

Jay Polra

jpolraa@gmail.com | +1 (219) 368-8271 | LinkedIn | GitHub | jaypolra.github.io

Publications

Visual Geometry Grounded Novel-View Acoustic Synthesis, *CVPR Workshop 2026*

Jay Polra, Dhwanil Chauhan, et al.

Accepted

Development of Trailing Image Detection for a Melt Shop Safety Tool, *AISTech 2026*

Kyle Toth, Jay Polra, et al.

Accepted & Presented

A Dual-Model Approach to Industrial Safety, *TMS / AIM 2025*

Kyle Toth, Jay Polra, et al.

Under Review

Research Experience

Graduate Research Assistant, *CIVS, Purdue University Northwest*

January 2025 – Present

- Existing acoustic synthesis methods required Structure-from-Motion geometry unavailable at inference. Redesigned around a feed-forward geometry encoder (VGGT) and cross-attention decoder routing on geometry but retrieving learned acoustic priors. Eliminated COLMAP with 10x faster preprocessing and improved results on all four metrics vs prior SOTA. Accepted at CVPR Workshop 2026.
- Camera-based safety systems fail when a vehicle disappears into a blind spot. Designed entry-exit state tracking as a conservative heuristic: assume worst case until a boundary camera confirms exit. Enables reliable zone isolation when the environment is physically invisible to the camera.
- Vision-language models produce fluent captions with no account of why. Built explainability tooling using Grad-CAM token attribution, region masking, and perturbation testing to verify causal influence of image regions on outputs, not just correlation.

Industry Experience

Site Reliability Engineer, *Asite Solutions Pvt Ltd*

January 2023 – July 2024

- Maintained 99% uptime SLA across 11 global production deployments (USA, UK, Canada, Australia, Hong Kong). Platform ran 25+ services on Azure with strict regional isolation requirements.
- JMS consumer queues were failing silently before users were affected. Wrote a Python monitor tracking consumer counts against baselines and triggering alerts before failures surfaced. Adopted as standard SRE team tooling.
- Built zero-downtime restart automation: verified a healthy instance was available before touching the degraded one. Safety check first, action second. Stack: Terraform, Jenkins, Ansible, Prometheus, Grafana, ELK, Azure Monitor.

Selected Projects

Geometry-Grounded Novel-View Acoustic Synthesis

PyTorch, VGGT, Transformers, STFT

- Feed-forward NVAS with separable K/V cross-attention: geometry routes attention, acoustic priors supply content. Outperformed AV-Cloud on MAG/ENV/LRE/DPAM with 3.24M vs 3.91M parameters.

AI Hazard Recognition System

YOLO, DeepSORT, FastAPI, React

- Multi-camera industrial safety monitoring with blind spot handling via entry-exit state tracking across 4 feeds. Real-time at 8 FPS with under 50ms zone state updates.

VLM Explainer: From Patches to Phrases

BLIP, CLIP, Grad-CAM, Streamlit

- Token-level Grad-CAM attribution with region masking and perturbation testing. Validated causal influence: masking the relevant region changed the caption. CLIP alignment scoring for cross-verification.

Education

Purdue University *MSc Computer Science* GPA: 3.3

Aug 2024 – Dec 2026

LJ Institute of Engineering and Technology *BE Information Technology* GPA: 3.72/4.0

Jun 2023

Technical Skills

Research PyTorch, Hugging Face, BLIP, CLIP, VGGT, YOLO, DeepSORT, OpenCV

Systems FastAPI, React, Streamlit, Docker, Kubernetes, Terraform, Jenkins

Languages Python, JavaScript, SQL, LaTeX, Bash, CUDA