Prem Gorde

949-302-1053 | premkgorde@gmail.com | linkedin.com/in/premkgorde | github.com/pkgorde

Experience

Software Development Engineer July 2025 – Present

Amazon Web Services (AWS) *Santa Clara, CA*

- Contributed to the **SageMaker Unified Studio** Data Engineering team, developing next-generation **machine learning workflow and model management features** integrated with AWS's latest service offerings.
- Spearheaded **UI development** for the new **SageMaker Model Hub: Foundation Models**, building scalable React/TypeScript components and integrating AWS APIs and CloudFormation resources, **improving model discovery and deployment efficiency by 40%**.
- Collaborated cross-functionally with **Product Managers, UX Designers, and Senoir/Principal Engineers** to deliver high-impact releases under strict deadlines, consistently meeting AWS's *Raising the Bar* engineering standards.
- Developed and maintained **data ingestion and monitoring pipelines** for SageMaker Unified Studio and EMR Studio using AWS Lambda, S3, CloudWatch, and Glue, ensuring high availability and operational excellence.
- Handled **on-call responsibilities**, triaging and resolving production incidents across distributed services to maintain **¿99.9% reliability SLAs** and deliver rapid customer issue resolution.
- Authored internal documentation, design proposals, and operational playbooks to improve team onboarding, knowledge sharing, and long-term system scalability.

Software Engineer April 2025 – July 2025

T-Mobile (Contract via IOPEX Technologies) *San Jose, CA*

- Developed and deployed multiple full-stack applications across Agile teams, using GitLab, Docker, and Portainer CI/CD pipelines.
- Architected a Statistical Anomaly Detection platform for time-series voice lane data using Python (pandas, scikit-learn, SQLAlchemy); delivered interactive dashboards via Streamlit for KPI monitoring.
- Led design and delivery of a data visualization dashboard platform with React.js/Vite frontend, Node.js/Express APIs, and PostgreSQL backend; managed end-to-end ETL and deployment workflows.

Founding Software Engineer October 2024 – January 2025

Tetsuwan Scientific *San Francisco, CA*

- Spearheading development of an AI-integrated platform for wet lab robotic devices, enabling automated workflows and enhancing efficiency in bio-chemistry research.
- Built and maintained frontend UI platform equipped with Langchain to deliver solutions. Integrated GCP and Azure registrations for new users, along with psql-database management, and Node.JS middleware.
- Integrated multi-agent autonomous AI system and voice transcription/summarization features (Whisper), transforming AI-generated outputs into actionable data structures for research workflows.
- Collaborated with Lab-automation and mechanical engineers and scientists to design intuitive interfaces, effectively communicating results and maintaining real time, low-latency workflow, resulting in 250% speedup in experimentation time

Graduate Student Researcher – Machine Learning & Data Science March 2023 – June 2024

UC Davis Health *Davis & Sacramento, CA*

- Built a Retrieval-Augmented Generation (RAG) system for clinical trial documents using fine-tuned LLMs (GPT-4, Claude Sonnet 3, Llama3), boosting retrieval performance by 30%. GitHub: Clinical-Trial-RAG.
- Engineered a full ETL pipeline for CT angiogram data, processing 8k+ DICOM sequences for ML-ready inputs.
- Implemented E3D-LSTM, ConvLSTM, and SimVP models via **OpenSTL**, improving aneurysm/internal bleed detection by 15±5%.
- Led model tuning, A/B testing, and dataset curation to enhance temporal prediction under noisy conditions.
- Partnered with radiologists and data scientists to ensure clinical relevance and interpretability of model outputs.

Data Science Research Intern July 2023 – September 2023

Lawrence Livermore National Laboratory *Livermore, CA*

- Performed statistical analysis on 443 CT scans using SSIM, variance, and pixel-wise metrics to model material panel degradation over 4 years, contributing to materials stability research, published in Technical Report.
- Identified stable "bad pixel" patterns and spatial burn-in effects; validated as predictive features for panel health monitoring.
- Analyzed scan frequency (duty cycle) trends, revealing strong correlation with dark current degradation in exposed regions.
- Collaborated with physicists and imaging scientists to deliver actionable insights, contributing to LLNL's structural health monitoring systems.

Projects

Kaggle x Google Gen AI Intensive Capstone: *GreenByteAI*

- Completed KaggleXGoogle's GenAI Intensive course. Content covered: foundation models, RAG, Agentic AI with LangGraph, embeddings, prompt engineering, Vertex AI fine-tuning.
- Built GreenByteAI: an Agentic AI pipeline that extracts nutrition and environmental (carbon, water) impact from meal images via tool calling, multi-modal (vision) models (gemini-2.0), and search/database lookups.

Machine Learning for Non-Majors (ECS 111) – Course Development & Instruction

- Designed and taught a Data Science/ML course for upper-division students, covering data science workflows, statistical analysis, and model implementation using NumPy, Pandas, Scikit-learn, and Seaborn.
- Led assignment design and group projects in data processing, model selection, and result interpretation. Guided student project work resulting in 2 journal submissions; average course grade: 87%.

Education

University of California - Davis Davis, CA

***Masters** of Science in **Computer Science**, Specialize in Data Science/Statistics*

University of California - Irvine Irvine, CA

***Bachelors** of Science in **Computer Science**, Specialize in Intelligent Systems; Minor in Management*

Technical Skills

Languages: Python, R, C/C++, Java, SQL (MySQL/pSQL/NoSQL), JavaScript/TypeScript, HTML/CSS
Frameworks: React, Node.JS, Next.JS, Flask, Django, FastAPI, TensorFlow, PyTorch, Scikit-Learn, Express
Developer Tools: Git, Docker, Kubernetes, Portainer, Google Cloud Platform (GCP), AWS, Azure Entra Suite, VSCode, Cursor
AWS Tools: Sagemaker Unified Studio, EMR Studio, Sagemaker AI, CloudWatch, CloudFormation,
Libraries: pandas, NumPy, SciPy, Matplotlib, Seaborn, Plotly, HF-Transformers, LangChain, LlamaIndex, SQLAlchemy
DS/ML alogrithms: RAG, Agentic AI, CNNs, Transformers, Neural Networks, Constraint Satisfaction, Core ML methodologies