

Suhas K M

(571)-237-1142 | suhaskm23@gmail.com | linkedin.com/in/suhaskm | github.com/suhas-km | [suhas.page](#)

PROFESSIONAL SUMMARY

Versatile engineer with 2 years of experience in AI research, software development, databases, and cloud infrastructure. Expertise in delivering software solutions across domains. Using cutting-edge technologies to solve complex challenges

EDUCATION

Northeastern University <i>Master of Science in Information Systems, GPA: 3.80</i>	September 2023 - May 2025 Arlington, VA
Visvesvaraya Technological University <i>Bachelors of Engineering in Computer Science</i>	August 2018 – July 2022 Bangalore, India

EXPERIENCE

Research Assistant - LLM Alignment, RLHF with Verifiable Rewards <i>Northeastern University</i>	January 2025 – May 2025 Arlington, VA
<ul style="list-style-type: none">Architected two adaptive attention neural networks, AILA(s) cut training time by 30%, cost by 20%, and improved long range sequence accuracies by 15% against baselines SOTA baselines Transformers, ViTs and LSTMsLSTMs v/s AILA on price forecasting, reduced MSE scores for USD from 15.55→2.85 and wheat from 5245→1536Ablated and pruned AILA to shrink model size 20% and accelerate inference 25% with negligible accuracy lossTrained AILA-2 to IMDB, achieving 84.3% accuracy with significantly faster training and inferenceDeveloped REALM, an augmented reward-model for RLHF that injects explicit relevance signals, boosting reward ranking accuracy by 4% and cutting off-topic generations	
Teaching Assistant - Data Science <i>Northeastern University</i>	January 2025 – May 2025 Arlington, VA
<ul style="list-style-type: none">Designed assignments, helping students think critically about data and math intuition, boosting real-world skillsHeld 40+ teaching hours and 15+ lab sessions on regression, classification, and clustering techniquesGuided projects on hyperparameter tuning, overfitting prevention, balancing bias-variance, and evaluation metricsCrafted containerization templates helping students master Docker and follow MLOps practices	
Cloud Engineer <i>Safecast Global</i>	May 2024 – January 2025 Arlington, VA
<ul style="list-style-type: none">Volunteered to design and document AWS architecture, leading to a significant boost in open-source contributionsOptimized AWS costs by analyzing resource utilization, achieving a 15% reduction in monthly cloud expensesMonitored and upgraded PostgreSQL, ensuring RDS versions stayed compliant with AWS support timelines	
Software Engineer <i>Runway Proptech LLC</i>	September 2022 – August 2023 Bangalore, India
<ul style="list-style-type: none">Led a Python-based AWS Serverless Architecture, cutting API costs by 30% and enhancing API scalabilityDeveloped 15+ Spring Boot/MVC and cloud serverless APIs, improving accessibility and ease of useAccelerated ETL pipeline for Salesforce CRM data migration, boosting speed by 35% with a fail-safe designImproved query performance for ACID compliance, reducing latency and enhancing data securityResolved CORS issues and conducted code analysis to improve API security and ensure compliance with standardsBoosted Agile and Scrum project efficiency by 20% through strict ticket handling and accelerated task completionEnhanced client retention and reduced churn by 15% with AWS SES for emails triggered based on user activityIncreased operational efficiency by 30% through a Tableau dashboard, enhancing data-driven decision-making	

TECHNICAL SKILLS

Programming Languages: Java, Python, SQL, R, Go, Shell, Object-Oriented Programming, Backend, Design Patterns
Frameworks/Tools: Spring, MVC, Microservices, React, FastAPI, Flask, Apache Spark, Kafka
Databases: PostgreSQL, MySQL, MongoDB, NoSQL, Oracle Database, Elasticsearch, Redis
Data Science/Machine Learning: PyTorch, scikit-learn, NLP, MLOps, Data Pipelines, Data Analytics, UNIX
Cloud Computing/DevOps: AWS, Google Cloud Platform (GCP), Azure, Serverless (AWS Lambda), Jenkins, GitHub Actions, CI/CD Pipelines, Oracle, Docker, Kubernetes, Terraform, Prometheus, Grafana, ELK Stack
Other Tools: Git, Linux, Tableau, PowerBI, Agile, Scrum, Data Structures, Algorithms, Project Management

LICENSES & CERTIFICATIONS

AWS Certified Cloud Practitioner (CLF-C02) | *Amazon Web Services*

RESEARCH

Waste Management in Urban Localities: ML & IoT | Springer | DOI: 10.1007/s13198-025-02872-8 July 2025

