Wang Bill Zhu

600 Central Ave APT 337
Riverside, CA, USA
92507

wangzhu@usc.edu
https://billzhu.me
+1 (360)-660-7466

Education

University of Southern California, Los Angeles, US — Jan 2021 - May 2026
Doctor of Philosophy, Computer Science
Advisors: Jesse Thomason, Robin Jia

University of Southern California, Los Angeles, US — Jan 2021 - Aug 2024
Master of Science, 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

I work on natural language processing applications on **cross-disciplinary** domains including **multimodal, robotics, healthcare, social science**, with an emphasis on building **verifiable systems** to solve real-world problems in these domains.

Employment

Researcher Intern, Meta Reality Labs — May 2024 - Dec 2024
Redmond, WA, US
Hosts: Xin Luna Dong, Kai Sun

Student Researcher, Google Deepmind — May 2023 - Nov 2023
Kirkland, WA, US
Hosts: Kristina Toutanova, Alekh Agarwal

Publications

Wang Bill Zhu*, Miaosen Chai*, Shangshang Wang, Yejia Liu, Song Bian, Honghua Dong, Willie Neiswanger, Robin Jia. **Precise Debugging Benchmark: Is Your Model Debugging or Regenerating?** *Findings of ACL, 2026.*

Liancheng Gong, **Wang Bill Zhu**, Jesse Thomason, Li Zhang. **Zero-Shot Iterative Formalization and Planning in Partially Observable Environments.** *Findings of ACL, 2026.*

Wang Bill Zhu, Miaosen Chai, Ishika Singh, Robin Jia, Jesse Thomason. **PSALM-V: Automating Symbolic Planning in Interactive Visual Environments with Large Language Models.** *IEEE International Conference on Robotics and Automation (ICRA), 2026.*

Wang Bill Zhu, Tianqi Chen, Xinyan Velocity Yu, Ching Ying Lin, Jade Law, Mazen Jizzini, Jorge J. Nieva, Ruishan Liu, Robin Jia. **Cancer-Myth: Evaluating AI Chatbot on Patient Questions with False Presuppositions.** *International Conference on Learning Representations (ICLR), 2026.*

Ang Li, Charles Wang, Deqing Fu, Kaiyu Yue, Zikui Cai, **Wang Bill Zhu**, Ollie Liu, Peng Guo, Willie Neiswanger, Furong Huang, Tom Goldstein, Micah Goldblum.

Zebra-CoT: A Dataset for Interleaved Vision Language Reasoning. *International Conference on Learning Representations (ICLR), 2026.*

Wang Bill Zhu, Deqing Fu, Kai Sun, Yi Lu, Zhaojiang Lin, Seungwhan Moon, Kanika Narang, Mustafa Canim, Yue Liu, Anuj Kumar, Xin Luna Dong. **VisualLens: Personalization through Visual History.** *Conference on Neural Information Processing Systems (NeurIPS), 2025.*

Ming Li, Jike Zhong, Shitian Zhao, Yuxiang Lai, Haoquan Zhang, **Wang Bill Zhu**, Kaipeng Zhang. **Think or Not Think: A Study of Explicit Thinking in Rule-Based Visual Reinforcement Fine-Tuning.** *Conference on Neural Information Processing Systems (NeurIPS), 2025.*

Wang Bill Zhu, Ishika Singh, Robin Jia, and Jesse Thomason. **Language Models can Infer Action Semantics for Classical Planners from Environment Feedback.** *North American Chapter of the Association for Computational Linguistics (NAACL), 2025.*

Deqing Fu, Tong Xiao, Rui Wang, **Wang Zhu**, Pengchuan Zhang, Guan Pang, Robin Jia, Lawrence Chen. **TLDR: Token-Level Detective Reward Model for Large Vision Language Models.** *International Conference on Learning Representations (ICLR), 2025.*

Jiacheng Chen, Tianhao Liang, Sherman Siu, Zhengqing Wang, Kai Wang, Yubo Wang, Yuansheng Ni, **Wang Zhu**, Ziyan Jiang, Bohan Lyu, Dongfu Jiang, Xuan He, Yuan Liu, Hexiang Hu, Xiang Yue, Wenhu Chen. **Mega-Bench: Scaling Multimodal Evaluation to Over 500 Real-World Tasks.** *International Conference on Learning Representations (ICLR), 2025.*

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.** *Open-Domain Reasoning Under Multi-Modal Settings at CVPR (O-DRUM, 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*, 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.*

