

Hamza Azeem

TECH LEAD | SOFTWARE ENGINEER

hamzaazeem1@gmail.com

azeem.ca

linkedin.com/in/hamzaazeem

650-426-8350

PROFESSIONAL EXPERIENCE

Amazon | Tech Lead

[Ads Supply Intelligence] - Scala • Java • Python • Spark • AWS

Dec. 2020 – Present

- Founding engineer/leader of 5-person Reach Engineering Data & Identity Platform team
- Built OTT Media Planning data pipeline processing Petabyte-scale data every day
- Designed org-wide Data Integrity platform to monitor 2PB+ daily across 20+ pipelines
- Spearheading scalable ML engineering innovations for Amazon Reach Forecasting
- Launched Amazon Ads' first Open Data API, adopted by top global HoldCos
- Built sub-200ms near real-time reach & frequency API from multi-TB Lucene index, powering core of Amazon DSP's Reach Forecasting abilities
- Secured investments from C-Suite leaders and engineered executive reporting tools

[Search ML Platform] - Scala • Java • Python • AWS

Aug. 2018 – Nov. 2020

- Drove fine-grained model training & analysis by creating Amazon's first low-latency, hourly-attributed dataset for shopping behaviour, impacting millions of customers
- Led and designed global GDPR automation project for Search, eliminating manual labor costs for servicing thousands of DSAR requests
- Established core Search marketplace datastore & automated onboarding system, cutting marketplace onboarding costs by over 70%
- Slashed petabyte-scale dataset costs by 90%: introduced Avro, optimized MapReduce logic, and migrated to AWS
- Mentored intern through shipping full-stack data debugging tool used by thousands of engineers/analysts across Amazon & attracted them back for full-time

[Visual Search (Intern)] - Python • Java • R

Aug. 2017 – Dec. 2017

- Developed Visual Search team's first machine-learned relevance ranking models for most frequent query-type in Amazon's camera search: barcodes
- Increased barcode match rate by 14% after deployment, exceeding S-Team goal

[Visual Search (Intern)] - Python • AWS

Aug. 2016 – Dec. 2016

- Created real-time pipeline reducing end-to-end delays from 3 days to 2-5 seconds
- Developed backend service for near real-time amendments of Amazon items from the Visual Search catalog through barcode search-keys

UKG | Software Engineer Intern | Toronto, Canada

[TouchBase Platform] - Java • ELK Stack • SQL • Android

Jan. 2016 – Apr. 2016

- Created real-time analytics platform for entire Time Clock platform web/mobile systems
- Allowed streaming of 100k+ logs on a single cluster under 1s search times

Oracle | Software Engineer Intern | Toronto, Canada

[SalesTools Platform] - Java • JavaScript • Chrome API

May. 2015 – Aug. 2015

- Invented & shipped new public-facing Oracle sales app under 3 months
- Presented app at OpenWorld Las Vegas with 2,100+ active users since product launch

LANGUAGES

- Java
- Python
- Scala
- SQL
- TypeScript

TECHNOLOGIES

- AWS
 - EMR/Glue/Sfn
 - EC2/IAM/Cfn
 - SQS/SNS + more
- Spark
- Lucene
- Git
- *nix

EDUCATION

University of Waterloo

B.A.Sc in Computer Engineering

2013 – 2018

Relevant Courses

- Fuzzy Inferencing Systems & AI
- Distributed Systems
- Operating Systems
- Compilers
- Databases
- Computer Networks

INTERESTS

- Travel
- Cooking
- Writing
- Fashion