

Sahil Pardasani

814-441-1036 | sahil.pardasani@gmail.com | <https://www.linkedin.com/in/sahil-pardasani/>
<https://github.com/sahilpardasani>

EDUCATION

The Pennsylvania State University | University Park, PA
Bachelor of Science in Computer Science
Minor in Mathematics, Engineering Leadership Development (ELD)

December 2025
GPA 3.40/4.00

EXPERIENCE

GenAIx | Remote

June 2025-Aug 2025

AI Consultant Intern

Built proof of concept pipelines combining structured and unstructured data for property risk evaluation and improved efficiency by 20%

Collaborated with product and research teams to identify solution gaps, benchmark existing AI tools, and translate client needs into a technical architecture and implementation plan

Helped translate business requirements into technical data architecture plans and implement solutions

OSSIG Lab | University Park, PA

Oct 2023- May 2025

Research Assistant

Built and tested evaluation pipelines using OpenAI and Hugging Face API for hallucination, alignment drift detection

Tuned and evaluated prompts and hyperparameters using Optuna and learnt about concepts like Test Time Compute, alignment drift and documenting insights

Published paper "Bridging the Know-Act Gap via Task Level Autoregressive Reasoning" on arXiv (submitted COLM 2026)

PROJECTS

Research Method Generation Project

April 2025

Fine-tuned T5-small and BART to generate research methodologies.

Tuned hyperparameters using the Optuna library to maximise ROUGE-L and BERT score performance.

Conducted semantic evaluation and model comparison to measure faithfulness and coherence.

Stock Market Price Prediction Using LSTM

June 2025

Developed a deep learning model using LSTM networks using PyTorch to predict the next day closing stock price using the past 90-day history.

Used residual correction techniques to improve RMSE accuracy by 12%

Designed a streamlit dashboard with an earnings calendar sidebar; integrated real-time logging into PostgreSQL

Zypher Kavach

May 2025

Developed a secure Streamlit chat UI around the Zephyr LLM and created an adversarial prompt test suite

Added multi-layer prompt-injection defences blocking 10+ types of attacks using regex and Hugging Face Moderation API

Every prompt is prefixed with a system-level instruction header to enforce safe model behaviour.

Logged all user prompts, model responses, and "blocked" flags to PostgreSQL for real-time auditing and analytics.

arXiv Conference Paper Tracker

March 2026

Built an automated daily pipeline using Python, arXiv API to fetch, filter; analyse ML/AI papers accepted to top conferences.

Integrated Llama 3.1 for structured paper analysis (problem statement, key metrics, technical breakdown)

Deployed FastAPI backend on Render with Docker using APScheduler cron jobs and GitHub REST API auto-commits to persist data across deployments

Built a responsive dashboard with date and category filtering, light/dark mode and live day-over-day metrics.

TECHNICAL SKILLS

Languages: Python, Java

Libraries & Frameworks: PyTorch, Hugging Face Transformers, Streamlit, LangChain, FastAPI, Tensor Flow, LangGraph, Pandas, Numpy, Scikit-Learn, pytest, NLTK

Tools & Platforms: Docker, GitHub, Postman, PromptLayer, n8n, Crew AI, Jira, Agile Methodologies

Databases: PostgreSQL, Weaviate, SQL, Snowflake

AI & Machine Learning Domains: Generative AI, NLP, Safety Evaluation, Retrieval Augmented Generation (RAG), BERT, GPT

Cloud: Microsoft Azure

Prompting Techniques: Zero Shot, Few Shot, Chain of Thought, Tree of Thought, System Prompting