Sumantrak Mukherjee

Scholar
Website: sumantrak.de

Email: mukherjeesumantrak@gmail.com
Mobile: +49 152 59931742
LinkedIn: Sumantrak

About

Researcher studying sequential decision-making under uncertainty, with a focus on multi-task bandits, robust information transfer, coordinated exploration, and statistical and computational limits of learning algorithms.

Education

BITS Pilani, India Jun 2019 – May 2023
B.E. in Electronics and Instrumentation
- Bachelor's thesis: Scalable Causal Bandits with Adequate Causal Discovery.
- Selected coursework: Probability and Statistics, Linear Algebra, Calculus, Deep Learning, Control Systems.

Research experience

Junior Researcher, DFKI Oct 2023 – Present
Kaiserslautern, Germany | Advised by Prof. Sebastian Vollmer
- Led the research consortium **AI4Nof1**, conducting research on multi-armed bandits with information sharing, coordinated exploration, and warm-starting, resulting in multiple research outputs.
- Coordinated research assistants and software developers to translate personalization methods into deployed systems, including a real-world study with university students and an IBS intervention app to be released on the App Store.
- Previously researched neural and spatiotemporal point processes in the **Eventful** project, investigating flexible neural modeling and causal inference.

Guest Researcher, RPTU Apr 2023 – Sep 2023
Kaiserslautern, Germany | Supervised by Prof. Sebastian Vollmer
- Researched active learning for relabeling, data correction, and noisy annotators.

Bachelor's Thesis Researcher, DFKI Jan 2023 – May 2023
Supervised by Prof. Sebastian Vollmer and Dr. Mengyan Zhang
- Integrated differentiable Bayesian structure learning with reward maximization and evaluated scalable causal-bandit methods.

Research Collaborator, Columbia University Jan 2022 – Aug 2022
Remote | Supervised by Dr. Zenna Tavares
- Implemented rational-agent models for (PO)MDPs with human biases and bounded rationality; developed JuliaProgrammingPuzzles.jl benchmarks.

Julia Summer of Code May 2021 – Jul 2021
Remote | Mentors: Dr. Sebastian Vollmer, Dr. Jiahao Chen, Dr. Moritz Schauer
- Implemented fairness in-processing algorithms in the MLJ ecosystem; studied adversarial-training stability and pairwise fairness.

Journal and main-conference publications

Neural Spatiotemporal Point Processes: Trends and Challenges
Mukherjee, S., Elhamdi, M., Mohler, G., Selby, D. A., Xie, Y., Vollmer, S., & Großmann, G. *Transactions on Machine Learning Research (TMLR), 2025.*

X-Hacking: The Threat of Misguided AutoML
Sharma, R., **Mukherjee, S.**, Sipka, A., Hüllermeier, E., Vollmer, S. J., Redyuk, S., & Selby, D. A. *International Conference on Machine Learning (ICML), 2025.*

Had Enough of Experts? Elicitation and Evaluation of Bayesian Priors from Large Language Models
Selby, D. A., Spriestersbach, K., Iwashita, Y., Bappert, D., Warrier, A., **Mukherjee, S.**, Asim, M. N., Kise, K., & Vollmer, S. *Stat 14(2): e70054, 2025.*

Chaotic and complex dynamics expose the limits of counterfactual reasoning
Aalaila, Y., Großmann, G., **Mukherjee, S.**, Wahl, J., & Vollmer, S. *Scientific Reports 16, Article 16978 (2026).*

Workshop publications

Linear-LLM-SCM: Benchmarking LLMs for Coefficient Elicitation in Linear-Gaussian Causal Models
Yamaoka, K., **Mukherjee, S.**, Gärtner, T., Selby, D. A., Konigorski, S., Hüllermeier, E., Bengs, V., & Vollmer, S. J. *Structured Data 4 Health Workshop at ICML, 2026.*

Co-Exploration and Co-Exploitation via Shared Structure in Multi-Task Bandits
Mukherjee, S., Lebedeva, S., Brandt, J., Margraf, V., Hanselle, J., Yamaoka, K., Bengs, V., Konigorski, S., Hüllermeier, E., & Vollmer, S. J. *European Workshop on Reinforcement Learning (EWRL), 2026.*

Distributionally Robust Warm-start for a Population of Bandits
Mukherjee, S., Basu, D., Brandt, J., Bengs, V., Hüllermeier, E., & Vollmer, S. J. *European Workshop on Reinforcement Learning (EWRL), 2026.*

Rethinking Counterfactuals: Hidden Assumptions and Practical Pitfalls
Großmann, G., Aalaila, Y., Selby, D. A., **Mukherjee, S.**, Vollmer, S., & Wahl, J. *CREST, 2026.*

