, Eden L, Li A, Uziel, Zhao, et al. (2025) Survey on evaluation of llm-based agents. arXiv preprint arXiv: 2503.16416.
18. Chowa S, Alvi R, Rahman SS, Rahman MA, Raiaan MAK, et al. (2026) From language to action: a review of large language models as autonomous agents and tool users. Artificial Intelligence Review.
19. Xu L, Sarkar M, Lonappan AI, Zubeldia I, Villanueva-Domingo, et al. (2025) Open-source planning & control system with language agents for autonomous scientific discovery. arXiv preprint arXiv: 2507.07257.
20. Contributors C (2024) Cmbagent benchmarks repository. https://github.com/cmbagent/Benchmarks.
21. Suzuki N, Rubin D, Lidman C, Aldering G, Amanullah R, et al. (2012) The hubble space telescope cluster supernova survey. v. improving the dark-energy constraints above z¿ 1 and building an early-type-hosted supernova sample. The Astrophysical Journal 746: 85.
22. Karukes E, Salucci P, Gentile G (2015) The dark mat-ter distribution in the spiral ngc 3198 out to 0.22 rvir. Astronomy & Astrophysics 578: A13.
23. Mu¨ ller, Baron J, Helled, R, Bouchy F, Parc L (2024) The mass-radius relation of exoplanets revisited. Astron-omy & Astrophysics 686: A296.
24. Bolton AS, Burles S, Koopmans LV, Treu T, Gavazzi R, et al. (2008) Schlegel, D. J. The sloan lens acs survey. v. the full acs strong-lens sample. The Astrophysical Journal 682: 964-984.
25. Bai Y, Jones A, Ndousse K, Askell A, Chen A, et al. (2022) Training a helpful and harmless assistant with rein-forcement learning from human feedback. arXiv preprint arXiv: 2204.05862.
26. Lin L, I.-K. A (1989) concordance correlation coefficient to evalu-ate reproducibility. Biometrics 45: 255-268.
27. Lin X, (2025) LLM-based agents suffer from hallucinations: A survey of taxonomy, methods, and directions. arXiv preprint arXiv: 2509.18970.

Copyright: ©2026 Shivam Rawat. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.