Scott Niekum

Associate Professor
College of Information and Computer Sciences
University of Massachusetts Amherst

Email: sniekum@cs.umass.edu
Phone: 413-545-2744
`https://people.cs.umass.edu/~sniekum`

Academic Employment

Associate Professor — Sept. 2022 –
University of Massachusetts Amherst,
College of Information and Computer Sciences

Associate Professor — Sept. 2021 – August 2022
The University of Texas at Austin, Department of Computer Science

Assistant Professor — Aug. 2015 – Aug. 2021
The University of Texas at Austin, Department of Computer Science

Postdoctoral Research Fellow — Aug. 2013 – Aug. 2015
Carnegie Mellon University, The Robotics Institute
Mentor: Prof. Christopher G. Atkeson

Industry Affiliations

Amazon Scholar — Feb. 2025 –
Amazon Robotics

Amazon Visiting Academic — Jul. 2024 – Jan. 2025
Amazon Robotics

Education

Doctor of Philosophy — Sept. 2009 – Sept. 2013
Computer Science Department, University of Massachusetts Amherst
Dissertation: *Semantically Grounded Learning from Unstructured Demonstrations*
Advisor: Prof. Andrew G. Barto

Bachelor of Science with Honors — Sept. 2001 – May 2005
School of Computer Science, Carnegie Mellon University
Additional major in Cognitive Science

Teaching

- **University of Massachusetts Amherst**

  – CS 690S: *Human-Centric Machine Learning (Graduate)*, Spring 2024
  – CS 383: *Artificial Intelligence (Undergraduate)*, Spring 2023, Fall 2023

- **The University of Texas at Austin**

  – CS 395T: *Robot Learning (Graduate)*, Fall 2015, Fall 2016, Fall 2017
  – CS 394R: *Reinforcement Learning (Graduate)*, Fall 2019, Spring 2022 (co-taught with Peter Stone)
  – CS 343(H): *Artificial Intelligence (Undergraduate)*, Spring 2017, Spring 2018, Fall 2018, Spring 2019, Spring 2020, Fall 2021
  – UT Austin Online MS: *Reinforcement Learning*, Spring 2020, Spring 2021

Awards and Honors

- RLDM 2022 Best Paper Award (For *Universal Off-Policy Evaluation*), 2022

- **AFOSR Young Investigator Award**, 2019

- College of Natural Sciences Teaching Excellence Award, 2019.

- **National Science Foundation CAREER Award**, 2018.

- Robocup@Home Domestic Standard Platform League, Third Place. Nagoya, Japan, July 2017. (Team Co-PI)

- Peter O'Donnell, Jr. Computer Sciences Endowed Faculty Fellowship, Sept 1, 2015 – Aug 31, 2020.

- NSF Graduate Research Fellowship Program, honorable mention, 2009, 2010.

- Carnegie Mellon Alumni Award for Excellence in Undergraduate Research, 2005.

Grants and Other Funding

- **Lead PI: Amazon**

  Foundation Model Reasoning with Domain-Specific APIs
  Award amount: $52,588
  Total PIs / Co-PIs: 1
  Dates: Sept 1, 2025 – Aug 31, 2026

- **Lead PI: Insulet Corporation**

  Diabetes Treatment via High-Confidence Reinforcement Learning
  Award amount: $908,706
  Total PIs / Co-PIs: 2
  Dates: Feb 1, 2025 – Jan 31, 2028

- **Lead PI: Effective Ventures Foundation**

  A New Paradigm for AI Alignment Based on Downstream Outcomes
  Award amount: $50,000
  Total PIs / Co-PIs: 1
  Dates: Aug 1, 2024 – July 31, 2025

- **Lead PI: Army Research Office**

  Modularity, Constraints and Multimodality in Learning for Complex, Long-Horizon, Sequential Decision-Making
  Award amount: $864,415
  Total PIs / Co-PIs: 3
  Dates: Jan 1, 2021 – Dec 31, 2024

- **Lead PI: AFOSR Young Investigator Award**

  Robot Learning from Demonstration with Auxiliary Contextual Data
  Award amount: $450,000
  Total PIs / Co-PIs: 1
  Dates: Jan 1, 2020 – Dec 31, 2024

- **Lead PI: Army Futures Command: Robotics Center of Excellence Grant**

  Human-AI Collaborative Search
  Award amount: $1,594,000
  Total PIs / Co-PIs: 1
  Dates: Sept 1, 2019 – Aug 31, 2023

- **Co-PI: NSF – National Robotics Initiative 2.0**

  IIS-1638107: *Improving Robot Learning from Feedback and Demonstration using Natural Language*
  Award amount: $749,411
  Total PIs / Co-PIs: 3
  Dates: Sept 1, 2019 – Aug 31, 2024

- **Lead PI: UT Austin Good Systems Grand Challenge**

  Probabilistically Safe and Correct Imitation Learning
  Award amount: $100,000
  Total PIs / Co-PIs: 1
  Dates: June 1, 2019 – May 31, 2020

- **Co-PI: Amazon Research Award**

  Language-Aided Learning from Demonstration
  Award amount: $55,000
  Total PIs / Co-PIs: 2
  Dates: Mar 1, 2019 – Feb 28, 2020

- **Lead PI: NSF CAREER Award**

  IIS-1749204: *Safe and Efficient Robot Learning from Demonstration in the Real World*
  Award amount: $524,605
  Total PIs / Co-PIs: 1
  Dates: June 1, 2018 – May 31, 2025

