

# Zhongquan (Peter) Cheng

St. Louis, MO · chengzhongquan0630@gmail.com · (315) 317-5618 · linkedin.com/in/petercheng · github.com/petercheng

---

## EDUCATION

---

### Washington University in St. Louis

August 2025 – May 2027 (Expected)

M.S. in Computer Science | GPA: 3.7/4.0

Coursework: Advanced Algorithms, Distributed Systems, Machine Learning, Computer Vision

### Syracuse University

August 2021 – May 2025

B.S. in Computer Science

---

## WORK EXPERIENCE

---

### Undergraduation.com — Backend Intern

September 2025 – Present

- Designed a distributed request pipeline using Server-Sent Events (SSE) and Vercel serverless functions, enabling asynchronous tool execution to support 50+ concurrent users with zero blocking
- Implemented a fault-tolerant LLM gateway with automated failover logic, utilizing latency-based A/B provider selection to maintain 99.9% uptime and optimize token costs during provider outages
- Optimized PostgreSQL performance by implementing a composite key caching strategy, eliminating 40% of redundant queries and reducing database load while ensuring real-time data consistency
- Enforced defense-in-depth security via Row-Level Security (RLS) and API middleware, strictly scoping Pinecone vector searches to counselor\_id to guarantee zero data leakage between tenants

Tech Stack: Node.js · Supabase · PostgreSQL · Vercel · Pinecone · Redis · LLM Orchestration

### Asia Pulp & Paper — Data Science Intern

May 2023 – June 2023

- Developed a time-series sales forecasting engine using XGBoost, engineering lag features to capture seasonal patterns and achieving 92% prediction accuracy for inventory optimization
- Automated ETL pipelines interacting with the Data Warehouse via optimized SQL queries (CTEs, indexing), reducing data extraction latency by 40% for downstream reporting
- Analyzed high-dimensional transaction data using Pandas vectorization, identifying key sales trends that informed production strategy for the Shanghai region

Tech Stack: Python · SQL · XGBoost · Scikit-learn · Pandas · ETL · Data Warehousing

---

## RESEARCH EXPERIENCE

---

### AI for Health Institute, WashU — Computer Vision Researcher

January 2025 – Present

- Built a medical imaging pipeline integrating the Meta Sapien model with OpenCV to generate robust segmentation masks for overhead patient monitoring across 10,000+ video frames
- Accelerated video data ingestion using FFmpeg for high-speed frame extraction, reducing preprocessing latency by 60% and enabling efficient processing of 50+ hours of clinical footage
- Created an image enhancement workflow that dynamically adjusts exposure levels, improving the human detection rate by 80% in low-light clinical environments

Tech Stack: Python · PyTorch · Meta Sapien · FFmpeg · OpenCV · Image Processing

### Syracuse University — Research Assistant (Advisor: Prof. Asif Salekin)

February 2024 – Present

- Developed a scalable Android data collection framework using the Factory Design Pattern to decouple sensor implementations, allowing seamless extension of multi-modal IMU streams
- Containerized the ML training pipeline using Docker Compose, cutting environment setup time from 2 hours to 5 minutes and ensuring reproducibility across compute nodes
- Designed a hybrid CNN-LSTM neural network to capture temporal anomalies in high-frequency IMU time-series data, achieving 87% accuracy in detecting depression biomarkers

Tech Stack: Android (Kotlin) · Docker Compose · Factory Pattern · CNN-LSTM · IMU Sensor Fusion

---

## TECHNICAL SKILLS

---

**Languages:** Python, JavaScript/TypeScript, Kotlin, SQL, Java, C/C++, C#, R, Haskell

**Frameworks & Tools:** PyTorch, TensorFlow, Node.js, React, Unity, Docker, OpenCV, FFmpeg, Git

**Data & AI:** XGBoost, Scikit-learn, Pandas, CNN-LSTM, LLM Fine-tuning & Orchestration, Meta Sapien, PEFT

**Platforms & Databases:** PostgreSQL, Supabase, Redis, Pinecone, Linux, Vercel, AWS