

Vedanth Prasanna Bharadwaj

(623) 210-8716 | vedanthpb1@gmail.com | linkedin.com/in/vedanth-p-bharadwaj | github.com/vedanthpbharadwaj

SUMMARY

Backend and cloud-focused software engineer with a Master's in Data Science & Engineering (ASU, GPA: 3.9) and hands-on production experience building serverless APIs, database systems, and ML pipelines on AWS. Currently building backend infrastructure end-to-end at MindSpark. Seeking full-time SWE roles in backend, cloud, or data engineering. Work-authorized on F1-OPT with STEM OPT extension eligibility.

EDUCATION

Arizona State University

Master of Science, Data Science, Analytics & Engineering – GPA: 3.9/4.0

Tempe, AZ

Jan 2024 – Dec 2025

B.N.M. Institute of Technology, Visvesvaraya Technological University

Bachelor of Engineering, Information Science & Engineering

Bengaluru, India

Aug 2018 – Jul 2022

TECHNICAL SKILLS

Languages: Python, Java, SQL (MySQL, PostgreSQL), JavaScript, HTML, CSS, R, C++

Frameworks & Libraries: Spring Boot, Angular, TensorFlow, Keras, Scikit-learn, Pandas, NumPy, OpenCV, MediaPipe

Cloud & Architecture: AWS (Lambda, API Gateway, S3, Aurora PostgreSQL, DynamoDB, Step Functions, CloudWatch), Serverless Architecture, Microservices, REST API Design, ETL Pipelines

AI/ML & Analytics: XGBoost, LSTM, CNN, Random Forest, Regression, Time-Series Forecasting, Feature Engineering, Prompt Engineering, Power BI, Tableau

Tools & Practices: Git, GitHub, OOP & Design Patterns, Data Structures & Algorithms, Agile/Scrum, System Design

EXPERIENCE

AWS Cloud Developer Intern, MindSpark, Mesa, AZ

Feb 2026 – Present

- Developed the core database gateway for the Kairos platform, enabling dynamic SQL operations with type-safe parameters and foreign-key-safe execution across **6 tables**; enforced referential integrity across **21 tables** with **zero data-integrity violations** in production
- Designed for reliability from the ground up: all multi-table writes execute in a single transaction with automatic rollback on failure, partial updates preserve existing data without overwriting, and custom retry logic handles database cold starts with stepped backoff
- Diagnosed and resolved a **25–30s cold-start latency** spike on a production Lambda — root-caused to eager config loading and inefficient SQL joins; refactored to lazy initialization and optimized queries, cutting response time to **150–200ms (~95% reduction)**
- Managed staged deployment pipeline across dev, staging, and prod using Lambda versioning and aliases; authored a **5-chapter engineering certification course** with graded assessments and production runbooks for full team knowledge transfer
- Leveraged Claude Code throughout backend development to analyze SQL logic, validate API payload structures, and surface edge cases in dynamic query generation — accelerating debugging cycles

Software Engineer Trainee, Cognizant, Bengaluru, India

Nov 2022 – Jul 2023

- Completed 8-month enterprise Java Full Stack training program (Spring Boot, Angular, RESTful APIs, microservices, AWS); maintained **95% average** across 20+ assessments in SQL, OOP design patterns, and cloud architecture
- Built 3 capstone projects simulating enterprise healthcare systems — covering microservices decomposition, relational schema design, REST API contracts, and end-to-end integration testing to professional development standards

PROJECTS

Urban Mobility Data Pipeline | Python, AWS Lambda, S3, DynamoDB, XGBoost, Power BI, SQL

Aug 2025 – Dec 2025

- Led a 5-person team to build a serverless AWS pipeline processing **41M+ NYC taxi records**; automated data validation and anomaly detection, improving data quality to **96%** and reducing manual effort by **80%**
- Built **58+ features** using time, location, and weather data to improve model accuracy; trained an XGBoost model achieving **R² = 0.90 (30% improvement)** over baseline
- Developed a Power BI dashboard tracking **12+ KPIs** across 265 city zones to identify peak demand and support fleet allocation decisions

Energy Consumption Forecasting | Python, TensorFlow, LSTM, Scikit-learn, Tableau

Jan 2024 – May 2024

- Led a 4-person team to build an energy demand forecasting system on **1M+ hourly records**; trained LSTM and Random Forest models achieving **97.8% accuracy (18% improvement)** over baseline
- Built an end-to-end ML pipeline covering preprocessing, feature generation, model evaluation, and anomaly detection; developed Tableau dashboards projecting **12% cost savings (\$50K/year)**

Human Emotion Recognition System | Python, TensorFlow, CNN, Bi-LSTM, OpenCV

Jan 2024 – May 2024

- Led a 5-person team to build a real-time emotion recognition system using video and audio data; achieved **86.3% accuracy (12% improvement)** over baseline
- Built a live inference pipeline using computer vision and audio processing; optimized for low-latency real-time performance in healthcare and HCI use cases

Indian Sign Language Recognition – Published | Python, OpenCV, MediaPipe, LSTM, TensorFlow

Oct 2021 – Jun 2022

- Built a real-time sign language interpreter using computer vision and sequence modeling; achieved **97% accuracy** and published in *IJRTI* (Vol. 7, Issue 6, 2022)

CERTIFICATIONS

Google Data Analytics Professional Certificate (SQL, Tableau, R, Data Cleaning)

Jun 2025

Microsoft: Data Analysis and Visualization with Power BI

Nov 2025