Wang (Bill) Zhu

600 Central Ave Apt 337
Riverside, CA, US
92507

wangzhu@usc.edu
https://billzhu.me
+1 (360) 600 7466

EDUCATION

University of Southern California, Los Angeles, US   **Jan 2021 - present**
Doctor of Philosophy, Computer Science

Simon Fraser University, Vancouver, Canada   **Sep 2017 - Aug 2020**
Bachelor of Applied Science, Computer Science

Zhejiang University, Hangzhou, China   **Sep 2015 - Aug 2020**
Bachelor of Engineering, Computer Science

RESEARCH INTEREST

Natural Language Processing
Vision-Language Grounding
Machine Learning

SELECTED PUBLICATIONS

Wang Zhu, Ishika Singh, Robin Jia, and Jesse Thomason. **Language Models can Infer Action Semantics for Classical Planners from Environment Feedback.** *arXiv:2406.02791.*

Wang Zhu, Alekh Agarwal, Mandar Joshi, Robin Jia, Jesse Thomason, Kristina Toutanova. **Efficient End-to-End Visual Document Understanding with Rationale Distillation.** *North American Chapter of the Association for Computational Linguistics (NAACL), 2024*

Wang Zhu, Jesse Thomason, Robin Jia. **Chain-of-Questions Training with Latent Answers for Robust Multistep Question Answering?** *Empirical Methods in Natural Language Processing (EMNLP), 2023.*

Wang Zhu, Ishika Singh*, Yuan Huan*, Robin Jia, Jesse Thomason. **VLN Pretraining Still Works with Nonsensical or Irrelevant Instructions.** *O-DRUM Workshop, CVPR, 2023.*

Jacob Krantz*, Shurjo Banerjee*, **Wang Zhu**, Jason J Corso, Peter Anderson, Stefan Lee, Jesse Thomason. **Iterative Vision-and-Language Navigation.** *Conference on Computer Vision and Pattern Recognition (CVPR), 2023.*

Wang Zhu, Jesse Thomason, Robin Jia. **Generalization Differences between End-to-End and Neuro-Symbolic Vision-Language Reasoning Systems.** *Findings of EMNLP, 2022.*

Yejia Liu*, **Wang Zhu***, Shaolei Ren. **Navigating Memory Construction by Global Pseudo-Task Simulation for Continual Learning.** *Conference on Neural Information Processing Systems (NeurIPS), 2022.*

Wang Zhu, Peter Shaw, Tal Linzen, Fei Sha. **Learning to Generalize Compositionally by Transferring Across Semantic Parsing Tasks.** *arXiv preprint: 2111.05013, 2021.*

Wang Zhu*, Hexiang Hu*, Jiacheng Chen, Zhiwei Deng, Vihan Jain, Eugene Ie, Fei Sha. **BabyWalk: Going Farther in Vision-and-Language Navigation by Taking Baby Steps.** *Annual Conference of the Association for Computational Linguisticss (ACL), 2020*

Guiliang Liu, Oliver Schulte, **Wang Zhu**, Qingcan Li. **Toward Interpretable Deep Reinforcement Learning with Linear Model U-Trees.** *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD), 2018*

Guiliang Liu, **Wang Zhu**, Oliver Schulte. Interpreting Deep Sports Analytics: Valuing Actions and Players in the NHL. *Machine Learning and Data Mining for Sports Analytics Workshop at ECML-PKDD (MLSA, ECML-PKDD), 2018*

TECHNICAL TALKS

"Chain-of-Questions Training with Latent Answers for Robust Multistep Question Answering?"
- Google Seattle NLP Lunch, Kirkland, WA, USA   Jul 2023

"VLN Pretraining Still Works with Nonsensical or Irrelevant Instructions"
- University of Southern California NLP Lunch, Los Angeles, CA, USA   Apr 2023

"Multistep Reasoning Transferability across Machine Reading Comprehension Benchmarks"
- University of Southern California NLP Lunch, Los Angeles, CA, USA   Nov 2022

"Generalization Differences between End-to-End and Neuro-Symbolic Vision-Language Reasoning Systems"
- Mila - Quebec AI Institute, Montral, QC, Canada   Oct 2022

"Evaluating the Robustness of Multi-Image Vision-and-Language Reasoning Systems"
- University of Southern California ML & Friends, Los Angeles, CA, USA   Mar 2022

PROFESSIONAL SERVICES

Reviewer of EMNLP 2021, ICML 2022, EMNLP 2022, CVPR 2023, ACL 2023, ICCV 2023, EMNLP 2023, ARR 2023 Dec, CVPR 2024, ECCV 2024, NeurIPS 2024, ARR 2024 Feb, Apr, Jun.

