

Ryan Sullivan

✉ rsulli@umd.edu • 🌐 <https://ryannavillus.github.io>
in <https://www.linkedin.com/in/ryan-navillus/> • 🐦 <https://x.com/RyanSullyvan>
📄 <https://github.com/RyanNavillus> • Last updated on March 30, 2025

I research **reinforcement learning** and **automated curriculum learning** to train intelligent agents in complex, **open-ended** environments and techniques to improve **large language models** with human feedback.

Education

Ph.D. in Computer Science , University of Maryland at College Park (3.8/4.0)	August 2020 – Present
B.S. in Computer Science , Purdue University (3.7/4.0) Concentrations in Software Engineering and Machine Learning	August 2016 – May 2020
B.S. in Applied Statistics , Purdue University (3.7/4.0)	August 2016 – May 2020
B.S. in Mathematics/Statistics , Purdue University (3.7/4.0)	August 2016 – May 2020

Industry Experience

Sony AI , Remote RL Research Intern for Game AI. Trained agents to play complex modern video games using state-of-the-art RL methods.	June 2024 – Present
Google AI , Mountain View, CA Student Researcher Developed new RLHF methods to personalize LLMs.	June 2023 – February 2024
Oak Ridge National Laboratory , Remote Part Time Research Associate Applied automatic curriculum learning and other RL methods to cybersecurity.	September 2022 – June 2023
Amazon , San Diego, CA Applied Science Intern Used offline and online RL to optimize security rules in Amazon sign-in systems.	May 2022 – September 2022
Swarm Labs , College Park, MD Graduate Student Researcher Developed a multi-agent evaluation service for RL agents using game-theoretic tournaments.	September 2021 – May 2022
MSU MIDI Lab , East Lansing, MI Undergraduate Research Intern	May 2019 – July 2019
Spensa , West Lafayette, IN iOS Development Intern	May 2017 – August 2017

Publications

 [[Google Scholar](#): 495 citations and an h-index of 7, [Semantic Scholar](#)]

Representative publications that I am a primary author on are **highlighted**.

2025

1. *Robust Multi-Objective Preference Alignment with Online DPO*
Ryan Sullivan, [Raghav Gupta](#), [Yunxuan Li](#), [Samrat Phatale](#), and [Abhinav Rastogi](#)
AAAI 2025

2024

2. *Gradient informed proximal policy optimization*
[Sanghyun Son](#), [Laura Zheng](#), **Ryan Sullivan**, [Yi-Ling Qiao](#), and [Ming Lin](#)
NeurIPS 2024
3. *Reward scale robustness for proximal policy optimization via DreamerV3 tricks*
Ryan Sullivan, [Akarsh Kumar](#), [Shengyi Huang](#), [John Dickerson](#), and [Joseph Suárez](#)
NeurIPS 2024
4. *Neural MMO 2.0: a massively multi-task addition to massively multi-agent learning*
[Joseph Suárez](#), [David Bloomin](#), [Kyoung Whan Choe](#), [Hao Xiang Li](#), **Ryan Sullivan**, [Nishaanth Kanna](#), [Daniel Scott](#), [Rose Shuman](#), [Herbie Bradley](#), [Louis Castricato](#), and others
NeurIPS 2024
5. *Open RL Benchmark: Comprehensive Tracked Experiments for Reinforcement Learning*
[Shengyi Huang](#), [Quentin Gallouédec](#), [Florian Felten](#), [Antonin Raffin](#), [Rousslan Fernand Julien Dossa](#), [Yanxiao Zhao](#), **Ryan Sullivan**, [Viktor Makoviychuk](#), [Denys Makoviichuk](#), [Mohamad H Danesh](#), and others
arXiv Preprint 2024

6. [Massively Multiagent Minigames for Training Generalist Agents](#)
Kyoung Whan Choe, **Ryan Sullivan**, and Joseph Suárez
arXiv Preprint 2024
7. [Conditional Language Policy: A General Framework for Steerable Multi-Objective Finetuning](#)
Kaiwen Wang, Rahul Kidambi, **Ryan Sullivan**, Alekh Agarwal, Christoph Dann, Andrea Michi, Marco Gelmi, Yunxuan Li, Raghav Gupta, Avinava Dubey, and others
EMNLP 2024
8. [Syllabus: Portable Curricula for Reinforcement Learning Agents](#)
Ryan Sullivan, Ryan Pégoud, Ameen Ur Rahman, Xincheng Yang, Junyun Huang, Aayush Verma, Nistha Mitra, and John P Dickerson
arXiv Preprint 2024

2022

9. [Cliff Diving: Exploring Reward Surfaces in Reinforcement Learning Environments](#)
Ryan Sullivan, Jordan K Terry, Benjamin Black, and John P Dickerson
ICML 2022
10. [Tracking large class projects in real-time using fine-grained source control](#)
Gustavo Rodriguez-Rivera, Jeff Turkstra, Jordan Buckmaster, Killian LeClainche, Shawn Montgomery, William Reed, **Ryan Sullivan**, and Jarett Lee
SIGCSE 2022

2021

11. [Pettingzoo: Gym for multi-agent reinforcement learning](#)
Jordan Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, **Ryan Sullivan**, Luis S Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, and others
NeurIPS 2021

2020

12. [Deep learning methods for segmentation of lines in pediatric chest radiographs](#)
Ryan Sullivan, Gregory Holste, Jonathan Burkow, and Adam Alessio
SPIE Medical Imaging 2020: Computer-Aided Diagnosis 2020
13. [Multi-class semantic segmentation of pediatric chest radiographs](#)
Gregory Holste, **Ryan Sullivan**, Michael Bindschadler, Nicholas Nagy, and Adam Alessio
SPIE Medical Imaging 2020: Image Processing 2020

Repositories

RyanNavillus/Syllabus ★31 <i>Portable curriculum learning infrastructure for RL agents.</i>	2023
NeuralMMO/environment ★519 <i>Complex Multi-agent RL research environment based on MMO videogames.</i>	2023
openrlbenchmark/openrlbenchmark ★209 <i>Comprehensive Tracked Experiments for Reinforcement Learning</i>	2023
RyanNavillus/reward-surfaces ★15 <i>Visualization tools for analyzing reward surfaces in RL.</i>	2021
Farama-Foundation/PettingZoo ★2.7k <i>Unified API for multi-agent RL environments.</i>	2020

Invited Talks

2023

1. Syllabus - Curriculum Learning Made Easy – Learning in Foundation Environments - LIFE Monthly Seminar ([slides](#))

2022

2. Exploring Reward Surfaces in Deep Reinforcement Learning – Kansas State University Knowledge Discovery in Databases Lab ([slides](#))
3. Cliff Diving - Exploring Reward Surfaces in Reinforcement Learning Environments – International Conference on Machine Learning ([slides](#))

Peer Review

Neural Information Processing Systems (NeurIPS) 2024
Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track 2024

Teaching

Object Oriented Programming (CMSC 131), Graduate Teaching Assistant Fall 2020, Spring 2021
<https://www.cs.umd.edu/class/fall2020/cmsc131-010X-030X/>

Extracurriculars

[Co-organized the Training Agents with Foundation Models workshop at RLC 2024](#) 2024
18 accepted papers and 4 invited talks on how foundation models can be trained for decision-making or improve tabula rasa RL agents.

[Co-organized the University of Maryland Multiagent Reinforcement Learning Reading Group](#) 2023
65 talks from industry and academia on multiagent RL, MARL, and related topics. We had over 1000 unique participants in 2023.