

Wangchunshu Zhou

<https://michaelzhouwang.github.io>

Education

ETH Zurich

Ph.D in Computer Science (quitted for launching my startup)
Advisor: Prof. Ryan Cotterell and Prof. Mrinmaya Sachan

Zurich, Switzerland

July. 2022 - April. 2023

Stanford University

Ph.D in Computer Science (quitted due to visa issue)
Advisor: in Rotation

LA, US

September. 2021 - July. 2022

Beihang University

Master in Computer Science and System Engineering
Diplome d'ingenieur (French Engineering Degree)
Advisor: Prof. Ke Xu

Beijing, China

Sep. 2018 - June. 2021

Beihang University

B.S. in Information and Computation Science (Sino-French Engineering School)
Major in Mathematics, Physics, and General Engineering.

Beijing, China

Sep. 2014 - Jul. 2018

Research Interests

My research interest lies in AGI. My current research goal is to bridge the gap between powerful pre-trained large language models and their impact on the lives of humankind, and develop **Language Technology for All**.

To achieve this goal, I identify two major research topics that I'm currently interested in: **content generation** with LLMs and **building LLM systems (Agents) to achieve AGI**.

Content generation with LLMs involves methods to better **control** and **customize** LLMs to generate desirable texts of **arbitrary length** that are **personalized**, **creative**, and **useful**. I am also interested in improving human-AI interaction in AI-assisted writing and building a next-generation human-AI **collaborative** writing platform.

As for AGI research, I am passionate about building algorithms that enable:

- *Data-centric* AI Agents that can *learn from experience* and *self-evolve* after being created and deployed in the wild.
- LLMs that can *self-evolve* or *self-improve* using large-scale web corpus and/or well-designed (multi-modal) workflows. Meanwhile, I also care about the superalignment problem which aims to align superhuman LLMs with human values. In particular, I am interested in making AGI technology and systems benefit everyone, which requires AGI personalization techniques. I shared my vision in this position paper.

I am also interested in or have worked in the following subjects:

- Efficient methods for LLMs [8,9,16,17,21,30] and Low resource NLP [6, 10, 11, 44].
- Generative AI: text generation and its evaluation [1,2-7,12-15,22-25,30,34,36,40,43]
- Language model pre-training and transfer learning for NLP [12,15,32,33,39-42].
- Commonsense reasoning and knowledge-based reasoning [4,12,13,31].
- Multi-modality [20,22,29,37,38].
- AI Agents [43,44,47,50,52,53,55].

Research/Work Experiences

○ Cofounder & CTO at AIWaves Inc.

Research & Development

Leading the R&D team on applications of LLMs on long text generation, HAI, and LLM agents.

HangZhou, China

April. 2023 -

○ Pre-doctoral Young Investigator at Allen Institute for AI

Mosaic team, Mentor: Dr. Ronan Le Bra, Prof. Yejin Choi

Research Internship, working on commonsense reasoning and efficient text generation.

Washington, U.S

July. 2021 - July 2022

○ Research Scientist at Bytedance AI Lab

Multi-modal team, Leader: Dr. Hang Li

Working on multi-modal large language model pre-training.

Beijing, China

October. 2021 - June 2022

○ Research Intern at Bytedance AI Lab

NLP group, Mentor: Dr. Jingjing Xu

Beijing, China

Mar. 2021 - Jun. 2021

Research Internship, working on green Deep Learning.

Research Intern at Ink Lab of USC

Los Angeles, U.S.

Jan. 2020 - Dec. 2020

○ *Ink Lab, Mentor: Prof. Xiang Ren*

Research Internship, working on two projects related to machine commonsense reasoning.

Research Intern at MILA

Quebec, Canada

Mar. 2020 - Jun. 2020

○ *Tang's Group, Mentor: Prof. Jian Tang*

Research Internship, working on Interpretable Natural Language Understanding.

Research Intern at Microsoft Research Asia

Beijing, China

Dec. 2018 - Present

○ *Natural Language Computing Group, Mentor: Dr. Tao Ge, Dr. Furu Wei*

Research Internship, working on natural language generation and efficient pre-trained language models.

