

# KENT BOURGOING

310-409-6668 | [kent1bp@berkeley.edu](mailto:kent1bp@berkeley.edu) | Los Angeles, CA | [linkedin.com/in/kent-bourgoing/](https://www.linkedin.com/in/kent-bourgoing/) | [github.com/kentbourgoing](https://github.com/kentbourgoing) | [kentbourgoing.com](http://kentbourgoing.com)

Recent data science master's graduate with experience building and deploying ML systems that deliver measurable business impact. Leverage ML, causal inference, hypothesis testing, and time series analysis, plus consulting skills to deliver stakeholder-ready results.

## EDUCATION

**University of California, Berkeley, School of Information**, Berkeley, CA

*December 2025*

*Master of Information and Data Science (MIDS)*

- **GPA:** 4.0 / 4.0
- **Relevant Coursework:** ML at Scale, NLP with Deep Learning, ML Systems Engineering, Experiments and Causal Inference, Statistical Methods for Discrete Response, Time Series, and Panel Data.

**University of California, Los Angeles (UCLA)**, Los Angeles, CA

*June 2023*

*Bachelor of Science, Chemical Engineering*

- **GPA:** 3.73 / 4.0

## SKILLS

- Languages: Python, R, SQL, MATLAB
- ML/NLP/LLMs: scikit-learn, PyTorch, TensorFlow/Keras, Transformers, RAG/GraphRAG (LangGraph)
- Data & Apps: PostgreSQL, Neo4j (Cypher), Redis, ETL, Docker, FastAPI, Streamlit
- Cloud/MLOps: AWS, Kubernetes (EKS), Istio, Grafana, k6
- Stats & Modeling: A/B testing, causal inference, experimental design, regression, time series (ARIMA/SARIMA), panel data; pandas, NumPy, RStudio
- Distributed: Spark, PySpark, Databricks, Hadoop/HDFS, MapReduce
- Tools/OS: Git/GitHub, Jupyter, VS Code; Linux, Bash, Windows
- Language: Spanish (native)

## PROJECTS

**AI-Powered Legal Citation Analyzer (Wolters Kluwer Client Capstone)**, AI Engineer

*August 2025 – December 2025*

- Delivered a dual-mode Streamlit legal research MVP (case lookup + chatbot) in a 14-week client capstone with weekly stakeholders; built a Neo4j knowledge graph (3.5K cases, 5.5K citations) via Python ETL from PDFs + CourtListener, including opinion chunking and Mistral-7B case summaries.
- Achieved 67% accuracy / 70% precision on citation treatment (pos/neg/neutral) using a 3-LLM ensemble (Claude 3.5, Mistral-7B, Llama 3-70B) with majority vote + prompt design, generating rationales for each citation.
- Built an explainable case-law labeling system: a time-weighted scoring model (recency/jurisdiction weights) producing Good/Bad/Moderate labels, plus an agentic GraphRAG chatbot in LangGraph (8 tools) using Neo4j queries + vector search/embeddings for grounded answers.

**Full End-to-End ML API Deployment on AWS EKS (MIDS)**, MLOps Engineer

*May 2025 – August 2025*

- Deployed a DistilBERT sentiment API with FastAPI/Docker/Redis at 70–100 req/s, zero downtime; preloaded 1GB weights to avoid cold starts; 40% cost cut with 95%+ cache hit rate.
- Ran production-style scaling on AWS EKS with p99 < 2s using HPA (1–70 pods), Istio routing, health checks, and k6 load tests (28K+ requests, 0 5xx) tracked in Grafana.

**Flight Delay Prediction at Scale (MIDS)**, Machine Learning Engineer

*January 2025 – May 2025*

- Built a distributed Spark/Databricks pipeline (PySpark; 5–10 workers) on 28M flight records to predict delay class; best MLP F1 = 54.6% (vs logistic regression/random forest).
- Implemented blocked time-series CV to prevent leakage with fold-specific transforms + imbalance handling (over/under-sampling, SMOTE); used parallel grid search for 3× faster tuning.
- Improved F1 by ~6% across 9 configs via probability recalibration to correct oversampling bias and improve class recall balance.

## EXPERIENCE

**Yorke Engineering, LLC**, Los Angeles, CA

Environmental Engineer Consultant I

*September 2023 – present*

- Reduced client compliance costs by \$100K+ annually by leading air quality compliance work for 100+ facilities across multiple industries.
- Supported a regulator-approved CY2023 GHG report revision saving \$5M and reducing reported emissions by ~134K metric tons by auditing boiler data in Excel and delivering regulator-ready analysis with stakeholders.
- Improved delivery efficiency 30% by using AI tools (Claude, ChatGPT, Perplexity) to speed drafting, checks, and analysis from scoping through final reports.