CMCL 2024

The 13th edition of the Workshop on Cognitive Modeling and Computational Linguistics

Proceedings of the Workshop

August 15, 2024

This workshop was supported by JST, PRESTO Grant Number JPMJPR21C2.

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

ISBN 979-8-89176-143-8

Introduction

Welcome to the 13th edition of the Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2024)!

CMCL has traditionally been the workshop of reference for research at the intersection between Computational Linguistics and Cognitive Science. After a blank in 2023, we are thrilled to be back, hosting this event once again after two years.

This year, CMCL has experienced multiple *firsts*, making it a landmark edition in its history. First, the organization team has transitioned to a younger generation and adopted modern logistics, such as using OpenReview and allowing commitments via ACL Rolling Review, for the first time. Second, this is the first CMCL held in the age of large language models (LLMs), prompting us to focus on fundamental scientific questions (e.g., their alignment with human cognition/perception) regarding artificial intelligence and cognitive science. Third, this is also the first CMCL held in Asia, marking a new geographical milestone for the workshop. Lastly, we received a record number of 55 submissions (37 regular submissions and 18 cross-submissions, including Findings papers), nearly doubling the submission number in the previous edition, providing a testament to the growing interest in this scientific, interdisciplinary field and the need for the dedicated workshop even in the age of somewhat engineeringly-oriented LLMs.

Out of 37 regular submissions, 34 papers are via direct submission (including 1 paper withdrawn before reviewing), and 3 papers are through the ARR commitment. We accepted 23 papers, resulting in an acceptance rate of 23/36=63.9%, slightly higher than in previous years. Additionally, 12 non-archival, cross-submissions were accepted and will be presented during the poster sessions. We are excited to have a diverse set of topics, including but not limited to, sentence processing, language acquisition, and new investigations powered by modern (multimodal) LLMs, covered in this year's program.

We extend our deepest gratitude to the Program Committee members; their dedication and expertise are the backbone of CMCL's success. We also express our sincere thanks to our invited speakers, Dr. Frank Keller, Dr. Aida Nematzadeh, and Dr. Sandro Pezzelle, for their valuable contributions to this year's program.

Lastly, we are immensely grateful to our sponsor, the Japan Science and Technology Agency. Their generous support allows us to subsidize the participation of our invited speakers.

The CMCL 2024 Organizing Committee

Organizers:

Tatsuki Kuribayashi, Mohammed bin Zayed University of Artificial Intelligence
Giulia Rambelli, University of Bologna
Ece Takmaz, University of Amsterdam
Philipp Wicke, Ludwig Maximilian University
Yohei Oseki, University of Tokyo

Invited speakers:

Frank Keller, University of Edinburgh
Aida Nematzadeh, Google DeepMind
Sandro Pezzelle, University of Amsterdam

Program committee:

Raquel G. Alhama, University of Amsterdam
Afra Alishahi, Tilburg University
Ted Briscoe, Mohamed bin Zayed University of Artificial Intelligence
Christos Christodoulopoulos, Amazon
Claudia Collacciani, University of Bologna
Aniello De Santo, University of Utah
Vera Demberg, Universität des Saarlandes
Dota Tianai Dong, Max Planck Institute for Psycholinguistics
Micha Elsner, Ohio State University
Wenxi Fei, Hong Kong Polytechnic University
Meaghan Fowlie, Utrecht University
Robert Frank, Yale University
Richard Futrell, University of California, Irvine
Shubham Garg, Amazon
John T. Hale, Johns Hopkins University, University of Georgia and DeepMind
Nora Hollenstein, University of Copenhagen
Samar Husain, Indian Institute of Technology Delhi
Go Inoue, Mohamed bin Zayed University of Artificial Intelligence
Cassandra L Jacobs, State University of New York, Buffalo
Carina Kauf, Massachusetts Institute of Technology
Yova Kementchedjhieva, Mohamed bin Zayed University of Artificial Intelligence
Snigdha Khanna, Indiana University
Fajri Koto, Mohamed bin Zayed University of Artificial Intelligence
Sandra Kübler, Indiana University at Bloomington
Yu Xi Li, Hong Kong Polytechnic University and Xi'an Jiaotong University
Tal Linzen, New York University and Google
Ziqian Luo, Oracle
Shrirang Mhalgi, Indiana University
James A. Michaelov, University of California, San Diego
Sathvik Nair, University of Maryland
Ludovica Pannitto, University of Bologna
Tiago Pimentel, ETH Zurich
Laurent Prevot, Université d'Aix-Marseille
Rachel Ryskin, University of California at Merced
William Schuler, Ohio State University, Columbus
Dylan Scott, National Taiwan Normal University

Cory Shain, Massachusetts Institute of Technology
Adina Williams, FAIR
Xinchen Yang, University of Maryland
Ryo Yoshida, The University of Tokyo

Table of Contents

Archival papers

Non-archival papers

Program

Thursday, August 15, 2024

08:45 - 09:00 *Opening Remarks*

09:00 - 09:40 *Keynote talk by Dr. Sandro Pezzelle*

09:40 - 10:40 *Session 1: Oral Presentations*

Hierarchical syntactic structure in human-like language models
Michael Wolfman, Donald Dunagan, Jonathan Brennan and John T. Hale

Do large language models resemble humans in language use?
Zhenguang Cai, Xufeng Duan, David Haslett, Shuqi Wang and Martin Pickering

Evaluating Vision-Language Models on Bistable Images
Artemis Panagopoulou, Coby Melkin and Chris Callison-Burch

10:40 - 11:00 *Break*

11:00 - 12:20 *Session 2: Poster Presentations*

BAMBINO-LM: (Bilingual-)Human-Inspired Continual Pre-training of BabyLM
Zhewen Shen, Aditya Joshi and Ruey-Cheng Chen

