

Findings 2026

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

**Findings of the Association for Computational Linguistics:
ACL 2026**


July 2-7, 2026

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

Diamond



Platinum



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 979-8-89176-395-1

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 Jurgens — 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>I Don't Need Solution. I Need Emotional Support : Empathetic LLMs based on Emotional Validation</i> Suhylene Son, Jungwoo Lim, Myunghoon Kang, Seongtae Hong, Yuna Hur, Evelyn H. Zi and Heuiseok Lim	1
<i>Extraction of Texters' Explicit Emotion Expressions in Crisis Conversations</i> Greg Buda, Ignacio J. Tripodi, Kelly L. Zuromski, Margaret Meagher and Elizabeth A. Olson	27
<i>WikiVideo: Article Generation from Multiple Videos</i> Alexander Martin, Reno Kriz, William Gantt Walden, Kate Sanders, Hannah Recknor, Eugene Yang, Francis Ferraro and Benjamin Van Durme	45
<i>Large Language Models Are Effective Human Annotation Assistants, But Not Good Independent Annotators</i> Feng Gu, Zongxia Li, Carlos R. Colon, Benjamin Evans, Ishani Mondal and Jordan Lee Boyd-Graber	71
<i>Accelerating LLM Fine-Tuning via Embedding Knowledge Transfer</i> Meishu Peng, Ziyue Zhang, Yi Zhang, Pengyang Wang, Zixuan Yuan and Denghui Zhang	90
<i>Understanding Secret Leakage Risks in Code LLMs: A Tokenization Perspective</i> Meifang Chen, Zhe Yang, Huang Nianchen, Yizhan Huang, Yichen LI, Zihan Li and Michael R. Lyu	108
<i>MALLM-GAN: Multi-Agent Large Language Model as Generative Adversarial Network for Synthesizing Tabular Data</i> Yaobin Ling, Xiaoqian Jiang and Yejin Kim	120
<i>EMCompress: Video-LLMs with Endomorphic Multimodal Compression</i> Zheyu Fan, Jiateng Liu, Yuji Zhang, Zihan Wang, Yi R. Fung, Manling Li and Heng Ji	137
<i>Learning from Contrastive Prompts: An Automated Prompt Optimization Framework</i> Mingqi Li, Karan Aggarwal, Yong Xie, Aitzaz Ahmad and Stephen Lau	163
<i>Private Seeds, Public LLMs: Realistic and Privacy-Preserving Synthetic Data Generation</i> Qian Ma and Sarah Rajtmajer	189
<i>ToolScope: An Agentic Framework for Vision-Guided and Long-Horizon Tool Use</i> Mengjie Deng, Guanting Dong and Zhicheng Dou	211
<i>VChain: Chain-of-Visual-Thought for Reasoning in Video Generation</i> Ziqi Huang, Ning Yu, Gordon Chen, Haonan Qiu, Paul Debevec and Ziwei Liu	226
<i>Diversity Collapse in Multi-Agent LLM Systems: Structural Coupling and Collective Failure in Open-Ended Idea Generation</i> Nuo Chen, Yicheng Tong, Yuzhe Yang, Yufei He, Xueyi Zhang, Zou Qingyun, Qian Wang and Bingsheng He	251
<i>Justice in Judgment: Unveiling (Hidden) Bias in LLM-assisted Peer Reviews</i> Sai Suresh Macharla Vasu, Ivaxi Sheth, Hui-Po Wang, Ruta Binkyte and Mario Fritz	307
<i>A Study of LLMs' Preferences for Libraries and Programming Languages</i> Lukas Twist, Jie M. Zhang, Mark Harman, Don Syme, Joost Noppen, Helen Yannakoudakis and Detlef Nauck	331

<i>Unveiling Inherent Visual Grounding in Multimodal LLMs for Text-Rich Images</i>	
Shijie Zhou, Jihyung Kil, Ming Li, Jiuxiang Gu, Curtis Wigington, Rajiv Jain, Changyou Chen and Ruiyi Zhang	352
<i>PUMA: Projected Universal Multilingual ASR for Low-Resource Settings. Application to Diverse African Languages</i>	
Ilyes Oukid, Bilal Faye, Hanane Azzag, Mustapha Lebbah and Said Yacine Boulahia	371
<i>MGPO: Thinking with Images via Multi-Turn Grounding-Based Reinforcement Learning</i>	
Xinyu Huang, Yuhao Dong, Weiwei Tian, Bo Li, Rui Feng and Ziwei Liu	383
<i>How Grounded is Wikipedia? A Study on Structured Evidential Support and Retrieval</i>	
William Gantt Walden, Kathryn Ricci, Miriam Wanner, Zhengping Jiang, Chandler May, Rongkun Zhou and Benjamin Van Durme	400
<i>Jailbreaking Attacks vs. Content Safety Filters: How Far Are We in the LLM Safety Arms Race?</i>	
Yuan Xin, Dingfan Chen, Linyi Yang, Michael Backes and Xiao Zhang	421
<i>BrowseConf: Confidence-Guided Test-Time Scaling for Web Agents</i>	
Litu Ou, Kuan Li, Huifeng Yin, Liwen Zhang, Zhongwang Zhang, Xixi Wu, Rui Ye, Zile Qiao, Yong Jiang, Pengjun Xie, Fei Huang and Jingren Zhou	446
<i>HyperEdit: Unlocking Instruction-based Text Editing in LLMs via Hypernetworks</i>	
Yiming Zeng, Jinghan Cao, Zexin Li, Wanhao Yu, Zhankai Ye, Dawei Xiang, Ting Hua, Xin Liu, Shangqian Gao and Tingting Yu	466
<i>Strategy-Induct: Task-Level Strategy Induction for Instruction Generation</i>	
Po-Chun Chen, Hen-Hsen Huang and Hsin-Hsi Chen	481
<i>Idiom Understanding as a Tool to Measure the Dialect Gap</i>	
David Beauchemin, Yan Tremblay, Mohamed Amine Youssef and Richard Khoury	505
<i>Logic Jailbreak: Efficiently Unlocking LLM Safety Restrictions Through Formal Logical Expression</i>	
Jingyu Peng, Maolin Wang, Nan Wang, Jiatong Li, Yuchen Li, Yuyang Ye, Wanyu Wang, Pengyue Jia, Kai Zhang and Xiangyu Zhao	523
<i>Two Heads Are Better Than One: Audio-Visual Speech Error Correction with Dual Hypotheses</i>	
Sungnyun Kim, Kangwook Jang, Sungwoo Cho, Joon Son Chung, Hoi-Rin Kim and Se-Young Yun	544
<i>Towards Hierarchical Multi-Step Reward Models for Enhanced Reasoning in Large Language Models</i>	
Teng Wang, Jiang Zhangyi, Zhenqi He, Hailei Gong, Shenyang Tong, Wenhan Yang, Zeyu Li, Yanan Zheng, Zifan He, Zewen Ye, Shengjie Ma and Jianping Zhang	565
<i>SciVQR: A Multidisciplinary Multimodal Benchmark for Advanced Scientific Reasoning Evaluation</i>	
Longteng Guo, Xuanxu Lin, Dongze Hao, Tongtian Yue, Pengkang Huo, Jiatong Ma, Yuchen Liu and Jing Liu	577
<i>UrBLiMP: A Benchmark for Evaluating the Linguistic Competence of Large Language Models in Urdu</i>	
Farah Adeeba, Brian Dillon, Hassan Sajjad and Rajesh Bhatt	602
<i>ANCHOR: LLM-driven Subject Conditioning for Text-to-Image Synthesis</i>	
Aashish Anantha Ramakrishnan, Sharon X Huang and Dongwon Lee	618
<i>Personalized Benchmarking: Evaluating LLMs by Individual Preferences</i>	
Cristina Garbacea, Heran Wang and Chenhao Tan	639

