Rishi Hazra

Research Scientist, Postdoctoral Researcher | FAIR, Meta Superintelligence Labs, London

rishihazra63@gmail.com rishihazra.github.io Scholar RishiHazra (+44) 7405 599904

EDUCATION

Ph.D. in Computer Science 2021 – 2025
Örebro University & WASP, Sweden
Thesis: Neurosymbolic decision-making with LLMs | Supervisors: Luc De Raedt, Pedro Zuidberg Dos Martires
Fully funded by WASP (Wallenberg AI, Autonomous Systems & Software Program), Sweden's largest research program

M.Tech in Artificial Intelligence 2017 – 2019
Indian Institute of Science, Bangalore
Thesis: Active learning in sequence tagging | Supervisor: Ambedkar Dukkipati | GPA: 8.10/10

B.Tech in Electrical Engineering 2013 – 2017
Birsa Institute of Technology, India
GPA: 8.03/10

RESEARCH & PROFESSIONAL EXPERIENCE

Research Scientist, Postdoctoral Researcher Jul 2025 – Present
FAIR, Meta Superintelligence Labs, London, UK

- **AIRA$_3$ (Kaggle Gold Medal):** Co-first author of an autonomous research agent that fine-tuned a 30B reasoning model end-to-end with no human in the loop, placing 8th of ∼4,000 teams in NVIDIA's live Nemotron challenge – the first Kaggle gold medal won by an autonomous system. It also reached gold-level performance in another Kaggle competition translating 4,000-year-old Akkadian clay tablets into English, showing it can generalize across domains. [announcement]
- **AIRA$_2$ (NeurIPS 2026):** Co-first author; identified the bottlenecks limiting AI research agents on MLE-bench and removed them by scaling agents across many GPUs. [pdf]
- **AgentBox (FAIR Engineering Award):** Built the execution infrastructure for running autonomous agents at scale across MSL, the backbone of multi-GPU AIRA experiments.
- **Production Impact:** Applied research agents to production GPU and accelerator kernel generation: 1.25–17× speedups on recommendation and LLM models, 27% lower kernel latency (internal benchmark), weeks-to-hours development.
- **Data Collection, Annotation & Vendor Management:** Led large-scale data collection and preference annotation with external vendors, used for post-training of models and LLM judges to improve the long-horizon ideation, planning, and judging abilities of AIRA$_2$.

Research Scientist Intern Aug 2024 – Feb 2025
FAIR, Meta AI, London, UK

- **AIRA$_1$ (NeurIPS 2025 Spotlight):** Developed research agents for MLE-bench, studying search policies, exploration, and generalization; open-sourced **aira-dojo**. [pdf] [code]
- **SAM 3 (ICLR 2026):** Owned the compositional concepts effort, leading large-scale data collection and building the evaluation for segmenting concepts with attributes and relations. [pdf] [code]

Research Scientist Intern Jul 2022 – Dec 2022
Meta Reality Labs Research, Redmond, USA

- **EgoTV (ICCV 2023):** First author; built a benchmark for verifying compositional multi-step tasks from natural-language descriptions in egocentric video, and proposed a neuro-symbolic grounding model for it. [web] [pdf] [code]

Data Scientist Apr 2020 – Sep 2020
Amazon Alexa-AI, Bangalore, India

- **Alexa Metrics:** Led an investigation into why offline evaluation metrics failed to predict live Natural Language Understanding error rates in production, and designed a better offline metric that tracks live errors, raising offline–online correlation across 4 locales and 2 device profiles over the last 10 model releases.

Research Associate Jun 2019 – Mar 2020
Statistics & Machine Learning Group, IISc Bangalore

- **Networked MARL (AAMAS 2020):** Co-first author; developed networked reinforcement learning agents that learn an emergent language to coordinate traffic signal control, and analysed how the language is grounded. [pdf]

HONORS & AWARDS

- **First Gold Medal in a Live Kaggle Competition achieved by an Autonomous System** 2026
  NVIDIA – Nemotron Model Reasoning Challenge (8^{th} / ∼4,000 teams, $AIRA_3$)
- **Featured in Forbes**, $AIRA_2$: Breaking Bottlenecks in AI Research Agents [article] 2026
- **FAIR Engineering Award**, AgentBox – infrastructure for scaling autonomous research agents 2025
- **Nominee, WASP Best Thesis Award** (annual, Sweden-wide; winner announced Jan 2027) 2026
- **Spotlight Presentation**, NeurIPS 2025 ($AIRA_1$) 2025
- **Top Reviewer**, NeurIPS 2022 and NeurIPS 2024 (top 8%) 2022, 2024
- **Best Project Award**, B.Tech thesis 2017
- **GATE 2017**, Graduate Aptitude Test in Engineering – top 0.1% 2017
  India's national exam for postgraduate engineering admissions

