EMNLP 2016

Conference on Empirical Methods in Natural Language Processing

Proceedings of the Workshop on Structured Prediction for Natural Language Processing

November 5, 2016
Austin, Texas, USA

Order copies of this and other ACL proceedings from:

Curran Associates
57 Morehouse Lane
Red Hook, New York 12571
USA
Tel: +1-845-758-0400
Fax: +1-845-758-2633
curran@proceedings.com

ISBN 978-1-945626-29-6

Preface

Welcome to the first workshop on structured prediction for NLP! Many prediction tasks in NLP involve assigning values to mutually dependent variables. For example, when designing a model to automatically perform linguistic analysis of a sentence or a document (e.g., parsing, semantic role labeling, or discourse analysis), it is crucial to model the correlations between labels. Many other NLP tasks, such as machine translation, textual entailment, and information extraction, can be also modeled as structured prediction problems.

In order to tackle such problems, various structured prediction approaches have been proposed, and their effectiveness has been demonstrated. Studying structured prediction is interesting from both NLP and machine learning (ML) perspectives. From the NLP perspective, syntax and semantics of natural language are clearly structured and advances in this area will enable researchers to understand the linguistic structure of data. From the ML perspective, the large amount of available text data and complex linguistic structures bring challenges to the learning community. Designing expressive yet tractable models and studying efficient learning and inference algorithms become important issues.

Recently, there has been significant interest in non-standard structured prediction approaches that take advantage of non-linearity, latent components, and/or approximate inference in both the NLP and ML communities. Researchers have also been discussing the intersection between deep learning and structured prediction through the DeepStructure reading group. This workshop intends to bring together NLP and ML researchers working on diverse aspects of structured prediction and expose the participants to recent progress in this area.

This year we have seven papers (six regular papers and one tutorial paper) covering various aspects of structured prediction, including neural networks, deep structured prediction software library, classical inside-outside algorithm, and imitation learning. We also invited four fantastic speakers and a great discussion panel. We hope you all enjoy the program!

Finally, we would like to thank all programming committee members, speakers, panelists, and authors. We are looking forward to seeing you in Austin.

Organizers:

Kai-Wei Chang, University of Virginia
Ming-Wei Chang, Microsoft Research
Alexander Rush, Harvard University
Vivek Srikumar, University of Utah

Program Committee:

Amir Globerson, Tel Aviv University (Israel)
Andre Martins, Unbabel (Portugal)
Chris Dyer, Carnegie Mellon University (USA)
Dan Roth, University of Illinois, Urbana-Champaign (USA)
David Sontag, New York University (USA)
Hal Daumé III, University of Maryland (USA)
Ivan Titov, University of Amsterdam (Netherlands)
Janardhan Rao Doppa, Washington State University (USA)
Jason Eisner, Johns Hopkins University (USA)
Kevin Gimpel, Toyota Technological Institute at Chicago (USA)
Luke Zettlemoyer, University of Washington (USA)
Matt Gormley, Carnegie Mellon University (USA)
Michael Collins, Columbia University (USA)
Mohit Bansal, UNC Chapel Hill (USA)
Ofer Meshi, Toyota Technological Institute at Chicago (USA)
Ryan McDonald, Google (USA)
Scott Yih, Microsoft Research (USA)
Sebastian Riedel, University College London (UK)
Shay Cohen, University of Edinburgh (UK)
Yoav Artzi, Cornell University (USA)
Yuan Zhang, University of Texas at Dallas (USA)
Tao Lei, Massachusetts Institute of Technology (USA)

Invited Speaker:

Kristina Toutanova, Microsoft Research
Andrew McCallum, University of Massachusetts, Amherst
Raquel Urtasun, University of Toronto
Dzmitry Bahdanau, University of Montreal

Panelists:

Hal Daumé III, University of Maryland (USA) (moderator)
Noah Smith, University of Washington (USA)
Sebastian Riedel, University College London (UK)
Zornitsa Kozareva, Amazon (USA)

Table of Contents

Workshop Program

Saturday, November 5, 2016

09:00–10:30 Section 1

9:00–9:15 *Welcome*
Organizers

9:15–10:00 *Invited Talk*
Kristina Toutanova

10:00–10:30 *Inside-Outside and Forward-Backward Algorithms Are Just Backprop (tutorial paper)*
Jason Eisner

10:30–11:00 *Coffee break*

11:00–12:30 Section 2

11:00–11:45 *Invited Talk*
Andrew McCallum

11:45–12:30 *Panel*
Panel: Moderator - Hal Daume

12:30–2:00 *Lunch*

Saturday, November 5, 2016 (continued)

14:00–15:30 Section 3

14:00–14:45 ***Poster***

Research on attention memory networks as a model for learning natural language inference
zhuang liu, Degen Huang, jing zhang and kaiyu huang

A Joint Model of Rhetorical Discourse Structure and Summarization
Naman Goyal and Jacob Eisenstein

Posterior regularization for Joint Modeling of Multiple Structured Prediction Tasks with Soft Constraints
Kartik Goyal and Chris Dyer

A Study of Imitation Learning Methods for Semantic Role Labeling
Travis Wolfe, Mark Dredze and Benjamin Van Durme

Introducing DRAIL – a Step Towards Declarative Deep Relational Learning
Xiao Zhang, María Leonor Pacheco, Chang Li and Dan Goldwasser

14:45–15:30 *Invited Talk*
Raquel Urtasun

15:30–16:00 ***Coffee Break***

Saturday, November 5, 2016 (continued)

16:00–17:50 Section 4

16:00–16:45 *Invited Talk*
Dzmitry Bahdanau

16:45–17:15 *Unsupervised Neural Hidden Markov Models*
Ke M. Tran, Yonatan Bisk, Ashish Vaswani, Daniel Marcu and Kevin Knight

17:15–17:30 ***Closing***