WASP 2026

# The 4th Workshop on Artificial Intelligence for Scientific Publications

# Proceedings of the 4th Workshop on Artificial Intelligence for Scientific Publications

November 10, 2026

©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 979-8-89176-456-9

# Preface

We are excited to welcome you to WASP at AACL-IJCNLP 2026, the Fourth Workshop for Artificial intelligence for Scientific Publications. This year the workshop is being held online on November 10, 2026.

Building on the success of the previous Workshops for Artificial Intelligence for Scientific Publications at IJCNLP-AACL 2022, 2023, and 2025, the Fourth Workshop on Artificial intelligence for Scientific Publications (WASP) at AACL-IJCNLP 2026 aims to establish itself as a pivotal platform for promoting discussions and research in the field of Natural Language Processing (NLP) and Artificial Intelligence (AI). This gathering brings together esteemed experts and renowned organizations with students and early-career researchers who are interested and invested in efforts to extract and mine the world's scientific knowledge from research papers. Their collaboration focuses on developing advanced algorithms, models, and tools that lay the foundation for future machine comprehension of scientific literature. The fourth iteration of WASP specifically concentrates on various topics related to Artificial Intelligence research for/with scientific publications.

We especially welcome participation from academic and research institutions, government and industry labs, publishers, and information service providers. Projects and organizations using NLP/ML techniques in their text mining and enrichment efforts are also welcome to participate. We strongly encourage the participation of students, researchers, and science practitioners from diverse backgrounds.

WASP 2026 includes one shared task where we invite teams (individuals and groups) to tackle the problem of recreating the partial citation graph from scientific figures, with AstroCLIMB: Astronomy Citation Linking from Illustrations - a Multimodal Benchmark.

WASP 2026 received 43 submissions of which 25 were accepted (10 long papers, 5 short papers, and 10 shared task system papers). We are thankful to our program committee members for helping us curate a strong WASP 2026 program.

On behalf of the program co-chairs:
Tirthankar Ghosal, Alberto Accomazzi, Kelly Lockhart, Felix Grezes, and Atilla Kaan Alkan.

# Organizing Committee

**Program Chairs**

Alberto Accomazzi, Center for Astrophysics | Harvard & Smithsonian
Tirthankar Ghosal, Oak Ridge National Laboratory
Felix Grezes, Center for Astrophysics | Harvard & Smithsonian
Kelly Lockhart, Center for Astrophysics | Harvard & Smithsonian
Atilla Kaan Alkan, Center for Astrophysics | Harvard & Smithsonian

# Program Committee

**Reviewers**

Antoine Gauquier, Ecole Normale Supérieure, Ecole Normale Supérieure de Paris
Alberto Accomazzi, Center for Astrophysics | Harvard & Smithsonian
Arne Binder, German Research Center for AI
Daniel E. Acuna, University of Colorado at Boulder
Joyita Chakraborty, National Institute of Technology Durgapur
Kelly Lockhart, Center for Astrophysics | Harvard & Smithsonian
Elyon Eyimife, Stevens Institute of Technology
Ibrahim Al Azher, Northern Illinois University
Mengyang Liu, Shanghai Jiao Tong University
Harsh Sakhrani, University of Southern California
Shuntaro Yada, Institute of Library, Information and Media Science, Tsukuba University, Tokyo Institute of Technology
Bhargab Choudhury, Department of Information Technology, Gauhati University
Sharmila Upadhyaya, Gesis leibniz Institute
Tosho Hirasawa, LayerX
Sridevi Wagle, Pacific Northwest National Laboratory
Sujit Pal, Elsevier
Atilla Kaan Alkan, Center for Astrophysics | Harvard & Smithsonian
Rahul Nagurtha, Google
Tirthankar Ghosal, Oak Ridge National Laboratory
Sergey Feldman, Allen Institute for Artificial Intelligence
Neil R. Smalheiser, University of Illinois at Chicago
Pawin Taechoyotin, University of Colorado at Boulder
Yagmur Ozturk, Université Grenoble Alpes
Wojtek Sylwestrzak, University of Warsaw
Taro Watanabe, Nara Institute of Science and Technology, Japan
Tim Schopf, National Institute of Informatics
Felix Grezes, Center for Astrophysics | Harvard & Smithsonian
Biswadip Mandal, Amazon
Ethan Seefried, Oak Ridge National Laboratory
Thomas S Allen, Center for Astrophysics | Harvard & Smithsonian
Sharmin Afrose, Oak Ridge National Laboratory
Ran Elgedawy, University of Tennessee, Knoxville
Swati Rajwal, Emory University
Sandeep Kumar, Mohamed bin Zayed University of Artificial Intelligence
Sanjay Das, Oak Ridge National Laboratory

