Bhavik Mehta

Seattle, USA mehtabhavik168@gmail.com linkedin.com/bhavikmehta1101 github.com/bhavik168 bhavikmehta.dev

Professional Summary

Software Engineer with 2 years building and shipping full-stack and AI-powered products across React, TypeScript, Node.js, Python, and AWS. Built multi-agent LLM pipelines, RAG systems, and async backend services, owning features end-to-end from design through post-launch on production PostgreSQL and Redis backends. Wired OpenAI and Amazon Bedrock APIs into retry-safe production workflows and collaborate closely with cross-functional product peers. Finishing an MS in CS at Seattle University, looking for full-stack, backend, or AI engineering roles.

Education

Seattle University **Sep 2024 - Jun 2026**
Masters, Computer Science *Seattle, USA*

Savitribai Phule Pune University **Jun 2019 - Aug 2022**
Bachelor of Engineering, Information Technology *Pune, India*

Technical Skills

Languages: Python, TypeScript, JavaScript, Java, Go, SQL
AI/Agents: OpenAI, Amazon Bedrock, LLMs, RAG, embeddings, vector databases, multi-agent pipelines
Backend: Node.js, FastAPI, Django, REST APIs, WebSocket, GraphQL, microservices
Frontend: React, Next.js, Angular, Redux, Tailwind
Cloud/Data: AWS, Lambda, Docker, CI/CD, GitHub Actions, PostgreSQL, Redis, DynamoDB

Professional Experience

Software Engineer | **Amdocs (AT&T)** **Sep 2022 – Jul 2024**

- Developed billing system features and backend functionalities on AT&T's enterprise platform in **Python** and **Java**, directly serving internal employees and affecting millions in monthly revenue.
- Owned 20+ billing applications across three US regions, independently managing feature releases, production stability, and cross-region deployments end-to-end.
- Diagnosed runtime and service-level failures across **Linux** environments by analyzing application logs, process behavior, and inter-service dependencies during high-severity production incidents.
- Automated **Python** and Java regression frameworks integrated into **Jenkins CI/CD** pipelines, catching 30% more defects before release and enabling faster, more consistent deployments.

Software Developer Intern | **Virtual Math Labs** **Aug 2021 – Dec 2021**

- Built modular backend components in **Python** and **JavaScript** using OOP and structured data modeling for a live exam platform serving 5,000+ students, cutting defect density 20% across testing cycles.
- Authored an end-to-end randomized question-generation pipeline in **Python** producing up to 50,000 unique exam questions per run, delivered in milliseconds under concurrent load.
- Provisioned and stabilized Linux backend services under real student load using shell scripting and containerized **Docker** workflows, improving staging-to-production parity.

Projects

CodeLens | github.com/CodeLens **May 2025 – Jun 2025**

- Built a multi-agent AI code review platform on Microsoft Semantic Kernel running four specialist **agents** concurrently via Task. When all, each returning a structured **JSON** output contract that merges into one scored review.
- Engineered dual execution modes behind a shared interface, calling **Azure OpenAI** via Semantic Kernel with exponential-backoff retry on 429/5xx, plus offline heuristic agents requiring zero cloud dependencies, switchable via one DI change.
- Shipped an **Angular** dashboard with live polling and trend charts over a **Redis**-cached **API**, surfacing team quality scores and weekly z-score regression alerts; GitHub Actions tested every push.

ARIA - Autonomous Response Intelligent Assistant | github.com/ARIA **May 2026**

- Built ARIA, a real-time AI co-pilot, on a six-agent **Amazon Bedrock** pipeline coordinated by Claude Sonnet 4, firing specialist agents asynchronously on keyword detection and surfacing one recommendation card in under 15 seconds; won 1st place at AWSHacks 2026.
- Grounded every agent output in a **RAG** layer over FEMA, AHA, and MPDS protocols using **Bedrock Knowledge Base** with Titan **Embeddings** v2 on an OpenSearch Serverless vector index, with Guardrails blocking unverified recommendations.
- Owned the project end-to-end on a **WebSocket**-driven **React** frontend and **Python** backend streaming each transcribed word in under 300ms, defined as infrastructure-as-code in AWS CDK with GitHub Actions auto-deploy.

peekwise.ai | peekwise.ai **Mar 2026 – Apr 2026**

- Built a competitive-intelligence **SaaS** on **Next.js** and **FastAPI** that enriches scraped reviews into sentiment, themes, and frustration scores via an **LLM**, decoupling ingestion, AI enrichment, and orchestration through an async queue.
- Sequenced lightweight and frontier **LLM** models across a nightly **inference** pipeline, processing roughly 2,075 reviews per day at zero **API** cost by sizing batches around free-tier rate limits with zero silent data loss.
- Separated hot analysis state in **Redis** from static config in **PostgreSQL**, removing real-time AI calls from the request path and keeping every dashboard read sub-millisecond regardless of competitor count.

Achievements & Awards

- **1st Place at AWSHacks 2026 for ARIA** (Amazon Bedrock Track) & **2nd Place at HackWithSeattle 2026** for Homebase
- **Open Source Contributor, Bruno, 2026** - Merged bugfix (42k+ stars)
- **1,000+ Organic Visits** - bhavikmehta.dev