

Daniel S. Brown

Assistant Professor, University of Utah

daniel.s.brown@utah.edu [Website](#) [Google Scholar](#)

Academic Employment

University of Utah Assistant Professor School of Computing, Robotics Center	2022 – Present
University of California, Berkeley Postdoctoral Scholar Mentors: Anca Dragan and Ken Goldberg	2020 – 2022

Education

University of Texas at Austin Ph.D. Computer Science Dissertation: <i>Safe and Efficient Inverse Reinforcement Learning</i> Advisor: Scott Niekum	Graduation: 2020
Brigham Young University M.S. Computer Science Thesis: <i>Toward Scalable Human Interaction with Bio-Inspired Robot Teams</i> Advisor: Michael A. Goodrich	Graduation: 2013
B.S. Mathematics Honors Thesis: <i>Learning and Control Techniques in Portfolio Optimization</i> Advisor: Sean Warnick	Graduation: 2011

Awards and Honors

AAAI New Faculty Highlight Award. Invited to give a spotlight talk at the AAAI Conference on Artificial Intelligence, 2025.

Best Paper Finalist. International Symposium on Medical Robotics (ISMR), 2024.

Best Paper Finalist. IEEE/ACM International Conference on Human Robot Interaction (HRI), 2024.

Price College of Engineering Top 15% teaching recognition, 2023, 2024.

NIH Trailblazer Award, 2024.

University of Utah VP NSF 2023 Cohort. Accepted to small cohort of pre-tenure faculty for grant writing mentorship.

Robotics: Science and Systems (RSS) Pioneer. Selected as one of 30 top early career researchers in robotics, 2021.

3rd Place RoboCup@Home. Member of UT Austin Villa team that won 3rd prize in the Domestic Standard Platform League in 2017 in Nagoya, Japan.

UT Austin Provost Fellowship. Four-year fellowship for graduate students in the College of Natural Sciences, 2016-2020.

Best Paper Finalist. Symposium on Distributed Autonomous Robot Systems (DARS), 2016.

Best Paper Finalist. IEEE/ACM International Conference on Human Robot Interaction (HRI), 2014.

Science, Mathematics and Research for Transformation (SMART) Scholarship. DoD scholarship awarded to top US students to pursue graduate school and work with DoD research labs, 2011-2013.

Teaching

The University of Utah

- **CS 5960/6960 Human-AI Alignment:** Fall 2022, Fall 2023
 - Fall 2022: Course Overall Rating: 5.57/6.00, Instructor Overall Rating: 5.86/6.00
 - Fall 2023 (5960): Course Overall Rating: 4.89/6.00, Instructor Overall Rating: 5.22/6.00
 - Fall 2023 (6960): Course Overall Rating: 5.67/6.00, Instructor Overall Rating: 6.00/6.00
- **CS 4300/6300 Artificial Intelligence:** Spring 2023
 - Spring 2023 (6300): Course Overall Rating: 5.59/6.00, Instructor Overall Rating: 5.61/6.00
 - Fall 2025 (4300/6300)
- **CS 5955/6955 Advanced Artificial Intelligence,** Spring 2025
 - Spring 2025 (5955): Course Overall Rating: 5.63/6.00, Instructor Overall Rating: 5.63/6.00
 - Spring 2025 (6955): Course Overall Rating: 5.86/6.00, Instructor Overall Rating: 5.90/6.00

Outreach and Undergraduate Mentoring

Undergraduate Research Opportunity Program (UROP) 2023 – Present
University of Utah

- Mentored 12 undergraduates as part of Utah’s Undergraduate Research Opportunity Program.
- Students attended weekly lab meetings with my graduate students and pursued independent research projects.

Engineering Scholars Program 2023 – Present
University of Utah

- Mentored 6 first-year undergraduate students.
- Students attended weekly lab meetings with my graduate students, learned about reinforcement learning, and presented their research at the Engineering Scholars Poster Symposium.

- Undergraduate Spring Scavenger Hunt** 2023
University of Utah
- Invited undergraduates to tour our lab and talk with graduate students about research.
- Cabrillo Community College Berkeley Robotics Tour** 2022
University of California, Berkeley
- Organized visit for members of the SACNAS chapter at Cabrillo College to visit UC Berkeley.
 - Led students on tour of robotics facilities at UC Berkeley and organized meetings with current UC Berkeley undergraduate and graduate students.
- Mentor for Transfer-to-Excellence REU** 2021
University of California, Berkeley
- Mentored underrepresented, first-generation California community college student during summer internship at UC Berkeley.
 - Helped student develop research, technical writing, and presentation skills.
 - Student's summer research was awarded **Best Poster** at SACNAS 2021.
- Mentor for BAIR Undergraduate Mentoring Program** 2020
University of California, Berkeley
- Matched with UC Berkeley undergraduates from underrepresented groups.
 - Met monthly and provided insight and advice regarding career paths in AI and how to get started with research.
- Austin Hour of Code Instructor** 2017–2019
Austin, TX
- Taught students in underserved elementary schools basic coding skills.
 - Introduced students to exciting career opportunities in computer science.
- Summer Intern Mentor** 2015, 2016
AFRL Information Directorate, Rome, NY
- Mentored three college students and one high school student during their summer internships.
 - Helped students improve their coding, research, and communication skills.
- STEM Robotics Coach** 2013
Staley Elementary School, Rome, NY
- Mentored student teams in building and competing in LEGO Mindstorms robotics challenges.
 - Helped inspire 4th and 5th graders to pursue STEM careers.

