

Ahan Bhatt

✉ bhattahan@gmail.com | [in LinkedIn](#) | [GitHub](#) | [Portfolio](#)

EDUCATION

University of California, Santa Cruz [3.82 GPA]

Master of Science in Computer Science and Engineering

Santa Cruz, California

Aug. 2025 – Jun 2026

Indus Institute of Technology and Engineering [3.94 GPA]

Bachelor of Technology in Computer Engineering

Ahmedabad, India

Aug. 2020 – July 2024

TECHNICAL SKILLS

Languages: Python, Java, C/C++, SQL, JavaScript, HTML/CSS, R

Machine Learning: PyTorch, TensorFlow, Keras, Transformers, scikit-learn, NumPy, Pandas, Matplotlib

LLM and Data Tools: LangChain, ChromaDB, RAG, vector databases, PyPDF

Backend and Deployment: Node.js, FastAPI, Flask, REST APIs, Docker, Kubernetes, Linux, Git, CI/CD, Nginx

Cloud: AWS EC2, AWS S3, AWS IAM, AWS Lambda, SageMaker, CloudWatch

Data Engineering: PostgreSQL, Redis, ETL, SQL optimization, schema design, data cleaning, time-series preprocessing

Software Engineering: TDD, API integration, version control, DSA, System Design, Performance Optimization, Code Reviews, Debugging

EXPERIENCE

Researcher

September 2025 – June 2026

AI Explainability and Accountability (AIEA) Lab @ UCSC

Santa Cruz, California

- Built ReliaGuard Studio and BreakPoint with Next.js, TypeScript, FastAPI, PostgreSQL, Docker, and MLflow to detect over/underreliance in human-AI decisions and compile real LLM failure modes into adversarial eval suites using neuro-symbolic modeling, conformal gating, multi-judge validation, and explainable dashboards.

Data Scientist

June 2024 – May 2025

AI Institute of South Carolina

Columbia, South Carolina (remote)

- Conducted research on multimodal toxicity mitigation and anomaly detection, using PostgreSQL-backed experiment tracking to manage dataset metadata, model outputs, evaluation metrics, and benchmarking results for projects like DE-HATE and Time Series Foundation Model evaluation.

Data Scientist

Jan. 2024 – May. 2024

Space Applications Centre, ISRO

Ahmedabad, Gujarat

- Built an end-to-end LSTM satellite clock-bias prediction platform with Apache Kafka live streaming and a Next.js/React/TypeScript UI, automating preprocessing and anomaly filtering to improve navigation timing forecasts.

Software Engineer (AI)

May 2022 – Jan. 2024

Indus Institute of Technology

Ahmedabad, Gujarat

- Developed Node.js-backed LLM pipelines for a university website chatbot, automated knowledge graph generation, and retrieval-augmented medical QA, evaluating GPT-4, LLaMA-2, and BERT for structured semantic extraction and grounded reasoning.

PROJECTS

[CacheIR](#)

- A narrow, inspectable compiler and runtime for decoder-only transformer inference. It imports a Llama/Mistral/Qwen-style model, lowers it into its own IR, runs transformer-specific optimization passes, plans memory and execution, selects backend kernels, and executes inference through a reference runtime.

[BreakPoint](#)

- A failure-to-eval compiler for LLM systems that mines real failure modes into versioned adversarial regression suites with DSL-defined traps/rubrics, multi-judge calibration, RAG/tool simulators, CI gates, FastAPI/Next.js dashboards, and 360 validated eval cases across 9 failure families.

[ReliaGuard Studio](#)

- Built a Next.js/TypeScript/FastAPI AI observability platform with SDKs, guardrail APIs, review queues, and conformal neuro-symbolic models to detect and audit over/underreliance in human-AI decision workflows.

[Research Forge – A Self-Improving Research Agent](#)

- A universal self-improving research agent that autonomously discovers papers, generates hypotheses, runs experiments in a sandbox Python environment, and updates its reasoning strategy using long-term graph memory.

[Meeting AI Assistant](#)

- Built a real-time AI meeting copilot desktop app (Electron + React + Express) that captures system audio, transcribes conversations, and streams context-aware responses using OpenAI API with screenshot and document context support.

[Time Series Foundation Model Evaluation for Anomaly Detection](#)

- Benchmarked Time Series Foundation Models such as TimeGPT, Chronos, and Time-MOE across multiple datasets, showing that traditional statistical and deep learning models can outperform TSFMs for anomaly detection and prediction.

[Knowledge Graph Generation from Large Language Models](#)

[Med-Bot: Retrieval Augmented Medical Assistant](#)

PUBLICATIONS

[Leveraging LSTM for Predictive Modeling of Satellite Clock Bias](#)

- *Presented at [CDMA 2025 conference](#) and available on [IEEE Xplore](#)*

[Generating Knowledge Graphs from Large Language Models](#)

[Time Series Foundational Models: Their Role in Anomaly Detection and Prediction](#)

- *Presented at [AAAI 2025 Anomaly Detection in Scientific Domains Workshop](#)*

[Med-Bot: An AI-Powered Assistant to Provide Accurate and Reliable Medical Information](#)