## Invited Talk

# Salience and Inquiry in Scientific Literature

**Jessy Li**
The University of Texas at Austin



**Tuesday, November 10, 2026 –**

**Abstract:** In this talk, we view the processing of scientific literature as an (implicitly) interactive discourse process: the reader is constantly abstracting out the most salient information, while staying inquisitive about the content. I will discuss the role of questions, first in deriving information salience from research documents. Using summarization as a behavioral probe, we reveal that across different lengths of summaries, the questions that models tend to answer remain intriguingly consistent. Next, we show how questions play important roles in connecting different pieces of a scientific document via generative discourse models, and extend existing frameworks to multimodal settings, modeling the connections across passages and figures. Finally, I will end the talk with brief discussions on AI review pilots at recent large conferences.

**Bio:** Jessy Li (https://jessyli.com/) is an associate professor in the Linguistics Department and (by courtesy) the Computer Science Department at the University of Texas at Austin. She received her Ph.D. from the Department of Computer and Information Science at the University of Pennsylvania in 2017. Her research interests are in computational linguistics and natural language processing, specifically discourse and pragmatics, generation and evaluation, and language & code. She is a recipient of an NSF CAREER Award, ACL and EMNLP Outstanding Paper Awards, an ACM SIGSOFT Distinguished Paper Award, among others. Jessy is an AI research lead at the NSF-Simons CosmicAI Institute. She is also the Secretary of NAACL, and a Program Chair of the EMNLP 2026 conference.

## Invited Talk

# From General-Purpose LLMs to Grounded Tools for Scientific Discovery

**Pavlos Protopapas**
Harvard University



**Tuesday, November 10, 2026 –**

**Abstract:** Large language models offer new ways to search, interpret, and synthesize scientific literature. Yet their responses can reflect biases in their training data, uneven coverage of research, and claims that are difficult to verify. How can we make these models useful while retaining control over the evidence and reasoning behind their outputs? Drawing on two recent research projects, I will discuss building applications on top of NASA's Astrophysics Data System (ADS) that connect language models to scholarly sources and guide their use through explicit workflows. The talk will explore how retrieval, tool use, and verification can help make outputs more transparent and assessable, while examining the limitations that remain. The central question is how to design AI tools that support scientific judgment and allow researchers to inspect and challenge their results.

**Bio:** Pavlos Protopapas is the Scientific Program Director for the master's programs in data science and computational science engineering at the Harvard John A. Paulson School of Engineering and Applied Sciences. As an educator, he currently teaches the CS109 course series for basic and advanced data science and a course in Machine Learning Operations at Harvard University. He previously taught the capstone courses for the master's programs, and also a course in deep reinforcement learning. His courses are also available through the Harvard Extension School and on HarvardX. In addition he has participated in data science schools around the world in Chile, Colombia, India, and Rwanda. Pavlos has a PhD in theoretical physics from the University of Pennsylvania and his research interests are at the intersection of astronomy, machine learning, and statistics.

# Table of Contents

*Overview of the Fourth Workshop on Artificial Intelligence for Scientific Publications (WASP 2026)*
Atilla Kaan Alkan, Felix Grezes, Kelly Lockhart, Alberto Accomazzi and Tirthankar Ghosal . . . 1

*AstroCLIMB: Astronomy Citation Linking from Illustrations - a Multimodal Benchmark and Overview of the WASP 2026 Shared Task*
Felix Grezes, Atilla Kaan Alkan, Kelly Lockhart, Alberto Accomazzi, Ethan Seefried and Tirthankar Ghosal . . . 10

*From Astronomical Circulars to Auditable Follow-up Records: Schema-Constrained LLM Extraction with Human Verification*
Fabian Schüssler, Sofia Bisero, M. Cellier, J. Cinosi, A. Ciric, Bernardo Cornejo, Ilja Jaroschewski, Atilla Kaan Alkan, R. Weizmann Kiendrebeogo and A. Saint-Paul . . . 17

*ACL-Verbatim: hallucination-free question answering for research*
Gábor Recski, Szilveszter Tóth, Nadia Verdha, István Boros and Ádám Kovács . . . 23

*Auditable Multi-View Hosted Multimodal Judging for AstroCLIMB*
Kewen Zhu and Junyuan Tan . . . 34

