

CHum 2026

The 2nd Workshop on Computational Humor

Proceedings of the Workshop

July 3, 2026

The CHum 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 979-8-89176-431-6

Preface

The first Computational Humor (CHum) Workshop that we organized took place at COLING 2025 in Abu Dhabi, UAE ([chum2025.github.io](https://github.com/chum2025)). It featured eleven selected papers and three invited keynotes that we thought represented the landscape of computational humor in the age of the “AI” turmoil that is still very much going on, but also pointed to future developments beyond prompting humor or humor analysis out of LLMs.

Then, as now, we take the position that in order to advance humor-contained interaction/generation/detection in computational systems, there is a need to understand where the systems fail and where they succeed. Relying solely on examples of impressively generated humorous texts is insufficient, as such outputs may stem from models’ capabilities in memorization and pattern matching rather than the acquisition of structured knowledge that demonstrates humor competence.

Because the aim of the workshop is to foster work on modeling the processes of humor with current methods in computational linguistics and natural language processing, papers to be accepted for presentation at the workshop had to be either based on current computational methods and informed by the existing humor research or advance the state of the art in humor research by applying current computational models. There are arguments that humor relies on the background knowledge that may not be available from a single modality, such as text, and as such is not machine-learnable from that modality in a straightforward manner. With this in mind, a principal goal of this workshop is to unite researchers who can together probe the limits of various meaning representations—symbolic, neural, hybrid, or unified—for humor processing.

Reliable and versatile humor-aware algorithms could have many uses. Commonly-cited examples, such as joke generation and the translation of humorous texts, tend to be application-oriented engineering tasks. But from another angle, computational humor, particularly if done in a symbolic (knowledge-based) fashion that captures the necessary and sufficient resources and algorithms, could help researchers in the social sciences and humanities better understand how human humor works. Along all these lines, knowledge-based work has progressed since the late nineties, from purely template-based to enhanced bag-of-word approaches to proposals for ontology-based systems. While this work has not yet left the arena of academic toy systems, there has been a slow but steady progression in the linguistic and computational complexity of the methods, as well as a broadening of application areas to include augmentative communication and computer-assisted translation.

As hinted at above, human humor is the result of a complex cognitive process that requires the experiencer to identify and act on culturally significant cues in the source material. Moreover, it depends on surprise and “violations of predictions”, making it difficult to capture in trained models that are at their core predictable and only pseudo-random. Nonetheless, recent work in conventional text domains on addressing the knowledge acquisition bottleneck, on highly fluent text generation, and on attention mechanisms and interpretable models, seem ripe for application to computational humor in supervised settings. And neurosymbolic hybrid approaches and unification approaches have recently entered the stage of computational humor research, where interesting new results should be expected in the near future.

This year, the second CHum (chumweb.org), again organized by the Semantic AI & Creativity Lab at East Texas A&M (etamu.edu/saicl) and the Computational Linguistics at Manitoba lab (clam.cs.umanitoba.ca), but now also sponsored by the Applied Knowledge Representation and Natural Language Understanding (AKRaNLU) Lab at Purdue University (polytechnic.purdue.edu/facilities/akranlu), will feature a similar number of papers that fit these aims. We also invited a keynote by a psychologist, and will round out the workshop with a discussion panel on computational humor and the law.

Finally, we would like to thank everyone who submitted a paper to the workshop, as well as the members of our Program Committee for their timely and insightful reviews.