- **Lead PI: NSF – Smart and Autonomous Systems**

  IIS-1724157: *Socially-Aware Autonomy for Long-Term Deployment of Always-On Heterogeneous Robot Teams*
  Award amount: $1,378,736
  Total PIs / Co-PIs: 3
  Dates: June 1, 2017 – May 31, 2021

- **Co-PI: Office of Naval Research**

  Off-Policy Evaluation for Grounded Simulation Learning
  Award amount: $900,000
  Total PIs / Co-PIs: 2
  Dates: May 1, 2017 – April 30, 2021

- **Lead PI: NSF – National Robotics Initiative**

  IIS-1638107: *Scalable Robot Autonomy through Remote Operator Assistance and Lifelong Learning*
  Award amount: $486,276
  Total PIs / Co-PIs: 2
  Dates: Sept 1, 2016 – Aug 31, 2019

- **Lead PI: NSF – Robust Intelligence**

  IIS-1617639: *High Confidence, Efficient Learning Under Rich Task Specifications*
  Award amount: $470,000
  Total PIs / Co-PIs: 2
  Dates: Aug 1, 2016 – July 31, 2019

Invited Talks and Panels

- Panelist, Reinforcement Learning Safety Workshop. Reinforcement Learning Conference. Amherst, MA. August 2024.

- R:SS Social Intelligence in Humans and Robots Workshop. *Models of Human Preference for AI Alignment*. Virtual. July 2023.

- San Francisco Alignment Workshop. *Models of human preference for learning reward functions*. OpenAI, San Francisco. February 2023.

- IROS RL-CONFORM Workshop. *The Role of Guarantees in Value Alignment*. Virtual. October 2022.

- RLDM 2022 Invited Speaker. *The Role of Guarantees in Value Alignment*. Providence, RI. June 2022.

- ICML Workshop on Human-AI Collaboration in Sequential Decision Making. *Quantifying Risk and Value Alignment of Policies*. Virtual. July 2021.

- University of Washington. *Scaling Probabilistically Safe Learning to Robotics*. Seattle, WA (virtual). January 2021.

- Simons Institute Deep Reinforcement Learning Workshop. *Is Safe Learning the Future of Reinforcement Learning?*. Berkeley, CA (virtual). October 2020.

- Carnegie Mellon University. *Scaling Probabilistically Safe Learning to Robotics*. Pittsburgh, PA (virtual). September 2020.

- ICRA Workshop on Interactive Robot Learning. *Scaling Probabilistically Safe Learning to Robotics*. Paris, France (virtual). June 2020.

- Stanford University. *Scaling Probabilistically Safe Learning to Robotics*. Palo Alto, CA (virtual). April 2020.

- McGill University. *Scaling Probabilistically Safe Learning to Robotics*. Montreal, Quebec (virtual). April 2020.

- University of Alberta. *Scaling Probabilistically Safe Learning to Robotics*. Edmonton, Alberta (virtual). March 2020.

- University of Southern California. *Scaling Probabilistically Safe Learning to Robotics*. Los Angeles, CA. February 2020.

- Brown University HCRI Seminar. *Scaling Probabilistically Safe Learning to Robotics*. Providence, RI. September 2019.

- Samsung AI Research. *Scaling Probabilistically Safe Learning to Robotics*. New York, NY. July 2019.

- University of California Berkeley. *Scaling Probabilistically Safe Learning to Robotics*. Berkeley, CA. April 2019.

- Toyota Technological Institute. *Scaling Probabilistically Safe Learning to Robotics*. Chicago, IL. November 2018.

- Time Machine AI Conference. *Where's my Robot: Challenges and Progress in Personal Robotics*. Austin, TX. November 2018.

- Army Research Lab. *Efficient and Safe Learning from Demonstration*. Adelphi, MD. November 2017.

- Distinguished Speaker, Association for Computing Machinery (ACM) Annual Banquet, Lamar University. *"But What About Skynet?": Separating Fact from Fiction in Artificial Intelligence*. Beaumont, TX. April 2017.

- Panelist, *Beyond BB-8: When Robots Start Acting Human*. South by Southwest Interactive Track. Austin, TX. March 2017.

- Panelist, *Science Writing: Inside AI*. National Association of Science Writers Conference. San Antonio, TX. October 2016.

- Northwestern University. *Discovering Structure in Robotics Tasks via Demonstrations and Active Learning.* Evanston, IL. October 2016.

- INRIA Workshop on Algorithmic Human-Robot Interaction. *Non-Policy Learning from Demonstration and Interaction.* Paris, France. July 2016.

- TTI / Vanguard Conference on "Big Understanding". *From Robot Learning to Embodied Understanding.* Austin, TX. February 2016.

- AI for Human-Robot Interaction Symposium, AAAI Fall Symposium Series. *Discovering Structure in Robotics Tasks via Demonstrations and Active Learning.* Arlington, VA. November 2015.

- Brown University. *Online Bayesian Changepoint Detection for Articulated Motion Models.* Providence, RI. October 2014.

- University of Michigan. *Online Bayesian Changepoint Detection for Articulated Motion Models.* Ann Arbor, MI. September 2014.

- Georgia Institute of Technology. *Online Bayesian Changepoint Detection for Articulated Motion Models.* Atlanta, GA. August 2014.

- University of Southern California. *Grounded Learning from Unstructured Demonstrations.* Los Angeles, CA. May 2014.

- Worcester Polytechnic Institute. *Online Bayesian Changepoint Detection for Articulated Motion Models.* Worcester, MA. May 2014.

