

Vineet Rao

+1 (734) 881 8969 | kemmannu@umich.edu | vineetrao25.github.io | github: vineetrao25 | linkedin: Vineet Rao

SUMMARY

Research Engineer with 3+ years of experience developing and deploying end-to-end autonomy and robotics systems, including AI-driven manufacturing automation. Open to roles in autonomous driving, robotics, or applied ML teams at the intersection of deep learning and real-world deployment.

EDUCATION

University of Michigan

M.S. in Electrical and Computer Engineering (ECE), Machine Learning | GPA: 4.00/4.00

Ann Arbor, MI

Aug 2021 – Apr 2023

EXPERIENCE

Magna International

Research Engineer

Troy, MI

June 2023 – Present

Last Mile Delivery Device (LMDD) – Reuters Coverage, Marvel Studios Demo

- Developed the perception and localization stack for an L4 autonomous delivery robot deployed in Toronto, taking the system from early-stage research to **90%** autonomous operation on public roads.
- Architected and trained an online map-generation transformer and multimodal localization transformer reducing average localization error by **20 cm** and achieving sub-**100 ms** inference on NVIDIA Jetson Orin.
- Designed an auto-labeling pipeline using CLIP and Qwen2.5-VL with human-in-the-loop verification on continuous data streams, reducing labeling costs by **5x (\$25K saved)** and cutting annotation cycle times from **2 weeks to 5 days**.

AI-Driven Manufacturing Automation

- Designing and building the autonomy stack for **AMRs** (Autonomous Mobile Robots); supporting deployment of **144 robots** across **14 factories**, saving **\$5M annually**.
- Developing **agentic AI systems** for manufacturing by integrating LLM-based planning and reasoning into factory automation workflows to enable more adaptive task planning and execution.
- Building **Vision-Language-Action (VLA) systems** to connect LLM-based task planning with physical robot execution, enabling closed-loop long horizon decision-making for manufacturing automation.

Amazon Lab126

Sunnyvale, CA

Applied Scientist Intern

May 2022 – Aug 2022

- Implemented a **radar-based deep learning system** for exercise recognition and repetition counting from **Doppler images**, achieving **98% accuracy** across 5 exercise types and enabling integration into health-monitoring devices.
- Built a **scalable multimodal data platform** to synchronize camera and radar streams and automate ground-truth labeling, reducing data preparation time by **50%** and accelerating model development.

RESEARCH

Research Associate — Prof. Justin Johnson's AI Lab (JAIL)

Ann Arbor, MI

Mentor: *Karan Desai*

Aug 2022 – Apr 2023

- Developed **SCAM (Segment and Classify Anything Model)** by integrating **CLIP** and **SAM** to address open-vocabulary instance segmentation, improving COCO mAP by **+7.5 AP** and sustaining performance with **<1 AP drop** using only **5% labels**.

PROJECTS

Self-Supervised Object Detection with Multimodal Image Captioning [report]

advisor: **Prof. Honglak Lee**

- Built an Self Supervised transformer-based detector using caption supervision for label-efficient object detection, reducing labeled data requirements by **20%** while maintaining comparable AP.

Visual Question Answering using Customized Prompts [report]

advisor: **Prof. Joyce Y. Chai**

- Developed a zero-shot **VQA** system using BLIP/CLIP + LLaMA to reduce dependence on labeled training data, enabling domain-adaptive question answering without fine-tuning.

SKILLS

Languages: Python, C++, Julia

ML/DL Frameworks: PyTorch, TensorFlow, NumPy, OpenCV, ONNX

Robotics/Autonomy: ROS, ROS2, Sensor Fusion (Radar + LiDAR + Camera), Localization

Architectures: Transformers, Vision Transformers (ViT), Multimodal Transformers (VLMs), CNNs, VLAs, Diffusion Models, VAEs

Tools: Docker, Git, CI/CD, Linux, SLURM