Zihao Jing

Western University, Canada

zihaoj24@gmail.com — Homepage — Github — Google Scholar

I received my M.Sc. in Computer Science from Western University, publishing three first-author papers at NeurIPS, ICLR, and ICML on structure-grounded multimodal reasoning. Prior to my M.Sc., I interned at SenseTime, where I fine-tuned a 100B-parameter LLM for production deployment at Sina Weibo and built the training pipeline behind Piccolo2 (C-MTEB top-1, May 2024). I am seeking industry research roles in multimodal reasoning, LLM post-training, and AI for Science, where I can translate research advances into scalable, production-grade systems.

EDUCATION

Western University, M.Sc. in Computer Science, Ontario, Canada — Sep. 2024–Apr. 2026
Beihang University, B.Eng. in Software Engineering, Beijing, China — Sep. 2020–Aug. 2024

SELECTED PUBLICATIONS

Top-Tier Conferences

[1] **Zihao Jing**, Qiuhao Zeng, Ruiyi Fang, Yan Yi Li, Yan Sun, Boyu Wang, Pingzhao Hu. *Scaling-Aware Adapter for Structure-Grounded LLM Reasoning*. **ICML 2026** [Code] [Paper]

- Proposed scaling-aware patching and a geometry-grounding adapter for structure-grounded LLM reasoning over variable-size spatial graphs.
- Achieved top-1 performance on *17/18* reasoning tasks from Mol-Instruction, RNA-QA, and DNA-Chat benchmarks.

[2] **Zihao Jing**, Qiuhao Zeng, Ruiyi Fang, Yan Sun, Boyu Wang, Pingzhao Hu. *Entropy-Guided Dynamic Tokens for Graph–LLM Alignment in Molecular Understanding*. In **ICLR 2026**. [Code] [Paper] [Poster] [Video][Slices]

- Proposed *EDT-Former*, an entropy-guided dynamic query connector that preserves variable-size substructure information for graph–LLM alignment while keeping the LLM backbone frozen.
- Achieved SOTA on *11/11* MoleculeQA and Mol-Instructions benchmarks and best results on *10/10* MoleculeNet/TDC property tasks under matched settings.

[3] **Zihao Jing**, Yan Sun, Yan Yi Li, Sugitha Janarthanan, Alana Deng, Pingzhao Hu. *Structure-Aware Fusion with Progressive Injection for Multimodal Molecular Representation Learning*. In **NeurIPS 2025**. [Code] [Paper] [Video&Poster]

- Proposed *MuMo*, a foundation-style multimodal molecular model with asymmetric Progressive Injection to stabilize token–structure fusion and mitigate modality collapse under unreliable 3D inputs.
- Ranked 1st on *22/29* downstream tasks with a *+2.7%* average improvement over strong baselines.

Selected Additional Works

[4] Junqin Huang, Zhongjie Hu, **Zihao Jing**, Mengya Gao, Yichao Wu. *Piccolo2: General Text Embedding with Multi-Task Hybrid Loss Training*. **SenseTime Technical Report, 2024** [SenseNova] [Code] [Report]

- Built training/evaluation pipelines and contributed to iterative optimization of a general-purpose embedding model; achieved top-1 on C-MTEB in May 2024.

Additional Co-authored Publications

Co-author on 3 additional papers at **ICML 2026** and **ICLR 2026** (×2).

RESEARCH GRANTS

Digital Research Alliance of Canada, RRG Competition 2026 — Secured 5×A100-80GB GPU-years on Canada's national supercomputing cluster (∼**US$80K** commercial value).

RESEARCH INTERNSHIPS

SenseTime — *LLM Research Intern* — 2023.09–2024.06

- **Text embeddings (Piccolo2):** Trained general-purpose embedding models with multi-task hybrid-loss objectives; built end-to-end training/evaluation pipelines and led iterative optimization of a generative embedding LLM; achieved top-1 ranking on C-MTEB (May 2024).
- **Domain LLM adaptation (100B):** Fine-tuned a 100B-parameter LLM for vertical livestream marketing; drove data/recipe iteration and productionized the model for deployment at Sina Weibo.
- **LLM research, engineering & scaling:** Gained hands-on experience with large-scale pretraining/fine-tuning codebases (SenseNova series), hyperparameter tuning, experiment tracking, and reproducible training workflows on multi-GPU infrastructure.

Jina AI — *AI Research Intern* — 2023.04–2023.09

- **LLM engineering:** Improved RunGPT interface and contributed solutions to the **Llama** open-source ecosystem.
- **Applied LLM analytics:** Implemented LLM-based denoising and sentiment pipeline for Budweiser public-opinion analysis; reduced operational cost by >13%.
- **Model commercialization:** Led evaluation/tuning of a super-resolution model; executed performance testing and produced pricing recommendations.

TECHNICAL SKILLS

- **LLMs:** pretraining, post-training, multimodal alignment, and agent workflow design.
- **Systems and HPC:** Linux, Slurm, Docker, Singularity, Conda/Mamba, Git, and reproducible environments.
- **Distributed Training:** PyTorch DDP/FSDP, DeepSpeed, multi-GPU/node training, profiling, and monitoring.
- **Programming:** Python, C/C++, Java, MATLAB, SQL, Bash, JavaScript, and full-stack development.

SELECTED AWARDS AND HONORS

- **Third Prize**, 13th National Undergraduate Mathematics Competition (National-level) 2021
- **Third Prize**, 32nd Beijing Undergraduate Mathematics Competition (Province-level) 2021
- **Honorable Mention (H Prize)**, Mathematical Contest In Modeling 2020