

Profile

PhD researcher in Computer Science at the University of Cambridge working on structured and algorithmic machine learning, with a focus on graph and hypergraph neural networks, multi-robot systems and neuro-algorithmic methods. My research develops efficient neural architectures that exploit problem structure while improving parameter efficiency, robustness, explainability and algorithmic guarantees.

Education and Academics

- 2023 – Present PhD in Computer Science at the **University of Cambridge (Trinity College)**
 Supervised by Prof. Amanda Prorok and Prof. Pietro Liò
- Research focus: attentional graph and hypergraph neural networks for multi-agent systems, multi-agent pathfinding (MAPF), certified robustness, causal explainability
 - Teaching: Supervised 3 MPhil projects, TA for Geometric Learning Course and Mobile Robot Systems Course, supervised students for C/C++, Prolog and Logic and Proof
- 2022 – 2023 Master's in Computer Science at the **University of Cambridge (Trinity College)**
- Obtained a **Distinction**; thesis supervised by Dr Petar Veličković and Prof. Pietro Liò
 - Courses: Deep Learning, GNNs, Machine Visual Perception, Advanced topics in ML
- 2019 – 2022 BA Honours in Computer Science at the **University of Cambridge (Trinity College)**
- Finished with **First Class Honours**

Work Experience

- Jun 2023 – Sep 2023 **Research Engineer Intern** at **Google DeepMind**, London
- LLM Lie Detection: Developed and evaluated behavioural probing methods for detecting deceptive/untruthful behaviours within LLMs, using causal analysis and semi-supervised learning
- Tech: Large Language Models (LLMs), behavioural probing, causal analysis, Python*
- Jun 2022 – Sep 2022 **Siri Intern** at **Apple**, Cambridge
- Analysed automatic speech recognition (ASR) confusion patterns and added disambiguation features to Siri
- Tech: Automatic Speech Recognition, NLP, error analysis, ambiguity resolution, Python, Swift*
- Dec 2020 – Apr 2022 **(ML) Part-Time Undergraduate Employee** at **ARM**, Cambridge
- Added support for ARM® microNPU in Apache TVM - An End-to-End Deep Learning Compiler Stack
- Tech: ML Compilers, hardware-aware optimization, edge deployment, Python, C++*
- Jul 2020 – Sep 2020 **ML Intern** at **ARM**, Cambridge
- Added support for TFLite operators in Apache TVM. This included piping TFLite operators to existing intermediate representation compilation code as well as adding new compilation code
- Tech: ML Compilers, hardware-aware optimization, edge deployment, Python, C++*

Publications and Preprints

- NeurIPS 2026 “**PANDA: Prior-guided Attentional Dual-path Architecture**” by [Rishabh Jain](#), Federico Belotti, Stefano Coniglio, Pietro Liò, Stefano Fiorini. **NeurIPS 2026**
- Developed a **dual-path attentional architecture** for directed hypergraph neural networks that can be integrated into different existing graph and hypergraph architectures, consistently improving performance across architectures and datasets, often leading to more **parameter-efficient** and **faster** models.

- NeurIPS 2026 **“Differentiable Retrieval-Augmented Generation for Predicting Cellular Responses to Gene Perturbation”** by Andrea Giuseppe Di Francesco, Andrea Rubbi, [Rishabh Jain](#), Pietro Liò. **NeurIPS 2026**
Co-developed PT-RAG, an end-to-end trainable **retrieval-augmented generation** module for **predicting cellular responses to gene perturbations**. Combined semantic search with cell-state-aware selection and integrated it with different foundational biological models; improved distributional similarity and prediction quality.
- ICLR 2026 **“Pairwise is Not Enough: Hypergraph Neural Networks for Multi-Agent Pathfinding”** by [Rishabh Jain](#), Keisuke Okumura, Michael Amir, Pietro Liò, Amanda Prorok. **ICLR 2026**
Designed a compact hypergraph neural network for modelling complex interactions in multi-agent pathfinding. **Our 1M-parameter model outperformed an 85M-parameter GPT-based baseline while using 100× less training data**, demonstrating parameter and data efficiency.
- AAAI 2026 **“Graph Attention-Guided Search for Dense Multi-Agent Pathfinding”** by [Rishabh Jain](#), Keisuke Okumura, Michael Amir, Amanda Prorok. **AAAI 2026 – Oral**
Developed a **neuro-algorithmic hybrid method for multi-agent pathfinding, outperforming the state-of-the-art algorithmic solvers for the first time**. The method preserves **completeness** and **eventual optimality** guarantees of algorithmic methods, while providing the **scalability** of neural methods, **breaking** the solution quality-solver speed **Pareto frontier**.
- Workshops at ICML, KDD 2023 **“Neural Priority Queues for Graph Neural Networks”** by [Rishabh Jain](#), Petar Veličković, Pietro Liò. **KLR ICML 2023 workshop – best paper award, DLG workshop at KDD 2023 – spotlight**
Master’s Thesis. Designed **differentiable priority-queue** memory modules for GNNs for greater algorithmic reasoning, as shown on CLRS-30 algorithmic reasoning benchmark. We also empirically demonstrate their effectiveness in modelling long-range interactions.
- Workshop at IJCAI 2023 **“Automatic Concept Embedding Models”** by [Rishabh Jain](#). **XAI Workshop at IJCAI 2023**
Developed an explainable image classifier that **automatically discovers visual concepts**, removing the need for manual concept annotations. Combined image segmentation, clustering and neural networks to support interpretable predictions and human correction of concept-level inputs.
- EMNLP 2022 **“Explaining NLP using Logic Explained Networks”** by [Rishabh Jain](#), Gabriele Ciravegna, Pietro Barbiero, Francesco Giannini, Davide Buffelli, Pietro Liò. **EMNLP 2022**
Undergraduate Thesis. Proposed an **explainable text classifier** that expresses individual predictions as human-readable logic rules. Improved explanation faithfulness and sensitivity over LIME; a user study found the rules more useful and user-friendly than feature-importance scores.

Awards & Academic Service

- 2022, 2023 Awarded **Examination Prize** by the College Council of **Trinity College, Cambridge**
- 2022 Elected to a **Senior Scholar** of **Trinity College, Cambridge** as a result of academic performance
- 2021 – 2022 **Secretary** of University of Cambridge Competitive Programming Society (**UCCPS**)
- Part of the Programme Committee for various machine learning conferences: **AAAI, ICLR, NeurIPS**.
 - Part of the Programme Committee for various robotics conferences: **IROS, ICRA, MRS**.

Technical Skills

Languages: Python | C++ | C

ML: GNNs | Hypergraph Neural Networks | LLMs | RAG | ML Compilers | PyTorch | JAX