

LISA ALAZRAKI

lisa.alazraki20@imperial.ac.uk • [Website](#) • [Scholar](#) • [GitHub](#) • [Linkedin](#) • [EU Citizen, UK Settled Status](#)

PhD student in ML/NLP at Imperial College London. Research interests: generalisable learning and reasoning, agent systems, robustness.

EDUCATION

Imperial College London , <i>PhD Computing</i>	Supervisor: Marek Rei	2027
Imperial College London , <i>MSc Computing (AI and Machine Learning)</i>	Classification: Distinction	2021
The Open University , <i>Grad. Cert. Theoretical Statistics and Probability</i>	Classification: Distinction	2020
The Open University , <i>BSc (Hons) Computing & IT and Mathematics</i>	Classification: 1st Class	2019

Scholarships and awards: IX-WAI Early Career Development Grant 2025 • Imperial Computing Conference 2024 Poster Competition First Prize IET Research Awards 2024 – Postgraduate Prize • Turing Institute Placement Award 2024/25 • Sir Richard Stapley Trust Annual Grant 2024 Imperial College Trust Grant 2023 • IET Travel Award 2023 • Sir Richard Stapley Trust Annual Grant 2023 • CogMI 2022 Best Student Paper EPSRC Doctoral Scholarship 2022 • Imperial College London Distinguished MSc Dissertation Award 2021 • DeepMind MSc Scholarship 2020/21 Open University Official Commendation from the Faculty of Maths, Computing and Technology 2017 • Leslie Walshaw Award 2016 in Mathematics

EXPERIENCE

Meta , <i>Research Scientist Intern</i> • London, UK	<i>Jun - Dec 2025</i>
• Manager: Akhil Mathur. Team: Post-training Reasoning, Agents.	
• Studied how LLM agent capabilities scale as a function of FLOPs and completion time horizon.	
• Designed a marketplace-inspired adaptive multi-agent system that coordinates small agents to push the accuracy/cost Pareto frontier.	
Cohere , <i>Research Intern</i> • London, UK	<i>Jun - Dec 2024</i>
• Manager: Max Bartolo. Team: Command Post-training.	
• Developed a reinforcement learning pipeline for reverse engineering human preferences that boosts LLM-as-a-judge evaluation.	
• Investigated implicit learning from mistakes, showing LLMs attain higher accuracy when not shown explicit corrective feedback.	
• Completed two distinct research projects at the same time, resulting in first-author papers at NeurIPS and EMNLP respectively.	
Google , <i>Research Intern</i> • Amsterdam, Netherlands	<i>Jun - Sep 2023</i>
• Manager: Thomas Mensink. Team: Perception.	
• Developed a model-ensembling framework for knowledge-intensive VQA that beats SOTA by 5% on Encyclopedic-VQA.	
• Presented the resulting publication at ICBINB at NeurIPS 2023.	
Google , <i>Student Researcher</i> • London, UK	<i>Oct - Dec 2022</i>
Research Intern • Zurich, Switzerland	<i>Jun - Sep 2022</i>
• Manager: Hamza Harkous. Team: Applied Privacy Research.	
• Developed a new pipeline for retrieval-augmented generation of user issues that was deployed to production.	
• Improved recall of existing issues by 10x over the previous model, with comparable semantic accuracy for new issue generation.	
• Granted a global patent as co-inventor of the overall system for navigating user feedback.	

SELECTED PAPERS

Scaling Small Agents Through Strategy Auctions , <i>In review</i>	2026
Lisa Alazraki, William F. Shen, Yoram Bachrach, Akhil Mathur	
Rethinking Rubric Generation for Improving LLM Judge and Reward Modeling for Open-ended Tasks , <i>In review</i>	2026
William F. Shen, Xinchu Qiu, Chenxi Whitehouse, Lisa Alazraki, Shashwat Goel, Francesco Barbieri, Timon Willi, Akhil Mathur, Ilias Leontiadis	
AgentCoMa: A Compositional Benchmark Mixing Commonsense and Math Reasoning in Real-World Scenarios , <i>In review</i>	2025
Lisa Alazraki, Lihu Chen, Ana Brassard, Joe Stacey, Hossein A. Rahmani, Marek Rei	
Improving the OOD Performance of Closed-Source LLMs on NLI Through Strategic Data Selection , <i>EACL 2026 Findings</i>	2025
Joe Stacey, Lisa Alazraki, Aran Ubhi, Beyza Ermis, Aaron Mueller, Marek Rei	
Reverse Engineering Human Preferences with Reinforcement Learning , <i>NeurIPS 2025 (Spotlight)</i>	2025
Lisa Alazraki, Yi Chern Tan, Jon Ander Campos, Maximilian Mozes, Marek Rei, Max Bartolo	
No Need for Explanations: LLMs Can Implicitly Learn from Mistakes In-context , <i>EMNLP 2025 (Oral)</i>	2025
Lisa Alazraki, Maximilian Mozes, Jon Ander Campos, Yi Chern Tan, Marek Rei, Max Bartolo	
Enhancing LLM Robustness to Perturbed Instructions: An Empirical Study , <i>ICLR 2025 BuildingTrust</i>	2025
Aryan Agrawal*, Lisa Alazraki*, Shahin Honarvar, Marek Rei (*Equal contribution)	
Meta-reasoning Improves Tool Use in Large Language Models , <i>NAACL 2025 Findings</i>	2024
Lisa Alazraki, Marek Rei	
How (not) to ensemble LVLMs for VQA , <i>NeurIPS 2023 ICBINB</i>	2023
Lisa Alazraki, Lluis Castrejon, Mostafa Dehghani, Fantine Huot, Jasper Uijlings, Thomas Mensink	

SKILLS

Programming languages	Python, TypeScript, JavaScript, Java, Lua, MATLAB/Octave, Maxima, Solidity, Prolog, Unix/Bash, HTML, CSS
Libraries / frameworks	PyTorch, TensorFlow, Keras, NumPy, Pandas, Scikit-learn, Transformers, NLTK, Jinja2, Matplotlib, React, Flask