Justin Lidard | Curriculum Vitae

Version: January 2024

PhD Candidate, Department of Mechanical and Aerospace Engineering, Princeton University

jlidard@princeton.edu • jlidard.github.io/site/ • Google Scholar

Research Interests

- **Game Theory**
- **Human-Robot interaction (HRI)**
- **Stochastic Optimal Control**
- **Deep Reinforcement and Imitation Learning**

Skills

Languages: Python
Technologies: PyTorch, Jax, CVX, Git, Slurm, AWS, High-Performance Computing
Simulators: Waymax, Habitat, PyBullet, Gym

Education

Princeton University — **Princeton, NJ**
Doctor of Philosophy in Mechanical and Aerospace Engineering — *Sept. 2020 – May. 2025*
National Science Foundation Graduate Research Fellowship
GPA: 3.95

University of Maryland, College Park — **College Park, MD**
Bachelor of Science in Aerospace Engineering, Summa Cum Laude (top 2%) — *Sept. 2016 – May 2020*
Minor: *Computer Science*
Aerospace Honors Research Thesis: *State and Output Feedback Control for Lift Maximization of a Pitching Airfoil*
GPA: 3.99

Research and Internship Experience

Intelligent Robot Motion Lab and Leonard Lab , Princeton University — **Princeton, NJ**
Graduate Research Assistant — *Jan. 2021 – Present*
Advisors: Prof. Anirudha Majumdar and Prof. Naomi Leonard

- Introduce a framework for sequence-level, distribution-free risk control for set-valued predictors, with a specific focus on optimal action selection in the face of highly multi-modal human-robot interactive scenarios. Show that an optimal policy that *conditions on human intent* can be learned using *privileged* knowledge of the latent human intent at during train time, and predict the latent human intent at test time. Using conformal prediction theory and statistical risk control, prove that multiple, non-monotonic risks can be controlled simultaneously for a certain class of sequence-level risk function. Propose instantaneous and sequence-level PAC bounds based on the Hoeffding-Bentkus inequality and a union bound over multiple hypothesis for risk control. Evaluate theoretical claims on simulated and hardware experiments: joint motion planning in PyBullet, social navigation in Habitat, and cooperative manipulation using a Franka Emika arm. Planning to submit to RSS 2024. Ongoing.
- Propose a game-theoretic planning framework (an extension of iLQGames) for autonomous driving that incorporates data-driven priors through a Kullback–Leibler (KL) regularization term, inducing a mixed strategy. Prove a closed form solution for the optimal multi-modal mixed strategy when the reference policy is a Gaussian mixture using the dual formulation of the well-known Donsker-Varadhan definition of the KL divergence. Prove that for a linear-quadratic game with a unimodal Gaussian reference policy, a local Nash Equilibrium is achieved in a single backward pass. Evaluate results in several autonomous racing and real scenarios from the Waymo Open Motion Dataset. Planning to submit to RSS 2024. Ongoing.
- Design and implement a novel game-theoretic framework for motion prediction for incentivizing coverage of diverse outcomes in the space of utilities. Submitted to IEEE Robotics and Automation Letters (RA-L). See 2022 TRI internship for more information.
- Design and provide theoretical analysis for a decentralized, online multi-agent Q-learning algorithm in which agents use message passing to share samples. Provide regret and probably-approximately-correct (PAC) bounds for provably faster learning than in the single-agent case. Published in American Control Conference 2022.

Toyota Research Institute (TRI) **Cambridge, MA**
Research Scientist Intern *May 2022 – Aug. 2022*
Mentor: Dr. Guy Rosman, **Program:** Human-Centric Driving Research

- Implement a 100M-parameter transformer model for game-theoretic joint motion prediction in autonomous driving. The quality and coverage of joint motion predictions is improved by regressing utilities of goal-conditioned trajectories. The local optima of the regressed utilities provide an additional mechanism for distinguishing predictions.
- Derive and implement novel loss to encourage sampling of *distinct* local Nash equilbria (LNE) during inference time. Also provide a novel sampling procedure, called *non-equilibrium suppression*, that samples from all available LNE without additional training. Other distance-based algorithms, such as non-maximum suppression (NMS), may repeatedly sample from the same local optima, limiting sample diversity.
- Validate findings on Waymo Interactive dataset (45k traffic scenes with vehicles, pedestrians, and cyclists), which provides especially challenging, multi-agent scenarios for prediction. Also validate on several high-order interactive splits of the training set (5k, 20k, 100k of 500k traffic scenes) exhibiting network yield interactions of varying complexity. Submit work to IEEE RA-L.