- Carnegie Mellon University. *Semantically Grounded Learning from Unstructured Demonstrations.* Pittsburgh, PA. April 2013.

- Massachusetts Institute of Technology. *Semantically Grounded Learning from Unstructured Demonstrations.* Cambridge, MA. March 2013.

Selected Press

- **BBC World Service** (Panel). The Real Story: Will artificial intelligence erode our rights?. Jan 26, 2024.

- New England Public Media (Interview). The Fabulous 413: Art and Science on the farms. Sept 11, 2023.

- WBUR Boston (Panel). What artificial intelligence could mean for art, content and race. August 31, 2023.

- **NPR Morning Edition** (Interview). Experts issue a dire warning about AI and encourage limits be imposed. June 1, 2023.

- Wired (Quote). Six Months Ago Elon Musk Called for a Pause on AI. Instead Development Sped Up. Sept 28, 2023.

- Wired (Quote). The Departure of 2 Google AI Researchers Spurs More Fallout. Mar 16, 2021.

Professional Activities

Event Organizing

- Member, Board of Directors, *Reinforcement Learning Conference*, August 2024—August 2029.

- **General chair and co-founder, *Reinforcement Learning Conference*, August 2024.**

- Organizer, *ICML 2024 Workshop on Models of Human Feedback for AI Alignment*, July 2024.

- Organizer, *NeurIPS 2022 Workshop on All Things Attention: Bridging Different Perspectives on Attention*, November 2022.

- Organizer, *RSS 2019 Workshop on Robust Autonomy: Safe Robot Learning and Control in Uncertain Real-World Environments*, July 2019.

- Co-chair, *AAAI 2017 video competition*, February 2017.

- Organizer, *RSS 2016 Workshop on Bootstrapping Manipulation Skills*, July 2016.

- Organizer, *AAAI 2015 Tutorial on Robot Learning from Demonstration*, January 2015.

- Organizer, *RSS 2013 Workshop on Human-Robot Collaboration*, June 2013.

Journal Action Editor

- Journal of Machine Learning Research (JMLR): 2023 –

Senior Area Chair

- International Conference on Machine Learning (ICML): 2025
- Neural Information Processing Systems (NeurIPS): 2023

Area Chair / Associate Editor

- Neural Information Processing Systems (NeurIPS): 2020
- International Conference on Learning Representations (ICLR): 2020
- IEEE International Conference on Robotics and Automation (ICRA): 2019
- AAAI Conference on Artificial Intelligence: 2018

Senior Program Committee

- AAAI Conference on Artificial Intelligence: 2017
- International Joint Conference on Artificial Intelligence (IJCAI): 2016

Conference Reviewing / Program Committee

- IEEE International Conference on Robotics and Automation (ICRA): 2023, 2022, 2021, 2018, 2017, 2016, 2015, 2014

- Neural Information Processing Systems (NeurIPS): 2019, 2018, 2015, 2011

- Robotics: Science and Systems (RSS): 2019, 2018, 2017, 2015, 2014

- The Multi-Disciplinary Conference on Reinforcement Learning and Decision Making (RLDM): 2019

- Conference on Robot Learning (CoRL): 2018

- International Conference on Intelligent Robots and Systems (IROS): 2024, 2018, 2017, 2015, 2014, 2013

- International Conference on Human Robot Interaction (HRI): 2017, 2012, 2013

- Autonomous Agents and Multi-Agent Systems (AAMAS): 2016, 2015, 2014

- International Conference on Humanoid Robots (Humanoids): 2014

- AAAI Special Robotics Track: 2013

- International Joint Conference on Artificial Intelligence (IJCAI): 2013

- AAAI Spring Symposium: 2013

- North East Student Colloquium on Artificial Intelligence: 2010

Journal Reviewing

- IEEE Transactions on Robotics (TRO): 2021

- International Journal of Robotics Research (IJRR): 2018, 2017, 2016

- IEEE Robotics and Automation Letters (RAL): 2015

- Artificial Intelligence (AIJ): 2014

- Frontiers in Computational Neuroscience: 2013

- IEEE Transactions on Autonomous Mental Development (TAMD): 2013

- IEEE Transactions on Systems, Man, and Cybernetics: 2013, 2012

- Journal of Machine Learning Research (JMLR): 2011

- Neurocomputing: 2010

Grant Reviewing

- Armstrong Fund for Science: 2023, 2024

- Army Research Office: 2020, 2025

- National Science Foundation: 2016, 2018

Departmental and University Service

University of Massachusetts Amherst

- Strategic Planning Committee, Co-Chair (2024–2025)

- Robotics Faculty Search Committee, Member (2024–2025)

- Strategic Planning Committee, Member (2023–2024)

- Promotion and Tenure Committee, Member (2023–2024)

- CAREER Cohort Advisor (2023–2024)

- Graduate Admission Committee, 50% Member (2023–2024)

- Graduate Program Committee, Member (2022–2023)

- Teaching Faculty Search Committee, Member (2022–2023)

University of Texas at Austin

- Graduate Advisor for Master of Data Science online program (2021–2022)

- Anna Hiss Robotics Building Space Committee, Member (2019–2022)

- Faculty Evaluation Committee, Member (2019–2020)

- Faculty Search Committee, Chair (2018–2019)

- Doctoral Admissions Committee, Member (2017–2018)

- Faculty Evaluation Committee, Member (2017)

- College of Natural Sciences 21st Century Curriculum Planning Implementation Task Force (2016–2017)

- Doctoral Admissions Committee, Co-chair (2016–2017)