*Can Generative AI Automate Data Extraction for Meta-Analysis? A Case Study on Intercropping Research*
Zehao Lu, Xingguo Xiong, Wopke van der Werf, Thijs L. van der Plas and Ioannis N. Athanasiadis 38

*HCMUS_Beyonders at AstroCLIMB: Same Paper or Citing Paper? Linking Astronomy Figures with OCR, Citation Graphs, and LLMs*
My-Linh Bang, Phu-Quy Nguyen-Lam, Duy-Minh Dao-Sy, Trung-Kiet Huynh, Phu-Hoa Pham and Chi-Nguyen Tran . . . 52

*Local Citation Recommendation Without Graph Signals*
Benjamin Basseri, Pavlos Protopapas, Kemeng Zhang, Alberto Accomazzi, Stephanie Jarmak and Anna Kelbert . . . 61

*CARDS: Citation-Aware Research Diffusion Summarization*
Autumn Toney-Wails, Katherine N. Quinn and Steph Batalis . . . 77

*Bridging the Data Bottleneck: A Multi-Modal Generative Extraction Pipeline for Legacy Scientific Literature*
Gaurav Kumar, Anasse Bari and David J. Nagel . . . 91

*GRAIL: Compact Multimodal Similarities for Astronomy Citation Linking*
Egor Fedosov . . . 99

*Parameter-Efficient Adaptation of Qwen3-VL for Astronomy Citation Linking at AstroCLIMB 2026*
Syed Mohaiminul Hoque . . . 104

*Modality-Conditioned Ensembles and Zero-Shot Cross-Encoder Reranking for Relation Classification Between Astronomy-Paper Objects*
M. Tahsan Zahid, Most. Atkia Farzana and Tonmoy Sarker Sourav . . . 112

*Fusing Sentence-Level and Document-Level Representations for Keyphrase Extraction*
Marko Pranjić, Kristijan Armeni and Senja Pollak . . . 122

*Beyond Fluency: A Framework for Analyzing Epistemic and Argumentative Shifts in AI-Assisted Scientific Text Reformulation*
Laura Centonze . . . . . 133

*TRACS: Telescope Reference & Astronomy Categorization Dataset and Shared Task*
Felix Grezes, Kelly Lockhart, Alberto Accomazzi, Jennifer Lynn Bartlett, Ethan Seefried, Anjali Pandiri and Tirthankar Ghosal . . . . . 142

*Public Scientific-Figure Adaptation for AstroCLIMB with a 4B Vision–Language Model*
Irving Ernesto Quezada Ramírez . . . . . 154

*Accelerating Literature Data Extraction using AI at the IPAC Archives*
Nicholas Susemiehl, Nima Chartab, Meca Lynn and Ricky Nilsson . . . . . 166

*Extracting Published Discussions of 13,616 Galaxies with Language Models*
Nolan Koblischke and Jo Bovy . . . . . 174

*Star Light, Star Bright: Can Multimodal Models Parse an Astronomy Visual Question Answering Dataset Right?*
J. P. Naiman . . . . . 180

*PSGS@2026: From Figures to Citation Graphs: QLoRA Based Multimodal Linking for AstroCLIMB*
Prajwal Ghimire, Shiv Ram Saud, Aashish Mahato, Sundeep Dawadi, Ramesh Saud, Sunil Regmi and Manoj Shakya . . . . . 192

*Towards Verifiable Answers: Reinforcement Learning with a Verifier Reduces Hallucinated References in Literature-Based Question Answering*
Rushank Goyal . . . . . 199

*Beyond Visual Similarity: Parameter-Efficient Vision–Language Models for Scientific Citation Linking on AstroCLIMB*
Ashmani Kumar, Shubham Sharma, K S Ramalakshmi and Ankush Chopra . . . . . 210

*Multimodal Feature Integration for Pairwise Relationship Classification in Astronomy Literatures*
Raj Mahesh Mane, Ananya Kalia, Aashna Shetty, Rutal Mahajan and Preksha Pareek . . . . . 216

*Scientific Table Generation: Structured Information Extraction from Scientific Literature*
Necva Bölücü, Huichen Yang, Roelien C. Timmer, Stephen Wan, Brian Jin, Changhyun Lee, Andreas Duenser and Jessica Irons . . . . . 224

*Recovering Method Parameters from Protocol Prose: Task Definition, Audited Extraction, and the Comparability It Licenses*
Fabio Rovai . . . . . 247

*Modality-Specific and Relation-Aware Modeling for Scientific Figure Linking*
Bodhisatta Maiti and Debshree Chowdhury . . . . . 262