

LM Playschool 2026

**The LM Playschool Workshop: Improving Language Models
through Learning from Dialogue Interaction**

Proceedings of the Workshop

October 28, 2026

The LM Playschool organizers gratefully acknowledge the support from the following sponsors.



©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

ISBN 979-8-89176-467-5

Introduction

This is the first edition of the LM Playschool Workshop, co-located with the 2026 Conference on Empirical Methods in Natural Language Processing on October 28, 2026 in Budapest, Hungary.

The LM Playschool Workshop aims to bring together researchers interested in language learning from social interaction in conversational, collaborative, task-oriented, multi-turn environments. Relevant research strands include: language games as a means of evaluating large language models and their ability to use language in interaction; connections between human language acquisition and machine language learning, with a particular focus on the role of social interaction in data-efficient learning; game-based and interactive environments for training large language models.

The workshop hosts a shared task, the LM Playschool Challenge. The shared task focuses on training collaborative, goal-oriented language agents through language interaction data. Submitted systems are evaluated for their ability to generalise to unseen “dialogue games”, i.e., goal-oriented tasks governed by multi-turn language interaction, as well as on traditional single-turn NLP benchmarks and cognitive tests.

We received 19 submissions across the three submission routes and accepted 14 papers: 6 to the main workshop track, 2 committed through ACL Rolling Review, and 6 shared task system descriptions.

The LM Playschool Workshop Organizing Committee: Alessandro Suglia, Mario Giulianelli, David Schlangen, Raffaella Bernardi, Raquel Fernández, Sherzod Hakimov, Sabrina McCallum, Davide Mazzaccara, Dieu-Thu Le, Oliver Lemon, Alexander Koller

Organizing Committee

Lead Organisers

Alessandro Suglia, University of Edinburgh, United Kingdom
Mario Giulianelli, University College London, United Kingdom

Steering Committee

David Schlangen, University of Potsdam, Germany
Raffaella Bernardi, Free University of Bozen-Bolzano, Italy
Raquel Fernández, University of Amsterdam, Netherlands

Shared Task Chair

Sherzod Hakimov, University of Potsdam, Germany

Diversity and Inclusion Chair

Sabrina McCallum, University of Edinburgh, United Kingdom

Web and Technical Chair

Davide Mazzaccara, University of Trento, Italy

OpenReview Chair

Dieu-Thu Le, Amazon - AGI, Germany

Publicity Chair

Oliver Lemon, Heriot-Watt University, United Kingdom

Finance Chair

Alexander Koller, Saarland University, Germany

Program Committee

Reviewers

Cengiz Acarturk

Raffaella Bernardi

Javier Chiyah-Garcia, Marko Culjak

Ondrej Dusek

Ambra Ferrari, Abdellah Fourtassi

Albert Gatt, Mario Giulianelli

Sherzod Hakimov, Julia Hockenmaier

Nikolai Ilinykh

Prashant Jayannavar

Alexander Koller, Udo Kruschwitz

Dieu-Thu Le

Davide Mazzaccara, Sabrina McCallum

Sandro Pezzelle, Massimo Poesio, Matthew Purver

Dimitrije Ristic

David Schlangen, Carina Silberer, Edwin Simpson, Alessandro Suglia, Michael Sullivan

Keynote Talk

On Causal & Social World Models in Language Models

Michael Franke
University of Tübingen



2026-10-28 – Time: 09:15 – 10:00

Abstract: Recent works have suggested that “world models” emerge in large language models as a side effect of the compression necessary for their high performance on the language modeling task. While robust evidence for such world models exists for relatively simple (e.g., finite-state) problems, it remains unclear whether we should expect LMs to also evolve (veridical / human-like) world models for human-relevant but highly abstract latent variables such as causal information, discourse states, or social cognition. I discuss in-principle considerations based on an augmented optimal learner approach that also takes architectural constraints of the learning system into account. This architecturally-constrained efficiency approach predicts that abstract function-sharing can be a reason for emergent representations of the causal data-generating process. To explore these predictions, I discuss two case studies, one on function-sharing between Theory of Mind reasoning and pragmatic reasoning in large language models, and one on the emergence of generalizable Bayesian inference of common causes in small transformers.

