

CoNLL 2022

**The 26th Conference on Computational Natural Language
Learning**

Proceedings of the Conference

December 7-8, 2022

©2022 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-959429-07-4

Introduction

Welcome to the 26th edition of the Conference on Computational Natural Language Learning (CoNLL). For the fifth time in a row, CoNLL is collocated and co-organized with EMNLP. The COVID-19 pandemic is still not behind us, but possibilities of traveling and meeting in person are increasing again. As such, this year’s edition is hybrid with both in-person and online talks and poster presentations.

CoNLL 2022 follows CoNLL 2020 and CoNLL 2021 in making this the third edition that specifically “focuses on theoretically, cognitively and scientifically motivated approaches to computational linguistics.” Just like in the previous two editions, this new focus was specified in the call for papers, in the instructions to reviewers and area chairs and is emphasized in publicity around the conference. Following EMNLP, we had a hybrid call accepting both direct and ARR submissions.

We received 102 direct submissions and 5 ARR submissions. From the direct submissions, 94 were sent out to reviewers (the other 8 being either retracted or desk rejected). The 5 ARR submissions were directly sent to our area chairs with their ARR reviews. We accepted 28 papers, 26 being direct submissions and 2 ARR submissions. In addition, CoNLL will feature two keynote talks by Noah Goodman and Allyson Ettinger. We thank both of them for accepting our invitation and are looking forward to their talks. We furthermore would like to thank all members of our program committee, listed on page iv, and our Area Chairs for many of whom the schedule overlapped with their Summer Break, in alphabetical order: Andrew Caines, Tanmoy Chakraborty, Kai-wei Chang, Ryan Cotterell, Dan Goldwasser, Micha Elsner, Rob van der Goot, Jena Hwang, Nora Hollenstein, Dieuwke Hupkes, Joseph Le Roux, Dipendra Misra, Preslav Nakov, Nanyun Peng, Maja Popovic, Emily Prud’hommeaux, Roi Reichart, Nathan Schneider, Kevin Small, Rui Wang, Adina Williams, Mark Yatskar.

Special thanks go to our publication chair R. Thomas McCoy, our publicity chair Jack Hessel and Webmaster Jens Lemmens. Without them, these proceedings could not have been completed or authors and other interested community members would have missed important information.

We received many useful tips and pieces of information from last year’s organizers, Arianna Bisazza and Omri Abend as well as from SiGNLL President Julia Hockemaier and SiGNLL Chair Afra Alishahi. Thank you for your support!

We hope you enjoy these proceedings.

Antske Fokkens and Vivek Srikumar
CoNLL 2022 conference co-chairs

Organizing Committee

Conference Chairs

Antske Fokkens, Vrije Universiteit Amsterdam and Eindhoven University of Technology, The Netherlands
Vivek Srikumar, University of Utah, USA

Publicity Chair

Jack Hessel, Allen Institute for AI, USA

Publication Chair

R. Thomas McCoy, Princeton University, USA

Area Chairs

Andrew Caines, University of Cambridge, UK
Tanmoy Chakraborty, IIT Delhi, India
Kai-wei Chang, University of California Los Angeles, USA
Ryan Cotterell, ETH Zürich, Switzerland
Dan Goldwasser, Purdue University, USA
Micha Elsner, The Ohio State University, USA
Rob van der Goot, University of Copenhagen, Denmark
Jena Hwang, Allen Institute for AI, USA
Nora Hollenstein, University of Copenhagen, Denmark
Dieuwke Hupkes, Facebook AI Research, France
Joseph Le Roux, Universite Sorbonne Paris Nord-CNRS UMR 70301, France
Dipendra Misra, Microsoft Research, USA
Preslav Nakov, Mohamed Bin Zayed University of Artificial Intelligence, UAE
Nanyun Peng, University of California Los Angeles, USA
Maja Popovic, Dublin City University, Ireland
Emily Prud'hommeaux, Boston College, USA
Roi Reichart, Technion - Israel Institute of Technology, Israel
Nathan Schneider, Georgetown University, USA
Kevin Small, Amazon, USA
Rui Wang, Shanghai Jiao Tong University, China
Adina Williams, FAIR, USA
Mark Yatskar, University of Pennsylvania, USA

Invited Speakers

Noah Goodman, Stanford University, USA
Allyson Ettinger, University of Chicago, USA

Program Committee

Reviewers

Omri Abend, Rodrigo Agerri, Željko Agić, Chris Alberti, Afra Alishahi, Aida Amini, Mark Anderson, Reut Apel, Jun Araki, Christoph Aurnhammer

