

40th AAAI Conference on Artificial Intelligence (AAAI-26), 38th Conference on Innovative Applications of Artificial Intelligence (IAAI-26), and 16th Symposium on Educational Advances in Artificial Intelligence (EAAI-26)

Volume 46: AAAI Special Track

- AI for Social Impact II
- Senior Member Presentation

Singapore
20-27 January 2026

ISBN: 979-8-3313-3420-8

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2026) by Association for the Advancement of Artificial Intelligence
All rights reserved.

Printed with permission by Curran Associates, Inc. (2026)

For permission requests, please contact Association for the Advancement of Artificial Intelligence
at the address below.

Association for the Advancement of Artificial Intelligence
2275 East Bayshore Road
Suite 160
Palo Alto, California 94303
USA

Phone: 1-650-328-3123
Fax: 1-650-321-4457

<https://aaai.org/Press/press.php>

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

TABLE OF CONTENTS

AAAI SPECIAL TRACK ON AI FOR SOCIAL IMPACT II

DeepGB-TB: A Risk-Balanced Cross-Attention Gradient-Boosted Convolutional Network for Rapid, Interpretable Tuberculosis Screening	38989
<i>Zhixiang Lu, Yulong Li, Feilong Tang, Zhengyong Jiang, Chong Li, Mian Zhou, Tenglong Li, Jionglong Su</i>	
GREAT: Generalizable Representation Enhancement Via Auxiliary Transformations for Zero-Shot Environmental Prediction	38998
<i>Shiyuan Luo, Chonghao Qiu, Runlong Yu, Yiqun Xie, Xiaowei Jia</i>	
A Causal Framework to Measure and Mitigate Non-Binary Treatment Discrimination	39007
<i>Ayan Majumdar, Deborah D. Kanubala, Kavya Gupta, Isabel Valera</i>	
A Benchmark Dataset for Spatially Aligned Road Damage Assessment in Small Uncrewed Aerial Systems Disaster Imagery	39016
<i>Thomas Manzini, Priyankari Perali, Raisa Karnik, Robin R. Murphy</i>	
Optimizing Urban Service Allocation with Time-Constrained Restless Bandits.....	39025
<i>Yi Mao, Andrew Perrault</i>	
Visual-Friendly Concept Protection Via Selective Adversarial Perturbations	39033
<i>Xiaoyue Mi, Fan Tang, You Wu, Juan Cao, Peng Li, Yang Liu</i>	
Improving Low-Resource Translation with Dictionary-Guided Fine-Tuning and RL: A Spanish-To-Wayuunaiki Study	39042
<i>Manuel Mosquera, Melissa Verónica Robles, Johan R. Portela, Rubén Manrique</i>	
Masked Mineral Modeling: Continent-Scale Mineral Prospecting Via Geospatial Infilling.....	39051
<i>Sujay Nair, Evan Austen Coleman, Sherrie Wang, Elsa Olivetti</i>	
Rescind: Countering Image Misconduct in Biomedical Publications with Vision-Language and State-Space Modeling.....	39061
<i>Soumyaroop Nandi, Prem Natarajan</i>	
Should You Use LLMs to Simulate Opinions? Quality Checks for Early-Stage Deliberation.....	39070
<i>Terrence Neumann, Maria De-Arteaga, Sina Fazelpour</i>	
Forecasting Clinical Risk from Textual Time Series: Structuring Narratives for Temporal AI in Healthcare.....	39080
<i>Shahriar Noroozizadeh, Sayantan Kumar, Jeremy C. Weiss</i>	
Culture Affordance Atlas: Reconciling Object Diversity Through Functional Mapping	39089
<i>Joan Nwatu, Longju Bai, Oana Ignat, Rada Mihalcea</i>	
Scalable Air-Quality Sensor Placement Via Gradient-Based Mutual Information Maximization.....	39098
<i>Zeel B Patel, Vinayak Rana, Nipun Batra</i>	
Debiasing Machine Learning Predictions for Causal Inference Without Additional Ground Truth Data: “One Map, Many Trials” in Satellite-Driven Poverty Analysis	39106
<i>Markus B. Pettersson, Connor Thomas Jerzak, Adel Daoud</i>	

