

Saket Tiwari

Website: saketirl.github.io

Google Scholar: [Saket Tiwari](#)

Email: sakett786@gmail.com

Mobile: +1-630-550-6016

EDUCATION

- **Brown University** Providence, RI
Ph.D. in Computer Science, advised by George Konidaris *Structure in High-Dimensional RL*
- **University of Massachusetts Amherst** Amherst, MA
Master of Science in Computer Science
- **Indian Institute of Technology Bombay** Mumbai, India
Bachelor of Technology in Computer Science with honors

PUBLICATIONS

- **Reinforcement Learning**
 - *Deep RL theory built to inform and improve methods for continuous state-action problems in a prescriptive manner.*
 - **Stiefel Manifold Optimization Accelerates Reinforcement Learning:** We show that RL with Stiefel manifold optimization has an implicit low-dimensional structure improving robotic control from pixels with a co-designed architecture and LLM post-training efficiency. **Under Review at Neurips 2026** [Link](#).
 - **Geometry of Neural Reinforcement Learning in Continuous State and Action Spaces:** We prove the *manifold hypothesis* for continuous control using neural networks and provide a neural network architecture that uses this insight for sample efficient actor-critic in high-dimensional robotic environments. Saket Tiwari, Omer Gottesman, and George Konidaris **ICLR Oral 2025, top 1.8%**. [Link](#).
 - **From Ticks to Flows: Dynamics of Neural Reinforcement Learning in Continuous Environments:** We use control theory principles to derive the effective non-parametric learning dynamics for Deep RL in continuous environments with application to robot learning. Saket Tiwari, Tejas Kotwal, George Konidaris **ICLR 2026**, [Link](#).
 - **Effects of Data Geometry in Early Deep Learning:** Connects deep learning theory to practical continuous datasets under the manifold hypothesis by bounding the density of linear regions. Saket Tiwari and George Konidaris. **NeurIPS 2022**. [Link](#)
 - **Spectral Collapse Drives Loss of Plasticity in Deep Continual Learning:** We study the trainability of deep NNs for continual reinforcement learning and a causal link between its “bulk” and neural plasticity for LLMs, and feed forwards networks. Naicheng He, Kaicheng Guo, Arjun Prakash, Saket Tiwari, Ruo Yu Tao, Tyrone Serapio, Amy Greenwald, George Konidaris **ICML 2026**. [Link](#)
 - **Meta Learning Parameterized Skills:** Introduced an algorithm that learns transferable skills for robotic control in simulation and uses a novel distance metric for skills and trajectories. Haotian Fu, Shangqun Yu, Saket Tiwari, Michael Littman, George Konidaris, **ICML 2023** [Link](#)
 - **Natural Option Critic:** Derived and demonstrated the advantages of a *natural-gradient* based learning algorithm on Atari environments with hierarchical RL in the options framework. Saket Tiwari and Philip S. Thomas, **AAAI 2019, spotlight**. [Link](#)
 - **A Domain-Agnostic Approach for Characterization of Lifelong Learning Systems:** I characterized the transferability of neural network policies and their features for continual learning in home navigation tasks from pixels (AI Habitat environment). Implemented a novel incremental natural actor-critic framework in the continual learning setting. Multiple authors. **Neural Networks Journal, 2023** [Link](#)
 - **Diagnosing and Mitigating Reward Collapse in GRPO for Text-to-Text Regression via Advantage Variance:** Designing a new reward function for regression using LLMs. Nour Abdelmoneim, Mehak Preet Dhaliwal, Andong Hua, Saket Tiwari, Wenbo Guo, Yao Qin **Under Review at ACL 2026**
- **Neural Networks and Causal Inference:**
 - **Neural Networks for Cache Miss-Rate Prediction:** Designed a deep neural network to predict cache miss rates in benchmark programs, inspired by the WaveNet model for audio generation. Saket Tiwari, Arjun K., Rishikesh Jha, and Eliot Moss, **NeurIPS 2018 Workshop on ML for Systems** [Link](#).
 - **Economic Status and Online Behavior: Causal Inference in Content Networks** Presented a model to predict causal effects of economic status on online behavior in a social-network setting with unobserved confounders and numerous covariates. Presentation at the **Conference on Network Science and Economics**, Saket Tiwari, Aarushi Kalra, and Debdoot Mukherjee.
 - **A Dataset of Geolocated Villages and Gram Panchayat Election Candidates in Uttar Pradesh** A data set for the causal analysis of election results. Aryan Srivastava, Aarushi Kalra, Saket Tiwari [Link](#).

