

JONATHAN BALLOCH, PHD



PhD-trained AI engineer specializing in learned autonomy and world models for robotics with a track record of shipping production AI systems to real hardware at scale.

WORK EXPERIENCE

Senior Software Engineer

Anduril Industries, Atlanta
04/2025 - Present

- Owned end-to-end design, development, testing strategy, and flight stack integration of a search autonomy system now deployed on dozens of assets and shipped to customers at scale
- Founded and led a new small UAS autonomy team; designed and shipped a learned search capability to a real platform within 6 months using hybrid of RL, imitation learning, and classical planning.
- Led technical design of perception, autonomy, and simulation systems across fixed-wing, multirotor, and rotary UAS platforms; drove integration from R&D prototype to production-ready software
- Leading productization and robustness efforts on the perception team, bridging the gap between research-level output and deployable flight software
- Work across the full stack: real-time C++ flight systems, model training, and simulation infrastructure

Graduate Research Assistant

Georgia Institute of Technology, Atlanta
08/2016 - 12/2024

- Thesis: Adapting reinforcement learning agents to changes in the environment.
Led teams of research engineers and undergrad, masters, PhD students in development of large-scale distributed training projects for reinforcement learning agents, supporting dozens of concurrent training workers across multiple GPUs on a SLURM cluster. Key development and research achievements:
- Production-grade ML training pipelines using libraries like PyTorch Lightning and JAX, improving training throughput in some cases by as much as 11x over single machine
 - Trained and evaluated world models and RL policies across MuJoCo, Habitat, and IsaacSim, with a focus on sim-to-real transfer and generalization under environment change.
 - Reduced training time for RL models by 65% through exploration and architectural optimization
 - Built and maintained multiple ML libraries with CI/CD pipelines, used by 20+ researchers
- Additionally worked on projects with on LLM-enabled planning agents and computer vision research.

PhD Intern

SRI International, Palo Alto
05/2021-12/2021

- Designed and implemented novel neural behavior trees approach for interpretable MA-PPO policies in multi-agent reinforcement learning, trained at-scale with Kubernetes on multi-gpu clusters.
- increased agent success by > 15% and convergence speed by 2x on LaserTag and internal multi-agent wargaming simulators.

PhD SWE Intern

Google, Seattle
05/2018-08/2018

- Developed scalable few-shot learning approach using active online sample selection for deep learning.
- Accelerated model training with distributing training across 200+ of TPUs and GPUs.
 - Achieved a 2x sample efficiency improvement in classification tasks (MNIST & CIFAR10).

Robotics Engineer

Intelligent Automation, Rockville, MD
07/2013-07/2016

- Specialized in design and development of computer vision, sensor fusion, and control systems in DARPA and DoD robotics research and development projects, collaborating with industry and academic groups. Developed experience with systems like ROS and strong familiarity with CI/CD infrastructure required to deliver robust software to customers. Key achievements:
- Optimized the C++ code for our realtime video feature tracking to <40ms latency and increase tracking persistence to 92%
 - Designed and implemented the all components for an omnidirectional 4 camera + 9-DoF IMU sensor head that provided 360 deg computer vision inputs at over 100 fps
 - Deployed dual-arm mobile manipulator full-body controller for DARPA Robotic Challenge Trials Expo

EDUCATION

PhD - Computer Science (Robotics)

Georgia Institute of Technology, Atlanta
08/2016 - 12/2024

M.S. - Robotics

University of Pennsylvania, Philadelphia
08/2011 - 12/2013

B.S. - Physics, Mathematics

Georgetown University, Washington, D.C.
08/2007 - 12/2011

CONTACT

- jon.balloch@gmail.com
- (203) 912-1793
- Atlanta, GA
- jballoch.com
- @balloch
- Jonathan Balloch
- 0000-0001-9998-8415
- Google Scholar

SKILLS

Expertise

- AI & Machine Learning (RL, CNN, LLM, RNN, VAE, +more)
- Software Engineering (OOP, Git, CI/CD, Debugging & Testing)
- Physics, Math & Statistics Engineering
- Entrepreneurship

