

CODI-CRAC 2026

**The 2nd Joint Workshop on Computational Approaches to  
Discourse, Context and Document-Level Inferences and  
Computational Models of Reference, Anaphora and  
Coreference (CODI-CRAC 2026)**

**Proceedings of the Workshop**

July 3, 2026

©2026 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)  
317 Sidney Baker St. S  
Suite 400 - 134  
Kerrville, TX 78028  
USA  
Tel: +1-855-225-1962  
[acl@aclweb.org](mailto:acl@aclweb.org)

ISBN 979-8-89176-400-2

## Preface

Welcome to the 2nd Joint Workshop on Computational Approaches to Discourse, Context and Document-Level Inferences (7th CODI) and Computational Models of Reference, Anaphora and Coreference (9th CRAC) — CODI-CRAC 2026, co-located with ACL 2026 in San Diego, California, USA!

Following the successful first joint edition at EMNLP 2025 in Suzhou, China, CODI-CRAC 2026 continues to bring together researchers working on all aspects of discourse, context, and document-level inferences, alongside research on computational models of reference, anaphora, and coreference. The CODI workshop has long provided a platform for a broad range of research at the discourse level, while CRAC has been a primary venue for exploring reference-related phenomena. By joining forces, CODI-CRAC aims to advance discourse-level research further, fostering discussions on the development of datasets, tools, theories and robust evaluation methods.

The workshop is planned as a one-day event, featuring two invited talks, regular papers, and the CRAC 2026 Shared Task on Multilingual Coreference Resolution. This year we received submissions spanning a wide range of topics, including discourse representation and parsing, connectives, entity tracking, discourse benchmarks, ambiguity, stance analysis, dialogue, applications, and large language models for long-text understanding and generation. All accepted papers are presented live during the workshop, with Q&A sessions.

We received a total of 32 submissions and accepted 8 regular long papers, 3 regular short papers, 1 demo paper, 8 shared task papers as well as 12 non-archival contributions (Findings papers and extended abstracts). Among the non-archival submissions, the program also includes presentations of papers accepted at other major conferences.

This year, the workshop hosts the CRAC 2026 Shared Task on Multilingual Coreference Resolution (<https://ufal.mff.cuni.cz/corefud/crac26>), whose submissions are part of the joint proceedings volume and cover 19 languages tackled by 10 LLM-based and unconstrained systems.

We thank our invited speakers: Ruihong Huang, Associate Professor at Texas A&M University and Philippe Laban, Research Scientist at Microsoft Research, New York. They helped us prepare an excellent and well-rounded workshop program. We also thank the ACL 2026 workshop chairs for their support.

Finally, we thank our sponsor, the Heidelberg Institute for Theoretical Studies (HITS) (<https://www.h-its.org>), for supporting the Student Diversity & Inclusion Grant.

The CODI-CRAC 2026 Organizers:

Chloé Braud, Christian Hardmeier, Chuyuan Li, Junyi Jessy Li, Sharid Loáiciga, Vincent Ng, Michal Novák, Maciej Ogrodniczuk, Massimo Poesio, Michael Strube and Amir Zeldes