RESEARCH EXPERIENCE

Researcher Intern, Meta Reality Labs   **May 2024 - Aug 2024**
AR-AI Team
Mentor: Luna Dong
- **Personalized Question Answering with Visual History**
  - ➜ Created a dataset for personalized question-answering with visual-only history.
  - ➜ Evaluated multiple VL models on multiple representations on the personalization task.

Student Researcher, Google Deepmind   **May 2023 - Nov 2023**
Wave Team
Mentor: Kristina Toutanova & Alekh Agarwal
- **Visual Document Understanding with Rationale Distillation**
  - ➜ Baked OCR and custom program generation capabilities in pixel-to-text models to improve their performance on visual document understanding.
  - ➜ Distilled rationales from large visual and language teacher models.

Research Assistant, University of Southern California   **Oct 2021 - May 2024**
Allegro Lab & GLAMOR Lab
Advisor: Prof. Robin Jia & Prof. Jesse Thomason
- **Robust Multistep Question Answering**
  - ➜ Proposed Chain-of-Questions, a framework that trains a model to robustly answer multistep questions by generating and answering sub-questions.

➔ Overcame the challenge of no annotated sub-answers by treating sub-answers as latent variables and training with a mixture of Hard-EM & RL.

➔ Outperformed GPT-3.5 and GPT-4 on several multistep QA benchmarks.

- **Evaluating Vision-and-Language Reasoning**
  ➔ Evaluated the robustness of widely-used multi-modal reasoning systems under both the multi-image and the single-image setups.
  ➔ Designed special robustness tests for multi-image QA.
  ➔ Re-designed a compositionally generalizable semantic parsing template for multi-image QA.

- **Vision-and-Language Navigation Pretraining**
  ➔ Analyzed VLN agents by adding noise to the pretraining data.
  ➔ Demonstrated word ordering and trajectory matching does not matter much in VLN pretraining.
  ➔ Demonstrated the data quantity is more important than data quality in VLN pretraining.
  ➔ Designed a nonsense-data augmentation method for effective VLN pretraining.

Research Assistant, University of Southern California   **Sep 2020 - Sep 2021**
Theoretical and Empirical Data Science Lab
Advisor: Prof. Fei Sha

- **Long Horizon Vision-and-Language Navigation**
  ➔ Observed the limitation in the transferability of the existing VLN model.
  ➔ Developed a baby-walk method to solve the problem.
  ➔ The first agent in the VLN community that generalizes towards long-horizon.

- **Compositional Generalization on NLP**
  ➔ Investigated learning representations that facilitate transfer learning from one compositional task to another.
  ➔ Designed an iterative dueling game for the learner to solve the problem.

Research Assistant, Simon Fraser University   **Oct 2017 - May 2019**
Vision and Media Lab & Computing Logic Lab
Advisor: Prof. Greg Mori & Prof. Oliver Schulte

- **Active Object Search**
  ➔ Developed a multi-object active object search method in a simulated environment.

- **Deep Reinforcement Learning on Sports**
  ➔ Adopted linear model tree for Ice-Hockey sports player ranking

TEACHING EXPERIENCE

- In the Fall of 2023, I was a TA for CS467, USC's undergraduate level machine learning course taught by Robin Jia, as one of the head TAs. I led three 50-minute discussion sections, assisted students, proofread all assignments and designed some questions for the midterm and final.
- In the Fall of 2022, I was a TA for CS567, USC's master level machine learning course, taught by Vatsal Sharan. I led three 50-minute discussion sections, assisted students, graded quizzes, and designed one assignment and some questions for the quizzes.

SKILLS

Operating Systems: Linux, macOS, Windows
Programming Languages: Python, LaTeX, Java, C/C++, Swift, Matlab
Tools/Framework: Pytorch, Tensorflow, Git
Languages: English (Proficient), Chinese (Native), Japanese (Beginner)

AWARDS &	USC Graduate School Fellowship	Jan 2021
SCHOLARSHIP	SFU Presidents Honour Roll	Oct 2018
	SFU Alumni Scholarship	Oct 2018
	SFU Open Scholarship	Oct 2018
	Undergraduate Student Research Awards	May 2018
	SFU Entrance Scholarship	Sep 2017
	The Mathematical Contest in Modeling, Meritorious Prize	Feb 2017