Brett Evan Barkley

bbarkley@utexas.edu | https://bebark.github.io/ | Graduating: 12/2026

Research Interests

Improving the reliability and scalability of synthetic data by understanding failure modes, developing diagnostic frameworks, and designing principled augmentation strategies across reinforcement learning, generative models, and foundation models.

Education

The University of Texas at Austin, PhD — **2022 - Present**
Electrical and Computer Engineering — **GPA:** 3.95

University of Maryland, M.S — **2015 - 2017**
Aerospace Engineering — **GPA:** 3.97

University of Maryland, B.S — **2010 - 2015**
Aerospace Engineering (Honors Program) — **GPA:** 4.00

Professional Experience

Research Scientist Intern, Waymo — 2026
AI Foundations Team

- Applying supervised fine-tuning (SFT) with LoRA to adapt multimodal foundation models for identifying and evaluating long-tail driving scenarios.
- Leveraging these models to develop scalable data curation approaches that reduce reliance on costly manual annotation and improve training data efficiency; research contribution targeted for publication.

Autonomy Aerospace Engineer, Johns Hopkins University Applied Physics Lab (JHU/APL) — 2017 - 2022

- Efforts culminated in the first ever combat tests between AI and human-piloted F-16s in 2023
- JHU/APL's Air Combat Evolution (ACE) deep reinforcement learning (DRL) lead for sub and full-scale aircraft
- Guidance, control, and aerospace simulation subject matter expert (SME) for JHU/APL ADT and ACE teams

Selected Research Projects

Scaling Pretrained Representations Enables Label-Free OOD Detection Without Fine-Tuning — 2026
First-author Paper in Submission & Google Sponsored Research Grant Lead

- Demonstrated across 59 vision and language backbone-task pairings that scaling frozen pretrained representations enables strong, task-agnostic, label-free OOD detection without fine-tuning. Paper.
- Showed that representation scaling reduces the importance of OOD detector complexity, with simple global and local estimators converging as backbone quality increases.
- Invited speaker at Google's AI Systems DevLabs, presenting this framework to Google's JAX, ML systems, and AI infrastructure teams.
- Led a 1-year **$100k Google Sponsored Research Grant** as Principal Investigator, directing research on scalable foundation-model OOD detection.

A Forensic Analysis of Synthetic Data in Reinforcement Learning — 2025
First-author Paper in Submission

- Developed a novel framework (FTFL) to diagnose and fix synthetic data failures in Model-Based Policy Optimization (MBPO). Paper.
- Identified two systemic issues - output–target scale mismatch and residual prediction variance inflation - and introduced lightweight remedies.
- Restored performance where MBPO previously collapsed, surpassing model-free baselines (SAC) in 5 of 7 DeepMind Control Suite tasks while maintaining strong results in OpenAI Gym.

SCOPED: Score–Curvature Out-of-distribution Proximity Evaluator for Diffusion — 2025
First-author ICLR 2026 Paper

- Developed a novel framework for out-of-distribution (OOD) detection that reduces inference cost by an order of magnitude compared to prior diffusion-based methods. Paper.
- Demonstrated strong generalization across vision and reinforcement learning benchmarks without retraining.
- Designed for deployment at scale, with fully parallelizable probes and low overhead suitable for real-time or compute-limited settings.
- Independently pitched this framework to leadership at NomadicML (a physical AI startup), securing a **$50k research grant** to scale and deploy these diffusion diagnostics for video OOD detection.

Stealing That Free Lunch: Exposing the Limits of Dyna-Style Model-Based Reinforcement Learning 2024 - 2025
First-author ICML 2025 Paper

- Developed a JAX-based training pipeline yielding up to a $40\times$ speedup in wall-clock time for MBPO. Code. Paper.
- Showed that model-based RL algorithms like MBPO exhibit strong performance in OpenAI Gym but often fail in DeepMind Control Suite tasks when trained from scratch.
- Investigated several potential explanations for this discrepancy, including model error and mitigation strategies; found that even modern techniques fail to close the gap.
- Presented framework optimizations and research findings to the Google JAX team, independently initiating the collaboration and authoring the successful proposal for a **$100k Google Sponsored Research Grant**.

Technical Skills

Languages: Python, C++, Cython, Bash, CUDA
Libraries/Software: JAX, Pytorch, Flax, Brax, Git, LaTeX

Research Service

- Reviewer: ICML 2025, ICRA 2025, IROS 2025, L4DC 2026, ICML 2026, ICLR 2026, NeurIPS 2026
- Gold Reviewer (Top 25%): ICML 2026

Selected Publications

1. Scaling Pretrained Representations Enables Label-Free Out-of-Distribution Detection Without Fine-Tuning
Brett Barkley, Preston Culbertson, David Fridovich-Keil | *In submission*
2. A Forensic Analysis of Synthetic Data in RL
Brett Barkley, David Fridovich-Keil | *In submission*
3. SCOPED: Score–Curvature Out-of-distribution Proximity Evaluator for Diffusion
Brett Barkley, Preston Culbertson, David Fridovich-Keil | ICLR 2026
4. Stealing That Free Lunch: Exposing the Limits of Dyna-Style Reinforcement Learning
Brett Barkley, David Fridovich-Keil | ICML 2025
5. An Investigation of Time Reversal Symmetry in Reinforcement Learning
Brett Barkley, Amy Zhang, David Fridovich-Keil | L4DC 2024