Bio: Michael Franke is a computational cognitive scientist, appointed as Professor of Pragmatics and Linguistics and coopted as a Professor of Computer Science at the University of Tübingen and head of the Cognitive Science & Pragmatics Lab. He is interested broadly in reasoning and language use in humans and machines. His work uses statistical, cognitive and evolutionary modeling, human-experimental methods, and simulation experiments. He is also interested in foundational questions concerning the relation between empirical data, abstract theory and explicit mathematical modeling. Currently, he is the coordinator of the DFG-funded Priority Area LaSTing and a co-speaker of the DFG-funded CRC Common Ground.

Keynote Talk

Learning to Interact, Interacting to Learn

Abdellah Fourtassi
Aix-Marseille University



2026-10-28 – Time: 11:00 – 11:45

Abstract: As children learn to coordinate turns, produce contingent responses, and use multimodal behavior, they become increasingly active participants in linguistic interaction. Their participation changes the evidence available to them. It elicits adapted input, feedback, alignment, and repair, which create new opportunities for learning. Interaction therefore provides a context in which different learning signals arise and become informative.

I present empirical and computational studies of this relationship. I use large-scale analyses of naturalistic child–caregiver interaction to examine the development of turn-taking, response contingency, and multimodal coordination. I then present work using child-like LMs and reinforcement learning to test whether and how different forms of naturally occurring feedback contribute to language learning.

This developmental perspective argues that learning to interact and acquiring functional linguistic knowledge are coupled objectives. It has implications for machine language learning, particularly for how models can ground their acquired knowledge in situated language use and shared goals.

Bio: Abdellah Fourtassi is an Associate Professor of Computer Science at Aix-Marseille University and a researcher at the Institute of Language, Communication and the Brain. His team studies how multimodal interaction shapes language learning. His research draws primarily on machine learning theory and methods to study human language development.

Keynote Talk

Language Learning through Social and Embodied Experience

Joyce Chai
University of Michigan



2026-10-28 – Time: 14:00 – 14:45

Abstract: Human language learning is shaped by social and embodied experience. Through social interaction, we acquire the use of language from communicative context, and through sensorimotor experience with the physical world, we learn words that are grounded to perception and action. In this talk, I will discuss several studies in my lab that explore social and embodied context for language learning and grounding. I will introduce an interactive learning paradigm and show how demonstrations and corrective feedback from teachers and active trials from students can facilitate efficient word learning in language models.

Bio: Joyce Chai is a Professor in the Department of Electrical Engineering and Computer Science at the University of Michigan. Her research interests span from natural language processing and embodied AI to human-AI collaboration. Her current work explores the intersection between language, vision, and action to enable situated communication with embodied agents. She is a recipient of the NSF Career Award and multiple awards with her students (e.g., Best Long Paper Award at ACL 2010, Outstanding Paper Awards at EMNLP 2021 and ACL 2023, and First Place in the Amazon Alexa AI Simbot Challenge 2023). She is a Fellow of ACL.

Keynote Talk

Data Efficiency in Children and Language Models

Michael C. Frank
Stanford University



2026-10-28 – Time: 16:00 – 16:45

Abstract: Large language models show intriguing emergent behaviors, yet they receive at least three to four – and sometimes as much as six – orders of magnitude more language data than human children. What accounts for this vast difference in sample efficiency? I will describe steps towards a paradigm in which we can address this question. In particular, I'll discuss the use of egocentric video (“baby head-cam”) data for model training, and the use of developmental data for model evaluation. This paradigm provides a model-based framework for exploring the nature of children’s early data efficiency.

Bio: Michael C. Frank is Benjamin Scott Crocker Professor of Human Biology in the Department of Psychology at Stanford University and Director of the Symbolic Systems Program. He received his PhD from MIT in Brain and Cognitive Sciences in 2010. He studies children’s language learning and development, with a focus on the use of large-scale datasets to understand the variability and consistency of learning across cultures. He is a founder of the ManyBabies Consortium, and has led open-data projects including Wordbank and the ongoing LEVANTE project. He has received awards including the Troland Award from the National Academy of Sciences and the FABBS Early Career Impact award. He served as President of the Cognitive Science Society, has edited for journals including *Cognition* and *Child Development*, and is current co-Editor in Chief of the *Open Encyclopedia of Cognitive Science*.

