

# Shrestha Mishra

(201) 253-9327 | [shresthamishra76@gmail.com](mailto:shresthamishra76@gmail.com) | [linkedin.com/in/shrestha-mishra07](https://www.linkedin.com/in/shrestha-mishra07) | [shresthamishra.vercel.app](https://shresthamishra.vercel.app)

## EDUCATION

### The University of Texas at Austin

Expected May 2028

Bachelor of Science, Computer Science and Mathematics

GPA: 3.93/4.0

**Relevant Coursework:** Computer Organization and Architecture, Data Structures and Algorithms, Machine Learning Text Analysis, Multivariable Calculus, Linear Algebra, Discrete Mathematics

**Activities and Societies:** Texas Convergent, Longhorn Developers, Texas Luminescence, Hook 'Em Hacks

## TECHNICAL SKILLS

**Programming Languages:** Python, JavaScript/TypeScript, Java, C, C++, SQL

**Frameworks & Libraries:** Next.js, Tailwind CSS, Node.js, Express, Flask, PyTorch, TensorFlow, Pandas, NumPy

**Developer Tools:** Git, Vercel, Supabase, Docker, AWS, GCP, Terraform, Jest

**AI Tooling:** Model Context Protocol (MCP), LLM orchestration, prompt engineering, agentic workflows, tool calling

## EXPERIENCE

### TANTV Studios

October 2025 – December 2025

Software Engineering Intern

Remote

- Developed core infrastructure for the **SyndexAI** platform, architecting a full-featured publisher dashboard across **2** Next.js applications with scalable UI components serving **500+** publishers
- Reduced page load times by **40%+** by implementing asynchronous API orchestration and intelligent caching layers for real-time content delivery
- Streamlined publisher onboarding by **60%** through secure OAuth authentication, AI-generated content summaries, and event-driven analytics pipelines

### Foot and Ankle Research Innovation Lab at Harvard Medical School

May 2024 – Present

Machine Learning Researcher, Lab Member

Boston, MA

- Implemented machine learning algorithms for anatomical landmark detection on **3000+** foot and ankle radiographs, leveraging CNNs and mechanisms to improve detection accuracy
- Achieved **94% Successful Detection Rate** by building and optimizing end-to-end ML training pipelines with PyTorch, integrating preprocessing, augmentation, and evaluation workflows
- Reduced surgical review time by **30%** and increased surgeon confidence by collaborating with **5+** surgical experts to implement **Grad-CAM** interpretability visualizations, translating model outputs into actionable insights

### AI Edge Lab at Harvard SEAS

May 2024 – July 2024

Prompt Engineer

Boston, MA

- Collaborated with research team to develop LLM for **text prompt-based chip architecture design** using a **Dockerized Label Studio** environment for data annotation and prompt parsing
- Reduced API costs by optimizing prompt tokenization with **NLTK** while establishing standardized labeling protocols across **1,550+** training examples for chip architecture design LLM

## PROJECTS

### Olympus (MCP-Powered Cloud Infrastructure) | Next.js, Flask, Terraform, Docker, AWS

November 2025

- Built platform leveraging NLP to drive **7 Terraform MCP tools** for automated AWS infrastructure provisioning
- Integrated NVIDIA NemoTron to deploy **3 AWS services (S3, EC2, Lambda)** through structured orchestration
- Dockerized Terraform runtime with **3 per-resource mutex locks** preventing concurrent state corruption

### BiteSwipe | Next.js, Tailwind CSS, Supabase, Yelp API

August 2025

- Built swipe-based restaurant discovery platform ingesting **5,000+** Yelp API listings with dynamic geo-filtering
- Engineered Supabase auth and **row-level storage policies** achieving **100%** session data retention
- Optimized query caching layer to sustain **>10k swipe requests/min** without degradation under load

### Converso (SaaS Machine Learning Companion) | Next.js, Tailwind CSS, Supabase, Stripe, Vercel

June 2025

- Architected real-time voice pipeline via **VAPI WebSockets** driving **100+** daily interactive AI lessons in beta
- Integrated **Stripe subscription billing** with Google/Email auth supporting **50+** active paying users
- Deployed PostgreSQL-backed Vercel runtime with **Sentry error telemetry** for real-time fault observability