# Organizing Committee

## General Chair

Chloé Braud, IRIT - CNRS - ANITI

Christian Hardmeier, IT University of Copenhagen

Junyi Jessy Li, University of Texas at Austin

Chuyuan Li, The University of British Columbia

Sharid Loáiciga, University of Gothenburg

Michal Novák, Charles University

Vincent Ng, Univeristy of Texas at Dallas

Maciej Odrodniczuk, Institute of Computer Science, Polish Academy of Sciences

Massimo Poesio, Queen Mary University of London and University of Utrecht

Michael Strube, Heidelberg Institute for Theoretical Studies

Amir Zeldes, Georgetown University

# Program Committee

## Program Committee

Jelke Bloem, University of Amsterdam  
Antoine Bourgois, Lattice (CNRS and ENS-PSL and Univ. Sorbonne Nouvelle)  
Giuseppe Carenini, university of british columbia  
Jackie Chi Kit Cheung, Mila / McGill University  
Loic De Langhe, Ghent University  
Elisa Ferracane, Abridge AI, Inc.  
Mark Finlayson, Florida International University  
Annemarie Friedrich, University of Augsburg  
Yulia Grishina, Amazon  
Wenwen Guan, University of Amsterdam  
Jie He, University of Edinburgh  
Lars Hellan, NTNU  
Veronique Hoste, LT3, Ghent University  
Hyewon Jang, University of Gothenburg  
Sungho Jeon, Samsung Electronics  
Yangfeng Ji, University of Virginia  
Miloslav Konopík, University of West Bohemia in Pilsen  
Sobha Lalitha Devi, AU-KBC Research Centre, Anna University  
Ekaterina Lapshinova-Koltunski, University of Hildesheim  
Jiří Mírovský, Charles University  
Costanza Navarretta, University of Copenhagen  
Anna Nedoluzhko, Charles University  
Constantin Orasan, University of Surrey  
Gerald Penn, University of Toronto  
Lucie Polakova, Charles University  
Martin Popel, Charles University  
Ian Porada, Mila, McGill University  
Ondrej Prazak, University of West Bohemia in Pilsen  
Hannah Rohde, University of Edinburgh  
Ahmed Ruby, Uppsala university  
Tatjana Scheffler, Ruhr University Bochum  
Hannah Seemann, Ruhr-University  
Manish Shrivastava, International Institute of Information Technology Hyderabad  
Jakub Sido, University of West Bohemia  
Manfred Stede, University of Potsdam  
Bonnie Webber, University of Edinburgh  
Suncheng Xiang, Shanghai Jiao Tong University  
Frances Yung, Saarland University  
Daniel Zeman, Charles University  
Deniz Zeyrek, Boğaziçi University and Middle East Technical University  
Zheng Zhao, University of Edinburgh  
Heike Zinsmeister, University of Hamburg  
Bartosz Żuk, Institute of Computer Science Polish Academy of Sciences  
Zdeněk Žabokrtský, Charles University

# Invited Talk

## Discourse Structure Guided NLP Models for Fine-grained Media Bias Analysis

Ruihong Huang  
Texas A&M University



**Fri, July 3, 2026 – Time: 09:10 – 10:00 – Room: Regatta**

**Abstract:** Thinking about more and more powerful pretrained models and an increasing context window size supported by recent LLMs, do we still need explicit discourse structures to guide semantic reasoning? In this talk, I will present our research on fine-grained sentence-level media bias analysis that shows that incorporating shallow discourse structures or event relation graphs enables NLP models to better understand broader context and recognize subtle sentence-level ideological bias. News media play a vast role in shaping public opinion not just by supplying information, but by selecting, packaging, and shaping that information to persuade as well. Sentence-level media bias analysis is challenging and aims to identify sentences within an article that can illuminate and explain the overall bias of the entire article. This talk will first show that understanding the discourse role of a sentence in telling a news story, as well as its discourse relation with nearby sentences, can help reveal the ideological leanings of the author even when the sentence itself appears merely neutral or factual. This talk will further show that analyzing events with respect to other events in the same document or across documents is critical for identifying bias sentences.

**Bio:** Dr. Ruihong Huang is an associate professor in the Department of Computer Science & Engineering at Texas A&M University (TAMU), College Station. She is also an adjunct associate professor in McWilliams School of Biomedical Informatics at UTHealth Houston. Huang received her PhD in computer science at the University of Utah and completed a postdoc at Stanford University. Her research is focused on event-centric NLP, discourse analysis, dialogue and pragmatics, LLM evaluation, LLM safety and moral reasoning. She is a recipient of the US National Science Foundation CAREER award (2020).

