Divyam Sengar

divyamsengar@gmail.com | Github Profile | LinkedIn | Website Portfolio

Education

Master of Science in Computer Science (AI Depth) Jan 2025 – June 2026
University of California, San Diego

Bachelor of Science in Computer Science Sep 2022 – Mar 2025
University of California, San Diego

Specialized Coursework:

- **AI/ML:** Theoretical Aspect of LLMs, LLM System Optimization, Probabilistic AI, Deep Learning, Statistical NLP
- **Math & Optimization:** Convex Optimization, Numerical Linear Algebra, Discrete & Continuous Optimization, Numerical Analysis
- **Advanced Theory & Compute:** Quantum Complexity & Cryptography, Lattice Algorithms & Applications, Parallel Computing, Physics Simulation, Computability & Complexity, Modern Cryptography, Data Center Systems

Work Experience

Teaching Assistant Jan 2026 – Present
UC San Diego Computer Science and Engineering *San Diego, CA*

- Sole TA for CSE 106: Discrete & Continuous Optimization, teaching 100+ students theory & foundational ML.

Software Development Engineer Intern Jun 2025 – Sep 2025
Amazon Web Services *Bellevue, WA*

- Fully architected and developed OSCAR, an open-source serverless multi-agent AI Chatbot for OpenSearch release automation. Highlighted by the Linux Foundation @ Open Source Summit, presented @ the OpenSearch Conference, and awarded formal recognition from AWS as an "AWS AI win".
- Engineered a multi-agent framework using AWS Bedrock to facilitate communication between supervisor and specialized agents. Architected a bespoke RAG system via OpenSearch VectorDB to retrieve context-specific knowledge from S3 data streams.
- Designed a highly scalable, event-driven backend using AWS CDK for Infrastructure as Code. The system leveraged an API Gateway, DynamoDB, and auto-scaling Lambda functions to process asynchronous Slack events without throttling.
- Implemented a novel dual-agent security architecture for secure, natural-language Jenkins job executions. Enabled users to query OpenSearch clusters and automate stakeholder communication, resulting in an 80% reduction in workflow management time, 90% faster metrics analysis, and 95% time saved on cross-team communication.

Software Engineer Intern Jun 2024 – Sep 2024
DataAnnotation *Remote*

- Improved AI chatbot performance through fine-tuning ensembled transformer models, resulting in improved code correctness and efficiency.
- Engineered training pipelines for LLMs, generating diverse and complex coding problems across multiple languages to enhance model reasoning.
- Optimized model architecture by eliminating key bottlenecks, implementing data-driven latency improvements, and boosting output quality.

Instructional Assistant Sep 2023 – Jun 2025
UC San Diego Computer Science and Engineering *San Diego, CA*

- Mentored 500+ student classes for Python Programming, Advanced Data Structures, Computer Organization, and Deep Learning.
- Conducted tutoring over 20+ weekly office hours, produced homework, exams, and dynamic question generators, improving assessment of student comprehension.

Full Stack Developer Intern Jun 2023 – Sep 2023
Computer Science and Engineering Society @ UCSD *San Diego, CA*

- Developed a dynamic membership platform website using React.js and Node.js, enabling secure user authentication and increasing membership engagement by 50%.
- Implemented the backend with Express.js and MongoDB to efficiently store and manage user & event data.

Projects

Privacy Impact Indicator | Chrome MV3, JavaScript, Python, Selenium, Scrapy, Pandas

- Engineered Privacy Impact Indicator, a Chrome extension with a real-time risk engine using a custom exponential scoring model, which successfully reduced user tracker exposure by 32% and third-party requests by 14%.
- Architected a Selenium-Scrapy pipeline to audit major AI platforms, analyzing 157 cookies to identify widespread use of long-lived persistence (averaging 293 days) and security flag vulnerabilities.
- Validated the tool's custom privacy metric across a 100-URL corpus, achieving an effective 0.91 AUROC against independent Privacy Badger grades to demonstrate strong discriminative power.

Diabetes Prediction | Python, sklearn, Pandas

- Codeveloped a diabetes prediction pipeline analyzing physiological and demographic data to evaluate and compare machine learning and transformer-based architectures.
- Tailored and ensembled SVM, XGBoost, Logistic Regression, and Random Forest models, establishing a robust baseline classifier that achieved 91.6% accuracy.
- Engineered a custom TabTransformer architecture for numerical and categorical features, yielding a 120% F1 score improvement over the baseline, an AUC of 0.97, and a peak accuracy of 96.57%.

Welp | Golang, Python, C, Docker, Kubernetes, GCP, gRPC

- Built a Yelp-like microservices platform with scalable Detail, Review, and Reservation services using Golang and gRPC, deployed via Docker/Kubernetes with automated CI/CD pipelines.
- Improved max throughput by 4.2x via Kubernetes pod scaling & eliminating frontend bottlenecks and reduced storage access latency by 85% via custom caching architecture & policies.
- Reduced cross-core memory latency by 10.8ms on GCP by engineering a thread scheduling benchmark to optimize thread-to-core affinity and cache locality.

QuoteGuard | Python, PyTorch

- Developed a transformer-classifier to verify speakers of political quotes & speeches, achieving 96.2% accuracy by leveraging custom attention and embedding logic.

Successorator | Java, SQL, Android, Git, Github Actions, XML

- Designed and implemented a multifaceted to-do list Android app using Java, with recurrent tasks, focus mode, categories, calendar synchronization, persistence, deletion, rollover, and task preemption.
- Used SQL & RoomDatabase for persistence & speedy data retrieval, customized MVP architecture to streamline app workflows, and engineered a robust CI/CD pipeline for project integrity.

Skills

Languages: Java, Python, C, C++, Go, R, SQL, MATLAB, HTML/CSS, JavaScript, CUDA

Frameworks & Libraries: React, Express, MongoDB, Node, Android, RoomDatabase, PyTorch, sklearn, TensorFlow, Pandas, gRPC, Numpy, Scipy

Tools & Cloud: Git, JUNIT, PlayCrypt, Roboelectric, Docker, GCP, Kubernetes, AWS, OpenSearch