

# KAINS P KAVULURI

MILA • UNIVERSITÉ DE MONTRÉAL • MONTREAL, QC  
+1 (514) 464-9006 • kainspraveen33@gmail.com • github.com/kainspraveen • linkedin.com/in/kainsprav

## EDUCATION

---

### UNIVERSITÉ DE MONTRÉAL & MILA - QUEBEC AI INSTITUTE

MONTREAL, CA

#### Master's Degree

Sept 2025 - Present

GPA: 4.3/4.3

Coursework: Fundamentals of Machine Learning, Data Science, Representation Learning

### NATIONAL INSTITUTE OF TECHNOLOGY KARNATAKA

KARNATAKA, IN

#### Bachelor's Degree

July 2017 - Dec 2021

Coursework: Computer Vision, Neural Networks, Deep Learning, Algorithms, Distributed Systems, Data Structures, High Performance Computing, Graph Theory

## RESEARCH PROJECTS

---

### RL FOR DRUG DISCOVERY (Ongoing)

Improving drug discovery using multiple agents acting on different policies to build molecules with multiple rewards.

- Finetuned a Diffusion Language Model (Dream-7B) as a generator followed by reinforcement learning post-training using a combination of multi-agent RL, GRPO, and GDPO.
- Improving molecule generation conditioned on natural language description compared to auto-regressive models (e.g., T5, LLaMA).

### REINFORCEMENT LEARNING IN NON-STATIONARY ENVIRONMENTS (Ongoing)

- Investigating loss of plasticity in neural networks and experimenting with evolutionary algorithms to overcome it.

### PROGRESSIVE CONFORMER FOR IMAGE CLASSIFICATION

- A hybrid architecture combining Transformers and CNNs at the attention level for progressive feature learning.

## SELECTED PROJECTS

---

### EXPECTED GOAL PREDICTION ON NHL RAW DATA

- Developed a robust goal probability prediction system for NHL gameplay data, achieving a peak AUC of 0.88 by benchmarking varied architectures.
- Engineered 18+ spatial and temporal features (e.g., shot angle vectors, rebound dynamics, power-play context).
- Conducted a rigorous comparative study between tree-based ensembles and neural networks; utilized Bayesian and Grid Search for hyperparameter tuning.

### AI AGENTS FOR COMPUTER USE

- Built an LLM agent enabling it to browse the web like a human, following natural language instructions.
- Processed cleaned HTML code of source pages into the LLM context, utilizing DSPy for few-shot prompting.

### ADDITIONAL PROJECTS

- **Deep UNet for Melanoma Segmentation:** Image segmentation model for Dermoscopy images used in downstream disease classification tasks.
- **Light-to-Camera Indoor Positioning:** Developed a novel indoor positioning system for mobile devices using light-based communication, prototyped with a Raspberry Pi.
- **AI for Dementia:** Built an iOS app assisting patients with Dementia using voice-to-text and text-to-voice models to remind people and items through tailored, evolved routines.

## PROFESSIONAL EXPERIENCE

---

### DEUTSCHE BANK GROUP

PUNE, IN

#### Associate, Chief Innovation Office

Mar 2024 – Sep 2025

- Architected and scaled LLM Agents for enterprise financial analysis, including a high-precision document QA system equipped with inline source-text highlighting to ensure zero-hallucination trust.
- Spearheaded the development of **DBLumina**, a strategic AI research platform built on Gemini and Vertex AI. Engineered hybrid semantic/keyword retrievers and automated evaluation frameworks for conversational QA.

- Integrated Reciprocal Rank Fusion (RRF) scoring and Maximal Marginal Relevance (MMR) re-ranking, accelerating analyst data discovery and search efficiency by 50%.
- Deployed HNSW indexing to dramatically accelerate retrieval latency across large-scale vector databases.
- Designed an end-to-end, multi-format (PDF, DOCX, PPTX) document parser for robust Gemini-powered RAG pipelines, fusing internal and FactSet data for real-time market insights.

**Senior Analyst, Chief Innovation Office**

**July 2022 – Mar 2024**

- Fine-tuned and instruction-tuned open-weight LLMs via LoRA for specialized banking applications, heavily optimizing for low-latency, secure cloud deployment.
- Built production-grade RAG pipelines for Transitional Risk and ESG analysis, establishing a bank-wide RAG evaluation framework utilizing LangChain, LangGraph, and MLflow.
- Hardened open-source MLOps frameworks (MLflow, Ludwig) to comply with rigorous internal security and data governance requirements.
- Engineered a predictive machine learning system (LightGBM, Random Forest) to forecast residential property energy consumption, directly supporting the bank’s global ESG compliance targets.

**Analyst**

**July 2021 – July 2022**

- Slashed trading book processing time by over 2 hours per book by deploying a predictive machine learning model estimating expected computational completion times.
- Engineered a highly concurrent microservice architecture capable of processing over 1,000,000 daily trade events with low latency.
- Provisioned and maintained a high-throughput, on-premise PostgreSQL database cluster utilizing Docker and OpenShift.

**PULSE SECURE LLC**

**BENGALURU, IN**

**Software Engineer Intern**

**Apr 2020 – Jun 2020**

- Developed a Deep LSTM neural network in TensorFlow to identify anomalous behavior and insider threats from enterprise log data, benchmarking against the CERT dataset.
- Shipped a Node.js web application for centralized log visualization, significantly accelerating debugging workflows for the engineering team.

**SKILLS**

---

- **Programming Languages:** Python, C++, Java, Bash Scripting
- **Tech Stack:** Google Cloud Platform (GCP), Vertex AI, Docker, Kubernetes (GKE), Kubeflow, TensorFlow, PyTorch, Scikit-learn, LangChain, FastAPI, Pandas, Git.
- **Soft Skills:** Project Leadership, Mentoring, Public Speaking, Team Collaboration