<i>Op-Fed: Opinion, Stance, and Monetary Policy Annotations on FOMC Transcripts Using Active Learning</i>	
Alisa Kanganis and Katherine A. Keith	676
<i>C³D: Enhancing LLM Reasoning via Confidence-Guided Contrastive Decoding</i>	
Yufeng Zhang, Xuepeng Wang, Lingxiang Wu and Jinqiao Wang	700
<i>Do LLMs Really Know What They Don't Know? Internal States Mainly Reflect Knowledge Recall Rather Than Truthfulness</i>	
Chi Seng Cheang, Hou Pong Chan, Wenxuan Zhang and Yang Deng	713
<i>Enhancing Hallucination Detection via Future Context</i>	
Joosung Lee, Cheonbok Park, Hwiyeol Jo, Jeonghoon Kim, Joonsuk Park and Kang Min Yoo	731
<i>MTO: A Multi-turn Conversational Text-to-OverpassQL Dataset for Enhanced OpenStreetMap Query Generation</i>	
Haodi Zhang, Xinrui Zhu, Mingze Kong, Zhidan Liu, Tao Fan, Kaishun Wu and Yuanfeng Song	750
<i>A Dual-Phase Self-Evolution Framework for Large Language Models</i>	
Haoran Sun, Zekun Zhang and Shaoning Zeng	772
<i>Preference-Aware Memory Update for Long-Term LLM Agents</i>	
Haoran Sun, Zekun Zhang and Shaoning Zeng	783
<i>Understanding In-Context Learning Beyond Transformers: An Investigation of State Space and Hybrid Architectures</i>	
Shenran Wang, Timothy Tin-Long Tse and Jian Zhu	794
<i>C²DLM: Causal Concept-Guided Diffusion Large Language Models</i>	
Kairong Han, Nuanqiao Shan, Ziyu Zhao, Zijing Hu, Xinpeng Dong, Ye Jun Jian, Lujia Pan, Fei Wu and Kun Kuang	821
<i>CMTD: Cognitive Modeling with Traits and Distortions for Multimodal Emotion Recognition in Conversations</i>	
Minh-Tien Nguyen, Huu-Loi Le, Manh-Cuong Phan and Hajime Hotta	839
<i>Bridging the Culture Gap: A Framework for LLM-Driven Socio-Cultural Localization of Math Word Problems in Low-Resource Languages</i>	
Israel Abebe Azime, Tadesse Destaw Belay, Dietrich Klakow, Philipp Slusallek and Anshuman Chhabra	855
<i>Stop Guessing When to Stop Testing: Efficient Model Evaluation with Just Enough Data</i>	
Ofir Arviv, Kristjan Greenewald, Yotam Perlitz, Hadar Mulian, Michal Shmueli-Scheuer and Leshem Choshen	871
<i>AnyGraph: Graph Foundation Model in the Wild</i>	
Lianghao Xia and Chao Huang	882
<i>DOSE: Data Selection for Multi-Modal LLMs via Off-the-Shelf Models</i>	
Biao Wu, Yiwu Zhong, Meng Fang and Ling Chen	897
<i>MergeIT: From Selection to Merging for Efficient Instruction Tuning</i>	
Hongyi Cai, Yuqian Fu, Hongming Fu and Bo Zhao	912
<i>The Illusion of Insight in Reasoning Models</i>	
Liv G. d'Aliberti and Manoel Horta Ribeiro	924

<i>Adaptive Layer Selection for Layer-Wise Token Pruning in LLM Inference</i>	Rei Taniguchi, Yuyang Dong, Makoto Onizuka and Chuan Xiao	967
<i>Aligning What LLMs Do and Say: Towards Self-Consistent Explanations</i>	Sahar Admoni, Ofra Amir, Assaf Hallak and Yftah Ziser	987
<i>TARo: Token-level Adaptive Routing for LLM Test-time Alignment</i>	Arushi Rai, Qiang Zhang, Hanqing Zeng, Yunkai Zhang, Dipesh Tamboli, Xiangjun Fan, Zhuokai Zhao and Lizhu Zhang	1004
<i>ProRank: Prompt Warmup via Reinforcement Learning for Small Language Models Reranking</i>	Xianming LI, Aamir Shakir, Rui Huang, Julius Lipp, Benjamin Clavié and Jing Li	1026
<i>ChemAmp: Amplified Chemistry Tools via Composable Agents</i>	Zhucong Li, Powei Chang, Jin Xiao, Zhijian Zhou, Qianyu He, Jiaqing Liang, Fenglei Cao, Xu Yinghui and Yuan Qi	1038
<i>CodeContests-O: Powering LLMs via Feedback-Driven Iterative Test Case Generation</i>	Jianfeng Cai, Jinhua Zhu, Ruopei Sun, Kangwen Zhao, Dongyun Xue, Mingxiao Feng, Wengang Zhou and Houqiang Li	1054
<i>Hybrid-Vector Retrieval for Visually Rich Documents: Combining Single-Vector Efficiency and Multi-Vector Accuracy</i>	Juyeon Kim, Geon Lee, Dongwon Choi, Taeuk Kim and Kijung Shin	1073
<i>Multi-Scale Spectral Selection and Entropy-Guided Uncertainty Fusion for Multimodal Rumor Detection</i>	Zongliang Han, Wenyu Guo, Guoqing Jin, Yang Liu, Fan Li, Dong Yu, Yan Song and Zhangfengzhen	1090
<i>Covariance Matrix-Driven Image Channel Allocation for Multimodal Fake News Detection</i>	Zongliang Han, Wenyu Guo, Guoqing Jin, Yang Liu, Yan Song, Dong Yu and Wang Min . .	1101
<i>Mechanistic Insights into Deferred Semantic Drift in LLMs</i>	Jingjie Zeng, Huayang Li, Liang Yang, Shaowu Zhang, Yuanyuan Sun and Hongfei Lin	1116
<i>JointCQ: Improving Factual Hallucination Detection with Joint Claim and Query Generation</i>	Fan Xu, Huixuan Zhang, Zhenliang Zhang, Jiahao Wang and Xiaojun Wan	1138
<i>Pseudo2Real: Task Arithmetic for Pseudo-Label Correction in Automatic Speech Recognition</i>	Yi-Cheng Lin, Yu-Hsuan Li Liang, Hsuan Su, Tzu-Quan Lin, Shang-Tse Chen, Yun-Nung Chen and Hung-yi Lee	1160
<i>SGA-MCTS: Decoupling Planning from Execution via Training-Free Atomic Experience Retrieval</i>	Xin Xie, Dongyun Xue, Wuguannan Yao, Mingxiao Feng, Wengang Zhou, Xiang Qi, Houqiang Li and Peng Zhang	1176
<i>Robustness via Referencing: Defending against Prompt Injection Attacks by Referencing the Executed Instruction</i>	Yulin Chen, Haoran Li, Yuan Sui, Yue Liu, Yufei He, Xiaoling Bai, Chi Fei, Li Yabo, Haozhe Ma, Yangqiu Song and Bryan Hooi	1194
<i>Differentially Private Synthetic Text Generation for Retrieval-Augmented Generation (RAG)</i>	Junki Mori, Kazuya Kakizaki, Taiki Miyagawa and Jun Sakuma	1216
<i>When Slower Isn't Truer: Inverse Scaling Law of Truthfulness in Multimodal Reasoning</i>	Sitong Fang, Wenjing Cao, Jiahao Li, Xuyao Wang, Chi-Min Chan, Sirui Han, Juntao Dai, Yike Guo, Yaodong Yang and Jiaming Ji	1236

