Vineet Jain

vinjain96@gmail.com ◇ Vineet Jain ◇ vineetjain96 ◇ Vineet Jain

Research Statement

My long-term goal is to build autonomous agents that *reason* about their environment, *adapt* through interaction, and *generalize* to unseen tasks. My research lies at the intersection of reinforcement learning, probabilistic inference, and generative modeling. My recent work investigates how reinforcement learning and inference-time search can improve LLM reasoning, long-horizon decision-making, and open-ended discovery.

Education

McGill University & Mila – Quebec AI Institute — Sep 2021 – Present
Ph.D. in Computer Science (expected early 2027) — GPA: 4.0 / 4.0
Advisor: Dr. Siamak Ravanbakhsh

Carnegie Mellon University — Aug 2019 – Dec 2020
M.S. in Machine Learning — GPA: 3.89 / 4.0

Indian Institute of Technology, Kharagpur — Aug 2014 – Aug 2019
B.Tech. in Electronics and Electrical Communication Engineering — GPA: 9.44 / 10
Minor in Computer Science
(Integrated M.Tech. in Visual Information Processing and Embedded Systems)

Selected Publications

* denotes equal contribution

1. **Recursive Self-Aggregation Unlocks Deep Thinking in Large Language Models** — Project webpage
Siddarth Venkatraman*, Vineet Jain*, Sarthak Mittal*, Vedant Shah, Johan Obando-Ceron, Yoshua Bengio, Brian Bartoldson, Bhavya Kailkhura, Guillaume Lajoie, Glen Berseth, Nikolay Malkin, Moksh Jain
Preprint, 2026

2. **A Comedy of Estimators: On KL Regularization in RL Training of LLMs** — arXiv
Vedant Shah*, Johan Obando-Ceron*, Vineet Jain*, Brian Bartoldson, Bhavya Kailkhura, Sarthak Mittal, Glen Berseth, Pablo Samuel Castro, Yoshua Bengio, Nikolay Malkin, Moksh Jain, Siddarth Venkatraman, Aaron Courville
COLM 2026

3. **Long-Horizon Model-Based Offline Reinforcement Learning without Conservatism** — OpenReview
Tianwei Ni, Esther Derman, Vineet Jain, Vincent Taboga, Siamak Ravanbakhsh, Pierre Luc-Bacon
ICML 2026

4. **Diffusion Tree Sampling: Scalable inference-time alignment of diffusion models** — Project webpage
Vineet Jain, Kusha Sareen, Mohammad Pedramfar, Siamak Ravanbakhsh
NeurIPS 2025

5. **Sampling from Energy-based Policies using Diffusion** — arXiv
Vineet Jain, Tara Akhound-Sadegh, Siamak Ravanbakhsh
RLC 2025

6. **Learning to Reach Goals via Diffusion** — PMLR
Vineet Jain, Siamak Ravanbakhsh
ICML 2024

7. **On Diffusion Modeling for Anomaly Detection** — OpenReview
Victor Livernoche*, Vineet Jain*, Yashar Hezaveh, Siamak Ravanbakhsh
ICLR 2024 (Spotlight)

8. **EqR: Equivariant Representations for Data-Efficient Reinforcement Learning** — PMLR
Arnab Mondal, Vineet Jain, Kaleem Siddiqi, Siamak Ravanbakhsh
ICML 2022

Research Experience

Meta Superintelligence Labs, Consumer Agents Aug 2026 – Present
Research Scientist Intern London, UK
- Research on post-training and RL for developing long-horizon agents.

Hosts: Alan Schelten

Google Research, Paradigms of Intelligence May 2026 – Aug 2026
Research Intern Montréal, QC, Canada
- Developed a feedback-driven LLM harness for closed-loop control of open-ended search.
- Applied the harness to automated algorithm discovery, enabling LLMs to interpret progress and adapt the search process.

Hosts: Roman Pogodin, Blake Richards

Valence Labs / Recursion Mar 2025 – July 2025
Research Intern Montréal, QC, Canada
- Developed diffusion-based Boltzmann samplers for 3D molecular conformations using only energy evaluations.
- Designed a reinforcement learning method that reduced gradient variance and improved sample quality.

Host: Emmanuel Bengio

Bosch Center for Artificial Intelligence Jan 2021 – Jul 2021
Machine Learning Research Intern Pittsburgh, PA, USA
- Investigated domain generalization under severe imbalance across source domains.
- Developed a generative adversarial network-based method to generate candidate samples for underrepresented domains.

Host: Devin Willmott

Université de Montréal and Mila – Quebec AI Institute May 2018 – Jul 2018
Research Intern Montréal, QC, Canada
- Investigated teacher-guided training for neural network classification.
- Developed a learning algorithm that generates student training examples using feedback from a pretrained teacher network.

Host: Yoshua Bengio

École polytechnique fédérale de Lausanne May 2017 – Jul 2017
Research Intern Lausanne, Switzerland
- Investigated memory-augmented and recurrent neural networks for sequence learning.
- Implemented a Neural Turing Machine and benchmarked it on copying, infilling, and language-generation tasks.

Host: Wulfram Gerstner

Invited Talks

Inference-time alignment of diffusion models June 2025
Sampling reading group, Valence Labs Montréal, QC, Canada

Learning to Reach Goals via Diffusion July 2024
Mila RL workshop Montréal, QC, Canada

Skills

Programming Languages: Python, C++, Bash
ML Frameworks: PyTorch, JAX, TensorFlow, vLLM
Tools: Hydra, DeepSpeed, SLURM, scikit-learn, Git, LaTeX
Human Languages: Hindi (native), English (fluent), French (intermediate)

Teaching & Academic Service

- Teaching Assistant, COMP 588, McGill University, Winter 2023.
- Teaching Assistant, Basic Electronics Laboratory, IIT Kharagpur, Fall and Winter 2019.
- Graduate Admissions Committee, Mila, 2023; Machine Learning Department, Carnegie Mellon University, 2020.
- Reviewer: NeurIPS 2022–2025; ICML 2023–2026 (Gold reviewer in 2026); ICLR 2024–2026; CVPR 2023–2024; ECCV 2022; ICCV 2023; selected ICML workshops