Dhruv Malik

2734 College Avenue, Berkeley, California 94705 United States
(425) 615-9054 | dhruvmalik@berkeley.edu

Education

University of California, Berkeley — Berkeley, CA
B.A. Applied Mathematics **GPA: 3.88/4.00** — Class of 2018
B.A Computer Science

Relevant Coursework: Linear Algebra, Probability Theory, Abstract Algebra, Algorithms, Artificial Intelligence, Real Analysis, Data Structures, Complex Analysis, Discrete Math, Multivariable Calculus, Structure & Interpretation of Computer Programs, Numerical Analysis*, Machine Learning*

*in progress

Publications

Malayandi Palaniappan*, **Dhruv Malik***, Jaime F. Fisac, Dylan Hadfield-Menell, Anca D. Dragan and Stuart Russell. Efficient Cooperative Inverse Reinforcement Learning. In *Proceedings of the Thiry Second AAAI Conference on Artificial Intelligence*, 2018. [currently in review]

Malayandi Palaniappan*, **Dhruv Malik***, Dylan Hadfield-Menell, Anca D. Dragan and Stuart Russell. Efficient Cooperative Inverse Reinforcement Learning. In *Reliable Machine Learning In The Wild Workshop* at *ICML*, 2017.

Jaime F. Fisac, Monica A. Gates, Jessica B. Hamrick, Chang Liu, Dylan Hadfield-Menell, Malayandi Palaniappan, **Dhruv Malik**, S. Shankar Sastry, Thomas L. Griffiths and Anca D. Dragan. Pragmatic Pedagogic Value Alignment. In *ISRR*, 2017.

*equal contribution

Research Experience

Center For Human Compatible Artificial Intelligence — Berkeley, CA
Research Intern — January 2017 – present

- Contributed as first author to major conference papers.
- Developed ideas for new algorithms to more efficiently compute policy equilibria in CIRL games and other multiagent systems.
- Wrote code and devised toy examples to experimentally test efficiency of our new algorithms against current state of the art algorithms.
- Advised by Anca Dragan and Stuart Russell.

Feldman Lab — Berkeley, CA
Research Intern — May 2014 – August 2016

- Developed a program which analyzes how well an animal performs on a stimulus recognition task. Algorithm returns a statistical measure for the probability that the animal may be using pattern recognition or other confounding cues to perform the task, as opposed to perfect sensory driven performance.

Projects

- **Partially Observable Monte Carlo Planning** – Implemented the POMCP algorithm for solving POMDPs, as described in Silver et al. (2010).
- **FV-POMCP** – Implemented the factored statistics version of the POMCP algorithm, as described in Amato et al. (2015), to solve a coordinator POMDP from a reduced CIRL game.
- **POMDP Value Iteration** – Implemented the dynamic programming exact value iteration algorithm for solving POMDPs.