- Developed a low-cost “smart-bin” integrating IoT sensors with a CNN classifier, achieving **83 %** waste-type accuracy and cutting collection mileage by **25 %**

Adaptive Integrated Layered Attention (AILA) | *arXiv* | [arXiv:2503.22742](https://arxiv.org/abs/2503.22742) May 2025

- Cut **S&P500** forecasting MSE by **52%** and IMDB train time by **71%** versus baselines with competitive results

PROJECTS

REALM – Reward Enhanced Alignment Learning Method | *RLHF, PPO/DPO* |  May 2025

- Built open-source **REALM**, a modular RLHF pipeline fusing reward scores with semantic similarity, boosting alignment to **95 %** RewardBench accuracy and higher TruthfulQA truthfulness

SparseRec: Memory-Optimized Recommender Engine | *Markov Chains, Social Graphs, MLOps* |  April 2025


- Engineered **SparseRec**, a memory-optimized Markov-chain recommender with social-graph boosts, delivering real-time next-item prediction on 50 k-item sparse catalogs, tracked end-to-end with MLflow-based MLOps

AlignAI – Microservice | *FastAPI, Next.js, Docker, PostgreSQL, GitHub Actions* |  December 2024

- Containerized a scalable policy & compliance AI Agentic service handling 10k RPS at <200 ms p95 latency; built with async REST endpoints, CI/CD pipelines, and zero-downtime rolling deploys, sustaining 99.9 % uptime

MatchPoint: The Resume Analyst | *React, FastAPI, AWS ECS, Docker, Llama Model* |  December 2024


- Developed a resume scoring app using React, FastAPI, and Llama models, with Dockerized deployment on AWS ECS for scalability, providing real-time NLP-based feedback and improvement suggestions

Safecast AWS Infrastructure Revamp | *AWS, IaC, Observability* |  2024


- Built AWS IaC stack (EC2-S3, RDS, Beanstalk) for Safecast, enabling autoscaling Go APIs, Grafana → Slack observability, and cost-optimized radiation data ingestion & visualization

D.C SafeHomeSeeker AI - aiXplain | *Generative AI, Vector Databases, Cloud, Maps* |  | [Blog](#) October 2024


- Won **hackathon** by developing an AI-powered app that aggregates housing data and generates property descriptions using Pinecone and OpenAI, with interactive maps displaying safety and amenities

Nutrition Tracker Database System | *SQL Server, Tableau, TSQL, Data Security* |  April 2024


- Developed a Nutrition Tracking System with SQL Server, Tableau dashboards, and secure handling of user-specific dietary and fitness data

word2vec Modelling | *NLP, GENSIM, PyTorch, TensorFlow, CBOW, Skip-Gram* |  April 2024


- Developed Word2Vec models with GENSIM, PyTorch, and TensorFlow, optimizing CBOW for analogy tests and cosine similarity while enhancing scalability for large datasets under constrained resources

QueryMaster-AI | *LLM, Prompt Engineering, LangChain, SQLite, Gemini Pro, Docker* |  January 2024

- Built an AI-powered app integrating Google’s Gemini Pro, achieving over 90% accuracy in translating natural language into SQL queries, enhancing data accessibility for non-technical users

Hotel PricePro | *Forecasting, Analytics, XGBoost, EDA, Feature Engineering* |  December 2023

- Built a price prediction model for hotel ADR using XGBoost, EDA, and hyperparameter tuning, achieving 71% accuracy and improving performance through feature selection techniques

Children’s DayCare Management System | *SpringBoot, Java, JDBC, ORM, JavaxMail* |  December 2023

- Built a backend web app to track vaccine schedules for children, integrating JavaxMail for notifications and ensuring HIPAA-compliant data sharing in the VFC program

LEADERSHIP, VOLUNTEERSHIP & EXTRACURRICULAR ACTIVITIES

Head of Student Developer Club and Soccer Team Captain

Visvesvaraya Technological University

May 2019 – August 2022

- Captained the soccer team to regional finals, enhancing teamwork and leadership skills
- Organized community service initiatives, including coding workshops

Community Service Volunteer

Bangalore Rural Development Organization - BRDO

January 2018 – May 2023

- Organized literacy programs, benefiting over 50 underprivileged students in rural areas
- Conducted computer literacy workshops focusing on MS Office and basic programming concepts