

# KARAN SINGH

## CONTACT INFORMATION

Tepper School of Business  
5000 Forbes Avenue  
Pittsburgh, PA, USA 15213

✉ [karansingh@cmu.edu](mailto:karansingh@cmu.edu)  
🌐 [i-am-karan-singh.github.io](https://i-am-karan-singh.github.io)  
☎ +1 (609) 516 5555

## PROFESSIONAL AFFILIATIONS

**CARNEGIE MELLON UNIVERSITY, PITTSBURGH** September 2022 –  
*Assistant Professor of Operations Research, Tepper School of Business*

**MICROSOFT RESEARCH, REDMOND** April 2021 – August 2022  
*Postdoctoral Researcher, Reinforcement Learning Group* Manager: Sham M. Kakade

**GOOGLE AI, PRINCETON** 2018-2020  
*Student Researcher* Host: Yoram Singer

## EDUCATION

**PRINCETON UNIVERSITY** January 2022  
*PhD in Computer Science.* GPA: 4.0/4  
ADVISER: Elad Hazan  
DISSERTATION: The Nonstochastic Control Problem  
Awarded the Jacobus Fellowship, Princeton University's top graduate student honor.

**INDIAN INSTITUTE OF TECHNOLOGY, KANPUR** June 2015  
*Bachelor of Technology in Computer Science.* GPA: 10.0/10  
Received the *President's Gold Medal*, for the best academic performance in the outgoing class.

## FIELDS OF INTEREST

Machine Learning, Control, Reinforcement Learning, Bandits, Online Learning, Convex and Nonconvex Optimization, Learning Theory.

Specifically, provably efficient algorithms and fundamental limits for *feedback-driven* learning, spanning both *prediction* and *control*. Drawing from the algorithmic toolkits of *optimization* and *sequential prediction*, together with techniques from *dynamical systems* and *control theory*, recent results include the first instance-optimal control algorithm, and provably efficient prediction in dynamical systems that exhibit long-term correlations.

## SELECTED AWARDS, HONORS & ACHIEVEMENTS

- 2020 IBM Herman Goldstine Postdoctoral Fellowship in Mathematical Sciences (declined)
- 2019 Jacobus Fellowship, Princeton University's top graduate student honor ([press](#), [more](#))
- 2019 Best Paper Award, Optimization for RL workshop at NeurIPS 2019
- 2019 Best Reviewer (top 5%), NeurIPS 2019
- 2018-19 Selected twice for an Oral Presentation at NeurIPS – top 0.5% of the submissions
- 2018 SEAS Award for Excellence by the Graduate School at Princeton
- 2018 Spotlight Prize, New York Academy of Sciences' 12<sup>th</sup> Annual ML Symposium
- 2017 ICML Travel Grant

- 2015 President's Gold Medal for the best academic performance at IIT Kanpur (**press**)  
Ranked 1st (among 820 students) at IIT Kanpur
- 2012-2015 Academic Excellence Award for 3 years at IIT Kanpur
- 2011 All-India Rank 14 in AIEEE 2011 among 1,050,000 students
- 2011 All-India Rank 140 in IIT-JEE 2011 among 485,000 students
- 2011 Gold Medal, the top 35 (0.1%) students at Indian National Physics Olympiad 2011
- 2009 Kishore Vaigyanik Protsahan Yojana Fellowship, Government of India

## PEER-REVIEWED PUBLICATIONS

All publications list authors in the alphabetical order, except those indicated with †.

A Boosting Approach to Reinforcement Learning

Nataly Brukhim, Elad Hazan, Karan Singh

*Neural Information Processing Systems (NeurIPS), 2022*

**Boosting for Online Convex Optimization**

Elad Hazan, Karan Singh

*International Conference on Machine Learning (ICML), 2021*

**A Regret Minimization Approach to Iterative Learning Control**

Naman Agarwal, Elad Hazan, Anirudha Majumdar, Karan Singh

*International Conference on Machine Learning (ICML), 2021*

**Improper Learning for Nonstochastic Control†**

Max Simchowitz, Karan Singh, Elad Hazan

*Conference on Learning Theory (COLT), 2020*

**No-Regret Prediction in Marginally Stable Systems**

Udaya Ghai, Holden Lee, Karan Singh, Cyril Zhang, Yi Zhang

*Conference on Learning Theory (COLT), 2020*

**The Nonstochastic Control Problem**

Elad Hazan, Sham Kakade, Karan Singh

*Algorithmic Learning Theory (ALT), 2020*

**Logarithmic Regret for Online Control**

Naman Agarwal, Elad Hazan, Karan Singh

*Neural Information Processing Systems (NeurIPS), 2019, Oral Presentation (<0.5% of submissions)*

(Also, **Best Paper Award** at the OptRL workshop at NeurIPS 2019)

