

Ahmad Abuomar

📍 New York, NY ✉ ahmad@abuomar.org ☎ (814) 689-9442 🔗 [linkedin.com/in/ahmadabuomar](https://www.linkedin.com/in/ahmadabuomar)

Work Experience

Amazon , Senior Software Engineer	Seattle, WA / New York, NY Oct 2020 – Jan 2026
• Built Scala/Spark data-quality guardrails on EMR Serverless for a multi-agent LLM product enrichment platform, cutting catalog failures from 90% to 40%, increasing attribute acceptance 50%, and supporting 10x throughput growth • Led end-to-end integration across 8+ engineering teams, resolving distributed ownership and dependencies to ship the product-data transformation capability • Cut deployment time 70% (2 hours to 36 minutes) and CPU usage 50% through distributed-systems tuning, including removal of 20.7M stale database entries and reducing host data loads from 22 minutes to 4 seconds • Migrated a legacy platform from 50M+ data mappings to ~8k (99.98%) with zero production incidents and cut downstream API traffic 97% (790 to 19 TPS) • Built and operated a catalog data normalization service processing 10-25% of automated product-matching traffic across multiple ingestion and enrichment workflows • Reduced off-hours maintenance from 200 to 10 hours/year (95%) through reliability and operational improvements	
Amazon , Software Engineer	Seattle, WA July 2014 – Oct 2020
• Designed and launched a product-data auto-normalization engine (DynamoDB and Elasticsearch) serving 15-20k requests/sec with ~11ms added p50 latency; increased normalization 25%, cut propagation from days to ~5 seconds, and eliminated manual operations with zero Sev-2 incidents after launch • Built data-classification infrastructure that reduced transformation-rule authoring from 10 hours to 30 minutes and supported 1,000+ annual transformations • Developed RESTful services handling 50k+ TPS with >99.999% availability and reduced infrastructure costs 80% through L2 caching and optimized data distribution • Built a product-variation sorting mechanism later adopted platform-wide and co-created a multi-parent product-classification hierarchy used across the catalog • Built a shadow fleet and model-endpoint instrumentation to measure deep-learning latency in a high-volume ingestion pipeline; performance dashboards were adopted by the ML team	

Projects

MotoWave
AI-powered motorcycle trip planning app for iOS 18+ (SwiftUI, Swift 6, SwiftData, Gemini API); shipped to TestFlight with automated fastlane + GitHub Actions CI

Technical Skills

Languages: Scala, Java, Python, TypeScript, Swift, SQL
Backend & Data: Distributed systems, Apache Spark 3.5, EMR Serverless, Large-scale ETL, RESTful services
Cloud: AWS (Lambda, S3, DynamoDB, Elasticsearch, CloudWatch, SNS/SQS), CDK v2
AI: LLM output validation at scale; AI-assisted development workflows (Claude Code, Codex, Gemini)

Education

MS	Georgia Institute of Technology , Computer Science (OMSCS)	Aug 2024 - Expected May 2028
BS	Pennsylvania State University , Computer Science and Mathematics (dual degree, Systems Analysis Option)	Aug 2010 – May 2014