<i>StressTest: Can YOUR Speech LM Handle the Stress?</i> Iddo Yosha, Gallil Maimon and Yossi Adi	1264
<i>EVADE: LLM-Based Explanation Generation and Validation for Error Detection in NLI</i> Longfei Zuo, Barbara Plank and Siyao Peng	1286
<i>AgenticRAGTracer: A Hop-Aware Benchmark for Diagnosing Multi-Step Retrieval Reasoning in Agentic RAG</i> Qijie You, Wenkai Yu, Hao Liang, Zhen Hao Wong and Wentao Zhang	1301
<i>Attention to Non-Adopters</i> Kaitlyn Zhou, Kristina Gligorić, Myra Cheng, Michelle S. Lam, Vyoma Raman, Boluwatife Aminu, Caeley Woo, Michael Brockman, Hannah Cha and Dan Jurafsky	1336
<i>Ready to Translate, Not to Represent? Bias and Performance Gaps in Multilingual LLMs Across Language Families and Domains</i> Md. Faiyaz Abdullah Sayeedi, Subhey Sadi Rahman, Md. Mahbub Alam, Md. Adnanul Islam, Jannatul Ferdous Deepti, Tasnim Mohiuddin, Md Mofijul Islam and Swakkhar Shatabda	1367
<i>Music Audio-Visual Question Answering Requires Specialized Multimodal Designs</i> Wenhao You, Xingjian Diao, Wenjun Huang, Chunhui Zhang, Keyi Kong, Weiyi Wu, Chiyu Ma, Zhongyu Ouyang, Tingxuan Wu, Ming Cheng, Soroush Vosoughi and Jiang Gui	1392
<i>TemporalVLM: Video LLMs for Temporal Reasoning in Long Videos</i> Fawad Javed Fateh, Umer Ahmed, Hamza Khan, Zeeshan Zia and Quoc-Huy Tran	1427
<i>Mitigating Judgment Preference Bias in Large Language Models through Group-Based Polling</i> Shuliang Liu, Zhipeng Xu, Zhenghao Liu, Yukun Yan, Minghe Yu, Yu Gu, Chong Chen, Huiyuan Xie and Ge Yu	1448
<i>Enhancing Long-Chain Reasoning Distillation through Error-Aware Self-Reflection</i> Zhuoyang Wu, Xinze Li, Zhenghao Liu, Yukun Yan, Zhiyuan Liu, Minghe Yu, Cheng Yang, Yu Gu, Ge Yu and Maosong Sun	1465
<i>New Compendium of a Myriad of Plants: A New Dataset Describing Ancient Chinese Plants</i> Xiaobin Shen, Zhongqing Wang, Shichen Li, Chu-Ren Huang and Guodong Zhou	1483
<i>Generative Interfaces for Language Models</i> Jiaqi Chen, Yanzhe Zhang, Yutong Zhang, Yijia Shao and Diyi Yang	1499
<i>EvoEdit: Evolving Null-space Alignment for Robust and Efficient Knowledge Editing</i> Sicheng Lyu, Yu Gu, Xinyu Wang, Jerry Huang, Sitao Luan, Yufei Cui, Xiao-Wen Chang and Peng Lu	1520
<i>Efficient Reasoning for LLMs through Speculative Chain-of-Thought</i> Jikai Wang, Juntao Li, Jianye Hou, Yan Bowen, Lijun Wu and Min Zhang	1541
<i>From Coarse to Fine: Self-Adaptive Hierarchical Planning for LLM Agents</i> Haoran Tan, Zeyu Zhang, Chen Ma, Tianze Liu, Quanyu Dai and Xu Chen	1554
<i>Data-Centric Perspectives on Agentic Retrieval-Augmented Generation: A Survey</i> Jingwen Deng, Jihao Huang, Zhen Hao Wong, Hao Liang, Quanqing Xu, Bin Cui and Wentao Zhang	1570
<i>WISCA: A Lightweight Model Transition Method to Improve LLM Training via Weight Scaling</i> Jiacheng Li, Jianchao Tan, Zhidong Yang, Pingwei Sun, Feiye Huo, Jiayu Qin, Xiangyu Zhang, Maoxin He, Guangming Tan, Weile Jia, Xunliang Cai and Tong Zhao	1589

<i>Disentangling Codemixing in Chats: The NUS ABC Codemixed Corpus</i>	
Svetlana Churina, Akshat Gupta, Nur Insyirah Binte Imam Mujtahid and Kokil Jaidka	1602
<i>When Thoughts Meet Facts: Reusable Reasoning for Long-Context LMs</i>	
Soyeong Jeong, Taehee Jung, Sung Ju Hwang, Joo-Kyung Kim and Dongyeop Kang	1625
<i>Infinity-Parser: Layout-Aware Reinforcement Learning with High-quality Document Parsing Dataset</i>	
Baode Wang, Biao Wu, Weizhen Li, Meng Fang, Zuming Huang, Jun Huang, Yanjie Liang, Haozhe Wang, Ling Chen, Wei Chu and Yuan Qi	1647
<i>Lunar-Bench: Towards Evaluating Task-Oriented Reasoning of LLMs in Lunar Exploration Scenarios</i>	
Xin-Yu Xiao, Ye Tian, Erwei Yin, Zhixian He, Shiqi Wang, Yalei Liu and Qianchen Xia . . .	1668
<i>BV-Blend: Uncertainty-Weighted Historical Baselines for Stable Critic-Free RL with Verifiable Rewards</i>	
Yupeng Chang, Yuan Wu and Yi Chang	1706
<i>Rethinking Research on Stereotypes: An Analysis through Social Psychological and Computational Perspectives</i>	
Kaustubh Shivshankar Shejole and Pushpak Bhattacharyya	1726
<i>Do LLMs Catch Their Own Mistakes? A Comprehensive Benchmark for Reflective Tool Use LLMs</i>	
Zheyuan Liu, Liqiang Xiao, Yang Li, Hyokun Yun, Lihong Li, Chao Zhang and Meng Jiang	1748
<i>Understanding Generalization in Role-Playing Models via Information Theory</i>	
Yongqi Li, Hao Lang, Fei Huang, Tiejun Qian and Yongbin Li	1774
<i>A Multilingual Dataset and Empirical Validation for the Mutual Reinforcement Effect in Information Extraction</i>	
Chengguang Gan, Sunbowen Lee, Qingyu Yin, Yunhao Liang, Xinyang He, Hanjun Wei, Younghun Lim, Shijian Wang, Hexiang Huang, QingHao Zhang, Shiwen Ni and Tatsunori Mori	1810
<i>TRACE: An Experiential Framework for Coherent Multi-hop Knowledge Graph Question Answering</i>	
Yingxu Wang, Jiaxin Huang, Mengzhu Wang and Nan Yin	1830
<i>Can LLMs See Without Pixels? Benchmarking Spatial Intelligence from Textual Descriptions</i>	
Zhongbin Guo, Zhen Yang, Yushan Li, Xinyue Zhang, Wenyu Gao, Jiacheng Wang, Chengzhi Li, Xiangrui Liu and Ping Jian	1852
<i>Lil: Less is Less When Applying Post-Training Sparse-Attention Algorithms in Long-Decode Stage</i>	
Junhao Hu, Fangze Li, Mingtao Xu, Feifan Meng, Shiju Zhao, Tiancheng Hu, Ting Peng, Anmin Liu, Wenrui Huang, Chenxu Liu, Ziyue Hua and Tao Xie	1898
<i>Static Models, Dynamic World: A Unified Perspective on Temporal Perception in Large Language Models</i>	
Chenhao Li, Dandan Song, Changzhi Zhou, Jun Yang, Yuhang Tian, Huipeng Ma, Guangyuan Feng, Luan Zhang, Xudong Li and Ke Duan	1913
<i>Decomposed Trust: Privacy, Adversarial Robustness, Ethics, and Fairness in Low-Rank LLMs</i>	
Daniel Agyei Asante, Md Mokarram Chowdhury and Yang Li	1933
<i>Beyond Token Length: Step Pruner for Efficient and Accurate Reasoning in Large Language Models</i>	
Canhui Wu, Qiong Cao, Chang Li, Zhenfang Wang, Chao Xue, Yuwei Fan, Wei Xi and Xiaodong He	1953
<i>Exploring Reasoning Reward Model for Agents</i>	
Kaixuan Fan, Kaituo Feng, Manyuan Zhang, Tianshuo Peng, Zhixun Li, Yilei Jiang, Shuang Chen and Xiangyu Yue	1975

