

XIANZHEN LUO

Harbin Institute of Technology, Heilongjiang, China 150001

@ xzluo@ir.hit.edu.cn 📞 (+86) 188 4556 0861 🌐 Luowaterbi.github.io 🏠 GitHub 🏠 Google Scholar

EDUCATION

Harbin Institute of Technology

Ph.D. Student in Code LLMs

Supervisor: Prof. Wanxiang Che and Assoc. Prof. Qingfu Zhu

- **Honors:** National Scholarship, ACL 2025 Outstanding Paper

Harbin, China

Sep 2022 – Present

Harbin Engineering University

Bachelor in Computer Science

- **Honors:** National Scholarship, 2 x ICPC Regional Silver Medals, 2 x ICPC EC-Final Bronze Medals.

Harbin, China

Sep 2018 – Jun 2022

EXPERIENCE

Kuaishou Technology

Kstar Research Intern

- Developed a Multi-Agent Framework (CVE-Factory) from scratch to scale terminal environments for code security. Fully automated collection, construction, testing, and evaluation without human intervention.
- Cross-validated with expert manual reproduction, achieving over 95% consistency. Supports asynchronous parallelism, generating 215 tasks in under 5 hours with 20 parallel on a single machine (vs. 10 hours per task for experts).
- Scaled to 4k+ high-quality vulnerability repair tasks. The performance of finetuned Qwen3-32B improved 5x, is comparable to Minimax-M2.7 and Claude Sonnet 4. Demonstrated strong generalization on TerminalBench and cross-language settings.

Beijing, China

Aug 2025 – Mar 2026

StepFun AI.

Pretrain Research Intern

- Developed and cleaned full-scale GitHub file & issue data from scratch. Further refined filter rules to remove garbled text, significantly reducing loss spikes.
- Conducted experiments on code data with FIM, Meta Info, MTP, and Focal Loss strategies.
- Explored and validated Scaling Laws for Code. Under certain constraints, prediction error was less than 0.1%.
- Discovered a logarithmic relationship between code compression ratio and downstream code task performance.

Beijing, China

Dec 2024 – Jul 2025

Du Xiaoman (Beijing) Science Technology Co., Ltd.

University-Industry Collaboration Researcher

- Synthesized code SFT data from the perspectives of correctness and detail sensitivity, achieving SOTA on small LLMs.
- Leveraged multilingual programming languages to assist reasoning, yielding consistent improvements across benchmarks and LLMs.
- Achieved 2x lossless acceleration with only 2MB extra storage by exploiting vocabulary distributions from LLM decoding, outperforming prior train-free speculative decoding methods by over 30%.

Beijing, China

Nov 2023 – Sep 2024

FIRST-AUTHOR PUBLICATIONS

[Arxiv](#) **CVE-Factory: Scaling Expert-Level Agentic Tasks for Code Security Vulnerability**

Xianzhen Luo*, Jingyuan Zhang*, Shiqi Zhou*, Rain Huang*, Chuan Xiao, Qingfu Zhu, Zhiyuan Ma, Xing Yue, Yang Yue, Wencong Zeng, Wanxiang Che

ACL 2026 **Scaling Laws for Code: A More Data-Hungry Regime**

Xianzhen Luo*, Wenzhen Zheng*, Qingfu Zhu, Rongyi Zhang, Houyi Li, Siming Huang, Yuantao Fan, Wanxiang Che

ICLR 2026 **How Many Code and Test Cases Are Enough? Evaluating Test Cases Generation from a Binary-Matrix Perspective**

Xianzhen Luo*, Jinyang Huang*, Wenzhen Zheng, Qingfu Zhu, Mingzheng Xu, Yiheng Xu, Yuantao Fan, Libo Qin, Wanxiang Che

🏆 ACL 2025 Outstanding Paper **Turning Trash into Treasure: Accelerating Inference of Large Language Models with Token Recycling**

Xianzhen Luo, Yixuan Wang, Qingfu Zhu, Zhiming Zhang, Xuanyu Zhang, Qing Yang, Dongliang Xu

ACL 2025 **ChartCoder: Advancing Multimodal Large Language Model for Chart-to-Code Generation**

Xuanle Zhao*, Xianzhen Luo*, Qi Shi, Chi Chen, Shuo Wang, Wanxiang Che, Zhiyuan Liu, Maosong Sun

EMNLP 2024 Python is Not Always the Best Choice: Embracing Multilingual Program of Thoughts
Xianzhen Luo, Qingfu Zhu, Zhiming Zhang, Libo Qin, Xuanyu Zhang, Qing Yang, Dongliang Xu, Wanxiang Che