LLM Targeted Underperformance Disproportionately Impacts Vulnerable Users	39116
<i>Elinor Poole-Dayana, Deb Roy, Jad Kabbara</i>	
CCD-Bench: Probing Cultural Conflict in Large Language Model Decision-Making.....	39125
<i>Hasibur Rahman, Hanan Salam</i>	
X-MuTeST: A Multilingual Benchmark for Explainable Hate Speech Detection and a Novel LLM- Consulted Explanation Framework	39134
<i>Mohammad Zia Ur Rehman, Sai Kartheek Reddy Kasu, Shashivardhan Reddy Koppula, Sai Rithwik Reddy Chirra, Shwetank Shekhar Singh, Nagendra Kumar</i>	
Toward Simulating Networked Societies with Formal Institutions Using AI Agents.....	39143
<i>Michael Richards, Danny Cowser, Daniel Nielson, Jacob W. Crandall</i>	
Navigation and Interaction for Blind Users Via a Cognitive Architecture	39151
<i>Oscar J. Romero, Anthony Tomasic, Elizabeth J. Carter, John Zimmerman, Aaron Steinfeld</i>	
A Retrieval Augmented Spatio-Temporal Framework for Traffic Prediction.....	39163
<i>Weilin Ruan, Xilin Dang, Ziyu Zhou, Sisuo Lyu, Yuxuan Liang</i>	
The Confidence Trap: Gender Bias and Predictive Certainty in LLMs.....	39173
<i>Ahmed Sabir, Markus Kängsepp, Rajesh Sharma</i>	
Active Learning for Animal Re-Identification with Ambiguity-Aware Sampling.....	39182
<i>Depanshu Sani, Mehar Khurana, Saket Anand</i>	
The Illusion of Fairness: Auditing Fairness Interventions in Algorithmic Hiring with Audit Studies.....	39191
<i>Disa Sariola, Patrick Button, Aron Culotta, Nicholas Mattei</i>	
Fast 3D Surrogate Modeling for Data Center Thermal Management.....	39201
<i>Soumyendu Sarker, Antonio Guillen-Perez, Zachariah J Carmichael, Avisek Naug, Refik Mert Cam, Vineet Gundecha, Ashwin Ramesh Babu, Sahand Ghorbanpour, Ricardo Luna Gutierrez</i>	
AURA: Development and Validation of an Augmented Unplanned Removal Alert System Using Synthetic ICU Videos	39211
<i>Junhyuk Seo, Hyeyoon Moon, Kyu-Hwan Jung, Namkee Oh, Taerim Kim</i>	
Machine Learning for Sustainable Rice Production: Region-Scale Monitoring of Water-Saving Practices in Punjab, India	39219
<i>Ando Shah, Rajveer Singh, Akram Zaytar, Girmaw Abebe Tadesse, Caleb Robinson, Negar Tafti, Stephen A. Wood, Rahul Dodhia, Juan M Lavista Ferres</i>	
AXON: Action Characterization Through Cross-Modal Knowledge Distillation for Neurodiverse Individuals.....	39228
<i>Siddhant Bikram Shah, Kristina T. Johnson</i>	
PlantTraitNet: An Uncertainty-Aware Multimodal Framework for Global-Scale Plant Trait Inference from Citizen Science Data.....	39239
<i>Ayushi Sharma, Johanna Trost, Daniel Lusk, Johannes Dollinger, Julian Schrader, Christian Rossi, Javier Lopatin, Etienne Laliberté, Simon Haberstroh, Jana Eichel, Daniel Mederer, Jose Miguel Cerda-Paredes, Shyam S. Phartyal, Lisa-Maricia Schwarz, Anja Linstädter, Maria Conceição Caldeira, Teja Kattenborn</i>	
OIDA-QA: A Multimodal Benchmark for Analyzing the Opioid Industry Documents Archive.....	39249
<i>Xuan Shen, Brian Wingenroth, Zichao Wang, Jason Kuen, Wanrong Zhu, Ruiyi Zhang, Yiwei Wang, Lichun Ma, Anqi Liu, Hongfu Liu, Tong Sun, Kevin S. Hawkins, Kate Tasker, G. Caleb Alexander, Jiuxiang Gu</i>	

