Ahmed Loughzali

Paris (open to relocation) • ahmed.loughzali@gmail.com • +33 7 80 77 74 94 • Website • GitHub • LinkedIn

MSc student in AI & Robotics, seeking end of studies 6-month research internship, starting January-March 2027.
Research interests: Embodied AI, representation learning and multimodal perception (text, audio, vision).

Education

Gwangju Institute of Science and Technology (GIST) – Exchange Semester – 3D Vision & Robotics *Fall 2026*
Gwangju, South Korea

Institut Mines-Télécom Nord Europe (Réseau IMT) – MSc in Computer Science – AI & Data Science *2022 – 2027*
Lille, France

Publications *(NeurIPS 2026 workshops)*

- *Adversarial Attacks Expose Path-Dependent Failures in Discrete Medical Image Tokenizers* – **first author**, accepted at **GenAI4Health** (work conducted at Algoverse)
- *Multi-Party Backchannel Prediction: a Diagnosis, a Benchmark, and a Ceiling* – **second author**, accepted at **ReMuCAI** (work conducted at Enchanted Tools)

Research & Industry Experience

Algoverse AI Research Fellowship – Research Fellow *May 2026 – September 2026*
Remote

- Studied adversarial robustness of discrete tokenizers in medical VLMs; first-author paper at GenAI4Health

Enchanted Tools (Humanoid startup) – AI & HRI Research Intern *April 2026 – August 2026*
Paris, France

- Designed turn-taking for Mirokaï, a social humanoid, fusing acoustic and linguistic cues: benchmarked VAP and SmartTurn, then combined the best with our fine-tuned multilingual Qwen 0.5B for semantic end-of-turn detection
- Built a custom lightweight barge-in detection model (synthetic data generation, training, quantization) deployed in real time on Jetson Orin NX under strict latency constraints
- Researched backchanneling in multi-party settings; co-authored a paper on the findings

Ferromobile (Autonomous car startup) – AI Engineer Intern *May 2025 – August 2025*
Paris, France

- Built and optimized a real-time TTS service (FastAPI, Piper TTS) on NVIDIA Jetson Xavier for autonomous vehicle voice guidance, achieving $\sim$0.6s latency
- Designed a local RAG pipeline (LangChain, ChromaDB, Mistral 7B) auto-generating validation test plans from requirements; deployed both services as Docker Compose stacks on embedded hardware

Selected Projects

Florigen Robotics × GIST Autonomous Driving Lab – Research Project – ongoing *Sept 2026 – Present*
South Korea

- Leading research on an JEPA world action model for robot collision avoidance: synthetic Isaac Sim data pipeline, model training, and benchmarking against diffusion-based baselines.

Surgical-JEPA – World Models Hackathon *June 2026*

- Adapted EB-JEPA into an action-conditioned latent world model for dVRK surgical video: DINOv3 encoder with a causal Transformer predictor conditioned on bimanual actions, trained via multi-GPU DDP
- Cut autoregressive drift to a third of the GRU baseline and improved 2-second LPIPS by $\sim$15%

OncoVIA – Mistral AI x GE Healthcare Hackathon – **1st place** *February 2026*

- Multimodal oncology pipeline combining 3D computer vision and agentic RAG to help radiologists generate accurate, patient-specific follow-up reports

Skills & Interests

Programming & Tools: Python, C++, Bash, Git, Docker, SLURM, gRPC, FastAPI
Robotics & Simulation: ROS 2, MuJoCo, Isaac Sim/Gym
ML & Edge AI: PyTorch, Hugging Face, TensorRT, quantization, NVIDIA Jetson (Orin/Xavier/Thor)
Languages: Arabic & French (native), English (C1), Spanish (B1), Russian (A2)
Interests: Hiking, reading and wrestling