Publications

* denotes equal contribution, † denotes corresponding author

- **MIMIR: A Customizable Agent Tuning Platform for Enhanced Scientific Applications [60]**
 - Xiangru Tang, Chunyuan Deng, hanminwang, Haoran Wang, Yilun Zhao, Wenqi Shi, Yi Fung, Wangchunshu Zhou, Jiannan Cao, Heng Ji, Arman Cohan, Mark Gerstein
 - in Proceedings of **EMNLP 2024 Demo**
 - We introduce MIMIR, a streamlined platform offering a customizable pipeline that enables users to leverage both private knowledge and publicly available, legally compliant datasets at scale for agent tuning.
- **Enhancing Data Quality via Training Dynamics from Private Domains for Collaborative Fine-Tuning of Large Language Models [59]**
 - Wanru Zhao, Hongxiang Fan, Shell Xu Hu, Wangchunshu Zhou, Nicholas Donald Lane
 - in Proceedings of **NeurIPS 2024**
 - We present a novel method to do data selection in the setting of collaborative fine-tuning using training dynamics.
- **MIO: A Foundation Model on Multimodal Tokens [58]**
 - Zekun Wang, King Zhu, Chunpu Xu, Wangchunshu Zhou, Jiaheng Liu, Yibo Zhang, Jiashuo Wang, Ning Shi, Siyu Li, Yizhi Li, Haoran Que, Zhaoxiang Zhang, Yuanxing Zhang, Ge Zhang, Ke Xu, Jie Fu, Wenhao Huang
 - arXiv preprint
 - We present HelloBench, a benchmark for long text generation with LLMs.
- **HelloBench: Evaluating Long Text Generation Capabilities of Large Language Models [57]**
 - Haoran Que, Feiyu Duan, Liqun He, Yutao Mou, Wangchunshu Zhou, Jiaheng Liu, Wenge Rong, Zekun Moore Wang, Jian Yang, Ge Zhang, Junran Peng, Zhaoxiang Zhang, Songyang Zhang, Kai Chen
 - arXiv preprint
 - We present HelloBench, a benchmark for long text generation with LLMs.
- **Towards LifeSpan Cognitive Systems [56]**
 - Yu Wang, Chi Han, Tongtong Wu, Xiaoxin He, Wangchunshu Zhou, Nafis Sadeq, Xiusi Chen, Zexue He, Wei Wang, Gholamreza Haffari, Heng Ji, Julian McAuley
 - arXiv preprint
 - We present an overview and roadmap towards lifespan cognitive systems – human-like systems that continuously interact with complex environments – whether simulated digital worlds or human society.
- **PositionID: LLMs can Control Lengths, Copy and Paste with Explicit Positional Awareness [55]**
 - Noah Wang, Feiyu Duan, Yibo Zhang, Wangchunshu Zhou, Ke Xu, Wenhao Huang, Jie Fu
 - in Proceedings of **EMNLP 2024 (Findings)**
 - We propose a novel prompting and fine-tuning method to allow LLMs to better control length and perform copy & paste operations with explicit position awareness.
- **Symbolic Learning Enables Self-Evolving Agents [54]**
 - Wangchunshu Zhou*†, Yixin Ou*, Shengwei Ding*, Long Li, Jialong Wu, Tiannan Wang, Jiamin Chen, Shuai Wang, Xiaohua Xu, Ningyu Zhang, Huajun Chen, Yuchen Eleanor Jiang†
 - arXiv preprint
 - We introduce agent symbolic learning, an pioneering work towards data-centric, self-evolving agents. We consider agents as symbolic networks and design a framework that uses language-based loss, gradients, and optimizers to optimize the prompts, tools, and workflow in an agent system.
- **Weaver: Foundation Models for Creative Writing [53]**
 - Tiannan Wang, Jiamin Chen, Qingrui Jia, Shuai Wang, Ruoyu Fang, Huilin Wang, Zhaowei Gao, Chunzhao Xie, Chuou Xu, Jihong Dai, Yibin Liu, Jialong Wu, Shengwei Ding, Long Li, Zhiwei Huang, Xinle Deng, Teng Yu, Gangan Ma, Han Xiao, Zixin Chen, Danjun Xiang, Yunxia Wang, Yuanyuan Zhu, Yi Xiao, Jing Wang, Yiru Wang, Siran Ding, Jiayang Huang, Jiayi Xu, Yilihamu Tayier, Zhenyu Hu, Yuan Gao, Chengfeng Zheng, Yueshu Ye, Yihang Li, Lei Wan, Xinyue Jiang, Yujie Wang, Siyu Cheng, Zhule Song, Xiangru Tang, Xiaohua Xu, Ningyu

