

Fengze Jia

jia.747@osu.edu • xanadum.github.io • github.com/Xanadum • huggingface.co/Xanadom1 • Columbus, OH

PROFILE

Final-year B.S. Mathematics student (Computer Science minor) at The Ohio State University, GPA 3.901/4.0, expected May 2027. Two parallel research tracks: directed topology and linear programming (with Prof. Sanjeevi Krishnan; talk at the Young Mathematicians Conference 2025), and robot learning—first-author workshop paper on compressed vision-language-action policies (IROS 2026 ScaleInfra workshop poster; arXiv:2609.23048), with hands-on real-robot experience.

EDUCATION

The Ohio State University — B.S. Mathematics, Computer Science minor 2024 – May 2027 (expected)
GPA 3.901/4.0
Xi'an Jiaotong-Liverpool University — Financial Mathematics 2022 – 2024
English-medium instruction; transferred to The Ohio State University in 2024

RESEARCH

Anatomy of a Closed-Loop Collapse — Compressed VLA Policy Failure and Repair 2026
Independent research • arXiv:2609.23048 • code and records: github.com/Xanadum/closed-loop-collapse

- Distilled Octo-Base (12→8 transformer layers by expert surgery; 86% of parameters retained) on the Open X-Embodiment mixture. The student passes the model family's own offline validation gates (0.996 / 1.000 teacher-ratios) yet collapses in closed-loop SimplerEnv evaluation on a WidowX pick-and-place task: 0/72 vs. the teacher's 40/72.
- Localized the failure to one stage boundary—graded degradation through early stages (object moved at 90%, grasped at 55% of the teacher's rate), categorical failure at transport (0% in every variant)—and isolated its signature by paired action-trace forensics: a persistent negative, late-heavy z residual.
- Falsified four standard therapies under matched controls (continued training, in-domain offline data, additive command compensation, phase-gated command clamping), excluding additive-bias and single-channel explanations.
- Restored teacher parity by a minimal-pair intervention—substituting half the training stream with 324 successful teacher rollouts from the deployment environment, every other setting held fixed—18/36 vs. 17/36 on held-out episodes, with the trace signature eliminated and a teacher-like perturbation-response profile recovered; the same 8-layer network reaching parity falsifies a capacity ceiling. Pre-committed checkpoint selection, execution-record verification, and a documented novelty search (1,001 candidates); full JAX/Flax pipeline and reproducibility package.

Directed Topology and Linear Programming 2025 – present
With Prof. Sanjeevi Krishnan, Department of Mathematics, The Ohio State University

- Theorem L (universal loose realization, 2026): a five-step cubical gluing construction (R0–R4) realizing every integer matrix A as a quotient of standard d -cubes in the category of precubical sets, whose nonnegative d -cycle cone is exactly $\ker A \cap \mathbb{R}_+^n$. Verified coefficient fidelity, row-space preservation (a single shadow row $r_\rho = -\mathbf{1}^\top A$), and categorical legality via congruence descent ($\partial^2 = 0$). Summary writeup compiled Aug 2026.
- LP for directed cycle detection (YMC 2025): formulated minimal directed cycles as the extremal rays of the positive kernel cone $\ker A \cap \mathbb{R}_+^n$ of a graph's incidence matrix; implemented iterative LP enumeration (SVD kernel basis, barrier-method solves) and contrasted it with Tarjan's DFS; proposed a directed homology dH_k restricted to the positive cone—the higher-dimensional direction that Theorem L subsequently realizes.

Buckeye AutoDrive — Perception Team Fall 2025 – Spring 2026
SAE/GM AutoDrive Challenge II (final year); advisor Prof. Wei-Lun Chao

- Collected closed-course driving data; built and labeled a traffic-light / speed-sign dataset; fine-tuned a lightweight YOLO detector for light/sign recognition. Model and dataset open-sourced on Hugging Face: AutoDrive_TrafficLight_Baseline_V1 and auto_drive_traffice_light.

Reinforcement Learning: PPO and Multi-Policy Gradient Summer 2025
With Prof. Maria Han Veiga, Department of Mathematics, The Ohio State University

- Implemented PPO baselines and an MPG variant in custom environments.

Denoising Potential on MNIST 2025
OSU Mathematics CYCLE program; co-first author with Chunyu Gu, with Daniel Tcheurekdjian

- Applied an energy-based denoising potential (Gaussian-mixture log-sum-exp, due to B. Vose and D. Mixon) as a gradient-ascent denoising stage between a feature extractor and a classifier (trainable step size, Mahalanobis components); >86% test accuracy on noisy MNIST with no overfitting.

PUBLICATIONS AND TALKS

- F. Jia. “Anatomy of a Closed-Loop Collapse: A Causal Case Study of a Compressed VLA Policy.” IROS 2026 Workshop on Building Scalable Infrastructure for Robot Learning (ScaleInfra), poster, Oct 2026. arXiv:2609.23048.
- Talk: “Linear Programming for Directed Cycle Detection.” Young Mathematicians Conference, The Ohio State University, July 30 – Aug 1, 2025.
- Poster: “Application of Denoising Potential on MNIST Dataset” (co-first author). OSU Mathematics CYCLE program, 2025.

SERVICE

- Reviewer, IROS 2026 ScaleInfra Workshop.

EXPERIENCE

Peer-Led Machine Learning Research Project — Columbus, OH

May – July 2025

- Self-directed research in model compression and knowledge distillation.
- Reviewed literature, adapted experimental code, configured software environments, and conducted experiments using shared Lambda Cloud GPU resources.
- Gained hands-on experience with a WidowX robotic platform, including research-oriented data collection and model evaluation.
- Explored model compression methods and troubleshooting workflows for machine-learning experiments.

Branch Officer — Agricultural Bank of China, Nanyang, China

June – July 2023

- Processed debit data and assisted with personal account opening.

SKILLS

Mathematics: topology, linear programming, graph theory, statistics, Markov processes, dynamical systems.

Computing: Python, JAX/Flax, PyTorch, Java; SimplierEnv/SAPIEN, RLDS/TFDS, Weights & Biases; Linux, Git, cloud GPU training (RunPod A100).