Miguel Ballesteros, Leslie Barret, Barend Beekhuizen, Robert Berwick, Yevgeni Berzak, Shohini Bhattasali, Eduardo Blanco, Bernd Bohnet, Gosse Bouma

José G. C. de Souza, Spencer Caplan, Giovanni Cassani, Rahma Chaabouni, Hanjie Chen, Xinchu Chen, Emmanuele Chersoni, Christos Christodouloupoulos, Alexander Clark, Ailís Cournane, Francisco M Couto

Forrest Davis, Angel Daza, Aniello De Santo, Steven Derby, Lucia Donatelli, Li Dong, Zi-Yi Dou, Rotem Dror, Andrew Drozdov, Li Du, Jonathan Dunn, Maria Pia di Buono

Akiko Espinosa Anke

Amir Feder, Andrea Fischer, Abdellah Fourtassi, Daniel Freudenthal, Yoshinari Fujinuma

Kim Gerdes, Kripabandhu Ghosh, Dan Goldwasser, Zebulon Goriely, Kartik Goyal, Onur Gungor, Ashim Gupta

Michael Hahn, Sadid A. Hasan, Daniel Hershcovich, Cong Duy Vu Hoang, Cuong Hoang, Mark Hopkins, Jennifer Hu, Kuan-Hao Huang

Ganesh Jawahar, Zhanming Jie, Lifeng Jin, Kristen Johnson, Jaap Jumelet

Diptesh Kanojia, Alina Karakanta, Casey Kennington, Tracy Holloway King, Amrith Krishna

John P. Lalor, Ni Lao, Alberto Lavelli, Phong Le, Jiaxuan Li, Tao Li, Marina Litvak, Nelson F. Liu, Zhengzhong Liu, Zoey Liu

Sean MacAvaney, Andreas Maletti, Stella Markantonatou, Bruno Martins, David Martins de Matos, Yevgen Matuskevych, Kate McCurdy, Stephen McGregor, Maitrey Mehta, Stephan Meylan, Petar Milin, Kanishka Misra, Daichi Mochihashi, Manuel Montes

Massimo Nicosia, Tong Niu

Tim O’Gorman, Fredrik Olsson, Yohei Oseki

Ankur Padia, Aishwarya Padmakumar, Alexandros Papangelis, Chan Young Park, Tiago Pimentel, Christopher Potts

Ella Rabinovich, Daniele Radicioni, Sagnik Ray Choudhury, Anthony Rios, Aiala Rosá

Elizabeth Salesky, Vicente Ivan Sanchez Carmona, Sashank Santhanam, Nikunj Saunshi, Thomas Schatz, William Schuler, Abu Awal Md Shoeb, Miikka Silfverberg, Efstathios Stamatatos, Mark Steedman, Karl Stratos, Yoshi Suhara, Alane Suhr

Anaïs Tack, Ran Tian, Amalia Toneva

Clara Vania, Shikhar Vashishth, Prashanth Vijayaraghavan, David Vilar, Esau Villatoro-Tello

Miaosen Wang, Rui Wang, Yizhong Wang, Taro Watanabe

Eduardo Xamena, Yang Xu

Fan Yin, Michael Yoder

Omnia Zayed, Rowan Zellers, Mike Zhang, Mozhi Zhang, Tianlin Zhang, Zhisong Zhang, Kai Zhao, Yichu Zhou, Vilém Zouhar