Zhang, Huajun Chen, Yuchen Eleanor Jiang[†], Wangchunshu Zhou[†]

- arXiv preprint
- We release Weaver, a series of foundation models in the domain of creative writing.
- **Agents: An Open-source Framework for Autonomous Language Agents [52]**
 - Wangchunshu Zhou[†], Yuchen Eleanor Jiang[†], Long Li*, Jialong Wu*, Tiannan Wang, Shi Qiu, Jintian Zhang, Jing Chen, Ruipu Wu, Shuai Wang, Shiding Zhu, Jiyu Chen, Wentao Zhang, Ningyu Zhang, Huajun Chen, Peng Cui, Mrinmaya Sachan
 - arXiv preprint
 - We release Agents, an open-sourced language agents framework that incorporates symbolic controls.
- **RecurrentGPT: Interactive Generation of (Arbitrarily) Long Text [51]**
 - Wangchunshu Zhou[†], Yuchen Eleanor Jiang[†], Peng Cui, Tiannan Wang, Zhenxin Xiao, Yifan Hou, Ryan Cotterell, Mrinmaya Sachan.
 - arXiv preprint
 - We propose RecurrentGPT, a framework that enables LLMs to generate arbitrarily long text while remaining consistent and coherent.
- **MAP-Neo: Highly Capable and Transparent Bilingual Large Language Model Series [50]**
 - Ge Zhang, Scott Qu, Jiaheng Liu, Chenchen Zhang, Chenghua Lin, Chou Leuang Yu, Danny Pan, Esther Cheng, Jie Liu, Qunshu Lin, Raven Yuan, Tuney Zheng, Wei Pang, Xinrun Du, Yiming Liang, Yinghao Ma, Yizhi Li, Ziyang Ma, Bill Lin, Emmanouil Benetos, Huan Yang, Junting Zhou, Kaijing Ma, Minghao Liu, Morry Niu, Noah Wang, Quehry Que, Ruibo Liu, Sine Liu, Shawn Guo, Soren Gao, Wangchunshu Zhou, Xinyue Zhang, Yizhi Zhou, Yubo Wang, Yuelin Bai, Yuhan Zhang, Yuxiang Zhang, Zenith Wang, Zhenzhu Yang, Zijian Zhao, Jiajun Zhang, Wanli Ouyang, Wenhao Huang, Wenhui Chen
 - arXiv preprint
 - We release MAP-Neo, a series of bilingual (Chinese and English) LLMs.
- **ML-Bench: Evaluating Large Language Models and Agents for Machine Learning Tasks on Repository-Level Code [49]**
 - Xiangru Tang, Yuliang Liu, Zefan Cai, Yanjun Shao, Junjie Lu, Yichi Zhang, Zexuan Deng, Helan Hu, Kaikai An, Ruijun Huang, Shuzheng Si, Sheng Chen, Haozhe Zhao, Liang Chen, Yan Wang, Tianyu Liu, Zhiwei Jiang, Baobao Chang, Yin Fang, Yujia Qin, Wangchunshu Zhou, Yilun Zhao, Arman Cohan, Mark Gerstein
