

ACL 2026

**The 64th Annual Meeting of the Association for
Computational Linguistics (ACL 2026)**

Proceedings of the Conference - Volume 1: Long Papers

July 2-7, 2026

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

Diamond

Bloomberg
Engineering




Platinum



Appen



ORACLE



Tencent



北京中关村学院 × 中关村人工智能研究院
Zhongguancun Academy × Zhongguancun Institute of Artificial Intelligence



Gold



JPMorganChase




Silver



Bronze



Academic Ally



©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 979-8-89176-390-6

Message from the General Chair

It is my great pleasure and honor to welcome you to the 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026), held in San Diego, California, USA, July 2–7, 2026. ACL 2026 continues our field’s tradition of excellence in scholarship, innovation, and inclusivity, and I am deeply grateful to the many volunteers who have worked tirelessly to bring this event to life.

It is exciting to see the many developments in the field reflected at our conference, such as the advancement of large language models (born out of our community) into many application directions. This is not without challenges, partly due to unsolved problems, and partly due to misuse of the technology. All this is reflected well at ACL 2026. The conference also reflects that the field is not only about advancing technology, but also the science of fundamentally understanding language and how it is used by many different speaker communities.

I want to express my deepest thanks to our Program Chairs — Maria Liakata, Viviane P. Moreira, Jijun Zhan, and David Jurgen — who have overseen the reviewing and selection process, and shaped a compelling and diverse scientific program. It is impossible to overstate how much work and dedication they put in to make this conference successful. Overall, we received more than 12,148 submissions and accepted 2,296 at the main conference and 2,163 as findings. This was made possible through close coordination with the ACL Rolling Review (ARR) team and I sincerely thank them. We are also deeply grateful to the many senior area chairs, area chairs, reviewers, and the Best Paper Committee (led by Antonios Anastasopoulos and Eunsol Choi), whose dedication ensured the high quality of our program. Isabelle Augenstein served as Conflict Chair to shepherd papers that were submitted by conference chairs.

I am also grateful to our Technical Open Review Chairs, Shuoyang Ding and Yi R. (May) Fun, for their behind-the-scenes work. For those joining virtually, our Virtual Infrastructure Chairs — Gael Guibon and Parisa Kordjamshidi — have worked to make the hybrid conference experience inclusive and engaging. I would also like to recognize the efforts of our Ethics Chairs — Margaret Mitchell and Paco Guzman — who have led the work aimed at ensuring that submissions meet the ethical standards of our field.

ACL 2026 hosts a rich set of 37 workshops and 6 tutorials, thanks to the dedicated efforts of our Workshop Chairs — Loic Barrault and Yang Zhao — and our Tutorial Chairs — Kenton Murray and Jacob Andreas. Our Demonstration Track Chairs — Greg Durrett and Ping Jian — have put together an impressive set of system demonstrations of cutting-edge innovations in NLP. I am proud to say that we also have a rich Industry Track, thanks to Yunyao Li, Georg Rehm, and Mei Tu who created a “conference in the conference”. Their experience from organizing the industry track already last year helped advance the visibility and impact of the event. We also hosted a special linguistic symposium entitled “Linguistics and NLP in the LLM Era” which featured Allyson Ettinger, Richard Futrell, Zoey Liu and Tom McCoy. As is customary, TACL and CL papers are presented at the conference. Thanks go to the TACL Editors-in-Chief — Asli Celikyilmaz, Roi Reichart, Dilek Hakkani-Tur, and Yun-Nung (Vivian) Chen — and the CL Editor-in-Chief, Wei Lu, for their coordination efforts.

This conference could only have been made possible thanks to the invaluable work of our ACL Business Manager and Director of Events, Jenn Rachford, and her team. Our Visa Chairs — David Chiang and Swabha Swayamdipta — have helped numerous participants with their visa needs and any logistical difficulties in attending the conference. I also thank our Student Volunteer Chairs — Jingbo Shang, Lianhui Qin, Xuezhe Ma, and Muhao Chen — for their behind-the-scenes work, which proved indispensable during the conference.

Additional thanks go to Santosh T.Y.S.S., Juan Diego Rodriguez, and Ona de Gibert, our Student Research Workshop Chairs, who, with the support of faculty advisors Yves Scherrer, Elias Stengel-Eskin, and Snigdha Chaturvedi, devoted their work to supporting and mentoring the next generation of researchers. They also secured additional funding from the Vienna Meeting Fund — congratulations!

Our Publication Chairs — David Stap, Haoran Xu, and Jan Niehues — have ensured timely and high-quality conference proceedings. I also thank the Handbook Chairs, Jia Xu and Ryan Cotterell, for assembling the handbook that will guide attendees through the conference program. Of course our conference

could not exist without the support of our sponsors! We are deeply grateful to our Sponsorship Chairs — Baotian Hu and Steve Richardson — and to Joel Tetrault, ACL’s Sponsorship Director, for securing the generous support that enables us to keep the conference accessible to as many participants as possible. Fostering an inclusive environment remains a key goal of ACL. Our Diversity and Inclusion Chairs — Ximena Gutierrez-Vasques, Anietie Andy, Gabrielle Gaudeau, and Felermino Ali — have worked to support diverse participation and organize activities that reflect the global nature of our vibrant community. An inclusive ACL allows researchers from all areas and backgrounds to contribute, collaborate, and thrive — enriching both the scientific conversation and the broader societal impact of our work. We remain committed to building a community where everyone feels welcomed and valued.

Thank you to the Publicity and Social Media Chairs — Raj Dabre, Wei Xu, Flor Miriam Plaza-del-Arco, Debora Nozza, and Yaping Zhang — for their communication work through social media, extending the conference’s reach to a wider audience. And this brings me to our Website and App Chairs — Abulhair Saparov and Najoung Kim — who kept our community informed with timely updates about all aspects of the conference. The huge number of emails and communications that we received was smoothly managed throughout the year by our Internal Communication Chairs, Ming Jiang and Margot Mieskes. We are very grateful!

Last but not least, I wish to thank the ACL Executive Committee for their guidance, and my fellow past General Chairs for generously sharing their experience and insights.

ACL 2026 is a community-wide effort. Whether you are presenting a paper, leading a workshop or tutorial, volunteering, mentoring, supporting or attending to learn and connect, your presence and participation make a difference at this conference. Thank you for being part of it!

Welcome to San Diego. Welcome to ACL 2026!

Philipp Koehn

Johns Hopkins University
General Chair, ACL 2026

Message from the Program Chairs

Welcome to the 64th Annual Meeting of the Association for Computational Linguistics! ACL 2026 will be held in a hybrid format, offering attendees the option to join us in person in San Diego, California, or to participate remotely from anywhere in the world.

Review Process. All submissions to ACL 2026 went through a two-stage review process. First, papers were submitted to the ACL Rolling Review (ARR), where they were reviewed by reviewers and received meta-reviews from Area Chairs. Then, authors had the option to commit their reviewed papers to ACL via a separate ACL 2026 commitment site. At this stage, Senior Area Chairs provided recommendations, and the Program Chairs made final acceptance decisions based on SAC recommendations, but also with a view to obtaining a balanced program and prioritizing work at a mature state of development. This process is consistent with previous conferences, including ACL 2025, EMNLP 2025, and EACL 2026. We worked closely with the ARR team, particularly the ARR January 2026 Editors-in-Chief, and served as guest Editors-in-Chief for this round. We helped recruit new reviewers and Area Chairs to ARR, resulting in 8,594 reviewers and 1,434 Area Chairs in the January ARR cycle, which received 10,518 submissions. ACL also recruited 255 Senior Area Chairs to oversee the review and meta-review process. Additionally, we collaborated with the ARR October 2025 cycle, which was shared with EACL 2026, and received 1,630 submissions listing ACL as their preferred venue. We recruited 23 Senior Area Chairs and also served as “super” Senior Area Chairs, overseeing the reviewing process of a large number of submissions.

Coordinating a review process of this scale required the efforts of a large number of volunteers. Delays at multiple stages of the process, including reviewing, meta-reviewing, and recommendation, placed additional burdens on the Program Chairs and ARR Editors. We are particularly grateful to those who absorbed this additional workload, serving as emergency Reviewers, Area Chairs, and Senior Area Chairs. For the ACL commitment phase, Senior Area Chairs made recommendations for 6,564 committed papers based on the reviews, meta-reviews, author responses, and the papers themselves, with final acceptance decisions made by the Program Chairs. The vast majority of papers committed to ACL 2026 came from the ARR January 2026 cycle (88%), followed by ARR October 2025 (10.4%). A small number of papers were committed directly from earlier cycles, including July 2025 (1.2%) and May 2025 (0.4%).

During the ARR review process, reviewers identified many non-existent references, which led to those submissions being desk rejected. During the final consistency checks of the camera-ready versions of papers accepted to ACL 2026, we identified over 100 papers that contained citations to non-existent literature. Since these papers were going to be published and were no longer subject to the anonymity requirements, WisPaper was used to identify potential citation issues, and then all flagged citations were reviewed by human experts (PCs and SACs) to confirm that references did not exist. Whether human-authored or LLM-generated, the inclusion of these non-existent references was a clear violation of the ACL Policy on Publication Ethics. For this reason, we had to reject such papers at this late stage. We recommend that future cycles should run such checks both at the submission stage and at the camera-ready stage to prevent the contamination of the scientific literature with false references and to identify such cases as early as possible.

Acceptance Rate. In total, there are 2,296 papers accepted to the Main Conference and 2,163 papers accepted to Findings. The acceptance rate calculation follows the precedent set by previous conferences that go through ACL Rolling Review (ARR), e.g., ACL 2025. The calculation takes into account the multi-stage process of ARR, in which a paper may be revised in ARR and later committed to the conference. In total, we had 12,148 unique submissions across the October 2025 (1,630) and January 2026 ARR (10,518). The acceptance rate is 18.9% for Main Conference papers and 17.8% for Findings papers.

Special Theme: Explainability of NLP Models. Following the precedent of theme tracks from previous events, ACL 2026 features a special theme designed to reflect on and stimulate discourse regarding

the current trajectory of the NLP field. Explainability encompasses the methodologies and techniques intended to render the internal mechanisms of complex NLP models, such as Large Language Models (LLMs), transparent and human-interpretable. It seeks to move beyond the “black box” paradigm, where predictions are accepted without scrutiny, and instead aims to elucidate the underlying reasoning of specific model outputs. Explainability is fundamental to establishing trust, ensuring fairness, and enabling responsible deployment. By identifying a model’s potential dependence on spurious correlations or inherent societal biases, explainability empowers developers to diagnose system errors, enhance robustness, and ensure accountability. This is particularly vital in high-stakes domains—including healthcare, finance, and law—where interpreting the rationale behind a decision is as essential as the decision itself. The theme track invited empirical and theoretical work, as well as surveys and position papers reflecting on the Explainability of NLP Models. Suggested topics of discussion included: (i) How do explainability methods need to be adapted for different model architectures? Can we develop a unified framework to evaluate explanations across these architectures? (ii) How can we rigorously and quantitatively evaluate the quality of an explanation? What metrics can reliably measure the faithfulness (accuracy of the model’s reasoning) and plausibility (human-perceived reasonableness) of an explanation? (iii) Can explanations be used to reliably detect when a model is making a biased prediction based on sensitive attributes? How can input-based explanations help mitigate social biases during model training? (iv) Can we use explanations to systematically find and fix problems in the training data itself, such as spurious correlations or annotation errors? How can explainability facilitate a human-in-the-loop process for iterative data refinement? (v) Can we identify specific directions, mechanisms, patterns, or “knobs” within a model’s internal activations that control high-level behaviors like abstaining from unanswerable questions? Can we design models that are inherently more interpretable?

We received 124 submissions to the special theme track during the review phase. Among these, 29 papers were accepted to the main conference and a further 27 to the Findings of ACL 2026. The conference will also feature a panel discussion on this theme, with leading experts from both academia and industry, with one of the keynote talks, by Tania Lombrozo, also focused on explainability.

Best Paper Selection. ACL 2026 implemented the updated ACL award policy that seeks to expand the pool of work recognized as outstanding. In total, 158 papers were nominated by the reviewers, area chairs, and senior area chairs for awards. The best paper committee assessed these papers to select the best papers, outstanding papers, and special awards for social impact and best resource. Separately, the senior area chairs also nominated their favorite papers as SAC Highlights. The selection process was double-blind, using anonymized camera-ready versions of the papers. The final selection was made by the best paper committee, and the winners will be announced during the opening ceremony to provide early recognition of the selected contributions and to maximize their visibility within the conference program. In addition, we have the 2026 ACL Dissertation Award, the CL high impact paper award, and the test of time awards.

Program Composition & Presentation Modes. Based on feedback from the conference support staff and the Underline team after ACL 2025, we decided to hold the virtual presentation sessions during the main conference. This enables us to align the virtual sessions with time slots that suit participants in Europe and Asia. This approach allows virtual attendees to participate concurrently with the physical event, avoiding the need for organizers and attendees to engage with the conference twice and separately. We recognize that many virtual attendees may not be able to attend keynotes during their scheduled times at ACL, so we added a third virtual-only keynote in a timezone aligned with Europe and Asia to allow greater live participation and visibility on this year’s theme.

The in-person conference will hold 7 Oral/Poster Sessions running in parallel across 10 rooms. This year, 504 main conference papers were selected for oral presentation by the program chairs. The goal was to create a balanced program that featured a diverse set of topics, rather than selecting papers solely on the basis of their review scores. To accommodate a larger number of oral presentations, presentation time was reduced to 8 minutes, followed by 2 minutes of questions. The in-person conference will have

2,078 poster presentations from Main and Findings papers. Rather than have all posters from a track in the same session, sufficiently-sized tracks were divided into two or more sessions to allow authors presenting in one session to attend the other poster sessions for their track. This year, Findings papers will be presented alongside the Main Conference posters, demos, and the Student Research Workshop. This follows the positive experience and feedback of ACL 2025 of holding Main and Findings posters concurrently. In addition to the main conference papers, the ACL program also includes 4 papers accepted by Computational Linguistics (CL) and 27 papers accepted by Transactions of the ACL (TACL). Among these, 18 journal papers will be presented in person as oral presentations, thematically distributed across appropriate sessions.

The virtual conference will feature over 1K presentations, including 752 Main Conference and 176 Findings Presentations, as well as Demos, Student Research Workshop, TACL, CL, Industry papers, and a virtual-only keynote talk.

Keynotes. We are happy to have three keynote talks at the conference, two in-person and one virtual.

- Philip Resnik (University of Maryland): A New Balancing Act: Reflections on the Relationship between Computational Linguistics and AI
- Yue Zhang (Westlake University): How Can AI Help Us Manage AI Conferences (where Troubles Are Caused by AI) – virtual
- Tania Lombrozo (Princeton University): Explanation and Understanding: A Perspective from Cognitive Science

Acknowledgments. Organizing ACL 2026 has been a collaborative effort, made possible by the dedication and hard work of thousands of people. We gratefully acknowledge the support and contributions of the following people:

- The General Chair, Philipp Koehn;
- The ARR Editors-in-Chief of the January 2026 cycle (Sarvnaz Karimi, Vincent Ng, Xiaodan Zhu, and Margot Mieskes) and the entire ARR team, especially Holy Lovenia, Freda Shi, and Yuntian Deng from the tech team;
- The technical OpenReview chairs, Shuoyang Ding, May Fung, and the OpenReview support team, in particular, Rachel, Emilia, and Harold, for technical help in setting up ACL 2026 on the OpenReview platform;
- The 255 Senior Area Chairs;
- The 1,434 Area Chairs and the 8,594 reviewers, and 171 ethics reviewers;
- The best paper committee chairs, Antonios Anastasopoulos and Eunsol Choi, and the best paper committee members;
- The ethics chairs, Margaret Mitchell and Paco Guzman;
- The conflicts chairs, Isabelle Augenstein and Joel Tetreault;
- The workshop chairs, Loic Barrault and Yang Zhao;
- The tutorial chairs, Kenton Murray and Jacob Andreas;
- The industry track chairs, Georg Rehm, Yunyao Li, and Mei Tu;
- The demonstration chairs, Greg Durrett and Ping Jian;

- The internal communications chairs, Ming Jiang and Margot Mieskes;
- The website and conference app chairs, Abulhair Saparov and Najoung Kim;
- The publication chairs, David Stap, Haoran Xu, and Jan Niehues;
- The handbook chairs, Jia Xu and Ryan Cotterell;
- The visa chairs, David Chiang and Swabha Swayamdipta;
- The publicity chairs, Raj Dabre, Wei Xu, Flor Miriam Plaza-del-Arco, Debora Nozza, and Yaping Zhang;
- The student research workshop chairs, Santosh T.Y.S.S., Juan Diego Rodriguez, and Ona de Gibert;
- The student research workshop chairs faculty advisors, Yves Scherrer, Elias Stengel-Eskin, and Snigdha Chaturvedi;
- The student volunteer chairs, Jingbo Shang, Lianhui Qin, Xuezhe Ma, and Muhao Chen;
- The diversity and inclusion chairs, Ximena Gutierrez-Vasques, Aniete Andy, Gabrielle Gaudeau, and Felermino Ali;
- The sponsorship chairs, Baotian Hu and Steve Richardson;
- The virtual infrastructure chairs, Gael Guibon, and Parisa Kordjamshidi;
- The ACL Anthology Director Matt Post and his team;
- The TACL editors-in-chief (Asli Celikyilmaz, Roi Reichart, Dilek Hakkani-Tur, and Vivian Chen) and CL Editor-in-Chief Wei Lu for coordinating TACL and CL presentations with us;
- Past program chairs for *CL conferences, in particular Christos Christodoulopoulos (EMNLP 2025), Vera Demberg (EACL 2026), and Ekaterina Shutova (ACL 2025), for information and support;
- Qi Zhang, developer of WisPaper, for checking camera-ready papers to ensure they did not contain hallucinated references;
- Damira Mrsic and the Underline team for help with the program;
- Jennifer Rachford and the entire conference support staff for their crucial help and patience as we learned our way through running such a complex process;
- All the authors of papers who submitted their papers for review in the ARR 2025 October and ARR 2026 January cycles, and those who committed to the ACL 2026 conference.

We hope you enjoy this year's huge and diverse program!

Maria Liakata (Queen Mary University of London, UK)
 Viviane Moreira (UFRGS, Brazil)
 Jiajun Zhang (Chinese Academy of Sciences, China)
 David Jurgens (University of Michigan, USA)
 ACL 2026 Program Co-Chairs

Organizing Committee

General Chair

Philipp Koehn, Johns Hopkins University

Program Chairs

Maria Liakata, Queen Mary University of London
Viviane P. Moreira, Institute of Informatics at UFRGS, Brasil
Jiajun Zhang, Chinese Academy of Sciences
David Jurgens, University of Michigan

Workshop Chairs

Loic Barrault, Meta
Yang Zhao, Chinese Academy of Sciences

Tutorial Chairs

Kenton Murray, Johns Hopkins University
Jacob Andreas, Massachusetts Institute of Technology

Demonstration Chairs

Greg Durrett, New York University
Ping Jian, Beijing Institute of Technology

Student Research Workshop Chairs

Santosh T.Y.S.S., Technical University of Munich & Amazon
Juan Diego Rodriguez, University of Texas at Austin
Ona de Gibert, University of Helsinki

Student Research Workshop Chairs: Faculty Advisors

Yves Scherrer, University of Oslo
Elias Stengel-Eskin, University of Texas at Austin
Snigdha Chaturvedi, University of North Carolina

Publication Chairs

David Stap, NXAI
Haoran Xu, Microsoft
Jan Niehues, Karlsruhe Institute of Technology

Handbook Chairs

Jia Xu, Stevens Institute of Technology

Ryan Cotterell, ETH Zürich

Sponsorship Chairs

Baotian Hu, Harbin Institute of Technology
Steve Richardson, Brigham Young University

Diversity and Inclusion Chairs

Ximena Gutierrez-Vasques, Universidad Nacional Autonoma de Mexico
Anietie Andy, Howard University
Gabrielle Gaudeau, University of Cambridge
Felermimo Ali, Microsoft Research Africa

Publicity Chairs

Raj Dabre, Google
Wei Xu, Georgia Tech
Flor Miriam Plaza-del-Arco, University of Leiden
Debora Nozza, Bocconi University
Yaping Zhang, Chinese Academy of Sciences

Website and Conference App Chairs

Abulhair Saparov, Purdue University
Najoung Kim, Boston University

Ethics Chairs

Margaret Mitchell, Huggingface
Paco Guzman, Meta

Student Volunteer Chairs

Jingbo Shang, University of California at San Diego
Lianhui Qin, University of California at San Diego
Xuezhe Ma, University of Southern California
Muhao Chen, University of California at Davis

Technical Open Review Chairs

Shuoyang Ding, NVIDIA
Yi R. (May) Fung, HKUST

Virtual Infrastructure Chairs

Gael Guibon, Université Sorbonne Paris Nord
Parisa Kordjamshidi, Michigan State University

Internal Communication Chairs

Ming Jiang, Indiana University Indianapolis
Margot Mieskes, University of Applied Sciences, Darmstadt

Industry Track Chairs

Georg Rehm, German Research Institute for Artificial Intelligence
Yunyao Li, Adobe
Mei Tu, Samsung

Best Paper Chairs

Antonios Anastasopoulos, George Mason University
Eunsol Choi, New York University

Visa Chairs

David Chiang, University of Notre Dame
Swabha Swayamdipta, University of Southern California

Conflict Chair

Isabelle Augenstein, University of Copenhagen

Keynote Talk

A New Balancing Act: Reflections on the Relationship between Computational Linguistics and AI

Philip Resnik
University of Maryland



Sunday, July 5, 2026 – Time: 09:30 – Room: Harbor Ballroom Level 2

Abstract: Beginning in the mid-1980s, a “statistical revolution” in our field redefined what it meant to work on language from a computational perspective. Steve Abney, in his essay opening *The Balancing Act: Combining Symbolic and Statistical Approaches to Language* (Klavans and Resnik, eds., MIT Press, 1996), wrote: “In the space of the last ten years, statistical methods have gone from being virtually unknown in computational linguistics to being a fundamental given”. We are again caught up in revolutionary times. Except this time around, the revolution is hurtling forward even faster, and it is redefining not only the research we do, but what it means to do research in the first place, the role of research in society, and society itself. What does this mean for ACL – what defines us as a community, what is our role, and where do we go from here?

Bio: Philip Resnik is a Professor at the University of Maryland with joint appointments in the Department of Linguistics and the Institute for Advanced Computer Studies. In 2020 he was named an ACL Fellow for significant contributions to symbolic-statistical methods for natural language processing, multilinguality, and the interdisciplinary study of language. Philip’s most recent research has focused in three main areas. The first is computational social science, with an emphasis on qualitative analysis and connecting the signal available in people’s language use with underlying mental state. The second is the computational cognitive neuroscience of language, using computational modeling in connection with brain imaging to look at the role of context and predictive processing during online language comprehension. The third involves fundamental questions about how current AI models relate to human cognition and to human society. Outside academia, Philip’s industry experience includes research at Bolt Beranek and Newman and Sun Microsystems Laboratories, as well as an internship at IBM T.J. Watson Research Center, and in entrepreneurial life he has been a technical co-founder of CodeRyte (clinical NLP, acquired by 3M in 2012), an advisor to FiscalNote (machine learning and analytics for government relations, went public in 2022), and he currently serves as an advisor to Trustible (technology provider for responsible AI governance). Philip was an undergrad in CS at Harvard and earned his PhD in Computer and Information Science at the University of Pennsylvania.

Keynote Talk

How Can AI Help Us Manage AI Conferences (where Troubles Are Caused by AI)

Yue Zhang
Westlake University



Monday, July 6, 2026 – Time: 13:00 – Room: Underline (Virtual)

Abstract: ACL has become 30+ times larger than two decades ago, and we face issues such as overwhelming participants, outdated papers, and low quality review. The reason can largely be attributed to increasing research in LLM, and AI accelerating research. Can AI help us address challenges in academic publication and evaluation? How can we accommodate AI-driven research? Could the research and publication paradigm be improved? In this talk, I will try to give some thoughts, sharing some of our recent work on AI review automation, paper recommendation, and AI arXiv, before briefly discussing a bit of AI research automation.

Bio: Yue Zhang is a Professor at Westlake University, and a fellow of ACL (<https://frcchang.github.io>). His research interests include fundamental NLP and its machine learning algorithms, and his recent research focuses on LLM reasoning and AI scientist. His major contributions to the field include machine learning algorithms for structured prediction (e.g., parsing and IE), neural NLP models (i.e., lattice and graph LSTM), and generalization for NLP/LM (e.g., OOD and logical reasoning). He co-authored the Cambridge University Press book *Natural Language Processing – a Machine Learning Perspective*. He served as a PC co-chair for CCL 2020, EMNLP 2022, and LMG 2025, and a general chair for AACL 2026 and NLPC 2026. He was a co-chair for the test-of-time award committee of ACL 2024 and 2025. He serves as editor in chief for the LLM journal, action editor for TACL, and associate editor for ACM/IEEE TASLP, ACM TALLIP, and IEEE TBD.

Keynote Talk
**Explanation and Understanding: A Perspective from
Cognitive Science**

Tania Lombrozo
Princeton University



Monday, July 6, 2026 – Time: 16:45 – Room: Harbor Ballroom Level 2

Abstract: What makes some explanations more satisfying than others? Why are humans so driven to explain? In this talk I'll take a cognitive scientific approach to explanation and understanding, drawing on decades of work in psychology, education, and philosophy to argue for a function-first approach to explanations—one on which we begin by characterizing what it is that explaining achieves. I'll suggest that the process of explaining plays a crucial role in discovery and generalization, and that we can understand the structure and phenomenology of explanations by first appreciating how explaining drives learning and helps construct our understanding of the world. I'll also consider some potential implications for explanation and understanding in AI systems.

Bio: Tania Lombrozo is the Arthur W. Marks '19 Professor of Psychology at Princeton University, as well as an Associate of the Department of Philosophy and the University Center for Human Values, the director of the Program in Cognitive Science, and co-director of Natural and Artificial Minds, a research initiative within the AI Lab. She received her Ph.D. in Psychology from Harvard University in 2006 after receiving a B.S. in Symbolic Systems and a B.A. in Philosophy from Stanford University. Dr. Lombrozo's research aims to address foundational questions about cognition using the empirical tools of cognitive psychology and the conceptual tools of analytic philosophy. Her work focuses on explanation and understanding, conceptual representation, categorization, social cognition, causal reasoning, and folk epistemology. She is the recipient of numerous early-career awards including the Stanton Prize from the Society for Philosophy and Psychology, the Spence Award from the Association for Psychological Science, a CAREER award from the National Science Foundation, and a James S. McDonnell Foundation Scholar Award in Understanding Human Cognition. Her trade book on explanation, *Why We Ask Why: The Science of Explanation and the Human Drive to Understand*, will be released in October.

Table of Contents

<i>OctoTools: A Multi-Agent Framework with Extensible Tools for Complex Reasoning</i> Pan Lu, Bowen Chen, Sheng Liu, Rahul Thapa, Joseph Boen and James Zou	1
<i>No Reader Left Behind: Multi-Agent Summaries Everyone Can Understand</i> Jimin Jung, MyoungJin Kim, Jaehyung Seo and Heuseok Lim	87
<i>Discover and Prove: An Open-source Agentic Framework for Hard Mode Automated Theorem Proving in Lean 4</i> Chengwu Liu, Yichun Yin, Ye Yuan, Jiakuan Xie, Botao Li, Siqi Li, Jianhao Shen, Yan Xu, Lifeng Shang and Ming Zhang	117
<i>Your Inference Request Will Become a Black Box: Confidential Inference for Cloud-based Large Language Models</i> Chung-ju Huang, Huiqiang Zhao, Yuanpeng He, Lijian Li, Wenpin Jiao, Zhi Jin, Peixuan Chen and Leye Wang	134
<i>Rhetorical Questions in LLM Representations: A Linear Probing Study</i> Louie Hong Yao, Vishesh Anand, Yuan Zhuang and Tianyu Jiang	155
<i>Dual-Axis Generative Reward Model Toward Semantic and Turn-taking Robustness in Interactive Spoken Dialogue Models</i> Yifu Chen, Shengpeng Ji, Zhengqing Liu, Qian Chen, Wen Wang, Ziqing Wang, Yangzhuo Li, Tianle Liang and Zhou Zhao	173
<i>Different types of syntactic agreement recruit the same units within large language models</i> Daria Kryvosheieva, Andrea Gregor de Varda, Evelina Fedorenko and Greta Tuckute	209
<i>Empowering Tabular Data Preparation with Language Models: Why and How?</i> Mengshi Chen, Yuxiang Sun, Tengchao Li, Jianwei Wang, Kai Wang, Xuemin Lin, Ying Zhang and Wenjie Zhang	228
<i>Learning Diverse Responses with Prefix-Conditioned Supervised Fine-Tuning</i> Zhiyuan Fan, Guanqiao Chen, Yanyi Huang, Mingkuan Zhao, Dadi Guo and Yi R. Fung	247
<i>EASE: Entity-Aware Sub-table Generation for Real-world Multi-table QA</i> Myunghoon Kang, Dahyun Jung, Suhyune Son, Seonmin Koo, Changwoo Chun, Daniel Rim, Haeyoung Kwon, Yuna Hur and Heuseok Lim	277
<i>Benchmarking LLM's Capability in Reasoning over Conflicting Web References</i> Yizhen Yuan, Rui Kong, Dongze Li, Yuanchun Li and Yunxin Liu	303
<i>Parallel Universes, Parallel Languages: A Comprehensive Study on LLM-based Multilingual Counterfactual Example Generation</i> Qianli Wang, Van Bach Nguyen, Yihong Liu, Fedor Splitt, Nils Feldhus, Christin Seifert, Hinrich Schuetze, Sebastian Möller and Vera Schmitt	323
<i>CLEAR: Cross-Lingual Enhancement in Retrieval via Reverse-training</i> Seungyeon Lee, Minhyuk Kim, Seongtae Hong, Youngjoon Jang, Dongsuk Oh and Heuseok Lim	347
<i>Aligning Language Models with Real-time Knowledge Editing</i> Chenming Tang, Yutong Yang, Kexue Wang and Yunfang Wu	363

<i>PosterForest: Hierarchical Multi-Agent Collaboration for Scientific Poster Generation</i> Jiho Choi, Seojeong Park, Seongjong Song and Hyunjung Shim	379
<i>SLR: Automated Synthesis for Scalable Logical Reasoning</i> Lukas Helff, Ahmad Omar, Felix Friedrich, Antonia Wüst, Hikaru Shindo, Rupert Mitchell, Tim Woydt, Patrick Schramowski, Wolfgang Stammer and Kristian Kersting	402
<i>Can Continual Pretraining Bridge the Performance Gap between General-purpose and Specialized Language Models in the Medical Domain?</i> Niclas Doll, Jasper Schulze Buschhoff, Shalaka Satheesh, Hammam Abdelwahab, Héctor Allende-Cid and Katrin Klug	427
<i>MINTQA: A Multi-Hop Question Answering Benchmark for Evaluating LLMs on New and Long-tail Knowledge</i> Jie He, Nan Hu, Wanqiu Long, Jiaoyan Chen and Jeff Z. Pan	445
<i>Sense and Sensitivity: Examining the Influence of Semantic Recall on Long Context Code Understanding</i> Adam Štorek, Mukur Gupta, Samira Hajizadeh, Prashast Srivastava and Suman Jana	480
<i>SenseRel: A Sense-Level Benchmark for Denotational and Connotational Meaning Relations</i> Pierluigi Cassotti, Naomi Baes, Stefano De Pascale, Jáder Martins Camboim de Sá, Francesco Periti, Nick Haslam, Dirk Geeraerts and Nina Tahmasebi	499
<i>WebSTAR: Scalable Data Synthesis for Computer Use Agents with Step-Level Filtering</i> Yifei He, Pranit Chawla, Yaser Souri, Subhojit Som and Xia Song	516
<i>PR-XAI: PageRank-Based Feature Attribution for Transformers</i> Behrooz Azarkhalili, Linyi Li and Maxwell W. Libbrecht	534
<i>CAPruner: Conceptual-Adjacent Scene Graph Pruner for Enhancing 3D Spatial Reasoning of Large Language Models</i> Shengli Zhou, Xiangchen Wang, Guanhua Chen and Feng Zheng	555
<i>MauBERT: Universal Phonetic Inductive Biases for Few-Shot Acoustic Units Discovery</i> Angelo Ortiz Tandazo, Manel Khentout, Youssef Bencheikroun, Thomas Hueber and Emmanuel Dupoux	568
<i>Your Reasoning Model Knows What Counts: Self-Guided Chain-of-Thought Pruning for Efficient Reasoning</i> Zi-Ao Ma, Xian-Ling Mao, Tian Lan, Chen Xu and Zhijing Wu	586
<i>ImReasoner: Improving Memory-based Language Models for Reasoning-in-a-Haystack Tasks</i> Ching-Yun Ko, Payel Das, Sihui Dai, Georgios Kollias, Subhjit Chaudhury, Aurelie C. Lozano and Pin-Yu Chen	606
<i>How Memory Management Impacts LLM Agents: An Empirical Study of Experience-Following Behavior</i> Zidi Xiong, Yuping Lin, Wenya Xie, Pengfei He, Zirui Liu, Jiliang Tang, Himabindu Lakkaraju and Zhen Xiang	623
<i>SACTOR: LLM-Driven Correct and Idiomatic C to Rust Translation with Static Analysis and FFI-Based Verification</i> Tianyang Zhou, Ziyi Zhang, Haowen Lin, Somesh Jha, Mihai Christodorescu, Kirill Levchenko and Varun Chandrasekaran	646

<i>IntraAgent: An LLM Agent for Content-Grounded Information Retrieval through Literature Review</i> Fengbo Ma, Zixin Rao, Xiaoting Li, Zhetao Chen, Hongyue Sun, Yiping Zhao, Xianyan Chen and Zhen Xiang	674
<i>Reasoning over Precedents Alongside Statutes: Case-Augmented Deliberative Alignment for LLM Safety</i> Can Jin, Rui Wu, Tong Che, Qixin Zhang, Hongwu Peng, Jiahui Zhao, Zhenting Wang, Wenqi Wei, Ligong Han, Zhao Zhang, Yuan Cao, Ruixiang Tang and Dimitris N. Metaxas	716
<i>Hybrid Autoregressive-Diffusion Model for Real-Time Sign Language Production</i> Maoxiao Ye, Xinfeng Ye and Sathiamoorthy Manoharan	750
<i>Augur: Modeling Covariate Causal Associations in Time Series via Large Language Models</i> Zhiqing Cui, Binwu Wang, Qingxiang Liu, Yeqiang Wang, Zhengyang Zhou, Yuxuan Liang and Yang Wang	764
<i>AgentRouter: A Knowledge-Graph-Guided LLM Router for Collaborative Multi-Agent Question Answering</i> Zheyuan Zhang, Kaiwen Shi, Zhengqing Yuan, Zehong Wang, Tianyi Ma, Keerthiram Murugesan, Vincent Galassi, Chuxu Zhang and Yanfang Ye	788
<i>MoCa: Modality-aware Continual Pre-training Makes Better Bidirectional Multimodal Embeddings</i> Haonan Chen, Hong Liu, Yuping Luo, Liang Wang, Nan Yang, Furu Wei and Zhicheng Dou	810
<i>ReEx-SQL: Reasoning with Execution-Aware Reinforcement Learning for Text-to-SQL</i> Yaxun Dai, Wenxuan Xie, Xialie Zhuang, Tianyu Yang, Ziyi Liu, Haiqin Yang, Yiyang Yang, Yuhang Zhao, Pingfu Chao and Wenhao Jiang	824
<i>DecIF: Improving Instruction-Following through Decomposition</i> Tingfeng Hui, Pengyu Zhu, Bowen Ping, Ling Tang, Guanting Dong, Yaqi Zhang and Sen Su	848
<i>NavA³: Understanding Any Instruction, Navigating Anywhere, Finding Anything</i> Lingfeng Zhang, Xiaoshuai Hao, Yingbo Tang, Haoxiang Fu, Xinyu Zheng, Pengwei Wang, Zhongyuan Wang, Wenbo Ding and Shanghang Zhang	868
<i>FastV-RAG: Towards Fast and Fine-Grained Video QA with Retrieval-Augmented Generation</i> Gen Li and Peiyu Liu	879
<i>Efficient Provably Secure Linguistic Steganography via Range Coding</i> Ruiyi Yan and Yugo Murawaki	890
<i>Uni-MMMU: A Massive Multi-discipline Multimodal Unified Benchmark</i> Kai Zou, Ziqi Huang, Yuhao Dong, Shulin Tian, Dian Zheng, Hongbo Liu, Jingwen He, Bin Liu, Yu Qiao and Ziwei Liu	908
<i>MemSearch-01: Empowering Large Language Models with Reasoning-Aligned Memory Growth in Agentic Search</i> Sheng Zhang, Junyi Li, Yingyi Zhang, Pengyue Jia, Yichao Wang, Xiaowei Qian, Wenlin Zhang, Maolin Wang, Yong Liu and Xiangyu Zhao	925
<i>CAPC-CG: A Large-Scale, Expert-Directed LLM-Annotated Corpus of Adaptive Policy Communication in China</i> Bolun Sun, Charles Chang, Yuen Yuen Ang, Ruotong Mu, Yuchen Xu, Zhengxin Zhang and Pingxu Hao	944

<i>MathFlow: Enhancing the Perceptual Flow of MLLMs for Visual Mathematical Problems</i> Shuhang Chen, Hangjie Yuan, Yunqiu Xu, Pengwei Liu, Tao Feng, Jun Cen, Zeying Huang and Yi Yang	967
<i>Anchored Sliding Window: Toward Robust and Imperceptible Linguistic Steganography</i> Ruiyi Yan, Shiao Meng and Yugo Murawaki	993
<i>What Makes Good Instruction-Tuning Data? An In-Context Learning Perspective</i> Guangzeng Han and Xiaolei Huang	1013
<i>RoBSA: RoPE-based Blockwise Sparse Multi-head Latent Attention</i> Xinyu Shi, Kairong Luo, Zhen Zheng and Wenguang Chen	1028
<i>XtraGPT: Context-Aware and Controllable Academic Paper Revision via Human-AI Collaboration</i> Nuo Chen, Andre Lin HuiKai, Jiaying Wu, Junyi Hou, Zining Zhang, Qian Wang, Xidong Wang and Bingsheng He	1045
<i>Learn to Relax with Large Language Models: Solving Constraint Optimization Problems via Bidirectional Coevolution</i> Beidan Liu, Zhengqiu Zhu, Chen Gao, Tianle Pu, Yong Zhao, Wei Qi and Quanjun Yin	1075
<i>A Game-Theoretic Negotiation Framework for Cross-Cultural Consensus</i> Guoxi Zhang, Jiawei Chen, Tianzhuo Yang, Jiaming Ji, Yaodong Yang and Juntao Dai	1095
<i>Focusing Condition: Inference-Time Self-Contrastive Steering Elicits Better Conditional Text Embeddings in LLMs</i> Zifeng Cheng, Lingyun Qian, Zhiwei Jiang, Cong Wang, Yafeng Yin, Fei Shen, Ao Zhou and Qing Gu	1135
<i>Exploring Attention Attractors in Large Language Models</i> Ziheng Wang, Zihao Yue, Wenxuan Wang and Qin Jin	1148
<i>Identifying the Periodicity of Information in Natural Language</i> Yulin OU, Yu Wang, Yang Xu and Hendrik Buschmeier	1161
<i>EmoHarbor: Evaluating Personalized Emotional Support by Simulating the User’s Internal World</i> Jing Ye, Lu Xiang, Yaping Zhang and Chengqing Zong	1176
<i>ReasonEmbed: Enhanced Text Embeddings for Reasoning-Intensive Document Retrieval</i> Jianlyu Chen, Junwei Lan, Chaofan Li, Defu Lian and Zheng Liu	1203
<i>SPAGBias: Uncovering and Tracing Structured Spatial Gender Bias in Large Language Models</i> Binxian Su, Haoye Lou, Shucheng Zhu, Weikang Wang, Ying Liu, Dong Yu and Pengyuan Liu	1222
<i>VALU: A Benchmark for Video Anomaly Temporal Localization and Understanding at Multiple Semantic Levels</i> Yixiao He, Menghao Zhang, Haifeng Sun, Jing Wang, Kangheng Lin, Jinghan Wang, Chenye Xu, Pengfei Ren, Qi Qi and Jingyu Wang	1252
<i>CoreGaze: Core Subgraph-Driven Visual Gaze Diffusion for Training-Free Referring Multimodal Large Language Models</i> Xiaoyang Yi, Jing Chen, Yuru Bao and Jian Zhang	1297
<i>RespiraMFM: A Multimodal Foundation Model with Contrastive Audio-Language Alignment for Respiratory Disease Identification</i> Shakhrul Iman Siam, Tiantian Feng, Jiankun Zhang, Shrikanth Narayanan and Mi Zhang ..	1316

<i>Beyond Detection: Evaluating Fallacy Awareness of LLMs in Interactive Scenarios</i> Conghui Niu, Ningxin Wu, Ziran Zhao, Dong Yu, Chen Kang and Pengyuan Liu	1331
<i>Crosscoding Through Time: Tracking Emergence & Consolidation Of Linguistic Representations Throughout LLM Pretraining</i> Deniz Bayazit, Aaron Mueller and Antoine Bosselut	1353
<i>Is EEG-to-Text Feasible in Real-World Scenarios? An In-Depth Analysis Using a Neuropsychology-Inspired Benchmark</i> Zihan Zhang, Yu Bao, Xiao Ding, Tianyi Jiang and Kai Xiong	1378
<i>ControlAudio: Tackling Text-Guided, Timing-Indicated and Intelligible Audio Generation via Progressive Diffusion Modeling</i> Yuxuan Jiang, Zehua Chen, Zeqian Ju, Yusheng Dai, Weibei Dou and Jun Zhu	1394
<i>PrefixNLI: Detecting Factual Inconsistencies as Soon as They Arise</i> Sapir Harary, Eran Hirsch, Aviv Slobodkin, David Wan, Mohit Bansal and Ido Dagan	1414
<i>RAM-SD: Retrieval-Augmented Multi-agent framework for Sarcasm Detection</i> Ziyang Zhou, Ziqi Liu, Yan Wang, Yiming Lin and Yangbin Chen	1434
<i>On the Emergence and Test-Time Use of Structural Information in Large Language Models</i> Michelle Chao Chen, Moritz Miller, Bernhard Schölkopf and Siyuan Guo	1449
<i>SAGE: A Search-AUGmented Evaluation of Large Language Models on Free-Form QA</i> Sher Badshah, Ali Emami and Hassan Sajjad	1466
<i>J4R: Learning to Judge with Equivalent Initial State Group Relative Policy Optimization</i> Austin Xu, Yilun Zhou, Xuan-Phi Nguyen, Caiming Xiong and Shafiq Joty	1492
<i>A Functionality-Grounded Benchmark for Evaluating Web Agents in E-commerce Domains</i> Xianren Zhang, Shreyas Prasad, Di Wang, Qihai Zeng, Suhang Wang, Wenbo Yan and Mat Hans	1512
<i>Reinforcement Learning for Self-Improving Agent with Skill Library</i> Jiongxiao Wang, Qiaojing Yan, Yawei Wang, Yijun Tian, Soumya Smruti Mishra, Zhichao Xu, Megha Gandhi, Panpan Xu and Lin Lee Cheong	1529
<i>CachePrune: Teaching LLMs What Not to Follow via KV-Cache Editing</i> Rui Wang, Junda Wu, Yu Xia, Tong Yu, Ruiyi Zhang, Ryan A. Rossi, Subrata Mitra, Lina Yao and Julian McAuley	1551
<i>Flip-Flop Consistency: Unsupervised Training for Robustness to Prompt Perturbations in LLMs</i> Parsa Hejabi, Elnaz Rahmati, Alireza Salkhordeh Ziabari and Morteza Dehghani	1571
<i>Leave My Images Alone: Preventing Multi-Modal Large Language Models from Analyzing Images via Visual Prompt Injection</i> Zedian Shao, Hongbin Liu, Yuepeng Hu and Neil Zhenqiang Gong	1588
<i>Logic Matters in Lightweight Hallucination Classification for RAG System</i> Ningyuan Yang and Kaizhu Huang	1605
<i>Harnessing Negative Signals: Reinforcement Distillation from Teacher Data for LLM Reasoning</i> Shuyao Xu, Cheng Peng, Jiangxuan Long, Weidi Xu, Wei Chu and Yuan Qi	1618
<i>Grammar Search for Multi-Agent Systems</i> Mayank Singh, Vikas Yadav, Shiva Krishna Reddy Malay, Shravan Nayak, Sai Rajeswar, Sathwik Tejaswi Madhusudhan and Eduardo Blanco	1640

<i>ZoomR: Memory Efficient Reasoning through Multi-Granularity Key Value Retrieval</i>	
David H. Yang, Yuxuan Zhu, Mohammad Mohammadi Amiri, Keerthiram Murugesan, Tejaswini Pedapati, Subhajit Chaudhury and Pin-Yu Chen	1656
<i>Explicit Trait Inference for Multi-Agent Coordination</i>	
Suhaib Abdurahman, Etsuko Ishii, Katerina Margatina, Divya Bhargavi, Monica Sunkara and Yi Zhang	1670
<i>SFTMix: Elevating Language Model Instruction Tuning with Mixup Recipe</i>	
Yuxin Xiao, Shujian Zhang, Marzyeh Ghassemi and Wenxuan Zhou	1705
<i>Psychological Steering in LLMs: An Evaluation of Effectiveness and Trustworthiness</i>	
Amin Banayeezade, Ala N. Tak, Fatemeh Bahrani, Anahita Bolourani, Leonardo Blas, Emilio Ferrara, Jonathan Gratch and Sai Praneeth Karimireddy	1719
<i>Hierarchical Policy Optimization for Simultaneous Translation of Unbounded Speech</i>	
Siqi Ouyang, Shuoyang Ding, Oleksii Hrinchuk, Vitaly Lavrukhin, Brian Yan, Boris Ginsburg and Lei Li	1772
<i>Act as you think: Reinforcing Consistent Reasoning in Medical Visual Question Answering</i>	
Songtao Jiang, Yuan Wang, Ruizhe Chen, Yan Zhang, Ruilin Luo, Bohan Lei, Yeying Jin, Sibao Song, ZhiBo Yang, Jimeng Sun, Jian Wu and Zuozhu Liu	1788
<i>CRISP: Persistent Concept Unlearning via Sparse Autoencoders</i>	
Tomer Ashuach, Dana Arad, Aaron Mueller, Martin Tutek and Yonatan Belinkov	1806
<i>LongSpec: Long-Context Lossless Speculative Decoding with Efficient Drafting and Verification</i>	
Penghui Yang, Cunxiao Du, Fengzhuo Zhang, Haonan Wang, Tianyu Pang, Chao Du and Bo An	1826
<i>Using Perspectival Words Is Harder Than Vocabulary Words for Humans —and Even More So for Multimodal Language Models</i>	
Dota Tianai Dong, Yifan Luo, Po-Ya Angela Wang, Asli Ozyurek and Paula Rubio-Fernandez	1845
<i>CHIMERA: A Knowledge Base of Scientific Idea Recombinations for Research Analysis and Ideation</i>	
Noy Sternlicht and Tom Hope	1871
<i>Towards Robust Real-World Spreadsheet Understanding with Multi-Agent Multi-Format Reasoning</i>	
Houxing Ren, Mingjie Zhan, Zimu Lu, Ke Wang, Yunqiao Yang, Haotian Hou and Hongsheng Li	1906
<i>GATE: Graph-based Adaptive Tool Evolution Across Diverse Tasks</i>	
Jianwen Luo, Yiming Huang, Jinxiang Meng, Fangyu Lei, Shizhu He, Xiao Liu, Shanshan Jiang, Bin Dong, Jun Zhao and Kang Liu	1934
<i>Erasing Without Remembering: Implicit Knowledge Forgetting in Large Language Models</i>	
Huazheng Wang, Yongcheng Jing, Haifeng Sun, Yingjie Wang, Jingyu Wang, Jianxin Liao and Dacheng Tao	1962
<i>Mitigating Hallucinations in Large Vision-Language Models without Performance Degradation</i>	
Xingyu Zhu, Junfeng Fang, Shuo Wang, Beier Zhu, Zhicai Wang, Yonghui Yang and Xiangnan He	1995
<i>APB-V: Accelerating Long-Video Understanding via Sequence-Parallelism-aware Approximate Attention</i>	
Yuxiang Huang, Mingye Li, Xu Han, Chaojun Xiao, Weilin Zhao, Ao Sun, Ziqi Yuan, Hao Zhou, Fandong Meng and Zhiyuan Liu	2010

<i>AraVQA: Building a New Arabic Factoid Visual Question Answering Dataset from Wikipedia</i> Sultan Alrowili, Younes Samih, Abed Alhakim Freihat and Mathan Kumar Eswaran	2026
<i>What is a protest anyway? Codebook conceptualization is still a first-order concern in LLM-era classification</i> Andrew Halterman and Katherine A. Keith	2043
<i>When Misinformation Speaks and Converses: Rethinking Fact-Checking in Audio Platforms</i> Chaewan Chun, Delvin Ce Zhang and Dongwon Lee	2060
<i>Contextual Relevance and Adaptive Sampling for LLM-Based Document Reranking</i> Jerry Huang, Siddarth Madala, Cheng Niu, Julia Hockenmaier and Tong Zhang	2076
<i>Quantifying Aleatoric Uncertainty of In-Context Learning for Robust Measure of LLM Prediction Confidence</i> Jinseok Chung, Minkyong Song, Hyunji Jung and Namhoon Lee	2090
<i>SATQuest: A Verifier for Logical Reasoning Evaluation and Reinforcement Fine-Tuning of LLMs</i> Yanxiao Zhao, Yaqian Li, Zi-Hao Bo, Rinyoichi Takezoe, Haojia Hui, Mo Guang, Renlei, Xiaolin Qin and Kaiwen Long	2109
<i>Not All Directions Matter: Towards Structured and Task-Aware Low-Rank Model Adaptation</i> Xi Xiao, Chenrui Ma, Yunbei Zhang, Chen Liu, Zhuxuanzi Wang, Yanshu Li, Lin Zhao, Guosheng Hu, Tianyang Wang and Hao Xu	2132
<i>ReContraster: Making Your Posters Stand Out with Regional Contrast</i> Peixuan Zhang, Zijian Jia, Ziqi Cai, Shuchen Weng, Si Li and Boxin Shi	2155
<i>Mechanistic Interpretability Should Prioritize Feature Consistency in Sparse Autoencoders</i> Xiangchen Song, Aashiq Muhamed, Yujia Zheng, Lingjing Kong, Zeyu Tang, Mona T. Diab, Virginia Smith and Kun Zhang	2172
<i>EquivPruner: Boosting Efficiency and Quality in LLM-Based Search via Action Pruning</i> Jiawei Liu, Qisi Chen, Jianshu Zhang, Quan Liu and Defu Lian	2211
<i>From Data-Centric to Sample-Centric: Enhancing LLM Reasoning via Progressive Optimization</i> Xinjie Chen, Minpeng Liao, Guoxin Chen, Chengxi Li, Biao Fu, Kai Fan and Xinggao Liu	2227
<i>CARES: Context-Aware Resolution Selector for VLMs</i> Moshe Kimhi, Nimrod Shabtay, Raja Giryes, Chaim Baskin and Eli Schwartz	2243
<i>ChatRI: Reinforcement Learning for Conversational Reasoning and Retrieval Augmented Question Answering</i> Simon Lupart, Mohammad Aliannejadi and Evangelos Kanoulas	2257
<i>Think in Sentences: Explicit Sentence Boundaries Enhance Language Model’s Capabilities</i> Zhichen Liu, Yongyuan Li and Yang Xu	2275
<i>MORPHOGEN: A Multilingual Benchmark for Evaluating Gender-Aware Morphological Generation</i> Mehul Agarwal, Aditya Aggarwal, Arnav Goel, Medha Hira and Anubha Gupta	2289
<i>TeRA: Vector-based Random Tensor Network for High-Rank Adaptation of Large Language Models</i> Yuxuan Gu, Wuyang Zhou, Giorgos Iacovides and Danilo Mandic	2314
<i>GenDis: Generative-Discriminative Dual-View Co-Training for Generalized Category Discovery</i> Xi Chen, Chuan Qin, Jinpeng Li, Shasha Hu, Chao Wang, Hengshu Zhu and Hui Xiong	2330

<i>CodeHacker: Automated Test Case Generation for Detecting Vulnerabilities in Competitive Programming Solutions</i>	
Jingwei Shi, Xinxiang Yin, Jing Huang, Shengyu Tao and Jinman Zhao	2352
<i>Identifying Bias in Machine-generated Text Detection</i>	
Kevin Stowe, Svetlana Afanaseva, Rodolfo C. Raimundo, Yitao Sun and Kailash Patil	2383
<i>DE-CLIP: Few-Shot Anomaly Detection via Difference-Guided Embedding Editing</i>	
Yage Zhang, Yukun Jiang, Michael Backes and Yang Zhang	2396
<i>Culinary Crossroads: A RAG Framework for Enhancing Diversity in Cross-Cultural Recipe Adaptation</i>	
Tianyi Hu, Andrea Morales-Garzón, Jingyi Zheng, Maria Maistro and Daniel Hershcovich	2408
<i>Mixture-of-Minds: Multi-Agent Reinforcement Learning for Table Understanding</i>	
Yuhang Zhou, Mingrui Zhang, Ke Li, Mingyi Wang, Qiao Liu, Qifei Wang, Jiayi Liu, Fei Liu, Serena Li, Weiwei LI, Mingze Gao, Abhishek Kumar, Xiangjun Fan, Zhuokai Zhao and Lizhu Zhang	2424
<i>MMCLIP: Cross-Modal Attention Masked Modelling for Medical Language-Image Pre-Training</i>	
Biao Wu, Yutong Xie, Zeyu Zhang, Vu Minh Hieu Phan, Qi Chen, Ling Chen and Qi Wu	2440
<i>MMSearch-RI: Incentivizing LMMs to Search</i>	
Jinming Wu, Zihao Deng, Wei Li, Yiding Liu, Bo You, Bo Li, Zejun MA and Ziwei Liu	2456
<i>Toward Secure Tuning: Mitigating Security Risks from Instruction Fine-Tuning</i>	
Yanrui Du, Fenglei Fan, Sendong Zhao, Jiawei Cao, Ming Ma, Danyang Zhao, Shuren Qi, Ting Liu and Bing Qin	2488
<i>SDE-SQL: Enhancing Text-to-SQL Generation in Large Language Models via Self-Driven Exploration with SQL Probes</i>	
Wenxuan Xie, Yaxun Dai and Wenhao Jiang	2507
<i>Splits! Flexible Sociocultural Linguistic Investigation at Scale</i>	
Eylon Caplan, Tania Chakraborty and Dan Goldwasser	2526
<i>Thinking beyond the anthropomorphic paradigm benefits LLM research</i>	
Lujain Ibrahim and Myra Cheng	2551
<i>Coordinating Search-Informed Reasoning and Reasoning-Guided Search in Claim Verification</i>	
Qisheng Hu, Quanyu Long and Wenya Wang	2564
<i>Breaking the Generator Barrier: Disentangled Representation for Generalizable AI-Text Detection</i>	
Xiao Pu, Zepeng Cheng, Lin Yuan, Yu Wu and Xiuli Bi	2586
<i>PUPPET: Neural-Symbolic Standardized Patients for Mental Health</i>	
Chen Xu, Yu ji, Zhenyu Lv, Yang Yi, Yizhe Yang, Luyao Ji, Chaoyi Chen, Xianyang Wang, Tian Lan, Zhihua Wang, Juan Wang, Xunde Dong, Fuze Tian, Qunxi Dong and Bin Hu	2599
<i>LightReasoner: Can Small Language Models Teach Large Language Models Reasoning?</i>	
Jingyuan Wang, Yankai Chen, Zhonghang Li and Chao Huang	2635
<i>OneRec-Think: In-Text Reasoning for Generative Recommendation</i>	
Zhanyu Liu, Shiyao Wang, Xingmei Wang, Rongzhou Zhang, Jiaxin Deng, Honghui Bao, Jinghao Zhang, Wuchao Li, Pengfei Zheng, Xiangyu Wu, Yifei Hu, Qigen Hu, Xinchun Luo, Lejian Ren, Zhang Zixing, Qianqian Wang, Kuo Cai, Yunfan Wu, Hongtao Cheng, Zexuan Cheng, Lu Ren, Huanjie Wang, Yi Su, Ruiming Tang, Kun Gai and Guorui Zhou	2664

<i>AwarenessBench: Assessing Cognitive Capabilities of Language Models</i>	
Xiaojian Li, Rongwu Xu, Tianyun Zhang, Yue Wang, Shuo Chen, Qiner Lyu, Briana Zhang, Peiran Yang, Kyle Xue Chen, Haoyuan Shi, Yu Wang and Wei Xu	2682
<i>Beyond Surface-Level Detection: Towards Cognitive-Driven Defense Against Jailbreak Attacks via Meta-Operations Reasoning</i>	
Rui Pu, Chaozhuo Li, Rui Ha, Litian Zhang, Lirong Qiu and Xi Zhang	2742
<i>SCAN: Structured Capability Assessment and Navigation for LLMs</i>	
Zongqi Wang, Tianle Gu, Chen Gong, Xin Tian, Siqi Bao and Yujiu Yang	2760
<i>EMCEE: Improving Multilingual Capability of LLMs via Bridging Knowledge and Reasoning with Extracted Synthetic Multilingual Context</i>	
Hamin Koo and Jaehyung Kim	2800
<i>Can LLMs Learn to Map the World from Local Descriptions?</i>	
Sirui Xia, Aili Chen, Xintao Wang, Tinghui Zhu, Yikai Zhang, Jiangjie Chen and Yanghua Xiao	2823
<i>LEASH: Adaptive Length Penalty and Reward Shaping for Efficient Large Reasoning Model</i>	
Yanhao Li, Lu Ma, Jiaran Zhang, Lexiang Tang, Wentao Zhang and Guibo Luo	2846
<i>ALIGN: Word Association Learning for Cultural Alignment in Large Language Models</i>	
Chunhua Liu, Kabir Manandhar Shrestha and Sukai Huang	2857
<i>RealTalk-CN: A Realistic Chinese Speech Task-Oriented Dialogue Benchmark with Cross-Modal Analysis</i>	
Enzhi Wang, Jiaming Zhou, Yuhang Jia, Aobo Kong, Qicheng Li and Yong Qin	2880
<i>HOPE: Hybrid Optimized Parallel Encoding with Supervised and Unsupervised Semantic Fusion for Depression Symptom Detection</i>	
Tu-Phuong Mai, Minh-Ha Le H., Duc-Luong Tran, Phuong-Anh Chu, Duy-Cat Can and Hoang-Quynh Le	2898
<i>Bias Fitting to Mitigate Length Bias of Reward Model in RLHF</i>	
Kangwen Zhao, Jianfeng Cai, Jinhua Zhu, Ruopei Sun, Dongyun Xue, Wengang Zhou, Li Li and Houqiang Li	2912
<i>Rethinking Sample Polarity in Reinforcement Learning with Verifiable Rewards</i>	
Xinyu Tang, Yuliang Zhan, Zhixun Li, Xin Zhao, Zhenduo Zhang, Zujie Wen, Zhiqiang Zhang and Jun Zhou	2928
<i>Beyond Examples: Towards Automated Thought-level In-Context Reasoning for Large Language Models</i>	
Jinyang Wu, Mingkuan Feng, Shuai Zhang, Feihu Che, Zhengqi Wen, Chonghua Liao, Ling Yang, Haoran Luo, Zheng Lian and Jianhua Tao	2955
<i>Two-Stage Regularization-Based Structured Pruning for LLMs</i>	
Mingkuan Feng, Jinyang Wu, Siyuan Liu, Shuai Zhang, Hongjian Fang, Ruihan Jin, Feihu Che, Pengpeng Shao, Zhengqi Wen and Jianhua Tao	2996
<i>K-Merge: Online Continual Merging of Adapters for On-device Large Language Models</i>	
Donald Shenaj, Ondrej Bohdal, Taha Ceritli, Mete Ozay, Pietro Zanuttigh and Umberto Michieli	3013

<i>SAC: Neural Speech Codec with Semantic-Acoustic Dual-Stream Quantization</i>	
Wenxi Chen, Ruiqi Yan, Yushen Chen, Zhikang Niu, Ziyang Ma, Xiquan Li, Yuzhe Liang, Wenhanlin, Shunshun Yin, Ming Tao, Xinsheng Wang and Xie Chen	3030
<i>CriticLean: Critic-Guided Reinforcement Learning for Mathematical Formalization</i>	
Zhongyuan Peng, Yifan Yao, Kaijing Ma, Shuyue Guo, Yizhe Li, Yichi Zhang, Chenchen Zhang, Yifan Zhang, Zhouliang Yu, Luming Li, Minghao Liu, Yihang Xia, Jiawei Shen, Yuchen Wu, Yixin Cao, Zhaoxiang Zhang, Wenhao Huang, Jiaheng Liu and Ge Zhang	3049
<i>TRACE: Evaluating Execution Efficiency of LLM-Based Code Translation</i>	
Zhihao Gong, Zeyu Sun, Dong Huang, Qingyuan Liang, Jie M. Zhang and Dan Hao	3089
<i>One Tokenizer To Rule Them All: Emergent Language Plasticity via Multilingual Tokenizers</i>	
Diana Abagyan, Alejandro R. Salamanca, Andres Felipe Cruz-Salinas, Kris Cao, Hangyu Lin, Acyr Locatelli, Marzieh Fadaee, Ahmet Üstün and Sara Hooker	3118
<i>It's Not What You Say, It's How You Say It: Evaluating LLM Responses to Expressions of Belief</i>	
Kevin Du, Clara Kümpel, Michelle Wastl and Alex Warstadt	3137
<i>CiPO: Counterfactual Unlearning for Large Reasoning Models through Iterative Preference Optimization</i>	
Junyi Li, Yongqiang Chen and Ningning Ding	3152
<i>DIA-HARM: Dialectal Disparities in Harmful Content Detection Across 50 English Dialects</i>	
Jason S Lucas, Matt Murtagh White, Ali Al-Lawati, Uchendu Uchendu, Adaku Uchendu and Dongwon Lee	3171
<i>Towards Scalable Lifelong Knowledge Editing with Selective Knowledge Suppression</i>	
Dahyun Jung, Jaewook Lee and Heuseok Lim	3215
<i>Optimizing Length Compression in Large Reasoning Models</i>	
Zhengxiang Cheng, Dongping Chen, Mingyang Fu and Tianyi Zhou	3232
<i>Bridging Language and Items for Retrieval and Recommendation: Benchmarking LLMs as Semantic Encoders</i>	
Yupeng Hou, Jiacheng Li, Xiangjun Fu, Zhankui He, An Yan, Xiusi Chen and Julian McAuley	3251
<i>Evolutionary Guided Decoding: Iterative Value Refinement for LLMs</i>	
Zhenhua Liu, Lijun Li, Ruizhe Chen, Yuxian Jiang, Tong Zhu, Zhaochen Su, Wenliang Chen and Jing Shao	3266
<i>CRAFT: Training-Free Cascaded Retrieval for Tabular QA</i>	
Adarsh Singh, Kushal Raj Bhandari, Jianxi Gao, Soham Dan and Vivek Gupta	3284
<i>Evaluating Legal Reasoning Traces with Legal Issue Tree Rubrics</i>	
Jinu Lee, Kyoung-Woon On, Sophia Simeng Han, Arman Cohan and Julia Hockenmaier . . .	3299
<i>KG-MuLQA: A Framework for KG-based Multi-Level QA Extraction and Long-Context LLM Evaluation</i>	
Nikita Tatarinov, Vidhyakshaya Kannan, Haricharana Srinivasa, Arnav Raj, Harpreet Singh Anand, Varun Singh, Aditya Luthra, Ravij Lade, Agam Shah and Sudheer Chava	3323
<i>COMPASS: Enhancing Agent Long-Horizon Reasoning with Evolving Context</i>	
Guangya Wan, Mingyang Ling, Xiaoqi Ren, Rujun Han, Sheng Li and Zizhao Zhang	3360

<i>BracketRank: Large Language Model Document Ranking via Reasoning-based Competitive Elimination</i>	
Abdelrahman Abdallah, Mohammed Ali, Bhawna Piryani and Adam Jatowt	3381
<i>Multi-component Causal Tracing in Large Language Models</i>	
Zirui Yan, Dennis Wei, Dmitriy A Katz, Prasanna Sattigeri and Ali Tajer	3398
<i>Sparse Feature Coactivation Reveals Causal Semantic Modules in Large Language Models</i>	
Ruixuan Deng, Xiaoyang Hu, Miles Gilberti, Shane Storks, Aman Taxali, Mike Angstadt, Chandra Sripada and Joyce Chai	3419
<i>TensorLens: End-to-End Transformer Analysis via High-Order Attention Tensors</i>	
Ido Andrew Atad, Itamar Zimmerman, Shahar Katz and Lior Wolf	3452
<i>Inferring Events from Time Series using Language Models</i>	
Mingtian Tan, Mike A Merrill, Zachary Gottesman, Tim Althoff, David Evans and Thomas Hartvigsen	3469
<i>OASIS: Online Sample Selection for Continual Instruction Tuning</i>	
Minjae Lee, Minhyuk Seo, Tingyu Qu, Tinne Tuytelaars and Jonghyun Choi	3491
<i>Make Mechanistic Interpretability Auditable: A Call to Develop Guidelines via Continuous Collaborative Reviewing</i>	
Michael Lan, Narmeen Fatimah Oozeer, Chaithanya Bandi, Philip Quirke, Austin Meek, Fazl Barez and Amir Abdullah	3516
<i>Mitigating Context Interference for Reliable and Efficient Search Agents</i>	
Boyang Xue, Bin Wu, Shuofei Qiao, Sheng Wang, Rui Wang, Yiming Du, Hongru Wang, Jeff Z. Pan, Emine Yilmaz, Kam-Fai Wong and Aldo Lipani	3541
<i>Chunks as Arms: Multi-Armed Bandit-Guided Sampling for Long-Context LLM Preference Optimization</i>	
Shaohua Duan, Pengcheng Huang, Xinze Li, Zhenghao Liu, Xiaoyuan Yi, Yukun Yan, Shuo Wang, Yu Gu, Ge Yu and Maosong Sun	3559
<i>Long Context Modeling with Ranked Memory-Augmented Retrieval</i>	
Ghadir Alselwi, Hao Xue, Shoaib Jameel, Basem Suleiman, Flora D. Salim and Imran Razzak	3576
<i>A Comprehensive Survey of Process Reward Models: Data Generation, Model Construction, and Usage</i>	
Congmin Zheng, Jiachen Zhu, Zhuoying Ou, Yuxiang Chen, Kangning Zhang, Rong Shan, Zeyu Zheng, Mengyue Yang, Jianghao Lin, Yong Yu and Weinan Zhang	3591
<i>CODERL+: Improving Code Generation via Reinforcement with Execution Semantics Alignment</i>	
Xue Jiang, Yihong Dong, Mengyang Liu, Deng Hongyi, Tian Wang, Yongding Tao, Zhi Jin, Wenpin Jiao and Ge Li	3608
<i>Saber: Efficient Sampling with Adaptive Acceleration and Backtracking Enhanced Remasking for Diffusion Language Model in Code Generation</i>	
Yihong Dong, Zhaoyu Ma, Xue Jiang, Zhiyuan Fan, Jiaru Qian, Yongmin Li, Jianha Xiao, Zhi Jin and Ge Li	3623
<i>Debiased Orthogonal Boundary-Driven Efficient Noise Mitigation</i>	
Hao Li, Jiayang Gu, Jingkuan Song, An Zhang and Lianli Gao	3643
<i>Collaborative Chain-of-Agents for Parametric-Retrieved Knowledge Synergy</i>	
Yi Jiang, Sendong Zhao, Jianbo Li, Haochun Wang, Lizhe Zhang, Yan Liu and Bing Qin	3663

<i>APPSI-139: A Parallel Corpus of English Application Privacy Policy Summarization and Interpretation</i> Pengyun Zhu, Qiheng Sun, Long Wen, Yanbo Wang, Yang Cao, Junxu Liu, Deyi Xiong, Jinfei Liu, Zhibo Wang and Kui Ren	3681
<i>Compressing then Matching: An Efficient Pre-training Paradigm for Multimodal Embedding</i> Da Li, Yuxiao Luo, Keping Bi, Jiafeng Guo, Wei Yuan, Biao Yang, Yan Wang, Fan Yang, Tingting Gao and Guorui Zhou	3707
<i>LLM-XTM: Enhancing Cross-Lingual Topic Models with Large Language Models</i> Minh Chu Xuan, Tien-Phat Nguyen, Linh Ngo Van, Dinh Viet Sang, Nguyen Thi Ngoc Diep and Trung Le	3719
<i>Open Your Model’s Eyes: Video and Context-Aware Multimodal Backchannel Prediction</i> Min-Jae Kim, Jun-Yeong Moon, Mujeeb Sung and Gyeong-Moon Park	3738
<i>Bringing Real-World Relations into Video Generation with Graph-Structured Knowledge</i> Joonhyung Park, Jaeyun Song, Sihwan Park and Eunho Yang	3756
<i>FC-TTS: Style and Timbre Control in Zero-Shot Text-to-Speech with Disentangled Speech Representations</i> Yoonhyung Lee, Hyunsin Park, Jinhwan Park and Jinkyu Lee	3772
<i>SAGE: Sparse Adaptive Guidance for Dependency-Aware Tabular Data Generation</i> Shuo Yang, Zheyu Zhang, Bardh Prenkaj and Gjergji Kasneci	3792
<i>Mask-to-Correct⁺: Leveraging Retriever Diversity for Masking-guided Faithful Fact Correction</i> Payel Santra, Lavisha Sharma, Madhusudan Ghosh and Partha Basuchowdhuri	3808
<i>Native Hybrid Attention for Efficient Sequence Modeling</i> Jusen Du, Jiayi Hu, Zhang Tao, Weigao Sun and Yu Cheng	3826
<i>UniversalRAG: Retrieval-Augmented Generation over Corpora of Diverse Modalities and Granularities</i> Woongyeong Yeo, Kangsan Kim, Soyeong Jeong, Jinheon Baek and Sung Ju Hwang	3843
<i>Crossing the Reward Bridge: Expanding Reinforcement Learning with Verifiable Rewards Across Diverse Domains</i> Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang and Dong Yu	3872
<i>From ID to LLM: Rethinking Representation Learning for Recommendation</i> Song-Li Wu, Zhaocheng Du, Weinan Gan, Jingyi Wang and Xianquan Wang	3893
<i>DMHM: Density-aware Manifold Learning and Hybrid Mahalanobis Energy for LLMs-generated Text Detection</i> Tianle Liu, Zhiliang Tian, Zhen Huang, Tianlun Liu, Jingyuan Huang, Zhaoning Zhang, Chengcheng Shao and Dongsheng Li	3906
<i>Dynamic Emotion and Personality Profiling for Multimodal Deception Detection</i> Li Zheng, Yanyi Luo, Hao Fei, Yuzhe Ding, Yujie Huang, Fei Li, Chong Teng and Donghong Ji	3930
<i>Can We Predict Before Executing Machine Learning Agents?</i> Jingsheng Zheng, Jintian Zhang, Yujie Luo, Yuren Mao, Yunjun Gao, Lun Du, Huajun Chen and Ningyu Zhang	3941
<i>IMPACT: Importance-Aware Activation Space Reconstruction</i> Md Mokarram Chowdhury, Daniel Agyei Asante, Ernie Chang and Yang Li	3975

<i>SOS-LoRA: Static Orthogonal-Subspace Low-Rank Adaptation with Fixed Multi-Scale Scaling</i> Yupeng Chang, Yuan Wu and Yi Chang	3993
<i>SDiaReward: Modeling and Benchmarking Spoken Dialogue Rewards with Modality and Colloquialness</i> Jingyu Lu, Yuhan Wang, Fan Zhuo, Xize Cheng, Changhao Pan, Xueyi Pu, Yifu Chen, Chenyuhao Wen, Tianle Liang and Zhou Zhao	4006
<i>Looking at Radiology Report Generation through a Causal Lens: A Survey</i> Satyam Kumar, Kaustubh Shivshankar Shejole and Pushpak Bhattacharyya	4029
<i>CTTA-T: Continual Test-Time Adaptation for Text Understanding via Teacher-Student with a Domain-aware and Generalized Teacher</i> Tianlun Liu, Zhiliang Tian, Zhen Huang, Xingzhi Zhou, Wanlong Yu, Tianle Liu, Feng Liu and Dongsheng Li	4058
<i>Adam’s Law: Textual Frequency Law on Large Language Models</i> Hongyuan Lu, Zixuan Li, Zefan Zhang, Bowen Cao and Wai Lam	4079
<i>Disco-RAG: Discourse-Aware Retrieval-Augmented Generation</i> Dongqi Liu, Hang Ding, Qiming Feng, Xurong Xie, Zhucun Xue, Chengjie Wang, Jian Li, Jiangning Zhang and Yabiao Wang	4106
<i>Stable Language Guidance for Vision–Language–Action Models</i> Zhihao Zhan, Yuhao Chen, Jiaying Zhou, Qinhan Lyu, Hao Liu, Keze Wang, Liang Lin and Guangrun Wang	4137
<i>Controlling Multimodal Conversational Agents with Coverage-Enhanced Latent Actions</i> Yongqi Li, Hao Lang, Tiejun Qian and Yongbin Li	4160
<i>Guaranteeing Knowledge Integration with Joint Decoding for Retrieval-Augmented Generation</i> Zhengyi Zhao, Shubo Zhang, Zezhong Wang, Yuxi Zhang, Huimin Wang, Yutian Zhao, Yefeng Zheng, Binyang Li, Kam-Fai Wong and Xian Wu	4181
<i>Are Emotion and Rhetoric Neurons in LLM? Neuron Recognition and Adaptive Masking for Emotion-Rhetoric Prediction Steering</i> Li Zheng, Xin Zhang, Shuyi He, Fei Li, Chong Teng, Jiang-Ming Yang, Donghong Ji and Zhuang Li	4206
<i>Towards Fine-Grained and Multi-Granular Contrastive Language-Speech Pre-training</i> Yifan Yang, Bing Han, Hui Wang, Wei Wang, Ziyang Ma, Long Zhou, Zengrui Jin, Guanrou Yang, Tianrui Wang, Xu Tan and Xie Chen	4217
<i>Are We Using the Right Benchmark: An Evaluation Framework for Visual Token Compression Methods</i> Chenfei Liao, Wensong Wang, Zichen Wen, Xu Zheng, Yiyu Wang, Haocong He, Yuanhuiyi Lyu, Lutao Jiang, Xin Zou, Yuqian Fu, Bin Ren, Linfeng Zhang and Xuming Hu	4236
<i>Retrieval as Generation: A Unified Framework with Self-Triggered Information Planning</i> Bo Li, Mingda Wang, Gexiang Fang, Shikun Zhang and Wei Ye	4254
<i>Between a Rock and a Hard Place: The Tension Between Ethical Reasoning and Safety Alignment in LLMs</i> Shei Pern Chua, Zhen Leng Thai, Kai Jun Teh, Xiao Li, Qibing Ren and Xiaolin Hu	4275
<i>EventWeave: A Dynamic Framework for Capturing Core and Supporting Events in Dialogue Systems</i> Zhengyi Zhao, Shubo Zhang, Yiming Du, Bin Liang, Baojun Wang, Zhongyang Li, Binyang Li and Kam-Fai Wong	4311