Aurora Flight Sciences Research and Development Center (Boeing) **Cambridge, MA**
Machine Learning Engineering Intern *Jun. 2020 – Aug. 2020*
Mentor: Dr. Sildomar Monteiro, **Program:** DARPA "Gamebreaker"

- Developed deep neural network, random forest, and logistic regression models in PyTorch for predicting outcomes of StarCraft II played by reinforcement learning agents.
- Implemented a performance-based skill rating system to rank both individual agents and teams in parameter sensitivity experiments. Showed that equilibria can be predicted with ranking system for both balanced and unbalanced parameter settings.

Collective Dynamics and Control Lab, University of Maryland **College Park, MD**
Undergraduate Research Assistant *Jun. 2018 – Jun. 2020*
Mentor: Prof. Derek Paley, **Program:** AFOSR

- Developed a first-order nonlinear dynamics model for an actuated pitching airfoil.
- Optimized periodic solution to produce 40% more lift output versus the optimal steady-state solution. Simulated full nonlinear system with output feedback using a recursive Bayesian filter.
- Published first-authored paper in AIAA Journal of Guidance, Control, and Dynamics.

Awards

1. NSF Graduate Research Fellowship, National Science Foundation, September 2022
2. Summerfield Second Year Fellowship, Princeton MAE, September 2021
3. 20 Twenties, Aviation Week Network, July 2021
4. DOD NDSEG Fellowship, Department of Defense, April 2020
5. NASA Space Technology Graduate Research Fellowship (declined), NASA, April 2020
6. Francis Robbins Upton Fellowship, Princeton University, January 2020

Publications

1. **J. Lidard**, A. Bachman, N. Leonard, and A. Majumdar. IntentMesh: Intent-Informed Human-Robot Cooperation with Generalization Guarantees. In preparation for RSS.
2. **J. Lidard**, **H. Hu**, Z. Zhang, A. Hancock, V. Modi, María Santos, N. Leonard, and J. Fisac. KL-Regularized General-Sum Dynamic Games: Incorporating Data-Driven Priors in Multi-Agent Planning. In preparation for RSS.
3. **J. Lidard**, O. So, Y. Zhang, J. DeCastro, X. Cui, X. Huang, Y. Kuo, J. Leonard, A. Balachandran, N. Leonard and G. Rosman. NashFormer: Leveraging Local Nash Equilibria for Semantically Diverse Trajectory Prediction. Submitted to RA-L.
4. **J. Lidard**, U. Madhushani, and N. E. Leonard. Provably efficient multi-agent reinforcement learning with fully decentralized communication. American Control Conference, Atlanta, GA, 2022.
5. **J. Lidard**, D. Goswami, D. Snyder, and D. A. Paley. Adaptive output feedback control for lift maximization of a pitching airfoil. AIAA Journal of Guidance, Control and Dynamics. January 2021.
6. **J. Lidard**, D. Goswami, D. Snyder, G. Sedky, A. Jones, and D. A. Paley. Output feedback control for lift maximization of a pitching airfoil. American Institute of Aeronautics and Astronautics (AIAA) Science

and Technology (SciTech) Forum, Orlando, FL. January 2020.

Patents

1. **J. Lidard**, O. So, Y. Zhang, J. DeCastro, X. Cui, X. Huang, Y. Kuo, J. Leonard, A. Balachandran, N. Leonard and G. Rosman. GAME-UP: Game-Aware Mode Enumeration and Understanding for Trajectory Prediction. Submitted.

Presentations

1. "Provably efficient multi-agent reinforcement learning with fully decentralized communication." American Control Conference, Atlanta, GA, 2022.
2. "Provably efficient decentralized communication for multi-agent RL." AAAI Workshop on Reinforcement Learning in Games (**Oral, Top 10% of submissions**), Vancouver, British Columbia, 2022.
3. "Output feedback control for lift maximization of a pitching airfoil." AIAA SciTech Forum, Orlando, Florida, January 2020.