Grants and Funding

- **National Artificial Intelligence Research Resource (NAIRR) Pilot Award.**
PI. *NSF*. \$48,760, 2024-2025.
- **SLES: High-Confidence Guarantees for Safe Reward and Policy Learning Under Uncertainty .**
PI. *NSF*. \$439,425, 2024-2027.
- **The 3-Armed Medic .**
Co-I. *University of Utah Remote and Austere Conditions Grand Challenge*. Total: \$100,000, Brown Share: \$34,000, 2024-2025.

- **Scaling Human-Guided Machine Learning to Multi-Human, Multi-Preference Settings** .
PI. *Army Research Lab*. \$100,000, 2024-2025.
- **Distributed, Interactive Human-Machine Teaming and Shared Autonomy** .
PI. *Army Research Lab*. Total: \$900,000, Brown Share: \$300,000, 2024-2027.
- **Transformative Autonomy at a Disruptively Less Invasive Scale in Surgical Robotics**.
Co-I. *ARPA-H*. Total: \$11,935,038, Brown Share: \$765,762 2024-2027.
- **A Gaze-Controlled Neck Exoskeleton for Dropped Head Syndrome**.
Co-PI *NIH Trailblazer*. \$295,000, 2023-2026.
- **Interactive Human-Agent Teaming via Multiple Feedback Types**.
PI. *Army Research Lab*. \$96,806, 2023-2024.
- **Collaborative Research: HCC: Medium: Aligning Robot Representations with Humans**.
Co-PI. *National Science Foundation: Human-Centered Computing*. \$305,523, 2023-2026.
- **Course Development Gift**.
PI *Open Philanthropy*. \$31,773, 2023-2024.
- **Value Alignment of AI Systems**.
PI *Open Philanthropy*. \$140,000, 2023-2024.
- **Robotics Certificates and Graduate Degree Programs**
Co-PI. *State of Utah System of Higher Education* \$545,460. 2022-2024.
- **Dynamic Multi-Agent Physical Search Problems with Probabilistic Knowledge**.
Co-PI. *AFOSR Laboratory Research Initiation Request*. \$630,000, 2015.
- **Adaptive Decision Making for Secure Bio-Inspired Computing**
Co-I *AFRL Commanders Research and Development Fund*. \$585,000, 2015.
- **Foundational Autonomy Demonstration and Evaluation**.
Co-PI. *AFRL Chief Scientist Fund*. \$50,000, 2015.
- **Swarm Intelligence for Multiagent Search and Reconnaissance**.
PI *Rome Laboratory Venture Research Funding*. \$50,000, 2014.

Invited Talks

Leveraging Human Input to Enable Robust, Interactive, and Aligned AI Systems. University of Texas at Austin, Oden Institute for Computational Engineering & Sciences. October, 2025.

Swarms, Emergent Behaviors, and Multi-Agent Systems. Utah Data Science & AI Lecture Series. September, 2025.

Toward Robust, Interactive, and Human-Aligned AI Systems. Utah AI Summit Keynote. June, 2025.

Human-in-the-Loop Robot Learning: Teaching, Correcting, and Adapting. Robotics: Science and Systems (RSS) Workshop Panelist. June, 2025.

Discovery and Deployment of Emergent Robot Swarm Behaviors via Representation Learning and Real2Sim2Real Transfer. Ploutos AI Forum. May, 2025.

Leveraging Human Input to Enable Robust, Interactive, and Aligned AI Systems. New Faculty Highlight Talk. AAAI Conference on Artificial Intelligence, February, 2025.

High-Confidence Guarantees for Safe Reward and Policy Learning Under Uncertainty . NSF Workshop on the Science of Safe AI, February, 2025.

Toward Interactive, Robust, and Aligned AI Systems. CX Tech Fest, Sandy, UT. January, 2025.

Toward Interactive, Robust, and Aligned AI Systems. Brigham Young University Computer Science Seminar. October, 2024.

Challenges and Progress Towards AI Alignment via Reinforcement Learning from Human Feedback. University of Utah Data Science Seminar. February, 2024.

Reinforcement Learning in and from the real-world. RL-CONFORM Workshop Panelist at the International Conference on Intelligent Robots and Systems (IROS). October, 2023.

Human-AI Alignment. Invited by Price College Engineering Deans office to present to Engineering National Advisory Council (ENAC). May, 2023.

Pitfalls and paths forward when learning rewards from human feedback. International Conference on Machine Learning (ICML) Workshop on Interactive learning with implicit human feedback. July, 2023.

Challenges and Progress Towards AI Alignment via Reinforcement Learning from Human Feedback. Marquette University Department Colloquium. March, 2023.