<i>MTMCS-Bench: Evaluating Contextual Safety of Multimodal Large Language Models in Multi-Turn Dialogues</i>	
Zheyuan Liu, Dongwhi Kim, Yixin Wan, Xiangchi Yuan, Zhaoxuan Tan, Fengran Mo and Meng Jiang	1992
<i>REAP: Towards Effective Training-Free Chemical Reasoning with Explicit Atomic Priors</i>	
Mingxu Zhang, Dazhong Shen, Qi Zhang and Ying Sun	2037
<i>Towards Self-Improving Error Diagnosis in Multi-Agent Systems</i>	
Jiazheng Li, Emine Yilmaz, Bei Chen and Thu Le	2063
<i>MulDimIF: A Multi-Dimensional Constraint Framework for Evaluating and Improving Instruction Following in Large Language Models</i>	
Junjie Ye, Caishuang Huang, Zhuohan Chen, Wenjie Fu, Chenyuan Yang, Leyi Yang, Yilong Wu, Peng Wang, Meng Zhou, Xiaolong Yang, Tao Gui, Qi Zhang, Zhongchao Shi, Jianping Fan and Xuanjing Huang	2078
<i>Text2Mem: A Unified Memory Operation Language for Memory Operating System</i>	
Leo Wang, Lihai Yang, Boyu Chen, Kerun Xu, Gongyi Zou, Bo Tang, Feiyu Xiong, Siheng Chen and Zhiyu li	2105
<i>AMO-Bench: Large Language Models Still Struggle in High School Math Competitions</i>	
Junlin Liu, Shengnan An, Shuang Zhou, Dan Ma, Yehao Lin, Xinxuan Lv, Xuanlin Wang, Xiaoyu Li, Ziwen Wang, Xuezhi Cao and Xunliang Cai	2120
<i>Solver-Independent Automated Problem Formulation via LLMs for High-Cost Simulation-Driven Design</i>	
Yuchen Li, Handing Wang, Bing Xue, Mengjie Zhang and Yaochu Jin	2138
<i>When Facts Change: Temporal Knowledge Conflict Resolution in LLMs</i>	
Jonas Wallat, Wolfgang Nejdl and Sandipan Sikdar	2154
<i>SCRIPT: A Subcharacter Compositional Representation Injection Module for Korean Pre-Trained Language Models</i>	
SungHo Kim, Juhyeong Park, Eda Atalay and SangKeun Lee	2185
<i>PibE-MPP: A Play-it-by-Ear Masking Performance Plug-in for LLMs</i>	
Mengwei Wang, Simin Niu, Xun Liang, Yuefeng Ma, Sensen Zhang, Jiawei Yang, Shichao Song, Hanyu Wang and Huayi Lai	2214
<i>RoleCDE: Benchmarking and Mitigating Role–Alignment Trade-offs in Role-Playing Agents</i>	
Huayi Lai, Shichao Song, Simin Niu, Hanyu Wang, Jiawei Yang, Zhouxing Wang, Zhiqiang Yin and Xun Liang	2226
<i>PolicyLLM: Towards Excellent Comprehension of Public Policy for Large Language Models</i>	
Han Bao, Penghao Zhang, Yue Huang, Zhengqing Yuan, Yanchi Ru, SU Rui, Yujun Zhou, Xiangqi Wang, Kehan Guo, Nitesh V Chawla, Yanfang Ye and Xiangliang Zhang	2249
<i>BOOKAGENT: Orchestrating Safety-Aware Visual Narratives via Multi-Agent Cognitive Calibration</i>	
Bo Gao, Chang Liu, Yuyang Miao, SiYuan Ma and Ser-Nam Lim	2275
<i>Feedback-Driven Tool-Use Improvements in Large Language Models via Automated Build Environments</i>	
Junjie Ye, Changhao Jiang, Zhengyin Du, Yufei Xu, Xuesong Yao, Zhiheng Xi, Xiaoran Fan, Qi Zhang, Tao Gui, Xuanjing Huang and Jiecao Chen	2293

<i>Do Domain-specific Experts exist in MoE-based LLMs?</i>	
Giang Do, Hung Le and Truyen Tran	2324
<i>AutoUE: Automated Generation of 3D Games in Unreal Engine via Multi-Agent Systems</i>	
Lei Yin, Wentao Cheng, Zhida Qin, Tianyu Huang, Yidong Li and Gangyi Ding	2341
<i>Comprehensive Benchmarking of Long-Form Speech Generation in Diverse Scenarios</i>	
Changhao Pan, Rui Yang, Han Wang, Zhuan Zhou, Xuming He, Wenxiang Guo, Ziyue Jiang, Ruiqi Li, Yu Zhang, Chenyuhao Wen, Ke Lei, Xiang Yin, Jingyu Lu, Zhiyuan Zhu and Zhou Zhao	2365
<i>CAST: Achieving Stable LLM-based Text Analysis for Data Analytics</i>	
Jinxiang Xie, Zihao Li, Wei He, Rui Ding, Shi Han and Dongmei Zhang	2401
<i>WavAlign: Enhancing Intelligence and Expressiveness in Spoken Dialogue Models via Adaptive Hybrid Post-Training</i>	
Yifu Chen, Shengpeng Ji, Qian Chen, Tianle Liang, Yangzhuo Li, Ziqing Wang, Wen Wang, Jingyu Lu, Haoxiao Wang, Xueyi Pu, Fan Zhuo and Zhou Zhao	2421
<i>OmniCompliance-100K: A Multi-Domain, Rule-Grounded, Real-World Safety Compliance Dataset</i>	
Wenbin Hu, Huihao Jing, Haochen Shi, Changxuan Fan, Haoran Li and Yangqiu Song	2445
<i>AutoMonitor-Bench: Evaluating the Reliability of LLM-Based Misbehavior Monitor</i>	
Shu Yang, Jingyu Hu, Tong Li, Hanqi Yan, Wenxuan Wang and Di Wang	2464
<i>Adaptive and Representative Multi-Interest Modeling for Recommendation with Large Language Model</i>	
Ziyan Wang, Yingpeng Du, Tianjun Wei, Haoyan Chua, Jieyi Bi, Jie Zhang and Zhu Sun	2484
<i>MathMixup: Boosting LLM Mathematical Reasoning with Difficulty-Controllable Data Synthesis and Curriculum Learning</i>	
Xuchen Li, Jing Chen, Xuzhao Li, Hao Liang, Xiaohuan Zhou, Taifeng Wang and Wentao Zhang	2497
<i>The Model Agreed, But Didn't Learn: Diagnosing Surface Compliance in Large Language Models</i>	
Xiaojie Gu, Ziying Huang, Weicong Hong, Jian Xie, Renze Lou and Kai Zhang	2514
<i>AWARE: Agentic Knowledge Warehousing for Contextual Intelligence</i>	
Hongjin Qian, Siqi Bao, Zhao Cao and Zheng Liu	2530
<i>Alignment Tuning for Large Language Models: A Data-Centric Lens on Alignment Data Pipelines</i>	
Hwanjun Song	2541
<i>Purging the Gray Zone: Latent-Geometric Denoising for Precise Knowledge Boundary Awareness</i>	
Hao An, Yibin Lou, Jiayi Guo and Yang Xu	2562
<i>Temporal Token Matters: Investigating and Interpreting the Consistency of Temporal Ordering in Large Language Models</i>	
Zhen Yang, Xinyue Zhang, Ping Jian, Chengzhi Li, Zhongbin Guo, Jiaping Feng and Wenpeng Lu	2577
<i>Towards LLM Agents for Earth Observation</i>	
Chia Hsiang Kao, Wenting Zhao, Cheryl Lam, Aarush Umap, Shreelekha Revankar, Samuel Speas, Snehal Bhagat, Rajeev Datta, Cheng Perng Phoo, Utkarsh Mall, Carl Vondrick, Kavita Bala and Bharath Hariharan	2597
<i>Jamendo-MT-QA: A Benchmark for Multi-Track Comparative Music Question Answering</i>	
Junyoung Koh, Jaeyun Lee, Soo Yong Kim, Gyu Hyeong Choi, Jung In Koh, Jordan Phillips, Yeonjin Lee and Min Song	2612

