

NISHANTH VANIPENTA

Perception & Computer Vision Engineer

Dallas–Fort Worth, TX (open to relocate) | nishi@terpmail.umd.edu | nishanth.ai | linkedin.com/in/nishanthvanipenta | github.com/Nishanth2708

SUMMARY

Perception engineer with 5 years shipping camera and laser 3D measurement systems end to end: optics and sensor selection, calibration, classical geometry, deep learning and GPU model serving. Systems I built scan 85,000+ tires a day across 1,200 drive-over systems and 8,000+ handheld scanners. 2 patents pending. Also build multi-agent LLM workflows.

EXPERIENCE

TraXtion | Computer Vision Engineer

Jan 2022 – Present

- Built the production vision pipeline for a drive-over tread scanner on 1,200 systems scanning 70,000+ tires a day: multi-camera laser triangulation measuring tread depth on moving vehicles to ± 0.5 mm of a manual gauge in about 15 s. Patent pending.
- Made laser-based 3D measurement reliable outdoors, in rain, glare and low light, by combining optical filtering with robust point clustering.
- Selected the imaging hardware (industrial cameras, lenses, line lasers, optical filters, IMU and motion sensors) and ran field-of-view and resolution analysis, including a dual-camera rig for commercial trucks with dual rear wheels.
- Led field evaluation of a 6-camera structured-light 3D scanner (point-cloud segmentation, shape-based tire classification); showed it could not meet accuracy outdoors and recommended the laser-triangulation approach that went to production.
- Built an in-house system that reads tire sidewall text from any viewpoint, including phone photos, at 7.5x lower cost per image than the outsourced vendor with 66% more accurate detections; it became a new product. Patent pending.
- Took Segment Anything (SAM) from a research script to a production GPU endpoint (Docker, infrastructure as code) behind a feature flag with automatic fallback; raised detection rate from 40% to 95%+ and removed cold starts for live traffic.
- Built the sensor-fusion engine for a handheld laser tread scanner used on 8,000+ devices scanning 10,000+ tires a day (0.01 mm resolution), fusing laser profiles with motion data.
- Built a ground-truth regression suite and dashboards that gate every release; hardened CI/CD and serverless production services with health checks, monitoring, retries and versioned APIs.
- Improved calibration processes for production units; built calibration, scan-review and 3D visualization tools used across the team.
- Guide technical work for a team of 5 and designed AI-assisted code review for the team. TrailBlazer Award 2025, the company's highest award.

ZD Metal Products (Tier 1 GM supplier) | Robotics Engineer

Mar 2021 – Jan 2022

- Led a team of 3 to build a vision system measuring tapered-hole diameters on the Hummer EV frame to 0.08 mm precision, raising defective-part detection by 14%.
- Trained a multi-class YOLOv5 detector in PyTorch (89% accuracy) and deployed an optimized CNN on an embedded device for presence checks of pins and bushings on an inspection jig.
- Integrated FLIR cameras (Spinnaker SDK) with AWS S3 and Textract; programmed FANUC and UR5e robots and diagnosed PLCs, pneumatics and motor controllers.

AGENTIC AI

- Built a multi-agent workflow for a paper-supply business (inventory, sales, accounting transactions, customer support, lead generation) as my Udacity Agentic AI Nanodegree capstone; automate daily engineering work with Claude Code, skills, hooks and MCP, and host local models with Ollama.

PATENTS & AWARDS

- Intelligent Drive-Over Tire Analytics System and Method (patent pending)
- System and Method for Automated Extraction of Tire Sidewall Information from Curved-Surface Images (patent pending)
- TrailBlazer Award 2025, TraXtion's highest honor

SKILLS

Computer vision: Camera calibration, homography, laser triangulation, structured light, point clouds, OCR, segmentation

ML / DL: PyTorch, TensorFlow, YOLO, Mask R-CNN, Segment Anything, scikit-learn

Deploy / cloud: GPU inference endpoints, Docker, CI/CD, serverless APIs, Azure, AWS, GCP

Robotics / hardware: Industrial cameras, line lasers, IMU, sensor fusion, ROS, FANUC, UR5e, PLCs

Languages: Python, C++, SQL; NumPy, SciPy, OpenCV, pandas

Agentic AI: Multi-agent workflows, tool calling, MCP, Claude Code skills/hooks, Ollama, RAG

EDUCATION

M.S. Robotics | University of Maryland, College Park (2021)

B.Tech. Mechanical Eng., SRM University (2018)

Project: ARIAC (Agile Robotics for Industrial Automation Competition), ROS and C++.