**Invited Talk**  
**Studying and Simulating Discourse Failures in Interactive AI  
Systems**

**Philippe Laban**  
Microsoft Research, New York



**Fri, July 3, 2026 – Time: 13:30 – 14:20 – Room: Regatta**

**Abstract:** In this talk, we will go over two projects that study human-AI communication in conversational interfaces (e.g., ChatGPT). With Lost in Conversation, we simulate multi-turn underspecified conversations that involve a user gradually revealing their full intent over time to the AI system, finding that though this is common in real-world interactions, this leads to large degradation in model performance. We then discuss the need for more realistic user simulation that can model user behaviors specifically in the case of User-AI conversational exchange, and discuss its role in evaluating AI systems.

**Bio:** Philippe is a Research Scientist at Microsoft Research, based in New York and a member of the AI Interaction and Learning (MSR AILL) group. His research lies at the intersection of two fields: NLP and HCI, with a general interest in understanding how NLP/AI technology can enable effective consumption and production of information for users (which has involved working on topics such as summarization, simplification, question generation, etc.). Recently, he has taken an interest in studying human-AI interaction in conversational settings.

## Table of Contents

<i>Baselines for Detection and Classification of Discourse Presentation in English Narrative</i> Reinaldo Di Polo, Mustafa Ocal and Mark Finlayson .....	1
<i>DiscoExplorer: An Open Interface for the Study of Multilingual Discourse Relations</i> Amir Zeldes .....	12
<i>Speech Disfluencies and LLM Confidence: Length Bias and Pragmatic Insensitivity in Brazilian Portuguese</i> Valeria Santos .....	24
<i>DM Doppelgangers: Implicit Connectives as eRST Signals</i> Lin Ai and Amir Zeldes .....	29
<i>What’s in a Bridge?: A Descriptive, Multi-Genre Analysis of the GUMBridge Corpus for Varieties of Bridging Anaphora</i> Lauren Levine and Amir Zeldes .....	40
<i>Dataset Cartography for Implicit Discourse Relation Recognition: Promises and Pitfalls</i> Daniil Ignatev, Denis Paperno and Massimo Poesio .....	53
<i>Universal Discourse Relations: A Proposal</i> Anna Latusek, Maciej Ogrodniczuk, Alina Wróblewska and Bartosz Żuk .....	65
<i>A First Step towards Dialog Simulation with Grounded Dialog Graphs</i> Michael Ginn, Matt Pauk, Tava Reese, Sameer Gupta, Giuseppe Castellucci, Kevin Small, Alessandro Moschitti, Derek Palmer, Martha Palmer, Alexis Palmer and Maria Pacheco .....	78
<i>Errors in coreference resolution in German: Effects of modality, simplification and heterogeneous training data</i> Sarah Jablotschkin, Ekaterina Lapshinova-Koltunski and Heike Zinsmeister .....	109
<i>Evaluating Pragmatic Reasoning in Large Language Models: Evidence from Scalar Diversity</i> Ye-Eun Cho .....	120
<b>CRAC 2026 Shared Task on Multilingual Coreference Resolution</b>	
<i>Findings of the Fifth Shared Task on Multilingual Coreference Resolution: Expanding Datasets for Long-Range Entities</i> Michal Novák, Miloslav Konopík, Anna Nedoluzhko, Martin Popel, Ondrej Prazak, Jakub Sido, Milan Straka, Zdeněk Žabokrtský and Daniel Zeman .....	130
<i>Generative Multilingual Coreference Resolution at CRAC 2026</i> Jakub Hejman, Ondrej Prazak and Miloslav Konopík .....	162
<i>Closing the Gap at CRAC 2026: Two-Stage Adaptation for LLM-Based Multilingual Coreference Resolution</i> Antoine Bourgois, Olga Semineck and Thierry Poibeau .....	167
<i>Landcore: Coreference Resolution with Language-Specific LLM-Enhanced Prompts and XML-Inspired Annotation Scheme</i> Jan Pavelka .....	184

