GGLab

Bachelor / Master Thesis

FAU

LEAN 4 · REINFORCEMENT LEARNING · LLMS

Two Heads, One Proof.

Can a prover and its inner critic learn to argue, and grow, together?

The Idea

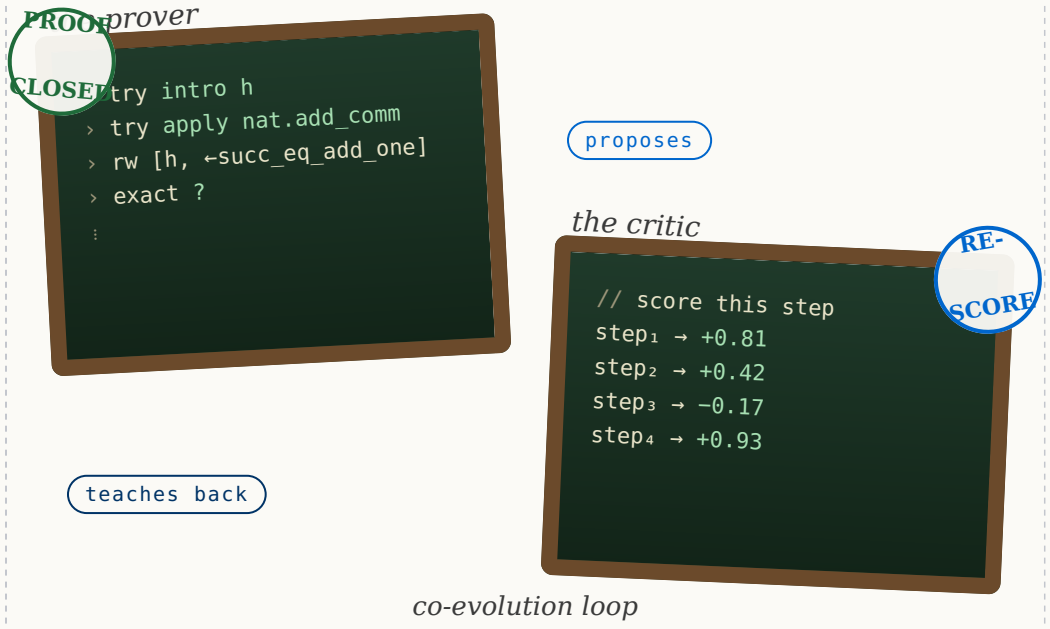

Two students, one blackboard.

One model proposes proofs in **Lean 4**. The other learns to judge them, and the better the judge gets, the better the prover gets, and so on. We train them *at the same time*, so each one's progress becomes the other's curriculum.

theorem.lean

```
-- a goal the system has to close on its own
theorem add_comm_succ (n m : ℕ) :
  n + (m + 1) = (n + m) + 1 := by
  -- ??  your model fills this in
```

Why it matters: a strong critic is the difference between a model that guesses proofs and a model that knows when it's wrong.

What You'll Do

- Set up a closed-loop training pipeline where a **policy** and a **reward model** improve one another over rounds of Lean 4 proof attempts.
- Design and ablate the signals that flow between the two. What makes the critic *useful*, not just confident?
- Run experiments on standard formal-math benchmarks and report what actually moves the needle.
- Help us turn the findings into a paper. Your name on it.

You'll come away knowing

- How modern RL pipelines are wired together end-to-end.
- What "alignment" looks like in a domain where answers are *provably* right or wrong.
- Enough Lean 4 to be dangerous.

Skills we're looking for

PYTHON	■■■■□
PYTORCH	■■■□□
RL / NLP	■■■□□
MATH (A +)	■■□□□
LEAN	□□□□□

filled = required · empty = we'll teach you

References

[1] Polu & Sutskever, *Generative Language Modeling for Automated Theorem Proving*, 2020.
[2] Yang et al., *LeanDojo*, NeurIPS 2023.
[3] Uluşan, Akbudak, Erer, Şahin, *FormalRewardBench: A Benchmark for Formal Theorem Proving Reward Models*, arXiv 2605.10141, 2026.

Supervisors

Zeynel Ulusan, MSc
zulusan23@ku.edu.tr

Prof. Dr. Gözde Gül Şahin
Intelligent Language Systems · Informatics · FAU Erlangen-Nürnberg
goezde.sahin@fau.de

Interested?

Send us an email! Include your CV, transcript, and a **brief** note on why this problem excites you.

subject line

"Thesis — Two Heads"