

Haimin Hu

CONTACT INFORMATION	3400 N Charles St Baltimore, MD 21218, United States +1 (510) 561-8285	haiminhu.org Google Scholar haimin@cs.jhu.edu
RESEARCH INTERESTS	My vision is to enable human-centered robotic systems that can be deployed, verified, and trusted. I work on new algorithms and theorems centered around dynamic game theory, integrating insights from control systems safety, numerical optimization, deep reinforcement learning, and generative AI. The core of my research is to close the robot's interaction and learning loop by imbuing the decision model with a representation of its own evolving uncertainty as well as that of others.	
PROFESSIONAL POSITIONS	Johns Hopkins University , Baltimore, MD Assistant Professor, Department of Computer Science	2026
PAST POSITIONS	University of Pennsylvania , Philadelphia, PA Postdoctoral Researcher, Department of Electrical and Systems Engineering, PRECISE Center	2025 - 2026
	Toyota Research Institute , Cambridge, MA PhD Research Intern, Human-Interactive Driving Team	2023
	Honda Research Institute USA, Inc. , San Jose, CA PhD Research Intern	2022
	National Institute for Nuclear Physics & University of Padova , Padova, Italy Machine Learning Research Intern	2017
EDUCATION	Princeton University , Princeton, NJ <i>Ph.D. & M.A. in Electrical and Computer Engineering</i> Advisor: Jaime Fernández Fisac Main collaborators: Naomi Ehrich Leonard, Bartolomeo Stellato, Andrea Bajcsy Thesis: Game-Theoretic Integration of Safety and Learning for Human-Centered Robotics	2020 - 2025
	University of Pennsylvania , Philadelphia, PA <i>M.S.E. in Electrical Engineering</i> Advisors: George Pappas, Manfred Morari, Nikolai Matni	2018 - 2020
	University of California, Berkeley , Berkeley, CA <i>Visiting student in Electrical Engineering and Computer Sciences</i> Advisor: Claire Tomlin	2017 - 2018
	ShanghaiTech University , Shanghai, China <i>B.E. in Electronic and Information Engineering</i> Advisor: Boris Houska	2014 - 2018
HONORS AND AWARDS	Penn AI Fellow Robotics: Science and Systems Pioneer "30 world's top early-career researchers in robotics." Cyber-Physical Systems Rising Star "30 outstanding PhD students and postdocs in Cyber-Physical Systems (CPS)." Human-Robot Interaction Pioneer "25 world's top student researchers in human-robot interaction (HRI)."	University of Pennsylvania, 2026 RSS Foundation, 2025 NSF, 2025 ACM/IEEE, 2024

Roberto Tempo Best CDC Paper Nominee	IEEE Control Systems Society, 2024
Best Presentation Award	RSS Safe Autonomy Workshop, 2024
Best Poster Award	Workshop on Optimization, Learning, and Control, 2024
Outstanding Presentation Award	Princeton Research Day, 2024
<i>Awarded to 1 group of graduate students across all majors.</i>	
Bede Liu Fund for Excellence (Travel Grant)	Princeton University, 2022 - 2025
Outstanding Research Award	University of Pennsylvania, 2020
<i>Awarded to 2 electrical engineering masters.</i>	
Outstanding Graduate	Shanghai Municipal Education Commission, 2018
President's Scholarship	ShanghaiTech University, 2017
Merit Student for Excellence in Research	ShanghaiTech University, 2017
<i>Awarded to 1 undergraduate across all majors.</i>	
Scholarship for Overseas Study	ShanghaiTech University, 2016 - 2018
Scholarship for Academic Excellence	ShanghaiTech University, 2015

JOURNAL
PUBLICATIONS

- [J4] **H. Hu**, D. Isele, S. Bae, and J. F. Fisac. [Active Uncertainty Reduction for Safe and Efficient Interaction Planning: A Shielding-Aware Dual Control Approach](#), *The International Journal of Robotics Research (IJRR)*, 2023.
- [J3] K. Hsu, **H. Hu**, and J. F. Fisac. [The Safety Filter: A Unified View of Safety-Critical Control in Autonomous Systems](#), *Annual Review of Control, Robotics, and Autonomous Systems*, 2023.
- [J2] **H. Hu**, K. Nakamura, and J. F. Fisac. [SHARP: Shielding-Aware Robust Planning for Safe and Efficient Human-Robot Interaction](#), *IEEE Robotics and Automation Letters (RA-L)*, 2022.
- [J1] M. Chen*, S. L. Herbert*, **H. Hu**, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. [FaSTrack: A Modular Framework for Real-Time Motion Planning and Guaranteed Safe Tracking](#), *IEEE Transactions on Automatic Control (TAC)*, 2021.