Preprints

Qiutong Tong Yi*, **Wang Bill Zhu***, Robin Jia, Jesse Thomason. **PDDL-Mind: Large Language Models are Capable on Belief Reasoning with Reliable State Tracking.** *arXiv: 2604.17819.*

Yuqing Yang, Tengxiao Liu, **Wang Bill Zhu**, Taiwei Shi, Linxin Song, Robin Jia. **Self-Evolving LLM Memory Extraction Across Heterogeneous Tasks.** *arXiv: 2604.11610.*

Amirmohammad Nazari, Sadra Sabouri, **Wang Bill Zhu**, Robin Jia, Souti Chattopadhyay, Mukund Raghothaman. Synthesizing Program Analyzers for Interactive Code Understanding.

Jacob Choi, Shuying Cao, Xingjian Dong, Amin Banayeeanzade, **Wang Bill Zhu**, Robin Jia, Sai Praneeth Karimireddy. **ContextLeak: Auditing Leakage in Private In-Context Learning Methods.** *arXiv:2507.16746.*

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

Awards & Grants

NSF ACCESS Computing Grant (CIS251157)	Oct 2025
USC Graduate School Fellowship	Jan 2021
SFU Presidents Honour Roll	Oct 2018
Undergraduate Student Research Awards	May 2018
The Mathematical Contest in Modeling, Meritorious Prize	Feb 2017

Mentorship

Yunkai Zhan (USC UG)	Apr 2026 - Present
Eric Huang (USC UG)	Aug 2025 - Present
Qiutong Tony Yi (USC UG)	Aug 2024 - Present
Miaosen Chai (USC UG → UChicago Ph.D.)	Aug 2024 - Jan 2026
Jacob Choi (USC MS)	Sep 2025 - Dec 2025
Shuying Cao (USC MS)	Sep 2025 - Dec 2025
Xingjian Dong (USC UG)	Sep 2025 - Dec 2025
Ang Li (Columbia University MS → NYU Ph.D.)	May 2025 - Sep 2025
Tianqi Chen (USC MS → Cognex)	Aug 2024 - Aug 2025
Jike Zhong (USC Ph.D.)	Jan 2025 - Jun 2025
Yuan Huang (USC MS → Corteva)	Jun 2022 - Oct 2023

Teaching

Invited Lecturer, CSCI 444 (Natural Language Processing) Fall 2024
Invited lecturer for a 50-student class, delivered a lecture on vision-language reasoning.

Head Teaching Assistant, CSCI 467 (Introduction to Machine Learning) Fall 2023
Head TA of 50-student class, managed a team of 6 TA and CPs, led three discussion sections, proofread all assignments and designed questions for the midterm and final.

Teaching Assistant, CSCI 567 (Machine Learning) Fall 2022
TA of 260-student class, led three discussion sections, graded quizzes, and designed one assignment and several questions for the quizzes.

Professional Services

Natural Language Processing:
Reviewer for EMNLP (2021, 2022, 2023), ACL (2023), ACL Rolling Review (2023, 2024, 2025, 2026), COLM (2026).

Machine Learning:
Reviewer for ICML (2022, 2026), NeurIPS (2024, 2025, 2026), ICLR (2026).

Computer Vision:
Area Chair for ECCV (2026); Reviewer for CVPR (2023, 2024, 2025), ICCV (2023, 2025), ECCV (2024).

Robotics:
Reviewer for CoRL (2025), ICRA (2026), HRI (2026).

Invited Talks

"Aligning LLMs with Human Well-Being: Medical, Social, and Cognitive Science Perspectives"
• UCB, Berkeley, CA, USA Apr 2026

"Domain-Adaptive Programming: Expanding the Boundaries of What LLMs Can Solve"
• EleutherAI Planning Reading Group, Los Angeles, CA, USA Mar 2026

"Domain-Adaptive Programming: Expanding the Boundaries of What LLMs Can Solve"
• USC ISI NLP Seminar, Los Angeles, CA, USA Jan 2026

"LLM Reasoning: A Mini Overview on Post-Training and Test-Time Compute"
• USC ML/NLP Reading Group, Los Angeles, CA, USA Mar 2025

"LLM Reasoning: A Mini Overview on Post-Training and Test-Time Compute"
• USC ML/NLP Reading Group, Los Angeles, CA, USA Mar 2025

"VisualLens: Personalization through Visual History."
• Meta AR-AI Reading Group, Redmond, WA, USA Aug 2024

"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