The Curious Case of Representational Alignment: Unravelling Visio-Linguistic Tasks in Emergent Communication
Tom Kouwenhoven, Max Peeperkorn, Bram Van Dijk and Tessa Verhoef

Do LLMs Agree with Humans on Emotional Associations to Nonsense Words?
Yui Miyakawa, Chihaya Matsuhira, Hirotaka Kato, Takatsugu Hirayama, Takahiro Komamizu and Ichiro Ide

Predict but Also Integrate: an Analysis of Sentence Processing Models for English and Hindi
Nina Delcaro, Luca Onnis and Raquel G. Alhama

Transformer Attention vs Human Attention in Anaphora Resolution
Anastasia Kozlova, Albina Akhmetgareeva, Aigul Khanova, Semen Kudriavtsev and Alena Fenogenova

Thursday, August 15, 2024 (continued)

Daily auditory environments in French-speaking infants: A longitudinal dataset
Estelle Hervé, Clément François and Laurent Prevot

What does Kiki look like? Cross-modal associations between speech sounds and visual shapes in vision-and-language models
Tessa Verhoef, Kiana Shahrasbi and Tom Kouwenhoven

Modeling Overregularization in Children with Small Language Models
Akari Haga, Saku Sugawara, Akiyo Fukatsu, Miyu Oba, Hiroki Ouchi, Taro Watanabe and Yohei Oseki

Language models' probability distributions are calibrated to cognitive profiles: An investigation of the predictive power of surprisal and entropy
Patrick Haller, Lena Sophia Bolliger and Lena Ann Jäger

What Makes Language Models Good-enough?
Daiki Asami and Saku Sugawara

Tree-Planted Transformers: Unidirectional Transformer Language Models with Implicit Syntactic Supervision
Ryo Yoshida, Taiga Someya and Yohei Oseki

How Much Does Non-verbal Communication Conform to Entropy Rate Constancy?: A Case Study on Listener Gaze in Interaction
Yu Wang, Yang Xu, Gabriel Skantze and Hendrik Buschmeier

VerbCLIP: Improving Verb Understanding in Vision-Language Models with Compositional Structures
Hadi Wazni, Kin Ian Lo and Mehrnoosh Sadrzadeh

So many design choices: Improving and interpreting neural agent communication in signaling games
Timothée Bernard and Timothee Mickus

12:20 - 14:00 *Lunch break*

14:00 - 15:00 *Session 3: Oral Presentations*

Large language models fail to derive atypicality inferences in a human-like manner
Charlotte Kurch, Margarita Ryzhova and Vera Demberg

Thursday, August 15, 2024 (continued)

Diachronic change in verb usage statistics predicts differences in sentence processing across the lifespan
Ellis Cain and Rachel Ryskin

How can large language models become more human?
Daphne Wang, Mehrnoosh Sadrzadeh, Miloš Stanojević, Wing-Yee Chow and Richard Breheny

15:00 - 15:40 *Invited talk by Dr. Frank Keller*

15:40 - 16:00 *Break*

16:00 - 17:20 *Session 4: Poster Presentations*

Locally Biased Transformers Better Align with Human Reading Times
Andrea Gregor De Varda and Marco Marelli

Evaluating Lexical Aspect with Large Language Models
Bolei Ma

Analysing and Validating Language Complexity Metrics Across South American Indigenous Languages
Felipe Ribas Serras, Miguel De Mello Carpi, Matheus Castello Branco and Marcelo Finger

Morphology Matters: Probing the Cross-linguistic Morphological Generalization Abilities of Large Language Models through a Wug Test
Dang Thi Thao Anh, Limor Raviv and Lukas Galke

Evaluating Grammatical Well-Formedness in Large Language Models: A Comparative Study with Human Judgments
Zhuang Qiu, Xufeng Duan and Zhenguang Cai

Evaluating Semantic Relations in Predicting Textual Labels for Images of Abstract and Concrete Concepts
Tarun Tater, Sabine Schulte Im Walde and Diego Frassinelli

How Useful is Context, Actually? Comparing LLMs and Humans on Discourse Marker Prediction
Emily Sadlier-Brown, Millie Lou, Miikka Silfverberg and Carla L. Hudson Kam

Thursday, August 15, 2024 (continued)

LLMs' morphological analyses of complex FST-generated Finnish words
Anssi Moisio, Mathias Creutz and Mikko Kurimo

An Eye Opener Regarding Task-Based Text Gradient Saliency
Guojun Wu, Lena Sophia Bolliger, David Robert Reich and Lena Ann Jäger

Improving Language Models for Emotion Analysis: Insights from Cognitive Science
Constant Bonard and Gustave Cortal

Structural Similarities Between Language Models and Neural Response Measurements
Antonia Karamolegkou, Jiaang Li, Yova Kementchedjhieva, Mostafa Abdou, Sune Lehmann and Anders Søgaard

Exploring Spatial Schema Intuitions in Large Language and Vision Models
Philipp Wicke and Lennart Wachowiak

Can You Learn Semantics Through Next-Word Prediction? The Case of Entailment
William Merrill, Zhaofeng Wu, Norihito Naka, Yoon Kim and Tal Linzen

The Emergence of High-Level Semantics in a Signaling Game
Timothée Bernard, Timothee Mickus and Hiroya Takamura

PUB: A Pragmatics Understanding Benchmark for Assessing LLMs' Pragmatics Capabilities
Settaluri Lakshmi Sravanthi, Meet Doshi, Pavan Kalyan Tankala, Rudra Murthy, Raj Dabre and Pushpak Bhattacharyya

17:20 - 18:00 *Invited talk by Dr. Aida Nematzadeh*

18:00 - 18:10 *Closing Remarks*