

# Akshay Gopalkrishnan

MS Computer Science at University of California, San Diego

[akshay.gopalkrishnan@gmail.com](mailto:akshay.gopalkrishnan@gmail.com) | [agopalkr@ucsd.edu](mailto:agopalkr@ucsd.edu) | +1 (408)-382-1263 | [Website](#) | [LinkedIn](#) | [Github](#) | US Citizen

## EDUCATION

---

### University of California, San Diego (Master's)

2023 – 2025 (Expected)

- Computer Science Master of Science with a focus in Artificial Intelligence
- Relevant Coursework: Generative AI, Advanced Computer Vision, Advances in Language Models, Data Systems for Machine Learning, Natural Language Processing, Probabilistic Reasoning & Learning

### University of California, San Diego (Bachelor's)

2019 – 2023

- Majoring in Electrical Engineering with a specialization in Machine Learning & Controls, **GPA: 3.946**
- *Honors & Awards:* Magna Cum Laude, Phi Beta Kappa, Tau Beta Pi, Provost Honors, ECE Henry Booker Award
- Relevant Coursework: Deep Learning, Recommender Systems, Object-Oriented Programming with C++, Advanced Data Structures, Design & Analysis of Algorithms, Machine Learning, Probabilistic Reasoning & Graphical Modeling, Intro to Computer Engineering, Python for Data Analysis

## RESEARCH

---

*Advised by Dr. Hao Su and Dr. Mohan Trivedi*

### Visual Grounding of Robotic Vision-Language Models

- Enhancing pretraining strategies for the visual grounding of robotic vision-language models by integrating OWL-ViT, Segment Anything 2, and Depth Anything 2 to track the 3D image location of relevant task objects and their distance from the robot end-effector.
- Finetuning pretrained VLMs with LLaMA-2 LLM backbones on large scale robotic action datasets and multi-node GPU clusters.

### Efficient Vision Language Models for end-to-end Autonomous Driving

- Enabling human interactivity and interpretability of autonomous driving (AD) systems by designing an efficient Vision Language Model (VLM) that performs visual question answering tasks for AD tasks such as perception and motion planning with multi-view traffic scene images.
- Developed a custom lightweight and efficient VLM framework that requires nearly twenty times less parameters and memory than most LLMs using a quantized T5 Language Model and a multi-view encoder that unifies single image ViT patch embeddings.
- Integrated multi-modal Retrieval Augmented Generation to enhance context for multi-view images and prompts.

### Analyzing Alcohol Impairment of Drivers

- Addressing traffic accidents involving alcohol impairment through designing a computer vision and deep learning framework to non-intrusively monitor a driver's blood alcohol concentration.
- Using a vision Transformer and deep learning multimodal data fusion techniques to create a model that predicts the blood alcohol concentration of a driver with camera and thermal image data.

### Salient Traffic Sign & Light Detection

- Thoroughly annotated traffic light and signs detections in addition to a salience property which determines whether a light or sign is relevant to the next immediate decision a driver is made.
- Trained the first detection network for traffic signs & lights that emphasizes detecting salient traffic lights and signs using a Deformable DETR model with focal loss.

## PUBLICATIONS

---

All my research publications can be found on my [Google Scholar profile](#)

- **Akshay Gopalkrishnan**, Ross Greer, Mohan Trivedi, "[Multi-Frame, Lightweight & Efficient Vision-Language Models for Question Answering in Autonomous Driving](#)", CVPR 2024 Vision and Language for Autonomous Driving and Robotics Workshop, 2024.
- Ross Greer, **Akshay Gopalkrishnan**, Maitrayee Keskar, Mohan Trivedi, "[Patterns of Vehicle Lights: Addressing Complexities in Curation and Annotation of Camera-Based Vehicle Light Datasets and Metrics](#)" Pattern Recognition Letters, 2024.
- Ross Greer, **Akshay Gopalkrishnan**, Nachiket Deo, Akshay Rangesh, Mohan Trivedi, "[Salient Sign Detection in Safe Autonomous Driving: AI Which Reasons Over Full Visual Context](#)" 27th Enhanced Safety of Vehicles Conference, 2023.

- Ross Greer, **Akshay Gopalkrishnan**, Jacob Landgren, Lulua Rakla, Anish Gopalan, Mohan, Trivedi, "[Robust Traffic Light Detection Using Saliency-Sensitive Loss: Computational Framework and Evaluations](#)" IEEE Intelligent Vehicles Symposium (IV), 2023.
- Ross Greer, Samveed Desai, Lulua Rakla, **Akshay Gopalkrishnan**, Afnan Alofi, Mohan Trivedi, "[Pedestrian Behavior Maps for Safety Advisories: CHAMP Framework and Real-World Data Analysis](#)" IEEE Intelligent Vehicles Symposium (IV), 2023.

## WORK EXPERIENCE

---

### Machine Learning Research Intern at CIPRA.ai

*Jun. 2024 – Current*

- Surpassed previous state-of-the-art machine learning models for daily blood pressure prediction by nearly 38% through researching and designing a time-series Transformer Encoder architecture that utilizes a sequence of daily wearable device data.
- Introduced to the company a new time-series deep learning architecture to predict and monitor patient blood glucose levels using past meal contents and blood glucose levels.
- Prototyped a LLM patient assistant using LangChain that can access patient data stored in SQL databases and generate text responses to questions from patients about their health.

### Software Development Engineer Intern at AWS

*Jun. 2022 – Sept. 2022*

- Decreased the latencies by up to 1000% for a comment microservice in an internal employee Amazon website used by over 30,000 Amazon employees.
- Designed the new, more efficient comment microservice that contains APIs allowing users to create, update, and delete comments stored on a database using Python and the AWS tools Lambda, DynamoDB, and API Gateway.

### Software Engineer Intern at Cisco

*Jun. 2021 – Sept. 2021*

- Wrote Python files to increase unit test coverage from 65% to 90% for a Django application which serves as a SD-WAN (Software Defined Wide Area Network) Self-Service Portal for Cisco customers.
- Developed a Python script that automatically syncs customer/organizational data from Cisco's Azure Cloud Portal to the Self-Service Portal application database.
- Added additional web APIs in Python allowing customers on the SD-WAN Self-Service Portal application to view all their virtual accounts stored on Cisco's database.

## TEACHING EXPERIENCE

---

**Research Mentor:** Polygence

*March. 2024 - Current*

**Teaching Assistant:** CSE 256, Natural Language Processing

*Sept. 2024 – Current*

**Teaching Assistant:** ECE 172A, Introduction to Intelligent Systems

*Jan. 2024 – Mar. 2024*

**Undergraduate Tutor:** CSE 101, Advanced Data Structures & Algorithms

*Sept. 2022 – Dec. 2022*

## SKILLS

---

**Programming Languages:** Python, Java, C++

**ML/AI Libraries/Tools:** Pytorch, Tensorflow, HuggingFace, scikit-learn, Numpy, Pandas, Ray, FSDP, Weights & Biases, OpenAI API, LangChain

**Engineering Skills/Tools:** Git, Linux, AWS, Docker