

BigPicture 2023

The Big Picture Workshop

Proceedings of the Workshop

December 7, 2023

The BigPicture organizers gratefully acknowledge the support from the following sponsors.

Gold



Silver



©2023 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 979-8-89176-051-6

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 EMNLP 2023, in Singapore, on December 7, 2023.

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 12 submissions, of which we accepted 10 for presentation at the workshop. Those 10 accepted papers are contained in this volume. We also accepted for presentation two additional papers to be included in Findings of EMNLP 2023.

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 one Best Paper session, one in-person poster session, and one virtual poster session.

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. We also thank Amazon, Google, and Hugging Face for generous sponsorship. Finally, we thank the EMNLP 2023 organizers for their hard work and support.

The Big Picture Workshop Organizers,

Yanai Elazar, Allyson Ettinger, Nora Kassner, Sebastian Ruder, Noah Smith

Organizing Committee

Program Chairs

Yanai Elazar, Allen Institute for AI & University of Washington, USA

Allyson Ettinger, Allen Institute for AI

Nora Kassner, Google DeepMind, United Kingdom

Sebastian Ruder, Google DeepMind, Germany

Noah A. Smith, Allen Institute for AI & University of Washington, USA

Program Committee

Reviewers

David Ifeoluwa Adelani, Jacob Andreas, Maria Antoniak, Akari Asai

Jonathan Berant, Ben Bogin

Eunsol Choi

Philipp Dufter, Greg Durrett

Matthias Gallé

Valentin Hofmann

Alankar Jain

Tal Linzen

Kyle Mahowald, Gaurav Mishra

Hao Peng, Jonas Pfeiffer, Barbara Plank

Machel Reid, Paul Roit

Ori Shapira, Vered Shwartz, Anders Søgaard

Ivan Vulić, Yogarshi Vyas

Adina Williams

Keynote Talk: The Vision Thing: Finding and Pursuing your Research Passion

Raymond J. Mooney

UT Austin

2023-12-07 09:15:00 – Room: Virgo 1 & 2

Abstract: A key contribution to being a successful researcher in natural language processing, as in any area, is having a clear overarching vision of what your body of research is trying to accomplish. Using my own 40-year career as an example, I will attempt to provide general advice on formulating and pursuing a coherent research vision. In particular, I will focus on formulating a unique, personal objective that exploits your specific talents, knowledge, and passions, and that is distinct from the current popular trends in the field. I will also focus on formulating a vision that bridges existing fields of study to produce an overarching agenda that unifies previously disparate ideas.

Bio: Raymond J. Mooney is a Professor in the Department of Computer Science at the University of Texas at Austin. He received his Ph.D. in 1988 from the University of Illinois at Urbana/Champaign. He is an author of over 200 published research papers, primarily in the areas of machine learning and natural language processing. He was the President of the International Machine Learning Society from 2008-2011, program co-chair for AAAI 2006, general chair for HLT-EMNLP 2005, and co-chair for ICML 1990. He is a Fellow of AAAI, ACM, and ACL and the recipient of the Classic Paper award from AAAI-19 and best paper awards from AAAI-96, KDD-04, ICML-05 and ACL-07.

Keynote Talk: Is Attention = Explanation and the Role of Interpretability in NLP

Sarah Wiegreffe

AI2 & UW

2023-12-07 11:00:00 – Room: Virgo 1 & 2

Abstract: Attention mechanisms have become a core component of neural models in Natural Language Processing over the past decade. These mechanisms not only deliver substantial performance improvements but also claim to offer insights into the models' inner workings. In this talk, we will highlight a series of contributions we have made that provided a critical perspective on the role of attention as a faithful explanation for model predictions, and sparked a larger conversation on the overarching goals of interpretability methods in NLP. We'll contrast our methodological approaches and findings to highlight that there is no one-size-fits-all answer to the question "Is attention explanation?". Finally, we'll explore the role of attention as an explanation mechanism in today's NLP landscape.

Bio: Sarah Wiegreffe is a postdoctoral researcher at the Allen Institute for AI (AI2), working on the Aristo project. She also holds a courtesy appointment in the Allen School of Computer Science and Engineering at the University of Washington. Her research focuses on understanding how language models make predictions in an effort to make them more transparent to human users. She received her PhD from Georgia Tech in 2022 advised by Professor Mark Riedl, during which time she interned at Google and AI2 and won the AI2 outstanding intern award. She frequently serves on conference program committees, receiving outstanding area chair award at ACL 2023.

