Sourav Chakraborty

Ph.D. Candidate in Computer Science | University of Colorado Boulder

souravchk25.github.io | mail.souravchk@gmail.com | +1-720-980-5863
Google Scholar | linkedin.com/in/souravchk | Boulder, Colorado, USA

Research Summary

I work on the mathematical foundations of **sequential decision-making under uncertainty**: how agents should learn and act when feedback is noisy, biased, delayed, or non-stationary, and how learning can **scale** to large networks of agents with no central controller. My research spans **multi-agent reinforcement learning at scale**, **multi-armed bandits**, and **online learning**, where I pose new problem settings and prove regret and scaling guarantees. I am now bringing these tools to **foundation models**: the geometry of efficient attention, and **post-training of large language models** with **reinforcement learning** for reasoning and alignment. Before graduate school I spent two years as a software engineer at **Flipkart**, working on search and ranking.

Education

Ongoing | **University of Colorado Boulder** — **Boulder, Colorado, USA**
Ph.D. in Computer Science | *Advisor: Prof. Lijun Chen*
2022 | **M.S.** in Computer Science | *GPA: 4.0/4.0*
Thesis: *Incentivized Exploration in Non-stationary Stochastic Bandits* [Thesis | Paper]
Lloyd Botway Award, the department's highest honor for a graduating M.S. student

2016 | **Birla Institute of Technology, Mesra** — **Ranchi, Jharkhand, India**
B.E. in Information Technology, Department of Computer Science

Research & Industry Experience

Present – May 2026 | **California Institute of Technology (Caltech)** — **Pasadena, California, USA**
Visiting Student Researcher, CMS Department | *Hosts: Prof. Steven Low, Prof. Adam Wierman*

- **Multi-agent RL over large networked systems.** Develop decentralized learning methods whose cost stays local as the number of agents grows, grounded in our spectral theory of locality.
- **Communication delay at scale.** Characterized how delayed information between agents affects learning in networked multi-agent RL, with Prof. Adam Wierman (paper under review).
- Invited talks on *Flickering Multi-Armed Bandits* and *locality in scalable MARL* to **RSRG** and **FALCON**, the hosts' research groups.

Present – Aug 2022 | **University of Colorado Boulder** — **Boulder, Colorado, USA**
Graduate Research Assistant, Dept. of Computer Science | *Advisor: Prof. Lijun Chen*

- **Scalable multi-agent RL.** Showed that locality in decentralized MARL is policy-dependent and gave a spectral condition under which agents plan with nearby agents only, with error decaying exponentially in the communication radius, avoiding the curse of dimensionality at scale (L4DC '26).
- **Bandits on dynamic graphs.** Introduced *Flickering Multi-Armed Bandits*, where an evolving random graph restricts which arms the learner can reach, and proved near-optimal sublinear regret for Erdős–Rényi and edge-Markovian graphs, together with a lower bound (Best Paper Runner-Up and Oral, L4DC '26). Related work: learnability under local movement (NeurIPS '25 ARLET) and local explore-then-commit under sliding-window mixing (IEEE CDC '26).
- **Continuous, multi-agent, and incentivized bandits.** Designed algorithms that pay myopic agents to explore despite payment-biased feedback, for non-stationary (IEEE ACC '24) and continuous Lipschitz actions (IEEE CDC '25); a communication-free multi-player protocol over continuous actions with near-optimal regret (AISTATS '26); and its extension to collisions (under review).
- **Foundation models and quantum bandits.** Contributed the mathematical analysis to *kernelized linear attention* over symmetric cones, which lifts linear attention's capacity limit; co-authored two papers on quantum bandit regret (all under review).
- Lead projects from problem formulation through proofs, writing, and rebuttals with my co-authors.

2019 – Dec 2016 | **Flipkart** — Software Engineer, Search & Ranking — **Bangalore, Karnataka, India**

- Built and shipped the end-to-end **query-suggestion pipeline** behind related search, improving product discovery while users type, for a search system serving 50M+ daily sessions (distributed batch jobs on Hadoop; Java, Cascading).
- Raised **retrieval quality and result ordering** with new ranking signals (recency, semantic similarity), validated by offline benchmarks and online **A/B tests** on a 10M-query set.
- Designed a pluggable MySQL backup platform that grew into **BRaaS**, Flipkart's central backup-and-recovery service, adopted by 20+ microservices and removing single points of failure.

Research Papers & Publications (* denotes equal contribution)

[12] **Kernelized Linear Attention: Breaking the Capacity Wall with Symmetric Cones** [arXiv]
Ayoub Ghriss, Sourav Chakraborty
Under review **[Preprint '26]**

[11] **Characterizing the Impact of Communication Delay in Networked Multi-Agent Reinforcement Learning**
Sourav Chakraborty*, Amit Kiran Rege*, Adam Wierman, Lijun Chen
Under review **[Preprint '26]**

[10] **Decentralized Lipschitz Bandits with Proximity Collisions**
Sourav Chakraborty*, Amit Kiran Rege*, Lijun Chen
Under review **[Preprint '26]**

[9] **Non-Stationary Quantum Bandits under a Variation Budget**
Jiarui Hu, Sourav Chakraborty, Lijun Chen
Under review **[Preprint '26]**

[8] **On Regret Bounds of Quantum Bandit Algorithms**
Jiarui Hu, Wenzhi Chen, Sourav Chakraborty, Lijun Chen
Under review **[Preprint '26]**

