

Stephane Hatgis-Kessell

Aspiring Scientist

Stanford, California — stephane@utexas.edu — American and Greek Citizen

EDUCATION

Ph.D., Computer Science

Starting in September 2024
Stanford University

B.S., Computer Science with Turing Scholars Honors

Year of completion: 2024
GPA: 3.79/4.0
University of Texas at Austin

Awards, Grants, and Recognition

National Science Foundation Graduate Research Fellowship

Awarded \approx \$150,000
2024 - 2029

Best Computer Science Honors Thesis Award

Selected as the top undergraduate computer science thesis among 32 undergraduates defending an honors thesis.
2024

UT Austin Dean's Honored Graduate in the College of Natural Sciences

Given to 1% of graduating students seniors who exemplify and embody the highest ideals of the college
2024

Honorable Mention of the 2024 CRA Outstanding Undergraduate Researcher Award

2024

University-Wide Award for Top Research Project, UT Co-Op George H. Mitchell Award

One of two undergraduate students awarded \$7,000 at a university with more than 40,000 students. 2023

Conference and Journal Publications

AAAI '23

W. Bradley Knox*, **Stephane Hatgis-Kessell***, Sigurdur Orn Adalgeirsson, Serena Booth, Anca Dragan, Peter Stone, and Scott Niekum, *Learning Optimal Advantage from Preferences and Mistaking it for Reward*, arxiv.org/abs/2310.02456, AAAI Conference on Artificial Intelligence.

TMLR '23

W. Bradley Knox*, **Stephane Hatgis-Kessell***, Serena Booth, Scott Niekum, Peter Stone, and Alessandro Allievi, *Models of Human Preference for Learning Reward Functions*, arxiv.org/abs/2206.02231, Transactions on Machine Learning Research.

RLDM '22

W. Bradley Knox*, **Stephane Hatgis-Kessell***, Serena Booth, Scott Niekum, Peter Stone, and Alessandro Allievi, *Extended Abstract: Partial Return Poorly Explains Human Preferences*, Multidisciplinary Conference on Reinforcement Learning and Decision Making. Selected for oral presentation.

* denotes equal contribution

Workshop Publications

ICML '23

W. Bradley Knox, **Stephane Hatgis-Kessell**, Sigurdur Orn Adalgeirsson, Serena Booth, Anca Dragan, Peter Stone, and Scott Niekum, *Learning Optimal Advantage from Preferences and Mistaking it for Reward*, 2023 ICML

Workshop on The Many Facets of Preference-based Learning (MFPL). Selected for oral presentation.

Teaching

Lead Mentor, UT Austin

Reinforcement Learning Directed Reading Group

1/22 - 5/22

Undergraduate Teaching Assistant, UT Austin

Honors Introduction to Computer Science Research

1/22 - 5/22

Lead Mentor, UT Austin

Computer Vision Directed Reading Group

3/21 - 1/22

Employment

Robotics Institute Summer Scholar, Carnegie Mellon University Robotics Institute

Worked in the Advanced Agent Laboratory on emergent communication in multi-agent reinforcement learning, collaborating with Prof. Katia Sycara.

6/22 - 9/22

Research Intern, Bosch R&D, Reinforcement Learning for Autonomous Driving Group

Worked on preference-based reinforcement learning, collaborating with Dr. Alessandro Allievi, Prof. Brad Knox, Prof. Peter Stone, and Prof. Scott Niekum.

3/21 - 6/22