



Luke Zettlemoyer

(University of Washington)

Friday, 1/25/19

2:30 pm Talk, DH 2315
4:00 pm Snacks, LTI 5th Floor

LABELING AND PREDICTING SEMANTIC ROLES DIRECTLY FROM TEXT.

I am an Associate Professor in the [Allen School of Computer Science & Engineering](#) at the [University of Washington](#).

I am also a [PECASE Awardee](#) and an [Allen Distinguished Investigator](#). Previously, I did postdoctoral research at the University of Edinburgh and was a Ph.D. student at MIT.

Semantic roles represent central aspects of the meaning of text, including roughly "who" did "what" to "whom," etc. In this talk, I will cover our recent efforts to building high quality semantic role labeling (SRL) systems, including advances in modeling and data annotation. The SRL models use relatively simple deep architectures that are trained end-to-end to jointly predict predicates and arguments, and can be run with no preprocessing (e.g. no POS tagger or syntactic parser). They also work extremely well, achieved nearly 40% relative error reductions over pre-neural methods on the PropBank benchmark. The data annotation is enabled by a new question-answer driven semantic role labeling (QA-SRL) formulation, which we show can represent most of the content provided by more traditional formulations while also enabling large scale crowdsourcing. Using this scheme, we were able to label over 60,000 sentences in a little over a week, and train high quality SRL models on this new data. The data and models are freely available online at [gasrl.org](#). Together, these advances make it possible for the first time to train highly accurate SRL models for any new domain at relatively modest cost.

This joint work was primarily led by Luheng He, Nicholas FitzGerald, and Julian Michael. Two of the projects received best paper honorable mentions at ACL 2018.