<i>Deliberative Searcher: Improving LLM Reliability via Reinforcement Learning with Constraints</i> Zhenyun Yin, Shujie Wang, Xuhong Wang, Xingjun Ma and Yingchun Wang	4340
<i>Thermometer of Thoughts: Enhancing LLM’s Exploration via Attention Temperature Modulation</i> Zhiyuan Yu, Shijian Xiao, Cam-Tu Nguyen, Zhangyue Yin, Lekai Xing, Wenzhong Li and Sanglu Lu	4355
<i>CodeFlowBench: A Multi-turn, Iterative Benchmark for Complex Code Generation</i> Sizhe Wang, Zhengren Wang, Dongsheng Ma, Yongan Yu, Rui Ling, Zhiyu li, Feiyu Xiong and Wentao Zhang	4369
<i>REFLEX: Self-Refining Explainable Fact-Checking via Verdict-Anchored Style Control</i> Chuyi Kong, Wei Gao, Jing Ma, Hongzhan Lin and Yuxi Sun	4403
<i>Illusions of Confidence? Diagnosing LLM Truthfulness via Neighborhood Consistency</i> Haoming Xu, Ningyuan Zhao, Yunzhi Yao, Weihong Xu, Hongru Wang, Xinle Deng, Shumin Deng, Jeff Z. Pan, Huajun Chen and Ningyu Zhang	4432
<i>Scaling Beyond Context: A Survey of Multimodal Retrieval-Augmented Generation for Document Understanding</i> Sensen Gao, Shanshan Zhao, Xu Jiang, Lunhao Duan, Yong Xien Chng, Qing-Guo Chen, Weihua Luo, Kaifu Zhang, Jia-Wang Bian and Mingming Gong	4458
<i>Beyond Literal Mapping: Benchmarking and Improving Non-Literal Translation Evaluation</i> Yanzhi Tian, Cunxiang Wang, Zeming Liu, Heyan Huang, Wenbo Yu, Dawei Song, Jie Tang and Yuhang Guo	4490
<i>G²RPO-A: Guided Group Relative Policy Optimization with Adaptive Guidance</i> Yongxin Guo, Wenbo Deng, Zhenglin Cheng and Xiaoying Tang	4525
<i>Efficient Low-Resource Language Adaptation via Multi-Source Dynamic Logit Fusion</i> Chen Zhang, Jiuheg Lin, Zhiyuan Liao and Yansong Feng	4540
<i>Open Schrödinger’s Closed Box: Identifying Retrieval Augmented Generation in API-Accessible Large Language Model Services</i> Yukun Jiang, Xinyue Shen, Michael Backes, Zheng Li and Yang Zhang	4558
<i>AFT-Tab: Adversarial Fine-Tuning for Tabular Data Synthesis with Long Text Columns</i> Yuhao Zhang, Liang Yan, Shaoming Duan, Xinyu Zha, Jinhang Su, Peiyi Han and Chuanyi Liu	4581
<i>Revisiting a Pain in the Neck: A Semantic Reasoning Benchmark for Language Models</i> Yang Liu, Hongming Li, Melissa Xiaohui Qin, Chao Huang and Qiankun Liu	4595
<i>Timely Machine: Awareness of Time Makes Test-Time Scaling Agentic</i> Yichuan Ma, Linyang Li, Yongkang Chen, Peiji Li, Xiaozhe Li, Qipeng Guo, Dahua Lin and Kai Chen	4619
<i>Multi-Granularity Semantic Revision for Large Language Model Distillation</i> Xiaoyu Liu, Yun Zhang, Wei Li, Simiao Li, Xudong Huang, Hanting Chen, Yehui Tang, Jie Hu, Zhiwei Xiong and Yunhe Wang	4637
<i>RADO: Reasoning Audit-Driven Optimization for Rigorous Reasoning in High-Stakes Domains</i> Zhijie Tan, Xu Chu, Guanyu Wang, Ziyu Li, Weiping Li and Tong Mo	4659
<i>Instruction Data Selection via Answer Divergence</i> Bo Li, Mingda Wang, Shikun Zhang and Wei Ye	4684

<i>Interpretable Safety Alignment via SAE-Constructed Low-Rank Subspace Adaptation</i>	
Dianyun Wang, Qingsen Ma, Yuhu Shang, Zhifeng Lu, Zhenbo Xu, Lechen Ning, Huijia Wu and Zhaofeng He	4703
<i>D²Plan: Dual-Agent Dynamic Global Planning for Complex Retrieval-Augmented Reasoning</i>	
Kangcheng Luo, Tinglang Wu and Yansong Feng	4722
<i>Instructions are all you need: Self-supervised Reinforcement Learning for Instruction Following</i>	
Qingyu Ren, Qianyu He, Powei Chang, Jie Zeng, Zeye Sun, Fei Yu, Jiaqing Liang and Yanghua Xiao	4755
<i>Analytical FFN-to-MoE Restructuring via Activation Pattern Analysis</i>	
Zehua Pei, Hui-Ling Zhen, Lancheng Zou, Xianzhi Yu, Wulong Liu, Sinno Jialin Pan, Mingxuan Yuan and Bei Yu	4777
<i>Towards Order Fairness: Mitigating LLMs Order Sensitivity through Dual Group Advantage Optimization</i>	
Xu Chu, Guanyu Wang, Zhijie Tan, Xinrong Chen, Ziyu Li, Tong Mo and Weiping Li	4790
<i>Revitalizing Black-Box Interpretability: Actionable Interpretability for LLMs via Proxy Models</i>	
Junhao Liu, Haonan Yu, Zhenyu Yan and Xin Zhang	4806
<i>MedMCP-Calc: Benchmarking LLMs for Realistic Medical Calculator Scenarios via MCP Integration</i>	
Yakun Zhu, Yutong Huang, Shengqian Qin, Zhongzhen Huang, Shaoting Zhang and Xiaofan Zhang	4845
<i>Uncovering Temporal Framing in the News</i>	
Tarek Mahmoud, Veronika Solopova, Premtim Sahitaj, Ariana Sahitaj, Max Upravitelev, Mervat Abassy, Hana Fatima Shaikh, Neda Foroutan, Vera Schmitt and Preslav Nakov	4874
<i>Simple-VGC: Enhancing Visual Grounding in Multimodal Reasoning via Adaptive Tool Composition</i>	
Ye Wang, Qianglong Chen, Siyuan Wang, Zejun Li, Shijie Guo, Zhirui Zhang and Zhongyu Wei	4903
<i>Enhancing Lexical Relation Mining with Structured Sememe Knowledge</i>	
Hansi Wang, Qiliang Liang, Yue Wang and Yang Liu	4930
<i>Look Less, Reason More: Rollout-Guided Adaptive Pixel-Space Reasoning</i>	
Xuchen Li, Xuzhao Li, Jiahui Gao, Renjie Pi, Shiyu Hu and Wentao Zhang	4948
<i>Conditional Memory via Scalable Lookup: A New Axis of Sparsity for Large Language Models</i>	
Xin Cheng, Wangding Zeng, Damai Dai, Qinyu Chen, Bingxuan Wang, Zhenda Xie, Kezhao Huang, Xingkai Yu, Zhewen Hao, Han Zhang, Yu-Kun Li, Huishuai Zhang, Dongyan Zhao and Wenfeng Liang	4968
<i>PII-Bench: Evaluating Query-Aware Privacy Protection Systems</i>	
Hao Shen, Zhouhong Gu, Haokai Hong, Weili Han and Hongfeng Chai	4991
<i>GenProve: Learning to Generate Text with Fine-Grained Provenance</i>	
Jingxuan Wei, Xingyue Wang, Yanghaoyu Liao, Jie Dong, Yuchen Liu, Caijun Jia, Bihui Yu and Junnan Zhu	5027
<i>Rethinking Meeting Effectiveness: A Benchmark and Framework for Temporal Fine-grained Automatic Meeting Effectiveness Evaluation</i>	
Yihang Li and Chenhui Chu	5049

<i>AgentOCR: Reimagining Agent History via Optical Self-Compression</i>	
Lang Feng, Fuchao Yang, Feng Chen, Xin Cheng, Haiyang Xu, Zhenglin Wan, Ming Yan and Bo An.....	5067
<i>MCP-Flow: Facilitating LLM Agents to Master Real-World, Diverse and Scaling MCP Tools</i>	
WenHao Wang, Peizhi Niu, Zhao Xu, Zhaoyu Chen, Jian Du, Yaxin Du, Xianghe Pang, Keduan Huang, Yanfeng Wang, Qiang Yan and Siheng Chen.....	5087
<i>Everyone is unique: Towards Behaviorally Heterogeneous Negotiation Dialogue Systems for Debt Collection</i>	
Yuhang Yang, Kai Tang, Chao Ye, Haobo Wang, Qiqi Luo, Jin Guang Zheng and Zhixin Zhang	5117
<i>Beyond the Crowd: LLM-Augmented Community Notes for Governing Health Misinformation</i>	
Jiaying Wu, Zihang Fu, Haonan Wang, Fanxiao Li, Jiafeng Guo, Preslav Nakov and Min-Yen Kan	5152
<i>BaseCal: Unsupervised Confidence Calibration via Base Model Signals</i>	
Hexiang Tan, Wanli Yang, Junwei Zhang, Xin Chen, Rui Tang, Du Su, Jingang Wang, Yuanzhuo Wang, Fei Sun and Xueqi Cheng	5172
<i>Beyond the Final Actor: Modeling the Dual Roles of Creator and Editor for Fine-Grained LLM-Generated Text Detection</i>	
Yang Li, Qiang Sheng, Zhengjia Wang, Yehan Yang, Danding Wang and Juan Cao	5188
<i>PILOT: Planning via Internalized Latent Optimization Trajectories for Large Language Models</i>	
Haoyu Zheng, Yun Zhu, Yuqian Yuan, Bo Yuan, Wenqiao Zhang, Siliang Tang and Jun Xiao	5204
<i>Miner: Mining Intrinsic Mastery for Data-Efficient RL in Large Reasoning Models</i>	
Shuyang Jiang, Yuhao Wang, Ya Zhang, Yanfeng Wang and Yu Wang	5223
<i>CEDAR: A Chinese Evaluation Dataset for Computational Argumentation</i>	
Tian Lan, Jiang Li, Rong Yan, Feilong Bao, Weihua Wang, Guanglai Gao and Xiangdong Su	5247
<i>MSEarth: A Multimodal Benchmark for Earth Science Phenomenon Discovery with MLLMs</i>	
Xiangyu Zhao, Wanghan Xu, Bo Liu, Yuhao Zhou, Fenghua Ling, Ben Fei, Xiaoyu Yue, Lei Bai, Wenlong Zhang and Xiao-Ming Wu.....	5270
<i>Reasoning Structure Matters for Safety Alignment of Reasoning Models</i>	
Yeonjun In, Wonjoong Kim, Sangwu Park and Chanyoung Park	5302
<i>Investigating the Representation of Backchannels and Fillers in Fine-tuned Language Models</i>	
Yu Wang, Leyi Lao, Langchu Huang, Gabriel Skantze, Yang Xu and Hendrik Buschmeier..	5319
<i>Retrievals Can Be Detrimental: Unveiling the Backdoor Vulnerability of Retrieval-Augmented Diffusion Models</i>	
Hao Fang, Xiaohang Sui, Hongyao Yu, Kuofeng Gao, Jiawei Kong, Sijin Yu, Bin Chen and Shu-Tao Xia	5349
<i>A Theoretically Grounded Approach to Summarizing Conversation Dynamics for Forecasting the Derailment of Online Conversations</i>	
Yingxue Fu and Anaïs Ollagnier.....	5368
<i>DPDV: Dual-Pathway and Dual-View Representation Learning for Bridging Information Asymmetry in Text-Video Retrieval</i>	
Zequn Xie, Xin Liu, Fangming Feng, Boyun Zhang and Tao Jin	5385

<i>Who Wrote This Line? Evaluating the Detection of LLM-Generated Classical Chinese Poetry</i>	
Jiang Li, Tian Lan, Shanshan Wang, Zdongxing, Dianqing Lin, Guanglai Gao, Derek F. Wong and Xiangdong Su	5395
<i>Representation Interventions Enable Lifelong Knowledge Memory Control in LLMs</i>	
Xuyuan Liu, Shengyu Chen, Xinshuai Dong, Yanchi Liu, Xujiang Zhao, Haoyu Wang, Yujun Yan, Haifeng Chen and Zhengzhang Chen	5414
<i>SLoRA: Balancing Plasticity and Forgetting in Large Language Models for Continual Learning</i>	
Lina Yang, Yusheng Liao, Yanfeng Wang and Yu Wang	5437
<i>G-Cap: A Game Character Caption Generator</i>	
Yang Yang, Feng Hu, Haiming Zhang, XU Cheng, Gui Zheng, Liang Yao and Wenqi Ren ..	5455
<i>Common to Whom? Regional Cultural Commonsense and LLM Bias in India</i>	
Sangmitra Madhusudan, Trush Shashank More, Steph Buongiorno, Renata Dividino, Jad Kabbara and Ali Emami	5474
<i>The Stackelberg Speaker: Optimizing Persuasive Communication in Social Deduction Games</i>	
Zhang Zheng, Deheng Ye, Peilin Zhao and Hao Wang	5520
<i>GuideDog: A Real-World Egocentric Multimodal Dataset for Blind and Low-Vision Accessibility-Aware Guidance</i>	
Junhyeok Kim, Jaewoo Park, Junhee Park, Sangeyl Lee, Jiwan Chung, Jisung Kim, Ji Hoon Joung and Youngjae Yu	5545
<i>LogicPoison: Logical Attacks on Graph Retrieval-Augmented Generation</i>	
Yilin Xiao, Jin Chen, Qinggang Zhang, Yujing Zhang, Chuang Zhou, Longhao Yang, Lingfei Ren, Xin Yang and Xiao Huang	5575
<i>Par-ITA: Benchmarking Seq2Seq and LLMs on a Human-Supervised Parallel Corpus for Italian Hyperpartisan Neutralization</i>	
Michele Joshua Maggini, Søren Fomsgaard, Michele Maestroni, Gaël Dias and Pablo Gamallo	5592
<i>Incentivizing Parametric Knowledge via Reinforcement Learning with Verifiable Rewards for Cross-Cultural Entity Translation</i>	
Jiang Zhou, Xiaohu Zhao, Xinwei Wu, Tianyu Dong, Hao Wang, Yangyang Liu, Heng Liu, Linlong Xu, Longyue Wang, Weihua Luo and Deyi Xiong	5616
<i>Writing-RL: Advancing Long-form Writing via Adaptive Curriculum Reinforcement Learning</i>	
Xuanyu Lei, Chenliang Li, Yuning Wu, Kaiming Liu, Weizhou Shen, Peng Li, Ming Yan, Fei Huang, Ya-Qin Zhang and Yang Liu	5639
<i>Spec-o3: A Tool-Augmented Vision-Language Agent for Rare Celestial Object Candidate Vetting via Automated Spectral Inspection</i>	
Minghui Jia, Qichao Zhang, Ali Luo, Linjing Li, Shuo Ye, Hailing Lu, Wen Hou and Dongbin Zhao	5662
<i>ROSE: An Intent-Centered Evaluation Metric for NL2SQL</i>	
Wenqi Pei, Shizheng Hou, Boyan Li, Chen Han, Zhichao Shi and Yuyu Luo	5682
<i>Compete to Complete: Co-opetition Adversarial Learning for Retrieval-Augmented Generation</i>	
Xin Liu, Yu-An Liu, Ruqing Zhang, Yixing Fan, Lixin Su, Jiafeng Guo and Xueqi Cheng ..	5710

<i>Demystifying Uncertainty in LLMs: Active Calibration between Concepts and Human Evaluations</i> Pengqi Li, Lizhong Ding, Zehao Zhou, Chunhui Zhang, Jiarun Fu, Hao Li, Ye Yuan and Guoren Wang	5726
<i>LLM-VA: Resolving the Jailbreak-Overrefusal Trade-off via Vector Alignment</i> Haonan Zhang, Dongxia Wang, Yi Liu, Kexin Chen and Wenhai Wang	5760
<i>TiKMIX: Efficient Semi-Dynamic Data Mixture via Data Influence for LLM Pre-training</i> Yifan Wang, Binbinliu, Fengze Liu, Yuanfan Guo, Jiayao Deng, Xuecheng Wu, Weidong Zhou, Xiaohuan Zhou and Taifeng Wang	5777
<i>Revisit What You See: Revealing Visual Semantics in Vision Tokens to Guide LVLM Decoding</i> Beomsik Cho and Jaehyung Kim	5794
<i>FlexGuard: Continuous Risk Scoring for Strictness-Adaptive LLM Content Moderation</i> Zhihao Ding, Jinming Li, Ze Lu and Jieming Shi	5825
<i>Modeling Human-Like Cognition for Stance Detection: Integrating Intuitive Judgment and Analytical Reasoning</i> Zhaodan Zhang, Jin Zhang, Jiafeng Guo and Xueqi Cheng	5852
<i>FinSight: Towards Real-World Financial Deep Research</i> Jiajie Jin, Yuyao Zhang, Yimeng Xu, Hongjin Qian, Yutao Zhu and Zhicheng Dou	5868
<i>Dual-Cluster Memory Agent: Resolving Multi-Paradigm Ambiguity in Optimization Problem Solving</i> Xinyu Zhang, Yuchen Wan, Boxuan Zhang, Zesheng Yang, Lingling Zhang, Bifan Wei and Jun Liu	5895
<i>AfriqueLLM: How Data Mixing and Model Architecture Impact Continued Pre-training for African Languages</i> Hao Yu, Tianyi Xu, Michael A. Hedderich, Wassim Hamidouche, Syed Waqas Zamir and David Ifeoluwa Adelani	5909
<i>What Does LLM Refinement Actually Improve? A Systematic Study on Document-Level Literary Translation</i> Shaomu Tan, Dawei Zhu, Ke Tran, Michael Denkowski, Sony Trenous, Leonardo F. R. Ribeiro, Bill Byrne and Felix Hieber	5929
<i>OctoBench: Benchmarking Scaffold-Aware Instruction Following in Repository-Grounded Agentic Coding</i> Deming Ding, Shichun Liu, Enhui Yang, Jiahang Lin, Ziyang Chen, Shihan Dou, Honglin Guo, Weiyu Cheng, Pengyu Zhao, Chengjun Xiao, Qunhong Zeng, Qi Zhang, Xuanjing Huang, Qidi Xu and Tao Gui	5958
<i>Lying with Truths: Open-Channel Multi-Agent Collusion for Belief Manipulation via Generative Montage</i> Jinwei Hu, Xinmiao Huang, Youcheng Sun, Yi Dong and Xiaowei Huang	5979
<i>EXCEEDS: Extracting Complex Events via Nugget-based Grid Modeling in Scientific Domain</i> Yi-Fan Lu, Xian-Ling Mao, Bo Wang, Xiao Liu and Heyan Huang	5997
<i>Heterogeneous Adaptive Policy Optimization: Tailoring Optimization to Every Token's Nature</i> Zheng Liu, Mengjie Liu, Siwei Wen, Mengzhang Cai, Bin Cui, Conghui He and Wentao Zhang	6023

<i>Mirroring Users: Towards Building Preference-aligned User Simulator with User Feedback in Recommendation</i>	
Tianjun Wei, Huizhong Guo, Yingpeng Du, Zhu Sun, Huang Chen, Dongxia Wang and Jie Zhang	
	6046
<i>Editing the Moving World: Model Editing for Video LLMs</i>	
Qian Zhang, Xinye Li, Xiaokai Wu, Junhao Xu, Zhanyue Qin, Qingbin Liu, Junxian Cai, Xi Chen, Bolin Zhang, Zhiying Tu, Dianhui Chu, Xiaoyan Yu and Dianbo Sui	6072
<i>Bridging the Sensory Gap: Visual Injection for Taxonomy Completion</i>	
Yuhang Niu, Hongyuan Xu, Ciyi Liu, Bofan Wei, Jiaqi Ye, Yanlong Wen and Xiaojie Yuan	6092
<i>OCP: Outlier-Centric Probing for Dynamic Structured Pruning of LLMs</i>	
Yang Ji and Ying Sun	6108
<i>Structured Episodic Event Memory</i>	
Zhengxuan Lu, Dongfang Li, Yukun Shi, Beilun Wang, Longyue Wang and Baotian Hu	6125
<i>MobileWorld: Benchmarking Autonomous Mobile Agents in Agent-User Interactive and MCP-Augmented Environments</i>	
Quyu Kong, Xu Zhang, Zhenyu Yang, Nolan Gao, Chen Liu, Panrong Tong, Chenglin Cai, Hanzhang Zhou, Jianan Zhang, Liangyu Chen, Zhidan Liu, Steven Hoi and Yue Wang	6142
<i>Measuring Human Contribution in AI-Assisted Content Generation</i>	
Yueqi Xie, Tao Qi, Jingwei Yi, Xiyuan Yang, Ryan Whalen, Junming Huang, Qian Ding, Yu Xie, Xing Xie and Fangzhao Wu	6168
<i>AutoTaskEval: Towards Domain-Specific and Fine-Grained Evaluation for LLMs</i>	
Qingqing Lyu, Linjuan Wu, Yongliang Shen, Hengwei Liu, Hao Li, Shengpei Jiang, Yin Zhang and Weiming Lu	6191
<i>ACE-Router: Generalizing History-Aware Routing from MCP Tools to the Agent Web</i>	
Zhiyuan Yao, Zishan Xu, Yifu Guo, Zhiguang Han, Cheng Yang, Shuo Zhang, Weinan Zhang, Xingshan Zeng and Weiwen Liu	6224
<i>CiteGuard: Faithful Citation Attribution for LLMs via Retrieval-Augmented Validation</i>	
Yee Man Choi, Xuehang Guo, Yi R. Fung and Qingyun Wang	6241
<i>From Insight to Action: A Novel Framework for Interpretability-Guided Data Selection in Large Language Models</i>	
Ling Shi, Xinwei Wu, Xiaohu Zhao, Hao Wang, Heng Liu, Yangyang Liu, Linlong Xu, Longyue Wang, Deyi Xiong and Weihua Luo	6258
<i>Unlocking the Potential of Diffusion Language Models through Template Infilling</i>	
Junhoo Lee, Seungyeon Kim and Nojun Kwak	6273
<i>UniSpec: Training-Free Speculative Decoding for Robust LLM Acceleration Across Languages and Hardware</i>	
Truong Dinh Do, Nguyen-Khang Le and Le-Minh Nguyen	6288
<i>Logical Consistency as a Bridge: Improving LLM Hallucination Detection via Label Constraint Modeling between Responses and Self-Judgments</i>	
Hao Mi, Qiang Sheng, Shaofei Wang, Beizhe Hu, Yifan Sun, Zhengjia Wang, Hengqi Zeng, Yang Li, Danding Wang and Juan Cao	6311

<i>UNIKIE-BENCH: Benchmarking Large Multimodal Models for Key Information Extraction in Visual Documents</i>	
Yifan Ji, Zhipeng Xu, Zhenghao Liu, Zulong Chen, Qian Zhang, Zhibo Yang, Junyang Lin, Yu Gu, Ge Yu and Maosong Sun	6331
<i>FS-Researcher: Test-Time Scaling for Long-Horizon Research Tasks with File-System-Based Agents</i>	
Chiwei Zhu, Benfeng Xu, Mingxuan Du, Shaohan Wang, Xiaorui Wang, Zhendong Mao and Yongdong Zhang	6353
<i>Annotating Dimensions of Social Perception in Text: A Sentence-Level Dataset of Warmth and Competence</i>	
Mutaz Ayesah, Saif M. Mohammad and Nedjma Ousidhoum	6374
<i>SCOUT: Selective Coupling via Optimal Unbalanced Transport for Interpretable Text Classification</i>	
Junhao Jia, Hanwen Zheng, Yueyi Wu, Huangwei Chen, Haishuai Wang, Jiajun Bu and Lei Wu	6413
<i>LLMs Enable Bag-of-Texts Representations for Short-Text Clustering</i>	
I-Fan Lin, Faegheh Hasibi and Suzan Verberne	6432
<i>TriEx: A Game-based Tri-View Framework for Explaining Internal Reasoning in Multi-Agent LLMs</i>	
Ziyi Wang, Chen Zhang, Wenjun Peng, Qi Wu and Xinyu Wang	6448
<i>What’s Left Unsaid? Detecting and Correcting Misleading Omissions in Multimodal News Previews</i>	
Fanxiao Li, Jiaying Wu, Tingchao Fu, Dayang Li, Herun Wan, Wei Zhou and Min-Yen Kan	6480
<i>ContextLens: Modeling Imperfect Privacy and Safety Context for Legal Compliance</i>	
Haoran Li, Yulin Chen, Huihao Jing, Wenbin Hu, Tsz Ho Li, Chanhou Lou, Hong Ting Tsang, Sirui Han and Yangqiu Song	6503
<i>StealthGraph: Exposing Domain-Specific Risks in LLMs through Knowledge-Graph-Guided Harmful Prompt Generation</i>	
Huawei Zheng, Xinqi Jiang, Sen Yang, Shouling Ji, Yingcai Wu and Dazhen Deng	6519
<i>Efficient Multi-Agent System Training with Data Influence-Oriented Tree Search</i>	
Wentao Shi, Zichun Yu, Fuli Feng, Xiangnan He and Chenyan Xiong	6540
<i>Integrating Data Validation with Large Language Models for Regulation-Guided Tabular Anomaly Detection</i>	
Haoliang Huang, Zihuang Cai, Zhuo Tang, Yifan Liu, Chen Tian, Kenli Li and Changjian Chen	6559
<i>Efficient Self-Evaluation for Diffusion Language Models via Sequence Regeneration</i>	
Linhao Zhong, Linyu Wu, Wen Wang, Yuling Xi, Chenchen Jing, Jiaheng Zhang, Hao Chen and Chunhua Shen	6582
<i>DeReA: Improving Idiom Translation with Detect-Retrieve-Arbitrate Reasoning</i>	
Rongqing Jiang, Xuebo Liu, Shengxin Liu, Yutong Wang, Min Zhang, Shimin Tao, Daimeng Wei and Min Zhang	6603
<i>Faithful Serum: Mitigating the Faithfulness Gap in Textual Explanations of LLM Decisions via Attribution Guidance</i>	
Bar Alon, Itamar Zimmerman and Lior Wolf	6622
<i>Visually-Guided Policy Optimization for Multimodal Reasoning</i>	
Zengbin Wang, Feng Xiong, Liang Lin, Xuecai Hu, Yong Wang, Yanlin Wang, Man Zhang and Xiangxiang Chu	6646

<i>Don't Be Misled by Style: A Style-Adaptive Reranker for Capturing Effective Knowledge in Retrieval-Augmented Generation</i>	
Ruwen Zhang, Bo Liu, Zhang Sheng Xiang, Yida Chen, Hantao Zhao, Ding Ding, Jiahui Jin and Jiuxin Cao	6665
<i>Text2Tabular – Reconstructing Tabular Research Data from Scientific Publications</i>	
Jonas Gottal and Florian Matthes	6682
<i>From Aha Moments to Controllable Thinking: Toward Meta-Cognitive Reasoning in LLMs via Decoupled Reasoning and Control</i>	
Rui Ha, Rui Pu, Chaozhuo Li, Li Sun and Sen Su	6698
<i>Calibrating Inference Time Alignment with Sequence-level Risk Accumulation</i>	
Shanwen Tan, Ziyang Dong, Wei Ju, Yiwei Fu, Hao Wu, Kun Wang, Yifan Wang and Ziyue Qiao	6711
<i>Diff4TST: Masked Diffusion Language Model for Text Style Transfer</i>	
Xinchen Ma, Gaole He, Yunshi Lan and Weining Qian	6736
<i>In-depth Research Impact Summarization through Fine-Grained Temporal Citation Analysis</i>	
Hiba Arnaout, Noy Sternlicht, Tom Hope and Iryna Gurevych	6749
<i>QuantileMark: A Message-Symmetric Multi-bit Watermark for LLMs</i>	
Junlin Zhu, Baizhou Huang and Xiaojun Wan	6790
<i>Standard-to-Dialect Transfer Trends Differ across Text and Speech: A Case Study on Intent and Topic Classification in German Dialects</i>	
Verena Blaschke, Miriam Winkler and Barbara Plank	6807
<i>Don't Click That: Teaching Web Agents to Resist Deceptive Interfaces</i>	
Yilin Zhang, Yingkai Hua, Chunyu Wei, Xin Wang and Yueguo Chen	6830
<i>Empirical Analysis of Decoding Biases in Masked Diffusion Models</i>	
Pengcheng Huang, Tianming Liu, Zhenghao Liu, Yukun Yan, Shuo Wang, Tong Xiao, Zulong Chen and Maosong Sun	6853
<i>LAFACT: Attribution-based Localization and Focused Sequential Analysis of Fact-Critical Tokens for Hallucination Detection</i>	
Xin Wang, Jiahao Li, Licheng Zhang and Zhendong Mao	6877
<i>DPC: Training-Free Text-to-SQL Candidate Selection via Dual-Paradigm Consistency</i>	
Boyan Li, Ou Ocean Kun Hei, Yue Yu and Yuyu Luo	6897
<i>Metaphor Reasoning is Meta-reasoning</i>	
Qianyu He, Junting Lu, Yikai Zhang, Siyu Yuan, Xiaojun Meng, Jiansheng Wei, Jiaqing Liang and Yanghua Xiao	6914
<i>Beyond Hard Masks: Progressive Token Evolution for Diffusion Language Models</i>	
Linhao Zhong, Linyu Wu, Bozhen Fang, Tianjian Feng, Chenchen Jing, Wen Wang, Jiaheng Zhang, Hao Chen and Chunhua Shen	6936
<i>MAS-Bench: A Unified Benchmark for Shortcut-Augmented Hybrid Mobile GUI Agents</i>	
Pengxiang Zhao, Guangyi Liu, Yaozhen Liang, Weiqing He, Zhengxi Lu, WenHao Wang, Yuehao Huang, Yuxiang Chai, Zhaolu Kang, Yaxuan Guo, Hao Wang, Kexin Zhang, Liang Liu and Yong Liu	6960

<i>HoWToBench: Holistic Evaluation for LLM’s Capability in Human-level Writing using Tree of Writing</i> Andrew Zhuoer Feng, Cunxiang Wang, Yu Luo, Lin Fan, Irene Zhou, Zikang Wang, Xiaotao Gu, Jie Tang, Hongning Wang and Minlie Huang	6986
<i>Hierarchical Reinforcement Learning with Augmented Step-Level Transitions for LLM Agents</i> Shuai Zhen, Yanhua Yu, Ruopei Guo, Nan Cheng and Yang Deng	7035
<i>Two Streams, One Sarcasm: Orthogonal Expert Tuning for Holistic Multimodal Sarcasm Understanding</i> Diandian Guo, Cong Cao, Fangfang Yuan, Pin Xu, Cheng Hu, Zhicheng Zhang, Yu Liu and Yanbing Liu	7054
<i>SeLaR: Selective Latent Reasoning in Large Language Models</i> Renyu Fu and Guibo Luo	7073
<i>STRIDE-ED: A Strategy-Grounded Stepwise Reasoning Framework for Empathetic Dialogue Systems</i> Hongru Ji, Yuyin Fan, Meng Zhao, Xianghua Li, Lianwei Wu and Chao Gao	7088
<i>GRAPHIA: Harnessing Social Graph Data to Enhance LLM-Based Social Simulation</i> Jiarui Ji, Zehua Zhang, Zhewei Wei, Bin Tong, Guan Wang and Bo Zheng	7103
<i>SharedRequest: Privacy-Preserving Model-Agnostic Inference for Large Language Models</i> Peihua Mai, Xuanrong Gao, Youlong Ding, Xianglong Du, Wei Liu and Yan Pang	7129
<i>Efficient Learned Data Compression via Dual-Stream Feature Decoupling</i> Huidong Ma, Xinyan Shi, Sun Hui, Xiaofei Yue, Xiaoguang Liu, Gang Wang and Wentong Cai 7151	
<i>Planning-Guided Tutoring with Assessment-Driven Memory for Pedagogical LLM Tutors</i> Zechen Li, Qiannan Zhu, Mei Wang, Jia Li and Hua Huang	7165
<i>MixKVQ: Query-Aware Mixed-Precision KV Cache Quantization for Long-Context Reasoning</i> Tao Zhang, Ziqian Zeng, Hao Peng, Huiping Zhuang and Cen Chen	7189
<i>When KV Cache Reuse Fails in Multi-Agent Systems: Cross-Candidate Interaction is Crucial for LLM Judges</i> Sichu Liang, Zhenglin Wang, Chujiajia, Pengfei Xia, Hui Zang and Deyu Zhou	7205
<i>Reducing Token Redundancy in LVLMS: A Systematic Review of Token Pruning Methods</i> Hanzhang Yuan, Mengxuan Hu, Wenhao Zhang, Tianlong Wang, Zhongliang Zhou, Jiasen Lu and Sheng Li	7231
<i>STEM: Structure-Tracing Evidence Mining for Knowledge Graphs-Driven Retrieval-Augmented Generation</i> Peng Yu, En Xu, Bin Chen, Haibiao Chen and Yinfei Xu	7252
<i>Evo-Attacker: Memory-Augmented Reinforcement Learning for Long-Horizon Tool Attacks on LLM-MAS</i> Bingyu Yan, Xiaoming Zhang, JinYu Hou, Chaozhuo Li, Ziyi Zhou, Yiming Hei and Litian Zhang 7286	
<i>Beyond Timestamps: Bridging Forward and Backward Reasoning in Temporal Numerical and Relational Understanding</i> Xinying Qian, Ying Zhang, Xuhui Sui, Yu Zhao, Baohang Zhou and Jeff Z. Pan	7301
<i>Reasoning Over Space: Enabling Geographic Reasoning for LLM-Based Generative Next POI Recommendation</i> Dongyi Lv, Qiuyu Ding, Heng-Da Xu, Zhaoxu Sun, Zhi Wang, Feng Xiong and Mu Xu	7322

<i>ET-Agent: Incentivizing Effective Tool-Integrated Reasoning Agent via Behavior Calibration</i> Yifei Chen, Guanting Dong and Zhicheng Dou	7337
<i>Spectral Disentanglement: Rank-Aware Task Adaptation for Rehearsal-free Continual Learning in LLMs</i> Huanxuan Liao, Shizhu He, Yupu Hao, Yequan Wang, Wenhao Teng, Xiangwen Liao, Jun Zhao and Kang Liu	7360
<i>DeepPlanning: Benchmarking Long-Horizon Agentic Planning with Verifiable Constraints</i> Yinger Zhang, Shutong Jiang, Renhao Li, Jianhong Tu, Yang Su, Lianghao Deng, Xudong Guo, ChenXu Lv and Junyang Lin	7377
<i>Arg-LLaDA: Argument Summarization via Large Language Diffusion Models and Sufficiency-Aware Refinement</i> Hao Li, Yizheng Sun, Viktor Schlegel, Kailai Yang, Riza Batista-Navarro and Goran Nenadic	7408
<i>AgencyBench: Benchmarking the Frontiers of Autonomous Agents in 1M-Token Real-World Contexts</i> Keyu Li, Junhao Shi, Yang Xiao, Mohan Jiang, Jie Sun, Yunze Wu, Dayuan Fu, Shijie Xia, Xiaojie Cai, Tianze Xu, Weiye Si, Wenjie Li, Dequan Wang and Pengfei Liu	7422
<i>All Languages Matter: Understanding and Mitigating Language Bias in Multilingual RAG</i> Dan Wang, Guozhao Mo, Yafei Shi, Cheng Zhang, Bo Zheng, Boxi Cao, Xuanang Chen, Yaojie Lu, Hongyu Lin, Ben He, Xianpei Han and Le Sun	7441
<i>Beyond Accuracy: Unveiling Inefficiency Patterns in Tool-Integrated Reasoning</i> Qisheng Su, Shiting Huang, Zhen Fang, Ziyang Chen, Zehui Chen and Feng Zhao	7456
<i>Safety-Utility Conflicts Are Not Global: Surgical Alignment via Head-Level Diagnosis</i> Wang Cai, Yilin Wen, Jinchang Hou, Du Su, Guoqiu Wang, Zhonghou Lv, Chenfu Bao and Yunfang Wu	7478
<i>VideoPro: Adaptive Program Reasoning for Long Video Understanding</i> Chenglin Li, Feng Han, Yikun Wang, Ruilin Li, Shuai Dong, Haowen Hou, Haitao Li, Qianglong Chen, Feng Tao, Jingqi Tong, Yin Zhang and Jiaqi Wang	7497
<i>Parity-Aware Byte-Pair Encoding: Improving Cross-lingual Fairness in Tokenization</i> Negar Foroutan, Clara Meister, Debjit Paul, Joel Niklaus, Sina Ahmadi, Antoine Bosselut and Rico Sennrich	7514
<i>Boundary-Guided Policy Optimization for Memory-efficient RL of Diffusion Large Language Models</i> Nianyi Lin, Jiajie Zhang, Lei Hou and Juanzi Li	7539
<i>ChartVerse: Scaling Chart Reasoning via Reliable Programmatic Synthesis from Scratch</i> Zheng Liu, Honglin Lin, Xiaoyang Wang, Xin Gao, Yu Li, Mengzhang Cai, Yun Zhu, Zhanping Zhong, Qizhi Pei, Zhuoshi Pan, Xiaoran Shang, Conghui He, Bin Cui, Wentao Zhang and Lijun Wu	7551
<i>NeoAMT: Neologism-Aware Agentic Machine Translation with Reinforcement Learning</i> Zhongtao Miao, Kaiyan Zhao, Masaaki Nagata and Yoshimasa Tsuruoka	7578
<i>arXiv2Table: Toward Realistic Benchmarking and Evaluation for LLM-Based Literature-Review Table Generation</i> Weiqi Wang, Jiefu Ou, Yangqiu Song, Benjamin Van Durme and Daniel Khashabi	7602

<i>Responsible Evaluation of AI for Mental Health</i>	
Hiba Arnaout, Anmol Goel, H. Andrew Schwartz, Steffen T. Eberhardt, Dana Atzil-Slonim, Gavin Doherty, Brian Schwartz, Wolfgang Lutz, Tim Althoff, Munmun De Choudhury, Hamidreza Jamalabadi, Raj Sanjay Shah, Flor Miriam Plaza-del-Arco, Dirk Hovy, Maria Liakata and Iryna Gurevych	7625
<i>JPU: Bridging Jailbreak Defense and Unlearning via On-Policy Path Rectification</i>	
Xi Wang, Songlei Jian, Shasha Li, Xiaopeng Li, Zhaoye Li, Bin Ji, Baosheng Wang and Jie Yu	7661
<i>SpeechLLM-as-Judges: Towards General and Interpretable Speech Quality Evaluation</i>	
Hui Wang, Jinghua Zhao, Yifan Yang, Shujie Liu, Junyang Chen, Yanzhe Zhang, Shiwan Zhao, Jinyu Li, Jiaming Zhou, Haoqin Sun, Yan Lu and Yong Qin	7675
<i>LePREC: Reasoning as Classification over Structured Factors for Assessing Relevance of Legal Issues</i>	
Fanyu Wang, Xiaoxi Kang, Paul Burgess, Aashish Srivastava, Chetan Arora, Adnan Trakic, Lay-Ki Soon, Md Khalid Hossain and Lizhen Qu	7701
<i>S2O: Early Stopping for Sparse Attention via Online Permutation</i>	
Yu Zhang, Songwei Liu, Chenqian Yan, Linsheng, Beichen Ning, Fangmin Chen and Xing Wang	7737
<i>DART: Disambiguation-Aware Reasoning for Video-guided Machine Translation</i>	
Boyu Guan, Chuang Han, Yang Zhao and Chengqing Zong	7752
<i>HCSpec: Two-Tier Horizontal Cascade Speculative Decoding for High-Efficiency Large Language Model Inference</i>	
Yizhou Zhang, Siming Chen, Hao Ye and Erhu Feng	7773
<i>Revisiting the Reliability of Language Models in Instruction-Following</i>	
Jianshuo Dong, Yutong Zhang, Liu Yan, Zhenyu Zhong, Tao Wei, Chao Zhang and Han Qiu	7784
<i>EfficientLLM: Unified Pruning-Aware Pretraining for Auto-Designed Compact Language Models</i>	
Xingrun Xing, Zheng Liu, Shitao Xiao, Boyan Gao, Yiming Liang, Haokun Lin, Xianlin Zeng, Guoqi Li and Jiajun Zhang	7813
<i>SPPO: Sequence-Level PPO for Long-Horizon Reasoning Tasks</i>	
Tianyi Wang, Yixia Li, Long Li, Yibiao Chen, Shaohan Huang, Yun Chen, Peng Li, Yang Liu and Guanhua Chen	7831
<i>Enhancing the Transferability of Jailbreak Attacks on Large Language Models via Exploiting Reparameterization Invariance</i>	
Ao Wang, Xinghao Yang, Yongshun Gong, Wei Liu, Bao-di Liu and Weifeng Liu	7854
<i>GIFT: Guided Fine-Tuning and Transfer for Enhancing Instruction-Tuned Language Models</i>	
Zhiwen Ruan, Yichao Du, Jianjie Zheng, Longyue Wang, Yun Chen, Peng Li, Jinsong Su, Yang Liu and Guanhua Chen	7866
<i>AMATA: Adaptive Multi-Agent Trajectory Alignment for Knowledge-Intensive Question Answering</i>	
Taolin Zhang, Dongyang Li, Chen Chen, Qizhou Chen, Jiuheg Wan, Xiaofeng He, Chengyu Wang and Richang Hong	7879
<i>ChessArena: A Chess Testbed for Evaluating Strategic Reasoning Capabilities of Large Language Models</i>	
Jincheng Liu, Sijun He, Jingjing Wu, Xiangsen Wang, Yang Chen, Zhaoqi Kuang, Siqi Bao and Yuan Yao	7901

<i>From Completion to Editing: Unlocking Context-Aware Code Infilling via Search-and-Replace Instruction Tuning</i>	
Jiajun Zhang, Zeyu Cui, Jiayi Yang, Lei Zhang, Yuheng Jing, Zeyao Ma, Tianyi Bai, Zilei Wang, Qiang Liu, Liang Wang, Binyuan Hui and Junyang Lin	7955
<i>What Makes an Ideal Quote? Recommending “Unexpected yet Rational” Quotations via Novelty</i>	
Poweï Chang, Jin Xiao, Guanglei Yue, Qianyu He, Yanghua Xiao, Deqing Yang and Jiaqing Liang	7985
<i>Do Not Step Into the Same River Twice: Learning to Reason from Trial and Error</i>	
Chenming Tang, Hsiu-Yuan Huang, Weijie Liu, Clive Bai, Saiyong Yang and Yunfang Wu .	8018
<i>iTAG: Inverse Design for Natural Text Generation with Accurate Causal Graph Annotations</i>	
Wenshuo Wang, Boyu Cao, Nan Zhuang and Wei Li	8033
<i>Dynamic Long Context Reasoning over Compressed Memory via End-to-End Reinforcement Learning</i>	
Zhuoen Chen, Dongfang Li, Meishan Zhang, Baotian Hu and Min Zhang	8064
<i>From Word to World: Can Large Language Models be Implicit Text-based World Models?</i>	
Yixia Li, Hongru Wang, Jiahao Qiu, Zhenfei Yin, Dongdong Zhang, Cheng Qian, Zeping Li, Xiaoteng Ma, Guanhua Chen and Heng Ji	8084
<i>GAMBIT: A Gamified Jailbreak Framework for Multimodal Large Language Models</i>	
Xiangdong Hu, Yangyang Jiang, Qin Hu and Xiaojun Jia	8112
<i>Large Language Model-Enhanced Multi-Armed Bandits</i>	
Jiahang Sun, Zhiyong Wang, Runhan Yang, Chenjun Xiao, John C.s. Lui and Zhongxiang Dai	8130
<i>Flattery in Motion: Benchmarking and Analyzing Sycophancy in Video-LLMs</i>	
Wenrui Zhou, Mohamed Hendy, Shu Yang, Qingsong Yang, Zikun Guo, Yuyu Luo, Lijie Hu and Di Wang	8146
<i>Mem2ActBench: A Benchmark for Evaluating Long-Term Memory Utilization in Task-Oriented Autonomous Agents</i>	
Yiting Shen, Kun Li, Wei Zhou and Songlin Hu	8173
<i>ARK: Answer-Centric Retriever Tuning via KG-augmented Curriculum Learning</i>	
Jiawei Zhou, Hang Ding and Haiyun Jiang	8191
<i>On the Emotion Understanding of Synthesized Speech</i>	
Yuan Ge, Haishu Zhao, AoKai Hao, Junxiang Zhang, Bei Li, Xiaoqian Liu, Chenglong Wang, Jianjin Wang, Bingsen Zhou, Bingyu Liu, JingBo Zhu, Zhengtao Yu and Tong Xiao	8209
<i>Taming ZombieAgents: A Markov State-Aware Framework for Resilient Multi-Agent Evolution</i>	
Taolin Zhang, Pukun Zhao, Qizhou Chen, Jiuheng Wan, Chen Chen, Xiaofeng He, Chengyu Wang and Richang Hong	8224
<i>A Survey of Large Language Model-Based Search Agents</i>	
Yunjia Xi, Jianghao Lin, Yongzhao Xiao, Zheli Zhou, Rong Shan, Te Gao, Jiachen Zhu, Weiwen Liu, Yong Yu and Weinan Zhang	8244
<i>RSDA: Restoring Stale Data Affinity via Dynamic Renovation Strategy for Mitigating Data Scarcity</i>	
Yidan Liang, Jia Zhu, Weijie Shi, Hanghui Guo, Yue Cui, Jiawei Shen, Guoqing Ma, Jingjiang Liu, Qingyu Niu, Yilin Wang, Shimin Di and Jiajie Xu	8280

<i>The Reasoning Trap: How Enhancing LLM Reasoning Amplifies Tool Hallucination</i> Chenlong Yin, Zeyang Sha, Shiwen Cui, Changhua Meng and Zechao Li	8310
<i>Protecting multimodal large language models against misleading visualizations</i> Jonathan Tonglet, Tinne Tuytelaars, Marie-Francine Moens and Iryna Gurevych	8329
<i>Dual Activation-Weight Sparsity: A Training-Free Framework for Efficient Large Language Model Compression</i> Luoyang Sun, Guangyan Li, Cheng Deng, Haifeng Zhang, Jian Zhao, Yongqiang Tang, Wensheng Zhang and Jun Wang	8350
<i>Repeated Sequences Reveal Gaps between Large Language Models and Natural Language</i> Kumiko Tanaka-Ishii	8367
<i>AgentCoMa: A Compositional Benchmark Mixing Commonsense and Mathematical Reasoning in Real-World Scenarios</i> Lisa Alazraki, Lihu Chen, Ana Brassard, Joe Stacey, Hossein A. Rahmani and Marek Rei ..	8383
<i>HERMES: KV Cache as Hierarchical Memory for Efficient Streaming Video Understanding</i> Haowei Zhang, Shudong Yang, Jinlan Fu, See-Kiong Ng and Xipeng Qiu	8411
<i>Reasoning with OmniThought: A Large CoT Dataset with Verbosity and Cognitive Difficulty Annotations</i> Wenrui Cai, Chengyu Wang, Junbing Yan, Jun Huang and Xiangzhong Fang	8431
<i>InferenceDynamics: Adaptive LLM Routing through Structured Capability and Knowledge Profiling</i> Haochen Shi, Tianshi Zheng, Weiqi Wang, Baixuan Xu, Chunyang Li, Chunkit Chan, Tao Fan and Yangqiu Song	8451
<i>ReportLogic: Evaluating Logical Quality in Deep Research Reports</i> Jujia Zhao, Zhaoxin Huan, Zihan Wang, Xiaolu Zhang, Jun Zhou, Suzan Verberne and Zhaochun Ren	8470
<i>Detecting RAG Extraction Attack via Dual-Path Runtime Integrity Game</i> Yuanbo Xie, Yingjie Zhang, Yulin Li, Shouyou Song, Xiaokun Chen, Zhihan Liu, Liya Su and Tingwen Liu	8503
<i>Beyond Monolingual Assumptions: A Survey on Code-Switched NLP in the Era of Large Language Models across Modalities</i> Rajvee Sheth, Samridhi Raj Sinha, Mahavir Patil, Himanshu Beniwal and Mayank Singh ..	8519
<i>Look Light, Think Heavy: What Multimodal Chain-of-Thought Reasoning Can and Cannot Do</i> Zhuoran Jin, Kejian Zhu, Hongbang Yuan, Yupu Hao, Pengfei Cao, Yubo Chen, Kang Liu and Jun Zhao	8567
<i>ARCQuant: Boosting NVFP4 Quantization with Augmented Residual Channels for LLMs</i> Haoqian Meng, Yilun Luo, Yafei Zhao, Wenyuan Liu, Peng Zhang and Xindian Ma	8609
<i>Revisiting Model Interpolation for Efficient Reasoning</i> Taiqiang Wu, Runming Yang, Tao Liu, Jiahao Wang and Ngai Wong	8624
<i>MirageBackdoor: A Stealthy Attack that Induces Think-Well-Answer-Wrong Reasoning</i> Yizhe Zeng, Wei Zhang, Yunpeng Li, Juxin Xiao, Xiao Wang and Yuling Liu	8639
<i>MulVul: Retrieval-augmented Multi-Agent Code Vulnerability Detection via Cross-Model Prompt Evolution</i> Zihan Wu, Jie Xu, Yun Peng, Chun Yong Chong and Xiaohua Jia	8660

<i>LiGen: Active Lipid Generation via a Molecular Language Model</i>	
Ying Zhan, Xiuqi Tang, Yan Zhang, Xiao Tan, Dian Shen, Zhou Yu and Beilun Wang	8674
<i>Gained in Translation: Privileged Pairwise Judges Enhance Multilingual Reasoning</i>	
Lintang Sutawika, Gokul Swamy, Steven Wu and Graham Neubig	8687
<i>Ted-Tok: Maintaining an Evolving Vocabulary for Lifelong Learning</i>	
Jiameng Huang, Zhi Zhang, Zhenyu He, Jiacheng Sun and Di He	8706
<i>When Correct Beliefs Collapse: Epistemic Resilience of LLMs under Clinical Pressure</i>	
Boyu Xiao, Xiuqi Tian, Xuwen Song, Haochun Wang, Guanchun Song, Sendong Zhao and Bing Qin	8720
<i>Beyond Transcripts: A Renewed Perspective on Audio Chaptering</i>	
Fabian Retkowski, Maïke Züfle, Thai Binh Nguyen, Jan Niehues and Alexander Waibel	8765
<i>SEA-BED: How Do Embedding Models Represent Southeast Asian Languages?</i>	
Wuttikorn Ponwitayarat, Peerat Limkonchotiwat, Raymond Ng, Jann Railey Montalan, Thura Aung, Jian Gang Ngui, Yosephine Susanto, William Chandra Tjhi, Panuthep Tasawong, Erik Cambria, Ekapol Chuangsuwanich and Sarana Nutanong	8788
<i>Is this chart lying to me? Automating the detection of misleading visualizations</i>	
Jonathan Tonglet, Jan Zimny, Tinne Tuytelaars and Iryna Gurevych	8823
<i>From Laboratory to Real-World Applications: Benchmarking Agentic Code Reasoning at the Repository Level</i>	
Jia Li, Yuxin Su and Michael R. Lyu	8845
<i>Privacy Collapse: Benign Fine-Tuning Can Break Contextual Privacy in Language Models</i>	
Anmol Goel, Cornelius Emde, Seong Joon Oh, Sangdoon Yun and Martin Gubri	8870
<i>Dissecting Failure Dynamics in Large Language Model Reasoning</i>	
Wei Zhu, Jian Zhang, Lixing Yu, Kun Yue and Zhiwen Tang	8893
<i>QBridge: Bridging Natural Language and SQL via Gold Query Rewriting with Agentic Refinement</i>	
Zhensheng Luo, Sai Wu, Yuan Qiu, Chang Yao, Gang Chen and Xiu Tang	8915
<i>SCVQ: Sparse-Compensated Vector Quantization for Large Language Models</i>	
Zixuan Zhou, Yujun Diao, Zicheng Kong, Dehua Ma, Zhenbo Xu, Pei Pei Li and Zhaofeng He	8934
<i>Shanks: Simultaneous Hearing and Thinking for Spoken Language Models</i>	
Cheng-Han Chiang, Xiaofei Wang, Linjie Li, Chung-Ching Lin, Kevin Lin, Shujie Liu, Zhendong Wang, Zhengyuan Yang, Hung-yi Lee and Lijuan Wang	8951
<i>Efficient KL Divergence Estimation via Truncated Top-K Integration for Large Language Models</i>	
Xinyuan Wang, Luozhijie Jin, Bo Wang, Yuan Li, Zhangyue Yin and Xipeng Qiu	8973
<i>Visual Attention Reasoning via Hierarchical Search and Self-Verification</i>	
Wei Cai, Jian Zhao, Yuchen Yuan, Tianle Zhang, Ming Zhu, Haichuan Tang and Xuelong Li	8986
<i>VFA: Empowering Multilingual MLLMs via Vision-Free Adaptation</i>	
Yixia Li, Yaqing Shi, Zhiwen Ruan, Dongdong Zhang, Lingjie Jiang, Shaohan Huang, Yun Chen, Guanhua Chen and Furu Wei	8998
<i>LearnerCoMPASS: Intelligent Tutoring System with Dynamic Cognitive Diagnosis and Multi-Model Path Planning</i>	
Ziji Sheng, Guiyao Tie, Weidong Wang, Pan Zhou and Daizong Liu	9016

<i>Investigating More Explainable and Partition-Free Compositionality Estimation for LLMs: A Rule-Generation Perspective</i>	
Ziyao Xu, Cong Wang and Houfeng Wang	9043
<i>LVLMS and Humans Ground Differently in Referential Communication</i>	
Peter Zeng, Weiling Li, Amie J. Paige, Zhengxiang Wang, Panagiotis Kaliosis, Dimitris Samaras, Gregory J. Zelinsky, Susan Brennan and Owen Rambow	9061
<i>Evaluating the Expressive Appropriateness of Speech in Rich Contexts</i>	
Tianrui Wang, Ziyang Ma, Yizhou Peng, Haoyu Wang, Zhikang Niu, Zikang Huang, Yihao Wu, Yi-Wen Chao, Yu Jiang, Yuheng Lu, Guanrou Yang, Xuanchen Li, Hexin Liu, Chunyu Qiang, Cheng Gong, Yifan Yang, Tianchi Liu, Junyu Wang, Nana Hou, Meng Ge, Fuming You, Yang Wei, Zhongqian Sun, Hu Haifeng, Xiaobao Wang, Eng Siong Chng, Xie Chen, Longbiao Wang and Jianwu Dang	9088
<i>UniMoE-Audio: Unified Speech and Music Generation with Dynamic-Capacity Mixture-of-Experts</i>	
Zhenyu Liu, Yunxin li, Xuanyu Zhang, Qixun Teng, Shenyuan Jiang, Xinyu Chen, Haoyuan Shi, Haolan Chen, Fanbo Meng, Mingjun Zhao, Yu Xu, Yancheng He, Baotian Hu, Haizhou Li and Min Zhang	9107
<i>Are LLMs Reliable Rankers? Rank Manipulation via Two-Stage Token Optimization</i>	
Tiancheng Xing, Jerry Li, Yixuan Du and Xiyang Hu	9120
<i>KCVR: Knowledge-Centric Video Reconstruction for Structured Pedagogical Summarization via Dynamic Graph Planning</i>	
Jingjiang Liu, Jia Zhu, Hanghui Guo, Weijie Shi, Yue Cui, Xiaokang Jin, Yilin Wang, Qingyu Niu, Jiawei Shen, Guoqing Ma, Yidan Liang, Shimin Di and Jiajie Xu	9133
<i>RepoGenesis: Benchmarking End-to-End Microservice Generation from Readme to Repository</i>	
Zhiyuan Peng, Xin Yin, Pu Zhao, Fangkai Yang, Lu Wang, Ran Jia, Xu Chen, Qingwei Lin, Saravan Rajmohan and Dongmei Zhang	9161
<i>CityCube: Benchmarking Cross-view Spatial Reasoning on Vision-Language Models in Urban Environments</i>	
Haotian Xu, Yue Hu, Zhengqiu Zhu, Chen Gao, Ziyou Wang, Junreng Rao, Wenhao Lu, Weishi Li, Qianjun Yin and Yong Li	9190
<i>ContrastKV: Robust KV Cache Eviction via Contrastive Signal Fusion for Multi-Query Generalization</i>	
Xingchi Chen, Peiyuan Zong, Ziqiang Gao, Qing Li, Yong Jiang, Fa Zhu and Hui Li	9216
<i>Teach a Reward Model to Correct Itself: Reward Guided Adversarial Failure Discovery for Robust Reward Modeling</i>	
Pankayaraj Pathmanathan and Furong Huang	9230
<i>Hierarchical Acoustic-Semantic Modeling: Modality Separation and Semantic Coherence for Full-Duplex SLMs</i>	
Zhenyu Liu, Xuanyu Zhang, Yunxin li, Qixun Teng, Shenyuan Jiang, Haolan Chen, Mingjun Zhao, Fanbo Meng, Yu Xu, Yancheng He, Baotian Hu, Haizhou Li and Min Zhang	9264
<i>VisPCO: Visual Token Pruning Configuration Optimization via Budget-Aware Pareto-Frontier Learning for Vision-Language Models</i>	
Huawei Ji, Yuanhao Sun, Yuan Jin, Cheng Deng, Jiabin Ding, Luoyi Fu and Xinbing Wang	9281
<i>CoSToM: Causal-oriented Steering for Intrinsic Theory-of-Mind Alignment in Large Language Models</i>	
Mengfan Li, Xuanhua Shi and Yang Deng	9302
<i>ConlangCrafter: Constructing Languages with a Multi-Hop LLM Pipeline</i>	
Morris Alper, Moran Yanuka, Raja Giryes and Gaspar Begus	9318

<i>XY-Tokenizer: Mitigating the Semantic-Acoustic Conflict in Low-Bitrate Speech Codecs</i>	
Yitian Gong, Luozhijie Jin, Kuangwei Chen, Dong Zhang, Ruifan Deng, Xiaogui Yang, Xin Zhang, Zhaoye Fei, Qinyuan Cheng, Shimin Li and Xipeng Qiu	9350
<i>Empathy in Diversity: Personalized Depression and Anxiety Therapy via Dialogue State Tracking and Patient-Aware Planning</i>	
Xinwei Yang, Junyi Fan, Yuqing Liu, Jiaxuan Wang, Jiashuai Zhang, Hongru Liang, Wenqiang Lei and Yao Song	9370
<i>VAPO: End-to-end Slide-Enhanced Speech Recognition with Omni-modal Large Language Models</i>	
Rui Hu, Delai Qiu, Yining Wang, Shengping Liu and Jitao Sang	9417
<i>Newspaper EatMeans Not Tasty": A Taxonomy and Benchmark for Coded Language in Real-World Chinese Online Reviews</i>	
Ruyuan Wan, Changye Li and Ting-Hao Kenneth Huang	9433
<i>Data Mixing Agent: Learning to Re-weight Domains for Continual Pre-training</i>	
Kailai Yang, Xiao Liu, Lei Ji, Hao Li, Xiao Liang, Zhiwei Liu, Yeyun Gong, Peng Cheng and Mao Yang	9447
<i>Experience-driven Multi-turn Reinforcement Learning for GUI Agents</i>	
Zhengxi Lu, Jiabo Ye, Fei Tang, Yongliang Shen, Haiyang Xu, Ziwei Zheng, Weiming Lu, Ming Yan, Fei Huang, Jun Xiao and Yueting Zhuang	9474
<i>UMMF: Protecting Copyright of Large Vision-Language Models through Unlearning-based Multi-modal Memorization Fingerprint</i>	
Xiaofan Zheng, Xinghao Wang and Xiaojun Wan	9497
<i>Automatic Correction of Writing Anomalies in Hausa Texts</i>	
Ahmad Mustapha Wali and Sergiu Nisioi	9514
<i>OS-Sentinel: Towards Safety-Enhanced Mobile GUI Agents via Hybrid Validation in Realistic Workflows</i>	
Qiushi Sun, Mukai Li, Zhoumianze Liu, Zhihui Xie, Fangzhi Xu, Zhangyue Yin, Kanzhi Cheng, Zehao Li, Zichen Ding, Qi Liu, Zhiyong Wu, Zhuosheng Zhang, Ben Kao and Lingpeng Kong ..	9529
<i>PAM: Enhancing General Alignment of Large Reasoning Models through Priority-Aware Metacognition</i>	
Zhihao Xu, Fuzhen Yang, Liang Lin and Xiting Wang	9554
<i>One Battle After Another: Probing LLMs' Limits on Multi-Turn Instruction Following with a Benchmark Evolving Framework</i>	
Qi Jia, Ye Shen, Xiujie Song, Kaiwei Zhang, Shibo Wang, Dun Pei, Xiangyang Zhu and Guangtao Zhai	9574
<i>AGSC: Adaptive Granularity and Semantic Clustering for Uncertainty Quantification in Long-text Generation</i>	
Guanran Luo, Wentao Qiu, Wanru Zhao, Wenhan Lv, Zhongquan Jian, Meihong Wang and Qingqiang Wu	9591
<i>Tracing the Roots: A Multi-Agent Framework for Uncovering Data Lineage in Post-Training LLMs</i>	
Yu Li, Xiaoran Shang, Qizhi Pei, Yun Zhu, Xin Gao, Honglin Lin, Zhanping Zhong, Zhuoshi Pan, Zheng Liu, Xiaoyang Wang, Conghui He, Dahua Lin, Feng Zhao and Lijun Wu	9606
<i>Understanding or Memorizing? A Case Study of German Definite Articles in Language Models</i>	
Jonathan Drechsel, Erisa Bytyqi and Steffen Herbold	9626

<i>Alloc-MoE: Budget-Aware Expert Activation Allocation for Efficient Mixture-of-Experts Inference</i>	Baihui Liu, Kaiyuan Tian, Wei Wang, Zhaoning Zhang, Linbo Qiao and Dongsheng Li	9653
<i>CodeEvo: Interaction-Driven Synthesis of Code-centric Data through Hybrid and Iterative Feedback</i>	Qiushi Sun, Jingyang Gong, Lei Li, Qipeng Guo and Fei Yuan	9668
<i>Benchmarking and Learning Real-World Customer Service Dialogue</i>	Tianhong Gao, Jundong Shen, Jiapeng Wang, Bei Shi, Ying Ju, Junfeng Yao and Huiyu Yu .	9688
<i>AdaJudge: Adaptive Multi-Perspective Judging for Reward Modeling</i>	Yongliang Miao, Yangyang Liang and Mengnan Du	9712
<i>IEvoAgent: Evolving Conversational Agent based on User Implicit Feedback</i>	Yichen Cai, Jiayang Li, Junyuan Qiu, Jingya Guo, Weitao You, Changyuan Yang, Lingyun Sun and Pei Chen	9725
<i>SGPVT: Self-Generated Proximal Visual Tokens for Mitigating Proximal Collateral Damage in MLLM Unlearning</i>	Jiaqi Li, Zhijing Zhang, Jiahui Geng, Sheng Bi, Chuanyi Zhang, Fan Liu and Guilin Qi	9744
<i>VIGIL: Defending LLM Agents Against Tool-Stream Injection via Verify-Before-Commit</i>	Junda Lin, Zhaomeng Zhou, Zhi Zheng, Shuochen Liu, Tong Xu, Yong Chen and Enhong Chen	9764
<i>“I See What You Did There”: Can Large Vision-Language Models Understand Multimodal Puns?</i>	Naen Xu, Jiayi Sheng, Changjiang Li, Chunyi Zhou, Yuyuan Li, Tianyu Du, Jun Wang, Zhihui Fu, Jinbao Li and Shouling Ji	9786
<i>TRACE: Traversal Retrieval-Augmented Chain of Evidence for Document Understanding</i>	Liqi He, Zuchao Li, Hao Huang and Ping Wang	9806
<i>FinKario: Event-Enhanced Automated Construction of Financial Knowledge Graph</i>	Xiang Li, Penglei Sun, Wanyun Zhou, Zikai Wei, Yongqi Zhang and Xiaowen Chu	9826
<i>EvoSci: A Bio-Inspired Multi-Agent Framework for the Evolution of Scientific Discovery</i>	Xiaoyu Xiong, Yuqi Ren and Deyi Xiong	9846
<i>DREAM: Deep Research Evaluation with Agentic Metrics</i>	Elad Ben Avraham, Changhao Li, Ron Dorfman, Roy Ganz, Oren Nuriel, Amir Dudai, Aviad Aberdam, Noah Flynn, Elman Mansimov, Aditya Kalyanpur and Ron Litman	9879
<i>ThinkPersona: Thinking with Persona Graphs for Faithful Individualized Role-Playing</i>	Yichen Cai, Pei Chen, Jiayang Li, Jingya Guo, Zejian Li, Changyuan Yang and Lingyun Sun	9905
<i>Sparrow: Text-Anchored Window Attention with Visual-Semantic Glimpsing for Speculative Decoding in Video LLMs</i>	Libo Zhang, Zhaoning Zhang, Hongwanyang, Peng Qiao and Dongsheng Li	9928
<i>SPEAK: Spiking Neurons as an Entropy-Aware Tokenizer for Large Language Models</i>	Ming Chen, Wenyaoy Li, Chao Liang, Shi Gu, Peng Lin, De Ma, Huajin Tang, Qian Zheng and Gang Pan	9943
<i>Unlocking Implicit Experience: Synthesizing Tool-Use Trajectories from Text</i>	Zhihao Xu, Rumei Li, Jiahuan Li, Rongxiang Weng, Jingang Wang, Xunliang Cai and Xiting Wang	9961
<i>A Layer-wise Analysis of Supervised Fine-Tuning</i>	Qinghua Zhao, Xueling Gong, Xinyu Chen, Zhongfeng Kang and Xinlu Li	9981

<i>Think in Latent Thoughts: A New Paradigm for Gloss-Free Sign Language Translation</i> Yiyang Jiang, Li Zhang, Xiao-Yong Wei and Li Qing	9993
<i>CTRAP: Embedding Collapse Trap to Safeguard Large Language Models from Harmful Fine-Tuning</i> Biao Yi, Tiansheng Huang, Baolei Zhang, Tong Li, Lihai Nie, Zheli Liu and Li Shen	10013
<i>Correct When Paired, Wrong When Split: Decoupling and Editing Modality-Specific Neurons in MLLMs</i> Tingchao Fu, Wenkai Wang, Fanxiao Li, Huadong Zhang, Jinhong Zhang, Dayang Li, Yunyun Dong, Renyang Liu and Wei Zhou	10031
<i>ACIArena: Toward Unified Evaluation for Agent Cascading Injection</i> Hengyu An, Minxi Li, Jinghuai Zhang, Naen Xu, Chunyi Zhou, Changjiang Li, Xiaogang Xu, Tianyu Du and Shouling Ji	10049
<i>PLAWBENCH: A Rubric-Based Benchmark for Evaluating LLMs in Real-World Legal Practice</i> Yuzhen Shi, Huanghai Liu, Yiran HU, Song Gaojie, Xu Xinran, Yubo Ma, Tianyi Tang, Li Zhang, Qingjing Chen, Feng Di, Wenbo Lv, Weiheng Wu, Kexin Yang, Sen Yang, Wei Wang, Rongyao Shi, Qiu Yuanyang, Yuemeng Qi, Zhang Jingwen, Sui Xiaoyu, Yifan Chen, Zhang Yi, An Yang, Bowen Yu, Dayiheng Liu, Junyang Lin, Weixing Shen, Bing Zhao, Charles L. A. Clarke and HU Wei	10067
<i>VoxMind: An End-to-End Agentic Spoken Dialogue System</i> Tianle Liang, Yifu Chen, Shengpeng Ji, Yijun Chen, Zhiyang Jia, Jingyu Lu, Fan Zhuo, Xueyi Pu, Yangzhuo Li and Zhou Zhao	10117
<i>MT³: A Synergistic Multi-Task RL Framework for Specializing MLLMs in Text Image Machine Translation</i> Zhaopeng Feng, Yupu Liang, Shaosheng Cao, Jiayuan Su, Jiahan Ren, Zhijie Zhou, Wenxuan Huang, Jian Wu and Zuozhu Liu	10140
<i>Do LLMs Capture Embodied Cognition and Cultural Variation? Cross-Linguistic Evidence from Demonstratives</i> Yu Wang, Emmanuele Chersoni and Chu-Ren Huang	10158
<i>Feeling Rules in Language Models: Mapping Norms of Emotional Appropriateness Across Roles, Institutions, and Intensity</i> Guangrui Fan, DanDan Liu, Aznul Qalid MD Sabri, Rui Zhang and Pan Lihu	10175
<i>SGT: Securing Open-Source LLMs Against Malicious Fine-tuning via Safety Guidance Trigger</i> Sunguk Shin, Fangzhao Wu, Byung-Jun Lee, Meeyoung Cha and Sungwon Park	10194
<i>Every Token Counts: Generalizing 16M Ultra-Long Context in Large Language Models</i> Xiang Hu, Zhanchao Zhou, Ruiqi Liang, Zehuan Li, Wei Wu and Jianguo Li	10208
<i>Self-Reflective Generation at Test Time</i> Jian Mu, Qixin Zhang, Zhiyong Wang, Menglin Yang, Shuang Qiu, Chengwei Qin, Zhongxiang Dai and Yao Shu	10221
<i>GKnow: Measuring the Entanglement of Gender Bias and Factual Gender</i> Leonor Veloso and Hinrich Schuetze	10240
<i>Reinforcing Agentic Search Via Reward Density Optimization</i> Kun Luo, Hongjin Qian, Zheng Liu, Ziyi Xia, Shitao Xiao, Zhao Cao, Siqi Bao, Jun Zhao and Kang Liu	10261
<i>Scaling External Knowledge Input Beyond Context Windows of LLMs via Multi-Agent Collaboration</i> Zijun Liu, Zhennan Wan, Peng Li, Ming Yan, Fei Huang and Yang Liu	10284

<i>M²PO: Multi-Perspective Multi-Pair Preference Optimization for Machine Translation</i>	
Hao Wang, Linlong Xu, Heng Liu, Yangyang Liu, Xiaohu Zhao, Bo Zeng, Liangying Shao, Yichen Dong, Xinwei Wu, Jiang Zhou, Tianyu Dong, Xiangxiang Zeng, Longyue Wang and Weihua Luo.....	10315
<i>An Exploration of Mamba for Speech Self-Supervised Models</i>	
Tzu-Quan Lin, Heng-Cheng Kuo, Tzu-Chieh Wei, Hsi-Chun Cheng, Chun Wei Chen, Hsien-Fu Hsiao, Yu Tsao and Hung-yi Lee	10337
<i>Ready Jurist One: Benchmarking Language Agents for Legal Intelligence in Dynamic Environments</i>	
Zheng Jia, Shengbin Yue, Wei Chen, Siyuan Wang, Yidong Liu, Zejun Li, Yun Song and Zhongyu Wei.....	10351
<i>When Morphology Hides in Plain Sight: Breaking the Isolation in Vietnamese and Beyond</i>	
Anh Trac Duc Dinh, Khang Hoang Nhat Vo, Tai Tien Ta, Vinh Cong Doan and Tho Quan .	10377
<i>FineLAP: Taming Heterogeneous Supervision for Fine-grained Language-Audio Pretraining</i>	
Xiquan Li, Xuenan Xu, Ziyang Ma, Wenxi Chen, Haolin He, Qiuqiang Kong and Xie Chen	10393
<i>OCR-Memory: Optical Context Retrieval for Long-Horizon Agent Memory</i>	
Jinze Li, Yang Zhang, Xin Yang, Jiayi QU, Jinfeng Xu, Shuo Yang, Junhua Ding and Edith Cheuk-Han Ngai.....	10409
<i>From Past To Path: Masked History Learning for Next-Item Prediction in Generative Recommendation</i>	
Kaiwen Wei, Kejun he, Xiaomian Kang, Jie Zhang, Ymyang, Li Jin, Zhenyang Li, Jiang Zhong, Richard He Bai and Junnan Zhu	10421
<i>MuSe: Multi-Stage Graph Reasoning via Vision-Language Models</i>	
Guanyu Wang, Xu Chu, Zhijie Tan, Xinrong Chen, Tong Mo and Weiping Li	10442
<i>Where and What: Reasoning Dynamic and Implicit Preferences in Situated Conversational Recommendation</i>	
Dongding Lin, Jian Wang, Yongqi Li and Wenjie Li.....	10463
<i>When Agents Look the Same: Quantifying Distillation-Induced Similarity in Tool-Use Behaviors</i>	
Chenghao Yang, Yuning Zhang, Zhoufutu Wen, Tao Gong, Jiaheng Liu, Qi Chu and Nenghai Yu	10482
<i>Mind the (DH) Gap! A Contrast in Risky Choices Between Reasoning and Conversational LLMs</i>	
Luise Ge, Yongyan Zhang and Yevgeniy Vorobeychik	10503
<i>Beyond Single-View Detection: A Dual-Space Reasoning Framework for Interpretable Harmful Meme Understanding</i>	
Wenqing Hou, Hongkui Tu, Ye Wang, Yue Zhang, Yuying Liu, Dong Zhu, Liqun Gao and Bin Zhou	10526
<i>From Experts to Bases: Orthogonal Subspace Mixture for Continual Multimodal Instruction Tuning</i>	
Pei Chen, Xilai Wang, Shiqixu, Zejian Li and Lingyun Sun	10545
<i>Distilling Large Embeddings via Hyperspherical Householder Quantization</i>	
Yihang Wang, Bin Wu, Yueyang Su, Tianfu Zhang, Yiqi Du, Lei Yu, Jiafeng Guo and Xueqi Cheng	10562
<i>Masked by Consensus: Disentangling Privileged Knowledge in LLM Correctness</i>	
Tomer Ashuach, Shai Gretz, Yoav Katz, Yonatan Belinkov and Liat Ein-Dor.....	10577