- Doctoral Admissions Committee, Member (2015–2016)

- Robotics Seminar Series, Co-organizer (2015–2018)

Outreach

- PI, UT Austin Freshman Research Initiative Stream: Robot Learning (2019–present)

- Faculty Volunteer, Women in CS Faculty Lunch Program (2019, 2017)

- Exhibitor, Explore UT (2018)

- Speaker, Code Longhorn computer science camp for underrepresented high school students (2016–2018)

- Speaker, First Bytes computer science camp for high school women (2016–2018)

Advising and Thesis Committees

PhD Supervisor: Completed

- Harshit Sikchi, Department of Computer Science, The University of Texas at Austin (Co-advised with Amy Zhang, Fall 2020–Spring 2025)
  Reinforcement Learning Beyond Rewards: Decision-Making in the Language of Visitations

- Caleb Chuck, Department of Computer Science, The University of Texas at Austin (Co-advised with Amy Zhang, Fall 2017–Fall 2024)
  Control-based Factorization through Causal Interactions and Hierarchical Reinforcement Learning

- Mincheol Kim, Department of Mechanical Engineering, The University of Texas at Austin (Co-advised with Ashish Deshpande, Fall 2019–Fall 2022)
  Reliable, Robust, and Adaptable Dexterous Manipulation via Mechanical and Computational Optimization

- Wonjoon Goo, Department of Computer Science, The University of Texas at Austin (Fall 2016–Fall 2022)
  Imitation Learning with Auxiliary, Suboptimal, and Task-Agnostic Data

- Prasoon Goyal, Department of Computer Science, The University of Texas at Austin (Co-advised with Ray Mooney, Spring 2017–Summer 2022)
  Using Natural Language to Aid Task Specification in Sequential Decision Making Problems

- Yuchen Cui, Department of Computer Science, The University of Texas at Austin (Fall 2015–Fall 2021)
  Efficient Algorithms for Low-Effort Human Teaching of Robots

- Akanksha Saran, Department of Computer Science, The University of Texas at Austin (Co-advised with Andrea Thomaz, Fall 2015–Fall 2021)
  Leveraging Multimodal Human Cues to Enhance Robot Learning from Demonstration

- Ajinkya Jain, Department of Mechanical Engineering, The University of Texas at Austin (Spring 2016—Summer 2021)
  Learning and Leveraging Kinematics for Robot Motion Planning Under Uncertainty

- Daniel Brown, Department of Computer Science, The University of Texas at Austin (Fall 2016–Summer 2020)
  Safe and Efficient Inverse Reinforcement Learning

PhD Supervisor: Current

- Abhranil Chandra, College of Information and Computer Sciences, University of Massachusetts Amherst (Fall 2025–)

- Tuhina Tripathi, College of Information and Computer Sciences, University of Massachusetts Amherst (Fall 2023–)

- Shripad Deshmukh, College of Information and Computer Sciences, University of Massachusetts Amherst (Fall 2023–)

- Yaswanth Chittepu, College of Information and Computer Sciences, University of Massachusetts Amherst (Fall 2022–)

Master's Thesis Supervisor

- Joel Iventosch, Department of Computer Science, The University of Texas at Austin
  A Deep Learning Framework for Model-free 6 Degree of Freedom Object Tracking (Fall 2015—Spring 2017)

Doctoral Committee Member

- Arushi Jain, School of Computer Science, McGill University. Supervisor: Doina Precup
  Towards Reliable Reinforcement Learning Algorithms.

- Abhinav Bhatia, College of Information and Computer Sciences, University of Massachusetts Amherst. Supervisor: Shlomo Zilberstein
  Learning To Think About Thinking: Metareasoning with Deep Reinforcement Learning for Efficient and Safe Decision Making.

- Dhawal Gupta, College of Information and Computer Sciences, University of Massachusetts Amherst. Supervisor: Bruno Castro da Silva
  Improving Temporal Credit Assignment in Reinforcement Learning.

- Joar Skalse, Department of Computer Science, Oxford University. Supervisor: Alessandro Abate
  Partial Identifiability and Misspecification in Inverse Reinforcement Learning.

- Zhang-Wei Hong, Electrical Engineering and Computer Science, Massachusetts Institute of Technology. Supervisor: Pulkit Agrawal
  Generative Discovery via Reinforcement Learning.

- Letian Chen, College of Computing, Georgia Tech. Supervisor: Matthew Gombolay
  Personalized, Safe, and Interactive Robot Programming via Human Demonstrations.

- Chris Nota, College of Information and Computer Sciences, University of Massachusetts Amherst. Supervisor: Philip Thomas
  Policy Gradient Method: Analysis, Misconceptions, and Improvements.

- Michael Lee, Robotics Institute, Carnegie Mellon University. Supervisor: Reid Simmons and Henny Admoni
  Improving the Transparency of Agent Decision Making to Humans Using Demonstrations.

- Harsh Satija, School of Computer Science, McGill University. Supervisor: Joelle Pineau
  Towards alignment of Reinforcement Learning agents; for consideration of safety, robustness and fairness.