<i>PortNLP at CRAC 2026: QLoRA Fine-Tuning with Bounded Entity Registry for Multilingual Coreference Resolution</i>	
Amber Shore, Russell Scheinberg, Malini Nagasundaram and Ameeta Agrawal . . . . .	193
<i>Lightweight Multilingual Coreference Resolution without LLMs @CRAC2026</i>	
Sobha Lalitha Devi, Aashik Ali S, Gopinath P, Vijay Sundar Ram and Pattabhi Rk Rao . . . . .	199
<i>CorPipe at CRAC 2026: Empty Nodes and Cross-Lingual Transfer in Multilingual Coreference Resolution</i>	
Milan Straka . . . . .	205
<i>Closing the Gap: Robust Multilingual Coreference Resolution with DAgger</i>	
Thomas Morton and Alex Warstadt . . . . .	217

# Program

## Friday, July 3, 2026

- 09:00 - 09:10     *Opening Remarks*
- 09:10 - 10:00     *Invited Talk 1 — Ruihong Huang: Discourse Structure Guided NLP Models for Fine-grained Media Bias Analysis*
- 10:00 - 10:30     *Shared Task Overview*
- Findings of the Fifth Shared Task on Multilingual Coreference Resolution: Expanding Datasets for Long-Range Entities*  
Michal Novák, Miloslav Konopík, Anna Nedoluzhko, Martin Popel, Ondrej Prazak, Jakub Sido, Milan Straka, Zdeněk Žabokrtský and Daniel Zeman
- 10:30 - 11:00     *Coffee Break*
- 11:00 - 12:00     *Session - Oral presentations*
- What's in a Bridge?: A Descriptive, Multi-Genre Analysis of the GUMBridge Corpus for Varieties of Bridging Anaphora*  
Lauren Levine and Amir Zeldes
- Universal Discourse Relations: A Proposal*  
Anna Latusek, Maciej Ogrodniczuk, Alina Wróblewska and Bartosz Żuk
- A First Step towards Dialog Simulation with Grounded Dialog Graphs*  
Michael Ginn, Matt Pauk, Tava Reese, Sameer Gupta, Giuseppe Castellucci, Kevin Small, Alessandro Moschitti, Derek Palmer, Martha Palmer, Alexis Palmer and Maria Pacheco
- Speech Disfluencies and LLM Confidence: Length Bias and Pragmatic Insensitivity in Brazilian Portuguese*  
Valeria Santos
- 12:00 - 13:30     *Lunch*
- 12:30 - 13:30     *Virtual poster session 1*
- Modeling Multimodal Discourse Coherence in Low-Resource Languages: Work-in-Progress*  
Mahule Roy and Subhas Roy

**Friday, July 3, 2026 (continued)**

*Baselines for Detection and Classification of Discourse Presentation in English Narrative*

Reinaldo Di Polo, Mustafa Ocal and Mark Finlayson

*Dataset Cartography for Implicit Discourse Relation Recognition: Promises and Pitfalls*

Daniil Ignatev, Denis Paperno and Massimo Poesio

*Errors in coreference resolution in German: Effects of modality, simplification and heterogeneous training data*

Sarah Jablotschkin, Ekaterina Lapshinova-Koltunski and Heike Zinsmeister

*BeDiscover: The Benchmark of Discourse Understanding in the Era of Reasoning Language Models*

Chuyuan Li and Giuseppe Carenini

*Multilingual Extraction and Recognition of Implicit Discourse Relations in Speech and Text*

Ahmed Ruby, Christian Hardmeier and Sara Stymne

*Evaluating Pragmatic Reasoning in Large Language Models: Evidence from Scalar Diversity*