Challenges and Progress Towards Efficient and Causal Preference-Based Reward Learning. OpenAI Alignment Workshop, February, 2023.

Latent Spaces and Learned Representation for Better Human Preference Learning. CoRL 2022 Workshop on Aligning Robot Representations with Humans, Dec 2023.

Interactive Imitation Learning. University of Texas at Austin PeARL Lab. July, 2022.

Leveraging Human Input to Enable Robust AI Systems. Semiautonomous Seminar at UC Berkeley. July 2022.

Leveraging Human Input to Enable Robust AI Systems. Stanford Robotics Seminar. May, 2022.

Efficient and Robust Robot Learning of Human Objectives. University of Alberta Artificial Intelligence Seminar. October, 2021.

Leveraging Human Input for Robust Robot Learning. Talking Robotics Virtual Seminar Series. October, 2021.

Harnessing Machine Learning for Science: Helping Robots Learn from, Predict, and Better Assist Humans. Berkeley Science at Cal Midday Science Cafe. May, 2021.

Safe and Efficient Inverse Reinforcement Learning. MIT AeroAstro Humans Interacting with Autonomy Workshop. January, 2021.

Safe and Efficient Imitation Learning. BAIR/CPAR/BDD Seminar, University of California, Berkeley. November, 2020.

Safe and Efficient Inverse Reinforcement Learning. University of Southern California. November, 2020.

Toward Safe and Efficient Inverse Reinforcement Learning. University of California Berkeley. September, 2019.

Toward Safe and Efficient Imitation Learning. Massachusetts Institute of Technology. August 2019.

Dumb and Dumber: Collective Intelligence through Simple Behaviors. AFRL/AFIT Autonomy Technical Interchange Meeting. Dayton, OH. September, 2015.

Controlling Bio-Inspired Swarms through Limited Interactions. Command, Control, Communications, Cyber and Intelligence (C4I) Technology Review Days. Utica, NY. June 2014.

Press

University of Utah Department of Mechanical Engineering (Dec 22, 2023) [NIH Trailblazer Award](#) [Highlights my NIH Trailblazer award with Prof. Zhang on intelligent neck exoskeletons.]

John and Marcia Price College of Engineering. (Sept 30, 2024) [Kahlert Researchers Join ARPA-H Project to Develop Fully Automated Surgical Robot](#) [Highlights my work on uncertainty quantification for surgical robotics.]

Blog. Eric Jang. (Oct 23, 2021) [“Just Ask for Generalization.”](#) [Highlights my work on D-REX with regards to generalizing from imperfect demonstrations].

Quanta Magazine. Natalie Wolchover. (January 30, 2020). [“Artificial Intelligence Will Do What We Ask. Thats a Problem.”](#) [Highlights my work on value alignment, including risk-aware active reward learning and T-REX].

Internal Service

Utah Robot Football Club Advisor 2025 – Present
Assist student-led club with sponsorship, fundraising, mentoring, technical and research advice, and help with team organization and recruiting.

Kahlert School of Computing Faculty Hiring Committee. Spring 2024
Screened, interviewed, and hosted AI, NLP, and robotics faculty candidates.

Upskilling in AI Course Developer and Instructor. 2023 – Present
Developed course content on Deep Learning, Reinforcement Learning, and Human-in-the-Loop Machine Learning and Alignment. Taught classes and workshops to groups of up to 30 faculty from across campus to help faculty incorporate AI into their grants and research.

University of Utah AI/ML/Robotics Course Teaching Liaison. 2023 – Present
Coordinated with all AI, ML, and robotics faculty to help determine class teaching assignments and make sure core courses are covered.

Moderator for Idaho National Labs Visit and Center of Excellence Workshop. 2024.
Moderated the discussion on AI and research interests of INL and Utah. Presented findings to all INL and Utah visitors to argue for an AI Center of Excellence.

University of Utah Robotics Track Director. 2022 – Present
In charge of administration and degree requirements for all robotics track students.

Utah Robotics Degree Creation Program Committee Member. 2022 – Present
Weekly meetings to design new Robotics degree at the University of Utah.

Grad visit day. 2023,2024
Gave overview talk on robotics research at the University of Utah to prospective students and helped organize social events.

University of Utah PhD and MS Admission Committee Member. 2023,2024
Reviewed PhD and MS admissions applications for in the areas of AI, ML, and Robotics.

External and Other Service

2024 RL Conference Workshop Co-Chair.

In charge of writing call for submissions, recruitment, review and selection, and coordination of workshops for the 2024 RL Conference.

Session Chair

- AAAI 2025, Session on “Humans and AI”
- HRI 2025, Session on “Humans Teaching Robots”

External Committee Member

- Kiarash Geraili, Montana State University. PhD Dissertation Committee, 2025.
- Ricardo Vega, George Mason University. PhD Dissertation Committee, 2024.
- Kevin Zhu, George Mason University. PhD Dissertation Committee, 2023.
- Dimitris Papdimitriou, UC Berkeley. PhD Dissertation Committee, 2022.