<i>Chinese Live-Streaming E-Commerce Morph Resolution: Datasets and Methods</i> Jipeng Qiang, Jiahao Zhu, Yi Zhu and Chaowei Zhang	2632
<i>MemoBrain: Executive Memory as an Agentic Brain for Reasoning</i> Hongjin Qian, Zhao Cao and Zheng Liu	2646
<i>Exploring the Capability Boundaries of LLMs in Mastering of Chinese Chouxiang Language</i> Dianqing Lin, Tian Lan, Jiali Zhu, Jiang Li, Wei Chen, Xu Liu, Aruukhan, Xiangdong Su, Hongxu Hou and Guanglai Gao	2663
<i>RECOR: Reasoning-focused Multi-turn Conversational Retrieval Benchmark</i> Mohammed Ali, Abdelrahman Abdallah, Amit Agarwal, Hitesh Laxmichand Patel and Adam Jatowt	2688
<i>Data Selection for Multi-turn Dialogue Instruction Tuning</i> Bo Li, Shikun Zhang and Wei Ye	2724
<i>AgentGC: Evolutionary Learning-based Lossless Compression for Genomics Data with LLM-driven Multiple Agent</i> Sun Hui, Ding Yanfeng, Huidong Ma, Chang Xu, Keyan Jin, Lizheng Zu, Cheng Zhong, Xiaoguang Liu, Gang Wang and Wentong Cai	2740
<i>Think Outside the Policy: In-Context Steered Policy Optimization</i> Hsiu-Yuan Huang, Chenming Tang, Weijie Liu, Clive Bai, Saiyong Yang and Yunfang Wu ..	2758
<i>Where CoT Reasoning Commits: Entropy Traces Identify Interpretable Attention Heads</i> Tianhe Zhang, Yonghong Deng, Ping Jian, Zhen Yang, Boyang Wang and Xinyue Zhang ...	2777
<i>From Coarse to Fine: Benchmarking and Reward Modeling for Writing-Centric Generation Tasks</i> Qingyu Ren, Tianjun Pan, Xingzhou Chen and Xuhong Wang	2796
<i>DentalGPT: Incentivizing Multimodal Reasoning in Dentistry</i> Zhenyang Cai, Jiaming Zhang, Junjie Zhao, Ziyi Zeng, Yanchao Li, Liang Jingyi, Junying Chen, Yunjin Yang, Jiajun You, Shuzhi Deng, Xieruiqiii, Yuanting Chen, Xiangyi Feng, Jianquan Li, Liangyi Chen, Junwen Wang, Shan Jiang and Benyou Wang	2811
<i>SARA: Unlocking Multilingual Knowledge in Mixture-of-Experts via Semantically Anchored Routing Alignment</i> Tianyu Dong, Yangyang Liu, Jiang Zhou, Xinwei Wu, Xiaohu Zhao, Hao Wang, Heng Liu, Linlong Xu, Longyue Wang, Weihua Luo, Shaolin Zhu and Deyi Xiong	2830
<i>Eliciting Medical Reasoning with Knowledge-enhanced Data Synthesis: A Semi-Supervised Reinforcement Learning Approach</i> Haolin Li, Shuyang Jiang, Ruipeng Zhang, Jiangchao Yao, Ya Zhang and Yanfeng Wang ...	2847
<i>AlphaQT-Bench: Diagnosing the Gap between Financial Code Generation and Quantitative Reasoning in LLMs</i> Sichun Luo, Yi Huang, Shichang Meng, Fengyuan Liu, Mukai Li, Qinghua Yao, Zefa Hu, Junlan Feng and Qi Liu	2871
<i>Fine-tuning with Hierarchical Prompting for Robust Propaganda Classification Across Annotation Schemas</i> Lukas Stähelin, Veronika Solopova, Max Upravitelev, David Kaplan, Premtim Sahitaj, Ariana Sahitaj, Charlott Jakob, Sebastian Möller and Vera Schmitt	2887

<i>ComfyFlow: Benchmarking LLMs for AIGC Workflow Generation</i>	
Zhenran Xu, Yiyu Wang, Yunxin li, Muyang Ye, Yangxue, Kai Chen, Longyue Wang, Weihua Luo, Baotian Hu and Min Zhang	2903
<i>SEA-Guard: Culturally Grounded Multilingual Safeguard for Southeast Asia</i>	
Panuthep Tasawong, Jian Gang Ngui, Alham Fikri Aji, Trevor Cohn and Peerat Limkonchotiwat	2917
<i>Prompt Injection in Automated Résumé Screening with Large Language Models: Single and Multi-Injection Settings</i>	
Preet Baxi, Jiannan Xu, Jane Yi Jiang and Stefanus Jasin	2942
<i>FuseSearch: Learning Adaptive Parallel Execution for Efficient Code Localization</i>	
Ke Xu, Siyang Xiao, Ming Liang, Yichen Yu, Zhixiang Wang, Jingxuan Xu, Dajun Chen, Wei Jiang and Yong Li	2954
<i>Revisiting Self-Play Preference Optimization: On the Role of Prompt Difficulty</i>	
Yao Xiao, Jung-jae Kim, Roy Ka-Wei Lee and Lidong Bing	2968
<i>Multi-Agent Comedy Club: Investigating Community Discussion Effects on LLM Humor Generation</i>	
Shiwei Hong, Lingyao Li, Ethan Z. Rong, Chenxinran Shen and Zhicong Lu	2981
<i>ComfyUI-RI: Exploring Reasoning Models for Workflow Generation</i>	
Zhenran Xu, Yiyu Wang, Yangxue, Longyue Wang, Weihua Luo, Kaifu Zhang, Baotian Hu and Min Zhang	2999
<i>Hard to Be Heard: Phoneme-Level ASR Analysis of Phonologically Complex, Low-Resource Endangered Languages</i>	
V.S.D.S.Mahesh Akavarapu, Michael Daniel and Gerhard Jäger	3014
<i>Beyond Surface-Level Pattern Trap: LLM Agents for Faster and Smarter Cross-Architecture Code Migration</i>	
Weijia Li, KE Gao, Pengfei Chen, Jiajie Li, Xinyu Wang, Yiran Le, Yize Wu and Ling Li . .	3029
<i>PairCoder: Pair Programming-Inspired Two-Agent Collaboration for Code Generation</i>	
Junhao Chen, Xiang Li, Yibin Xu, Yuehan Cui, Fangsheng Weng, Hao Zhao, Fei Ma and Qi Tian	3043
<i>OptiVerse: A Comprehensive Benchmark towards Optimization Problem Solving</i>	
Xinyu Zhang, Boxuan Zhang, Yuchen Wan, Lingling Zhang, YiXing Yao, Bifan Wei, Yaqiang Wu and Jun Liu	3059
<i>DiagnosisArena: Benchmarking Diagnostic Reasoning for Large Language Models</i>	
Yakun Zhu, Zhongzhen Huang, Linjie Mu, Yutong Huang, Wei Nie, Jiaji Liu, Shaoting Zhang, Pengfei Liu and Xiaofan Zhang	3074
<i>AMA: Adaptive Memory via Multi-Agent Collaboration</i>	
Weiquan Huang, Zixuan Wang, Hehai Lin, Sudong Wang, Bo Xu, Qian Li, Beier Zhu, Linyi Yang and Chengwei Qin	3099
<i>WESR: A Benchmark and Strong Baseline for Word-level Event-Speech Recognition</i>	
Chenchen Yang, Kexin Huang, Liwei Fan, Qian Tu, Botian Jiang, Dong Zhang, Linqi Yin, Shimin Li, Zhaoye Fei, Qinyuan Cheng and Xipeng Qiu	3121
<i>Applicability Condition Extraction for Therapeutic Drug-Disease Relations</i>	
Guanting Luo, Noriki Nishida, Yuji Matsumoto and Yuki Arase	3135

