

# Wangchunshu Zhou

<https://michaelzhouwang.github.io>

## Education

---

### ETH Zurich

*Ph.D in Computer Science*

Advisor: Prof. Ryan Cotterell and Prof. Mrinmaya Sachan

Zurich, Switzerland

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 Interest

---

My primary research goal is to bridge the gap between the advances on state-of-the-art numbers on leaderboards and the impact of language technology on the lives of humankind, and develop **Language Technology for All**. To achieve this goal and make language technology accessible in most people's lives, I identify two major research topics that I'm interested in: **efficiency** and **trustworthiness** of NLP models. Efficiency involves both the amount of **computation** and **data** required for (pre-)training and using NLP models. Trustworthiness involves the **interpretability**, **fairness**, and **robustness** with respect to adversarial attacks and out-of-distribution samples, as well as reliable **evaluation** of NLP models beyond numbers on leaderboards.

I am also interested in or working in the following subjects:

- o Green NLP[8,9,16,17], Low resource NLP[6,11], and Learning NLP models from high-level supervision[10].
- o Language model pre-training and transfer learning for NLP[12,15].
- o Commonsense reasoning and knowledge-based reasoning[4,12,13].
- o Robust NLP models for OOD samples and reducing spurious dataset biases.
- o Interpretability[10], biases, and fairness of NLP models.
- o Natural language generation[3,5,14], evaluation for NLG models[2].

## Publications & Preprints

---

\* denotes equal contribution, links in numbers

- o **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.
- o **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.
- o **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.
- o **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.
  - 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.
- o **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.
- o **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.
- o **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.
- o **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.
- o **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.
- o **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.
- o **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.

## Research Experiences

---

- o **Pre-doctoral Young Investigator at Allen Institute for AI** **Washington, U.S**  
*Mosaic team, Mentor: Dr. Ronan Le Bra, Prof. Yejin Choi*  
 Research Internship, working on commonsense reasoning and efficient text generation. *July. 2021 -*
- o **Research Intern at Bytedance AI Lab** **Beijing, China**  
*NLP group, Mentor: Dr. Jingjing Xu*  
 Research Internship, working on green Deep Learning. *Mar. 2021 - Jun. 2020*
- o **Research Intern at Ink Lab of USC** **Los Angeles, U.S.**  
*Ink Lab, Mentor: Prof. Xiang Ren*  
 Research Internship, working on two projects related to machine commonsense reasoning. *Jan. 2020 - Dec. 2020*
- o **Research Intern at MILA** **Quebec, Canada**  
*Tang's Group, Mentor: Prof. Jian Tang*  
 Research Internship, working on Interpretable Natural Language Understanding. *Mar. 2020 - Jun. 2020*
- o **Research Intern at Microsoft Research Asia** **Beijing, China**  
*Natural Language Computing Group, Mentor: Dr. Tao Ge, Dr. Furu Wei*  
 Research Internship, working on natural language generation and efficient pre-trained language models. *Dec. 2018 - Present*

## Services

---

- Reviewer: NeurIPS 2020-2022, AAAI 2021-2023, ICML 2021-2022, ACL 2021, EMNLP 2021-2022, ICLR 2022-2023, ACL Rolling Review.
- Student Volunteer: ACL 2019, ICLR 2020, EMNLP 2020
- Review Assistant: EMNLP 2019, ACL 2020

## Honors and 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%)