

# Joshua Gomes

[gomesjoshua@gmail.com](mailto:gomesjoshua@gmail.com) | <https://www.linkedin.com/in/codewitty> | <https://joshuagomes.vercel.app/>  
Seattle, WA | (408) 891-6172

## WORK EXPERIENCE

### Sr Full-Stack Engineer | Aurelian, Seattle, WA

May 2024–Present

*Conversational AI automating non-emergency calls, triaging, and creating dispatch reports in real-time.*

- Spearheaded the development of a conversational AI system leveraging Azure's Text-to-Speech (TTS) and Speech-to-Text (STT) services, integrated with LLMs to autonomously process non-emergency calls. Focused on dispatcher needs, reducing workload by 80%, and allowing prioritization of critical 911 emergencies
- Architected a GIS-based backend system using Python, FastAPI, and Azure to create and manage custom jurisdictional boundaries, enabling dynamic triaging of non-emergency calls based on jurisdiction-specific law enforcement protocols. Implemented an optimized data processing pipeline to translate geographic boundaries into queryable jurisdiction rules, stored in PostgreSQL, allowing for seamless real-time adjustments and fine-grained, location-based call routing
- Built an LLM prompt testing dashboard featuring version control, performance metrics tracking, and experimentation capabilities for different models, prompts, and parameters to fine-tune responses. Achieved a 40% reduction in response latency, significantly enhancing prompt quality and consistency
- Designed and developed a real-time dashboard using Python and FastAPI for backend APIs, and React, TypeScript, and Next.js for the front end. Integrated features like call recording, transcripts, translation, and caller insights, providing dispatchers with a comprehensive view of all call information. Enabled dispatchers to verify caller details, access transcripts, and take action with a single click, streamlining workflows and significantly enhancing efficiency. Reduced issue resolution time by 70% and improved overall system response accuracy by 15%, fostering high stakeholder trust in system reliability

### Full-Stack Engineer | Codex Health, Palo Alto, CA

Sep. 2022–May 2024

*Revolutionizing Cardiometabolic healthcare with AI-enabled Virtual Care Management*

- Implemented multitenancy within GCP's serverless architecture, reducing infrastructure costs by 30% and enhancing scalability. Leveraged Terraform IAC and a Terragrunt wrapper for streamlined resource provisioning along with Golang, cutting deployment time by 60% and bolstering system maintainability. These optimizations improved efficiency and cost-effectiveness while accommodating diverse customer requirements
- Utilized fine-tuning techniques on healthcare-specific data to generate SOAP notes using an OpenAI Large Language Model (LLM), reducing clinician time spent on documentation tasks by approximately 40%. Resulted in an annual time saving of 5,000 hours across clinical staff, enhancing operational efficiency and patient care
- Achieved a 15% decrease in bug incidents during new feature rollouts by implementing feature flags using Launchdarkly for targeted feature rollouts and A/B testing

### Software Engineer, Big Data | Affinity Solutions, San Jose, CA

Aug. 2019–Aug. 2022

*Telling impactful customer stories with Big Data.*

- Led a project to develop and deploy a DSAR CCPA compliance framework on the Identity Engine, leveraging PySpark, Hadoop, and Hive for processing millions of records, ensuring company-wide data privacy compliance
- Migrated the DSAR ETL pipeline to AWS using Hive on EMR, S3, Glue, and RedShift, streamlining data processing and storage and improving performance
- Configured and managed Apache Airflow on AWS EC2 instances to automate and orchestrate data workflows, enabling scheduling, monitoring, and alerting of data pipeline tasks, resulting in a 50% reduction in manual labor for pipeline management

## TECHNICAL SKILLS

- **Programming Languages:** Python, Golang, Java, C++, C, TypeScript
- **Cloud Platforms & Data Management:** GCP, AWS, Azure, Spark, Airflow, Hadoop, Hive, SQL, NoSQL
- **Web Development:** HTML5, CSS3, React, Next.js, Node.js, FastAPI, REST, Bootstrap
- **Infrastructure & Monitoring:** Linux/Unix, Shell Scripting, Terraform, Datadog, Tableau, GitHub CI/CD, Asana, Jira, Bitbucket, Pytest, Mixpanel

## EDUCATION

University of California, Irvine - Irvine, CA  
Bachelor of Science in Computer Science