<i>Unified Thinker: A General Reasoning Core for Image Generation</i>	
Sashuai Zhou, Qiang Zhou, Jijin Hu, Hanqing Yang, Yue Cao, Junpeng Ma, Yinchao Ma, Jun Song, Tiezheng Ge, Cheng Yu, Bo Zheng and Zhou Zhao	10597
<i>Visual and Memory-Augmented Soccer Commentary Generation</i>	
Haoran Sun, Natthawut Kertkeidkachorn and Kiyooki Shirai	10614
<i>InstructDiff: Domain-Adaptive Data Selection via Contrastive Entropy for Efficient LLM Fine-Tuning</i>	
Junyou Su, He Zhu, Xiao Luo, Liyu Zhang, Hong-Yu Zhou, Yun Chen, Peng Li, Yang Liu and Guanhua Chen.....	10630
<i>Jakiro: Boosting Speculative Decoding via Decoupled MoE</i>	
Haiduo Huang, Fuwei Yang, Zhenhua Liu and Pengju Ren	10649
<i>Beyond Word Boundaries: A Hebrew Coreference Benchmark and an Evaluation Protocol for Morphologically Complex Text</i>	
Refael Shaked Greenfeld and Reut Tsarfaty	10669
<i>Reasoning with Ontology Graph: Toward Type-Constrained Knowledge Graph Question Answering</i>	
Yongxue Shan, Jie Peng, Zixuan Dong, Fei Hu and Xiaodong Wang	10684
<i>Visual Self-Fulfilling Alignment: Shaping Safety-Oriented Personas via Threat-Related Images</i>	
Qishun Yang, Shu Yang, Lijie Hu and Di Wang	10698
<i>GRASPrune: Global Gating for Budgeted Structured Pruning of Large Language Models</i>	
Ziyang Wang, Jiangfeng Xiao, Chuan Xiao, Ruoxiang LI, Rui Mao and Jianbin Qin	10719
<i>All That Glitters Is Not Gold: A Benchmark for Reference-Free Counterfactual Financial Misinformation Detection</i>	
Yuechen Jiang, Zhiwei Liu, Yupeng Cao, Yueru He, Ziyang Xu, Chen Xu, Zhiyang Deng, Prayag Tiwari, Xi Chen, Alejandro Lopez-Lira, Jimin Huang, Junichi Tsujii and Sophia Ananiadou	10737
<i>Musical Score Understanding Benchmark: Evaluating Large Language Models' Comprehension of Complete Musical Scores</i>	
Congren Dai, Yue Yang, Krinos Li, Huichi Zhou, Shijie Liang, Zhang Bo, Enyang Liu, Ge Jin, Hongran An, Haosen Zhang, Peiyuan Jing, KinHei Lee, Zhenxuan Zhang, Xiaobing Li and Maosong Sun.....	10777
<i>EchoVLM: Dynamic Mixture-of-Experts Vision-Language Model for Universal Ultrasound Intelligence</i>	
Chaoyin She, Ruifang Lu, Lida Chen, Wei Wang and Qinghua Huang	10800
<i>Critical-CoT: A Robust Defense Framework against Reasoning-Level Backdoor Attacks in Large Language Models</i>	
Vu Tuan Truong and Long Bao Le	10823
<i>How Do LLMs and VLMs Understand Viewpoint Rotation Without Vision? An Interpretability Study</i>	
Zhen Yang, Ping Jian, Zhongbin Guo, Zuming Zhang, Chengzhi Li, Yonghong Deng, Xinyue Zhang and Wenpeng Lu	10842
<i>Sycophants in the Courtroom: Are LLMs Fragile to Juridical Authority and Evolving Legal Standards?</i>	
Lorenzo Molfetta, Alessio Cocchieri, Luca Ragazzi, Ilaria Bartolini, Marco Patella and Gianluca Moro	10865
<i>Valid Survey Simulations with Limited Human Data: The Roles of Prompting, Fine-Tuning, and Rectification</i>	
Stefan Krsteski, Giuseppe Russo, Serina Chang, Robert West and Kristina Gligorić.....	10887

<i>ODUTQA-MDC: A Task for Open-Domain Underspecified Tabular QA with Multi-turn Dialogue-based Clarification</i>	
Zhensheng Wang, ZhanTeng Lin, Wenmian Yang, Kun Zhou, Yiquan Zhang and Weijia Jia	10907
<i>CuBridge: An LLM-Based Framework for Understanding and Reconstructing High-Performance Attention Kernels</i>	
Xing Ma, Yangjie Zhou, Wu Sun, Zihan Liu, Jingwen Leng, Yun Lin, Shixuan Sun, Minyi Guo and Jin Song Dong	10929
<i>Learning from Contrasts: Synthesizing Reasoning Paths from Diverse Search Trajectories</i>	
Peiyang Liu, Zhirui Chen, Xi Wang, Di Liang, Youru Li, Zhi Cai and Wei Ye	10947
<i>Defenses Against Prompt Attacks Learn Surface Heuristics</i>	
Li Li, Chenxiao Yu, Zhiyu Ni, Hao Li, Charith Peris, Chaowei Xiao and Yue Zhao	10970
<i>Memory-Augmented LLM-based Multi-Agent System for Automated Feature Generation on Tabular Data</i>	
FengXian Dong, Zhi Zheng, Xiao Han, Wei Chen, Jingqing Ruan, Tong Xu, Yong Chen and Enhong Chen	10988
<i>Learning from the Irrecoverable: Error-Localized Policy Optimization for Tool-Integrated LLM Reasoning</i>	
Qiao Liang, Yuke Zhu, Chao Ge, Lei Yang, Ying Shen, Bo Zheng and Sheng Guo	11008
<i>Learning from Evolving Training Dynamics: An Entropy-Maximizing Data Curation Strategy for LLM Supervised Post-Training</i>	
Mengxiang Zhang and Lingyuan Liu	11029
<i>Reinforcement Learning on Pre-Training Data</i>	
Siheng Li, Kejiao Li, Zenan Xu, Guanhua Huang, Kun Li, Haoyuan Wu, Wujiajia, Zihao Zheng, Chenchen Zhang, Kun Shi, Xue Gong, Qi Yi, Ruibin Xiong, Tingqiang Xu, Yuhao Jiang, Jianfeng Yan, Yuyuan Zeng, Guanghui Xu, Jinbao Xue, Zhijiang xu, Zheng Fang, Shuai LI, Qibin Liu, Xiaoxue Li, Zhuoyu Li, Yangyu Tao, Fei Gao, Cheng Jiang, Bochao Wang, Kai Liu, Jianchen Zhu, Wai Lam, Bo Zhou and Di Wang	11046
<i>Limited Linguistic Diversity in Embodied AI Datasets</i>	
Selma Liliane Wanna, Agnes Luhtaru, Jonathan Salfity, Ryan Barron, Juston Moore, Cynthia Matuszek and Mitch Pryor	11058
<i>Towards Mitigating Modality Bias in Vision-Language Models for Temporal Action Localization</i>	
Jiaqi Li, Guangming Wang, Shuntian Zheng, Minzhe Ni, Xiaoman Lu, Guanghui Ye and Yu Guan	11087
<i>CreditDecoding: Accelerating Parallel Decoding in Diffusion Large Language Models with Trace Credit</i>	
Kangyu Wang, Zhiyun Jiang, Haibo Feng, Weijia Zhao, Lin Liu, Jianguo Li, Zhenzhong Lan and Weiyao Lin	11105
<i>Question Difficulty Estimation for Large Language Models via Answer Plausibility Scoring</i>	
Jamshid Mozafari, Bhawna Piryani and Adam Jatowt	11124
<i>Vector Calligrapher: Generating Scalable Vector Graphics via Structured Linguistic Supervision</i>	
Bo Zhou, Xikang Chen, Yan Gong and Yin Zhang	11152
<i>CoQuIR: A Comprehensive Benchmark for Code Quality-Aware Information Retrieval</i>	
Jiahui Geng, Fengyu Cai, Shaobo Cui, Qing Li, Liangwei Chen, Chenyang Lyu, Haonan Li, Derui Zhu, Alexander Pretschner, Heinz Koeppel and Fakhri Karray	11169

<i>Sherry: Hardware-Efficient 1.25-Bit Ternary Quantization via Fine-grained Sparsification</i>	
Hong Huang, Decheng Wu, Qiangqiang Hu, Guanghua Yu, Jinhai Yang, Jianchen Zhu, Xue Liu and Dapeng Wu	11186
<i>DMN: A Compositional Framework for Jailbreaking Multimodal LLMs with Multi-Image Inputs</i>	
Wenzhuo Xu, Zhipeng Wei, Zonghao Ying, Deyue Zhang, Dongdong Yang, Xiangzheng Zhang and Quanchen Zou	11205
<i>PIAST: Rapid Prompting with In-context Augmentation for Scarce Training data</i>	
Paweł Batorski and Paul Swoboda	11222
<i>ProgressLM: Towards Progress Reasoning in Vision-Language Models</i>	
Jianshu Zhang, Chengxuan Qian, Haosen Sun, Haoran Lu, Dingcheng Wang, Letian Xue and Han Liu	11243
<i>Forest Before Trees: Latent Superposition for Efficient Visual Reasoning</i>	
Yubo Wang, Juntian Zhang, Yichen Wu, Yankai Lin, Nils Lukas and Yuhan Liu	11272
<i>Beyond Pedagogical Principles: Multi-Horizon Preference Optimization for Efficient Socratic Tutoring</i>	
Xin Shi, Chao Zhang, Yifan Zhu, Xueqiao Zhang and Yawei Luo	11289
<i>Protecting Language Models Against Unauthorized Distillation through Trace Rewriting</i>	
Xinhang Ma, William Yeoh, Ning Zhang and Yevgeniy Vorobeychik	11307
<i>The Confidence Dichotomy: Analyzing and Mitigating Miscalibration in Tool-Use Agents</i>	
Weihao Xuan, Qingcheng Zeng, Heli Qi, Yunze Xiao, Junjue Wang and Naoto Yokoya	11325
<i>XOXO: Stealthy Cross-Origin Context Poisoning Attacks against AI Coding Assistants</i>	
Adam Štorek, Mukur Gupta, Noopur Bhatt, Aditya Gupta, Janie Kim, Prashast Srivastava and Suman Jana	11350
<i>SafeMCP: Proactive Power Regulation for LLM Agent Defense via Environment-Grounded Look-Ahead Reasoning</i>	
Lichao Wang, ZhaoXing Ren, Tianzhuo Yang, Jiaming Ji, Chi Harold Liu, Yaodong Yang and Juntao Dai	11374
<i>C2: Scalable Rubric-Augmented Reward Modeling from Binary Preferences</i>	
Akira Kawabata and Saku Sugawara	11397
<i>LLaDA 1.5: Variance-Reduced Preference Optimization for Large Language Diffusion Models</i>	
Fengqi Zhu, Rongzhen Wang, Shen Nie, Xiaolu Zhang, Chunwei Wu, Jun Zhou, Yankai Lin, Ji-Rong Wen and Chongxuan Li	11425
<i>HarmRLVR: Weaponizing Verifiable Rewards for Harmful LLM Alignment</i>	
Yuexiao Liu, Lijun Li, Xingjun Wang and Jing Shao	11461
<i>Just Pass Twice: Efficient Token Classification with LLMs for Zero-Shot NER</i>	
Ahmed Ewais, Ahmed Hashish and Amr Ali	11480
<i>Lifting Optimized Binaries to Canonical Compiler IR via Structure-Aware Retrieval and Iterative Verification</i>	
Xiaobao Zhu, Jie Ren, Zhiqiang Li, Jie Zheng, Zhanyong Tang and Zheng Wang	11498
<i>Provably Safe Offline-to-Online RL: Decoupling Learning from Data-Driven Safety Enforcement</i>	
Kaitong Cai, Jusheng Zhang and Keze Wang	11517
<i>SHAPE: Unifying Safety, Helpfulness and Pedagogy for Educational LLMs</i>	
Sihang Zhao, Kangrui Yu, Youliang Yuan, Pinjia He and Hongyi Wen	11537

<i>Evolving Sparsity: Leveraging Token Importance Dynamics for Efficient LLM Decoding with Sparse Attention</i>	Ruizi Han, Miao Zhang, Ziyue Qiao and Liqiang Nie	11554
<i>MOA: Multi-Objective Alignment for Role-Playing Agents</i>	Chonghua Liao, Ke Wang, Yuchuan Wu, Ruoran Li, Fei Huang and Yongbin Li	11567
<i>Losing our Tail, Again: (Un)Natural Selection & Multilingual LLMs</i>	Eva Vanmassenhove	11589
<i>From Verbatim to Gist: Distilling Pyramidal Multimodal Memory via Semantic Information Bottleneck for Long-Horizon Video Agents</i>	Niu Lian, Yuting Wang, Hanshu Yao, Jinpeng Wang, Bin Chen, Yaowei Wang, Min Zhang and Shu-Tao Xia	11601
<i>Chain-of-Memory: Lightweight Memory Construction with Dynamic Evolution for LLM Agents</i>	Xiucheng Xu, Bingbing Xu, Tian Xueyun, Zihe Huang, Rongxin Chen, Li Yunfan and Huawei Shen	11618
<i>WebCoderBench: Benchmarking Web Application Generation with Comprehensive and Interpretable Evaluation Metrics</i>	Chenxu Liu, Yingjie Fu, Wei Yang, Ying Zhang and Tao Xie	11632
<i>Efficient Test-Time Scaling of Multi-Step Reasoning by Probing Internal States of Large Language Models</i>	Jingwei Ni, Ekaterina Fadeeva, Tianyi Wu, Mubashara Akhtar, Jiaheng Zhang, Elliott Ash, Markus Leippold, Timothy Baldwin, See-Kiong Ng, Artem Shelmanov and Mrinmaya Sachan .	11667
<i>Adaptive Prompt Structure Factorization: A Framework for Self-Discovering and Optimizing Compositional Prompt Programs</i>	Haoyue Liu, Zhichao Wang, Yongxin Guo, Haoran Shou and Xiaoying Tang	11690
<i>Cognitive Alpha Mining via LLM-Driven Code-Based Evolution</i>	Fengyuan Liu, Yi Huang, Sichun Luo, Yuqi Wang, Yazheng Yang, Xinye Li, Zefa Hu, Junlan Feng and Qi Liu	11715
<i>Why Do Emotions Change? Appraisal-Guided Reasoning for Emotion–Cause Triplet Extraction in Conversations</i>	Qiao Liang, Ying Shen, Yao Liu, Tiantian Chen and Lin Zhang	11750
<i>KV-Embedding: Training-free Text Embedding via Internal KV Re-routing in Decoder-only LLMs</i>	Yixuan Tang and Yi Yang	11773
<i>Fingerprinting LLMs via Prompt Injection</i>	Yuepeng Hu, Zhengyuan Jiang, Mengyuan Li, Osama Ahmed, Zhicong Huang, Cheng Hong and Neil Zhenqiang Gong	11795
<i>Exploring Layer Activation Dynamic of CoT via Knowledge Probe</i>	Chuanxin Zhang, Jiajun Liu, Yao He, Wenjun Ke, Peng Wang, Yankun Le, Sirui Liu and Zhaoyu Yang	11811
<i>From Trust to Compromise: Outcome-Verified LLM Phishing Simulation and Real-Time Defense</i>	Tulika Tewari, Nalin Asanka Gamagedara Arachchilage, Jagat Sesh Challa and Dhruv Kumar	11831
<i>EvoNarrator: Modeling Scientific Evolution for Feasible Hypothesis Generation</i>	Xiaoying Le, Pengfei Qian, Yuanzhao Zhai, Xu Zhang, Qian Liu, Feng Dawei and Bo Ding	11846

<i>DecoCal: Decoding with Calibration in Diffusion Large Language Models</i> Fan Xu, Huixuan Zhang and Xiaojun Wan	11866
<i>LiveFact: A Dynamic, Time-Aware Benchmark for LLM-Driven Fake News Detection</i> Cheng Xu, Changhong Jin, Yingjie Niu, Nan Yan, Yuke Mei, Shuhao Guan, Liming Chen and Tahar Kechadi	11881
<i>Beyond I Don't Know": Evaluating LLM Self-Awareness in Discriminating Data and Model Uncertainty</i> Jingyi Ren, Ante Wang, Yunghwei Lai, Xiaolong Wang, Linlu Gong, Weitao Li, Weizhi Ma and Yang Liu	11911
<i>Unlearners Can Lie: Evaluating and Improving Honesty in LLM Unlearning</i> Renjie Gu, Jiazhen Du, Yihua Zhang and Sijia Liu	11930
<i>MatchTIR: Fine-Grained Supervision for Tool-Integrated Reasoning via Bipartite Matching</i> Changle Qu, Sunhao Dai, Hengyi Cai, Jun Xu, Shuaiqiang Wang and Dawei Yin	11953
<i>Knowledge is Not Enough: Injecting RL Skills for Continual Adaptation</i> Pingzhi Tang, Yiding Wang and Muhan Zhang	11969
<i>BhashaSutra: A Task-Centric Unified Survey of Indian NLP Datasets, Corpora, and Resources</i> Raghvendra Kumar, Devankar Raj and Sriparna Saha	11998
<i>A Dual-Task Paradigm to Investigate Sentence Comprehension Strategies in Language Models</i> Rei Emura and Saku Sugawara	12065
<i>Diffusion-CAM: Faithful Visual Explanations for dMLLMs</i> Haomin Zuo, Yidi Li, Luoxiao Yang and Xiaofeng Zhang	12085
<i>Beyond Noise: Characterizing Creative Potential in Unverifiable LLM Hallucinations</i> Yu Yan, Chunhong Zhang, Haiyu Zhao, Ziyang Zeng, Zihao Liu, Yongkang Wu, Jianzhou Diao, 陈奕杰, Shujie Wang and Zheng Hu	12102
<i>LLM-SLM Collaborative Framework of Idiomatic Expression Generation</i> Hui Gao, Changhao Song, Peng Zhang, Jing Zhang, Chang Yang and Liuxian Ge	12125
<i>Focus-dLLM: Accelerating Long-Context Diffusion LLM Inference via Confidence-Guided Context Focusing</i> Lingkun Long, Yushi Huang, Shihao Bai, Ruihao Gong, Jun Zhang, Ao Zhou and Jianlei Yang	12146
<i>Table Question Answering in the Era of Large Language Models: A Comprehensive Survey of Tasks, Methods, and Evaluation</i> Wei Zhou, Bolei Ma, Annemarie Friedrich and Mohsen Mesgar	12161
<i>RouteMoA: Dynamic Routing without Pre-Inference Boosts Efficient Mixture-of-Agents</i> Jize Wang, Han Wu, Zhiyuan You, Yiming Song, Yijun Wang, Zifei Shan, Yining Li, Songyang Zhang, Xinyi Le, Cailian Chen, Xinpeng Guan and Dacheng Tao	12190
<i>Patches of Nonlinearity: Instruction Vectors in Large Language Models</i> Irina Bigoulaeva, Jonas Rohweder, Subhabrata Dutta and Iryna Gurevych	12209
<i>Reasoning Gets Harder for LLMs Inside A Dialogue</i> Ivan Kartáč, Mateusz Lango and Ondrej Dusek	12263
<i>Capability Decomposition for Unified Information Extraction via Hierarchical Mixture-of-Experts</i> Jing Zhou, Peng Wang, Wenjun Ke, Jiajun Liu and Yao He	12304

<i>PDR: A Plug-and-Play Positional Decay Framework for LLM Pre-training Data Detection</i> Jinhan Liu, Yibo Yang, Ruiying Lu, Piotr Piękos, Yimeng Chen, Peng Wang and Dandan Guo 12319	
<i>PRISM: Probabilistic Reward Model with Inherent Structural Modeling</i> Yuhang Zhou, Yixin Cao, Yuchen Ni, Shihan Dou, Xutian Chen, Ge Zhang, Xiang Liu and Guangnan Ye	12345
<i>Multilingual Idioms in Sentences and Conversations Across High-, Medium-, and Low-Resource Languages</i> Saeed Almheiri, Bilal Elbouardi, Salsabila Zahirah Pranida, Irina Nikishina, Ashwath Rao B, Parameswari Krishnamurthy, Muhammad Cendekia Airlangga, Rifo Ahmad Genadi, Nguyen Phan Gia Bao, Amir Hossein Yari, Hawau Olamide Toyin, Nurdaulet Mukhituly, Mena Attia, Beshar Hassan, Ahmad Fathan Hidayatullah, Tatsuki Kuribayashi, Haonan Li, Suma Bhat and Fajri Koto	12363
<i>Author-in-the-Loop Response Generation and Evaluation: Integrating Author Expertise and Intent in Responses to Peer Review</i> Qian Ruan and Iryna Gurevych	12390
<i>WeightLoRA: Keep Only Necessary Adapters</i> Andrey Veprikov, Vladimir Solodkin, Zyl Alexander, Andrey Savchenko and Aleksandr Beznosikov 12419	
<i>Reward Modeling for Scientific Writing Evaluation</i> Furkan Şahinuç, Subhabrata Dutta and Iryna Gurevych	12438
<i>Deriving Character Logic from Storyline as Codified Decision Trees</i> Letian Peng, Kun Zhou, Longfei Yun, Yupeng Hou and Jingbo Shang	12480
<i>GASim: A Graph-Accelerated Hybrid Framework for Social Simulation</i> Xuan Zhou, Yanhui Sun, Hantao Yao, Allen He, Yongdong Zhang and Wu Liu	12510
<i>Mock Worlds, Real Skills: Building Small Agentic Language Models with Synthetic Tasks, Simulated Environments, and Rubric-Based Rewards</i> Yuanjie Lyu, Chengyu Wang, Lei Shen, Jun Huang and Tong Xu	12529
<i>Investigating Cross-Modal Skill Injection: Scenarios, Methods, and Hyperparameters</i> Zhiyu Xu, Lean Wang, Yuanxin Liu, Lei Li, Hao Zhou, Fandong Meng, Jie Zhou and Xu Sun 12546	
<i>The Bidirectional Process Reward Model</i> Lingyin Zhang, Jun Gao, Xiaoxue Ren and Ziqiang Cao	12564
<i>AgentMark: Utility-Preserving Behavioral Watermarking for Agents</i> Kaibo Huang, Jin Tan, Yukun Wei, Wanling Li, Zipei Zhang, Hui Tian, Zhongliang Yang and Linna Zhou	12581
<i>Generative Gamer: Learning Equilibrium Strategy by LLM-driven Dynamic Deduction</i> Yadong Zhang, Xinshu Shen, Yupei Ren, Shangqing Zhao and Man Lan	12604
<i>Probing for Reading Times</i> Eleftheria Tsipidi, Samuel Kiegeand, Francesco Ignazio Re, Tianyang Xu, Mario Giulianelli, Karolina Stanczak and Ryan Cotterell	12618
<i>No More Stale Feedback: Co-Evolving Critics for Open-World Agent Learning</i> Zhicong Li, Lingjie Jiang, Yulan Hu, Xingchen Zeng, Yixia Li, Xiangwen Zhang, Guanhua Chen, Zheng Pan, Xin Li and Yong Liu	12643

<i>Continuous Interpretive Steering for Scalar Diversity</i> Ye-eun Cho	12661
<i>A Scalable Multi-LLM Collaboration System with Retrieval-based Selection and Exploration-Exploitation-Driven Enhancement</i> Shengji Tang, Jianjian Cao, Weihao Lin, Jiale Hong, Bo Zhang, Shuyue Hu, Lei Bai, Tao Chen, Wanli Ouyang and Peng Ye	12679
<i>From Form to Logic: Masked Reconstruction and Reasoning Distillation for Short Video Fake News Detection</i> Qingyan Wang, Lianwei Wu, Botao Wang, Wangkang and Yaxiong Wang	12698
<i>UR² : Unify RAG and Reasoning through Reinforcement Learning</i> Weitao Li, Boran Xiang, Xiaolong Wang, Jingyi Ren, Ante Wang, Zhinan Gou, Weizhi Ma and Yang Liu	12712
<i>LLM-as-Scheduler: Agentic Workflow Dynamic Scheduling</i> Dawei Xiang, Kexin Chu, Wenyan Xu, Wenhui Zhang and Wei Zhang	12752
<i>Mathematical Proof as a Litmus Test: Revealing Failure Modes of Advanced Large Reasoning Models</i> Dadi Guo, Jiayu Liu, Zhiyuan Fan, Zhitao He, Haoran Li, Yuxin Li, Yumeng Wang and Yi R. Fung	12764
<i>Memory-R1: Enhancing Large Language Model Agents to Manage and Utilize Memories via Reinforcement Learning</i> Sikuan Yan, Xiufeng Yang, Zuchao Huang, Ercong Nie, Zifeng Ding, Zonggen Li, Xiaowen Ma, Jinhe Bi, Kristian Kersting, Jeff Z. Pan, Hinrich Schuetze, Volker Tresp and Yunpu Ma	12805
<i>CostBench: Evaluating Multi-Turn Cost-Optimal Planning and Adaptation in Dynamic Environments for LLM Tool-Use Agents</i> Jiayu Liu, Cheng Qian, Zhaochen Su, Qing Zong, Shijue Huang, Bingxiang He and Yi R. Fung	12826
<i>Decoding Scientific Experimental Images: The SPUR Benchmark for Perception, Understanding, and Reasoning</i> Junpeng Ding, Zichen Tang, Haihong E, Mengyuan Ji, Yang Liu, Haolin Tian, Haiyang Sun, Pengqi Sun, Yang Xu, Yichen Liu, Haocheng Gao, Zijie Xi, Ruomeng Jiang, Peizhi Zhao, Rongjin Li, Yuanze Li, Jiacheng Liu, Zhongjun Yang, Jintong Chen and Siying Lin	12859
<i>SegTune: Structured and Fine-Grained Control for Song Generation</i> Yuejiao Wang, Zihao Ji, Pengfei Cai, Xu Li, Haorui Zheng, Zewen Song, Zhongliang Liu, Chen Zhang and Pengfei Wan	12883
<i>Making Large Language Models Efficient Dense Retrievers</i> Yibin Lei, Shwai He, Ang Li and Andrew Yates	12898
<i>Lightweight LLM Agent Memory with Small Language Models</i> Jiaquan Zhang, Chaoning Zhang, Shuxu Chen, Zhenzhen Huang, Pengcheng Zheng, Zhicheng Wang, Ping Guo, Fan Mo, Sung-Ho Bae, Jie Zou, Jiwei Wei and Yang Yang	12914
<i>HeteroSpec: Leveraging Contextual Heterogeneity for Efficient Speculative Decoding</i> Siran Liu, Yang Ye, Qianchao Zhu, Zane Cao and Yongchao He	12930
<i>A Model of the Language Process</i> Brandon Duderstadt and Hayden Helm	12948

<i>Rethinking Table Pruning in TableQA: From Sequential Revisions to Gold Trajectory-Supervised Parallel Search</i>	
Yu Guo, Shenghao Ye, Shuangwu Chen, Zijian Wen, Tao Zhang, Bai Qirui, Dong Jin, Yunpeng Hou, Huasen He, Jianyang and Xiaobin Tan	12960
<i>Interleaved Tool-Call Reasoning for Protein Function Understanding</i>	
Chuanliu Fan, Zicheng Ma, Huanran Meng, Aijia Zhang, Wenjie Du, Jun Zhang, Ziqiang Cao and Guohong Fu	12977
<i>Among Us: Language of Conspiracy Theorists on Mainstream Reddit</i>	
Francesco Corso, Giuseppe Russo, Francesco Pierri and Gianmarco De Francisci Morales	12996
<i>HybridKV: Hybrid KV Cache Compression for Efficient Multimodal Large Language Model Inference</i>	
Bowen Zeng, Feiyang Ren, Jun Zhang, Xiaoling Gu, Ke Chen, Lidan Shou and Huan Li ..	13018
<i>Breaking Down and Building Up: Mixture of Skill-Based Vision-and-Language Navigation Agents</i>	
Tianyi Ma, Yue Zhang, Zehao Wang and Parisa Kordjamshidi	13035
<i>Bridging Distance and Spectral Positional Encodings via Anchor-Based Diffusion Geometry Approximation</i>	
Yan Zimo, Zheng Xie, Runfan Duan, Chang Liu and Wumei Du	13066
<i>A Goal Without a Plan Is Just a Wish: Efficient and Effective Global Planner Training for Long-Horizon Agent Task</i>	
Shuzheng Si, Haozhe Zhao, Kangyang Luo, Gang Chen, Fanchao Qi, Minjia Zhang, Baobao Chang and Maosong Sun	13086
<i>Stereotype Bias in a Bilingual Setting: A Culturally Grounded Evaluation in Kazakhstan</i>	
Nurkhan Laiyk, Daniil Orel, Ayana Mussabayeva, Maiya Goloburda, Kamila Kuishibekova, Liya Goloburda, Diana Turmakhan, Preslav Nakov, Yuxia Wang and Fajri Koto	13114
<i>Arguments that Alter Minds: LLM Rationales Sway Human (and LLM) Notions of Plausibility</i>	
Shramay Palta, Peter A. Rankel, Sarah Wiegrefe and Rachel Rudinger	13132
<i>GLARE: Agentic Reasoning for Legal Judgment Prediction</i>	
Xinyu Yang, Chenlong Deng and Zhicheng Dou	13153
<i>Vista-LLM: Decoupled Query-Guided Visual Token Pruning for Efficient Long-Video Large Language Models</i>	
Zhenyu Li, Zuchao Li, Ping Wang, Lefei Zhang and Haojun Ai	13171
<i>RoboFailRing: Retrieval-Augmented and Language Grounding Failure Detection for VLM-enabled Robotic Manipulation</i>	
Chenduo Ying, Linkang Du, Yuanchao Shu and Peng Cheng	13188
<i>Scaling Law for Multimodal Large Language Model Supervised Fine-Tuning</i>	
YiFan Zhang, Tao Yu, Feng Li, Chaoyou Fu, Yibo Hu, Kun Wang, Qingsong Wen, Zhang Zhang, Liang Wang and Rong Jin	13203
<i>CrochetBench: Can Vision-Language Models Move from Describing to Doing in Crochet Domain?</i>	
Peiyu Li, Xiaobao Huang, Ting Hua and Nitesh V Chawla	13229
<i>DORA: A Dual-Objective Reinforcement Learning Framework for Effective and Efficient Multimodal Agentic Search</i>	
Guangming Qin, Yuhao Deng, Yukun Zhao, Zhenyang Li, Junfeng Wang, Dawei Yin, Ye Yuan, Guoren Wang, Yizhou Yan, Chengliang Chai and Lei Cao	13252

<i>IUQ: Interrogative Uncertainty Quantification for Long-Form Large Language Model Generation</i> Haozhi Fan, Jinhao Duan and Kaidi Xu	13273
<i>CODESTRUCT: Code Agents over Structured Action Spaces</i> Myeongsoo Kim, Chao-Chun Hsu, Dingmin Wang, Shweta Garg, Varun Kumar and Murali Krishna Ramanathan	13290
<i>Answering the Wrong Question: Reasoning Trace Inversion for Abstention in LLMs</i> Abinitha Gourabathina, Inkit Padhi, Manish Nagireddy, Subhajit Chaudhury and Prasanna Sattigeri	13307
<i>Beyond Single-shot Writing: Deep Research Agents are Unreliable at Multi-turn Report Revision</i> Bingsen Chen, Boyan Li, Ping Nie, Yuyu Zhang, Xi Ye and Chen Zhao	13325
<i>FinCall-Surprise: A Large Scale Multi-modal Benchmark for Earning Surprise Prediction</i> Dong Shu, Yanguang Liu, Huopu Zhang and Mengnan Du	13357
<i>ALDEN: Reinforcement Learning for Active Navigation and Evidence Gathering in Long Documents</i> Tianyu Yang, Terry Ruas, Yijun Tian, Jan Philip Wahle, Daniel Kurzawe and Bela Gipp ..	13371
<i>Reinforced Informativeness Optimization for Long-Form Retrieval-Augmented Generation</i> Yuhao Wang, Ruiyang Ren, Yucheng Wang, Xin Zhao, Jing Liu, Hua Wu and Haifeng Wang	13393
<i>Android Coach: Improve Online Agentic Training Efficiency with Single State Multiple Actions</i> Guo Gan, Yuxuan Ding, Cong Chen, Yuwei Ren, Yin Huang and Hong Zhou	13407
<i>Inside Out: Evolving User-Centric Core Memory Trees for Long-Term Personalized Dialogue Systems</i> Jihao Zhao, Ding Chen, Zhaoxin Fan, Kerun Xu, Mengting Hu, Bo Tang, Feiyu Xiong and Zhiyu li	13429
<i>FinChart-Bench: Benchmarking Financial Chart Comprehension in Vision-Language Models</i> Dong Shu, Haoyang Yuan, Yuchen Wang, Yanguang Liu, Huopu Zhang and Mengnan Du.	13447
<i>From Charts to Code: A Hierarchical Benchmark for Multimodal Models</i> Jiahao Tang, Henry Hengyuan Zhao, Lijian Wu, Zijian Zhang, Yifei Tao, Dongxing Mao, Yang Wan, Jingru Tan, Min Zeng, Min Li and Alex Jinpeng Wang	13467
<i>Learning While Staying Curious: Entropy-Preserving Supervised Fine-Tuning via Adaptive Self-Distillation for Large Reasoning Models</i> Hao Wang, Hao Gu, Hongming Piao, Kaixiong Gong, Yuxiao Ye, Xiangyu Yue, Sirui Han, Yike Guo and Dapeng Wu	13567
<i>Toward Consistent World Models with Multi-Token Prediction and Latent Semantic Enhancement</i> Qimin Zhong, Hao Liao, Haiming Qin, Mingyang Zhou, Rui Mao, Wei Chen and Naipeng Chao	13582
<i>SpiderFlow: Efficient Topology-Aware Scheduling for LLM Training Across Decentralized GPU Clusters</i> Zihan Chang, Shuibing He, Bo Zhou, Sheng Xiao, Siling Yang, Rui Wang and Zhe Pan ...	13603
<i>Unleashing Spatial Reasoning in Multimodal Large Language Models via Textual Representation Guided Reasoning</i> Jiacheng Hua, Yishu Yin, Yuhang Wu, Tai Wang, Yifei Huang and Miao Liu	13616
<i>UniCorn: Towards Self-Improving Unified Multimodal Models through Self-Generated Supervision</i> Zhen Fang, Ruiyan Han, XinYu Sun, Yuchen Ma, Ziheng Wang, Yu Zeng, Zehui Chen, Lin Chen, Wenxuan Huang, Wei-Jie Xu, Yi Cao and Feng Zhao	13638

<i>Why Multimodal In-Context Learning Lags Behind? Unveiling the Inner Mechanisms and Bottlenecks</i> Yu Wang and Sharon Li	13670
<i>Current Agents Fail to Leverage World Model as Tool for Foresight</i> Cheng Qian, Emre Can Acikgoz, Bingxuan Li, Xiusi Chen, Yuji Zhang, Bingxiang He, Qinyu Luo, Gokhan Tur, Dilek Hakkani-Tür, Yunzhu Li and Heng Ji	13686
<i>Language Model as Planner and Formalizer under Constraints</i> Cassie Huang, Stuti Mohan, Ziyi Yang, Stefanie Tellex and Li Zhang	13724
<i>HeLa-Mem: Hebbian Learning and Associative Memory for LLM Agents</i> Jinchang Zhu, Jindong Li, Cheng Zhang, Jiahong Liu and Menglin Yang	13757
<i>Locket: Robust Feature-Locking Technique for Language Models</i> Lipeng He, Vasisht Duddu and N. Asokan	13770
<i>Can Factual Opinions Be Edited (Manipulated) in Large Language Models?</i> Yuanpu Cao, Ziyi Yin, Fenglong Ma and Jinghui Chen	13785
<i>Language of Thought Shapes Output Diversity in Large Language Models</i> Shaoyang Xu and Wenxuan Zhang	13802
<i>Aligning Backchannel and Dialogue Context Representations via Contrastive LLM Fine-Tuning</i> Livia Qian and Gabriel Skantze	13817
<i>Closing the Spatial Execution Gap in Digital Whiteboards via Verifiable Reinforcement Learning</i> Chang Liu, Benjamin Wagley, Zibo Wang, Mehmet E. Belviranli and Bo Wu	13834
<i>ExecVerify: White-Box RL with Verifiable Stepwise Rewards for Code Execution Reasoning</i> Lingxiao Tang, He Ye, Zhaoyang Chu, Muyang Ye, Zhongxin Liu, Xiaoxue Ren and Lingfeng Bao	13850
<i>Prune as You Generate: Online Rollout Pruning for Faster and Better RLVR</i> Haobo Xu, Sirui Chen, Ruizhong Qiu, Yuchen Yan, Chen Luo, Monica Xiao Cheng, Jingrui He and Hanghang Tong	13876
<i>Mitigating Legal Hallucinations via Symbolic Constraints and Analogical Precedents</i> Zixuan Huang, Yanxiang Ma, Luhan Wang, Yunke Wang, Duo Shi and Chang Xu	13894
<i>How Tokenization Limits Phonological Knowledge Representation in Language Models and How to Improve Them</i> Disen Liao and Freda Shi	13921
<i>Illusions of the Gold Standard: A Large-scale Analysis of Human Evaluation Protocols for Long-form Text Generation</i> Katelyn X. Mei, Yi-Li Hsu, Minjoon Choi, Zongwan Cao, Chenjun Xu, Bingbing Wen, Su Lin Blodgett and Lucy Lu Wang	13939
<i>Omni-RewardBench: Toward a Comprehensive Evaluation of Generative Reward Models Across Modalities</i> Chi-Min Chan, Yujin Zhou, Pengcheng Wen, Boqin Yin, Jiaming Ji, Juntao Dai, Wei Xue, Sirui Han and Yike Guo	13962
<i>From Answers to Arguments: Toward Trustworthy Clinical Diagnostic Reasoning with Toulmin-Guided Curriculum Goal-Conditioned Learning</i> Chen Zhan, Xiaoyu Tan, Gengchen Ma, Yu-Jie Xiong, Xiaoyan Jiang and Xihe Qiu	13985

<i>Verifiable LLM-Generated Text Detection via Projected Semantic-Structural Distributions</i>	
Ruochong Xiong, Qien Li, Wangwang Lian, Yulong Wan, Hanlin Xue, Zhouxing Tan, Han Yang, Fengyu Lu and Junfei Liu	14005
<i>Is Human-Like Text Liked by Humans? Multilingual Human Detection and Preference Against AI</i>	
Yuxia Wang, Rui Xing, Jonibek Mansurov, Giovanni Puccetti, Zhuohan Xie, Minh Ngoc Ta, Jiahui Geng, Jinyan Su, Mervat Abassy, Saadeldine Eletter, Kareem Elozeiri, Nurkhan Laiyk, Maiya Goloburda, Tarek Mahmoud, Raj Vardhan Tomar, Alexander Aziz, Ryuto Koike, Masahiro Kaneko, Artem Shelmanov, Ekaterina Artemova, Vladislav Mikhailov, Akim Tsvigun, Alham Fikri Aji, Nizar Habash, Iryna Gurevych and Preslav Nakov	14043
<i>Powerful Training-Free Membership Inference Against Fine-Tuned Autoregressive Language Models</i>	
David Ilić, David Stanojević and Kostadin Cvejovski	14077
<i>Communication-Efficient Desire Alignment for Proactive Embodied Human-Agent Interaction</i>	
Yuanfei Wang, Xinju Huang, Fangwei Zhong, Yaodong Yang, Yizhou Wang, Yuanpei Chen and Hao Dong	14094
<i>When Seeing Overrides Knowing: Disentangling Knowledge Conflicts in Vision-Language Models</i>	
Francesco Ortu, Zhijing Jin, Diego Doimo and Alberto Cazzaniga	14109
<i>ICDAGENT: Empowering Agentic Large Language Models for Explainable Medical Coding</i>	
Ziyi Yin, Yuanpu Cao, Ting Wang, Jinghui Chen and Fenglong Ma	14131
<i>Hyperion: Private Token Sampling with Homomorphic Encryption</i>	
Lawrence Lim, Jiaming Liu, Vikas Kalagi, Divyakant Agrawal and Amr El Abbadi	14150
<i>Map of Encoders – Mapping Sentence Encoders using Quantum Relative Entropy</i>	
Gaifan Zhang and Danushka Bollegala	14160
<i>Mango: Multi-Agent Web Navigation via Global-View Optimization</i>	
Weixi Tong, Yifeng Di and Tianyi Zhang	14209
<i>DPWriter: Reinforcement Learning with Diverse Planning Branching for Creative Writing</i>	
Qian Cao, Yahui Liu, Wei Bi, Yi Zhao, Ruihua Song, Xiting Wang, Ruiming Tang, Guorui Zhou and Han Li	14224
<i>Iterative Dual-Model Alignment for Story Evaluation</i>	
Bruce Qin and Dan Goldwasser	14251
<i>Discovery and Reinforcement of Tool-Integrated Reasoning Chains via Rollout Trees</i>	
Kun Li, Zenan Xu, Junan Li, Zengrui Jin, Jinghao Deng, Zexuan Qiu and Bo Zhou	14265
<i>When Identity Skews Debate: Anonymization for Bias-Reduced Multi-Agent Reasoning</i>	
Hyeong Kyu Choi, Jerry Zhu and Sharon Li	14284
<i>Talking to a Know-It-All GPT or a Second-Guesser Claude? How Repair reveals distinct Multi-Turn Behavior in LLMs</i>	
Clara Lachenmaier, Hannah Bultmann and Sina Zarriëß	14312
<i>HistLens: Mapping Idea Change across Concepts and Corpora</i>	
Yi Jing, Weiyun Qiu, Yihang Peng and Zhifang Sui	14326
<i>ReFEree: Reference-Free and Fine-Grained Method for Evaluating Factual Consistency in Real-World Code Summarization</i>	
Suyoung Bae, CheolWon Na, Jaehoon Lee, Yumin Lee, YunSeok Choi and Jee-Hyong Lee	14352

<i>MeanAudio: Fast and Faithful Text-to-Audio Generation with Mean Flows</i>	
Xiquan Li, Junxi Liu, Yuzhe Liang, Zhikang Niu, Wenxi Chen and Xie Chen	14378
<i>ModeX: Evaluator-Free Best-of-N Selection for Open-Ended Generation</i>	
Hyeong Kyu Choi and Sharon Li	14394
<i>What Makes LLMs Effective Sequential Recommenders? A Study on Preference Intensity and Temporal Context</i>	
Zhongyu Ouyang, Qianlong Wen, Chunhui Zhang, Yanfang Ye and Soroush Vosoughi	14417
<i>TeamFusion: Supporting Open-ended Teamwork with Multi-Agent Systems</i>	
Jiale Liu, Victor Bursztyn, Lin Ai, Haoliang Wang, Sunav Choudhary, Saayan Mitra and Qingyun Wu	14435
<i>FormulaSPIN: Self-Play Fine-Tuning for Natural Language to Spreadsheet Formula Generation</i>	
Cy Xie	14457
<i>Stable Signer: Hierarchical Sign Language Generative Model</i>	
Sen Fang, Yalin Feng, Hongbin Zhong, Yanxin Zhang and Dimitris N. Metaxas	14472
<i>DisCal: Distribution-Aware Calibration for Mathematical Reasoning Under Character-Level Noisy Inputs</i>	
Bo Zhang, Jiawei Zhang, Cong Gao, Bingxu Han, Minghao Hu, Jun Zhang, Yunbo Cao, Zhunchen Luo, Wen Yao, Guotong Geng and Zhong Wang	14484
<i>SARA: Selective and Adaptive Retrieval-augmented Generation with Context Compression</i>	
Yiqiao Jin, Kartik Sharma, Vineeth Rakesh, Yingtong Dou, Menghai Pan, Mahashweta Das and Srijan Kumar	14508
<i>FinChain: A Symbolic Benchmark for Verifiable Chain-of-Thought Financial Reasoning</i>	
Zhuohan Xie, Daniil Orel, Rushil Thareja, Dhruv Sahnan, Hachem Madmoun, Fan Zhang, Debopriyo Banerjee, Georgi Nenkov Georgiev, Xueqing Peng, Lingfei Qian, Jimin Huang, Jinyan Su, Aaryamonvikram Singh, Rui Xing, Rania Elbadry, Chen Xu, Haonan Li, Fajri Koto, Ivan Koychev, Tanmoy Chakraborty, Yuxia Wang, Salem Lahlou, Veselin Stoyanov, Sophia Ananiadou and Preslav Nakov	14529
<i>AI use in American newspapers is widespread, uneven, and rarely disclosed</i>	
Jenna Russell, Marzena Karpinska, Destiny Akinode, James Zhou, Katherine Thai, Bradley Emi, Max Spero and Mohit Iyyer	14554
<i>For-Value: Efficient Forward-Only Data Valuation for finetuning LLMs and VLMs</i>	
Wenlong Deng, Qi Zeng, Jiaming Zhang, Minghui Chen, Zixin Ding, Christos Thrampoulidis, Boying Gong and Xiaoxiao Li	14581
<i>The Illusion of Specialization: Unveiling the Domain-Invariant Standing Committee in Mixture-of-Experts Models</i>	
Yan Wang, Yitao Xu, Nanhan Shen, Jinyan Su, Jimin Huang and Zining Zhu	14601
<i>Find Your Optimal Teacher: Personalized Data Synthesis via Router-Guided Multi-Teacher Distillation</i>	
Hengyuan Zhang, Shiping Yang, Xiao Liang, Chenming Shang, Yuxuan Jiang, Chaofan Tao, Jing Xiong, Hayden Kwok-Hay So, Ruobing Xie, Angel X Chang and Ngai Wong	14619
<i>Single-Pass, Depth-Selective Reading for Multi-Aspect Sentiment Analysis</i>	
Yan Xia, Zhuangzhuang Pan, Amirrudin Kamsin and Chee Seng Chan	14638
<i>PARASITE: Conditional System Prompt Poisoning to Hijack LLMs</i>	
Viet Pham and Thai Le	14657

<i>Generating Literature-Driven Scientific Theories at Scale</i>	
Peter Jansen, Peter Clark, Doug Downey and Daniel S Weld	14677
<i>Controllable Memory Usage: Balancing Anchoring and Innovation in Long-Term Human-Agent Interaction</i>	
Zisu Huang, Muzhao Tian, Xiaohua Wang, Jingwen Xu, Zhengkang Guo, Qi Qian, Kaitao Song, Jiakang Yuan, Changze Lv and Xiaoqing Zheng	14699
<i>Digitizing Nepal’s Written Heritage: A Comprehensive HTR Pipeline for Old Nepali Manuscripts</i>	
Anjali Sarawgi, Esteban Garces Arias and Christof Zotter	14720
<i>XMark: Reliable Multi-Bit Watermarking for LLM-Generated Texts</i>	
Jiahao Xu, Rui Hu, Olivera Kotevska and Zikai Zhang	14747
<i>Believing without Seeing: Quality Scores for Contextualizing Vision-Language Model Explanations</i>	
Keyu He, Tejas Srinivasan, Brihi Joshi, Xiang Ren, Jesse Thomason and Swabha Swayamdipta	14764
<i>MONETA: Multimodal Industry Classification through Geographic Information with Multi Agent Systems</i>	
Arda Yüksel, Gabriel Thiem, Susanne Walter, Patrick Felka, Gabriela Alves Werb and Ivan Habernal	14790
<i>CASPER in the Machine: Insights into Character Variety in LLM-Generated Stories</i>	
Anneliese Brei, Abhishek Sharma, Nicholas Sanaie, Lu Wang and Snigdha Chaturvedi ...	14815
<i>Language Models Struggle to Use Representations Learned In-Context</i>	
Michael A. Lepori, Tal Linzen, Ann Yuan and Katja Filippova	14841
<i>SlideAgent: Hierarchical Agentic Framework for Multi-Page Visual Document Understanding</i>	
Yiqiao Jin, Rachneet Kaur, Zhen Zeng, Sumitra Ganesh and Srijan Kumar	14858
<i>YIELD: A Large-Scale Dataset and Evaluation Framework for Information Elicitation Agents</i>	
Victor De Lima and Grace Hui Yang	14882
<i>Spatial-Agent: Agentic Geo-spatial Reasoning with Scientific Core Concepts</i>	
Riyang Bao, Cheng Yang, Dazhou Yu, Zhexiang Tang, Gengchen Mai and Liang Zhao ...	14896
<i>Rethinking Evaluation for LLM Hallucination Detection: A Desiderata, A New RAG-based Benchmark, New Insights</i>	
Wenbo Chen, Veena Padmanabhan, Tootiya Giyahchi, Elaine Wong and Leman Akoglu ..	14912
<i>Min-k Sampling: Decoupling Truncation from Temperature Scaling via Relative Logit Dynamics</i>	
Yuanhao Ding, Meimingwei Li, Esteban Garces Arias, Matthias Aßenmacher, Christian Heumann and Chongsheng Zhang	14932
<i>Bridging the Memorization-Utilization Gap: Near-Lossless Context Compression via Reinforcement Learning</i>	
Yujan Ting, Xu Tang, Terrence Chen and Weijing Huang	14949
<i>PRIME: A Process-Outcome Alignment Benchmark for Verifiable Reasoning in Mathematics and Engineering</i>	
Xiangfeng Wang, Hangyu Guo, Yanlin Lai, Mitt Huang, Liang Zhao, Chengyuan Yao, Yinmin Zhang, Qi Han, Xiaoxiaoren, Chun Yuan, Tong Xu, Zheng Ge, Xiangyu Zhang and Daxin Jiang	14973
<i>ZARA: Training-Free Motion Time-Series Reasoning via Evidence-Grounded LLM Agents</i>	
Zechen Li, Baiyu Chen, Hao Xue and Flora D. Salim	14986

<i>AutoJuder: An Agent-Driven Framework for Efficient Benchmarking of MLLMs</i>	
Xuanwen Ding, Chengjun Pan, Zejun Li, Jiwen Zhang, Siyuan Wang and Zhongyu Wei . . .	15009
<i>Implicit Representations of Grammaticality in Language Models</i>	
Yingshan Susan Wang, Linlu Qiu, Zhaofeng Wu, Roger P. Levy and Yoon Kim	15035
<i>Discovering and Causally Validating Emotion-Sensitive Neurons in Large Audio-Language Models</i>	
Xiutian Zhao, Björn Schuller and Berrak Sisman	15056
<i>Uncovering Intervention Opportunities for Suicide Prevention with Language Model Assistants</i>	
Jaspreet Ranjit, Hyundong Justin Cho, Claire J. Smerdon, Yoonsoo Nam, Myles Phung, Jonathan May, John R. Blosnich and Swabha Swayamdipta	15072
<i>The Imperfective Paradox in Large Language Models</i>	
Bolei Ma and Yusuke Miyao	15093
<i>MalruleLib: Large-Scale Executable Misconception Reasoning with Step Traces for Modeling Student Thinking in Mathematics</i>	
Xinghe Chen, Naiming Liu and Shashank Sonkar	15112
<i>How to Improve LLMs' Performance on Specific Languages: A Perspective on LLM-Derived Language Similarity</i>	
Xinhe Shi, Qingcheng Zeng, Weihao Xuan and Linchao Zhu	15139
<i>MATCH: Modulating Attention via In-Context Retrieval for Long-Context Transformers</i>	
Linrui Ma, Chun Hei Lo, Xinyu Wang, Peng Lu, Xihao Yuan, Hanting Chen, Kai Han, Xinghao Chen, Chengjun Zhan, Hanlin xu, Yichun Yin, Lifeng Shang, Feng Wen, Boxing Chen and Yufei Cui	15165
<i>Protecting Bystander Privacy via Selective Hearing in Audio LLMs</i>	
Xiao Zhan, Guangzhi Sun, Jose Such and Phil Woodland	15180
<i>LiveCLKTBench: Towards Reliable Evaluation of Cross-Lingual Knowledge Transfer in Multilingual LLMs</i>	
Pei-Fu Guo, Yun-Da Tsai, Chun-Chia Hsu, Kai-Xin Chen, Ya An Tsai, Kai-Wei Chang, Nanyun Peng, Mi-Yen Yeh and Shou-De Lin	15193
<i>To Lie or Not to Lie? Investigating The Biased Spread of Global Lies by LLMs</i>	
Zohaib Khan, Mustafa Dogan, Ifeoma Okoh, Pouya Sadeghi, Siddhartha Shrestha, Sergius Justus Chesami Nyah, Mahmoud O. Mokhiamar, Michael J Ryan and Tarek Naous	15209
<i>Robertha: Eigenspectrum Regularized Attention for Robust Natural Language Understanding</i>	
Andreia Podasca and Anup Das	15229
<i>Agentic Rubrics as Contextual Verifiers for SWE Agents</i>	
Mohit Raghavendra, Anisha Gunjal, Bing Liu and Yunzhong He	15265
<i>Evaluating the Impact of Verbal Multiword Expressions on Machine Translation</i>	
Linfeng Liu, Saptarshi Ghosh and Tianyu Jiang	15291
<i>MedVerse: Efficient and Reliable Medical Reasoning via DAG-Structured Parallel Execution</i>	
Jianwen Chen, Xinyu Yang, Peng Xia, Arian Azarang, Yueh Z Lee, Gang Li, Hongtu Zhu, Yun Li, Beidi Chen and Huaxiu Yao	15320
<i>MetFuse: Figurative Fusion between Metonymy and Metaphor</i>	
Saptarshi Ghosh and Tianyu Jiang	15337

<i>Adaptive Constraint Propagation: Scaling Structured Inference for Large Language Models via Meta-Reinforcement Learning</i>	
Ibne Farabi Shihab, Sanjeda Akter and Anuj Sharma	15351
<i>STReasoner: Empowering LLMs for Spatio-Temporal Reasoning in Time Series via Spatial-Aware Reinforcement Learning</i>	
Juntong Ni, Shiyu Wang, Qi He, Ming Jin and Wei Jin.....	15369
<i>When Vision-Language Models Judge Without Seeing: Exposing Informativeness Bias</i>	
Xiaohan Zou, Roshan Sridhar, Mohammadtaher Safarzadeh and Dan Roth	15417
<i>Hallucination Detection in LLMs with Topological Divergence on Attention Graphs</i>	
Alexandra Bazarova, Andrei Volodichev, Aleksandr Yugay, Andrey Shulga, Alina Ermilova, Konstantin Polev, Julia Belikova, Rauf Parchiev, Dmitry Simakov, Maxim Savchenko, Andrey Savchenko, Serguei Barannikov and Alexey Zaytsev	15449
<i>Exploring Concreteness Through a Figurative Lens</i>	
Saptarshi Ghosh and Tianyu Jiang	15471
<i>Mediocrity is the key for LLM as a Judge Anchor Selection</i>	
Shachar Don-Yehiya, Asaf Yehudai, Leshem Choshen and Omri Abend	15491
<i>When CorrectIs Not Safe: Can We Trust Functionally Correct Patches Generated by Code Agents?</i>	
Yibo Peng, James Song, Lei Li, Xinyu Yang, Mihai Christodorescu, Ravi Mangal, Corina S. Pasareanu, Haizhong Zheng and Beidi Chen	15514
<i>CoG: Controllable Graph Reasoning via Relational Blueprints and Failure-Aware Refinement over Knowledge Graphs</i>	
Yuanxiang Liu, Songze Li, Xiaoke Guo, Zhaoyan Gong, Qifei Zhang, Huajun Chen and Wen Zhang	15547
<i>Latent Agents: A Post-Training Procedure for Internalized Multi-Agent Debate</i>	
John Seon Keun Yi, Aaron Mueller and Dokyun Lee	15569
<i>Act-Adaptive Margin: Dynamically Calibrating Reward Models for Subjective Ambiguity</i>	
Feiteng Fang, Dingwei Chen, Xiang Huang, Ting-En Lin, Yuchuan Wu, Xiong Liu, Jing Ye, Ziqiang Liu, Haonan Zhang, Liang Zhu, Hamid Alinejad-Rokny, Min Yang and Yongbin Li....	15603
<i>To Think or Not to Think: The Hidden Cost of Meta-Training with Excessive CoT Examples</i>	
Vignesh Kothapalli, Ata Fatahibaarzi, Hamed Firooz and Maziar Sanjabi	15619
<i>VIGNETTE: Socially Grounded Bias Evaluation for Vision-Language Models</i>	
Chahat Raj, Bowen Wei, Aylin Caliskan, Antonios Anastasopoulos and Ziwei Zhu	15645
<i>BiMind: A Dual-Head Reasoning Model with Attention-Geometry Adapter for Incorrect Information Detection</i>	
Zhongxing Zhang, Emily K. Vraga, Jisu Huh and Jaideep Srivastava	15674
<i>Autonomous Knowledge Graph Exploration with Adaptive Breadth-Depth Retrieval</i>	
Joaquin Polonuer, Lucas Vittor, Iñaki Arango, Ayush Noori, David A. Clifton, Luciano Del Corro and Marinka Zitnik	15694
<i>Retrieval Heads are Dynamic</i>	
Yuping Lin, Zitao Li, Yue Xing, Pengfei He, Yingqian Cui, Yaliang Li, Bolin Ding, Jingren Zhou and Jiliang Tang	15710

<i>SynthAgent: Adapting Web Agents with Synthetic Supervision</i>	
Zhaoyang Wang, Yiming Liang, Xuchao Zhang, Qianhui Wu, Siwei Han, Anson Bastos, Rujia Wang, Chetan Bansal, Baolin Peng, Jianfeng Gao, Saravan Rajmohan and Huaxiu Yao	15730
<i>More Than Meets the Eye: Measuring the Semiotic Gap in Vision-Language Models via Semantic Anchorage</i>	
Wei He	15753
<i>LLM Agents in Law: Taxonomy, Applications, and Challenges</i>	
Shuang Liu, Ruijia Zhang, Ruoyun Ma, Yujia Deng, Lanyi Zhu, Jiayu Li, Zelong Li, Zhibin Shen and Mengnan Du	15768
<i>BenchMarker: An Education-Inspired Toolkit for Highlighting Flaws in Multiple-Choice Benchmarks</i>	
Nishant Balepur, Bhavya Rajasekaran, Hyunjin Jane Oh, Michael Xie, Atrey Desai, Vipul Gupta, Steven James Moore, Eunsol Choi, Rachel Rudinger and Jordan Lee Boyd-Graber	15793
<i>Model Internal Sleuthing: Finding Lexical Identity and Inflectional Features in Modern Language Models</i>	
Michael Li and Nishant Subramani	15825
<i>HAT: Hallucination Annotation for Translation</i>	
Rajen Chatterjee, Xintong Li, Paisarn Charoenpornasawat and Allen Lee	15865
<i>DMRetriever: A Family of Models for Improved Text Retrieval in Disaster Management</i>	
Kai Yin, Xiangjue Dong, Chengkai Liu, Allen Lin, Lingfeng Shi, Ali Mostafavi and James Caverlee	15889
<i>Language Models Don't Know What You Want: Evaluating Personalization in Deep Research Needs Real Users</i>	
Nishant Balepur, Malachi Hamada, Varsha Kishore, Sergey Feldman, Amanpreet Singh, Pao Siangliulue, Joseph Chee Chang, Eunsol Choi, Jordan Lee Boyd-Graber and Aakanksha Naik . .	15910
<i>LAD-RAG: Layout-aware Dynamic RAG for Visually-Rich Document Understanding</i>	
Zhivar Sourati, Zheng Wang, Marianne Menglin Liu, Yazhe Hu, Mengqing Guo, Sujeeth Bharadwaj, Kyu J. Han, Tao Sheng, Sujith Ravi, Morteza Dehghani and Dan Roth	15945
<i>Among Us: Measuring and Mitigating Malicious Contributions in Model Collaboration Systems</i>	
Ziyuan Yang, Wenxuan Ding, Shangbin Feng and Yulia Tsvetkov	15969
<i>Not All Tokens Matter: Towards Efficient LLM Reasoning via Token Significance in Reinforcement Learning</i>	
Hanbing Liu, Lang Cao, Yuanyi Ren, Mengyu Zhou, Haoyu Dong, Xiaojun Ma, Shi Han and Dongmei Zhang	15989
<i>Query-Efficient Agentic Graph Extraction Attacks on GraphRAG Systems</i>	
Shuhua Yang, Jiahao Zhang, Yilong Wang, Dongwon Lee and Suhang Wang	16017
<i>Function Words as Statistical Cues for Language Learning</i>	
Xiulin Yang, Heidi R. Getz and Ethan Gotlieb Wilcox	16042
<i>Evaluating Robustness of Large Language Models Against Multilingual Typographical Errors</i>	
Raoyuan Zhao, Yihong Liu, Lena Altinger, Hinrich Schuetze and Michael A. Hedderich . .	16059
<i>Listening Like Humans: Semantics-Guided Noise-Robust Multimodal Speech Recognition</i>	
Yan Fang, Jun Chen, Yian Yao, Shuxin Zhong, Min Sun and Kaishun Wu	16079

<i>Measuring What Matters!! Assessing Therapeutic Principles in Mental-Health Conversation</i>	
Abdullah Mazhar, Het Riteshkumar Shah, Aseem Srivastava, Smriti Joshi and Md Shad Akhtar	
	16094
<i>ParaCodex: A Profiling-Guided Autonomous Coding Agent for Reliable Parallel Code Generation and Translation</i>	
Erel Kaplan, Tomer Bitan, Lian Ghrayeb, Le Chen, Tom Yotam, Niranjan Hasabnis and Gal Oren	
	16113
<i>CoLA: A Choice Leakage Attack Framework to Expose Privacy Risks in Subset Training</i>	
Qi Li, Cheng-Long Wang, Yinzhi Cao and Di Wang	16137
<i>MediEval: A Unified Medical Benchmark for Patient-Contextual and Knowledge-Grounded Reasoning in LLMs</i>	
Zhan Qu and Michael Färber	16150
<i>The Path Not Taken: Duality in Reasoning about Program Execution</i>	
Eshgin Hasanov, Md. Mahadi Hassan, Santu Karmaker and Aashish Yadavally	16165
<i>Accommodation and Epistemic Vigilance: A Pragmatic Account of Why LLMs Fail to Challenge Harmful Beliefs</i>	
Myra Cheng, Robert D. Hawkins and Dan Jurafsky	16181
<i>Every Response Counts: Quantifying Uncertainty of LLM-based Multi-Agent Systems through Tensor Decomposition</i>	
Tiejun Chen, Huaiyuan Yao, Jia Chen, Evangelos E. Papalexakis and Hua Wei	16204
<i>Uncertainty Quantification in LLM Agents: Foundations, Emerging Challenges, and Opportunities</i>	
Changdae Oh, Seongheon Park, To Eun Kim, Jiatong Li, Wendi Li, Samuel Yeh, Sean Du, Hamed Hassani, Paul Bogdan, Dawn Song and Sharon Li	16219
<i>Guided by Gut: Efficient Test-Time Scaling with Reinforced Intrinsic Confidence</i>	
Amirhosein Ghasemabadi, Keith G. Mills, Baochun Li and Di Niu	16251
<i>Knowing When Not to Answer: Lightweight KB-Aligned OOD Detection for Safe RAG</i>	
Ilias Triantafyllopoulos, Renyi Qu, Salvatore Giorgi, Brenda Curtis, Lyle Ungar and João Sedoc	
	16266
<i>Guidelines as Environments: A World Model Approach to Rule Following</i>	
Haiqing Li, Wenliang Zhong, Yin hao Wu, Hehuan Ma, Yuzhi Guo, Thao M. Dang and Junzhou Huang	
	16302
<i>Cross-Modal Taxonomic Generalization in (Vision-) Language Models</i>	
Tianyang Xu, Marcelo Sandoval-Castañeda, Karen Livescu, Greg Shakhnarovich and Kanishka Misra	
	16319
<i>Biomedical Question Answering via Multi-Level Summarization on a Local Knowledge Graph</i>	
Lingxiao Guan, Yuanhao Huang and Jie Liu	
	16338
<i>Sigmoid Head for Quality Estimation under Language Ambiguity</i>	
Tu Anh Dinh and Jan Niehues	
	16365
<i>RExBench: Can coding agents autonomously implement AI research extensions?</i>	
Nicholas Edwards, Yukyung Lee, Yujun Audrey Mao, Yulu Qin, Sebastian Schuster and Najoung Kim	
	16380

<i>Narrative License and Model Sycophancy in LLM Summaries of Scientific Work</i> Calvin Isch and Grace Jennings	16418
<i>AlignUSER: Human-Aligned LLM Agents via World Models for Recommender System Evaluation</i> Nicolas Bougie, Gian Maria Marconi, Xiaotong Ye and Narimawa Watanabe	16433
<i>Domain Generalizable AI Guardrails with Augmented Policy Training</i> Minqian Liu, Ioana Baldini, David Rabinowitz, David S Rosenberg, Sebastian Gehrmann and Mark Dredze	16452
<i>APEX-MEM: Agentic Semi-Structured Memory with Temporal Reasoning for Long-Term Conversational AI</i> Pratyay Banerjee, Masud Moshtaghi, Shivashankar Subramanian, Amita Misra and Ankit Chadha	16470
<i>emg2speech: synthesizing speech from electromyography using self-supervised speech models</i> Harshvardhana T Gowda, Daniel C Comstock and Lee M. Miller	16490
<i>CoAct: Co-Active LLM Preference Learning with Human-AI Synergy</i> Ruiyao Xu, Mihir Parmar, Tiankai Yang, Zhengyu Hu, Yue Zhao and Kaize Ding	16508
<i>Label and Explanation Variation in LLM-Based Annotation: a Case Study in Natural Language Inference</i> Artur Kulmizev, Erika Lombart, Patrick Watrin and Marie-Catherine de Marneffe	16526
<i>Text-Attributed Knowledge Graph Enrichment with Large Language Models for Medical Concept Representation</i> Mohsen Nayebi Kerdabadi, Arya Hadizadeh Moghaddam, Chen Chen, Dongjie Wang and Zijun Yao	16544
<i>Evolving Agents</i> Leonardo Ranaldi	16561
<i>ViDoRe V3: A Comprehensive Evaluation of Retrieval Augmented Generation in Complex Real-World Scenarios</i> António Loison, Quentin Macé, Antoine Edy, Victor Xing, Tom Balough, Gabriel de Souza P. Moreira, Bo Liu, Manuel Faysse, Celine Hudelot and Gautier Viaud	16570
<i>Why LLM Safety Guardrails Collapse After Fine-tuning: A Similarity Analysis Between Alignment and Fine-tuning Datasets</i> Lei Hsiung, Tianyu Pang, Yung-Chen Tang, Linyue Song, Tsung-Yi Ho, Pin-Yu Chen and Yaoqing Yang	16601
<i>TabEmb: Joint Semantic-Structure Embedding for Table Annotation</i> Ehsan Hoseinzade, Ke Wang and Anandharaju Durai Raju	16620
<i>Text-to-Distribution Prediction with Quantile Tokens and Neighbor Context</i> Yilun Zhu, Yuan Zhuang, Nikhita Vedula, Dushyanta Dhyani, Shaoyuan Xu, Mohsen Bayati, Bryan Wang and Shervin Malmasi	16632
<i>Check Your Work: Structured Checklist Feedback for Improving Large Language Models</i> Jonathan Cook, Tim Rocktäschel, Jakob Nicolaus Foerster, Dennis Aumiller and Alex Wang	16649
<i>LOKA: Conflict-Aware LLM Knowledge Update with Adaptive Knowledge Memory</i> Binchi Zhang, Zhengzhang Chen, Zaiyi Zheng, Jundong Li and Haifeng Chen	16689

<i>Student Guides Teacher: Weak-to-Strong Inference via Spectral Orthogonal Exploration</i>	
Dayu Wang, Jiaye Yang, Weikang Li, Jiahui Liang, Yang Li, Deguo Xia and Jizhou Huang	16716
<i>When LLMs Read Tables Carelessly: Measuring and Reducing Data Referencing Errors</i>	
Yuqing Yang, Qi Zhu, Zhen Han, Boran Han, Zhengyuan Shen, Shuai Wang, Vassilis N. Ioannidis and Huzefa Rangwala	16734
<i>PubMed Reasoner: Dynamic Reasoning-based Retrieval for Evidence-Grounded Biomedical Question Answering</i>	
Yiqing Zhang, Xiaozhong Liu and Fabricio Murai	16753
<i>Decoupling Task-Solving and Output Formatting in LLM Generation</i>	
Haikang Deng, Po-Nien Kung and Nanyun Peng	16764
<i>GS-Quant: Granular Semantic and Generative Structural Quantization for Knowledge Graph Completion</i>	
Qizhuo Xie, Yunhui Liu, Yu Xing, Qianzi Hou, Xudong Jin, Tao Zheng and Tieke He	16782
<i>Mind the Gap in Cultural Alignment: Task-Aware Culture Management for Large Language Models</i>	
Binchi Zhang, Xujiang Zhao, Jundong Li, Haifeng Chen and Zhengzhang Chen	16798
<i>Analyzing and Internalizing Complex Policy Documents for LLM Agents</i>	
Jiateng Liu, Zhenhailong Wang, Xiaojiang Huang, Yingjie Li, Xiang Li, Chenlei Guo, Xing Fan, Ruhi Sarikaya and Heng Ji	16814
<i>In-Context Representation Hijacking</i>	
Itay Yona, Amir Sarid, Michael Karasik and Yossi Gandelsman	16852
<i>Quantifying Metric and Model Agreement in Bias Evaluation of Large Language Models</i>	
Arash Asgari, Huan Wu, Amirreza Naziri, Mojtaba Kolahdouzi and Laleh Seyyed-Kalantari	16868
<i>MultiFinBen: Benchmarking Large Language Models for Multilingual and Multimodal Financial Application</i>	
Xueqing Peng, Lingfei Qian, Yan Wang, Ruoyu Xiang, Yueru He, Yang Ren, Mingyang Jiang, Vincent Jim Zhang, Yuqing Guo, Jeff Zhao, Huan He, Yi Han, Yun Feng, Yuechen Jiang, Yupeng Cao, Haohang Li, Yangyang Yu, Xiaoyu Wang, Penglei Gao, Shengyuan Lin, Keyi Wang, Shanshan Yang, Yilun Zhao, Zhiwei Liu, Peng Lu, Jerry Huang, Suyuchen Wang, Triantafillos Papadopoulos, Polydoros Giannouris, Efstathia Soufleri, Nuo Chen, Zhiyang Deng, Heming Fu, Yijia Zhao, Mingquan Lin, Meikang Qiu, Kaleb E Smith, Arman Cohan, Xiao-Yang Liu, Jimin Huang, Guojun Xiong, Alejandro Lopez-Lira, Xi Chen, Junichi Tsujii, Jian-Yun Nie, Sophia Ananiadou and Qianqian Xie	16934
<i>Bidirectional LMs are Better Knowledge Memorizers? A Benchmark for Real-world Knowledge Injection</i>	
Yuwei Zhang, Wenhao Yu, Shangbin Feng, Yifan Zhu, Letian Peng, Jayanth Srinivasa, Gaowen Liu and Jingbo Shang	16964
<i>Mapping the Circumplex of Affect: Geometric Analysis of Emotion Representations via Hyperspherical Contrastive Learning</i>	
Yusuke Yamauchi and Akiko Aizawa	16981
<i>Do LLMs Really Need 10+ Thoughts for “Find the Time 1000 Days Later”? Towards Structural Understanding of LLM Overthinking</i>	
Xinliang Frederick Zhang, Anhad Mohanane, Alexandra Chronopoulou, Pinelopi Papalampidi, Somit Gupta, Tsendsuren Munkhdalai, Lu Wang and Shyam Upadhyay	17005
<i>Anchor: Branch-Point Data Generation for GUI Agents</i>	
Jinbiao Wei, Yilun Zhao, Kangqi Ni and Arman Cohan	17031

<i>When One LLM Drools, Multi-LLM Collaboration Rules</i>	
Shangbin Feng, Wenxuan Ding, Alisa Liu, Zifeng Wang, Weijia Shi, Yike Wang, Shannon Zejiang Shen, Xiaochuang Han, Hunter Lang, Chen-Yu Lee, Tomas Pfister, Yejin Choi and Yulia Tsvetkov	17048
<i>From Query to Counsel: Structured Reasoning with a Multi-Agent Framework and Dataset for Legal Consultation</i>	
Mingfei Lu, Yi Zhang, Mengjia Wu and Yue Feng	17064
<i>Putting Captions to the Test: Evaluating Video Caption Quality through Multiple-Choice Question Answering</i>	
Zizhen Wang, Bo Feng, Zhengfeng Lai, Shiyu Li, Yang Lu, Meng Cao, Ping Huang and Xiaoming Simon Wang	17079
<i>Subject-level Inference for Realistic Text Anonymization Evaluation</i>	
Myeong Seok Oh, Dong-Yun Kim, Hanseok Oh, Chaeon Kang, Joeun Kang, Xiaonan Wang, Hyunjung Park, Young Cheol Jung and Hansaem Kim	17100
<i>SceneAlign: Aligning Multimodal Reasoning to Scene Graphs in Complex Visual Scenes</i>	
Chuhan Wang, Xintong Li, Jennifer Yuntong Zhang, Junda Wu, Chengkai Huang, Lina Yao, Julian McAuley and Jingbo Shang	17136
<i>Select Before Use: On the Importance of Reference Model Selection in Preference Alignment</i>	
Muyang Li, Runze Wu, Xiangyu Zhao, Bo Han, Daoyi Dong and Tongliang Liu	17151
<i>Lost in Translation, and Found: Detecting and Interpreting Translation Effects</i>	
Shira Wein, Anna Serbina, Jiyuan Ji, Nathan Wolf, Jason DeGraaff, Prajakta Kini and Maria Leonor Pacheco	17172
<i>From Isolation to Entanglement: When Do Interpretability Methods Identify and Disentangle Known Concepts?</i>	
Aaron Mueller, Andrew Lee, Shruti Joshi, Ekdeep Singh Lubana, Dhanya Sridhar and Patrik Reizinger	17188
<i>Paraphrasing as Zero-shot Translation with Feature-guided Diversity Enhancement</i>	
Ziyue Yan, Hongying Zan, Xinglin Lyu and Hongfei Xu	17211
<i>VOYAGER: A Training Free Approach for Generating Diverse Datasets using LLMs</i>	
Avinash Amballa, Yashas Malur Saidutta, Chi-Heng Lin, Vivek Kulkarni and Srinivas Chappidi	17224
<i>CaseFacts: A Benchmark for Legal Fact-Checking and Precedent Retrieval</i>	
Akshith Reddy Putta, Jacob Devasier and Chengkai Li	17246
<i>SpecAgent: A Speculative Retrieval and Forecasting Agent for Code Completion</i>	
George Ma, Anurag Koul, Qi Chen, Yawen Wu, Sachit Kuhar, Yu Yu, Aritra Sengupta, Varun Kumar and Murali Krishna Ramanathan	17266
<i>From Inheritance to Saturation: Disentangling the Evolution of Visual Redundancy for Architecture-Aware MLLM Inference Acceleration</i>	
Jiaqi Shi, Xulong Zhang, Yuechan Li, Xiaoyang Qu and Jianzong Wang	17328
<i>Your Reasoning Model is Secretly a Reward Model - Optimization-Free Verification from Experience</i>	
Zhenwen Liang, Ruosen Li, Yujun Zhou, Linfeng Song, Dian Yu, Xinya Du, Haitao Mi and Dong Yu	17358

