

# Ayush Sekhari

Web: [ayush.sekhari.com](http://ayush.sekhari.com)

Email: [sekhari@mit.edu](mailto:sekhari@mit.edu)

Phone: +1 607 (379) 4346

## RESEARCH INTERESTS

Sequential Decision Making and Interactive Learning, Theoretical Foundations of Machine Learning and Machine Unlearning, Stochastic Optimization.

## CURRENT POSITION

Massachusetts Institute of Technology, USA

2022 - current

Postdoctoral Associate, Advisor: Prof. Alexander (Sasha) Rakhlin

## EDUCATION

Cornell University, Ithaca, NY, USA

2017 - 22

Ph.D. in Computer Science

Advisors: Prof. Karthik Sridharan (Primary), Prof. Robert Kleinberg (Co-chair)

Thesis: [Non-convex and Interactive Learning via Stochastic Optimization](#)

Indian Institute of Technology (IIT) Kanpur, India

2012 - 16

B.Tech. in Computer Science and Engineering (GPA: 10/10)

Institute Rank 1, **Presidents Gold Medal** for the best academic performance in the graduating batch of 2016.

## SELECTED AWARDS

- **Finalist for Meta AI PhD Fellowship** in Statistics, 2022.
- **Best Student Paper Award** at COLT 2019.
- **Best Talk Award, Honorable Mention**, New York Academy of Science (NYAS) 2020.
- **Presidents Gold Medal**, IIT Kanpur, 2016.

## PUBLICATIONS AND PREPRINTS

### Conference Publications

Following the tradition in theoretical and mathematical sciences, authors are ordered alphabetically by their last name in papers marked with: [\(alphabetical\)](#).

#### [Offline Reinforcement Learning: Role of State Aggregation and Trajectory Data](#)

Zeyu Jia, Alexander Rakhlin, **Ayush Sekhari**, Chen-Yu Wei [\(alphabetical\)](#)

*Conference on Learning Theory 2024 (COLT 2024)*.

#### [Random Latent Exploration for Deep Reinforcement Learning](#)

Srinath Mahankali, Zhang-Wei Hong, **Ayush Sekhari**, Alexander Rakhlin, Pulkit Agrawal

*International Conference on Machine Learning 2024 (ICML 2024)*.

#### [Harnessing Density Ratios for Online Reinforcement Learning](#)

Philip Amortila, Dylan J. Foster, Nan Jiang, **Ayush Sekhari**, Tengyang Xie [\(alphabetical\)](#)

*International Conference on Learning Representations 2024 (ICLR 2024)*. **Spotlight**

#### [Offline Data Enhanced On-Policy Policy Gradient with Provable Guarantees](#)

Yifei Zhou\*, **Ayush Sekhari**\* (equal contribution), Yuda Song and Wen Sun

*International Conference on Learning Representations 2024 (ICLR 2024)*.

#### [When is Agnostic Reinforcement Learning Statistically Tractable?](#)

Zeyu Jia, Gene Li, **Ayush Sekhari**, Nati Srebro and Alexander Rakhlin [\(alphabetical\)](#)

*Neural Information Processing Systems 2023 (NeurIPS 2023).*

### **Contextual Bandits and Imitation Learning via Preference-Based Active Queries**

**Ayush Sekhari**, Karthik Sridharan, Wen Sun, and Runzhe Wu (alphabetical)  
*Neural Information Processing Systems 2023 (NeurIPS 2023).*

### **Selective Sampling and Imitation Learning via Online Regression**

**Ayush Sekhari**, Karthik Sridharan, Wen Sun, and Runzhe Wu (alphabetical)  
*Neural Information Processing Systems 2023 (NeurIPS 2023).*

### **Hidden Poison: Machine Unlearning Enables Camouflaged Poisoning Attacks**

Jimmy Z. Di, Jack Douglas, Jayadev Acharya, Gautam Kamath, **Ayush Sekhari**  
*Neural Information Processing Systems 2023 (NeurIPS 2023).*

