

Toney Zhen

✉ toney.zhen@sjsu.edu | ☎ (562)673-5846 | 📍 San Jose, CA | 🔗 [linkedin.com/in/toney-zhen](https://www.linkedin.com/in/toney-zhen)

Education

San Jose State University

2025 - 2027 (Expected)

M.S. in Artificial Intelligence

- **Relevant Coursework:** Artificial Intelligence, Machine Learning, Deep Learning, Reinforcement Learning

University of California, Irvine

2020 - 2024

B.S. in Computer Science

GPA: 3.76/4.00

- **Relevant Coursework:** Web Apps, Data Management, Algorithms, Operating Systems, System Design
- **Awards & Activities:** Regents' Scholar; Officer of Chess Club & Pokemon GO Club; member of ACM & Hiking Club

Skills

Programming Languages: Python, Java, JavaScript, C, C++, SQL

AI/ML: PyTorch, torchvision, Scikit-learn, NumPy, Pandas, Matplotlib, CNNs, Vision Transformers

Frameworks & Tools: Django, Flask, Node.js, React, JUnit, Git, Docker, Kubernetes, AWS, VS Code

Experience

Pan Research Group at UCI

Oct. 2023 - Jan. 2025

AI/ML Research Assistant

Irvine, CA

- Trained a Pytorch model on a dataset of **300+** labeled images, achieving **79%** accuracy in grain boundary detection, replacing slow manual process and freeing researchers to focus on analysis
- Leveraged NumPy, SciPy, and Matplotlib to predict TEM to HRTEM image transformations, improving data quality
- Collaborated weekly with faculty and grad students, incorporating their feedback to refine models and validate results
- Documented experiments and shared reproducible code so other lab members could extend the work, supporting ongoing material science research

Federal Geographic Data Committee

Sep. 2023 - May 2024

Applied Science Research Intern

Remote

- Automated management of extensive spreadsheet data, saving hours of manual effort while ensuring data integrity
- Worked under a remote supervisor and adapted deliverables based on feedback

Projects

DeepFake Image Classifier

Python, Scikit-learn, PyTorch, torchvision

- Fine-tuned pretrained ResNet and EfficientNet CNNs in PyTorch to classify real vs. AI-generated images across two datasets, establishing baselines to detect synthetic media
- Co-developed advanced models (ConvNeXt-Tiny, MaxViT, dual-branch MaxViT), tuning the top model to **~0.99** validation AUC to improve detection over the CNN baselines
- Ran cross-dataset and Grad-CAM analysis to identify where models failed to generalize

FabFlix MovieMart

JavaScript, jQuery, Java, Jakarta Servlet, MySQL

- A scalable full-stack web app providing support for an e-commerce catalogue of **20,000+** movies and **70,000+** stars
- Set up an AWS EC2 instance configured with MySQL and Apache Tomcat to host and deploy the app
- Provided essential features such as sessions, fuzzy search with autocomplete, advanced search options, secure login, PreparedStatements to prevent SQL injections, and load balancing for optimizing website performance
- Employed MySQL replication to synchronize data, Docker to rapidly deploy containers, and configured a Kubernetes cluster on AWS to manage pods, enabling exponential app scaling and support for **~9,000** requests/min

LLM Safety Guardrails

Python, Scikit-learn, NumPy, Pandas, Matplotlib

- Built a sensitive-data (PII) input guardrail to block private information before it reaches an LLM, training and comparing multiple classifiers on a **300k** example PII dataset
- Raised the top model to **92.9%** accuracy, **97.6%** recall, and **96.1%** F1 by replacing TF-IDF features with text embeddings and PCA, improving detection while cutting training runtime
- Coordinated with a **4-person** team, building four complementary input/output guardrails, aligning on shared precision and recall targets so results stayed comparable across models