CONFERENCE
PUBLICATIONS

- [C17] D. D. Oh*, D. P. Nguyen*, **H. Hu**, J. F. Fisac. [Provably Optimal Reinforcement Learning under Safety Filtering](#), *Conference of the International Association for Safe and Ethical Artificial Intelligence (IASEAI)*, 2026.
- [C16] D. D. Oh*, J. Lidard*, **H. Hu**, H. Sinhmar, E. Lazarski, D. Gopinath, E. Sumner, J. DeCastro, G. Rosman, N. E. Leonard, J. F. Fisac. [Safety with Agency: Human-Centered Safety Filter with Application to AI-Assisted Motorsports](#), *Robotics: Science and Systems (RSS)*, 2025.
- [C15] **H. Hu**, J. F. Fisac, N. E. Leonard, D. Gopinath, J. DeCastro, G. Rosman. [Think Deep and Fast: Learning Neural Nonlinear Opinion Dynamics from Inverse Dynamic Games for Split-Second Interactions](#), *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [C14] J. Wang*, **H. Hu***, D. P. Nguyen, and J. F. Fisac. [MAGICS: Adversarial RL with Minimax Actors Guided by Implicit Critic Stackelberg for Convergent Neural Synthesis of Robot Safety](#), *World Symposium on the Algorithmic Foundations of Robotics (WAFR) XVI*, 2024.
- [C13] **H. Hu***, G. Dragotto*, Z. Zhang, K. Liang, B. Stellato, and J. F. Fisac. [Who Plays First? Optimizing the Order of Play in Stackelberg Games with Many Robots](#), *Robotics: Science and Systems (RSS)*, 2024.
- [C12] J. Lidard*, **H. Hu***, A. Hancock, Z. Zhang, A. G. Contreras, V. Modi, J. DeCastro, D. Gopinath, G. Rosman, N. E. Leonard, M. Santos, and J. F. Fisac. [Blending Data-Driven Priors in Dynamic Games](#), *Robotics: Science and Systems (RSS)*, 2024.

* indicates equal contribution.

- [C11] **H. Hu***, Z. Zhang*, K. Nakamura, A. Bajcsy, and J. F. Fisac. [Deception Game: Closing the Safety-Learning Loop in Interactive Robot Autonomy](#), *Conference on Robot Learning (CoRL)*, 2023. **Best Presentation Award** at the 2024 RSS Safe Autonomy Workshop.
- [C10] **H. Hu**, K. Nakamura, K. Hsu, N. E. Leonard, and J. F. Fisac. [Emergent Coordination Through Game-Induced Nonlinear Opinion Dynamics](#), *IEEE Conference on Decision and Control (CDC)*, 2023. **Nominated for the Roberto Tempo Best CDC Paper Award**.
- [C9] **H. Hu** and J. F. Fisac. [Active Uncertainty Reduction for Human-Robot Interaction: An Implicit Dual Control Approach](#), *World Symposium on the Algorithmic Foundations of Robotics (WAFR) XV*, 2022. **Invited extension for IJRR special issue**.
- [C8] L. Lindemann, **H. Hu**, A. Robey, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni. [Learning Hybrid Control Barrier Functions from Data](#), *Conference on Robot Learning (CoRL)*, 2020.
- [C7] **H. Hu**, M. Fazlyab, M. Morari, and G. J. Pappas. [Reach-SDP: Reachability Analysis of Closed-Loop Systems with Neural Network Controllers via Semidefinite Programming](#), *IEEE Conference on Decision and Control (CDC)*, 2020.
- [C6] A. Robey*, **H. Hu***, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni. [Learning Control Barrier Functions from Expert Demonstrations](#), *IEEE Conference on Decision and Control (CDC)*, 2020.
- [C5] **H. Hu**, K. Gatsis, M. Morari, and G. J. Pappas. [Non-Cooperative Distributed MPC with Iterative Learning](#), *21st IFAC World Congress*, 2020.
- [C4] **H. Hu**, K. Gatsis, M. Morari, and G. J. Pappas. [Tuning Communication Latency for Distributed Model Predictive Control](#), *IFAC Conference on Networked Systems (NecSys)*, 2019.
- [C3] X. Feng, **H. Hu**, M. E. Villanueva, and B. Houska. [Min-Max Differential Inequalities for Polytopic Tube MPC](#), *American Control Conference (ACC)*, 2019.
- [C2] **H. Hu**, Y. Pu, M. Chen, and C. J. Tomlin. [Plug and Play Distributed Model Predictive Control for Heavy Duty Vehicle Platooning and Interaction with Passenger Vehicles](#), *IEEE Conference on Decision and Control (CDC)*, 2018.
- [C1] **H. Hu**, X. Feng, R. Quirynen, M. E. Villanueva, and B. Houska. [Real-Time Tube MPC Applied to a 10-State Quadrotor Model](#), *American Control Conference (ACC)*, 2018.

