

Arun Verma

✉ arunverma100@gmail.com ☎ +65 8898 0672 🏠 <https://arunv3rma.github.io>

Research Interests

My research focuses on developing adaptive, modular, and resource-efficient AI systems that safely learn, evolve, and make complex decisions in dynamic real-world environments. I am particularly interested in reinforcement learning, scalable system design, and the reliable deployment of AI to solve real-world problems at scale.

Experience

Singapore-MIT Alliance for Research and Technology Centre Sep 2024 – Present

Postdoctoral Associate, M3S Program

- Designing novel resource-efficient machine learning algorithms for efficient and effective AI systems.
- Formulate novel research directions and lead high-impact, long-term research projects.
- Mentoring and supporting junior researchers and interns with their research projects.

National University of Singapore May 2021 – Aug 2024

Research Fellow, Department of Computer Science

- Developed machine learning algorithms for sequential decision-making problems related to multi-armed bandits, Bayesian optimization, active learning, and reinforcement learning, with applications in AutoML and LLMs.
- Mentored Ph.D. students with their research projects.

Indian Institute of Technology Bombay Jan 2021 – Apr 2021

Research Associate, Industrial Engineering and Operations Research (IEOR)

- Studied the impact of auxiliary information associated with reward to derive a tighter variance bound for the mean reward estimator and exploited it to design more efficient algorithms for sequential decision-making problems.

Teaching Assistant (TA)

Jan 2016 – Dec 2020

- IE 613: Online Machine Learning Head TA for Spring 2018, 2019, 2020 offerings
- IE 605: Engineering Statistics TA for Fall 2020 offering
- IE 611: Introduction to Stochastic Models TA for Fall 2017 offering
- IE 616: Decision Analysis and Game Theory TA for Spring 2017 offering
- *Non-Academic Assistance:* Department System Administrator for IEOB Jul 2016 – Dec 2020

INRIA Lille Jun 2019 – Nov 2019

Raman-Charpak Fellow, SequeL Team (now SCOOOL)

- Worked on the sequential decision-making problem with time-series data, motivated by health-related scenarios where a person has to decide whether to visit a physician based on their health history.

Conduent Labs Aug 2018 – Nov 2018

PhD Intern, Machine Learning & Statistics Group

- Designed and analyzed sequential decision-making algorithms for efficient resource allocation under censored data constraints, where true data is accessible only under specific conditions.

Indian Institute of Science Jul 2015 – Dec 2015

Research Assistant, Department of Computational and Data Sciences

- Developed a Big Data analytics platform for resource management using open-source frameworks (Spark, Storm, HBase, HDFS, Apollo) and data visualization tools (Node.js, D3) to improve monitoring and decision-making.

Indian Institute of Technology Bombay Jul 2014 – Jun 2015

Project Research Assistant, Department of Computer Science & Engineering

- Worked on BodhiTree, an online classroom learning platform. Collaborated with students on projects to add new features and improve existing functionalities.

Project Research Intern, Department of Civil Engineering

Jan 2014 – Jun 2014

- Developed an Android app for real-time tracking of IIT Bombay campus buses.
- Designed an algorithm to compute and display the shortest route on Google Maps using stored geographical data for offline use.

Education

Indian Institute of Technology Bombay

Jan 2016 – Dec 2020

Doctor of Philosophy (Ph.D.)

- **Thesis awards:** Naik and Rastogi Excellence in Ph. D. Thesis Award and COMSNETS Best Ph.D. Thesis Award.

Shobhit University

Jul 2010 – Jun 2014

Bachelor of Technology in Computer Science and Engineering

- **Award:** Received Academic Excellence Award (Gold Medal) for having the highest cumulative GPA in my batch.