<i>What Do LLMs Learn First? Asymmetric Learning Dynamics of Input Complexity and Output Ambiguity in Preference Alignment</i>	Mengyang Li, Jingwen Wang and Pinlong Zhao	17373
<i>Multi-Agent-as-Judge: Aligning LLM-Agent-Based Automated Evaluation with Multi-Dimensional Human Evaluation</i>	Jiaju Chen, Yuxuan Lu, Xiaojie Wang, Huimin Zeng, Jing Huang, Jiri Gesi, Ying Xu, Bingsheng Yao and Dakuo Wang	17389
<i>OpenRubrics: Towards Scalable Synthetic Rubric Generation for Reward Modeling and LLM Alignment</i>	Tianci Liu, Ran Xu, Tony Yu, Ilgee Hong, Carl Yang, Tuo Zhao and Haoyu Wang	17417
<i>Challenging the Boundaries of Reasoning: An Olympiad-Level Math Benchmark for Large Language Models</i>	Haoxiang Sun, Yingqian Min, Zhipeng Chen, Xin Zhao and Ji-Rong Wen	17438
<i>LENS: LLM-Enabled Narrative Synthesis for Mental Health by Aligning Multimodal Sensing with Language Models</i>	Wenxuan Xu, Arvind Pillai, Subigy Nepal, Amanda C. Collins, Daniel M Mackin, Michael V. Heinz, Tess Z Griffin, Nicholas C. Jacobson and Andrew Campbell	17458
<i>FedProxy: Federated Fine-Tuning of LLMs via Proxy SLMs and Heterogeneity-Aware Fusion</i>	Tao Fan, Guoqiang Ma, Yuanfeng Song, Lixin Fan, Kai Chen and Qiang Yang	17482
<i>DiZiNER: Disagreement-guided Instruction Refinement via Simulating Pilot Annotation for Zero-shot Named Entity Recognition</i>	Siun Kim and Hyung-Jin Yoon	17498
<i>A Data-Centric Approach to Generalizable Speech Deepfake Detection</i>	Wen Huang, Yuchen Mao and Yanmin Qian	17520
<i>Free Energy-Driven Reinforcement Learning with Adaptive Advantage Shaping for Unsupervised Reasoning in LLMs</i>	Yiming Huang, Zhenbo Shi, Xin-Cheng Wen, Jichuan Zeng, Cuiyun Gao, Peiyi Han and Chuanyi Liu	17540
<i>PaperRegister: Boosting Flexible-grained Paper Search via Hierarchical Register Indexing</i>	Zhuoqun Li, Xuanang Chen, Hongyu Lin, Yaojie Lu, Xianpei Han, Shanshan Jiang, Bin Dong and Le Sun	17561
<i>ETR: Entropy Trend Reward for Efficient Chain-of-Thought Reasoning</i>	Xuan Xiong, Huan Liu, Li Gu, Zhixiang Chi, Yue Qiu, Yuanhao YU and Yang Wang	17576
<i>All Changes May Have Invariant Principles: Improving Ever-Shifting Harmful Meme Detection via Design Concept Reproduction</i>	Ziyong Jiang, Mingyang Li, Junjie Wang, Yuekai Huang, Jie Huang, Zhiyuan Chang, Zhaoyang Li and Qing Wang	17595
<i>Activation-Guided Local Editing for Jailbreaking Attacks</i>	Jiecong Wang, Haoran Li, Hao Peng, Ziqian Zeng, Zihao Wang, Haohua Du and Zhengtao Yu	17614
<i>Controlling Distributional Bias in Multi-Round LLM Generation via KL-Optimized Fine-Tuning</i>	Yanbei Jiang, Amr Keleg, Ryandito Diandaru, Jey Han Lau, Lea Frermann, Biaoyan Fang and Fajri Koto	17634

<i>On the Continued Value of Universal Dependencies in the Era of Large Language Models</i> Wenxi Li and Jingyu Peng	17650
<i>Adaptive Text2GQL: Integrating Structural Twig Linking and Evolutionary In-Context Learning</i> Fang Niu, Chaokun Wang, Hang Zhang and Songyao Wang	17663
<i>XToM: Exploring the Multilingual Theory of Mind for Large Language Models</i> Chunkit Chan, Yauwai Yim, Hongchuan Zeng, Zhiying Zou, Xinyuan Cheng, Zhifan Sun, Zheyue Deng, Kawai Chung, Yuzhuo Ao, Fan Yixiang, Cheng Jiayang, Ercong Nie, Ginny Wong, Helmut Schmid, Hinrich Schuetze, Simon See and Yangqiu Song	17681
<i>Cognitive Policy-Driven LLM for Diagnosis and Intervention of Cognitive Distortions in Emotional Support Conversation</i> Lin Zhong, Renjin Zhu, Shujuan Ma, Jinhao Cui, Lingzhi Wang, Hao Chen and Qing Liao	17717
<i>SciFlow-Bench: Evaluating Structure-Aware Scientific Diagram Generation via Inverse Parsing</i> Tong Zhang, Honglin Lin, Zhou Liu, Chong Chen and Wentao Zhang	17747
<i>Immediate Inference: The Missing Foundation in Large Language Model Logical Reasoning</i> Sihang Jiang, Zhiyu Lu, Keyi Wang, Jiaqing Liang, Yanghua Xiao, Xiaojun Meng and Jiansheng Wei	17766
<i>Tracing Logit Trajectories Across Layer Depth: Dataset-Level Explainability for Language Models</i> Jeesu Jung and Sangkeun Jung	17800
<i>REZE: Representation Regularization for Domain-adaptive Text Embedding Pre-finetuning</i> Seungmin Lee, Jeonghwan Lee, Hyunkuk Lim, Sejoon Kim and Mingi Sung	17824
<i>TAMA: Target-Aware Multilingual Abuse Detection by Cascaded Conditional Multi-Task Learning</i> JiYan Liu, Youzheng Liu, Taihang Wang, Yimin Wang, Ye Jiang and Diana Maynard	17842
<i>Reference Attack: A New Cross-Modal Jailbreaking Attack against Multimodal Large Language Models</i> Yulong Wang, Yifei Fu and Jiayi Gao	17860
<i>POWSM: A Phonetic Open Whisper-Style Speech Foundation Model</i> Chin-Jou Li, Kalvin Chang, Shikhar Bharadwaj, Eunjung Yeo, Kwanghee Choi, Jian Zhu, David R. Mortensen and Shinji Watanabe	17882
<i>LCR-RAG: Enhancing Logical Consistency in Retrieval-Augmented Generation via Neuro-symbolic Reinforcement Learning</i> Wenxiang Zheng, Guo Tang, Shixin Jiang, Liangyu Huo, Xiyuan Zhang, Jian Xie and Ming Liu	17900
<i>CIA: Inferring the Communication Topology from LLM-based Multi-Agent Systems</i> Yongxuan Wu, Xixun Lin, He Zhang, Nan Sun, Kun Wang, Chuan Zhou, Shirui Pan and Yanan Cao	17917
<i>Efficiently Learning To Reason or Not to Reason: Root-token Policy Optimization for Adaptive Thinking</i> Taehyeon Kim, Hyunsoo Lee, Youngsoo Jang and Moontae Lee	17934
<i>SAGE: Synergistic Adaptive Gating of Experts for Hateful Video Detection</i> Jie Huang, Xin Liao, Junjie Wang, Mingyang Li, Wenshuo Wang, Ziyong Jiang, Shoubin Li and Qing Wang	17950
<i>HiKEY: Hierarchical Multimodal Retrieval for Open-Domain Document Question Answering</i> Joongmin Shin, Gyuhoo Shim, Jeongbae Park, Jaehyung Seo and Heuseok Lim	17967

<i>Towards Efficient and Effective Diffusion Language Model Inference via Semantic-Aware Adaptive Denoising</i>	
Fan Li, Yu Gu, Zhigang Wang, Fangling Leng, Zhenghao Liu and Ge Yu	17988
<i>AdvancedIF: Rubric-Based Benchmarking and Reinforcement Learning for Advancing LLM Instruction Following</i>	
Yun He, Wenzhe Li, Hejia Zhang, Songlin Li, Karishma Mandyam, Sopan Khosla, Yuanhao Xiong, Nanshu Wang, Xiaoliang Peng, Beibin Li, Shengjie Bi, Shishir G Patil, Qi Qi, Shengyu Feng, Julian Katz-Samuels, Richard Yuanzhe Pang, Sujan Kumar Gonugondla, Hunter Lang, Yue Yu, Yundi Qian, Maryam Fazel-Zarandi, Licheng Yu, Amine Benhalloum, Hany Hassan Awadalla and Manaal Faruqi	18003
<i>Improving Autoformalization Using Direct Dependency Retrieval</i>	
Shaoqi Wang, Lu Yu, Siwei Lou, Feng Yan, Chunjie Yang, Qing Cui and Jun Zhou	18023
<i>Long-Chain Reasoning Distillation via Adaptive Prefix Alignment</i>	
Zhenghao Liu, Zhuoyang Wu, Xinze Li, Yukun Yan, Shuo Wang, Zulong Chen, Yu Gu, Ge Yu and Maosong Sun	18041
<i>TRN-R1-Zero: Text-rich Network Reasoning via LLMs with Reinforcement Learning Only</i>	
Yilun Liu, Ruihong Qiu and Zi Huang	18060
<i>Threshold Differential Attention for Sink-Free, Ultra-Sparse, and Non-Dispersive Language Modeling</i>	
Xingyue Huang, Xueying Ding, Mingxuan Ju, Yozen Liu, Neil Shah and Tong Zhao	18073
<i>PRiSM: Benchmarking Phone Realization in Speech Models</i>	
Shikhar Bharadwaj, Chin-Jou Li, Yoonjae Kim, Kwanghee Choi, Eunjung Yeo, Ryan Soh-Eun Shim, Hanyu Zhou, Brendon Boldt, Karen Rosero, Kalvin Chang, Darsh Agrawal, Keer Xu, Chao-Han Huck Yang, Jian Zhu, Shinji Watanabe and David R. Mortensen	18093
<i>Your Reasoning Benchmark May Not Test Reasoning: Revealing Perception Bottleneck in Abstract Reasoning Benchmarks</i>	
Xinhe Wang, Jin Huang, Xingjian Zhang, Tianhao Wang and Jiaqi W. Ma	18113
<i>Model-Based Imaginative Planning for Embodied Agents</i>	
Junru Song, Hengzhe Jin, Yucong Huang, Tingsong Jiang, Weien Zhou, Feifei Wang, Yang Yang, Ying Wen and Wen Yao	18125
<i>Persona-Grounded Safety Evaluation of AI Companions in Multi-Turn Conversations</i>	
Perna Juneja and Lika Lomidze	18148
<i>Beyond Chunking: Discourse-Aware Hierarchical Retrieval for Long Document Question Answering</i>	
Huiyao Chen, Yi Yang, Yinghui Li, Meishan Zhang, Baotian Hu and Min Zhang	18176
<i>Beyond Value Benchmarks: Measuring Value-Structure Alignment in Large Language Models via Symmetric Q-Sorts</i>	
Jingting Zheng, Yuqi Ren, Linhao Yu, Yongqi Leng and Deyi Xiong	18199
<i>Would LLMs be Good Historical Linguists and Chinese Dialect Learners?</i>	
Yicheng Liu, Shumin Shi, Youchao Zhou and Xingchen Zhang	18231
<i>ProActor: Timing-Aware Reinforcement Learning for Proactive Task Scheduling Agents</i>	
Lei Ding, Bin He, Chenguang Wang and Yang Liu	18257
<i>Seeing No Evil: Blinding Large Vision-Language Models to Safety Instructions via Adversarial Attention Hijacking</i>	
Jingru Li, Wei Ren and Tianqing Zhu	18304

<i>CADMate: Generating CAD Assembly Plan with Geometric Chain-of-Thought and Spatial Physical Rewards</i>	
Jiali Chen, DingBa Fu, Xusen Hei, Yuhang Liu, Yiyang Chen, Jiayuan Xie, Wenqi Fan and Yi Cai	
	18324
<i>Semantic-Aware Logical Reasoning via a Semiotic Framework</i>	
Yunhao Zhang, Xinglang Zhang, Junxi Sheng, Wenbing Li, Junqing Yu, Yi-Ping Phoebe Chen, Wei Yang and Zikai Song	
	18349
<i>Calibration-Aware Policy Optimization for Reasoning LLMs</i>	
Ziqi Wang, Xingzhou Lou, Meiqi Wu, Zhengqi Wen and Junge Zhang	
	18375
<i>GBV-SQL: Guided Generation and SQL2Text Back-Translation Validation for Multi-Agent Text2SQL</i>	
Daojun Chen, Xi Wang, Shenyuan Ren, Qingzhi Ma, Pengpeng Zhao and An Liu	
	18391
<i>Branch-and-Browse: Efficient and Controllable Web Exploration with Tree-Structured Reasoning and Action Memory</i>	
Shiqi He, Yue Cui, Xinyu Ma, Yaliang Li, Bolin Ding and Mosharaf Chowdhury	
	18407
<i>Beyond End-to-End: Dynamic Chain Optimization for Private LLM Adaptation on the Edge</i>	
Yebo Wu, Jingguang Li, Chunlin Tian, KaHou Tam, Zhijiang Guo and Li Li	
	18419
<i>TalkLoRA: Communication-Aware Mixture of Low-Rank Adaptation for Large Language Models</i>	
Lin Mu, Haiyang Wang, Li Ni, Lei Sang, Zhize Wu, Peiquan Jin and Yiwen Zhang	
	18436
<i>TARE: Lightweight Token-Aware Representation Editing for Fine-tuning Transformer-like Models</i>	
Yulong Wang and Siyu Zhao	
	18452
<i>Soft Orthogonal Low-Rank Adaptation for Knowledge Sharing in Large Language Model Continual Learning</i>	
Yitong Wang, Xue Han, WenChun Gao, Qian Hu, Jiahui Wang, Ziqing Wang, Lijun Mei and Junlan Feng	
	18472
<i>ChartAgent: A Multimodal Agent for Visually Grounded Reasoning in Complex Chart Question Answering</i>	
Rachneet Kaur, Nishan Srishankar, Zhen Zeng and Sumitra Ganesh	
	18488
<i>Curing Miracle Steps in LLM Mathematical Reasoning with Rubric Rewards</i>	
Youliang Yuan, Qiuyang Mang, Jingbang Chen, Hong Wan, Xiaoyuan Liu, Junjielong Xu, Jen-tse Huang, Wenxuan Wang, Wenxiang Jiao and Pinjia He	
	18556
<i>Adaptive Planning for Multi-Attribute Controllable Summarization with Monte Carlo Tree Search</i>	
Sangwon Ryu, Heejin Do, Yunsu Kim, Gary Lee and Jungseul Ok	
	18578
<i>CPT-Agent: A Cognitive Process Theory-driven Framework for Student Simulation in Writing Development</i>	
Yuhan Chen, Zizhuo Shen, Miaomiao Cheng, Xu Han, Jiefu Gong, Shijin Wang and Wei Song	
	18596
<i>Team-Based Self-Play With Dual Adaptive Weighting for Fine-Tuning LLMs</i>	
Wu Li, Yigeng Zhou, Zesheng Shi, Yequan Wang, Min Zhang and Jing Li	
	18617
<i>SearchGym: Bootstrapping Real-World Search Agents via Cost-Effective and High-Fidelity Environment Simulation</i>	
Xichen Zhang, Ziyi He, Yinghao Zhu, Sitong Wu, Shaozuo Yu, Meng Chu, Wenhui Zhang, Haoru Tan and Jiaya Jia	
	18633

<i>Learning to Select: Query-Aware Adaptive Dimension Selection for Dense Retrieval</i> Zhanyu Wu, Richong Zhang and Zhijie Nie	18666
<i>MeepleLM: A Virtual Playtester Simulating Diverse Subjective Experiences</i> Zizhen Li, Chuanhao Li, Yibin Wang, Jianwen Sun, Yukang Feng, Jiabin Ai, Fanrui Zhang, Mingzhu Sun, Yifei Huang and Kaipeng Zhang	18678
<i>MSMO-ABSA: Multi-Scale and Multi-Objective Optimization for Cross-Lingual Aspect-Based Sentiment Analysis</i> Chengyan Wu, Bolei Ma, Ningyuan Deng, He Yanqing, Yun Xue and Liu Xiaoyong	18723
<i>FineSteer: A Unified Framework for Fine-Grained Inference-Time Steering in Large Language Models</i> Zixuan Weng, Jinghui Zhang, Kunlin Cai, Ying Li, Peiran Wang and Yuan Tian	18736
<i>Is the Attention Matrix Really the Key to Self-Attention in Multivariate Long-Term Time Series Forecasting?</i> Xinyu Li, Kexi Chen, Jiajie Shen, Ying Zheng, Hong Lu, Jin Zhao and Xin Wang	18757
<i>Difficulty-Controllable Cloze Question Distractor Generation</i> Seokhoon Kang, Yejin Jeon, Seonjeong Hwang and Gary Lee	18774
<i>ToolPRM: Fine-Grained Inference Scaling of Structured Outputs for Function Calling</i> Jianghao Lin, Yuanyuan Shi, Xin Peng, Renjie Ding, Hairui Wang, Yuxuan Peng, Bizhe Bai, Weixi Song, Fengshuo Bai, Huacan Chai, Weinan Zhang, Fei Huang and Ying Wen	18792
<i>GTA: Generating Long-horizon Tasks for Web Agents at Scale</i> Tenghao Huang, Kung-Hsiang Huang, Prafulla Kumar Choubey, Yilun Zhou, Muhao Chen, Jonathan May and Chien-Sheng Wu	18805
<i>Closing the Modality Reasoning Gap for Speech Large Language Models</i> Chaoren Wang, Heng Lu, Xueyao Zhang, Shujie Liu, Yan Lu, Jinyu Li and Zhizheng Wu .	18821
<i>Logical Phase Transitions: Understanding Collapse in LLM Logical Reasoning</i> Xinglang Zhang, Yunyao Zhang, ZeLiang Chen, Junqing Yu, Wei Yang and Zikai Song ...	18836
<i>SpecCache: Speculative KV Cache Reuse for Efficient RAG Serving</i> Zijian Wen, Tao Zhang, Shuangwu Chen, Shenghao Ye, Yu Guo, Qirui Chen, Jingxian Shuai, Yunpeng Hou, Huasen He and Jianyang	18861
<i>REG: Retrieval via Emotion Similarity for Guiding Empathetic Dialogue Generation</i> Xu Wang, Bo Wang, Yang Xiang, Yihong Tang, Dongming Zhao, Yuzifei and Yuexian Hou	18872
<i>TellWhisper: Tell Whisper Who Speaks When</i> Yifan Hu, Peiji Yang, Zhisheng Wang, Yicheng Zhong and Rui Liu	18884
<i>Taming System Complexity: Demystifying Software Engineering Agents in Diagnosing Linux Kernel Faults</i> Zhenhao Zhou, Zhuochen Huang, Yike He, Chong Wang, Jiajun Wang, Yijian Wu, Xin Peng and Yiling Lou	18899
<i>DisCo_Speech: Controllable Zero-Shot Speech Generation with A Disentangled Speech Codec</i> Tao Li, Wenshuo Ge, Zhichao Wang, Zihao Cui, Yong Ma, Yingying Gao, Chao Deng, Shilei Zhang and Junlan Feng	18917
<i>WildReward: Learning Reward Models from In-the-Wild Human Interactions</i> Hao Peng, Yunjia Qi, Xiaozhi Wang, Zijun Yao, Lei Hou and Juanzi Li	18932

<i>Mitigating Catastrophic Forgetting in Target Language Adaptation of LLMs via Source-Shielded Updates</i>	Atsuki Yamaguchi, Terufumi Morishita, Aline Villavicencio and Nikolaos Aletras	18949
<i>SEE: Signal Embedding Energy for Quantifying Noise Interference in Large Audio Language Models</i>	Yuanhe Zhang, Jiayu Tian, Yibo Zhang, Shilinlu Yan, Liang Lin, Zhenhong Zhou, Li Sun and Sen Su	18976
<i>Know the Known and the Unknown: Reasonable Answer Generation with Knowledge-Informed Citations</i>	Yichi Zhang, Zhuo Chen, Lingbing Guo, Jun Xu, Mengshu Sun, Zhizhen Liu, Lei Liang, Wen Zhang and Huajun Chen	18996
<i>FocalOrder: Focal Preference Optimization for Reading Order Detection</i>	Fuyuan Liu, Dianyu Yu, He Ren, Nayu Liu, Xiaomian Kang, Delai Qiu, Fa Zhang, Genpeng Zhen, Shengping Liu, Liang Jiaen, Weihuang, Yining Wang and Junnan Zhu	19012
<i>TLSA: LLM-Guided Text-Label Space Alignment with Contrastive Learning for Generalized Category Discovery</i>	Wenxi Xu, Chuan Qin, Xi Chen, Chuyu Fang, Yuanchun Zhou and Hengshu Zhu	19030
<i>Hierarchical Token Prepending: Enhancing Information Flow in Decoder-based LLM Embeddings</i>	Xueying Ding, Xingyue Huang, Mingxuan Ju, Liam Collins, Yozen Liu, Leman Akoglu, Neil Shah and Tong Zhao	19047
<i>Suggest-Verify-Revise: A Three-Stage Document-Level Event Causality Identification with Narrative Consistency</i>	Ya Su, Hu Zhang, Dan Qiao, YuJie Wang, Yunxiao Zhao, Yue Fan, Shike Li, Ru Li and Hongye Tan	19067
<i>StoryCoder: Narrative Reformulation for Structured Reasoning in LLM Code Generation</i>	Geonhui Jang, Dongyoon Han and YoungJoon Yoo	19092
<i>ODL-TempLLM: Ontology-Guided and Description Logic-Reasoned Temporal Reasoning with LLMs</i>	Jinshuo Liu, Cheng Bi, Meng Wang, Juan Deng, Donghong Ji and Jeff Z. Pan	19113
<i>FACTrial: Factorized Clinical Contrastive Training for Scalable Patient-Trial Retrieval</i>	Xuanren Chen, Chongyang Tao, Tao Shen and Shuai Ma	19127
<i>Your Students Don't Use LLMs Like You Wish They Did</i>	Sebastian Kobler, Matthew Clemson, Angela Sun and Jonathan K. Kummerfeld	19146
<i>Cut Your Losses! Learning to Prune Paths Early for Efficient Parallel Reasoning</i>	Jiayi Bi, Tongxu Luo, Wenyu Du, Zhengyang Tang and Benyou Wang	19171
<i>Discourse Coherence and Response-Guided Context Rewriting for Multi-Party Dialogue Generation</i>	Zhiyu Cao, Peifeng Li and Qiaoming Zhu	19190
<i>EvolvR: Self-Evolving Pairwise Reasoning for Story Evaluation to Enhance Generation</i>	Xinda Wang, Zhengxu Hou, Yangshijie Zhang, Yanbingren, Jialin Liu, ChenZhuo Zhao, Zhibo Yang, Bin-Bin Yang and Feng Xiao	19208
<i>Double: Breaking the Acceleration Limit via Double Retrieval Speculative Parallelism</i>	Yuhao Shen, Tianyu Liu, Junyi Shen, Jinyang Wu, Quan Kong, Huan Li and Cong Wang	19242
<i>BWLA: Breaking the Barrier of WIAx Post-Training Quantization for LLMs</i>	Zhixiong Zhao, Zukang Xu and Dawei Yang	19264

<i>Progressive Multimodal Search and Reasoning for Knowledge-Intensive Visual Question Answering</i> Changin Choi, Wonseok Lee, Jungmin Ko and Wonjong Rhee	19291
<i>Dynamics of Cognitive Heterogeneity: Investigating Behavioral Biases in Multi-Stage Supply Chains with LLM-Based Simulation</i> Jiuyun Jiang, Yuecheng Hong, Bo Yang, Jin Yang, Guangxin Jiang, Xiaomeng Guo and Guang Xiao	19316
<i>Triviality Corrected Endogenous Reward</i> Xinda Wang, Zhengxu Hou, Yangshijie Zhang, Yanbingren, Jialin Liu, ChenZhuo Zhao, Zhibo Yang, Bin-Bin Yang and Feng Xiao	19334
<i>Identity-Robust Language Model Generation via Content Integrity Preservation</i> Miao Zhang, Kelly Chen, Mehrab Tanjim and Rumi Chunara	19356
<i>RFS-Guard: Detecting Reasoning Hallucinations via Cross-Phase Routing Focus in Large Reasoning Models</i> Zihang Liu, Fang Zhouhua, Hui Liu, Zhiwei Liu, Yong Li and Haishuai Wang	19371
<i>Causal-ESC: Reliable Policy Learning for Emotional Support Conversation via Causal Inference</i> Xv Wang, Wang Zhenyu, Guanyu Zheng and Rui Zhang	19386
<i>Chain-of-Thought Compression Should Not Be Blind: V-Skip for Efficient Multimodal Reasoning via Dual-Path Anchoring</i> Dongxu Zhang, Yiding Sun, Cheng Tan, Wenbiao Yan, Ning Yang, Jihua Zhu and Haijun Zhang	19402
<i>CodeJudgeBench: Benchmarking LLM-as-a-Judge for Coding Tasks</i> Hongchao Jiang, Yiming Chen, Yushi Cao, Hung-yi Lee and Robby T. Tan	19416
<i>Towards a Mechanistic Understanding of Large Reasoning Models: A Survey of Training, Inference, and Failures</i> Yi Hu, Jiaqi Gu, Ruxin Wang, Zijun Yao, Hao Peng, Xiaobao Wu, Jianhui Chen, Muhan Zhang and Liangming Pan	19449
<i>Modeling LLM Unlearning as an Asymmetric Two-Task Learning Problem</i> Zeguan Xiao, Siqing Li, Yong Wang, Xuetao Wei, Jian Yang, Yun Chen and Guanhua Chen	19467
<i>MVP: Enhancing Video Large Language Models via Self-supervised Masked Video Prediction</i> Xiaokun Sun, Zezhong Wu, Zewen Ding and Linli Xu	19478
<i>Knowledge Poisoning Attacks on Medical Multi-Modal Retrieval-Augmented Generation</i> Peiru Yang, Haoran Zheng, Tong Ju, Shiting Wang, Wanchun Ni, Jiajun Liu, Shangguang Wang, Yongfeng Huang and Tao Qi	19494
<i>Maximizing Local Entropy Where It Matters: Prefix-Aware Localized LLM Unlearning</i> Naixin Zhai, Pengyang Shao, Binbin Zheng, Yonghui Yang, Fei Shen, Long Bai and Xun Yang	19514
<i>When Bigger Isn't Better: A Comprehensive Fairness Evaluation of Political Bias in Multi-News Summarisation</i> Nannan Huang, Iffat Maab and Junichi Yamagishi	19532
<i>Language Acquisition Device in Large Language Models</i> Masato Mita, Taiga Someya, Ryo Yoshida and Yohei Oseki	19564

<i>CAML: A Conflict-Aware Molecular Language Model Merging Framework for Multi-Constraint Molecular Generation</i>	
Xuanbai Ren, Luoda Tan, Pei Liu, Tengfei Ma, Xiangzheng Fu, Longyue Wang, Yiping Liu and Xiangxiang Zeng	19578
<i>Stratagem: Learning Transferable Reasoning via Trajectory-Modulated Game Self-Play</i>	
Xiachong Feng, Deyi Yin, Xiaocheng Feng, Yi Jiang, Libo Qin, Yangfan Ye, Lei Huang, Weitao Ma, Qiming Li, Yuxuan Gu, Bing Qin and Lingpeng Kong	19595
<i>NOSE: Neural Olfactory-Semantic Embedding with Tri-Modal Orthogonal Contrastive Learning</i>	
Yanyi Su, Hongshuai Wang, Zhifeng Gao and Jun Cheng	19615
<i>ReFL: Reflective Feedback Learning for Hallucination Detection of Large Language Models</i>	
Cunhang Fan, Jun Zhang, Xue Zhang, Shuai Zhang, Zhao Lv, Jianhua Tao and Zhengqi Wen	19648
<i>Fine-Mem: Fine-Grained Feedback Alignment for Long-Horizon Memory Management</i>	
Weitao Ma, Xiaocheng Feng, Lei Huang, Xiachong Feng, Zhanyu Ma, Jun Xu, Jiuchong Gao, Jinghua Hao, Renqing He and Bing Qin	19666
<i>NaturalSloth: Revisiting Denial-of-Service Attacks on Large Language Models</i>	
Yiming Chen, Zexin Li, Xianghu Yue, Robby T. Tan and Haizhou Li	19685
<i>HiGoE: Hierarchical Graph of Evidence to Enhance Retrieval-Augmented Generation for Long-context Summarization</i>	
Long Yuan, Kaiwen Tian, Zi Chen, Bolong Zheng and Chuan Ma	19703
<i>Region-Grounded Report Generation for 3D Medical Imaging: A Fine-Grained Dataset and Graph-Enhanced Framework</i>	
Cong Huy Nguyen, Son Dinh Nguyen, Guanlin Li, Tuan Dung Nguyen, Aditya Narayan Sankaran, Mai Huy Thong, Thanh Trung Nguyen, Mai Hong Son, Reza Farahbakhsh, Phi Le Nguyen and Noel Crespi	19725
<i>UI-Copilot: Advancing Long-Horizon GUI Automation via Tool-Integrated Policy Optimization</i>	
Zhengxi Lu, Fei Tang, Guangyi Liu, Jin Ma, Kaitao Song, Xu Tan, Wenqi Zhang, Weiming Lu, Jun Xiao, Yueting Zhuang and Yongliang Shen	19741
<i>STK-Adapter: Incorporating Evolving Graph and Event Chain for Temporal Knowledge Graph Extrapolation</i>	
Shuyuan Zhao, Wei Chen, Weijie Zhang, Xinrui Hou, Junfeng Shen, Boyan Shi, Shengnan Guo, Youfang Lin and Huaiyu Wan	19763
<i>Learning What Matters: Dynamic Dimension Selection and Aggregation for Interpretable Vision-Language Reward Modeling</i>	
Qiyuan Chen, Hongsen Huang, Jiahe Chen, Qian Shao, Jintai Chen, Hongxia Xu, Renjie Hua, Ren Chuan and Jian Wu	19779
<i>DeepGuard: Secure Code Generation via Multi-Layer Semantic Aggregation</i>	
Li Huang, Zhongxin Liu, Yifan Wu, Tao Yin, Dong li, Jichao Bi, Nankun Mu, Hongyu Zhang and Meng Yan	19793
<i>Distillation Traps and Guards: A Calibration Knob for LLM Distillability</i>	
Weixiao Zhan, Yongcheng Jing, Leszek Rutkowski and Dacheng Tao	19814
<i>DVMap: Fine-Grained Pluralistic Value Alignment via High-Consensus Demographic-Value Mapping</i>	
Pengyun Zhu, Yuqi Ren, Zhen Wang, Lei Yang and Deyi Xiong	19834

<i>HTMR: Hybrid Token Masking Reinforcement Learning with Verifiable Rewards for Event Argument Extraction with Multi-Perspective Reasoning</i>	Jianwen Luo, Yongkang Jin, Yu Hong and Jianmin Yao	19853
<i>IS-CoT: Breaking the Long-form Generation Collapse via Interleaved Structural Thinking</i>	Zechen Sun, Yuyang Sun, Zecheng Tang, Juntao Li, Wenpeng Hu, Wenliang Chen, Zhunchen Luo, Guotong Geng and Min Zhang	19874
<i>Seeing the Whole Elephant: A Benchmark for Failure Attribution in LLM-based Multi-Agent Systems</i>	Mengzhuo Chen, Junjie Wang, Fangwen Mu, Yawen Wang, Zhe Liu, Huanxiang Feng and Qing Wang	19888
<i>MTAVG-Bench: A Diagnostic Benchmark for Multi-Talker Dialogue-Centric Audio-Video Generation</i>	Yanghao Zhou, Haitian Li, Rexar Lin, Heyan Huang, Jinxing Zhou, Changsen Yuan, Tian Lan, Ziqin Zhou, Yudong Li, Jiajun Xu, Jingyun Liao, YiMing Cheng, Xuefeng Chen, Xian-Ling Mao and Yousheng Feng	19906
<i>Why LLMs Hallucinate on Structured Knowledge: A Mechanistic Analysis of Reasoning over Linearized Representations</i>	Shanghao Li, Jinda Han, Yibo Wang, Yuanjie Zhu, Zihe Song, Langzhou He, Kenan Kamel A Alghythee and Philip S. Yu	19943
<i>SubTokenTest: A Practical Benchmark for Real-World Sub-token Understanding</i>	Shuyang Hou, Yi Hu and Muhan Zhang	19957
<i>Benchmarking and Enabling Efficient Chinese Medical Retrieval via Asymmetric Encoders</i>	Angqing Jiang, Jianlyu Chen, Zhefang, Yongcan Wang, Xinpeng Li, Keyu Ding and Defu Lian	20000
<i>Merlin’s Whisper: Enabling Efficient Reasoning in Large Language Models via Black-box Persuasive Prompting</i>	Heming Xia, Cunxiao Du, Rui Li, Chak Tou Leong, Yongqi Li and Wenjie Li	20021
<i>CAKE: Causal-Guided Adaptive Knowledge Editing for LLMs</i>	Shuxin Liu and Jianhao Zhang	20040
<i>Preconditioned Test-Time Adaptation for Out-of-Distribution Debiasing in Narrative Generation</i>	Hanwen Shen, Ting Ying, Jiajie Lu and Shanshan Wang	20071
<i>Backdoor Collapse: Eliminating Unknown Threats Via Known Backdoor Aggregation In Language Models</i>	Liang Lin, Miao Yu, Moayad Aloqaily, Zhenhong Zhou, Kun Wang, Linsey Pang, Prakhar Mehrotra and Qingsong Wen	20105
<i>Less is More: Improving LLM Reasoning with Minimal Test-Time Intervention</i>	Zhen Yang, Mingyang Zhang, Feng Chen, Ganggui Ding, Liang Hou, Xin Tao and Ying-Cong Chen	20124
<i>League of LLMs: A Benchmark-Free Paradigm for Mutual Evaluation of Large Language Models</i>	Qianhong Guo, Wei Xie, Xiaofang Cai, Enze Wang, Shuoyoucheng Ma, Xiaobing Sun, Tian Xia, Kai Chen, Xiaofeng Wang and Baosheng Wang	20138
<i>RARE: Redundancy-Aware Retrieval Evaluation Framework for High-Similarity Corpora</i>	Hanjun Cho and Jay-Yoon Lee	20160

<i>Leibniz: Theory-of-Mind Driven Neuro-Symbolic Logical Reasoning via Multi-Agent Collaboration</i> Yue Fan, Hu Zhang, Yunxiao Zhao, Guangjun Zhang, Hao Zhan, Ru Li, Hongye Tan and Yuanlong Wang	20186
<i>The Paradox of Outcome Optimization: A Causal Information-Theoretic Bound on Reasoning Shortcuts in LLMs</i> Zihan Chen, Yiming Zhang, Wenxiang Geng, Zenghui Ding and Yining Sun	20211
<i>SPENCE: A Syntactic Probe for Detecting Contamination in NL2SQL Benchmarks</i> Mohammadtaher Safarzadeh, Hitesh Laxmichand Patel, Afshin Oroojlooy, Graham Horwood and Dan Roth	20222
<i>PaTaRM: Bridging Pairwise and Pointwise Signals via Preference-Aware Task-Adaptive Reward Modeling</i> Ai Jian, Jingqing Ruan, Xing Ma, Dailin Li, Weipeng Zhang, Ke Zeng and Xunliang Cai ..	20240
<i>Uncovering Sentiment Analysis Circuit in Large Language Model</i> Shichen Li, Zhouyang Wang, Zhongqing Wang and Peifeng Li	20269
<i>Think Before Writing: Feature-Level Multi-Objective Optimization for Generative Citation Visibility</i> Zikang Liu and Peilan Xu	20290
<i>Zero-Shot Multimodal Retrieval with Multi-Scale Contextual Representations</i> Sourajit Saha and Tejas Gokhale	20304
<i>ReEfBench: Quantifying the Reasoning Efficiency of LLMs</i> Zhizhang Fu, Yuancheng Gu, Chenkai Hu, Hanmeng Liu and Yue Zhang	20325
<i>Beyond Ranking: Fine-Grained Diagnostics and Self-Improvement for MLLMs</i> Mingze Xu, Zijing Zhao, Qiming Peng, Houwen Peng, Han Hu, Zhanhui Kang and Yuxing Han	20347
<i>Bridging Internal Consistency and External Alignment: A Causal and Dynamic Interpretability Framework for LLM Generation</i> Shuyao Xiao, Shengling Wang and Ke Chao	20378
<i>Beyond the Last Frame: Process-aware Evaluation for Generative Video Reasoning</i> Yifan Li, YuKai Gu, Yingqian Min, Zikang Liu, Yifan Du, Kun Zhou, Min Yang, Xin Zhao and Minghui Qiu	20393
<i>THOR: A Theta-Gamma Hierarchical Oscillatory Reasoning Framework for Multi-hop QA</i> Ziyang Ling, Ronald X. Xu and Mingzhai Sun	20410
<i>How Should We Enhance the Safety of Large Reasoning Models: An Empirical Study</i> Zhexin Zhang, Xian Qi Loye, Victor Shea-Jay Huang, Junxiao Yang, Qi Zhu, Shiyao Cui, Fei Mi, Lifeng Shang, Yingkang Wang, Hongning Wang and Minlie Huang	20438
<i>HSCoDeComp: A Realistic and Expert-level Agent Benchmark for Hierarchical Rule Application</i> Tian Lan, Yiqian Yang, Qianghuai Jia, Li Zhu, Hui Jiang, Hang Zhu, Weihua Luo and Longyue Wang	20458
<i>AnchorSeg: Language Grounded Query Banks for Reasoning Segmentation</i> Rui Qian, Chuanhang Deng, Qiang Huang, Jian Xiong, Mingxuan Li, Yingbo Zhou, Wei Zhai, Jintao Chen and Dejing Dou	20490
<i>R³ AG : Retriever Routing for Retrieval – Augmented Generation</i> Tong Zhao, Yutao Zhu, Yucheng Tian and Zhicheng Dou	20506

<i>ClusterRAG: Cluster-Based Collaborative Filtering for Personalized Retrieval-Augmented Generation</i> Gibson Nkhata, Uttamasha Anjally Oyshi, Quan Mai and Susan Gauch	20523
<i>Toward Robust In-Context Learning: Leveraging Out-of-distribution Proxies for Target Inaccessible Demonstration Retrieval</i> Hao Xu, Rite Bo, Fausto Giunchiglia, Yingji Li and Rui Song	20540
<i>AutoSchemaKG: Autonomous Knowledge Graph Construction through Dynamic Schema Induction from Web-Scale Corpora</i> Jiaxin Bai, Wei Fan, Qi Hu, Qing Zong, Chunyang Li, Hong Ting Tsang, Hongyu Luo, Yauwai Yim, Haoyu Huang, Xiao Zhou, Feng Qin, Tianshi Zheng, Xi Peng, Xin Yao, Huiwen Yang, Leijie Wu, Ji Yi, Gong Zhang, Renhai Chen and Yangqiu Song	20557
<i>Aligned Multi-View Scripts for Universal Chart-to-Code Generation</i> Zhihan Zhang and Lizi Liao	20585
<i>SpeakerSleuth: Can Large Audio-Language Models Judge Speaker Consistency across Multi-turn Dialogues?</i> Jonggeun Lee, Junseong Pyo, Gyuhyeon Seo and Yohan Jo	20612
<i>Breaking Block Boundaries: Anchor-based History-stable Decoding for Diffusion Large Language Models</i> Shun Zou, Yong Wang, Zehui Chen, Lin Chen, Chongyang Tao, Feng Zhao and Xiangxiang Chu	20637
<i>HiSVD: Principled Low-Rank Approximation of LLMs via Hierarchical Modeling of Information Capacity and Spectral Structure</i> Zhuo Chen, Minghao Li, Xiaoqian Ma, Siqi Fan, Xiusheng Huang, Zhang Liujie and Weihang Chen	20659
<i>Unveiling the Unknown: Open-Set Entity Typing via Two-Stage Generation</i> Hu Chen, Binhan Yang and Wei Shen	20678
<i>Don't Adapt Small Language Models for Tools; Adapt Tool Schemas to the Models</i> Jonggeun Lee, Woojung Song, Jongwook Han, Haesung Pyun and Yohan Jo	20695
<i>Language Reconstruction with Brain Predictive Coding from fMRI Data</i> Congchi Yin, Ziyi Ye and Piji Li	20720
<i>Chaining the Evidence: Robust Reinforcement Learning for Deep Search Agents with Citation-Aware Rubric Rewards</i> Jiajie Zhang, Xin Lv, Ling Feng, Lei Hou and Juanzi Li	20744
<i>Think Natively: Unlocking Multilingual Reasoning with Consistency-Enhanced Reinforcement Learning</i> Xue Zhang, Yunlong Liang, Fandong Meng, Songming Zhang, Kaiyu Huang, Yufeng Chen, Xu Jinan and Jie Zhou	20766
<i>Mem²Evolve: Towards Self-Evolving Agents via Co-Evolutionary Capability Expansion and Experience Distillation</i> Zihao Cheng, Zeming Liu, Yingyu Shan, Xinyi Wang, Xiangrong Zhu, Yunpu Ma, Hongru Wang, Yuhang Guo, Wei Lin and Yunhong Wang	20784
<i>VideoCuRL: Video Curriculum Reinforcement Learning with Orthogonal Difficulty Decomposition</i> Hongbo Jin, Kuanwei Lin, Wenhao Zhang, Yichen Jin and Ge Li	20832

<i>Visual Inception: Compromising Long-term Planning in Agentic Recommenders via Multimodal Memory Poisoning</i>	
Jiachen Qian	20846
<i>Language on Demand, Knowledge at Core: Composing LLMs with Encoder-Decoder Translation Models for Extensible Multilinguality</i>	
Mengyu Bu and Yang Feng	20863
<i>GeoArena: Evaluating Open-World Geographic Reasoning in Large Vision-Language Models</i>	
Pengyue Jia, Yingyi Zhang, Xiangyu Zhao and Sharon Li	20886
<i>Fisher-Driven Adaptive Locating for Knowledge Editing in Large Language Models</i>	
Chenghao Xu, Jiexi Yan, Guangtao Lyu, Qi Liu, Muli Yang and Cheng Deng	20902
<i>From AR to Diffusion: Efficiently Adapting Large Language Models with Strictly Causal and Elastic Horizons</i>	
Xiangyu Ma, Teng Xiao, Zuchao Li and Lefei Zhang	20914
<i>ExPerT: Personalizing LLM Responses to Users' Domain Expertise via Query-Wise Semantic and Keystroke Behavioral Cues</i>	
Yeji Park, Jiwon Tark and Taesik Gong	20928
<i>Measure Twice, Click Once: Co-evolving Proposer and Visual Critic via Reinforcement Learning for GUI Grounding</i>	
Wenkai Wang, Xiyun Li, Hongcan Guo, Wenhao Yu, Tianqing Fang, Haitao Mi, Dong Yu and Shengyu Zhang	20964
<i>Retrieving to Recover: Towards Incomplete Audio-Visual Question Answering via Semantic-consistent Purification</i>	
Jiayu Zhang, Shuo Ye, Qilang Ye, Zihan Song, Jiajian Huang and Zitong YU	20985
<i>ChatHLS: Towards Systematic Design Automation and Optimization for High-Level Synthesis</i>	
Runkai Li, Jia Xiong, Xiuyuan He, Jieru Zhao, Jiaqi Lv, Haowen Fang, Lei Qi and Xi Wang	20996
<i>Cultural Benchmarking of LLMs in Standard and Dialectal Arabic Dialogues</i>	
Muhammad Dehan Al Kautsar, Saeed Almheiri, Momina Ahsan, Bilal Elbouardi, Younes Samih, Sarfraz Ahmad, Amr Keleg, Omar El Herraoui, Kareem Elzeky, Abed Alhakim Freihat, Mohamed Anwar, Zhuohan Xie, Junhong Liang, Mohammad Rustom Al Nasar, Preslav Nakov and Fajri Koto	21016
<i>ReRec: Reasoning-Augmented LLM-based Recommendation Assistant via Reinforcement Fine-tuning</i>	
Jiani Huang, Shijie Wang, Liang-Bo Ning, Wenqi Fan and Li Qing	21040
<i>MoA: Heterogeneous Mixture of Adapters for Parameter-Efficient Fine-Tuning of Large Language Models</i>	
Jie Cao, Tianwei Lin, Bo Yuan, Rolan Yan, Hongyang He, Wenqiao Zhang, Juncheng Li, Dongping Zhang, Siliang Tang and Yueting Zhuang	21056
<i>Beyond the Context Window: Scaling Agentic RL via End-to-end Optimized Context Compression</i>	
Miao Lu, Weiwei Sun, Weihua Du, Zhan Ling, Xuesong Yao, Kang Liu and Jiecao Chen ..	21074
<i>Probing the Safety Robustness of LLMs in Latent Space</i>	
Tianle Gu, Kexin Huang, Zongqi Wang, Yixu Wang, Jie Li, Xin Wang, Yang Yao, Yujiu Yang, Yan Teng and Yingchun Wang	21126

<i>LoopTool: Closing the Data–Training Loop for Robust LLM Tool Calls</i>	
Kangning Zhang, Weiwen Liu, Wenxiang Jiao, Kounianhua Du, Yuan Lu, Weinan Zhang and Yong Yu	21144
<i>Placing Puzzle Pieces Where They Matter: A Question Augmentation Framework for Reinforcement Learning</i>	
Yangyi Fang and Haolin Shi	21165
<i>USB: A COMPREHENSIVE AND UNIFIED SAFETY EVALUATION BENCHMARK FOR MULTI-MODAL LARGE LANGUAGE MODELS</i>	
Baolin Zheng, Guanlin Chen, Qingyang Teng, Hongqiong Zhong, Yingshui Tan, Zhendong Liu, Weixun Wang, Jiaheng Liu, Jian Yang, Huiyun Jing, Jincheng Wei, Wenbo Su, Xiaoyong Zhu, Bo Zheng and Kaifu Zhang	21184
<i>FlowSearch: Advancing Deep Research with Dynamic Structured Knowledge Flow</i>	
Yusong Hu, Runmin Ma, Yue Fan, Jinxin Shi, Zongsheng Cao, Yuhao Zhou, Jiakang Yuan, Shuaiyu Zhang, Shiyang Feng, Xiangchao Yan, Shufei Zhang, Wenlong Zhang, Lei Bai and Bo Zhang	21212
<i>LLMEval-Fair: A Large-Scale Longitudinal Study on Robust and Fair Evaluation of Large Language Models</i>	
Ming Zhang, Yujiong Shen, Jingyi Deng, Yuhui Wang, Huayu Sha, Kexin Tan, Qiyuan Peng, Yue Zhang, Junzhe Wang, Shichun Liu, Yueyuan Huang, Jingqi Tong, Changhao Jiang, Yilong Wu, Zhihao Zhang, Mingqi Wu, Mingxu Chai, Zhiheng Xi, Shihan Dou, Tao Gui, Qi Zhang and Xuanjing Huang	21246
<i>SAME: Signer-Aware Mixture-of-Experts for Test-Time Adaptation in Sign Language Translation</i>	
Lujia Yang, Weicai Yan, Yongbo He, Qifei Zhang, Tao Jin, Jinshan Zhang, Meng Xi and Jianwei Yin	21274
<i>Controllable LLM Reasoning via Sparse Autoencoder-Based Steering</i>	
Yi Fang, Wenjie Wang, Mingfeng Xue, Boyi Deng, Fengli Xu, Dayiheng Liu and Fuli Feng	21290
<i>LEDOM: Reverse Language Model</i>	
Xunjian Yin, Sitao Cheng, Yuxi Xie, Xinyu Hu, Li Lin, Xinyi Wang, Liangming Pan, William Yang Wang and Xiaojun Wan	21306
<i>AEGIS: A Holistic Benchmark for Evaluating Forensic Analysis of AI-Generated Academic Images</i>	
Bo Zhang, Tzu-Yen Ma, Zichen Tang, Junpeng Ding, Zirui Wang, Yizhuo Zhao, Peilin Gao, Zijie Xi, Zixin Ding, Haiyang Sun, Haocheng Gao, Yuan Liu, Liangjia Wang, Yiling Huang, Yujie Wang, Yuyue Zhang, Ronghui Xi, Yuanze Li, Jiacheng Liu, Zhongjun Yang and Haihong E.....	21327
<i>The GaoYao Benchmark: A Comprehensive Framework for Evaluating Multilingual and Multicultural Abilities of Large Language Models</i>	
Yilun Liu, Chunguang Zhao, Mengyao Piao, Lingqi Miao, Shimin Tao, Minggui HE, Chenxin Liu, Zhang Li, Mahongxia, Jiabin Guo, Chen Liu, Liqun Deng, Jiansheng Wei, Xiaojun Meng, Fanyi Du, Daimeng Wei and Yanghua Xiao	21364
<i>METRO: Towards Strategy Induction from Expert Dialogue Transcripts for Non-collaborative Dialogues</i>	
Haofu Yang, Jiayi Liu, Chen Huang, Faguo Wu, Wenqiang Lei and See-Kiong Ng	21385
<i>TamEdit: Trajectory-Aware Meta-Learning for Specificity-Preserving Continual Knowledge Editing</i>	
Shiqiang Tian, Cheng Ding, Qin Chen, Jie Zhou and Liang He	21417

<i>LexRel: Benchmarking Legal Relation Extraction for Chinese Civil Cases</i>	
Yida Cai, Ranjuexiao Hu, Huiyuan Xie, Chenyang Li, Yun Liu, Yuxiao Ye, Zhenghao Liu, Weixing Shen and Zhiyuan Liu	21440
<i>Agentic Memory: Learning Unified Long-Term and Short-Term Memory Management for Large Language Model Agents</i>	
Yi Yu, Liuyi Yao, Yuexiang Xie, Qingquan Tan, Jiaqi Feng, Yaliang Li and Libing Wu	21457
<i>Profiling-Free Mixed-Precision Quantization for MoE LLMs via Fuzzy Rule Interpolation</i>	
Huachen Qi, Ruiyu Zhuo, Bowen Shi, Xiang Chang, Fei Chao, Changjing Shang and Qiang Shen	21484
<i>Tailoring Diagnostic Modeling to Individual Learners: Personalized Distractor Generation via MCTS-Guided Reasoning Reconstruction</i>	
Tao Wu, Jingyuan Chen, Wang Lin, Jian Zhan, Mengze Li, Fangzhou Jin, Min Zhang, Kun Kuang and Fei Wu	21500
<i>MTR-Bench: A Comprehensive Benchmark for Multi-Turn Reasoning Evaluation</i>	
Xiaoyuan Li, Keqin Bao, Yubo Ma, Moxin Li, Wenjie Wang, Rui Men, Yichang Zhang, Fuli Feng and Dayiheng Liu	21525
<i>Community-Aware Assessment of Social Textual Engagement and Resonance: A Human-Centric Perspective on User-Generated Content Evaluation</i>	
Tianjiao Li, Kai Zhao, Xiang Li, Yang Liu and Huyang Sun	21578
<i>Query-Aware Knowledge Retrieval via Hyperbolic Structuring</i>	
Chuang Zhou, Junnan Dong, Yilin Xiao, Shengyuan Chen, Su Dong, di Yin, Xing Sun, Zhaozhuo Xu and Xiao Huang	21601
<i>EduBench: A Comprehensive Benchmarking Dataset for Evaluating Large Language Models in Diverse Educational Scenarios</i>	
Bin Xu, Yu Bai, Huashan Sun, Yiguan Lin, Siming Liu, Xinyue Liang, Yaolin Li, Zhuangzhi Dong, Jingren Zhang, Yufan Deng, Xinyu Zou, Yang Gao and Heyan Huang	21615
<i>WebClipper: Efficient Evolution of Web Agents with Graph-based Trajectory Pruning</i>	
Junjie Wang, Zequn Xie, Dan Yang, Jie Feng, Yue Shen, Duolin Sun, Meixiu Long, Yihan Jiao, Zhehao Tan, Jian Wang, Peng Wei and Jinjie Gu	21646
<i>MMAC: A Multilingual, Multimodal Alignment Framework for Cultural Grounding Evaluation</i>	
Weihua Zheng, Zhengyuan Liu, Tanmoy Chakraborty, Weiwen Xu, Xiaoxue Gao, Bryan Chen Zhengyu Tan, Bowei Zou, Chang Liu, Yujia Hu, Xing Xie, Xiaoyuan Yi, Jing Yao, Chaojun Wang, Long Li, Rui Liu, Huiyao Liu, Koji Inoue, Ryuichi Sumida, Tatsuya Kawahara, Fan Xu, Lingyu Ye, Wei Tian, Dongjun Kim, Jimin Jung, Jaehyung Seo, Nadya Yuki Wangsajaya, Pham Minh Duc, Ojasva Saxena, Palash Nandi, Xiyan Tao, Wiwik Karlina, Tuan Luong, Keertana Arun Vasan, Roy Ka-Wei Lee and Nancy F. Chen	21667
<i>Step-GRPO: Internalizing Dynamic Early Exit for Efficient Reasoning</i>	
Benteng Chen, Weida Wang, Shufei Zhang, Mingbao Lin and Min Zhang	21710
<i>Gold-Medal-Level Olympiad Geometry Solving with Efficient Heuristic Auxiliary Constructions</i>	
Boyan Duan, Xiao Liang, Shuai Lu, Yaoxiang Wang, Yelong Shen, Kai-Wei Chang, Ying Nian Wu, Mao Yang, Weizhu Chen and Yeyun Gong	21725
<i>Rethinking Jailbreak Detection of Large Vision Language Models with Representational Contrastive Scoring</i>	
Peichun Hua, Hao Li, Shanghao Shi, Zhiyuan Yu and Ning Zhang	21748

<i>The Digital Dunning-Kruger Effect: Decoupling Hallucinations via Geometric Hidden-state Observation for Semantic Truthfulness</i>	
Yueheng Mao, Min Yu, Gengwang Li, Jianguo Jiang, Gang Li, Meng Zhang, Zhen Xu, Weiqing Huang and Ming Liu	21786
<i>d-TreeRPO: Towards More Reliable Policy Optimization for Diffusion Language Models</i>	
Leyi Pan, Shuchang Tao, Yunpeng Zhai, Zheyu Fu, Liancheng Fang, Minghua He, Lingzhe Zhang, Zhaoyang Liu, Bolin Ding, Aiwei Liu and Lijie Wen	21801
<i>Traffic-RL: Reinforced LLMs Bring Human-Like Reasoning to Traffic Signal Control Systems</i>	
Xingchen Zou, Yuhao Yang, Zheng Chen, Xixuan Hao, Yiqi Chen, Chao Huang and Yuxuan Liang	21823
<i>Nirvana: A Specialized Generalist Model With Task-Aware Memory Mechanism</i>	
Yuhua Jiang, Shuang Cheng, Yihao Liu, Ermo Hua, Che Jiang, Weigao Sun, Yu Cheng, Feifei Gao, Biqing Qi and Bowen Zhou	21839
<i>BoYaEval: Evaluating Multimodal Large Language Models on Understanding Ancient Chinese Musical Scores</i>	
Jiajia Li, Weizhi Xue, Yao Yao, Qiwei Li, Chenchong, Zuchao Li, Ping Wang and Hai Zhao	21858
<i>Rectifying the Emotional Flow: Aligning Priors and Dynamic Guidance for High-Arousal Text-to-Speech</i>	
Fangming Feng, Dongjie Fu, Zequn Xie, Yu Zhang, Yangyang Wu, Zhou Zhao and Tao Jin	21874
<i>Perplexity-Aware Data Scaling Law: Perplexity Landscapes Predict Performance for Continual Pre-training</i>	
Lei Liu, Hao Zhu, Xiaoyan Yang, Yue Shen, Zhixuan Chu, Jian Wang, Jinjie Gu and Kui Ren	21889
<i>Benchmarking Large Vision-Language Models on CFMME: A Comprehensive Chinese Financial Multimodal Evaluation Dataset</i>	
Qian Chen, Xianyin Zhang, Yanzhi Liu, Lifan Guo, Feng Chen and Chi Zhang	21902
<i>AutoReproduce: Automatic AI Experiment Reproduction with Paper Lineage</i>	
Xuanle Zhao, Zilin Sang, Yuxuan Li, Qi Shi, Weilun Zhao, Shuo Wang, Duzhen Zhang, Xu Han, Zhiyuan Liu and Maosong Sun	21920
<i>How Can Synthetic Data Improve Multilingual Language Model Pretraining? A Data Quality Perspective</i>	
Tongyao Zhu, Qian Liu, Chang Ma, Jinghan Zhang, Longxu Dou, Junxian He and Shiqi Chen	21943
<i>Can Large Language Models Infer Causal Relationships from Real-World Text?</i>	
Ryan Saklad, Aman Chadha, Oleg V. Pavlov and Raha Moraffah	21957
<i>Distributional Clarity: The Hidden Driver of RL-Friendliness in Large Language Models</i>	
Shaoning Sun, Mingzhu Cai, Huang He, Bingjin Chen, Siqi Bao, Yujiu Yang, Hua Wu and Haifeng Wang	21990
<i>ReasonRank: Empowering Passage Ranking with Strong Reasoning Ability</i>	
Wenhan Liu, Xinyu Ma, Weiwei Sun, Yutao Zhu, Yuchen Li, Dawei Yin and Zhicheng Dou	22007
<i>ATIR: Towards Audio-Text Interleaved Contextual Retrieval</i>	
Tong Zhao, Chenghao Zhang, Yutao Zhu and Zhicheng Dou	22032

<i>Task-Aware LLM Routing with Multi-Level Task-Profile-Guided Data Synthesis for Cold-Start Scenarios</i>	Hui Liu, Bin Zou, Kecheng Chen, Jie Liu, Wenya Wang and Haoliang Li	22047
<i>SelFusion: Self-distillation for Diffusion Language Models</i>	Hyeonngsoo Lim, Jinyoung Kim, Eun Seo Seo, Min Ho Jang and Ji Won Yoon	22077
<i>Benchmarking Web Agent Safety under E-commerce Deceptive Interfaces</i>	Zijing Shi, Meng Fang and Ling Chen	22090
<i>Counteracting the Matthew Effect in Self-Improvement of LVLMs through Head-Tail Re-balancing</i>	Xin Guo, Zhiheng Xi, Yiwen Ding, Yitao Zhai, Xiaowei Shi, Xunliang Cai, Tao Gui, Qi Zhang and Xuanjing Huang	22104
<i>MAGIC: Deep Geometric Evolution with Structural Consensus for Temporal Knowledge Graph Reasoning</i>	Chengao Liu, Yuan Li, Yingze Wang and Jianbin Jiao	22122
<i>MACS: Modality-Aware Capacity Scaling for Efficient Multimodal MoE Inference</i>	Bo Li, Chuan Wu and Shaolin Zhu	22134
<i>Efficient Prior-Guided Reasoning for Robust Retrieval-Augmented Generation under Conflicts</i>	Xiaowei Yuan, Ziyang Huang, Zhao Yang, Yequan Wang, Jun Zhao and Kang Liu	22149
<i>Markovian Linguistic-Temporal Bridge: Unlocking the Potential of LLMs for Time Series Forecasting</i>	Siming Sun, Kai Zhang, Xuejun Jiang, Wenchao Meng and Qinmin Yang	22165
<i>Template-assisted Contrastive Learning of Task-oriented Dialogue Sentence Embeddings</i>	Minsik Oh, Jiwei Li and Guoyin Wang	22181
<i>MIRTH: Mutual-Information Reasoning with Temporal Hubs for Vision-Language-Action Agents</i>	Hao Sun, Yu Song, Teng Shiyu, Ziwei Niu and Yen-wei Chen	22199
<i>IRIS: Interleaved Reinforcement with Incremental Staged Curriculum for Cross-Lingual Mathematical Reasoning</i>	Navya Gupta, Rishitej Reddy Vyalla, Avinash Anand, Chhavi Kirtani, Erik Cambria, Zhengchen Zhang, Zhengkui Wang, Timothy Liu, Aik Beng Ng, Simon See and Rajiv Ratn Shah	22216
<i>FLARE: Fine-Grained Length-Aware Routing for Resource-Efficient Heterogeneous LLM Serving</i>	Yujia Fu, Heming Zhong, Dan Huang and Yutong Lu	22249
<i>MAESTRO: Meta-learning Adaptive Estimation of Scalarization Trade-offs for Reward Optimization</i>	Yang Zhao, Hepeng Wang, Xiao Ding, Yangou Ouyang, Bibo Cai, Kai Xiong, Jinglong Gao, Zhouhao Sun, Li Du, Bing Qin and Ting Liu	22267
<i>Bit-by-Bit: Progressive QAT Strategy with Outlier Channel Splitting for Stable Low-Bit LLMs</i>	Binxing Xu, Hao Gu, Lujun Li, Hao Wang, Bei Liu, Jiacheng Liu, Qiyuan Zhu, Xintong Yang, Chao Li, Sirui Han and Yike Guo	22284
<i>OS-Symphony: A Holistic Framework for Robust and Generalist Computer-Using Agents</i>	Bowen Yang, Kaiming Jin, Zhenyu Wu, Zhaoyang Liu, Qiushi Sun, Zehao Li, JingJing Xie, Zhoumianze Liu, Fangzhi Xu, Kanzhi Cheng, Yian Wang, Qingyun Li, Yu Qiao, Zun Wang and Zichen Ding	22300
<i>Learning from Cognition: Enhancing RL Efficiency for LLM Reasoning via Hierarchical Metacognitive Decomposition and Refinement</i>	Zexu Sun, Yongcheng Zeng, Erxue Min, Heyang Gao, Bokai Ji, Dugang Liu, Xing Tang, Xiuqiang He and Xu Chen	22331

<i>BrowseComp-Plus: A Fair and Disentangled Evaluation Benchmark for Deep Search Agents</i>	
Zijian Chen, Xueguang Ma, Shengyao Zhuang, Ping Nie, Kai Zou, Sahel Sharifymoghaddam, Andrew Liu, Joshua Green, Kshama Patel, Ruoxi Meng, Mingyi Su, Yanxi Li, Haoran Hong, Xinyu Shi, Xuye Liu, Hosna Oyarhoseini, Nandan Thakur, Crystina Zhang, Luyu Gao, Wenhu Chen and Jimmy Lin	22349
<i>LitVISTA: A Benchmark for Narrative Orchestration in Literary Text</i>	
Mingzhe Lu, Yiwen Wang, Yanbing Liu, Qi You, Chong Liu, Ruize Qin, Haoyu Dong, Wenyu Zhang, JiaRui Zhang, Yue Hu and Yunpeng Li	22371
<i>RST-Guarder: Enhancing Long-Context Robustness for Safeguards via RST Parsing and Probabilistic Inference</i>	
Xu Zhang and Xiaojun Wan	22397
<i>EpiCaR: Knowing What You Don't Know Matters for Better Reasoning in LLMs</i>	
Jewon Yeom, Jaewon Sok, Seonghyeon Park, Jeongjae Park and Taesup Kim	22414
<i>MoEC: A Memory-Routed Mixture-of-Experts Controller for Adaptive Minecraft Control</i>	
Hui Wu, Chao Xu, Jianghui Wang, Ziqiong Liu, Dong Li, Yiwei Dai and Emad Barsoum .	22444
<i>MECH: A Cost-Effective Multi-Task Cascade Framework for Classroom Opinion Evolution Recognition</i>	
Yancui Li, Xiaoyu Zhou, Guoyi Miao and Fang Kong	22460
<i>Debiasing Reward Models via Causally Motivated Inference-Time Intervention</i>	
Kazutoshi Shinoda, Kosuke Nishida and Kyosuke Nishida	22474
<i>Steganography Beyond Pixels: Reimagining Image Steganography as Cross-Modal Linguistic Communication</i>	
Ren Lijing and Denghui Zhang	22490
<i>Hard2Verify: A Step-Level Verification Benchmark for Open-Ended Frontier Math</i>	
Shrey Pandit, Austin Xu, Xuan-Phi Nguyen, Yifei Ming, Caiming Xiong and Shafiq Joty .	22502
<i>Agent-based Substructure Counting under Local Differential Privacy</i>	
Yuting Zhang, Kai Wang, Wei Ni, Ying Zhang and Wenjie Zhang	22518
<i>NeuReasoner: Towards Explainable, Controllable, and Unified Reasoning via Mixture-of-Neurons</i>	
Haonan Dong, Kehan Jiang, Haoran Ye, Wenhao Zhu, Zhaolu Kang and Guojie Song	22534
<i>Red Teaming Large Reasoning Models</i>	
Jiawei Chen, Yang Yang, Chao Yu, Yu Tian, Zhi Cao, Xue Yang, Linghao Li, Hang Su and Zhaoxia Yin	22559
<i>InsLogicBench: An Argumentation Logic Grounded Benchmark for Complex Insurance Claims Adjudication</i>	
Jin Liu, Yunpeng Liu, Keyi Wang, Jie Shi, Xiao Xu, Wenkang Huang, Xingzhong Xu, Xin Liang and Yanghua Xiao	22592
<i>LinkQA: Synthesizing Diverse QA from Multiple Seeds Strongly Linked by Knowledge Points</i>	
Xuemiao Zhang, Can Ren, Chengying Tu, Rongxiang Weng, Hongfei Yan, Jingang Wang and Xunliang Cai	22620
<i>Comparative Analysis of the Intrinsic Metrics for Tokenizers and their effect on Downstream Tasks for Hindi and Marathi</i>	
Shagun Dwivedi and Kaushik Gopalan	22652

<i>Omni-Embed-Audio: Leveraging Multimodal LLMs for Robust Audio-Text Retrieval</i>	
HaeJun Yoo, Yongseop Shin, Insung Lee, Myoung-Wan Koo and Du-Seong Chang	22664
<i>SPIO: Ensemble and Selective Strategies via LLM-Based Multi-Agent Planning in Automated Data Science</i>	
Wonduk Seo, Juhyeon Lee, Yanjun Shao, Qingshan Zhou, Seunghyun Lee and Yi Bu	22681
<i>Comparing human and language models sentence processing difficulties on complex structures</i>	
Samuel Joseph Amouyal, Aya Meltzer-Asscher and Jonathan Berant	22704
<i>Phun-Bench: Evaluating LLMs on Phonological Understanding in Chinese</i>	
Xing Yue, Yongliang Shen and Weiming Lu	22723
<i>Causal2Vec: Improving Decoder-only LLMs as Embedding Models through a Contextual Token</i>	
Ailiang Lin, Zhuoyun Li, Yusong Wang, Kotaro Funakoshi and Manabu Okumura	22769
<i>Enhancing Multimodal Large Language Models for Ancient Chinese Character Evolution Analysis via Glyph-Driven Fine-Tuning</i>	
Rui Song, Lida Shi, Ruihua Qi, Yingji Li and Hao Xu	22789
<i>See2Refine: Vision-Language Feedback Improves LLM-Based eHMI Action Designers</i>	
Ding Xia, Xinyue Gui, Mark Colley, Fan Gao, Zhongyi Zhou, Dongyuan Li, Renhe Jiang and Takeo Igarashi	22805
<i>A Survey of Reinforcement Learning for Large Language Models under Data Scarcity: Challenges and Solutions</i>	
Zhiyin Yu, Yuchen Mou, Juncheng Yan, Junyu Luo, Chunchun Chen, Xing Wei, Yunhui Liu, Hongru Sun, Yuxing Zhang, Jun Xu, Yatao Bian, Ming Zhang, Wei Ye, Tieke He, Jie Yang, Guanjie Zheng, Zhonghai Wu, Bo Zhang, Lei Bai and Xiao Luo	22823
<i>SADA: Bridging In-Context Learning and Fine-Tuning via State-Aligned Distillation Adapters</i>	
Wenhao Gao, Tianlong Wang, Wei Jia, Linhao Zhang, Aiwei Liu, Miao Fan and Zhou Xiao	22847
<i>Anti-Length Shift: Dynamic Outlier Truncation for Training Efficient Reasoning Models</i>	
Wei Wu, Liyi Chen, Congxi Xiao, Tianfu Wang, Qimeng Wang, Chengqiang Lu, Yan Gao, Yiwu, Yao Hu and Hui Xiong	22863
<i>The Proxy Presumption: From Semantic Embeddings to Valid Social Measures</i>	
Baishi Li, Ta Yu, Kelvin J.I. Koa and Ke-Wei Huang	22892
<i>Nested Browser-Use Learning for Agentic Information Seeking</i>	
Baixuan Li, Jialong Wu, Wenbiao Yin, Kuan Li, Zhongwang Zhang, Huifeng Yin, Zhengwei Tao, Liwen Zhang, Pengjun Xie, Jingren Zhou, Yong Jiang, Wentao Zhang and Zhiqiang Gao	22911
<i>CamoQuery: Language-Guided Reasoning Camouflaged Object Segmentation</i>	
Tianxin Han, Qing Dong, Xingwei Wang, Jie Jia, Gang Wu, Bowen Yang and Fu Zhang	22924
<i>Large Language Models Are Bad Dice Players: LLMs Struggle to Generate Random Numbers from Statistical Distributions</i>	
Minda Zhao, Yilun Du and Mengyu Wang	22942
<i>PROMPRINT: Prompt Fingerprinting via First-Token Response for LLM App Cloning Detection</i>	
Jungmin Lee, Peizhuo Lv and Yeonjoon Lee	22960
<i>MAB-DQA: Addressing Query Aspect Importance in Document Question Answering with Multi-Armed Bandits</i>	
Yixin Xiang, Yunshan Ma, Xiaoyu Du, Yibing Chen, Yanxin Zhang and Jinhui Tang	22973

<i>DVCQR: Dual-View Conversational Query Rewriting with Stage-wise Reinforcement Learning</i> Chenyi Li, Xinhui Tu and Zaixiang Wang	22993
<i>CoEvolve: Training LLM Agents via Agent-Data Mutual Evolution</i> Shidong Yang, Ziyu Ma, Tongwen Huang, Yiming Hu, Yong Wang and Xiangxiang Chu..	23015
<i>Make LLMs See Like Investigators, Not Just Think More: The Role of Structured Analysis in Investigative Reasoning</i> Jaewook Lee, Myeong-Cheol Kang and Jong-hun Shin	23037
<i>FormalScience: Scalable Human-in-the-Loop Autoformalisation of Science with Agentic Code Generation in Lean</i> Jordan Meadows, Lan Zhang and Andre Freitas	23059
<i>MUR: Momentum Uncertainty guided Reasoning for Large Language Models</i> Hang Yan, Fangzhi Xu, Rongman Xu, Yifei Li, Jian Zhang, Haoran Luo, Xiaobao Wu, Anh Tuan Luu, Haiteng Zhao, Qika Lin and Jun Liu	23078
<i>BiMol-Diff: A Unified Diffusion Framework for Molecular Generation and Captioning</i> Aditya Hemant Shahane, Anuj Kumar Sirohi, Devansh Arora, Nitin Kumar, Prathosh AP and Sandeep Kumar	23104
<i>PEAP: Proactive Embodied Action Sequence Planning with Joint Understanding of Vision and Audio Perception</i> Tianwei Lan, Jiaqi Wu, Zeming Liu, Zhaoxin Fan, Haifeng Wang and Yuhang Guo	23118
<i>Automated Creativity Evaluation of Language Models Across Open-Ended Tasks</i> Tan Min Sen, Zachary Choy Kit Chun, Syed Ali Redha Alsagoff, Nadya Yuki Wangsajaya, Banerjee Mohor, Swaagat Bikash Saikia and Alvin Chan	23139
<i>Aligning Agents via Planning: A Benchmark for Trajectory-Level Reward Modeling</i> Jiaxuan Wang, Yulan Hu, Wenjin Yang, Zheng Pan, Xin Li and Lan-Zhe Guo	23174
<i>From Synthesis to Clinical Assistance: A Strategy-Aware Agent Framework for Autism Intervention based on Real Clinical Dataset</i> Junhong Lai, Shuzhong Lai, Yanhao Yu, Wanlin Chen, Chenyu Yan, Haifeng Li, Lin Yao and Yueming Wang	23201
<i>How Training Data Shapes the Use of Parametric and In-Context Knowledge in Language Models</i> Minsung Kim, Dong-Kyum Kim, Jea Kwon, Nakyeong Yang, Kyomin Jung and Meeyoung Cha	23242
<i>PAR: Training-Free Positional Perturbation and Attention Recycling for Faithful OCR</i> Yao Yao, Manwen Liao, Weitian Zhang, Zuchao Li and Hai Zhao	23258
<i>BTC-LLM: Efficient Sub-1-Bit LLM Quantization via Learnable Transformation and Binary Codebook</i> Hao Gu, Lujun Li, Hao Wang, Lei Wang, Zheyu Wang, Bei Liu, Jiacheng Liu, Qiyuan Zhu, Sirui Han and Yike Guo	23274
<i>Learning What to Ignore: Mitigating Negative Transfer in Medical Knowledge Fusion via Clinical Task-Adaptive Selection</i> Xinyan Deng, Shoubin Dong and Xiaorou Zheng	23295
<i>MMTutorBench: The First Multimodal Benchmark for AI Math Tutoring</i> Tengchao Yang, Sichen Guo, Mengzhao Jia, Jiaming Su, Yuanyang Liu, Zhihan Zhang and Meng Jiang	23310