**Online Control with Adversarial Disturbances**

Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, Karan Singh

*International Conference on Machine Learning (ICML), 2019*

**Provably Efficient Maximum Entropy Exploration**

Elad Hazan, Sham Kakade, Karan Singh, Abby Van Soest

*International Conference on Machine Learning (ICML), 2019*

**Efficient Full-Matrix Adaptive Regularization**

Naman Agarwal, Brian Bullins, Xinyi Chen, Elad Hazan, Karan Singh, Cyril Zhang, Yi Zhang

*International Conference on Machine Learning (ICML), 2019*

### Spectral Filtering for General Linear Dynamical Systems

Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, Yi Zhang

*Neural Information Processing Systems (NeurIPS), 2018, Oral Presentation* (<0.5% of submissions)

### Learning Linear Dynamical Systems via Spectral Filtering

Elad Hazan, Karan Singh, Cyril Zhang

*Neural Information Processing Systems (NeurIPS), 2017, Spotlight* (<5% of submissions)  
(Also, **Spotlight Prize** at New York Academy of Sciences' ML Symposium 2018)

### The Price of Differential Privacy for Online Learning

Naman Agarwal, Karan Singh

*International Conference on Machine Learning (ICML), 2017*

### Efficient Regret Minimization in Non-Convex Games

Elad Hazan, Karan Singh, Cyril Zhang

*International Conference on Machine Learning (ICML), 2017*

## PREPRINTS, PATENTS AND TECHNICAL REPORTS

All publications list authors in the alphabetical order, except those indicated with †.

### Variance-Reduced Conservative Policy Iteration

Naman Agarwal, Brian Bullins, Karan Singh

### Best of Both Worlds in Online Control: Competitive Ratio and Policy Regret†

Gautam Goel, Naman Agarwal, Karan Singh, Elad Hazan

### Dynamic Learning System

Elad Hazan, Karan Singh, Cyril Zhang

US Patent 11,138,513 B2, approved Oct 2021

### Machine Learning for Mechanical Ventilation Control†

Daniel Suo, Cyril Zhang, Paula Gradu, Udaya Ghai, Xinyi Chen, Edgar Minasyan, Naman Agarwal, Karan Singh, Julianne LaChance, Tom Zajdel, Manuel Schottdorf, Daniel Cohen, Elad Hazan

Machine Learning for Health (ML4H), 2021 Workshop Track

**Featured** in Princeton Engineering news

### Deluca – A Differentiable Control Library: Environments, Methods, and Benchmarking†

Paula Gradu, John Hallman, Daniel Suo, Alex Yu, Naman Agarwal, Udaya Ghai, Karan Singh, Cyril Zhang, Anirudha Majumdar, Elad Hazan

NeurIPS Workshop on Differentiable Computer Vision & Physics, 2020 Oral Presentation

### Towards Provable Control for Unknown Linear Dynamical Systems

Sanjeev Arora, Elad Hazan, Holden Lee, Karan Singh, Cyril Zhang, Yi Zhang

International Conference on Learning Representations (ICLR), Workshop Track, 2018

### Dynamic Task Allocation for Crowdsourcing†

Angela Zhou, Irineo Cabrerros, Karan Singh

ICML Workshop on Data Efficient Machine Learning, 2016

## INTERNSHIPS

MICROSOFT RESEARCH, REDMOND

Undergrad Research Intern

*Summer 2014*

Host: Sumit Gulwani

## TEACHING

ECONOMICS AND COMPUTATION (COS 445)

Princeton University

*Fall 2017*

Assistant Instructor

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING (COS 402)

Princeton University

*Spring 2016*

Assistant Instructor

INTRODUCTION TO MACHINE LEARNING (COS 324)

Princeton University

*Fall 2016*

Assistant Instructor

DATA STRUCTURE & ALGORITHMS (ESO 207)

IIT Kanpur

*Fall 2014*

Teaching Assistant

## SERVICE

### PROGRAM COMMITTEE

Conference on Learning Theory (COLT) 2021, 2022

Algorithmic Learning Theory (ALT) 2021, 2022

AAAI Conference on Artificial Intelligence (AAAI), 2020

### REVIEWER

Conference on Learning Theory (COLT) 2017, 2018, 2020

International Conference on Machine Learning (ICML) 2018, 2019, 2020, 2021, 2022

Neural Information Processing Systems (NeurIPS) 2018, 2019, 2020, 2021, 2022

International Conference on Learning Representations (ICLR) 2020, 2021

Journal of Machine Learning Research

Mathematical Programming

OptRL Workshop, NeurIPS 2019

### GRADUATE STUDENTS ADMISSIONS COMMITTEE

*2018-2020*

Computer Science, Princeton University

### ORGANIZER

*2017-2019*

Alg-ML Reading Group, Princeton University

### WEBMASTER

*2017-2019*

ML Theory Website, Princeton University

September 26, 2022