

# Justin M. Lidard

Email: [jlidard@princeton.edu](mailto:jlidard@princeton.edu)

Homepage: [jlidard.github.io](http://jlidard.github.io)

My research interests lie in improving generalization performance and safety of robotic systems around humans and allowing them to reason about uncertain outcomes in the presence of new, potentially out-of-distribution tasks, environments, and decision makers. I am excited to apply theoretical and algorithmic tools from machine learning, statistical analysis, and optimal control to solve real-world robotic problems: in particular, robust execution of real manipulation policies, uncertainty quantification of neural-network-based policies and LLM/VLM-based task planners, and multi-modal, game-theoretic interactions.

## Education

---

**Ph.D. in Mechanical and Aerospace Engineering, Princeton University** 2020 – Present

**Bachelor of Science in Aerospace Engineering, University of Maryland, College Park** 2016–2020  
Minor in Computer Science. GPA: 3.99/4.00

## Experiences

---

**Graduate Researcher, Princeton University** 2019 – Present

I work in the Intelligent Robot Motion (IRoM) Lab advised by Prof. Anirudha Majumdar and Prof. Naomi Leonard. My work centers on enabling robotic policies, particularly in human-robot interaction tasks, to effectively reason about uncertainty from environmental factors, task dynamics, and interactions with other decision-making agents. Additionally, I focus on making these policies more human-like by anticipating a variety of intents and incentives.

**Research Intern, Toyota Research Institute (TRI)** 2022 Summer

Worked on improving coverage of rare events for large-scale trajectory prediction models for autonomous vehicles using techniques from game theory and stochastic optimal control. My work allowed these models to better capture the distribution of closed-loop behaviors by ensuring that the predictions are safe, and smooth, and interactive.

**Machine Learning Intern, Aurora Flight Sciences** 2020 Summer

Worked on developing neural network models to predict the outcome of multi-agent games.

**Undergraduate Researcher, University of Maryland** 2018 – 2020

I worked in the Collective Dynamics and Control Lab advised by Prof. Derek Paley. My research involved developing a nonlinear dynamics model for an actuated pitching airfoil, optimizing its periodic solution for enhanced lift using techniques from bifurcation theory.

## Ongoing Projects

---

**Enabling Human-Guided Robotic Construction with Large Vision-Language Models** 2024–Present

This project seeks to enable rapid design iteration for on-site robotic construction guided by a human. Our system allows a robot to leverage in-context learning, wherein a large vision-language model queries examples of real construction scenarios, from initial design to finished product. The robot interacts with the human to finalize the design. We provide a guarantee on the human's satisfaction with the final design with Conformal Prediction.

**Diffusion Policy Policy Optimization** 2024–Present

Diffusion Policies have become state-of-the-art for learning complex robot trajectories from diverse data sources. However, since they are learned from imperfect demonstrations, they may be unreliable, and the datasets they are trained on may fail to cover the optimal behavior for the task. In this project, we show that Diffusion Policies can be fine-tuned using reinforcement learning (RL). Our method, which includes

a set of best practices, is able to stably and reliably train long-horizon and sparse-reward manipulation tasks to a high success rate.

### **Blending Data-Driven Priors in Dynamic Games**

2024–Present

Large-scale imitation learning has become commonplace in many robot learning domains, especially in autonomous driving. However, imitation-pretrained policies can fail in closed-loop due to the non-interactive nature of their training datasets. In this project, we introduce a framework wherein autonomous vehicle planners can benefit from broadly generalizable imitation models but leverage runtime game-theoretic planning for tuneably interactive behavior.

### **Fast, Flexible, and Transparent Strategy Alignment for AI-Assisted Racing**

2024–Present

While AI agents can reach superhuman prowess in fully autonomous racing in just a few days of training time, humans take much longer to build their proficiency. In this project (a collaboration with TRI), we seek to enable AI agents to reliably “coach” human drivers to reach a higher level of performance. We construct a two-layer safety filter to ensure safety during the learning process: an outer layer “nudges” the human towards optimal strategies, and an inner layer maintains the strict physical safety of the vehicle.

## **Publications**

---

### **[7] Diffusion Policy Policy Optimization**

A. Z. Ren, J. Lidard, L. L. Ankile, A. Simeonov, P. Agrawal, A. Majumdar, B. Burchfiel, H. Dai, M. Simchowitz. *Under review for ICLR, 2024.*

### **[6] Risk-Calibrated Human-Robot Interaction via Set-Valued Intent Prediction**

J. Lidard, H. Pham, A. Bachman, B. Boateng, A. Majumdar. *Robotics: Science and Systems (RSS), 2024.*

### **[5] Blending Data-Driven Priors in Dynamic Games**

J. Lidard\*, H. Hu\*, A. Hancock, Z. Zhang, A. G. Contreras, V. Modi, J. DeCastro, D. Gopinath, G. Rosman, N. E. Leonard, M. Santos, J. F. Fisac. *Robotics: Science and Systems (RSS), 2024.*

### **[4] NashFormer: Leveraging Local Nash Equilibria for Semantically Diverse Trajectory Prediction**

J. Lidard, O. So, Y. Zhang, J. DeCastro, X. Cui, X. Huang, Y. L. Kuo, J. Leonard, A. Balachandran, N. E. Leonard, G. Rosman. *Under review, 2023.*

