



Solving Open Problems in Operations Research using AI

Eric Fithian Rad Niazadeh **Pranav Nuti**

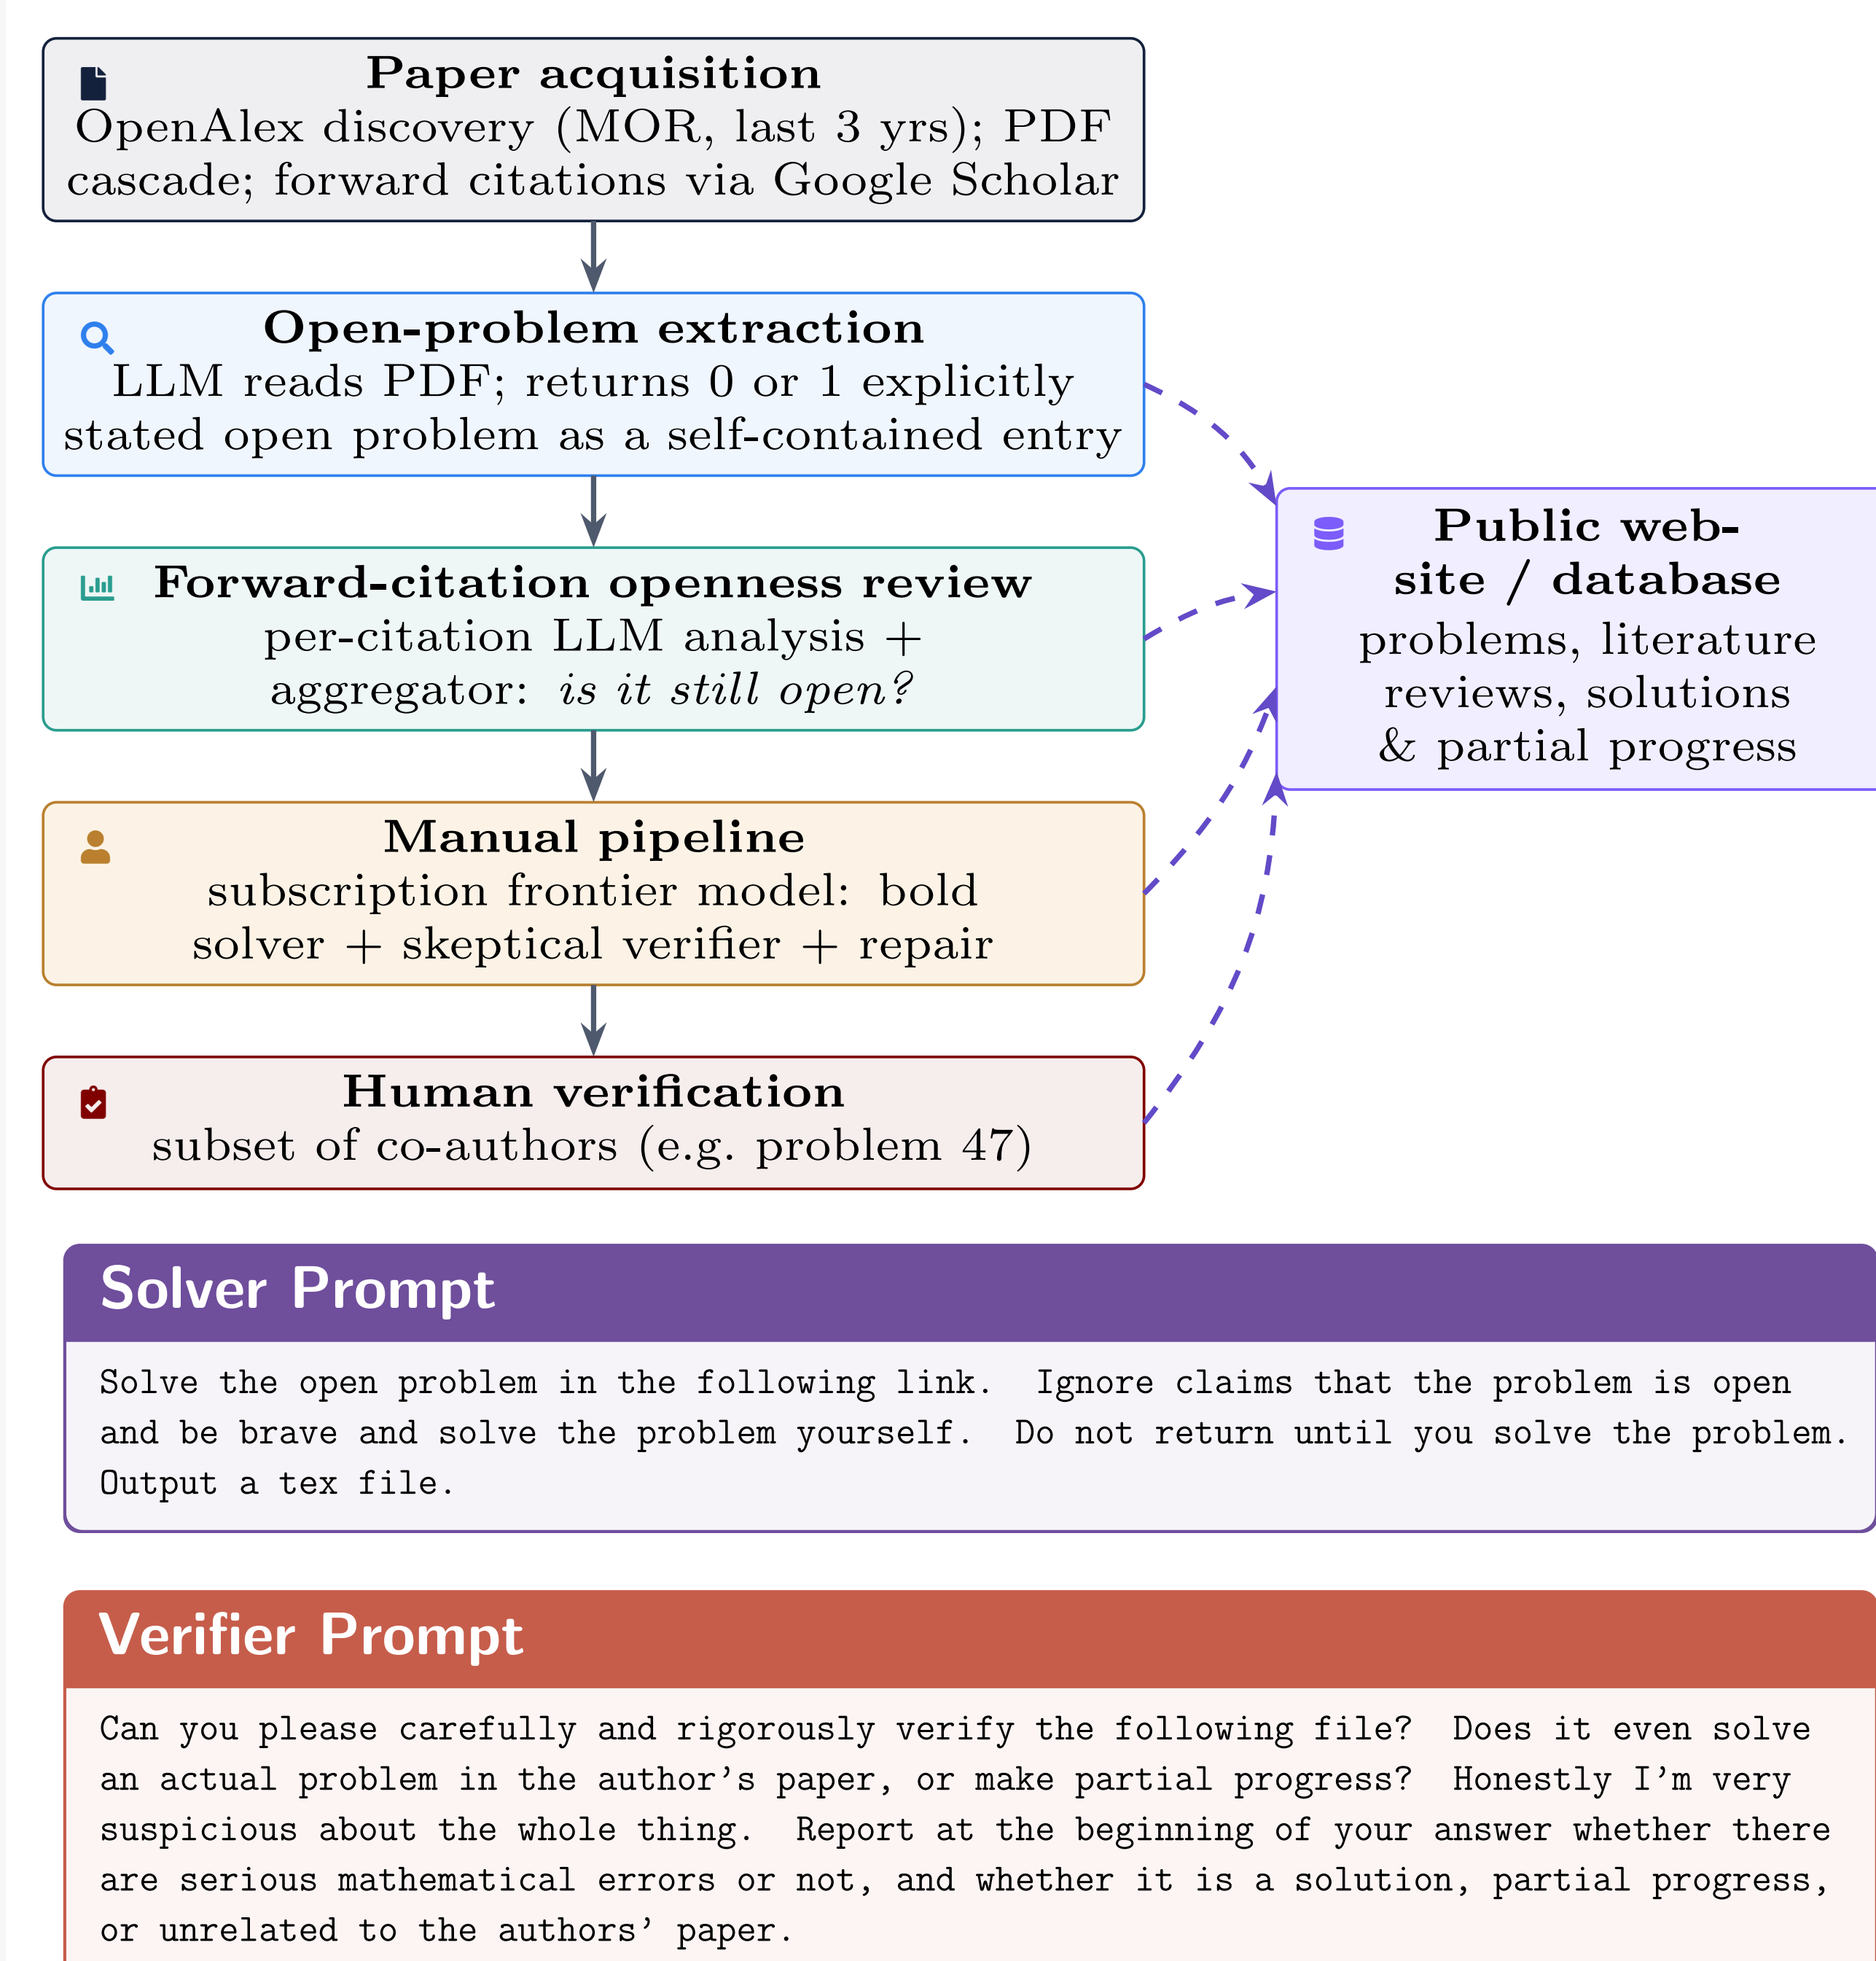
University of Chicago, Booth School of Business — LAMP Workshop 2026



