

HARSHIT KUMAR GUPTA

Mobile: +1-8582678960

Amazon Alias: gphars@

E-Mail: harshit.knit@gmail.com

LinkedIn: in.linkedin.com/in/harshit786

GitHub: github.com/harshitkgupta

Portfolio: harshitgupta.info

PROFESSIONAL SUMMARY

- Have worked on building scalable and responsible AI Applications using ML models and evaluating them on various KPI(Key Performance Indicator), then deploying inference endpoints as microservices.
- I have expertise in Java and Python, Spring, PostgreSQL, DynamoDB and Redis databases.
- I also have experience in Hadoop Ecosystem, Hive, Spark, Containerization(Docker, Rancher), and AWS.
- I have good hands on Machine Learning Algorithms, Libraries and Generative AI(LLM, Agent, MCP etc)
- Pursued M.Tech in Computer Technology from Indian Institute of Technology, Delhi with 8.43/10 CGPA.
- My M.Tech thesis involves work for classification of Autism Spectrum Disorder using ML models.
- I have a total 10+ years of experience in designing Scalable systems with deep dive and innovation.

Professional Experience

Amazon	Senior Software Engineer	since Sep 2020
--------	--------------------------	----------------

1. Project: BSOR (BUSINESS SYSTEM OF RECORD) - Gen AI AGENT

Entry point and orchestrator for all Gen AI use cases- New event onboarding, Simplified Onboarding Migration for already onboarded events, E2E testing across multiple Services, Customer Asks and Operations

Roles and Responsibilities:

- Building unified knowledgebase and Retrieval system using RAG
- Creating MCP Gateway service for existing Coral based Services(4+)
- Creating MCP server for creating prompts and tools for SO Migration and E2E testing etc.
- Deploying Gen AI Agents using Strands in AgentCore

Technologies: LLM, Agent with (Memory, Observability and Evaluation), MCP servers, Strands, Agent Core

Impact: Reduced SO migration onboarding time from 20+ weeks to less than 2 weeks per event

2. Project: Transactional Deferral Anomaly Detection

real-time anomaly detection system to detect incorrect deferral type(computed by SPYDER withholding engine) for Selling partner financial transactions

Roles and Responsibilities:

- Designing guardrail to check incorrectly released or deferred transactions.
- Utilized event features from Seller Event Aggregator and Deferral Policies
- Built end-to-end ML pipeline with SageMaker training, inference endpoints

Technologies: Python, SageMaker, Lambda, Andes, Iceberg, Step Functions, Glue, and Auto ML(AutoGluon)

Impact: Prevention of ~530MM delayed transactions to keep customer trust

Detection of ~5MM prematurely released transaction causing less cash flow for Amazon

3. Project : Smart Charging Time Recommendation System

ML-powered charge timing optimization system for Amazon Seller Debt Manager (SDM) to improve debt recovery rates and charge success rates through intelligent retry scheduling

Roles and Responsibilities:

- Designed ensemble classifier model to predict optimal charge retry times for failed transactions
- Utilized temporal and regional features to build model for increasing charge success rate
- Built end-to-end ML pipeline with SageMaker training, inference endpoints

Technologies: Python, SageMaker, Lambda, Ensemble models, Herd Workflows, Step Functions, Glue and AppConfig

Impact: Improved charge success rates by 410 bps for scheduled charge failures and 390 bps for Partial charge failures.

4. Project : Disbursement Gatekeeper

Intelligent Request Gating and Anomaly Detection of Seller Disbursement

Roles and Responsibilities:

- Designing of scalable system predicting anomaly detection for 300K Monthly Disbursement

- Training of machine learning model for predicting anomalous score base on past pattern
- Building Bookkeeper to keep track of failure at various indexes
- Automated Logic to cut tickets to downstream for any identified issue

Technologies: Java, Lambda, Glue, Step Function, Herd, Redis, Dynamodb, Hubble, Redshift, SageMaker, Logistic Regression, NeuralProphet, Top2Vec, FuzzyWuzzy

Impact: Intelligent circuit breakers that achieved a 99.75% success rate across \$327B in annual payouts.

5. Project : SPURSH

Aggregates transaction details into journal entries that are posted to the General Ledger

Roles and Responsibilities:

- Designing of scalable system for JournalPostmanService and JournalPostmanPreprocessor
- Removing Journal status publishing and ticketing dependency from OFA
- Log Aggregation, Monitoring and Alarming of NAWS based services using CloudWatch
- Extended SNS and HTML Email Notification Improvement in SPURSH Outbound
- JournalPostmanService and JournalPostmanPreprocessor Migration to Native AWS

Technologies: Java, Fargate, Batch, OFA, Docker, Oracle, CDK

Loyalty Juggernaut

Product Engineering Architect

since June 2019

Project : Gravy (Cloud Platform for Customer Engagement and Loyalty Management)

Designing offers and campaigns for customer attraction with cashback and points management

Customer targeting using on fly offer creation and dynamic rules

Building Predictive and analytics model for customer, location retain and product, offer recommendation.

Roles and Responsibilities:

- Designing of scalable transaction layer for realtime and batch processing
- Designing Drools based execution engine to execute custom rules created from blockly
- Designing Incremental Data Processing Pipeline in Data Lake
- Creating various batch and analytics processing pipelines in AWS.
- Created features for KPI based limit for offer execution.
- Member segmentation and product targeting for different segments

Technologies: Java, Spring Boot, DynamoDB, Lambda, Fargate, Drools, Blockly, Apache Hudi, PySpark, AWS

CognitiveScale

Senior Software Development Engineer

Nov 2017-June 2019

1. Project: Debt Risk Hotspots ML Model

Discover patterns of features of accounts that are conducive to bad debt.

