

The University of Chicago Computer Science Department
PRESENTS:

“Self-Supervised Natural Language Processing”



William Wang

University of California, Santa Barbara

Abstract:

Learning to reason and understand the world’s knowledge is a fundamental problem in Artificial Intelligence (AI). While it is always hypothesized that the learning models should be generalizable and flexible, in practice, most of the progress is still in classic supervised learning settings that require a large amount of annotated training data and heuristic objectives.

With the vast amount of language data available in digital form, now is a good opportunity to move beyond traditional supervised learning methods. The core research question that I will address in this talk is the following: how can we design self-supervised deep learning methods to operate over rich language and knowledge representations? In this talk, I will describe some examples of my work in advancing the state-of-the-arts in methods of deep reinforcement learning for NLP, including: 1) Reinforced Co-Training, a new semi-supervised learning framework that is driven by a reinforced performance-driven data selection policy agent ; 2) AREL, a self-adaptive inverse reinforcement learning agent for visual storytelling; and 3) DeepPath, an explainable path-based reasoning agent for inferring unknown facts. I will conclude this talk by describing my other research interests and my future research plans in the interdisciplinary field of AI and data science.

Bio:

William Wang is an Assistant Professor in the Department of Computer Science at the University of California, Santa Barbara. He received his PhD from School of Computer Science, Carnegie Mellon University in 2016. He has broad interests in machine learning approaches to data science, including natural language processing, statistical relational learning, information extraction, computational social science, dialogue, and vision. He directs UCSB's NLP Group (nlp.cs.ucsb.edu): in two years, UCSB advanced in the NLP area from an undefined ranking position to top 3 in 2018 according to CSRankings.org. He has published more than 60 papers at leading NLP/AI/ML conferences and journals, and received best paper awards (or nominations) at ASRU 2013, CIKM 2013, and EMNLP 2015, a DARPA Young Faculty Award (Class of 2018), two IBM Faculty Awards in 2017 and 2018, a Facebook Research Award in 2018, an Adobe Research Award in 2018, and the Richard King Mellon Presidential Fellowship in 2011. He served as an Area Chair for NAACL, ACL, EMNLP, and AAAI. He is an alumnus of Columbia University, Yahoo! Labs, Microsoft Research Redmond, and University of Southern California. In addition to research, William enjoys writing scientific articles that impact the broader online community: his microblog Weibo has 111,000+ followers and more than 2,000,000 views each month. His work and opinions frequently appear at major international media outlets such as Wired, VICE, Fast Company, NASDAQ, Scientific American, The Next Web, The Brookings Institution, Law.com, and Mental Floss.

Wednesday, March 6, 2019

3:30 pm

JCL 390

Host: Ben Zhao