

# RESEARCH STATEMENT

Zeyu Zheng

Carnegie Mellon University

zeyuzhen@andrew.cmu.edu

**Overview** For millennia, mathematics has been built upon human intuition, insight, and the gradual development of rigorous proof systems. Throughout this process, new technological paradigms have repeatedly reshaped mathematical practice. When oral traditions and string-and-knot records could no longer sustain increasingly complex ideas, symbolic notation provided a shared language for abstract reasoning, enabling profound structures to be systematically recorded and transmitted.

Thousands of years later, the advent of the computer initiated another profound transformation. Researchers began leveraging large-scale computation for systematic enumeration and experimentation, uncovering patterns, proposing conjectures, and testing their validity. Complex problems in algebra, Euclidean geometry, and number theory became tractable through automated approaches, and century-old challenges like the Four Color Theorem and the Kepler Conjecture were thereby resolved. These developments demonstrated that once new tools are integrated into mathematical practice, they can fundamentally alter both the pace and manner of mathematical discovery.

Today, we may stand at the threshold of yet another transformation. Over the past century, the corpus of human mathematical knowledge has expanded at an unprecedented rate. Formal languages such as Lean4 are now offering a new framework for expressing and verifying this vast system. Coincidentally, artificial intelligence represented by large language models has developed rapidly in recent years, demonstrating surprising potential in creativity and reasoning capabilities. Their combination promises to enable mathematicians, for the first time, to conduct “experiments” at the level of logical propositions, much like experimental scientists: generating candidate conjectures, verifying their feasibility, and gradually revealing new structures and patterns through iterative exploration.

My research is dedicated to addressing both the opportunities and challenges of this frontier. Situated at the intersection of pure mathematics, computer-assisted proof, and artificial intelligence, I seek to explore and develop new paradigms for mathematical discovery.

**Past and Current Research** My work to date demonstrates a deep commitment to both pure mathematics and the technologies that advance mathematical research. In mathematics, I have delved into the field of extremal combinatorics, making contributions to Turán-type problems on graphs and hypergraphs, coding theory, and extremal set theory. Together with my collaborators, I completed the determination of the linear Turán number for linear 3-graphs with at most four edges, initiated the study of the plane-saturation number and generalized different codes, and for the first time in decades, advanced the bounds for the classical Oddtown problem modulo composite numbers.

Alongside these mathematical contributions, I have been actively integrating technology into my mathematical research. My work includes explorations into the application of computational techniques in mathematics, such as using the occupancy method to study the number of independent sets in hypergraphs and applying the flag algebra method to investigate configuration problems in vector spaces over finite fields. These methodological explorations not only yielded new mathematical results but also equipped me with a broad, interdisciplinary perspective.

I have also developed a strong foundation in cutting-edge technology in AI and formalization. I have contributed to the Google DeepMind’s formal conjectures repository, an open source collection of formalized statements of math conjectures in Lean. Furthermore, I proposed the Planner-Prover neural theorem proving paradigm and, with my collaborators, implemented it at industry scale. Our efforts produced BFS-Prover-V2, a state-of-the-art step-level AI theorem prover trained in Lean 4 via a novel reinforcement learning pipeline, now released as an open-source tool to empower the formal math community. By combining the planning capabilities of large language models with the rigor of Lean’s formal verification, we achieved significant breakthroughs on mathematical benchmarks like miniF2F and ProofNet. This approach demonstrates how to combine AI creativity with the reliability of formal systems, providing a promising pathway towards automated mathematical discovery.

**Future Directions** Looking ahead, I intend to remain deeply engaged with the pure mathematics community, closely following problems of central interest and making original contributions. In particular, I aim to further advance the techniques we developed for the Oddtown problem and to pursue several of the classic extremal combinatorics conjectures posed by Erdős and his collaborators. Throughout these explorations, I will continue to reflect on how computational tools can enrich and accelerate the processes of discovery and proof.

In parallel, I will actively follow the rapid advances in artificial intelligence and investigate how they can be harnessed to accelerate mathematical discovery. With neural theorem provers now capable of solving Putnam-level and PhD qualifying exam problems, the frontier has shifted toward even more demanding domains, where data sparsity poses an acute challenge. To push this boundary, I plan to develop large-scale auto-formalization systems that translate natural language mathematics into formal language, and to design more scalable reinforcement learning algorithms for formal proof generation. Also, I will maintain the BFS-Prover series for the community and contribute actively to open-source Lean4 formalization repositories, fostering broader access to tools that integrate AI with formal mathematics.

Together, these efforts aim to contribute to a form of mathematical discovery where the boundary between human insight and computational methods becomes increasingly fluid, yet mutually reinforcing.