Keynote Talk: Is Attention = Explanation and the Role of Interpretability in NLP

Sarthak Jain

AWS AI Labs

2023-12-07 11:00:00 – Room: Virgo 1 & 2

Abstract: Attention mechanisms have become a core component of neural models in Natural Language Processing over the past decade. These mechanisms not only deliver substantial performance improvements but also claim to offer insights into the models' inner workings. In this talk, we will highlight a series of contributions we have made that provided a critical perspective on the role of attention as a faithful explanation for model predictions, and sparked a larger conversation on the overarching goals of interpretability methods in NLP. We'll contrast our methodological approaches and findings to highlight that there is no one-size-fits-all answer to the question "Is attention explanation?". Finally, we'll explore the role of attention as an explanation mechanism in today's NLP landscape.

Bio: Sarthak Jain is an Applied Scientist working on generative AI models at AWS. He received his PhD in 2022 from Northeastern University, where he was advised by Byron Wallace. Before this, he completed his BTech in Computer Engineering from Delhi Technological University. His current research interests include the interpretability and analysis of deep learning models.

Keynote Talk: On the Outcomes of Scientific Disagreements of Machine Morality

Liwei Jiang

University of Washington

2023-12-07 13:30:00 – Room: **Virgo 1 & 2**

Abstract: Disagreements and conflict are vital for driving scholarly progress, social and scientific alike. In research, we often identify gaps in others’ and our own work, to present new ideas that remedy them. Disagreements are often small in nature: We disagree on methods rather than the research programme itself. In this talk, we discuss a disagreement of a different nature: namely one in which the substance of the disagreement is the existence of the task itself. We reflect on the experience of the conflict, how it was resolved, and what outcomes it has had.

In particular, Liwei will share her current interdisciplinary research journey on AI + humanity sparked by the Delphi experience. She will introduce Value Kaleidoscope—a novel computational system aiming to model potentially conflicting, pluralistic human values interwoven in human decision-making. Finally, she will talk about an exciting co-evolution opportunity unfolding between frontier AI technology and humanity fields.

Zeerak will go over ongoing work that considers the foundations and limits of machine learning and NLP with regard to ethically appropriate work. Specifically, they will discuss the use of the distributional hypothesis, and what particular visions of our societies it offers, and how machine learning seeks to construct our future in the vision of the past.

Bio: Liwei Jiang is a Ph.D. student in the Paul G. Allen School of Computer Science and Engineering at the University of Washington, specializing in Artificial Intelligence (AI) and Natural Language Processing (NLP). She is intrigued to tackle real-world needs with AI and understand the charms, mysteries, and peculiarities of humans. Thus, Her current research focuses on the co-evolution of AI and humanity: how to build better AI by taking inspiration from humans and how to gain valuable insights into humans by advancing AI. She has published at many NLP and AI venues (e.g., ACL, EMNLP, NAACL, NeurIPS, AAAI). Her work has been featured in many media outlets, including the New York Times, Wired, the Guardian, the Verge, IEEE Spectrum, and Nature Outlook. She works as a student researcher at Allen Institute for Artificial Intelligence (AI2).

Keynote Talk: On the Outcomes of Scientific Disagreements of Machine Morality

Zeerak Talat

Mohamed Bin Zayed University of Artificial Intelligence

2023-12-07 13:30:00 – Room: Virgo 1 & 2

Abstract: Disagreements and conflict are vital for driving scholarly progress, social and scientific alike. In research, we often identify gaps in others’ and our own work, to present new ideas that remedy them. Disagreements are often small in nature: We disagree on methods rather than the research programme itself. In this talk, we discuss a disagreement of a different nature: namely one in which the substance of the disagreement is the existence of the task itself. We reflect on the experience of the conflict, how it was resolved, and what outcomes it has had.

In particular, Liwei will share her current interdisciplinary research journey on AI + humanity sparked by the Delphi experience. She will introduce Value Kaleidoscope—a novel computational system aiming to model potentially conflicting, pluralistic human values interwoven in human decision-making. Finally, she will talk about an exciting co-evolution opportunity unfolding between frontier AI technology and humanity fields.