<i>ACR: Adaptive Context Refactoring via Context Refactoring Operators for Multi-Turn Dialogue</i>	
Jiawei Shen, Jia Zhu, Hanghui Guo, Weijie Shi, Yue Cui, Qingyu Niu, Guoqing Ma, Jingjiang Liu, Yidan Liang, Yilin Wang, Shimin Di and Jiajie Xu	3149
<i>ALGOGEN: Tool-Generated Verifiable Traces for Reliable Algorithm Visualization</i>	
Liaokunpeng, Yuexiao Ma, Yisheng Lin, Hualin Zeng, Xiawu Zheng and Rongrong Ji	3168
<i>PivotAttack: Rethinking the Search Trajectory in Hard-Label Text Attacks via Pivot Words</i>	
Yuzhi Liang, Shiliang Xiao, Jingsong Wei, Qiliang Lin and Xia Li	3190
<i>MiMIC: Mitigating Visual Modality Collapse in Universal Multimodal Retrieval While Avoiding Semantic Misalignment</i>	
Li Juan, Chuanghao Ding, Xujie Zhang and Cam-Tu Nguyen	3208
<i>Towards Proactive Personalization through Profile Customization for Individual Users in Dialogues</i>	
Xiaotian Zhang, Yuan Wang, Ruizhe Chen, Zeya Wang, Runchen Hou and Zuozhu Liu	3221
<i>Beyond Stochastic Exploration: What Makes Training Data Valuable for Agentic Search</i>	
Chuzhan Hao, Wenfeng Feng, Guochao Jiang, Guofeng Quan, Guohua Liu and Yuewei Zhang	3241
<i>Can LLMs Hear the Dogwhistle?</i>	
Yifan Liu, Yi Lin, Xinwei Guo, Ziwei Wang, Jiabin Zhang, Guanhua Chen, Haiyan Wu, Xiangyu Zhao, Xin Yao and Xuetao Wei	3256
<i>HiPrune: Hierarchical Attention for Efficient Token Pruning in Vision-Language Models</i>	
Jizhihui Liu, Guangdao Zhu, Feiyi Du, Niu Lian, Jun Li, Bin Chen, Weili Guan and Yaowei Wang	3274
<i>MIThinker: A Plug-and-Play Policy-Optimized Thinker For Motivational Interviewing Counseling</i>	
Yizhe Yang, Palakorn Achananuparp, Heyan Huang, Jing Jiang and Ee-Peng Lim	3292
<i>The Art of (Mis)alignment: How Fine-Tuning Methods Effectively Misalign and Realign LLMs in Post-Training</i>	
Rui Zhang, Hongwei Li, Yun Shen, Xinyue Shen, Wenbo Jiang, Guowen Xu, Yang Liu, Michael Backes and Yang Zhang	3329
<i>ADaPT: Token-Level Decoupling for Efficient Large Reasoning Models</i>	
Tingyun li, Zishang Jiang, Jinyi Han, Xinyi Wang, Sihang Jiang, Han Xia, Zhaoqian Dai, Ma Shuguang, Fei Yu, Jiaqing Liang and Yanghua Xiao	3355
<i>Talk to Your Slides: High-Efficiency Slide Editing via Language-Driven Structured Data Manipulation</i>	
Kyudan Jung, Hojun Cho, Jooyeol Yun, Soyoung Yang, Jaehyeok Jang and Jaegul Choo	3370
<i>Verifiable Parameterization of Bayesian Networks from Scientific Literature: Unlocking Unstructured Empirical Evidence</i>	
Jonas Gottal and Florian Matthes	3400
<i>PRIME: Ultra-Low-Rank Principal-Residual Model Merging</i>	
Seung-Ho Lee, Kyungsu Lee, Bazarvaani Zuchi, Jeongmin Ahn, Insuk Seo, Donghyeon Jeon, Inho Kang and Seung-Hoon Na	3415
<i>Don't Tell the Answer, Truly Guide the Reasoning During RL Rollouts</i>	
Xinyi Wang, Jinyi Han, Zishang Jiang, Tingyun li, Jiaqing Liang, Sihang Jiang, Zhaoqian Dai, Ma Shuguang, Fei Yu and Yanghua Xiao	3437

<i>Mask Tokens as Prophet: Fine-Grained Cache Eviction for Efficient dLLM Inference</i>	
Jianuo Huang, Yaojie Zhang, Yicun Yang, Benhao Huang and Linfeng Zhang	3456
<i>SPaCe: Unlocking Sample-Efficient Large Language Models Training With Self-Pace Curriculum Learning</i>	
Van Dai Do, Manh Nguyen, Svetha Venkatesh and Hung Le	3480
<i>Crossroads of Optimization under Uncertainty: How to Choose the Optimal Model</i>	
Chengxi She, Zhiqiang Chen, Xingyu Lu, Caihua Chen, Piaoyang Zhao and Xuedong Wang	3508
<i>Know Thy Enemy: Securing LLMs Against Prompt Injection via Diverse Data Synthesis and Instruction-Level Chain-of-Thought Learning</i>	
Zhiyuan Chang, Mingyang Li, Yuekai Huang, Ziyou Jiang, Xiaojun Jia, Qian Xiong, Junjie Wang, Zhaoyang Li and Qing Wang	3540
<i>ModularMoE: Fast LLM Customization with Parameter-Sharing Mixture-of-Experts for Low-Resource Settings</i>	
Jiaxing Liu, Qi Qi, Haifeng Sun, Dunjun Li, Zirui Zhuang, Bo He, Xiang Yang, Cong Liu, Jianxin Liao and Jingyu Wang	3562
<i>UCGRec: User-Centric Graph Learning for LLM-based Sequential Recommendation</i>	
HanBeul Kim, CheolWon Na, Suyoung Bae, YunSeok Choi and Jee-Hyong Lee	3576
<i>HeteroRAG: A Heterogeneous Retrieval-Augmented Generation Framework for Medical Vision Language Tasks</i>	
Zhe Chen, Yusheng Liao, Zhiyuan Zhu, Haolin Li, Hongcheng Liu, Yanfeng Wang and Yu Wang	3592
<i>Subgraph-Guided Executable Logical Form Generation for Knowledge Base Question Answering</i>	
Yuhang Tian, Dandan Song, Zhijing Wu, Changzhi Zhou, Jun Yang, Huipeng Ma, Chenhao Li, Luan Zhang, Yading Li, Xudong Li, Shenxi Liu and Jing Jiang	3614
<i>Towards Reliable Paper Contributions Annotation in the ACL Rolling Review</i>	
Julien Aubert-Bédouchaud, Florian Boudin, Akiko Aizawa, Beatrice Daille and Richard Dufour	3636
<i>Memory Dial: A Training Framework for Controllable Memorization in Language Models</i>	
Xiangbo Zhang and Ali Emami	3654
<i>SEAD: Self-Evolving Agent for Multi-Turn Service Dialogue</i>	
Yuqin Dai, Ning Gao, Wei Zhang, Jie Wang, Luozichen, Ruiyuan Wu, Jinpeng Wang and Chaozheng Wang	3674
<i>DataSciBench: An LLM Agent Benchmark for Data Science</i>	
Dan Zhang, Sining Zhoubian, Min Cai, Fengzu Li, Lekang Yang, Wei Wang, Tianjiao Dong, Ziniu Hu, Jie Tang and Yisong Yue	3685
<i>PVPO: Pre-Estimated Value-Based Policy Optimization for Agentic Reasoning</i>	
Wenfeng Feng, Penghong Zhao, Guochao Jiang, Chuzhan Hao, Guohua Liu and Yuewei Zhang	3729
<i>Momoka-RAG: MCTS-Organized Mapping of Knowledge Associations for Long-Document Retrieval Augmented Generation</i>	
Wenyu Tao, Xiaofen Xing, Zeliang Li and Xiangmin Xu	3749

<i>A3: Android Agent Arena for Mobile GUI Agents with Essential-State Procedural Evaluation</i>	
Yuxiang Chai, Shunye Tang, Han Xiao, Weifeng Lin, Hanhao Li, Jiayu Zhang, Liang Liu, Pengxiang Zhao, Guangyi Liu, Guozhi Wang, Shuai Ren, Rongduo Han, Haining Zhang, Siyuan Huang and Hongsheng Li	3774
<i>Intrinsic Mutual Information as a Modulator for Preference Optimization</i>	
Peng Liao, Peijia Zheng, Lingbo Li, Shangsong Liang and Lin Chen	3790
<i>EVGeoQA: Benchmarking LLMs on Dynamic, Multi-Objective Geo-Spatial Exploration</i>	
Jianfei Wu, Zhichun Wang, Zhensheng Wang and Zhiyu He	3809
<i>FinMRAGBench: A Realistic and Complex Benchmark for Multi-Modal RAG in Financial Document Analysis</i>	
Shouqing Yang, Qi Zhang, Yuhang Yang, Ruikang Xu, Yuwei Hou, Zhulin Jia, Lirong Gao, Haobo Wang, Jinglei Chen, Jiexiang Wang, Sheng Guo, Bo Zheng and Gang Chen	3822
<i>Apeiron: A Scalable LLM-agentic Framework for Autonomous Full-lifecycle Demand-optimized Application Synthesis</i>	
Junyan Cheng, Ankit Srivastava, Jessie Zeng, Milenko Drinic and Jack W. Stokes	3868
<i>Eval-RAR: Evaluation-Driven Retrieval-Augmented Reasoning via Reinforcement Learning</i>	
Heng Yu, Rui Li, Qi Liu, Wenjun Feng, Junfeng Kang and Yi Zhan	3900
<i>LUMINA: Long-horizon Understanding for Multi-turn Interactive Agents</i>	
Amin Rakhsha, Thomas Hehn, Pietro Mazzaglia, Fabio Valerio Massoli, Arash Behboodi and Tribhuvanesh Orekondy	3913
<i>Travel on the ICD Tree: Benchmarking Agentic Reasoning for ICD Coding from Chinese Electronic Medical Records</i>	
Xinjie Xu, Yongqi Fan, Shuang-shuang Chen, Qi Ye, Weibin Guo and Xinxuan Hu	3927
<i>RiT: Rubrics-in-Thinking Reinforcement Learning for Improved Reasoning in Large Language Models</i>	
Xiaobin Tian, Shuai Yuan, Muyun Ding, Haonan Chen and Xiaoxi Jiang	3944
<i>From If-Statements to ML Pipelines: Revisiting Bias in Code-Generation</i>	
Minh Duc Bui, Xenia Heilmann, Mattia Cerrato, Manuel Mager and Katharina Von Der Wense	3958
<i>SEA-SafeguardBench: Culturally Grounded Safety Benchmark for Southeast Asian Languages</i>	
Panuthep Tasawong, Jian Gang Ngui, Alham Fikri Aji, Trevor Cohn and Peerat Limkonchotiwat	3973
<i>COCOGEC: Counterfactual Generation for Robust Grammatical Error Correction</i>	
Qianyu Wang, Xiaoman Wang, Yuanyuan Liang, Xinyuan Li and Yunshi Lan	4004
<i>DRP: Distilled Reasoning Pruning with Mathematical Skill-aware Step Decomposition for Efficient Large Reasoning Models</i>	
Yuxuan Jiang, Dawei Li and Francis Ferraro	4020
<i>Bridging the Pose-Semantic Gap: A Cascade Framework for Text-Based Person Anomaly Search</i>	
Zequan Xie, Guijin Luo, Chuxin Wang, Sihang Cai, Tao Jin, Zhou Zhao and Yixuan Tang ...	4040
<i>TemplateRL: Structured Template-Guided Reinforcement Learning for LLM Reasoning</i>	
Jinyang Wu, Chonghua Liao, Mingkuan Feng, Shuai Zhang, Zhengqi Wen, Haoran Luo, Ling Yang, Huazhe Xu and Jianhua Tao	4050

