

SUBHOJYOTI MUKHERJEE

Almaden Tower,
Adobe Systems,
San Jose, CA 95113

Email: subhomuk@adobe.com,
subhojyotimukherjee22@gmail.com
Website: <https://subhojyoti.github.io/>

Research Interests Reinforcement Learning ([RL](#)), Large Language Models ([LLM](#)), Reinforcement Learning with Human Feedback ([RLHF](#)), Incontext Learning ([ICL](#)), Optimal Design ([OD](#)).

Work Experience **Research Scientist, Adobe Research**, San Jose, USA 2025, March – present

Products

1. Led pre-training of a small on-device English language model for Document Overview in Acrobat Reader (launched Sept. 2025); patent filed.
2. Led RL post-training of a small on-device multilingual language model for Spanish Document Overview in Acrobat Reader (planned 2026).
3. Contributed to synthetic data generation and led post-training for the Adobe Express Agent Orchestrator (ongoing); patent filed.
4. Contributed to RL post-training for the Firefly image editing model (ongoing).

Education **University of Wisconsin-Madison**, Madison, USA Fall 2019 – 2025, Feb
Ph.D., Electrical & Computer Engineering
Adviser: Dr. Robert Nowak, Dr. Josiah Hanna and Dr. Qiaomin Xie

University of Wisconsin-Madison, Madison, USA Fall 2019 – 2021
M.S., Electrical Engineering
Adviser: Dr. Robert Nowak

Indian Institute of Technology Madras, India 2015–2018
M.S (Research), Computer Science
Advisers: Dr. Balaraman Ravindran and Dr. Nandan Sudarsanam

West Bengal University of Technology, Kolkata, India 2009–2013
Bachelor of Technology, Computer Science & Engineering

Publications

1. **Subhojyoti Mukherjee**, Viet Dac Lai, et.al. "*Offline RL by Reward-Weighted Fine-Tuning for Conversation Optimization*". **NeurIPS 2025, main conference** [Paper] ([LLM](#), [RL](#))
2. **Subhojyoti Mukherjee***, Yu Xia*, et.al. "*From Selection to Generation: A Survey of LLM-based Active Learning*". **ACL 2025, main conference** [Paper] ([LLM](#))
3. Gaurush Hiranandani, Haolun Wu, **Subhojyoti Mukherjee**, Sanmi Koyejo "*Log-its are All We Need to Adapt Closed Models*". **ICML 2025, main conference** [Paper] ([LLM](#), [RL](#))
4. **Subhojyoti Mukherjee**, Josiah Hanna, Qiaomin Xie, Robert Nowak, "*Pretraining Decision Transformers with Reward Prediction for In-Context Multi-task Structured Bandit Learning*". **RLC 2025** [Paper] ([LLM](#), [ICL](#))