<i>GrACE: A Generative Approach to Better Confidence Elicitation and Efficient Test-Time Scaling in Large Language Models</i>	Zhaohan Zhang, Ziquan Liu and Ioannis Patras	23333
<i>AutoGraph-RL: End-to-End Reinforcement Learning for Knowledge Graph Construction</i>	Hong Ting Tsang, Jiabin Bai, Haoyu Huang, Qiao Xiao, Tianshi Zheng, Baixuan Xu, Shujie Liu and Yangqiu Song	23351
<i>Understanding Structured Financial Data with LLMs: A Case Study on Fraud Detection</i>	Xuwei Tan, Yao Ma and Xueru Zhang	23375
<i>Gap-K%: Measuring Top-1 Prediction Gap for Detecting Pretraining Data</i>	Minseo Kwak and Jaehyung Kim	23391
<i>Experience Retrieval-Augmentation with Electronic Health Records Enables Accurate Discharge QA</i>	Justice Ou, Tinglin Huang, Yilun Zhao, Ziyang Yu, Peiqing Lu, Yifei Shen and Rex Ying .	23406
<i>Are they lovers or friends? Evaluating LLMs’ Social Reasoning in English and Korean Dialogues</i>	Eunsu Kim, Junyeong Park, Juhyun Oh, Kiwoong Park, Seyoung Song, A. Seza Dođruöz, Alice Oh and Najoung Kim	23431
<i>SAIR-Comb : A Structure-Aware Iterative Refinement Framework for Combinatorics Autoformalization</i>	Weijie Jiang, Gaolei He, Beibei Xiong, Jianlin Wang and Zhengfeng Yang	23452
<i>Stop Hardening Everything: A Training-Free Neuron-Level Defense for Neural Ranking Models</i>	Yu-An Liu, Ruqing Zhang, Hongru Song, Jiafeng Guo, Yixing Fan and Xueqi Cheng	23474
<i>TED-TTS: Training-Free Intra-Utterance Emotion and Duration Control for Text-to-Speech Synthesis</i>	Qifan Liang, Yuansen Liu, Ruixin Wei, Nan Lu, Junchuan Zhao and Ye Wang	23485
<i>VIB-Probe: Detecting and Mitigating Hallucinations in Vision-Language Models via Variational Information Bottleneck</i>	Feiran Zhang, Yixin Wu, Zhenghua Wang, Xiaohua Wang, Changze Lv, Xuanjing Huang and Xiaoqing Zheng	23509
<i>Detoxification for LLM: From Dataset Itself</i>	Wei Shao, Yihang Wang, Gao yu Zhu, Ziqiang Cheng, Lei Yu, Jiafeng Guo and Xueqi Cheng	23522
<i>When Is Thinking Enough? Early Exit via Sufficiency Assessment for Efficient Reasoning</i>	Yang Xiang, Yixin Ji, Ruotao Xu, Dan Qiao, Zheming Yang, Juntao Li and Min Zhang . . .	23541
<i>Instant Personalized Large Language Model Adaptation via Hypernetwork</i>	Zhaoxuan Tan, Zixuan Zhang, Haoyang Wen, Zheng Li, Rongzhi Zhang, Pei Chen, Fengran Mo, Zheyuan Liu, Qingkai Zeng, Qingyu Yin and Meng Jiang	23557
<i>From Off-Policy to On-Policy: Enhancing GUI Agents via Bi-level Expert-to-Policy Assimilation</i>	Zezhou Wang, Ziyun Zhang, Xiaoyi Zhang, Zhuzhong Qian and Yan Lu	23581
<i>Safeguarding LLM Fine-tuning via Push-Pull Distributional Alignment</i>	Haozhong Wang, Zhuo Li, Yibo Yang, He Zhao, Hongyuan Zha and Dandan Guo	23624
<i>I²B-LPO: Latent Policy Optimization via Iterative Information Bottleneck</i>	Huilin Deng, Hongchen Luo, Yue Zhu, Long Li, Zhuoyue Chen, Xinghao Zhao, Ming LI, Chuyang Zhao, Jihai Zhang, MengChang Wang, Yang Cao and Yu Kang	23647

<i>What Do Prosody and Text Convey? Characterizing How Meaningful Information is Distributed Across Multiple Channels</i>	
Aditya Yadavalli, Tiago Pimentel, Tamar I Regev, Ethan Gotlieb Wilcox and Alex Warstadt	23665
<i>Right at My Level: A Unified Multilingual Framework for Proficiency-Aware Text Simplification</i>	
Jinhong Jeong, Junghun Park and Youngjae Yu	23680
<i>See the Forest for the Trees: Loosely Speculative Decoding via Visual-Semantic Guidance for Efficient Inference of Video LLMs</i>	
Yicheng Ji, Jun Zhang, Jinpeng Chen, Cong Wang, Lidan Shou, Gang Chen and Huan Li	23707
<i>LaoBench: A Large-Scale Multidimensional Lao Benchmark for Large Language Models</i>	
Jian Gao, Richeng Xuan, Zhaolu Kang, Dingshi Liao, Wenxin Huang, Zongmou Huang, Yangdi Xu, Bowen Qin, Zheqi He, Xi Yang, Changjinli and Yonghua Lin	23727
<i>Gradient-Guided Multi-Judge Prompt Optimization</i>	
ChenZhuo Zhao, Xinda Wang, Pu Zhao, Yue Huang, Junting Lu, Ziqian Liu, Qingwei Lin, Saravan Rajmohan and Dongmei Zhang	23744
<i>ASTRA: Adaptive Semantic Tree Reasoning Architecture for Complex Table Question Answering</i>	
Xiaoke Guo, Songze Li, Zhiqiang Liu, Zhaoyan Gong, Yuanxiang Liu, Huajun Chen and Wen Zhang	23774
<i>DEBAR: Mitigating Contextual Bias in Cross-Document Relation Extraction via Dual-Stream Decoupling</i>	
Zhixuan Yang, Fu Zhang, Huangming Xu and Jingwei Cheng	23804
<i>IF-RewardBench: Benchmarking Judge Models for Instruction-Following Evaluation</i>	
Bosi Wen, Yilin Niu, Cunxiang Wang, Xiaoying Ling, Ying Zhang, Pei Ke, Hongning Wang and Minlie Huang	23816
<i>LegalChainReasoner: Grounding Criminal Judicial Opinion Generation via Structured Legal Chains</i>	
Weizhe Shi, Qiqi Wang, Yihong Pan, Qian Liu and Kaiqi Zhao	23844
<i>InsideOut: Measuring and Mitigating Insider–Outsider Bias in Interview Script Generation</i>	
Yixin Wan, Xingrun Chen and Kai-Wei Chang	23864
<i>Beyond Meta-Reasoning: Metacognitive Consolidation for Self-Improving LLM Reasoning</i>	
Ziqing Zhuang, Linhai Zhang, Jiasheng Si, Deyu Zhou and Yulan He	23884
<i>Mnemis: Dual-Route Retrieval on Hierarchical Graphs for Long-Term LLM Memory</i>	
Zihao Tang, Xin Yu, Ziyu Xiao, Zengxuan Wen, Zelin Li, Jiayi Zhou, Hualei Wang, Haohua Wang, Haizhen Huang, Weiwei Deng, Feng Sun and Qi Zhang	23914
<i>Consolidation or Adaptation? PRISM: Disentangling SFT and RL Data via Gradient Concentration</i>	
Yang Zhao, Yangou Ouyang, Xiao Ding, Hepeng Wang, Bibo Cai, Kai Xiong, Jinglong Gao, Zhouhao Sun, Li Du, Bing Qin and Ting Liu	23929
<i>ADVICE: Answer-Dependent Verbalized Confidence Estimation</i>	
KiJung Seo, Sehun Lim and Taeuk Kim	23942
<i>Powering Verifiable Learning via Automated Evolutionary Data Synthesis</i>	
He Du, Bowen Li, Aijun Yang, Siyang He, Qipeng Guo, Kai Chen and Dacheng Tao	23963
<i>SPARK: Strategic Policy-Aware Exploration via Dynamic Branching for Long-Horizon Agentic Learning</i>	
Jinyang Wu, Shuo Yang, Yuhao Shen, Shuai Zhang, Zhengqi Wen and Jianhua Tao	23981

<i>Scaling Laws for Code: A More Data-Hungry Regime</i>	
Xianzhen Luo, Wenzhen Zheng, Qingfu Zhu, Rongyi Zhang, Houyi Li, Siming Huang, YuanTao Fan and Wanxiang Che	24005
<i>When TableQA Meets Noise: A Dual Denoising Framework for Complex Questions and Large-scale Tables</i>	
Shenghao Ye, Yu Guo, Dong Jin, Yuxiang Wang, Yikai Shen, Yunpeng Hou, Shuangwu Chen, Jianyang and Xiaofeng Jiang	24022
<i>VRPO: Rethinking Value Modeling for Robust RL under Noisy Supervision in LLM Post-Training</i>	
Dingwei Zhu, Shihan Dou, Zhiheng Xi, Senjie Jin, Guoqiang Zhang, Jiazheng Zhang, Junjie Ye, Mingxu Chai, Enyu Zhou, Ming Zhang, Yuhui Wang, Caishuang Huang, Chenhao Huang, Yunke Zhang, Yuran Wang, Tao Gui, Qi Zhang, Xipeng Qiu and Xuanjing Huang	24046
<i>Taming Actor-Observer Asymmetry in Agents via Dialectical Alignment</i>	
Bobo Li, Wu Rui, Zibo Ji, Meishan Zhang, Hao Fei, Min Zhang, Mong-Li Lee and Wynne Hsu	24068
<i>Knowledge Beyond Language: Bridging the Gap in Multilingual Machine Unlearning Evaluation</i>	
Kyomin Hwang, Hyeonjin Kim, Sangyeon Cho and Nojun Kwak	24085
<i>AT²PO: Agentic Turn-based Policy Optimization via Tree Search</i>	
Zefang Zong, Dingwei Chen, Yang Li, Qi Yi, Bo Zhou, Chengming Li, BO Qian, Peng Chen and Jie Jiang	24120
<i>SecureVibeBench: Benchmarking Secure Vibe Coding of AI Agents via Reconstructing Vulnerability-Introducing Scenarios</i>	
Junkai Chen, Huihui Huang, Yunbo Lyu, Junwen An, Jieke Shi, Chengran Yang, Ting Zhang, Haoye Tian, Yikun Li, Zhenhao Li, Xin Zhou, Xing Hu and David Lo	24144
<i>Text-Guided Multi-Scale Frequency Representation Adaptation</i>	
Weicai Yan, Xinhua Ma, Wang Lin and Tao Jin	24169
<i>ADAPT: Benchmarking Commonsense Planning under Unspecified Affordance Constraints</i>	
Pei-An Chen, Yongching Liang, Jia-Fong Yeh, Hung-Ting Su, Yi-Ting Chen, Min Sun and Winston H. Hsu	24188
<i>GeoRA: Geometry-Aware Low-Rank Adaptation for RLVR</i>	
Jiaying Zhang, Lei Shi, Jiguo Li, Jun Xu, Jiuchong Gao, Jinghua Hao and Renqing He ...	24207
<i>KoCo: Conditioning Language Model Pre-training on Knowledge Coordinates</i>	
Yudong Li, Jiawei Cai and Linlin Shen	24222
<i>Truth or Sophistry? LoFa: A Benchmark for LLM Robustness Against Logical Fallacies</i>	
Xudong Shen, li Yuan, Ye Chen, Xin Wu, Yi Cai and Zhiyong Wu	24236
<i>Merging Triggers, Breaking Backdoors: Defensive Poisoning for Instruction-Tuned Language Models</i>	
San Kim and Gary Lee	24269
<i>SHAPE: Stage-aware Hierarchical Advantage via Potential Estimation for LLM Reasoning</i>	
Zhengyang Ai, Zikang Shan, Xiaodong Ai, Jingxian Tang, Hangkai Hu and Pinyan Lu ...	24288
<i>Looking Beyond the One: Operationalizing and Eliciting Visual Ambiguity in VLLMs</i>	
Yuchong Chen, Bowei Zou, Yuhan Chen, Yifan Fan, Xinyu Li, Shujun Cao and Yu Hong .	24306
<i>JanusMM: A Benchmark for Self-Deprecation Understanding in Real-World Multimodal Conversations</i>	
Xinyi Xu, Bingguang Hao, Yongyi Xiong, Zimo Chen, Xinchen Liu, Hongxin Guo, Xuelong Wang, Silin Zhou and Shihan Dou	24324

<i>Attention as Selector: Unlocking VLM Attention for Long Document Page Retrieval</i> Minfeng Zhu, Linxin Bao, Wei Chen and Linchao Zhu	24344
<i>LAMCL: A Length-aware Momentum Contrastive Learning Framework for Multiscale Machine-Revised Text Detection</i> Bing Zhou, Zhe Huang, Shilei Tan, Kai Zhao and Zhou Yongcheng	24366
<i>Ascending the Infinite Ladder: Benchmarking Spatial Deformation Reasoning in Vision-Language Models</i> Jiahuan Zhang, Shunwen Bai, Tianheng Wang, KaiWen Guo, Zijia Song, Hanqing WU, Guozheng Rao, Kai Han and Kaicheng Yu.....	24381
<i>Safe-FedLLM: Delving into the Safety of Federated Large Language Models</i> Mingxiang Tao, Yu Tian, Wenxuan Tu, Yue Yang, Xue Yang and Xiangyan Tang.....	24405
<i>TEMA: Anchor the Image, Follow the Text for Multi-Modification Composed Image Retrieval</i> Zixu Li, Yupeng Hu, Zhiheng Fu, Zhiwei Chen, Yongqi Li and Liqiang Nie	24421
<i>ImCoref-CeS: An Improved Lightweight Pipeline for Coreference Resolution with LLM-based Checker-Splitter Refinement</i> Kangyang Luo, Yuzhuo Bai, Shuzheng Si, Cheng Gao, Zhitong Wang, Yingli Shen, Wenhao Li, Zhu Liu, Yufeng Han, Jiayi Wu, Cunliang Kong and Maosong Sun.....	24443
<i>When Efficiency Meets Safety: A Benchmark Security Analysis of KV Cache Compression in Large Language Models</i> Xiaoxiao Ma, Kuofeng Gao, Zeyi Lu, Wenxi Jiang, Hao Fang, Hao Wu, Bin Chen and Shu-Tao Xia.....	24472
<i>WebAggregator: Enhancing Compositional Reasoning Capabilities of Deep Research Agent Foundation Models</i> Rui Wang, Ce Zhang, Jun-Yu Ma, Jianshu Zhang, Hongru Wang, Yi Chen, Boyang Xue, Tianqing Fang, Zhisong Zhang, Hongming Zhang, Haitao Mi, Dong Yu and Kam-Fai Wong	24486
<i>SEARL: Joint Optimization of Policy and Tool Graph Memory for Self-Evolving Agents</i> Xinshun Feng, Xinhao Song, Lijun Li, Gongshen Liu and Jing Shao	24518
<i>Good Arguments Against the People Pleasers: How Reasoning Mitigates (Yet Masks) LLM Sycophancy</i> Zhaoxin Feng, Zheng Chen, Jianfei Ma, Yip Tin Po, Emmanuele Chersoni and Bo Li	24536
<i>Can Persona-Prompted LLMs Emulate Subgroup Values? An Empirical Analysis of Generalisability and Fairness in Cultural Alignment</i> Bryan Chen Zhengyu Tan, Zhengyuan Liu, Xiaoyuan Yi, Jing Yao, Xing Xie, Nancy F. Chen and Roy Ka-Wei Lee.....	24571
<i>FoE: Forest of Errors Makes the First Solution the Best in Large Reasoning Models</i> Kehan Jiang, Haonan Dong, Zhaolu Kang, Zhengzhou Zhu and Guojie Song	24591
<i>Lingua-Graph: A Unified Representation of Cross-Task Common Substructures for Analytic Language Processing</i> Mingming Sun, Runze Jiang, Zhu Zhangchenxi, Minlong Peng and Yunfeng Cai	24630
<i>Adaptive Spatial and Temporal Redundancy Optimization for Efficient Reasoning in Large Language Models</i> Tianle Chen, Pengyu Cheng, Qiyuan Zhu, Jiacheng Wang, Bei Liu, Hao Gu, Ruijie Shen, Xiaofeng Hou, Sirui Han and Jiacheng Liu	24647

<i>AAPO: Enhancing the Reasoning Capabilities of LLMs with Advantage Margin</i>	
Jian Xiong, Jingbo Zhou, Jingyong Ye, Qiang Huang and Dejing Dou	24663
<i>Outcome-Grounded Advantage Reshaping for Fine-Grained Credit Assignment in Mathematical Reasoning</i>	
Ziheng Li, Liu Kang, Feng Xiao, Luxi Xing, Qingyi Si, Zhuoran Li, Weikang Gong, Deqing Yang, Yanghua Xiao and Hongcheng Guo	24681
<i>Learning to Edit Knowledge via Instruction-based Chain-of-Thought Prompting</i>	
Jinhu Fu, Yan Bai, Longzhu He, Yihang Lou, Yanxiao Zhao, Li Sun and Sen Su	24694
<i>BadScientist: Can a Research Agent Write Convincing but Unsound Papers that Fool LLM Reviewers?</i>	
Fengqing Jiang, Yichen Feng, Yuetai Li, Luyao Niu, Basel Alomair and Radha Poovendran	24712
<i>DFAMS: Dynamic-flow guided Federated Alignment based Multi-prototype Search</i>	
Zhibang Yang, Xinke Jiang, Rihong Qiu, Ruiqing Li, Yihang Zhang, Yue Fang, Yongxin Xu, Hongxin Ding, Xu Chu, Junfeng Zhao and Yasha Wang	24728
<i>PARIF: Pushing the Pareto Frontier of Instruction Following and Reasoning with Curriculum Reinforcement Learning</i>	
Rongchuan Mu, Zexin Wang, Qianyu Wang, MingHua Ma, Zekun Wang, Ming Liu and Bing Qin	24753
<i>GeoLauX: A Benchmark for Evaluating MLLMs' Geometry Performance on Long-Step Problems Requiring Auxiliary Lines</i>	
Yumeng Fu, Jiayin Zhu, Lingling Zhang, Wenjun Wu, Bo Zhao, Shaoxuan Ma, Yushun Zhang and Jun Liu	24784
<i>Unlocking Multilingual Reasoning Capability of LLMs and LVLMs through Representation Engineering</i>	
Qiming Li, Xiaocheng Feng, Yixuan Ma, Ruihan Chen, Zihe Tong, Zekai Ye, Xiachong Feng, Libo Qin, Haoyu Ren, Kun Chen, Yunfei Lu, Dandan Tu and Bing Qin	24810
<i>Into the Gray Zone: Domain Contexts Can Blur LLM Safety Boundaries</i>	
Ki Sen Hung, Xi Yang, Chang Liu, Haoran Li, Kejiang Chen, Changxuan Fan, Tsun On Kwok, Weiming Zhang, Xiaomeng Li and Yangqiu Song	24830
<i>MED-COREASONER: Reducing Language Disparities in Medical Reasoning via Language-Informed Co-Reasoning</i>	
Fan Gao, Sherry T. Tong, Jiwoong Sohn, Jiahao Huang, Junfeng Jiang, Ding Xia, Piyalitt Ittichaiwong, Kanyakorn Veerakanjana, Hyunjae Kim, Qingyu Chen, Edison Marrese-Taylor, Kazuma Kobayashi, Akiko Aizawa and Irene Li	24868
<i>Psyche-RI: Towards Reliable Psychological LLMs through Unified Empathy, Expertise, and Reasoning</i>	
Chongyuan Dai, Jinpeng Hu, Hongchang Shi, Zhuo Li, Dan Guo, Xun Yang and Meng Wang	24889
<i>Mitigating Action-Relation Hallucinations in LVLMs via Relation-aware Visual Enhancement</i>	
Zhenxin Qin, Qiang Li, Qingzhuo Wang, Ruiyang Qin, Zhihua Wei and Wen Shen	24907
<i>Agentic Oversight via Dialectic Reasoning</i>	
Leonardo Ranaldi and Federico Ranaldi	24924
<i>FOREVER: Forgetting Curve-Inspired Memory Replay for Language Model Continual Learning</i>	
Yujie Feng, Hao Wang, Jian Li, Xu Chu, Zhaolu Kang, Yiran Liu, Yasha Wang, Philip S. Yu and Xiao-Ming Wu	24945

<i>CharTide: Data-Centric Chart-to-Code Generation via Tri-Perspective Tuning and Inquiry-Driven Evolution</i>	
Xiangxi Zheng, Kuang He, Jiayi Hu, Ping Yu, Rui Yan, Yuan Yao, Peng Hou, Anxiang Zeng and Alex Jinpeng Wang	24966
<i>PRISP: Privacy-Safe Few-Shot Personalization via Lightweight Adaptation</i>	
Junho Park, Dohoon Kim and Taesup Moon	24986
<i>IF-CRITIC: Towards a Fine-Grained LLM Critic for Instruction-Following Evaluation</i>	
Bosi Wen, Yilin Niu, Cunxiang Wang, Pei Ke, Xiaoying Ling, Ying Zhang, Aohan Zeng, Hongning Wang and Minlie Huang	25004
<i>Know Your Place: Diagnosing Implicit Social Adaptation Failures in Chinese Large Language Models</i>	
Yu Tian, Jie Xing, Ziming Li, Jiang Li, Zehua Duo, Tian Lan, Xu Liu, Guanglai Gao and Xiangdong Su	25030
<i>Agent Newsroom: Efficient Chronological Report Generation via Dynamic Multi-Agent Collaboration</i>	
Zhenhua Wang, Chunlei Wang, Yue Geng and Bang Wang	25063
<i>Locomo-Plus: Beyond-Factual Cognitive Memory Evaluation Framework for LLM Agents</i>	
Yifei Li, Weidong Guo, Lingling Zhang, Rongman Xu, Muye Huang, Hui Liu, Lijiao Xu, Yu Xu and Jun Liu	25085
<i>CUB: Benchmarking Context Utilisation Techniques for Language Models</i>	
Lovisa Hagström, Youna Kim, Haeun Yu, Sang-goo Lee, Richard Johansson, Hyunsoo Cho and Isabelle Augenstein	25101
<i>VLN-NF: Feasibility-Aware Vision-and-Language Navigation with False-Premise Instructions</i>	
Hung-Ting Su, Ting-Jun Wang, Jia-Fong Yeh, Min Sun and Winston H. Hsu	25134
<i>NiuTrans.LMT: Toward Inclusive and Scalable Multilingual Machine Translation with LLMs</i>	
Yingfeng Luo, Ziqiang Xu, Yuxuan Ouyang, MuRun Yang, DingYang Lin, Kaiyan Chang, Tong Zheng, Bei Li, Peinan Feng, Quan Du, Tong Xiao and JingBo Zhu	25151
<i>ChipSeek: Optimizing Verilog Generation via EDA-Integrated Reinforcement Learning</i>	
Zhirong Chen, Kaiyan Chang, Zhuolin Li, Cangyuan Li, Xinyang He, Chujie Chen, Mengdi Wang, Haobo Xu, Yinhe Han, Huawei Li and Ying Wang	25180
<i>MERIT Feedback Elicits Better Bargaining in LLM Negotiators</i>	
Jihwan Oh, Murad Aghazada, Yooju Shin, Se-Young Yun and Taehyeon Kim	25202
<i>Collision to Cognition: Hash-Driven Graph Construction for Efficient RAG</i>	
Chuang Zhou, Zheng Yuan, Linhao Luo, Zhaozhuo Xu, Yilin Xiao, Junnan Dong, Siyu An, di Yin, Xing Sun and Xiao Huang	25224
<i>WebSynthesis: World Model-Guided Monte Carlo Tree Search for Efficient WebAgent Trajectory Synthesis</i>	
Yifei Gao, Junhong Ye, Yifan Yang, Jiaqi Wang, Yi Zhang, Zhang Ruichen and Jitao Sang	25241
<i>EyeMulator: Improving Code Language Models by Mimicking Human Visual Attention</i>	
Yifan Zhang, Chen Huang, Yueke Zhang, Jiahao Zhang, Toby Jia-Jun Li, Collin McMillan, Kevin Leach and Yu Huang	25259
<i>TPA: Next Token Probability Attribution for Detecting Hallucinations in RAG</i>	
Pengqian Lu, Jie Lu, Anjin Liu and Guangquan Zhang	25273

<i>GMFL: Efficient Global Masking for Federated LLM Fine-tuning</i> Xin Huang, Yan Hu, Yue-Jiao Gong and Xinglin Zhang	25293
<i>AgentGL: Towards Agentic Graph Learning with LLMs via Reinforcement Learning</i> Yuanfu Sun, Kang Li, Dongzhe Fan, Jiajin Liu and Qiaoyu Tan	25313
<i>MM-JudgeBias: A Benchmark for Evaluating Compositional Biases in MLLM-as-a-Judge</i> Sua Lee, Sanghee Park and Jinbae Im	25336
<i>CSPO: Alleviating Reward Ambiguity for Structured Table-to-LaTeX Generation</i> Yunfan Yang, Cuiling Lan, Jitao Sang and Yan Lu	25374
<i>Chart-MRAG: Benchmarking Multimodal Retrieval Augmented Generation on Chart-based Documents</i> Ymyang, Jiang Zhong, Li Jin, Xiao Sun, Jingwang Huang, Gaojinpeng, Qing Liu, Yang Bai, Jingyuan Zhang, Rui Jiang, Qin Lei and Kaiwen Wei	25392
<i>Towards Stable and Effective Reinforcement Learning for Mixture-of-Experts</i> Di Zhang, Xun Wu, Shaohan Huang, Lingjie Jiang, Yaru Hao, Li Dong, Zewen Chi, Zhifang Sui and Furu Wei	25446
<i>AIRCoder: Adaptive Integration of Multi-dimensional Retrieval for Repository-level Code Completion</i> Chuanqi Shi, Miao Gao and Zhiqiang Gao	25458
<i>ReLook: Vision-Grounded RL with a Multimodal LLM Critic for Agentic Web Coding</i> Yuhang Li, Chenchen Zhang, Ruilin Lv, Ao Liu, Ken Deng, Yuanxing Zhang, Jiaheng Liu and Bo Zhou	25471
<i>SafetyMem: Adaptive Jailbreak Defense via Dual-Component Safety Memory</i> Hao Wang, Ziyi Ni, Huacan Wang, Pin Lyu and Lei Sha	25486
<i>Enhancing Agentic Textual Graph Retrieval with Synthetic Stepwise Supervision</i> Ge Chang, Jinbo Su, Jiacheng Liu, Pengfei Yang, Yuhao Shang, Huiwen Zheng, Hongli Ma, Yan Liang, Yuanchun Li and Yunxin Liu	25510
<i>Cognitive Scaffold: From Fluid Context to Crystallized Memory for Long-Horizon DeepResearch Agents</i> Qiuyuan Ai, Zenghuang Fu, Zhaoyang Li, Ping Jiang, Haoyu Wu, Jie Song and Guannan He	25526
<i>Revisiting Metric Reliability for Fine-grained Evaluation of Machine Translation and Summarization in Indian Languages</i> Amir Hossein Yari, Kalmit Kulkarni, Ahmad Raza Khan and Fajri Koto	25543
<i>Learning to Think on Hypergraph: HyperCoT for Structure-Guided N-ary Knowledge Graph Completion</i> Mengxue Yang, Jinming Li, Chun Yang, Jiaqi Zhu, Jiafan Li, Guanhua Zhang and Ying Li	25562
<i>Two Pathways to Truthfulness: On the Intrinsic Encoding of LLM Hallucinations</i> Wen Luo, Guangyue Peng, Wei Li, Shaohang Wei, Feifan Song, Liang Wang, Nan Yang, Xingxing Zhang, Jing Jin, Furu Wei and Houfeng Wang	25580
<i>LLM-ForcedAligner: A Non-Autoregressive and Accurate LLM-Based Forced Aligner for Multilingual and Long-Form Speech</i> Bingshen Mu, Xian Shi, Xiong Wang, Hexin Liu, Jin Xu and Lei Xie	25627
<i>Forget What Matters, Keep the Rest: Selective Unlearning of Informative Tokens</i> Seunghee Koh, Sunghyun Baek, Youngdong Kim and Junmo Kim	25639

<i>Latent-Condensed Transformer for Efficient Long Context Modeling</i>	
Zeng You, Yaofu Chen, Qiuwu Chen, Ying Sun, Shuhai Zhang, Yingjian Li, Yaowei Wang and Mingkui Tan	25656
<i>MedForge: Interpretable Medical Deepfake Detection via Forgery-aware Reasoning</i>	
Zhihui Chen, Kai He, Qingyuan Lei, Bin Pu, Jian Zhang, Yuling Xu and Mengling Feng ..	25672
<i>CrossGuard: Safeguarding MLLMs against Joint-Modal Implicit Malicious Attacks</i>	
Xu Zhang, Hao Li and Zhichao Lu	25693
<i>Can Spectral-Clipping Enable Better Learning While Forgetting Less for Low-Rank Adaptation?</i>	
Hyowon Wi and Noseong Park	25708
<i>Re³: Relevance & Recency Retrieval for Mitigating Temporal Hallucination</i>	
Jiawei Cao, Jie Ouyang, Mingyue Cheng, Zhaomeng Zhou, Chunli Liu, Yupeng Li, Zirui Liu and Shijin Wang	25735
<i>TeCES: Collaborative Geometric Knowledge Representation Framework under Evolving Fact Snapshots</i>	
Jiujiang Guo, Zhengliang Guo, Kai Wang, Meiyang Wang, Dehua Peng, Shaozu Yuan, Chengyin Hu, Shuan Ai and Yiwei Wei	25761
<i>MicroC-KT: Modeling Community Effect via Learning Micro-Environment for Evidence-Grounded Explainable Knowledge Tracing</i>	
Zhiyi Duan, Zixing Shi, Bing Jia and Qi Wang	25781
<i>ImF: Embedding an Implicit Fingerprint in Your Large Language Models</i>	
Jiaxuan Wu, Wanli Peng, Hang fu, Xue Yiming and Juan Wen	25804
<i>Temp-RL: A Unified Autonomous Agent for Complex Temporal KGQA via Reverse Curriculum Reinforcement Learning</i>	
Zhaoyan Gong, Zhiqiang Liu, Songze Li, Xiaoke Guo, Yuanxiang Liu, Xinle Deng, Zhizhen Liu, Lei Liang, Huajun Chen and Wen Zhang	25826
<i>S2G-RAG: Structured Sufficiency and Gap Judging for Iterative Retrieval-Augmented QA</i>	
Minghan Li, Junjie Zou, Xinxuan Lv, Chao Zhang and Guodong Zhou	25846
<i>Jailbreaking Multimodal Large Language Models using Multi-Clip Video</i>	
Choongwon Kang, Seungjong Sun, Hyunmin Jun and Jang Hyun Kim	25863
<i>HAG: Hierarchical Demographic Tree-based Agent Generation for Topic-Adaptive Simulation</i>	
Rongxin Chen, Tianyu Wu, Bingbing Xu, JiaTang Luo, Xiucheng Xu and Huawei Shen ..	25890
<i>Stable-RAG: Mitigating Retrieval-Permutation-Induced Hallucinations in Retrieval-Augmented Generation</i>	
Qianchi Zhang, Hainan Zhang, Liang Pang, Hong-Wei Zheng and Zhiming Zheng	25907
<i>From Logical to Computational Sparsity: Structure-Aware Block-Sparse Attention for Long-Code Completion</i>	
Yanli Wang, Yanlin Wang, Bowen Zhang, Yiwei Zhang, Daya Guo, Jiachi Chen, Hongyu Zhang and Zibin Zheng	25927
<i>CogEvolve: A Multimodal Benchmark for Evaluating Relational Reasoning in Semantic Extension</i>	
Jingjie Zeng, Huayang Li, Liang Yang, Yuanyuan Sun, Shaowu Zhang and Hongfei Lin ..	25943
<i>Towards Privacy-Preserving Large Language Model: Text-free Inference Through Alignment and Adaptation</i>	
Jeongho Yoon, Chanhee Park, Yongchan Chun, Hyeonseok Moon and Heuseok Lim	25961

<i>VisRet: Visualization Improves Knowledge-Intensive Text-to-Image Retrieval</i> Di Wu, Yixin Wan and Kai-Wei Chang	25982
<i>Selective Knowledge Distillation: Fusing LLM Semantic Strengths with DNN Efficiency for Binary Code Similarity Detection</i> Shize Zhou, Peiyu Liu, Lirong Fu, Tong Ye and Wenhai Wang	26002
<i>Mixture-of-Experts with Intermediate CTC Supervision for Accented Speech Recognition</i> Wonjun Lee, Hyounghun Kim and Gary Lee	26015
<i>TransLLM: A Unified Multi-Task Large Language Model for Urban Transportation via Learnable Prompting</i> Jiaming Leng, Yunying Bi, Chuan Qin, Zhenya Huang, Bing Yin, Haojie Ren, Yanyong Zhang and Chao Wang	26028
<i>CoreCodeBench: Decoupling Code Intelligence via Fine-Grained Repository-Level Tasks</i> Lingyue Fu, Hao Guan, Bolun Zhang, Haowei Yuan, Yaoming Zhu, Lin Qiu, Zongyu Wang, Xuezhi Cao, Xunliang Cai, Weiwen Liu, Weinan Zhang and Yong Yu	26043
<i>Stop Fixating on Prompts: Reasoning Hijacking and Constraint Tightening for Red-Teaming LLM Agents</i> Yanxu Mao, Peipei Liu, Tiehan Cui, Congying Liu, Mingzhe Xing and Datao You	26068
<i>Attention Basin: Why Contextual Position Matters in Large Language Models</i> Zihao Yi, Zhenqing Ling, Delong Zeng, Haohao Luo, Zhe Xu, Wei Liu, Jian Luan, Wanxia Cao and Ying Shen	26086
<i>RRAttention: Dynamic Block Sparse Attention via Per-Head Round-Robin Shifts for Long-Context Inference</i> Siran Liu, Guoxia Wang, Sa Wang, Jinle Zeng, Haoyang Xie, Siyu Lou, Jiabin Yang, Dianhai Yu, Haifeng Wang and Chao Yang	26107
<i>Beyond Surface Features: Advancing Medical Vision-Language Alignment via Dynamic Evidence-Guided Preference Optimization</i> Zixuan Huang, Zhihong Zhu, Xiaolong Liu, Yanchao Hao, Manman Zhang, Zheng Wei, Bowen Xing, Xian Wu, Ye Li, Fen Miao and Yefeng Zheng	26125
<i>SAFE-QAQ: End-to-End Slow-Thinking Audio-Text Fraud Detection via Reinforcement Learning</i> Peidong Wang, Zhiming Ma, Xin Dai, Yongkang Liu, Shi Feng, Xiaocui Yang, Wenxing Hu, Zhihao Wang, Mingjun Pan, Li Yuan and Daling Wang	26138
<i>MaDS: Long-Horizon GUI Automation via Synergizing Dual-Layer Memory and Multi-Round Debate</i> Pengchen Chen, Shi Chen, Qiming Ye, Xinli Chen, Xinran Li and Wei Xiang	26158
<i>CIRAG: Construction–Integration Retrieval and Adaptive Generation for Multi-hop Question Answering</i> Zili Wei, Yilin Wang, Xiaocui Yang, Shi Feng, Weidong Bao, Daling Wang, Zihan Wang and Yifei Zhang	26181
<i>TinyJudge: Unverifiable Constraint Alignment via Lightweight Specialist Ensembles</i> Yirong Zeng, Yufei Liu, Xiao Ding, Yutai Hou, Yuxian Wang, Wu Ning, Haonan Song, Dandan Tu, Qixun Zhang, Yuxiang He, Bibo Cai and Ting Liu	26198
<i>Dynamic Infilling Anchors for Format-Constrained Generation in Diffusion Large Language Models</i> Boyan Han, Yiwei Wang, Yi Song, Yujun Cai and Chi Zhang	26213

<i>Compiling Activation Steering into Weights via Null-Space Constraints for Stealthy Backdoors</i>	
Rui Yin, Tianxu Han, Naen Xu, Changjiang Li, Ping He, Chunyi Zhou, Jun Wang, Zhihui Fu, Tianyu Du, Jinbao Li and Shouling Ji	26228
<i>Inhibitory Attacks on Backdoor-based Fingerprinting for Large Language Models</i>	
Hang fu, Wanli Peng, Yinghan Zhou, Jiakuan Wu, Juan Wen and Xue Yiming	26246
<i>TimeSAF: Towards LLM-Guided Semantic Asynchronous Fusion for Time Series Forecasting</i>	
Fan Zhang, Shiming Fan and Hua Wang	26267
<i>Pruning Unsafe Tickets: A Resource-Efficient Framework for Safer and More Robust LLMs</i>	
Wai Man Si, Mingjie Li, Michael Backes and Yang Zhang	26285
<i>EQUIP: EQUIvariant preserving In-Place updates for Efficient Token Pruning</i>	
Arun Ramachandran, R Govindarajan, Murali Annavaram and Prakash Raghavendra	26303
<i>Exons-Detect: Identifying and Amplifying Exonic Tokens via Hidden-State Discrepancy for Robust AI-Generated Text Detection</i>	
Xiaowei Zhu, Yubing Ren, Fang Fang, Shi Wang, Yanan Cao and Li Guo	26324
<i>Reasoning Fails Where Step Flow Breaks</i>	
Xiaoyu Xu, Yulan Pan, Xiaosong Yuan, Zhihong Shen, Minghao Su, Yuanhao Su and Xiaofeng Zhang	26337
<i>Cat-MoD: Accelerating Multimodal Alignment via Caption Token Guided Asymmetric Mixture-of-Depths</i>	
YiJie Huang, Xiaocui Yang, Shi Feng, Wen Zhang, Kaisong Song, Yifei Zhang and Daling Wang	26354
<i>Beyond Scaling: Measuring and Predicting the Upper Bound of Knowledge Retention in Language Model Pre-Training</i>	
Changhao Jiang, Ming Zhang, Yifei Cao, Junjie Ye, Xiaoran Fan, Shihan Dou, Zhiheng Xi, Jiajun Sun, Yi Dong, Yujiong Shen, Jingqi Tong, Baoyu Fan, Tao Gui, Qi Zhang and Xuanjing Huang .	26377
<i>Efficient Paths and Dense Rewards: Probabilistic Flow Reasoning for Large Language Models</i>	
Yan Liu, Feng Zhang, Zhanyu Ma, Jun Xu, Jiuchong Gao, Jinghua Hao, Renqing He, Han Liu and Yangdong Deng	26397
<i>TriPlay-RL: Tri-Role Self-Play Reinforcement Learning for LLM Safety Alignment</i>	
Zhewen Tan, Wenhan Yu, Jianfeng Si, Tongxin Liu, Kaiqi Guan, Huiyan Jin, Jiawen Tao, Xiaokun Yuan, Xiangzheng Zhang, Duohe Ma, Tong Yang and Lin Sun	26415
<i>Cross-Modal Coreference Alignment: Enabling Reliable Information Transfer in Omni-LLMs</i>	
Hongcheng Liu, Yuhao Wang, Zhe Chen, Pingjie Wang, Zhiyuan Zhu, Yixuan Hou, Yanfeng Wang and Yu Wang	26430
<i>Inflated Excellence or True Performance? Rethinking Medical Diagnostic Benchmarks with Dynamic Evaluation</i>	
Xiangxu Zhang, Lei Li, Yanyun Zhou, Xiao Zhou, Yingying Zhang and Xian Wu	26454
<i>TrendFact: A Benchmark Towards Hotspot Perception in Automatic Fact-Checking</i>	
Xiaocheng Zhang, Xi Wang, Yifei Lu, Jianing Wang, Zhuangzhuang Ye, Mengjiao Bao, Peng Yan and Xiaohong Su	26494
<i>MPBoCo: Multimodal Prompt-based Boundary-enhanced Continual Framework for Joint Entity and Relation Extraction</i>	
Guanglu Sun, Xinyu Liu, Lili Liang, Yang Yu, Fei Lang, Suxia Zhu and Ming Liu	26514

<i>ConfSpec: Efficient Step-Level Speculative Reasoning via Confidence-Gated Verification</i> Siran Liu, Zane Cao and Yongchao He	26525
<i>Personalizing LLMs with Binary Feedback: A Preference-Calibrated Optimization Framework</i> Xilai Ma, Liye Zhao, Weijun Yao, Haibing Di, Wenya Wang and Jing Li	26539
<i>Localized Low-Rank Adaptation within Clustered Parameter Subspaces</i> Jiahao Xiong, Yihe Liu, Xianming Hu, Hongbo Zhao, Nuoyi Chen, Jie Zhang and Kai Zhang	26556
<i>SGVEF-LOOP: Coverage-Guided Progressive Topological Exploration and Fact-Grounded Metamorphic Evaluation for MCP Agents</i> Zenghao Liu and Yansong Zhang	26573
<i>MentalSeek-Dx: Towards Progressive Hypothetico-Deductive Reasoning for Real-world Psychiatric Diagnosis</i> Xiao Sun, Ymyang, Xinyi Jiang, Yu Tian, Junnan Zhu, Jiang Zhong, Qin Lei, Jingwang Huang, Haoyang Zeng, Xinyu Zhou, Xin Xiao and Kaiwen Wei	26600
<i>Think Parallax: Solving Multi-Hop Problems via Multi-View Knowledge-Graph-Based Retrieval-Augmented Generation</i> Jinliang Liu, Jiale Bai and Shaoning Zeng	26637
<i>AIM-CoT: Active Information-driven Multimodal Chain-of-Thought for Vision-Language Reasoning</i> Xiping Li and Jianghong Ma	26656
<i>Benchmarking Egocentric Clinical Intent Understanding Capability for Medical Multimodal Large Language Models</i> Shaonan Liu, Guo Yu, Xiaoling Luo, Shiyi Zheng, Jie Liu, Wenting Chen and Linlin Shen	26682
<i>Data Pollination: An Emergent Ecological Process Driving AI Population Evolution</i> Shufang Xie, Qizhi Pei, Ang Lv, Jingyang Hu, Lijun Wu and Rui Yan	26698
<i>Extending First-Order Logic for Factual Reasoning over Knowledge Graphs</i> Yuanzhen Hao and Desheng Wu	26722
<i>Diffuse Thinking: Exploring Diffusion Language Models as Efficient Thought Proposers for Reasoning</i> Chenyang Shao, Sijian Ren, Fengli Xu and Yong Li	26739
<i>Does Memory Need Graphs? A Unified Framework and Empirical Analysis for Long-Term Dialog Memory</i> Sen Hu, Yuxiang Wei, Jiaxin Ran, Xueran Han, Zhiyuan Yao, Huacan Wang, Ronghao Chen and Lei Zou	26758
<i>False Friends in the Shell: Unveiling the Emoticon Semantic Confusion in Large Language Models</i> Weipeng Jiang, Xiaoyu Zhang, Juan Zhai, Shiqing Ma, Chao Shen and Yang Liu	26783
<i>DocLens: A Tool-Augmented Multi-Agent Framework for Long Visual Document Understanding</i> Dawei Zhu, Rui Meng, Jiefeng Chen, Sujian Li, Tomas Pfister and Jinsung Yoon	26804
<i>Stability Implies Redundancy: Delta Attention Selective Halting for Efficient Long-Context Prefilling</i> Yujie Chen, Tailai Chen, Yifeng Gao, Zoe Wanying He, Yijue Xu, Shaobo Wang and Linfeng Zhang	26830
<i>Beyond Self-Report: Bridging the Intention-Behavior Gap in Critical Thinking Assessment via Interpretable Multi-Agent System</i> Zekun Li, Jifan Yu, Haoxuan Li, Ye He, Daniel Zhang-Li, Shangqing Tu, Joy Jia Yin Lim, Yikun Jiang, Jiaxin Yuan and Yu Zhang	26849

<i>LLM Reasoning as Trajectories: Step-Specific Representation Geometry and Correctness Signals</i>	
Lihao Sun, Hang Dong, Bo Qiao, Qingwei Lin, Dongmei Zhang and Saravan Rajmohan . . .	26872
<i>A Data-Efficient Path to Multilingual LLMs: Language Expansion via Post-training PARAMΔ Integration into Upcycled MoE</i>	
Hao Zhou, Tianhao Li, Zhijun Wang, Shuaijie She, Linjuan Wu, Hao-Ran Wei, Baosong Yang, Jiajun Chen and Shujian Huang	26888
<i>The Evolution of Thought: Tracking LLM Overthinking via Reasoning Dynamics Analysis</i>	
Zihao Wei, Liang Pang, Jiahao Liu, Wenjie Shi, Jingcheng Deng, Shicheng Xu, Zenghao Duan, Jingang Wang, Fei Sun, Huawei Shen and Xueqi Cheng	26905
<i>CoCoNUTS: Concentrating on Content while Neglecting Uninformative Textual Styles for AI-Generated Peer Review Detection</i>	
Yihan Chen, Jiawei Chen, Guozhao Mo, Xuanang Chen, Ben He, Xianpei Han and Le Sun	26921
<i>Can LLMs Estimate Cognitive Complexity of Reading Comprehension Items?</i>	
Seonjeong Hwang, Hyounghun Kim and Gary Lee	26951
<i>ReCoQA: A Benchmark for Tool-Augmented and Multi-Step Reasoning in Real Estate Question and Answering</i>	
Yindong Zhang, Wenmian Yang, Yiquan Zhang and Weijia Jia	26967
<i>Human or LLM as Standardized Patients? A Comparative Study in Medical Education</i>	
Bingquan Zhang, Xiaoxiao Liu, Yuchi Wang, Zhou Lei, Qianqian Xie and Benyou Wang .	26988
<i>Harmonizing the Past, Present, and Future: A Null-Space Constrained Region-Specific Method for Continual Learning in LLMs</i>	
Jinhui Chen, Shizhu He, Xingchang Yang, Huanxuan Liao, Yequan Wang, Xiangwen Liao, Wenhao Teng, Kang Liu and Jun Zhao	27013
<i>Attention Weights as an Indicator: Analyzing and Improving Document Utilization in Retrieval-Augmented Generation</i>	
Jing Jin, Yuhan Song, Wen Luo and Houfeng Wang	27034
<i>Inertia in Moral and Value Judgments of Large Language Models</i>	
Bruce W. Lee, Yeongheon Lee and Hyunsoo Cho	27053
<i>LiveCultureBench: a Multi-Agent, Multi-Cultural Benchmark for Large Language Models in Dynamic Social Simulations</i>	
Viet Thanh Pham, Lizhen Qu, Thuy-Trang Vu, Gholamreza Haffari and Dinh Phung	27076
<i>Enabling Agents to Communicate Entirely in Latent Space</i>	
Zhuoyun Du, Runze Wang, Huiyu Bai, Zouying Cao, Xiaoyong Zhu, Yu Cheng, Bo Zheng, Wei Chen and Haochao Ying	27106
<i>DR-Arena: an Automated Evaluation Framework for Deep Research Agents</i>	
Yiwen Gao, Ruochen Zhao, Yang Deng and Wenxuan Zhang	27130
<i>Question Tells You Where the Answer Is: Intention-aware Long-Context KV Cache Compression</i>	
Liang Zhao, Xiaocheng Feng, Weihong Zhong, Lei Huang, Kun Zhu, Baoxin Wang, Dayong Wu, Guoping Hu, Ting Liu and Bing Qin	27153
<i>LADR: Locality-Aware Dynamic Rescue for Efficient Text-to-Image Generation with Diffusion Large Language Models</i>	
Chenglin Wang, Yucheng Zhou, Shuang Chen, Tao Wang and Kai Zhang	27170

<i>PACE: Predictive Adaptive Context Extraction for Long-Horizon LLM Agents</i>	
Lei Wei, Xiao Peng, TT, Guannan Zhang, Chenhao Jiang, Hongyu Li, Lanbo Lin, Yuanwu Xu, Jiayao Liu, Kesu Wang and Bin Wang	27184
<i>PaCoRe: Learning to Scale Test-Time Compute with Parallel Coordinated Reasoning</i>	
Jingcheng Hu, Yinmin Zhang, Shijie Shang, Xiaobo Yang, Yue Peng, Zhewei Huang, Hebin Zhou, Xin Wu, Jie Cheng, Fanqi Wan, Xiangwen Kong, Chengyuan Yao, Kaiwen Yan, Ailin Huang, Hongyu Zhou, Qi Han, Zheng Ge, Xiangyu Zhang and Heung-Yeung Shum	27200
<i>Draft, Verify, Restore: Self-Refining Historical Inscription Restoration with a Unified MLLM</i>	
Yuyi Zhang, Junle Liu, Peirong Zhang, Jianliang Liu, Zhenhua Yang and Lianwen Jin	27216
<i>SAM3-I: Segment Anything with Instructions</i>	
Jingjing Li, Yue Feng, Yuchen Guo, Jincai Huang, Wei Ji, Qi Bi, Yongri Piao, Miao Zhang, Xiaoqi Zhao, Qiang Chen, Shihao Zou, Huchuan Lu and Li Cheng	27232
<i>Stop When Enough: Adaptive Early-Stopping for Chain-of-Thought Reasoning</i>	
Renliang Sun, Wei Cheng, Dawei Li, Haifeng Chen and Wei Wang	27250
<i>Improving Long-Context Translation via Self-Supervised Dual Learning</i>	
Shanbo Cheng, Shuaijie She, Yu Bao, Jianbing Zhang, Jiajun Chen and Shujian Huang	27269
<i>CaRL-EM: Cost-Aware Reinforcement Learning for Entity Matching with LLMs</i>	
Chaohui Guo, Michel C. A. Klein and Zhisheng Huang	27281
<i>Audio Jailbreak: An Open Comprehensive Benchmark for Jailbreaking Large Audio-Language Models</i>	
Zirui Song, Qian Jiang, Mingxuan Cui, Mingzhe Li, Lang Gao, Zeyu Zhang, Zixiang Xu, Yanbo Wang, Guangxian Ouyang, Zhenhao Chen and Xiuying Chen	27294
<i>When Personalization Legitimizes Risks: Uncovering Safety Vulnerabilities in Personalized Dialogue Agents</i>	
Jiahe Guo, Xiangran Guo, Yulin Hu, Zimo Long, Xingyu Sui, Xuda Zhi, Yongbo Huang, Hao He, Weixiang Zhao, Yanyan Zhao and Bing Qin	27309
<i>AlphaContext: An Evolutionary Tree-based Psychometric Context Generator for Creativity Assessment</i>	
Yixuan Wang, Yue Huang, Hong Qian, Yunzhao Wei, Yifei Ding, Wenkai Wang, Zhi Liu, Zhongjing Huang, Aimin Zhou and Jiajun Guo	27336
<i>Demystifying Data Organization for Enhanced LLM Training</i>	
Yalun Dai, Yangyu Huang, Tongshen Yang, Yonghan Wang, Xin Zhang, Wenshan Wu, Qihao Zhao, Hao Li, Yuanyuan Gao, Kim-Hui Yap and Scarlett Li	27358
<i>Beyond Single View: A Comprehensive Benchmark for Medical Multimodal Large Language Models on Multi-Image Understanding</i>	
Dexuan Xu, Yuan Jiayin, Jianing Wang, Yanyuan Chen, Hanpin Wang and Yu Huang	27376
<i>When Seeing Is not Enough: Revealing the Limits of Active Reasoning in MLLMs</i>	
Hongcheng Liu, Pingjie Wang, Yuhao Wang, Siqun Ou, Yanfeng Wang and Yu Wang	27395
<i>CuMA: Aligning LLMs with Sparse Cultural Values via Demographic-Aware Mixture of Adapters</i>	
Ao Sun, Xiaoyu Wang, Zhe Tan, Yu Li, Zhu Jiachen, Yuheng Jia and Shu Su	27423
<i>Neo-Classic: A Benchmark for Evaluating Linguistic-Aesthetic Reasoning in Classical Chinese Poetry</i>	
Han Zhang, Zihan Gu, Zhiyuan Wang, Tianyi Ma, Jiacheng Lu, Xinyan Zhang, Yuhao Wei and Cheng Hua	27442

<i>A Multi-Agent Framework for Feature-Constrained Difficulty Control in Reading Comprehension Item Generation</i>	Seonjeong Hwang, Jun Seo, Hyoungun Kim and Gary Lee	27466
<i>ToMMer - Efficient Entity Mention Detection from Large Language Models</i>	Victor Morand, Nadi Tomeh, Josiane Mothe and Benjamin Piwowarski	27489
<i>Shuttle Between Symbolic Instructions and Neural Parameters of Large Language Models</i>	Wangtao Sun, Haotian Xu, Huanxuan Liao, Xuanqing Yu, Zhongtao Jiang, Shizhu He, Jun Zhao and Kang Liu	27510
<i>Mitigating Spurious Correlations in Text Classification Using Latent Space Geometry</i>	Jiasen Gao, Xiaoliang Chen, Duoqian Miao, Xu Gu, Xianyong Li and Yajun Du	27525
<i>Efficient Hyperparameter Optimization for LLM Reinforcement Learning</i>	Minping Chen, Bowen Xiao, Du Liang, Chuxuan Zeng and Zeyi Wen	27540
<i>Hetero-Designer: Automated Design of Multi-Agent Systems with Heterogeneous LLMs</i>	Zhiheng Zhang, Yuanzhe Zhang, Bohan Yu, Daojian Zeng, Kang Liu and Jun Zhao	27553
<i>Achieving Multi-Hop Calculation and Safe Abstention in Financial Numerical Reasoning by Metric Graph Constrained LLMs</i>	Aoyuan Jiang, Liang Hong, Haoxuan Liu and Rui Wang	27575
<i>SOCIA-EVO: Automated Simulator Construction via Dual-Anchored Bi-Level Optimization</i>	Yuncheng Hua, Sion Weatherhead, Mehdi Jafari, Hao Xue and Flora D. Salim	27596
<i>SoundBreak: A Systematic Study of Audio-Only Adversarial Attacks on Trimodal Models</i>	Aafiya Shamshad Hussain, Gaurav Srivastava, Alvi Md Ishmam, Zaber Ibn Abdul Hakim and Chris Thomas	27635
<i>Trust Within? Seek Beyond? Knowledge Boundary Aware Policy Optimization for Agentic Search</i>	Tao Feng, Xinke Jiang, Xinyan Hu, Yonggang Zhang, Zhen Tao, Wentao Zhang, Boyang Liu, Wenhao Jiang and Chao Wu	27664
<i>RPC-Bench: A Fine-grained Benchmark for Research Paper Comprehension</i>	Yelin Chen, Fanjin Zhang, Suping Sun, Yunhe Pang, Yuanchun Wang, Jian Song, XiaoYan Li, Lei Hou, Shu Zhao, Jie Tang and Juanzi Li	27683
<i>FusionFlow: Enabling Deep Structural Exploration for Automated Agentic Workflow Generation</i>	Xiang Wang, Zongtao Yang, Zhuojian Hong, Shuhao Zhang and Wei Wei	27718
<i>CompTab: A Comprehensive Benchmark for Real-World TableQA with Complex Reasoning and Irregular Tables</i>	Zhen Yang, Wei Du, Jie Wang, Wenze Zhou, Xiangfeng Meng, Zhengyang Wang, Suping Sun, Ziwei Du, Haodong Zou, Jie Chen, Yongbin Liu, Shicheng Tan, Jiahao Ying and Shu Zhao	27761
<i>Learning More from Less: Exploiting Counterfactuals for Data-Efficient Chart Understanding</i>	Jianzhu Bao, Haozhen Zhang, Kuicai Dong, Bozhi Wu, Sarthak Ketanbhai Modi, Zi Pong Lim, Yon Shin Teo and Wenya Wang	27775
<i>Video-MMMU: Evaluating Knowledge Acquisition from Multidisciplinary Professional Videos</i>	Kairui Hu, Penghao Wu, Fanyi Pu, Wang Xiao, Xiang Yue, Bo Li, Yuanhan Zhang and Ziwei Liu	27798
<i>Semantically Comprehensive Token Pruning in LVLMS via Maximizing Concept Coverage</i>	Xueting Li, Qi Liu, Chenghao Xu, Xu Yang, Guangtao Lyu, Jiahua Li and Cheng Deng	27829

<i>OptiCo: Adaptive Distributed Training Optimization via Collaborative Agent Reasoning</i>	
Sheng Chen, Tang Zhe, Weixing Zhang, Fei Yang, Yuanyuan. Wang, Tianlin Li and Yang Liu	
	27847
<i>MemBuilder: Reinforcing LLMs for Long-Term Memory Construction via Attributed Dense Rewards</i>	
Zhiyu Shen, Ziming Wu, Fuming Lai, Shaobing Lian and Yanghui Rao	27868
<i>Towards Fine-grained Audio Captioning with Multimodal Contextual Fusion</i>	
Shunian Chen, Xinyuan Xie, Zheshu Chen, Owen Lee, Liyan Zhao, Zhan Su, Qilin Sun and Benyou Wang	27888
<i>Probing Semantic Alignment, Lexical Invariance, and Syntactic Influence in LLM Metaphor Processing</i>	
Fengying Ye, Shanshan Wang, Lidia S. Chao and Derek F. Wong	27914
<i>MathCanvas: Intrinsic Visual Chain-of-Thought for Multimodal Mathematical Reasoning</i>	
Weikang Shi, Aldrich Yu, Rongyao Fang, Houxing Ren, Ke Wang, Aojun Zhou, Changyao Tian, Xinyu Fu, Yuxuan Hu, Zimu Lu, Linjiang Huang, Si Liu, Rui Liu and Hongsheng Li	27933
<i>Rethinking the Role of Entropy in Optimizing Tool-Use Behaviors for Large Language Model Agents</i>	
Zeping Li, Hongru Wang, Yiwen Zhao, Guanhua Chen, Yixia Li, Keyang Chen, Yixin Cao, Guangnan Ye, Hongfeng Chai and Zhenfei Yin	27955
<i>Collaboration of Fusion and Independence: Hypercomplex-driven Robust Multi-Modal Knowledge Graph Completion</i>	
Zhiqiang Liu, Yichi Zhang, Mengshu Sun, Lei Liang and Wen Zhang	27968
<i>LLM-MC-Affect: LLM-Based Monte Carlo Modeling of Affective Trajectories and Latent Ambiguity for Interpersonal Dynamic Insight</i>	
Yu-Zheng Lin, Bono Po-Jen Shih, John Paul Martin Encinas, Elizabeth Victoria Achom, Karan Patel, Jesus Horacio Pacheco, Sicong Shao, Jyotikrishna Dass, Soheil Salehi and Pratik Satam	27984
<i>Multilingual and Cross-Lingual Citation Needed Detection on Wikipedia for Lower-Resource Languages</i>	
Gerrit Quaremba, Amy Rechkemmer, Elizabeth Black, Denny Vrandečić and Elena Simperl	28001
<i>Process Reward Models Meet Planning: Generating Precise and Scalable Datasets for Step-Level Rewards</i>	
Raffaele Pisano and Roberto Navigli	28022
<i>UniSonate: A Unified Model for Speech, Music, and Sound Effect Generation with Text Instructions</i>	
Chunyu Qiang, Xiaopeng Wang, Kang Yin, Yuzhe Liang, Yuxin Guo, Teng Ma, Ziyu Zhang, Tianrui Wang, Cheng Gong, Yushen Chen, Ruibo Fu, Longbiao Wang and Jianwu Dang	28043
<i>AgentAsk: Multi-Agent Systems Need to Ask</i>	
Bohan Lin, Kuo Yang, Zelin Tan, Yingchuan Lai, Chen Zhang, Guibin Zhang, Xinlei Yu, Miao Yu, Xu Wang, Yudong Zhang and Yang Wang	28055
<i>HopWeaver: Cross-Document Synthesis of High-Quality and Authentic Multi-Hop Questions</i>	
Zhiyu Shen, Jiyuan Liu, Yunhe Pang, Yanghui Rao, Fu Lee Wang and Jianxing Yu	28078
<i>REST: Stress Testing Large Reasoning Models by Asking Multiple Problems at Once</i>	
Zhuoshi Pan, Qizhi Pei, Yu Li, Zinan Tang, QiYao Sun, H. Vicky Zhao, Conghui He and Lijun Wu	28110
<i>You Can Have a Second Chance: Unbiased and Multi-bit Watermarking for Diffusion Language Models with Regret-based Remasking</i>	
Ke Yang, Dongyang Liang, Jing Yu, Shuguang Yuan and Chi Chen	28141

<i>VCORE: Variance-Controlled Optimization-based Reweighting for Chain-of-Thought Supervision</i> Xuan Gong, Senmiao Wang, Hanbo Huang, Ruoyu Sun and Shiyu Liang	28161
<i>MAGNET: Towards Adaptive GUI Agents with Memory-Driven Knowledge Evolution</i> Libo Sun, Jiwen Zhang, Siyuan Wang and Zhongyu Wei	28181
<i>VLN-MME: Diagnosing MLLMs as Language-guided Visual Navigation Agents</i> Xunyi Zhao, Gengze Zhou and Qi Wu	28207
<i>ImplicitMemBench: Measuring Unconscious Behavioral Adaptation in Large Language Models</i> Chonghan Qin, Xiachong Feng, Weitao Ma, Xiaocheng Feng and Lingpeng Kong	28232
<i>Semantic Hardness Is Not Visual Hardness: Sign-Aware Hard Negative Mining for Sign Language Retrieval</i> Junmyeong Lee, Chan Hur, ChangSu Choi, Sukmin Cho, Fitsum Gaim, Eui Jun Hwang, Hoyun Song and KyungTae Lim	28262
<i>Difference in Task Performance on Sparse Speech Representations</i> Wenjie Peng, Chen Chen and Thomas Hain	28278
<i>VerilogLAVD: LLM-Aided Pattern Generation for Verilog CWE Detection</i> Xiang Long, Yingjie Xia, Li Kuang, Yao Wan and ZiHao Liu	28292
<i>Temporal Sampling for Forgotten Reasoning in LLMs</i> Yuetai Li, Zhangchen Xu, Fengqing Jiang, Bhaskar Ramasubramanian, Luyao Niu, Bill Yuchen Lin, Xiang Yue and Radha Poovendran	28309
<i>Conflict-Aware Memory for Embodied Agents: Enhancing Vector Data Quality via Detection Rules</i> Kexin Ma, Haotian Wang, Shenglin Chen, Yishuai Cai, Huangyuyu and Ruochun Jin	28328
<i>Benchmarking Deflection and Hallucination in Large Vision-Language Models</i> Nicholas Moratelli, Christopher Davis, Leonardo F. R. Ribeiro, Bill Byrne and Gonzalo Iglesias	28348
<i>STAPO: Selective Trajectory-Aware Policy Optimization for LLM Agent Training</i> Qiuyi Qi, Tian Liang, Mutian Bao, Jinjian Zhang, Dongnan Liu, Wei Zhou, Linjian Mo, Ming Kong, Jie Liu, Feng Zhang and Qiang Zhu	28371
<i>Read As Human: Compressing Context via Parallelizable Close Reading and Skimming</i> Jiwei Tang, Shilei Liu, Zhicheng Zhang, Qingsong Lv, Runsong Zhao, Tingwei Lu, Langming Liu, Haibin Chen, Yujin Yuan, Hai-Tao Zheng, Wenbo Su and Bo Zheng	28393
<i>OASIS: Mitigating Harmful Fine-tuning Attacks on LLMs via Orthogonal and Adaptive Safety Alignment Strategy</i> Jiayu Tang, Guowei Peng, Qiu hao Xie, Yuning Yang, Xiurui Xie and Guisong Liu	28407
<i>KoCo-Bench: Can Large Language Models Leverage Domain Knowledge in Software Development?</i> Xue Jiang, Ge Li, Jiaru Qian, Xianjie Shi, Chenjie Li, Hao Zhu, Ziyu Wang, Jielun Zhang, Zeyu Zhao, Kechi Zhang, Jia Li, Wenpin Jiao, Zhi Jin and Yihong Dong	28422
<i>Enhancing Two Steps Textual Anomaly Detection through Anisotropy Mitigation</i> Pierre Fihey, Matthieu Labeau and Pavlo Mozharovskiy	28442
<i>InfiniteWeb: Scalable Web Environment Synthesis for GUI Agent Training</i> Ziyun Zhang, Zezhou Wang, Xiaoyi Zhang, Zongyu Guo, Jiahao Li, Bin Li and Yan Lu	28465
<i>MagicBench: Diagnosing Visual Agency Loss and Semantic Dependency in Multimodal LLMs</i> Tang Da Huang, Weidong Tang, Wen Qi Xu and Xianpeng Guo	28493

<i>Bloom-Eval: A Hierarchical Evaluation Benchmark for Automatic Survey Generation Based on Bloom’s Taxonomy</i>	
Fei Zhang, Zhe Zhao, HaiBin Wen, Tianshuo Wei, Zaixi Zhang, Chao Yang and Ye Wei . .	28512
<i>Cultivating Forensic Reasoning for Generalizable Multimodal Manipulation Detection</i>	
Yuchen Zhang, Yaxiong Wang, Kecheng Han, Yujiao Wu, Lianwei Wu, Li Zhu and Zhedong Zheng	28545
<i>Fair-CCD: Mitigating Bias in Large Language Models for Tabular Classification Through Context-Contrastive Decoding</i>	
Donghan Liu, Han Sun, Zhaohui Wang, Qin Li and Min Zhang	28561
<i>DVI-DTM: Dual-View Representation Learning for Interpretable Short Text Dynamic Topic Modeling</i>	
Di Liu, Zheng Fang and Bin Wu	28576
<i>GenesisFunc: Multi-Agent Data Generation for Accurate and Generalizable Function-Calling</i>	
Hao-Xiang Xu, Chong Deng, Jiaqing Liu, Wen Wang, Qian Chen, Lujia Bao, Xiangang Li and Zhen-Hua Ling	28594
<i>Selective Test-Time Debiasing for CLIP via Reward Gating</i>	
Jaeho Han, Jisoo Yang, Hyeondong Woo, Mingyu Jeon, Sunjae Yoon and Junyeong Kim .	28617
<i>CURE: Critique-Driven Unified Reinforcement Learning for Test-Time Self-Improvement</i>	
Guirong Chen, Shuqi Ye, Wenkai Yang, Shiqi Shen, Guangyao Shen and Yankai Lin	28632
<i>LGSA: Label Geometry Structuring and Aligning for Hierarchical Text Classification</i>	
Shuai Zhang, Weibo Xu, Jiahao Nie and Kecheng Huang	28654
<i>FactVerse: A Benchmark for Factual Consistency in Interleaved Image–Text Generation</i>	
Yubo Shan, Kun Zhang, Qiming Xu, Liping Cao, Yingying Cao, Jian Zhang, Yu Wang, Jingyuan Li and Yuanzhuo Wang	28666
<i>GMSA: Enhancing Context Compression via Group Merging and Layer Semantic Alignment</i>	
Jiwei Tang, Zhicheng Zhang, Shunlong Wu, Jingheng Ye, Lichen Bai, Zitai Wang, Tingwei Lu, Lin Hai, Yiming Zhao, Hai-Tao Zheng and Hong-Gee Kim	28690
<i>SpidR-Adapt: A Universal Speech Representation Model for Few-Shot Adaptation</i>	
Mahi Luthra, Jiayi Shen, Maxime Poli, Angelo Ortiz Tandazo, Yosuke Higuchi, Youssef Bencheikroun, Martin Gleize, Charles-Éric Saint-James, Dongyan Lin, Phillip Rust, Angel Villar-Corrales, Surya, Vanessa Stark, Rashel Moritz, Juan Pino, Yann LeCun and Emmanuel Dupoux	28705
<i>EIFFEL: a novel benchmark to measure bias of English heavy training on French idiomatic expressions</i>	
Charlotte Noel, Nicholas Asher, Olivier Gouvert, Farah Benamara and Julie Hunter	28729
<i>Latent Attention Denoising: A Training-Free Energy-Based Framework for Mitigating Hallucinations in Vision-Language Models</i>	
Zhiwen Luo, Siyu Jiang, Weilong Jiang and Kun He	28748
<i>From Knowing to Teaching: Scaffolding Pedagogical Decisions for LLM Agent</i>	
Yucheng Wang, Shen Yang, Jifan Yu, Haoxuan Li, Joy Jia Yin Lim, Daniel Zhang-Li, Huiqin Liu, Lei Hou, Juanzi Li and Bin Xu	28778
<i>Learn Like Humans: Use Meta-cognitive Reflection for Efficient Self-Improvement</i>	
Xinmeng Hou, Bohao Qu, Wuqi Wang, Peiliang Gong, Qing Guo and Yang Liu	28802
<i>Think Faster Than Words: Efficient LLM Chain-of-Thought Reasoning via Dynamic Shortcut Decoding</i>	
Fan Liu, Yanhao Wang, Min Zhang, Zhikang Chen, Zeyuan Li, Lewei He and Jiahui Pan . .	28825

<i>AEA: Adaptive Expert Allocation Improves Sentence Embeddings from Mixture-of-Experts LLM</i>	
Shufan Yang, Zifeng Cheng, Zhiwei Jiang, Qingfeng Qi, Yafeng Yin, Cong Wang, Ao Zhou and Qing Gu	28837
<i>It's High Time: A Survey of Temporal Question Answering</i>	
Bhawna Piryani, Abdelrahman Abdallah, Jamshid Mozafari, Avishek Anand and Adam Jatowt	28852
<i>SDAR-VL: Stable and Efficient Block-wise Diffusion for Vision-Language Understanding</i>	
Shuang Cheng, Yuhua Jiang, Zineng Zhou, Dawei Liu, Tao Wang, Linfeng Zhang, Biqing Qi and Bowen Zhou	28882
<i>N-GLARE: An Non-Generative Latent Representation-Efficient LLM Safety Evaluator</i>	
Zheyu Lin, Jirui Yang, Yukui Qiu, Yubing Bao, Hengqi Guo and Yao Guan	28902
<i>Don't Act Blindly: Robust GUI Automation via Action-Effect Verification and Self-Correction</i>	
Yuzhe Zhang, Xianwei Xue, Xingyong Wu, Mengke Chen, Chen Liu, Xinran He, Run Shao, Feiran Liu, Huanmin Xu, Qitong Pan and Haiwei Wang	28924
<i>Beyond Static Persona Consistency: Dynamic Persona Coherence in LLM Role-Playing</i>	
Yirui QI, Xiaoming Zhang, Ruilin Zeng, Mengyao Liu, Ziyi Zhou, Dezhuang Miao, Bingyu Yan and Zhenyu Guan	28942
<i>What If Consensus Lies? Selective-Complementary Reinforcement Learning at Test Time</i>	
Dong Yan, Jian Liang, Yanbo Wang, Shuo Lu, Ran He and Tieniu Tan	28957
<i>Exploring and Distilling Multi-Dimensional Clues for Interpretable Social Bot Detection</i>	
Yi Han, Haiqi Lu, Lizi Liao, Shuhan Zhou, Yuanxing Liu, Weinan Zhang and Ting Liu ...	28971
<i>Why Are We Moral? An LLM-based Agent Simulation Approach to the Study of Moral Evolution</i>	
Zhou Ziheng, Huacong Tang, Mingjie Bi, Wanying He, Fang Sun, Yizhou Sun, Ying Nian Wu, Demetri Terzopoulos, Yipeng Kang and Fangwei Zhong	28993
<i>Systematicity between Forms and Meanings across Languages Supports Efficient Communication</i>	
Doreen Osmelak, Yang Xu, Michael Hahn and Kate McCurdy	29047
<i>On the (In-)Security of the Shuffling Defense in the Transformer Secure Inference</i>	
Zhengyi Li, Yakai Wang, Jingwen Leng, Kang Yang, Yu Yu, Jiaping Gui, Yu Feng, Ning Liu and Minyi Guo	29084
<i>LaCo: Layer-wise Compensation for Pruned Large Language Models</i>	
Yingen Liu, Fan Wu, Panxuyan, Ruihui Li, Zhuo Tang and Kenli Li	29099
<i>QiMeng-PRepair: Precise Code Repair via Edit-Aware Reward Optimization</i>	
Changxin Ke, Rui Zhang, Jiaming Guo, Yuanbo Wen, Li Ding, Shuo Wang, Xuyuan Zhu, Xiong Peng, Di Huang, Zidong Du, Xing Hu, Qi Guo and Yunji Chen	29114
<i>Location Not Found: Exposing Implicit Local and Global Biases in Multilingual LLMs</i>	
Guy Mor-Lan, Omer Goldman, Matan Eyal, Adi Mayrav Gilady, Sivan Eiger, Idan Szpektor, Avinatan Hassidim, Yossi Matias and Reut Tsarfaty	29130
<i>Incorporating Temporal Coherence to Cross-Document Event Coreference Resolution</i>	
Xinyu Chen, Peifeng Li and Qiaoming Zhu	29151
<i>LLMs in Sarcasm Detection? It's elementary! (Or is it?)</i>	
Priyanshu Mahato, Aniket Santosh Mishra and Kripabandhu Ghosh	29169

