Adrian Rodriguez

axr190042@utdallas.edu | axr2718.github.io | axr2718 | arod2718

Research interests: energy-based models for reasoning, planning, and causal world models.

Education

The University of Texas at Dallas
Master of Science in Computer Science (GPA: 3.75/4.0) — Expected May 2027
Thesis: *Learning Time-Invariant Energy Landscapes for Compositional Reasoning and Planning*
Advisor: Yunhui Guo

Bachelor of Science in Computer Science (GPA: 3.45/4.0) — Aug 2020 – May 2025

Northeast Texas Community College
Associate of Science in Physics (GPA: 3.89/4.0) — Aug 2015 – Dec 2018

Publications

Adrian Rodriguez, Angelica Kim, Yunhui Guo, and Yilun Du (2026). "Unifying Reasoning and Planning through Energy Minimization". In: *Advances in Neural Information Processing Systems (NeurIPS)*. [project] [code].

Adrian Rodriguez, Benjamin Hoover, Yunhui Guo, and Yilun Du (2026). "Generative Associative Memory via Equilibrium Matching". In: *ICLR 2026 Workshop on New Frontiers in Associative Memory*. [paper].

Sarthak Kumar Maharana, Saksham Singh Kushwaha, Baoming Zhang*, **Adrian Rodriguez***, Songtao Wei*, Yapeng Tian, and Yunhui Guo (2025). "AVROBUSTBENCH: Benchmarking the Robustness of Audio-Visual Recognition Models at Test-Time". In: *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*. [paper] [project] [code].

Brighton Nuwagira, **Adrian Rodriguez**, Qiwei Li, and Baris Coskunuzer (2025). "Breast Cancer Detection with Topological Deep Learning". In: *2025 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. [paper].

In preparation

Adrian Rodriguez*, Junxuan Shen*, Alex Oarga*, and Yilun Du (2026). "Compositional Energy-based Models for Reasoning and Planning". Manuscript in preparation.

*Equal contribution.

Research Experience

Embodied Minds Lab, Harvard University – *Visiting Researcher* — Dec 2025 – Present

- Led the development of Energy Minimization Search (EMS) with Yilun Du. On graph coloring, 8-Queens, 3-SAT, and crosswords, EMS matched or beat DiffEBM + PEM on each task's main metric while using about 52–10,700× fewer model evaluations. On graph coloring, it produced 86% fewer conflicts.
- Trained Energy Transformers with Equilibrium Matching for generative associative memory. With fixed positional encodings, increasing depth from one to eight blocks lowered CIFAR-10 FID from 75.83 to 28.56.
- Running experiments for a co-first-authored manuscript that compares compositional and full-instance training for reasoning and planning, contrasts diffusion and time-independent training objectives, and evaluates methods that learn the compositional structure from data.

Data Efficient Intelligence Learning Lab, UT Dallas – *Student Researcher* — Nov 2024 – Present

- Conducting MS thesis research on learning time-invariant energy landscapes for compositional reasoning and planning, advised by Yunhui Guo.
- Built AudioSet-2C for AVRobustBench by applying the benchmark's 15 audio-visual corruptions to AudioSet, and evaluated each audio-visual model on it. Was solely responsible for the public code release.

Cognitive Informatics and Statistics Lab, UT Dallas – *Student Researcher* May 2024 – Present

– Built a working open-source implementation of Richard M. Golden's algorithm for compressing neural networks through weight pruning.

– Investigating whether Fisher information can replace simulation in statistical inference, in collaboration with Dr. Golden.

Topological Machine Learning Group, UT Dallas – *Student Researcher* Oct 2024 – May 2025

– Implemented Topo-VT with Baris Coskunuzer. This SwinV2 transformer combines ultrasound features with Betti-vector embeddings through cross-attention to detect breast cancer. It raised SwinV2's accuracy by 1.70 percentage points on average across four datasets (10.09 on BUS-BRA) and beat every CNN baseline on two.

Industry Experience

Ensemble AI – *Machine Learning Research Intern* May 2025 – Sep 2025

– Reimplemented Qwen3 with NdLinear layers for model compression. Built a pipeline for pretraining on FineWeb, fine-tuning, and reinforcement learning, and evaluated the resulting models on HellaSwag.

Deca Defense – *Machine Learning Research Intern* Dec 2024 – Apr 2025

– Integrated open multimodal language models into drone-simulator pipelines to control simulated drones, and investigated model compression for deployment on resource-constrained hardware.

Awards

Lambda Research Grant Oct 2026
$2,000 in GPU credits for *Unifying Reasoning and Planning through Energy Minimization.*

Google Research Grant Jan 2026
$5,000 in GPU credits for MS thesis research on energy-based models for reasoning and planning.

Lambda Research Grant Dec 2025
$2,000 in GPU credits for energy-based modeling research.

Teaching and Mentoring

Course materials, The University of Texas at Dallas 2026 – 2027
Developing PyTorch and Keras tutorials for Richard M. Golden's undergraduate courses in Intelligent Systems Analysis (CS/CGS 4314) and Intelligent Systems Design (CS/CGS 4315), and graduate courses in Neural Net Math (ACN 6348) and Statistical Machine Learning (ACN 6349).

Student mentoring, Embodied Minds Lab
Mentored two Harvard University students in energy-based models for reasoning and planning.

– **Angelica Kim**, MS in Data Science 2026
Co-author of the NeurIPS 2026 paper on Energy Minimization Search.

– **Archit Tamhane**, BS in Computer Science 2026
Energy-based models for recursive reasoning.

Technical Skills

Languages: Python, C++, R, Bash
Machine Learning: PyTorch, JAX, Hugging Face Transformers, Keras, scikit-learn
Tools: Git, Docker, Weights & Biases