CLAM: Causal Spatial Disaggregation to Infer Local Effects From Coarse Data
Großmann, G.*, **Mukherjee, S.***, & Vollmer, S. *NeurIPS CausScien Workshop, 2025. *Equal contribution.*

SQUID: A Bayesian Approach for Physics-Informed Event Modeling
Mukherjee, S., Vollmer, S., & Großmann, G. *NeurIPS Workshop on Differentiable Systems and Scientific ML, 2025.*

Peculiarities of Counterfactual Point Process Generation
Großmann, G., **Mukherjee, S.**, & Vollmer, S. *ACM SIGSPATIAL STCausal Workshop, 2024, pp. 11–22.*

Graph Agnostic Causal Bayesian Optimisation
Mukherjee, S.*, Zhang, M.*, Flaxman, S., & Vollmer, S. *NeurIPS BDU Workshop, 2024. *Equal contribution.*

Had Enough of Experts? Elicitation and Evaluation of Bayesian Priors from Large Language Models
Selby, D. A., Spriestersbach, K., Iwashita, Y., Bappert, D., Warrier, A., **Mukherjee, S.**, Asim, M. N., Kise, K., & Vollmer, S. *NeurIPS BDU Workshop, 2024; related to 2025 Stat article.*

Flexible Group Fairness Metrics for Survival Analysis
Sonabend, R., Pfisterer, F., Mishler, A., Schauer, M., Burk, L., **Mukherjee, S.**, & Vollmer, S. *DSHealth Workshop at KDD, 2022.*

Preprints

HawkesNest: A Multi-Axis Benchmark for Spatio-Temporal Pattern Complexity
Aalaila, Y., **Mukherjee, S.**, Großmann, G., & Vollmer, S. *Preprint, 2025.*

Challenges and Solutions for Bandits in the Wild: Warm-Started Mixture Bandits for Cross-Cohort Slate Recommendation
Lebedeva, S.*, **Mukherjee, S.***, Sadal, A. A., Ekşi, I., Sharma, R., Mueller, J., Dombrowski, T., Karolus, J., Bengs, V., Hüllermeier, E., & Vollmer, S. *Preprint, 2026. *Equal contribution.*

Research presentations

- Peculiarities of Counterfactual Point Process Generation — ACM SIGSPATIAL STCausal Workshop (2024).
- Graph Agnostic Causal Bayesian Optimisation — Presented at the Workshop on Safe and Causally Informed Deep Reinforcement Learning (2024).
- CLAM: Causal Spatial Disaggregation to Infer Local Effects From Coarse Data — Presented at the STAR Space-Time Causality Reading Group.
- Presented Eventful project research to DFKI's Scientific Advisory Board; received positive feedback.

Selected honors and awards

- Selected for Reinforcement Learning Summer School (RLSS) 2026 and Machine Learning Summer School (MLSS) 2026.
- Third place, MLSS Hackathon: NO-BEAR (bandit-based LLM resource allocation for automated science).
- RPTU Research Stipend (EUR 5,400), Apr–Sep 2023.
- Julia Summer of Code award (USD 1,200), 2021; CSS@ICLR project funding (USD 500), 2022.
- Hult Prize India winner, 2021; represented India at the Castle Accelerator.

Teaching

Guest Lecturer, Collaborative Intelligence (RPTU) Jan – May 2024
Kaiserslautern, Germany
- Lectured on human-in-the-loop learning and active learning; designed coursework and mentored projects.

Head Teaching Assistant, Neural Networks and Fuzzy Logic (BITS Pilani) Aug – Dec 2022
Pilani, India
- Led 12 teaching assistants and coordinated laboratories, grading, and assignments.

Teaching Assistant, Neural Networks and Fuzzy Logic (BITS Pilani) Jan – May 2022
Pilani, India
- Assessed work for 120 students, led practical workshops, and mentored three student teams.

Academic service and leadership

- Reviewer: AAAI 2025; GIScience 2025 STCausal Workshop.
- Co-organizer: STAR Space-Time Causality Reading Group (2025–present); DFKI DSA Research Exchange Group (2024–present).
- Chair, BITS ACM (2022–2023); previously ML SIG and Special Interest Group Coordinator (2020–2022).

Technical Skills

Mathematics & Statistics: Probability, statistical inference, optimization, statistical learning theory.
Machine Learning: Reinforcement learning, sequential decision-making, probabilistic modeling, Bayesian machine learning.
Causal Inference: Causal graphical models, do-calculus, causal discovery and inference, counterfactual reasoning.
Computational Tools: Python, Julia, LaTeX, Lean (basic), numerical simulations, and experimental evaluation.