

David Chang

Senior Software Engineer | AI/ML Engineer | Founding Engineer

david.chang.ca@outlook.com <https://www.linkedin.com/in/david-chang-o> Playa Vista, CA

SUMMARY

Senior Full Stack and AI Engineer with 12 years of experience designing and scaling distributed backend systems, AI/ML pipelines, and workflow orchestration platforms using Python, FastAPI, Node.js, Kubernetes, and AWS.

I've led development of production ML infrastructure at Google, and spent 6+ years in healthcare at Zing Health and Ambience Health building EHR integrations (FHIR/HL7), medical coding automation (ICD-10/CPT), and clinical workflow systems that reduce administrative burden and improve care delivery, most recently architecting AI workflow platforms at Salt AI.

I'm now seeking a role in a fast-moving healthcare or startup environment where I can apply deep technical expertise and product thinking to build reliable, high-impact systems.

EXPERIENCE

Senior Full Stack Engineer

01/2025 - 03/2026

[Salt AI](#)

- Built and scaled the node runtime and orchestration layer for AI workflows using Python, FastAPI, Celery, and Redis, enabling reliable execution, retries, and dependency scheduling across distributed model pipelines
- Designed and implemented workflow deployment infrastructure using Docker, Kubernetes, and AWS, converting visual pipelines into production-ready APIs and webhooks while improving deployment reliability and environment isolation
- Developed a connector framework integrating enterprise systems such as Slack, Notion, Google Drive, and PostgreSQL, standardizing data ingestion pipelines and reducing integration time for new connectors by over 60%
- Architected secure private-cloud and on-prem deployments using AWS VPC, IAM, and Kubernetes, enabling enterprise customers to run AI workflows within isolated environments while maintaining strict data governance and compliance requirements
- Built execution services and platform APIs using Python, Django/FastAPI, and GraphQL, enabling scalable workflow orchestration, asynchronous task processing, and external integrations for production AI pipelines
- Led improvements to workflow reliability and performance, implementing idempotent job execution, retry policies, dead-letter queues, and monitoring with Prometheus/Grafana, significantly reducing pipeline failures and improving runtime observability

Senior Software Engineer

08/2021 - 12/2024

[Ambience Health](#)

- Designed and implemented medical coding automation services using Python, FastAPI, and transformer-based NLP models to extract ICD-10 and CPT codes from clinical notes, improving coding accuracy and reducing manual billing review effort
- Built patient chart intelligence pipelines using LLM-based summarization, Python microservices, and PostgreSQL to generate structured patient summaries from EHR records, reducing physician chart review time and improving clinical decision context
- Developed clinical workflow automation services with Node.js, Kafka event streams, and distributed workers to orchestrate referrals, lab orders, and discharge documentation, significantly reducing administrative workload across clinical operations
- Engineered scalable EHR integration infrastructure using FHIR and HL7 APIs, TypeScript backend services, and secure data pipelines to enable real-time synchronization with Epic and Cerner systems while maintaining HIPAA-compliant data processing
- Improved platform reliability by implementing idempotent workflow execution, retry strategies, and distributed tracing using Kubernetes, Redis queues, and Open Telemetry, enhancing system observability and reducing incident resolution time across AI workflow services

Senior Software Engineer

05/2018 - 08/2021

[Zing Health](#)

- Built and maintained full-stack features for the Member Portal using React, Node.js, and PostgreSQL, enabling members to view benefits, check claims, access EOBS, and find providers, improving digital self-service adoption
- Designed and implemented REST APIs and backend services for the Care Coordination & Case Management platform, enabling nurses and care managers to manage care plans, tasks, and social determinant workflows more efficiently
- Developed data pipelines and analytics services using Python, SQL, and cloud infrastructure to support the Population Health platform, helping identify high-risk members and enabling proactive care interventions
- Collaborated with product managers, clinicians, and operations teams to translate complex Medicare Advantage workflows into scalable platform services, improving system reliability and accelerating delivery of new member features
- Optimized database queries and backend services, implemented monitoring and automated tests, and improved API performance, helping the platform support growing member enrollment while maintaining reliable access to critical healthcare data

EXPERIENCE

Senior Software Engineer

02/2015 - 04/2018

[Google](#)

- Led backend development of scalable machine learning pipelines using TensorFlow and distributed data processing tools to support recommendation and computer vision models used across several internal Google AI services
- Designed and implemented production-grade ML infrastructure for training and deployment of speech recognition and vision models, improving inference latency and reliability through optimized data pipelines and GPU-accelerated model execution
- Collaborated with research scientists to productionize experimental ML models into scalable services, building reusable APIs and monitoring systems that enabled teams to deploy and iterate on AI features safely
- Mentored junior engineers and contributed to architecture decisions for distributed model training workflows, gaining deep experience in large-scale machine learning systems, data engineering, and production AI reliability practices
- Built and optimized machine learning training pipelines using TensorFlow, supporting large-scale experimentation in speech recognition and recommendation models while improving distributed training performance and dataset preprocessing workflows
- Implemented backend services and data processing components that enabled researchers to train and evaluate computer vision models at scale using internal Google infrastructure and distributed storage systems

Junior Software Engineer

06/2013 - 02/2015

[Google](#)

- Contributed to internal developer productivity tools used by engineering teams, improving build automation, debugging workflows, and development environments through backend services and UI improvements
- Implemented features and performance optimizations for internal tooling that supported large codebases, learning best practices in scalable software architecture, code quality standards, and collaborative engineering workflows

SKILLS

Backend:

Python, FastAPI, Django, Node.js, Express.js, REST APIs, GraphQL, Microservices, Distributed Systems, Asynchronous Task Processing, Celery, Kafka

Frontend: React, TypeScript, JavaScript, HTML, CSS

Database: PostgreSQL, SQL, Redis

AI / Machine Learning:

TensorFlow, LLM Applications, NLP, Transformer Models, Medical Coding Automation, Speech Recognition, Computer Vision, Recommendation Systems, Model Training Pipelines, ML Infrastructure, AI Workflow Orchestration

EDUCATION

Bachelor's Degree in Computer Science

Illinois Institute of Technology

Chicago, USA

09/2009 - 05/2013