Daniel S. Brown

Contact Information

EECS Department
University of California, Berkeley
Berkeley, CA 94720

dsbrown@berkeley.edu
www.cs.utexas.edu/~dsbrown

Academic Employment

University of California, Berkeley 2020 – Present
Postdoctoral Scholar
Mentors: Anca Dragan & Ken Goldberg

Education

University of Texas at Austin Graduation: 2020
Ph.D. Computer Science
Dissertation: *Safe and Efficient Inverse Reinforcement Learning*
Advisor: Scott Niekum

Brigham Young University
M.S. Computer Science Graduation: 2013
Thesis: *Toward Scalable Human Interaction with Bio-Inspired Robot Teams*
Advisor: Michael A. Goodrich

B.S. Mathematics Graduation: 2011
Honors Thesis: *Learning and Control Techniques in Portfolio Optimization*
Advisor: Sean Warnick

Awards and Honors

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

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

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

Best Paper Nomination. 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.

Publications

[Google Scholar Page]

Journal Articles

1. **Two Invariants of Human-Swarm Interaction.**
   Daniel S. Brown*, Michael A. Goodich, Shin-Young Jung, and Sean Kerman.*
   Journal of Human-Robot Interaction, 2016.
2. **Exact and Heuristic Algorithms for Risk-Aware Stochastic Physical Search.**
   Daniel S. Brown*, Jeffrey Hudack, Nathaniel Gemelli, Bikramjit Banerjee.*
   Computational Intelligence, 2016.

Refereed Conference Proceedings

1. **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.
2. **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.
3. **Value Alignment Verification.**
   Daniel S. Brown*, Jordan Schneider*, Anca D. Dragan, Scott Niekum.*
   International Conference on Machine Learning (ICML), 2021.
4. **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.
5. **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.
6. **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.
7. **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.
8. **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.
9. **Bayesian Robust Optimization for Imitation Learning.**
   Daniel S. Brown*, Scott Niekum, Marek Petrik.*
   Neural Information Processing Systems (NeurIPS), 2020.
10. **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.

11. **Better-than-Demonstrator Imitation Learning via Automatically-Ranked Demonstrations.**
***Daniel S. Brown**, Wonjoon Goo, Scott Niekum.*
Conference on Robot Learning (CoRL), 2019.

12. **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.

13. **Machine Teaching for Inverse Reinforcement Learning: Algorithms and Applications.**
***Daniel S. Brown**, Scott Niekum.* AAAI Conference on Artificial Intelligence (AAAI), 2019.

14. **Risk-Aware Active Inverse Reinforcement Learning.**
Daniel S. Brown, Yuchen Cui*, Scott Niekum.*
Conference on Robot Learning (CoRL), 2018.

15. **Efficient Probabilistic Performance Bounds for Inverse Reinforcement Learning.**
***Daniel S. Brown**, Scott Niekum.*
AAAI Conference on Artificial Intelligence (AAAI), 2018.

16. **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.
Nominated for Best Paper Award

17. **Classifying Swarm Behaviors via Compressive Subspace Learning.**
Matthew Berger, Lee M. Seversky, ***Daniel S. Brown**.*
International Conference on Robotics and Automation (ICRA), 2016.

18. **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.

19. **k-Agent Sufficiency for Multiagent Stochastic Physical Search Problems.**
***Daniel S. Brown**, Steven Loscalzo, Nathaniel Gemelli.*
International Conference on Algorithmic Decision Theory (ADT), 2015.

20. **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.

21. **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.

22. **Limited Bandwidth Recognition of Collective Behaviors in Bio-Inspired Swarms.**
***Daniel S. Brown** and Michael A. Goodrich.*
Autonomous Agents and Multiagent Systems (AAMAS), 2014.

23. **Human-Swarm Interactions Based on Managing Attractors.**
***Daniel S. Brown**, Sean Kerman, and Michael A. Goodrich.*

International Conference on Human-Robot Interaction (HRI), 2014.
Nominated for Best Paper Award

24. **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.
25. **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. **Offline Preference-Based Apprenticeship Learning.**
Daniel Shin, ***Daniel S. Brown****.*
ICML Workshop on Human-AI Collaboration in Sequential Decision-Making, 2021.
2. **Deep Bayesian Reward Learning from Preferences.**
Daniel S. Brown*, Scott Niekum.*
NeurIPS Workshop on Safety and Robustness in Decision Making, 2019.
3. **LAAIR: A Layered Architecture for Autonomous Interactive Robots.**
Yuchian Jian, Nick Walker, Minkyu Kim, Nicolas Brissonneau, ***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.
4. **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.
5. **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
Department of Computer Science, 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.

