DIBYA GHOSH

dibya@berkeley.edu (925) 725-5956

EDUCATION

B.A. Computer Science, Applied Mathematics Expected May 2019
University of California, Berkeley
Overall GPA: 4.00
Honors: Quantedge Award, Robert Kraft Award, Deans List (6 semesters)

Graduate Coursework: Theoretical Statistics, Computer Vision, Probability, Information Theory
Selected Undergraduate Coursework - Machine Learning, Convex Optimization, Algorithms, Randomized Algorithms, Stochastic Processes, Game Theory, Operating Systems, Computer Architecture

RESEARCH

Berkeley Artificial Intelligence Research Lab May 2017 - Present
Advised by Professor Sergey Levine

My research focuses on building flexible and powerful deep reinforcement learning algorithms better suited for solving complex real-world robotic problems.

Divide-and-Conquer Reinforcement Learning
Published at the International Conference on Learning Representations 2018

Problems that exhibit considerable environment diversity make direct policy optimization with deep RL algorithms challenging. In this work, we present a novel algorithm which is designed to solve such tasks. The algorithm partitions the initial state space into "slices", optimizing local policies on each slice, which are gradually unified into a single policy that succeeds on the full task. We evaluate on challenging grasping, manipulation, and locomotion tasks, and show that our method fully solves the task, where prior methods fail to learn completely. I am the lead contributor on this project.

Variational Inverse Control with Events
Published at the Conference on Neural Information Processing Systems 2018

The design of a reward function often poses a major practical challenge to real-world applications of reinforcement learning. In this work, we propose an algorithm to learn reward functions that correspond to reaching a provided set of goal images. Our method generalizes inverse reinforcement learning methods to requiring only goal images instead of full demonstrations, which is much simpler to provide in practice. We demonstrate the effectiveness of our methods on continuous control tasks from images where rewards are hard or even impossible to specify.

Learning Actionable Representations with Goal Conditioned Policies
Under review at the International Conference on Learning Representations 2019

Effective and functional representations in RL have the potential to tremendously accelerate learning progress and solve more challenging problems. In this paper, we aim to learn functionally salient representations: representations that capture *only* those factors of variation that are important for decision making – that are "actionable". We show how these learned representations can be useful to improve exploration for sparse reward problems, to enable long horizon hierarchical reinforcement learning, and as a state representation for learning policies for downstream tasks. I am the lead contributor on this project.

Lawrence Berkeley National Laboratory March 2016 - Present
Advised by Professor Ben Brown

My research focuses on developing interpretable and explainable machine learning algorithms for analyzing high-dimensional scientific systems, such as genomics and agriculture.

Iterative Random Forests

Understanding how high-order interactions in transcriptomes drive gene expression presents a substantial statistical challenge. Building on Random Forests, and Random Intersection Trees, we developed the iterative Random Forest algorithm (iRF), which trains a feature-weighted ensemble of decision trees to detect stable, high-order interactions without additional computational overhead. The first iteration of this paper was published in the Proceedings of the National Academy of Sciences (PNAS).

Interpretable Density Estimation in Genomics Data.
Presented at the Platform for Advanced Scientific Computing 2018

Deep generative models cannot tractably be used to perform search or optimize auxiliary objectives. We introduce a novel adversarial training algorithm for explicit density models, which iterates between fitting a discriminator random forest, and approximating a density using feature interactions discovered in the forest. This parallelizable method can find high-order interactions and and produce density models for these interactions that can be used in constrained optimization problems. I am the lead contributor on this project.

PUBLICATIONS

Learning Actionable Representations with Goal-Conditioned Policies.
Dibya Ghosh, Abhishek Gupta, Sergey Levine.
Under review at *International Conference on Learning Representations (ICLR) 2019.*

Variational Inverse Control with Events: A General Framework for Data-Driven Reward Definition.
Justin Fu*, Avi Singh*, **Dibya Ghosh**, Larry Yang, Sergey Levine.
Conference on Neural Information Processing Systems (NIPS) 2018.

Divide-and-Conquer Reinforcement Learning.
Dibya Ghosh, Avi Singh, Aravind Rajeswaran, Vikash Kumar, Sergey Levine.
International Conference on Learning Representations (ICLR) 2018.

Theory Meets Data.
Ani Adhikari, **Dibya Ghosh**, et al.
Online textbook.

TALKS

Interpretable Density Estimation in Genomics Data.
Contributed talk at the *Platform for Advanced Scientific Computing (PASC) Conference 2018.*

Iterative Random Forests for Explainable Machine Learning.
Invited talk at the *Machine Learning for Science (ML4Sci) Workshop* 2018.

TEACHING

University of California, Berkeley

Head Teaching Assistant, Stat 140 - Probability for Data Science — Spring 2017,Spring 2018
Head Teaching Assistant, Stat 134 - Concepts in Probability — Fall 2017
Teaching Assistant, Stat 140 - Probability for Data Science — Fall 2018

Course Developer, Stat 140 - Probability for Data Science — Fall 2016 - Present
Course Development Assistant, Stat C8 - Foundations of Data Science — Spring 2016

INDUSTRY

Preminon January 2017 - December 2017
Machine Learning Developer