Ye-Eun Cho

*Generative Multilingual Coreference Resolution at CRAC 2026*

Jakub Hejman, Ondrej Prazak and Miloslav Konopík

*Lightweight Multilingual Coreference Resolution without LLMs @CRAC2026*

Sobha Lalitha Devi, Aashik Ali S, Gopinath P, Vijay Sundar Ram and Pattabhi Rk Rao

13:30 - 14:20 *Invited Talk 2 — Philippe Laban: Studying and Simulating Discourse Failures in Interactive AI Systems*

14:20 - 15:30 *Poster Session 1*

*Anthropomorphic Uncertainty: What Verbalized Uncertainty in Language Models is Missing*

Dennis Ulmer, Alexandra Lorson, Ivan Titov and Christian Hardmeier

*Expect the Unexpected? Testing the Surprisal of Salient Entities*

Jessica Lin and Amir Zeldes

**Friday, July 3, 2026 (continued)**

*Do Language Models Use Logophoric Cues? Evidence from Mandarin Chinese Long-Distance Reflexive*

Yunfang Dong

*Speech Disfluencies and LLM Confidence: Length Bias and Pragmatic Insensitivity in Brazilian Portuguese*

Valeria Santos

*The lady doth protest too much: discourse expectations affect judgments of argument strength*

S.r. Pulapura, Hannah Rohde and Catherine Lai

*Strategic Dialogue Assessment: The Crooked Path to Innocence*

Anshun Zheng, Junyi Jessy Li and David Beaver

*DiscoExplorer: An Open Interface for the Study of Multilingual Discourse Relations*

Amir Zeldes

*Closing the Gap at CRAC 2026: Two-Stage Adaptation for LLM-Based Multilingual Coreference Resolution*

Antoine Bourgois, Olga Seminck and Thierry Poibeau

*CorPipe at CRAC 2026: Empty Nodes and Cross-Lingual Transfer in Multilingual Coreference Resolution*

Milan Straka

15:30 - 16:00 *Coffee Break*

16:00 - 17:00 *Poster Session 2*

*DM Doppelgangers: Implicit Connectives as eRST Signals*

Lin Ai and Amir Zeldes

*What's in a Bridge?: A Descriptive, Multi-Genre Analysis of the GUMBridge Corpus for Varieties of Bridging Anaphora*

Lauren Levine and Amir Zeldes

*Universal Discourse Relations: A Proposal*

Anna Latusek, Maciej Ogródniczuk, Alina Wróblewska and Bartosz Żuk

**Friday, July 3, 2026 (continued)**

*Investigating discourse-level lexical semantic adaptation: A definition modeling approach*

Bill Noble and Brandon Roy

*A First Step towards Dialog Simulation with Grounded Dialog Graphs*

Michael Ginn, Matt Pauk, Tava Reese, Sameer Gupta, Giuseppe Castellucci, Kevin Small, Alessandro Moschitti, Derek Palmer, Martha Palmer, Alexis Palmer and Maria Pacheco

*Semantic-pragmatic Annotations in the Prague Dependency Treebank*

Marie Mikulová, Eva Hajicova, Jiří Mírovský, Anna Nedoluzhko, Michal Novák, Pavlína Synková, Jan Štěpánek, Barbora Štěpánková and Jan Hajič

*Landcore: Coreference Resolution with Language-Specific LLM-Enhanced Prompts and XML-Inspired Annotation Scheme*

Jan Pavelka

*PortNLP at CRAC 2026: QLoRA Fine-Tuning with Bounded Entity Registry for Multilingual Coreference Resolution*

Amber Shore, Russell Scheinberg, Malini Nagasundaram and Ameeta Agrawal

*Closing the Gap: Robust Multilingual Coreference Resolution with DAgger*

Thomas Morton and Alex Warstadt

17:00 - 17:20 *Session - Best paper presentation*

*DM Doppelgangers: Implicit Connectives as eRST Signals*

Lin Ai and Amir Zeldes

17:20 - 17:30 *Closing Remarks*