Teaching Experience

Middle School Python Bootcamp Instructor 2018
Austin, TX
- Taught beginning and intermediate Python classes to middle school students.

- Personally developed new curriculum, lectures, and projects to engage students and teach key coding skills.

Teaching Assistant
CS 343H Honors Artificial Intelligence, UT Austin Computer Science Dept. 2019
- Helped students master the fundamentals of artificial intelligence and apply them to an "AI for Social Good" final project.
- Taught class while instructor was absent.
- Held regular office hours to assist students with homework, programming project, and exam preparation. Wrote and graded midterms and final exam.

CS 412 Modeling and Optimization, BYU Computer Science Dept. 2011
- Helped students master linear programming and convex optimization techniques.
- Taught class when instructor was absent.
- Held regular office hours to assist students with homework, programming projects, and exam preparation.

Mathematics Tutor 2005, 2007 – 2008
Math Tutoring Lab, BYU, Provo, UT
- Tutored 20-30 college students daily in advanced mathematics and problem solving techniques.
- Increased test scores by planning and teaching exam reviews for classes of 30-40 students.

Teaching Interests

Intro to Programming	Human-Robot Interaction
Data Structures	Robot Learning
Algorithms	Imitation Learning
Introduction to Artificial Intelligence	Human-in-the-Loop Machine Learning
Introduction to Machine Learning	Collective Intelligence and Bio-Inspired Swarms
Reinforcement Learning	

Outreach

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.

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
- Lawrence Chan: L4DC 2021
- Ryan Hoque: CASE 2021
- Ashwin Balakrishna: ICML 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
- 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
- Russell Coleman: ICML 2020

Grants and Funding

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

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 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.

Research and Professional Experience

InterACT Laboratory University of California, Berkeley
Postdoctoral Scholar 2020 – Present
Active preference learning, confidence-based human-robot interactions, and planning, and robust imitation learning.

AUTOLab University of California, Berkeley
Postdoctoral Scholar 2020 – Present
Dexterous manipulation, self-supervised robotic learning, human-in-the loop imitation learning.

Personal Autonomous Robotics Lab University of Texas at Austin
Graduate Research Assistant 2016 – 2020
Safe learning from demonstration, reward learning from suboptimal demonstrations, machine teaching.

Air Force Research Lab, Information Directorate Rome, NY
Computer Scientist 2013 – 2016
Emergent behaviors, evolutionary search, robot swarms, physical search problems.

Human Centered Machined Intelligence Lab Brigham Young University
Graduate Research Assistant 2011 – 2013
Control and observability for bio-inspired robot swarms,.

EPIS Inc. Portland, OR
Energy Market Analyst Intern 2010
Optimization techniques to determine optimal portfolios in the energy market.

Information and Decision Algorithms Lab Brigham Young University
Undergraduate Research Assistant 2008-2011
Computational finance, portfolio optimization.

Academic Service

Workshop Organization

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

Journal Reviewing

IEEE Transactions on Robotics (T-RO), 2019-21.

IEEE Robotics and Automation Letters (RA-L), 2017-21.

Transactions on Emerging Topics in Computing, 2018.

Transactions on Human-Robot Interaction, 2018.

IEEE Intelligent Systems, 2016-17.

IEEE Transactions on Human-Machine Systems, 2015.

Conference Reviewing

Conference on Robot Learning (CoRL), 2021.

Neural Information Processing Systems (NeurIPS), 2020, 2021.

Robotics: Science and Systems (RSS), 2021.

International Joint Conference on Artificial Intelligence (IJCAI), 2021.

International Conference on Machine Learning (ICML), 2021.

IEEE International Conference on Robotics and Automation (ICRA), 2021.

International Conference on Automation Science and Engineering (CASE), 2021.

International Symposium on Multi-Robot and Multi-Agent Systems (MRS), 2021.

AAAI Conference on Artificial Intelligence (AAAI), 2020.

International Conference on Human Robot Interaction (HRI), 2015-16, 2020.

International Conference on Learning Representations (ICLR), 2019-20.

American Control Conference (ACC), 2016.

Autonomous Agents and Multi-Agent Systems (AAMAS), 2015.

Workshop Reviewing

R:SS Pioneers Workshop, 2021.

ICRA Workshop: Social Intelligence in Humans and Robots, 2021.