ZEYU HUANG

zeroyuhuang.github.io/ | Google Scholar Page | zeroy.huang@gmail.com

Education

School of Informatics, The University of Edinburgh — Edinburgh, UK
PhD in NLP; **2026 Apple PhD Scholar in AI/ML** — June 2027 expected

Beihang University — Beijing, China
Master of Electronic Information; GPA: 3.9/4.0; Ranking: 2/86 — June 2023

Beihang University — Beijing, China
Bachelor of Information and Computing Science; GPA: 3.7/4.0; Ranking: 6/87 — June 2020

Experience

AliStar Research Intern, Qwen Pretraining Team — Jul. 2026 - Present
Advisors: Dr. Bo Zheng & Zihan Qiu — Beijing

TopSeed Research Intern, ByteDance Seed Speech Team — Mar. 2026 - Jul. 2026
Advisors: Dr. Wenhao Zhu & Shanbo Cheng — Shanghai

Student Researcher, Google DeepMind Continual Learning Team — Aug. 2025 - Feb. 2026
Advisors: Dr. Marc'Aurelio Ranzato & Dr. Adhiguna Kuncoro — London

Research Intern at INF Technology — Mar. 2024 - Aug. 2024
Advisor: Zili Wang — Remote

Research Intern at BAAI — Sep. 2022 - Jun. 2023
Advisor: Dr. Jie Fu — Beijing, China

Research Intern at WeChat AI — Nov. 2021 - May. 2022
Advisor: Dr. Yikang Shen — Beijing, China

Research Interests

My research focuses on how large language models learn, including (1) learning paradigms: how models learn from corrections, demonstrations, rewards, and context; (2) optimization objectives: how training objectives shape the behavioral patterns models develop; and (3) model architectures: how modularity, routing in MoE models, and components such as attention and residual connections affect training dynamics.

Publications

* denotes equal contribution. Selection rates are relative to submissions in the corresponding track.

Zihan Qiu*, **Zeyu Huang***, Kaiyue Wen*, Peng Jin*, Bo Zheng*,..., Ivan Titov, Dayiheng Liu, Jingren Zhou, Junyang Lin. A Unified View of Attention and Residual Sinks: Outlier-Driven Rescaling is Essential for Transformer Training - **COLM 2026 Oral (24/2,939; 0.82%)**.

Zeyu Huang, Tianhao Cheng, Zihan Qiu, Zili Wang, Yinghui Xu, Edoardo M Ponti, Ivan Titov. Blending Supervised and Reinforcement Fine-Tuning with Prefix Sampling - **ICML 2026**

Yijun Yang*, **Zeyu Huang***, Wenhao Zhu, Zihan Qiu, Fei Yuan, Jeff Z Pan, Ivan Titov. A Controllable Examination for Long-Context Language Models - **NeurIPS 2025 D&B Spotlight (56/1,995; 2.81%)**

Zihan Qiu*, Zekun Wang*, Bo Zheng*, **Zeyu Huang***, Kaiyue Wen, Songlin Yang, ... & Junyang Lin. Gated Attention for Large Language Models: Non-linearity, Sparsity, and Attention-Sink-Free. **NeurIPS 2025 Best Paper Award (3/21,575; 0.014%)**

Zihan Qiu*, **Zeyu Huang***, Bo Zheng, Kaiyue Wen, Zekun Wang, ... & Junyang Lin. Demons in the Detail: On Implementing Load Balancing Loss for Training Specialized Mixture-of-Expert Models - **ACL 2025**

Zeyu Huang, Zihan Qiu, Zili Wang, Edoardo M Ponti, Ivan Titov. Post-hoc Reward Calibration: A Case Study on Length Bias - **ICLR 2025**

Zihan Qiu*, **Zeyu Huang***, Shuang Cheng, Yizhi Zhou, Zili Wang, Ivan Titov, Jie Fu. Layerwise Recurrent Router for Mixture-of-Experts - **ICLR 2025**

Ka Man Lo*, **Zeyu Huang***, Zihan Qiu*, Zili Wang, Jie Fu. A Closer Look into Mixture-of-Experts in Large Language Models. **NAACL 2025 Findings**

Wenyu Du, Tongxu Luo, Zihan Qiu, **Zeyu Huang**, Yikang Shen, Reynold Cheng, Yike Guo, Jie Fu. Stacking Your Transformers: A Closer Look at Model Growth for Efficient LLM Pre-Training - **NeurIPS 2024 Spotlight (326/15,671; 2.08%)**.

