

Daniel S. Brown

Contact Information

School of Computing
University of Utah
Salt Lake City, UT

dsbrown@berkeley.edu
<https://www.cs.utah.edu/~dsbrown>

Academic Employment

University of Utah Assistant Professor School of Computing	2022 – Present
University of California, Berkeley Postdoctoral Scholar Mentors: Anca Dragan & 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

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. 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

Outreach

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.

Student Mentoring

PhD Students

- Ryan Hoque: CASE 2021, CoRL 2021
- Ashwin Balakrishna: CoRL 2020, ICML 2021, CoRL 2021
- Dimitris Papadimitriou: NeurIPS Workshop 2021.
- Lawrence Chan: L4DC 2021
- Daniel Seita: CASE 2021
- Brijen Thananjeyan: CASE 2021
- Andreea Bobu: ICRA 2021
- Michael Danielczuk: CoRL 2020, CASE 2021
- Jordan Schneider: ICML 2021

Masters Students

- Shivin Devgon: CASE 2021
- Michael Luo: CASE 2021
- Matthew Johnson: BICT 2015

Undergraduate Students

- Gerard Donahue: NeurIPS Workshop 2021
- Daniel Shin: ICML Workshop 2021
- Emanuel Navarro-Ortiz: SACNAS 2021
- Zaynah Javed: ICML 2021
- Satvik Sharma: ICML 2021
- Jerry Zhu: ICML 2021
- Avik Jain: L4DC 2021
- Carl Putterman: CASE 2021
- Arjun Sripathy: ICRA 2021
- Matthew Zurek: ICRA 2021
- Daniel Shin: NeurIPS 2021 Workshop
- Albert Wilcox: CoRL 2021
- Russell Coleman: ICML 2020

Grants and Funding

- **Future Fund.** Principal Investigator. “Efficient Value Alignment Verification.” \$280,000 total, 2022-2024.
- **AFOSR Laboratory Research Initiation Request.** Co-Principal Investigator. “Dynamic Multi-Agent Physical Search Problems with Probabilistic Knowledge.” \$630,000 total, 2015.
- **AFRL Commanders Research and Development Fund.** Investigator. “Adaptive Decision Making for Secure Bio-Inspired Computing.” \$585,000 total, 2015.
- **AFRL Chief Scientist Fund.** Co-Principal Investigator. “Foundational Autonomy Demonstration and Evaluation.” \$50,000 total, 2015.

- **Rome Laboratory Venture Research Funding.** Principal Investigator. “Swarm Intelligence for Multiagent Search and Reconnaissance.” \$50,000 total, 2014.

Invited Talks

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

Blog. Eric Jang. (Oct 23, 2021) “Just Ask for Generalization.” <https://evjang.com/2021/10/23/generalization.html> [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. That’s a Problem.” <https://www.quantamagazine.org/artificial-intellig>

[ence-will-do-what-we-ask-thats-a-problem-20200130](#) [Highlights my work on value alignment, including risk-aware active reward learning and T-REX].

Academic Service

UC Berkeley PhD Admission Committee Member

Reviewed PhD admissions applications for UC Berkeley in the areas of AI and Robotics. Learned how to assess student applications for diversity and potential for research success.

Workshop Organization

Co-Organizer for NeurIPS Workshop on Safe and Robust Control of Uncertain Systems, 2021.

Journal Reviewing

Journal of Machine Learning Research (**JMLR**), IEEE Transactions on Robotics (**T-RO**), IEEE Robotics and Automation Letters (**RA-L**), ACM Transactions on Human-Robot Interaction (**THRI**), IEEE Intelligent Systems, IEEE Transactions on Human-Machine Systems.

Conference Reviewing

Neural Information Processing Systems (**NeurIPS**), Robotics: Science and Systems (**R:SS**), Conference on Robot Learning (**CoRL**), International Conference on Machine Learning (**ICML**), International Conference on Learning Representations (**ICLR**), International Conference on Robotics and Automation (**ICRA**), International Conference on Human Robot Interaction (**HRI**), AAAI Conference on Artificial Intelligence (**AAAI**), , International Joint Conference on Artificial Intelligence (**IJCAI**), International Conference on Automation Science and Engineering (**CASE**), Autonomous Agents and Multi-Agent Systems (**AAMAS**).