Zeerak will go over ongoing work that considers the foundations and limits of machine learning and NLP with regard to ethically appropriate work. Specifically, they will discuss the use of the distributional hypothesis, and what particular visions of our societies it offers, and how machine learning seeks to construct our future in the vision of the past.

Bio: Zeerak Talat (formerly known as Zeerak Waseem) is a Research Fellow at Mohamed Bin Zayed University of Artificial Intelligence. Zeerak holds a Ph.D. in Computer Science from the University of Sheffield, with a focus on natural language processing. Zeerak’s work examines the assumptions that underpin NLP and machine learning (ML) technologies. Drawing on research from anthropology, discard studies, science and technology studies, and media studies, their work seeks to consider NLP and ML technologies through the lens of content moderation technologies to understand how they can cause harm to individuals and societies.

Keynote Talk: The Role of Demonstrations: What In-Context Learning actually does

Sewon Min

University of Washington

2023-12-07 16:00:00 – Room: Virgo 1 & 2

Abstract: In-Context Learning (ICL) enables a language model (LM) to learn a new correlation between inputs and outputs during inference, without explicit gradient updates. In this talk, we show a series of work centered around the research question: whether or not the correctness of demonstrations is needed for good performance of ICL. Through a series of experiments and analyses, we delve into the nuances of this relationship across various experimental setups, models (plain LMs or instruction-tuned ones), and tasks (classification or generation). Our findings contribute to a broader understanding of how LMs engage in in-context learning, shedding light on what new correlations they can or cannot learn, and leading to a new line of research in discovering unexpected behaviors of LMs.

Bio: Sewon Min is a final year Ph.D. candidate at the University of Washington, advised by Luke Zettlemoyer and Hannaneh Hajishirzi. Her research is in language modeling, focusing on new dimensions in modeling, scaling, and efficiency, and their extensions for information-seeking, legality, and privacy. She co-instructed and co-organized multiple tutorials and workshops at ACL, EMNLP, NAACL and NeurIPS. She is a recipient of the J.P. Morgan Fellowship, and was at Meta AI, Google Research, and Salesforce Research.

Keynote Talk: The Role of Demonstrations: What In-Context Learning actually does

Junyeob Kim

Seoul National University

2023-12-07 16:00:00 – Room: **Virgo 1 & 2**

Abstract: In-Context Learning (ICL) enables a language model (LM) to learn a new correlation between inputs and outputs during inference, without explicit gradient updates. In this talk, we show a series of work centered around the research question: whether or not the correctness of demonstrations is needed for good performance of ICL. Through a series of experiments and analyses, we delve into the nuances of this relationship across various experimental setups, models (plain LMs or instruction-tuned ones), and tasks (classification or generation). Our findings contribute to a broader understanding of how LMs engage in in-context learning, shedding light on what new correlations they can or cannot learn, and leading to a new line of research in discovering unexpected behaviors of LMs.

Bio: Sewon Min is a final year Ph.D. candidate at the University of Washington, advised by Luke Zettlemoyer and Hannaneh Hajishirzi. Her research is in language modeling, focusing on new dimensions in modeling, scaling, and efficiency, and their extensions for information-seeking, legality, and privacy. She co-instructed and co-organized multiple tutorials and workshops at ACL, EMNLP, NAACL and NeurIPS. She is a recipient of the J.P. Morgan Fellowship, and was at Meta AI, Google Research, and Salesforce Research.

Keynote Talk: The Role of Demonstrations: What In-Context Learning actually does

Kang Min Yoo

NAVER Cloud, NAVER AI Lab

2023-12-07 16:00:00 – Room: Virgo 1 & 2

Abstract: In-Context Learning (ICL) enables a language model (LM) to learn a new correlation between inputs and outputs during inference, without explicit gradient updates. In this talk, we show a series of work centered around the research question: whether or not the correctness of demonstrations is needed for good performance of ICL. Through a series of experiments and analyses, we delve into the nuances of this relationship across various experimental setups, models (plain LMs or instruction-tuned ones), and tasks (classification or generation). Our findings contribute to a broader understanding of how LMs engage in in-context learning, shedding light on what new correlations they can or cannot learn, and leading to a new line of research in discovering unexpected behaviors of LMs.

