

Michael Cao

651-492-2346 | michaelcao@stanford.edu | [linkedin.com/in/cao-michael](https://www.linkedin.com/in/cao-michael) | github.com/mcao0620

EDUCATION

- Stanford University** Stanford, CA
Master of Science in Computer Science (AI) *Jan. 2022 – June 2024*
- Stanford University** Stanford, CA
Bachelor of Science in Computer Science (Information), GPA: 3.99 *Sept. 2019 – June 2023*

EXPERIENCE

- Software Engineering Intern** June 2023 – Sept. 2023
Databento *New York, NY*
- Implemented and deployed a reliable WebSocket API in Python for real-time message streaming to client portals
 - Improved flexibility with topic subscriptions and established backend-wide event streaming with Redis pub/sub
 - Saved millions of API calls/day and ensured staying under Sentry monitoring quota by eliminating busy-polling
 - Collaborated closely with frontend and product leads to align project requirements with business objectives
- Python Development Intern** June 2022 – Aug. 2022
Akuna Capital *Chicago, IL*
- Built full-stack web app using Python and React, enabling traders to configure parameters for new expirations
 - Introduced an asynchronous per-parameter handler strategy to allow for efficient parameter creation
 - Established robust input validation, error handling, and error reporting and achieved 80% test coverage
 - Designed a flexible and effective PostgreSQL database schema for storing parameter configurations
- Software Engineering Intern** June 2021 – Sept. 2021
DoorDash *San Francisco, CA*
- Deployed gRPC microservice endpoints using Kotlin, optimizing store menu and data fetching for SEO webpages
 - Implemented a read-through caching scheme for Redis with partial caching and request deduplication
 - Reduced average endpoint latency by 10x from 200ms to 20ms and lifted SEO page conversion rates by 3%
 - Managed Kafka consumer to process menu change events and keep cached menus up to date
- Software Engineering Intern** June 2020 – Sept. 2020
UnitedHealth Group *Minnetonka, MN*
- Implemented CycleGAN in Tensorflow for document cleanup, improving OCR performance by over 5%
 - Developed a React-based frontend used by 10+ teams to streamline and automate image classification workflows

PROJECTS

- Parabolic - Knowledge-Enriched Software Tickets** | *LlamaIndex, GPT-4, Tree-sitter, Streamlit* Aug. 2023
- Implemented generative AI assistant that helps users create codebase-aware software tickets
 - Wrote indexer to generate embeddings as well as Abstract Syntax Trees (ASTs) for code chunks
 - Utilized semantic search to identify relevant code snippets and NetworkX/Graphviz to visualize their relationships
 - Generated issue titles and descriptions with GPT-4 and integrated with the Linear API to upload enriched tickets
- Droplet - Soundbite Social Media Platform** | *React Native, Redux, Firebase* Sept. 2020 – Jan. 2021
- Developed a social media app for shortform audio and launched beta with 200 downloads on iOS
 - Implemented and owned key pages such as the main feed, profile page, and audio recording page
- Retro QANet - CS224N Final Project** | *PyTorch* Feb. 2021 – Mar. 2021
- Developed machine reading comprehension (MRC) model to accurately solve question answering problems from the Stanford Question Answering Dataset (SQuAD)
 - Combined two leading non-PCE model architectures to achieve an F1/EM score of 66.10/62.28

TECHNICAL SKILLS

Programming Languages: Python, JavaScript, Kotlin, C, C++, SQL, HTML, CSS, Swift, Java, C#
Frameworks: React, Redux, Node.js, FastAPI, Express, SwiftUI
Tools: Git, Docker, Kubernetes, Postgres, Firebase, MongoDB, Kafka, Unity
Libraries: PyTorch, NumPy, TensorFlow, pandas, matplotlib, scikit-learn, seaborn
Certifications: Coursera Deep Learning Specialization