CoNLL 2017

The SIGNLL Conference on Computational Natural Language Learning

Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies

August 3-4, 2017
Vancouver, Canada

Sponsors:

Google, Inc.

DFKI Berlin

CRACKER project

text & form

ÚFAL

Charles University

LINDAT/CLARIN

ISBN 978-1-945626-70-8

Introduction

This volume contains papers describing systems submitted to the *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies* and an overview paper summarizing the task, its features, evaluation methodology for the main and additional metrics, and some interesting observations about the submitted systems and the task as a whole.

This Shared Task (`http://universaldependencies.org/conll17/`) can be seen as an extension of the CoNLL 2007 Shared Task on parsing, but there are many important differences that make this year's task unique with several "firsts". Most importantly, the data for this task come from the Universal Dependencies project (`http://universaldependencies.org`), which provides annotated treebanks for a large number of languages using the same annotation scheme for all of them. In the shared task setting, this allows for more meaningful comparison between systems as well as languages, since differences are much more likely due to true parser differences rather than differences caused by annotation schemes. In addition, the number of languages for which training data were available is unprecedented for a single shared task: a total of 64 treebanks in 45 languages have been provided for training the systems. Additional data have been provided too, as were some baseline systems for those who wanted to try only some particular aspect of parsing. Overall, the task can be described as "closed", since only pre-approved data could be used.

For evaluation, there were 81 datasets (standard datasets for the treebank languages provided for training, plus more test sets in known languages, but based on a specially created and annotated parallel corpus, and four surprise language test sets). Participants had to process all the test sets. The TIRA platform has been used for evaluation, as was the case already for the CoNLL 2015 and 2016 Shared Tasks, meaning that participants had to provide their code on a designated virtual machine to be run by the organizers to produce official results. However, test data have been published after the official evaluation period, and participants could run their systems at home to produce additional results they were allowed to include in the system description papers. There was one main evaluation metric – Labeled Attachment Score – for the main ranking table evaluating dependency parsing performance, plus additional metrics for tokenization, word and sentence segmentation, POS tagging, lemmatization and disambiguation of morphological features, and separate metrics computed for interesting subsets of the evaluation data.

A total of 32 systems ran successfully and have been ranked (`http://universaldependencies.org/conll17/results.html`). While there are clear overall winners, we would like to thank all participants for working hard on their submissions and adapting their systems not only to the datasets available, but also to the evaluation platform. We would like to thank all of them for their effort, since it is the participants who are the core of any shared task's success.

We would like to thank the CoNLL organizers for their support and the reviewers for helping to improve the submitted system papers. Special thanks go to Martin Potthast of the TIRA platform for handling such a large number of systems, running often for several hours each, and for being very responsive and helpful to us and all system participants, round the clock during the evaluation week and beyond. We also thank to the 200+ people working on the Universal Dependencies project during the past three years, without whom there would be no data.

Jan Hajič, Daniel Zeman, Joakim Nivre, Filip Ginter, Slav Petrov, Milan Straka, Martin Popel, Eduard Bejček

Organizers of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies

Prague, June 2017

Chair:

Jan Hajič, Charles University

Management group:

Daniel Zeman, Charles University
Joakim Nivre, Uppsala University
Filip Ginter, University of Turku
Slav Petrov, Google
Milan Straka, Charles University
Martin Popel, Charles University
Eduard Bejček, Charles University (Shared Task Proceedings Publication Chair)

TIRA support:

Martin Potthast, University of Weimar

Support group:

Christopher Manning, Stanford University
Marie-Catherine de Marneffe, Ohio State University
Yoav Goldberg, Bar Ilan University
Reut Tsarfaty, Open University of Israel
Sampo Pyysalo, University of Cambridge
Francis Tyers, UiT Norgga árktalaš universitehta
Jenna Kanerva, University of Turku
Çağrı Çöltekin, University of Tübingen
Juhani Luotolahti, University of Turku

Program Committee:

Lauriane Aufrant, LIMSI-CNRS, DGA
Miguel Ballesteros, IBM Research
Tiberiu Boroș, Romanian Academy
Agnieszka Falenska, University of Stuttgart
Marcos Garcia, Universidade da Coruña
Filip Ginter, University of Turku
Giuseppe Celano, Leipzig University
Johannes Heinecke, Orange Labs
James Henderson, Xerox Research Centre Europe
Jenna Kanerva, University of Turku
Ömer Kırnap, Koç University
Adam Lopez, University of Edinburgh
Christopher Manning, Stanford University
Paola Merlo, University of Geneva
Dat Quoc Nguyen, Macquarie University
Joakim Nivre, Uppsala University
Hiroshi Noji, Nara Institute of Science and Technology
Jungyeul Park, University of Arizona
Martin Popel, Charles University
Sudeshna Sarkar, IIT Kharagpur
Sebastian Schuster, Stanford University
Djamé Seddah, Université Paris Sorbonne
Tianze Shi, Cornell University
Pavel Sofroniev, University of Tübingen
Francis Tyers, UiT Norgga árktalaš universitehta
David Vilares, Universidade da Coruña
Hao Wang, Shanghai JiaoTong University
Katsumasa Yoshikawa, IBM Research
Deniz Yuret, Koç University
Zdeněk Žabokrtský, Charles University
Lilja Øvrelid, University of Oslo

