Alvyn Kwon

646-719-6569 | alvyn.kwon@gmail.com | linkedin.com/in/alvynkwon | github.com/AlvynK123

EDUCATION

Georgia Institute of Technology   **August 2023 - May 2027**
Bachelor of Science in Computer Science (AI & Systems), GPA: 3.75   *Atlanta, Georgia*
- Relevant Coursework: Machine Learning, Artificial Intelligence, System and Networks, Computer Organization and Systems, Design and Analysis of Algorithms, Probability and Statistics, Objects and Design, Linear Algebra

SKILLS AND HONORS

Languages: Python, Java, C/C++, SQL, JavaScript/TypeScript
Systems / Infrastructure: AWS (EC2, EKS, S3), Docker, Kubernetes, Linux, CI/CD (Jenkins, GitHub Actions)
Data / APIs: PostgreSQL, MySQL, MongoDB, RESTful APIs, ETL pipelines
Honors: USACO Gold Division (2022), AIME Qualifier (2021, 2022)

EXPERIENCE

Samsung Electronics America   **June 2025 – August 2025**
Software Engineer Intern   *Mountain View, CA*
- Designed and automated metrics collection, alerting, and anomaly detection pipelines for large-scale eCommerce services using Python, improving incident detection speed by **10x**
- Migrated production monitoring workloads from AWS EC2 to Kubernetes (EKS), implementing Jenkins-based CI/CD pipelines to improve deployment reliability and rollback safety
- Built reusable internal observability APIs and dashboards (TypeScript/Next.js, REST, SQL) adopted across teams to reduce manual diagnostics by **80%**

Lead Software Engineer, Platform Systems   **March 2025 – Present**
Georgia Tech Student Government Association IT   *Atlanta, GA*
- Leading the development Course Critique, serving **40K users/semester**, guiding roadmap execution and reviews
- Built user-facing GPA dashboards and instructor search interfaces (Next.js/TypeScript) integrated with Firebase Auth/Firestore for real-time, personalized data exploration
- Automated ETL workflows in Python to transform institutional grade datasets into analytics-ready SQL tables, reducing manual processing by **90%** and ensuring reliable semester updates

Lead AI/ML Undergraduate Researcher   **Aug. 2024 – May 2025**
Georgia Institute of Technology - College of Computing   *Atlanta, GA*
- Led a **5-person** team fine-tuning PyTorch LLMs on large-scale materials datasets, boosting predictive accuracy by **90%**
- Designed and iteratively refined prompts for conditional crystal structure generation, applying Chain of Thought (CoT) prompting to improve output reliability
- Applied transfer learning and hyperparameter optimization to enhance generalization across scientific benchmarks
- Developed regression metrics and error analysis tools with Python/NumPy/Pandas, enabling rapid diagnosis of model bottlenecks and accelerating experiment cycles by **2x**

Software Engineer Intern   **May 2024 - Aug. 2024**
MasterKlass Academy   *Palisades Park, NJ*
- Developed and maintained **backend microservices** (Java, Spring Boot) supporting **500+ users**, with emphasis on clean APIs and service reliability
- Implemented containerized CI/CD pipelines with Docker and GitHub Actions, adding automated testing to reduce production defects by **30%**

PROJECTS

Airline Management GUI | *SQL, Express, TypeScript, React, MySQL*
- Developed an airline management app by translating airline operations into a **relational schema** with **14 tables**, **19 stored procedures**, and **6 views**, enabling full simulations and passing **1000+** unit tests
- Reduced backend code by **70%** via modular Express routes and reusable Next.js components

AI-Driven Music Personalization App | *Python, Django, PostgreSQL, OpenAI API*
- Created an AI personalization system integrating OpenAI API with Spotify data for dynamic playlists and user insights
- Implemented asynchronous data pipelines and PostgreSQL storage for scalable recommendations and efficient handling of high-volume listening histories

Fake eCommerce Review Prediction Model | *Python, Scikit-learn, Pytorch, Kaggle*
- Built an eCommerce fake review detection model using **40K labeled product reviews** and **200+** textual and metadata features from Kaggle to identify fraudulent content
- Achieved **88% Logistic Regression** and **90% SVM** accuracies via data scaling, hyperparameter tuning, and optimization techniques