Flow-based Extremal Mathematical Structure Discovery

Baran Hashemi

The discovery of extremal structures in mathematics requires navigating vast and nonconvex landscapes where analytical methods offer little guidance and brute-force search becomes intractable. We introduce FlowBoost, a closed-loop generative framework that learns to discover rare and extremal combinatorial structures by combining three components: (i) a geometry-aware conditional flow-matching model that learns to sample high-quality configurations, (ii) reward-guided policy optimization with action exploration that directly optimizes the generation process toward the objective while maintaining diversity, and (iii) stochastic local search (Stochastic Relaxation with Perturbations) for both training-data generation and final refinement. Unlike prior open-loop approaches, such as PatternBoost that either retrains on filtered discrete samples, or AlphaEvolve where they rely on frozen Large Language Models (LLMs) as evolutionary mutation operators, FlowBoost enforces geometric feasibility during sampling, and propagates reward signal directly into the generative model, closing the optimization loop and requiring much smaller training sets and shorter training times, and reducing the required outer-loop iterations by orders of magnitude, while eliminating dependence on LLMs. We demonstrate the framework on four geometric optimization problems: sphere packing in hypercubes, circle packing maximizing sum of radii, the Heilbronn triangle problem, and star discrepancy minimization. In several cases, FlowBoost discovers configurations that match or exceed the best known results. For circle packings, we improve the best known lower bounds, surpassing the LLM-based system AlphaEvolve while using substantially fewer computational resources. For the Heilbronn problem, FlowBoost improves the minimum triangle area approaching the best known numerical values. For sphere packing in dimension \$12\$, our method finds configurations denser than those produced by classical heuristics. To our knowledge, FlowBoost is the first systematic application of flow-based generative models with Reinforcement Learning to extremal structure discovery in pure mathematics. We term this paradigm de novo mathematical structure design. Our results demonstrate that lightweight, domain-specific generative models, when equipped with geometric inductive biases and reward-guided learning, can match or exceed the performance of LLM-based systems at a fraction of the computational cost. The code is available at \url{https://github.com/berczig/FlowBoost}.