Organizing Committee

Tristan Miller, University of Manitoba

Ori Amir, Fulbright University Vietnam

Julia Rayz, Purdue University

Tiansi Dong, Alan Turing Institute & University of Cambridge

Christian F. Hempelmann, East Texas A&M University

Program Committee

Adam Jatowt, University of Innsbruck
Anne-Gwenn Bosser, École Nationale d'Ingénieurs de Brest
Aram Sinnreich, American University
Elena Mikhalkova, Tyumen State University
Hongfei Lin, Dalian University of Technology
Joe Toplyn, Twenty Lane Media
Kory Mathewson, Google DeepMind
Larry Lefkowitz, Working Knowledge LLC
Liana Ermakova, University of Western Brittany
Liang Yang, Dalian University of Technology
Mirella Manfredi, Zurich University
Monika, JP Morgan
Nadezhda Ganzherli, Tyumen State University
Nathaniel Laywine, York University
Piotr Mirowski, Google DeepMind
Rada Mihalcea, University of Michigan
Sergio Spaccavento, Osservatorio Feldman
Sophie Jentzsch, German Aerospace Center (DLR)
Tatiana Ringenberg, Purdue University
Thomas Winters, KU Leuven
Tony Veale, University College Dublin
Vladislav Maraev, University of Gothenburg
Yi Zhang, Purdue University

Cognition Without Evolution: An Evolutionary Perspective on AI Humor

Gil Greengross
Aberystwyth University

Abstract: Recent advances in large language models have transformed research on computational humor and raised fundamental questions about the nature of humor itself. Human humor is widely viewed as an evolutionary adaptation that promotes social bonding, cooperation, status, and mate attraction. AI systems, however, lack the evolutionary history, emotions, and social motives that underlie these functions, yet they can generate, recognize, and evaluate humor. This talk explores AI humor through the lens of evolutionary psychology. I argue that AI provides a unique natural experiment for distinguishing between aspects of humor that depend on evolved human adaptations and those that reflect more general information-processing mechanisms. Drawing on research in computational humor, evolutionary psychology, and humor studies, I examine similarities and differences between biological and artificial systems, discuss evolved psychological adaptations that shape the way we assess and experience AI humor, and suggest ways in which computational approaches to humor can benefit from invoking an evolutionary perspective. The talk is exploratory and while it may not provide many answers, it is intended to identify new questions for interdisciplinary research.

Panel Computational Humor and the Law

As humor-aware AI systems move from research prototypes into public-facing applications, they raise complex questions at the intersection of language, law, creativity, and governance. This multidisciplinary panel brings together digital rights advocate Brad Templeton, television writer and Witscript creator Joe Toplyn, and legal scholar Laura E. Little to examine the social, ethical, and legal implications of computational humor. Topics may include free speech and content moderation, intellectual property and comedic style, liability for AI-generated jokes, and the challenges of humor across cultures and legal jurisdictions. Drawing on perspectives from online community governance, professional comedy writing, and legal doctrine, the panel will explore how humor-aware AI can be developed responsibly while preserving creativity, expression, and cultural nuance. The session will conclude with reflections on guiding principles for future computational humor research.

Panelists:

Brad Templeton, Electronic Frontier Foundation and rec.humor.funny

Joe Toplyn, Twenty Lane Media, LLC

Laura E. Little, Temple Law School

Table of Contents

<i>One Joke to Rule them All? On the (Im)possibility of Generalizing Humor Detection</i> Mor Turgeman, Chen Shani and Dafna Shahaf	1
<i>Timing In stand-up Comedy: Text, Audio, Laughter, Kinesics (TIC-TALK): Pipeline and Database for the Multimodal Study of Comedic Timing</i> Yaelle Zribi, Florian Cafiero, Vincent Lépinay and Chahan Vidal-Gorène	29
<i>Arabic Humor as a Diagnostic Probe for Large Language Models</i> Wajdi Zaghouani	39
<i>Cards Against LLMs: Benchmarking Humor Alignment in Large Language Models</i> Yousra Fettach, Guillaume Bied, Hannu Toivonen and Tijl De Bie	51
<i>The Roast of GPT4o: Experiments in Generating, Detecting and Evaluating Celebrity Roast Comedy</i> Jens Lemmens, Jérémy Genette, Tony Veale and Walter Daelemans	65
<i>Phonetic Cues Improve LLM-Based Pun Detection in Short Text</i> Adith Santosh Thaniserikaran and Govind Harikrishnan	72
<i>Does Bigger Mean Funnier? Evaluating Humor Generation Across the Qwen3 Model Family</i> Jatin Agrawal and Radhika Mamidi	81
<i>Navigating the Joke Space: Towards Automated Originality Assessment of AI-Generated Humor</i> Ori Amir, Huyen Ngo, Joe Toplyn and Kevin Hickerson	95