

Aditya Raj

Mohali, India

adityakeshrissm123@gmail.com | LinkedIn | GitHub

Education

Indian Institute of Technology Madras 2025–2029
B.S. in Data Science and Applications

Chandigarh University 2025–2029
B.E. in Computer Science Engineering (Artificial Intelligence & Machine Learning)

Research Interests

Model Based Reinforcement Learning, World Models, Exploration and Uncertainty, Decision Making, AI Agents, Scientific Machine Learning

Research Experience

GroupBrainAI Sep. 2026–Present
Research Intern

- Investigating how reinforcement learning agents can use learned world models, uncertainty, and information gain to select informative experience and improve sample efficiency.
- Studying VIME, Plan2Explore, OPAX, and LOVE while designing controlled experiments to compare general model uncertainty with uncertainty that is relevant to an agent’s decisions.

Research Manuscripts

Beyond Accuracy for Evaluating Evidence-Aligned Multimodal Reasoning

Co-author | Manuscript under revision

- Proposed **EviBench**, a seven-axis benchmark for evaluating evidence-grounded reasoning in vision-language models beyond answer accuracy.
- Evaluated state-of-the-art open- and closed-source vision-language models across multiple benchmarks using evidence alignment, confidence calibration, and counterfactual robustness.

Photonic Quantum Reservoir Computing on Real Financial Surfaces: Identifying the Structural Ceiling of Near-Random-Walk Data

Co-author | Accepted for publication in Physica Scripta

- Investigated the predictive limits of photonic quantum reservoir computing for financial time-series exhibiting near-random-walk behavior.

Research Projects

Physics-Informed World Models for Learning Dynamical Systems 2026

- Developed a Physics-Informed World Model using Latent Neural ODEs for long-horizon orbital prediction.
- Compared MLP, PINN and PIWM on a two-body benchmark; achieved $25\times$ extrapolation after training on only 20% of trajectories.
- Open-sourced implementation: GitHub.

Investigating Reward Design and the Alignment Problem in Reinforcement Learning 2026

- Studied reward misspecification and specification gaming using PPO.
- Demonstrated reward hacking in MountainCar and LunarLander environments.
- Open-sourced implementation: GitHub.

Technical Skills

Languages: Python, C++, C

ML/DL: PyTorch, TensorFlow, Scikit-Learn, NumPy, Pandas, Vision-Language Models

RL: Gymnasium, Stable-Baselines3

Quantum: Qiskit, PennyLane, pyQPanda3

Databases: MySQL, PostgreSQL

Tools: Git, GitHub, Jupyter Notebook

Awards & Competitions

- CIC Wukong Cup International Quantum Computing Competition (2026): Third Prize** among 232 finalist teams.
- MIT iQuHACK (2026): Team Leader;** developed a hybrid quantum-classical optimization pipeline using QAOA and CUDA-Q.
- Q-volution Hackathon (2026): Honorable Mention** for a photonic quantum kernel model.
- Yale Peaked Hackathon (2026): Placed 21st** among 549 teams.