

Yixuan Wang

Ph.D. Student, Mechanical & Aerospace Engineering · University of Florida · Expected May 2027

Gainesville, FL · wang.yixuan@ufl.edu · nomushroomw.github.io · Google Scholar · GitHub

RESEARCH INTERESTS

I study machine learning systems, especially generative models, through the lens of nonlinear control and dynamical systems. Using Lyapunov analysis, geometry, and optimization, I analyze their behavior and underlying manifold structure, and turn this understanding into new algorithms and model structures with guarantees on stability, convergence, and safety.

EDUCATION

University of Florida — Ph.D., Mechanical & Aerospace Engineering 2024 – Present

Nonlinear Controls and Robotics (NCR) Lab · Advisor: Prof. Warren E. Dixon · Coadvisor: Prof. Dan P. Guralnik

Johns Hopkins University — M.S., Computer Science 2023 – 2024

Research on off-dynamics reinforcement learning with Prof. Anqi Liu's group, leading to a NeurIPS 2024 paper

University of Florida — M.S., Mechanical Engineering 2020 – 2022

Tongji University — B.S., Mechanical Design, Manufacture and Automation 2015 – 2020

SELECTED PEER-REVIEWED PUBLICATIONS

- Effective Model Pruning: Measuring the Redundancy of Model Components**
Yixuan Wang, Dan P. Guralnik, Saiedeh Akbari, Warren E. Dixon
ICML 2026 Spotlight · arXiv:2509.25606
- Riemannian Lyapunov Framework: Optimization as Closed-Loop Control on Riemannian Manifolds**
Yixuan Wang, Omkar Sudhir Patil, Warren E. Dixon
NeurIPS 2026 · arXiv:2601.22284
- Quantifying Trade-Offs Between Stability and Goal-Obfuscation**
Yixuan Wang, Dan P. Guralnik, Warren E. Dixon
IEEE Conference on Decision and Control (CDC) 2026 · arXiv:2605.06630
- Rao-Blackwellized particle filter for agent's intention inference**
Yixuan Wang, Dan P. Guralnik, Warren E. Dixon
IFAC World Congress 2026 · arXiv:2512.09269
- Off-Dynamics Reinforcement Learning via Domain Adaptation and Reward Augmented Imitation**
Yihong Guo, **Yixuan Wang**, Yuanyuan Shi, Pan Xu, Anqi Liu
NeurIPS 2024 · arXiv:2411.09891

PREPRINTS & MANUSCRIPTS

- Skeletons in Flow: Graph Structured Flow Matching for Human Motion Prediction**
Yixuan Wang, Brandon C. Fallin, Warren E. Dixon
arXiv:2609.32231
- Learn What's Left, Not What's Mastered: Saturation Aware Advantage Reweighting for Multi-Reward Policy Optimization**
Yixuan Wang, Yifei Chen, Haichao Zhang, Haozheng Luo, Xander Wu, Jie Ni, Yun Fu, Nuno Vasconcelos, Yijiang Li
arXiv:2608.16072
- Energy Generative Modeling: A Lyapunov-based Energy Matching Perspective**
Yixuan Wang, Wenqian Xue, Warren E. Dixon
arXiv:2605.05530
- Lyapunov-Based Kolmogorov-Arnold Network (KAN) Adaptive Control**
Xuehui Shen, Wenqian Xue, **Yixuan Wang**, Warren E. Dixon
arXiv:2512.21437

RESEARCH EXPERIENCE

University of Florida — Graduate Research Assistant, NCR Lab2024 – Present

Advisor: Prof. Warren E. Dixon · Coadvisor: Prof. Dan P. Guralnik

Common thread: view learning, generation, and inference algorithms as dynamical systems, and use tools from nonlinear control, Lyapunov theory, and Bayesian estimation to understand and redesign them.

- **Optimization as feedback control.**
Formulated optimization as a closed-loop controlled dynamical system on a Riemannian manifold coupling model parameters with internal optimizer states; a Lyapunov construction links objective reduction to regulation of the optimizer's internal dynamics, giving a common description of SGD, Adam, and Lion and a blueprint for designing optimizers around explicit stability objectives (NeurIPS 2026).
- **Generation as controlled transport.**
Developed graph structured flow matching that transports the entire future skeleton trajectory through one conditional velocity field while 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 (preprint). Cast energy-based sampling as a control problem with a Kullback–Leibler Lyapunov function, yielding a finite-step stopping rule for Langevin sampling (preprint).
- **Belief dynamics in inference and privacy.**
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 (IFAC 2026; IEEE CDC 2026).
- **Allocating model capacity and optimization effort.**
Given any importance-score vector, set an adaptive sparsity level via its effective sample size, with a provable loss-change bound; the method works with any importance criterion and architecture, validated on MLPs, CNNs, Transformers, LLaMA (7B–70B), and KANs, and even beyond networks on image pixels (ICML 2026 Spotlight). Reweighted each reward's advantage by batch-level saturation so multi-objective policy optimization stops spending effort on saturated objectives: +3.8 pts average over GDPO on five math benchmarks (preprint).

RESEARCH PROPOSAL DEVELOPMENT

Contributed to advisor-led proposals funded by the National Science Foundation (NSF), Air Force Office of Scientific Research (AFOSR), and Air Force Research Laboratory (AFRL).

- **Research aim development.** Led the development and writing of 1-2 research aims per proposal, from research questions and technical innovations to technical approaches and research plans.
- **Preliminary research and proposal authorship.** Conducted preliminary experiments and authored the corresponding proposal sections to establish the proposed aims' motivation and feasibility.

PROFESSIONAL SERVICE

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

TECHNICAL SKILLS

Machine Learning	PyTorch; generative models (energy-based, diffusion, flow matching); reinforcement learning; model compression
ML Systems	Multi-GPU distributed training (DDP) on NVIDIA B200 GPUs; Slurm-based HPC workflows (UF HiPerGator)
Theory	Nonlinear & adaptive control, Lyapunov analysis, dynamical systems, optimization on manifolds, Bayesian estimation & particle filtering
Programming Languages	Python, C/C++, MATLAB / Simulink, Git, LaTeX English (proficient), Mandarin (native), German (B2)