# MADHUR THAREJA

Dual Major ES + CS • IIT Madras, BITS Pilani ◇ madhurthareja1105@gmail.com • Github • LinkedIn

## RESEARCH PROFILE

Undergraduate researcher spanning Computer Vision, Generative Modeling, Multimodal Reasoning, and Embedded Systems. My work focuses on building end-to-end intelligent systems that connect sensing, representation, reasoning, and actuation: multimodal grounding with uncertainty-aware reasoning (PAL, MediLLM), symbolic–neural integration for domain-constrained decision making, signal processing and perception under extreme noise (SONAR), and hardware–software co-design for efficient inference (RISC-V, FPGA, robotics platforms). I work across the full stack: algorithms, optimization, systems, and embedded deployment.

## PUBLICATIONS & DEMONSTRATIONS

- Thareja, M., et al. *"PAL: Personalised Adaptive Learning System."* Accepted as Demonstration Paper, AAAI Conference on Artificial Intelligence (AAAI 2026). Contributed the multi-agent multimodal reasoning module integrating VLM grounding, uncertainty estimation, and RL-based adaptation.
- *CANDI: Context-Aware NeuroSymbolic Domain Intelligence* — Under Revision. Developed multimodal alignment components and evaluation pipelines for role-aware and task-aware prompting across multiple LLMs using BERTScore, BLEURT, MoverScore, and Cosine Similarity.

## RESEARCH EXPERIENCE

**University of Pittsburgh** November 2025 – Present
*School of Medicine, Research Intern* *Remote*

· Contributing to LivecellX, a comprehensive deep learning live-cell analysis framework written in Python, designed specifically for segmenting, tracking, and analyzing single-cell trajectories in live-cell imaging datasets. reasoning pathways for diagnostic tasks.
· Currently working on identifying and testing the pre-existing tracking models for single-cell analysis. Ultrack: Ultrack is a versatile and scalable cell tracking method designed to address the challenges of tracking cells across 2D, 3D, and multichannel timelapse recordings, especially in complex and crowded tissues where segmentation is often ambiguous. By evaluating multiple candidate segmentations and employing temporal consistency, Ultrack ensures robust performance under segmentation uncertainty. developed by RoyerLab

**Carnegie Mellon University** August 2025 – Present
*Research Intern* *Remote*

· Contributing to MediLLM by coupling MedGemma VLM perceptual features with clinically structured reasoning pathways for diagnostic tasks.
· Extended MMedRAG through symbolic filtering and controlled retrieval to improve factual grounding and reduce hallucinations in clinical QA pipelines.
· Designed a triage-oriented multi-speciality agent system and an Easy to Hard (E2H) trainer module modelling diagnosis as progressive refinement under uncertainty.

**AI Institute of South Carolina** June 2025 – December 2025
*Research Intern* *Remote*

· Developed the multi-agent multimodal reasoning architecture for PAL by integrating VLM-based semantic grounding with uncertainty-weighted reinforcement learning policies.
· Analyzed structured prompting strategies (role-aware, task-aware) and their impact on semantic faithfulness for T1/T2 domain-specific reasoning tasks.
· Co-author of the AAAI 2026 demonstration paper.

**Samsung Semiconductor Research & IIIT Bangalore** Jan 2025 – Feb 2025
*Chip Design Studio — National Rank #2* *Bengaluru*

- Designed and verified single-cycle and 5-stage pipelined RISC-V microarchitectures, analysing timing, hazards, and throughput–area tradeoffs across FPGA platforms (ZCU104, DE0-Nano, Pynq-Z1).
- Completed full RTL-to-GDSII physical design using Cadence Genus/Innovus (45nm PDK), with functional equivalence validated via Conformal LEC.
- Implemented UART, PWM, frequency-scaling, and color-detection IP blocks and integrated them into a robotics control pipeline.

**IIT Guwahati** Summer 2025
*Research Intern* *Guwahati, India*

- Investigated object detection in low-SNR sonar imagery using CLAHE-based contrast enhancement, Gaussian denoising, and sonar-specific augmentations.
- Benchmarked YOLOv8n, Faster R-CNN, SSD, and AquaYOLO to assess noise sensitivity and preprocessing alignment effects.
- Achieved **mAP@0.5 = 0.9477** and **F1 = 0.9299**; identified dataset leakage and reran controlled experiments for validated results.

**Autonomous Smart Wheelchair @IIT Madras** 2024 – 2025
*Undergraduate Researcher* *Chennai, India*

- Engineered a semi-autonomous mobility system integrating LiDAR, stereo cameras, GPS, and IMU to construct a unified 360° environmental representation.
- Built a distributed control stack (Raspberry Pi 5 + ESP32) supporting low-latency multi-camera streaming (V4L2), remote teleoperation, and real-time control loops.
- Implemented obstacle detection, emergency override, and dead-reckoning navigation using IMU-corrected trajectories for GPS-degraded environments.

## OPEN SOURCE CONTRIBUTIONS

**Intel OpenVINO Toolkit** Contributor

Implemented the Equal elementwise operator for RISC-V RVV1.0, including fp32 RVV JIT emitter design and vectorization extensions within the CPU backend.

## SYSTEMS BACKGROUND

Coursework and independent work covering computer networking (link to transport layers), operating systems, distributed systems, and cloud-native deployment. Experience with relational and graph-structured data (SQL, NoSQL, knowledge graphs) and algorithmic foundations through data structures and algorithms. Combined with ML and hardware experience, this supports a complete sensor-to-cloud understanding of intelligent systems.

## EDUCATION

**IIT Madras** — B.S. Electronic Systems (CGPA: 8.43) 2023–2027
**BITS Pilani** — B.Sc.(H) Computer Science (CGPA: 9.87) 2023–2027

## SKILLS

**ML/AI:** multimodal alignment, reinforcement learning, retrieval systems, model compression, PyTorch
**Signal/Vision:** sonar imaging, low-SNR enhancement, VLMs, feature extraction
**Hardware:** Verilog, SystemVerilog, RISC-V ISA, FPGA prototyping, RTL-to-GDSII
**Systems:** computer networks, distributed systems, cloud deployment, Linux environments
**Data/Algorithms:** DSA, SQL/NoSQL databases, knowledge graphs
**Tools:** ROS2, Vivado, Quartus, Cadence Genus/Innovus, MATLAB, Bash
**Languages:** Python, C++, C, Java