GRADUATE RESEARCH OPENINGS

Smart Systems, Built with Modern Multimodal AI

Vision–Language Models · Agentic AI · Cross-Modal Learning · 3D Perception · Edge Deployment

We are a **smart systems design lab**: we take the newest techniques in AI — vision–language models, cross-modal learning, LLM agents, 3D perception — and engineer them into **complete intelligent systems that real end users depend on**. Our current flagship teaches **millimeter-wave radar to see**, using vision foundation models as teachers and aligning radar with **language**, deployed on real robotic platforms. Anyone can fine-tune a model on a crowded RGB benchmark; here you build the whole system — sensing, learning, reasoning, language, deployment — and you can be **first** at something rather than the ten-thousandth.

WHAT WE WORK WITH

VISION–LANGUAGE MODELS	FOUNDATION-MODEL DISTILLATION	POINT-CLOUD TRANSFORMERS	LLM FINE-TUNING
BEV / OCCUPANCY PREDICTION	AGENTIC AI & TOOL USE	RAG & SPATIAL MEMORY	EDGE AI (JETSON)

PyTorch · OpenCV · Hugging Face · KPConv & U-Net architectures · OneFormer / SAM-class teachers · mmWave radar + RGB-D + IMU sensor suites · on-device Llama-class models

IEEE MASS 2026
regular paper accepted (29.7% acceptance rate)

2 systems
RESA indoor radar guidance · MyPath city platform (beta)

800+
multimodal recording sessions (radar · RGB · depth · IMU)

Deployed
real hardware + live beta, with city & univ. partners

*Led by **Dr. Vaskar Raychoudhury** — Humboldt & DAAD Fellow · Sr. Member, ACM & IEEE · 2,000+ citations · Graduate Program Co-director*

WHAT THE RESEARCH LOOKS LIKE

Two flagship systems anchor the lab. **RESA** (IEEE MASS 2026) turns raw 77 GHz radar into real-time spoken indoor guidance — point-cloud segmentation, free-space U-Nets, occupancy memory, an on-device LLM — trained by *cross-modal supervision*: a vision foundation model teaches the radar to see. **MyPath** (v1.0 beta, with UMBC) is a community-powered accessibility platform: accessible routing plus a living city map that every journey improves. The next generation of both is **language-first**: assistants you simply ask — “Is the hallway to my left clear?” “Which entrance works for my chair?”

[wheelchair-mounted sensing rig — 77 GHz radar, RGB-D camera, Jetson Orin Nano]

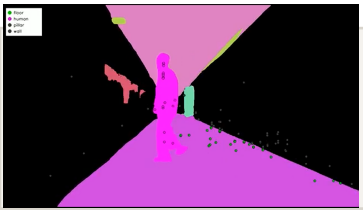

[Object Detection using Radar Point Clouds and Depth Camera Images]

Models that ship to real users

Our systems power accessible navigation for wheelchair users and people with low vision, developed with the Miami University Center for Assistive Technology within a federally funded research program. Your thesis models don’t stop at a benchmark — they run on hardware, in cities, for people — and that deployment story is what reviewers consistently single out.

Open Thesis Directions

Every direction below follows the lab's pattern: a modern AI technique, engineered into a real deployed system, scoped to become an MS thesis with a publication target — and each can grow into a PhD track. You start from a working system, a calibrated multimodal dataset, and hardware already on the bench. Directions 1–5 center on RESA; page 3 opens the city-scale track built on MyPath.

1

Vision–Language Models for Spatial Understanding

Ground language in 3D scene representations built from radar: scene question answering ("What's ahead of me?"), radar-token alignment with LLMs, and instruction-following spatial dialogue. Adjacent to VQA and multimodal LLM research — on a modality where almost nothing exists yet.

2

Cross-Modal Learning & Foundation-Model Distillation

Use vision foundation models (OneFormer, SAM/DINOv2-class) as teachers for radar perception: contrastive radar–vision pretraining, cross-modal label transfer, and representation alignment across calibrated sensor pairs.

