

Cameron Kettler

camjamesk@gmail.com | (832)-776-1598 | www.linkedin.com/in/cameronkettler | <https://github.com/cameronkettler> | <https://cameronkettler.github.io/>

EDUCATION

The University of Texas at Austin - B.S. Electrical and Computer Engineering

May 2025

SKILLS

Languages: Python, Java, JavaScript, C/C++, SQL, R

Frameworks & Tools: React, Next.js, FastAPI, Flask, LangGraph, SQLAlchemy, Playwright, GitHub Actions

Cloud & Data: AWS (ECS, SQS, S3, DynamoDB), PostgreSQL, Supabase, Snowflake, MongoDB, dbt, Docker, Datadog

EXPERIENCE

PMG | AI & Software Engineer I | Dallas, TX

June 2025 - Present

- One of two engineers leading development of Ask Alli, PMG's natural-language analytics assistant, owning its backend architecture and AI workflows while contributing to frontend development for a platform used weekly by 25% of users
- Led the migration of Ask Alli from a monolithic backend to a modular LangGraph orchestration architecture with intent-based routing, reducing latency, improving query quality, and providing run-level observability for faster debugging and iteration
- Built and deployed a reusable AI insights service on AWS (ECS, SQS, S3) with REST APIs that power automated insight generation across multiple PMG applications
- Designed AI-powered insight-generation nodes that automatically surfaced anomalies, trends, top-performing segments, and cross-dimensional relationships across analytics datasets, reducing manual analysis for media and analytics teams
- Modernized the Ask Alli frontend using PMG's internal design system, introducing real-time workflow status updates that improved visibility into long-running AI workflows
- Built the organization's first automated API smoke-testing framework using Playwright and GitHub Actions, proactively detecting breaking API changes before production deployments and improving release reliability

Paylocity | Software Engineering Intern | Chicago, IL (Remote)

May 2024 - August 2024

- Designed and implemented incremental Snowflake loading pipelines, eliminating full reloads of millions of records and reducing refresh times by approximately 70%
- Built dbt data pipelines to migrate raw S3 data into Snowflake dimensional models, reducing manual analyst data preparation by approximately 40% while improving reporting accuracy

Advantest | Software Engineering Intern | Austin, TX

May 2023 - August 2023

- Developed a Python application using the MinHash algorithm to identify and remove duplicate Java files, reducing data redundancy and improving AI model training efficiency by 20%
- Automated SFTP data transfers with Python, ensuring reliable delivery of datasets for machine learning model training

PROJECTS

Congress for Normal People | Full-Stack AI Application | [Live](#) | [GitHub](#)

June 2026 - Present

- Built and deployed a production AI application that helps users understand how federal legislation affects them by generating personalized, source-backed reports using congressional, campaign finance, lobbying, and demographic data
- Engineered LangGraph workflows that retrieve bills, aggregate context from Congress.gov, OpenFEC, Lobbying Disclosure, Census, and web search APIs, generate structured LLM reports, and stream results to users in real time
- Implemented authenticated multi-user workflows with personalized legislative monitoring, profile-based district resolution, and report caching to reduce repeated LLM and external API calls
- Designed and deployed the application's production infrastructure across Next.js, FastAPI, PostgreSQL, Supabase, Vercel, Render, and GitHub Actions with CI/CD, automated testing, report caching, and resilient external API integrations

University of Texas Center for Identity | PrivacyCheck

August 2024 - May 2025

- Developed PrivacyCheck, a Chrome extension that scores website privacy policies against GDPR and user-control metrics, helping users make informed data decisions
- Built and integrated a large language model (GPT-4o mini) backend to automate scoring, improving GDPR accuracy by 13% and User Control accuracy by 6.6% compared to legacy models
- Engineered a breach notification system using AWS Lambda and DynamoDB to aggregate/query public breach repositories, alerting users to historical security incidents tied to visited domains