Curriculum Vitae, Nishant Sharma nishant.sharma@mila.quebec · Website · GitHub · LinkedIn

About

First-year Research MSc (in Machine Learning) student at Mila. I am broadly interested in AI safety and the geometry of deep learning. I use that lens for continual learning (how a network can keep learning without losing plasticity), for building aligned AI systems, and for the optimization of distributed training. I've previously worked on data attribution and computational neuroscience.

Education

Research MSc, Machine Learning, Université de Montréal, Montréal, Canada *2026 – Present*
Department of Computer Science and Operations Research (DIRO) and Mila – Quebec AI Institute
Advised by Professor Sarath Chandar; co-supervised by Professor Amal Zouaq

B.Tech, Physics, Indian Institute of Technology Delhi, New Delhi, India *2020 – 2024*
Thesis: Synchronization and Dynamics of Neural Phase Oscillators
Advised by Professor Saurabh Gandhi and Professor Pradipta Ghosh

Selected Publications and Preprints

*C: conference, P: preprint, W: workshop. * implies co-first authorship*

[C1] Shubham, K.*, **Sharma, N.***, Tiwari, K., Prathosh, A.P. (2026) Enhancing Trustworthiness of Fine-Tuned LLMs via Regularized Subset Selection (*Published* at ICLR 2026)

[P1] **Sharma, N.**, Gomezjurado Gonzalez, L., Talluri, G., Chandar, S., Naganuma, H. (2026) Understanding Stability of Outer Nesterov Momentum on Stochastic Pseudo-Gradients (*Preprint* on arXiv)

[P2] Quamar, A.*, Areeb, M.*, **Sharma, N.**, Shreekumar, A., Rosenthal, J., Ozmen, M. O., Kuznetsov, M., Celik, Z. B. (2025) Adaptive Blockwise Search: Inference-Time Alignment for Large Language Models (*Preprint* on arXiv)

[W1] **Sharma, N.**, Talluri, G., Naganuma, H. (2026) From Inner Randomness to Outer Stability (*Accepted* at the 18th Annual Workshop on Optimization for Machine Learning, NeurIPS 2026)

[W2] Shubham, K.*, **Sharma, N.***, Tiwari, K., Prathosh, A.P. (2025) Trust, But Attribute: Tracing Impact of Data on Trustworthiness in Supervised LLM Fine-Tuning (*Published* at the Reliable ML from Unreliable Data Workshop, NeurIPS 2025)

[W3] **Sharma, N.**, Quamar, A., Xie, P. (2025) Decoding Histone Modification Signatures of Non-Coding RNAs via Foundation Models (*Published* at the Multi-modal Foundation Models and Large Language Models for Life Sciences Workshop, NeurIPS 2025)

Research Experience

Graduate Researcher, Mila – Quebec AI Institute *Aug 2026 – Present*
Montréal, Canada Advisors: Professor Sarath Chandar & Professor Amal Zouaq

- Working on continual learning (at Chandar Research Lab), AI safety (at LAMA-WeST Lab) and the optimization of deep neural networks. Studied the stability of outer Nesterov momentum on stochastic pseudo-gradients in two-phase optimizers – [P1].

Research Assistant, Representation Lab, IISc Bangalore *May 2025 – Dec 2025*
Bangalore, India Advisor: Professor Prathosh A.P.

- Proposed a two-stage, compute-efficient repair of fine-tuned LLMs that improves trustworthiness (bias, ethics, truthfulness) while preserving downstream performance. Extended influence functions to trust metrics, selected a diverse subset with DPP-based regularization, and repaired the model by natural gradient ascent under the proximal Bregman response function (PBRF) framework – [C1], [W2].

Research Intern, PurSec Lab, Purdue University *May 2025 – Oct 2025*
West Lafayette, Indiana, USA Advisor: Professor Z. Berkay Celik

- Proposed AdaSearch, a decoding method that prioritizes early tokens, with >10 pp gains across alignment tasks (reasoning, sentiment, helpfulness & harmlessness) – [P2].

