

Drew Colmenar

drew_colmenar@berkeley.edu | [linkedin.com/in/drew-colmenar](https://www.linkedin.com/in/drew-colmenar) | drewcolmenar.org

Education

University of California, Berkeley

August 2021 - May 2025

B.A. in Data Science, Domain Emphasis in Economics

Experience

[RecallR AI](#)

San Francisco, California

Product & Engineering Lead

June 2025 – November 2025

- Designed stateless APIs for long-term memory in LLM agents, enabling structured writes, retrieval, and conflict resolution
- Led development of hybrid RAG system combining semantic search (Milvus vector db), knowledge graph (MongoDB adjacency list), and heuristic scoring (temporal + BM25) achieving 97.5% accuracy on LongMemEval tasks
- Reduced context size by >90% on 1m+ token sessions vs naive full-context, avoiding quadratic context scaling
- Defined Postgres schemas for structured user and memory metadata (timestamps, entities), enabling retrieval and filtering

Tesla

Palo Alto, California

Data Engineer and Analytics Intern

January 2024 – May 2024

- Identified fragmentation in analytics pipelines and refactored Airflow + BigQuery pipelines from ETL to ELT, consolidating 100+ pipelines into 6 analyst-driven datasets
- Organized migration of legacy Tableau dashboards to Tesla's in-house analytics tool, ensuring continuity and consistency
- Coordinated 40+ stakeholder sessions to resolve data gaps and align metric definitions across teams for new dashboards
- Analyzed ETL pipelines and produced documentation with a python script, improving analytics workflows transparency
- Enabled accurate metric definitions and debugging by building a SQL parser to trace column lineage in complex queries

EVgo

Remote

Product & ML Intern

January 2023 – July 2023

- Led creation of internal ops platform replacing Excel workflows, consolidating station metadata, feedback, and analytics
- Analyzed 200k+ free-text entries with a custom NLP model for structured product insights, identifying key user pain points
- Delivered market and sentiment analysis to product and ops teams via Tableau, informing feature and expansion decisions

Kaiser Permanente

Los Angeles, California

Data Analytics Intern

May 2022 – August 2022

- Built and monitored PowerBI dashboards to track 150+ QA reports, reducing weekly review meeting time by ~30 minutes
- Engineered a SQLite backend for PowerBI dashboards, automating consolidation of 20+ Excel files into a unified dataset
- Queried MySQL to extract and structure operational data, delivering actionable insights to cross-functional QA teams

SAAS (Student Association of Applied Statistics @ UC Berkeley)

Berkeley, California

Project Manager, Executive Leadership

August 2023 – May 2025

- Led a six-person team to automate shipment-data extraction using OCR and CNN pipelines.
 - Delivered lectures on Statistics, SQL, and Probability to 100 Berkeley students with custom Jupyter Notebook materials
 - Organized multiple machine learning case competitions, sourcing datasets and developing outlines for 60+ projects
-

Projects

[FineAnts](#) – AI-Powered Personal Finance Platform

Python, FastAPI, Supabase (Postgres), React, Plaid API, OpenAI

- Drove development of core product to 300+ beta users, shaping features through user feedback and usage analytics
 - Designed LLM orchestration system routing user messages to parameter driven tools for data retrieval (transactions, trends, goals), goal creation, graph generation, validation, and web search to produce final contextual responses
 - Built secure transaction ingestion layer via Plaid, designing financial data models in Postgres to support scalable analytics
 - Developed cost-aware model routing and rate limiting to optimize token usage, feature injection, and latency
-

Skills & Abilities

Languages & Frameworks: Python, SQL, Tableau, FastAPI, React, Next.js, PostgreSQL, BigQuery, Git, Docker

ML & Data Libraries: PyTorch, TensorFlow, Scikit-learn, Pandas, Transformers, Tokenizers, Milvus, LangChain

Skills: AI integration, RAG pipelines, A/B testing, data engineering, data analysis, prompt engineering