Short version at *ML Safety Workshop, and at Trustworthy and Socially Responsible Machine Learning (TSRML) at NeurIPS 2022.*

### **Model-Free Reinforcement Learning with the Decision-Estimation Coefficient**

Dylan J. Foster, Noah Golowich, Jian Qian, Alexander Rakhlin, **Ayush Sekhari** (alphabetical)  
*Neural Information Processing Systems 2023 (NeurIPS 2023).*

### **Ticketed Learning-Unlearning Schemes**

Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, **Ayush Sekhari**, Chiyuan Zhang (alphabetical)  
*Conference on Learning Theory 2023 (COLT 2023).*

### **Computationally Efficient PAC RL in POMDPs with Latent Determinism and Conditional Embeddings**

Masatoshi Uehara, **Ayush Sekhari**, Jason D. Lee, Nathan Kallus, Wen Sun  
*International Conference on Machine Learning (ICML 2023).*

### **Hybrid RL: Using Both Offline and Online Data Can Make RL Efficient**

Yuda Song\*, Yifei Zhou\* (equal contribution), **Ayush Sekhari**, Andrew Bagnell, Akshay Krishnamurthy and Wen Sun  
*International Conference on Learning Representations 2023 (ICLR 2023).*  
Short version at *Offline RL Workshop at NeurIPS 2022.*

### **On the Complexity of Adversarial Decision Making**

Dylan J. Foster, Alexander Rakhlin, **Ayush Sekhari**, Karthik Sridharan (alphabetical)  
*Neural Information Processing Systems 2022 (NeurIPS 2022).* **Oral Presentation.**

### **From Gradient Flow on Population Loss to Learning with Stochastic Gradient Descent**

with Satyen Kale, Jason D. Lee, Chris De Sa, **Ayush Sekhari**, Karthik Sridharan. (alphabetical)  
*Neural Information Processing Systems 2022 (NeurIPS 2022).*

### **Provably Efficient Reinforcement Learning in Partially Observable Dynamical Systems**

Masatoshi Uehara, **Ayush Sekhari**, Jason D. Lee, Nathan Kallus, Wen Sun.  
*Neural Information Processing Systems 2022 (NeurIPS 2022).*

### **Guarantees for Epsilon-Greedy Reinforcement Learning with Function Approximation**

Christoph Dann, Yishay Mansour, Mehryar Mohri, **Ayush Sekhari**, Karthik Sridharan. (alphabetical)  
*International Conference on Machine Learning 2022 (ICML 2023).*  
Short version at *Reinforcement Learning and Decision Making (RLDM) 2022.*

### **SGD: The role of Implicit Regularization, Batch-size and Multiple Epochs**

Satyen Kale, **Ayush Sekhari**, Karthik Sridharan. (alphabetical)

*Neural Information Processing Systems 2021 (NeurIPS 2021).*

### **Agnostic Reinforcement Learning with Low-Rank MDPs and Rich Observations**

Christoph Dann, Yishay Mansour, Mehryar Mohri, **Ayush Sekhari**, Karthik Sridharan. (alphabetical)

*Neural Information Processing Systems 2021 (NeurIPS 2021).* **Spotlight**

### **Remember What You Want to Forget: Algorithms for Machine Unlearning**

**Ayush Sekhari**, Jayadev Acharya, Gautam Kamath, Ananda Theertha Suresh.

*Neural Information Processing Systems 2021 (NeurIPS 2021).*

Short version at *Theory and Practice of Differential Privacy (TPDP) 2021.*

### **Neural Active Learning with Performance Guarantees**

Pranjal Awasthi, Christoph Dann, Claudio Gentile, **Ayush Sekhari**, Zhilei Wang.