Role and Responsibilities:

- Interpreting model to understand similar bad debt accounts
- Defining similarities of accounts quantitatively by using feature importance.
- Stratification to understand the characteristics of clusters
- Defining metrics of clusters for quantitative and qualitative features.

Technologies: Java, Python, scikit learn, Apache Livy, Spark, CatBoost, SHAP

2. Project: Patient Scheduling Constraint Optimizer

Recommend ideal patient schedules based on patient preferences and optimal use of resources.

Role and Responsibilities:

- Building a Patient Scheduling model to efficiently schedule medical appointments
- Using patient preferences, medical rules, time slots and facility availability as constraints
- Solving complex combinatorial optimization using OR-Tools

Technologies: Java, Spring boot, Python, Pandas, OR-Tools, Docker

3. Project: Bad Debt Risk Advisor ML Model

Building predictive model to determine bad debt for Hospital Billing and Professional Billing

Role & Responsibilities:

- Building a model for predicting bad debt accounts

- Developing KPI for measuring model performance
- Building jobs for distributed data processing in Hive Using Spark
- Using CatBoost for scaling and improving the accuracy of the model

Technologies: Java, Python, pandas, scikit learn, Apache Livy, Spark, CatBoost

4. Project: Predictive User Engagement on Tax Filing System

Improve user engagement and increase Tax Filing by using the ML model.

Role & Responsibilities:

- Building model to handle historical data and click through data both
- Building a predictive model for user intervention
- Improving model response time in Production to handle peak season load
- Designing KPI for measuring ML model performance

Technologies: Java, Python, XgBoost, MongoDB, docker, AWS

Cadence Design System	Senior Software Development Engineer	Feb 2017-Oct 2017
------------------------------	---	--------------------------

Project: Allegro EDM Solutions (Collaborative Library and design data management)

Role & Responsibilities:

- Involved in development of EDM (Enterprise Data Management)
- Designing DAO layer to support SQL, NoSQL and Graph Databases
- Done POC on designing SQLg connector for Tinker Pop to support MySQL
- Hands on with Design Pattern, REST, MVC

Technologies: J2SE, Swing, Tinker Pop, SQLg, MySQL, Neo4j

Intel Security	Senior Software Development Engineer	July 2015-Feb 2017
-----------------------	---	---------------------------

Project: Application Control and Change Control with Global Threat Intelligence

Role & Responsibilities:

- Drove several features on my own and contributed in their end to end delivery
- Lead feature to support SHA-256 for Rule Groups and Policy Discovery
- Done POC's for various feature improvements (like Advanced Filter)
- Hands on with Design Pattern, REST, MVC

Technologies: J2EE, MFS(Spring based Framework), JavaScript, MS SQL

Academics & Value Addition			
Year	Degree / Exam	Institute	GPA/Marks
2013-15	M.Tech in Computer Technology	Indian Institute of Technology, Delhi	8.43/10
2009-13	B.Tech in Computer Science	Kamla Nehru Institute of Technology, Sultanpur	81.16%
2007-09	Intermediate (U.P. Board)	Onkar Nath Dhawan Inter College, Kanpur	83.20%
2007	High School (U.P. Board)	Saraswati Vidya Mandir Inter College, Hamirpur	80.17%

Academic Awards

- 3rd prize winner in Mind-Hunters event, National Level Techfest , Effluence
- 2nd prize winner in Fill Up The Code event, Tech Carnival By Computer Society Of India
- Consolation prize in Paper Presentation on "WEB 3.0", organized by I.E.I.
- Consolation prize in Technical Wordsworth event organized by Computer Society of India
- State Level project on Water Resource Conservation in National Children's Science Congress

M.Tech Thesis

Title - Assessment of Autism Spectrum Disorder in Toddlers using Speech Features

- **Supervisor:** Dr. Santanu Chaudhury, Dept. of Electrical Engg, IIT Delhi
- **Description:** Designing an Android App to collect voice sample and store in the cloud
 - Analysis of Speech Samples using Spectrogram and Scalogram
 - Feature Extraction using Discrete Wavelet Transform and Discrete Wavelet Packet Analysis
 - Classification of speech samples using SVM, Random Forest, HMM, CNN classifiers
 - Application of Deep Learning Convolutional Neural Network for Feature Learning and Classification
- **Technologies:** Python, numpy, scipy, scikit-learn, scipy, caffe, pylearn

Academic Projects:

- Image Encryption & Decryption and Transformation: Secure & Compressed image transfer
- Aayush : Web platform for interaction of Doctors and Patients using J2EE with Struts Framework
- Weather Forecasting Application: Using Yahoo Weather API and support for Offline Queries
- Sentimental Analysis for Facebook posts: Sentiment Detection and Analysis of posts
- Illuminance Correction: Android App for Illuminance Correction on Image using OpenCV for Android

Certifications

- Algorithms: Design and Analysis, Part 1, Stanford University, Coursera
- Algorithms: Design and Analysis, Part 2, Stanford University, Coursera
- Image and video processing, Duke University, Coursera
- Machine Learning, Stanford University, Coursera