Elvin Li

469-732-0613 | li.elvin739@gmail.com | linkedin.com/in/li-elvin | github.com/ElvinLit | https://elvinlit.github.io/

Research Interests

AI/ML Robustness, Adversarial Machine Learning, LLM/VLM Safety & Security, Continual Learning

Education

University of California - San Diego — San Diego, CA

Mathematics - Computer Science (B.S.) | GPA: 3.85/4.00 — *Expected: Jun 2026*

Relevant Coursework: Advanced Data Structures, Design and Analysis of Algorithms, Computer Organization, Machine Learning, AI Algorithms, Natural Language Processing, Stochastic Processes, Convex Optimization, Mathematical Statistics, Linear Algebra, Numerical Analysis, Vector Calculus, Differential Equations, Abstract Algebra, Real Analysis, Graph Theory

Technical Skills

Languages: Python, C/C++, Java, JavaScript, SQL, HTML/CSS
Machine Learning: PyTorch, TensorFlow, Scikit-Learn, Pandas, NumPy, Matplotlib, vLLM, HuggingFace
Developer Frameworks: FastAPI, Flask, Electron, ReactJS, Amazon Web Services (AWS)
Developer Tools: Git, Jupyter Notebook, Visual Studio Code, Vim, Kubernetes, Docker

Research Experience

Systems Energy and Efficiency Lab at UC San Diego — Oct 2023 – Present

Undergraduate Researcher, primarily supervised by Dr. Onat Gungor — *San Diego, CA*

SAFE: Self-Supervised Anomaly Detection Framework for Intrusion Detection [AAAI-25 AICS] **(Accepted)**

- First author lead of the conceptual development of a self-supervised intrusion-detection framework designed for realistic intrusion detection environments with scarce labeled data. Accepted to AAAI-25 AICS as an oral paper.
- Designed and implemented the full pipeline with a focus on deployment constraints, resulting in a 26% improvement over published state-of-the-art intrusion detectors on the same benchmark datasets.
- Pioneered a tabular-to-image transformation pipeline and constructed a masked autoencoder to extract high-quality latent features, establishing a new approach for applying image-based autoencoders to tabular anomaly detection.

CITADEL: Continual Anomaly Detection for Enhanced Learning in IoT Intrusion Detection [IoTJ'25] **(Under Review)**

- Extended self-supervised intrusion detection into a continual-learning setting to address practical limitations of finite model storage deployed classifiers. Spearheaded project orchestration as first author.
- Built a complete continual-learning experimental suite, including data-stream generation, benchmarking pipeline, and evaluation protocol, to support reproducible comparative studies.
- Developed a memory-efficient sample retention approach that preserves past attack sensitivity without storing all historical inputs, improving backward transfer by 72% over standard continual-learning baselines.

Rigorous Evaluation of Machine Learning-Based Intrusion Detection Against Adversarial Attacks [CSR'24] **(Published)**

- Contributed as second author to a study on intrusion-detection model robustness under adversarial perturbations.
- Developed a benchmarking framework enabling systematic cross-model comparison under multiple adversarial attack algorithms, implementing traditional ML classifiers and novel deep learning methods.
- Established the experimental protocol including data preparation, adversarial generation, and robustness metrics to ensure reproducible evaluation.

DYNAMITE: Dynamic Defense Selection for Machine Learning-based Intrusion Detection [SafeThings'25] **(Published)**

- Contributed to the development of an adaptive defense selector that selectively applies defensive strategies per incoming data sample. Awarded Best Paper.
- Reduced defensive computational overhead by 96% compared to earlier adaptive methods, producing a substantial improvement over existing adversarial defense methods.
- Demonstrated improvement of up to 65.8% in model robustness over the best static defense system, validating the system as a practical and effective defense mechanism for real-world intrusion-detection workloads.

Scripps Institution of Oceanography — Sep 2023 – Jun 2024

Undergraduate Researcher, supervised by Dr. Tom Corringham — *San Diego, CA*

- Developed ML-based tools for climate researchers, primarily to analyze large scale textual climate data.
- Implemented statistical learning models (KNN, XGBoost, etc.) and fine-tuned large language models (BERT, GPT) on climate corpora, creating topic classifiers for regional analysis of prevalent climate issues.
- Employed In-context learning (ICL) techniques to augment output generation for LLMs, fostering consistent classification abilities given zero-shot, one-shot, and few-shot examples.

Awards

Best Paper Award, SafeThings'25: Awarded for R&D on a practical defense mechanism for adversarial attacks.
Locke Lord Award, DFWCIA: Received $500 for developing an intelligent travel agent with booking capabilities.

Work Experience

Amazon Jun 2025 – Sep 2025
Software Development Engineering Intern *Seattle, WA*

- Designed and developed an end-to-end network traffic prediction system to anticipate high-traffic events and rescale Alexa AI services, minimizing latency risks across 20+ services to reduce customer throttling.
- Accelerated scaling workflows by 90% (20 days to 2 days) and drove org-wide adoption by 50+ engineers.
- Built system using React and a Java backend with AWS RDS (PostgreSQL); designed 17 REST APIs, managed a productionized CI/CD pipeline, integrated a deep learning model and deployed on ECS Fargate.

Projects

Intelligent Travel Planner | *Python, Flask, HTML/CSS, OpenAI, Selenium*

- Led a team of five to build a full-stack intelligent travel assistant with real-time flight search, lodging and restaurant retrieval, and personalized itinerary recommendations.
- Integrated LangChain with OpenAI's API to enable context-aware conversational planning and dynamic task execution.
- Awarded 3rd place overall at the DFWCIA Hackathon and received the $500 Locke Lord Prize

Dictate: Voice Controlled System Overlay | *FastAPI, LangChain/LangGraph, React, MCPs*

- Developed a system application for hands-free system interaction via NLP and real-time action execution.
- Built service using FastAPI Websockets for low-latency API communication; Powered by agentic LLMs through LangGraph and Claude, integrating 15 Model Context Protocol tools for automated computer tasks.

Activities

Stanford University Code in Place | *Section Leader* Apr 2023 - Jun 2023

- Hosted live weekly programming lessons for CS106A (Programming Methodologies), facilitating a learning environment for a cohort of 15 students on introductory data structures and programming principles.

Triton NeuroTech | *Machine Learning Team* Nov 2023 – Mar 2024

- Developed an LSTM with 90% accuracy for the Neural Prosthetics Group, effectively leveraging EMG technology to translate muscle signals into robotic movements for prosthetic limbs.
- The resulting prosthetic-arm demonstration was featured at the California NeuroTechX conference.

CSES Innovate | *Research Engineer* Jun 2025 - Present

- Collaborating with a team of eight developers to build a voice-to-code IDE designed to improve accessibility for programmers with motor impairments.
- Designing and implementing the Automatic Speech Recognition (ASR) module to facilitate natural language understanding.