

Jalaj Bhandari

Columbia University
345, SEAS, New York, NY - 10027

jb3618@columbia.edu
<http://www.columbia.edu/~jb3618/>

Research Statement

My thesis work explores the foundations of Reinforcement learning algorithms from an optimization perspective. I have also done research work in approximate Bayesian inference, specifically in designing efficient Markov Chain Monte Carlo (MCMC) methods for posterior sampling.

Education

Columbia University, New York (2013 - Present)

PhD, Operations Research

Advisors: Garud Iyengar and Daniel Russo.

GPA: 4.0/4.0

Indian Institute of Technology, Delhi (2008 - 2012)

Bachelor of Technology (B.Tech), Industrial Engineering and Operations Research

GPA: 9.1/10

Department Rank 1

Publications

Global Optimality Guarantees For Policy Gradient Methods

Jalaj Bhandari and Daniel Russo

OptRL workshop, NeurIPS 2019; RLDM, 2019

Journal version in preparation

A finite time analysis of Temporal Difference learning with linear function approximation

Jalaj Bhandari, Daniel Russo and Raghav Singal

Conference on Learning Theory (COLT), 2018; OptRL workshop, NeurIPS 2019; RLDM, 2019

Journal version accepted to Operations Research with minor revisions.

Annular Augmentation sampling

Francois Fagan, Jalaj Bhandari and John Cunningham

Oral presentation (top 5%), Conference on Artificial Intelligence and Statistics (AISTATS), 2017

Elliptical Slice Sampling with Expectation Propagation

Francois Fagan, Jalaj Bhandari and John Cunningham

Conference on Uncertainty in Artificial Intelligence (UAI), 2016

On the tightness of an LP relaxation for rational optimization and its applications

Vashist Avadhanula, Jalaj Bhandari, Vineet Goyal and Assaf Zeevi

Operations Research Letters, 2016.

Current research projects

Tractable Deep Q-learning for Continuous Control

Prof. Garud Iyengar, Columbia University and Krzysztof Choromanski, Google, NYC

We consider Reinforcement learning for high dimensional continuous action spaces and propose a tractable approximate Q-learning framework using neural networks with polynomial activations.

Professional Experience

Research Intern, INVITAE, NYC (June - September, 2019)

Proposed a simple posterior sampling based approach to balance between exploration and exploitation in order to recommend (genetic testing) tasks to scientists at INVITAE.

Research Intern, Adobe, California (June - September, 2018)

Developed a Reinforcement learning framework for personalized marketing using Deep Q-learning, with the goal to maximize long-run customer engagement.

Awards and Scholarships

Travel awards: RLDM, 2019

Fellowship, Columbia University, 2013-2018

Institute Silver Medal (Indian Institute of Technology, Delhi)

Institute Merit Scholarship (top 5 % of all students) for 6/8 semesters (Indian Institute of Technology, Delhi)

Coursework

Dynamic programming and Reinforcement learning, Deep Reinforcement learning, Deep learning for computer vision, Computational aspects of robotics, Gaussian processes, Bayesian methods, Convex optimization, Graphical models.

Additional Skills

Programming: Python, Tensorflow, R, C++ (undergraduate work)