

# Siddharth Padmakumar

spadmaku@andrew.cmu.edu | 412-863-9664 | [linkedin.com/in/siddharthpadmakumar/](https://www.linkedin.com/in/siddharthpadmakumar/) | [GitHub](#)

## EDUCATION

### Carnegie Mellon University (CMU)

Master of Science in Artificial Intelligence & Innovation (MSAI)

Relevant Coursework: Computer Systems, Machine Learning, AI Law & Future Markets, Deep Learning

Pittsburgh, Pennsylvania

May 2028

### PES University (PESU)

Bachelor of Technology in Computer Science and Engineering

Machine Intelligence and Data Science Specialization

GPA: 9.18 / 10.00

Bengaluru, India

May 2024

## HONORS

- 3rd Place, IBM InCITE 2024 Hackathon: production-ready MVP integration using Watsonx Orchestrate
- Best Overall Startup: PES University Centre for Innovation & Entrepreneurship (CIE)
- Dr. MRD Scholarship, PES University: Awarded to top 20% of Computer Science students

## SKILLS

**AI & Machine Learning:** TensorFlow, LLMs, RAG, GANs, GCNs, RNNs, Computer Vision, Differential Privacy, NLP

**Programming Languages:** Python, Java, C, C++, JavaScript

**Frameworks:** FastAPI, Flask, Docker, Spring Boot, REST APIs, WebSockets, React, Firebase, SQL, MongoDB, AWS

## EXPERIENCE

### IBM India Software Labs

Software Developer

Bengaluru, India

Jul 2024 – Dec 2025

- Proposed and engineered a knowledge base consolidating 300+ docs for a 70-member team, extending it into a vector database with embeddings for LLM-based RAG search, cutting retrieval time from ~1 day to seconds.
- Refactored internal enterprise Java tools, fixing memory leaks and stabilizing crashes, improving speed by 50%.
- Discovered and reported 20 critical and 50+ other OWASP security vulnerabilities in IBM software products, all remediated before release.

Software Developer Intern

Jan 2024 – Jun 2024

- Developed a differential-privacy toolkit for healthcare analytics (OpenDP SmartNoise, diffprivlib) with PII detection and synthetic data generation, earning Significant Contributor honors at IBM's 2024 Meraki Intern Projects.
- Implemented a holiday-management and leave-scheduling feature (Java/JSP, SQL) with full CRUD, a scheduler, and location-based access control, streamlining scheduling for 70+ users.

### Communityfarm

Software Intern

Bengaluru, India

Jun 2022 – Sep 2022

- Developed a Python order-rerouting system (Flask, React, MySQL) clustering drivers by K-means and Haversine distance, with CRUD and stored procedures, cutting manual effort from 5+ hours to under 3 minutes.

### PES University Centre of IoT

Summer Intern

Bengaluru, India

Jun 2022 – Aug 2022

- Trained YOLOv4 and R-CNN models for weed and crop detection and developed a React Native app integrated with Raspberry Pi for automated farmed control, with the live bed view updating in real time via Firestore listeners as the camera reports detections.

## PUBLICATIONS

[Statistically Inspired Discrepancy Detection for Anomalous Spatio-Temporal Graphs](#),

*International Journal of Data Science and Analytics (Springer)*, vol. 20, pp. 3267-3281, 2025.

[Extended Abstract](#) presented at [IEEE DSAA 2024](#), San Diego.

## PROJECTS

[Bill2Go](#): VLM based Receipt Parsing and Real-Time Collaborative Bill Splitting (<https://bill2go.app>)

- Built and deployed a React Native mobile app (Google Play) and a Render-hosted React web app + FastAPI backend for real-time bill splitting, using Gemini Vision to parse receipts and Firebase (Auth, Firestore) for no-login shared sessions, secured with OWASP-aligned rate limiting, input validation, and access-control checks.

**Audio Emotion Style Transfer using Generative Adversarial Networks** - Machine Intelligence @ PESU

- Applied a CycleGAN (TensorFlow) to learn the mapping between emotional speech styles (angry to sad) for audio style transfer. Built an audio-to-mel-spectrogram pipeline (librosa), adapted the model for single-channel input and reconstructed audio through Griffin-Lim inversion.

**Cross-Language Part-of-Speech Tagging for Telugu and Kannada languages** - Research Project @ PESU

- Built an 87M-token Telugu corpus from books and news data, then fine-tuned mBERT for part-of-speech tagging and transferred the model to Kannada, achieving 89% tagging accuracy.