

# MADHAVAN PANNEERSELVAM KUMAR

Chicago, IL | (312) 687-9437 | [pkmadhavan9802@gmail.com](mailto:pkmadhavan9802@gmail.com) | [linkedin.com/in/madhavanpk](https://www.linkedin.com/in/madhavanpk) | [github.com/PKMadhavan](https://github.com/PKMadhavan) | [pkmadhavan.vercel.app](https://pkmadhavan.vercel.app)

## SUMMARY

M.S. Artificial Intelligence student (2026) who builds production-grade agentic systems — RAG pipelines, multi-agent orchestration, and RLHF fine-tuning — from scratch. Differentiates through eval-first development: golden datasets, LLM-as-judge scoring, and CI quality gates that confirm systems work before scaling them.

## EDUCATION

|                                                               |          |
|---------------------------------------------------------------|----------|
| M.S. Artificial Intelligence — DePaul University, Chicago, IL | Aug 2026 |
| M.S. Applied Data Science — SRM University, India             | Apr 2024 |
| B.S. Data Science — Sri Ramachandra University, India         | Apr 2022 |

## TECHNICAL SKILLS

**LLM & NLP:** Claude API (Anthropic SDK), OpenAI API, prompt engineering, LangGraph (cyclic agents, state machines), LangChain, MCP servers

**Evals & Observability:** RAGAS, LLM-as-a-judge, LangSmith (tracing, metrics), CI/CD quality gates, regression testing

**ML & Modeling:** PyTorch, Hugging Face (transformers, trl), RLHF (SFT, reward modeling, PPO), scikit-learn, XGBoost, SHAP

**RAG & Knowledge Retrieval:** RAG pipeline design, Qdrant, pgvector, BM25 + dense hybrid retrieval

**Engineering & Infra:** Python, FastAPI, Docker, PostgreSQL, GCP, AWS, GitHub Actions (CI/CD), SQL

## RESEARCH

**ML-Powered Corporate Bond Liquidity Scoring: A Multi-Measure XGBoost Framework on FINRA TRACE Data with SHAP Interpretability**

*Accepted — 2nd International Online Conference on Risk and Financial Management (IOCRF 2026), JRFM/MDPI*

- Developed an XGBoost-based liquidity scoring model on FINRA TRACE bond transaction data, engineering multi-measure features across bid-ask spread, turnover, and price impact to capture composite market liquidity.
- Applied SHAP for model interpretability — identifying which liquidity signals drive predictions most, enabling transparent, auditable scoring.

## KEY PROJECTS

**RLHF Fine-Tuning Pipeline — SFT, Reward Modeling, PPO *Personal Project***

**Stack:** Python, PyTorch, GPT-2, trl, Hugging Face

- Built a from-scratch RLHF pipeline on GPT-2 (SFT → reward model on 160K preference pairs → PPO), running systematic ablations on reward-model capacity and KL-penalty coefficients to characterize convergence and stability.

**Multi-Agent System with MCP Tool Integration & A2A Orchestration *Personal Project***

**Stack:** Python, LangChain, GPT-4o-mini, ChromaDB, MCP servers, A2A, OpenAI API

- Built an orchestrator that routes tasks to specialist agents over A2A, each backed by sandboxed MCP tools (RAG, calculator, web search), with tool-level authorization boundaries.

**Production RAG System with Evaluation Framework *Personal Project***

**Stack:** Python, FastAPI, pgvector, Hugging Face, BM25, Redis, GCP, GitHub Actions, LangSmith, RAGAS

- Built an end-to-end RAG pipeline (ingestion → embedding → hybrid dense + sparse retrieval → FastAPI serving) with full LangSmith observability and RAGAS evals enforced via CI/CD quality gates.

**Dreamer-V1 World-Model RL — Autonomous Highway Driving *Personal Project***

**Stack:** Python, PyTorch, Docker, GitHub Actions (CI)

- Built a Dreamer-V1 world-model RL agent that learns highway driving policies by training entirely on imagined latent-space rollouts; containerized with Docker and CI, with a live demo.

**Agentic Cleantech AI System *Personal Project***

**Stack:** Python, LangChain, ChromaDB, Tavily, FastAPI

- Built a hybrid RAG agent for cleantech news using a custom ReAct agent over ChromaDB with web-search fallback, evaluated with an LLM-as-judge harness.

## WORK EXPERIENCE

**Machine Learning Intern | Global Techno Solutions, India**

Jun – Sep 2023

- Built production Python ML pipelines (scikit-learn, TensorFlow) automating ingestion, validation, training, and inference; reduced runtime 30% through GPU workload profiling and restructuring.

**Data Analytics & Process Automation Intern | Alteryx, India**

Jun – Aug 2022

- Turned 5+ messy manual workflows into automated, validated ETL pipelines across 3 departments (60% effort reduction); wrote SOPs so the team could own them independently.