
PERSONAL INFORMATION	Full Name: Ryan Anthony Bahlous-Boldi Email: rbahlousbold@umass.edu
RESEARCH INTERESTS	Artificial Intelligence, Cognitive Science, Deep Learning, Reinforcement Learning, Robotics, Program Synthesis, Evolutionary Computation, Human Computer Interaction, Artificial Life, Generative World Models, Open Endedness, and Exploration.
EDUCATION	University of Massachusetts Amherst Amherst, MA <i>Bachelor of Science, Computer Science</i> Exp. May 2025 GPA: 4.0 Member of the Commonwealth Honors College Minors: Philosophy, Psychology
COMPUTER SKILLS	<i>Languages & Frameworks</i> Python, Clojure, C++, Java, JavaScript, R, Numpy, PyTorch, Jax, Flax
RESEARCH EXPERIENCE	<i>Advanced Agent – Robotics Technology (AART) Lab</i> 5/2024– Carnegie Mellon University Pittsburgh, PA Advisor: Katia Sycara <ul style="list-style-type: none">→ Member of the Robotics Institute Summer Scholars program.→ Encouraging emergent communication between diverse agents in LLM-based multi-agent reinforcement learning systems. <i>Programs Under Selection and Heredity (PUSH) Lab</i> 9/2021– University of Massachusetts Amherst & Amherst College Amherst, MA Advisor: Lee Spector <ul style="list-style-type: none">→ Applying lexibase selection to evolutionary computation systems such as genetic programming and evolutionary reinforcement learning, with a focus on exploration and diversity. <i>Safe, Confident and Aligned Learning + Robotics (SCALAR) lab</i> 9/2022– Manning College of Information and Computer Sciences, Amherst, MA University of Massachusetts Amherst Advisor: Scott Niekum <ul style="list-style-type: none">→ Improving safety and alignment of reinforcement learning from human feedback (RLHF) systems through learning a diverse distribution of reward functions or policies from preferences with hidden context. <i>Interactive and Collaborative Autonomous Robotics (ICAROS) lab</i> 5/2023–3/2024 Viterbi School of Engineering, University of Southern California Los Angeles, CA Advisor: Stefanos Nikolaidis <ul style="list-style-type: none">→ Integrating Quality Diversity algorithms such as Covariance Matrix Adaptation MAP Annealing with deep reinforcement learning.→ Developed a novel technique to create a goal-conditioned policy from a diverse set of evolved Q-functions. <i>Biologically Inspired Neural & Dynamical Systems Lab (BINDs) lab</i> 9/2021–9/2022 Manning College of Information and Computer Sciences Amherst, MA University of Massachusetts Amherst Advisor: Cooper Sigrist

- Selected for the [BINDslings](#) program where I explored consequences of and ways to improve the modularity of neural networks.

TEACHING AND MENTORSHIP

Teacher, *X-Camp Academy* 9/2021-

- Helped students to participate in the USA Computing Olympiad (USACO).
- Developed, lead and taught X-Camp's first AI Classes, that cover a broad range of topics such as computer vision, natural language processing and AI decision making.

Teaching Research Team, *X-Camp Academy* 6/2022-

- Communicate expectations to and from development and operation teams.
- Led the migration to a new teaching platform facilitating effective teaching and scaling of the company.

President, *MassAI* - UMass Amherst's AI Student Organization 1/2023-
Team Leader, *Team UMass: ProjectX ML Research Comp. Winners* 9/2022-1/2023

PUBLICATIONS

under "Ryan Boldi" and "Ryan Bahlous-Boldi"

* denotes equal contribution

Journal Articles

Ryan Boldi*, Martin Briesch*, Dominik Sobania, Alexander Lalejini, Thomas Helmuth, Franz Rothlauf, Charles Ofria, and Lee Spector. 2023. Informed Down-Sampled Lexicase Selection: Identifying productive training cases for efficient problem solving. <https://arxiv.org/abs/2301.01488>. In *Evolutionary Computation*. MIT Press.

Book Chapters

Lee Spector, Li Ding and **Ryan Boldi**. 2024. Particularity. In Genetic Programming Theory and Practice XX. New York: Springer. https://doi.org/10.1007/978-981-99-8413-8_9. Preprint posted at <https://arxiv.org/abs/2306.06812>

Conference and Workshop Papers

Ryan Boldi, Li Ding, Lee Spector, and Scott Niekum. 2024. Pareto-Optimal Learning from Preferences with Hidden Context. In the RL Safety Workshop at the First Annual Reinforcement Learning Conference (RLSW @ RLC 2024). Posted at <https://arxiv.org/abs/2406.15599>.

Ryan Boldi, Ashley Bao, Martin Briesch, Thomas Helmuth, Dominik Sobania, Lee Spector, Alexander Lalejini. 2024. Untangling the Effects of Down-Sampling and Selection in Genetic Programming. In Proceedings of ALIFE 2024: the 2024 Conference on Artificial Life. (ALIFE '24). https://doi.org/10.1162/isal_a_00832. MIT Press. Posted at <https://arxiv.org/abs/2304.07089>.