 - arXiv preprint
 - We release ML-Bench, a benchmark for evaluating the abilities of LLMs and agents for machine learning related tasks.
- **How Many Unicorns Are in This Image? A Safety Evaluation Benchmark for Vision LLMs [48]**
 - Haoqin Tu, Chenhang Cui, Zijun Wang, Yiyang Zhou, Bingchen Zhao, Junlin Han, Wangchunshu Zhou, Huaxiu Yao, Cihang Xie
 - in Proceedings of ECCV 2024
 - We release a safety evaluation benchmark for vision-language models.
- **Prioritizing Safeguarding Over Autonomy: Risks of LLM Agents for Science [47]**
 - Xiangru Tang, Qiao Jin, Kunlun Zhu, Tongxin Yuan, Yichi Zhang, Wangchunshu Zhou, Meng Qu, Yilun Zhao, Jian Tang, Zhuosheng Zhang, Arman Cohan, Zhiyong Lu, Mark Gerstein
 - arXiv preprint
 - We release a shared task on terminology machine translation.
- **Findings of the WMT 2023 Shared Task on Machine Translation with Terminologies [46]**
 - Kirill Semenov, Vilém Zouhar, Tom Kocmi, Dongdong Zhang, Wangchunshu Zhou, Yuchen Eleanor Jiang
 - **WMT 2023 Shared Task**
 - We organized a shared task on terminology machine translation.
- **Interactive Natural Language Processing [45]**
 - Zekun Wang, Ge Zhang, Kexin Yang, Ning Shi, Wangchunshu Zhou, Shaochun Hao, Guangzheng Xiong, Yizhi Li, Mong Yuan Sim, Xiuying Chen, Qingqing Zhu, Zhenzhu Yang, Adam Nik, Qi Liu, Chenghua Lin, Shi Wang, Ruibo Liu, Wenhui Chen, Ke Xu, Dayiheng Liu, Yike Guo, Jie Fu
 - arXiv preprint
 - We release a survey on interactive natural language processing.
- **AutoAct: Automatic Agent Learning from Scratch via Self-Planning [44]**
 - Shuofei Qiao, Ningyu Zhang, Runnan Fang, Yujie Luo, Wangchunshu Zhou, Yuchen Eleanor Jiang, chengfei lv, Huajun Chen.
 - in Proceedings of ACL 2024
 - We propose AutoAct, a framework to automatically synthesize training data for agentic LLM fine-tuning.
- **RoleLLM: Benchmarking, Eliciting, and Enhancing Role-Playing Abilities of Large Language Models [43]**
 - Zekun Wang, Z.Y. Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong

Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhao Huang, Wenhui Chen, Jie Fu, Junran Peng.

- in Proceedings of **ACL 2024 (Findings)**
- We propose datasets and methods to train and use role-playing LLMs.
- **CIF-Bench: A Chinese Instruction-Following Benchmark for Evaluating the Generalizability of Large Language Models [42]**
 - Yizhi Li, Ge Zhang, Xingwei Qu, Jiali Li, Zhaoqun Li, Zekun Wang, Hao Li, Ruibin Yuan, Yinghao Ma, Kai Zhang, Wangchunshu Zhou, Yiming Liang, Lei Zhang, Lei Ma, Jiajun Zhang, Zuowen Li, Wenhao Huang, Chenghua Lin, Wenhui Chen, Jie Fu.
 - in Proceedings of **ACL 2024 (Findings)**
 - We release a Chinese instruction dataset for training and evaluating Chinese LLMs.
- **LoraRetriever: Input-Aware LoRA Retrieval and Composition for Mixed Tasks in the Wild [41]**
 - Ziyu Zhao, Leilei Gan, Guoyin Wang, Wangchunshu Zhou, Hongxia Yang, Kun Kuang, Fei Wu.
 - in Proceedings of **ACL 2024 (Findings)**
 - We propose LoraRetriever, a retrieve-then-compose framework that adaptively retrieves and composes multiple LoRAs according to the input prompts.
- **Struc-Bench: Are Large Language Models Good at Generating Complex Structured Tabular Data? [40]**
 - Xiangru Tang, Yiming Zong, Jason Phang, Yilun Zhao, Wangchunshu Zhou, Arman Cohan, Mark Gerstein.
 - in Proceedings of **NAACL 2024 (Oral)**
 - We propose a benchmark on structured data generation with large language models and methods to improve structured data generation.
- **OpenMoE: An Early Effort on Open Mixture-of-Experts Language Models [39]**
 - Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, Yang You
 - in Proceedings of **ICML 2024**
 - We open-source a framework and a series of MOE models.
- **SmartTrim: Adaptive Tokens and Attention Pruning for Efficient Vision-Language Models [38]**
 - Zekun Wang, Jingchang Chen, Wangchunshu Zhou, Haichao Zhu, Jiafeng Liang, Liping Shan, Ming Liu, Dongliang Xu, Qing Yang, Bing Qin
 - in Proceedings of **COLING 2024**
 - We propose an adaptive pruning method for efficient vision-language models.
- **X²-VLM: All-In-One Pre-trained Model For Vision-Language Tasks [37]**
 - Jiao Sun*, Yufei Tian*, Wangchunshu Zhou*, Nan Xu*, Qian Hu, John Frederick Wieting, Xuezhe Ma
 - in Proceedings of **EMNLP 2023**
 - We conduct a throughout evaluation of the abilities of LLMs on controlled text generation.
- **Evaluating Large Language Models on Controlled Generation Tasks [36]**
 - Jiao Sun*, Yufei Tian*, Wangchunshu Zhou*, Nan Xu*, Qian Hu, John Frederick Wieting, Xuezhe Ma
 - in Proceedings of **EMNLP 2023**
 - We conduct a throughout evaluation of the abilities of LLMs on controlled text generation.
- **Towards a Mechanistic Interpretation of Multi-Step Reasoning Capabilities of Language Models [35]**
 - Yifan Hou, Jiaoda Li, Yu Fei, Alessandro Stolfo, Wangchunshu Zhou, Guangtao Zeng, Antoine Bosselut, Mrinmaya Sachan
 - in Proceedings of **EMNLP 2023**
 - We conduct an analysis of the step-by-step reasoning behavior of large language models.
- **Doolittle: Benchmarks and Corpora for Academic Writing Formalization [34]**
 - Shizhe Diao, Yongyu Lei, Liangming Pan, Tianqing Fang, Wangchunshu Zhou, Sedrick Scott Keh, Min-Yen Kan, Tong Zhang
 - in Proceedings of **EMNLP 2023**
 - We release a new benchmark for the task of academic writing formalization and set up the baselines.
- **Let's Synthesize Step by Step: Iterative Dataset Synthesis with Large Language Models by Extrapolating Errors from Small Models [33]**
 - Ruida Wang, Wangchunshu Zhou[†], Mrinmaya Sachan
 - in Proceedings of **EMNLP 2023 (Findings)**
 - We propose a novel method to use LLMs to synthesize training data for smaller models by extrapolating errors smaller models make on a validation set.
- **To Repeat or Not To Repeat: Insights from Scaling LLM under Token-Crisis [32]**
 - Fuzhao Xue, Yao Fu, Wangchunshu Zhou, Zangwei Zheng, Yang You
 - in Proceedings of **NeurIPS 2023**
 - We investigate the problem of insufficient tokens available for pre-training larger LLMs and the impact of repeating pre-training data.