(alphabetical)

*Neural Information Processing Systems 2021 (NeurIPS 2021).*

### **Reinforcement Learning with Feedback Graphs**

Christoph Dann, Yishay Mansour, Mehryar Mohri, **Ayush Sekhari**, Karthik Sridharan. (alphabetical)

*Neural Information Processing Systems 2020 (NeurIPS 2020).*

Short version at *ICML 2020 Theoretical foundations of RL workshop.*

### **Second-Order Information in Non-Convex Stochastic Optimization: Power and Limitations.**

Yossi Arjevani, Yair Carmon, Dylan Foster, **Ayush Sekhari**, Karthik Sridharan. (alphabetical)

*Conference on Learning Theory 2020 (COLT 2020).*

### **The Complexity of Making the Gradient Small in Stochastic Convex Optimization.**

Dylan J. Foster, **Ayush Sekhari**, Ohad Shamir, Nathan Srebro, Karthik Sridharan, and Blake Woodworth. (alphabetical)

*Conference on Learning Theory 2019 (COLT 2019).* **Best Student Paper Award.**

### **Uniform Convergence of Gradients for Non-Convex Learning and Optimization.**

Dylan Foster, **Ayush Sekhari**, and Karthik Sridharan. (alphabetical)

*Neural Information Processing Systems 2018 (NeurIPS 2018).*

Short version at *ICML 2018 Non-convex Optimization workshop.*

## **Workshop Publications**

### **A Brief Study of In-Domain Transfer and Learning from Fewer Samples using a Few Simple Priors.**

Marc Pickett, **Ayush Sekhari**, and James Davidson.

*ICML 2017 workshop: Picky Learners - Choosing Alternative Ways to Process Data.*

Awarded the **Best Paper Award, Honorable Mention** among the workshop submissions.

## **Papers Currently Under Submission**

The papers listed below are currently undergoing review at various conferences (as of October 2024):

### **On the Unlearnability of the Learnable**

Yeshwanth Cherapanamjeri, Sumegha Garg, Nived Rajaraman, **Ayush Sekhari**, Abhishek Shetty (alphabetical).

### **System Aware Unlearning Algorithms: Use Lesser, Forget Faster**

Linda Lu, Ayush Sekhari, Karthik Sridharan (alphabetical).

### **Unstable Unlearning: The Hidden Risk of Concept Resurgence in Diffusion Models**

Vinith Menon Suriyakumar<sup>†</sup>, Rohan Alur<sup>†</sup> (equal contribution), **Ayush Sekhari**, Manish Raghavan, Ashia C. Wilson.

### Langevin Dynamics: A Unified Perspective on Optimization via Lyapunov Potentials

August Y. Chen, **Ayush Sekhari**, Karthik Sridharan (alphabetical).

### Machine Unlearning Fails to Remove Data Poisoning Attacks

Martin Pawelczyk<sup>†</sup>, Jimmy Z. Di<sup>†</sup>, Yiwei Lu, Gautam Kamath\*, **Ayush Sekhari\***, Seth Neel\* (†-equal contribution, \*-equal advising).

Preliminary version at *2nd Workshop on Generative AI and Law at International Conference of Machine Learning (ICML) 2024*. **Oral Presentation**.

### Computationally Efficient RL under Linear Bellman Completeness for Deterministic Dynamics

Runzhe Wu, **Ayush Sekhari**, Akshay Krishnamurthy, Wen Sun (alphabetical).