- Tushar Nagarajan, Department of Computer Science, The University of Texas at Austin. Supervisor: Kristen Grauman
  Learning affordance, environment and interaction representations by watching people in video

- Ziyang Tang, Department of Computer Science, The University of Texas at Austin. Supervisor: Qiang Liu
  Efficient And Safe Off-Policy Evaluation: From Point Estimation to Interval Estimation

- Yinan Zhao, Information School, The University of Texas at Austin. Supervisor: Danna Gurari
  Learning to manipulate images with image segmentation, search and synthesis

- Farzan Memerian, Institute for Computational Engineering and Sciences, The University of Texas at Austin. Supervisor: Ufuk Topcu
  Side Information, Robustness and Self-Supervision in Imitation Learning

- Mahsa Ghasemi, Institute for Computational Engineering and Sciences, The University of Texas at Austin. Supervisor: Ufuk Topcu
  Efficient, Reliable, and Interpretable Decision Making for Human Autonomy Co-Existence

- Saurabh Arora, Department of Computer Science, University of Georgia. Supervisor: Prashant Doshi
  Framework and Algorithms for Inverse Reinforcement Learning

- Aishwarya Padmakumar, Computer Science, The University of Texas at Austin. Supervisor: Raymond Mooney
  Dialog as a Vehicle for Lifelong Learning of Grounded Language Understanding Systems

- Yu-Chuan Su, Computer Science, The University of Texas at Austin. Supervisor: Kristen Grauman
  Learning for 360 Video Compression, Recognition, and Display

- Bo Xiong, Computer Science, The University of Texas at Austin. Supervisor: Kristen Grauman
  Learning to Compose Photos from Passive Cameras

- Josiah Hanna, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Data Efficient Reinforcement Learning with Off-Policy and Simulated Data

- Matthew Horn, Mechanical Engineering, The University of Texas at Austin. Supervisor: Sheldon Landsberger
  Trust and Safety in Robotic Applications

- Adam Allevato, Electrical and Computer Engineering, The University of Texas at Austin. Supervisor: Andrea Thomaz

- Pravesh Ranchod, Computer Science, University of the Witwatersrand. Supervisor: George Konidaris
  Skill Discovery from Multiple Related Demonstrators

- Lijia Liu, Computer Science, The University of Texas at Austin. Supervisor: Dana Ballard
  Cognitive Control of Motor Synergies

- Alex Broad, Department of Electrical Engineering and Computer Science, Northwestern University. Supervisor: Brenna Argall
  Model-Based Shared Control of Human-Machine Systems with Unknown Dynamics

- Jesse Thomason, Computer Science, The University of Texas at Austin. Supervisor: Raymond Mooney *Continuously Improving Natural Language Understanding for Robotic Systems through Semantic Parsing, Dialog, and Multi-modal Perception*

- Sanmit Narvekar, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Curriculum Learning in Reinforcement Learning

- Elad Liebman, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Sequential Decision Making in Artificial Musical Intelligence

- Dinesh Jayaraman, Computer Science, The University of Texas at Austin. Supervisor: Kristen Grauman
  Embodied Learning For Visual Recognition

- Jake Menashe, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Intrinsically-motivated Hierarchical Reinforcement Learning

- Suyog Jain, Computer Science, The University of Texas at Austin. Supervisor: Kristen Grauman
  Active Image and Video Segmentation

- Patrick MacAlpine, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Multilayered Skill Learning and Movement Coordination for Autonomous Robotic Agents in Spatial Domains

- Kai-Yang Chiang, Computer Science, The University of Texas at Austin. Supervisor: Inderjit Dhillon
  Analysis of Dyadic Interactions Using Machine Learning Methods

- Piyush Khandelwal, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  On-Demand Coordination of Multiple Service Robots

- Reza Mahjourian, Computer Science, The University of Texas at Austin. Supervisor: Risto Miikkulainen
  Neuroevolutionary Planning for Robotic Control

- Aditya Rawal, Computer Science, The University of Texas at Austin. Supervisor: Risto Miikkulainen
  Evolving Neural Networks for Sequence Processing

- Parham Pournazari, Mechanical Engineering, The University of Texas at Austin. Supervisor: Eric van Oort
  Real-time Learning of Dynamical Drilling Models for Event Detection and Robust Optimal Control

Master's Committee Member

- Brahma Pavse, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Reducing Sampling Error in Batch Temporal Difference Learning

- Prabhat Nagarajan, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Nondeterminism as a Reproducibility Challenge for Deep Reinforcement Learning

Undergraduate Thesis Committee Member

- Antony Yun, Computer Science, The University of Texas at Austin. Supervisor: Raymond Mooney
  Evaluating the Robustness of Natural Language Reward Shaping Models to Spatial Relations

- David Wang, Computer Science, The University of Texas at Austin. Supervisor: Philipp Krahenbuhl
  Using Deep Neural Networks to Predict Chess Player Skill

- Brahma Pavse, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Reinforced Inverse Dynamics Modeling for Learning from a Single Observed Demonstration

- Harsh Goyal, Computer Science, The University of Texas at Austin. Supervisor: Peter Stone
  Holistic Action Transforms

- Kamil Ali, Computer Science, The University of Texas at Austin. Supervisor: Philipp Krahenbuhl
  Can We Predict Actions Independent of the Actor?

Publications

Strongly Peer-Reviewed Conference and Journal Publications

1. H. Sikchi, A. Tirinzoni, A. Touati, Y. Xu, A. Kanervisto, **S. Niekum**, A. Zhang, A. Lazaric, M. Pirotta. *Fast Adaptation with Behavioral Foundation Models.* Reinforcement Learning Conference (RLC), August 2025.

2. R. Boldi, L. Ding, L. Spector, **S. Niekum**. *Pareto-Optimal Learning from Preferences with Hidden Context.* Reinforcement Learning Conference (RLC), August 2025.