PREPRINTS

- [P2] K. Liang, **H. Hu**, X. Zhao, D. Song, T. L. Griffiths, and J. F. Fisac. [Machine Bullshit: Characterizing the Emergent Disregard for Truth in Large Language Models](#), 2025.
- [P1] K. Liang, **H. Hu**, R. Liu, T. L. Griffiths, and J. F. Fisac. [RLHS: Mitigating Misalignment in RLHF with Hindsight Simulation](#), 2024.

TEACHING

- Safety-Critical Robotic Systems** (Graduate) Princeton University, Fall 2021, 2022
Head Teaching Assistant; Delivered guest lectures on robust motion planning; Instructor: Jaime F. Fisac
- Model Predictive Control** (Graduate) University of Pennsylvania, Spring 2020
Head Teaching Assistant; Delivered a guest lecture on invariant sets; Instructor: Manfred Morari
- Science and Art Summer School** (Undergraduate) University of Padova, Summer 2024
Lecturer in Robotics; Delivered a guest lecture on HRI; Course Director: Ernesto Caraffoli

INVITED TALKS

- UPenn GRASP Student Faculty Industry (SFI) Seminar** 2025
- Johns Hopkins University, CS, ECE, & LCSR Joint Seminar**
- Duke University, Thomas Lord Department of MEMS Seminar**
- Texas A&M University, CS & ECE Joint Seminar**
- George Washington University, MAE Department Seminar**
- University of Colorado Boulder, ECEE Department Seminar**

North Carolina State University, ECE Department Seminar	
UCSD Safe Autonomous Systems Lab	2024
From Gambits to Assurances: Game-Theoretic Integration of Safety and Learning for Human-Centered Robotics	
IEEE CDC Workshop on Bifurcations for Control	2024
Constructing Nonlinear Opinion Dynamics for Split-Second Game-Theoretic Interactions	
Scania CV AB KTH Royal Institute of Technology	2024
The Safety Filter: A Unified View of Safety-Critical Control in Autonomous Systems	
University of Padova Science and Art Summer School	2024
Robotics: Toward Safer and More Efficient Interaction with Humans	
Northeast Systems and Control Workshop	2024
Human-Aware Safety Filters: A Unified View of Safety-Critical Control, Learning, and Interaction	
UPenn Vijay Kumar Lab CMU Safe AI Lab UIUC Human-Centered Autonomy Lab	2024
Who Plays First? Optimizing the Order of Play in Stackelberg Games with Many Robots	
Creative Convergence Workshop	2023
Nonlinear Opinion Dynamics Meets Game Theory: Split-second Decisions for Interactive Robotics	
Toyota Research Institute	2023
Learning to Break Indecisions in Multi-Agent Interaction	
University of Colorado Boulder Honda Research Institute USA, Inc.	2022
Active Uncertainty Reduction for Human-Robot Interaction	
Princeton University, Robotics Seminar	2021
Shielding-Aware Robust Planning for Safe and Efficient Human-Robot Interaction	
ETH Zürich, Melanie Zeilinger Lab	2020
Distributed Model Predictive Control for the Internet-of-Things	
University of Pennsylvania, George Pappas Lab	2018
Plug and Play Distributed MPC for Heavy Duty Vehicle Platooning	
Chinese Academy of Sciences, Shanghai Branch	2018
Real-time Tube-based Model Predictive Control	
UC Berkeley, Hybrid Systems Laboratory	2017
Model Predictive Control for Fast and Guaranteed Safe Motion Planning	

ADVISING AND
MENTORSHIP

PhD students

Wei Wang (PhD student at Penn ESE)
Donggeon David Oh (PhD student at Princeton ECE)
Justin Wang (PhD student at Princeton ECE)
Madison Bland (PhD student at Princeton ECE)

