BigPicture 2026

The Big Picture v2: Crafting a Research Narrative

Proceedings of the Workshop

July 4, 2026

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 979-8-89176-416-3

Introduction

Welcome to the Proceedings of the first iteration of the Big Picture Workshop (The Big Picture: Crafting a Research Narrative). The workshop is hosted at ACL 2026, in San Diego, USA, on July 4th, 2026.

The Big Picture Workshop provides a dedicated venue for exploring and distilling broader NLP research narratives. All research exists within a larger context, and progress is made by standing on the shoulders of giants: building on the foundations laid by earlier researchers. In light of rapid publication rates and concise paper formats, it has become increasingly difficult, however, to recognize the larger story to which a paper is connected. The Big Picture Workshop invites researchers to reflect on how their individual contributions fit within the overall research landscape and what stories they are telling with their bodies of research. The goals of the workshop are to enhance communication and understanding between different lines of work, highlight how works connect and build on each other, generate insights that are difficult to glean without combining and reconciling different research narratives, encourage broader collaboration and awareness of prior work in the NLP community, and facilitate understanding of trajectories and insights within the field of NLP.

We received 15 submissions, of which we accepted 12 for presentation at the workshop. Those 12 accepted papers are contained in this volume.

The workshop schedule features one standard invited talk, and three special invited presentations designed to foster live engagement between different lines of related work. In these special presentations, two to three invited presenters speak on their individual lines of work and the connections between them, followed by a moderated discussion further exploring the overall narrative that emerges from these works in aggregate. In addition to invited presentations, the workshop features an in-person poster session, and spotlight talks.

We extend heartfelt thanks to our program committee, our participants, and all authors who submitted papers for consideration—your engagement has been critical to the success of the workshop. Finally, we thank the ACL 2026 organizers and workshop chairs for their hard work and support.

The Big Picture Workshop Organizers,

Yanai Elazar, Allyson Ettinger, Nora Kassner, Sebastian Ruder

Organizing Committee

Organizers

Yanai Elazar, Bar-Ilan University
Allyson Ettinger, Allen Institute for AI
Nora Kassner, Google DeepMind
Sebastian Ruder, Meta Superintelligence Labs

Program Committee

Reviewers

David Ifeoluwa Adelani, Maria Antoniak

Alexandra Chronopoulou, Marta R. Costa-jussà

Matthias Gallé

Tiago Pimentel, Barbara Plank

Philip Resnik, Kyle Richardson

Vered Shwartz

Yogarshi Vyas

Xinyi Wang, Chenxi Whitehouse

Keynote Talk

Where it Hurts: Finding Durable Questions While Moving Fast

Noah A. Smith
University of Washington & Allen Institute for Artificial Intelligence
2026-04-07 09:30:00 – Room: **TBD**

Abstract: In a fast-moving field, it can be hard to tell which problems are urgent, which are merely loud, and which are worth building a research life around. This talk considers how researchers can stay responsive to rapid change without letting the field's volatility set their agenda. I will discuss tools for identifying which questions remain meaningful across shifts in methods, data, benchmarks, and institutions. The goal is a practical vocabulary for finding direction when everything seems to be moving at once.

Bio: Noah A. Smith is the inaugural Vice Provost for Artificial Intelligence and Charles and Lisa Simonyi Endowed Chair for Artificial Intelligence and Emerging Technologies at the University of Washington, where he is also a Professor in the Paul G. Allen School of Computer Science & Engineering. He is Senior Director of NLP Research at the Allen Institute for Artificial Intelligence, directs the OLMo open language modeling effort, and leads the NSF- and NVIDIA-supported project "Open Multimodal AI Infrastructure to Accelerate Science." His research spans language and music technologies, multimodal AI, and multifaceted evaluation of AI systems.

Keynote Talk

Does mechanistic interpretability need interventions?

Aaron Mueller
Boston University
2026-04-07 10:50:00 – Room: **TBD**

Abstract: Mechanistic interpretability often treats interventions as the gold standard of evidence, relying on circuit ablations and representation steering to support claims about how models actually work. But are interventions really sufficient, or even necessary, for making mechanistic claims? In this debate-style talk, we trace the history that led the field to embrace interventions, and argue that the answer to both questions is a contentious no.We start by showing that interventions alone are not sufficient to explain model behavior by highlighting cases where causal methods can produce misleading or outright spurious explanations. Then, we debate whether interventions are necessary at all, exploring how alternative notions of causality and carefully designed behavioral evidence may also support strong mechanistic claims without directly intervening on a model. We conclude by discussing future directions for mechanistic interpretability, and how we can draw inspiration from other scientific disciplines to ask what should count as a good explanation.

Bio: Aaron Mueller is an Assistant Professor of Computer Science and, by courtesy, of Data Science at Boston University. His research centers on developing interpretability and evaluation methods inspired by causal and linguistic principles, and applying these to precisely control and improve the generalization of language technologies. He completed his Ph.D. at Johns Hopkins University.