Resilience Inference for Supply Chains with Hypergraph Neural Network.....	39259
<i>Zetian Shen, Hongjun Wang, Jiyuan Chen, Xuan Song</i>	
Talk, Snap, Complain: Validation-Aware Multimodal Expert Framework for Fine-Grained Customer Grievances.....	39268
<i>Rishu Kumar Singh, Navneet Shreya, Sarmistha Das, Apoorva Singh, Sriparna Saha</i>	
SpiderGen: Towards Procedure Generation for Carbon Life Cycle Assessments with Generative AI.....	39277
<i>Anupama Sitaraman, Bharathan Balaji, Yuvraj Agarwal</i>	
EvalQAG: A Framework for Automatic Complex QA Generation and a Benchmark QA Dataset for Policy Documents.....	39285
<i>Kirtan Brijeshbhai Soni, Krish Rupapara, Arpit Rana, Ghanshyam Verma, Paul Buitelaar</i>	
Investigating Vaccine Buyer’s Remorse: Post-Vaccination Decision Regret in COVID-19 Social Media Using Politically Diverse Human Annotation	39294
<i>Miles Stanley, Soumyajit Datta, Ashutosh Kumar, Ashiqur R. Khudabukhsh</i>	
What Are They Filtering Out? an Experimental Benchmark of Filtering Strategies for Harm Reduction in Pretraining Datasets	39303
<i>Marco Antonio Stranisci, Christian Hardmeier</i>	
BSAN: Behavioral State Attention Network for Modeling Mosquito Host-Seeking Behavior.....	39314
<i>Haotian Sun, Jessie Zixin Li, John M. Marshall</i>	
ESG-Bench: Benchmarking Long-Context ESG Reports for Hallucination Mitigation	39322
<i>Siqi Sun, Ben Peng Wu, Mali Jin, Peizhen Bai, Hanpei Zhang, Xingyi Song</i>	
Achieving Fairness Without Harm Via Selective Demographic Experts.....	39331
<i>Xuwei Tan, Yuanlong Wang, Thai-Hoang Pham, Ping Zhang, Xueru Zhang</i>	
AI in the Wild: A Meta-Analytic Evaluation of Depression Detection from Social Media Data	39340
<i>Xianglu Tang, Joyee W. Jin, Emily Ma, Xingyu Li</i>	
ForeSWE: Forecasting Snow-Water Equivalent with an Uncertainty-Aware Attention Model	39348
<i>Krishu K Thapa, Supriya Savalkar, Bhupinderjeet Singh, Trong Nghia Hoang, Kirti Rajagopalan, Ananth Kalyanaraman</i>	
BATIS: Bayesian Approaches for Targeted Improvement of Species Distribution Models	39358
<i>Catherine Villeneuve, Benjamin Akera, Mélisande Teng, David Rolnick</i>	
RipAlert: A Future-Frame-Aware Framework for Rip Current Forecasting and Early Alerting.....	39368
<i>Meng Wan, Qi Su, Zhixin Xia, Kanglin Chen, Jue Wang, Tiantian Liu, Rongqiang Cao, Hui Cui, Peng Shi, Yangang Wang, Liqiang Feng, Zhenbing Zhao</i>	
CARE-Bench: A Benchmark of Diverse Client Simulations Guided by Expert Principles for Evaluating LLMs in Psychological Counseling	39378
<i>Bichen Wang, Yixin Sun, Junzhe Wang, Hao Yang, Xing Fu, Yanyan Zhao, Si Wei, Shijin Wang, Bing Qin</i>	
Multi-Agent VLMs Guided Self-Training with PNU Loss for Low-Resource Offensive Content Detection	39387
<i>Han Wang, Deyi Ji, Junyu Lu, Lanyun Zhu, Hailong Zhang, Haiyang Wu, Liqun Liu, Peng Shu, Roy Ka-Wei Lee</i>	
EA: Managing Green Data Centers Using Deep Reinforcement Learning Without Discounting	39397
<i>Peijian Wang, Thu Nguyen</i>	