Master's students

Albert Gimó Contreras (visiting student at Princeton, recipient of the "La Caixa" Fellowship)
Xinyi Liu (UMich/Honda Research Institute → PhD at UNC/NCSU Joint Department of BME)

Undergraduate students

Kensuke Nakamura (Princeton ECE → PhD student at CMU)
Vikash Modi (Princeton ECE → SpaceX)
Elle Lazarski (Princeton ECE)

OUTREACH **Program Co-organizer, Princeton AI4ALL** 2022 - 2024
The summer camp teaches AI technology and policy to high school students, to broaden access to AI and inspire our future leaders to use it ethically.

Lecturer in Robotics, University of Padova Science and Art Summer School 2024
The summer school offers broad topics in science and art to undergraduate students in China and Italy.

OPEN-SOURCE SOFTWARE **BranchAndPlay**: An efficient solver that finds the socially-optimal Stackelberg (leader-follower) equilibrium for general-sum dynamic games.

KLGame: An efficient solver for non-cooperative dynamic games with Kullback-Leibler (KL) regularization with respect to a general, stochastic, and possibly multi-modal reference policy.

DeceptionGame: Approximate HJI Reachability analysis for robot safety in joint physical-belief spaces using adversarial reinforcement learning.

OpinionGame: Automatic synthesis of nonlinear opinion dynamics using dynamic games.

JaxRacing: JAX-based implementation of ILQR for safe autonomous racing.

DualControlHRI: Active uncertainty reduction for human-robot interaction using dual control.

SHARP: A human-aware motion planner that preempts costly safety-enforcing maneuvers triggered by unlikely human actions.

LearningCBFs: Learning a provably valid control barrier function (CBF) from expert data.

MEDIA COVERAGE **Game-theoretic integration of safety, interaction and learning** AIhub, 2025

Machine Bullshit: Disregard for truth in LLMs New Scientist & IEEE Spectrum, 2025

Data drives quicker, safer decisions for race cars and robots Princeton EQuad News, 2024

Who plays first? Using game theory to coordinate many robots Princeton Research, 2024

PROFESSIONAL ACTIVITIES **Associate Editor & Program Committee**

- IEEE Robotics and Automation Letters (RA-L) 2024 - 2027
- IEEE Conference on Advanced Robotics and its Social Impact (ARSO) 2026
- IEEE Workshop on Physical Intelligence: Systems and Applications (PISA) 2026
- **HRI Pioneers**, Program Chair 2024

Workshop & Seminar Organization

- **3rd Workshop on Multi-Agent Dynamic Games (MAD-Games)**, Main Organizer 2026
- **ICRA Workshop on Public Trust in Autonomous Systems**, Main Organizer 2025
- **Princeton Robotics Seminar**, Main Organizer 2024 - 2025

External Reviewer for Journals

- International Journal of Robotics Research (IJRR)
- IEEE Transactions on Robotics (T-RO)
- Automatica
- IEEE Transactions on Automatic Control (TAC)
- IEEE Transactions on Control Systems Technology (TCST)

- IEEE Transactions on Control of Network Systems (TCNS)
- IEEE Robotics and Automation Letters (RA-L)
- IEEE Control Systems Letters (L-CSS)
- International Journal of Robust and Nonlinear Control
- Optimal Control Applications and Methods
- IET Control Theory & Applications
- Chemical Engineering Science
- IEEE Access

External Reviewer for Conferences and Workshops

- Robotics: Science and Systems (RSS)
- RSS Pioneers – Faculty Reviewer
- Conference on Robot Learning (CoRL)
- World Symposium on the Algorithmic Foundations of Robotics (WAFR)
- ACM/IEEE International Conference on Human–Robot Interaction (HRI)
- IEEE International Conference on Robotics and Automation (ICRA)
- IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
- IEEE Conference on Decision and Control (CDC)
- American Control Conference (ACC)
- IFAC World Congress
- Annual Learning for Dynamics & Control Conference (L4DC)
- International Association for Safe & Ethical AI (IASEAI)
- International Conference on Neuro-symbolic Systems (NeuS)
- IFAC Conference on Networked Systems (NecSys)
- Annual Conference on Information Sciences and Systems (CISS)
- RSS Safe Autonomy Workshop
- ICRA Workshop on Bridging the Lab-to-Real Gap