Table of Contents

LM Playschool Workshop

<i>Beyond Sally-Anne: Evaluating Theory of Mind in Large Language Models using Epistemic Schelling Points</i>	
Roberta Rocca, Sami Boukourt, Geoff Keeling and Winnie Street	1
<i>Can Faithful CoT Improve Multi-turn LLM Agents? Effectiveness, Boundaries, and Paradoxes</i>	
Tianle Gu, Xinqi Yang and Yan Teng	19
<i>ProcessPrefBench: Evaluating Cross-Task Transfer of Individual Process Preferences from Naturalistic Human–AI Interaction</i>	
Yiwen Yang and Guanming Liu	30
<i>Role-Specific Strengths and Alignment in LLM Cross-Play and Self-Play</i>	
Kata Naszadi and Christof Monz	41
<i>InterviewArena: Estimating Uninformative Participants through Interactive Dialogue under Incomplete Information</i>	
Raika Koki, Karina Karasawa, Ryoki Kanayama, Yuki Suzumura, Haruto Fujita, Ryoma Obara, Yusuke Sakai, Hidetaka Kamigaito, Katsuhiko Hayashi and Shogo Matsuno	49
<i>Protocol Competence of Small Language Models in Interactive Language Game Environments</i>	
Suchir Salhan and Paula Buttery	57
<i>Pareto-Dominant Clarification: Post-Training Coding LLMs via PPO-Lagrangian Budget Constraints</i>	
Abhinav Rajput and Acey Vogelstein	80
<i>An Empirical Study on Improving an LLMs based Interviewer Against Human Interviewees</i>	
Rikuto Tsuchida, Ryo Hasegawa, Tetsuya Sonoda and Takehito Utsuro	97

LM Playschool Challenge

<i>The LM Playschool Challenge 2026</i>	
Sherzod Hakimov, Dimitrije Ristic, Philipp Sadler, Davide Mazzaccara, Sabrina McCallum, Filippo Momentè, Maria Manina, Karl Osswald, Dieu-Thu Le, Oliver Lemon, Alexander Koller, Raffaella Bernardi, Raquel Fernández, Mario Giulianelli, Alessandro Suglia and David Schlangen	110
<i>Enhancing LLM Performance via LoRA SFT and DPO for the LM Playschool Challenge</i>	
Vasamsetti Nihal Tej, Velineni Harshavaishnav and Shreya Ghosh	127
<i>Adapting SCoRe to Single-Player Dialogue Games: Initial Experiments on PLAYPEN</i>	
Neda Foroutan and Valhari Meshram	132
<i>First Make It Playable, Then Make It Good: Staged Interaction Learning for Small Dialogue-Game Agents</i>	
Syed Mahbulul Huq and Pranava Madhyastha	140
<i>Raising a Small Language Model: From Imitation to Curiosity-Driven Learning in Dialogue Games</i>	
Varad Srivastava	148
<i>Teaching or Sharpening? An Exploration of the Potential of DPO Post-Training for Dialogue Games</i>	
Marika Sarzotti and Simone Baratella	160

Acquire, Repair, Preserve: A Diagnosis-Guided Post-Training Recipe for Small-Model Dialogue Game Agents

Nan Li 170

Program

Wednesday, October 28, 2026

09:00 - 09:15	<i>Welcome & Introduction</i>
09:15 - 10:00	<i>Invited Talk: Michael Franke</i>
10:00 - 10:30	<i>Flash Talks Non-Challenge Papers</i>
10:30 - 11:00	<i>Coffee Break</i>
11:00 - 11:45	<i>Invited Talk: Abdellah Fourtassi</i>
11:45 - 12:00	<i>Flash Talks Shared Task Submissions</i>
12:00 - 12:15	<i>Top 2 Shared Task Submissions</i>
12:15 - 12:30	<i>Top 1 Shared Task Submissions</i>
12:30 - 14:00	<i>Lunch</i>
14:00 - 14:45	<i>Invited Talk: Joyce Chai</i>
14:45 - 15:30	<i>Posters</i>
15:30 - 16:00	<i>Coffee Break</i>
16:00 - 16:45	<i>Invited Talk: Michael Frank</i>
16:45 - 17:30	<i>Panel Discussion</i>