

KUN ZHOU 周昆

(+86)15851856250 [◇ francis_kun_zhou@163.com](mailto:francis_kun_zhou@163.com)

[Homepage](#) || [Google Scholar](#)

MAIN RESEARCH INTERESTS

Large Language Models, Natural Language Process, Recommender System

EDUCATION

Renmin University of China

Aug 2020 - July 2024 (Expected)

- Ph.D - *Advisor: Wayne Xin Zhao and Ji-Rong Wen*

Peking University

Aug 2017 - July 2020

- Master - *Advisor: Jingsong Yu*

Hohai University

Aug 2013 - July 2017

- Bachelor

ONGOING PROJECTS

- **1. Co-Leading a Team for Training a Bilingual LLM: YuLan-Chat (玉兰)**

- Release 13B and 65B LLMs that can well follow English or Chinese instructions, outperforming contemporary LLaMA-based bilingual LLMs on Chinese (CEval and AGIEval) and an English benchmark (MMLU).
- Propose an automatic way to augment and sample high-quality bilingual instructions, and a multi-stage tuning strategy for enhancing instruction following.
- Just start pre-training our 12B backbone model based on carefully collected pre-training data.

- **2. Steering LLMs during Inference to Handle Long-tail Questions.**

RECENT PREPRINTS (LARGE LANGUAGE MODEL)

- Wayne Xin Zhao, **Kun Zhou**, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, Ji-Rong Wen. A survey of large language models. Arxiv. *A very comprehensive survey with 85 pages and 610 references.*
- Jinhao Jiang[†], **Kun Zhou**, Zican Dong, Keming Ye, Wayne Xin Zhao, Ji-Rong Wen. StructGPT: A General Framework for Large Language Model to Reason over Structured Data. Arxiv.
- Zhipeng Chen[†], **Kun Zhou**, Beichen Zhang, Zheng Gong, Wayne Xin Zhao, Ji-Rong Wen. ChatCoT: Tool-Augmented Chain-of-Thought Reasoning on Chat-based Large Language Models. Arxiv.
- Beichen Zhang[†], **Kun Zhou**, Xilin Wei, Wayne Xin Zhao, Jing Sha, Shijin Wang, Ji-Rong Wen. Evaluating and Improving Tool-Augmented Computation-Intensive Math Reasoning. Arxiv.
- Yifan Li[†], Yifan Du[†], **Kun Zhou**, Jinpeng Wang, Wayne Xin Zhao, Ji-Rong Wen. Evaluating Object Hallucination in Large Vision-Language Models. Arxiv.

SELECTED PUBLICATIONS AND PREPRINTS (NATURAL LANGUAGE PROCESS)

[†] indicates students mentored by me at RUC.

1. **Kun Zhou**, Yifan Li, Wayne Xin Zhao, Ji-Rong Wen. Diffusion-NAT: Self-prompting Discrete Diffusion for Non-Autoregressive Text Generation. In *Arxiv*.
2. Ze-Feng Gao, **Kun Zhou**, Peiyu Liu, Wayne Xin Zhao and Ji-Rong Wen. Small Pre-trained Language Models Can be Fine-tuned as Large Models via Over-Parameterization. In *ACL 2023*. *Nomination for Outstanding Papers*.
3. Hangyu Guo[†], **Kun Zhou**, Wayne Xin Zhao, Qinyu Zhang and Ji-Rong Wen. Visually-augmented Pretrained Language Models for NLP Tasks without Images. In *ACL 2023*.
4. **Kun Zhou**, Xiao Liu, Yeyun Gong, Wayne Xin Zhao, Daxin Jiang, Nan Duan, Ji-Rong Wen. MASTER: Multi-task Pre-trained Bottlenecked Masked Autoencoders are Better Dense Retrievers. In *ECML-PKDD 2023*.
5. Jinhao Jiang[†], **Kun Zhou**, Wayne Xin Zhao, Ji-Rong Wen. UniKGQA: Unified Retrieval and Reasoning for Solving Multi-hop Question Answering Over Knowledge Graph. In *ICLR 2023*.
6. Wayne Xin Zhao, **Kun Zhou**, Beichen Zhang, Zheng Gong, Zhipeng Chen, Yuanhang Zhou, Ji-Rong Wen, Jing Sha, Shijin Wang, Cong Liu, Guoping Hu. JiuZhang 2.0: A Unified Chinese Pre-trained Language Model for Multi-task Mathematical Problem Solving. In *KDD 2023*.
7. Wayne Xin Zhao, **Kun Zhou**, Zheng Gong, Beichen Zhang, Yuanhang Zhou, Jing Sha, Zhigang Chen, Shijin Wang, Cong Liu, Ji-Rong Wen. JiuZhang: A Chinese Pre-trained Language Model for Mathematical Problem Understanding. In *KDD 2022*. *Deployed in iFLYTEK Zhixue app*.
8. **Kun Zhou**, Yeyun Gong, Xiao Liu, Wayne Xin Zhao, Yelong Shen, Anlei Dong, Jingwen Lu, Rangan Majumder, Ji-Rong Wen, Nan Duan, Weizhu Chen. SimANS: Simple Ambiguous Negatives Sampling for Dense Text Retrieval. In *EMNLP 2022*. *Deployed in Bing search*.
9. **Kun Zhou**, Beichen Zhang, Wayne Xin Zhao and Ji-Rong Wen. Debaised Contrastive Learning of Unsupervised Sentence Representations. In *ACL 2022*.
10. Zheng Gong[†], **Kun Zhou**, Xin Zhao, Jing Sha, Shijin Wang and Ji-Rong Wen. Continual Pre-training of Language Models for Math Problem Understanding with Syntax-Aware Memory Network. In *ACL 2022*.

