Riyaz Ahuja

riyaza@andrew.cmu.edu 404-719-1661 riyazahuja.com riyaz-ahuja riyazahuja

Education

Carnegie Mellon University *Aug 2022 – May 2026*
M.S. Mathematics, B.S. Mathematics, B.S. Computer Science

- GPA: 3.82/4.0
- 2025-2026 ACS Scholar (Top 40 seniors in the university)
- 2024-2025 Richard A. Moore Award Winner
- Dean's List with High Honors
- Mathematics Honors Program
- **Relevant Coursework:** Interactive Theorem Proving, Logic and Mechanized Reasoning, Differential Topology, Algebraic Geometry, Model Theory, Categorical Logic, Category Theory, Complex Analysis, Parallel Computer Architecture, Machine Learning

Research Experience

Cofounder & Chief Researcher *Pittsburgh, PA*
Stagira Labs *May 2025 – Present*

- Cofounded Stagira Labs to investigate multi-agent AI and game-theoretic dynamics in automated scientific discovery.
- Architected Agora, a decentralized marketplace where human and AI agents trade formal definitions, conjectures, and proofs as verifiable smart contracts within a shared Lean 4 library.
- Conducted experiments on miniCTX and related benchmarks, demonstrating that interacting populations of specialized agents (provers, verifiers, critics) outperform monolithic baselines under fixed compute budgets.

Researcher *Pittsburgh, PA*
Carnegie Mellon University *May 2024 – Present*

- Conducting research in automated theorem proving and neurosymbolic AI under Profs. Jeremy Avigad, Prasad Tetali, and Sean Welleck.
- Designed ImProver, an agentic framework utilizing a novel Chain-of-States technique and symbolic context retrieval to optimize proofs, achieving ∼21% length reduction and ∼9% declarativity gain (ICLR 2025).
- Developing ImProver^2, a pipeline training Small Language Models (SLMs) via bootstrapped reinforcement learning to iteratively self-improve on structural metrics (modularity, dependency) where human data is scarce.
- Built LeanM2, a type-consistent bridge handling complex types (e.g., nested polynomial quotient rings) to generate formally verified Gröbner basis witnesses, significantly outperforming `polyrith`.
- Initiated cLean, a GPU verification framework featuring a custom Lean elaborator that translates do-blocks into DeviceIR to prove functional correctness, race-freedom, and barrier-safety.
- Founded the first joint Stanford-CMU neural theorem proving lab, leading development on game-theoretic approaches to automated theorem discovery.

Professional Experience

Business Development Analyst *Pittsburgh, PA*
Expii *May 2023 – Dec 2023*

- Established strategic partnerships with educational institutions, nonprofits, and political figures to integrate the LIVE program into standard curricula across the US.
- Co-organized, marketed, and presented at 112 public events in 63 cities, driving a 3x growth in customer acquisition and generating over $500,000 in revenue in three months.
- Collaborated with international NGOs to launch a Philippines office and expanded program adoption to top-tier schools in Saudi Arabia, India, Hong Kong, and Romania.

Software Engineer *Pittsburgh, PA*
Expii *Sept 2022 – May 2023*

- Led a 15-person engineering and design team for the LIVE project, delivering a real-time P2P math education platform with near-zero marginal operating costs.
- Architected AWS infrastructure (DynamoDB, Lambda) to secure sensitive customer data, tripling development velocity and generating $50,000+ in value via automated pipelines.
- Fine-tuned LLaMA and GPT models on mathematical corpora to automate content generation, producing over 15,000 USAMO-level competition problems with solutions.

Founder *Cleveland, OH*
The Purple Hand [Defunct] *Dec 2020 – Aug 2022*

- Directed a team of 6 to deliver human trafficking awareness workshops, securing $20,000 in funding as a registered 501(c)(3) nonprofit.
- Partnered with law enforcement and state task forces to provide career development resources and support systems for survivors.

