JATIN KULKARNI

Seattle, WA | 832-623-2238 | j10.kulkarni@gmail.com
linkedin.com/in/jatin-kulkarni | github.com/jatinkulkarni | jatinkulkarni.com

EDUCATION

Cornell Tech (Cornell University), New York, NY — May 2025
Master of Engineering in Computer Science | GPA: 3.863 |Merit Scholar
Relevant Coursework: Applied Machine Learning, Machine Learning Engineering, Computer Vision, Natural Language Processing

University of Texas, Austin, TX — May 2024
Bachelor of Science in Computer Science | GPA: 3.72
Concentration in Machine Learning & Artificial Intelligence
Minors: Business, Design Strategies, Entrepreneurship
University Honors (August 2020 – May 2024)
Relevant Coursework: Natural Language Processing, Neural Networks, Intro to Artificial Intelligence, Principles of Machine Learning

TECHNICAL SKILLS

Coding Language and Tools: Python, Java, C, C++, JavaScript, TypeScript, HTML, CSS, Git, GitHub, Docker, Figma
Machine Learning and AI: TensorFlow, PyTorch, Scikit-Learn, LlamaIndex, LangChain, HuggingFace, Azure AI Search

EXPERIENCE

Amazon Web Services (AWS) – Software Development Engineer, AI Platforms, Bellevue, WA — July 2025 – Present
- Identified a scalability bottleneck in configuration storage and led migration to AWS AppConfig, reducing feature deployment time from 6–10 days to 1–2 days
- Reduced customer friction by ~96% by analyzing 290+ limit increase tickets and redesigning default capacity limits
- Enabled zero-touch region expansion by implementing dynamic region configuration across multiple team codebases, eliminating manual coordination and enabling automatic region launches
- Designed a capacity validation system across customer-facing paths to prevent exposure of restricted EC2 capacity while enabling internal tooling
- Integrated new instance types into SageMaker Training Plans, unblocking HyperPod and Training Job teams to begin testing and rollout
- Built support for new Inference resource types in SageMaker Training Plans, enabling flexible training and inference workloads through redesigned API validation
- Authored design documents and led implementation of capacity-related features from development through production deployment

Aristocrat Technologies, AI Research & Development Intern, Austin, TX — January 2024 - June 2024
- Built LLM-based systems using Retrieval-Augmented Generation (LlamaIndex, LangChain) to improve test case generation and root cause analysis
- Developed defect prediction models using data analysis techniques to improve software reliability and reduce flaky tests
- Collaborated across teams to deploy AI-driven quality assurance solutions, improving testing workflows at scale

Kellogg Brown & Root International, Inc. – Technical Intern, Houston, TX — May 2023 - August 2023
- Implemented automation pipeline to migrate enterprise document systems to Azure Cognitive Search
- Configured cloud security protocols using AWS IAM and Azure Active Directory for new projects and roles

Responsible Artificial Intelligence Institute – Software Engineering Intern, Austin, TX — January 2022 - December 2022
- Scaled RAI Collab platform to improve performance and reduce maintenance overhead through system redesign
- Built AI Regulatory Tracker (React + Firebase) tracking 170+ global AI policies used across multiple countries

Research and Projects

Geolocation Research, Cornell University, Prof. Hadar Averbuch-Elor | Visual Computing Group — Fall 2024 – Spring 2025
- Developed vision-language models for image-based geolocation using hierarchical embeddings and ViCLIP
- Designed hierarchical geocell representations to improve prediction granularity and model generalization
- Conducted large-scale experiments and ablations to evaluate embedding strategies and geographic resolution tradeoffs

Handwritten Equation → LaTeX (CNN, ViT, Seq2Seq Transformers), Cornell University — Fall 2024
- Built multi-stage pipeline combining CNN/ViT-based symbol recognition with T5-based sequence generation
- Achieved 72.87% symbol accuracy and 66.19% exact match rate (BLEU 0.84)
- Evaluated model architectures and training strategies to optimize sequence generation accuracy and robustness

Multimodal Ocular Disease Classification (CLIP, ResNet), Cornell University — Spring 2024
- Developed multimodal classification system using CLIP and ResNet, integrating imaging and patient metadata
- Achieved 92.52% accuracy (ResNet) and 89.31% (CLIP), analyzing tradeoffs between multimodal and image-only models
- Applied Grad-CAM and embedding analysis to evaluate interpretability and model behavior