Keynote Talk

Does mechanistic interpretability need interventions?

Tiago Pimentel
ETH Zurich
2026-04-07 10:50:00 – Room: **TBD**

Abstract: Mechanistic interpretability often treats interventions as the gold standard of evidence, relying on circuit ablations and representation steering to support claims about how models actually work. But are interventions really sufficient, or even necessary, for making mechanistic claims? In this debate-style talk, we trace the history that led the field to embrace interventions, and argue that the answer to both questions is a contentious no.We start by showing that interventions alone are not sufficient to explain model behavior by highlighting cases where causal methods can produce misleading or outright spurious explanations. Then, we debate whether interventions are necessary at all, exploring how alternative notions of causality and carefully designed behavioral evidence may also support strong mechanistic claims without directly intervening on a model. We conclude by discussing future directions for mechanistic interpretability, and how we can draw inspiration from other scientific disciplines to ask what should count as a good explanation.

Bio: Tiago Pimentel is a Postdoctoral Researcher at ETH Zürich, working in machine learning interpretability and psycholinguistics. His long-term goal is to understand how humans and machines process language. To this end, his research adopts an interdisciplinary approach, leveraging information theory and causality to study the mechanisms behind model behaviour and human cognition.

Table of Contents

From Natural Language to Certified Geometry Proofs: A Survey of LLM-Augmented Verification and Neuro-Symbolic Theorem Proving
Ioannis Tzachristas and Georgios Tzachristas 1

Open Problems Solved by LLMs? A Survey of Verifiable Mathematical Discovery
Ioannis Tzachristas, Georgios Tzachristas and Aifen Sui 10

Beyond Hallucination: Reframing LLM Quality Assessment as Task-Output Alignment
Andrew Hoblitzell 22

Challenging the Myth: A Research Arc on LLMs as Human Simulacra
Simon Münker, Achim Rettinger and Damian Trilling 31

Towards Trustworthy AI-Mediated Communication Across Languages and Cultures
Dayeon Ki 45

Challenging Quadratic Attention - A Holistic View On the Rise of Alternative Language Model Architectures
Alexander M. Fichtl, Jeremias Bohn, Josefin Kelber, Edoardo Mosca and Georg Groh 60

Why Low-Resource NLP Needs More Than Cross-Lingual Transfer: Lessons Learned from Luxembourgish
Fred Philippy, Siwen Guo, Jacques Klein and Tegawendé F. Bissyandé 82

Building Arabic NLP from the Ground Up: Twenty Years of Lessons, Failures, and Open Problems
Wajdi Zaghouani 94

Speaking of Language: Reflections on Metalanguage Research in NLP
Nathan Schneider and Antonios Anastasopoulos 107

Harnessing the Latent Space: From Steering Vectors to Model Calibrators for Control and Trust
Nishant Subramani 119

Language Models as Measurement Apparatus for Culture
Kent K. Chang 131

Memorisation Meets Compositionality in Natural Language Processing
Verna Dankers 144

Program

Saturday, July 4, 2026

09:00 - 09:10 *Opening Remarks*

09:10 - 10:30 *Paper Presentations*

From Natural Language to Certified Geometry Proofs: A Survey of LLM-Augmented Verification and Neuro-Symbolic Theorem Proving
Ioannis Tzachristas and Georgios Tzachristas

Open Problems Solved by LLMs? A Survey of Verifiable Mathematical Discovery
Ioannis Tzachristas, Georgios Tzachristas and Aifen Sui

Beyond Hallucination: Reframing LLM Quality Assessment as Task-Output Alignment
Andrew Hoblitzell

Challenging the Myth: A Research Arc on LLMs as Human Simulacra
Simon Münker, Achim Rettinger and Damian Trilling

Towards Trustworthy AI-Mediated Communication Across Languages and Cultures
Dayeon Ki

Challenging Quadratic Attention - A Holistic View On the Rise of Alternative Language Model Architectures
Alexander M. Fichtl, Jeremias Bohn, Josefin Kelber, Edoardo Mosca and Georg Groh

10:30 - 11:00 *Break*

11:00 - 12:30 *Paper Presentations*

Why Low-Resource NLP Needs More Than Cross-Lingual Transfer: Lessons Learned from Luxembourgish
Fred Philippy, Siwen Guo, Jacques Klein and Tegawendé F. Bissyandé

Building Arabic NLP from the Ground Up: Twenty Years of Lessons, Failures, and Open Problems
Wajdi Zaghouani

Saturday, July 4, 2026 (continued)

Speaking of Language: Reflections on Metalanguage Research in NLP
Nathan Schneider and Antonios Anastasopoulos

Harnessing the Latent Space: From Steering Vectors to Model Calibrators for Control and Trust
Nishant Subramani

Language Models as Measurement Apparatus for Culture
Kent K. Chang

Memorisation Meets Compositionality in Natural Language Processing
Verna Dankers

12:30 - 12:45 *Closing Remarks*