Publications

Peer-Reviewed Papers | * denotes equal contributions

1. Tang Zhenwei*, Arun Verma*, Zijian Zhou, Zhaoxuan Wu, Alok Prakash, Daniela Rus, Bryan Low, 'MineDraft: A Framework for Batch Parallel Speculative Decoding,' ICML 2026.
2. Arun Verma, Indrajit Saha, Makoto Yokoo, Bryan Low, 'Keep Everyone Happy: Online Fair Division of Numerous Items with Few Copies,' ICML 2026. (Also presented at ICLR 2026 Workshop on AI for Mechanism Design and Strategic Decision Making.)
3. Ze Yu Zhang*, Arun Verma*, Finale Doshi-Velez, Bryan Low, 'Understanding the Relationship between Prompts and Response Uncertainty in Large Language Models,' ACL Findings 2026. (Also presented at ICLR 2025 Workshop on Quantify Uncertainty and Hallucination in Foundation Models: The Next Frontier in Reliable AI.)
4. Xiaoqiang Lin, Arun Verma, Dai Zhongxiang, Daniela Rus, Bryan Low, 'ActiveDPO: Active Direct Preference Optimization for Sample-Efficient Alignment,' ICLR 2026. (Also presented at ICML 2025 Workshop on Models of Human Feedback for AI Alignment.)
5. Arun Verma, Manjesh K Hanawal, Arun Rajkumar, 'Stochastic Multi-Armed Bandits with Limited Control Variate,' COMSNETS 2026 (**Full Presentation**).
6. Jiangwei Chen, Pham Kieu Thao Nguyen, Rachael Hwee Ling Sim, Arun Verma, Zhaoxuan Wu, Chuan-Sheng Foo, Bryan Low, 'Incentivizing Time-Aware Fairness in Data Sharing,' NeurIPS 2025.
7. Arun Verma, Zhaoxuan Wu, Zijian Zhou, Xiaoqiang Lin, Zhiliang Chen, Rachael Hwee Ling Sim, Rui Qiao, Jingtian Wang, Nhung Bui, Xinyuan Niu, Wenyang Hu, Gregory Kang Ruey Lau, Zi-Yu Khoo, Zitong Zhao, Xinyi Xu, Apivich Hemachandra, See-Kiong Ng, Bryan Low, 'Uncovering Scaling Laws for Large Language Models via Inverse Problems,' EMNLP Findings 2025.
8. Arijit Mukherjee, Mrinmoy Mazumder, Arun Verma, Hitesh Tikariha, Raktim Bhattacharya, Ooi Qi En, Sanjay Swarup, 'A cell-free bacterial signal orchestrates trans-kingdom fitness tradeoff to enhance sulfur deficiency tolerance in plants,' Cell Host & Microbe, 2025.
9. Zhaoxuan Wu*, Zijian Zhou*, Arun Verma, Alok Prakash, Daniela Rus, Bryan Low, 'TETRIS: Optimal Draft Token Selection for Batch Speculative Decoding,' ACL 2025.
10. Arun Verma*, Dai Zhongxiang*, Xiaoqiang Lin, Patrick Jaillet, Bryan Low, 'Neural Dueling Bandits: Preference-Based Optimization with Human Feedback,' ICLR 2025. (Also presented at ICML 2024 Workshop on Foundations of Reinforcement Learning and Control – Connections and Perspectives).
11. Xinyi Xu, Zhaoxuan Wu, Rui Qiao, Arun Verma, Yao Shu, Jingtian Wang, Xinyuan Niu, Zhenfeng He, Jiangwei Chen, Zijian Zhou, Gregory Kang Ruey Lau, Hieu Dao, Lucas Agussurja, Rachael Hwee Ling Sim, Xiaoqiang Lin, Wenyang Hu, Zhongxiang Dai, Pang Wei Koh, Bryan Low, 'Position Paper: Data-Centric AI in the Age of Large Language Models,' EMNLP Findings 2024.
12. Srijith Balakrishnan, Beatrice Cassottana, Arun Verma, 'Application of Clustering Algorithms for Dimensionality Reduction in Infrastructure Resilience Prediction Models,' Structure and Infrastructure Engineering 2024.
13. Arun Verma, Zhongxiang Dai, Shu Yao, Bryan Low, 'Exploiting Correlated Auxiliary Feedback in Parameterized Bandits,' NeurIPS 2023.
14. Zhongxiang Dai, Gregory Kang Ruey Lau, Arun Verma, Shu Yao, Bryan Low, Patrick Jaillet, 'Quantum Bayesian Optimization,' NeurIPS, 2023.
15. Thanh Lam, Arun Verma, Bryan Low, Patrick Jaillet, 'Risk-Aware Reinforcement Learning with Coherent Risk Measures and Non-Linear Function Approximation,' ICLR 2023.
16. Zhongxiang Dai, Shu Yao, Arun Verma, Bryan Low, Patrick Jaillet, 'Federated Neural Bandit,' ICLR 2023.
17. Shu Yao, Zhongxiang Dai, Weicong Sng, Arun Verma, Patrick Jaillet, Bryan Low, 'Zeroth-Order Optimization with Trajectory-Informed Derivative Estimation,' ICLR 2023.
18. Xinyi Xu, Zhaoxuan Wu, Arun Verma, Chuan Sheng Foo, Bryan Low, 'FAIR: Fair Collaborative Active Learning with Individual Rationality for Scientific Discovery,' AISTATS 2023.