5. **Subhojyoti Mukherjee**, Qiaomin Xie, Robert Nowak, "*Multi-task Representation Learning for Fixed Budget Pure-Exploration in Linear and Bilinear Bandits*". **RLC 2025** [Paper] ([RL](#))
6. **Subhojyoti Mukherjee**, Anusha Lalitha, Kousha Kalantari, Aniket Anand Deshmukh, Ge Liu, Yifei Ma, Branislav Kveton, "*Optimal Design for Human Preference Elicitation*". (**NeurIPS 2024, main conference**) [Paper] ([LLM](#), [RLHF](#))
7. **Subhojyoti Mukherjee**, Josiah Hanna, Robert Nowak, "*SaVeR: Optimal Data Collection Strategy for Safe Policy Evaluation in Tabular MDP*". (**ICML 2024, main conference**) [Paper] ([RL](#))
8. **Subhojyoti Mukherjee**, Anusha Lalitha, Kousha Kalantari, Aniket Anand Deshmukh, Ge Liu, Yifei Ma, Branislav Kveton, "*Optimal Design for K-Way Human Feedback*". (**Models of Human Feedback for AI Alignment workshop ICML 2024**) [Paper] ([LLM](#), [RLHF](#), [OD](#))
9. Aniruddha Bhargava, Lalit Jain, Branislav Kveton, Ge Liu, **Subhojyoti Mukherjee**, "*Off-Policy Evaluation from Logged Human Feedback*". (**Models of Human Feedback for AI Alignment workshop ICML 2024**) [Paper] ([LLM](#))
10. **Subhojyoti Mukherjee**, Qiaomin Xie, Josiah Hanna, Robert Nowak, "*SPEED: Optimal Experimental Design for Policy Evaluation in Linear Heteroscedastic Bandits*". (**AISTATS 2024**) [Paper] ([RL](#), [OD](#))
11. **Subhojyoti Mukherjee**, Qiaomin Xie, Josiah Hanna, Robert Nowak, "*Multi-task Representation Learning for Pure Exploration in Bilinear Bandits*", Neural Information Processing Systems. (**NeurIPS 2023**) [Paper] ([RL](#), [OD](#))
12. **Subhojyoti Mukherjee**, Josiah Hanna, Robert Nowak, "*ReVar: Strengthening Policy Evaluation via Reduced Variance Sampling*". Uncertainty in Artificial Intelligence. (**UAI-22**) [Paper] ([RL](#))
13. **Subhojyoti Mukherjee**, "*Safety Aware Changepoint Detection for Piecewise i.i.d. Bandits*". Uncertainty in Artificial Intelligence (**UAI-22**). [Paper] ([RL](#))
14. **Subhojyoti Mukherjee***, Ardhendu Tripathy*, Robert Nowak, "*Chernoff Sampling for Active Testing and Extension to Active Regression*". The 25th International Conference on Artificial Intelligence and Statistics (**AISTATS-22**). [Paper] ([RL](#), [OD](#))
15. Blake Mason, Romain Camilleri, **Subhojyoti Mukherjee**, Kevin Jamieson, Robert Nowak, Lalit Jain, "*Nearly Optimal Algorithms for Level Set Estimation*". The 25th International Conference on Artificial Intelligence and Statistics (**AISTATS-22**). [Paper] ([RL](#), [OD](#))
16. Samarth Gupta, Shreyas Chaudhari, **Subhojyoti Mukherjee**, Gauri Joshi, Osman Yagan, "*A Unified Approach to Translate Classical Bandit Algorithms to the Structured Bandit Setting*", *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP-21)*. [Paper] ([RL](#))
17. Samarth Gupta, Shreyas Chaudhari, **Subhojyoti Mukherjee**, Gauri Joshi, Osman Yagan, "*A Unified Approach to Translate Classical Bandit Algorithms to the Structured Bandit Setting*", *IEEE Journal on Selected Areas in Information Theory (2020)*. [Paper] ([RL](#))
18. **Subhojyoti Mukherjee**, and Odalric-Ambrym-Maillard, "*Distribution-dependent and Time-uniform Bounds for Piecewise i.i.d Bandits*", *Thirty-sixth International Conference on Machine Learning (ICML-19)*, Workshop on Reinforcement Learning for Real Life 2019 track [Poster]. [Paper] ([RL](#))

19. **Subhojyoti Mukherjee**, K.P. Naveen, Nandan Sudarsanam, and Balaraman Ravindran, “Efficient UCBV: An Almost Optimal Algorithm using Variance Estimates”, *Proceedings of the Thirty-Second Association for the Advancement of Artificial Intelligence (AAAI-18)*, main conference track [Oral]. [Paper] ([RL](#))
20. **Subhojyoti Mukherjee**, K.P. Naveen, Nandan Sudarsanam, and Balaraman Ravindran, “Thresholding Bandits with Augmented UCB”, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, main conference track [Poster]. [Paper] ([RL](#))

Preprints

1. **Subhojyoti Mukherjee**, Stefano Petrangeli, Branislav Kveton, Trung Bui, Franck Dernoncourt, Arko Mukherjee, “Agentic Planning with Reasoning for Image Styling via Offline RL”. (in submission to ICML 2026)
2. Puneet Mathur, Branislav Kveton, **Subhojyoti Mukherjee**, Viet Dac Lai, “Partial Policy Gradients for RL in LLMs”. (in submission to ICML 2026)
3. Daeun Lee, **Subhojyoti Mukherjee**, Branislav Kveton, Ryan A. Rossi, Viet Dac Lai, Seunghyun Yoon, Trung Bui, Franck Dernoncourt, Mohit Bansal, “StreamGaze: Gaze-Guided Temporal Reasoning and Proactive Understanding in Streaming Videos”. (in submission to CVPR 2026)
4. Yash Savani, Branislav Kveton, Yuchen Liu, Yilin Wang, Jing Shi, **Subhojyoti Mukherjee**, Nikos Vlassis, Krishna Kumar Singh, “Stepwise Credit Assignment for GRPO on Flow-Matching Models”. (in submission to CVPR 2026)
5. **Subhojyoti Mukherjee**, Anusha Lalitha, Sailik Sengupta, Aniket Anand Deshmukh, Branislav Kveton, “Multi-Objective Alignment of Large Language Models Through Hypervolume Maximization”. [Paper] ([LLM](#), [RLHF](#))
6. **Subhojyoti Mukherjee**, Ge Liu, Aniket Anand Deshmukh, Anusha Lalitha, Yifei Ma, Branislav Kveton, “Optimal Design for Adaptive In-Context Prompt Tuning in Large Language Models”. [Paper] ([LLM](#), [ICL](#), [OD](#))
7. **Subhojyoti Mukherjee**, Ruihao Zhu, Branislav Kveton, “Efficient and Interpretable Bandit Algorithms”. [Paper] ([RL](#), [OD](#))