<i>Dual-Reasoner: Bridging Interleaved Atomicity and Streaming Latency via Thinking-while-Talking</i> Yangzhuo Li, Shengpeng Ji, Yifu Chen, Tianle Liang, Haoyu Yang, Junboli, Jun Fang, Lin Li and Qingyang Hong	4081
<i>FBS: Modeling Native Parallel Reading inside a Transformer</i> Tongxi Wang	4106
<i>HqKV: Towards Hybrid Quantization and Eviction for KV Cache in Long-Context LLM Inference</i> He Wang, Yu Gu, Fangfang Li, Zhigang Wang, Zhenghao Liu, Ning Wang, Xiaohua Li and Ge Yu	4138
<i>Task Matters: Knowledge Requirements Shape LLM Responses to Context–Memory Conflict</i> Kaiser Sun, Fan Bai and Mark Dredze	4154
<i>Adaptive Data Collection for Latin-American Community-sourced Evaluation of Stereotypes (LACES)</i> Guido Ivetta, Pietro Palombini, Sofia Martinelli, Marcos J Gomez, M Emilia Echeveste, Sunipa Dev, Vinodkumar Prabhakaran and Luciana Benotti	4177
<i>MR-ALIGN: Meta-Reasoning Informed Factuality Alignment for Large Reasoning Models</i> Xinming Wang, Jian Xu, Bin Yu, Sheng Lian, Yi Chen, Boran Wang, Yingjian Zhu, Hongzhu Yi, Hong-Ming Yang, Han Hu, Cheng-Lin Liu and Xu-Yao Zhang	4191
<i>Chameleons and Guardians: Unveiling the Divergence in Personality Plasticity and Cognitive Resistance across LLMs</i> Jiaxin Wu, Jie Zhao, Huangxinrong, Cuihong Zhang and Xiuzhu Wu	4215
<i>Entropy Scheduling in Reinforcement Learning for Large Language Models</i> Xingjin Wang, Howe Tissue, Lu Wang, Linjing Li and Daniel Dajun Zeng	4239
<i>TopoDIM: One-shot Topology Generation of Diverse Interaction Modes for Multi-Agent Systems</i> Rui Sun, Jie Ding, Chenghua Gong, Tianjun Gu, Yihang Jiang, Juyuan Zhang, Liming Pan and Linyuan Lü	4252
<i>AVA: Attentive VLM Agent for Mastering StarCraft II</i> Weiyu Ma, Yuqian Fu, Zecheng Zhang, Bernard Ghanem and Guohao Li	4270
<i>Structure-Aware Quantized Retrieval for Long-Document Question Answering</i> Hui Huang, Julien Velcin and Yacine Kessaci	4291
<i>GRPO-CARE: Consistency-Aware Reinforcement Learning for Multimodal Reasoning</i> Yi Chen, Yuying Ge, Rui Wang, Yixiao Ge, Junhao Cheng, Ying Shan and Xihui Liu	4305
<i>Silencing the Guardrails: Inference-Time Jailbreaking via Dynamic Contextual Representation Ablation</i> Wenpeng Xing, Moran Fang, Guangtai Wang, Changting Lin and Meng Han	4321
<i>Diffusion with Truncated Blocks: Fast and High-Quality Text Generation using Truncated Block Generation</i> Yuyan Zhou, Weiyu Chen and James Kwok	4335
<i>GigaCheck: Detecting LLM-generated Content via Object-Centric Span Localization</i> Irina Tolstykh, Aleksandra Tsybina, Sergey Yakubson, Aleksandr Gordeev, Vladimir Dokholyan and Maksim Kuprashevich	4349
<i>Chain-in-Tree: Back to Sequential Reasoning in LLM Tree Search</i> Xinzhe Li	4365

<i>Addressing Overthinking in Large Vision-Language Models via Gated Perception-Reasoning Optimization</i>	Xingjian Diao, Zheyuan Liu, Chunhui Zhang, Weiyi Wu, Keyi Kong, Lin Shi, Kaize Ding, Soroush Vosoughi and Jiang Gui.....	4393
<i>CSMCIR: CoT-Enhanced Symmetric Alignment with Memory Bank for Composed Image Retrieval</i>	Zhipeng Qian, Zihan Liang, Yufei Ma, Ben Chen, Huangyu Dai, Yiwei Ma, Jiayi Ji, Chenyi Lei, Han Li and Xiaoshuai Sun	4411
<i>RepoDistill: Distilling Repository Knowledge through Compression-Aware Budget Allocation and Policy Optimization</i>	Xin Yin, Zixiang Ding, Yiang Zhang, Qiang Wang, Rui Wang, Chao Ni and Zhe Cui	4425
<i>Finding RELIEF: Shaping Reasoning Behavior without Reasoning Supervision via Belief Engineering</i>	Chak Tou Leong, Dingwei Chen, Heming Xia, Qingyu Yin, Sunbowen Lee, Jian Wang and Wenjie Li	4444
<i>Bridging SFT and RL: Dynamic Policy Optimization for Robust Reasoning</i>	Taojie Zhu, Dongyang Xu, Ding Zou, Sen Zhao, Qiaobo Hao, Zhiguo Yang and Yonghong He	4468
<i>FrontCoder: Scaling Visual Fidelity in Front-End Code Generation</i>	Jun Feng, Jian Yang, Wei Zhang, Jing Wang, Keyi Chen, Xiaokun Yang, Weicheng Gu, Yihang Lou, Yan Bai and Xianglong Liu	4485
<i>Bridging Cognition and Affect: Emotion-Aware Opinion Summarization using LLMs</i>	Arnav Attri and Anuj Attri	4506
<i>The Double Bind: Revisiting Preprinting and Peer Review Two Years After the Removal of the ACL Anonymity Period</i>	A Pranav, Shane Storks and Anne Lauscher	4566
<i>TinyAlign: Boosting Lightweight Vision-Language Models by Mitigating Modal Alignment Bottlenecks</i>	Yuanze Hu, Xinyu Wang, Zhichao Yang, Gen Li, Ye Qiu, Zhaoxin Fan, Yifan Sun, Wenjun wu, Jin Dong and Xiaotie Deng	4585
<i>Heterogeneity in Formal Linguistic Competence of Language Models: Is Data the Real Bottleneck?</i>	H S V N S Kowndinya Renduchintala and Sumit Bhatia	4598
<i>Real Men are Tough: Evaluating Gender Bias and Sensitivity to Masculinity Norms in LLMs</i>	Elisa Leonardelli, Camilla Casula, Boglarka Nyul and Sara Tonelli.....	4609
<i>RISER: Orchestrating Latent Reasoning Skills for Adaptive Activation Steering</i>	Wencheng Ye, Xiaoyang Yuan, Yi Bin, Hengyu Jin, Liang Peng, Pengpeng Zeng and Heng Tao Shen.....	4627
<i>ReCoT-NER: Enhancing Zero-Shot Named Entity Recognition through Chain-of-Thought Prompting and Recall-Oriented Loss Optimization</i>	Dabin Fu and Fanghong Zhang.....	4645
<i>SyncThink: A Training-Free Strategy to Align Inference Termination with Reasoning Saturation</i>	Gengyang Li, Wang Cai, Yifeng Gao and Yunfang Wu.....	4657
<i>Reason in Chains, Learn in Trees: Self-Rectification and Grafting for Multi-turn Agent Policy Optimization</i>	Yu Li, Sizhe Tang and Tian Lan	4673

