

BHUVAN KODURU

(412) 657-3872 | Pittsburgh, PA | bkoduru@andrew.cmu.edu | [linkedin](#) | [website](#)

EDUCATION

Carnegie Mellon University, M.S. in AI/NLP (Intelligent Information Systems) Pittsburgh, PA | Dec 2026
GPA: **4.11/4.00** | Coursework: Multimodal Machine Learning, Speech Technology for Conversational AI, Deep Learning
JSS Science and Technology University, B.E. in Information Science & Engineering Mysore, India | Jun 2024
GPA: **9.44/10.00** | Leadership: ML Lead, Google Developer Student Clubs | Editor, Computer Society of India Editorial Board

EXPERIENCE

Simular.ai — *Research Intern* Palo Alto, CA | Jun 2026 – Present
• Advancing **AI safety** and evolving harnesses of computer-use agents; advised by Xin Eric Wang, Jiachen Yang, and Ang Li (publication accepted at **CoLM Lifelong Agents Workshop 2026**; more in progress).
• Reverse-engineering internal representations and failure modes of agentic **vision-language models** to detect and steer away from unsafe behavior before it executes.
• Enabling **real-time intervention** on VLM Computer Use Agents on live desktop environments.

CMU, Robotics Institute — *Graduate Research Assistant* Pittsburgh, PA | Jan – May 2026
• Engineered an evaluation harness simulating **8,000+ multi-agent Among Us games** across **21 LLMs** (vLLM + APIs) to quantify deception and Theory of Mind.
• Formalized a game-theoretic model of strategic deception and released a novel benchmark dataset with a moderator agent.

Adobe Inc. — *Associate Technical Consultant* Bengaluru, India | Aug 2024 – Apr 2025
• Architected a **multi-agent pipeline** on Adobe's orchestration library — sequential agents for requirement extraction, module retrieval, and scenario stitching — automating plaintext-to-Fusion conversion; advanced to **top-24 semifinalists worldwide**.
• Built a Workfront Fusion documentation tool that parsed scenario JSON to flag orphan nodes and generated AI summaries into an auto-populated Word template, cutting documentation time **70%**.

Nosh Robotics — *Computer Vision Intern* Bengaluru, India | Feb – Jun 2024
• Trained and quantized **ResNet, Inception, and MobileNet** classifiers for real-time edge estimation of ingredient cooking levels.
• Applied zero-shot **CLIP** for oil/water pan detection and built a RAG food Q&A chatbot integrating **FAISS** and **ChromaDB**.

Spinal Cord Lab, IIT Hyderabad — *Research Intern* Remote | Oct 2023 – Mar 2024
• Analyzed motion-capture and mobile-camera data across **17 subjects / 23 activities**, using LED-blink sequences for time synchronization.
• Diagnosed and resolved a critical pipeline bug affecting all experiments, then computed per-joint offsets to benchmark estimation accuracy.

Adobe Inc. — *Technical Consultant Intern* Bengaluru, India | May – Jul 2023
• Deployed an AI agent natively inside Workfront Fusion with @-mention commands (generate, edit, caption) calling **ChatGPT, Photoshop, and Stable Diffusion** APIs, reducing campaign turnaround from weeks to hours.

PUBLICATIONS

• Heard but Not Heeded: Paralinguistic Information Encoding and Loss in Audio-Language Models — **B. Koduru**, D. Alharthi, R. Singh, B. Raj (**CoLM 2026**)
• Subliminal Prosody Learning: Auxiliary Emotion Supervision Redistributes Affective Representations Across ALM Layers — **B. Koduru**, D. Alharthi, R. Singh, B. Raj (**ICML Mechanistic Interpretability Workshop 2026**)
• RIVET: Robust Idempotent Voice Attribute Editing — D. Alharthi, **B. Koduru**, B. Raj, R. Singh (**INTERSPEECH 2026**)
• Computer Use with Evolving Software — S. Agashe, G. Gonzalez-Pumariaga, **B. Koduru**, J. Yang, A. Li, X. E. Wang (**Lifelong Agents Workshop @ CoLM 2026**)
• AmongUs-X: Benchmarking Strategic Deception of LLM Agents via Theory-of-Mind Metrics — K. Mirakhor*, **B. Koduru***, A. Kontogiannis, R. Zabounidis, K. Sycara, W. Kim (**NeurIPS 2026**, under review) **equal contribution*

PROJECTS

• **Persona-Induced Response Modulation** — Engineered an activation-steering pipeline inducing confident personas without fine-tuning; **95% uncertainty reduction**, ablated on Llama 3.1 8B & Gemma 2 2B. (*Oct – Dec 2025*)
• **Arabic TTS Fine-Tuning** — Supervised fine-tuning of NeuTTS-Air on a proprietary Arabic dataset (**30% WER**); Prof. Bhiksha Raj's lab. (*Sep – Dec 2025*)
• **FoodieSpot** — **Restaurant Booking Agent** — Built a Llama 3.1 8B agent with tool use, MongoDB reservations, and RAG retrieval. (*Feb 2025*)

HACKATHON WINS

• **Hackmitten** (MIT Thandavapura) — Devised a Portuguese pronunciation scorer using acoustic-waveform similarity & DTW (no STT). *2023*
• **Prayog** (VVCE Mysuru) — Built a real-time yoga pose estimator with joint-angle feedback & a ChatGPT assistant. *2023*
• **SwarmIT** (HPE & GSSIETW) — Developed a hospital scheduler pairing CNN skin-disease diagnosis with traffic-aware allocation. *2023*
• **Hackfest** (T-Hub & ICE Hyderabad) — Crafted a dev-tool suite: code-to-flowchart, complexity metrics & malicious-commit detection. *2023*

SKILLS

Languages: Python, C/C++, SQL

ML / DL: PyTorch, Hugging Face Transformers, fine-tuning (LoRA/SFT), quantization, distributed training, NumPy, Pandas
LLMs & Agents: multi-agent systems, RAG, vLLM, activation steering, mechanistic interpretability, CLIP, FAISS, ChromaDB
Tools: Git, Linux, Docker, CUDA, MongoDB, Flask