Table of Contents

<i>A Multilingual Bag-of-Entities Model for Zero-Shot Cross-Lingual Text Classification</i> Sosuke Nishikawa, Ikuya Yamada, Yoshimasa Tsuruoka and Isao Echizen	1
<i>Collateral facilitation in humans and language models</i> James Michaelov and Benjamin Bergen	13
<i>How Hate Speech Varies by Target Identity: A Computational Analysis</i> Michael Yoder, Lynnette Ng, David West Brown and Kathleen Carley	27
<i>Continual Learning for Natural Language Generations with Transformer Calibration</i> Peng Yang, Dingcheng Li and Ping Li	40
<i>That’s so cute!: The CARE Dataset for Affective Response Detection</i> Jane Yu and Alon Halevy	50
<i>A Fine-grained Interpretability Evaluation Benchmark for Neural NLP</i> Lijie Wang, Yaozong Shen, Shuyuan Peng, Shuai Zhang, Xinyan Xiao, Hao Liu, Hongxuan Tang, Ying Chen, Hua Wu and Haifeng Wang	70
<i>Towards More Natural Artificial Languages</i> Mark Hopkins	85
<i>Causal Analysis of Syntactic Agreement Neurons in Multilingual Language Models</i> Aaron Mueller, Yu Xia and Tal Linzen	95
<i>Combining Noisy Semantic Signals with Orthographic Cues: Cognate Induction for the Indic Dialect Continuum</i> Niyati Bafna, Josef van Genabith, Cristina España-Bonet and Zdeněk Žabokrtský	110
<i>Detecting Unintended Social Bias in Toxic Language Datasets</i> Nihar Sahoo, Himanshu Gupta and Pushpak Bhattacharyya	132
<i>Incremental Processing of Principle B: Mismatches Between Neural Models and Humans</i> Forrest Davis	144
<i>Parsing as Deduction Revisited: Using an Automatic Theorem Prover to Solve an SMT Model of a Minimalist Parser</i> Sagar Indurkha	157
<i>Entailment Semantics Can Be Extracted from an Ideal Language Model</i> William Merrill, Alex Warstadt and Tal Linzen	176
<i>On Neurons Invariant to Sentence Structural Changes in Neural Machine Translation</i> Gal Patel, Leshem Choshen and Omri Abend	194
<i>Shared knowledge in natural conversations: can entropy metrics shed light on information transfers?</i> Eliot Maës, Philippe Blache and Leonor Becerra	213
<i>Leveraging a New Spanish Corpus for Multilingual and Cross-lingual Metaphor Detection</i> Elisa Sanchez-Bayona and Rodrigo Agerri	228
<i>Cognitive Simplification Operations Improve Text Simplification</i> Eytan Chamovitz and Omri Abend	241

<i>On Language Spaces, Scales and Cross-Lingual Transfer of UD Parsers</i>	
Tanja Samardžić, Ximena Gutierrez-Vasques, Rob van der Goot, Max Müller-Eberstein, Olga Pelloni and Barbara Plank	266
<i>Visual Semantic Parsing: From Images to Abstract Meaning Representation</i>	
Mohamed Ashraf Abdelsalam, Zhan Shi, Federico Fancellu, Kalliopi Basioti, Dhaivat Bhatt, Vladimir Pavlovic and Afsaneh Fazly	282
<i>Syntactic Surprisal From Neural Models Predicts, But Underestimates, Human Processing Difficulty From Syntactic Ambiguities</i>	
Suhas Arehalli, Brian Dillon and Tal Linzen	301
<i>OpenStance: Real-world Zero-shot Stance Detection</i>	
Hanzi Xu, Slobodan Vucetic and Wenpeng Yin	314
<i>Optimizing text representations to capture (dis)similarity between political parties</i>	
Tanise Ceron, Nico Blokker and Sebastian Padó	325
<i>Computational cognitive modeling of predictive sentence processing in a second language</i>	
Umesh Patil and Sol Lago	339
<i>PIE-QG: Paraphrased Information Extraction for Unsupervised Question Generation from Small Corpora</i>	
Dinesh Nagumothu, Bahadorreza Ofoghi, Guangyan Huang and Peter Eklund	350
<i>Probing for targeted syntactic knowledge through grammatical error detection</i>	
Christopher Davis, Christopher Bryant, Andrew Caines, Marek Rei and Paula Buttery	360
<i>An Alignment-based Approach to Text Segmentation Similarity Scoring</i>	
Gerardo Ocampo Diaz and Jessica Ouyang	374
<i>Enhancing the Transformer Decoder with Transition-based Syntax</i>	
Leshem Choshen and Omri Abend	384
<i>Characterizing Verbatim Short-Term Memory in Neural Language Models</i>	
Kristijan Armeni, Christopher Honey and Tal Linzen	405

Program

Wednesday, December 7, 2022

09:00 - 09:10 *Opening Remarks*

09:10 - 10:30 *Keynote 1: Noah Goodman*

10:30 - 11:00 *Coffee Break*

11:00 - 12:30 *Oral Session 1: Machine Learning for NLP, Model Interpretation*

Continual Learning for Natural Language Generations with Transformer Calibration

Peng Yang, Dingcheng Li and Ping Li

Towards More Natural Artificial Languages

Mark Hopkins

Probing for targeted syntactic knowledge through grammatical error detection

Christopher Davis, Christopher Bryant, Andrew Caines, Marek Rei and Paula Buttery

Enhancing the Transformer Decoder with Transition-based Syntax

Leshem Choshen and Omri Abend

12:30 - 14:00 *Lunch Break*

14:00 - 15:30 *Oral Session 2: Multilingual Work and Translation*

A Multilingual Bag-of-Entities Model for Zero-Shot Cross-Lingual Text Classification