Publications

ImProver: Agent-based Automated Proof Optimization *Oct 2024*

Riyaz Ahuja, Jeremy Avigad, Prasad Tetali, Sean Welleck

In Proceedings of the Thirteenth International Conference on Learning Representations (ICLR) pp. 29521–29543, 2025.

10.48550/arXiv.2410.04753

Projects

cLean *Nov 2025-Present*

- Engineered a domain-specific language and transpiler in Lean 4 that utilizes macro expansion to convert functional Lean code into a semantic intermediary and compiles to PTX.
- Implemented a CPU simulator, GPU launcher and verification harness to prove functional correctness, race-freedom, and barrier-safety of GPU kernels, and execute them on GPU directly from Lean.

CohesiveGalois *Nov 2025-Present*

- Formalized a minimal 10-axiom system for cohesive ∞-topoi and proved the logical independence of descent axioms and minimality of the axiom set.
- Constructed a counterexample showing the non-conservativity of the Shape functor, establishing that finite covers can arise from infinitesimal refinements.

ImProver^2 *Jun 2025-Present*

- Built a generalized neurosymbolic pipeline that bootstraps Small Language Models (SLMs) using iterative filtering and self-improvement.
- Bootstraps small language models via the RL pipeline to outperform much larger proprietary foundation models on research-level theorems by directly optimizing user-defined structural value functions (e.g., modularity) with a high degree of stability. To be submitted to ICML 2026.

LeanM2 *Apr 2025*

- Developed the first type-consistent interface between Lean 4 and Macaulay2, handling noncomputable dependent types and nested rings.
- Created custom metavariable synthesis tactic for simplified and modular DSL generation.
- Automates Grobner basis ideal membership proofs by retrieving computation witnesses from M2 and reconstructing them as formally verified Lean certificates.

Keraunos *Jan 2024*

- Engineered an autonomous drone swarm capable of formation flight and dynamic target tracking using Arduino and custom flight controllers.
- Modified hardware for voltage rerouting to enable live video transmission; implemented stereoscopic trian-

gulation and YOLOv5 object detection for real-time depth perception and leader-follower navigation.

Whitney-Graustein Theorem *Dec 2023*

- Complete formalization of the Whitney-Graustein theorem in Lean 4, classifying regular closed curves in the plane up to regular homotopy.
- Led the formalization of turning numbers and proving their invariance under regular homotopy.

Baire Category Theorem *Oct 2023*

- Formalized the Baire Category Theorem for complete metric spaces in Lean 4, establishing that countable unions of nowhere dense sets have empty interior.
- Developed foundational topological definitions including meager sets, non-meager sets, and Baire spaces, now integrated into the library.

Orbisol *Sept 2023*

- Developed a high-performance orbital debris tracking engine using C++ and OpenGL, integrating datasets from NASA, ESA, and NORAD.
- Implemented a custom octree spatial partitioning system to optimize collision risk calculations and rendering for thousands of tracked objects. (1st Place, CMUHacks 2023).

Technologies

Languages: Python, Lean, C++, C, Java, C#, JavaScript, Typescript, HTML/CSS, SML, Haskell, Rust, x86 Assembly, Verilog, OCaml, SQL

Technologies: ReactJS, NodeJS, NextJS, CUDA, cuDNN, AWS (S3, Dynamo/RDS, Lambda, EC2), Pytorch, Keras, Langchain/graph/smith, Macaulay2, Matlab, Mathematica, Firebase, MongoDB, Git

References

Dr. Jeremy Avigad: Professor of Philosophy and Mathematical Sciences, Director of the Hoskinson Center of Formal Mathematics (CMU) — avigad@cmu.edu, 412-268-8568

Dr. Prasad Tetali: Professor of Mathematical Sciences, Department Head (CMU) — ptetali@cmu.edu, 412-268-2545

Dr. Sean Welleck: Assistant Professor, Language Technologies Institute (CMU) — swelleck@andrew.cmu.edu

Dr. Po-Shen Loh: Professor of Mathematical Sciences (CMU), CEO of Expii — ploh@cmu.edu, 412-268-2548