

AMIT NAMBURI

📍 La Jolla, San Diego | 📞 +1(858)319-6075 | ✉ anamburi@ucsd.edu

🌐 linkedin.com/in/amit-namburi | 🐙 github.com/namburiamit | 🎓 Google Scholar | 🌐 namburiamit.com

EDUCATION

University of California, San Diego

M.S. Computer Science

Mar 2025 - June 2026

GPA: 4.0/4.0

Relevant Coursework : Generative AI, Advanced Deep Learning, Machine Learning Systems, Large Scale Data Systems, Probabilistic Modeling, Optimization for Machine Learning Hardware, Advanced Algorithm Design

Graduate Teaching Assistant: Courses CSE 153/253 (ML for Music) and CSE 158/258 (ML for Recommender Systems and Web Mining)

B.S. Computer Science

Sep 2021 - Mar 2025

GPA: 3.7/4.0 **Awards/Honors:** Provost Honors

Relevant Coursework : Machine Learning, Computer Vision, Recommender Systems, Natural Language Processing, Operating Systems, Data Structures & Algorithms, Databases, Web Development, Computer Architecture

EXPERIENCE

Software Engineer Intern

Jun 2025 – Sep 2025

Apple Inc.

San Diego, CA

Engineered a **SoC analytics dashboard** with **FastAPI**, **Streamlit**, **Spark**, integrating **datalakes** for regression and performance monitoring. Optimized **data pipelines**, cutting multi-hour processing into **analysis completed within 4 minutes**.

Leveraged an **MCP server** combining deterministic **statistics** with regression context to power an **interactive chatbot**.

Developed **Audio LLMs-as-a-judge** with the **Siri team**, applying **LLM fine-tuning** for pronunciation detection and evaluation.

iContest Finalist (Top 10 of 750 interns) for innovative intern project proposal.

Tech Stack: Python, Spark, FastAPI, Streamlit, Splunk, Tableau, MCP Servers, Authentication, LLM Fine-tuning, Evaluation Pipelines

Graduate Research Assistant, Multimodal AI

Jan 2024 – Present

McAuley Lab, UC San Diego

San Diego, CA

Contributed to and **led research projects** on multimodal AI, with publications accepted at **ICASSP, EMNLP, and ISMIR**.

Designed **multimodal pipelines** integrating **audio, vision, and language** for retrieval, captioning, and recommendation.

Optimized **LoRA/QLoRA fine-tuning** for 7B+ LLM backbones on **multi-GPU clusters**, improving memory and throughput.

Exploring **RL policies for LLM data selection** and efficient **visual prompting** for multimodal grounding.

Tech Stack: PyTorch, Hugging Face, LangChain, LoRA/QLoRA, Docker, NVIDIA GPUs, RLHF, Visual Prompting

Data Engineer

Mar 2023 – Mar 2025

Fair Data Informatics Lab

San Diego, CA

Redesigned the **Foundry Dashboard**, a scalable **ETL platform** supporting biomedical data curation and analytics.

Engineered an automated **resource ingestion protocol** for AWS EC2 clients, enabling reliable batch and streaming workflows.

Built **RAG pipelines** with **LlamaParse, LlamaIndex, and NVIDIA NIM** to mine PubMed-scale literature for **antibody specificity**.

Tech Stack: Angular.js, AWS EC2, Python, LlamaIndex, NVIDIA NIM, ETL Systems

Software Engineer Intern

Jun 2023 – Sep 2023

Orektic Solutions LLP

Remote

Engineered automated **UI regression suites** with **TypeScript, Selenium, Cypress**, reducing test cycles by **70%**.

Integrated **Jira-driven CI/CD pipelines** to enforce quality gates and streamline **release deployments**.

Tech Stack: TypeScript, Selenium, Cypress, Jira, CI/CD Automation

Machine Learning Intern

Jul 2022 – Sep 2022

IPMD Inc. (Start-up)

Berkeley, CA

Built multimodal **ML pipelines** for **emotion analysis**, integrating audio, vision, and text signals in Project M.

Implemented **Node.js REST APIs** with **React/TypeScript** frontends to serve ML inference in real time.

Tech Stack: Python, Node.js, React, TypeScript, ML toolkits

PROJECTS

SixthSense - Full Stack Fortune Prediction Web App

Spring 2023

UCSD Team Project

Co-led a **10-member team** to build a scalable full-stack web app, following Agile development practices.

Developed a **modular CRUD architecture** with persistent data storage and responsive UI.

Implemented automated **CI/CD pipelines** using JUnit and Puppeteer for reliable testing and deployment.

TECHNICAL SKILLS

Languages: Python (expert), C/C++, Java, JavaScript/TypeScript, SQL, ARM Assembly, R

ML/AI: PyTorch, TensorFlow, Hugging Face, Transformers, LangChain, LlamaIndex, LoRA/QLoRA, Keras

Data/Systems: Spark, Apache Airflow, Hadoop, Dask, CUDA, REST APIs

Cloud/DevOps: AWS (EC2/S3, SageMaker), Docker, Kubernetes, Linux, Git

Web/Apps: FastAPI, Streamlit, React, Angular, Node.js, Express, Firebase, Django, D3.js

MusiCRS: Benchmarking Audio-Centric Conversational Recommendation

Rohan Surana*, **Amit Namburi***, Gagan Mundada*, Abhay Lal*, Zachary Novack, Julian McAuley, Junda Wu.
arXiv preprint, September 2025. [arXiv]

Introduces a benchmark for conversational recommendation systems that integrate audio grounding for better context modeling.

WildScore: Benchmarking MLLMs in-the-Wild Symbolic Music Reasoning

Gagan Mundada, Yash Vishe, **Amit Namburi**, Xin Xu, Zachary Novack, Julian McAuley, Junda Wu.
Accepted at EMNLP 2025. [arXiv]

Evaluates multimodal LLMs for symbolic music reasoning under noisy, real-world distribution shifts.

FUTGA-MIR: Enhancing Fine-grained and Temporally-aware Music Understanding with MIR

Junda Wu, Zachary Novack, **Amit Namburi**, Hao-Wen Dong, Carol Chen, Jiaheng Dai, Julian McAuley.
Accepted at ICASSP 2025. [IEEE Xplore]

Proposes temporally aware MIR augmentations that improve fine-grained music retrieval and captioning.

CoLLAP: Contrastive Long-form Language-Audio Pretraining with Musical Temporal Structure Augmentation

Junda Wu, Warren Li, Zachary Novack, **Amit Namburi**, Carol Chen, Julian McAuley.
Accepted at ICASSP 2025. [IEEE Xplore]

Introduces contrastive pretraining strategies for long-form audio-text alignment using temporal structure cues.