

MITESH ADAKE

Los Angeles, CA, USA | +1 (901) 651-7876 | adake@usc.edu | [linkedin.com/in/mitesh-adake](https://www.linkedin.com/in/mitesh-adake) | mitadake.github.io/

EDUCATION

University of Southern California

Master's, Computer Science (GPA: 3.66/ 4.0)

- **Coursework:** Applied Natural Language Processing, Advanced Computer Vision, Adversarial and Trustworthy Foundation Models

Aug 2024 - May 2026

Los Angeles, CA, USA

Pune Institute of Computer Technology

Bachelor's, Computer Engineering (GPA: 3.84/ 4.0)

- **Coursework:** Data Structures & Algorithms, Machine Learning, Natural Language Processing, Deep Learning

Jul 2019 - May 2023

Pune, MH, India

SKILLS

- **Programming Languages:** Python, Java
- **Databases / Query Languages:** MongoDB, MySQL, Oracle (SQL)
- **Cloud Platforms:** Amazon Web Services, Google Cloud Platform
- **AI / Machine Learning:** Neural Networks, Natural Language Processing, Large Language Models, PyTorch, Tensorflow, Keras
- **Tools & Frameworks:** Ab Initio, Dataiku, Docker, Flask, FastAPI, Django, Power BI

PROFESSIONAL EXPERIENCE

FedEx | *Revenue Management - Intern*

Sep 2025 - Present

- Built 3 automated systems leveraging data engineering to automate pricing analyst workflows, cutting processing time from ~5 days to ~1 hour
- Partnered with cross-functional teams on PowerBI to create dashboards and AI proof-of-concepts for data-driven pricing systems

Revenue Management - Summer Intern

Jun 2025 - Aug 2025

- Designed and developed an automated system using Python, SQL, Streamlit, and Docker to process user inputs and perform complex, business-critical database operations
- Reduced manual processing time from ~10 days to ~10 minutes to optimize workflow logic, significantly increasing operational efficiency

Software Engineer

Aug 2023 - Jul 2024

- Migrated Dataiku data pipeline to Ab Initio (ETL) data pipeline for Harmonized Code search engine, resulting in annual cost savings of \$ 1M for FedEx and accelerating data processing time
- Executed an optimization strategy in Python for Natural Language Processing (NLP) code, significantly reducing data discrepancies from 25% to negligible levels and eliminating data shuffling from 75% of text data

Persistent Systems | *Machine Learning Intern*

Jun 2022 - Aug 2022

- Implemented NLP-based question deduplication on company's community portal, improving efficiency and reducing duplicate postings to 10%
- Leveraged Python, Flask API, FAST API, and Docker to create a scalable and robust solution

NICE | *Data Engineer Intern*

Nov 2021 - Apr 2022

- Built scalable, serverless real-time file ingestion using AWS (S3, DynamoDB, Lambda) and MongoDB for efficient file processing
- Streamlined data management tasks by implementing a user-friendly Python Flask interface

PROJECTS & OUTSIDE EXPERIENCE

TokenGen: Visualize Token Prediction Evolution & Attention | [Link to project](#)

Mar 2025 - Apr 2025

- Developed visualization tool to identify layer-wise token evolution patterns across 12+ transformer layers for GPT2 & OPT

PriceNet: Stock Price Prediction using Large Language Models

Sep 2024 - Nov 2024

- Designed and implemented PriceNet, using Llama and Gemini for explainable financial time series forecasting with accuracy of 50% and ROGUE-2 of 0.546 on reasoning of predictions

Hate Speech Unlearning for LLMs

Jan 2025 - Mar 2025

- Implemented Task Vector Arithmetic, Contrastive Learning, and Fine-tuning to remove hate speech from LLaMA 3 models, resulting in a drop of harmful rate from 62% to 34% on PKU-SafeRLHF dataset