Ryan Boldi, Li Ding and Lee Spector. 2023. Objectives Are All You Need: Solving Deceptive Problems Without Explicit Diversity Maintenance. In the Workshop on Agent Learning in Open-Endedness at NeurIPS (ALOE @ NeurIPS 2023). Posted at <https://arxiv.org/pdf/2311.02283>

Ryan Boldi and Lee Spector. 2023. Can the Problem-Solving Benefits of Quality Diversity Be Obtained Without Explicit Diversity Maintenance? In Genetic and Evolutionary Computation Conference Companion (GECCO '23).

Ryan Boldi, Thomas Helmuth, and Lee Spector. 2022. The environmental discontinuity hypothesis for down-sampled lexicase selection. In The 2022 Conference on Artificial Life - Why it Didn't Work-Shop (ALIFE '22). Posted at <https://arxiv.org/abs/2208.11111>

[org/pdf/2205.15931](https://arxiv.org/pdf/2205.15931)

Li Ding, **Ryan Boldi**, Thomas Helmuth, and Lee Spector. 2022. Lexicase selection at scale. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '22).

Posters and Poster Papers

Ryan Boldi, Li Ding and Lee Spector. 2023. Solving Deceptive Problems without Explicit Diversity Maintenance. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '24).

Ryan Boldi, Matthew Fontaine, Sumeet Batra, Gaurav Sukhatme and Stefanos Nikolaidis. 2024. Generating Diverse Induced Policies for Conditioned Policy Distillation. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '24)

Ryan Boldi, Charles Zhang, Lee Spector. 2023. Encouraging Diversity in Reinforcement Learning with Lexicase Selection. RL at Harvard Workshop 2023.

Ryan Boldi, Ashley Bao, Martin Briesch, Thomas Helmuth, Dominik Sobania, Lee Spector, Alexander Lalejini. 2023. The Problem Solving Benefits of Down-Sampling Vary by Selection Scheme. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '23).

Ryan Boldi, Alexander Lalejini, Thomas Helmuth, Lee Spector. 2023. A static analysis of informed down-samples. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '23).

Li Ding, **Ryan Boldi**, Thomas Helmuth, and Lee Spector. 2022. Going faster and hence further with lexicase selection. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '22).

UNDER REVIEW **Ryan Bahlous-Boldi**, Li Ding, Lee Spector, and Scott Niekum. 2024. Pareto-Optimal Learning from Preferences with Hidden Context.

Ryan Bahlous-Boldi*, Maxence Faldor*, Luca Grillotti, Hannah Janmohamed, Lisa Coiffard, Lee Spector, and Antoine Cully. Dominated Novelty Search: Rethinking Local Competition in Quality-Diversity. Posted at <https://arxiv.org/abs/2502.00593>

PRE-PRINTS **Ryan Boldi***, Aadam Lokhandwala*, Edward Annatone, Yuval Schechter, Alexander Lavrenenko, Cooper Sigrist. 2023. Improving Recommendation System Serendipity Through Lexicase Selection. Posted at <https://arxiv.org/abs/2305.11044>

PRESENTATION (*In addition to those listed as conference/workshop papers and posters above*)

Conference Encouraging Diversity in Reinforcement Learning with Lexicase Selection
Poster: RL at Harvard Workshop 2023 Cambridge, MA

Think Before You Act: Generating High-Quality Diverse Reasoning Policies
Poster: SoCal Undergraduate Research Symposium 2023 Los Angeles, CA

	The Emergence of Diversity Emerging Researchers in Artificial Life Lightning Talk 2023 Conference on Artificial Life	Sapporo, Japan
Invited	Evolutionary Computation UMass Amherst Guest Lecture COMPSCI 389 - Introduction to Machine Learning	Spring 2023 Amherst, MA
	Lexicase Selection and Reinforcement Learning Personal Autonomous Robotics Lab (PeARL), UT Austin Autonomous Learning Laboratory, UMass Amherst	Fall 2022 Austin, Texas Amherst, MA
	Lexicase Selection and the Diversity of Quality Adaptive and Intelligent Robotics Lab, Imperial College London	Summer 2022 London, UK
	Evolutionary Algorithms United Arab Emirates Ministry of Artificial Intelligence	Fall 2020 Dubai, UAE
AWARDS	Selected as an <i>HRI Pioneer</i> HRI'25: IEEE/ACM International Conference on Human-Robot Interaction	
	<i>Goldwater Scholarship</i> Barry Goldwater Scholarship & Excellence in Education Foundation, 2024	\$7,500
	<i>Alumni Achievement Scholarship</i> University of Massachusetts Amherst, 2024	\$750
	<i>ProjectX ML Research Competition Winner</i> University of Toronto, 2023	\$20,000
	<i>Dean's Merit Scholarship</i> Manning College of Information and Computer Sciences, 2022	\$1,500
	<i>John E. and Alice M. Flynn Scholarship</i> University of Massachusetts Amherst, 2022	\$3,300
	<i>Imagine Cup Junior Winner</i> Microsoft, 2020	
MEMBERSHIP	<i>International Society for Artificial Life</i> <i>ACM SIGEVO, Special Interest Group for Genetic and Evolutionary Computation</i> <i>IEEE Robotics & Automation Society</i>	
SERVICE	<i>Reviewer, ICLR 2025</i> <i>Volunteer, RLC 2024</i> <i>Reviewer, GECCO 2024 Graph-based GP Workshop</i> <i>Volunteer, GECCO 2023</i>	Singapore Amherst, MA Melbourne, Australia Lisbon, Portugal