

Siddharth Bayapureddy

+91 7075286119 | siddharthbayapureddy@gmail.com | siddharth-bayapureddy.vercel.app | linkedin.com/in/siddharth-bayapureddy | github.com/SiddharthBayapureddy

EDUCATION

BITS Pilani, Hyderabad Campus

Bachelor of Engineering in Computer Science - 9.13 CGPA

Hyderabad, TS

Aug. 2024 – July 2028

Aakash Institute

Intermediate Board - 96.5%

Hyderabad, TS

July. 2022 – May 2024

Genesis School

CBSE 10th Board - 96%

Hyderabad, TS

July. 2015 – May 2022

EXPERIENCE

ACM Member - GenAI & ML

Sep. 2025 – Present

BITS Pilani, Hyderabad Campus

- Designed a Face-Anonymizer using Flow-matching (lesser steps than diffusion) to achieve real-time masking with low latency
- Developed a Resume Screener application that helps students gain insights from their resume and maps them to the expected PS-II or Placement statistics.

Lynq Member - Agentic AI

Sep. 2025 – Present

BITS Pilani, Hyderabad Campus

- Built a fest-budget tracker that leverages the Agentic workflow to verify the purchases and bills.
- Integrated crawl4ai to autonomously scrape real-time market prices and cross-validate them against submitted receipts, flagging inconsistencies in the budget tracker.

Technical Team Member - Innovation Cell

Sep. 2024 – May 2025

BITS Pilani, Hyderabad Campus

- Led a backend development workshop on FastAPI and Python covering RESTful APIs, database integration, and deployments.
- Assisted in a Python Workshop and mentored students through hands-on-coding exercises.

PROJECTS

vorp | *Python, LiteLLM, RAG, VectorDB, Typer, Rich, PyPI*

- Developed a terminal-based autonomous AI agent enabling developers to manage files, query codebases through natural language commands
- Used LiteLLM to switch between multiple LLMs depending on the workflow, and implemented Chat persistence and conversation storage.
- Leveraged ChromaDB as a vector database to implement a RAG-powered agent.

SparseSynth | *Python, Stable Diffusion, LoRA, Meta-Learning, PyTorch, Transformers*

- Built a meta-learning pipeline that takes 30–40 reference images and synthesizes 500+ training samples to address low-data constraints.
- Designed a Transformer-based aggregator that encodes reference images into a unified representation and predicts task-specific LoRA weights, eliminating the need for per-concept fine-tuning.
- Injected generated LoRA adapters into a Stable Diffusion backbone to condition image generation on unseen concepts at inference time, enabling few-shot generalization.

GAN from scratch | *PyTorch, CNNs, DCGAN*

- Implemented a DCGAN from scratch in PyTorch, replicating the Generator-Discriminator adversarial architecture from the original research paper.
- Trained on the MNIST dataset, iteratively tuning hyperparameters to improve generated image quality.

TECHNICAL SKILLS

Languages: Python, Java, C, C++ , SQL (Postgres)

AI/ML Tools: PyTorch, LangChain, LangGraph, LiteLLM, ChromaDB, Stable Diffusion, LoRA, crawl4ai, scikit-learn, Hugging Face

Developer Tools: Git, GitHub, VS Code, Visual Studio, Supabase, Jupyter Notebooks, Postman, Docker