- **Commonsense Knowledge Transfer for Pre-trained Language Models [31]**
 - Wangchunshu Zhou, Ronan Le Bra, Yejin Choi
 - in Proceedings of **ACL 2023 (Findings)**
 - We propose commonsense knowledge transfer, a self-supervised method for transferring the commonsense knowledge stored in a commonsense knowledge model to a general large language model.
- **Modular Transformers: Compressing Transformers into Modularized Layers for Flexible Efficient Inference [30]**
 - Wangchunshu Zhou, Ronan Le Bra, Yejin Choi
 - in Proceedings of **ACL 2023 (Findings)**
 - We propose to train modularized transformers with multi-grained module replacing. Modularized transformers can be flexibly assembled according to different performance-efficiency trade-offs.
- **Cross-View Language Modeling: Towards Unified Cross-Lingual Cross-Modal Pre-training [29]**
 - Yan Zeng*, Wangchunshu Zhou*, Ao Luo, Ziming Cheng, Xinsong Zhang
 - in Proceedings of **ACL 2023**
 - We propose CCLM, a unified framework for multi-lingual multi-modal pre-training.
- **Don't Group, Just Rescore: A Simpler Alternative to Constrained Beam Search [28]**
 - Yuchen Eleanor Jiang, Tianyu Liu, Wangchunshu Zhou, Kexun Zhang, Vilém Zouhar, Alex Warstadt, Mrinmaya Sachan, Ryan Cotterell
 - We propose dynamic beam rescoring (DBR), a simple yet effective constrained decoding method that only rescore candidates instead of grouping them.
- **EfficientVLM: Fast and Accurate Vision-Language Models via Knowledge Distillation and Modal-adaptive Pruning [27]**
 - Tiannan Wang*, Wangchunshu Zhou*[†], Yan Zeng, Xinsong Zhang, Lei Yu
 - in Proceedings of **ACL 2023 (Findings)**
 - We propose a framework that combines general distillation and task/modal-specific pruning to compress vision-language models.
- **Learning to Predict Persona Information for Dialogue Personalization without Explicit Persona Description [26]**
 - Wangchunshu Zhou[†], Qifei Li, Chenle Li
 - in Proceedings of **ACL 2023 (Findings)**
 - We propose to train a persona predictor to enable dialogue personalization without explicit persona information.
- **Controlled Text Generation with Natural Language Instructions [25]**
 - Wangchunshu Zhou, Yuchen Eleanor Jiang, Ethan Wilcox, Ryan Cotterell, Mrinmaya Sachan
 - in Proceedings of **ICML 2023**
 - We introduce a INSTRUCTCTG, a controlled text generation framework that incorporates different constraints by conditioning on natural language descriptions and demonstrations of the constraints.
- **Automatic Educational Question Generation with Difficulty Level Controls [24]**
 - Ying Jiao, Kumar Shridhar, Peng Cui, Wangchunshu Zhou and Mrinmaya Sachan
 - in Proceedings of **AIED 2023**
 - We propose a constrained text generation framework for automatic educational question generation with difficulty level controls.
- **Predicting Reference-Based MT Metrics Without the Reference [23]**
 - Vilém Zouhar, Shehzaad Dhuliawala, Wangchunshu Zhou, Nico Daheim, Tom Kocmi, Yuchen Eleanor Jiang, Mrinmaya Sachan.
 - in Proceedings of **EACL 2023**
 - We show that it is possible to train a model to predict reference-based MT metrics without the reference, which can benefit quality estimation for machine translation systems.
- **Write and Paint: Generative Vision-Language Models are Unified Modal Learners [22]**
 - Shizhe Diao, Wangchunshu Zhou, Xinsong Zhang, Jiawei Wang.
 - in Proceedings of **ICLR 2023**
 - We show that generative pre-training on image and text modality can be unified and results in a powerful foundation model capable of both vision, text, and vision-language tasks.
- **Efficiently Tuned Parameters Are Task Embeddings [21]**
 - Wangchunshu Zhou*, Canwen Xu*, Julian McAuley.
 - in Proceedings of **EMNLP 2022**
 - We show that inter-task transferability can be efficiently predicted by parameter-efficient fine-tuning on source and target tasks.
- **VLUE: A Multi-Task Benchmark for Evaluating Vision-Language Models [20]**
 - Wangchunshu Zhou*[†], Yan Zeng*, Shizhe Diao*, Xinsong Zhang*.