SELECTED PUBLICATIONS (RECOMMENDER SYSTEM)

-
11. **Kun Zhou**, Hui Wang, Ji-Rong Wen, Wayne Xin Zhao. Enhancing Multi-View Smoothness for Sequential Recommendation Models. In *TOIS 2023*.
 12. Zhipeng Zhao[†], **Kun Zhou**, Xiaolei Wang, Wayne Xin Zhao, Fan Pan, Zhao Cao and Ji-Rong Wen. Alleviating the Long-Tail Problem in Conversational Recommender Systems. In *RecSys 2023*.
 13. Xiaolei Wang[†], **Kun Zhou**, Ji-Rong Wen and Wayne Xin Zhao. Towards Unified Conversational Recommender Systems via Knowledge-Enhanced Prompt Learning. In *KDD 2022*.
 14. Yuanhang Zhou[†], **Kun Zhou**, Wayne Xin Zhao, Cheng Wang, Peng Jiang and He Hu. C2-CRS: Coarse-to-Fine Contrastive Learning for Conversational Recommender System. In *WSDM 2022*.
 15. **Kun Zhou**, Hui Yu, Wayne Xin Zhao and Ji-Rong Wen. Filter-enhanced MLP is All You Need for Sequential Recommendation. In *WWW 2022*. *Selected as the most influential WWW paper by PaperDigest*.²
 16. Shuqing Bian, Xin Zhao, **Kun Zhou**, Xu Chen, Jing Cai, Yancheng He, Xingji Luo and Ji-Rong Wen. A Novel Macro-Micro Fusion Network for User Representation Learning on Mobile Apps. In *WWW 2021*. *Deployed in Tencent Kuaibao online system*.
 17. **Kun Zhou**, Xiaolei Wang, Yuanhang Zhou, Chenzhan Shang, Yuan Cheng, Wayne Xin Zhao, Yaliang Li and Ji-Rong Wen. CRSLab: An Open-Source Toolkit for Building Conversational Recommender System.

²<https://www.paperdigest.org/2023/04/most-influential-www-papers-2023-04/>

In *ACL 2021 Demo Track*.

18. **Kun Zhou**, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang and Ji-Rong Wen. Towards Topic-Guided Conversational Recommender System. In *COLING 2020*.

19. **Kun Zhou**, Wayne Xin Zhao, Shuqing Bian, Yuanhang Zhou, Ji-Rong Wen and Jingsong Yu. Improving Conversational Recommender Systems via Knowledge Graph based Semantic Fusion. In *KDD 2020*. *Selected as most influential KDD papers by PaperDigest.*³

20. **Kun Zhou**, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang and Ji-Rong Wen. S³-Rec: Self-Supervised Learning for Sequential Recommendation with Mutual Information Maximization. In *CIKM 2020*. *Selected as the most influential CIKM paper by PaperDigest.*⁴ *Deployed in Meituan online recommender system.*

AWARDS AND HONORS

2022 Baidu Scholarship (10 PhD Students Worldwide).

2022 MSRA Fellowship (12 PhD Students in Asia-Pacific-region).

2022 Bytedance Scholarship (10 PhD Students in China Mainland).

2022 Baogang Scholarship (12 Students in RUC).

2022 National Scholarship.

2018 Jane Street Electronic Trading Challenges

Ranked 1st

2018 The Data Open Challenges Citadel Datathen

Ranked 2nd

2015 China Undergraduate Mathematical Contest in Modeling

National Second Prize

2015 National Zhou Peiyuan college student mechanics competition

National Third Prize

2015 Jiangsu Province Undergraduate Mechanical Competition

The First Prize

RESEARCH EXPERIENCE

MSRA NLC Group

Apr. 2022 - May 2023

- Pre-trained Language Model for Sentence Representation Learning

- Advisor: *Yeyun Gong and Nan Duan*

iFLYTEK Research

Sep. 2021 - Mar. 2022

- Pre-trained Language Model for Mathematical Reasoning

- Advisor: *Jing Sha and Shijin Wang*

Meituan NLP Center

Dec 2019 - May 2021

- Recommender System

- Advisor: *Fuzheng Zhang and Zhongyuan Wang*

Microsoft XiaoIce Group

Aug 2019 - June 2020

- Open-domain Dialog System

- Advisor: *Kai Zhang and Yu Wu*

SERVICE

Senior PC Reviewer

IJCAI 2021

PC Reviewer *TOIS, TIP, TASLP, TORS, ACL, EMNLP, COLING, NIPS, KDD, WWW, WSDM, AAAI*

³<https://www.paperdigest.org/2023/04/most-influential-kdd-papers-2023-04/>

⁴<https://www.paperdigest.org/2023/04/most-influential-cikm-papers-2023-04/>