

Yifan Wang

✉ yifanwang.career@gmail.com | ☎ +44 7900610965 | London, UK

🌐 [linkedin: yifan-wang-58203a2a8](https://www.linkedin.com/in/yifan-wang-58203a2a8) | 🌐 <https://yifan-wang.yifanwang-career.workers.dev/>

PROFILE

MSc AI student at UCL with experience in LLM agent workflows, embodied AI, and reliability analysis, with an interest in interpretable and trustworthy AI systems.

EDUCATION

University College London

MSc Biomedical and Healthcare Artificial Intelligence Sep 2025 – Sep 2026
Reinforcement Learning · Deep Representations · NLP · Responsible AI · Probabilistic Modelling

University of Nottingham

BEng (Hons) Electrical and Electronic Engineering Sep 2021 - Sep 2025
Artificial Intelligence · Robotics and Control · Digital Communications · Engineering Software Design

WORK EXPERIENCE & PROJECTS

Vision-Language-Action (VLA) & Embodied AI Interpretability

Affiliate | Holistic AI Jan 2026 – Present

- **Evaluated VLA reliability** by analysing policy behaviour, action consistency, and failure modes across SmolVLA, Pi0.5, and Qwen-VLA in embodied manipulation tasks.
- **Built closed-loop testbeds** in MuJoCo and NVIDIA Isaac Sim to evaluate VLA robustness, task stability, and **sim-to-real transfer** to the SO-ARM101 manipulator.
- **Analysed internal failure signals** through interpretability-driven experiments, linking multimodal perturbations to downstream action errors and policy degradation.

Reward Modelling and DPO for Retrieval-Augmented Medical Reasoning

Research Project | Group Member Dec 2025 - Present

- **Built evidence-grounded medical QA pipeline** combining retrieval, evidence compression, reward-model verification, and multi-step reasoning over MedQA, MedMCQA, and PubMedQA.
- **Constructed DPO preference pairs** by generating Llama-based reasoning traces and scoring them with a process reward model for reward-guided post-training.
- **Improved answer reliability** by combining DPO, retrieval, and reward-model reranking, increasing accuracy from 0.6340 to 0.7867 and improving calibration across Brier, NLL, and overconfidence metrics.

Go-To-Market Multi-Agent Analysis System

Independent Project Mar 2026 – Apr 2026

- **Built AI market intelligence platform** for SWOT analysis, market research, executive reporting, and AIGC-assisted business insight generation.
- **Implemented hybrid RAG system** with real-time web search, dense retrieval, source grounding, and post-generation citation checking to improve factual reliability.
- **Deployed full-stack backend** with FastAPI, PostgreSQL, Redis, Docker, JWT authentication, and monitoring, reducing report generation time from 45 minutes to 3.2 minutes and P95 latency from 8.2s to 350ms.

SKILLS

Programming: Python, C++

ML & AI: LLM applications, RAG, RL, reward modelling, mechanistic interpretability, evaluation design

Frameworks and Tools: PyTorch, Transformers, FastAPI, MuJoCo, NVIDIA Isaac Sim, Git, Docker

Systems: Retrieval pipelines, agent workflows, experiment design, monitoring, backend deployment