Lanhao Niu

+44 07962 687742 | s2798445@ed.ac.uk | lanhao.niu.official@gmail.com

Education

- **University of Edinburgh** *09/2025 – 08/2026*
  MSc Cognitive Science Edinburgh, UK
  - **Relevant Courses:** Accelerated Natural Language Processing; Text Technologies for Data Science; Machine Learning Systems; Seminar in Cognitive Modelling; Research Methods in Developmental Linguistics
- **Nanjing Agricultural University** *09/2021 – 06/2025*
  B.Eng. in Data Science and Big Data Technology Nanjing, China
  - **Honours & Awards:** First Prize, Chinese Collegiate Computing Competition (2024); Grand Prize, Jiangsu Province Collegiate Computing Competition (2024)

Publications

C=Conference, S=In Submission

[C.1] Zhimin Hu, **Lanhao Niu**, Sashank Varma (2026). **The Representational Geometry of Number**. Accepted by *CogSci 2026*. arXiv:2602.06843 [cs.CL, cs.AI].

[C.2] Zhimin Hu, **Lanhao Niu**, Sashank Varma (2026). **Representational Geometry Reveals How Context Structures Concept Spaces in Language Models**. Accepted by *ICML 2026 Mech Interp Workshop* (Spotlight).

Research Projects

- **LLMs for Probabilistic Inference** *01/2026 – Present*
  Advisor: Prof. Siddharth Narayanaswamy, University of Edinburgh
  - Designed a data representation framework for probabilistic programming traces and built evaluation metrics to assess LLM output quality, implemented as a modular, reusable codebase in Python.
  - Investigated LLMs as learned proposal mechanisms for probabilistic inference; conducted systematic experiments comparing proposal strategies and analysed failure modes to guide model improvement. This work has also informed my thinking on how sampling-based reasoning could support interpretable decision traces in LLM-based agents.
- **Representational Geometry of Concepts in LLMs** *08/2025 – 05/2026*
  Advisor: Prof. Sashank Varma, Georgia Institute of Technology
  - Designed and ran reproducible experiments to quantify how context and task variation affect internal representations across a large set of LLM architectures, using similarity metrics and visualisation tools.
  - Built an evaluation pipeline comparing cross-model alignment of representations; analysed results to identify stable structure shared across architectures.
  - Demonstrated that both contextual transformation fields and task-specific embeddings share stable relational structure across architectures, providing mechanistic evidence for representational universality in LLMs; contributed to two accepted papers (CogSci 2026; ICML 2026 Workshop Spotlight).

Work Experience

- **Z.AI (Zhipu AI)** [🌐] *05/2024 – 12/2024*
  Research Intern Beijing, China
  - Built and optimised LLM-powered agents for AI search and enterprise knowledge Q&A, integrating RAG pipelines with structured knowledge bases; designed tool-use workflows combining NER, Cypher query generation, and LLM-based reasoning for knowledge graph querying.
  - Fine-tuned LLMs for structured knowledge extraction using LoRA via Llama-Factory, achieving high accuracy on entity-relation triple extraction and SQL generation tasks.
  - Built multimodal document parsing pipelines using Vision-Language Models, designing visual-prompting strategies that significantly outperformed traditional OCR pipelines.
- **DeepLang AI** [🌐] *04/2024 – 05/2024*
  Data Engineering Intern Beijing, China
  - Built scalable data processing pipelines for LLM training dataset construction, developing processors to extract key-value pairs and convert multi-format data into structured pre-training and SFT datasets; applied NER models and tokenisation tools for part-of-speech tagging and proper noun extraction.

Skills

- **LLM:** RAG, SFT (LoRA via Llama-Factory), LangChain, LlamaIndex, Dify; tool-use agent design, prompt engineering
- **Machine Learning:** PyTorch, HuggingFace Transformers, Vision-Language Models, Representational Geometry
- **Engineering:** Python, C++, Git, Docker, Google Cloud, Linux/Bash, Django, Gradio