[7] **Learning from Local Walks on Dynamic Graphs with Bandit Feedback** [arXiv]
Sourav Chakraborty*, Amit Kiran Rege*, Claire Monteleoni, Lijun Chen
65th IEEE Conference on Decision and Control, Honolulu **[IEEE CDC '26]**

[6] **[Best Paper Runner-Up]** & (Oral) **Flickering Multi-Armed Bandits** [arXiv]
Sourav Chakraborty*, Amit Kiran Rege*, Claire Monteleoni, Lijun Chen
8th Annual Learning for Dynamics and Control Conference, Los Angeles, USA. **[L4DC '26]**

[5] **A Unified Framework for Locality in Scalable MARL** [arXiv]
Sourav Chakraborty*, Amit Kiran Rege*, Claire Monteleoni, Lijun Chen
8th Annual Learning for Dynamics and Control Conference, Los Angeles, USA **[L4DC '26]**

[4] **Multi-Agent Lipschitz Bandits** [arXiv]
Sourav Chakraborty*, Amit Kiran Rege*, Claire Monteleoni, Lijun Chen
29th International Conference on Artificial Intelligence and Statistics, Tangier, Morocco. **[AISTATS '26]**

[3] **Bandit Learning on Dynamic Graphs** [OpenReview]
Sourav Chakraborty*, Amit Kiran Rege*, Claire Monteleoni, Lijun Chen
NeurIPS 2025 Workshop on Aligning Reinforcement Learning Experimentalists and Theorists, San Diego **[NeurIPS '25]**

[2] **Incentivized Lipschitz Bandits** [IEEE | arXiv]
Sourav Chakraborty*, Amit Kiran Rege*, Claire Monteleoni, Lijun Chen
64th IEEE Conference on Decision and Control, Rio de Janeiro **[IEEE CDC '25]**

[1] **Incentivized Exploration in Non-stationary Stochastic Bandits** [IEEE | arXiv]
Sourav Chakraborty, Lijun Chen
2024 American Control Conference, Toronto **[IEEE ACC '24]**

Awards, Honors & Fellowships

Ralph J. Slutz Student Excellence Award Dept. of CS, for academic excellence, research, and innovation 2026
Best Paper Runner-Up, L4DC 2026 **1 of 2** recipients, for *Flickering Multi-Armed Bandits* 2026
Outstanding Research Paper Award **(2×)** Dept. of Comp. Sci., for the L4DC '26 and IEEE ACC '24 papers 2024, 2026
Summer Research Fellowship Dept. of Computer Science, funding full-time summer research 2026
Graduate Research Assistant Fellowship funding full-time research for Spring 2026 2025
Conference Travel Awards **(3×)** IEEE CDC '26, IEEE CDC '25, and IEEE ACC '24 (full fellowship) 2024–26
Outstanding Teaching Assistant Award Dept. of Computer Science, for teaching and TA leadership 2025
Outstanding Service Award Dept. of Computer Science, for Lead TA and graduate committee service 2025
Publication Recognition Award for peer-reviewed publication in AY 2023–24 (IEEE ACC '24) 2024
CU Research Expo Poster Award **(2×)** University of Colorado Boulder's annual research showcase 2022, 2024
Early Career Development Fellowship Dept. of Computer Science, at the start of the Ph.D. 2022
Lloyd Botway Award the department's **highest honor** for a graduating M.S. student 2022

Talks & Presentations

Caltech CMS Talks to the RSRG and FALCON research groups: *Flickering MAB*; *Locality in Scalable MARL* 2026
L4DC 2026, Los Angeles Oral Talk: *Flickering Multi-Armed Bandits* Jun 2026
IEEE CDC 2025, Rio de Janeiro Paper talk: *Incentivized Lipschitz Bandits* Dec 2025
IEEE ACC 2024, Toronto Paper talk: *Incentivized Exploration in Non-stationary Stochastic Bandits* Jul 2024
University of Colorado Boulder Guest lectures: *Bandit Models* (Advanced ML); *Graphs* (Data Structures) 2023

Academic Service

Conference Reviewer ICML, NeurIPS, AAAI, AISTATS, IEEE CDC, IEEE ACC
Journal Reviewer IEEE Transactions on Control of Network Systems (TCNS)
Student Representative Elected to the CS Graduate Committee, University of Colorado Boulder (2023)

Teaching & Leadership

Department Lead Teaching Assistant, Dept. of Computer Science, University of Colorado Boulder 2022 – 2025
- Selected through a teaching review and interviews; reappointed for **three** consecutive years.
- Ran TA hiring interviews and annual onboarding for 280+ instructional staff; organized training workshops and mentored new graduate TAs.
- Acted as the TAs' liaison with department leadership.

Instructor of Record, University of Colorado Boulder 2020 – 2021
- *Introduction to Computational Thinking* (CSCI 1200), Fall 2021: about 200 students.
- *Introduction to Data Science with Probability & Statistics* (CSCI 3022), Summer 2020: about 50 students.

Graduate Teaching Assistant, University of Colorado Boulder 2020 – Present
- *Algorithms*, *Data Structures*, *Starting Computing*; **10+** semesters. **Outstanding TA Award** (2025).

Technical Skills

Methods Reinforcement learning, multi-agent RL, bandits, online learning, stochastic optimization
Programming Python, C++, Java, SQL, Shell, R, Julia | PyTorch, NumPy, Pandas, scikit-learn
Systems Docker, Kubernetes, (Py)Spark, Hadoop (Cascading), MySQL, Git, Linux