3. H.T. Tse, P.S. Thomas, **S. Niekum**. *High-Confidence Policy Improvement from Human Feedback.* Reinforcement Learning Conference (RLC), August 2025.

4. Y. Chittepu, B. Metevier, W. Schwarzer, **S. Niekum**, P.S. Thomas. *Reinforcement Learning from Human Feedback with High-Confidence Safety Guarantees.* Reinforcement Learning Conference (RLC), August 2025.

5. C. Chuck, F. Feng, C. Qi, C. Shi, S. Agarwal, A. Zhang, **S. Niekum**. *Null Counterfactual Factor Interactions for Goal-Conditioned Reinforcement Learning.* International Conference on Learning Representations (ICLR), April 2025.

6. H. Xu, S. Li, H. Sikchi, **S. Niekum**, A. Zhang. *Iterative Dual-RL: An Optimal Discriminator Weighted Imitation Perspective for Reinforcement Learning.* International Conference on Learning Representations (ICLR), April 2025.

7. R. Rafailov, Y. Chittepu, R. Park, H. Sikchi, J. Hejna, W. Knox, C. Finn, **S. Niekum**. *Scaling Laws for Reward Model Overoptimization in Direct Alignment Algorithms.* Neural Information Processing Systems (NeurIPS), December 2024.

8. Z. Wang, J. Hu, C. Chuck, S. Chen, R. Martín-Martín, A. Zhang, **S. Niekum**, P. Stone. *SkiLD: Unsupervised Skill Discovery Guided by Local Dependencies.* Neural Information Processing Systems (NeurIPS), December 2024.

9. S. Chung, **S. Niekum**, D. Krueger. *Predicting Future Actions of Reinforcement Learning Agents.* Neural Information Processing Systems (NeurIPS), December 2024.

10. H. Sikchi, C. Chuck, A. Zhang, **S. Niekum**. *A Dual Approach to Imitation Learning from Observations with Offline Datasets.* Conference on Robot Learning (CoRL), November 2024.

11. J. Hanna, Y. Chandak, P.S. Thomas, M. White, P. Stone, **S. Niekum**. *Data-Efficient Policy Evaluation Through Behavior Policy Search.* Journal of Machine Learning Research (JMLR), October 2024.

12. P. Singhal, N. Lambert, **S. Niekum**, T. Goyal, G. Durrett. *D2PO: Discriminator-Guided DPO with Response Evaluation Models.* Conference on Language Modeling (COLM), October 2024.

13. M. Rudolph, C. Chuck, K. Black, M. Lvovsky, **S. Niekum**, and A. Zhang *Learning Action-based Representations Using Invariance.* Reinforcement Learning Conference (RLC), August 2024.

14. H. Sikchi, R. Chitnis, A. Touati, A. Geramifard, A. Zhang, **S.Niekum** *Score Models for Offline Goal-Conditioned Reinforcement Learning.* International Conference on Learning Representations (ICLR), May 2024.

15. H. Sikchi, A. Zhang, **S. Niekum**. *Dual RL: Unification and New Methods for Reinforcement and Imitation Learning.* International Conference on Learning Representations (ICLR), May 2024.

16. J. Hejna, R. Rafailov, H. Sikchi, C. Finn, **S. Niekum**, W. Bradley Knox, D. Sadigh *Contrastive Preference Learning: Learning from Human Feedback without RL.* International Conference on Learning Representations (ICLR), May 2024.

17. C. Chuck, K. Black, A. Arjun, Y. Zhu, **S. Niekum**. *Granger-Causal Hierarchical Skill Discovery.* Transactions on Machine Learning Research (TMLR), March 2024.

18. W. Knox, S. Hatgis-Kessell, S.O. Adalgeirsson, S. Booth, A. Dragan, P. Stone, **S. Niekum**. *Learning Optimal Advantage from Preferences and Mistaking it for Reward.* AAAI Conference on Artificial Intelligence, February 2024.

19. W. Knox, S. Hatgis-Kessell, S. Booth, **S. Niekum**, P. Stone, A. Allievi. *Models of Human Preference for Learning Reward Functions.* Transactions on Machine Learning Research (TMLR), January 2024.

20. S. Booth, W. Knox, J. Shah, **S. Niekum**, P. Stone, and A. Allievi. *The Perils of Trial-and-Error Reward Design: Misdesign through Overfitting and Invalid Task Specifications.* AAAI Conference on Artificial Intelligence, February 2023.

21. H. Sikchi, A. Saran, W. Goo, and **S. Niekum**. *A Ranking Game for Imitation Learning.* Transactions on Machine Learning Research (TMLR), January 2023.

22. A. Saran, K. Desai, M.L. Chang, R. Lioutikov, A. Thomaz, and **S. Niekum**. *Understanding Acoustic Patterns of Human Teachers Demonstrating Manipulation Tasks to Robots.* International Conference on Intelligent Robots and Systems (IROS), October 2022.

23. Y. Cui, **S. Niekum**, A. Gupta, V. Kumar, and A. Rajeswaran. *Can Foundation Models Perform Zero-Shot Task Specification For Robot Manipulation?* Learning for Dynamics and Control Conference (L4DC), June 2022.

24. S. Giguere, B. Metevier, B. Castro da Silva, Y. Brun, P.S. Thomas, **S. Niekum**. *Fairness Guarantees under Demographic Shift.* International Conference on Learning Representations (ICLR), April 2022.

25. C. Yuan, Y. Chandak, S. Giguere, P.S. Thomas, and **S. Niekum**. *SOPE: Spectrum of Off-Policy Estimators.* Neural Information Processing Systems (NeurIPS), December 2021.