3

Generative AI at the Edge

Distill large language models into small, on-device models for real-time spoken guidance: synthetic instruction-dataset generation, fine-tuning, and rigorous evaluation of safety-critical language (beyond BLEU — rubric, judge-model, and human-comprehension evaluation).

4

Agentic & Embodied AI

Build tool-using AI agents over a live perception stack: an assistant that queries occupancy memory, re-scans free space, maintains long-term spatial memory with retrieval (RAG over places), and decides when to speak — with a deterministic safety layer holding veto power.

5

3D Perception & Occupancy

Point-cloud transformers, BEV representations, dense free-space prediction, and traversability estimation ("can a 70 cm-wide platform fit through that gap?") from sparse, multipath-prone returns — core 3D vision on genuinely hard data.

BEYOND NAVIGATION — MORE ACTIVE RESEARCH LINES

Not every thesis here is about navigation. The lab also runs active lines in **multi-agent reinforcement learning for urban logistics** (QuikDel, DeliverAI, Last Mile — IEEE T-ITS, IEEE Access, IJCNN), **federated learning** (FedAccess, COMPSAC 2024), **open-world & continual learning** (ACM Computing Surveys — 218 citations), and **quantum computing** for delivery optimization. If one of these is your direction, there is a thesis-ready problem and a publication trail waiting.

Why a systems lab, not another benchmark thesis?

First-mover problems. On radar–language models, radar-grounded VLMs, and city-scale accessibility AI, a strong MS thesis can be the first published work of its kind — not an increment on a saturated benchmark.

Everything is ready. Two working systems (RESA and MyPath), calibrated multimodal datasets, open repositories, a live beta with real users, and hardware on the bench from day one.

A publication track record. Theses are scoped around concrete venues (IEEE MASS, PerCom, sensing, multimodal-AI, and accessibility venues), with follow-up papers already mapped.

Skills employers read. PyTorch, transformers, VLMs, LLM fine-tuning, RAG/agents, geospatial AI, and edge deployment — plus the rarer skill of integrating them into a working system, which is what industry actually hires for.

From Campus to City Scale: MyPath

MyPath is the lab's outdoor, city-scale system: a community-powered accessibility platform (with UMBC, now in invite-only beta) that combines wheelchair-aware accessible routing with live, community-verified barrier data — a *living map* where every user journey improves the information for the next person, and becomes actionable intelligence for cities. It is the same lab pattern at a different scale: modern AI, engineered into a real system, for real end users. And its research roadmap is where several of the most exciting open problems live:

City-Scale Geospatial Vision

Vision and vision–language models that audit aerial, satellite, and street-level imagery for curb cuts, ramps, surface quality, and path widths — producing a full accessibility baseline for an entire district without street-by-street surveys. Segmentation, VLM-based visual question answering ("is this entrance step-free?"), and remote sensing meet a problem cities will actually pay attention to.

An Agentic Accessibility Copilot

A tool-using LLM agent with retrieval over the living map: it plans multi-stop accessible journeys, checks barriers before departure, matches routes to a user's individual mobility profile, and answers plain-language accessibility questions. Opt-in by design, grounded in verified map data — an agent architecture with real users waiting for it.

Trustworthy Crowd Intelligence

Community reports conflict, decay, and vary in reliability. Open problems: truth inference over crowdsourced barrier reports, reporter-reliability modeling, fusing human reports with sensor evidence, and *barrier forecasting* — spatiotemporal ML that predicts elevator outages and surface degradation before users encounter them.

The Indoor–Outdoor Continuum

Unify the lab's two systems: MyPath routes you across the city, RESA's radar guidance takes over inside the building, one conversational assistant spans both. Cross-system handoff, a shared user profile, privacy-preserving mmWave sensing of public barriers — a seamless accessible navigation continuum no other group can currently build.

Interested? Come see them run.

The fastest way to know if this is your lab is a 15-minute demo — the wheelchair rig speaking live guidance, and MyPath's living map. Email me or stop by; bring your transcript and one paper you found exciting.

