Chenyang Li

✉chenyang.li@cs.ox.ac.uk Google Scholar

Education

University of Oxford MSc
Advanced Computer Science Oxford, UK • Oct 2026 – Present
Expected October 2027

Queen Mary University of London (Joint Programme) BSc • First Class Honours
Intelligent Science and Technology London, UK • 2022 – 2026

Beijing University of Posts and Telecommunications (Joint Programme) BEng • 3.89/4.0
Intelligent Science and Technology Beijing, China • 2022 – 2026

- Rank: 1/97
- GRE: 332 (Quantitative: 170, Verbal: 162)

Research Interests

LLM Interpretability **LLM Reasoning and Evaluation** **AI Fairness and Safety**

Publications

LexChain: Modeling Legal Reasoning Chains for Chinese Tort Case Analysis
AAAI 2026
Huiyuan Xie*, **Chenyang Li***, Huining Zhu, Chubin Zhang, Yuxiao Ye, Zhenghao Liu, Zhiyuan Liu

*Equal contribution

LexRel: Benchmarking Legal Relation Extraction for Chinese Civil Cases
ACL 2026
Yida Cai, Ranjuexiao Hu, Huiyuan Xie, **Chenyang Li**, Yun Liu, Yuxiao Ye, Zhenghao Liu, Weixing Shen, Zhiyuan Liu

LexKairos: Benchmarking Legal Temporal Capabilities in LLMs
Manuscript under review at ACL Rolling Review
Chenyang Li, Zejia Feng, Yuqin Huang, Yuxiao Ye, Huiyuan Xie

Rethinking Reasoning Quality in Large Language Models through Enhanced Chain-of-Thought via RL
arXiv preprint, 2025
Haoyang He, Zihua Rong, Kun Ji, **Chenyang Li**, Qing Huang, Chong Xia, Lan Yang, Honggang Zhang

Research Experience

Tsinghua NLP Group (THUNLP), Tsinghua University Beijing, China
Research Intern in LLM Interpretability Sep 2026 – Present
Research Intern in Legal AI Sep 2024 – Jun 2026

Mechanistic Interpretability (Ongoing)

- Conducted preliminary activation vector extraction and steering experiments to investigate demographic representations in LLMs.

Temporal Reasoning & Evaluation

- **Led the development of a 9-task benchmark to evaluate LLM temporal capabilities**, spanning temporal knowledge, time modelling, and complex temporal reasoning in **time-sensitive legal settings.**
- Systematically evaluated **8 LLMs across multiple settings**, revealing persistent limitations in temporal knowledge and complex temporal reasoning.

Multi-Step Reasoning & Post-Training

- **Investigated multi-step reasoning in LLMs** using human-annotated reasoning chains from civil court rulings, constructing datasets for **post-training and evaluation**.
- Fine-tuned LLMs using **SFT and DPO** and systematically evaluated reasoning quality across seven dimensions, achieving a **38% improvement** over Llama-3.1-8B on the constructed reasoning benchmark.

Fine-Grained Legal Relation Extraction

- Constructed **LLM-distilled training data** for fine-grained legal relation extraction.
- Applied **SFT** to improve legal argument extraction, **increasing micro-F1 from 0.048 to 0.323** over the InternLM3-8B baseline.

Legal LLM Evaluation — *Collaboration with University of Cambridge*

- Developed evaluation datasets and systematically evaluated **open- and closed-source LLMs** on reasoning over UK legislation and case law.

Deep Learning Lab, Westlake University Hangzhou, China
Summer Intern in AI and Psychology Jul 2024 – Aug 2024

- Analyzed counselling dialogues between clients and LLM/human counsellors to examine differences in **counselling strategies** used by LLMs and human counsellors.
- Developed a **Retrieval-Augmented Generation (RAG)** system to align LLM-generated responses with human counsellors' response patterns.

Skills

Programming
Python, Java, C

ML & LLM
PyTorch, LLaMA-Factory, vLLM, LangGraph, AutoGen

Awards

IET Best Student Award Jun 2026
The Institution of Engineering and Technology (IET)
S**ole recipient** from the joint degree programme's 2026 graduating cohort for outstanding academic achievement in engineering and technology, following institutional nomination and IET approval.

Principal's Prize Jun 2026
Queen Mary University of London
One of 18 recipients **from over 600 graduating students** in the joint programme for outstanding academic achievement.

National Scholarship Nov 2025
Ministry of Education, People's Republic of China
China's highest national merit-based scholarship awarded to the **top 0.4%** of undergraduate students nationwide.

National Second Prize, CUMCM Nov 2024
China Society for Industrial and Applied Mathematics
Top **2%** of over 60,000 teams nationwide in the China Undergraduate Mathematical Contest in Modeling (CUMCM).

Outstanding Bachelor's Thesis Award Jun 2026
Beijing University of Posts and Telecommunications
Awarded for excellence in undergraduate thesis research, with **79 individual theses selected university-wide**.

Selected Technical Project

Transformer Language Model from Scratch — Stanford CS336

- Implemented a **Transformer language model from scratch in PyTorch**, including BPE tokenization, self-attention with RoPE, RMSNorm, SwiGLU, and AdamW, without pre-built torch.nn layers.
- Built an end-to-end **training and generation pipeline**, including learning-rate scheduling, gradient clipping, checkpointing, and autoregressive decoding.

Volunteering

Peer Mentor Nov 2024 – Jun 2026
Queen Mary University of London
• Mentored 4 first-year students on study planning and well-being, providing continuous academic and personal support.

• Delivered two experience-sharing sessions for 70+ first-year students; received highly positive feedback with 87% rating it "Very Helpful".