26. I. Durugkar, M. Tec, **S. Niekum**, and P. Stone. *Adversarial Intrinsic Motivation for Reinforcement Learning.* Neural Information Processing Systems (NeurIPS), December 2021.

27. Y. Chandak, **S. Niekum**, B. Castro da Silva, E. Learned-Miller, E. Brunskill, and P.S. Thomas. *Universal Off-Policy Evaluation.* Neural Information Processing Systems (NeurIPS), December 2021. ***Also received the RLDM 2022 Best Paper Award***

28. A. Jain, S. Giguere, R. Lioutikov, and **S. Niekum**. *Distributional Depth-Based Estimation of Object Articulation Models.* Conference on Robot Learning (CoRL), November 2021.

29. M. Kim, **S. Niekum**, and A. Deshpande. *SCAPE: Learning Stiffness Control from Augmented Position Control Experiences.* Conference on Robot Learning (CoRL), November 2021.

30. W. Goo and **S. Niekum**. *You Only Evaluate Once: A Simple Baseline Algorithm for Offline RL.* Conference on Robot Learning (CoRL), November 2021.

31. F. Memarian, W. Goo, R. Lioutikov, **S. Niekum**, and U. Topcu. *Self-Supervised Online Reward Shaping in Sparse-Reward Environments* IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), September 2021.

32. Y. Cui, P. Koppol, H. Admoni, **S. Niekum**, R. Simmons, A. Steinfeld, and T. Fitzgerald. *Understanding the Relationship between Interactions and Outcomes in Human-in-the-Loop Machine Learning.* International Joint Conference on Artificial Intelligence (IJCAI), August 2021.

33. D.S. Brown, J. Schneider, A. Dragan, and **S. Niekum**. *Value Alignment Verification.* International Conference on Machine Learning (ICML), July 2021.

34. J.P. Hanna, **S. Niekum**, and P. Stone. *Importance sampling in reinforcement learning with an estimated behavior policy.* Machine Learning Journal (MLJ), June 2021.

35. A. Jain, R. Lioutikov, C. Chuck, and **S. Niekum**. *ScrewNet: Category-Independent Articulation Model Estimation From Depth Images Using Screw Theory.* IEEE International Conference on Robotics and Automation (ICRA), June 2021.

36. A. Saran, R. Zhang, E.S. Short, and **S. Niekum**. *Efficiently Guiding Imitation Learning Algorithms with Human Gaze.* International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2021.

37. O. Kroemer, **S. Niekum**, and G. Konidaris. *A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms.* Journal of Machine Learning Research, 22(30):1-82, January 2021.

38. D.S. Brown, **S. Niekum**, and M. Petrik. *Bayesian Robust Optimization for Imitation Learning.* Neural Information Processing Systems (NeurIPS), December 2020.

39. Y. Cui, Q. Zhang, A. Allievi, P. Stone, **S. Niekum**, and W. Knox. *The EMPATHIC Framework for Task Learning from Implicit Human Feedback.* Conference on Robot Learning (CoRL), November 2020.

40. P. Goyal, **S. Niekum**, and R. Mooney. *PixL2R: Guiding Reinforcement Learning Using Natural Language by Mapping Pixels to Rewards.* Conference on Robot Learning (CoRL), November 2020.

41. C. Chuck, S. Chockchowwat, and **S. Niekum**. *Hypothesis-Driven Skill Discovery for Hierarchical Deep Reinforcement Learning.* IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October 2020.

42. A. Jain and **S. Niekum**. *Learning Hybrid Object Kinematics for Efficient Hierarchical Planning Under Uncertainty.* IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October 2020.

43. D.S. Brown, R. Coleman, R. Srinivasan, and **S. Niekum**. *Safe Imitation Learning via Fast Bayesian Reward Inference from Preferences.* International Conference on Machine Learning (ICML), July 2020.

44. R. Zhang, A. Saran, B. Liu, Y. Zhu, S. Guo, **S. Niekum**, D. Ballard, M. Hayhoe. *Human Gaze Assisted Artificial Intelligence: A Review.* International Joint Conference on Artificial Intelligence (IJCAI), July 2020.

45. D.S. Brown, W. Goo, and **S. Niekum**. *Better-than-Demonstrator Imitation Learning via Automatically-Ranked Demonstrations.* Conference on Robot Learning (CoRL), October 2019.

46. A. Saran, E.S. Short, A.L. Thomaz, and **S. Niekum**. *Understanding Teacher Gaze Patterns for Robot Learning.* Conference on Robot Learning (CoRL), October 2019.

47. P. Goyal, **S. Niekum**, and R. Mooney. *Using Natural Language for Reward Shaping in Reinforcement Learning.* International Joint Conference on Artificial Intelligence (IJCAI), August 2019.

48. D.S. Brown, W. Goo, P. Nagarajan, and **S. Niekum**. *Extrapolating Beyond Suboptimal Demonstrations via Inverse Reinforcement Learning from Observations.* International Conference on Machine Learning (ICML), June 2019.

49. J.P. Hanna, **S. Niekum**, and P. Stone. *Importance Sampling Policy Evaluation with an Estimated Behavior Policy.* International Conference on Machine Learning (ICML), June 2019.

50. W. Goo and **S. Niekum**. *One-Shot Learning of Multi-Step Tasks from Observation via Activity Localization in Auxiliary Video.* IEEE International Conference on Robotics and Automation (ICRA), May 2019.