Shaping Human–AI Collaboration in Education: Effects of AI-Assisted Decision-Making Paradigms and Human–AI Decision Consistency on Pre-Service Teachers’ Psychological States and Performance	39406
<i>Yingying Wang, Qin Ni, Haoxin Xu, Jiaqi Yin, Tingjiang Wei</i>	
Trustworthy Classification for Complex Social Surveys: A Memory-Enhanced Hierarchical Framework with Calibrated Uncertainty	39415
<i>Zeqiang Wang, Rebecca Oldroyd, Yuqi Wang, Jiageng Wu, Jie Yang, Wei Wang, Nishanth R. Sastry, Jon Johnson, Suparna De</i>	
Fair Graph Learning with Limited Sensitive Attribute Information	39423
<i>Zichong Wang, Jie Yang, Jun Zhuang, Puqing Jiang, Mingzhe Chen, Ye Hu, Wenbin Zhang</i>	
EgoEMS: A High-Fidelity Multimodal Egocentric Dataset for Cognitive Assistance in Emergency Medical Services	39432
<i>Keshara Weerasinghe, Xueren Ge, Tessa Heick, Lahiru Nuwan Wijayasingha, Anthony Cortez, Abhishek Satpathy, John Stankovic, Homa Alemzadeh</i>	
DuoCast: Duo-Probabilistic Diffusion for Precipitation Nowcasting.....	39442
<i>Penghui Wen, Mengwei He, Patrick Filippi, Na Zhao, Feng Zhang, Thomas Francis Bishop, Zhiyong Wang, Kun Hu</i>	
Error Correction in Radiology Reports: A Knowledge Distillation-Based Multi-Stage Framework.....	39451
<i>Jinge Wu, Zhaolong Wu, Ruizhe Li, Tong Chen, Abul Hasan, Yunsoo Kim, Jason Pui-Yin Cheung, Teng Zhang, Honghan Wu</i>	
Explainable Oracle Bone Script Recognition Via Multimodal Pictographic Reasoning	39460
<i>Yin Wu, Zhengxuan Zhang, Jiayu Chen, Chang Xu, Yuyu Luo, Nan Tang, Hui Xiong</i>	
LLM Safety in Judicial AI: A Stress Test of Social Media Influence on Real-World Judgments.....	39468
<i>Yixuan Xie, Yang He, Xiaoyu Yang, Xu Gai, Pan Hui</i>	
Judging by the Rules: Compliance-Aligned Framework for Modern Slavery Statement Monitoring.....	39477
<i>Wenhao Xu, Akshatha Arodi, Jian-Yun Nie, Arsène Fansi Tchango</i>	
BCWildfire: A Long-Term Multi-Factor Dataset and Deep Learning Benchmark for Boreal Wildfire Risk Prediction	39486
<i>Zhengsen Xu, Sibò Cheng, Lanying Wang, Hongjie He, Wentao Sun, Jonathan Li, Lincoln Linlin Xu</i>	
STEAMROLLER: A Multi-Agent System for Inclusive Automatic Speech Recognition for People Who Stutter.....	39495
<i>Ziqi Xu, Yi Liu, Yuekang Li, Ling Shi, Kailong Wang, Yongxin Zhao</i>	
PandemIQ Llama: A Domain-Adapted Foundation Model for Enhanced Pandemic Intelligence.....	39504
<i>Jingmei Yang, Mahtab Talaei, Britta Lassmann, Nahid Bhadelia, Ioannis Ch. Paschalidis</i>	
Multivariate Gaussian Representation Learning for Medical Action Evaluation	39513
<i>Luming Yang, Haoxian Liu, Siqing Li, Alper Yilmaz</i>	
HSKBenchmark: Modeling and Benchmarking Chinese Second Language Acquisition in Large Language Models Through Curriculum Tuning	39522
<i>Qihao Yang, Xuelin Wang, Jiale Chen, Xuelian Dong, Yuxin Hao, Tianyong Hao</i>	
A Human-AI Collaboration Mechanism Study on AIGC Assisted Image Production for Special Coverage.....	39531
<i>Yajie Yang, Yuqing Zhao, Xiaochao Xi, Yinan Zhu</i>	

ChatCLIDS: Simulating Persuasive AI Dialogues to Promote Closed-Loop Insulin Adoption in Type 1 Diabetes Care.....	39539
<i>Zonghai Yao, Talha Chafekar, Junda Wang, Shuo Han, Feiyun Ouyang, Junhui Qian, Lingxi Li, Hong Yu</i>	
Advancing Safe Mechanical Ventilation Using Offline RL with Hybrid Actions and Clinically Aligned Rewards	39548
<i>Muhammad Hamza Yousuf, Jason Li, Sahar Vahdati, Raphael Theilen, Jakob Wittenstein, Jens Lehmann</i>	
TrustEnergy: A Unified Framework for Accurate and Reliable User-Level Energy Usage Prediction	39558
<i>Dahai Yu, Rongchao Xu, Dingyi Zhuang, Yuheng Bu, Shenhao Wang, Guang Wang</i>	
Uncovering Bias Paths with LLM-Guided Causal Discovery: An Active Learning and Dynamic Scoring Approach	39567
<i>Khadija Zanna, Akane Sano</i>	
Enhancing Binary Encoded Crime Linkage Analysis Using Siamese Network	39576
<i>Yicheng Zhan, Fahim Ahmed, Amy Burrell, Matthew Tonkin, Sarah Galambos, Jessica Woodhams, Dalal Alrajeh</i>	
Improving Day-Ahead Grid Carbon Intensity Forecasting by Joint Modeling of Local-Temporal and Cross-Variable Dependencies Across Different Frequencies	39585
<i>Bowen Zhang, Hongda Tian, Adam Berry, A. Craig Roussac</i>	
Scaling Equitable Reflection Assessment in Education Via Large Language Models and Role-Based Feedback Agents	39594
<i>Chenyu Zhang, Xiaohang Luo</i>	
Game Theory Based Community-Aware Opinion Dynamics	39603
<i>Shanfan Zhang, Yongyi Lin, Xiaoting Shen, Zhan Bu, Yuan Rao</i>	
Realistic Curriculum Reinforcement Learning for Autonomous and Sustainable Marine Vessel Navigation	39612
<i>Xiaocai Zhang, Zhe Xiao, Maohan Liang, Tao Liu, Haijiang Li, Wenbin Zhang</i>	
Adaptive Regulation Via Dual-Layer Evolution (ARDE): A Multi-Agent Approach to Balancing Efficiency, Fairness, and Diversity in Crowdsourced Platforms	39620
<i>Xuwen Zhang, Xiao Xue, Xia Xie, Qun Ma</i>	
HierarNet: Independent Interactive Hierarchical Disease Outbreak Forecasting	39628
<i>Zichi Zhang, Phi Hung Nguyen, Ngoc Phu Doan, Viet-Hung Tran, Xuan Hoang Nguyen, Hui Wang, Hans Vandierendonck, Son Thai Mai</i>	
Investigating Social Bias Propagation in Federated Fine-Tuning of Large Language Models.....	39637
<i>Jiaxu Zhao, Meng Fang, Mingze Zhong, Shunfeng Zheng, Ling Chen, Mykola Pechenizkiy</i>	
Assessing Automated Fact-Checking for Medical LLM Responses with Knowledge Graphs.....	39646
<i>Shasha Zhou, Mingyu Huang, Jack Cole, Charles Britton, Ming Yin, Jan Wolber, Ke Li</i>	
The Fixed-Point of Distrust: A Formal Theory of Perceived Systemic Incompetence	39657
<i>Weilin Zhou</i>	

