

# Yixuan Wang

Ph.D. Student, University of Florida · Expected May 2027 · Gainesville, FL  
wang.yixuan@ufl.edu · nomushroomw.github.io · github.com/noMushroomw · Google Scholar · LinkedIn

## RESEARCH PROFILE

Machine learning researcher who views learning and generative algorithms as dynamical systems, using nonlinear control and Lyapunov analysis to understand them and to improve their algorithms and structure with guarantees on stability, convergence, and safety. First-author papers at ICML (Spotlight), NeurIPS, IEEE CDC, and IFAC World Congress. Seeking research scientist roles starting summer 2027.

## SELECTED RESEARCH CONTRIBUTIONS

- **Effective Model Pruning (ICML 2026 Spotlight).** Given any importance-score vector, sets an adaptive sparsity level via its effective sample size, with a provable loss-change bound. Works with any importance criterion and any architecture: validated on MLPs, CNNs, Transformers, LLaMA (7B–70B), and KANs, and even beyond networks on image pixels, so it applies wherever one must decide how many elements to keep.
- **Riemannian Lyapunov framework (NeurIPS 2026).** Formulated optimization as a closed-loop controlled dynamical system on a Riemannian manifold coupling model parameters with optimizer states; a Lyapunov construction links objective reduction to regulation of internal optimizer dynamics, giving a common description of SGD, Adam, and Lion and a blueprint for designing optimizers around explicit stability objectives.
- **Saturation-aware multi-reward RL (SA-MRPO, arXiv:2608.16072).** Saturated objectives keep absorbing optimization effort despite having no headroom; SA-MRPO reweights each reward's advantage by batch-level saturation to redirect effort to unmastered objectives: +3.8 pts average over GDPO on five math benchmarks (up to +9.2 on AMC23).
- **Graph structured flow matching (GSFM, arXiv:2609.32231).** To predict diverse human motion that respects the body's articulated structure, transports the whole future skeleton trajectory through one conditional velocity field, preserving bone lengths by construction. A compact 4.47M-parameter model reaches near-best AMASS precision and transfers zero-shot to the unseen Human3.6M skeleton, beating every retargeting baseline.
- **Bayesian intent inference and privacy (IFAC 2026, IEEE CDC 2026).** Built a Rao-Blackwellized particle filter that infers an agent's goal online from noisy trajectories, and control policies that regulate the observer's belief to bound goal-information disclosure via probabilistic control barrier functions.

## TECHNICAL SKILLS

|                  |                                                                                      |
|------------------|--------------------------------------------------------------------------------------|
| Machine Learning | PyTorch; generative modeling (EBMs, diffusion, flow matching); RL; model compression |
| ML Systems       | Multi-GPU DDP training on NVIDIA B200 GPUs; Slurm-based HPC (UF HiPerGator)          |
| Theory           | Lyapunov analysis, Riemannian optimization, constrained generation                   |
| Programming      | Python, C/C++, MATLAB / Simulink, Git, LaTeX                                         |

## SELECTED PUBLICATIONS

|                                                                                                                        |                            |
|------------------------------------------------------------------------------------------------------------------------|----------------------------|
| 1. Effective Model Pruning: Measuring the Redundancy of Model Components                                               | ICML 2026 <b>Spotlight</b> |
| 2. Riemannian Lyapunov Framework: Optimization as Closed-Loop Control on Riemannian Manifolds                          | NeurIPS 2026               |
| 3. Off-Dynamics Reinforcement Learning via Domain Adaptation and Reward Augmented Imitation                            | NeurIPS 2024               |
| 4. Learn What's Left, Not What's Mastered: Saturation Aware Advantage Reweighting for Multi-Reward Policy Optimization | Preprint                   |
| 5. Skeletons in Flow: Graph Structured Flow Matching for Human Motion Prediction                                       | Preprint                   |

## PROFESSIONAL SERVICE

Reviewer: NeurIPS 2025–2026; ICML 2026; ICLR 2026

## EDUCATION

University of Florida — Ph.D., Mechanical & Aerospace Engineering 2024 – 2027 (exp.)  
M.S. CS, Johns Hopkins University (2024); M.S. ME, University of Florida (2022); B.S. ME, Tongji University (2020)