51. Y. Cui, D. Isele, **S. Niekum**, and K. Fujimura. *Uncertainty-Aware Data Aggregation for Deep Imitation Learning.* IEEE International Conference on Robotics and Automation (ICRA), May 2019.

52. D.S. Brown and **S. Niekum**. *Machine Teaching for Inverse Reinforcement Learning: Algorithms and Applications.* AAAI Conference on Artificial Intelligence, February 2019.

53. A. Jain and **S. Niekum**. *Efficient Hierarchical Robot Motion Planning Under Uncertainty and Hybrid Dynamics.* Conference on Robot Learning (CoRL), October 2018.

54. D.S. Brown, Y. Cui, and **S. Niekum**. *Risk-Aware Active Inverse Reinforcement Learning.* Conference on Robot Learning (CoRL), October 2018.

55. A. Saran, S. Majumdar, E.S. Short, A.L. Thomaz, and **S. Niekum**. *Human Gaze Following for Human-Robot Interaction.* IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October 2018.

56. Y. Cui and **S. Niekum**. *Active Reward Learning from Critiques.* IEEE International Conference on Robotics and Automation (ICRA), May 2018.

57. R.A. Gutierrez, V. Chu, A.L. Thomaz, and **S. Niekum**. *Incremental Task Modification via Corrective Demonstrations.* IEEE International Conference on Robotics and Automation (ICRA), May 2018.

58. D.S. Brown and **S. Niekum**. *Efficient Probabilistic Performance Bounds for Inverse Reinforcement Learning.* AAAI Conference on Artificial Intelligence, February 2018.

59. M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, **S. Niekum**, and U. Topcu. *Safe Reinforcement Learning via Shielding.* AAAI Conference on Artificial Intelligence, February 2018.

60. A. Saran, B. Lakic, S. Majumdar, J. Hess, and **S. Niekum**. *Viewpoint Selection for Visual Failure Detection.* IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), September 2017.

61. H.A. Poonawala, M. Alshiekh, **S. Niekum**, and U. Topcu. *Classification Error Correction: A Case Study in Brain-Computer Interfacing.* IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), September 2017.

62. J.P. Hanna, P.S. Thomas, P. Stone, and **S. Niekum**. *Data-Efficient Policy Evaluation Through Behavior Policy Search.* International Conference on Machine Learning (ICML), August 2017.

63. J.P. Hanna, P. Stone, and **S. Niekum**. *Bootstrapping with Models: Confidence Intervals for Off-Policy Evaluation.* International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2017.

64. P. Khandelwal, E. Liebman, **S. Niekum**, and P. Stone. *On the Analysis of Complex Backup Strategies in Monte Carlo Tree Search.* International Conference on Machine Learning (ICML), June 2016.

65. P.S. Thomas, **S. Niekum**, G. Theocharous, and G.D. Konidaris. *Policy Evaluation Using the Omega-Return.* Advances in Neural Information Processing Systems (NeurIPS), December 2015.

66. **S. Niekum**, S. Osentoski, C.G. Atkeson, and A.G. Barto. *Online Bayesian Changepoint Detection for Articulated Motion Models.* IEEE International Conference on Robotics and Automation (ICRA), May 2015.

67. K. Hausman, **S. Niekum**, S. Osentoski, and G. Sukhatme. *Active Articulation Model Estimation through Interactive Perception.* IEEE International Conference on Robotics and Automation (ICRA), May 2015.

68. **S. Niekum**, S. Osentoski, G.D. Konidaris, S. Chitta, B. Marthi, and A.G. Barto. *Learning Grounded Finite-State Representations from Unstructured Demonstrations.* International Journal of Robotics Research (IJRR), January 2015.

69. **S. Niekum**, S. Osentoski, S. Chitta, B. Marthi, and A.G. Barto. *Incremental Semantically Grounded Learning from Demonstration.* Robotics: Science and Systems (RSS), June 2013.

70. **S. Niekum**, S. Osentoski, G.D. Konidaris, and A.G. Barto. *Learning and Generalization of Complex Tasks from Unstructured Demonstrations.* IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October 2012.

71. **S. Niekum** and A.G. Barto. *Clustering via Dirichlet Process Mixture Models for Portable Skill Discovery.* Advances in Neural Information Processing Systems (NeurIPS), December 2011.

72. G.D. Konidaris, **S. Niekum**, and P.S. Thomas. *TD_γ: Reevaluating Complex Backups in Temporal Difference Learning.* Advances in Neural Information Processing Systems (NeurIPS), December 2011.

73. **S. Niekum**, A.G. Barto, L. Spector. *Genetic Programming for Reward Function Search.* IEEE Transactions on Autonomous Mental Development, vol.2, no.2, pp.83-90, June 2010.

74. D.R. Thompson, **S. Niekum**, T. Smith, and D. Wettergreen. *Automatic Detection and Classification of Geological Features of Interest.* IEEE Aerospace Conference, March 2005.

75. T. Smith, **S. Niekum**, D.R. Thompson, and D. Wettergreen. *Concepts for Science Autonomy During Robotic Traverse and Survey.* IEEE Aerospace Conference, March 2005.

Dissertations and Theses

76. **S. Niekum**. *Semantically Grounded Learning from Unstructured Demonstrations.* Doctoral Dissertation, Computer Science Department, University of Massachusetts Amherst, September 2013.

77. **S. Niekum**. *Reliable Rock Detection and Classification for Autonomous Science.* Undergraduate Senior Thesis, School of Computer Science, Carnegie Mellon University, April 2005.