

INDIVIDUAL OPPORTUNITY THESIS

Recursive Asterisms

Explore the frontiers of mathematics.
Build on verifiable foundations.

$$P \stackrel{?}{=} NP$$

The question that opened the inquiry

01 · Original formula


$$(\neg x \vee y) \wedge x$$



Two satisfying assignments

02 · Delete a blocked clause


$$x$$



Four satisfying assignments

03 · Delete the remaining unit

$$\top$$



Eight satisfying assignments

Checked deletion transforms the formula.

Protected answers stay the same: $z \in \{0, 1\}$

THE OPPORTUNITY

Recursive Asterisms would connect AI-guided mathematical exploration, visual explanation, and formal verification in one research workspace. A question becomes an evolving mathematical object: users can unfold its definitions, test an assumption, inspect a counterexample, and follow the evidence behind a result. The opportunity is to make the discoveries along a difficult investigation understandable and useful to the next person—even when the original problem remains unsolved.

THE PROBLEM AND INITIAL CUSTOMER

Our research exposed a recurring internal cost: reconstructing what had been claimed, which assumptions applied, and what had actually been checked across conversations, proof files, and notes. A valid theorem can become unreliable in practice when a summary loses its conditions. The initial customer hypothesis is a small mathematics or formal-methods research group already using AI and Lean. Researchers and reviewers would use RA; a laboratory leader could sponsor adoption. The prevalence and cost of this problem outside our own work still require validation. [1]

ANOTHER CHECKED RESULT

For Boolean tables of arity ≤ 3 : B is the selector's valid output; $\kappa(I)$ is the minimum valid set size.

$$|B| \leq 9\kappa(I)$$

EVIDENCE: DISCOVERIES ON THE PATH TO P VERSUS NP

The investigation informing RA produced four selected Lean-checked result families: a factor-nine size guarantee for a selector over specified Boolean constraints; preservation of protected variable values during simplification, with witness reconstruction; a correct Horn decision component with a guarded input bridge; and a theorem identifying when appended output cannot interfere with specified machine execution. A recorded local replay accepted 17 positive runs and rejected eight deliberately false controls. These scoped results do not resolve P versus NP. [1–3]

One example makes the value visible. Over variables (x, y, z) , the formula $(\neg x \vee y) \wedge x$ can undergo two checked deletions, leaving x and then true. Its complete solution count grows from two to four to eight, while the possible values of protected variable z remain unchanged. Reversing the deletion trace reconstructs a valid original solution: $001 \rightarrow 101 \rightarrow 111$, with z fixed. The wrong repair order ends at an invalid answer. RA would let a user see why the order matters and inspect the exact conditions behind it. [2]

The record also preserves adverse evidence. In a separate finite experiment, three residual formulas shrank from 49, 121, and 241 clauses to one each, yet coverage stayed unchanged and a full comparison ran slower. The value of the method includes discovering what fails. [1]

RECURSIVE ASTERISMS


Preserve. Reconstruct. Verify.

$$\pi_P \text{ Models}(F) = \pi_P \text{ Models}(R)$$

F: original formula · R: reduced formula · P: protected variable set.

A checked trace preserves the exact model projection.


Start with a residual solution



001

Coordinates: x, y, z


Restore the last deletion



101

Protected $z = 1$ throughout

Recover an original solution



111

Wrong order: 101 · fails original F

The reconstruction runs backward through the deletion trace. The variable the user protected never changes.

ALTERNATIVES AND WHY NOW

Researchers can already combine ChatGPT, Lean/mathlib, an editor, and Notion; Aristotle also offers AI-assisted formal proof work. RA must earn adoption through better understanding and reuse than these tools provide when used well. AI assistance and established proof infrastructure make the proposed integration plausible now. Our hypothesis is that connecting exploration to evidence can reduce the review and reconstruction burden created by increasingly productive research workflows. [4–6]

RISKIEST ASSUMPTION

Researchers will save more effort through trustworthy explanation and reuse than they spend formalizing, reviewing, and maintaining the shared material—and a buyer will fund that benefit. Technical results establish a substantive internal case study; they do not establish customer demand.

WHAT WOULD DISCONFIRM THE OPPORTUNITY?

Interview 12–15 prospective users about recent work, lost time, and actual consequences. Then compare three independent groups' reuse tasks against their existing workflows, counting expert review time. Before testing, set a provisional target of 25% less total time without increased mathematical error, plus successful reuse of another group's result. Little recurring need, misleading explanations, review costs exceeding savings, or failure to obtain two sponsor-backed pilot commitments would weaken or disconfirm the thesis. These tests have not yet occurred.

FIRST INVESTMENT

Build one private Boolean-logic pilot connecting search, visual explanation, assumptions, proof inspection, and reuse. Test paid team collaboration and managed verification as funding hypotheses. Expand toward broader mathematics and chemistry after evidence supports adoption. The business should create value through useful intermediate discoveries without depending on a Millennium Prize solution.

SOURCES

- 1 Research exchange
- 2 Lean programme and replay records
- 3 Stephen Cook: P versus NP
- 4 Lean: Validating a Lean Proof
- 5 Mathlib contribution and review
- 6 Aristotle

AI assistance: ChatGPT/Codex assisted with source synthesis, drafting, critique, and visual preparation. Interpretation and final recommendations remain the author's responsibility.



Explore the research
recursiveasterisms.com