SENIOR MEMBER PRESENTATION

On Trustworthy, Explainable, and Verifiable High-Level Autonomy Via Hierarchical Planning	39666
<i>Roman Barták</i>	
The Shepherd Test: How Will Super Intelligent Agents Balance Care and Control in Asymmetric Relationships?	39675
<i>Djallel Bouneffouf, Matthew Riemer, Kush R. Varshney</i>	
Multi-Armed Bandits Meet Large Language Models.....	39682
<i>Djallel Bouneffouf, Raphael Feraud</i>	
Multi-Agent Path Finding with Unassigned Agents (MAPFUA).....	39691
<i>Ariel Felner, Roni Stern</i>	
Aligning Generative Music AI with Human Preferences: Methods and Challenges.....	39699
<i>Dorien Herremans, Abhinaba Roy</i>	
Democratizing LLM Efficiency: From Hyperscale Optimizations to Universal Deployability	39707
<i>Hen-Hsen Huang</i>	
Knowledge-Guided Machine Learning: A Paradigm Shift in AI for Science.....	39715
<i>Anuj Karpatne, Xiaowei Jia, Vipin Kumar</i>	
Complex Reasoning Over Vision and Language --Leveraging Neurosymbolic AI.....	39718
<i>Parisa Kordjamshidi</i>	
The Tatort Test of Intelligence: Towards Narrative Comprehension as a Benchmark for AI	39722
<i>Stefan Kramer, Lennart Baur, Lars Reinhardt</i>	
System L: Toward System 2-Style Legal Reasoning.....	39728
<i>Chuanyi Li, Yi Feng, Vincent Ng</i>	
Rethinking AI: From Functions to Functors	39737
<i>Sridhar Mahadevan</i>	
Responsible Mechanism Design.....	39745
<i>Pavel Naumov</i>	
Is Word Sense Disambiguation Dead in the LLM Era?.....	39753
<i>Roberto Navigli</i>	
Why Isn't Relational Learning Taking Over the World?	39763
<i>David Poole</i>	
10 Open Challenges Steering the Future of Vision-Language-Action Models.....	39771
<i>Soujanya Poria, Navonil Majumder, Chia-Yu Hung, Amir Ali Bagherzadeh, Chuan Li, Kenneth Kwok, Ziwei Wang, Cheston Tan, Jiajun Wu, David Hsu</i>	
Implications for AI Research: Applying Lessons from the Expert Systems Boom and Bust to the Current Large-Language Model Boom	39780
<i>Ted Senator</i>	
Scaling Up AI Alignment	39785
<i>Aarti Singh</i>	

Voting in Divisible Settings: A Survey 39789
Warut Suksompong, Nicholas Teh

Right Looks, Wrong Reasons: Compositional Fidelity in Text-To-Image Generation 39797
Mayank Vatsa, Aparna Bharati, Richa Singh

Author Index