Students present their work internationally — recent venues include Osaka, Melbourne, and Guangzhou; IEEE MASS 2026 is in Hong Kong this October. And you inherit the lab's network: UMBC, UW–Milwaukee, Dayton Children's Hospital, and Miami Biology.

Open to MSCS thesis students and PhD applicants; MS-to-PhD transitions welcome.

Dr. Vaskar Raychoudhury
Humboldt & DAAD Fellow · Sr. Member, ACM & IEEE
Graduate Program Co-director, CSE
Miami University, Oxford, OH

Email: raychov@miamioh.edu
MyPath: mypath@miamioh.edu

Scan for code, papers & the RESA repo:
github.com/MU-Smart

SMART SYSTEMS · GRADUATE RESEARCH & STUDENT OUTCOMES

Our Students Publish. Then They Launch.

Graduate research outcomes, 2018–2026 — Dr. Vaskar Raychoudhury, Miami University

In this lab, a graduate degree is a research career launchpad: students first-author papers at IEEE journals and international conferences, and leave with the publication record and deployed-systems portfolio that R1 PhD programs and top employers select for.

24+	10+	3	Meta · EY
papers co-authored with recent grad students since 2018	first-authored by students (IEEE Access, IEEE TAI, MASS…)	alumni advanced to PhD programs at R1 universities	and other top employers hired the rest

EIGHT RECENT STUDENTS, AT A GLANCE

Haoxiang (Steven) Yu
MS, Computer Science

8 papers — first author of an **IEEE Trans. on Artificial Intelligence** journal article and an ICDCN paper; co-author on IEEE Access, COMPSAC, PAAMS, MobiQuitous.

→ **PhD, UT Austin → now at Meta**

Nadim Mahmud
MS, Computer Science

6 papers in two years — first author at **ECML PKDD 2026**; core author on RESA (IEEE MASS 2026), OmniPath (COMPSAC 2026), SmartComp 2026, MASS 2025, and an AAAI workshop paper.

→ **PhD student, University of Cincinnati**

Rochishnu Banerjee
Graduate researcher

4 papers — first author at **IEEE COMPSAC 2024** (FedAccess) and **IEEE MSN 2022**; co-author on MobiQuitous 2023 and RESNA/Assistive Technology.

Md Nahid Hasan
MS, Computer Science

First author of **QuikDel** (IEEE Access, 2026), a hierarchical multi-agent Q-learning delivery system, plus a multi-agent systems preprint under review.

→ **PhD student, UNC Charlotte**

Valeria Mokrenko
BS/MS researcher

First author at **IEEE COMPSAC 2021** (transfer learning for surface detection); co-author of the ACM BuildSys 2019 paper that seeded the MyPath surface-sensing line.

→ **EY (Ernst & Young)**

Robert Kilgore
MS, Computer Science

Co-author of **QuikDel** (IEEE Access, 2026); reinforcement learning and routing systems for urban delivery.

Shrawani Silwal
MS, Computer Science

3 papers — first author at **IEEE MASS 2020** (particle-swarm taxi ride-sharing) and **IEEE SMARTCOMP 2019** (ride-sharing architectures survey); co-author at ICDCN 2020.

→ **World Fuel Services (Fortune 500)**

John Hata
BS/MS researcher

3 papers — first author at **PAAMS 2022** (crowd evacuation with wheelchair users) and **ACM HotMobile 2019** (CARE: campus-wide accessible routing); co-author of MiamiMapper (ACM Q2SWinet 2019).

→ **Software Engineer, Perforce Software**

The pattern is deliberate.

Every thesis is scoped around a publishable contribution to a real deployed system — venue papers, an open-source footprint, and a story that stands out in PhD and industry applications alike. Across all co-authors, lab students have co-authored 25+ papers since 2018. **Your name is the next card:** thesis projects in radar-grounded vision–language models, agentic copilots, and city-scale geospatial vision are scoped and waiting — raychov@miamioh.edu.