<i>Beyond Itinerary Planning—A Real-World Benchmark for Multi-Turn and Tool-Using Travel Tasks</i>	
Xiang Cheng, Yulan Hu, Xiangwen Zhang, Lu Xu, Lide Tan, Zheng Pan, Xin Li and Yong Liu	
	29200
<i>TLoRA: Task-aware Low Rank Adaptation of Large Language Models</i>	
Weicheng Lin, Yi Zhang, Jiawei Dang and Liang-Jie Zhang	29252
<i>MARS-RA: Rank Aggregation for Credit Assignment via Multimodal Comparisons in Embodied Multi-Agent Cooperation</i>	
Dawei Wang, Di Zhao, Xinyuan Liu, Marci Chi Ma, Xiaoyang Liu, Chengming Zhou, Gary Ushaw and Richard Davison	29270
<i>Persona-E²: A Human-Grounded Dataset for Personality-Shaped Emotional Responses to Textual Events</i>	
Yuqin Yang, Haowu Zhou, Haoran Tu, Zhiwen Hui, Shiqi Yan, HaoYang Li, Dong She, Xianrong Yao, Yang Gao and Zhanpeng Jin	29290
<i>Interleaved Latent Visual Reasoning with Selective Perceptual Modeling</i>	
Shuai Dong, Siyuan Wang, Xingyu Liu, Chenglin Li, Haowen Hou and Zhongyu Wei	29316
<i>Think before Go: Hierarchical Reasoning for Image-goal Navigation</i>	
Pengna Li, Kangyi Wu, Shaoqing Xu, Fang Li, Lin Zhao, Long Chen, Zhi-Xin Yang and Nanning Zheng	29336
<i>LLMs as Knowledge Graph Refiners: Mitigating Factual Inconsistencies in Generative Knowledge Extraction</i>	
Donghyun Kim, Hyeongjun Yang, Seokju Hwang, Kyong-Ho Lee and Chanhee Lee	29358
<i>SILO-BENCH: A Scalable Environment for Evaluating Distributed Coordination in Multi-Agent LLM Systems</i>	
Yuzhe Zhang, Feiran Liu, Yi Shan, Xinyi Huang, Xin Yang, Yueqi Zhu, Xuxin Cheng, Cao Liu, Ke Zeng, Terry Jingchen Zhang and Wenyuan Jiang	29379
<i>Mitigating Structural Knowledge Collapse in Domain-Specific LLMs via Morpheme-Aware KV-Aggregation</i>	
Yuxuan Si, Zheqi Lv, Chengxi Zang, Zhengyu Chen and Fei Wu	29399
<i>Zero-shot Jianzi Recognition as Structured Visual Information Extraction in Open Compositional Symbolic Systems</i>	
Zehan Li, Fu Zhang, Zhijun Liu and Jingwei Cheng	29414
<i>Whose Facts Win? LLM Source Preferences under Knowledge Conflicts</i>	
Jakob Schuster, Vagrant Gautam and Katja Markert	29430
<i>CLARity: Reasoning Consistency Alone Can Teach Reinforced Experts</i>	
Jiuheng Lin, Cong Jiang, Zirui Wu, Jiarui Sun and Yansong Feng	29460
<i>EGSS: Entropy-guided Stepwise Scaling for Reliable Software Engineering</i>	
Chenhui Mao, Yuanting Lei, Zhixiang Wei, Ming Liang, Zhixiang Wang, Jingxuan Xu, Dajun Chen, Wei Jiang and Yong Li	29481
<i>Detecting What Queries Seek: Steering LLM Safety with FFN Output Activation Monitoring</i>	
Xiaohao Luo, Ying Wei and Rui Zhao	29500
<i>Temporal Evidence Chain for Temporal Knowledge Graph Question Answering with Large Language Models</i>	
Shihao Liu, Xiaofei Zhou, Bo Wang and Geyuan Zhang	29515

<i>Modal Dependency Parsing as Structured Prediction over Source-Cue Scope</i>	
Jayeol Chun and Nianwen Xue	29530
<i>DEFT: Demystifying VLN Failures via a Unified Dual-View Explainability Framework for LLM-based Agents</i>	
Yawen Wang, Yihan Dai, Jianming Chen, Junjie Wang and Qing Wang	29541
<i>GLIER: Generative Legal Inference and Evidence Ranking for Legal Case Retrieval</i>	
Minghan Li, Tianrui Lv, Chao Zhang and Guodong Zhou	29561
<i>LLMs (Almost) Never Abstain Under Medical Uncertainty</i>	
Alessio Cocchieri, Luca Ragazzi, Giuseppe Tagliavini and Gianluca Moro	29573
<i>TOWER+: Bridging Generality and Translation Specialization in Multilingual LLMs</i>	
Ricardo Rei, Nuno M Guerreiro, José Pombal, João Alves, Amin Farajian, Pedro Teixeira and Andre Martins	29614
<i>Culture-Aware Machine Translation in Large Language Models: Benchmarking and Investigation</i>	
Zekun Yuan, Yangfan Ye, Xiaocheng Feng, Baohang Li, Qichen Hong, Yunfei Lu, Dandan Tu and Bing Qin	29636
<i>Learning from Near-Misses: Error-Aware Contrastive Few-Shot Learning for NL2Formula</i>	
Zhihao Shuai, Yiyun Chen, Maolin Ma, Yutong Chen, Hanjia Qiu, Jing Xu, Ziyue Chen and Weikai Yang	29662
<i>Calibrated Speculative Decoding: Frequency-Guided Candidate Selection for Efficient Inference</i>	
Zhouxuwen, Fangxin Liu, Chao Wang, Xiao Zheng, Hao Zheng, Min He, Li Jiang and Haibing Guan	29677
<i>Learning from Mistakes: Negative Reasoning Samples Enhance Out-of-Domain Generalization</i>	
Tian Xueyun, Minghua Ma, Bingbing Xu, Nuoyan Lyu, Wei Li, Heng Dong, Zheng Chu, Yuanzhuo Wang and Huawei Shen	29689
<i>LongTutor: Benchmarking Large Language Models for Long-term Personalized Tutoring</i>	
Ning Li, Zheng Zhang, Zhenya Huang, Rui Li, Yi Zhan, Yinbo Luo, Qi Liu and Enhong Chen	29712
<i>HiChunk: Evaluating and Enhancing Retrieval Augmented Generation with Hierarchical Chunking</i>	
Wensheng Lu, Keyu Chen, Zhifeng Shen, Ruizhi Qiao and Xing Sun	29738
<i>How Good is Your Wikipedia? Auditing Data Quality for Low-resource and Multilingual NLP</i>	
Kushal Tatariya, Artur Kulmizev, Wessel Poelman, Esther Ploeger, Marcel Bollmann, Johannes Bjerva, Jiaming Luo, Heather Lent and Miryam de Lhoneux	29754
<i>Understanding and Mitigating Political Stance Cross-topic Generalization in Large Language Models</i>	
Jiayi Zhang, Shu Yang, Junchao Wu, Derek F. Wong and Di Wang	29775
<i>MPR-GUI: Benchmarking and Enhancing Multilingual Perception and Reasoning in GUI Agents</i>	
Ruihan Chen, Qiming Li, Xiaocheng Feng, Weihong Zhong, Xiaoliang Yang, Yuxuan Gu, Zekun Zhou, Yunfei Lu, Haoyu Ren, Kun Chen, Dandan Tu and Bing Qin	29798
<i>CoVerRL: Breaking the Consensus Trap in Label-Free Reasoning via Generator-Verifier Co-Evolution</i>	
Teng Pan, Yuchen Yan, Zixuan Wang, Ruiqing Zhang, Guiyang Hou, Wenqi Zhang, Weiming Lu, Jun Xiao and Yongliang Shen	29833
<i>From Weights to Activations: Is Steering the Next Frontier of Adaptation?</i>	
Simon Ostermann, Daniil Gurgurov, Tanja Baeumel, Michael A. Hedderich, Sebastian Lapuschkin, Wojciech Samek and Vera Schmitt	29854

Can LLMs Act as Historians? Evaluating Historical Research Capabilities of LLMs via the Chinese Imperial Examination

Lirong Gao, Zeqing Wang, Yuyan Cai, Jiayi Deng, Yanmei Gu, Yiming Zhang, Jia Zhou, Yanfei Zhang and Junbo Zhao 29880

Less Languages, Less Tokens: An Efficient Unified Logic Cross-lingual Chain-of-Thought Reasoning Framework

Chenyuan Zhang, Qiguang Chen, Xie Chen, Zhuotao Tian, Bowen Xing, Meishan Zhang, Libo Qin, Baotian Hu and Min Zhang 29912

Embracing Anisotropy: Turning Massive Activations into Interpretable Control Knobs for Large Language Models

Youngji Roh, Hyunjin Cho and Jaehyung Kim 29930

From Language to Driving: A Dual-Loop SLM-Enhanced Framework for Multi-Planner Scheduling via a Domain-Specific Language

Jiawei Liu, Xun Gong, Muli Yang, Xingrui Yu, Fen Fang, Xulei Yang, Ivor Tsang, Yunfeng hu, Hong Chen and Qing Guo 29957

One Cognitive Loop Is Enough: SODA unlocks Pure-Text Spatial Reasoning in Large Language Models

Shunwen Bai, Jiahuan Zhang, Haoran Huang, Yurun Wang, Jiale Liu, Yanxi Wu, Ningzhe Yu, Yudong Gao and Mingjun Cheng 29974

Compatibility-Aware Dynamic Fine-Tuning for Large Language Models

Yucheng Zhou, Junwei Sheng, Qianning Wang and Jianbing Shen 29997

Spectral Characterization and Mitigation of Sequential Knowledge Editing Collapse

Chi Zhang, Mengqi Zhang, Xiaotian Ye, Runxi Cheng, Zisheng Zhou, Ying Zhou, Pengjie Ren and Zhumin Chen 30009

SSA: Improving Performance With a Better Scoring Function

Omar Naim, Swarnadeep Bhar, Jerome Bolte and Nicholas Asher 30033

Think Better, Not Longer: Token-Level Marginal Utility for Efficient Reasoning in Large Reasoning Models

Jiawei Li, Yang Gao, Huashan Sun and Chong Feng 30052

AgentSlimming: Towards Efficient and Cost-Aware Multi-Agent Systems

Yulang Chen, Haoxuan Peng, Jinyan Liu, Zichen Wen, Dongrui Liu and Linfeng Zhang .. 30064

Reusable Experiences: Latent Routing and Modular Composition in LLMs

Shuai Ling, Lizi Liao, Dongmei Jiang and Weili Guan 30087

Communicating in Emergent Language with an Induced Morphological Phrasebook

Brendon Boldt and David R. Mortensen 30101

Controllable Contamination Detection for Reliable LLM Evaluation with Statistical Guarantees

Zheng Zhang, Qi Liu, Siyuan Liang, Ning Li, Zirui Hu, Weibo Gao, Rui Li, Zhenya Huang, Leszek Rutkowski, Baosheng Yu and Dacheng Tao 30122

From 1,000,000 Users to Every User: Scaling Up Personalized Preference for User-level Alignment

Jia-Nan Li, Jian Guan, Songhao Wu, Wei Wu and Rui Yan 30144

Beyond Atomic Characters: Glyph-Aware Sub-character Alignment for Low-Resource Multilingual OCR

Mengxiao Zhu, Haixu Chen, Jiu Sha, Jie Liu and Ge Shi 30169

<i>Why Supervised Fine-Tuning Fails to Learn: A Systematic Study of Incomplete Learning in Large Language Models</i>	
Chao Xue, Yao Wang, Mengqiao Liu, Di Liang, Xingsheng Han, Peiyang Liu, Xianjie Wu, Chenyao Lu, Lei Jiang, Yu Lu, Haibo Shi, Shuang Liang, Minlong Peng and Flora D. Salim . . .	30186
<i>KnowMe-Bench: Benchmarking Person Understanding for Lifelong Digital Companions</i>	
Tingyu Wu, Zhisheng Chen, Ziyang Weng, Shuhe Wang, Shuo Zhang, Sen Hu, Silin Wu, Qizhen Lan, Huacan Wang and Ronghao Chen	30214
<i>PersonalityDBench: A Dataset for Personality Disorders - from Modeling to Controlled Generation</i>	
Federico Ravenda, Seyed Ali Bahrainian, Daniele Montagnani, Antonietta Mira and Andrea Rabballo	30239
<i>Towards Interpretable Tabular Reasoning: Enhancing LLM Reasoning on Tabular Data with Pre-Constructed Logic Graph</i>	
Lirong Gao, Zewei Yu, Zhongrui Yin, Qi Zhang, Yuke Zhu, Bo Zheng, Haobo Wang, Junbo Zhao, Gang Chen and Sheng Guo	30260
<i>Beyond Explicit Refusals: Soft-Failure Attacks on Retrieval-Augmented Generation</i>	
Wentao Zhang, Yan Zhuang, ZhuHang Zheng, Mingfei Zhang, Jiawen Deng and Fuji Ren	30281
<i>Compressing LLM Knowledge into Graph Representations for Text-attributed Graphs Learning</i>	
Runhuai Chen, Dian Shen, Dandan Zhang, Kaihong Huang, Linghui Meng and Beilun Wang	30303
<i>End-to-End Optimization of LLM-Driven Multi-Agent Search Systems via Heterogeneous-Group-Based Reinforcement Learning</i>	
Guanzhong Chen, Shaoxiong Yang, Chao Li, Wei Liu, Jian Luan and Zenglin Xu	30319
<i>Translation with Thought: Difficulty-Adaptive Reasoning via Reinforcement Learning for Multi-Domain Machine Translation</i>	
Yongshi Ye, Biao Fu, Chongxuan Huang, Yidong Chen and Xiaodong Shi	30339
<i>Segment, Embed, and Align: A Universal Recipe for Aligning Subtitles to Signing</i>	
Zifan Jiang, Youngjoon Jang, Liliane Momeni, Gül Varol, Sarah Ebling and Andrew Zisserman	30371
<i>Understanding Emergent Misalignment via Feature Superposition Geometry</i>	
Gouki Minegishi, Hiroki Furuta, Takeshi Kojima, Yusuke Iwasawa and Yutaka Matsuo . . .	30385
<i>MavenCoder: Competitive Code Generation via Model Adaptive Planning Strategies and Multi-Perspective Verification Enhancement</i>	
ZhenChun Xu, Yi Cai, Dajun Zheng, li Yuan, Mengchen Zhao, Qixiang Wang and Jiexin Wang	30415
<i>Evidence-Augmented Policy Optimization with Reward Co-Evolution for Long-Context Reasoning</i>	
Xin Guan, Zijian Li, Shen Huang, Pengjun Xie, Jingren Zhou and Jiuxin Cao	30440
<i>Generating Attribution Reports for Manipulated Facial Images: A Dataset and Baseline</i>	
Jingchun Lian, Lingyu Liu, Yaxiong Wang, Yujiao Wu, Lianwei Wu, Li Zhu and Zhedong Zheng	30455
<i>Optimizing RAG Rerankers with LLM Feedback via Reinforcement Learning</i>	
Yuhang Wu, Xiangqing Shen, Fanfan Wang, Cangqi Zhou, Zhen Wu, Xinyu Dai and Rui Xia	30474

<i>Explainable and Fine-Grained Safeguarding of LLM Multi-Agent Systems via Bi-Level Graph Anomaly Detection</i>	
Junjun Pan, Yixin Liu, Rui Miao, Kaize Ding, Yu Zheng, Quoc Viet Hung Nguyen, Alan Wee-Chung Liew and Shirui Pan	30491
<i>Understanding and Mitigating Bias Inheritance in LLM-based Data Augmentation on Downstream Tasks</i>	
Miaomiao Li, Hao Chen, Yang Wang, Tingyuan Zhu, Weijia Zhang, Kaijie Zhu, Kam-Fai Wong and Jindong Wang	30507
<i>Beyond the Panorama: Training-Free Hierarchical Perception-Reasoning for Fine-Grained Vision in MLLMs</i>	
Xiaoyang Yi, Jing Chen, Li Peng, Yuru Bao and Jian Zhang	30534
<i>To Judge or Not to Judge: Can Large Language Models Leverage the Dispute Focus in Legal Judgment?</i>	
Luoming Hu, Liang Yang, Jingjie Zeng and Zijie Xing	30550
<i>CPR-RAG: Clinical Prior-Regularized Retrieval for Anatomy-Aware 3D CT Report Generation</i>	
Sungkyu Yang, Kang-Min Kim and Mansu Kim	30572
<i>Evaluating Answer Leakage Robustness of LLM Tutors against Adversarial Student Attacks</i>	
Jin Zhao, Marta Knežević and Tanja Käser	30588
<i>Expect the Unexpected? Testing the Surprisal of Salient Entities</i>	
Jessica Lin and Amir Zeldes	30618
<i>ArgGenBench: Benchmarking the Complex Controlled Argument Generation Capability of Large Language Models</i>	
Bojun Jin, Jianzhu Bao, Yang Sun, Yice Zhang and Ruifeng Xu	30630
<i>Language Models Learn Universal Representations of Numbers and Here’s Why You Should Care</i>	
Michal Štefánik, Timothee Mickus, Marek Kadlčík, Bertram Højer, Michal Spiegel, Raúl Vázquez, Aman Sinha, Josef Kuchař, Philipp Mondorf and Pontus Stenetorp	30663
<i>CURE-Med: Curriculum-Informed Reinforcement Learning for Multilingual Medical Reasoning</i>	
Eric Onyame, Akash Ghosh, Subhadip Baidya, Sriparna Saha, Xiuying Chen and Chirag Agarwal	30682
<i>Inverting the Shield: Systematically Generating Safety Tests from Policy Specifications</i>	
Xiaoyue Lu, Xianglin Yang, Haijun Liu, Jiahao Liu, Kuntai Cai, Yan Xiao and Jin Song Dong	30704
<i>HEALing Entropy Collapse: Enhancing Exploration in Few-Shot RLVR via Hybrid-Domain Entropy Dynamics Alignment</i>	
Zhanyu Liu, Qingguo Hu, Ante Wang, Chenqing Liu, Zhishang Xiang, Hui Li, Delai Qiu and Jinsong Su	30732
<i>Do LLMs Encode Functional Importance of Reasoning Tokens ?</i>	
Janvijay Singh and Dilek Hakkani-Tür	30749
<i>EA-Agent: A Structured Multi-Step Reasoning Agent for Entity Alignment</i>	
Yixuan Nan, Xixun Lin, Yanmin Shang, Ge Zhang, Zheng Fang, Fang Fang and Yanan Cao	30774
<i>Too Nice to Tell the Truth: Quantifying Agreeableness-Driven Sycophancy in Role-Playing Language Models</i>	
Arya Shah, Deepali Mishra and Chaklam Silpasuwanchai	30788

<i>Mobile-RI: Towards Interactive Capability for VLM-Based Mobile Agent via Systematic Training</i>	
Jihao Gu, Qihang Ai, Yingyao Wang, Pi Bu, Jingxuan Xing, Yue Cao, Zekun Zhu, Wei Jiang, Ziming Wang, Yingxiu Zhao, Ming-Liang Zhang, Jun Song, Yuning Jiang and Bo Zheng	30802
<i>QuantumQA: Enhancing Scientific Reasoning via Physics-Consistent Dataset and Verification-Aware Reinforcement Learning</i>	
Songxin Qu, Tai-Ping Sun, Yun-Jie Wang, Huan-Yu Liu, Cheng Xue, Xiao-Fan Xu, Han Fang, Yang Yang, Yu-Chun Wu, Guo-Ping Guo and Zhao-Yun Chen	30821
<i>The “Knowledge–Behavior Gap” in Cultural Taboo Safety of Large Language Models</i>	
Ying He, Sihang Jiang, Xingzhou Chen, Zhouhong Gu, Yiwei Gu, Minggui HE, Shimin Tao, Mahongxia and Yanghua Xiao	30846
<i>OSCBench: Benchmarking Object State Change in Text-to-Video Generation</i>	
Xianjing Han, Bin Zhu, Shiqi Hu, Franklin Mingzhe Li, Patrick Carrington, Roger Zimmermann and Jingjing Chen	30867
<i>Evaluation Pitfalls and Challenges in Multimedia Event Extraction</i>	
Philipp Seeberger, Steffen Freisinger, Tobias Bocklet and Korbinian Riedhammer	30885
<i>UMPIRE: Unveiling LLM-generated Posts via Redundant Expressions</i>	
Xiaoquan Yi, Haixing Wu, Haozhao Wang, Yichen Li, Yuhua Li, Rui Zhang and Ruixuan Li	30901
<i>SpeechMedAssist: Efficiently and Effectively Adapting Speech Language Models for Medical Consultation</i>	
Sirry Chen, Jieyi Wang, Wei Chen and Zhongyu Wei	30914
<i>An Experimental Study on the Influence of Culture on Cross-Lingual Sentiment Transfer</i>	
Ahao Liu, Haitong Yang and Wang Chuanrong	30936
<i>Attribution, Citation, and Quotation: A Survey of Evidence-based Text Generation with Large Language Models</i>	
Tobias Schreieder, Tim Schopf and Michael Färber	30956
<i>MirrorCAPTCHA: Wild CAPTCHA, Wild Distribution, Wild Web-based Platform Meet Multimodal LLM Agents</i>	
Xiangyu Wu, Yuwei Hu, Tianyu Cui, Yueying Tian, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, Yang Yang and Jianfeng Lu	31001
<i>ReCreate: Reasoning and Creating Domain Agents Driven by Experience</i>	
Zhezhen Hao, Hong Wang, Jian Luo, Jianqing Zhang, Yuyan Zhou, Qiang Lin, Can Wang, Hande Dong and Jiawei Chen	31018
<i>Can AI-Generated Persuasion Be Detected? Persuaficial Benchmark and AI vs. Human Linguistic Differences</i>	
Arkadiusz Modzelewski, Paweł Golik, Anna Kołos and Giovanni Da San Martino	31047
<i>GASE: Graph-Aware Semantic Embedding Learning with Frozen LLMs for Text-Attributed Graphs</i>	
Mingqian Ding, Jianjun Li, Wenqi Yang, Zhibo Zhang and Yushen Fang	31073
<i>Compositional Steering of Large Language Models with Steering Tokens</i>	
Gorjan Radevski, Kiril Gashteovski, Giwon Hong, Carolin Lawrence and Goran Glavaš . .	31087
<i>Rethinking Entropy Interventions in RLVR: An Entropy Change Perspective</i>	
Zhezhen Hao, Hong Wang, Haoyang Liu, Jian Luo, Jiarui Yu, Hande Dong, Qiang Lin, Can Wang and Jiawei Chen	31105

<i>SwissGov-RSD: A Human-annotated, Cross-lingual Benchmark for Token-level Recognition of Semantic Differences Between Related Documents</i>	
Michelle Wastl, Jannis Vamvas and Rico Sennrich	31134
<i>Seeing but Not Thinking: Routing Distraction in Multimodal Mixture-of-Experts</i>	
Haolei Xu, Haiwen Hong, Hongxing Li, Rui Zhou, Yang Zhang, Longtao Huang, Hui Xue, Yongliang Shen, Weiming Lu and Yueting Zhuang	31164
<i>RubricBench: Aligning Model-Generated Rubrics with Human Standards</i>	
Junyi Zhou, Qiyuan Zhang, Yufei Wang, Fuyuan Lyu, Yidong Ming, Can Xu, Qingfeng Sun, Kai Zheng, Peng Kang, Xue Liu and Chen Ma	31179
<i>Verify Before You Commit: Towards Faithful Reasoning in LLM Agents via Self-Auditing</i>	
Wenhao Yuan, Chenchen Lin, Jian Chen, Jinfeng Xu, Xuehe Wang and Edith Cheuk-Han Ngai	31201
<i>GiLT: Augmenting Transformer Language Models with Dependency Graphs</i>	
Tianyu Huang, Yida Zhao, Chuyan Zhou and Kewei Tu	31226
<i>TAMAS: Benchmarking Adversarial Risks in Multi-Agent LLM Systems</i>	
Ishan Kavathekar, Hemang Jain, Ameya Rathod, Ponnurangam Kumaraguru and Tanuja Ganu	31238
<i>How Controllable Are Large Language Models? A Unified Evaluation across Behavioral Granularities</i>	
Ziwen Xu, Kewei Xu, Haoming Xu, Haiwen Hong, Longtao Huang, Hui Xue, Ningyu Zhang, Yongliang Shen, Guozhou Zheng, Huajun Chen and Shumin Deng	31269
<i>Scaling Behaviors of LLM Reinforcement Learning Post-Training: An Empirical Study in Mathematical Reasoning</i>	
Zelin Tan, Hejia Geng, Xiaohang Yu, Mulei Zhang, Guancheng Wan, Yifan Zhou, Qiang He, Xiangyuan Xue, Heng Zhou, Yutao Fan, Zhong-Zhi Li, Zaibin Zhang, Guibin Zhang, Chen Zhang, Zhenfei Yin, Philip Torr and Lei Bai	31300
<i>RubricHub: A Comprehensive and Highly Discriminative Rubric Dataset via Automated Coarse-to-Fine Generation</i>	
Sunzhu Li, Jiale Zhao, Huimin Ren, Zhenlin Wei, Yang Zhou, Jingwen Yang, Shunyu Liu, Kaike Zhang and Chen Wei	31320
<i>MIND Your Reasoning: A Meta-Cognitive Intuitive-Reflective Network for Dual-Reasoning in Multi-modal Stance Detection</i>	
Bingbing Wang, Zhengda Jin, Bin Liang, Wenjie Li, Jing Li, Ruifeng Xu and Min Zhang .	31345
<i>A Survey of Inductive Reasoning for Large Language Models</i>	
Kedi Chen, Dezhao Ruan, Yuhao Dan, Yaoting Wang, Siyu Yan, Xuecheng Wu, Yinqi Zhang, Qin Chen, Jie Zhou, Liang He, Biqing Qi, Linyang Li, Qipeng Guo, Xiaoming Shi and Wei Zhang .	31356
<i>CogToM: A Comprehensive Theory of Mind Benchmark inspired by Human Cognition for Large Language Models</i>	
Haibo Tong, Zeyang Yue, Feifei Zhao, Erliang Lin, Lu Jia, Ruolin Chen, Yinqian Sun, Qian Zhang and Yi Zeng	31383
<i>CITE: Benchmarking Heterogeneous Text-Attributed Graph Models</i>	
Chenghao Zhang, Qingqing Long, Ludi Wang, Wenjuan Cui, Jianjun Yu and Yi Du	31426
<i>MM-StanceDet: Retrieval-Augmented Multi-modal Multi-agent Stance Detection</i>	
Weihai Lu, Zhejun Zhao, Yanshu Li and Huan He	31449

<i>Disentangling Reasoning Logic to Resolve Explicit Knowledge Conflicts</i>	
Xianda Zheng, Zijian Huang, Meng-Fen Chiang, Jiamou Liu, Yuan Fang, Michael J. Witbrock and Kaiqi Zhao	31465
<i>UniVocal: Unified Speech-Singing Code-Switching Synthesis</i>	
YuFei Shi, Qian Chen, Wen Wang, Xiangang Li, Zhen-Hua Ling and Yang Ai	31479
<i>ReasoningGuard: Safeguarding Large Reasoning Models with Inference-time Safety Aha Moments</i>	
Yuquan Wang, Mi Zhang, Yining Wang, Geng Hong, Mi Wen, Xiaoyu You and Min Yang	31497
<i>VecInfer: Efficient LLM Inference with Low-Bit KV Cache via Outlier-Suppressed Vector Quantization</i>	
Dingyu Yao, Chenxu Yang, Zhengyang Tong, Zheng Lin, Wei Liu, Jian Luan and Weiping Wang	31527
<i>When Does Mixing Help? Analyzing Query Embedding Interpolation in Multilingual Dense Retrieval</i>	
Tongyao Zhu, Huang Chao Ming and Min-Yen Kan	31544
<i>WIST: Web-Grounded Iterative Self-Play Tree for Domain-Targeted Reasoning Improvement</i>	
Fangyuan Li, Pengfei Li, Shijie Wang, Junqi Gao, Jianxing Liu, Biqing Qi and Yuqiang Li	31563
<i>Frankentext: Stitching random text fragments into long-form narratives</i>	
Chau Minh Pham, Jenna Russell, Dzung Pham and Mohit Iyyer	31586
<i>SAFO: Stable Adaptive Fairness Optimization for LLM-Based Social Survey Simulation</i>	
Chenxi Lin, Zhuoren Jiang, Kaisong Song and Yiquan Wu	31626
<i>The Prosody of Emojis</i>	
Giulio Zhou, Tsz Kin Lam, Alexandra Birch and Barry Haddow	31655
<i>Piece of Table: A Divide-and-Conquer Approach for Selecting Subtables in Table Question Answering</i>	
Wonjin Lee, Kyumin Kim, Sungjae Lee, Jihun Lee and Kwang In Kim	31672
<i>QA-MoE: Towards a Continuous Reliability Spectrum with Quality-Aware Mixture of Experts for Robust Multimodal Sentiment Analysis</i>	
Yitong Zhu, Yuxuan Jiang, Guanxuan Jiang, Bojing Hou, Peng Yuan Zhou, Ge Lin and Yuyang Wang	31689
<i>Enhancing LLM-based Search Agents via Contribution Weighted Group Relative Policy Optimization</i>	
Junzhe Wang, Zhiheng Xi, Yajie Yang, Hao Luo, Shihan Dou, Tao Gui and Qi Zhang	31704
<i>Why Steering Works: Toward a Unified View of Language Model Parameter Dynamics</i>	
Ziwen Xu, Chenyan WU, Hengyu Sun, Haiwen Hong, Mengru Wang, Yunzhi Yao, Longtao Huang, Hui Xue, Shumin Deng, Zhixuan Chu, Huajun Chen and Ningyu Zhang	31719
<i>EmoMAS: Emotion-Aware Multi-Agent System for High-Stakes Edge-Deployable Negotiation with Bayesian Orchestration</i>	
Yunbo Long, Yuhan Liu and Liming Xu	31737
<i>Decisive: Guiding User Decisions with Optimal Preference Elicitation from Unstructured Documents</i>	
Akriti Jain, Anish Mulay, Divyansh Verma, Aishani Pandey, Pritika Ramu and Aparna Garimella	31774
<i>From Naturalness to Norms: Interactional Cultural Competence for SpeechLMs</i>	
Santosh T.y.s.s.	31787
<i>Optimizing User Profiles via Contextual Bandits for Retrieval-Augmented LLM Personalization</i>	
Linfeng Du, Ye Yuan, Zichen Zhao, Fuyuan Lyu, Emiliano Penalosa, Xiuying Chen, Zipeng Sun, Jikun Kang, Laurent Charlin, Xue Liu and Haolun Wu	31803

<i>Teaching Language Models to Check Grounded Claim Factuality with Human Test-Taking Strategies</i> Yuxuan Ye, Raul Santos-Rodriguez and Edwin Simpson	31818
<i>Evo-PI: Aligning Medical Reasoning via Evolving Principle-Guided Supervision</i> Xianda Zheng, Huan Gao, Meng-Fen Chiang, Michael J. Witbrock, Kaiqi Zhao and Shangyang Li	31835
<i>Understanding LLM Performance Degradation in Multi-Instance Processing: The Roles of Instance Count and Context Length</i> Jingxuan Chen, Mohammad Taher Pilehvar and Jose Camacho-Collados	31851
<i>Gated Differentiable Working Memory for Long-Context Language Modeling</i> Lingrui Mei, Shenghua Liu, Yiwei Wang, Yuyao Ge, Baolong Bi, Jiayu Yao, Jun Wan, Ziling Yin, Jiafeng Guo and Xueqi Cheng	31885
<i>Stress Testing Factual Consistency Metrics for Long-Document Summarization</i> Zain Muhammad Mujahid, Dustin Wright and Isabelle Augenstein	31914
<i>Do LLMs Know Tool Irrelevance? Demystifying Structural Alignment Bias in Tool Invocations</i> Yilong Liu, Xixun Lin, Pengfei Cao, Ge Zhang, Fang Fang and Yanan Cao	31934
<i>PASs-MoE: Mitigating Misaligned Co-drift among Router and Experts via Pathway Activation Subspaces for Continual Learning</i> ZhiYan Hou, Haiyun Guo, Haokai Ma, Yandu Sun, Yonghui Yang and Jinqiao Wang	31959
<i>LLM-Generated Text May Harm Your Retrieval! A Robust Detection Strategy for Retrieval-Augmented Generation</i> Zhaoheng Huang, Yutao Zhu, Ji-Rong Wen and Zhicheng Dou	31973
<i>GQLBench: A Large-Scale Cross-Domain, Cross-Dialect Benchmark for NL2GQL</i> Yanning Su, Yuhang Zhou, Yang Fang, Sen Liu, Guangnan Ye and Hongfeng Chai	31989
<i>ReviewGrounder: Improving Review Substantiveness with Rubric-Guided, Tool-Integrated Agents</i> Zhuofeng Li, Yi Lu, Dongfu Jiang, Haoxiang Zhang, Yuyang Bai, Chuan Li, Yu Wang, Shuiwang Ji, Jianwen Xie and Yu Zhang	32015
<i>The Retrieval Bottleneck: Scaling Laws for Reinforcement Learning in RAG</i> Shu Zhou, Jinman Leng, Yufei Song, Xin Wang, Tao Fan and Hao Wang	32045
<i>Generating then Refining for Reliable Knowledge Base Question Answering</i> Jianqi Gao, Hang Yu, Jian Cao, Ranran Bu, Jinghua Tang, Nengjun Zhu and Yonggang Zhang	32070
<i>EvoSpark: Endogenous Interactive Agent Societies for Unified Long-Horizon Narrative Evolution</i> Shiyu He, Minchi Kuang, Mengxian Wang, Bin Hu and Tingxiang Gu	32088
<i>SeCuRepair: Semantics-Aligned, Curriculum-Driven, and Reasoning-Enhanced Vulnerability Repair Framework</i> Chengran Yang, Ting Zhang, Jinfeng Jiang, Xin Zhou, Haoye Tian, Mingzhe Du, Jieke Shi, Junkai Chen, Yikun Li, Eng Lieh Ouh, Lwin Khin Shar and David Lo	32108
<i>HiddenGuard: Fine-Grained Safe Generation with Specialized Representation Router</i> Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Ruibin Yuan and Xueqi Cheng	32124
<i>Why Do LLM-based Web Agents Fail? A Hierarchical Planning Perspective</i> Mohamed Aghzal, Gregory J. Stein and Ziyu Yao	32157

<i>Backdoors in RLVR: Jailbreak Backdoors in LLMs From Verifiable Reward</i>	
Weiyang Guo, Zesheng Shi, Zeen Zhu, Yuan Zhou, Min Zhang and Jing Li	32181
<i>On the Proper Treatment of Units in Surprisal Theory</i>	
Samuel Kiegeleland, Vésteinn Snæbjarnarson, Tim Vieira and Ryan Cotterell	32202
<i>ThreadSumm: Summarization of Nested Discourse Threads Using Tree of Thoughts</i>	
Olubusayo Olabisi, Ekata Mitra and Ameeta Agrawal	32225
<i>InquireMobile: Teaching VLM-based Mobile Agent to Request Human Assistance via Reinforcement Fine-Tuning</i>	
Qihang Ai, Pi Bu, Yue Cao, Yingyao Wang, Jihao Gu, Jingxuan Xing, Zekun Zhu, Wei Jiang, Zhicheng Zheng, Jun Song and Yuning Jiang	32241
<i>Token-Level Policy Optimization: Linking Group-Level Rewards to Token-Level Aggregation via sequence-level likelihood</i>	
Xingyu Lin, Yilin Wen, Du Su, En Wang, Wenbin Liu, Zhonghou Lv, Jinchang Hou and Chenfu Bao	32256
<i>De-Anonymization at Scale via Tournament-Style Attribution</i>	
Lirui Zhang and Huishuai Zhang	32270
<i>Cross-Modal Masked Compositional Concept Modeling for Enhancing Visio-Linguistic Compositionality</i>	
Wei Li, Zhen Huang and Xinmei Tian	32284
<i>StratMem-Bench: Evaluating Strategic Memory Use in Virtual Character Conversation Beyond Factual Recall</i>	
Yerong Wu, Tianxing Wu, Minghao Zhu, Hangyu Sha and Haofen Wang	32309
<i>RLSeek: Evidence-Grounded Reasoning for RAG Hallucination Detection</i>	
Zhaoheng Huang, Dacheng Wen, Yutao Zhu, Xiaoying Lian, Yushi Liang, Kai Hao, Nan Li, Liangjie Zhang, Qi Zhang, Ji-Rong Wen, Zhicheng Dou and Fangzhao Wu	32329
<i>CO-EVO: Co-evolving Semantic Anchoring and Style Diversification for Federated DG-ReID</i>	
Fengchunzhang, Qiang Ma, Liuyu Xiang, Jinshan Lai, Tingxuan Huang and Jianwei Hu ..	32348
<i>Discovering Properties of Inflectional Morphology in Neural Emergent Communication</i>	
Miles Gilberti, Shane Storcks and Huteng Dai	32360
<i>Label Effects: Shared Heuristic Reliance in Trust Assessment by Humans and LLM-as-a-Judge</i>	
Xin Sun, Di Wu, Sijing Qin, Isao Echizen, Abdallah El Ali and Saku Sugawara	32378
<i>Self-SoftCoT: A Self-Consistent Framework via Position-Aware Latent Space Reinforcement Learning</i>	
Liangliang Dong, Lianlei Shan and Shuaimin Li	32393
<i>Bayesian Social Deduction with Graph-Informed Language Models</i>	
Shahab Rahimirad, Guven Gergerli, Lucia Romero, Angela Qian, Matthew Lyle Olson, Simon Stepputtis and Joseph Campbell	32415
<i>Routing with Generated Data: Annotation-Free LLM Skill Estimation and Expert Selection</i>	
Tianyi Niu, Justin Chen, Genta Indra Winata, Shi-Xiong Zhang, Supriyo Chakraborty, Sambit Sahu, Yue Zhang, Elias Stengel-Eskin and Mohit Bansal	32441
<i>Injecting Context via Situation Working Memory for Logical Reasoning with LLMs</i>	
Jieun Kim, Seoha Lim, YoungHae Choi and Sung-Bae Cho	32467

<i>ProMed: Shapley Information Gain Guided Reinforcement Learning for Proactive Medical LLMs</i>	
Hongxin Ding, Baixiang Huang, Yue Fang, Weibin Liao, Xinke Jiang, Jinyang Zhang, Yinghao Zhu, Zheng Li, Liantao Ma, Junfeng Zhao and Yasha Wang	32481
<i>SafeAgent: Safeguarding LLM Agents via an Automated Risk Simulator</i>	
Xueyang Zhou, Weidong Wang, Lin Lu, Jiawen Shi, Guiyao Tie, Xu Yongtian, Lixing Chen, Pan Zhou, Neil Zhenqiang Gong and Lichao Sun	32516
<i>Diagnosing Hidden Instabilities in Model Editing via Uncertainty Quantification</i>	
Zihan Gu, TianYi Zhang, Xinyan Zhang, Zhiyuan Wang, Han Zhang, Yuhao Wei, Jiacheng Lu, Tianyi Ma, Xingsheng Zhang, Hua Zhang and Yue Hu	32544
<i>Alexandria: A Multi-Domain Dialectal Arabic Machine Translation Dataset for Culturally Inclusive and Linguistically Diverse LLMs</i>	
Abdellah EL Mekki, Samar M. Magdy, Houdaifa Atou, Ruwa AbuHweidi, Baraah Qawasmeh, Omer Nacar, Thikra Al-hibiri, Razan Saadie, Hamzah A. Alsayadi, Nadia Ghezaiel Hammouda, Alshima Mohammed Alkhazimi, Aya Hamod, Al-Yas Yaqoob Al-Ghafri, Wesam El-Sayed, Asila Ismail al Sharji, Mohamad Ballout, Anas Belfathi, Karim Ghaddar, Serry Sibae, Alaa Aoun, Aeej Mohammed Aseri, Lina Abureesh, Ahlam Bashiti, Majdal Yousef, Abdulaziz Hafiz, Yehdih Mohamed, Emira Hamedtou, Brakehe Emegah, Rahaf Alhamouri, Youssef Nafea, Aya El Aatar, Walid Al-Dhabyani, Emhemed S. Hamed, Sara Shatnawi, Fakhraddin Alwajih, Khalid Elkhidir, Ashwag Alasmari, Abdurrahman Gerrio, Omar Said Alshahri, AbdelRahim A. Elmadany, Ismail Berrada, Amir Azad Adli Al-kathiri, Fadi Zaraket, Mustafa Jarrar, Yahya Mohamed EL Hadj, Hassan Alhuzali and Muhammad Abdul-Mageed	32567
<i>Can AI Be a Good Peer Reviewer? A Survey of Peer Review Process, Evaluation, and the Future</i>	
Sihong Wu, Owen Jiang, Yilun Zhao, Tiansheng Hu, Yiling Ma, Kaiyan Zhang, Manasi Patwardhan and Arman Cohan	32593
<i>How Much Would a Clinician Edit This Draft? Evaluating LLM Alignment for Patient Message Response Drafting</i>	
Parker Seegmiller, Joseph Gatto, Sarah E. Greer, Ganza Belise Isingizwe, Rohan Ray, Timothy E. Burdick and Sarah Masud Preum	32620
<i>MetaBench: A Multi-task Benchmark for Assessing LLMs in Metabolomics</i>	
Yuxing Lu, Xukai Zhao, J. Ben Tamo, Micky C. Nnamdi, Rui Peng, Shuang Zeng, Xingyu Hu, Jinzhuo Wang and May Dongmei Wang	32647
<i>MTA: Multi-Granular Trajectory Alignment for Large Language Model Distillation</i>	
Pham Khanh Chi, Quoc Phong Dao, Thuat Nguyen, Linh Ngo Van, Trung Le and Thanh Hong Nguyen	32669
<i>Measuring Distribution Shift in User Prompts and Its Effects on LLM Performance</i>	
Parker Seegmiller and Sarah Masud Preum	32685
<i>TALAS: Teacher-Anchored Layer Alignment with Adaptive Sharpness-Aware Minimization for Embedding Distillation</i>	
Quoc Phong Dao, Hoang Son Nguyen, Pham Khanh Chi, Linh Ngo Van, Nguyen Thi Ngoc Diep, Thien Huu Nguyen and Trung Le	32705
<i>Vocabulary Shapes Cross-Lingual Variation of Word-Order Learnability in Language Models</i>	
Jonas Mayer Martins, Jaap Jumelet, Viola Priesemann and Lisa Beinborn	32724
<i>Quantifying and Understanding Uncertainty in Large Reasoning Models</i>	
Yangyi Li, Chenxu Zhao and Mengdi Huai	32741

<i>SimPBL: A Multi-Agent Framework for Project-Based Learning</i>	
Daniel Zhang-Li, Joy Jia Yin Lim, Binglin Liu, Shangqing Tu, Zijun Yao, Hao Peng, Jifan Yu, Haoxuan Li, Zhanxin Hao, Ye He, Zekun Li, Jiangyi Wang, Lei Hou, Bin Xu, Xin Cong, Zhiyuan Liu, Huiqin Liu, Yu Zhang and Juanzi Li	32755
<i>Schoenfeld’s Anatomy of Mathematical Reasoning by Language Models</i>	
Ming Li, Chenrui Fan, Yize Cheng, Soheil Feizi and Tianyi Zhou	32773
<i>Diagnosing Spatial Consistency across Perspectives and Viewpoints in Large Vision-Language Models</i>	
Yoonji Kim, Jieun Kim, Yujin Jeong and Sung-Bae Cho	32803
<i>A Multi-Agent Framework for High-Interaction Terminal Simulation</i>	
Kai Wei, Yuwen Cui, Kehan Shen, Hua Wei and Guangjing Wang	32828
<i>Parameter Importance is Not Static: Evolving Parameter Isolation for Supervised Fine-Tuning</i>	
Zekai Lin, Chao Xue, Di Liang, Xingsheng Han, Peiyang Liu, Xianjie Wu, Lei Jiang, Yu Lu, Bob Simons, Shuang Liang and Minlong Peng	32846
<i>Enhanced Reasoning for Biomedical Document-Level Relation Extraction via a Novel Cascade Language Model Framework</i>	
Haohua Song, Wenhao Gu, Zhijing Li, Yunwenyu, Tiantian Zhu, Xiao Yang and Zexuan Zhu	32865
<i>Multimodal Large Language Models for Multi-Subject In-Context Image Generation</i>	
Yucheng Zhou, Dubing Chen, Huan Zheng and Jianbing Shen	32880
<i>RSMem: Knowledge-Enhanced Memory Evolution for Remote Sensing Agents with Systematic Evaluation</i>	
Bingxian Wu, Yu Zhang, Zonghao Guo, Tang Liu, Chen Qian, Yuxiang Lu, Xingbo Du, Yanghao Li, Yidan Zhang, Chi Chen, Ling Yao, Chenghu Zhou and Maosong Sun	32899
<i>Synthetic Data Generation for Training Diversified Commonsense Reasoning Models</i>	
Tianhui Zhang, Bei Peng and Danushka Bollegala	32916
<i>Discourse Realization of Generics in Human and LLM-generated Texts</i>	
Søren Kirkegaard Fomsgaard, Martial Pastor, Gaël Dias and Nelleke Oostdijk	32938
<i>SRA: Span Representation Alignment for Large Language Model Distillation</i>	
Quoc Phong Dao, Hoang Son Nguyen, Pham Khanh Chi, Tung Nguyen, Linh Ngo Van, Nguyen Thi Ngoc Diep and Trung Le	32961
<i>The GDN-CC Dataset: Automatic Corpus Clarification for AI-enhanced Democratic Citizen Consultations</i>	
Pierre-Antoine Lequeu, Léo Labat, Laurène Cave, Gaël Lejeune, François Yvon and Benjamin Piwowarski	32976
<i>Behavior Knowledge Merge in Reinforced Agentic Models</i>	
Xiangchi Yuan, Dachuan Shi, Chunhui Zhang, Zheyuan Liu, Shenglong Yao, Soroush Vosoughi and Wenke Lee	33007
<i>Leveraging Outline-Optimized Generative Interactions and Critique for Self-Refining Outlines with Reinforcement Learning</i>	
Hengwei Liu, Haoyuan Ma, Qingqing Lyu, Daoxin Zhang, Yao Hu, Yongliang Shen, Yin Zhang and Weiming Lu	33029
<i>Neuron-Aware Active Few-Shot Learning for LLMs</i>	
Zhuowei Chen, Liwei Chen, Christian Schunn, Raquel Coelho and Xiang Lorraine Li	33047

CommonLID: Re-evaluating State-of-the-Art Language Identification Performance on Web Data

Pedro Ortiz Suarez, Laurie Burchell, Catherine Arnett, Rafael Mosquera, Sara Hincapié Monsalve, Thom Vaughan, Damian Stewart, Malte Ostendorff, Idris Abdulmumin, Vukosi Marivate, Shamsuddeen Hassan Muhammad, Atnafu Lambebo Tonja, Hend Al-Khalifa, Nadia Ghezaiel Hammouda, Verrah Akinyi Otiende, Tack Hwa Wong, Jakhongir Saydaliev, Melika Nobakhtian, Muhammad Ravi Shulthan Habibi, Chalamalasetti Kranti, Carol Muchemi, Khang Nguyen, Faisal Muhammad Adam, Luis Frentzen Salim, Reem Alqifari, Cynthia Jayne Amol, Joseph Marvin Imperial, Ilker Kesen, Ahmad Mustafid, Pavel Stepachev, Leshem Choshen, David Anugraha, Hamada Nayel, Seid Muhie Yimam, Vallerie Alexandra Putra, My Chiffon Nguyen, Azmine Toushik Wasi, Gouthami Vadithya, Rob Van Der Goot, Lanwenn ar C'horr, Karan Dua, Andrew Yates, Mithil Bangera, Yeshil Bangera, Hitesh Laxmichand Patel, Shu Okabe, Fenal Ashokbhai Ilasariya, Dmitry Gaynullin, Genta Indra Winata, Yiyuan Li, Juan Pablo Martínez, Amit Agarwal, Ikhlusal Akmal Hanif, Raia Abu Ahmad, Esther Adenuga, Filbert Aurelian Tjiaranata, Weerayut Buaphet, Michael Anugraha, Sowmya Vajjala, Benjamin L Rice, Azril Hafizi Amirudin, Jesujoba Oluwadara Alabi, Srikant Panda, Yassine Toughrai, Bruhan Kyomuhendo, Daniel Ruffinelli, Akshata, Manuel Goulão, Ej Zhou, Ingrid Gabriela Franco Ramirez, Cristina Aggazzotti, Konstantin Dobler, Jun Kevin, Quentin Pagès, Nicholas Andrews, Nuhu Ibrahim, Mattes Ruckdeschel, Amr Keleg, Mike Zhang, Casper Rufaro Muziri, Saron Samuel, Sotaro Takeshita, Kun Kerdthaisong, Luca Foppiano, Rasul Dent, Tommaso Green, Ahmad Mustapha Wali, Kamohe-lo Makaaka, Vicky Feliren, Inshirah Idris, Hande Celikkanat, Abdulhamid Abubakar, Jean Maillard, Benoît Sagot, Thibault Clérice, Kenton Murray and Sarah K. K. Luger 33063

Rethinking Expert Trajectory Utilization in LLM Post-training for Mathematical Reasoning

Bowen Ding, Yuhan Chen, Jiayang Lyu, Jiyao Yuan, Qi Zhu, Shuangshuang Tian, Dantong Zhu, Futing Wang, Heyuan Deng, Fei Mi, Lifeng Shang and Tao Lin 33081

LoPT: Lossless Parallel Tokenization Acceleration for Long Context Inference of Large Language Model

Wei Shao, Zheng Lingchao, Pengyu Wang, Peizhen Zheng, Li Jun and Yuwei Fan 33107

SSSD: Simply-Scalable Speculative Decoding

Michele Marzollo, Jiawei Zhuang, Niklas Roemer, Niklas Zwingenberger, Lorenz K Muller and Lukas Cavigelli 33123

RedCoder: Automated Multi-Turn Red Teaming for Code LLMs

Wenjie Jacky Mo, Qin Liu, Xiaofei Wen, Dongwon Jung, Hadi Askari, Wenxuan Zhou, Zhe Zhao and Muhao Chen 33140

Scaling Test-Time Compute to Achieve IOI Gold Medal with Open-Weight Models

Mehrzad Samadi, Aleksander Ficek, Sean Narenthiran, Siddhartha Jain, Wasi Uddin Ahmad, Somshubra Majumdar, Vahid Noroozi and Boris Ginsburg 33156

PIArena: A Platform for Prompt Injection Evaluation

Runpeng Geng, Chenlong Yin, Yanting Wang, Ying Chen and Jinyuan Jia 33170

Clozing the Gap: Exploring Why Language Model Surprisal Outperforms Cloze Surprisal

Sathvik Nair and Byung-Doh Oh 33193

ChemReason-Bench: Benchmarking Large Language Models for Procedural Reasoning in Experimental Chemistry

Jinwei Zhang, Xucheng Liang, Yu Zhang, Ruijie Yu, Xiaokang Yang, Yaohui Jin and Yanyan Xu 33211

How Instruction and Reasoning Data shape Post-Training: Data Quality through the Lens of Layer-wise Gradients

Ming Li, Yanhong Li, Ziyue Li and Tianyi Zhou 33249

<i>Tailored Primitive Initialization is the Secret Key to Reinforcement Learning</i>	
Yihang Yao, Guangtao Zeng, Raina Wu, Yang Zhang, Ding Zhao, Zhang-Wei Hong and Chuang Gan	33300
<i>MARS²: Scaling Multi-Agent Tree Search via Reinforcement Learning for Code Generation</i>	
Pengfei Li, Shijie Wang, Fangyuan Li, Yikun Fu, Kaifeng Liu, Kaiyan Zhang, Dazhi Zhang, Yuqiang Li, Biqing Qi and Bowen Zhou	33319
<i>LoVeC: Reinforcement Learning for Better Verbalized Confidence in Long-Form Generation</i>	
Caiqi Zhang, Xiaochen Zhu, Chengzu Li, Nigel Collier and Andreas Vlachos	33336
<i>Mitigating Lost in Multi-turn Conversation via Curriculum RL with Verifiable Accuracy and Abstention Rewards</i>	
Ming Li, Pei Chen, Zhenhao Zhang, Tao Yang, Xinyang Zhang, Han Li, Tianyu Cao, Ming Zeng, Zhuofeng Wu, Meng Jiang, Huasheng Li, Lihong Li and Bing Yin	33364
<i>AdabNER: Arabic Digital Archive Books with Nested Entity Recognition</i>	
Aya Mourad and Mustafa Jarrar	33382
<i>Anchoring the Cache: Mitigating Contextual Hallucination in KV-Compressed Long-Context Summarization</i>	
Yu Fu, Chen Luo, Josef Valvoda, Xin Zhang, Xuejing Lei, Xiao Pan, Hui Liu and Yue Dong	33397
<i>NeuralFSM: Adaptive Multi-Agent Coordination via Learning Finite-State Execution Policy</i>	
Jiye Wang, Yu Wang, Jianbin Li, Shiduo Yang, Kenan Guo and Yuanhe Zhao	33414
<i>Ranking Reasoning LLMs under Test-Time Scaling</i>	
Mohsen Hariri, Michael Hinczewski, Jing Ma and Vipin Chaudhary	33437
<i>Quantifying and Improving the Robustness of Retrieval-Augmented Language Models Against Spurious Features in Grounding Data</i>	
Shiping Yang, Jie Wu, Wenbiao Ding, Ning Wu, Shining Liang, Ming Gong, Hongzhi Li, Hengyuan Zhang, Angel X Chang and Dongmei Zhang	33479
<i>CL²GEC: A Multi-Discipline Benchmark for Continual Learning in Chinese Literature Grammatical Error Correction</i>	
Shang Qin, Jingheng Ye, Yinghui Li, Hai-Tao Zheng, Qi Li, Jinxiao Shan, Zhixing Li and Hong-Gee Kim	33500
<i>PEARL: Self-Evolving Assistant for Time Management with Reinforcement Learning</i>	
Bingxuan Li, Jeonghwan Kim, Cheng Qian, Xiusi Chen, Eitan Anzenberg, Niran Kundapur and Heng Ji	33522
<i>Beyond Static Artifacts: An Evolutionary Framework for Synthetic Claim Generation</i>	
Yeqing Teng, Jiasheng Si, Shuxia Lin, Linhai Zhang, Weiyu Zhang, Wenpeng Lu, Deyu Zhou and Xiaoming Wu	33540
<i>CloneMem: Benchmarking Long-Term Memory for AI Clones</i>	
Sen Hu, Zhiyu Zhang, Yuxiang Wei, Xueran Han, Zhenheng Tang, Ronghao Chen and Huacan Wang	33571
<i>Memory efficiency and resource-rational encoding in sentence processing</i>	
Weijie Xu, Brian Dillon and Richard Futrell	33603
<i>PRISM: Probing Reasoning, Instruction, and Source Memory in LLM Hallucinations</i>	
Yuhe Wu, Guangyu Wang, Yuran Chen, Jiatong Zhang, Yutong Zhang, Yujie Chen, Jiaming Shang, Guang Zhang and Zhuang Liu	33619

<i>VC-Inspector: Advancing Reference-free Evaluation of Video Captions with Factual Analysis</i> Shubhashis Roy Dipta, Tz-Ying Wu and Subarna Tripathi	33657
<i>MQM Re-Annotation: A Technique for Collaborative Evaluation of Machine Translation</i> Parker Riley, Daniel Deutsch, Mara Finkelstein, Colten DiIanni, Juraj Juraska and Markus Freitag 33673	
<i>Multilingual Language Models Encode Script Over Linguistic Structure</i> Aastha A K Verma, Anwoy Chatterjee, Mehak Gupta and Tanmoy Chakraborty	33685
<i>ChiKhaPo: A Large-Scale Multilingual Benchmark for Evaluating Lexical Comprehension and Generation in Large Language Models</i> Emily Chang and Niyati Bafna	33720
<i>UniDataBench: Evaluating Data Analytics Agents Across Structured and Unstructured Data</i> Han Weng, Zhou Liu, Yuanfeng Song, Xiaoming Yin, Xing Chen and Wentao Zhang	33755
<i>Beyond Code Pairs: Dialogue-Based Data Generation for LLM Code Translation</i> Le Chen, Nuo Xu, Winson Chen, Bin Lei, Pei-Hung Lin, Dunzhi Zhou, Rajeev Thakur, Caiwen Ding, Ali Jannesari and Chunhua Liao	33781
<i>MM-PoisonRAG: Disrupting Multimodal RAG with Local and Global Knowledge Poisoning Attacks</i> Hyeonjeong Ha, Qiusi Zhan, Jeonghwan Kim, Dimitrios Bralios, Saikrishna Sanniboina, Nanyun Peng, Kai-Wei Chang, Daniel Kang and Heng Ji	33804
<i>IterCOMP: Reasoning-aware Adaptive Prompt Compression for Multi-hop Question Answering</i> JungMin Yun and YoungBin Kim	33827
<i>Do LLMs Really Memorize Personally Identifiable Information? Revisiting PII Leakage with a Cue-Controlled Memorization Framework</i> Xiaoyu Luo, Yiyi Chen, Qiongxiu Li and Johannes Bjerva	33841
<i>Observations and Remedies for Large Language Model Bias in Self-Consuming Performative Loop</i> Yaxuan Wang, Zhongteng Cai, Yujia Bao, Xueru Zhang and Yang Liu	33862
<i>RAG-on-a-Diet: A Reinforcement Learning-Based Dynamic Resource Optimization Framework for RAG</i> Hongwen Ding and Yizheng Zhao	33883
<i>MTSQL-R1: Towards Long-Horizon Multi-Turn Text-to-SQL via Agentic Training</i> Taicheng Guo, Hai Wang, Chaochun Liu, Mohsen Golalikhani, Xin Chen, Xiangliang Zhang and Chandan K. Reddy	33905
<i>The Subjectivity of Respect in Police Traffic Stops: Modeling Community Perspectives in Body-Worn Camera Footage</i> Prene Golazizian, Elnaz Rahmati, Jackson Trager, Zhivar Sourati, Nona Ghazizadeh, Georgios Chochlakis, Jose J. Alcocer, Kerby Bennett, Aarya Vijay Devnani, Parsa Hejabi, Harry G. Muttram, Akshay Kiran Padte, Mehrshad Saadatinia, Chenhao Wu, Alireza Salkhordeh Ziabari, Michael Sierra-Arévalo, Nicholas Weller, Shrikanth Narayanan, Benjamin A.t. Graham and Morteza Dehghani	33939
<i>LLM-Based Multi-Task Bangla Hate Speech Detection: Type, Severity, and Target</i> Md Arid Hasan, Firoj Alam, Md Fahad Hossain, Usman Naseem and Syed Ishtiaque Ahmed	33962
<i>MMSciCode: Real-world Evaluation of Multilingual Multi-Discipline Scientific Research Coding</i> Xue Xia, Zheyuan Yang, Arman Cohan and Yilun Zhao	33981

<i>CURA: Clinical Uncertainty Risk Alignment for Language Model–Based Risk Prediction</i> Sizhe Wang, Ziqi Xu, Claire Najjuuko, Charles Alba and Chenyang Lu	34000
<i>Unlocking the Black Box of Latent Reasoning: An Interpretability-Guided Approach to Intervention</i> Shuochen Chang, Tong Bai, Xiaofeng Zhang, Qianli Ma, Qingyang Liu, Zhaohe Liao, Yibo Miao and Li Niu	34019
<i>EthicMind: A Risk-Aware Framework for Ethical-Emotional Alignment in Multi-Turn Dialogue</i> Jiawen Deng, Wei Li, Wentao Zhang, Ziyun Jiao and Fuji Ren	34033
<i>Resolving the Security-Auditability Dilemma with Auditable Latent Chain-of-Thought Alignment</i> Guan Wang, Biyu Zhou, Xuehai Tang, Jizhong Han and Songlin Hu	34051
<i>ART: Attention Replacement Technique to Improve Factuality in LLMs</i> Ziqin Luo, Yihao Quan, Xiaofeng Zhang, Xiaosong Yuan and Chen Shen	34068
<i>More Thinking, Less Talking: Internalizing Deliberative Safety into LLM Parameters</i> Guan Wang, Xuehai Tang, Biyu Zhou, Jizhong Han and Songlin Hu	34079
<i>ToolScope: Enhancing LLM Agent Tool Use through Tool Merging and Context-Aware Filtering</i> Marianne Menglin Liu, Daniel Garcia, Fjona Parllaku, Vikas Upadhyay, Fahad Shah and Dan Roth	34095
<i>PRInTS: Reward Modeling for Long-Horizon Information Seeking</i> Jaewoo Lee, Archiki Prasad, Justin Chen, Zaid Khan, Elias Stengel-Eskin and Mohit Bansal	34120
<i>Follow the Flow: On Information Flow Across Textual Tokens in Text-to-Image Models</i> Guy Kaplan, Michael Toker, Yuval Reif, Yonatan Belinkov and Roy Schwartz	34139
<i>Fiction Flows: A Replication and Reinterpretation of Narrative Sequentiality</i> Andrew Piper, Sil Hamilton, Haiqi Zhou and Federico Pianzola	34158
<i>Conjunctive Prompt Attacks in Multi-Agent LLM Systems</i> Nokimul Hasan Arif, Qian Lou and Mengxin Zheng	34175
<i>Evaluating Visual Narrative Coherence in Story Visualization via Diversified Storylines</i> Minha Jhang, Kyeongman Park, Hyukhun Koh and Kyomin Jung	34192
<i>Rethinking the Idiomaticity Decomposability Hypothesis: Evidence from Distributional Learning</i> Maggie Mi, Golzar Atefi, Atsuki Yamaguchi, Felix Gers, Aline Villavicencio and Nafise Sadat Moosavi	34208
<i>KinyaProp: Fine-Grained Propaganda Annotation in Kinyarwanda</i> Manzi Fabrice Niyigaba, Ivory Yang and Soroush Vosoughi	34230
<i>Revisiting Audio-language Pretraining for Learning General-purpose Audio Representation</i> Wei-Cheng Tseng, Xuanru Zhou, Mingyue Huo, Yiwen Shao, Hao Zhang and Dong Yu ..	34244
<i>Trajectory Signatures of Deception in Large Language Models</i> Virraji Mothukuri and Reza M. Parizi	34264
<i>Improving the Distributional Alignment of LLMs using Supervision</i> Gauri Kambhatla, Sanjana Gautam, Angela Zhang, Alexander Liu, Ravi Srinivasan, Junyi Jessy Li and Matthew Lease	34277
<i>SMARTER: A Data-efficient Framework to Improve Toxicity Detection with Explanation via Self-augmenting Large Language Models</i> Huy Nghiem, Advik Sachdeva and Hal Daumé Iii	34306

<i>GameplayQA: A Benchmarking Framework for Decision-Dense POV-Synced Multi-Video Understanding of 3D Virtual Agents</i>	
Yunzhe Wang, Runhui Xu, Kexin Zheng, Tianyi Zhang, Jayavibhav Niranjana Kogundi, Soham Hans and Volkan Ustun	34333
<i>DeepFact: Co-Evolving Benchmarks and Agents for Deep Research Factuality</i>	
Yukun Huang, Leonardo F. R. Ribeiro, Momchil Hardalov, Bhuwan Dhingra, Markus Dreyer and Venkatesh Saligrama	34356
<i>RI-RE: Cross-Domain Relation Extraction with RLVR</i>	
Runpeng Dai, Tong Zheng, Run Yang, Kaixian Yu and Hongtu Zhu	34387
<i>Training LLMs for Divide-and-Conquer Reasoning Elevates Test-Time Scalability</i>	
Xiao Liang, Zhong-Zhi Li, Zhenghao Lin, Eric Hanchen Jiang, Hengyuan Zhang, Yelong Shen, Kai-Wei Chang, Ying Nian Wu, Yeyun Gong and Weizhu Chen	34402
<i>ClaimDB: A Fact Verification Benchmark over Large Structured Data</i>	
Michael Theologitis, Preetam Prabhu Srikanth Dammu, Chirag Shah and Dan Suciu	34428
<i>PluRule: A Benchmark for Moderating Pluralistic Communities on Social Media</i>	
Zoher Kachwala, Bao Tran Truong, Rasika Muralidharan, Haewoon Kwak, Jisun An and Filippo Menczer	34452
<i>What Makes a Good Curriculum? Disentangling the Effects of Data Ordering on LLM Mathematical Reasoning</i>	
Yaning Jia, Chunhui Zhang, Xingjian Diao, Xiangchi Yuan, Zhongyu Ouyang, Chiyu Ma and Soroush Vosoughi	34472
<i>CASS: Nvidia to AMD Transpilation with Data, Models, and Benchmark</i>	
Ahmed Heakl, Gustavo Bertolo Stahl, Sarim Hashmi, Seung Hun Eddie Han, Mukul Ranjan, Arina Kharlamova, Salman Khan and Abdulrahman Mahmoud	34489
<i>SAHM: A Benchmark for Arabic Financial and Shari’ah-Compliant Reasoning</i>	
Rania Elbadry, Sarfraz Ahmad, Ahmed Heakl, Dani Bouch, Momina Ahsan, Muhra AlMahri, Marwa Elsaid Khalil, Yuxia Wang, Salem Lahlou, Sophia Ananiadou, Veselin Stoyanov, Jimin Huang, Xueqing Peng, Preslav Nakov and Zhuohan Xie	34509
<i>Confident, Calibrated, or Complicit: Safety Alignment and Ideological Bias in LLM Hate Speech Detection</i>	
Sanjeevan Selvaganapathy and Mehwish Nasim	34537
<i>Aligning Large Language Models via Fully Self-Synthetic Data</i>	
Shangjian Yin, Zhepei Wei, Xinyu Zhu, Wei-Lin Chen and Yu Meng	34553
<i>Corpus-Dependent Subcharacter Encoding via HMM-Guided Code Assignment</i>	
Tatsuya Hiraoka	34569
<i>DRInQ: Evaluating Conversational Implicature with Controlled Context Variation</i>	
Hirona Jacqueline Arai and Xiang Ren	34594
<i>From Individual to Common: An Early Exploration of Consensus in Non-verifiable Data for Balanced Preference Optimization</i>	
Shangjian Yin and Zhouxing Shi	34612
<i>RIMRULE: Improving Tool-Using Language Agents via MDL-Guided Rule Learning</i>	
Xiang Gao, Yuguang Yao, Qi Zhang, Kaiwen Dong, Avinash Baidya, Ruocheng Guo, Hilaf Hassan and Kamalika Das	34631

<i>GAM: Hierarchical Graph-based Agentic Memory for LLM Agents</i>	
Zhaofen Wu, Hanrong Zhang, Fulin Lin, Wujiang Xu, Xinran Xu, Yankai Chen, Henry Peng Zou, Shaowen Chen, Weizhi Zhang, Xue Liu, Philip S. Yu and Hongwei Wang	34647
<i>CypherSmith: Transforming Text-to-Cypher Generation for LLMs with Synthetic Data</i>	
Zeyu Zhang, Kexuan Sun, Zheng Tang, Jens-S. Vöckler, Thien Huu Nguyen and Thuy Vu	34665
<i>New Terms, New Toxicity: Consensus-based Chinese Neologism Toxicity Detection via Search-Augmented LLMs</i>	
Shiyao Cui, QingLin Zhang, Di Wang, Yida Lu, Zhexin Zhang, Jinhua Gao, Jinglin Yang, Min He, Han Qiu and Minlie Huang.....	34683
<i>ATGL: An Adaptive-Threshold Global Loss for Document-level Relation Extraction</i>	
Huangming Xu, Fu Zhang, Zhixuan Yang, Lu Zhang and Jingwei Cheng	34702
<i>L2Dir: Integrating L₂-Norm and Directional Alignment for Unsupervised Contrastive Representation Learning in Multimodal Retrieval</i>	
Tianyu Zong, Rui Dai, Hongzhu Yi, Yuanxiang Wang, Zhenghao Zhang, Zhenyu Guan, Yujia Yang, Bingkang Shi, Yueyang Ding, Xiangxiang Chu, Kaikui Liu and Jungang Xu	34717
<i>Towards Intrinsic Interpretability of Large Language Models: A Survey of Design Principles and Architectures</i>	
Yutong Gao, Qinglin Meng, Yuan Zhou and Liangming Pan.....	34741
<i>ScholaWrite: A Dataset of End-to-End Scholarly Writing</i>	
Khanh Chi Le, Linghe Wang, Minhwa Lee, Ross Volkov, Luan Tuyen Chau and Dongyeop Kang	34755
<i>What Deserves Memory: Adaptive Memory Distillation for LLM Agents</i>	
Wenquan Ma, Jiayan Nan and WenLong Wu.....	34789
<i>SOAPTriage: SOAP-Guided Multi-View Clinical Text Modeling Framework for Automated ESI Prediction</i>	
Enming Wang, Jianlei Wang, Xueping Peng, Hongjiao Guan, Yinglong Wang, Sibao Wei, Jianbin Guo, Ruifeng Xu and Wenpeng Lu	34813
<i>MoRI: Learning Motivation-Grounded Reasoning for Scientific Ideation in Large Language Models</i>	
Chenyang Gu, Jiahao Cheng, Meicong Zhang, Pujun Zheng, Jinquan Zheng and Guoxiu He	34838
<i>WSDPO: A Generative Word Sense Disambiguation Framework with Chain-of-Thought and Preference Optimization</i>	
Kunpeng Kang, Shuaimin Li, Kaiyuan Zhang, Luyang Zhang, Jiasheng Si, Bing Xu, Kehai Chen, Muyun Yang and Wenpeng Lu	34870
<i>From Regulatory Approvals to Patents: Cross-Domain Linking for Cardiovascular Device Traceability</i>	
Qingqing Yang, Haijiang Liu and Moyan Li	34890
<i>TrustTable: A Neuro-Symbolic Auditing Framework for Faithful Table QA</i>	
Guangzhen Zhao, Dechang Kong, Tongyu Wu and Zhenjiang Dong	34907
<i>Lizard: An Efficient Linearization Framework for Large Language Models</i>	
Chien Van Nguyen, Huy Huu Nguyen, Ruiyi Zhang, Hanieh Deilamsalehy, Puneet Mathur, Viet Dac Lai, Haoliang Wang, Jayakumar Subramanian, Ryan A. Rossi, Trung Bui, Nikos Vlassis, Franck Dernoncourt and Thien Huu Nguyen.....	34935
<i>TPS-Bench: Evaluating AI Agents' Tool Planning & Scheduling Abilities in Compounding Tasks</i>	
Hanwen Xu, Xuyao Huang, Yuzhe Liu and Zhijie Deng	34949

<i>S2S-Arena: Evaluating Paralinguistic Instruction Following in Speech-to-Speech Models</i>	
Feng Jiang, Zhiyu Lin, Yiyang Liu, Liumeng Xue, Fan Bu, Yuhao Du, Xiangying Chen, Benyou Wang and Haizhou Li	34962
<i>Can LLM Safety Be Ensured by Constraining Parameter Regions?</i>	
Zongmin Li, Jian Su, Farah Benamara and Aixin Sun	34979
<i>CityVG: Contrastive Fine-Tuning and Reward-Based Chain-of-Thought Reasoning for Zero-Shot City-Scale 3D Visual Grounding</i>	
Jianjun Zhang and Hanli Wang	35012
<i>E2EDev: Benchmarking Large Language Models in End-to-End Software Development Task</i>	
Jingyao Liu, Chen Huang, Zhizhao Guan, Wenqiang Lei and Yang Deng	35032
<i>Reasoning While Asking: Transforming Reasoning Large Language Models from Passive Solvers to Proactive Inquirers</i>	
Xin Chen, Feng Jiang, Yiqian Zhang, Hardy Chen, Shuo Yan, Wenya Xie, Min Yang and Shujian Huang	35069
<i>Utility-Oriented Visual Evidence Selection for Multimodal Retrieval-Augmented Generation</i>	
Weiqing Luo, Zongye Hu, Xiao Wang, Zhiyuan Yu, Haofeng Zhang and Ziyi Huang	35091
<i>Mitigating Selection Bias in Large Language Models via Permutation-Aware GRPO</i>	
Jinquan Zheng, Jia Yuan, Jiacheng Yao, Chenyang Gu, Pujun Zheng and Guoxiu He	35125
<i>KG-ViP: Bridging Knowledge Grounding and Visual Perception in Multi-modal LLMs for Visual Question Answering</i>	
Zhiyang Li, Ao Ke, Yukun Cao and Xike Xie	35144
<i>Reinforcement Learning–Guided Adaptive Tuning for Out-of-Distribution Harmful Text Detection</i>	
Mengyu Xiang, Tinghao Chen, Boxu Han, Qiudan Li, Shu Wu and Daniel Dajun Zeng ...	35158
<i>SOAR: Supervision from Observation for Agentic Reinforcement Learning</i>	
Meng Li, Lei Li, Xiting Wang, Yi Yuan, Zheng Wei, Brucebian and Zang Li	35175
<i>Business as Rulesual: A Benchmark and Framework for Business Rule Flow Modeling with LLMs</i>	
Chen Yang, Ruping Xu, Ruizhe Li, Bin Cao and Jing Fan	35198
<i>Specializing Large Models for Oracle Bone Script Interpretation via Component-Grounded Multimodal Knowledge Augmentation</i>	
Jianing Zhang, Runan Li, Honglin Pang, Ding Xia, Zhou Zhu, Qian Zhang, Chuntao Li and Xi Yang	35221
<i>HyperMem: Hypergraph Memory for Long-Term Conversations</i>	
Juwei Yue, Chuanrui Hu, Jiawei Sheng, Zuyi Zhou, Wenyan Zhang, Tingwen Liu, Li Guo and Yafeng Deng	35237
<i>LeLoRA: Learnable Low-Rank Adaptation of Large Language Models</i>	
Xiaoling Zhou, Mingjie Zhang, Zhemg Lee, Wei Ye and Shikun Zhang	35255
<i>Gated Tree Cross-Attention for Checkpoint-Compatible Syntax Injection in Decoder-Only LLMs</i>	
Xinyu Gao, Shaonan Wang and Nai Ding	35274
<i>DiningBench: A Hierarchical Multi-view Benchmark for Perception and Reasoning in the Dietary Domain</i>	
Song Jin, Juntian Zhang, Xun Zhang, Zeying Tian, Fei Jiang, Guojun Yin, Wei Lin, Yong Liu and Rui Yan	35289

