

Gavin Abrigo

Merced, CA | (408)-888-8732 | abrigogavino5@gmail.com | [LinkedIn](#) | [GitHub](#)

SUMMARY

Undergraduate computer science student with hands-on experience developing Python and R data pipelines, statistical models, and machine learning workflows. Skilled in data preparation, SQL, vector embeddings, image classification, knowledge graph retrofitting, and retrieval-augmented generation using PyTorch and pandas. Seeking a machine learning engineering internship focused on data-driven analysis, workflow automation, and production-oriented ML systems.

EDUCATION

University of California Merced

Bachelor of Science, Computer Science & Engineering

May 2027

Merced, CA

TECHNICAL SKILLS

- **Programming Languages:** TypeScript, JavaScript, GoLang, Java, Python, C/C++, HTML/CSS, LaTeX, MIPS Assembly, SQL, MATLAB, SPARQL, R
- **Data & Databases:** PostgreSQL, SQLite, MongoDB, Excel
- **Principles & Tools:** Shell, React.js, FastAPI, Flask, REST APIs, Git, GitHub, Node.js, Next.js, Vite, StreamLit, Docker, CI/CD, JSON, System Architecture, GCP
- **AI, Machine Learning & Statistical Methods:** Nvidia NIM, Tesseract, OpenCV, PyTorch, Pandas, LLM API, Claude API, OpenAI API, Prompt Design, Large Language Models, Foundational Models, Retrieval Augmented Generation, Statistical Modeling, Machine Learning, Causal Inference, Recommender Systems, Regression, Classification Trees, Unsupervised Learning, XGBoost, Scikit-learn, NumPy

EXPERIENCE

LLM Knowledge Graph Retrofitter (KGR) - Prof. Pandey

Sep 2025 - Present

Undergraduate Research Intern

Merced, CA

- Improve factual accuracy and reduce hallucinations in NLP models by building a Python-based Knowledge Graph Retrofitting (KGR) and Retrieval-Augmented Generation (RAG) pipeline that integrates Wikidata through GraphRAG, resulting in a noticeable reduction in hallucination occurrences on internal test sets
- Fine-tune entity-linking systems and evaluate LLM verification performance on benchmark datasets using exact-match, F1, and classification-tree metrics with scikit-learn, achieving higher verification scores than the baseline
- Perform statistical modeling, including regression and causal inference, on knowledge-graph data with R to uncover key entity relations, streamline the retrofitting workflow, and improve overall efficiency
- Experiment with transformer-based deep learning architectures in PyTorch to enhance retrieval-augmented generation, leading to faster response times and higher relevance in generated outputs

BART - Software Engineering Capstone

Jan 2026 - May 2026

Software Engineering Associate

Merced, CA

- Automatically classified infrastructure incident images submitted through a web and mobile application using structured metadata and unsupervised learning techniques, routing submissions to appropriate maintenance teams
- Developed image-embedding and similarity-detection workflows with Google Gemini to deduplicate submissions, store vector embeddings in PostgreSQL, extract contextual location data, and reduce manual review overhead
- Collaborate with product management and engineering teams to refine feature requirements and deliver new app functionalities in an agile environment

PROJECTS

NYC Taxi Trip Duration Regression and Digit Classification [Github](#)

Oct 2025 - Dec 2025

- Built end-to-end machine learning pipelines in Python with scikit-learn to predict NYC taxi trip durations and classify MNIST digits, applying statistical modeling and machine learning techniques.
- Implemented logistic regression with hyperparameter tuning and an XGBoost model with 5-fold stratified cross-validation on large-scale datasets, using Pandas for data preprocessing and outlier removal.
- Presented model performance and analytical findings to non-technical stakeholders, translating technical results into actionable recommendations.

ML-ZERO [Github](#)

Jul 2026 - Aug 2026

- Implemented character-level language models from scratch in PyTorch (bigram, MLP, batch-normalized nets) following Andrej Karpathy's makemore curriculum to build foundational-model intuition beyond high-level APIs.
- Manually derived backpropagation gradients and built a WaveNet-inspired convolutional architecture, reinforcing the gradient-based optimization concepts applied in the KGR research pipeline.
- Built a nanoGPT-style transformer with self-attention in PyTorch to replicate core LLM mechanics, extending prompt design and LLM API experience into low-level model architecture.