## INVITED TALKS

### Offline Reinforcement Learning: Role of State Aggregation and Trajectory Data

COLT, Edmonton, CA

Jun 2024

Adaptive Learning in Complex Environments Workshop, TTIC, Chicago, USA

Apr 2024

### Offline Data Enhanced On-Policy Policy Gradient

Virtual RL theory Seminar Series

Apr 2024

CSA theory seminar, IISc Bangalore, India

Apr 2024

### Ticketed Learning-Unlearning Schemes

Max Planck Institute for Intelligent Systems, Tubingen, Germany

May 2024

Conference on Learning Theory (COLT), Bangalore, India

Jul 2023

CS Theory Seminar, Cornell University, Ithaca, USA

Nov 2023

Annual Conference on Information Sciences and Systems (CISS), Princeton, USA

Mar 2024

CS Theory Seminar, University of Pennsylvania, Philadelphia, USA

Mar 2024

CSA theory seminar, IISc Bangalore, India

Apr 2024

### Machine Unlearning: Algorithms, complexity, and new challenges

Meta AI Research, USA

Apr 2023

### Hybrid RL: Using Both Offline and Online Data Can Make RL Efficient

ImprobableAI (Prof. Pulkit Agarwal's lab) meeting, MIT

Mar 2023

### On the Complexity of Adversarial Decision Making

Virtual RL Theory Seminar Series

Jul 2023

BLISS seminar, UC Berkeley, USA

Feb 2023

Information Theory and Applications (ITA) Workshop, San Diego, USA

Feb 2023

Theory seminar, UCSD, San Diego, USA

Feb 2023

Microsoft research NYC, USA

Feb 2023

ML Tea, Massachusetts Institute of Technology, USA

Apr 2023

### When does SGD learn?

Prof. Dan Roy's lab, University of Toronto, CA

Oct 2022

Mathematical Foundations of deep learning reading group, ETH Zurich

Nov 2022

### Remember What You Want to Forget: Algorithms for Machine Unlearning

AI Seminar, Cornell University

Feb 2022

Prof. Jiantao Jiao's lab, UC Berkeley	Aug 2021
<b>SGD: The role of implicit regularization, batch-size and multiple-epochs</b>	
Mathematical foundations of deep learning reading group, ELLIS, ETH Zurich	May 2022
Foundations of Data Science (FODS) seminar, IISC (Tsinghua University), China	Apr 2022
Collaboration on the theoretical foundations of deep learning (MODL) monthly meeting	Feb 2022
Theory Seminar, Cornell University	May 2021
Algorithms and Theory Seminar, University of Waterloo (CA)	Nov 2021
Learning Theory Seminar, Google Research NY	Nov 2021
<b>Agnostic Reinforcement Learning with Low-Rank MDPs and Rich Observations</b>	
Artificial Intelligence (AI) Seminar, Cornell University	Mar 2021
RL reading group, Cornell University	Jun 2021
<b>Second-Order Information in Non-Convex Stochastic Optimization: Power and Limitations</b>	
Highlights beyond SIGMETRICS 2021, Beijing, China (Virtual)	Jun 2021
Spotlight talk, Annual ML Symposium, New York Academy of Sciences (NYAS)	Mar 2020
<b>Best Talk Award, Honorable Mention</b>	
Conference on Learning Theory (COLT), Conference talk	Jul 2020
Learning Theory Seminar, Google NYC	Nov 2020
Theory Tea, Cornell University	Nov 2020
<b>The Complexity of Making the Gradient Small in Stochastic Convex Optimization</b>	
Intern Talk Series, Google Research, New York	Jul 2019
Theory Seminar, Cornell University	Nov 2019
<b>Uniform Convergence of Gradients for Non-Convex Learning and Optimization</b>	
ICML workshop on Modern Trends in Non-convex Optimization for ML	Jun 2018
Annual ML Symposium, New York Academy of Sciences (NYAS)	Feb 2019

## PROFESSIONAL EXPERIENCE

### Research Internships During Ph.D.

<b>Google Research, Mountain View, USA</b>	Summer 2022
<i>Mentors: Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, and Chiyuan Zhang</i>	
Summer internship focusing on machine unlearning.	
<b>University of Alberta, Edmonton, Canada</b>	Summer 2021
<i>Mentors: Prof. Csaba Szepesvári</i>	
Summer internship focusing on Reinforcement Learning (RL) with partial observability.	
<b>Google Research, New York City, USA</b>	Spring 2020 - 21
<i>Mentors: Prof. Mehryar Mohri</i>	
Student researcher with Learning Theory team working on RL and non-convex optimization.	
<b>Google Research, New York City, USA</b>	Summer 2019
<i>Mentors: Prof. Mehryar Mohri, Chris Dann, Claudio Gentile</i>	
Summer internship with the learning theory group focusing on active learning and RL.	
<b>Toyota Technological Institute, Chicago, USA</b>	Summer 2018
<i>Mentors: Prof. Natan Srebro, Srinadh Bhojanapalli</i>	
Summer internship focusing on non-convex optimization and theory for deep learning.	

