

Karthik Pullalarevu

kpullala@andrew.cmu.edu | (412) 608-6231 | karthikpullalarevu.github.io/ | linkedin.com/in/karthikpullalarevu

EDUCATION

Carnegie Mellon University

Master of Science in Computer Vision | GPA: 4.22/4.0

Coursework: Learning for 3D Vision, Advanced Computer Vision, Intro to Robot Learning, Computer Graphics

Pittsburgh, PA

December 2026

Vellore Institute of Technology

Bachelor of Science in Computer & Electronics Engineering | GPA: 3.67/4.00

Chennai, India

June 2022

SKILLS

Programming Languages: Advanced: Python; Basic: C++, C, Java

Libraries: Advanced: PyTorch, NumPy, Pandas, Huggingface Transformers, Docker, ONNX; Basic: Tensorflow

Computer Vision: Advanced: Vision Transformers, Diffusion models, Novel View Synthesis, Neural Radiance Fields, Semantic Segmentation, 3D Reconstruction, Monocular Depth Estimation, Multimodal Models, GANs

PROFESSIONAL AND RESEARCH EXPERIENCE

Aurora

Computer Vision Student Researcher

Pittsburgh, PA

January 2026 - Present

- Analyzing rasterized scene representations on the Waymo Open Dataset (WOD) and designing generative visual interventions (Gemini Nano) to study end-to-end (E2E) planning models for autonomous driving in long-tail scenarios.
- Evaluating model behavior using metrics like Average Displacement Error (ADE) and Rater Feedback Score (RFS).
- Generating long-tail driving scenarios via agent interventions (e.g., pedestrian jaywalking) in the nuPlan simulator.

Hyperverge

Senior Machine Learning Engineer (MLE 3)

Bangalore, India

January 2022 - August 2025

- Developed custom multi-modal vision transformer models for face-spoof and deepfake detection, achieving **0.15% FAR** and **0.06% FRR**, serving **20 million verifications monthly** in production and resulting in three US patents.
- Designed a dense embedding retrieval pipeline using contrastive learning and nearest neighbor search with **20ms latency**, enabling real-time fraud-pattern clustering and reducing sudden fraud-inflow detection time to **15 minutes**.
- Reduced model size by **91% (335MB to 28MB)** and improved CPU **inference latency by 60%** through post-training quantization, knowledge distillation and pruning; deployed optimized models on-device via ONNX Runtime / TFLite.
- Built late-fusion anomaly detection models using mobile sensor data, depth and RGB for real-time fraud detection.
- Created an in-house synthetic dataset for eyewear detection by overlaying transparent frames onto facial images and deployed a lightweight multi-branch MobileNet model with **99.8% precision, 100% recall** and **5.8ms latency**.
- Automated hiring process by building an LLM powered agent for parsing resumes, scoring candidates and conducting automated AI interviews using speech to text and text to speech, **cutting hiring time by 95%** and **cost by 65%**.

PUBLICATIONS & PATENTS

[1] Wiggle and Go! System Identification for Zero-Shot Dynamic Rope Manipulation [2026 - Under Review]

[2] **Karthik, P.**, et al. Semantic segmentation for plant phenotyping using advanced deep learning pipelines. **Multimedia Tools & Applications** [10.1007/s11042-021-11770-7]

[3] Method and System For Determining Liveness of a Subject. **[3 US Patents]** [USPTO 20250104480]

PROJECTS

3D Shape Completion for Robotic Grasping

August 2025 - December 2025

- Explored test-time 3D reconstruction methods for shape completion to support downstream robotic grasping tasks.
- Implemented LoRA-based adaptation on Co3D dataset for object classes including *bottle, shoe and broccoli*.
- Investigated flow-matching techniques to reduce edge bleeding artifacts in reconstructed outputs.

Multiview pose synthesis from 2D objects

September 2024 - October 2024

- Developed an end-end pipeline for novel pose estimation using text-guided segmentation, image to 3D view generation followed by Stable Diffusion Inpainting using azimuth and polar values.
- Proposed Grounding SAM model for prompt-based object detection leading to 6% higher accuracy on segmentation.

Spectral alignment of brain mesh surfaces

August 2021 - December 2021

- Awarded the MITACS Fellowship and interned at ETS Montreal, Canada. Contributions to research on Graph CNNs for cortical shape analysis, increasing brain parcellation accuracy by 3%.