

Victor Zhong

Seattle WA, United States

✉ victor@victorzhong.com | 🏠 www.victorzhong.com | 📷 vzhong | 🌐 victorzhong

Research Interests

My interests are in natural language processing and machine learning. I am particularly interested in how to use language understanding to facilitate learning. My recent work has focused on using machine reading to produce zero-shot policies for new environment dynamics, using conditional language models to adapt existing semantic parsers to new databases, and using extract-and-edit techniques to build dialogue agents for new domains. Previously, I worked on dialogue, question answering, semantic parsing, and knowledge base population.

Education

University of Washington

PhD in Computer Science

Advisor: Prof. Luke Zettlemoyer

Apple Scholar in ML/AI Fellowship

2018 - Present

Seattle, WA

Stanford University

Master of Science in Computer Science, CGPA 3.96

Advisor: Prof. Christopher D. Manning

2014 - 2016

Stanford, CA

University of Toronto

Bachelor of Applied Science in Computer Engineering with Honours, CGPA 3.82

Advisors: Prof. Joseph Euzebe (Zeb) Tate and Prof. Jonathan Rose

2009 - 2014

Toronto, Canada

Dean's List, Natural Sciences and Engineering Research Council of Canada (NSERC) Undergraduate Research Award

Publications

Grounding Language to Entities and Dynamics for Generalization in Reinforcement Learning. H. J. Austin Wang, **Victor Zhong**, Karthik Narasimhan. *International Conference on Machine Learning (ICML)*, 2021

LEWIS: Levenshtein Editing for Unsupervised Text Style Transfer. Machel Reid, **Victor Zhong**. *Findings of the Association for Computational Linguistics (ACL Findings)*, 2021.

Grounded Adaptation for Zero-shot Executable Semantic Parsing. **Victor Zhong**, Mike Lewis, Sida I. Wang, Luke Zettlemoyer. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.

RTFM: Generalising to Novel Environment Dynamics via Reading. **Victor Zhong**, Tim Rocktäschel, Edward Grefenstette. *International Conference on Learning Representations (ICLR)*, 2020.

E3: Entailment-driven Extracting and Editing for Conversational Machine Reading. **Victor Zhong**, Luke Zettlemoyer. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.

Multi-hop Reading Comprehension through Question Decomposition and Rescoring. Sewon Min, **Victor Zhong**, Luke Zettlemoyer, Hananeh Hajishirzi. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.

Coarse-grain Fine-grain Coattention Network for Multi-evidence Question Answering. **Victor Zhong**, Caiming Xiong, Nitish Shirish Keskar, Richard Socher. *International Conference on Learning Representations (ICLR)*, 2019.

Global-Locally Self-Attentive Encoder for Dialogue State Tracking. **Victor Zhong**, Caiming Xiong, Richard Socher. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.

Efficient and Robust Question Answering from Minimal Context over Documents. Sewon Min, **Victor Zhong**, Richard Socher, Caiming Xiong. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.

DCN+: Mixed Objective and Deep Residual Coattention for Question Answering. **Victor Zhong***, Caiming Xiong*, Richard Socher (*equal contribution). *International Conference on Learning Representations (ICLR)*, 2018.

Seq2SQL: Generating Structured Queries from Natural Language using Reinforcement Learning. **Victor Zhong**, Caiming Xiong, Richard Socher. *Preprint CoRR abs/1709.00103*, 2017.

Position-aware Attention and Supervised Data Improve Slot Filling (Oral, Best Resource Paper runner-up). Yuhao Zhang, **Victor Zhong**,

Danqi Chen, Gabor Angeli, Christopher D. Manning. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.

Dynamic Coattention Networks for Question Answering. Victor Zhong* and Caiming Xiong*, Richard Socher (*equal contribution). *International Conference on Learning Representations (ICLR)*, 2017.

Ask Me Anything: Dynamic Memory Networks for Natural Language Processing. Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, Richard Socher. *International Conference on Machine Learning (ICML)*, 2016.

Bootstrapped Self Training for Knowledge Base Population. Gabor Angeli, Victor Zhong, Danqi Chen, Arun Chaganty, Jason Bolton, Melvin J. Premkumar, Panupong Pasupat, Sonal Gupta, Christopher D. Manning. *Text Analysis Conference (TAC)*, 2015.

Professional Services

Workshop organizing committee: 1st Workshop on Interactive Executable Semantic Parsing (EMNLP 2020), 1st Workshop on Language and Reinforcement Learning (ICML 2020).

Reviewer: NeurIPS (2018-present), ACL (2018-present), ICML (2019-present), ICLR (2018-present), EMNLP (2018-present), NAACL (2018-present), CoRL (2019-present), CoNLL (2018-present)

Invited talks: Microsoft Research (September 2020), Microsoft Dialogue Research (April 2019), Target Data Science (May 2018), Salesforce IQ (April 2017), Stanford NLP (March 2017)

Research and Work Experience

University of Washington

RESEARCH ASSISTANT

Seattle, WA

2018-present

Facebook AI Research

VISITING RESEARCHER

Seattle, WA

2021-present

Google Brain Robotics

RESEARCH INTERN

Mountainview, CA

2020 Summer

Facebook AI Research

RESEARCH INTERN

London, UK

2019 Summer

Salesforce Research

RESEARCH SCIENTIST

Palo Alto, CA

2016-2018

Stanford Natural Language Processing Group

RESEARCH ASSISTANT

Stanford, CA

2015-2016

Metamind Inc.

RESEARCH INTERN

Palo Alto, CA

2016 Winter

Bloomberg Labs

SOFTWARE ENGINEERING INTERN

New York City, NY

2014 Summer

Stanford AI Lab

RESEARCH ASSISTANT

Stanford, CA

2013 Summer

Altera Corporation

SOFTWARE ENGINEERING INTERN

San Jose, CA

2012-2013

University of Toronto

RESEARCH ASSISTANT

Toronto, ON

2011 Summer

National University of Singapore

RESEARCH ASSISTANT

Singapore

2010 Summer

Open Source Projects

RTFM, Code for our ICLR paper *RTFM: Generalising to Novel Environment Dynamics via Reading*.

E3, Code for our ACL paper *E3: Entailment-driven Extracting and Editing for Conversational Machine Reading*.

WikiSQL, Large scale, crowd-sourced dataset for developing natural language interface for databases.

Dashboard, Experiment management and visualization package with web app.

Embeddings, Library that automatically downloads pretrained word embeddings and populates a local SQL database for fast access.

Stanza, Stanford NLP deep learning infrastructure repository for Theano and Tensorflow.

Torchlib, Torch libraries for machine learning and NLP.

CoreNLP.lua, A Lua client for Stanford CoreNLP.

[facebookresearch/rtfm](#) 

[vzhong/e3](#) 

[salesforce/wikisql](#) 

[vzhong/dashboard](#) 

[vzhong/embeddings](#) 

[stanfordnlp/stanza](#) 

[vzhong/torchlib](#) 

[vzhong/corenlp.lua](#) 