### **[3] Provably Efficient Multi-Agent Reinforcement Learning with Fully Decentralized Communication**

J. Lidard, U. Madhushani, N. E. Leonard. *American Control Conference, 2022.*

### **[2] Feedback Control and Parameter Estimation for Lift Maximization of a Pitching Airfoil**

J. M. Lidard, D. Goswami, D. Snyder, G. Sedky, A. R. Jones, D. A. Paley. *Journal of Guidance, Control, and Dynamics, 2021.*

### **[1] Output Feedback Control for Lift Maximization of a Pitching Airfoil**

J. Lidard, D. Goswami, D. Snyder, G. Sedky, A. R. Jones, D. A. Paley. *AIAA Scitech Forum, 2020.*

## **Workshops**

---

### **[W3] Risk-Calibrated Human-Robot Interaction via Set-Valued Intent Prediction**

J. Lidard, H. Pham, A. Bachman, B. Boateng, A. Majumdar. *Back to the Future: Robot Learning Going Probabilistic (Best Paper Award), IEEE International Conference on Robotics and Automation (ICRA), 2024.*

### **[W2] Blending Data-Driven Priors in Dynamic Games**

J. Lidard\*, H. Hu\*, A. Hancock, Z. Zhang, A. G. Contreras, V. Modi, J. DeCastro, D. Gopinath, G. Rosman, N. E. Leonard, M. Santos, F. Fisac. *Northeast Systems and Control Workshop (Oral), 2024.*

### **[W1] Provably Efficient Multi-Agent Reinforcement Learning with Fully Decentralized Communication**

J. Lidard, U. Madhushani, N. E. Leonard. *Workshop on Reinforcement Learning in Games (Oral), AIAA, 2022.*

## Awards and Honors

---

- **NSF Graduate Research Fellowship, National Science Foundation** 2022  
The GRFP provides three years of funding for graduate research.
- **Summerfield Fellowship, Princeton University** 2021  
One year funding award for performance in the second year.
- **NDSEG Fellowship, Air Force Office of Scientific Research** 2020  
The NDSEG provides up to three years of funding for graduate research.
- **Space Technology Graduate Research Fellowship (declined), NASA** 2020  
The NSRF provides up to four years of funding for graduate research in space technology.
- **Francis Robbins Upton Fellowship, Princeton University** 2020  
Granted to incoming students based on research and academic achievement.

## Teaching Experience

---

### Teaching Assistant

- Fall 2023, MAE 577 Multi-robot Systems, Princeton MAE department.
- Fall 2022, MAE 545 Collective Intelligence, Princeton MAE department.
- Spring 2022, MAE 434 Modern Control, Princeton MAE department.
- Spring 2019, ENAE380: Flight Software, UMCP Aero department.

## Academic Service

---

### Undergraduate research mentoring

- William Li (MAE), Princeton 2023 - Present
- Christina Zhang (COS), Princeton 2023 - Present
- Hang Pham (COS), Princeton 2023 - Present
- Bryan Boateng (MAE), Princeton 2023 - Present
- Ariel Bachman (MAE), Princeton 2023 - 2024
- Christine Ohenzuwa (MAE), Princeton 2022 - 2023

### Reviewing

- IEEE Robotics and Automation Letters (RA-L) 2024 - Present
- IEEE International Conference on Robotics and Automation (ICRA) 2024 - Present

### Miscellaneous

- Co-organizer, Princeton Symposium on Safe Deployment of Foundation Models Nov 2024

## Key Skills

---

**Topics:** Robotics (Kinematics, Dynamics, Motion Planning, Feedback Control, Grasping), Machine Learning (Regression, Classification, SVM), Deep Learning (Neural Networks, CNNs, RNNs, Transformers), Reinforcement Learning, Imitation Learning (Diffusion Policies, Boltzmann Models, Inverse Reinforcement Learning), Large (Vision-)Language Models, Dynamic Game Theory (Nash Equilibria, Stackelberg Equilibria, Potential Games, Mixed Strategies).

**Software:** Python, C/C++, MATLAB, HTML, Unix scripting; Pytorch, Robot Operating System (ROS), MeshLab; PyBullet, Mujoco, IsaacGym, High-performance computing (HPC).

## References

---

**Anirudha Majumdar**

Associate Professor  
Princeton University  
Mechanical and Aerospace Engineering  
ani.majumdar@princeton.edu  
(609) 258-0854

**Jaime F. Fisac**

Assistant Professor  
Princeton University  
Electrical and Computer Engineering  
jfisac@princeton.edu  
(609) 258-2017

**Sildomar Monteiro**

Manager  
Aurora Flight Sciences  
AI & ML Group (Lecturer, MIT)  
smonteiro@mit.edu  
(703) 369-3633

**Naomi E. Leonard**

Edwin S. Wilsey Professor, Department Chair  
Princeton University  
Mechanical and Aerospace Engineering  
naomi@princeton.edu  
(609) 258-5129

**Guy Rosman**

Manager  
Toyota Research Institute (TRI)  
Human Aware Interaction & Learning  
guy.rosman@tri.global  
(650) 673-2225