<i>Octopus: Gated Selective Attention for Memory-Bounded Long-Context Inference in Large Language Models</i>	
Chien Van Nguyen, Ryan A. Rossi, Linh Ngo Van, Franck Dernoncourt and Thien Huu Nguyen	
	35311
<i>SAME: Safety-Aware Model Editing Guided by Safety Transformation</i>	
Jiayi Wang, Shipeng Wang, Ji Wu and Jian Sun	35324
<i>Where the Cat Sat: A Multilingual Framework for Spatial Language Understanding</i>	
Demian Inostroza, Ekaterina Vylomova, Charles Kemp, Mae Carroll, Wanchun Li and Meladel Mistica	35344
<i>LCMA-SRT: Language-Conditional Mixture-of-Experts Adapters for Joint Multilingual Speech Recognition and Translation</i>	
Nanjie Li, Xiaoyong Guo, Hao Huang, Xu Haihua and Wei Shi	35363
<i>DefGen-Bench: A Benchmark for Chinese Criminal Defence Opinion Generation in LegalAI</i>	
Senbo Zhang, Qiqi Wang, Fanghao Lou, Guanyu Chen, Yihong Pan, Huijia Li and Qian Liu	35378
<i>Escaping the Echo Trap: On Credit Assignment Failure in Multi-turn LLM Self-Reflection</i>	
Linxuan Du, Guangquan Xue, Xiaobo Liang, Qipeng Huang, Yuyang Ding, Xinyu Shi, Zhang Yijun, Ji Qi, Wenpeng Zhu, Juntao Li and Min Zhang	35393
<i>ARF-RLHF: Adaptive Reward-Following for RLHF through Emotion-Driven Self-Supervision and Trace-Biased Dynamic Optimization</i>	
Yuxuan Zhang	35406
<i>SMART: Evaluating LLMs' Mathematical Reasoning via a Human Cognitive Process-Inspired Benchmark</i>	
Yujie Hou, Mei Wang, Yaoyao Zhong, Ting Zhang, Xuetao Ma and Hua Huang	35426
<i>Decentralized Arena: Towards Democratic and Scalable Automatic Evaluation of Language Models</i>	
Yanbin Yin, Kun Zhou, Zhen Wang, Xiangdong Zhang, Yifei Shao, Shibo Hao, Yi Gu, Jieyuan Liu, Somanshu Singla, Tianyang Liu, Eric P. Xing, Zhengzhong Liu, Haojian Jin and Zhiting Hu	35453
<i>LBLLM: Lightweight Binarization of Large Language Models via Three-Stage Distillation</i>	
Siqing Song, Chuang Wang, Yong Lang, Yi Yang and Xu-Yao Zhang	35470
<i>SeDev: Structured Semantic Exploration for LLM-Driven Code Generation</i>	
Ronghui Yang, Jie Liu, Jiajie Zeng, Jiexin Wang, Jiuchuan Jiang, Bo An, Yi Cai and Mengchen Zhao	35485
<i>EVM-QuestBench: An Execution-Grounded Benchmark for Natural-Language Transaction Code Generation</i>	
Pei Yang, Wanyi Chen, Ke Wang, Lynn Ai, Eric Yang and Tianyu Shi	35513
<i>TIGER: Text-Informed Generalized Enzyme-Reaction Retrieval</i>	
Yuhang Zhang, Keyan Ding, Peilin Chen, Han Liu, Can Lin, Ruixi Chen, Shiqi Wang and Qi Song	35530
<i>MUSEG: Reinforcing Video Temporal Understanding via Timestamp-Aware Multi-Segment Grounding</i>	
Fuwen Luo, Shengfeng Lou, Chi Chen, Ziyue Wang, Chenliang Li, Weizhou Shen, Jiyue Guo, Peng Li, Ming Yan, Ji Zhang, Fei Huang and Yang Liu	35549
<i>User Perceptions vs. Proxy LLM Judges: Privacy and Helpfulness in LLM Responses to Privacy-Sensitive Scenarios</i>	
Xiaoyuan Wu, Roshni Kaushik, Wenkai Li, Lujo Bauer and Koichi Onoue	35562

<i>HARPO: Hierarchical Agentic Reasoning for User-Aligned Conversational Recommendation</i> Subham Raj, Aman Vaibhav Jha, Mayank Anand and Sriparna Saha	35580
<i>ShopSimulator: Evaluating and Exploring RL-Driven LLM Agent for Shopping Assistants</i> Pei Wang, Yanan Wu, Xiaoshuai Song, Weixun Wang, Gengru Chen, Zhongwen Li, Kezhong Yan, Qi Liu, Ken Deng, Shuaibing Zhao, Shaopan Xiong, Xuepeng Liu, Xuefeng Chen, Wanxi Deng, Wenbo Su and Bo Zheng	35600
<i>NL \Rightarrow Schedule: Evaluate Multitask Scheduling Capability of Large Language Models</i> Wenrui Liao, Weihong Du, Yi Li, Hongru Liang and Wenqiang Lei	35620
<i>SLICEFORMER: Static Program Slicing Using Language Models With Dataflow-Aware Pretraining and Constrained Decoding</i> Pengfei He, Shaowei Wang, Tse-Hsun Chen and Muhammad Asaduzzaman	35641
<i>SCOPE: Boosting LLM Efficiency with Scoped Position Encoding</i> Qingguo Qi, Hongyang Chen and Zhao Li	35653
<i>Debate-of-Thoughts: Resolving Knowledge Conflicts in LLMs Through Internal Deliberation</i> Guocong Li, Qirui Hu, Ping Wang, Guofeng Zhang, Jian Wu and Hongxia Xu	35674
<i>Projecting Out the Malice: A Global Subspace Approach to LLM Detoxification</i> Zenghao Duan, Zhiyi Yin, Zhichao Shi, Liang Pang, Shaoling Jing, Zihe Huang, Jiayi Wu, Yu Yan, Jingcheng Deng, Huawei Shen and Xueqi Cheng	35697
<i>From Local to Global: Revisiting Structured Pruning Paradigms for Large Language Models</i> Ziyan Wang, Enmao Diao, Qi Le, Pu Wang, Minwoo Lee, Shu-ping Yeh, Evgeny Stupachenko, Hao Feng and Li Yang	35720
<i>Audio MultiChallenge: A Multi-Turn Evaluation of Spoken Dialogue Systems on Natural Human Interaction</i> Advait Gosai, Tyler Vuong, Utkarsh Tyagi, Steven Li, Wenjia You, Miheer Bavare, Arda Uçar, Zhongwang Fang, Brian Jang, Bing Liu and Yunzhong He	35740
<i>GenomeQA: Benchmarking General Large Language Models for Genome Sequence Understanding</i> Weicai Long, Yusen Hou, Junning Feng, Houcheng su, Shuo Yang, Donglin Xie and Yanlin Zhang	35771
<i>VideoStir: Understanding Long Videos via Spatio-Temporally Structured and Intent-Aware RAG</i> Honghao Fu, Miao Xu, Yiwei Wang, Dailing Zhang, Jun Liu and Yujun Cai	35793
<i>Massively Multilingual Joint Segmentation and Glossing</i> Michael Ginn, Lindia Tjuatja, Enora Rice, Ali Marashian, Maria Valentini, Jasmine Xu, Graham Neubig and Alexis Palmer	35806
<i>TexOCR: Advancing Document OCR Models for Compilable Page-to-LaTeX Reconstruction</i> Chengye Wang, Lin Fu, Zexi Kuang and Yilun Zhao	35821
<i>If an LLM Were a Character, Would It Know Its Own Story? Evaluating Lifelong Learning in LLMs</i> Siqi Fan, Xiusheng Huang, Yiqun Yao, Xuezhi Fang, Kang Liu, Peng Han, Shuo Shang, Aixin Sun and Yequan Wang	35846
<i>Graph Reasoning Paradigm: Structured and Symbolic Reasoning with Topology-Aware Reinforcement Learning for Large Language Models</i> Runxuan Liu, Xianhao Ou, Xinyan Ma, Jiyuan Wang, Jiafeng Liang, Jiaqi Li, Tao He, Zheng Chu, Rongchuan Mu, Zekun Wang, Baoxin Wang, Dayong Wu, Ming Liu, Shijin Wang, Guoping Hu and Bing Qin	35859

<i>LLM-Powered Benchmark Factory: Reliable, Generic, and Efficient</i>	
Peiwen Yuan, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Yueqi Zhang, Jiayi Shi, Chuyi Tan, Boyuan Pan, Yao Hu and Kan Li	35882
<i>COSMOS: Connectivity-Oriented Submodular Maximization for Optimal Subgraph Retrieval</i>	
Boci Peng, Xiao Liu, Boren Hu, Yun Zhu, Xuanbo Fan, Yanwei Yue, Chunyu Yang and Yan Zhang	35904
<i>Attn-GS: Attention-Guided Context Compression for Efficient Personalized LLMs</i>	
Shenglai Zeng, Tianqi Zheng, Chuan Tian, Dante Everaert, Yau-Shian Wang, Yupin Huang, Michael J. Morais, Rohit Patki, Jinjin Tian, Xinnan Dai, Kai Guo, Monica Xiao Cheng and Hui Liu	35925
<i>Structured Semantic Information Helps Retrieve Better Examples for In-Context Learning Applied to Few-Shot Relation Extraction</i>	
Aunabil Chakma, Mihai Surdeanu and Eduardo Blanco	35947
<i>DUD: Decoupled Update Dynamics for Reliable Uncertainty Quantification in Large Language Models</i>	
Yixin Bu, Runze Xia, Guanyun Zou, Yupeng Ji, Hongliang Dai, Haodong Liu and Piji Li .	35972
<i>JurisBench: A Deep Benchmark for Assessing Large Language Models in Professional Legal Practice</i>	
Ziang Chen, Guannan Li, Fanlin Ji, Yipeng Kang, Jiaqi Li, Muhan Zhang, Yangtao Zhang, Li Tianjiao, Jiannan Wang, Xin Guo, Song-Chun Zhu and Bin Ling	35994
<i>LeCoDe: A Benchmark Dataset for Interactive Legal Consultation Dialogue Evaluation</i>	
Weikang Yuan, Kaisong Song, Zhuoren Jiang, Junjie Cao, Yujie Zhang, Jun Lin, Kun Kuang, Ji Zhang and Xiaozhong Liu	36019
<i>METER: Evaluating Multi-Level Contextual Causal Reasoning in Large Language Models</i>	
Pengfeng Li, Chen Huang, Chaoqun Hao, Hongyao Chen, Xiao-Yong Wei, Wenqiang Lei and See-Kiong Ng	36049
<i>PersonalAlign: Hierarchical Implicit Intent Alignment for Personalized GUI Agent with Long-Term User-Centric Records</i>	
Yibo Lyu, Gongwei Chen, Rui Shao, Weili Guan and Liqiang Nie	36074
<i>Empowering GUI Agents via Autonomous Experience Exploration and Hindsight Experience Utilization for Task Planning</i>	
Tianyi Men, Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu and Jun Zhao	36090
<i>Excuse me, may I say something...CoLabScience, A Proactive AI Assistant for Biomedical Discovery and LLM-Expert Collaborations</i>	
Yang Wu, Jinhong Yu, Jingwei Xiong, Zhimin Tao and Xiaozhong Liu	36109
<i>Knowledge-driven Augmentation and Retrieval for Integrative Temporal Adaptation</i>	
Weisi Liu, Guangzeng Han and Xiaolei Huang	36130
<i>SAD: A Large-Scale Strategic Argumentative Dialogue Dataset</i>	
YongKang Liu, Jiayang Yu, Mingyang Wang, Yiqun Zhang, Ercong Nie, Shi Feng, Daling Wang, Kaisong Song and Hinrich Schuetze	36144
<i>PICTURE: Enhancing Theory-of-Mind in Large Language Models by Revealing, Not Hiding, Characters' Lack of Knowledge</i>	
Eojin Jeon and SangKeun Lee	36165
<i>Union-of-Experts: Neurons in Mixture-of-Experts are Secretly Routers</i>	
Songhao Wu, Ang Lv, Ruobing Xie, Samm Sun, Di Wang, Rui Yan and Yankai Lin	36193

<i>ReflectRM: Boosting Generative Reward Models via Self-Reflection within a Unified Judgment Framework</i>	
Kai Qin, Liangxin Liu, Yu Liang, Longzheng Wang, Wangyan, Zhang Yueyang, Long Xia, Zhiyuan Sun, Houde Liu and Daiting Shi	36207
<i>SQL-Trail: Multi-Turn Reinforcement Learning with Interleaved Feedback for Text-to-SQL</i>	
Harper Hua, Zhen Han, Zhengyuan Shen, Meng-Chieh Lee, Sheng Guan, Qi Zhu, Sullam Jeoung, Yueyan Chen, Yunfei Bai, Shuai Wang, Vassilis N. Ioannidis and Huzefa Rangwala	36224
<i>Bootstrapping Code Translation with Weighted Multilanguage Exploration</i>	
Yuhan Wu, Huan Zhang, Wei Cheng, Chen Shen, Jingyue Yang and Wei Hu	36247
<i>Flow-Based Page Unique Semantic Mapping Architecture for Document Visual Question Answering</i>	
Haosen Wang, Jing Xiao, Chaochao Du, Xiaowang Zhang and Zhiyong Feng	36260
<i>Pub-LawBench: Public-Oriented Benchmarking for LegalAI</i>	
Qiaoyu Zheng, Zehan Ma, Yijing Zhang, Qiqi Wang, Huijia Li and Qian Liu	36281
<i>SWAN: Semantic Watermarking with Abstract Meaning Representation</i>	
Ziping Ye, Gourab Dey, Christos Christodoulopoulos, Charith Peris, Anil Ramakrishna, Weitong Ruan, Aram Galstyan, Kai-Wei Chang, Rahul Gupta and Ninareh Mehrabi	36304
<i>Turning Failures into Value: Negative Experience Replay for RLVR via Confidence Gating and Boundary Failure Sampling</i>	
Jialiang Guo, Fucheng Xiong, Xu He, Haodong Zhao, Xingyang li, Ke Zeng and Xunliang Cai	36316
<i>LazyEviction: Lagged KV Eviction with Attention Pattern Observation for Efficient Long Reasoning</i>	
Haoyue Zhang, Hualei Zhang, Xiaosong Ma, Jie Zhang and Song Guo	36335
<i>Diversity in Unity, Theory in Practice: Hierarchical Multitask Benchmarks for Chinese Minority Languages</i>	
Yijie Li, Xi Cao, Yuan Sun, Quulgan Minggad, Abdulla Ablikim and Jia Qing Cai Wang . .	36353
<i>LangSAE Editing: Improving Multilingual Information Retrieval via Post-hoc Language Identity Removal</i>	
Dongjun Kim, Jeongho Yoon, Chanjun Park and Heuseok Lim	36374
<i>Interpretable Traces, Unexpected Outcomes: Investigating the Disconnect in Trace-Based Knowledge Distillation</i>	
Siddhant Bhambri, Upasana Biswas and Subbarao Kambhampati	36390
<i>SpecMind: Cognitively Inspired, Interactive Multi-Turn Framework for Postcondition Inference</i>	
Cuong Chi Le, Minh V.t. Pham, Tung D. Vu, Van Duc Cuong, Phan Nhat Huy, Phan Nhat Hoang and Tien N. Nguyen	36409
<i>Memorization, Emergence, and Explaining Reversal Failures: A Controlled Study of Relational Semantics in LLMs</i>	
Yihua Zhu, Qianying Liu, Jiaxin Wang, Fei Cheng, Chaoran Liu, Akiko Aizawa, Sadao Kurohashi and Hidetoshi Shimodaira	36425
<i>Framing Political Bias in Multilingual LLMs Across Pakistani Languages</i>	
Afrozah Nadeem, Mark Dras and Usman Naseem	36444
<i>InsAT: Instance-aware Semantic Alignment and Transfer from Human–Object Keypoints for Zero-to-Few-shot Action Understanding</i>	
Kazuki Tsutsukawa	36487

<i>Multi-View Attention Multiple-Instance Learning Enhanced by LLM Reasoning for Cognitive Distortion Detection</i>	Jun Seo Kim, Hyemi Kim, Woo Joo OH, Hongjin Cho, Hochul Lee and Hye Hyeon Kim .	36505
<i>CHAIRO: Contextual Hierarchical Analogical Induction and Reasoning Optimization for LLMs</i>	Haotian Lu, Yuchen Mou and Bingzhe Wu	36526
<i>Test of Time: Rethinking Temporal Signal of Benchmark Contamination</i>	Terry Jingchen Zhang, Gopal Dev, Ning Wang, Max Obreiter, Wenyuan Jiang, Punya Syon Pandey, Keenan Samway, Yinya Huang, Bernhard Schölkopf, Mrinmaya Sachan and Zhijing Jin	36538
<i>An Existence Proof for Neural Language Models That Can Explain Garden-Path Effects via Surprisal</i>	Ryo Yoshida, Shinnosuke Isono, Taiga Someya, Yohei Oseki and Tatsuki Kuribayashi	36556
<i>How Context Shapes Truth: Geometric Transformations of Statement-level Truth Representations in LLMs</i>	Shivam Adarsh, Maria Maistro and Christina Lioma	36570
<i>Mitigating Tokenization-Induced Distance Distortion in Long-Context Multilingual Machine Translation</i>	Khotso Selialia, Antoine Nzeyimana and Fatima M. Anwar	36591
<i>Probing Audio-Visual Reasoning in Multimodal Language Models through the Lens of Audio</i>	Kaixiong Gong, Kaituo Feng, Bohao Li, Yibing Wang, Mofan Cheng, Shijia Yang, Jiaming Han, Benyou Wang, Yutong Bai, Zhuoran Yang and Xiangyu Yue	36603
<i>Reasoning Hijacking: The Fragility of Reasoning Alignment in Large Language Models</i>	Yuansen Liu, Yixuan Tang and Anthony Kum Hoe Tung	36646
<i>The Dominance of Text Space: Unveiling the Asymmetric Nature of Cross-Modal Alignment in Large Language Models</i>	Linqing Chen, Hanmeng Zhong, Wentao Wu and Peng Zhou	36666
<i>Towards Explainable Diagnosis: A Self-learned Explanatory Knowledge Base Approach</i>	Dongqi Huang, Tong Zhou, Zhuoran Jin, Shenghui Shi, Maoyujiao, Kang Liu, Jun Zhao and Yubo Chen	36679
<i>WildFeedback: Aligning LLMs With In-situ User Interactions And Feedback</i>	Taiwei Shi, Zhuoer Wang, Longqi Yang, Ying-Chun Lin, Zexue He, Mengting Wan, Pei Zhou, Sujay Kumar Jauhar, Sihao Chen, Shan Xia, Hongfei Zhang, Jieyu Zhao, Xiaofeng Xu, Xia Song and Jennifer Neville	36701
<i>SSG: Logit-Balanced Vocabulary Partitioning for LLM Watermarking</i>	Chenxi Gu, Xiaoning Du and John C. Grundy	36726
<i>PaT: Planning-after-Trial for Efficient Test-Time Code Generation</i>	Youngsik Yoon, Sungjae Lee, Seockbean Song, Siwei Wang, Wei Chen and Jungseul Ok .	36738
<i>Joint Knowledge Base Completion and Question Answering by Combining Large Language Models and Small Language Models</i>	Yinan Liu, Dongying Lin, Sigang Luo, Xiaochun Yang and Bin Wang	36756
<i>Rethinking Reasoning-Intensive Retrieval: Evaluating and Advancing Retrievers in Agentic Search Systems</i>	Yilun Zhao, Jinbiao Wei, Tingyu Song, Siyue Zhang, Chen Zhao and Arman Cohan	36776

<i>The Role of Mixed-Language Documents for Multilingual Large Language Model Pretraining</i> Jiandong Shao, Raphael Tang, Crystina Zhang, Karin Sevegnani, Pontus Stenetorp, Jianfei Yang and Yao Lu	36807
<i>ARCHITECT: Uncertainty-Aware Dynamic Tool Learning via Causal Intervention for Open-World Agents</i> Zhangyi Wang, Jiexiang Xu, Bingnan Yu and Zongze Li	36819
<i>ES4R: Speech Encoding Based on Prepositive Affective Modeling for Empathetic Response Generation</i> Zhuoyue Gao, Xiaohui Wang, Xiaocui Yang, Wen Zhang, Daling Wang, Shi Feng and Yifei Zhang	36830
<i>MAGMA: A Multi-Graph based Agentic Memory Architecture for AI Agents</i> Dongming Jiang, Yi Li, Guanpeng Li and Bingzhe Li	36848
<i>The Side Effects of Being Smart: Safety Risks in MLLMs' Multi-Image Reasoning</i> Renmiao Chen, Yida Lu, Shiyao Cui, Xuan Ouyang, Victor Shea-Jay Huang, Shumin Zhang, Chengwei Pan, Han Qiu and Minlie Huang	36866
<i>Action Boundary Blindness: When LLM Agents Cannot Tell Where One Action Ends and Another Begins</i> Zhangyi Wang, Bingnan Yu, Jiexiang Xu and Zongze Li	36883
<i>Verbal-R3: Verbal Reranker as the Missing Bridge between Retrieval and Reasoning</i> Sangkwon Park, Donghun Kang, Jisoo Mok and Sungroh Yoon	36900
<i>IGenBench: Benchmarking the Reliability of Text-to-Infographic Generation</i> Yinghao Tang, Xueding Liu, Boyuan Zhang, Tingfeng Lan, Yupeng Xie, Jiale Lao, Yiyao Wang, Haoxuan Li, Tingting Gao, Bo Pan, Luoxuan Weng, Xiuqi Huang, Minfeng Zhu, Yingchaojie Feng, Yuyu Luo and Wei Chen	36919
<i>ProMedical: Hierarchical Fine-Grained Criteria Modeling for Medical LLM Alignment via Explicit Injection</i> He Geng, Yangmin Huang, Lixian Lai, Qianyun Du, Hui Chu, Zhiyang He, Jiaxue Hu and Xiaodong Tao	36955
<i>R2IF: Aligning Reasoning with Decisions via Composite Rewards for Interpretable LLM Function Calling</i> Aijia Cheng, Kailong Wang, Ling Shi and Yongxin Zhao	36995
<i>From Detection to Understanding: Multi-Turn Reasoning for Video Misinformation Analysis</i> Zhi Zeng, Jiaying Wu, Minnan Luo, Di Zhang, Yifei Yang, Xiangzheng Kong, Herun Wan and Zihan Ma	37009
<i>Frozen LLMs are Native Decoders for High-Norm Semantic Vectors</i> Yunsheng Zeng and Yongmei Tan	37028
<i>Enhancing Reinforcement Learning for Radiology Report Generation with Evidence-aware Rewards and Self-correcting Preference Learning</i> Qin Zhou, Guoyan Liang, Qianyi Yang, Jingyuan Chen, Sai Wu, Chang Yao and Zhe Wang	37044
<i>VisAidMath: Benchmarking Visual-Aided Mathematical Reasoning</i> Jingkun Ma, Runzhe Zhan, Yang Li, Di Sun, Hou Pong Chan, Lidia S. Chao and Derek F. Wong	37057

<i>RISK: A Framework for GUI Agents in E-commerce Risk Management</i>	
Renqi Chen, Zeyin Tao, Jianming Guo, Jingzhe Zhu, Yiheng Peng, Qingqing Sun, Tianyi Zhang and Shuai Chen	37104
<i>MiniRAG: A Lightweight RAG system with Small Language Models</i>	
Tianyu Fan, Jingyuan Wang, Xubin Ren and Chao Huang	37121
<i>Glyph: Scaling Context Windows via Visual-Text Compression</i>	
Jiale Cheng, Yusen Liu, Xinyu Zhang, Yulin Fei, Wenyi Hong, Ruiliang Lyu, Weihang Wang, Zhe Su, Xiaotao Gu, Xiao Liu, Yushi Bai, Jie Tang, Hongning Wang and Minlie Huang	37145
<i>TA-GRPO-d: Trajectory-Aware GRPO for Optimizing Denoising Trajectories in Diffusion LLMs</i>	
Gyanyeop Kim and Sangwoo Kang	37159
<i>AG-GRPO: Answer-Guided GRPO for Masked Diffusion Language Models</i>	
Juhyeong Kim, Gyanyeop Kim and Sangwoo Kang	37175
<i>Grammar as Control: Modular Language Generation for the Long Tail</i>	
Ndapa Nakashole	37192
<i>MTR-Suite: A Framework for Evaluating and Synthesizing Conversational Retrieval Benchmarks</i>	
Junhao Ruan, Abudukeymu Abudula, Bei Li, Yongjing Yin, Xinyu Liu, Kechen Jiao, Xin Chen, Jingang Wang, Xunliang Cai, Tong Xiao and JingBo Zhu	37223
<i>Beyond Majority Voting: Towards Fine-grained and More Reliable Reward Signal for Test-Time Reinforcement Learning</i>	
Weiqin Wang, Yile Wang, Kehao Chen and Hui Huang	37251
<i>Beyond Markovian Forgetfulness: Episodic Memory for Reasoning-Intensive Retrieval</i>	
Dohyeon Lee, Yeonseok Jeong and Seung-won Hwang	37266
<i>DUAL RM: Beyond Rule-based Preference Reward Modeling via Meta-Reward</i>	
Xiaobo Liang, Wanfu Wang, Qipeng Huang, Yuyang Ding, Zecheng Tang, Yixin Ji, Qianben Chen, Zhe Zhao, Kehai Chen, Juntao Li and Min Zhang	37281
<i>Evolutionary Negative Module Pruning for Better LoRA Merging</i>	
Anda Cao, Zhuo Gou, Yi Wang, Kaixuan Chen, Yu Wang, Can Wang, Mingli Song and Jie Song	37297
<i>Example Quality Matters: Multi-Aspects Example Augmentation for Private Library Programming</i>	
Yuhao Li, Haifeng Sun, Xuesong Zhang, Shu Yao, Haoyu Zheng, Yvchuan Wang, Huazheng Wang, Zirui Zhuang, Qi Qi, Jianxin Liao and Jingyu Wang	37311
<i>Resonating with RoPE: Spectral Quantization for High-Fidelity Key Cache Compression</i>	
Xuefei Wang, Haoyu Tang, Tianyuan Liang, Zhibin Wang, Yupeng Hu and Weili Guan ...	37328
<i>CSRP: Chain-of-Thought Reasoning for Chinese Text Correction via Reinforcement Learning with Efficiency-Aware Rewards</i>	
Wei Tian, Yuhao Zhou and Man Lan	37349
<i>Adaptive Retrieval for Reasoning</i>	
Jongho Kim, Jaeyoung Kim, Jihyuk Kim, Yu Jin Kim, Seung-won Hwang and Moontae Lee	37366
<i>On Emergent Social World Models — Evidence for Functional Integration of Theory of Mind and Pragmatic Reasoning in Language Models</i>	
Polina Tsvilodub, Jan-Felix Klumpp, Amir Pour, Jennifer Hu and Michael Franke	37382

<i>ToolOmni: Enabling Open-World Tool Use via Agentic learning with Proactive Retrieval and Grounded Execution</i>	Shouzheng Huang, Meishan Zhang, Baotian Hu and Min Zhang	37421
<i>From TDMA to CDMA: A Multi-bit Watermark for Diffusion Language Models</i>	Baizhou Huang and Xiaojun Wan	37440
<i>LegalGraphRAG: Multi-Agent Graph Retrieval-Augmented Generation for Reliable Legal Reasoning</i>	Zerui Chen, Qinggang Zhang, Zhishang Xiang, Zhimin Wei, Linfeng Gao, Xiao Huang, Zhihong Zhang and Jinsong Su	37455
<i>Characterizing the Expressivity of Local Attention in Transformers</i>	Jiaoda Li and Ryan Cotterell	37485
<i>TH-RAG : Topic-Based Hierarchical Knowledge Graphs for Robust Multi-hop Reasoning in Graph-based RAG Systems</i>	JungHyoun Kim, Soohyeong Kim, Seok Jun Hwang, Jeonghyeon Park and Yong Suk Choi	37508
<i>Beyond Static Benchmarks: Synthesizing Harmful Content via Persona-based Simulation for Robust Evaluation</i>	Huije Lee, Jisu Shin, Hoyun Song, Changgeon Ko and Jong C. Park	37532
<i>DiVE: Decoupling Intra-layer Visual Evidence for Mitigating Hallucinations in Large Vision-Language Models</i>	Xinwei Li, Li Lin, Hui Jiao, Li Yao, Tien-Tsin Wong and Hanqian Wu	37552
<i>Impatient Users Confuse AI Agents: High-fidelity Simulations of Human Traits for Testing Agents</i>	Muyu He, Anand Kumar, Soumyadeep Bakshi, James Zou and Nazneen Rajani	37569
<i>Efficient Process Reward Modeling via Contrastive Mutual Information</i>	Nakyung Lee, Sangwoo Hong and Jungwoo Lee	37603
<i>R³: End-to-End Reasoning-based Planning for Multi-step Retrosynthesis via Reinforcement Learning</i>	YiFei Wang, Qizhi Pei, Jiangtao Feng, Yuntian Shi, Yi Duan, Lihao Wang, Lei Bai, Lijun Wu, Wei-Ying Ma and Hao Zhou	37618
<i>River-LLM: Large Language Model Seamless Exit Based on KV Share</i>	Yingtao Shen and An Zou	37633
<i>Execution as Verification: Fine-Grained Self-Correcting Reasoning for Complex KBQA</i>	Minghan Zhang, Zhen Yang, Haodong Zou, Jie Chen, Zhen Duan and Shu Zhao	37646
<i>Zero-Shot Detection of LLM-Generated Text using Temperature Sensitivity</i>	Shixuan Ma, Jiahao Li, Zhendong Mao and Quan Wang	37664
<i>MARD: Module-Aware Reasoning Distillation for Language Models with Adaptive Supervision</i>	Wenqi Yang, Jianjun Li, Zhibo Zhang, Mingqian Ding and Yushen Fang	37680
<i>Dashboard2Code: Evaluating Multimodal Models on Reconstructing Interactive Dashboards</i>	Tianhao Niu, Ziyu Han, Qiguang Chen, Shiqi Zhou, Baocai Shan, Hengjie Fang, Qingfu Zhu and Wanxiang Che	37696
<i>Can Compact Language Models Search Like Agents? Distillation-Guided Policy Optimization for Preserving Agentic RAG Capabilities</i>	Rikuto Kotoge, Mai Nishimura and Jiaxin Ma	37733

<i>CE-GPPO: Coordinating Entropy via Gradient-Preserving Clipping Policy Optimization in Reinforcement Learning</i>	
Zhenpeng Su, Leiyu Pan, Minxuan Lv, Yuntao Li, Wenping Hu, Fuzheng Zhang, Kun Gai and Guorui Zhou	37747
<i>Omni-I2C: A Holistic Benchmark for High-Fidelity Image-to-Code Generation</i>	
Jiawei Zhou, Chi Zhang, Xiang Feng, Qiming Zhang, Haibo Qiu, Lihuo He, Dengpan Ye, Xinbo Gao and Jing Zhang	37761
<i>Revealing Procedural Reasoning Structures in Chain-of-Thought Training via Span-Level Gradient Organization</i>	
Jia Liu, Jiaxin Luo, Weiwen Xu, Jonathan M. Garibaldi, Xiao-Kun Wu, Yixue Hao and Min Chen	37799
<i>StepHint: Multi-level Stepwise Hints Enhance Reinforcement Learning to Reason</i>	
Kaiyi Zhang, Ang Lv, Jinpeng Li, Yongbo Wang, Feng Wang, Haoyuan hu and Rui Yan ..	37846
<i>Social Dynamics as Critical Vulnerabilities that Undermine Objective Decision-Making in LLM Collectives</i>	
Changgeon Ko, Jisu Shin, Hoyun Song, Huije Lee, Eui Jun Hwang and Jong C. Park	37865
<i>FairQE: Multi-Agent Framework for Mitigating Gender Bias in Translation Quality Estimation</i>	
Jinhee Jang, Juhwan Choi, Dongjin Lee, Seunguk Yu and YoungBin Kim	37891
<i>AdaDPI: Document-level Translation Adaptive Agent via Dynamic Parametric Internalization</i>	
Hong Ren, Liting Deng, Shaolin Zhu and Deyi Xiong	37912
<i>Mitigating Over-Refusal in Aligned Large Language Models via Inference-Time Activation Energy</i>	
Eric Hanchen Jiang, Weixuan Ou, Run Liu, Shengyuan Pang, Guancheng Wan, Ranjie Duan, Wei Dong, Kai-Wei Chang, XiaoFeng Wang, Ying Nian Wu and Xinfeng Li	37930
<i>Refining and Reusing Annotation Guidelines for LLM Annotation</i>	
Kon Woo Kim, Jin-Dong Kim and Akiko Aizawa	37951
<i>Cross-Examination Framework: A Task-Agnostic Diagnostic for Information Fidelity in Text-to-Text Generation</i>	
Tathagata Raha, Clement Christophe, Nada Saadi, Hamza A Javed, Marco AF Pimentel, Ronnie Rajan and Praveenkumar Kanithi	37965
<i>AlignCultura: Towards Culturally Aligned Large Language Models?</i>	
Gautam Siddharth Kashyap, Mark Dras and Usman Naseem	37986
<i>UPDESH: Synthesizing Grounded Instruction Tuning Data for 13 Indic Languages</i>	
Pranjal A Chitale, Varun Gumma, Sanchit Ahuja, Prashant Kodali, Manan Uppadhyay, Deepthi Sudharsan and Sunayana Sitaram	37997
<i>Dynamic Generation of Multi LLM Agents Communication Topologies with Graph Diffusion Models</i>	
Eric Hanchen Jiang, Levina Li, Frank Wan, Xiao Liang, Sophia Yin, Yuchen Wu, Xinfeng Li, Yizhou Sun, Wei Wang, Kai-Wei Chang and Ying Nian Wu	38042
<i>TaxPraBen: A Scalable Benchmark for Structured Evaluation of LLMs in Chinese Real-World Tax Practice</i>	
Gang Hu, Yating Chen, Haiyan Ding, Wang Gao, Huang Jiajia, Min Peng, Qianqian Xie and Kun Yue	38061

<i>Think How to Think: Mitigating Overthinking with Autonomous Difficulty Cognition in Large Reasoning Models</i>	
Yongjiang Liu, Haoxi Li, Xiaosong Ma, Jie Zhang and Song Guo	38105
<i>Theory-optimal Quantization Based on Flatness</i>	
Xiusheng Huang, Zhe Li, Xuanwu Yin, Lu Wang, Yequan Wang, Dong Li, Emad Barsoum and Kang Liu	38127
<i>When Background Matters: Breaking Medical Vision Language Models by Transferable Attack</i>	
Akash Ghosh, Subhadip Baidya, Sriparna Saha and Xiuying Chen	38143
<i>Tears or Cheers? Benchmarking LLMs via Culturally Elicited Distinct Affective Responses</i>	
Chongyuan Dai, Yaling Shen, Zihan Gao, Jia Li, Yishun Jiang, Yaxiong Wang, Liu Liu, Zongyuan Ge and Jinpeng Hu	38171
<i>Emergent Misalignment via In-Context Learning: Narrow in-context examples can produce broadly misaligned LLMs</i>	
Nikita Afonin, Nikita Andriianov, Vahagn Hovhannisyanyan, Nikhil Bageshpura, Kyle Liu, Kevin Zhu, Sunishchal Dev, Ashwinee Panda, Oleg Rogov, Elena Tutubalina, Alexander Panchenko and Mikhail Seleznyov	38197
<i>EvoRoute: Experience-Driven Self-Routing LLM Agent Systems</i>	
Guibin Zhang, Haiyang Yu, Kaiming Yang, Bingli Wu, Fei Huang, Yongbin Li and Shuicheng Yan	38213
<i>Foresight Optimization for Strategic Reasoning in Large Language Models</i>	
Jessie Wang, Jiawen Duan, Jian Wang, Kaitao Song, Chunpu Xu, Johnny K. W. Ho, YU Fenggang, Johan F. Hoorn and Wenjie Li	38226
<i>DetectRL-X: Towards Reliable Multilingual and Real-World LLM-Generated Text Detection</i>	
Junchao Wu, Yefeng Liu, Chenyu Zhu, Hao Zhang, Zeyu Wu, Tianqi Shi, Yichao Du, Longyue Wang, Weihua Luo, Jinsong Su and Derek F. Wong	38247
<i>ImmersiveTTS: Environment-Aware Text-to-Speech with Multimodal Diffusion Transformer and Domain-Specific Representation Alignment</i>	
Jun-Hak Yun, Seung-Bin Kim and Seong-Whan Lee	38295
<i>When Efficiency Becomes a Vulnerability: Computational Cost Attacks on WebAgents</i>	
Liang-Bo Ning, Yuchen Zhu, Heqing Huang, Xin Wang, Yi Chang, Li Qing and Wenqi Fan	38315
<i>Evolving Beyond Snapshots: Harmonizing Structure and Sequence via Entity State Tuning for Temporal Knowledge Graph Forecasting</i>	
Siyuan Li, Yunjia Wu, Yiyong Xiao, Pingyang Huang, Peize Li, Ruitong Liu, Yan Wen and Te Sun	38336
<i>CodeRipple: Wavelet-Based Detection of LLM-Generated Code</i>	
Xingyu Yao, Zhendong Mao and Quan Wang	38351
<i>PHOTON: Hierarchical Autoregressive Modeling for Lightspeed and Memory-Efficient Language Generation</i>	
Yuma Ichikawa, Naoya Takagi, Takumi Nakagawa, Yuzi Kanazawa and Akira Sakai	38365
<i>CoMeT: Collaborative Memory Transformer for Efficient Long Context Modeling</i>	
Runsong Zhao, Shilei Liu, Jiwei Tang, Langming Liu, Haibin Chen, Weidong Zhang, Yujin Yuan, Tong Xiao, JingBo Zhu, Wenbo Su and Bo Zheng	38384

<i>CheckRLM: Effective Knowledge–Thought Coherence Checking in Retrieval-Augmented Reasoning</i>	
Dingling Xu, Ruobing Wang, Qingfei Zhao, Yukun Yan, Zhichun Wang, Daren Zha, Shi Yu, Zhenghao Liu, Shuo Wang, Xu Han and Maosong Sun	38403
<i>CAMEC: Complexity-Aware Multi-Expert Collaboration for Reliable Chinese Medical Question Answering</i>	
Yukang Wu, Xiyuan Jia, Jiayi Wu, HongChen Yu, Yuhan Qiu and Wu Guohua	38427
<i>Vocabulary Hijacking in LVLMS: Unveiling Critical Attention Heads by Excluding Inert Tokens to Mitigate Hallucination</i>	
Yangneng Chen, Junlin Li, Weijun Yao, Xilai Ma, Guodong DU, Wenya Wang and Jing Li	38458
<i>HumanLLM: Benchmarking and Improving LLM Anthropomorphism via Human Cognitive Patterns</i>	
Xintao Wang, Jian Yang, Weiyuan Li, Rui Xie, Jen-tse Huang, Jun Gao, Shuai Huang, Yueping Kang, Yuanli Guo, Hongwei Feng and Yanghua Xiao	38486
<i>Vision-Language Introspection: Mitigating Overconfident Hallucinations in MLLMs via Interpretable Bi-Causal Steering</i>	
Shuliang Liu, Songbo Yang, Dong Fang, Sihang Jia, Yuqi Tang, Lingfeng Su, Ruoshui Peng, Yibo Yan, Xin Zou and Xuming Hu	38518
<i>v-HUB: A Benchmark for Video Humor Understanding from Vision and Sound</i>	
Zhengpeng Shi, Yanpeng Zhao, Jianqun Zhou, Yuxuan Wang, Qinrong Cui, Wei Bi, Song-Chun Zhu, Bo Zhao and Zilong Zheng	38544
<i>Beyond Single Representations: Multi-Model Embedding Fusion for Stable Text Classification</i>	
Jiho Gwak and Yuchul Jung	38568
<i>Confidence Should Be Calibrated More Than One Turn Deep</i>	
Zhaohan Zhang, Chengzhengxu Li, Xiaoming Liu, Chao Shen, Ziquan Liu and Ioannis Patras	38578
<i>NaviMaster: Learning a Unified Policy for GUI and Embodied Navigation Tasks</i>	
Zhihao Luo, Wentao Yan, Jingyu Gong, Min Wang, Zhizhong Zhang, Xuhong Wang, Yuan Xie and Xin Tan	38595
<i>Teaching LLMs Human-Like Editing of Inappropriate Argumentation via Reinforcement Learning</i>	
Timon Ziegenbein, Maja Stahl and Henning Wachsmuth	38616
<i>FLAIR: Steering LLM Mathematical Problem Solving based on A Fuzzy-Logic-Assisted Reasoner</i>	
Hao Wu, Hongru Sun, Wanqing Li, Xinguo Yu, Hao Ming, Xiao Luo, Wenbin Zhang, Jiahong Zhao, Yi Guo and Jie Yang	38638
<i>QuDAR: Query-Wise Dual-Perspective Adaptive Retrieval</i>	
Joeun Kim, Seunghyoun Yoon, Xuan-Bach Le, Youngeun Nam, Doyoung Kim, Hwanjun Song and Jae-Gil Lee	38662
<i>Trait Activation in Silicon: A Situation-Aware Framework for Psychologically Grounded Role-Playing</i>	
Zuolong Li, Pingyu Wu, Xianwen Huang, Tianyi Wei and Wenbo Zhou	38680
<i>SPARKLE: A Structured and Plug-and-play Agentic Retrieval Policy for Adaptive RAG Models</i>	
Jinyuan Fang, Zaiqiao Meng and Craig Macdonald	38699
<i>G-IdiomAlign: A Gloss-Pivoted Benchmark for Cross-Lingual Idiom Alignment</i>	
Fengying Ye, Yanming Sun, Runzhe Zhan, Lidia S. Chao, Zheqi Zhang and Derek F. Wong	38720

<i>SciCoQA: Quality Assurance for Scientific Paper–Code Alignment</i>	
Tim Baumgärtner and Iryna Gurevych	38740
<i>ImpRIF: Stronger Implicit Reasoning Leads to Better Complex Instruction Following</i>	
Yuan Cheng Yang, Lin Yang, Xu Wang, Chao Tong and Haihua Yang	38771
<i>HalluAudio: A Comprehensive Benchmark for Hallucination Detection in Large Audio-Language Models</i>	
Feiyu Zhao, Yiming Chen, Wenhuan Lu, Daipeng Zhang, Xianghu Yue and Jianguo Wei ..	38797
<i>ReTraceQA: Evaluating Reasoning Traces of Small Language Models in Commonsense Question Answering</i>	
Francesco Maria Molfese, Luca Moroni, Ciro Porcaro, Simone Conia and Roberto Navigli	38817
<i>BEFT: Bias-Efficient Fine-Tuning of Language Models in Low-Data Regimes</i>	
Baichuan Huang, Ananth Balashankar and Amir Aminifar	38833
<i>JoPR: Joint Emotion Perception and Reasoning for Conversational Emotion Recognition</i>	
Yumeng Fu, Weitao Huang, Junjie Wu, Hao Teng, Meishan Zhang and Bingquan Liu	38852
<i>Uncertainty-Aware Routing for Principled Alignment with MoE Dynamics</i>	
Yilong Chen, Junyuan Shang, Yuchen Feng, Zhenyu Zhang, Naibin Gu, Ziqi Wang, Tingwen Liu, Shuohuan Wang, Yu Sun, Hua Wu and Haifeng Wang	38865
<i>S⁴: Operationalizing Speech Act Theory for Strategic Semi-Structured Psychiatric Interview</i>	
Guanqun Bi, Zhoufu Liu, Zhuang Chen, Dazhen Wan, Xiyao Xiao and Minlie Huang	38881
<i>More Aligned, Less Diverse? Analyzing the Grammar and Lexicon of Two Generations of LLMs</i>	
Adrián Gude, Roi Santos-Rios, Francis Bond, Dan Flickinger, Carlos Gómez-Rodríguez and Olga Zamaraeva	38900
<i>MemCoRL: Alternating Co-Optimization of Memory Retrieval and Utilization via Collaborative Reinforcement Learning</i>	
Yuewen Liu, Peng Xu, Muxi Diao, Anyi Zhang, Yang Li and Yutong Zhang	38912
<i>Fast and Accurate Fisher-Guided Quantization via Efficient Kronecker Factorization</i>	
Viktoriia A. Chekalina, Gerasin Timofey, Andrey Kuznetsov and Evgeny Frolov	38925
<i>Discovering a Shared Logical Subspace: Steering LLM Logical Reasoning via Alignment of Natural-Language and Symbolic Views</i>	
Feihao Fang, My T. Thai and Yuanyuan Lei	38936
<i>One Refiner to Unlock Them All: Inference-Time Reasoning Elicitation via Reinforcement Query Refinement</i>	
Yixiao Zhou, Dongzhou Cheng, Zhiliang wu, Yi Yang, Yu Cheng and Hehe Fan	38957
<i>Why Does Reinforcement Learning Generalize? A Feature-Level Mechanistic Study of Post-Training in Large Language Models</i>	
Dan Shi, Zhuowen Han, Simon Ostermann, Renren Jin, Josef Van Genabith and Deyi Xiong	38979
<i>Polymorphic Universal Transformer</i>	
Yilong Chen, Zitian Gao, Yihao Xiao, Jason Klein Liu, Xinyu Yang, Yifan Luo, Haoming Luo, Zhengmao Ye, Tingwen Liu, Ran Tao and Bryan Dai	39001
<i>Can Reasoning Path still be Effective as Input? Bridging Post-Reasoning to Chain-of-Thought Compression</i>	
Chengzhengxu Li, Xiaoming Liu, Zhaohan Zhang, Shengchao Liu, Guoxin Ma, Yu Lan, Cong Wang and Chao Shen	39014

<i>REAL: REtrieval-reAsoning and Logic-constructed Attention Behaviors for Long-Context KV Cache Compression</i>	
Mengjie Li, Yuan Feng, Xike Xie and William J. Song	39035
<i>From Where Words Come: Efficient Regularization of Code Tokenizers Through Source Attribution</i>	
Pavel Chizhov, Egor Bogomolov and Ivan P. Yamshchikov	39053
<i>EmoS: A High-Fidelity Multimodal Benchmark for Fine-grained Streaming Emotional Understanding</i>	
Pengze Guo, Jingxi Liang, Zhiwen Xie, Qifeng Wang and Derek F. Wong	39074
<i>Surprisal Minimisation over Goal-directed Alternatives Predicts Production Choice in Dialogue</i>	
Thomas P. Utting, Mario Giulianelli and Arabella Sinclair	39090
<i>HowToNarrate: A General-Domain Benchmark for Synchronized Video Narration with External Knowledge</i>	
Xueyan Wang, Dingyi Yang and Qin Jin	39110
<i>Putting HUMANS first: Efficient LAM Evaluation with Human Preference Alignment</i>	
Woody Haosheng Gan, William Barr Held and Diyi Yang	39132
<i>Towards Fast and Accurate Modeling for Cross-Lingual Label Projection</i>	
Thang Le, Huy Huu Nguyen, Anh Tuan Luu, Thamar Solorio and Thien Huu Nguyen	39175
<i>False Friends or Cognates? A Cross-lingual Semantic Ambiguity Evaluation for Galician, Portuguese and Spanish</i>	
Marta Vázquez Abuín, Jose Camacho-Collados and Marcos Garcia	39199
<i>BlindGuard: Safeguarding LLM-based Multi-Agent Systems under Unknown Attacks</i>	
Rui Miao, Yixin Liu, Yili Wang, Xu Shen, Yue Tan, Yiwei Dai, Shirui Pan and Xin Wang	39215
<i>Black-Box Membership Inference Attacks for Video Training Data in Multimodal Large Language Models</i>	
Jinrui Wang, Zhenfeng Gao, Wendan Wang, Huili Wang, Zichen Qin, Linjie Zhu, Hongke Fu, Shangguang Wang and Tao Qi	39235
<i>Mitigating Safety Context Amnesia in Multimodal Reasoning Models via Intent-Guided Safety Reasoning</i>	
Xiyao Dong, Guangsheng Cheng, YiLong Chen, Xiaojin Zhang and Kun He	39249
<i>On-policy Reinforcement Fine-tuning with Offline reward for Multi-step Embodied Planning</i>	
Di Wu, Jiaxin Fan, Chloe Gu, Guanbo Wang, Wei Yin, Wenhao Li and Bo Jin	39277
<i>Please refuse to answer me! Mitigating Over-Refusal in Large Language Models via Adaptive Contrastive Decoding</i>	
Yupeng Qi, Ziyu Lyu, Lixin Cui, Lu Bai and Feng Xia	39308
<i>Knowledge Injection Exists in MoE? Exploring Expert-Aware Contrast Decoding in MoE for Mitigating LLMs' Hallucinations</i>	
Xinyue Fang, Zhiliang Tian, Zhen Huang, Ziyi Pan, Zhihua Wen, Xi Wang, Quntian Fang and Dongsheng Li	39326
<i>ERCThinker: Fast-Slow Thinking for Emotion Recognition in Conversation</i>	
Yumeng Fu, Weitao Huang, Junjie Wu, Hao Teng, Shouduo Shang, Meishan Zhang and Bingquan Liu	39344
<i>REaR : Retrieve, Expand and Refine for Effective Multitable Retrieval</i>	
Rishita Agarwal, Himanshu Singhal, Peter Baile Chen, Manan Roy Choudhury, Dan Roth and Vivek Gupta	39360

<i>Figure It Out: Improve the Frontier of Reasoning with Executable Visual States</i>	
Meiqi Chen, Fandong Meng and Jie Zhou	39375
<i>MARCH: Multi-Agent Reinforced Check for Hallucination</i>	
Zhuo Li, Yupeng Zhang, Pengyu Cheng, Jiajun Song, Mengyu Zhou, Hao Li, Shujie Hu, Yu Qin, Erchao.zec, Xiaoxi Jiang and Guanjunjiang	39389
<i>A Survey of Deep Learning for Geometry Problem Solving</i>	
Jianzhe Ma, Wenxuan Wang and Qin Jin	39416
<i>ConsistRM: Improving Generative Reward Models via Consistency-Aware Self-Training</i>	
Yu Liang, Liangxin Liu, Longzheng Wang, Wangyan, Zhang Yueyang, Long Xia, Zhiyuan Sun and Daiting Shi	39449
<i>Beyond Experience Retrieval: Learning to Generate Utility-Optimized Structured Experience for Frozen LLMs</i>	
Xuancheng Li, Haitao Li, Yujia Zhou, Yiqun Liu and Qingyao Ai	39467
<i>Zero-shot Large Language Models for Automatic Readability Assessment</i>	
Riley Grossman and Yi Chen	39483
<i>NoisyCausal: A Benchmark for Evaluating Causal Reasoning Under Structured Noise</i>	
Zhi Xu and Yun Fu	39500
<i>Chain-of-Thought as a Lens: Evaluating Structured Reasoning Alignment between Human Preferences and Large Language Models</i>	
Boxuan Wang, Zhuoyun Li, Xinmiao Huang, Xiaowei Huang and Yi Dong	39514
<i>Robust Membership Inference for Large Language Models under Adversarial Generative Corruption</i>	
Yuanhong Huang, Huili Wang, Xueying Bai, Jinrui Wang, Jiajun Liu, Ziqin Wang, Wanchun Ni, Shangguang Wang and Tao Qi	39531
<i>Towards Trustworthy Smart Contract Synthesis: A Multi-Agent Framework with Lean-Based Verification</i>	
Bowei Zhang, Hanbing Liu, Qixin Tian, Siyu Chen, Ziyuan Wang and Qi Qi	39548
<i>LoRA on the Go: Instance-level Dynamic LoRA Selection and Merging</i>	
Seungeon Lee, Soumi Das, Manish Gupta and Krishna P. Gummadi	39583
<i>Deep Supervised Contrastive Learning of Pitch Contours for Robust Pitch Accent Classification in Seoul Korean</i>	
Hyunjung Joo and GyeongTaek Lee	39602
<i>DARM: Distribution-Aware Reward Modeling by Alleviating Biases from Low Preference-Context Dependency Data</i>	
Shaofan Liu, Guoqiang Zhang, Shihan Dou, Huiyuan Zheng, Yiming Zhou, Junjie Ye, Shaowen Wang, Shichun Liu, Jiazheng Zhang, Tao Gui, Qi Zhang and Xuanjing Huang	39622
<i>KnowRL: Exploring Knowledgeable Reinforcement Learning for Factuality</i>	
Baochang Ren, Shuofei Qiao, Ningyu Zhang, Da Zheng and Huajun Chen	39640
<i>Rethinking Depression Prediction from a Fine-Grained Subscore Modeling Perspective via Multi-Task Learning</i>	
Zhenguang Wang, Bo Li, Wenhui Tan, Peng Cao, Yang Wang, Jia Duan, Fei Wang and Osmar Zaiane	39659

<i>PDTrim: Targeted Pruning for Prefill-Decode Disaggregation in Inference</i>	
Hao Zhang, Lyu Mengsi, Zhuo Chen, Yulong Ao and Yonghua Lin	39673
<i>ASTRA: An Automated Framework for Strategy Discovery, Retrieval, and Evolution for Jailbreaking LLMs</i>	
Xu Liu, Yan Chen, Kan Ling, Yichi Zhu, Hengrun Zhang, Guisheng Fan and Huiqun Yu ..	39691
<i>LLM Safety From Within: Detecting Harmful Content with Internal Representations</i>	
Difan Jiao, Yilun Liu, Ye Yuan, Zhenwei Tang, Linfeng Du, Haolun Wu and Ashton Anderson	39711
<i>AttnPO: Attention-Guided Process Supervision for Efficient Reasoning</i>	
Shuaiyi Nie, Dingsiyu, Wenyuan Zhang, Linhao Yu, Tianmeng Yang, Yao Chen, Weichong Yin, Yu Sun, Hua Wu and Tingwen Liu	39728
<i>Evaluating LLMs on Large-Scale Graph Property Estimation via Random Walks</i>	
Sunil Kumar Maurya and Xin Liu	39749
<i>MedEinst: Benchmarking the Einstellung Effect in Medical LLMs through Counterfactual Differential Diagnosis</i>	
Wenting Chen, Guolin Huang, Wenxuan Wang and Zhongrui Zhu	39778
<i>Libra-VLA: Achieving Learning Equilibrium via Asynchronous Coarse-to-Fine Dual-System</i>	
Yifei Wei, Linqing Zhong, Yi Liu, Yuxiang Lu, Xindong He, Maoqing Yao and Guanghui Ren	39799
<i>Language Statistics and False Belief Reasoning: Evidence from 41 Open-Weight LMs</i>	
Sean Trott, Samuel M. Taylor, Cameron Robert Jones, James A. Michaelov and Pamela D. Rivière	39816
<i>Decoding-Unlearning: Fact Forgetting via Entropy-Guided Inference</i>	
Jingwen Pu, Mingjun Shi, Xinrui Ren, Yizhe Wang, Xinyu Zhang, Zhaokun Wang and Kun She	39834
<i>Global Adaptive Momentum Meets Local Personalized Perturbation: Efficient Federated LLM Fine-Tuning with Zeroth-Order Gradients</i>	
Zihan Chen, Howard Hao Yang, Tony Quek and Kai Fong Ernest Chong	39861
<i>Beyond Query Memorization: Large Language Model Routing with Query Decomposition and Historical Matching</i>	
Bo Lv, Jingbo Sun, Jianwei Lv, Chen Tang, Shaojie Zhang, Nayu Liu, Guoxin Yu, Zihao Li, Qichao Zhang, Dongbin Zhao, Ping Luo and Yue Yu	39876
<i>Market-Bench: Benchmarking Large Language Models on Economic and Trade Competition</i>	
Yushuo Zheng, Huiyu Duan, Zicheng Zhang, Yucheng Zhu, Xionghuo Min and Guangtao Zhai	39893
<i>Benchmarking Post-Training Quantization of Large Language Models under Microscaling Floating Point Formats</i>	
Manyi Zhang, Ji-Fu Li, Zhongao Sun, Haoli Bai, Hui-Ling Zhen, Zhenhua Dong and Xianzhi Yu	39907
<i>HiEdit: Lifelong Model Editing with Hierarchical Reinforcement Learning</i>	
Yangfan Wang, Tianyang Sun, Chen Tang, Jie Liu, Wei Cai and Jingchi Jiang	39924

<i>Too Long, Do Re-weighting for Efficient LLM Reasoning Compression</i>	
Zhong-Zhi Li, Xiao Liang, Zihao Tang, Lei Ji, Peijie Wang, Haotian Xu, Xing W, Haizhen Huang, Weiwei Deng, Yeyun Gong, Zhijiang Guo, Xiao Liu, Fei Yin and Cheng-Lin Liu	39943
<i>Measuring Social Bias in Vision-Language Models with Face-Only Counterfactuals from Real Photos</i>	
Haodong Chen, Qiang Huang, Jiaqi Zhao, Qiuping Jiang, Xiaojun Chang and Jun Yu	39963
<i>KASER: Knowledge-Aligned Student Error Simulator for Open-Ended Coding Tasks</i>	
Zhangqi Duan, Nigel Fernandez and Andrew Lan	39988
<i>GroupToM-Bench: Benchmarking Group Theory of Mind and Nonlinear Social Emergence in MLLMs</i>	
Weidong Tang, Jierui Li, Yueling Hou, Zihan Mei, Can Zhang, Xinyan Wan, Zhiyuan Liang, Pengfei Zhou, Yang You and Wangbo Zhao	40007
<i>Assessing the Belief Consistency of Large Language Models on the Logical Conversation Process</i>	
Tomoki Tsujimura, Matīss Rikters, Masaki Asada, Shusaku Egami, Tatsuya Ishigaki, Ken Yano and Hiroya Takamura	40032
<i>Textual Steering Vectors Can Improve Visual Understanding in Multimodal Large Language Models</i>	
Woody Haosheng Gan, Deqing Fu, Julian Asilis, Ollie Liu, Vatsal Sharan, Robin Jia and Willie Neiswanger	40056
<i>Learning Uncertainty from Sequential Internal Dispersion in Large Language Models</i>	
Ponhvoan Srey, Xiaobao Wu, Cong-Duy T Nguyen and Anh Tuan Luu	40088
<i>On-the-Fly VLA Adaptation via Test-Time Reinforcement Learning</i>	
Changyu Liu, Yiyang Liu, Taowen Wang, Qiao Zhuang, James Chenhao Liang, Wenhao Yang, Renjing Xu, Qifan Wang, Dongfang Liu and Cheng Han	40107
<i>AdaMix: Adaptive Mixing for Short and Long Reasoning Adapters</i>	
Hao Luo, Xiao Yan, Xinyan Li, Qiming Zeng, Yuhao Lin, Shanshan Feng, Hao Wang and Jiawei Jiang	40126
<i>Quantifying and Mitigating Socially Desirable Responding in LLMs: A Desirability-Matched Graded Forced-Choice Psychometric Study</i>	
Kensuke Okada, Yui Furukawa and Kyosuke Bunji	40148
<i>From Curated Data to Scalable Models: Continual Pre-training of Dense and MoE Large Language Models for Tibetan</i>	
Lei Yang, Lei Yu Pan, Bojian Xiong, Renren Jin, Shaowei Zhang, Yue Chen, Ling Shi, Jiang Zhou, Junru Wu, Zhen Wang, Jianxiang Peng, Juesi Xiao, Tianyu Dong, Zhuowen Han, Zhuo Chen, Yuqi Ren and Deyi Xiong	40167
<i>UrbanGeoEval: A City-Scale Benchmark for Evaluating Large Language Models in Geospatial Reasoning</i>	
Mutian Bao, Qiuyi Qi, Tian Liang, Jinjian Zhang, Wei Zhou, Ming Kong, Linjian Mo and Qiang Zhu	40183
<i>TAIGR: Towards Modeling Influencer Content on Social Media via Structured, Pragmatic Inference</i>	
Nishanth Sridhar Nakshatri, Eylon Caplan, Rajkumar Pujari and Dan Goldwasser	40224
<i>Afri-MCQA: Multimodal Cultural Question Answering for African Languages</i>	
Atnafu Lambebo Tonja, Srija Anand, Emilio Villa-Cueva, Israel Abebe Azime, Jesujoba Oluwadara Alabi, Muhidin A. Mohamed, Debela Desalegn Yadeta, Negasi Haile Abadi, Abigail Oppong, Nnaemeka Casmir Obiefuna, Idris Abdulmumin, Naome A Etori, Eric Peter Wairagala, Kanda Patrick Tshinu, Imanigirimbabazi Emmanuel, Gabofetswe Malema, Alham Fikri Aji, David Ifeoluwa Adelani and Tamar Solorio	40249

<i>E-ViC: Reasoning Beyond Text via Embodied Visual Chain for Spatial Intelligence</i>	
Junbo Qi, Yi Zhang, Hanchu Ni, Che Liu, Zhimin Yao, Ruilin Yang, Xiancong Ren, Liangjian Wen, Wei Ge, Yuya Ieiri, Osamu Yoshie, Yong Dai and Xiaozhu Ju	40283
<i>MessToClean: Evidence-Grouped Structure-Preserving Reconstruction for Real-World Degraded Exam Paper Images</i>	
Jiayi Tuo, Cheng Tang, Zihan Wang, Chenyue Zhou, Yao Li, Yanbiao Ma, Chao Wang, Wei Dai, Mingxuan Wang, Shitong Qin and Ziwei Zhao	40304
<i>Phonemes to the Rescue: Multilingual Tokenization Based on International Phonetic Alphabet</i>	
Milan Miletić, Julie Kallini and Ekaterina Shutova	40323
<i>Structure Guided Retrieval-Augmented Generation for Factual Queries</i>	
Miao Xie, Xiao Zhang, Yi Li and Chunli Lv	40350
<i>GROKE: Vision-Free Navigation Instruction Evaluation via Graph Reasoning on OpenStreetMap</i>	
Farzad Shami, Subhrasankha Dey, Nico Van de Weghe and Henrikki Tenkanen	40371
<i>Emotion-Wheel-Guided Audio-Referred Text Representation for Multimodal Emotion Recognition in Conversation</i>	
Eunseon Seong, Harim Lee, Dahye Kim, Changhyun Kim and Dong-Kyu Chae	40392
<i>LongVideoAgent: Multi-Agent Reasoning with Long Videos</i>	
Runtao Liu, Ziyi Liu, Jiaqi Tang, Yue Ma, Renjie Pi, Jipeng Zhang and Qifeng Chen	40404
<i>BOSCH: Black-Box Binary Optimization for Short-Context Attention-Head Selection in LLMs</i>	
Abbas Ghaddar, Ivan Kobzyev, Boxing Chen and Yufei Cui	40417
<i>Scaling Reasoning, Losing Control: Evaluating Instruction Following in Large Reasoning Models</i>	
Tingchen Fu, Yafu Li, Jiawei Gu, Xiaoye Qu and Yu Cheng	40445
<i>MirrorQA: Benchmarking Multimodal LLMs on Mirror-Orientation Reasoning</i>	
Jingping Liu, Xingchen Peng, Yan Zhou, Ziyang Liu, Jie Zhai, Ronghao Chen, Huacan Wang and Xiaofeng Jia	40464
<i>Robust Tool Use via Fission-GRPO: Learning to Recover from Execution Errors</i>	
Zhiwei Zhang, Fei Zhao, Rui Wang, Zezhong Wang, Bin Liang, Jiakang Wang, Yao Hu, Shaosheng Cao and Kam-Fai Wong	40477
<i>DimABSA: Building Multilingual and Multidomain Datasets for Dimensional Aspect-Based Sentiment Analysis</i>	
Lung-Hao Lee, Liang-Chih Yu, Natalia V Loukachevitch, Ilseyar Alimova, Alexander Panchenko, Tzu-Mi Lin, Zhe-Yu Xu, Jian-Yu Zhou, Guangmin Zheng, Jin Wang, Sharanya Awasthi, Jonas Becker, Jan Philip Wahle, Terry Ruas, Shamsuddeen Hassan Muhammad and Saif M. Mohammad	40492
<i>CAP: Controllable Alignment Prompting for Unlearning in LLMs</i>	
Zhaokun Wang, Jinyu Guo, Jingwen Pu, Hongli Pu, Meng Yang, Xunlei Chen, Jie Ou, Wenyi Li, Guangchun Luo and Wenhong Tian	40519
<i>GeoRC: A Benchmark for Geolocation Reasoning Chains</i>	
Mohit Talreja, Joshua Diao, Jim James, Radu Casapu, Tejas Santanam, Ethan Mendes, Alan Ritter, Wei Xu and James Hays	40540
<i>SemCSE-Multi: Multifaceted and Decodable Embeddings for Aspect-Specific and Interpretable Scientific Domain Mapping</i>	
Marc Brinner and Sina Zarriß	40565

<i>Rendering Data Unlearnable by Exploiting LLM Alignment Mechanisms</i>	
Ruihan Zhang and Jun Sun	40587
<i>CAR-bench: Evaluating the Consistency and Limit-Awareness of LLM Agents under Real-World Uncertainty</i>	
Johannes Kirmayr, Lukas Stappen and Elisabeth Andre	40599
<i>Tiny Scales, Great Challenges: The Limits of Multimodal LLMs in Scale Recognition</i>	
Jihang Jin, Ronghao Chen, Hao Zhang, Ziyang Liu, Huacan Wang, Qi Ye and Jingping Liu	40619
<i>M³-VQA: A Benchmark for Multimodal, Multi-Entity, Multi-Hop Visual Question Answering</i>	
Jiatong Ma, Longteng Guo, Yuchen Liu, Zijia Zhao, Dongze Hao, Xuanxu Lin and Jing Liu	40633
<i>From Selection to Refinement: Iterative Optimization for Instruction Data</i>	
Hang Hu, Ziyang Liu, Rujie Wen, Ruihui Hou, Xueyan Wu, Mu Zhang, Jianxing Yu, Tong Ruan and Jingping Liu	40670
<i>VISTA: Verification In Sequential Turn-based Assessment</i>	
Ashley Lewis, Andrew Perrault, Eric Fosler-Lussier and Michael White	40685
<i>Knowledge-to-Verification: Exploring RLVR for LLMs in Knowledge-Intensive Domains</i>	
Zhonghang Yuan, Zhefan Wang, Fang Hu, Zihong Chen, Jinzhe Li, Gang Li, Jie Ying, Huanjun Kong, Songyang Zhang and Nanqing Dong	40715
<i>Mem-Gallery: Benchmarking Multimodal Long-Term Conversational Memory for MLLM Agents</i>	
Yuanchen Bei, Tianxin Wei, Xuying Ning, Yanjun Zhao, Zhining Liu, Xiao Lin, Yada Zhu, Hendrik Hamann, Jingrui He and Hanghang Tong	40750
<i>HarDBench: A Benchmark for Draft-Based Co-Authoring Jailbreak Attacks for Safe Human-LLM Collaborative Writing</i>	
EunTae Kim, Soomin Han and Buru Chang	40785
<i>Mind Reader: Latent User Demand-Guided Content Optimization for Generative Search Engine</i>	
Tong Chen, JiaWei Guo, Yuxi Li, Baiming Chen, Houxing Ren, Zhang Zhiwei, Yunxiang Zhang, Hanyang Xia, Kun Liang and Zhaoran Fan	40832
<i>Retrieval-Augmented Defense: Adaptive and Controllable Jailbreak Prevention for Large Language Models</i>	
Guangyu Yang, Jinghong Chen, Jingbiao Mei, Weizhe Lin and Bill Byrne	40849
<i>CNSL-bench: Benchmarking the Sign Language Understanding Capabilities of MLLMs on Chinese National Sign Language</i>	
Rui Zhao, Xuewen Zhong, Xiaoyun Zheng, Jinsong Su and Yidong Chen	40869
<i>Thinking in Schemas: Robust Syllogistic Reasoning in LLMs</i>	
Federico Ranaldi, Leonardo Ranaldi, Fabio Massimo Zanzotto and Shay B Cohen	40891
<i>Amadeus: Autoregressive Model with Bidirectional Attribute Modelling for Symbolic Music</i>	
Hongju Su, Ke Li, Lan Yang, Honggang Zhang and Yi-Zhe Song	40910
<i>FedMental: Evaluating Federated Learning for Mental Health Detection from Social Media Data</i>	
Nureddin Ali Abdelkadir, Anjali Ratnam, Zeerak Talat and Stevie Chancellor	40929
<i>Edit-Aware Reward Modeling for Chinese Grammatical Error Correction</i>	
Yilin Li and Xiaojun Wan	40945

<i>GenPT: Beyond Self-Report for Reliable LLM Psychometrics via Generative Projective Testing</i> Ming Wang, Shuang Wu, Bixuan Wang, Lu Lin, Yuxin Chen, Xiaocui Yang, Daling Wang, Shi Feng, Yifei Zhang and Yufan Sun	40958
<i>Still Between Us? Evaluating and Improving Voice Assistant Robustness to Third-Party Interruptions</i> Dongwook Lee, Eunwoo Song, Che Hyun Lee, Heeseung Kim and Sungroh Yoon	40975
<i>Learning to Conceal Risk: Controllable Multi-turn Red Teaming for LLMs in the Financial Domain</i> Gang Cheng, Haibo Jin, Wenbin Zhang, Haohan Wang and Jun Zhuang	41005
<i>Beyond Static Alignment: Adaptive Arbitration for Semantic Incongruence in Semi-Supervised Multi-modal Sentiment Analysis</i> Huicong Li, Xiangbo Ji and Wei Wu	41021
<i>LLM Beliefs Are in Their Heads</i> Alessandro Corona Mendoza and Anders Søgaard	41033
<i>ProactiveEval: A Unified Evaluation Framework for Proactive Dialogue Agents</i> Tianjian Liu, Fanqi Wan, Jiajian Guo and Xiaojun Quan	41068
<i>The Visual Iconicity Challenge: Evaluating Vision-Language Models on Sign Language Form–Meaning Mapping</i> Onur Keleş, Asli Ozyurek, Gerardo Ortega, Kadir Gökğöz and Esam Ghaleb	41101
<i>From Competition to Synergy: Unlocking Reinforcement Learning for Subject-Driven Image Generation</i> Ziwei Huang, Ying Shu, Fanghao, Quanyu Long, Wenya Wang, Qiushi Guo, Tiezheng Ge and Leilei Gan	41117
<i>Deep-Reporter: Deep Research for Grounded Multimodal Long-Form Generation</i> Fangda Ye, Kuicai Dong, Xie Zhifei, Yuxin Hu, Yihang Yin, Shurui Huang, Shikai Dong, Chen Zhang, Jianzhu Bao and Shuicheng Yan	41137
<i>Embodied-Reasoner: Synergizing Visual Search, Reasoning, and Action for Embodied Interactive Tasks</i> Wenqi Zhang, Mengna Wang, Gangao Liu, Huixin Xu, Yiwei Jiang, Yongliang Shen, Guiyang Hou, Zhe Zheng, Hang Zhang, Xin Li, Jiajun Liu, Weiming Lu, Peng Li and Yueting Zhuang	41178
<i>Temporal Precision Matters: Brain-Tuning Speech Language Models with Millisecond-Resolution Neural Signals</i> Zhejun Zhang, Wenqing Zhou, Haozhe Xu, Lin Zhang and Lei Li	41208
<i>One Pair Suffices: Unlocking Universal Zero-Shot Translation via Cross-Architecture Alignment</i> Hao Zong, Cong Hu Yuan, Chao Bei, Wentao Chen, Huan Liu, Kaiyu Huang and Degen Huang	41227
<i>LASA: Language-Agnostic Semantic Alignment at the Semantic Bottleneck for LLM Safety</i> Junxiao Yang, Haoran Liu, Jinzhe Tu, Jiale Cheng, Zhexin Zhang, Shiyao Cui, Jiaqi Weng, Jialing Tao, Hui Xue, Hongning Wang, Han Qiu and Minlie Huang	41238
<i>Modeling Annotator Disagreement with Demographic-Aware Experts and Synthetic Perspectives</i> Yinuo Xu, Veronica Derricks, Allison Earl and David Jurgens	41260
<i>Multimodal Safety Evaluation in Generative Agent Social Simulations</i> Alhim Adonai Vera Gonzalez, Carlos Hinojosa, Karen Sanchez, Haidar Bin Hamid, Donghoon Kim and Bernard Ghanem	41295

<i>Quantifying the Impact of Translation Errors on Multilingual LLM Evaluation</i>	
Klaudia Thellmann, Bernhard Stadler, Michael Färber and Jens Lehmann	41311
<i>Long-Context Reasoning Through Proxy-Based Chain-of-Thought Tuning</i>	
Miao Li, Irina Saporina, Alexander Gurung and Mirella Lapata	41335
<i>MASPO: Unifying Gradient Utilization, Probability Mass, and Signal Reliability for Robust and Sample-Efficient LLM Reasoning</i>	
Xiaoliang Fu, Jiaye Lin, Yangyi Fang, Binbin Zheng, Chaowen Hu, Zekai Shao, Cong Qin, Lu Pan, Ke Zeng and Xunliang Cai	41348
<i>Towards A Scanpath-Conditioned Surprisal Theory: Modeling Reader Information States</i>	
Michael Mooney and Edmond S. L. Ho	41366
<i>SafeMT: Multi-turn Safety for Multimodal Language Models</i>	
Han Zhu, Juntao Dai, Jiaming Ji, Haoran Li, Chengkun Cai, Pengcheng Wen, Chi-Min Chan, Boyuan Chen, Yaodong Yang, Sirui Han and Yike Guo	41381
<i>From $\log \pi$ to π: Taming Divergence in Soft Clipping via Bilateral Decoupled Decay of Probability Gradient Weight</i>	
Xiaoliang Fu, Jiaye Lin, Yangyi Fang, Chaowen Hu, Cong Qin, Zekai Shao, Binbin Zheng, Lu Pan and Ke Zeng	41410
<i>Paramanu: Compact and Competitive Monolingual Language Models for Low-Resource Morphologically Rich Indian Languages</i>	
Mitodru Niyogi, Eric Gaussier and Arnab Bhattacharya	41431
<i>Linear-Time and Constant-Memory Text Embeddings Based on Recurrent Language Models</i>	
Tobias Grantner, Emanuel Sallinger and Martin Flechl	41459
<i>Outcome Accuracy is Not Enough: Aligning the Reasoning Process of Reward Models</i>	
Binghai Wang, Yantao Liu, Yuxuan Liu, Tianyi Tang, Shenzhi Wang, Chang Gao, Chujie Zheng, Yichang Zhang, Le Yu, Shixuan Liu, Tao Gui, Qi Zhang, Xuanjing Huang, Bowen Yu, Fei Huang and Junyang Lin	41482
<i>CIS-BWE: Chaos-Informed Speech Bandwidth Extension</i>	
Tarikul Islam Tamiti, Tonmoy Das, Nursadul Mamun and Anomadarshi Barua	41509
<i>The Pitfalls of KV Cache Compression</i>	
Alex Chen, Renato Geh, Aditya Grover, Guy Van Den Broeck and Daniel Mingyi Israel ..	41530
<i>Survey Response Generation: Generating Closed-Ended Survey Responses In-Silico with Large Language Models</i>	
Georg Ahnert, Anna-Carolina Haensch, Barbara Plank and Markus Strohmaier	41554
<i>Sign-Language Datasets at Scale: A Comprehensive Survey on Resources, Benchmarks, and Annotation Standards</i>	
Yiming Ni, Zhi-Qi Cheng, Jiayu Li and Wei Cheng	41578
<i>Anchoring the Affective Manifold: Learning Canonical and Disentangled Representations via Generative Cross-Modal Alignment</i>	
Weibin Li, Jintao Cheng, Xiaoyu Tang and Chi Man Vong	41605
<i>Who Plays Which Role When? Communication Role Dynamics for Peer Recognition and Team Performance Prediction</i>	
Yifan Song, Wenxuan Wendy Shi, Brian Bailey and Tal August	41615

