# Harita Parikh

626-366-5605 | haritaparikh050902@gmail.com | linkedin.com/in/haritaparikh | github.com/haritaparikh05 | haritaparikh.me

## EDUCATION

**California State University - Los Angeles** 08/2024 - 05/2026
**Master of Science - Computer Science** **GPA: 4.0/4.0**
Coursework: Advanced AI, Database Systems, Cloud Computing, Software Architecture, Operating Systems, Computer Networks

**Charotar University of Science and Technology, India** 10/2020 - 05/2024
**Bachelor of Technology - Information Technology** **GPA: 3.92/4.0**
Coursework: Data Structures & Algorithms, Machine Learning, Computer Vision, Full Stack Development, Object-Oriented Design

## TECHNICAL SKILLS

**Languages:** Python, Java, C/C++, JavaScript, TypeScript, SQL, R
**AI/ML:** Deep Learning, Natural Language Processing (NLP), BERT, Model Deployment, MLOps, A/B Testing
**Full Stack & Systems:** React.js, Node.js, Spring Boot, Flask, REST APIs, Microservices, JWT/OAuth, Testing (JUnit, PyTest)
**ML Frameworks:** PyTorch, TensorFlow, Scikit-learn, Pandas, NumPy, SciPy, NLTK, PySpark
**Data Engineering:** ETL/ELT, Feature Engineering, Data Modeling, Apache Spark, Databricks, Snowflake, Tableau, Power BI
**Databases:** PostgreSQL, MySQL, MongoDB, Elasticsearch
**Cloud & DevOps:** AWS (EC2, S3, Lambda), GCP, Azure, Docker, Kubernetes, CI/CD (GitHub Actions, Jenkins), Git, Linux

## EXPERIENCE

**Graduate Research Assistant |** ***California State University - Los Angeles*** 09/2024 - 05/2026
- Processed 4,536 open codes from 175 interview transcripts by developing ATA-LLM, a framework combining 1,536-dimensional LLM embeddings, UMAP/DenseMAP dimensionality reduction, and HDBSCAN clustering, with full PyTest coverage.
- Achieved 0.893 topic coherence and 77% cosine-similarity alignment with human-coded themes across 200+ Optuna Bayesian optimization trials, outperforming t-SNE baselines (0.817 TC) and confirming results via statistical significance testing.

**Data Science Intern |** ***Collabera Digital - Gujarat, India*** 01/2024 - 05/2024
- Delivered 83% accuracy in automated skill extraction by fine-tuning and deploying BERT-based NLP models (PyTest-covered) via production REST APIs, reducing manual screening effort by 40% across large-scale recruitment workflows.
- Improved candidate-job matching precision by 25%, validated through A/B and statistical significance testing, by building GPT-based prompt engineering pipelines and curating labeled training datasets for scalable, automated talent mapping.

**Android Developer Intern |** ***iTeam Technology - Gujarat, India*** 05/2023 - 07/2023
- Reduced backend response latency by 35% by building and verifying production-ready Android features using Java, JUnit/Espresso, RESTful APIs, and MySQL, enabling real-time data interactions for concurrent users.
- Strengthened data consistency by 20% by designing normalized relational schemas and optimizing SQL queries to support real-time inventory tracking and transaction processing.

## PROJECTS

**Neuz Now - AI-Powered News Aggregation Platform -** ***React Native, Node.js, MySQL, LLMs, JWT/OAuth***
- Enabled secure, role-based access control across 3 user roles (Reader, Author, Admin) by architecting a news aggregation platform using React Native, Node.js, and MySQL with JWT authentication and RBAC.
- Decreased reading time by 60% by integrating Google Generative AI into a Node.js/Express.js pipeline to auto-summarize news articles, boosting engagement through personalized content delivery.

**Agentic RPA Platform -** ***Claude/Gemini Tool Use, Playwright, Pydantic, Flask***
- Eliminated 100% of LLM inference cost and latency on repeat runs by architecting an agentic system that separates LLM-driven discovery from a deterministic, model-free replay engine, reused across tenants via a versioned override layer.
- Enabled safe operation on an IT-helpdesk console via a policy engine with allowlisting, risk classification, and PII redaction, plus an agent-facing catalog with a stability gate catching 2 real bugs, validated by 35 tests.

**Point Cloud Processing with Point Transformer -** ***PyTorch, Deep Learning, Multi-GPU, 3D Vision***
- Achieved 92.40% classification accuracy (3.46% loss) on the ModelNet10 benchmark (4,899 shapes, 10 categories) by implementing a Point Transformer with self-attention and positional encoding.
- Optimized spatial feature retention across scales by engineering a 4-layer offset-attention U-Net decoder with 5 downsampling and 5 upsampling stages, benchmarked against PointNet, voxel-based, and projection-based architectures.

## PUBLICATIONS

- "Exploring Manifold-Based Clustering Techniques for Enhanced Inductive Thematic Analysis" - Springer Nature
- "Comprehensive Study and Analysis of Point Transformer for Point Cloud Data" - Journal of Propulsion Technology
- "Number Plate Detection and Recognition using OpenCV" - IEEE Xplore