Wenyu Du, Shuang Cheng, Tongxu Luo, Zihan Qiu, **Zeyu Huang**, Ka Chun Cheung, Reynold Cheng, Jie Fu. Unlocking Continual Learning Abilities in Language Models - **EMNLP 2024 Findings**

Zihan Qiu*, **Zeyu Huang***, Jie Fu. Unlocking Emergent Modularity in Large Language Models. **NAACL 2024**. **Outstanding Paper Award (6/2,434; 0.25%)**

Zihan Qiu*, **Zeyu Huang***, Youcheng Huang*, Jie Fu. Empirical Study on Updating Key-Value Memories in Transformer Feed-forward Layers. Tiny Papers @ **ICLR 2024**

Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, Zhang Xiong. Transformer-Patcher: One Mistake Worth One Neuron. **ICLR 2023.**

Zeyu Huang, Wenge Rong, Xiaofeng Zhang, Yuanxin Ouyang, Chenghua Lin, Zhang Xiong. Token Relation Aware Chinese Named Entity Recognition. **ACM TALLIP**

Zeyu Huang, Xiaofeng Zhang, Jun Bai, Wenge Rong, Yuanxin Ouyang, Zhang Xiong. Solving Math Word Problems Following Logically Consistent Template. **IJCNN 2023**

Xiaofeng Zhang*, Yikang Shen*, **Zeyu Huang**, Jie Zhou, Wenge Rong, Zhang Xiong. Mixture of Attention Heads: Selecting Attention Heads Per Token. **EMNLP 2022**

Pre-prints

Zeyu Huang, Adhiguna Kuncoro, Qixuan Feng, Jiajun Shen, Lucio Dery, Arthur Szlam, Marc'Aurelio Ranzato. Context Training with Active Information Seeking

Tianhao Cheng*, **Zeyu Huang***, Zihan Qiu, Yu Cheng, Edoardo Ponti, Yinghui Xu, Ivan Titov, Zenglin Xu. The Cancellation Hypothesis in Critic-Free RL: From Outcome Rewards to Token Credits

Honors and Awards

- **2026 Apple PhD Scholar in AI/ML** (20 worldwide)
- **NeurIPS 2025 Best Paper Award** (3/21,575 main-track submissions; 0.014%)
- **NAACL 2024 Outstanding Paper Award** (6/2,434 submissions; 0.25%)
- EPSRC DTA Scholarship (international flexibility) - 2024-2027
- National Scholarship for Graduate Students Award - 2021-2022 - (3/291)
- Outstanding Graduate Student of Beihang University - 2020-2021

Extra-curricular Activities

International Exchange — Casablanca, Morocco
Ecole Centrale Casablanca — Nov 2019, 15 days

International Exchange — Paris, France
CentraleSupélec, Saclay — Feb 2018, 8 days

Student works — Beijing, China
President of the Graduate Student Union of ECPK — Dec. 2020 - Dec. 2021

Orange Lab & Ecole Centrale de Pekin — Beijing, China
Oct 2020 - May 2021 — Group Leader of the Project

- **Federated Recommendation** - A recommendation system trained in Federated Learning manner.
- This project is co-supervised by Orange Lab of Beijing and Ecole Centrale de Pekin.
- The product of this project includes a survey report and deployment of a Federated news recommender

Skills Summary

- **Programming**: Python, C, Shell, JAVA, Latex, MATLAB
- **Languages**: Chinese(Native), English (IELTS 7.5), French (B2)
- **ML/AI Frameworks**: PyTorch, JAX, TensorFlow, PyTorch Lightning, Fairseq, Hugging Face, Flax