Reviewer for Proposal Panels

- NSF Panel Reviewer, 2023, 2024, 2025.
- NSF TRAILBLAZER Review Panel - Behavioral Modeling & Robotics, 2024.
- Army Research Office (ARO) Proposal Review, 2024.
- DOE Advancements in Artificial Intelligence for Science Panel, 2024.

Workshop Organization

- Co-Organizer for RSS 2025 Workshop on Human-in-the-Loop Robot Learning: Teaching, Correcting, and Adapting.
- Co-Organizer for IROS 2023 Workshop on Human-Multi-Robot Interaction.
- Co-Organizer for NeurIPS 2021 Workshop on Safe and Robust Control of Uncertain Systems.

Journal Reviewing

- Journal of Artificial Intelligence Research (JAIR), 2025.
- Transactions on Machine Learning Research (TMLR), 2022-25.

- IEEE Robotics and Automation Letters (**RA-L**), 2017-25.
- Journal of Machine Learning Research (**JMLR**), 2021-23.
- IEEE Transactions on Robotics (**T-RO**), 2019-22.
- ACM Transactions on Human-Robot Interaction (**THRI**), 2018, 2022, 2024.
- AI Review, 2022.

Conference and Journal Reviewing Leadership

- Senior Program Committee, AAAI Conference on AI, 2025.
- Area Chair, Conference on Robot Learning, 2025.
- Senior Reviewer, Reinforcement Learning Conference, 2025.

Other Conference Reviewing

- Neural Information Processing Systems (**NeurIPS**), 2020, 2021, 2022, 2025.
- Learning for Dynamics & Control Conference (L4DC), 2025.
- International Symposium of Robotics Research (**ISRR**), 2024.
- Robotics Science and Systems (**RSS**), 2024
- International Symposium on Medical Robotics (**ISMR**), 2024
- IEEE International Conference on Intelligent Robots and Systems (**IROS**), 2022-24
- Conference on Robot Learning (**CoRL**), 2021-24.
- IEEE International Conference on Robotics and Automation (**ICRA**), 2021-23.
- AAAI Conference on Artificial Intelligence (**AAAI**), 2021-23.
- Robotics: Science and Systems (**RSS**), 2021.
- International Joint Conference on Artificial Intelligence (**IJCAI**), 2021, 2023.
- International Conference on Machine Learning (**ICML**), 2021.
- International Conference on Automation Science and Engineering (**CASE**), 2021.
- International Symposium on Multi-Robot and Multi-Agent Systems (**MRS**), 2021.
- International Conference on Human Robot Interaction (**HRI**), 2015-16, 2020, 2022, 2023.
- International Conference on Learning Representations (**ICLR**), 2019-20, 2024.
- American Control Conference (**ACC**), 2016.
- Autonomous Agents and Multi-Agent Systems (**AAMAS**), 2015.

Workshop Reviewing

- HRI Pioneers reviewer, 2023.
- CoRL 2022 Workshop on Aligning Robot Representations with Humans, 2022.
- NeurIPS Workshop on Safe and Robust Control of Uncertain Systems, 2021.
- RSS Pioneers Workshop, 2021, 2024.
- ICRA Workshop: Social Intelligence in Humans and Robots, 2021.

Undergraduate Students Mentored

University of Utah Undergraduate Research Opportunities Program (UROP).

- Zhengyao Li, Spring 2026
- Joy Zhuo, Spring 2026
- Joy Zhuo, Fall 2025
- Zhengyao Li, Fall 2025
- Jerry Wu, Fall 2025
- Duc La, Summer 2025
- Christian Rogers, Summer 2025
- Tom Fenno, Spring 2025
- Blake Lawlor, Spring 2025
- Tom Fenno, Fall 2024
- Ethan Gilmore, Fall 2024
- Joey Cai, Spring 2024
- Jeremy Clark, Spring 2024
- Emily Erickson, Spring 2024
- Seth Doubek, Spring 2024

University of Utah Engineering Scholars Program.

- Emilie Parra, Spring 2023
- Danny Garringer, Spring 2023
- Brock Perreca, Spring 2024
- Noah Zaffos, Spring 2024
- Christian Rogers, Spring 2025
- Wallace McCarthy, Spring 2025

Publications

[Google Scholar Page](#)

Preprints

1. **How Vulnerable Is My Learned Policy? Universal Adversarial Perturbation Attacks On Modern Behavior Cloning Policies**
Akansha Kalra, Basavasagar Patil, Guanhong Tao, Daniel S Brown
arXiv preprint arXiv:2502.03698

Refereed Journal and Conference Proceedings

1. **R2BC: Multi-Agent Imitation Learning from Single-Agent Demonstrations.**
Connor Mattson, Varun Raveendra, Ellen Novoseller, Nicholas Waytowich, Vernon J. Lawhern, Daniel S. Brown
International Conference on Robotics and Automation (ICRA), 2026.