Bio: Kang Min Yoo is actively engaged in the fields of artificial intelligence and computational linguistics. He currently holds key roles as a Research and Applied Scientist at NAVER Cloud and as a Visiting Professor at Seoul National University's AI Institute. With an Integrated M.S. and Ph.D. in Computer Science from Seoul National University, his primary areas of expertise include large language models and natural language processing. At NAVER Cloud, he has spearheaded projects focused on developing Korean-centric LLM-based chat agents and the HyperT5 Seq2Seq HyperCLOVA. Additionally, Kang Min Yoo contributes to the academic community through his roles as an area chair and program committee member.

Table of Contents

<i>Where are We in Event-centric Emotion Analysis? Bridging Emotion Role Labeling and Appraisal-based Approaches</i>	
Roman Klinger	1
<i>Working Towards Digital Documentation of Uralic Languages With Open-Source Tools and Modern NLP Methods</i>	
Mika Härmäläinen, Jack Rueter, Khalid Alnajjar and Niko Tapio Partanen	18
<i>Computational Narrative Understanding: A Big Picture Analysis</i>	
Andrew Piper	28
<i>The Case for Scalable, Data-Driven Theory: A Paradigm for Scientific Progress in NLP</i>	
Julian Michael	40
<i>Thesis Distillation: Investigating The Impact of Bias in NLP Models on Hate Speech Detection</i>	
Fatma Elsafoury	53
<i>Large Language Models as SocioTechnical Systems</i>	
Kaustubh Dhole	66
<i>Towards Low-resource Language Generation with Limited Supervision</i>	
Kaushal Kumar Maurya and Maunendra Sankar Desarkar	80
<i>Transformers as Graph-to-Graph Models</i>	
James Henderson, Alireza Mohammadshahi, Andrei Catalin Coman and Lesly Miculicich	93
<i>It's MBR All the Way Down: Modern Generation Techniques Through the Lens of Minimum Bayes Risk</i>	
Amanda Bertsch, Alex Xie, Graham Neubig and Matthew R. Gormley	108
<i>Analyzing Pre-trained and Fine-tuned Language Models</i>	
Marius Mosbach	123

Program

Thursday, December 7, 2023

10:05 - 09:15 *Break*

09:00 - 09:15 *Opening Remarks*

09:15 - 10:05 *The Vision Thing: Finding and Pursuing your Research Passion*

10:05 - 10:30 *Talk*

The Case for Scalable, Data-Driven Theory: A Paradigm for Scientific Progress in NLP

Julian Michael

11:00 - 10:30 *Break*

11:00 - 12:00 *Is Attention = Explanation and the Role of Interpretability in NLP*

12:00 - 13:30 *Lunch Break*

11:00 - 12:00 *From Machine Morality to Pluralistic Human Values: When AI Interfaces with Humanity*

14:30 - 15:30 *Poster Session*

Where are We in Event-centric Emotion Analysis? Bridging Emotion Role Labeling and Appraisal-based Approaches

Roman Klinger

Working Towards Digital Documentation of Uralic Languages With Open-Source Tools and Modern NLP Methods

Mika Hämmäläinen, Jack Rueter, Khalid Alnajjar and Niko Tapio Partanen

Computational Narrative Understanding: A Big Picture Analysis

Andrew Piper

Thesis Distillation: Investigating The Impact of Bias in NLP Models on Hate Speech Detection

Fatma Elsafoury

Thursday, December 7, 2023 (continued)

Large Language Models as SocioTechnical Systems

Kaustubh Dhole

Towards Low-resource Language Generation with Limited Supervision

Kaushal Kumar Maurya and Maunendra Sankar Desarkar

Transformers as Graph-to-Graph Models

James Henderson, Alireza Mohammadshahi, Andrei Catalin Coman and Lesly Miculicich

It's MBR All the Way Down: Modern Generation Techniques Through the Lens of Minimum Bayes Risk

Amanda Bertsch, Alex Xie, Graham Neubig and Matthew R. Gormley

Analyzing Pre-trained and Fine-tuned Language Models

Marius Mosbach

15:30 - 16:00

Break

16:00 - 17:00

The Role of Demonstrations: What In-Context Learning actually does

17:00 - 17:15

Closing Remarks