- in Proceedings of **ICML 2022**
- We release VLUE, a multi-task benchmark for evaluating the “true” transfer/generalize ability of pre-trained vision-language models.
- **Contextual Representation Learning beyond Masked Language Modeling [19]**
 - Zhiyi Fu*, Wangchunshu Zhou*, Jingjing Xu*, Hao Zhou, Lei Li.
 - in Proceedings of **ACL 2022**
 - We present a learning to teach method to improve knowledge distillation with meta learning.
- **BERT Learns to Teach: Knowledge Distillation with Meta Learning [18]**
 - Wangchunshu Zhou*, Canwen Xu*, Julian McAuley.
 - in Proceedings of **ACL 2022**
 - We present a learning to teach method to improve knowledge distillation with meta learning.
- **A Survey on Green Deep Learning . [17]**
 - Jingjing Xu*, Wangchunshu Zhou*, Zhiyi Fu*, Hao Zhou, Lei Li.
 - arXiv preprint
 - We present a systematic review of the development of Green deep learning technologies.
- **Beyond Preserved Accuracy: Evaluating Loyalty and Robustness of BERT Compression. [16]**
 - Canwen Xu*, Wangchunshu Zhou*, Tao Ge, Ke Xu, Julian McAuley, Furu Wei.
 - in Proceedings of **EMNLP 2021 (Oral)**
 - We propose to evaluate BERT/model compression techniques in terms of loyalty and robustness in addition to accuracy.
- **Improving Sequence-to-Sequence Pre-training via Sequence Span Rewriting. [15]**
 - Wangchunshu Zhou, Tao Ge, Canwen Xu, Ke Xu, Furu Wei.
 - in Proceedings of **EMNLP 2021**
 - We introduce a sequence span rewriting objective which bridge the gap between pre-training and fine-tuning of sequence-to-sequence pre-trained models and provide more informative pre-training signals.
- **Learning from Perturbations: Diverse and Informative Dialogue Generation with Inverse Adversarial Training. [14]**
 - Wangchunshu Zhou*[†], Qifei Li*, Chenle Li.
 - in Proceedings of **ACL 2021 (Oral)**
 - We introduce an inverse adversarial learning framework that encourages neural dialogue agents to generate diverse and informative responses that are sensitive to perturbations in the dialogue history.
- **Blow the Dog Whistle: A Chinese Dataset for Cant Understanding with Common Sense and World Knowledge [13]**
 - Canwen Xu*, Wangchunshu Zhou*, Tao Ge, Ke Xu, Julian McAuley, Furu Wei.
 - in Proceedings of **NAACL 2021**
 - We introduce a large and diverse Chinese dataset for creating and understanding cant from a computational linguistics perspective.
- **Pre-training Text-to-Text Transformers for Concept-centric Commonsense. [12]**
 - Wangchunshu Zhou*, Dong-Ho Lee*, Ravi Selvam, Seyeon Lee, Bill Yuchen Lin, Xiang Ren
 - in Proceedings of **ICLR 2021**.
 - We introduce novel self-supervised pre-training objectives in both generative and contrastive form and a joint training framework to improve the commonsense reasoning ability of pre-trained text-to-text transformers and train CALM, a Concept-Aware Language Model.
- **Connecting the Dots Between Fact Verification and Fake News Detection. [11]**
 - Qifei Li*, Wangchunshu Zhou*[†].
 - in Proceedings of **COLING 2020 (Oral)**
 - We propose to connect the dots between Fact Verification, which is better explored, and Fake News Detection, which is less developed, by using a pre-trained summarization model to summarize a news article into a short claim and then use a pre-trained fact verification model to classify the trustworthiness of the news.
- **Towards Interpretable Natural Language Understanding with Explanations as Latent Variables. [10]**
 - Wangchunshu Zhou*, Jinyi Hu*, Hanlin Zhang*, Xiaodan Liang, Maosong Sun, Chenyan Xiong, Jian Tang.
 - in Proceedings of **NeurIPS 2020**
 - We develop ELV, a general framework for interpretable natural language understanding that requires only a small set of human annotated explanations for training.
- **BERT Loses Patience: Fast and Robust Inference with Early Exit. [9]**
 - Wangchunshu Zhou*, Canwen Xu*, Tao Ge, Julian McAuley, Ke Xu, Furu Wei.
 - in Proceedings of **NeurIPS 2020**
 - We propose PABEE, a straightforward yet effective inference method that can be used as a plug-and-play technique to simultaneously improve the efficiency and robustness of a pretrained language model (PLM).