2. **A Multi-Layer Sim-to-Real Framework for Gaze-Driven Assistive Neck Exoskeletons.**
Colin Rubow, Eric Brewer, Ian Bales, Haohan Zhang, Daniel S. Brown
International Conference on Robotics and Automation (ICRA), 2026.
3. **ProbeMDE: Uncertainty-Guided Active Proprioception for Monocular Depth Estimation in Surgical Robotics.** *Britton Jordan, Jordan Thompson, Jesse F. d’Almeida, Hao Li, Nithesh Kumar, Susheela Sharma Stern, James Ferguson, Ipek Oguz, Robert J. Webster III, Daniel S. Brown, Alan Kuntz*
International Conference on Robotics and Automation (ICRA), 2026.
4. **Dataset Poisoning Attacks on Behavioral Cloning Policies.**
Akansha Kalra, Soumil Datta, Ethan Gilmore, Duc La, Guanhong Tao, Daniel S. Brown
EAI-International Conference on Security and Privacy in Cyber-Physical Systems and Smart Vehicles, 2025.
5. **Agreement Volatility: A Second-Order Metric for Uncertainty Quantification in Surgical Robot Learning.**
Jordan Thompson, Britton Jordan, Daniel S. Brown, and Alan Kuntz
Conference on Robot Learning (CoRL), 2025.
6. **Toward Robust, Interactive, and HumanAligned AI Systems.**
Daniel S. Brown
AI Magazine, 2025.
7. **Leveraging Human Input to Enable Robust, Interactive, and Aligned AI Systems.**
Daniel S. Brown
AAAI Conference on Artificial Intelligence, 2025.
8. **Discovery and Deployment of Emergent Robot Swarm Behaviors via Representation Learning and Real2Sim2Real Transfer.**
Connor Mattson, Varun Raveendra, Ricardo Vega, Cameron Nowzari, Daniel S. Drew, and Daniel S. Brown
International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2025.
9. **Toward Zero-Shot User Intent Recognition in Shared Autonomy.**
Atharv Belsare, Zohre Karimi, Connor Mattson, Daniel S. Brown
ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2025.
10. **A dataset of paired head and eye movements during visual tasks in virtual environments**
Colin Rubow, Chia-Hsuan Tsai, Eric Brewer, Connor Mattson, Daniel S Brown, Haohan Zhang
Scientific Data, Nature Publishing Group, 2024.
11. **Uncertainty Aware Forward and Inverse Kinematics for Tendon-Driven Continuum Robots via Mixture Density Networks**
Jordan Thompson, Brian Y Cho, Daniel S Brown, Alan Kuntz
Journal of Medical Robotics Research (JMRR), 2024
12. **Can Differentiable Decision Trees Learn Interpretable Reward Functions?**
Akansha Kalra, Daniel S Brown
Reinforcement Learning Conference (RLC), 2024
13. **Representation Alignment from Human Feedback for Cross-Embodiment Reward Learning from Mixed-Quality Demonstrations**
Connor Mattson, Anurag Aribandi, Daniel S Brown
Reinforcement Learning Conference (RLC), 2024

14. **ROIL: Robust Offline Imitation Learning without Trajectories**
Gersi Doko, Guang Yang, Daniel S Brown, Marek Petrik
Reinforcement Learning Conference (RLC), 2024
15. **Spiking Neural Networks as a Controller for Emergent Swarm Agents**
Kevin Zhu, Connor Mattson, Shay Snyder, Ricardo Vega, Daniel S Brown, Maryam Parsa, Cameron Nowzari
International Conference on Neuromorphic Systems (ICONS), 2024.
16. **Modeling Kinematic Uncertainty of Tendon-Driven Continuum Robots via Mixture Density Networks.**
Jordan Thompson, Brian Y. Cho, Daniel S. Brown, Alan Kuntz
International Symposium on Medical Robotics (ISMR), 2024.
17. **Reward Learning from Suboptimal Demonstrations with Applications in Surgical Electrocautery.**
Zohre Karimi, Shing-Hei Ho, Bao Thach, Alan Kuntz, Daniel S. Brown
International Symposium on Medical Robotics (ISMR), 2024.
Best Paper Finalist Award
18. **Autonomous Assessment of Demonstration Sufficiency via Bayesian Inverse Reinforcement Learning.**
Tu Trinh, Haoyu Chen, Daniel S. Brown
ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2024.
Best Paper Finalist Award
19. **Bayesian Constraint Inference from User Demonstrations Based on Margin-Respecting Preference Models.**
Dimitris Papadimitriou, Daniel S. Brown
IEEE International Conference on Robotics and Automation (ICRA), 2024.
20. **Quantifying Assistive Robustness Via the Natural-Adversarial Frontier.**
Jerry Zhi-Yang He, Daniel S Brown, Zackory Erickson, Anca Dragan
Conference on Robot Learning (CoRL), 2023.
21. **Exploring Behavior Discovery Methods for Heterogeneous Swarms of Limited-Capability Robots**
Connor Mattson, Jeremy C Clark, Daniel S Brown
IEEE International Symposium on Multi-Robot & Multi-Agent Systems (MRS), 2023.
22. **Contextual Reliability: When Different Features Matter in Different Contexts.**
Gaurav Rohit Ghosal, Amrith Setlur, Daniel S Brown, Anca Dragan, Aditi Raghunathan
International Conference on Machine Learning (ICML), 2023.
23. **Leveraging Human Feedback to Evolve and Discover Novel Emergent Behaviors in Robot Swarms.**
Connor Mattson, Daniel S. Brown
Genetic and Evolutionary Computation Conference (GECCO), 2023.
24. **Efficient Preference-Based Reinforcement Learning Using Learned Dynamics Models.**
Yi Liu, Gaurav Datta, Ellen Novoseller, Daniel S. Brown
International Conference on Robotics and Automation (ICRA), 2023.
25. **Causal Confusion and Reward Misidentification in Preference-Based Reward Learning.**
Jeremy Tien, Jerry Zhi-Yang He, Zackory Erickson, Anca D. Dragan, Daniel S. Brown
International Conference on Learning Representations (ICLR), 2023.
26. **SIRL: Similarity-based Implicit Representation Learning**
Andreea Bobu, Yi Liu, Rohin Shah, Daniel S. Brown, Anca D. Dragan
International Conference on Human Robot Interaction (HRI), 2023.