19. [Arun Verma*](#), Dai Zhongxiang*, Bryan Low, 'Bayesian Optimization under Stochastic Delayed Feedback,' ICML 2022 (**Spotlight**).
20. [Arun Verma](#) and Manjesh K Hanawal, 'Stochastic Multi-Armed Bandits with Control Variates,' NeurIPS 2021.
21. [Arun Verma](#), Manjesh K Hanawal, Csaba Szepesvári, Venkatesh Saligrama, 'Online Algorithm for Unsupervised Sequential Selection with Contextual Information,' NeurIPS 2020.
22. [Arun Verma](#), Manjesh K Hanawal, N. Hemachandra, 'Thompson Sampling for Unsupervised Sequential Selection,' ACML 2020.
23. Debamita Ghosh, [Arun Verma](#), Manjesh K Hanawal, 'Learning and Fairness in Energy Harvesting: A Maximin Multi-Armed Bandits Approach,' IEEE SPCOM 2020.
24. [Arun Verma](#) and Manjesh K Hanawal, 'Stochastic Network Utility Maximization with Unknown Utility: Multi-Armed Bandits Approach,' IEEE INFOCOM 2020.
25. [Arun Verma](#), Manjesh K Hanawal, Arun Rajkumar, Raman Sankaran, 'Censored Semi- Bandits: A Framework for Resource Allocation with Censored Feedback,' NeurIPS 2019.
26. [Arun Verma](#), Manjesh K Hanawal, Rahul Vaze 'Distributed Algorithms for Efficient Learning and Coordination in Ad Hoc Networks,' WiOPT 2019.
27. [Arun Verma](#), Manjesh K Hanawal, Csaba Szepesvári, Venkatesh Saligrama, 'Online Algorithm for Unsupervised Sensor Selection,' AISTATS 2019.
28. Yogesh Simmhan, Anshu Shukla, [Arun Verma](#), 'Benchmarking Fast Data Platforms for the Aadhaar Biometric Database,' WBDB 2015.

Peer-Reviewed Extended Abstract and Workshops Papers

1. Thanh Q. Tran*, [Arun Verma*](#), Kiwan Wong, Bryan Low, Daniela Rus, Wei Xiao, 'BarrierSteer: LLM Safety via Learning Barrier Steering,' ICLR 2026 Workshop on Principled Design for Trustworthy AI - Interpretability, Robustness, and Safety across Modalities (**Spotlight**).
2. [Arun Verma](#), Indrajit Saha, Makoto Yokoo, Bryan Low, 'COBRA: Contextual Bandit Algorithm for Ensuring Truthful Strategic Agents,' ICLR 2026 Workshop on AI for Mechanism Design and Strategic Decision Making.
3. Ze Yu Zhang*, Nhung Bui*, [Arun Verma*](#), Bolin Ding, Bryan Low, 'MMT: Achieving Immediate Exact Federated Unlearning with Improved Post-Unlearning Performance,' ICLR 2026 Workshop on Test-Time Updates.
4. [Arun Verma](#), Xiaoqiang Lin, Dai Zhongxiang, Daniela Rus, Bryan Low, 'Active Human Feedback Collection via Neural Contextual Dueling Bandits,' ICLR 2025 Workshop on Bidirectional Human-AI Alignment.
5. Xiaoqiang Lin, Dai Zhongxiang, [Arun Verma](#), See-Kiong Ng, Patrick Jaillet, Bryan Low, 'Prompt Optimization with Human Feedback,' ICML 2024 Workshop on Models of Human Feedback for AI Alignment (**Oral**).
6. Srijith Balakrishnan, Beatrice Cassottana, [Arun Verma](#), 'Application of Clustering Algorithms for Dimensionality Reduction in Infrastructure Resilience Prediction Models,' TRB 2024.
7. Srijith Balakrishnan, Beatrice Cassottana, [Arun Verma](#), 'A Network Clustering Approach for Dimensionality Reduction in Machine Learning Models for Infrastructure Resilience Analysis,' ESREL 2022.
8. [Arun Verma](#), Manjesh K Hanawal, N. Hemachandra, 'Unsupervised Online Feature Selection for Cost-Sensitive Medical Diagnosis,' NetHealth, COMSNETS 2020.
9. [Arun Verma](#), 'Online Cost Sensitive Feature Selection in Cascaded Features Problem,' in *Asian Universities Alliance Postgraduate Academic Forum, Tsinghua University, Beijing, China, October 20-22, 2018. Best Poster Award.*
10. [Arun Verma](#) and Manjesh K Hanawal, 'Unsupervised Cost-Sensitive Predictions with Side Information,' CoDS-COMAD 2018.
11. [Arun Verma](#), Yogesh Simmhan, N. Hemachandra, 'Scalable Online Analytics for IoT Applications using Big Data Platforms,' in Graduate Research Workshop, CoDS 2017.