SELECTED PUBLICATIONS

1. **$AIRA_3$: Autonomous AI Research Agents in Live Competition**
   M Josifoski*, K Hambardzumyan*, E Toledo*, R Hazra*, N Baldwin*, M Kuchnik*
   Live Kaggle Gold Medal, 2026 [announcement]
2. **$AIRA_2$: Overcoming Bottlenecks in AI Research Agents**
   K Hambardzumyan*, N Baldwin*, E Toledo*, R Hazra*, M Kuchnik*, M Josifoski* et al.
   NeurIPS 2026 [pdf]
3. **AI Research Preference Models**
   TS Foster*, B Al Omari*, T Fu*, T Mann, C Domond et al.
   Preprint 2026 [pdf]
4. **Training AI Co-Scientists Using Rubric Rewards**
   S Goel, R Hazra, D Jayalath, T Willi, P Jain, WF Shen, I Leontiadis, F Barbieri, Y Bachrach, J Geiping, C Whitehouse
   ICML 2026 [pdf] [dataset]
5. **COvolve: Adversarial Co-Evolution of Large-Language-Model-Generated Policies and Environments via Two-Player Zero-Sum Game**
   A Sygkounas, R Hazra, A Persson, PZD Martires, A Loutfi
   GECCO 2026 [pdf]
6. **SAM 3: Segment Anything with Concepts**
   N Carion et al.
   ICLR 2026 [pdf] [code] [website]
7. **AI Research Agents for ML: Search, Exploration, and Generalization in MLE-bench**
   E Toledo, K Hambardzumyan, M Josifoski, R Hazra et al.
   NeurIPS 2025 (Spotlight) [pdf] [code]
8. **Combining Code Generating Large Language Models and Self-Play to Iteratively Refine Strategies in Games**
   Y Bachrach, E Toledo, K Hambardzumyan, D Magka et al., R Hazra
   IJCAI 2025 [pdf]
9. **LexiCon: Planning under Temporal Constraints in Natural Language**
   P Mantenoglou, R Hazra, PZD Martires, L De Raedt
   NeurIPS 2025 (DB Track) [pdf] [code]
10. **Have Large Language Models Learned to Reason? A Characterization via 3-SAT**
    R Hazra, G Venturato, PZD Martires, L De Raedt
    COLM 2025 [pdf] [code]
11. **REvolve: Reward Evolution with Large Language Models using Human Feedback**
    R Hazra*, A Sygkounas*, A Persson, A Loutfi, PZD Martires
    ICLR 2025 [web] [pdf] [code]
12. **SayCanPay: Heuristic Planning with LLMs using Learnable Domain Knowledge**
    R Hazra, PZD Martires, L De Raedt
    AAAI 2024 [web] [pdf] [code]
13. **EgoTV: Egocentric Task Verification from Natural Language Task Descriptions**
    R Hazra, B Chen, A Rai, N Kamra, R Desai
    ICCV 2023 [web] [pdf] [code]
14. **Deep Explainable Relational Reinforcement Learning: A Neuro-Symbolic Approach**
    R Hazra, L De Raedt
    ECML PKDD 2023 [pdf]
15. **Active^2 Learning: Reducing Redundancies in Active Learning for Sequence Tagging**
    R Hazra, P Dutta, S Gupta, MA Qaathir, A Dukkipati
    NAACL-HLT 2021 [pdf] [code]

16. **Networked Multi-Agent Reinforcement Learning with Emergent Communication**
S Gupta*, R Hazra*, A Dukkipati
AAMAS 2020 [pdf]

Workshop Papers

W1. **LLM-Driven Adaptability or Pre-programmed Efficiency? A Comparative Study for Short Interactions**
T Schreiter*, JV Ruppel*, R Hazra, A Rudenko, M Magnusson, AJ Lilienthal
HRI 2025 (Late-Breaking Report) [pdf]

W2. **Bidirectional Intent Communication: A Role for Large Foundation Models**
T Schreiter*, R Hazra*, J Ruppel, A Rudenko
LLMs in the RO-MAN Age @ IEEE RO-MAN 2024 [pdf]

W3. **Intrinsically Motivated Compositional Language Emergence**
R Hazra, S Dixit, S Sen
ViGIL @ NAACL 2021 [pdf]

W4. **gComm: An Environment for Investigating Generalization in Grounded Language Acquisition**
R Hazra*, S Dixit*
ViGIL @ NAACL 2021 [pdf]

* Equal contribution.

TALKS

- **A Computational Theory Perspective on Reasoning** – invited talks [slides] 2025
  Spring Workshop on Mining and Learning (SMiLe 2025) | University of Trento | UCSB NLP Group
- **SayCanPay: Planning with LLMs** – Workshop on LLMs and Reasoning [event] Apr 2024
  TAILOR – Foundations of Trustworthy AI: Integrating Learning, Optimization and Reasoning
- **Planning & Robotics with LLMs** – invited talk, Bosch Robotics Seminar [slides] Mar 2024
- **EgoTV: Egocentric Task Verification** – Embodied AI Seminar, Meta AI (FAIR) May 2023
- **A (Large) Language Model is All You Need** – position talk, AASS, Örebro University Feb 2023

MENTORSHIP

Master Thesis Co-supervision 2024 – 2025
TU Chemnitz
Jens V Rüppel | with Tim Schreiter

ACADEMIC SERVICE

Session Chair: AAMAS 2022
Program Committee: NeurIPS 2022–26 | ICML 2023–25 | ICLR 2024–25 | KR 2024 | EACL 2023 | TMLR