- **BERT-of-Theseus: Compressing BERT by Progressive Module Replacing.** [8]
 - Canwen Xu*, Wangchunshu Zhou*, Tao Ge, Furu Wei, Ming Zhou.
 - in Proceedings of **EMNLP 2020**
 - We propose a novel model compression approach that progressively replace each component in a large model by smaller modules. We apply our proposed approach to compress BERT and achieved state-of-the-art performance in comparable settings.
- **Pseudo Bidirectional Decoding for Local Sequence Transduction.** [7]
 - Wangchunshu Zhou, Tao Ge, Ke Xu.
 - in Proceedings of **EMNLP 2020 (Findings)**
 - We propose a pseudo bidirectional decoding approach which utilize the characteristic of local sequence transduction tasks like grammatical error correction and spell correction.
- **Improving Grammatical Error Correction with Machine Translation Pairs.** [6]
 - Wangchunshu Zhou, Tao Ge, Chang Mu, Ke Xu, Furu Wei, Ming Zhou.
 - in Proceedings of **EMNLP 2020 (Findings)**
 - We propose to use a pair of Machine Translation models with different qualities to synthesize pseudo-parallel data for pretraining Grammatical Error Correction models.
- **Scheduled DropHead: A Regularization Method for Transformer Models.** [5]
 - Wangchunshu Zhou, Tao Ge, Ke Xu, Furu Wei, Ming Zhou.
 - in Proceedings of **EMNLP 2020 (Findings)**
 - We introduce DropHead, a structured dropout mechanism for the multi-head attention mechanism in the transformer models, and a specifically designed dropout rate schedule for DropHead.
- **CommonGen: A Constrained Text Generation Challenge for Generative Commonsense Reasoning** [4]
 - Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, , Pei Zhou, Chandra Bhagavatula, Yejin Choi, Xiang Ren
 - in Proceedings of **EMNLP 2020 (Findings)**
 - We propose a novel constrained text generation task and released a dataset to test the generative commonsense reasoning ability of different NLG models.
- **Self-Adversarial Learning with Comparative Discrimination for Text Generation.** [3]
 - Wangchunshu Zhou, Tao Ge, Ke Xu, Furu Wei, Ming Zhou.
 - in Proceedings of **ICLR 2020**
 - We propose to integrate the self-play mechanism, which is commonly used in the RL community, into training of GANs to reduce the reward sparsity and mode collapse problem and make training more stable.
- **Learning to Compare for Better Training and Evaluation of Open Domain Text Generation Models.** [2]
 - Wangchunshu Zhou, Ke Xu.
 - in Proceedings of **AAAI 2020 (Oral)**.
 - We proposed a novel "Learning to Compare" paradigm and employ the skill rating system, which is commonly used to evaluate human chess players' skill, to evaluate the performance of open domain text generation systems.
- **BERT-based Lexical Substitution.** [1]
 - Wangchunshu Zhou, Tao Ge, Ke Xu, Furu Wei, Ming Zhou.
 - in Proceedings of **ACL 2019**
 - We proposed a novel lexical substitution based on pretrained masked language models (e.g. BERT) to automatically propose substitute candidates and rank them without relying on external lexical resources.

Services

- Area Chair (Action Editor): ACL/ARR: 2024-
- Reviewer: NeurIPS 2020-2024, AAAI 2021-2023, ICML 2021-2024, ACL 2021-2023, EMNLP 2021-2023, ICLR 2022-2024, ACL/ARR: 2021-2024
- Student Volunteer: ACL 2019, ICLR 2020, EMNLP 2020

Awards

- Baidu PhD Fellowship, class of 2022 (10 recipients worldwide)
- XiaoMi Fellowship, Beihang University, 2020 (25/3000)
- National Scholarship (Graduate Student), 2020 (4/240)
- Student Travel Grant: AAAI 2020, ICLR 2020
- Scholarship of Academic Excellence (Master), Beihang University. 2018, 2019 (top 15%)
- The CASC Award, Beihang University. 2017 (3/115)
- Scholarship of Academic Excellence (Bachelor), Beihang University. 2016, 2017 (top 20%)

Hobbies and Miscellaneous

- Video game: FIFA23: Ranked Top50 (worldwide) among all XBOX players, 20/20 wins in the Weekend League; League of Legends; Red Alert 2; Age of Empires 2.
- Reading sci-fic novels and watching sci-fic movies. Huge fan of Harry Potter series and The Wheel of Time series.
- Football (Real Madrid and CR7), watching F1 (Lewis Hamilton and Mercedes) and playing Sim Racing games.