Publications

[Google Scholar Page](#)

Refereed Journal and Conference Proceedings

1. **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.
2. **Monte Carlo Augmented Actor-Critic for Sparse Reward Deep Reinforcement Learning from Suboptimal Demonstrations.**
Albert Wilcox, Ashwin Balakrishna, Daniel S. Brown, Jules Dedieu, Wyame Benslimane, Ken Goldberg
Neural Information Processing Systems (NeurIPS), 2022.
3. **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.
4. **Learning Switching Criteria for Sim2Real Transfer of Robotic Fabric Manipulation Policies.**
Satvik Sharma, Ellen Novoseller, Vainavi Viswanath, Zaynah Javed, Rishi Parikh, Ryan

- Hoque, Ashwin Balakrishna, Daniel S. Brown, Ken Goldberg*
International Conference on Automation Science and Engineering (CASE), 2022.
5. **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.
 6. **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.
 7. **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.
 8. **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.
 9. **Value Alignment Verification.**
Daniel S. Brown, Jordan Schneider*, Anca D. Dragan, Scott Niekum.*
International Conference on Machine Learning (ICML), 2021.
 10. **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.
 11. **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.
 12. **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.
 13. **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.
 14. **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.
 15. **Bayesian Robust Optimization for Imitation Learning.**
Daniel S. Brown, Scott Niekum, Marek Petrik.
Neural Information Processing Systems (NeurIPS), 2020.
 16. **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.
 17. **Better-than-Demonstrator Imitation Learning via Automatically-Ranked Demonstrations.**
Daniel S. Brown, Wonjoon Goo, Scott Niekum.
Conference on Robot Learning (CoRL), 2019.

18. **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.
19. **Machine Teaching for Inverse Reinforcement Learning: Algorithms and Applications.**
Daniel S. Brown, *Scott Niekum*.
AAAI Conference on Artificial Intelligence (AAAI), 2019.
20. **Risk-Aware Active Inverse Reinforcement Learning.**
*Daniel S. Brown**, *Yuchen Cui**, *Scott Niekum*.
Conference on Robot Learning (CoRL), 2018.
21. **Efficient Probabilistic Performance Bounds for Inverse Reinforcement Learning.**
Daniel S. Brown, *Scott Niekum*.
AAAI Conference on Artificial Intelligence (AAAI), 2018.
22. **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 Award Finalist
23. **Classifying Swarm Behaviors via Compressive Subspace Learning.**
Matthew Berger, *Lee M. Seversky*, *Daniel S. Brown*.
International Conference on Robotics and Automation (ICRA), 2016.
24. **Two Invariants of Human-Swarm Interaction.**
Daniel S. Brown, *Michael A. Goodrich*, *Shin-Young Jung*, and *Sean Kerman*.
Journal of Human-Robot Interaction, 2016.
25. **Exact and Heuristic Algorithms for Risk-Aware Stochastic Physical Search.**
Daniel S. Brown, *Jeffrey Hudack*, *Nathaniel Gemelli*, *Bikramjit Banerjee*.
Computational Intelligence, 2016.
26. **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.
27. **k-Agent Sufficiency for Multiagent Stochastic Physical Search Problems.**
Daniel S. Brown, *Steven Loscalzo*, *Nathaniel Gemelli*.
International Conference on Algorithmic Decision Theory (ADT), 2015.
28. **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.
29. **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.
30. **Limited Bandwidth Recognition of Collective Behaviors in Bio-Inspired Swarms.**
Daniel S. Brown and *Michael A. Goodrich*.
Autonomous Agents and Multiagent Systems (AAMAS), 2014.

31. **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 Award Finalist
32. **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.
33. **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

1. **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.
2. **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.
3. **Bayesian Inverse Constrained Reinforcement Learning.**
Dimitris Papadimitriou, Usman Anwar, Daniel S. Brown, .
NeurIPS Workshop on Safe and Robust Control of Uncertain Systems, 2021.
4. **Offline Preference-Based Apprenticeship Learning.**
Daniel Shin, Daniel S. Brown.
ICML Workshop on Human-AI Collaboration in Sequential Decision-Making, 2021.
5. **Deep Bayesian Reward Learning from Preferences.**
Daniel S. Brown, Scott Niekum.
NeurIPS Workshop on Safety and Robustness in Decision Making, 2019.
6. **LAAIR: A Layered Architecture for Autonomous Interactive Robots.**
Yuchian Jian, Nick Walker, Minkyu Kim, Nicolas Brisseonneau, 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.
7. **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.
8. **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.