Table of Contents

Conference Program

Overview

CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies
Daniel Zeman, Martin Popel, Milan Straka, Jan Hajic, Joakim Nivre, Filip Ginter, Juhani Luotolahti, Sampo Pyysalo, Slav Petrov, Martin Potthast, Francis Tyers, Elena Badmaeva, Memduh Gokirmak, Anna Nedoluzhko, Silvie Cinkova, Jan Hajic jr., Jaroslava Hlavacova, Václava Kettnerová, Zdenka Uresova, Jenna Kanerva, Stina Ojala, Anna Missilä, Christopher D. Manning, Sebastian Schuster, Siva Reddy, Dima Taji, Nizar Habash, Herman Leung, Marie-Catherine de Marneffe, Manuela Sanguinetti, Maria Simi, Hiroshi Kanayama, Valeria dePaiva, Kira Droganova, Héctor Martínez Alonso, Çağrı Çöltekin, Umut Sulubacak, Hans Uszkoreit, Vivien Macketanz, Aljoscha Burchardt, Kim Harris, Katrin Marheinecke, Georg Rehm, Tolga Kayadelen, Mohammed Attia, Ali Elkahky, Zhuoran Yu, Emily Pitler, Saran Lertpradit, Michael Mandl, Jesse Kirchner, Hector Fernandez Alcalde, Jana Strnadová, Esha Banerjee, Ruli Manurung, Antonio Stella, Atsuko Shimada, Sookyoung Kwak, Gustavo Mendonca, Tatiana Lando, Rattima Nitisaroj and Josie Li

Systems

Stanford's Graph-based Neural Dependency Parser at the CoNLL 2017 Shared Task
Timothy Dozat, Peng Qi and Christopher D. Manning

Combining Global Models for Parsing Universal Dependencies
Tianze Shi, Felix G. Wu, Xilun Chen and Yao Cheng

IMS at the CoNLL 2017 UD Shared Task: CRFs and Perceptrons Meet Neural Networks
Anders Björkelund, Agnieszka Falenska, Xiang Yu and Jonas Kuhn

The HIT-SCIR System for End-to-End Parsing of Universal Dependencies
Wanxiang Che, Jiang Guo, Yuxuan Wang, Bo Zheng, Huaipeng Zhao, Yang Liu, Dechuan Teng and Ting Liu

A System for Multilingual Dependency Parsing based on Bidirectional LSTM Feature Representations
KyungTae Lim and Thierry Poibeau

Adversarial Training for Cross-Domain Universal Dependency Parsing
Motoki Sato, Hitoshi Manabe, Hiroshi Noji and Yuji Matsumoto

Parsing with Context Embeddings
Ömer Kırnap, Berkay Furkan Önder and Deniz Yuret

Tokenizing, POS Tagging, Lemmatizing and Parsing UD 2.0 with UDPipe
Milan Straka and Jana Straková

Systems (continued)

Systems (continued)

Corpus Selection Approaches for Multilingual Parsing from Raw Text to Universal Dependencies
Ryan Hornby, Clark Taylor and Jungyeul Park

From Raw Text to Universal Dependencies - Look, No Tags!
Miryam de Lhoneux, Yan Shao, Ali Basirat, Eliyahu Kiperwasser, Sara Stymne, Yoav Goldberg and Joakim Nivre

Initial Explorations of CCG Supertagging for Universal Dependency Parsing
Burak Kerim Akkuş, Heval Azizoglu and Ruket Cakici

CLCL (Geneva) DINN Parser: a Neural Network Dependency Parser Ten Years Later
Christophe Moor, Paola Merlo, James Henderson and Haozhou Wang

A Fast and Lightweight System for Multilingual Dependency Parsing
Tao Ji, Yuanbin Wu and Man Lan

The ParisNLP entry at the ConLL UD Shared Task 2017: A Tale of a #Parsing-Tragedy
Eric De La Clergerie, Benoît Sagot and Djamé Seddah

Universal Joint Morph-Syntactic Processing: The Open University of Israel's Submission to The CoNLL 2017 Shared Task
Amir More and Reut Tsarfaty

A Semi-universal Pipelined Approach to the CoNLL 2017 UD Shared Task
Hiroshi Kanayama, Masayasu Muraoka and Katsumasa Yoshikawa

A rule-based system for cross-lingual parsing of Romance languages with Universal Dependencies
Marcos Garcia and Pablo Gamallo