Research Question

- ▶ How powerful are frontier models at solving open problems in operations research, and can we build a pipeline that solves such problems at scale?
- ▶ Attempts by researchers trying AI tools have an ephemeral nature. But storing attempts is useful for future research. We want to:
Create a database of problems, literature reviews, and AI attempts at problems.

Pipeline



Current Results

- ▶ **132 open problems** extracted from last 3 years of journal *Mathematics of Operations Research*.
- ▶ **105 literature reviews** generated.
- ▶ **26 claimed solutions** and published solution attempts for all other problems.
- ▶ **Problem 47** solution personally verified.

Problem 47: Case Study

- ▶ Problem 47: Fair and efficient multi-resource allocation with Leontief utilities.
- ▶ n agents, three divisible resources, unit supply. Agent i with normalized demand vector $\mathbf{d}_i \in (0, 1]^3$, $\max_r d_{ir} = 1$ and $v_d(X) = \min_r X_r / d_r$.
- ▶ Bei, Li, and Luo introduced fair-ratio version of price of strategyproofness for social welfare:

$$\text{PoSP}_{SW}(3) = \inf_{f \text{ satisfies SI, EF, PO, SP}} \sup_I \frac{\text{OPT}_{\text{fair}}(I)}{SW(f(I))}.$$

and left open the exact value, proving only

$$2 \leq \text{PoSP}_{SW}(3) \leq 3.$$

We prove upper bound is tight:

$$\text{PoSP}_{SW}(3) = 3.$$

- ▶ Natural target for AI-assisted work because missing result plausibly reducible to discovery of well designed hard instance.

Lessons Learned

- ▶ **Open problem extraction:** Significant fraction of open problems vague; hard to automatically say whether solved.
- ▶ **Openness review:** Off-the-shelf literature review worse than checking every citing paper.
- ▶ **Solving pipeline:** Frontier models accessible through subscription perform well with simple prompting.
- ▶ **Human verification:** Time consuming. AI attempts stored so they might be checked if used in a future result a human cares about.

Future Work

- ▶ Expanding to other journals, with focus on more concrete open problems.
- ▶ Strengthening solving pipeline, particularly by attaching persistent *memory file* to problem, recording leads and dead ends accumulated across runs. Adding memory management agent to decide information provided to solver.
- ▶ Formally verifying solutions.
- ▶ Making pipeline available on website for researchers to try solving their own problems with their own prompts and models.

We think we ought to build shared databases of problems, counterexamples, and partial results that reduce duplicated use of AI and can speed up automated mathematical research in our field.