NICHOLAS EDWARDS

Kolingasse 14-16, 1090 Wien, Austria
nicholas.edwards@univie.ac.at · Personal Website · LinkedIn · Google Scholar

EDUCATION

University of Vienna, PhD Computer Science *2025–present (expected 2028)*
(Supervisor: Sebastian Schuster)
Research Focus: evaluating and improving LLM agents as collaborators (currently coding agents), including when agents should ask for clarification vs. proceed autonomously, and how they handle users interrupting and correcting them mid-task.

University of Edinburgh, MSc Speech & Language Processing *2021–2022*
Result: Distinction

University of Cambridge, BA Linguistics *2018–2021*
Result: First Class

PROFESSIONAL EXPERIENCE

University College London/Google DeepMind July 2025–September 2025
Course Developer (contractor) *London, UK*

- Contributed to the development of Google DeepMind's AI Research Foundations course.
- Developed the majority of educational materials, including coding lab notebooks and articles, for Course 7 (Accelerate Your Model).
- Assisted in developing coding lab notebooks for Course 3 (Design And Train Neural Networks).

University College London September 2024–August 2025
Research Assistant *London, UK*

- Investigated LLMs as autonomous research agents in collaboration with Boston University, leading to **RExBench** (ACL 2026; co-first author), a benchmark evaluating whether coding agents can implement novel AI research extensions.
- Built the evaluation infrastructure for RExBench: orchestrated per-task sandboxed runs (Docker/Apptainer, ephemeral cloud VMs, A100 GPUs) to benchmark 14 agent configurations (2 harnesses × 7 LLM backbones) on 12 expert-curated research-extension tasks.
- Later integrated RExBench into the **Harbor** evaluation framework (Harbor Adapters, NeurIPS 2026).
- Supervisors: Sebastian Schuster (UCL) & Najoung Kim (Boston University)

Cohere January 2024–August 2025
Senior Data Quality Specialist – RAG Data Science (contractor) *London, UK*

- Curated high-quality human-annotated training data to improve the data science-related capabilities of Cohere's LLMs.
- Combined skills in coding (Python, SQL) and writing to create training examples encompassing database exploration and understanding.

Sense Street September 2022–September 2023
NLP Data Scientist *London, UK*

- Worked on all aspects of the machine learning pipeline, from data annotation through to model finetuning, working with state-of-the-art language models for novel classification and information extraction tasks on trader conversations in the financial domain.

PUBLICATIONS

Peer-Reviewed Publications

1. **Harbor Adapters and Harbor-Index: Infrastructure and a Curated Meta-Dataset for Large-Scale Agentic Evaluation**
   Shi, L., ..., **Edwards, N.**, ..., Schmidt, L., & Shaw, A.
   *Advances in Neural Information Processing Systems, Evaluations & Datasets Track (**NeurIPS** 2026)*

2. **Ask or Assume? Uncertainty-Aware Clarification-Seeking in Coding Agents**
   Edwards, N., & Schuster, S.
   *Findings of the Association for Computational Linguistics (**EMNLP** 2026)*

3. **RExBench: Can coding agents autonomously implement AI research extensions?**
   Edwards, N.*, Lee, Y.*, Mao, Y. A., Qin, Y., Schuster, S.†, & Kim, N.†
   *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (**ACL** 2026)*
   *Equal contribution. †Equal senior authorship.

4. **Anaphoric Structure Emerges Between Neural Networks**
   Edwards, N., Rohde, H., & Coxe-Conklin, H.
   *Proceedings of the 45th Annual Meeting of the Cognitive Science Society (**CogSci** 2023)*

Preprints & Under Review

P1. **From Models to Systems: A Survey of Explainability for Tool-Augmented Language Models and AI Agents**
   Roth, B., **Edwards, N.**, Hong, P., Schoenegger, L., & Schuster, S.
   University of Vienna (SSRN preprint), 2026

PRESENTATIONS

Ask or Assume? Uncertainty-Aware Clarification-Seeking in Coding Agents *October 2026*
Poster Presentations, Conference on Empirical Methods in Natural Language Processing (EMNLP Findings) & REALM Workshop, Budapest

RExBench: Can coding agents autonomously implement AI research extensions? *July 2026*
Poster Presentation, Annual Meeting of the Association for Computational Linguistics (ACL), San Diego

RExBench: Can coding agents autonomously implement AI research extensions? *September 2025*
Invited Talk, Prof. Ludwig Schmidt's Research Group, Stanford/UW

ACADEMIC SERVICE

Volunteer: EMNLP 2026
Conference Reviewer: NeurIPS 2026 Evaluations & Datasets Track

TECHNICAL SKILLS

Programming Languages	Python, SQL, Bash, R
Libraries	PyTorch, vLLM, Transformers, Pandas, scikit-learn
Agentic Tools	OpenHands, Harbor
Infrastructure	Git, Docker, Slurm, Modal, AWS, Apptainer

OTHER SKILLS

Natural Languages	English (native), Russian (fluent), German (advanced), French (intermediate)