Software & Tools

- Deep Learning Frameworks (PyTorch, JAX, TensorFlow, ONNX)
- Infrastructure & Scaling (Docker & Kubernetes, SLURM)
- SciPy Data Stack (numpy, pandas, matplotlib)
- Cloud & Data (AWS, GCP, SQL, Redis)

Programming

- Python
- Bash
- C++
- Matlab
- Java
- LaTeX

Operating Systems

- Linux
- MacOS
- Windows

Languages

- English
- Russian
- Spanish

SELECT PUBLICATIONS

Concept Bottleneck World Models

 **Balloch, J. C.**, Eagan, K., Loza, R., Riedl, M.
 2025  Under Review

Dual Objective Priority Sampling in Model-based Reinforcement Learning

 **Balloch, J. C.**, Riedl, M.
 2025  Under Review

Characteristics of Effective Exploration for Transfer in Reinforcement Learning

 **Balloch, J. C.**, Bhagat, R., Zollicoffer, G., Jia R., Kim, J., Riedl, M.
 2024  Reinforcement Learning Conference (RLC): Finding the Frame Workshop

Neuro-Symbolic World Models for Adapting to Open World Novelty

 **Balloch, J. C.**, Lin, Z., Peng, X., Hussain, M., Srinivas, A., Wright, R., Kim, J.M. and Riedl, M.
 2023  Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 2848-2850.

NovGrid: A Flexible Grid World for Evaluating Agent Response to Novelty

 **Balloch, J. C.**, Lin, Z., Hussain, M., Srinivas, A., Wright, R., Peng, X., Kim, J., Riedl, M.
 2022  AAAI Spring Symposium on Designing Artificial Intelligence for Open Worlds [Long Oral]

Memory-efficient semi-supervised continual learning: The world is its own replay buffer

 Smith, J., **Balloch, J. C.**, Hsu, Y. C., Kira, Z.
 2021  International Joint Conference on Neural Networks (IJCNN), IEEE

Always be Dreaming: A new approach for data-free class-incremental learning

 Smith, J., Hsu, Y. C., **Balloch, J. C.**, Shen, Y., Jin, H., and Kira, Z.
 2021  In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 9374-9384.

Tell Me A Story Like I'm Five: Story Generation via Question Answering

 Castricato, L., Frazier, S., **Balloch, J. C.**, Riedl, M.
 2021  Proc of the 3rd Workshop on Narrative Understanding

Detecting and Adapting to Novelty in Games

 Peng, X., **Balloch, J. C.**, Riedl M.
 2020  AAAI Workshop on Reinforcement Learning in Games

Taking Recoveries to Task: Recovery-Driven Development for Recipe-based Robot Tasks

 Banerjee, S., Daruna, A., Kent, D., Liu, W., **Balloch, J. C.**, Jain, A., Krishnan, A., Chernova, S.
 2019  IEEE International Symposium on Robotics Research (ISRR)

The MacGyverbot: Tool Construction by Autonomous Agents

 Nair, Lakshmi, **Balloch, J. C.**, Chernova, S. “.”, 2019.
 2019  IEEE International Conference on Robotics and Automation (ICRA)

Unbiasing Semantic Segmentation for Robot Perception using Synthetic Data Feature Transfer

 **Balloch, J. C.**, Aggraval, V., Essa, I., Chernova, S.
 2018  ArXiv:1809.03676

An RGBD segmentation model for robot vision learned from synthetic data

 **Balloch, J. C.**, Chernova, S.
 2017  Robotics Science and Systems (RSS): Workshop on Spatial-Semantic Representations in Robotics

Landmark-Based Robust Navigation for Tactical UGV Control in GPS-Denied Communication-Degraded Environments

 Endo, Y., **Balloch, J.**, Grushin, A., Lee, M.W., Handelman, D.
 2016  SPIE Unmanned Systems Technology XVIII

The Evolution of Titan's detached haze layer near equinox in 2009

 West, R. A., **Balloch, J.**, Dumont, P., Lavvas, P., Lorenz, R., Rannou, P., Turtle, E. P., Ray, T.
 2011  Geophysical Research Letters, 38, doi: 10.1029/2011GL046843