27. **Benchmarks and Algorithms for Offline Preference-Based Reward Learning.**
Daniel Shin, Anca D. Dragan, Daniel S. Brown.
Transactions on Machine Learning Research, 2023.
28. **The Effect of Modeling Human Rationality Level on Learning Rewards from Multiple Feedback Types.**
Gaurav R. Ghosal, Matthew Zurek, Daniel S. Brown, Anca D. Dragan
AAAI Conference on Artificial Intelligence (AAAI), 2023.
29. **Bayesian Methods for Constraint Inference in Reinforcement Learning.**
Dimitris Papadimitriou, Usman Anwar, Daniel S. Brown.
Transactions on Machine Learning Research, 2022.
30. **Learning Representations that Enable Generalization in Assistive Tasks.**
Jerry Zhi-Yang He, Zackory Erickson, Daniel S. Brown, Aditi Raghunathan, Anca D. Dragan
Conference on Robot Learning (CoRL), 2022.
31. **Monte Carlo Augmented Actor-Critic for Sparse Reward Deep Reinforcement Learning from Suboptimal Demonstrations.**
Albert Wilcox, Ashwin Balakrishna, Jules Dedieu, Wyame Benslimane, Daniel S. Brown, Ken Goldberg
Neural Information Processing Systems (NeurIPS), 2022.
32. **Teaching Robots to Span the Space of Functional Expressive Motion.**
Arjun Sripathy, Andreea Bobu, Zhongyu Li, Koushil Sreenath, Daniel S. Brown, Anca D. Dragan
International Conference on Robot and Systems (IROS), 2022.
33. **Learning Switching Criteria for Sim2Real Transfer of Robotic Fabric Manipulation Policies.**
Satvik Sharma, Ellen Novoseller, Vainavi Viswanath, Zaymah Javed, Rishi Parikh, Ryan Hoque, Ashwin Balakrishna, Daniel S. Brown, Ken Goldberg
International Conference on Automation Science and Engineering (CASE), 2022.
34. **LEGS: Learning Efficient Grasp Sets for Exploratory Grasping.**
Letian Fu, Michael Danielczuk, Ashwin Balakrishna, Daniel S. Brown, Jeffrey Ichnowski, Eugen Solowjow, Ken Goldberg
International Conference on Robotics and Automation (ICRA), 2022.
35. **ThriftyDagger: Budget-Aware Novelty and Risk Gating for Interactive Imitation Learning.**
Ryan Hoque, Ashwin Balakrishna, Ellen Novoseller, Albert Wilcox, Daniel S. Brown, Ken Goldberg.
Conference on Robot Learning (CoRL), 2021.
36. **LazyDagger: Reducing Context Switching in Interactive Imitation Learning.**
Ryan Hoque, Ashwin Balakrishna, Carl Putterman, Michael Luo, Daniel S. Brown, Daniel Seita, Brijen Thananjeyan, Ellen Novoseller, Ken Goldberg.
IEEE Conference on Automation Science and Engineering (CASE), 2021.
37. **Kit-Net: Self-Supervised Learning to Kit Novel 3D Objects into Novel 3D Cavities.**
Shivin Devgon, Jeffrey Ichnowski, Michael Danielczuk, Daniel S. Brown, Ashwin Balakrishna, Shirin Joshi, Eduardo M. C. Rocha, Eugen Solowjow, Ken Goldberg.
IEEE Conference on Automation Science and Engineering (CASE), 2021.
38. **Value Alignment Verification.**
*Daniel S. Brown**, *Jordan Schneider**, *Anca D. Dragan*, *Scott Niekum.*
International Conference on Machine Learning (ICML), 2021.

