Siddharth Banerjee

+1 (555) - 555 - 5555 | sidbanerjee[at]berkeley.edu | https://sidbanerjee.me | linkedin.com/in/sidbanerjee00

Education

University of California, Berkeley — Graduating May 2027
B.S. Electrical Engineering & Computer Science — *GPA: 3.97*

Relevant Coursework: Data Structures and Algorithms, Machine Learning, Advanced Algorithms, Optimization Models in Engineering, Probability and Random Processes, Signals and Systems, Computer Architecture, Computer Security
Activities & Organizations: Teaching Assistant for Signals, Dynamics, & Information class, DiversaTech Consulting, Berkeley AI Research (BAIR), Center for New Music & Audio Technologies Research, Public School Records

Experience

Center for New Music and Audio Technologies (CNMAT) — January 2025 - Present
Undergraduate Machine Learning Researcher — *Berkeley, CA*

- Built a chord-aware evaluation judge for a generative jazz Transformer model by combining melody and harmony models, fitted across 73 artists and 178K notes. Judge predicts real solos better than melody alone (3.29 vs 3.36 bits/note) and showed the model's solos follow the chords about twice as closely as human players' solos.
- Designed a "justified risk" metric to score the model's improvised risks against human practice, calibrated on 8.9K human risks. It caught corrupted solos at AUC 0.84 and found the model's risk-taking matches human players' (AUC 0.52).

Apple — January - August 2026
Machine Learning Research Intern — *Cupertino, CA*

- Benchmarked SOTA audio quality metrics (Zimmos, SongEval, etc.) against perceived quality on audio enhancement model; engineered a DSP-based suppression detection algorithm to identify perceptually noticeable artifacts.
- Deployed an internal Python package and full-stack platform (React, Flask) for dataset-level metric analysis, weighing metric performance to guide targeted listening, cutting evaluation time by 80% across hundreds of samples.
- Created a synthetic dataset via acoustic room simulation; designed and trained a no-reference oversuppression detection model using band-split RoPE transformer architecture, achieving 97% test accuracy.

Amazon — May - August 2025
Software Development Engineer Intern — *Sunnyvale, CA*

- Created Cerberus, an AI agent that analyzes AdTech publisher-region business metrics, integrating codebases as retrieval-augmented generation (RAG) sources for enhanced debugging insights & recommendations.
- Engineered a multi-threaded metric aggregation and analysis pipeline across 80+ publishers, replacing manual data validation with automated analysis—reducing debugging time by 20% and preventing potential financial loss.
- Implemented an MCP (Model Context Protocol) client-server architecture to enable on-demand analysis of business metrics, reducing time needed to pull and assess metrics by 80%.

Berkeley Artificial Intelligence Research (BAIR) Lab — June 2025 - January 2026
Undergraduate NLP/Machine Learning Researcher — *Berkeley, CA*

- Built a streaming, causal, hierarchical framework for real-time full-duplex dialogue behavior prediction, outperforming baselines (e.g. F1 improvement of 40%) at 0.101s inference latency, 20x faster than comparable models.
- Created ConversationGoT-120h, a 120-hour causal-annotated benchmark dataset, and a Graph-of-Thoughts rationale module that beat GPT-4o on reasoning quality by 15% at 4x lower latency.

Projects

VegaWatcher | *yFinance, React, Next.js, Llama3, FinBert, SupaBase*

- Built full-stack application to track stock sentiment from social media sources, including an algorithm to calculate weighted implied volatility, with web-scraped data processed through fine-tuned LLMs and codified in an SQL database.
- Developed frontend UI/UX using React for real-time data visualization, featuring interactive tables and graphs that enable dynamic, user-friendly insights into stock trends and sentiment patterns across platforms.

Journalist Red Room | *React, Flask, Stripe, SupaBase*

- Agents collaborate to critique research, reports, and essays in a fullstack app for students and professionals.
- Developed a novel agent-agent-human orchestration scheme with unique agent roles (fact-checking, bias-checking, etc.).

Technical Skills

Languages/Frameworks: Python, Java, TypeScript, Golang, C, C++, C#, SQL, JavaScript, React, Next.js, Node.js, Vite, Flask, Django, HTML, CSS, LaTeX
Technologies: Git, CI/CD DevOps, AWS (CDK, Bedrock, Lambda, s3, IAM, Step Functions), Docker, SupaBase, Render, Pandas, PyTorch, Selenium, Max/MSP, Ableton, TouchDesigner