Research Intern, PXie Lab, UC San Diego *Feb 2025 – May 2025*
San Diego, California, USA Advisor: Professor Pengtao Xie

- Built a simple transfer-learning pipeline using foundation models that jointly screens 50 histone-ncRNA interactions, reaching an AUROC of 0.95 (a 7 pp lead on baselines). Extended the DeepHistone marker window from 1 kbp to 30 kbp with a mean pooling head, showing that a simple CNN tower can discriminate histones from long genomic context – [W3].

Research Assistant, Complex Systems & Neural Dynamics Lab, IIT Delhi *Aug 2023 – Oct 2023*
New Delhi, India Advisors: Professor Saurabh Gandhi & Professor Pradipta Ghosh

- Simulated seizure dynamics in the brain as a network of resource-constrained Kuramoto phase oscillators to uncover explosive synchronization triggers in epileptic seizures. Identified critical weight percentiles in the Mouse Brain Network that reveal bistability conditions, and showed that selective pruning enhances transient explosive synchronization.

Academic Service and Talks

Reviewer, AISTATS 2027, ICLR 2027, NeurIPS 2026, ICLR 2026, AAAI 2026, ICML 2025, NeurIPS 2025 Workshop on Foundation Models and Large Language Models for Life Sciences, NeurIPS 2025 Workshop on Machine Learning and the Physical Sciences *2025 – 2026*

Program Committee, AAAI AI Alignment Track *2025*

Talk, Using LLMs for Simulating Attacks in Provenance Graphs. Presented at *Amazon Research* *2025*

Talk, Certifying GNNs Against Semantic-Preserving Mutations. Presented at *PurSec Lab, Purdue* *2025*

Talk, Delivering Business Value Using LLM Agents in 2024. Presented at *Aptivaa Global R&D* *2024*

Editor, IIT Tech Ambit—National Pan-IIT Technology Magazine *2023*

Achievements & Honors

Fellow, Alberta Machine Intelligence Institute (Amii) Accelerator Program *2026*

Graduate Research Scholarship, CAD $95K (stipend and full tuition coverage) *2026 – 2028*

Member, NVIDIA Inception Program; received $125K in compute *2024*

Top 10 Deep-Tech Startups in India, NASSCOM (1% nationwide) *2024*

Finalist, Samsung R&D Solve for Tomorrow (3% nationwide) *2023*

99.15 percentile in JEE Advanced; 99.58 percentile in JEE Main among 1.2M+ candidates *2020*

2nd Runner-up, NAEST Conference (0.3% nationwide) *2019*

Industry Experience

Member of Technical Staff, Vamana AI *Aug 2025 – Jul 2026*
San Francisco, California, USA

- Worked on inference optimization for VLMs and multi-modal search with temporal localization, and built scalable GPU indexing on Modal. *Advised by Professor Abhinav Shrivastava (University of Maryland).*

Co-Founder & CEO, Eigenfrequency Technologies *Jul 2023 – Mar 2025*
New Delhi, India

- Co-founded and scaled an AI startup to $250K ARR. Pioneered AutoLLM with hierarchical RAG architectures, driving enterprise adoption. Notable clients included Mizuho Bank, PNC Bank and the Supreme Court of India.

Selected Projects

1. *2025, Kronfluence.* Contributor to Anthropic's EKFAC implementation for large language models. Extended approximations of inverse Hessian–vector products and implemented modules for data attribution and trust-based evaluation.
2. *2025, FAdamPyTorch.* Contributor to the FAdam PyTorch library. Implemented natural gradient ascent with Fisher approximation and integrated TorchTitan-compatible logging for large-scale training/unlearning.
3. *2025, Load-Balanced Expert Rerouting for MoE Inference.* Developed a lightweight inference-time rerouting strategy that balances expert loads, reducing latency while preserving memory efficiency.
4. *2023, Hierarchical Reinforcement Learning for Weapon Target Assignment.* Implemented a state-split flow Deep Q-Network to perform real-time threat assessment, reducing computational overhead by 50%.
5. *2023, Neural Contextual Multi-Armed Bandit Algorithms.* Implemented a neural UCB algorithm with Thompson sampling for stochastic contextual bandits with non-linear rewards, matching LinUCB's $O(\sqrt{T})$ regret bound.