

Devansh Agarwal

thedevanshagarwal@gmail.com | 408.333.1264 | San Jose, CA, 95136

[LinkedIn](#) | [Google Scholar](#) | [GitHub](#) | [Website](#)

EDUCATION

Cornell University

Master's in Computer Science

August 2022 - May 2023

GPA: 4.08/4.0

Indian Institute of Technology (IIT) Hyderabad

Bachelor of Technology in Computer Science and Engineering

July 2016 - May 2020

GPA: 9.14/10 (Dept Rank: 1)

Advanced Coursework: Natural Language Processing (LLMs) • Applied Machine Learning • Deep Learning for Vision • Special Topics in Machine Learning • Probability of Computing • Computational Topology • Computer Graphics • Reinforcement Learning

TECHNICAL SKILLS

Machine Learning: AWS SageMaker • Bedrock • PyTorch • PyTorch Lightning • Scikit-Learn • LangGraph • LangChain • Transformers

Languages: C • C++ • C# • Java • Python • JS • TS • Matlab • Swift • PHP • SQL

SDKs/Environments: Unreal Engine • Unity • Visual Studio • Eclipse • VS Code • ARKit • Google VR • PlatformIO • Arduino

Other Technologies: AWS Lambda • SNS • SQS • ECS • EC2 • DynamoDB • S3 • Kafka • Git • Next.js • Node.js • Angular • React • CI/CD

WORK EXPERIENCE

AWS (Amazon), Machine Learning Engineer II

September 2024 - Present

- Built AWS Transform, a B2B agentic service that modernizes .NET enterprise applications. Leveraged advanced prompting, fine-tuned LLMs, self-correcting loops, and RAG to transform 50M+ lines of code and modernize thousands of workloads.
- Developed a GNN-based clustering system to analyze and group processes running in customers' EC2 instances, enabling automated modernization recommendations at scale.

EchoTrack, Founder (Cornell Ignite Fellow)

July 2023 - September 2024

- Built an end-to-end machine learning platform for acoustic image data, including data collection, processing, developing a new foundational model, and real-time inference.
- Developed ConvNets, RNNs, and Transformers for recognizing human body poses, speech, facial expressions, and activities using acoustic signals. Employed data augmentation and masking techniques to create state-of-the-art systems, achieving 90%+ accuracy. Deployed these models on cloud, smartphones and CNN accelerator hardware.
- Secured \$240,000 in funding from Cornell University and \$50,000 from National Science Foundation.

Amazon, Software Engineer II

June 2020 - July 2022

- Proposed and developed a machine learning recommendation tool for auto-filling customers' preferred delivery dates, reducing clicks from 5 to 1 for single-click checkout.
- Created a widget for Amazon's mobile homepage that is selectively shown to customers based on machine learning models. It serves 1600 transactions per second with 20 ms latency, decreasing customer anxiety for service orders by 2% and saving ~\$1 million in support costs.
- Migrated notifications platform to AWS, implemented an S3-based request and response tracking framework, ensuring timely and accurate notifications.

Amazon, Software Engineer Intern

May 2019 - July 2019

- Designed and engineered an invoice generating tool for a new B2B2C Amazon business. Created the back-end using DynamoDB and EC2, and front-end using Angular, HTML, CSS and JS.
- This tool helped save around ~40 SDE hours per month.

SmartVizX, Software Engineer Intern

May 2018 - July 2018

- Integral part of the team that developed India's 1st VR platform for the construction Industry, Trezi.
- Built the cloud stack and integrated it with the VR app for Trezi, which allowed users to download textures and models in real-time. Enabling Trezi to create a multi-million-dollar VR marketplace.

TEACHING ASSISTANT

- Cornell University:** Computer Vision (CS 5670), and Computer Networks (CS 5456).
- IIT Hyderabad:** Statistics (MA 2140), Probability (MA 2110), Vector Calculus (MA 1130), and Data Structures (CS 1353).

SELECTED PUBLICATIONS

- EchoGuide: Active Acoustic Guidance for LLM-Based Eating Event Analysis from Egocentric Videos**, V. Parikh, S. Mahmud, D. Agarwal, ..., and C. Zhang. (ISWC), 2024, **Best Paper Honorable Mention**.
- HPSpeech: Silent speech interface for commodity headphones**, R. Zhang, H. Chen, D. Agarwal, ..., C. Zhang. In ACM International Symposium on Wearable Computers (ISWC), 2023, **Best Paper Honorable Mention**.
- On Adversarial Robustness: A Neural Architecture Search perspective**, C.Devaguptapu, D. Agarwal, G. Mittal, P. Gopalani, and V.N. Balasubramanian. Workshop on Adversarial Robustness in Real-World, ICCV-21, also accepted in 4 ICLR-21 workshops, **Contributed Talk and Spotlight**.

AWARDS

- Ignite Fellowship, received \$240,000 as startup funding from Cornell University.
- Outstanding Reviewer ACM IMWUT 2024, and ACM UIST 2024.
- Institute Silver Medal for securing the highest GPA in my department at IIT, Hyderabad.
- Top 0.4% in IIT Joint Entrance Examination, 2016; taken by more than 1 million students.