39. **Policy Gradient Bayesian Robust Optimization for Imitation Learning.**
Zaynah Javed, Daniel S. Brown*, Ashwin Baladrishna, Satvik Sharma, Jerry Zhu, Marek Petrik, Anca D. Dragan, Ken Goldberg.*
International Conference on Machine Learning (ICML), 2021.
40. **Optimal Cost Design for Model Predictive Control.**
Avik Jain, Lawrence Chan, Daniel S. Brown, Anca D. Dragan.
Learning for Dynamics and Control Conference (L4DC), 2021.
41. **Situational Confidence Assistance for Lifelong Shared Autonomy.**
Matthew Zurek, Andreea Bobu, Daniel S. Brown, Anca D. Dragan.
International Conference on Robotics and Automation (ICRA), 2021.
42. **Dynamically Switching Human Prediction Models for Efficient Planning.**
Arjun Sripathy, Andreea Bobu, Daniel S. Brown, Anca D. Dragan.
International Conference on Robotics and Automation (ICRA), 2021.
43. **Exploratory Grasping: Self-Supervised Asymptotically Optimal Algorithms for Grasping and Re-Grasping Polyhedral Objects.**
Michael Danielczuk, Ashwin Balakrishna, Daniel S. Brown, Ken Goldberg.
Conference on Robot Learning (CoRL), 2020.
44. **Bayesian Robust Optimization for Imitation Learning.**
Daniel S. Brown, Scott Niekum, Marek Petrik.
Neural Information Processing Systems (NeurIPS), 2020.
45. **Safe Imitation Learning via Fast Bayesian Reward Inference from Preferences.**
Daniel S. Brown, Russell Coleman, Ravi Srinivasan, Scott Niekum.
International Conference on Machine Learning (ICML), 2020.
46. **Better-than-Demonstrator Imitation Learning via Automatically-Ranked Demonstrations.**
Daniel S. Brown, Wonjoon Goo, Scott Niekum.
Conference on Robot Learning (CoRL), 2019.
47. **Extrapolating Beyond Suboptimal Demonstrations via Inverse Reinforcement Learning from Observations.**
Daniel S. Brown, Wonjoon Goo*, Prabhat Nagarajan, Scott Niekum.*
International Conference on Machine Learning (ICML), 2019.
48. **Machine Teaching for Inverse Reinforcement Learning: Algorithms and Applications.**
Daniel S. Brown, Scott Niekum.
AAAI Conference on Artificial Intelligence (AAAI), 2019.
49. **Risk-Aware Active Inverse Reinforcement Learning.**
Daniel S. Brown, Yuchen Cui*, Scott Niekum.*
Conference on Robot Learning (CoRL), 2018.
50. **Efficient Probabilistic Performance Bounds for Inverse Reinforcement Learning.**
Daniel S. Brown, Scott Niekum.
AAAI Conference on Artificial Intelligence (AAAI), 2018.
51. **Discovery and Exploration of Novel Swarm Behaviors given Limited Robot Capabilities.**
Daniel S. Brown, Ryan Turner, Oliver Hennigh, Steven Loscalzo.
International Symposium on Distributed Autonomous Robotic Systems (DARS), 2016.
Best Paper Finalist Award
52. **Classifying Swarm Behaviors via Compressive Subspace Learning.**
Matthew Berger, Lee M. Seversky, Daniel S. Brown.
International Conference on Robotics and Automation (ICRA), 2016.

53. **Two Invariants of Human-Swarm Interaction.**
Daniel S. Brown, Michael A. Goodrich, Shin-Young Jung, and Sean Kerman.
Journal of Human-Robot Interaction, 2016.
54. **Exact and Heuristic Algorithms for Risk-Aware Stochastic Physical Search.**
Daniel S. Brown, Jeffrey Hudack, Nathaniel Gemelli, Bikramjit Banerjee.
Computational Intelligence, 2016.
55. **Evolving and Controlling Perimeter, Rendezvous, and Foraging Behaviors in a Computation-Free Robot Swarm.**
Matthew Johnson, Daniel S. Brown.
International Conference on Bio-inspired Information and Communications Technologies (BICT), 2015.
56. **k-Agent Sufficiency for Multiagent Stochastic Physical Search Problems.**
Daniel S. Brown, Steven Loscalzo, Nathaniel Gemelli.
International Conference on Algorithmic Decision Theory (ADT), 2015.
57. **Multiobjective Optimization for the Stochastic Physical Search Problem.**
Jeffrey Hudack, Nathaniel Gemelli, Daniel S. Brown, Steven Loscalzo, Jae C. Oh.
International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, 2015.
58. **Balancing Human and Inter-Agent Influences for Shared Control of Bio-Inspired Collectives.**
Daniel S. Brown, Shin-Young Jung, and Michael A. Goodrich.
Proceedings of IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2014.
59. **Limited Bandwidth Recognition of Collective Behaviors in Bio-Inspired Swarms.**
Daniel S. Brown and Michael A. Goodrich.
Autonomous Agents and Multiagent Systems (AAMAS), 2014.
60. **Human-Swarm Interactions Based on Managing Attractors.**
Daniel S. Brown, Sean Kerman, and Michael A. Goodrich.
International Conference on Human-Robot Interaction (HRI), 2014.
***Best Paper Finalist Award ***
61. **Shaping Couzin-like Torus Swarms through Coordinated Mediation.**
Shin-Young Jung, Daniel S. Brown, and Michael A. Goodrich. International Conference on Systems, Man, and Cybernetics (SMC), 2013.
62. **Supporting Human Interaction with Robust Robot Swarms.**
Sean Kerman, Daniel S. Brown, and Michael A. Goodrich. Proceedings of the International Symposium on Resilient Control Systems, 2012.

