

Ryan Sullivan

rsulli@umd.edu

<https://ryannavillus.github.io/>

973-738-7415

Profile

Applying to research internships for reinforcement learning. Currently focusing on automatic curriculum learning, generalizable exploration methods, and open-ended learning.

Education

University of Maryland at College Park | GPA: 3.8 / 4.0 | August 2020 – Present

- Pursuing a Ph.D. in Computer Science, researching reinforcement learning and artificial intelligence.

Purdue University | GPA: 3.7 / 4.0 | August 2016 – May 2020

- Bachelor of Science in Computer Science Honors | Major GPA: 3.8 / 4.0
 - Concentrations in Software Engineering and Machine Learning
- Bachelor of Science in Statistics - Math Emphasis
- Bachelor of Science in Mathematics/Statistics
- Honors College Curriculum

Honors and Awards

- National Merit Scholarship
- Northrop Grumman, Raytheon, and Cisco Corporate Scholarships
- Intel Scholar in the Semiconductor Research Corporation's Undergraduate Research Program

Research

Reward Scale Robustness for Proximal Policy Optimization | February 2023 – May 2023

- Added robustness tricks from DreamerV3 to PPO to improve its robustness to various reward scales.
- Improved performance on poorly normalized reward functions and **full paper will appear at NeurIPS 2023**

Automatic Curriculum Learning | January 2022 – Present

- Developing automatic curriculum methods to solve procedural content generation (PCG) environments.
- Created [Syllabus](#), a library with high quality implementations of curriculum learning methods and infrastructure to easily apply them to existing RL agent training code.

Reward Surfaces of Policy Networks | January 2021 – January 2022

- Wrote library and performed experiments plotting the reward surface of RL agent policy networks.
- Identified new source of optimization challenges in the policy gradient direction via visualizations.
- Published at **ICML 2022**. Also appeared at the RL for Games AAAI22 workshop.

TCAV Explanations for Reinforcement Learning Agents | May 2021 – October 2021

- Using TCAV to explain which concepts in the state space that RL agents use for decision making.

Multiagent Evaluation for Real World Conditions | January 2021 – May 2021

- Tested state of the art methods for game theoretic evaluation of multiagent performance in a population.
- Researched impact of sampling noise and agent collusion on DeepMind's AlphaRank and Meta-Nash.

Computer Vision Research at Michigan State University | May 2019 – July 2019

- Used a U-Net-style convolutional neural network to automatically segment tubes and catheters in pediatric chest X rays, achieving visually accurate results on an unsolved medical computer vision task.

Select Publications

- **Sullivan, R.**, Kumar, A., Huang, S., Dickerson, J. P., Suarez, J., (2023). “Reward Scale Robustness for Proximal Policy Optimization via DreamerV3 Tricks”. (To appear) In *Advances in Neural Information Processing Systems, 2023*.
- Son, S., Zheng, L., **Sullivan, R.**, Qiao, Y., Lin, M., (2023). “Gradient Informed Proximal Policy Optimization”. (To appear) In *Advances in Neural Information Processing Systems, 2023*.
- Suarez, J., Bloomin, D., Choe, K., Li, H., **Sullivan, R.**, Ravichandran, N., Scott, D., Shuman, R., Bradley, H., Castricatto, L., Isola, P., You, K., Jiang, Y., Li, Q., Chen, J., Zhu, X., (2023). “Neural MMO 2.0: A Massively Multi-task Addition to Massively Multi-agent Learning”. (To appear) In *Advances in Neural Information Processing Systems, 2023*.
- **Sullivan, R.**, Terry, J. K., Black, B., (2022). “Cliff Diving: Exploring Reward Surfaces in Reinforcement Learning Environments”. In *The International Conference on Machine Learning, 2022*.
- Terry, J. K., Black, B., Grammel, N., Jayakumar, M., Hari, A., **Sullivan, R.**, Santos, L., Perez, R., Horsch, C., Dieffendahl, C., Williams, N., Lokesh, Y., Ravi, P. (2021). “PettingZoo: Gym for Multi-Agent Reinforcement Learning”. In *Advances in Neural Information Processing Systems, 2021*.
- Rodriguez-Rivera G., Turkstra J., Buckmaster J., LeClainche K., Montgomery S., Reed W., **Sullivan R.** and Lee J., “Tracking Large Class Projects in Real-Time Using Fine-Grained Source Control,” In *SIGCSE '22: Proceedings of the 52nd ACM Technical Symposium on Computer Science Education, 2022*.
- Holste, G., **Sullivan, R.**, Bindschadler, M., Nagy, N., & Alessio, A. (2020). “Multi-class semantic segmentation of pediatric chest radiographs”. In *SPIE, Medical Imaging 2020: Image Processing (Vol 11313, bll 323–330)*. doi:10.1117/12.2544426
- **Sullivan, R.**, Holste, G., Burkow, J., & Alessio, A. (2020). “Deep learning methods for segmentation of lines in pediatric chest radiographs”. In *SPIE, Medical Imaging 2020: Computer-Aided Diagnosis (Vol 11314, bll 577–583)*. doi:10.1117/12.2550686

Professional Experience

Student Researcher at Google Research | June 2023 – Present

- Developing techniques for using reinforcement learning from human feedback to personalize LLMs.
- Researching finetuning and prompting methods to optimize multi-objective RLHF problems at test-time.
- Writing methods, reviewing literature, and running large scale experiments using JAX on TPUs.

Part Time Research Associate at Oak Ridge National Lab | September 2022 – June 2023

- Worked part time to apply automatic curriculum learning and other RL methods to cybersecurity.
- Developed infrastructure to train and evaluate agents in complex network simulation.

Applied Scientist Intern at Amazon | May 2022 – September 2022

- Used offline and online reinforcement learning to optimize security rules Amazon sign-in systems.
- Worked with science and product teams to formulate problem as a sequential decision-making process.
- Created efficient simulator of sign-in process to generate new experiences for exploration.
- Developed models that will be deployed globally to reduce friction in the user sign in experience.

Open-Source Developer of PettingZoo | August 2020 – January 2022

- Maintainer and developer of PettingZoo, a unified API for multi-agent reinforcement learning environments.
- Updated and rewrote environments, improved documentation, and patched significant bugs.
- Since I joined the team, PettingZoo became the third most installed RL library, and is supported by many major RL frameworks including Stable Baselines 3, RLLib, and the Autonomous Learning Library.