WORK EXPERIENCE

- **University of California, Santa Barbara** Santa Barbara, CA
Post-Doctoral Researcher with Prof. Yao Qin Feb 2026 (Ongoing)
 - **Continual Reinforcement Learning:** Developing a geometric perspective for optimization in continual robot learning using world models for manipulation using a robot arm.
 - **Predicting AI Agent Behavior:** Insulin delivery for diabetes management using an RL controller.
- **Data Science for the Common Good** Amherst, MA
Machine Learning Research Intern
 - **Causality Detection in Social Networks:** Researched the causal effects of user connections in social networks using deep variational auto-encoders on a dataset with millions of connections.
- **LinkedIn** Bangalore, India
Senior Machine Learning Engineer
 - **Video Classification at Scale:** Led the video-classification project. Trained deep-learning models and deployed them to production, handling more than 1,000 FPS using asynchronous and distributed processing on GPUs.
 - **Hierarchical Text Embeddings:** Developed distributed training for hierarchical neural auto-encoders by introducing a neural-machine-translation type architecture in TensorFlow, achieving significant speed-up.
 - **Company-User Recommendation Engine:** Applied a Thompson-sampling based explore-exploit model.
- **LinkedIn** Bangalore, India
Machine Learning Engineer
 - **Spam Filtering:** Designed, developed and maintained a high-QPS service that classified all content created on LinkedIn (millions of entities daily) to filter spam. Designed training pipelines for machine-learning models, deployment, and synchronous and asynchronous classification.
 - **Ads Auto-Approval:** Turned a hack-day project that auto-classified advertisements on LinkedIn using text classification into a full fledged ads auto-approval pipeline, leading to a reduced wait time for ads to go live.
- **TimeMyTask** Mumbai, India
Co-Founder
 - **Co-Founder:** Employed more than 50 people; raised capital and achieved positive EBITDA within the first six months. Introduced health insurance and self-defense classes for employees.
- **Amazon** Hyderabad, India
Software Engineering Intern
 - **Eclipse Plugin:** Worked on the Marketplace team to fix a common issue synchronizing source code across geographically distributed machines, improving developer productivity. Developed an IDE plugin.

RELEVANT SKILLS

- **Agentic AI Systems and Frameworks:** Experienced with Claude Code, Hermes Agent, OpenAI Codex, Cursor Agent, GitHub Copilot Agent Mode.
- **Robotic Frameworks and Software:** ROS/ROS 2, Isaac Sim, MuJoCo, ManiSkill, Hugging Face LeRobot, Franka Software Stack, libfranka, Moveit, Gymnasium-Robotics
- **Cloud Based Training:** Proficient at multi-gpu training and cloud GPU tools.

RELEVANT COURSEWORK

- **Graduate:** Reinforcement Learning; Robotics; Non-Linear Dynamical Systems I and II; Stochastic Optimal Control; Probabilistic Graphical Models; Advanced Algorithms; Reintegrating AI; Machine-Learning Theory; Real Analysis I and II; Artificial Intelligence; Operating Systems; Program Derivation;

SCHOLARSHIPS AND ACHIEVEMENTS

- **Tilde Research Fellowship:** Received 10,000\$ worth of GPU compute credits for research on continual learning.
- **All-India Rank 116:** Ranked 116 out of 0.5 million people who appeared for the Joint Entrance Examination in 2010.
- **KVPY Scholarship:** Awarded the KVPY fellowship by the Government of India, granted to exceptionally motivated students to pursue research in the sciences. Attended a four-day conference at IISc Bangalore.
- **Sri Chaitanya Scholarship:** Received a 100% fee waiver for academic performance.