### Other Work Experience

<b>Goldman Sachs, Hong Kong</b> <i>Mentors: Elie Franko, Dunstan Marris</i>	Summer 2017
Desk-Strat for Fixed Income, Currency and Commodities (FICC) macro trading group. Developed tools and techniques for analyzing mean reverting stochastic processes to predict the evolution of interest rate swaps.	
<b>AI Residency, Google Brain, Mountain View, USA</b> <i>Mentors: James Davidson, Vikas Sindhwani</i>	2016 - 17
Worked on adaptive reinforcement learning (RL), simulation to real transfer and autonomous navigation with Brain-robotics research group.	
<b>Undergraduate Internship, CBL, University of Cambridge, UK</b> <i>Mentors: James Lloyd, Prof. Zoubin Ghahramani</i>	Summer 2015
Worked on shape constrained gaussian processes and their applications to auto-ML at Computational and Biological Learning lab (CBL), University of Cambridge.	
<b>Summer Undergraduate Research Grant for Excellence, IIT Kanpur, India</b> <i>Mentors: Prof. Amey Karkare (IIT-K), Sumit Gulwani (MSR Redmond)</i>	Summer 2014
Worked on automatic program synthesis and translation of natural language to first order logic (FOL) as a part of summer research grant for excellence (SURGE) program.	
<b>Tata Institute of Fundamental Research (TIFR), India</b> <i>Mentors: Prof. Rajeve Bhalerao</i>	Summer 2013
Worked on theoretical modeling of relativistic and dissipative fluid dynamics as a NIUS (National Initiative for Undergraduate Sciences) scholar at the Department of Theoretical Physics at Tata Institute of Fundamental Research (TIFR), India.	

## RESEARCH ADVISING AND MENTORING

I find great pleasure in mentoring. Below are the students whom I have closely advised through research projects, which have resulted in published papers and arxiv-preprints (currently under review):

### PhD Students

- Zeyu Jia, MIT
- Gene Li, TTIC → Two Sigma
- August Y. Chen, Cornell University
- Linda Lu, Cornell University
- Runzhe Wu, Cornell University
- Nived Rajaraman, UC Berkeley
- Yiwei Lu, University of Waterloo

### Masters / Undergraduate Students

- Srinath Mahankali, MIT, Class of 2025
- Jimmy Z. Di, University of Waterloo, Class of 2024
- Jack Douglas, University of Waterloo, Class of 2025
- Yifei Zhou, Cornell → PhD at UC Berkeley, Class of 2023

## TEACHING AND LEARNING

I have worked as a Teaching Assistant (TA) for the following courses:

- |  |             |
|--|-------------|
| • <b>CS6783: Machine Learning Theory, Cornell University</b><br>TA for Prof. Karthik Sridharan (Graduate level course) | Fall 2018   |
| • <b>CS4820: Introduction to Analysis of Algorithms, Cornell University</b><br>Head TA for Prof. Robert Kleinberg      | Spring 2018 |

- **CS4786/5786: Machine Learning for Data Science, Cornell University** Fall 2017  
Head TA for Prof. Karthik Sridharan
- **ESC101: Fundamentals of Computing, IIT Kanpur, India** Fall 2015  
TA for Prof. Nitin Saxena

I believe that focused reading groups and seminars are essential for forming and nurturing research communities. Towards that, have initiated and organized the following:

- **Learning Theory Seminar, Google Research NYC** 2019-2022  
Co-organized with Prof. Mehryar Mohri and Chris Dann.
- **Learning Theory Seminar, Cornell University** Fall 2019  
Co-organized under Prof. Karthik Sridharan and Prof. Nika Haghtalab.
- **Concentration Inequalities Seminar, Cornell University** Spring 2019  
Organized with Prof. Karthik Sridharan.
- **Theory Tea, Cornell University** Fall 2018  
Weekly gathering of PhD students to present and discuss recent trends in theoretical computer science and machine learning.
- **Algorithmic Game Theory Reading Group, Cornell University** Fall 2018  
Graduate student reading group organized under Prof. Eva Tardos.

