

# Yago Romano Martínez

+1 (931)-510-1872 | yagoromano1738@gmail.com  
linkedin.com/in/yago-romano-martinez-3b1249192 | github.com/yagoromano

## SUMMARY

AI Engineer with a track record of shipping agentic AI systems end-to-end, from designing multimodal VLM benchmarks and RL training pipelines to deploying production LLM applications that integrate directly into enterprise workflows. Published research across LLM agents, autonomous vehicle perception, and robotics security.

## EDUCATION

### Tennessee Technological University

Cookeville, TN

M.S. in Computer Science (Concentration in Data Science/ Machine learning), GPA: 3.8

08/2023 – 08/2025

- Relevant courses: Adv. Topics in AI, Deep Learning, ML & Data Mining, Software & Systems Security

B.S. in Computer Science; Minor in Science and Math, GPA: 3.7

08/2019 – 05/2023

## SKILLS

- **Programming:** Python, C++, Java, MySQL, React, Flask
- **ML & AI:** PyTorch, TensorFlow, Hugging Face Transformers, LLMs, Multimodal Systems, YOLOv8, Reinforcement Learning, NLP, Distributed ML, NeMo RL
- **Tools & Systems:** Docker, Google Cloud Platform (GCP), Vertex AI, Google ADK, ROS2, SLURM, Git/GitHub, REST APIs, CI/CD
- **Languages:** English (Fluent), Spanish (Fluent), Japanese (Conversational)

## WORK EXPERIENCE

### HCA Healthcare

Cookeville, TN

AI Engineering Intern

01/2026 – Current

- Built and maintain an internal GUI-agent benchmark platform that containerizes a Linux desktop environment in Docker to evaluate multimodal VLM agents on realistic enterprise HR workflows, with deterministic checkpoint-based scoring and full trajectory logging.
- Developed a hybrid perception system combining DOM extraction for browser-based tasks with OCR fallback for native desktop interactions, enabling visual grounding and action execution for VLM agents across UI environments.
- Extended the platform into a reinforcement learning training environment integrated with NVIDIA NeMo RL, supporting parallel sandbox rollouts, task-shaped reward signals, and continuous multimodal screenshot feedback so agents observe live interface state after every action during training.
- Developed and deployed a context-aware data governance chatbot embedded in a production enterprise web app, featuring LLM-driven grounded SQL retrieval, multi-turn session memory, intent-based question routing, and automated dashboard filter control from natural language queries.

### Tennessee Technological University

Cookeville, TN

Graduate Research Assistant

05/2024 – 08/2025

- Deployed YOLOv8 object detection in CARLA, extending perception coverage to 180° via a three-camera rig and LiDAR-camera fusion, improving distance estimation and braking response time under adversarial driving conditions.
- Uncovered a supply chain vulnerability in SROS2 by crafting a malicious package that exfiltrates credentials at startup without triggering anomalous network traffic, exposing weaknesses in ROS2's package trust model.

### Tennessee Technological University

Cookeville, TN

Graduate Teaching Assistant

01/2024 – 07/2025

- Mentored students in algorithms, ML, and data analysis; designed course assignments using live LLM inference and real-world API datasets to bridge theory and applied practice.

## PROJECTS

### LLMs as Agentic Cooperative Players in Multiplayer Uno

08/2024 – 08/2025

- Architected an LLM-powered cooperative agent for multiplayer UNO that assists a human player in real time by reasoning over game state and recommending context-aware moves.
- Scaled evaluation to 10,000 distributed games by extending RLCARD UNO with PyTorch Distributed, DeepSpeed, and SLURM, yielding a statistically significant 4.2% win-rate improvement over baseline.

## PUBLICATIONS

- Romano Martinez Y “Mitigation of Camouflaged Adversarial Attacks in Autonomous Vehicles.” *ISCAS* 2025.
- Romano Martinez Y, Roberts J. “LLMs as Agentic Cooperative Players in Multiplayer UNO.” *ICMLA* 2025.
- Hasan Sakib T, Romano Martinez Y. “Supply Chain Exploitation of Secure ROS 2 Systems.” *MILCOM* 2025.