

[← Go to NeurIPS 2026 Workshop AgenticLS homepage](#)

Position: Toward a Computational Theory of Scientific Discovery

Odysseus Kaloidas¹

¹ University of California, Berkeley (berkeley.edu)

27 Sept 2026 (modified: 27 Sept 2026)

NeurIPS 2026 Workshop AgenticLS Submission AgenticLS, Reviewers, Authors

Revisions CC BY 4.0

Keywords: scientific discovery, AI agents, capability forecasting, computational complexity, scaling laws, biomedical discovery, experimental validation

TL;DR: Predicting which scientific problems AI will solve next requires connecting capability growth to target success, procedure availability and validated completion.

Abstract:

As AI improves, which scientific problems become solvable next, and when? We argue that a computational theory of scientific discovery should connect capability growth to target success, the availability of a suitable procedure, and validated completion. A worked assay scenario shows the consequence: the first procedure to become available can have a lower chance of delivering a result by a common deadline. Two retrospective studies test what information improves target predictions. Across 73 models and two reasoning families, object count improves pooled held-out prediction but worsens larger-task deduction; it does not isolate computational complexity. In 550 chemical-table replays, early outcome summaries help only in some endpoint-budget settings. A retrospective diagnostic finds that separate budget-specific fits do not remove the longer-budget reversal, while accurate aggregate success rates hide opposing errors across sources. Exact conditional calculations explain which measurements constrain availability dates and why validation and evaluation costs matter. Together, these comparisons motivate a research programme for choosing informative measurements and comparing research procedures. They establish conditional consequences and retrospective prediction boundaries; prospective completion dates remain unvalidated.

Email Sharing: We authorize the sharing of all author emails with Program Chairs.

Data Release: We authorize the release of our submission and author names to the public in the event of acceptance.

Submission Number: 129

Filter by reply type...

Filter by author...

Search keywords...

Sort: Newest First

Everyone

Program Chairs

Submission129 Authors

1 / 1 reply shown



Paper Decision

Decision by Program Chairs  30 Sept 2026, 01:23 (modified: 30 Sept 2026, 07:34)

 Program Chairs, Authors  [Revisions](#)

Decision: Accept (Poster)

Comment:

The paper provides a rigorous position framework distinguishing model capability from target-specific transfer, observation access, and validated scientific completion. Conditional examples and candid retrospective failures sharpen an important research agenda. The contribution is conceptual and diagnostic, not a validated calendar forecast or demonstrated allocation policy. Its explicit assumptions and counterexamples make that bounded position useful.

Add: [Official Comment](#)

[About OpenReview](#)

[FAQ](#)

[Sponsors](#)

[All Venues](#)

[Contact](#)

[Hosting a Venue](#)

[Terms of Use / Privacy Policy](#)

[Donate](#)

[OpenReview](#) is a long-term project to advance science through improved peer review with legal nonprofit status. We gratefully acknowledge the support of the [OpenReview Sponsors](#). © 2026 OpenReview