

Siddhartha Venkatayogi

Austin, TX (Open to Remote or Relocation) | 737-977-5454 | sidvenkatayogi@utexas.edu
[linkedin.com/in/sidvenkatayogi](https://www.linkedin.com/in/sidvenkatayogi) | github.com/sidvenkatayogi | sidvenkatayogi.github.io

Education

The University of Texas at Austin

Austin, TX

Bachelor of Science, Computer Science; Intended Double Major in Mathematics

May 2029

Overall GPA: 3.92 | SAT: 1550

- **Coursework:** Computer Architecture, Data Structures, Discrete Math, Probability, Linear Algebra, Calculus III
- **Activities:** ECLAIR Robotics, Texas Luminescence, Texas Asian Business Student Association

Experience

Texas Luminescence

September 2025 – Present

Software Engineer

Austin, TX

- Designed a RESTful API with Flask to serve generative color palettes, integrating ChromaDB and LLM inference
- Executed an RLVR fine-tuning protocol with deterministic reward metrics for aesthetic output, using serverless LoRA fine-tuning on the Fireworks.ai platform, improving Qwen3-8B palette similarity scores by 13% over baseline
- Deployed a backend server on Render using Docker for containerization, serving a React + Vite frontend

UT Computational Linguistics Research Group

May 2025 – January 2026

Research Assistant

Austin, TX

- Implemented a RAG pipeline to test improvements in factual accuracy in LLM medical question answering tasks
- Designed systematic perturbation framework across 809 test cases finding LLMs fail to maintain safety guardrails when presented with manipulated medical evidence, reporting up to 80% failure rates on toxic perturbations
- Built automated evaluation pipelines with LLM-as-judge, achieving 90%+ accuracy on safety assessment tasks

Code2College

June 2025 – July 2025

Volunteer Instructor

Austin, TX

- Delivered an 8-week online Python curriculum to 30+ high school students with a 96% pass rate
- Supported students to grasp fundamental Python concepts like syntax, OOP, and Data Structures

Projects

Transformer Language Model from Scratch | *Natural Language Processing*

Transformer model for understanding and generating text

Python, PyTorch, Transformers

- Implemented multi-head self attention architecture, positional encodings, and chunking for seq2seq in PyTorch
- Visualized attention maps to analyze learned sequence representations and model interpretability patterns

PIXIE (PIcture eXploration and Inference Engine) | *Computer Vision, UI/UX*

Visual, color, and subject similarity image recommendation system

Python, Pillow, PyQt, FAISS

- Engineered a novel color extraction and distance metric algorithm with Pillow, leading to a 40% increase in user ratings for visual cohesiveness over RGB-Euclidean baseline
- Leveraged visual transformer models for multimodal vector embedding, with semantic and visual similarity
- Designed algorithms and a GUI for scalable hexagonal and circular image organization, enabling intuitive UX

FlightSense | *Blockchain, Predictive Analytics*

Decentralized flight delay prediction markets on Solana

FastAPI, Typescript, Next.js, Solana, PostgreSQL

- Architected decentralized prediction markets on the Solana blockchain using the Anchor framework, enabling users to trade YES/NO outcome shares on flight delays with real-time price signals.

Research

Faithfulness vs. Safety: Evaluating LLM Behavior Under Counterfactual Medical Evidence

[arXiv:2601.11886](https://arxiv.org/abs/2601.11886) (currently in review)

Technical Skills

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

ML/AI Frameworks: PyTorch, NumPy, Pandas, Scikit-learn, FAISS, ChromaDB

Tools & Technologies: Git, Matplotlib, Pillow, GCP, AWS, OpenAI API, Selenium, BeautifulSoup, SQLite