NAACL HLT 2019

Structured Prediction for NLP

Proceedings of the Third Workshop

June 7, 2019
Minneapolis, Minnesota, USA

Introduction

Welcome to the Third Workshop on Structured Prediction for NLP!

Structured prediction has a strong tradition within the natural language processing (NLP) community, owing to the discrete, compositional nature of words and sentences, which leads to natural combinatorial representations such as trees, sequences, segments, or alignments, among others. It is no surprise that structured output models have been successful and popular in NLP applications since their inception. Many other NLP tasks, including, but not limited to: semantic parsing, slot filling, machine translation, or information extraction, are commonly modeled as structured problems, and accounting for said structure has often lead to performance gain.

Of late, continuous representation learning via neural networks has been a significant complementary direction, leading to improvements in unsupervised and semi-supervised pre-training, transfer learning, domain adaptation, *etc.* Using word embeddings as features for structured models such as part-of-speech taggers count among the very first uses of continuous embeddings in NLP, and the symbiosis between the two approaches is an exciting research direction today.

The five papers (as well as three additional non-archival papers) accepted for presentation in this edition of the workshop, after double-blind peer review, all explore this interplay between structure and neural data representations, from different, important points of view. The program includes work on structure-informed representation learning, transfer learning, partial supervision, and parallelization of computation in structured computation graphs. Our program also includes six invited presentations from influential researchers.

Our warmest thanks go to the program committee – for their time and effort providing valuable feedback, to all submitting authors – for their thought-provoking work, and to the invited speakers – for doing us the honor of joining our program. We are looking forward to seeing you in Minneapolis!

Zornitsa Kozareva
Julia Kreutzer
Gerasimos Lampouras
André Martins
Vlad Niculae
Sujith Ravi
Andreas Vlachos

Organizers:

    Zornitsa Kozareva, Google, USA
    Julia Kreutzer, Heidelberg University, Germany
    Gerasimos Lampouras, University of Cambridge, UK
    André F.T. Martins, Unbabel & University of Lisbon, Portugal
    Vlad Niculae, University of Lisbon, Portugal
    Sujith Ravi, Google, USA
    Andreas Vlachos, University of Cambridge, UK

Program Committee:

    Yoav Artzi, Cornell University, USA
    Wilker Aziz, University of Amsterdam, Netherlands
    Joost Bastings, University of Amsterdam, Netherlands
    Xilun Chen, Cornell University, USA
    Shay Cohen, University of Edinburgh, UK
    Hal Daumé, Microsoft & University of Maryland, USA
    Kevin Gimpel, TTI Chicago, USA
    Matt Gormley, CMU, USA
    Arzoo Katiyar, Cornell University, USA
    Yoon Kim, Harvard University, USA
    Parisa Kordjamshidi, Tulane University, USA
    Amandla Mabona, University of Cambridge, UK
    Pranava Madhyastha, Imperial College London, UK
    Sebastian Mielke, Johns Hopkins University, USA
    Roi Reichart, Technion - Israel Institute of Technology, Israel
    Marek Rei, University of Cambridge, UK
    Stefan Riezler, Heidelberg University, Germany
    Hiko Schamoni, Heidelberg University, Germany
    Tianze Shi, Cornell University, USA
    Artem Sokolov, Amazon, Germany
    Vivek Srikumar, University of Utah, USA
    Ivan Titov, University of Edinburgh, Scotland
    Luke Zettlemoyer, University of Washington, USA

Invited Speakers:

    Claire Cardie, Cornell University, USA
    Chris Dyer, DeepMind, UK
    Jason Eisner, Johns Hopkins University, USA
    Hannaneh Hajishirzi, University of Washington, USA
    He He, Stanford University, USA
    Andrew McCallum, University of Massachusetts Amherst, USA

Table of Contents

Workshop Program

Friday, June 7, 2019

9:00–9:10   *Opening Remarks*

9:10–9:50   *Invited Talk: Andrew McCallum*

9:50–10:30   *Invited Talk: Hannaneh Hajishirzi*

10:30–11:00   *Coffee Break*

11:00–11:40   *Invited Talk by He He*

11:40–12:10   *Session 1: Contributed Talks*

11:40–11:55   *Parallelizable Stack Long Short-Term Memory*
Shuoyang Ding and Philipp Koehn

11:55–12:10   *Tracking Discrete and Continuous Entity State for Process Understanding*
Aditya Gupta and Greg Durrett

12:10–14:00   *Lunch*

14:00–14:40   *Invited Talk: Chris Dyer*

14:40–15:00   *Poster Spotlight*

15:00–16:00   *Poster Session (and non-archival papers)*

SPARSE: Structured Prediction using Argument-Relative Structured Encoding
Rishi Bommasani, Arzoo Katiyar and Claire Cardie

Friday, June 7, 2019 (continued)

Lightly-supervised Representation Learning with Global Interpretability
Andrew Zupon, Maria Alexeeva, Marco Valenzuela-Escárcega, Ajay Nagesh and
Mihai Surdeanu

Semi-Supervised Teacher-Student Architecture for Relation Extraction
Fan Luo, Ajay Nagesh, Rebecca Sharp and Mihai Surdeanu

15:30–16:00   *Coffee Break (With Posters)*

16:00–16:40   *Invited Talk: Claire Cardie*

16:40–17:20   *Invited Talk: Jason Eisner*