

Zhengxuan (Zen) Wu
Phone: 216-551-7046
Email: wuzhengx@stanford.edu
Website: <https://zen-wu.social/>

EDUCATION

- Stanford University** *2020/09 - 2022/05*
M.S. in Symbolic Systems Program *GPA: 4.00/4.00*
Focus: *Cognitive Science and AI*
Advisors: *Christopher Potts; Desmond C. Ong*
- University of Pennsylvania** *2015/09 - 2017/05*
M.S. in Computer Science
- Case Western Research University** *2012/09 - 2015/05*
B.S. in Aerospace Engineering

RESEARCH EXPERIENCE

- Stanford AI Lab (SAIL) - Graduate Researcher** *2020/08 - present*
· Working on inducing symbolic causal structures on neural networks through interchange interventions.
· Developing program synthesis pipeline to solve ARC task using Knowledge Graph and RL.
· Built the next-gen sentiment analysis benchmark DynaSent ([NAACL21](#), [ACL21](#)).
· Investigating fine-tuning under know label distribution shifts, and compositional generalization with symbolic-neural systems ([NeurIPS21](#)).

- Stanford Social Neurscience Lab - Graduate Researcher** *2018/03 - present*
· Advancing feature importance attribution methods in BERT-like models for better interpretations.
· Led the development of attention and relevance tracing for the Transformer model ([BlackboxNLP@EMNLP20](#)).
· Led the development of context-guided BERT by proposing novel quasi-attention mechanism ([AAAI21](#)).
· Built deep learning models for sentiment analysis tasks, including LSTM, VRNN and BERT ([IEEE ACII19](#)).
· Jointly led the collection of a large story-telling sentiment analysis dataset, SEND ([IEEE TAC19](#)).

- Stanford HCI Lab - Graduate Researcher** *2018/03 - 2020/08*
· Enhanced HabitLab, a personalized productivity intervention system on Chrome browser.
· Contributed to study user behavioral changes through online intervention systems ([CSCW18](#), [CHI19](#), [21](#)).

PROFESSIONAL EXPERIENCE

- VMware, Inc.** - Software Engineer III *2017/07 - 2020/09*
· Developed scalable data-center management platform.
- Swift Capital (Paypal, Inc.)** - Machine Learning Intern *2016/05 - 2016/09*
· Developed machine learning systems to predict the credit scores of loan applicants.

MANUSCRIPTS AND PUBLICATIONS¹

- [Zhengxuan Wu*](#), Atticus Geiger*, Hanson Lu, Josh Rozner, Elisa Kreiss, Thomas Icard, Noah Goodman, Christopher Potts, “Causal Distillation for Language Models”, M.s., Stanford University, https://zen-wu.social/papers/ACL22_CausalDistill.pdf.
- Atticus Geiger*, [Zhengxuan Wu*](#), Hanson Lu, Josh Rozner, Elisa Kreiss, Thomas Icard, Noah Goodman, Christopher Potts, “Inducing Symbolic Causal Structures to Produce Systematic and Interpretable Neural Networks”, M.s., Stanford University, https://zen-wu.social/papers/ICML22_IIT.pdf.
- [Zhengxuan Wu](#), Nelson F. Liu, Christopher Potts, “Identifying the Limits of Cross-Domain Knowledge Transfer for Pretrained Models”, M.s., Stanford University, <https://arxiv.org/abs/2104.08410>.
- [Zhengxuan Wu](#), Desmond C. Ong, “On Explaining Your Explanations of BERT: An Empirical Study with Sequence Classification”, M.s., Stanford University and National University of Singapore, <https://arxiv.org/abs/2101.00196>.
- [Zhengxuan Wu*](#), Elisa Kreiss*, Desmond C. Ong, Christopher Potts, “ReaSCAN: Compositional Reasoning in Language Grounding”, ([NeurIPS21](#)), <https://arxiv.org/abs/2109.08994>.