#### ADDITIONAL HONORS AND AWARDS

- **General Proficiency Medal**, Computer Science and Engineering, IIT Kanpur, 2016.
- **V. Rajaraman Scholarship**, Department of Computer Science and Engineering, IIT Kanpur, 2016.
- **Cambridge-Tubingen PhD Fellowship** 2016 (Declined).
- **Academic Excellence Award**, IIT Kanpur, 2013, 2014, 2015.
- **Shmt. Dharam Vati Garg Donor Scholarship**, IIT Kanpur, 2015.
- **Summer Undergraduate Research Grant for Excellence (SURGE)**, IIT Kanpur, 2014.
- **National Initiative for Undergraduate Sciences (NIUS) Scholar** for fundamental research in Physics, Government of India, 2013.
- **Din Dayal's Gold Medal** for excellence in Mathematics, Delhi Public School, Faridabad, India, 2012.
- **Kishore Vaigyanic Protsahan Yojana (KVPY) scholarship**, Department of Science and Technology, Government of India, 2011.
- **All India Rank 11** (amongst 1,140,000 students) in All India Engineering Entrance Examination, 2012.

#### PROFESSIONAL SERVICE

##### Conference / Workshop Organization

- **Updatable Machine Learning Workshop (UpML)**, July 2022.

Organized a day long research workshop at International Conference of machine learning (ICML) 2022, Baltimore, USA, on addressing post-deployment issues in ML. Co-organizers: Prof. Jayadev Acharya and Prof. Gautam Kamath.

Workshop website <https://upml2022.github.io/>.

##### Area Chair / PC Member

Algorithmic Learning Theory, 2024.

Conference on Learning Theory, 2021-24.

## Reviewing

### **Conferences:**

- Algorithmic Learning Theory (ALT), 2021-23.
- Neural Information Processing Systems (NeurIPS), 2019-24.
- International Conference on Machine Learning (ICML), 2019-21.
- International Conference on Learning Representation, 2019, 2023.
- Foundations of Responsible Computing (FORC), 2021.
- Artificial Intelligence and Statistics Conference (AISTATS), 2019, 2023.
- Innovations in Theoretical Computer Science (ITCS), 2020.
- International Symposium on Information Theory (ISIT), 2020.

### **Journals:**

- Journal of Complexity, 2021.
- Journal of Machine Learning Research, 2021-22.

### **Workshops:**

- Workshop on Understanding and Improving Generalization in Deep Learning at ICML 2019.
- Updatable Machine Learning Workshop at ICML 2022.

## Diversity, Equity, and Inclusion Efforts

- **Part of Communications Committee at Learning Theory Alliance**, Spring 2024-Current.  
Website - <https://let-all.com/>.
- **Learning Theory Alliance**, Fall 2023.  
Volunteered for organization of mentoring tables for Fall 2023 Mentorship Workshop focused on communicating one's research verbally.
- **Mentor for Project SHORT**, Fall 2021-current.  
A student led organization working to shrink the socioeconomic gap in graduate school applications.
- **Girls Adventures in Math (GAIM), Cornell University**, 2018, 2019.  
Volunteered to help organize an in-person Math Olympiad for upper elementary and middle school girls in NY state. GAIM has moved fully online post-COVID, thus making it accessible to teams all over USA.
- **Institute Counselling Service, IIT Kanpur**, 2013-15.  
Worked as an Academic Mentor (2013), Student Guide (2013), Senior Academic Mentor (2014), and Link Student (2014) to advise and mentor peer undergraduate students.

## Miscellaneous

- **Reviewer for Ph.D. Admissions**, CS Department, Cornell University, Fall 2018.
- **TA training for CIS undergraduates**, Cornell University, Fall 2020.