Publication data: Google Scholar / DBLP, August 2026, plus accepted papers not yet indexed (IEEE MASS 2026, ECML PKDD 2026, COMPSAC 2026, SmartComp 2026). Bracketed amber items are placements pending confirmation.

SMART SYSTEMS · UNDERGRADUATE RESEARCH

Undergrads Here Don’t Just Assist. They Author.

Capstones, Summer Scholars & independent research, 2018–2026 — Dr. Vaskar Raychoudhury, Miami University

Undergraduates in this lab work on real systems — radar rigs, city-scale accessibility platforms, quantum and agentic-AI prototypes — and their work routinely ends in peer-reviewed papers, best-capstone awards, and launched careers.

80+	6	2	IEEE MASS
undergraduates mentored: 65 capstone + summer scholars	papers with UG co-authors (2 with UG first authors)	best-capstone awards — and both teams published	ACM BuildSys (CORE A), ICDCN, MobiQuitous & more

TWO AWARD-WINNING CAPSTONES — BOTH PUBLISHED

RESA — Radar Wheelchair Guidance

2ND-BEST CAPSTONE 2026 · IEEE MASS 2026 (29.7% ACCEPTANCE)

Noah Emmert, Demetrius Hullum Scott, Joseph Sanuk, and Ryan Hardek built a radar-only indoor guidance system for blind and low-vision wheelchair users — and are the **first four authors** of the resulting IEEE conference paper.

→ **Demetrius: continuing the research · Noah: medical school · Joseph: Progressive**

MiamiMapper — Wi-Fi Crowd Analysis

BEST CAPSTONE 2019 · ACM Q2SWINET @ MSWIM 2019

Nicholas Jarvis, John Hata, and Nicholas Wayne turned passive Wi-Fi probe monitoring into an indoor crowd-analysis system, published at an ACM workshop (13 citations).

→ **John Hata: MS with the lab → Software Engineer, Perforce — the undergrad-to-career arc in one person**

MORE UNDERGRAD-AUTHORED PAPERS

Haoxiang Yu — UG **first author**, dynamic taxi ride-sharing (ICDCN 2020, 16 citations); later MS → UT Austin PhD → Meta. · **Thomas Nguyen** — UG **first author**, MyPath accessible routing (MobiQuitous 2023). · **Zheng Cao & Ce Zhang** — co-authors, smart surface classification (ACM BuildSys 2019 main track, CORE A, 28 citations). · **Manh Nguyen** — personal-data accessibility survey (PerAwareCity @ IEEE PerCom 2019).

THE 2025–26 SLATE: QUANTUM TO RADAR TO AGENTIC AI

RESA (4 students)	mmWave radar + embedded AI wheelchair guidance — the award-and-publication story above.
Quantum Computing for Food Delivery (2)	Miranda Walker & Colin Kula — quantum approaches to urban delivery optimization, extending the lab’s multi-agent delivery research line.
Multi-Sensor Data Analysis & Visualization Toolkit (5)	Zachary Soryal, Nick Vandegriff, Eliot Seifrit, William Tran & Sam Adu-Gyapong — an AI-agent toolkit for health research data.
Cross-College Biology Collaboration	With Prof. Tomoyasu (Biology): computer-vision detection of planar cell polarity in beetle wings — UG researcher Aabha Phadke (Biology/Pre-med), with capstone alum Nick Vandegriff building the OpenCV pipeline.

Your capstone can be a paper. Your summer can be research.

Capstone teams, Undergraduate Summer Scholars, and independent researchers all plug into the same live systems — radar hardware, the MyPath platform, quantum and agentic-AI prototypes. One project this year spanned computer vision for a biology lab’s microscopy. If you want your name on real work: raychov@miamioh.edu.

Counts: 65 capstone students (2018–2026) plus Undergraduate Summer Scholars (2018, 2019, 2021, 2022) and CADS researchers. Publication data: Google Scholar / DBLP, August 2026.