EMNLP 2024 Make Some Noise: Unlocking Language Model Parallel Inference Capability through Noisy Training
Yixuan Wang, Xianzhen Luo*, Fuxuan Wei, Yijun Liu, Qingfu Zhu, Xuanyu Zhang, Qing Yang, Dongliang Xu, Wanxiang Che*

Arxiv Is Compression Really Linear with Code Intelligence?
Shijie Xuyang, Xianzhen Luo*, Zheng Chu, Houyi Li, Siming Huang, Qiufeng Wang, Wanxiang Che, Qingfu Zhu, Shuigeng Zhou*

Arxiv Success is in the Details: Evaluate and Enhance Details Sensitivity of Code LLMs through Counterfactuals
Xianzhen Luo, Qingfu Zhu, Zhiming Zhang, Mingzheng Xu, Tianhao Cheng, Yixuan Wang, Zheng Chu, Shijie Xuyang, Zhiyuan Ma, YuanTao Fan, Wanxiang Che

Arxiv Semi-Instruct: Bridging Natural-Instruct and Self-Instruct for Code Large Language Models
Xianzhen Luo, Qingfu Zhu, Zhiming Zhang, Xu Wang, Qing Yang, Dongliang Xu, Wanxiang Che

OTHER PUBLICATIONS

ACL 2026 (Findings) Format-Adapter: Improving Reasoning Capability of LLMs by Adapting Suitable Format
Dingzirui Wang, Xuanliang Zhang, Rongyu Cao, Longxu Dou, Xianzhen Luo, Yingwei Ma, Qingfu Zhu, Wanxiang Che, Binhua Li, Fei Huang, Yongbin Li

Survey From Code Foundation Models to Agents and Applications: A Practical Guide to Code Intelligence
Core Contributor

Tech Report Step-3 is Large yet Affordable: Model-system Co-design for Cost-effective Decoding
Core Contributor

ACL 2025 Oral OpenCoder: The Open Cookbook for Top-Tier Code Large Language Models
Siming Huang, Tianhao Cheng, Jason Klein Liu, Weidi Xu, Jiaran Hao, Liuyihan Song, Yang Xu, Jian Yang, Jiaheng Liu, Chenchen Zhang, Linzheng Chai, Ruifeng Yuan, Xianzhen Luo, Qiufeng Wang, YuanTao Fan, Qingfu Zhu, Zhaoxiang Zhang, Yang Gao, Jie Fu, Qian Liu, Houyi Li, Ge Zhang, Yuan Qi, Xu Yinghui, Wei Chu, Zili Wang

KDD 2025 Advancing Tool-Augmented Large Language Models via Meta-Verification and Reflection Learning
Zhiyuan Ma, Jiayu Liu, Xianzhen Luo, Zhenya Huang, Qingfu Zhu, Wanxiang Che

ACL 2025 (Findings) ChartEdit: How Far Are MLLMs From Automating Chart Analysis? Evaluating MLLMs' Capability via Chart Editing
Xuanle Zhao, Xuexin Liu*, Yang Haoyue*, Xianzhen Luo, Fanhu Zeng, Jianling Li, Qi Shi, Chi Chen*

COLING 2024 A Survey on Natural Language Processing for Programming
Qingfu Zhu, Xianzhen Luo, Fang Liu, Cuiyun Gao, Wanxiang Che

ACL 2022 (Findings) Inverse is Better! Fast and Accurate Prompt for Few-shot Slot Tagging
Yutai Hou, Cheng Chen, Xianzhen Luo, Bohan Li, Wanxiang Che.

AI Open 2022 Augmented and challenging datasets with multi-step reasoning and multi-span questions for Chinese judicial reading comprehension
Qingye Meng, Ziyue Wang, Hang Chen, Xianzhen Luo, Baoxin Wang, Zhipeng Chen, Yiming Cui, Dayong Wu, Zhigang Chen, Shijin Wang

Arxiv Automated Snippet-Alignment Data Augmentation for Code Translation
Zhiming Zhang, Qingfu Zhu, Xianzhen Luo, Yixuan Wang, Bohan Li, Wanxiang Che

PROJECTS

Huozhi: An Open-Source Universal LLM | 🗣️ 12 ☆ 200 Mar 2023 – May 2023

- Served as a core developer, responsible for the collection and curation of data for code pretraining and post-training.
- Led the collection and organization of Chinese pretraining datasets.

Abacus: A Lightweight Code LLM | 🗣️ 2 ☆ 47 Oct 2023 – Sep 2024

- Abacus with 2.7B parameters, outperforms other Code LLMs with parameters ≤3B such as Stable Code-3B and Granite-3B-Code on both coding and general language tasks.
- Led post-training data construction.