

Vasily Ilin

Website: vilin97.github.io | Google Scholar: [Vasily Ilin](#) | Github: [Vilin97](#) | email: vasilin97@gmail.com

Education

University of Washington | 2020 - 2026

Applied Mathematics, PhD, \$15,000 in awards.
Director of the [Math AI Lab](#), \$86,000 in funding.

Boston University | 2015 - 2020

Computer Science, MS.
Mathematics, BA & MA (Magna Cum Laude).

Skills

Technical Skills: Python, Jax, HPC, PyTorch, WandB, SQL, Julia, Java, Lean, cluster computing

CS & Math Skills: AI for Math, deep learning, plasma modeling, data engineering, diffusion, sampling, formalization

Work Experience

AI Intern at Google Cloud | Summer 2025

Improved a Slides2Video model with a multi-agentic APO algorithm. Achieved F1 score of 0.9. Collaborated with two teams in Google DeepMind. Gave 6 research talks. Was acknowledged in a memo to 4 Vice Presidents and 3 Directors.

ML Intern at Google Cloud (Python, SQL) | Summer 2024

Trained XGBoost to predict defective TPUs from HBM ECC telemetry, achieved 80% accuracy and improved test efficiency 4x.

ML Intern at YouTube (Python, C++, SQL) | Summer 2023

Trained model to achieve a 2x Egress/Ingress improvement in a CDN. Proposed and implemented a load balancing algorithm.

Data Engineering Intern at Android, Google (Flume Java, SQL) | Summer 2022

Built a distributed pipeline for Quick Share, from logs to dashboards. Sped up dashboards by 30x using approximate aggregation.

Google Summer of Code (Julia) | Summer 2021 | [code](#), [blog post](#)

Implemented, tested, benchmarked, and optimized algorithms to simulate jump processes. Sped up simulations by 70%.

Leadership

Director of UW Math AI Lab (Lean, LLMs) | 2022-present | [Github org](#), [website](#)

- Three papers from the Lab accepted to ICLR 2026 workshops; three papers submitted to ICML 2026.
- As Director, lead the 60-member Lab: setting research direction, organizing project groups, securing funding and compute.
- Mentored or administrated 63 projects involving 185 undergraduates, 17 graduate mentors, and 10 faculty mentors.
- Built UW's Lean community from zero to 30+ active users; taught Lean to 200+ students and 3 faculty members.
- Contributed nine lemmas to `mathlib` with Lab students. PRs: [19798](#), [19896](#), [19886](#), [32824](#).
- Created the Math AI Seminar and organized 15 talks; delivered invited talks at JMM 2025, Lean workshop, and Google.
- Secured \$86,000 in Lab funding (grants from UW, Google, Nebius, and individual donations).
- Organized Lean Together, a weekly meeting to learn Lean.

Organizer of Math AI, plasma and ML seminars | 2022-present

- Organized the [Math AI Seminar](#); delivered three talks.
- Organized PlaSMoS (plasma simulation and modeling seminar); delivered three talks.
- Helped organize MLJC (machine learning journal club); delivered one talk.

Publications

Semantic Search over 9 Million Mathematical Theorems | 2026 | ICLR workshop | [project page](#), [pre-print](#)

A dataset of 9 million theorems from arXiv, and semantic search outperforming Google and LLM search by 70%.

Learning to Repair Lean Proofs from Compiler Feedback | 2026 | ICLR workshop | [pre-print](#)

A dataset of 260k erroneous Lean proofs, compiler feedback and corrections. Outperforming Goedel-32B in error correction.

A Neural Score-Based Method for Deterministic Collisional Plasma Simulation | 2026 | ICLR workshop

A neural network method for deterministic simulation of collisional plasma, inspired by score-based generative modeling.

DiScoFormer: Plug-In Density and Score Estimation with Transformers | 2026 | ICLR workshop | [pre-print](#)

A symmetry-equivariant KDE-inspired transformer for density and score estimation, significantly outperforming KDE.

CayleyPy-4: AI-Holography. Towards analogs of holographic string dualities for AI tasks | 2026

A holographic duality framework mapping Cayley graph complexity to geometry. New conjectures and AI-assisted algorithms.

Stability of the homogeneous Landau equation in relative entropy and applications to score-based methods | 2025 | [pre-print](#)

A proof of relative entropy stability of the Landau equation with applications to accuracy of score-based particle solvers.

Score-Based Deterministic Density Sampling | 2025 | ICLR workshop, CPAA | [paper](#)

Deterministic sampling of an unnormalized density using on-the-fly score estimation with a neural network.

RealEdit: Reddit Edits As a Large-scale Empirical Dataset for Image Transformations | CVPR, 2025 | [project page](#)

A dataset and diffusion model to perform text-guided image edits. The first global edit model trained on real data.

Transport Based Particle Methods for the Fokker-Planck-Landau Equation | CMS, 2025 | [paper](#)

An algorithm for simulation of plasma using a neural network, inspired by score-based generative modeling.

Community and Mentorship Through the Experimental Lean Lab | JMM, 2025 | [abstract](#)

We share lessons we've learned in building community and mentoring undergraduate research projects in Lean.

Extending `JumpProcesses.jl` for Fast Point Process Simulation with Time-Varying Intensities | JuliaCon, 2024 | [paper](#)

An algorithm to efficiently simulate any point process on the real line with a continuous intensity rate.

Catalyst: Fast Biochemical Modeling with Julia | PLOS Comp Bio, 2024 | [paper](#)

Julia library for modeling and high-performance simulation of chemical reaction networks.