Pre-prints

1. [Arun Verma](#), Manjesh K Hanawal, Arun Rajkumar, Raman Sankaran, 'Censored Semi- Bandits for Resource Allocation.' arXiv pre-print.

Invited Talks

- Talk on 'Resource-Efficient Machine Learning: Toward Sustainable AI,' at MIT-Singapore Symposium on Embodied and Scalable AI: Interfaces, Systems, and Applications on May 28, 2025.
- Talk on 'Sequential Decision Problems with Weak Feedback,' as part of the Multi-Armed Bandits (CS6046) course at IIT Madras on May 4, 2021.
- RIKEN seminar on 'Sequential Decision Problems with Weak Feedback,' at RIKEN AIP, Japan on November 27, 2020.
- SequeL seminar on 'Censored Semi- Bandits' at INRIA Lille, France on November 24, 2019.

- Talk on 'Online Algorithm for Cost-Sensitive Unsupervised Learning,' at RIKEN AIP, Japan during my visit to Approximate Bayesian Inference (ABI) Team from April 9-16, 2019.
- Lighting talk on 'Unsupervised Cost Sensitive Predictions with Contextual Information' at IEOR Day 2018.

Selected Awards

- Awarded Microsoft Travel Grant, COMSNETS Association and LRN Foundation Travel Award, and INFOCOM 2020 Student Conference Award for attending IEEE INFOCOM 2020.
- Travel Award from LinkedIn for attending FATE-ML workshop 2020 at IISc, Bengaluru, India.
- Awarded Google Travel Grant and NeurIPS Student Travel Award for attending NeurIPS 2019.
- Awarded Raman-Charpak Fellowship 2018 for visiting SequeL Team, INRIA Lille, France.
- Awarded ACM-India/IARCS and Microsoft Travel Grant for attending AISTATS 2019.
- Awarded Technical Color of Hostel-14, IIT Bombay for 2017-18 and 2018-19 for technical contributions towards hostel activities during 2017-18 and 2018-19.
- Awarded Sports Color (Cricket) of Hostel-14, IIT Bombay for year 2017-18.
- Conference Travel Awards from HiPC 2016, CoDS 2017, CoDS-COMAD 2018, AUAPAF 2018.
- IITB Research Internship Award 2013-14 from Indian Institute of Technology Bombay.

Academic Services

- Reviewer: NeurIPS 2020-2026 | ICLR 2021-2026 | ICML 2021-2026 | ACL 2025-2026 | EMNLP 2025 | AISTATS 2022-2026 | AAAI 2021-2024, 2026 | IJCAI 2021-2025 | AAMAS 2023-2025 | TMLR 2022-2024 | JMLR 2025-2026.
- Best or Top Reviewer Award: ICML 2022, NeurIPS 2022.
- Volunteer: NeurIPS 2020, AISTATS 2021, ICLR 2021, ICML 2021, NeurIPS 2021, ICLR 2023, ICLR 2026.