Workshop Proceedings and Late Breaking Reports

1. **ConeCut: Scalable Removal of Preference Redundancy.**
Purbid Bambrro, Daniel S. Brown, Ana Marasović
CoLM Workshop on Socially Responsible Language Modelling Research (SoLaR), 2025.
2. **Early Failure Detection in Autonomous Surgical Soft-Tissue Manipulation via Uncertainty Quantification.**
Jordan Thompson, Anthony Le, Daniel S. Brown, Alan Kuntz
RSS Workshop on Out-of-Distribution Generalization in Robotics, 2025.
3. **R2BC: Multi-Agent Imitation Learning from Single-Agent Demonstrations.**
Connor Mattson, Daniel S. Brown
RSS Workshop on Human-in-the-Loop Robot Learning: Teaching, Correcting, and Adapting, 2025.

4. **Emergent Swarm Behavior Prediction using 1D-Convolution for Autonomous Closed-Loop Behavior Control.**
Varun Raveendra, Connor Mattson, Daniel S. Brown
AAMAS Workshop on Autonomous Robots and Multirobot Systems, 2025.
5. **Learner and Teacher Perspectives on Learning Rewards from Multiple Types of Human Feedback.**
Ali Larian, Atharv Belsare, Zifan Wu, Daniel S. Brown
RSS Workshop on Human-in-the-Loop Robot Learning: Teaching, Correcting, and Adapting, 2025.
6. **Player-Centric Procedural Content Generation: Enhancing Runtime Customization by Integrating Real-Time Player Feedback.**
Nancy Blackburn, Michael Gardone, Daniel S. Brown
CHI Play Work-In-Progress, 2023.
7. **Towards a Gaze-Driven Assistive Neck Exoskeleton via Virtual Reality Data Collection.**
Jordan Thompson, Haohan Zhang, Daniel S. Brown
HRI Workshop on Virtual, Augmented, and Mixed-Reality for Human-Robot Interactions, 2023.
8. **Interpretable Reward Learning via Differentiable Decision Trees.**
Akasha Kalra, Daniel S. Brown
NeurIPS Workshop on ML Safety, 2022.
9. **Autonomous Assessment of Demonstration Sufficiency via Bayesian Inverse Reinforcement Learning.**
Tu Trinh and Daniel S. Brown.
AAAI Fall Symposium on Lessons Learned for Autonomous Assessment of Machine Abilities, 2022.
10. **A Study of Causal Confusion in Preference-Based Reward Learning.**
Jeremy Tien, Jerry Zhi-Yang He, Zackory Erickson, Anca D. Dragan, Daniel S. Brown
RSS Workshop on Overlooked Aspects of Imitation Learning: Systems, Data, Tasks, and Beyond, 2022.
11. **Unbiased Efficient Feature Counts for Inverse RL.**
Gerard Donahue, Brendan Crowe, Marek Petrik, Daniel S. Brown.
NeurIPS Workshop on Safe and Robust Control of Uncertain Systems, 2021.
12. **Bayesian Inverse Constrained Reinforcement Learning.**
Dimitris Papadimitriou, Usman Anwar, Daniel S. Brown, .
NeurIPS Workshop on Safe and Robust Control of Uncertain Systems, 2021.
13. **Offline Preference-Based Apprenticeship Learning.**
Daniel Shin, Daniel S. Brown.
ICML Workshop on Human-AI Collaboration in Sequential Decision-Making, 2021.
14. **Deep Bayesian Reward Learning from Preferences.**
Daniel S. Brown, Scott Niekum.
NeurIPS Workshop on Safety and Robustness in Decision Making, 2019.
15. **LAAIR: A Layered Architecture for Autonomous Interactive Robots.**
Yuchian Jian, Nick Walker, Minkyu Kim, Nicolas Brisseon, Daniel S. Brown, Justin W. Hart, Scott Niekum, Luis Sentis, Peter Stone.
AAAI Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy, 2018.
16. **Toward Probabilistic Safety Bounds for Robot Learning from Demonstration.**
Daniel S. Brown, Scott Niekum.
AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction, 2017.

17. **Algorithms for Stochastic Physical Search on General Graphs.**
Daniel S. Brown, Jeffrey Hudack, Bikramjit Banerjee.
AAAI Workshop on Planning, Search, and Optimization Workshop, 2015.

Dissertation and Theses

1. **Safe and Efficient Inverse Reinforcement Learning.**
Daniel S. Brown
Doctoral Dissertation
University of Texas at Austin, 2020.
2. **Toward Scalable Human-Swarm Interaction with Bio-Inspired Robot Teams.**
Daniel S. Brown
Masters Thesis
Brigham Young University, 2013.
3. **Learning and Control Techniques for Portfolio Optimization.**
Daniel S. Brown.
Honors Thesis
Brigham Young University, 2011.