¹*equal contribution

6. Christopher Potts*, Zhengxuan Wu*, Atticus Geiger, Douwe Kiela, “DynaSent: A Dynamic Benchmark for Sentiment Analysis”, (ACL21), <https://arxiv.org/abs/2012.15349>.
7. Zhengxuan Wu, Desmond C. Ong, “Context-Guided BERT for Targeted Aspect-Based Sentiment Analysis”, (AAAI21), <https://arxiv.org/abs/2010.07523>.
8. Zhengxuan Wu, Desmond C. Ong, “Pragmatically Informative Color Generation by Grounding Contextual Modifiers”, (SCiL21), <https://arxiv.org/abs/2010.04372>.
9. Zhengxuan Wu, Thanh-Son Nguyen and Desmond C. Ong, “Structured Self-Attention Weights Encode Semantics in Sentiment Analysis”, (BlackboxNLP@EMNLP20), <https://arxiv.org/abs/2010.04922>.
10. Zhengxuan Wu, Xiyu Zhang, Zhi-Xuan Tan, Jamil Zaki, Desmond C. Ong, “Attending to Emotional Narratives”, (IEEE ACII19), <https://arxiv.org/abs/1907.04197>.
11. Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts and Adina Williams, “Dynabench: Rethinking Benchmarking in NLP”, (NAACL21), <https://arxiv.org/abs/2104.14337>.
12. Geza Kovacs, Zhengxuan Wu and Michael S. Bernstein, “Not Now, Ask Later: Users Weaken Their Behavior Change Regimen Over Time, But Believe They Will Imminently Re-Strengthen It”, (CHI21), <https://arxiv.org/abs/2101.11743>.
13. Desmond C. Ong, Zhengxuan Wu, Zhi-Xuan Tan, Marianne Reddan, Isabella Kahhale, Alison Mattek and Jamil Zaki, “Modeling emotion in complex stories: the Stanford Emotional Narratives Dataset”, (IEEE TAC19), <https://arxiv.org/abs/1912.05008>.
14. Geza Kovacs, Drew Mylander Gregory, Zilin Ma, Zhengxuan Wu, Golrokh Emami, Jacob Ray and Michael S. Bernstein, “Conservation of Procrastination: Do Productivity Interventions Save Time or Just Redistribute It?”, (CHI19), <https://dl.acm.org/doi/10.1145/3290605.3300560>.
15. Geza Kovacs, Zhengxuan Wu and Michael S. Bernstein, “Rotating Online Behavior Change Interventions Increases Effectiveness But Also Increases Attrition”, (CSCW18), <https://dl.acm.org/doi/10.1145/3274364>.
16. Erik J. Stalcup, James S. T’ien, Jonathan Jordan, Zhengxuan Wu, Gabriel Nastac and Chengyao Li, “Upward Flame Spread and Extinction over Wavy Solids”, (CST20), <https://www.tandfonline.com/doi/abs/10.1080/00102202.2020.1738411>.

ACADEMIC EXPERIENCE

- Reviewer for CHI 2019
- Invited Abstract Presentation in IC2S2 2019, University of Amsterdam, Netherlands

OPEN SOURCE PROJECTS

Dynabench @ Facebook, Inc. - Contributor

Deep Learning ◊ PyTorch ◊ React ◊ Python

Developing as an individual contributor to the Dynamic Adversarial Benchmarking (Dynabench) Platform launched by Facebook, Inc.

Kaggle - ARC Challenge

Deep Learning ◊ PyTorch ◊ GNN ◊ RL ◊ Program Synthesis

Building artificial general intelligent agents to solve reasoning tasks.

CSI @ Kubernetes - Contributor

Go ◊ C++ ◊ VMware ◊ Kubernetes

Developing large-scale open-source data-center management platform on VMware cloud.

HabitLab @ Stanford HCI - Contributor

HCI ◊ Intervention ◊ Chrome App ◊ RL

Contributed more than 10k+ lines of code to the HabitLab, a app in Chrome for better work efficiency.

TECHNICAL STRENGTHS

- **Program Languages:** Python, C++/C, C#, Java, R, Matlab, Haskell, Bash.

- **Machine Learning:** Discriminative and Generative Models (CNN/RNN/LSTM/VAE/GAN/HMM on CUDA); Reinforcement Learning; Multi-task Learning; Graph Neural Networks.
- **AI + Big Data:** PyTorch, scikit-learn, Keras, TensorFlow, NumPy, Pandas, H2O, MapReduce (Hadoop).
- **Data Mining:** PyData, SciPy, SNAP, SQL, NoSQL (Mongo), NetworkX, Jupyter.
- **Data Science:** Mixed Linear Model, Hierarchical Logistic Regression, A/B Testings, Crowdsourcing (MTurk).
- **Server + Database:** Node.js, Flask, MongoDB, PostgreSQL, Kubernetes, Docker, Google Cloud, AWS EC2, Heroku, Azure, Jenkins CICD.
- **Web + Mobile:** HTML/CSS/JS, Polymer, React, Webpack, Apache, Android (Java), Xcode.