Sosuke Nishikawa, Ikuya Yamada, Yoshimasa Tsuruoka and Isao Echizen

Combining Noisy Semantic Signals with Orthographic Cues: Cognate Induction for the Indic Dialect Continuum

Niyati Bafna, Josef van Genabith, Cristina España-Bonet and Zdeněk Žabokrtský

On Neurons Invariant to Sentence Structural Changes in Neural Machine Translation

Gal Patel, Leshem Choshen and Omri Abend

Wednesday, December 7, 2022 (continued)

On Language Spaces, Scales and Cross-Lingual Transfer of UD Parsers

Tanja Samardžić, Ximena Gutierrez-Vasques, Rob van der Goot, Max Müller-Eberstein, Olga Pelloni and Barbara Plank

15:30 - 16:00 *Coffee Break*

16:00 - 17:30 *Poster Session 1: Virtual*

How Hate Speech Varies by Target Identity: A Computational Analysis

Michael Yoder, Lynnette Ng, David West Brown and Kathleen Carley

OpenStance: Real-world Zero-shot Stance Detection

Hanzi Xu, Slobodan Vucetic and Wenpeng Yin

Characterizing Verbatim Short-Term Memory in Neural Language Models

Kristijan Armeni, Christopher Honey and Tal Linzen

Parsing as Deduction Revisited: Using an Automatic Theorem Prover to Solve an SMT Model of a Minimalist Parser

Sagar Indurkha

An Alignment-based Approach to Text Segmentation Similarity Scoring

Gerardo Ocampo Diaz and Jessica Ouyang

A Fine-grained Interpretability Evaluation Benchmark for Neural NLP

Lijie Wang, Yaozong Shen, Shuyuan Peng, Shuai Zhang, Xinyan Xiao, Hao Liu, Hongxuan Tang, Ying Chen, Hua Wu and Haifeng Wang

Thursday, December 8, 2022

09:10 - 10:30 *Keynote 2: Allyson Ettinger*

10:30 - 11:00 *Coffee Break*

11:00 - 12:30 *Poster session 2: In-person*

That's so cute!: The CARE Dataset for Affective Response Detection
Jane Yu and Alon Halevy

Causal Analysis of Syntactic Agreement Neurons in Multilingual Language Models
Aaron Mueller, Yu Xia and Tal Linzen

Detecting Unintended Social Bias in Toxic Language Datasets
Nihar Sahoo, Himanshu Gupta and Pushpak Bhattacharyya

Leveraging a New Spanish Corpus for Multilingual and Cross-lingual Metaphor Detection
Elisa Sanchez-Bayona and Rodrigo Agerri

Cognitive Simplification Operations Improve Text Simplification
Eytan Chamovitz and Omri Abend

Optimizing text representations to capture (dis)similarity between political parties
Tanise Ceron, Nico Blokker and Sebastian Padó

PIE-QG: Paraphrased Information Extraction for Unsupervised Question Generation from Small Corpora
Dinesh Nagumothu, Bahadorreza Ofoghi, Guangyan Huang and Peter Eklund

12:30 - 14:00 *Lunch Break*

14:00 - 15:30 *Oral Session 3: Psycholinguistics and Language Models*

Collateral facilitation in humans and language models
James Michaelov and Benjamin Bergen

Thursday, December 8, 2022 (continued)

Incremental Processing of Principle B: Mismatches Between Neural Models and Humans

Forrest Davis

Syntactic Surprisal From Neural Models Predicts, But Underestimates, Human Processing Difficulty From Syntactic Ambiguities

Suhas Arehalli, Brian Dillon and Tal Linzen

Computational cognitive modeling of predictive sentence processing in a second language

Umesh Patil and Sol Lago

15:30 - 16:00 *Coffee Break*

16:00 - 17:00 *Oral Session 4: Semantics and Grounding*

Entailment Semantics Can Be Extracted from an Ideal Language Model

William Merrill, Alex Warstadt and Tal Linzen

Shared knowledge in natural conversations: can entropy metrics shed light on information transfers?

Eliot Maës, Philippe Blache and Leonor Becerra

Visual Semantic Parsing: From Images to Abstract Meaning Representation

Mohamed Ashraf Abdelsalam, Zhan Shi, Federico Fancellu, Kalliopi Basioti, Dhairat Bhatt, Vladimir Pavlovic and Afsaneh Fazly

17:00 - 17:15 *Best Paper Awards and Closing*