Mentorship:

Primary Mentor: Daeun Lee — StreamGaze project (paper submitted to CVPR 2026)

Co-Mentor: Junghyun Lee; Yash Savani; Chima Adiole; Reuben A. Luera; Mohammed Elmmoghany; Bo Ni

Research Internships

1. **Amazon AWS AI, Santa Clara, USA:** Summer 2024 (Full-time), Host: Branislav Kveton, Yifei Ma, Anusha Lalitha, Kousha Kalantiri, Aniket Deshmukh. Working on Alignment for Multi-objective optimization with LLMs.
2. **Amazon AWS AI, Santa Clara, USA:** Fall 2023 (Part-time), Host: Branislav Kveton, Yifei Ma, Anusha Lalitha, Kousha Kalantiri, Ge Liu, Aniket Deshmukh, Anoop Deoras. Working on RLHF with LLMs
3. **Amazon AWS AI, Santa Clara, USA:** Summer 2023 (Full-time), Host: Branislav Kveton, Yifei Ma, Anusha Lalitha, Ge Liu, Aniket Deshmukh, Anoop Deoras. Worked on Active In-Context Learning with LLMs
4. **CMU, ECE Dept., Pittsburgh, USA:** Summer 2019, Host: Gauri Joshi. Worked on Structured Bandits

5. **Adobe Research, San Jose, USA:** Spring 2018. Host: Branislav Kveton. Worked on Item recommendation with Ranking and Bandits
6. **INRIA, SequeL Lab, Lille, France:** Fall 2017, Host: Odalric Maillard. Worked on Non-stationary Bandits

PhD Thesis (ECE, UW-Madison) Adaptive Data Collection for Policy Evaluation, Multi-Task Learning and LLM Alignment [Thesis] ([RL](#), [OD](#), [LLM](#))

Master's Thesis (EE, UW-Madison) Active Sequential Hypothesis Testing with Extension to Active Regression and Multi-armed Bandits [Thesis] ([RL](#), [OD](#))

Master's Thesis (CS, IIT Madras) Finite-time Analysis of Frequentist Strategies for Multi-armed Bandits [Thesis] ([RL](#))

Teaching Experience	Teaching/Research Assistant , UW-Madison <i>Matrix Methods in Machine Learning</i> - Prof. Robert Nowak <i>Mathematical Foundation in Machine Learning</i> - Prof. Robert Nowak	2019–2025
	Teaching Assistant , UMass Amherst <i>Natural Language Processing</i> - Prof. Mohit Iyyer <i>Design of Algorithms</i> - Prof. Daniel Sheldon	2018–2019
	Teaching Assistant , IIT Madras <i>Introduction to Programming</i> - Prof. Raghavendra Rao B. V. <i>Reinforcement Learning</i> (twice) - Prof. Balaraman Ravindran	2015–2018

Reviewer and Service

1. AISTATS, UAI, AAAI, ICML, ICLR, NeurIPS, TMLR, KDD, RLC, CVPR
2. Main Co-ordinator of SILO seminar at UW-Madison

Award Grants and Fellowship

1. Top reviewer award for UAI 2023, Neurips 2023
2. Student Scholarship for AAAI 2018, UAI 2022, Neurips 2023
3. UW-Madison nominee for Apple PhD fellowship and Two-sigma PhD fellowship. Received UW-Madison Chancellor's Opportunity Fellowship 2019-20, UW-Madison ECE Welcome Award of USD 3000.
4. IIT Madras student travel grant of USD 2300, Google travel grant of USD 1700, Microsoft travel grant of USD 1435 (declined).

Other Achievements

Ranked 1150/155190 candidates in Graduate Aptitude Test in Engineering (**GATE**) 2014.
Secured 98.93 percentile in Common Admission Test (**CAT**) 2014 among 196988 candidates.

References Available Upon Request.