<i>Incomplete In-context Learning</i>	
Wenqiang Wang, Wen Yujia, Yan Xiao, Zhifeng Chen, Yangshijie Zhang, Peng Chen, Mingbo Yang and Xiaochun Cao	41629
<i>Fine-tuning vs. In-context Learning in Large Language Models: A Formal Language Learning Perspective</i>	
Bishwamittra Ghosh, Soumi Das, Till Speicher, Qinyuan Wu, Mohammad Aflah Khan, Deepak Garg, Krishna P. Gummadi and Evimaria Terzi	41651
<i>Where Paths Split: Localized, Calibrated Control of Moral Reasoning in Large Language Models</i>	
Chenchen Yuan, Zheyu Zhang and Gjergji Kasneci	41698
<i>Social Story Frames: Contextual Reasoning about Narrative Intent and Reception</i>	
Joel Mire, Maria Antoniak, Steven R Wilson, Zexin Ma, Achyutarama R Ganti, Andrew Piper and Maarten Sap	41722
<i>Lightweight and Faithful Visual Condition Checking in Behavior Trees via Expert-Regularized Reinforcement Learning</i>	
Hyosik Moon and Eldan Cohen	41764
<i>When in Doubt, Consult: Expert Debate for Sexism Detection via Confidence-Based Routing</i>	
Anwar Alajmi and Gabriele Pergola	41800
<i>Failure Modes in Multi-Hop QA: The Weakest Link Effect and the Recognition Bottleneck</i>	
Meiru Zhang, Zaiqiao Meng and Nigel Collier	41823
<i>TagSpeech: End-to-End Multi-Speaker ASR and Diarization with Fine-Grained Temporal Grounding</i>	
Mingyue Huo, Yiwen Shao and Yuheng Zhang	41847
<i>On the Effect of Hyperparameters in Language Modeling for Computational Linguistics</i>	
Ruoxi Ning, Yongpeng Zhu, Qingcheng Zeng, Tatsuki Kuribayashi and Freda Shi	41863
<i>Prosody as Supervision: Bridging the Non-Verbal–Verbal for Multilingual Speech Emotion Recognition</i>	
Girish, Mohd Mujtaba Akhtar and Muskaan Singh	41881
<i>Mechanisms of Prompt-Induced Hallucination in Vision–Language Models</i>	
William Rudman, Michal Golovanevsky, Dana Arad, Yonatan Belinkov, Carsten Eickhoff, Ritambhara Singh and Kyle Mahowald	41894
<i>LLM4Cell: Taxonomy and Evaluation of LLM and Agentic Models for Single-Cell Biology</i>	
Sajib Acharjee Dip, Adrika Zafor, Bikash Kumar Paul, Uddip Acharjee Shuvo, Muhit Islam Emon, Xuan Wang and Liqing Zhang	41913
<i>Wait! There’s a Way Out: A Decision Mechanism for Forecasting Conversational Derailment</i>	
Laerdon Kim, Vivian Nguyen and Cristian Danescu-Niculescu-Mizil	41955
<i>From Nodes to Narratives: Explaining Graph Neural Networks with LLMs and Graph Context</i>	
Peyman Baghersahi, Gregoire Fournier, Pranav Nyati and Sourav Medya	41975
<i>RealChart2Code: Bridging the Gap in Real-World Chart-to-Code Generation via Multi-Task Evaluation</i>	
Jiajun Zhang, Yuying Li, Zhixun Li, Xingyu Guo, Jingzhuo Wu, Leqi Zheng, Yiran Yang, Jianke Zhang, Qingbin Li, Shannan Yan, Changguo Jia, Junfei Wu, Zilei Wang, Qiang Liu and Liang Wang	41995
<i>Success and Cost Elicit Convention Formation for Efficient Communication</i>	
Saujas Vaduguru, Yilun Hua, Yoav Artzi and Daniel Fried	42033

<i>Measuring User’s Mental Models of Speech Translation in Human-AI Collaboration</i>	
HyoJung Han, Nishant Balepur, Jordan Lee Boyd-Graber and Marine Carpuat	42051
<i>VeriTaS: The First Dynamic Benchmark for Multimodal Automated Fact-Checking</i>	
Mark Rothemmel, Marcus Kornmann, Marcus Rohrbach and Anna Rohrbach	42066
<i>A Survey of Reasoning-Intensive Retrieval: Progress and Challenges</i>	
Yiyang Wei, Tingyu Song, Siyue Zhang and Yilun Zhao	42101
<i>Which Reasoning Trajectories Teach Students to Reason Better? A Simple Metric of Informative Alignment</i>	
Yuming Yang, Mingyoung Lai, Wanxu Zhao, Xiaoran Fan, Zhiheng Xi, Mingqi Wu, Chiyue Huang, Jun Zhao, Haijun Lv, Jian Tong, Yunhua Zhou, Yicheng Zou, Qipeng Guo, Tao Gui, Qi Zhang and Xuanjing Huang	42123
<i>From Proof to Program: Characterizing Tool-Induced Reasoning Hallucinations in Large Language Models</i>	
Farima Fatahi Bayat, Pouya Pezeshkpour and Estevam Hruschka	42151
<i>Paper Circle: An Open-source Multi-agent Research Discovery and Analysis Framework</i>	
Komal Kumar, Aman Chadha, Salman Khan, Fahad Shahbaz Khan and Hisham Cholakkal	42170
<i>PEAR: Pairwise Evaluation for Automatic Relative Scoring in Machine Translation</i>	
Lorenzo Proietti, Roman Grundkiewicz and Matt Post	42189
<i>AV-Dialog: Spoken Dialogue Models with Audio-Visual Input</i>	
Tuochoo Chen, Bandhav Veluri, Hongyu Gong and Shyamnath Gollakota	42208
<i>A BERTology View of LLM Orchestrations: Token- and Layer-Selective Probes for Efficient Single-Pass Classification</i>	
Gonzalo Ariel Meyoyan and Luciano Del Corro	42226
<i>Can Large Language Models Keep Up? Benchmarking Online Adaptation to Continual Knowledge Streams</i>	
Jiyeon Kim, Hyunji Lee, Dylan Zhou, Sue Hyun Park, Seunghyun Yoon, Trung Bui, Franck Dernoncourt, Sungmin Cha and Minjoon Seo	42240
<i>Evaluating Language Model Pluralism through In-the-wild Crowd Discussions</i>	
Gagan Mundada, Rohan Surana, Nandhini Swaminathan, Bodhisattwa Prasad Majumder, Junda Wu, Julian McAuley and Zhouhang Xie	42273
<i>PRBench: Large-Scale Expert Rubrics for Evaluating High-Stakes Professional Reasoning</i>	
Afra Feyza Akyürek, Advait Gosai, Chen Bo Calvin Zhang, Vipul Gupta, Jaehwan Jeong, Anisha Gunjal, Tahseen Rabbani, Maria Mazzone, David Randolph IV, Mohammad Mahmoudi Meymand, Gurshaan Chattha, Paula Rodriguez, Diego A. Mares Buendia, Pavit Singh, Michael Liu, Subodh Chawla, Peter Cline, Lucy Ogaz, Ernesto Gabriel Hernández Montoya, Zihao Wang, Pavi Bhattar, Marcos Ayestaran, Bing Liu and Yunzhong He	42297
<i>Constructing Interpretable Features from Compositional Neuron Groups</i>	
Or David Shafran, Atticus Geiger and Mor Geva	42326
<i>Simulated Students in Tutoring Dialogues: Substance or Illusion?</i>	
Alexander Scarlatos, Jaewook Lee, Simon Woodhead and Andrew Lan	42349
<i>A Survey of Multimodal Mathematical Reasoning: From Perception, Alignment to Reasoning</i>	
Tianyu Yang, Sihong Wu, Yilun Zhao, Zhenwen Liang, Lisen Dai, Chen Zhao, Minhao Cheng, Arman Cohan and Xiangliang Zhang	42386

<i>Empowering Multi-Turn Tool-Integrated Agentic Reasoning with Group Turn Policy Optimization</i> Yifeng Ding, Hung Le, Songyang Han, Kangrui Ruan, Zhenghui Jin, Varun Kumar, Zijian Wang and Anoop Deoras	42409
<i>HypoEval: Hypothesis-Guided Evaluation for Natural Language Generation</i> Mingxuan Li, Hanchen Li and Chenhao Tan	42424
<i>LLMSurgeon: Diagnosing Data Mixture of Large Language Models</i> Yaxin Luo, Jiacheng Cui, Xiaohan Zhao, Xinyi Shang, Jiacheng Liu, Xinyue Bi, Zhaoyi Li and Zhiqiang Shen	42444
<i>The African Languages Lab: A Collaborative Approach to Advancing Low-Resource African NLP</i> Sheriff Issaka, Keyi Wang, Yinka Ajibola, Oluwatuminu Samuel-Ipaye, Zhaoyi Zhang, Nichte Aguillon Jimenez, Evans Kofi Ageyi, Abraham Lin, Rohan Ramachandran, Sadick Abdul Mumin, Faith Nchifor, Mohammed Shuraim Issah, Erick Rosas Gonzalez, Lieqi Liu, Sylvester Kpei, Jemimah Kusi Osei, Carlene Ajeneza, Persis Boateng, Prisca Adwoa Dufie Yeboah and Saadia Gabriel	42460
<i>LogosKG: Hardware-Optimized Scalable and Interpretable Knowledge Graph Retrieval</i> He Cheng, Yifu Wu, Saksham Khatwani, Maya Kruse, Dmitriy Dligach, Timothy A. Miller, Majid Afshar and Yanjun Gao	42478
<i>GCA Framework: A GCC Countries–Grounded Dataset and Agentic Pipeline for Climate Decision Support</i> Muhammad Umer Sheikh, Khawar Shehzad, Salman Khan, Fahad Shahbaz Khan and Muhammad Haris Khan	42494
<i>Effects of Collaboration on the Performance of Interactive Theme Discovery Systems</i> Alvin Po-Chun Chen, Rohan Das, Dananjay Srinivas, Alexandra Barry, Maksim Seniw and Maria Leonor Pacheco	42507
<i>GraphSynth: Resolving the Diversity-Reliability Trade-off with Probabilistic Factor Graphs</i> Zehua Cheng, Wei Dai, Jiahao Sun and Thomas Lukasiewicz	42527
<i>A Structured Clustering Approach for Inducing Media Narratives</i> Rohan Das, Advait Deshmukh, Alexandria Leto, Zohar Naaman, I-Ta Lee and Maria Leonor Pacheco	42544
<i>ACBQ: Adaptive Cross-Block Quantization of Large Language Models</i> Hailing Wang, Jianglin Lu, Yitian Zhang, Huimin Zeng and Yun Fu	42578
<i>SecureGate: Learning When to Reveal PII Safely via Token-Gated Dual-Adapters for Federated LLMs</i> Mohamed Shaaban and Mohamed Elmahallawy	42593
<i>EDUMATH: Generating Standards-aligned Educational Math Word Problems</i> Bryan R Christ, Penelope Molitz, Beau LeBlond, Zachary Gottesman, Jonathan Kropko and Thomas Hartvigsen	42611
<i>AdaFuse: Adaptive Ensemble Decoding for Large Language Models</i> Chengming Cui, Tianxin Wei, Ziyi Chen, Ruizhong Qiu, Zhichen Zeng, Zhining Liu, Xuying Ning, Duo Zhou and Jingrui He	42644
<i>Uncertainty-Aware Test-Time Search for Optimization Problem Solving</i> Linlin Yu, Xujiang Zhao, Dong Li, Yanchi Liu, Wei Cheng, Zhengzhang Chen, Chen Zhao, Feng Chen and Haifeng Chen	42658
<i>TLPO: Token-Level Policy Optimization for Mitigating Language Confusion in Large Language Models</i> Jinho Choo, JunSeung Lee, Jimyeong Kim, Yeeho Song, S. K. Hong and Yeong-Dae Kwon	42670

<i>JW-SVD: Bridging the Cross-Modal Mismatch in Post-Training MLLM Compression</i>	
Runchao Li, Yao Fu, Mu Sheng, Haotian Yu, Xianxuan Long and Kenneth A. Loparo	42691
<i>Advancing Reasoning in Diffusion Language Models with Denoising Process Rewards</i>	
Shaoan Xie, Lingjing Kong, Xiangchen Song, Xinshuai Dong, Guangyi Chen, Eric P. Xing and Kun Zhang	42703
<i>Estimating the Black-box LLM Uncertainty with Distribution-Aligned Adversarial Distillation</i>	
Huizi Cui, Huan Ma, Qilin Wang, Yuhang Gao and Changqing Zhang	42721
<i>CoDial: Interpretable Task-Oriented Dialogue Systems Through Dialogue Flow Alignment</i>	
Radin Shayanfar, Chu Fei Luo, Rohan V Bhambhoria, Samuel Dahan and Xiaodan Zhu . .	42741
<i>BioProAgent: Neuro-Symbolic Grounding for Constrained Scientific Planning</i>	
Yuyang Liu, Jingya Wang, Liuzhenghao Lv and Yonghong Tian	42764
<i>MDP-GRPO: Stabilized Group Relative Policy Optimization for Multi-Constraint Instruction Following</i>	
Mohammad Mahdi Salmani-Zarchi, Zahra Rahimi, Hesham Faili and Mohammad Javad Dousti	42784
<i>TokenTiming: A Dynamic Alignment Method for Universal Speculative Decoding Model Pairs</i>	
Sibo Xiao, Fu Jinyuan, Zhongle Xie and Lidan Shou	42798
<i>GUIDE: Towards Scalable Advising for Research Ideas</i>	
Yaowenqi Liu, BingXu Meng, Rui Pan, Yuxing Liu, Jerry Huang, Jiaxuan You and Tong Zhang	42813
<i>ARES: Adaptive Red-Teaming and End-to-End Repair of Policy-Reward System</i>	
Jiacheng Liang, Yao Ma, Tharindu Kumarage, Satyapriya Krishna, Rahul Gupta, Kai-Wei Chang, Aram Galstyan and Charith Peris	42835
<i>Reasoning Traces Shape Outputs but Models Won't Say So</i>	
Yijie Hao, Lingjie Chen, Ali Emami and Joyce C. Ho	42852
<i>Value of Information: A Framework for Human-Agent Communication</i>	
Yijiang River Dong, Tiancheng Hu, Zheng Hui, Caiqi Zhang, Ivan Vulić, Andreea Bobu and Nigel Collier	42879
<i>AutoRAN: Automated Hijacking of Safety Reasoning in Large Reasoning Models</i>	
Jiacheng Liang, Tanqiu Jiang, Yuhui Wang, Rongyi Zhu, Fenglong Ma and Ting Wang . . .	42897
<i>Thinking Alignment of Scenario-Oriented User Simulation</i>	
Xiaoting Wu, Yi Huang, Chunyang Gao, Mengfei Guo, Jingyu Yao and Junlan Feng	42932
<i>AdapShot: Adaptive Many-Shot In-Context Learning with Semantic-Aware KV Cache Reuse</i>	
Jie Ou, Jinyu Guo, Shiyao Guo, Yuang Li, Ruiqi Wu, Zhaokun Wang, Wenyi Li and Wenhong Tian	42946
<i>Credal Concept Bottleneck Models for Epistemic-Aleatoric Uncertainty Decomposition</i>	
Tanmoy Mukherjee, Thomas Bailleux, Pierre Marquis and Zied Bouraoui	42961
<i>How do Role Models Shape Collective Morality? Exemplar-Driven Moral Learning in Multi-Agent Simulation</i>	
Junjie Liao, Huacong Tang, Zhou Ziheng, Yizhou Wang and Fangwei Zhong	42981
<i>ReActR: Reasoning through Error-Activated Reflection for LLM Post-Training</i>	
Lina Sun	43017

<i>RL-PLUS: Countering Capability Boundary Collapse of LLMs in Reinforcement Learning with Hybrid-policy Optimization</i>	
Yihong Dong, Xue Jiang, Yongding Tao, Huanyu Liu, Kechi Zhang, Lili Mou, Rongyu Cao, Yingwei MA, Jue Chen, Binhua Li, Zhi Jin, Fei Huang, Yongbin Li and Ge Li	43029
<i>Explain the Synth: Interpretable Evaluation of LLM Data Synthesis</i>	
Yue Yang, Fan Yang, Yu Bai and Hao Wang	43054
<i>Beyond the Leaderboard: Rethinking Medical Benchmarks for Large Language Models</i>	
Wenxuan Wang, Zizhan Ma, Guo Yu, Yiu-Fai Cheung, Meidan Ding, Jie Liu, Wenting Chen and Linlin Shen	43078
<i>Speech-Hands: A Self-Reflection Voice Agentic Approach to Speech Recognition and Audio Reasoning with Omni Perception</i>	
Zhen Wan, Chao-Han Huck Yang, Jinchuan Tian, Hanrong Ye, Ankita Pasad, Szu-Wei Fu, Arushi Goel, Ryo Hachiuma, Shizhe Diao, Kunal Dhawan, Sreyan Ghosh, Yusuke Hirota, Zhehuai Chen, Rafael Valle, Chenhui Chu, Shinji Watanabe, Boris Ginsburg and Yu-Chiang Frank Wang	43124
<i>When Personalization Tricks Detectors: The Feature-Inversion Trap in Machine-Generated Text Detection</i>	
Lang Gao, Xuhui Li, Chenxi Wang, Mingzhe Li, Wei Liu, Zirui Song, Jinghui Zhang, Rui Yan, Preslav Nakov and Xiuying Chen	43143
<i>HeteroCache: A Dynamic Retrieval Approach to Heterogeneous KV Cache Compression for Long-Context LLM Inference</i>	
Zhiyuan Shi, Qibo Qiu, Xuefeng, Zhonglin Jiang, Li Yu, Jian Jiang, Xiaofei He and Wenxiao Wang	43172
<i>Sparse-RL: Breaking the Memory Wall in LLM Reinforcement Learning via Stable Sparse Rollouts</i>	
Sijia Luo, Xiaokang Zhang, Yuxuan Hu, Bohan Zhang, Ke Wang, Jinbo Su, Mengshu Sun, Lei Liang and Jing Zhang	43188
<i>C-World: A Computer Use Agent Environment Creator</i>	
Ziqiao Xi, Shuang Liang, Qi Liu, Jiaqing Zhang, Letian Peng, Fang Nan, Meshal Nayim, Tianhui Zhang, Rishika Mundada, Lianhui Qin, Biwei Huang and Kun Zhou	43202
<i>Do Morals Guide How LLMs Think? The Role of Ethical Perspectives in General Problem Solving</i>	
Iseo Kim, Eunjin Hong and Juae Kim	43217
<i>ViLL-E: Video LLM Embeddings for Retrieval</i>	
Rohit Gupta, Jayakrishnan Unnikrishnan, Fan Fei, Sheng Liu, Son Tran and Mubarak Shah	43239
<i>From Trajectories to Graphs: Contract-Checked Editing for Verifier-Guided LLM Reasoning</i>	
Rui Li and Shuang Cao	43259
<i>SiLP: Enhancing Non-Dominant Language Capabilities with a Selective Bidirectional Language Projection Framework</i>	
Junpeng Liu, Jiuyi Li, Kaiyu Huang, Bo Jin, Degen Huang and Hui Xiong	43307
<i>Simple Agents, Biased Judges: Efficient Multi-Party Dialogue Generation & The Evaluation Gap</i>	
Kunal Samanta, Faisal Tareque Shohan, Amine Trabelsi and Richard Khoury	43326
<i>EdgeFormer: Latency-Aware Collaborative Multi-Head Attention of Transformer Inference in Edge Networks</i>	
Yiming Yao, Jianwei Niu, Bin Dai and Tao Ren	43346

<i>LEAF: Knowledge Distillation of Text Embedding Models with Teacher-Aligned Representations</i> Robin Vujanic and Thomas Rückstieß	43362
<i>SAT: Balancing Reasoning Accuracy and Efficiency with Stepwise Adaptive Thinking</i> Weiyang Huang, Xuefeng Bai, Kehai Chen, Xinyang Chen, Yibin Chen, Weili Guan and Min Zhang	43384
<i>Deep Research with Open-Domain Evaluation and Multi-Stage Guardrails for Safety</i> Wei-Chieh Huang, Henry Peng Zou, Yaozu Wu, Dongyuan Li, Yankai Chen, Weizhi Zhang, Yangning Li, Angelo Zangari, Jizhou Guo, Chunyu Miao, Liancheng Fang, Langzhou He, Yinghui Li, Renhe Jiang and Philip S. Yu	43403
<i>P-Check: Advancing Personalized Reward Model via Learning to Generate Dynamic Checklist</i> Kwangwook Seo and Dongha Lee	43447
<i>Locate and Explain: Joint Multimodal Emotion Cause Extraction and Summarization in Conversation</i> Jikun Wan, Chen Gong and Guohong Fu	43472
<i>Understanding the Behaviors of Environment-aware Information Retrieval</i> Ruifeng Yuan, Chaochao Yuan, David Dai, Yu Rong, Hong Cheng, Hou Pong Chan and Chenghao Xiao	43490
<i>ServImage: An Image Generation and Editing Benchmark from Real-world Commercial Imaging Services</i> Fengxian Ji, Jingpu Yang, Zirui Song, Lang Gao, Junhong Liang, Zhenhao Chen, Jinghui Zhang and Xiuying Chen	43504
<i>FAIRGAMER: Evaluating Social Biases in LLM-Based Video Game NPCs</i> Bingkang Shi, Jen-tse Huang, Luo Long, Tianyu Zong, Hongzhu Yi, Yuanxiang Wang, Songlin Hu, Xiaodan Zhang and Zhongjiang Yao	43530
<i>EVOTOOL: Self-Evolving Tool-Use Policy Optimization in LLM Agents via Blame-Aware Mutation and Diversity-Aware Selection</i> Shuo Yang, Caren Han, Xueqi Ma, Yan Li, Mohammad Reza Ghasemi Madani and Eduard Hovy	43553
<i>A Lightweight Explainable Guardrail for Prompt Safety</i> Md Asiful Islam and Mihai Surdeanu	43573
<i>Explainable Disentangled Representation Learning for Generalizable Authorship Attribution in the Era of Generative AI</i> Hieu Man, Van-Cuong Pham, Nghia Trung Ngo, Franck Dernoncourt and Thien Huu Nguyen	43592
<i>PodBench: A Comprehensive Benchmark for Instruction-Aware Audio-Oriented Podcast Script Generation</i> Chenning Xu, Mao Zheng, Mingyu Zheng and Mingyang Song	43605
<i>MIND: From Passive Mimicry to Active Reasoning through Capability-Aware Multi-Perspective CoT Distillation</i> Jin Cui, Jiaqi Guo, Jiepeng Zhou, Ruixuan Yang, Jiayi Lu, Jiajun Xu, Jiangcheng Song, Boran Zhao and Pengju Ren	43622
<i>Knowledge Vector of Logical Reasoning in Large Language Models</i> Zixuan Wang and Yuanyuan Lei	43637

<i>BanHADEX: Towards Explainable HATE Speech Detection in Bangla Using Human Annotated EXplanation</i>	
Faisal Hossain Raquib, Akm Moshir Rahman Mazumder, Md Fahim, Md Tahmid Hasan Fuad, Md Farhan Ishmam, Faria Sultana, M Ashraf Amin, Amin Ahsan Ali and Akmmahbubur Rahman	43652
<i>SMILE-Next: Teaching Large Language Models to Detect, Classify, and Reason about Laughter</i>	
Lee Jung-Mok, Kim Sung-Bin, Joohyun Chang, Lee Hyun and Tae-Hyun Oh	43675
<i>MolMem: Memory-Augmented Agentic Reinforcement Learning for Sample-Efficient Molecular Optimization</i>	
Ziqing Wang, Yibo Wen, Abhishek Pandey, Han Liu and Kaize Ding	43694
<i>Activation Decomposition and Steering for LLM Backdoor Remediation</i>	
Lingfeng Zhong, Qiongfai Xu and Usman Naseem	43713
<i>A Survey of Large Language Models for Text-Guided Molecular Discovery: From Molecule Generation to Optimization</i>	
Ziqing Wang, Kexin Zhang, Zihan Zhao, Yibo Wen, Abhishek Pandey, Han Liu and Kaize Ding	43738
<i>Reward Alignment Optimization: A Direct Point-wise Alignment Approach</i>	
Zelin Li, Jia Leng, Dawei Song and Yangen Hu	43770
<i>Programming over Thinking: Efficient and Robust Multi-Constraint Planning</i>	
Derrick Goh Xin Deik, Quanyu Long, Zhengyuan Liu, Nancy F. Chen and Wenya Wang	43785
<i>LANG: Reinforcement Learning for Multilingual Reasoning with Language-Adaptive Hint Guidance</i>	
Yuchun Fan, Bei Li, Peiguang Li, Yilin Wang, Yongyu Mu, Jian Yang, Xin Chen, Rongxiang Weng, Jingang Wang, Xunliang Cai, JingBo Zhu and Tong Xiao	43838
<i>Can You Make It Sound Like You? Post-Editing LLM-Generated Text for Personal Style</i>	
Connor Baumler, Calvin Bao, Huy Nghiem, Xinchun Yang, Marine Carpuat and Hal Daumé Iii	43867
<i>ReCode: Reinforcing Code Generation with Reasoning-Process Rewards</i>	
Lishui Fan, Yu Zhang, Mouxiang Chen and Zhongxin Liu	43896
<i>PPA-Plan: Proactive Pitfall Avoidance for Reliable Planning in Long-Context LLM Reasoning</i>	
Byeongjin Kim, Gyuwan Kim and Seo Yeon Park	43915
<i>OPeRA: A Dataset of Observation, Persona, Rationale, and Action for Evaluating LLMs on Human Online Shopping Behavior Simulation</i>	
Ziyi Wang, Yuxuan Lu, Wenbo Li, Amirali Amini, Bo Sun, Yakov Bart, Weimin Lyu, Jiri Gesi, Tian Wang, Jing Huang, Yu Su, Upol Ehsan, Malihe Alikhani, Toby Jia-Jun Li, Lydia Chilton and Dakuo Wang	43942
<i>Can LLM Agents Simulate Multi-Turn Human Behavior? Evidence from Real Online Customer Behavior Data</i>	
Yuxuan Lu, Jing Huang, Yan Han, Bingsheng Yao, Sisong Bei, Yaochen Xie, Yisi Sang, Qi He and Dakuo Wang	43961
<i>RATE: Reviewer Profiling and Annotation-free Training for Expertise Ranking in Peer Review Systems</i>	
Weicong Liu, Zixuan Yang, Yibo Zhao and Xiang Li	43978

<i>The Bitter Lesson of Diffusion Language Models for Agentic Workflows: A Comprehensive Reality Check</i>	
Qingyu Lu, Liang Ding, Kanjian Zhang, Jinxia Zhang and Dacheng Tao	43997
<i>Trajectory2Task: Training Robust Tool-Calling Agents with Synthesized Yet Verifiable Data for Complex User Intents</i>	
Ziyi Wang, Yuxuan Lu, Yimeng Zhang, Pei Chen, Ziwei Dong, Jing Huang, Jiri Gesi, Xianfeng Tang, Chen Luo, Qun Liu, Yisi Sang, Hanqing Lu, Manling Li, Jin Lai and Dakuo Wang	44021
<i>CT-FineBench: A Diagnostic Fidelity Benchmark for Fine-Grained Evaluation of CT Report Generation</i>	
Ruifeng Yuan, Wanxing Chang, Weiwei Cao, Bowen Shi, Zhongyu Wei, Ling Zhang and Jianpeng Zhang	44045
<i>Lost in Execution: On the Multilingual Robustness of Tool Calling in Large Language Models</i>	
Zheng Luo, T Pranav Kutralingam, Ogochukwu N. Okoani, Wanpeng Xu, Hua Wei and Xiyang Hu	44059
<i>MTA: A Merge-then-Adapt Framework for Personalized Large Language Models</i>	
Xiaopeng Li, Yuanjin Zheng, Wanyu Wang, Wenlin Zhang, Pengyue Jia, Yingyi Zhang, Haiying He, Mengyang Ma, Yiqi Wang, Maolin Wang, Xuetao Wei and Xiangyu Zhao	44078
<i>Investigating Counterfactual Unfairness in LLMs towards Identities through Humor</i>	
Shubin Kim, Yejin Son, Junyeong Park, Keummin Ka, Seungbeen Lee, Jaeyoung Lee, Hyeju Jang, Alice Oh and Youngjae Yu	44092
<i>Response-G1: Explicit Scene Graph Modeling for Proactive Streaming Video Understanding</i>	
Ke Ma, Jiaqi Tang, Bin Guo, Xueting Han, Ruonan Xu, Qingfeng He, Ziheng Wang, Xu Wang, Qifeng Chen, Zhiwen Yu and Yunhao Liu	44139
<i>POLYCHARTQA: Benchmarking Large Vision-Language Models with Multilingual Chart Question Answering</i>	
Yichen Xu, Liangyu Chen, Liang Zhang, Zihao Yue, Jianzhe Ma, Wenxuan Wang and Qin Jin	44154
<i>Nature-Inspired Population-Based Evolution of Large Language Models</i>	
Yiqun Zhang, Peng Ye, Xiaocui Yang, Shi Feng, Shufei Zhang, Lei Bai, Wanli Ouyang and Shuyue Hu	44187
<i>MTRouter: Cost-Aware Multi-Turn LLM Routing with History-Model Joint Embeddings</i>	
Yiqun Zhang, Hao Li, Zihan Wang, Shi Feng, Xiaocui Yang, Daling Wang, Bo Zhang, Lei Bai and Shuyue Hu	44206
<i>Feeling Right vs. Being Right: How AI Sycophancy Affects Value-Laden Deliberation</i>	
Jeongwoo Ryu, Soomin Kim, Jinsu Eun, Kyusik Kim, Changhoon Oh and Bongwon Suh	44227
<i>ParaSuite: Boosting LLM Reasoning via Paradox Resolution</i>	
Bin Chen, Yu Zhang, Hongfei Ye, Huiyang Wang, Wenxi Liu and Hongyang Chen	44246
<i>ReTRE: Benchmarking LLM Transfer Robustness with Structure-Preserving Variants</i>	
ZhongDong Li, Weijie Shi, Yue Cui, Haolun MA, Yuanjun Liu, Jiawei Li, An Liu, Jia Zhu and Jiajie Xu	44257
<i>Do MLLMs Capture How Interfaces Guide User Behavior? A Benchmark for Multimodal UI/UX Design Understanding</i>	
Jaehyun Jeon, Min Soo Kim, Janghan Yoon, Sumin Shim, Yejin Choi, Hanbin Kim, Dae Hyun Kim and Youngjae Yu	44269

<i>DarwinTOD: LLM-Driven Lifelong Self-evolution for Task-oriented Dialog Systems</i>	
Shuyu Zhang, Yujie Liu, Xinru Wang, Cheng Zhang, Yanmin Zhu and Bin Li	44295
<i>Route to Rome Attack: Directing LLM Routers to Expensive Models via Adversarial Suffix Optimization</i>	
Haochun Tang, Yuliang Yan, Jiahua Lu, Huaxiao Liu and Enyan Dai	44333
<i>GUI0: Self-Evolving Foundational GUI Agents in Super App Ecosystems</i>	
Xinyi Wang, Wei Dai, Kyle Qiao, Ke Wang, Peng Chen, Gang Cao, Kangqin, Zhongpu Wang, Xiaode Zhang, Yanming Liu, Jihao Gu, Jingtao Xu and Gong Zhi	44348
<i>Understanding the Prompt Sensitivity</i>	
Yang Liu and Chenhui Chu	44362
<i>rSIM: Incentivizing Reasoning Capabilities of LLMs via Reinforced Strategy Injection</i>	
Sijia Chen, Baochun Li and Di Niu	44389
<i>Automatic and Reliable Evaluation for Academic Caption-to-Figure Generation with LMMs</i>	
Guanghui Ye, Huan Zhao, Qin Zhu, Fengnan Li, Jiaqi Li, Yixian Shen, Zhonghao Ren and Zhihua Jiang	44406
<i>SafeConstellations: Mitigating Over-Refusals in LLMs Through Task-Aware Representation Steering</i>	
Utsav Maskey, Sumit Yadav, Mark Dras and Usman Naseem	44424
<i>AIPO: Adaptive Information Guided Token-Level Reinforcement Learning for Large Language Model Reasoning</i>	
Bin Chen, Hongfei Ye, Huiyang Wang, Wenxi Liu, Yu Zhang and Furui Liu	44441
<i>AgentGym2: Benchmarking Large Language Model Agents in De-Idealized Real-World Environments</i>	
Zhiheng Xi, Dingwen Yang, Jiaqi Liu, Jixuan Huang, Honglin Guo, Baodai Huang, Tinggang Chen, Qi Zhang, Zhonghang Lu, Chenyu Liu, Jiajun Sun, Jiazheng Zhang, Dingwei Zhu, Xin Guo, Junzhe Wang, Zhihao Zhang, Yuming Yang, Junjie Ye, Minghe Gao, Dongrui Liu, Jiaming Ji, Guohao Li, Tao Gui, Qi Zhang and Xuanjing Huang	44451
<i>Through the Magnifying Glass: Adaptive Perception Magnification for Hallucination-Free VLM Decoding</i>	
Shunqi Mao, Chaoyi Zhang and Weidong Cai	44480
<i>Beyond Fully Random Masking: Attention-Guided Denoising and Optimization for Diffusion Language Models</i>	
Jia Deng, Junyi Li, Xin Zhao, Jinpeng Wang, Hongyu Lu and Ji-Rong Wen	44502
<i>MemRec: Collaborative Memory-Augmented Agentic Recommender System</i>	
Weixin Chen, Yuhan Zhao, Jingyuan Huang, Zihe Ye, Mingxuan Ju, Tong Zhao, Neil Shah, Li Chen and Yongfeng Zhang	44515
<i>INFACT: A Diagnostic Benchmark for Induced Faithfulness and Factuality Hallucinations in Video-LLMs</i>	
Yangjunqi, Yuecong Min, Jie Zhang, Shiguang Shan and Xilin Chen	44545
<i>CLAOCS-TX: Cross-Lingual Triplet Extraction with Aspect-Opinion-Aware Code-Switched Prompting and LLM-Guided Contrastive Distillation</i>	
Lipika Dewangan and Chandresh Kumar Maurya	44561
<i>DAC-Bench: A Decision-Aware Benchmark for Compositional Mobile GUI Tasks</i>	
Yuqing Zhang, Honghui Sheng, Xueyu Hu, Shengyu Zhang and Fei Wu	44578

<i>GMoE: Global Mixture of Experts with Logit Propagation</i> Geonwoo Hong and Taehwan Kim	44599
<i>When Does Language Matter? Multilingual Instructions Reveal Step-wise Language Sensitivity in Vision-Language-Action Models</i> Xuan Dong, Zhe Han, Tianhao Niu, Qingfu Zhu and Wanxiang Che	44615
<i>DialDefer: A Framework for Detecting and Mitigating LLM Dialogic Deference</i> Parisa Rabbani, Priyam Sahoo, Ruben Mathew, Aishee Mondal, Harshita Ketharaman, Nimet Beyza Bozdag and Dilek Hakkani-Tür	44630
<i>Benchmarking Fine-Grained Error Detection in Multimodal Reasoning</i> Chi-Min Chan, Han Zhu, Chunyang Jiang, Jiaming Ji, Juntao Dai, Wei Xue, Sirui Han and Yike Guo	44672
<i>Parallel Test-Time Scaling for Latent Reasoning Models</i> Runyang You, Yongqi Li, Meng Liu, Wenjie Wang, Liqiang Nie and Wenjie Li	44703
<i>SciMDR: Advancing Scientific Multimodal Document Reasoning</i> Ziyu Chen, Yilun Zhao, Chengye Wang, Rilyn R. Han, Manasi Patwardhan and Arman Cohan	44718
<i>When Benchmarks Leak: Inference-Time Decontamination for LLMs</i> Jianzhe Chai, YU Zhe and Jun Sakuma	44743
<i>HSGraphAgent: Knowledge-Graph-Guided Large Language Models for Harmonized System Code Classification</i> Qiang Xia, Zijian Zhang, Ao Wang, Wenhan Wang, Xiangyu Wang and Jian Li	44761
<i>Truth as a Trajectory: What Internal Representations Reveal About Large Language Model Reasoning</i> Hamed Damirchi, Imezadelajara, Ehsan Abbasnejad, Afshar Shamsi, Zhen Zhang and Javen Qin-feng Shi	44774
<i>VL-Calibration: Decoupled Confidence Calibration for Large Vision-Language Models Reasoning</i> Wenyi Xiao, Xinchu XU and Leilei Gan	44791
<i>Don't Corrupt the Fact: A Trustworthy RAG Watermarking Framework based on Dual Factual Shield</i> Hao Huang, JiaTang Luo, Ruihua Zhou, Yunpeng Li and Yuling Liu	44816
<i>I2E: From Image Pixels to Actionable Interactive Environments for Text-Guided Image Editing</i> Jinghan Yu, Junhao Xiao, Chenyu Zhu, Jiaming Li, Jia Li, HanMing Deng, Xirui Wang, Guoli Jia, Jianjun Li, Xiang Bai, Bowen Zhou and Zhiyuan Ma	44827
<i>TInR: Exploring Tool-Internalized Reasoning in Large Language Models</i> Qiancheng Xu, Yongqi Li, Fan Liu, Hongru Wang, Min Yang and Wenjie Li	44851
<i>Fine-Grained Analysis of Shared Syntactic Mechanisms in Language Models</i> Ryoma Kumon and Hitomi Yanaka	44866
<i>One Persona, Many Cues, Different Results: How Sociodemographic Cues Impact LLM Personalization</i> Franziska Weeber, Vera Neplenbroek, Jan Batzner and Sebastian Padó	44892
<i>Lost in the Mix: Evaluating LLM Understanding of Code-Switched Text</i> Amr Mohamed, Yang Zhang, Michalis Vazirgiannis and Guokan Shang	44922
<i>Building LLMs Like LEGO: Two-dimensional Architecture Reassembly of Large Language Models</i> Xingyu Wu, Yu Zhou and KC Tan	44939

<i>How Long Reasoning Chains Influence LLMs' Judgment of Answer Factuality</i> Minzhu Tu, Shiyu Ni and Keping Bi	44957
<i>MMError: A Benchmark for Erroneous Reasoning in Vision-Language Models</i> Yang Shi, Yifeng Xie, Minzhe Guo, Liangsi Lu, Mingxuan Huang, Jingchao Wang, Zhihong Zhu, Boyan Xu and Zhiqi Huang	44973
<i>Learning How and What to Memorize: Cognition-Inspired Two-Stage Optimization for Evolving Memory</i> Derong Xu, Shuochen Liu, Pengfei Luo, Pengyue Jia, Yingyi Zhang, Yi Wen, Yimin Deng, Wenlin Zhang, Enhong Chen, Xiangyu Zhao and Tong Xu	44987
<i>PolitNuggets: Benchmarking Agentic Discovery of Long-Tail Political Facts</i> Yifei Zhu	45012
<i>Diagnosing and Remediating Representation Deficiencies for Deterministic Reasoning in KGQA</i> Gewen Liang, Mufan Xu, Kehai Chen, Wei Wang, Yuwei Wang, Muyun Yang, Tiejun Zhao and Min Zhang	45036
<i>Revealing the Seen, Imagining the Beyond: A Survey of Image-Grounded Chain-of-Thought Reasoning in Multimodal LLMs</i> Qihua Dong, Yitian Zhang, Huimin Zeng, Yizhou Wang, Jianglin Lu, Kuo Yang and Yun Fu	45055
<i>Learning Faster with Better Tokens: Parameter-Efficient Vocabulary Adaptation for Specialized Text Summarization</i> Gunjan Balde, Soumyadeep Roy, Mainack Mondal and Niloy Ganguly	45071
<i>Adversarial Metric Learning for Fine-Grained Emotion Classification</i> Junfan Chen, Sizhe Wu, Richong Zhang and Chunming Hu	45087
<i>OMIBench: Benchmarking Olympiad-Level Multi-Image Reasoning in Large Vision-Language Models</i> Qiguang Chen, Chengyu Luan, Jiajun Wu, Qiming Yu, Yi Yang, Yizhuo Li, Jingqi Tong, Xiachong Feng, Libo Qin and Wanxiang Che	45100
<i>Authorship Attribution in Multilingual Machine-Generated Texts</i> Lucio La Cava, Dominik Macko, Robert Moro, Ivan Srba and Andrea Tagarelli	45136
<i>Layer-Wise High-Impact Parameter Ratio Optimization in Post-Training Quantization for Large Language Models</i> Cuong Pham, Anh Dung Hoang, Cuong C. Nguyen, Trung Le, Gustavo Carneiro and Thanh-Toan Do	45153
<i>On the Role of Discriminative Models in Generative Relation Extraction</i> Guozheng Li, Peng Wang, Zijie Xu, Jing Zhou, Jiajun Liu and Ziyu Shang	45167
<i>Speculative End-Turn Detector for Efficient Speech Chatbot Assistant</i> Hyunjong Ok, Suho Yoo and Jaeho Lee	45184
<i>TwUSD: A Benchmark Dataset and Structure-Aware LLM Framework for User Stance Detection</i> Fuqiang Niu, Zini Chen, Zhiyu Xie, Hu Huang, Qing Liao, Qianlong Wang, Genan Dai and Bowen Zhang	45198
<i>Breaking the Impasse: Dual-Scale Evolutionary Policy Training for Social Language Agents</i> Minzheng Wang, Run Luo, Yanbo Wang, Zichen Liu, Yuqiao Tan, Tao Tan, Nan Xu, Lu Wang and Wenji Mao	45213

<i>Render-of-Thought: Rendering Textual Chain-of-Thought as Images for Visual Latent Reasoning</i> Yifan Wang, Shiyu Li, Peiming Li, Xiaochen Yang, Zheng Wei and Yang Tang	45236
<i>Legal Judgment Prediction: A Reflection on the State of the Art</i> Yi Feng, Chuanyi Li and Vincent Ng.....	45254
<i>GRAD: Generalizing RAG Adaptation with Decoding</i> Youngwon Lee, Seung-won Hwang, Zhewei Yao and Yuxiong He	45274
<i>SHARP: Self-adaptive Harmful Category-aware Prompt Generation for Black-box Jailbreaking</i> Yingjie Xue, Xingyou Xia, Jun Zhang, Yunbo Cao, Dengpan Ye, Guotong Geng and Fei Li	45291
<i>Unveiling the Limits of Large Language Models in Inferring Pragmatic Meaning from Non-Verbal Responses</i> Sugyeong Eo and Heuseok Lim.....	45304
<i>Does RLVR Extend Reasoning Boundaries? Investigating Capability Expansion in Vision-Language Models</i> Minghe Shen, Zhuo Zhi, Chonghan Liu, Shuo Xing, Zhengzhong Tu and Che Liu.....	45317
<i>Perception, Understanding and Reasoning: A Multimodal Benchmark for Video Fake News Detection</i> Cui Yakun, Peng Qi, Fushuo Huo, Hang Du, Weijie Shi, Juntao Dai, Zhenghao Zhu and Sirui Han	45332
<i>RV-HATE: Reinforced Multi-Module Voting for Implicit Hate Speech Detection</i> Yejin Lee, Hyeseon An and Yo-Sub Han	45364
<i>Cross-Prompt Automated Essay Scoring of Multiple Traits: Making Sense of the State of the Art</i> Shengjie Li and Vincent Ng	45384
<i>JARVIS or Ultron? A Survey on the Safety and Security Threats of Computer-Using Agents</i> Ada Chen, Yongjiang Wu, Junyuan Zhang, Jingyu Xiao, Shu Yang, Jen-tse Huang, Kun Wang, Wenxuan Wang and Shuai Wang	45407
<i>Rose-SQL: Role-State Evolution Guided Structured Reasoning for Multi-Turn Text-to-SQL</i> Le Zhou, Feng Yao, Fengcai Qiao, bo xu, Fangyuan Wang and Boyan Xu	45442
<i>Why Do More Experts Fail? A Theoretical Analysis of Model Merging</i> Zijing Wang, Xingle Xu, YongKang Liu, Yiqun Zhang, Peiqin Lin, Shi Feng, Daling Wang, Xiaocui Yang and Hinrich Schuetze.....	45460
<i>Combining Distantly Supervised Models with In Context Learning for Monolingual and Cross-Lingual Relation Extraction</i> Vipul Kumar Rathore, Malik Hammad Faisal, Parag Singla and Mausam	45483
<i>Explaining Sources of Uncertainty in Automated Fact-Checking</i> Jingyi Sun, Greta Warren, Irina Shklovski and Isabelle Augenstein.....	45510
<i>TRACE: A Corpus of Team Creative Discussions</i> Yixuan Jiang, Tiancheng Hu, Jose Hernandez-Orallo, David Stillwell and Luning Sun	45535
<i>Evaluating Structure-Aware Retrieval and Safety in Statute-Centric Legal QA</i> Kyubyung Chae, Jewon Yeom, Jeongjae Park, Seunghyun Bae, Ijun Jang, Hyunbin Jin, Jinkwan Jang and Taesup Kim.....	45553
<i>PRISMA: Preference-Reinforced Self-Training Approach for Interpretable Emotionally Intelligent Negotiation Dialogues</i> Prajwal Vijay Kajare, Priyanshu Priya, Bikash Santra and Asif Ekbal	45574

<i>SkMTEB: Slovak Massive Text Embedding Benchmark and Model Adaptation</i>	
Marek Suppa, Andrej Ridzik, Daniel Hládek, Natália Kňážeková and Viktória Ondrejová .	45597
<i>A Linguistics-Aware LLM Watermarking via Syntactic Predictability</i>	
Shinwoo Park, Hyejin Park, Hyeseon An and Yo-Sub Han	45629
<i>Selective Contrastive Learning For Gloss Free Sign Language Translation</i>	
Chang Hao Lai, Rui Zhao, Xuewen Zhong, Jinsong Su and Yidong Chen	45648
<i>SciCustom: A Framework for Custom Evaluation of Scientific Capabilities in Large Language Models</i>	
Yiyang Gu, Junwei Yang, Junyu Luo, Ye Yuan, Bin Feng, Yingce Xia, Shufang Xie, Kaili Liu, Bohan Wu, Qi Shi, Haoran Li, Beier Xiao, Zhiping Xiao, Xiao Luo, Weizhi Zhang, Philip S. Yu, Zequn Liu and Ming Zhang	45661
<i>NSF-SciFy: Mining the NSF Awards Database for Scientific Claims</i>	
Delip Rao, Weiqiu You, Eric Wong and Chris Callison-Burch	45679
<i>Learning Invariant Modality Representation for Robust Multimodal Learning from a Causal Inference Perspective</i>	
Sijie Mai and Shiqin Han	45698
<i>Multi-Task Representation Alignment on Language Understanding: A Mutual Information Perspective</i>	
Dou Hu, Lingwei Wei, Hongjiang Xiao, Songlin Hu and Yuan Zhang	45722
<i>EconProver: Towards More Economical Test-Time Scaling for Automated Theorem Proving</i>	
Mukai Li, Linfeng Song, Zhenwen Liang, Jiahao Xu, Shansan Gong, Qi Liu, Haitao Mi and Dong Yu	45741
<i>Thinking-Based Non-Thinking: Solving the Reward Hacking Problem in Training Hybrid Reasoning Models via Reinforcement Learning</i>	
Siyuan Gan, Jiaheng Liu, Boyan Wang, Tianpei Yang, Runqing Miao, Yuyao Zhang, Fanyu Meng, Junlan Feng, Linjian Meng, Jing Huo and Yang Gao	45754
<i>LLM-Based Multi-Agent Systems for Clinical Workflows: A Survey of AI Hospitals</i>	
Zonghai Yao and Hong yu	45772
<i>Mind's Eye: A Benchmark of Visual Abstraction, Transformation and Composition for Multimodal LLMs</i>	
Rohit Sinha, Aditya Sanjiv Kanade, Sai Srinivas Kancheti, Vineeth N. Balasubramanian and Tanuja Ganu	45794
<i>EverMemOS: A Self-Organizing Memory Operating System for Structured Long-Horizon Reasoning</i>	
Chuanrui Hu, Xingze Gao, Zuyi Zhou, Dannong Xu, Yi Bai, Xintong Li, Hui Zhang, Tong Li, Chong Zhang, Lidong Bing and Yafeng Deng	45836
<i>Interpretable Coreference Resolution Evaluation Using Explicit Semantics</i>	
Bruno Gatti, Giuliano Martinelli and Roberto Navigli	45854
<i>SURE or Not? Investigating Semantic Understanding in Dense Retrieval Models</i>	
Lingdi Kong, Xuanang Chen, Ben He and Le Sun	45873
<i>When High Accuracy Hides Poor Calibration: Rethinking Confidence Evaluation in Transformer-Based Text Classification with Balanced Brier Score</i>	
Guilherme Fonseca, Gabriel Prenassi, Washington Cunha, Leonardo Chaves Dutra da Rocha and Marcos André Gonçalves	45888

<i>Doc-V*: Coarse-to-Fine Interactive Visual Reasoning for Multi-Page Document VQA</i>	
Yuanlei Zheng, Pei Fu, Hang Li, Ziyang Wang, Yuyi Zhang, Wenyu Ruan, Xiaojin Zhang, Zhongyu Wei, Zhenbo Luo, Jian Luan, Wei Chen and Xiang Bai	45901
<i>Privacy-RI: Privacy-Aware Multi-LLM Agent Collaboration via Reinforcement Learning</i>	
Zheng Hui, Yijiang River Dong, Sanhanat Sivapiromrat, Ehsan Shareghi and Nigel Collier	45924
<i>Reinforcement Learning for Diffusion LLMs via Energy-Based Gibbs Alignment</i>	
Yijia Fan, Jing Yang, Mingyu Liu, Kaitong Cai, Jian Wang, Keze Wang and Jusheng Zhang	45938
<i>CxMP: A Linguistic Minimal-Pair Benchmark for Evaluating Constructional Understanding in Language Models</i>	
Miyu Oba and Saku Sugawara	45949
<i>Evaluating Temporal Consistency in Multi-Turn Language Models</i>	
Yash Kumar Atri, Steven L. Johnson and Thomas Hartvigsen	45964
<i>Synthia: Scalable Grounded Persona Generation from Social Media Data</i>	
Vahid Rahimzadeh, Erfan Moosavi Monazzah, Mohammad Taher Pilehvar and Yadollah Yaghoobzadeh	45983
<i>From Recognition to Reasoning: Benchmarking and Enhancing MLLMs on Real-World Receipt Document Understanding</i>	
Yandi Wang, Libin Zhan, Ziwei Huang, Tiancheng Luo, Yuxuan Jiang, Wang Dong, Leilei Gan and Jun Chen	46007
<i>LogicEval: A Systematic Framework for Evaluating Automated Repair Techniques for Logical Vulnerabilities in Real-World Software</i>	
Syed Md Mukit Rashid, Abdullah Al Ishtiaq, Kai Tu, Yilu Dong, Tianwei Wu, Ali Ranjbar, Tianchang Yang, Najrin Sultana, Shagufta Mehnaz and Syed Rafiul Hussain	46025
<i>Mind the Pause: Disfluency-Aware Objective Tuning for Multilingual Speech Correction with LLMs</i>	
Deepak Kumar, Baban Gain and Asif Ekbal	46050
<i>BioTool: A Comprehensive Tool-Calling Dataset for Enhancing Biomedical Capabilities of Large Language Models</i>	
Xin Gao, Ruiyi Zhang, Meixi Du, Peijia Qin and Pengtao Xie	46071
<i>COMPASS: A Framework for Evaluating Organization-Specific Policy Alignment in LLMs</i>	
Dasol Choi, DongGeon Lee, Brigitta Jessica Kartono, Helena Berndt, Taeyoun Kwon, Joonwon Jang, Haon Park, Hwanjo Yu and Minsuk Kahng	46087
<i>Paper2Rebuttal: A Multi-Agent Framework for Transparent Author Response Assistance</i>	
Qianli Ma, Chang Guo, Zhiheng Tian, Siyu Wang, Jipeng Xiao, Yuanhao Yue and Zhipeng Zhang	46134
<i>Data Efficient RLVR via Off-Policy Influence Guidance</i>	
Erle Zhu, Dazhi Jiang, Yuan Wang, Xujun Li, Jiale Cheng, Yuxian Gu, Yilin Niu, Aohan Zeng, Jie Tang, Minlie Huang and Hongning Wang	46167
<i>CEBC: Conformal Evidence-Bounded Control for Low-Hallucination Vision–Language Generation</i>	
Ashish Mishra, Tarun Kumar, Arpit Shah, Suparna Bhattacharya and Martin Foltin	46193
<i>Dual Alignment Between Language Model Layers and Human Sentence Processing</i>	
Tatsuki Kuribayashi, Alex Warstadt, Yohei Oseki and Ethan Gotlieb Wilcox	46207

<i>Rethinking Composed Image Retrieval Evaluation: A Fine-Grained Benchmark from Image Editing</i> Tingyu Song, Yanzhao Zhang, Mingxin Li, Zhuoning Guo, Dingkun Long, Pengjun Xie, Siyue Zhang, Yilun Zhao and Shu Wu	46224
<i>EDSD: Entropy-Driven Design for Faster Speculative Decoding</i> Longkai Cheng, Ximing Wang, Jiangcai Zhu, Kailai Shao, Chao Chen and Haixiang Hu ..	46243
<i>MUTANT: A Recipe for Multilingual Tokenizer Design</i> Souvik Rana, Ashish Kulkarni, Arul Menezes, Chandra Khatri and Shubham Agarwal	46261
<i>FactCorrector: A Graph-Inspired Approach to Long-Form Factuality Correction of Large Language Models</i> Javier Carnerero-Cano, Massimiliano Pronesti, Radu Marinescu, Tigran T. Tchrakian, James Barry, Jasmina Gajcin, Yufang Hou, Alessandra Pascale and Elizabeth M. Daly	46278
<i>MADE: A Living Benchmark for Multi-Label Text Classification with Uncertainty Quantification of Medical Device Adverse Events</i> Raunak Agarwal, Markus A. Wenzel, Simon Baur, Jonas Zimmer, George Harvey and Jackie Ma	46308
<i>Overcoming Copyright Barriers in Corpus Distribution Through Non-Reversible Hashing</i> Arthur Amalvy, Vincent Labatut, Xavier Bost and Hen-Hsen Huang	46329
<i>UniSRM: A Unified Speech Reward Model for Reasoning-Based Fine-grained Assessment</i> Yuanyuan Wang, Dongchao Yang, Yayue Deng, Zhiyong Wu, Steven Y. Guo, Helen M. Meng and Xixin Wu	46346
<i>DPEPO: Diverse Parallel Exploration Policy Optimization for LLM-based Agents</i> JunShuo Zhang, Chengrui Huang, Feng Guo, Zihan Li, Ke Shi, Menghua Jiang, Jiguo Yu, Shuo Shang and Shen Gao	46367
<i>TEA-Bench: A Systematic Benchmarking of Tool-enhanced Emotional Support Dialogue Agent</i> Xingyu Sui, Yanyan Zhao, Yulin Hu, Jiahe Guo, Weixiang Zhao and Bing Qin	46390
<i>An Information-Theoretic Foundation for the Subregular Hierarchy</i> Mai Phan Quoc Hung, Khanh Nguyen Quoc, Đoàn Minh Luong, Duong Thu Ngan, Duong Thi Phuong Thao and Tuan Do	46417
<i>ReGATE: Learning Faster and Better with Fewer Tokens in MLLMs</i> Chaoyu Li, Yogesh Kulkarni and Pooyan Fazli	46430
<i>From Words to Pixels: A Comprehensive Survey on Large Language Models in Visual Segmentation</i> Yizhou Wang, Mang Tik Chiu, Lingzhi Zhang, Xuan Shen, Sohrab Amirghodsi and Yun Fu	46447
<i>Red-Bandit: Test-Time Adaptation for LLM Red-Teaming via Bandit-Guided LoRA Experts</i> Christos Ziakas, Nicholas Loo, Nishita Jain and Alessandra Russo	46462
<i>Graph-Based Alternatives to LLMs for Human Simulation</i> Joseph Suh, Suhong Moon and Serina Chang	46479
<i>Minimal Free Resolution Guided Adaptive Tree Reasoning</i> Dezhao Tang, Meihan Liu, Yulai Tong, Guan Yuan and Qiuyan Yan	46507
<i>GrAInS: Gradient-based Attribution for Inference-Time Steering of LLMs and VLMs</i> Duy Nguyen, Archiki Prasad, Elias Stengel-Eskin and Mohit Bansal	46523
<i>DiNO: Disinformation Narrative Observer</i> Witold Sosnowski, Arkadiusz Modzelewski, Kinga Skorupska and Adam Wierzbicki	46544

<i>Agentic Very Long Video Understanding</i>	
Aniket Rege, Arka Sadhu, Yuliang Li, Kejie Li, Ramya Korlakai Vinayak, Yuning Chai, Yong Jae Lee and Hyo Jin Kim	46575
<i>OLA: Output Language Alignment in Code-Switched LLM Interactions</i>	
Juhyun Oh, Haneul Yoo, Faiz Ghifari Haznitrana and Alice Oh	46603
<i>Generalizing Trust: Weak-to-Strong Trustworthiness in Language Models</i>	
Lillian Sun, Martin Pawelczyk, Zhenting Qi, Aounon Kumar and Himabindu Lakkaraju . .	46625
<i>Reasoning Is Not All You Need: Examining LLMs for Multi-Turn Mental Health Conversations</i>	
Mohit Chandra, Siddharth Sriraman, Harneet Singh Khanuja, Yiqiao Jin and Munmun De Choudhury	46648
<i>BatonVoice: An Operationalist Framework for Enhancing Controllable Speech Synthesis with Linguistic Intelligence from LLMs</i>	
Yue Wang, Ruotian Ma, Xingyu Chen, Zhengliang Shi, Morunliu Yang, Wanshun Chen, Huang Liu, Jiadi Yao, Xin He, Qu Yang, Qingxuan Jiang, Fanghua Ye, Juntao Li, Zhaopeng Tu, Xiaolong Li, Liefeng Bo and Min Zhang	46683
<i>Monotonic Scaffolding as a Diagnostic Lens for Legal Reasoning in LLMs</i>	
Pedro Calais, Janderson Santos, Anisio Lacerda and Wagner Meira Jr.	46698
<i>Jailbreak-Zero: A Path to Pareto Optimal Red Teaming for Large Language Models</i>	
Kai Hu, Abhinav Aggarwal, Mehran Khodabandeh, David Zhang, Eric Hsin, Li Chen, Ankit Jain, Matt Fredrikson and Akash Bharadwaj	46720
<i>Factual Retrieval in LLMs Is a Redundant, Distributed and Non-Contiguous Process</i>	
Hail Hochman, Natalie Shapira and Yoav Goldberg	46747
<i>What About the Scene With the Hitler Reference? HAUNT: A Framework to Probe LLMs' Self-consistency in Closed Domains Via Adversarial Nudge</i>	
Arka Dutta, Sujan Dutta, Rijul Magu, Soumyajit Datta, Munmun De Choudhury and Ashiqur R. KhudaBukhsh	46769
<i>GitChameleon 2.0: Evaluating AI Code Generation Against Python Library Version Incompatibilities</i>	
Diganta Misra, Nizar Islah, Victor May, Brice Rauby, Zihan Wang, Justine Gehring, Antonio Orvieto, Muawiz Sajjad Chaudhary, Eilif B. Muller, Irina Rish, Samira Ebrahimi Kahou and Massimo Caccia	46792
<i>Knowledge without Wisdom: Measuring Misalignment between LLMs and Intended Impact</i>	
Michael Hardy and Yunsung Kim	46832
<i>Apertus: Democratizing Open and Compliant LLMs for Global Language Environments</i>	
Alejandro Hernández-Cano, Alexander Hägele, Allen Hao Huang, Angelika Romanou, Antoni-Joan Solergibert, Barna Pásztor, Bettina Messmer, Dhia Garbaya, Eduard Frank Ďurech, Ido Hakimi, Juan Garcia Giraldo, Mete Ismayilzada, Negar Foroutan, Skander Moalla, Tiancheng Chen, Vinko Sabolčec, Yixuan Xu, Michael Aerni, Badr AlKhamissi, Inés Altemir Marinas, Mohammad Hossein Amani, Matin Ansaripour, Ilia Badanin, Harold Benoit, Emanuela Boros, Nicholas John Browning, Fabian Bösch, Maximilian Böther, Niklas Canova, Camille Challier, Clément Charmillot, Jonathan Coles, Jan Milan Deriu, Arnout Devos, Lukas Drescher, Daniil Dzenhaliou, Maud Ehrmann, Dongyang Fan, Simin Fan, Silin Gao, Miguel Gila, María Grandury, Diba Hashemi, Alexander Miserlis Hoyle, Jiaming Jiang, Mark Klein, Andrei Kucharavy, Anastasiia Kucherenko, Frederike Lübeck, Roman Machacek, Theofilos Ioannis Manitaras, Andreas Marfurt, Kyle Matoba, Simon Matrenok, Henrique Mendonça, Fawzi Roberto Mohamed, Syrielle Montariol, Luca Mouchel, Sven Najem-Meyer, Jingwei	

Ni, Gennaro Oliva, Matteo Pagliardini, Elia Palme, Andrei Panferov, Léo Paoletti, Marco Passerini, Ivan Pavlov, Auguste Poiroux, Kaustubh Ponkshe, Nathan Ranchin, Javier Rando, Mathieu Sauser, Jakhongir Saydaliev, Mukhammadali Sayfiddinov, Marian Schneider, Stefano Schuppli, Marco Scialanga, Andrei Semenov, Kumar Shridhar, Raghav Singhal, Anna Sotnikova, Alexander Sternfeld, Ayush Kumar Tarun, Paul Teiletche, Jannis Vamvas, Xiaozhe Yao, Hao Zhao, Alexander Ilic, Ana Klimovic, Andreas Krause, Caglar Gulcehre, David Rosenthal, Elliott Ash, Florian Tramèr, Joost VandeVondele, Livio Veraldi, Martin Rajman, Thomas C. Schulthess, Torsten Hoefler, Antoine Bosselut, Martin Jaggi and Imanol Schlag	46877
<i>Reframing Responsibility: Framing-Aware Event Causality Identification</i>	
Jin Zhao, Jiayi Yao, Xinrui Hu and Nianwen Xue	46956
<i>Adaptive Instruction Composition for Automated LLM Red-Teaming</i>	
Jesse Zymet, Andy Luo, Swapnil Shinde, Sahil Wadhwa and Emily Chen	46978
<i>CHOIR: Harmonizing Structured Persona Diversity for Robust Collaborative LLM Reasoning</i>	
Xiangjue Dong, Cong Wang, Maria Teleki, Millennium Bismay, Ruihong Huang and James Caverlee	46997
<i>TabReX: Tabular Referenceless eXplainable Evaluation</i>	
Tejas Anvekar, Junha Park, Aparna Garimella and Vivek Gupta	47015
<i>DREAM-S: Speculative Decoding with Searchable Drafting and Target-Aware Refinement for Multi-modal Generation</i>	
Zining Liu, Yunhai Hu, Tianhua Xia, BO Bao, Eric Sather, Vithursan Thangarasa and Sai Qian Zhang	47031
<i>Revisiting Non-Verbatim Memorization in Large Language Models: The Role of Entity Surface Forms</i>	
Yuto Nishida, Naoki Shikoda, Yosuke Kishinami, Ryo Fujii, Makoto Morishita, Hidetaka Kamigaito and Taro Watanabe	47046
<i>ChatAnime: Towards User-Centered Emotional Support in LLM-based Virtual Character Chat</i>	
Lanlan Qiu, Sophia Xiao Pu, Yeqi Feng, Wenchang Gao and Tianxing He	47062
<i>DyBBT: Dynamic Balance via Bandit-inspired Targeting for Dialog Policy with Cognitive Dual Systems</i>	
Shuyu Zhang, Yifan Wei, Jialuo Yuan, Xinru Wang, Yanmin Zhu, Yujie Liu and Bin Li . . .	47079
<i>SciPedia: Unlocking the Value of Scientific Data for Pre-training</i>	
Yiwei Qin, Zhen Huang, Tiantian Mi, Weiye Si, Qipeng Guo, Siyuan Feng and Pengfei Liu	47112
<i>From Word Sequences to Behavioral Sequences: Adapting Modeling and Evaluation Paradigms for Longitudinal NLP</i>	
Adithya V Ganesan, Vasudha Varadarajan, Oscar Kjell, Whitney Ringwald, Scott M. Feltman, Benjamin J. Luft, Roman Kotov, Ryan L. Boyd and H. Andrew Schwartz	47160
<i>Why Mean Pooling Works: Quantifying Second-Order Collapse in Text Embeddings</i>	
Tomomasa Hara, Hiroto Kurita, Masaaki Imaizumi, Kentaro Inui and Sho Yokoi	47180
<i>Self-Guided Alignment: Adaptive Preference Sensing for Multi-Objective Generation</i>	
Ning Wang, Zhanyang Liu, Taotao Zhou, Xinrui Zhang, Zongru Shao and Haojie Zhou . . .	47202
<i>ReasMark: A Robust Watermark for Attributing LLM Reasoning Under Knowledge Distillation Attacks</i>	
Peizhuo Lv, Ruihua Zhou, Yunpeng Li, Ruigang Liang, Xingshuo Han, XiaoFeng Wang, Wei Dong and Yuling Liu	47221
<i>One Single Hub Text Breaks CLIP: Identifying Vulnerabilities in Cross-Modal Encoders via Hubness</i>	
Hiroyuki Deguchi, Katsuki Chousa and Yusuke Sakai	47242

<i>Awakening Dormant Experts: Counterfactual Routing to Mitigate MoE Hallucinations</i>	
Wentao Hu, Yanbo Zhai, Xiaohui Hu, Mingkuan Zhao, Shanhong yu, Xue Liu, Kaidong Yu, Shuangyong Song and Xuelong Li	47257
<i>Characterizing and Evaluating Working Emotion Vocabularies in Multilingual Large Language Models</i>	
Nicholas Deas, Ivan Ernesto Perez Mejia, Ellie Yang and Kathleen McKeown	47272
<i>HalluCitation Matters: Revealing the Impact of Hallucinated References with 300 Hallucinated Papers in ACL Conferences</i>	
Yusuke Sakai, Hidetaka Kamigaito and Taro Watanabe	47295
<i>Solve-Detect-Verify: Inference-Time Scaling with Flexible Generative Verifier</i>	
Jianyuan Zhong, Zeju Li, Zhijian Xu, Xiangyu Wen, Kezhi Li and Qiang Xu	47377
<i>SPARD: Self-Paced Curriculum for RL Alignment via Integrating Reward Dynamics and Data Utility</i>	
Xuyang Zhi, Peilun Zhou, Chengqiang Lu, Hang Lv, Yiwei Liang, Rongyang Zhang, Yan Gao, Yiwu, Yao Hu, Hongchao Gu, Defu Lian, Hao Wang and Enhong Chen	47402
<i>Lost in Simulation: LLM-Simulated Users are Unreliable Proxies for Human Users in Agentic Evaluations</i>	
Preethi Seshadri, Samuel Cahyawijaya, Ayomide Odumakinde, Sameer Singh and Seraphina Goldfarb-Tarrant	47423
<i>Toward Robust Evaluation for Multilingual Grammatical Error Correction: Can Large Language Models Replace Human References?</i>	
Alla Rozovskaya and Dan Roth	47440
<i>Cell-Based Representation of Relational Binding in Language Models</i>	
Qin Dai, Benjamin Heinzerling and Kentaro Inui	47464
<i>LOTUS: Evolving Multimodal Unlearning via Hyperbolic Entailment and Lorentz Transport</i>	
Zekun Wang, Jingjie Zeng, Yingxu Li, Hongfei Lin and Liang Yang	47525
<i>KARL: Reinforcement Learning for LLM Agents on Multi-Turn Knowledge-Intensive Agentic Tasks</i>	
Xueqiao Sun, Xiao Liu, Bowen Lv, Hanchen Zhang, Bohao Jing, Zehan Qi, Yifan Xu, Yuxiao Dong and Jie Tang	47539
<i>FIGMA: Towards FIne-Grained Music retrievAl</i>	
Nishit Anand, Ashish Seth, Sreyan Ghosh, Dinesh Manocha and Ramani Duraiswami	47559
<i>MathSight: A Benchmark Exploring Have Vision-Language Models Really Seen in University-Level Mathematical Reasoning?</i>	
Yuandong Wang, Yao Cui, Yuxin Zhao, Zhen Yang, Yangfu Zhu and Zhenzhou Shao	47573
<i>TokDrift: When LLM Speaks in Subwords but Code Speaks in Grammar</i>	
Yinxi Li, Yuntian Deng and Pengyu Nie	47605
<i>REVEALER: Reinforcement-Guided Visual Reasoning for Element-Level Text-Image Alignment Evaluation</i>	
FuLin Shi, Wenyi Xiao, Leilei Gan, Liang Ding and Binchen	47630
<i>ReasonAny: Incorporating Reasoning Capability to Any Model via Simple and Effective Model Merging</i>	
Junyao Yang, Chen Qian, Wen Shen, Yong Liu, Jing Shao and Dongrui Liu	47650
<i>Do LLM Agents Mirror Socio-Cognitive Effects in Power-Asymmetric Conversations?</i>	
Anvesh Rao Vijjini, Sagar B. Manjunath and Snigdha Chaturvedi	47676

<i>CIG: Measuring Conversational Information Gain in Deliberative Dialogues with Semantic Memory Dynamics</i>	Ming-Bin Chen, Jey Han Lau and Lea Frermann	47702
<i>A Multilingual Social Bias Benchmark Incorporating Thinking Processes</i>	Masahiro Kaneko, Danushka Bollegala and Timothy Baldwin	47726
<i>Preference Heads in Large Language Models: A Mechanistic Framework for Interpretable Personalization</i>	Weixu Zhang, Ye Yuan, Changjiang Han, Yuxing Tian, Zipeng Sun, Linfeng Du, Jikun Kang, Hong Kang, Xue Liu and Haolun Wu	47742
<i>Influence-based Online Experience Selection for Effective RLHF</i>	Yifan Gong, Jing Yao, Xiting Wang, Xunlong Wang, Xiaoyuan Yi and Xing Xie	47755
<i>NaturalGAIA: A Verifiable Benchmark and Hierarchical Framework for Long-Horizon GUI Tasks</i>	Zihan Zheng, Tianle Cui, Taoran Wang, Fengtao Wang, Jiahui Pan, Lewei He and Qianglong Chen	47772
<i>Look Within or Beyond? A Theoretical Comparison Between Parameter-Efficient and Full Fine-Tuning</i>	YongKang Liu, Xingle Xu, Ercong Nie, Zijing Wang, Shi Feng, Daling Wang, Qian Li and Hinrich Schuetze	47800
<i>Tackling the Root of Misinformation by Teaching Laypeople about Logical Fallacies via Socratic Questioning and Critical Argumentation</i>	Minjing Shi, Junling Wang, Jingwei Ni, Sankalan Pal Chowdhury and Mrinmaya Sachan	47821
<i>TRAC: Teacher-Guided Token Reward with Adaptive Calibration for Robust Policy Optimization</i>	Sitong Wu, Haoru Tan, Xichen Zhang, Bin Xia, Wenhui Zhang, Xiaojuan Qi, Bei Yu and Jiaya Jia	47869
<i>Macaron: Controlled, Human-Written Benchmark for Multilingual and Multicultural Reasoning via Template-Filling</i>	Alaa Elsetohy, Sama Hadhoud, Haryo Akbarianto Wibowo, Chenxi Whitehouse, Genta Indra Winata, Fajri Koto and Alham Fikri Aji	47885
<i>Typology-Aware Multilingual Morphosyntactic Parsing with Joint Abstract Node Modeling</i>	Kutay Acar and Gülşen Eryiğit	47907
<i>SharVeT: Similarity-aware Parameter Sharing with Vector-based Tuning for Efficient LLM Compression</i>	Jeongin Yun, Jaeri Lee, Jongjin Kim, Minjun Kim, Jinho Song and U Kang	47921
<i>Beyond Variance: Knowledge-Aware LLM Compression via Fisher-Aligned Subspace Diagnostics</i>	Ibne Farabi Shihab, Sanjeda Akter and Anuj Sharma	47938
<i>REMIND: Memorization and Unlearning in LLMs Through the Lens of Input Loss Landscapes</i>	Liran Cohen, Yaniv Nemcovsky and Avi Mendelson	47955
<i>Reinforced Efficient Reasoning via Semantically Diverse Exploration</i>	Ziqi Zhao, Zhaochun Ren, Jiahong Zou, Liu Yang, Zhiwei Xu, Xuri Ge, Zhumin Chen, Xinyu Ma, Daiting Shi, Shuaiqiang Wang, Dawei Yin and Xin Xin	47994
<i>Is Chain-of-Thought Really Not Explainability? Chain-of-Thought Can Be Faithful without Hint Verbalization</i>	Kerem Zaman and Shashank Srivastava	48008

<i>MegaRAG: Multimodal Knowledge Graph-Based Retrieval Augmented Generation</i>	
Chi-Hsiang Hsiao, Yi-Cheng Wang, Tzung-Sheng Lin, Yi-Ren Yeh and Chu-song Chen . .	48031
<i>SELEcting over Tokens: Curating Pre-training Data at Scale via Token Classification</i>	
Xin Tong, Weidong Zhang, Jiaang Li, Haibin Chen, Shilei Liu, Langming Liu, Kangtao Lv, Yujin Yuan, Wenbo Su and Bo Zheng	48060
<i>TROJail: Trajectory-Level Optimization for Multi-Turn Large Language Model Jailbreaks with Process Rewards</i>	
Xiqiao Xiong, Ouxiang Li, Zhuo Liu, Moxin Li, Wentao Shi, Fengbin Zhu, Qifan Wang and Fuli Feng	48086
<i>Accurate and Efficient Statistical Testing for Word Semantic Breadth</i>	
Yo Ehara	48110
<i>Constructing coherent spatial memory in LLM agents through graph rectification</i>	
Puzhen Zhang, Xuyang Chen, Yu Feng, Yuhan Jiang and Liqiu Meng	48126