<i>FlashMem: Distilling Intrinsic Latent Memory via Computation Reuse</i> Yubo Hou, Zhisheng Chen, Tao Wan and Zengchang Qin	4687
<i>QueueEDIT: Structural Self-Correction for Sequential Model Editing in LLMs</i> Taolin Zhang, Haidong Kang, Dongyang Li, Qizhou Chen, Xiaofeng He, Chengyu Wang and Richang Hong	4706
<i>CAGenMol: Condition-Aware Diffusion Language Model for Goal-Directed Molecular Generation</i> Yanting Li, Zhuoyang Jiang, Enyan Dai, Lei Wang, Wen-Cai Ye and Li Liu	4720
<i>MAC-Reasoner: A Multi-Agent Collaborative Framework for Enhancing Logical Reasoning in Large Language Models</i> Yehua Lin, Liping Zheng and Yin Chen	4739
<i>Once Correct, Still Wrong: Counterfactual Hallucination in Multilingual Vision-Language Models</i> Basel Mousi, Fahim Dalvi, Shammur Absar Chowdhury, Firoj Alam and Nadir Durrani	4763
<i>DIFFA-2: A Practical Diffusion Large Language Model for General Audio Understanding</i> Jiaming Zhou, Xuxin Cheng, Shiwan Zhao, Yuhang Jia, Cao Liu, Ke Zeng, Xunliang Cai and Yong Qin	4789
<i>Beyond Marginal Distributions: A Framework to Evaluate the Representativeness of Demographic-Aligned LLMs</i> Tristan Williams, Franziska Weeber, Sebastian Padó and Alan Akbik	4806
<i>CORES: Code-Oriented Reasoning for Complex Text-to-SQL and Generalizable TableQA</i> Meng Zhang, Ruochun Jin, Yuanxi Peng, Wenjing Yang, Haotian Wang, Liting Sun, Kun Hu, Silin Yang and Zhang Ke-di	4827
<i>Do Not Guess, Verify: Logic-Guided Adaptive Reasoning for Multimodal Misinformation Detection</i> Kun Huang, Rui Qiu, Xiaoming Li and Salah Uddin	4848
<i>Hi-ZFO: Hierarchical Zeroth- and First-Order LLM Fine-Tuning via Importance-Guided Tensor Selection</i> Feihu Jin and Ying Tan	4862
<i>MCP-Guard: A Multi-Stage Defense-in-Depth Framework for Securing Model Context Protocol in Agentic AI</i> Wenpeng Xing, Zhonghao Qi, Yupeng Qin, Yilin Li, Caini Chang, Jiahui Yu, Changting Lin, Zhenzhen Xie and Meng Han	4877
<i>Half-S: Halving the Scale for Near-Lossless 4-Bit LLM Training</i> Jinyang Du, Ruihao Gong, Linghan Ai, Zining Wang, Yunke Peng, Yao Wang, Lei Yan, Wxuefei, Yaoyuan Wang, Jinyang Guo, Dahua Lin and Xianglong Liu	4890
<i>F-Actor: Controllable Conversational Behavior in Full-Duplex Models</i> Maike Züfle, Ondrej Klejch, Nicholas Sanders, Jan Niehues, Alexandra Birch and Tsz Kin Lam	4904
<i>APEX: Learning Adaptive Priorities for Multi-Objective Alignment in Vision-Language Generation</i> Dongliang Chen, Xinlin Zhuang, Junjie Xu, Luojuan Xie, Zehui Wang, Jiayi Zhuang, Haolin Yang, Liang Dou, Xiao He, Xingjiao Wu and Ying Qian	4922
<i>DART: Mitigating Harm Drift in Difference-Aware LLMs via Distill-Audit-Repair Training</i> Ziwen Pan, Zihan Liang, Jad Kabbara and Ali Emami	4940

<i>Experience-Driven Reflective Co-Evolution of Prompts and Heuristics for Autonomous Algorithm Design</i>	Yihong Liu, Junyi Li, Hongyu Lu, Xin Zhao and Ji-Rong Wen	4981
<i>Information Representation Fairness in Long-Document Embeddings: The Peculiar Interaction of Positional and Language Bias</i>	Elias Schuhmacher, Andrianos Michail, Juri Opitz, Rico Sennrich and Simon Clematide . . .	4996
<i>RA-RRG: Multimodal Retrieval-Augmented Radiology Report Generation with Key Phrase Extraction</i>	Jonggwon Park, Byungmu Yoon, Soobum Kim and Kyoyun Choi	5029
<i>Context-Value-Action Architecture for Value-Driven Large Language Model Agents</i>	TianZe Zhang, Sirui Sun, Yuhang Xie, Xin Zhang, Zhiqiang Wu and Guojie Song	5049
<i>Uncovering Strategic Egoism Behaviors in Large Language Models</i>	Yaoyuan Zhang, Zonghao Ying, Aishan Liu, Jian Yang, Tianlin Li, Yaodong Yang and Xianglong Liu	5074
<i>EMA: An Episodic Memory Agent for Efficient and Selective Memory</i>	Hongyi Lan, Jiaqi Song, Zhengjia Zhong, Hui Li, Hong Liu, Xianming Lin and Rongrong Ji	5088
<i>ChildTalk: A Multi-Dialect Chinese Child Speech Corpus with Full-Length Child–Caregiver Conversations for Speech Recognition</i>	Jiaming Zhou, Yujie Guo, Shiwan Zhao, Yao Lu, Jianye Wang, Haoqin Sun, Hui Wang and Yong Qin	5103
<i>SrDetection: A Self-Referential Framework for Data Leakage Detection in Code Large Language Models</i>	Shuaimin Li, Liyang Fan, Zeyang li, Zhuoyue Wan, Yufang Lin, Shiwen Ni, Feiteng Fang, Hamid Alinejad-Rokny, Yuanfeng Song, Kun Jing, Chen Jason Zhang and Min Yang	5117
<i>Memo-SQL: Structured Decomposition and Experience-Driven Self-Correction for Training-Free NL2SQL</i>	Yang Zerui, Weichuan Wang, Yanwei Xu, Linqi Song, Yudai Matsuda, Wei Han and Bo Bai	5130
<i>Do We Always Need Query-Level Workflows? Rethinking Agentic Workflow Generation for Multi-Agent Systems</i>	Zixu Wang, Bingbing Xu, Yige Yuan, Huawei Shen and Xueqi Cheng	5149
<i>How Chain-of-Thought Works? Tracing Information Flow from Decoding, Projection, and Activation</i>	Hao Yang, Qinghua Zhao, Lei Li, Lingyi Meng and Mengda Yu	5166
<i>A Survey of Retentive Network</i>	Haiqi Yang, Zhiyuan Li, Yi Chang and Yuan Wu	5200
<i>RAG-KT: Cross-platform Explainable Knowledge Tracing with Multi-view Fusion Retrieval Generation</i>	Zhiyi Duan, Hongyu Yuan and Rui Liu	5217
<i>Conformal Event Prediction with Temporal Knowledge Graph</i>	Cheng Hu, Cong Cao, Fangfang Yuan, Diandian Guo, Pin Xu, Yu Liu and Yanbing Liu	5233
<i>Infinite Babble: Inflating 3D Vision-Language Model Inference Overhead via Adversarial Geometric Perturbation</i>	Shuoyang Sun, Jiaxin Hong, Yv Zhang, Kuofeng Gao, Hao Fang, Fan Mo, Bin Chen and Shu-Tao Xia	5249
<i>Uncertainty-Calibrated Elastic Alignment for Multimodal Sentiment Analysis with Missing Modalities</i>	Kang He, Yuzhe Ding, Rao Fu, Yukang Feng, Kaipeng Zhang, Yiming Liu, Fei Li, Chong Teng and Donghong Ji	5268

<i>On the Representation Geometry of LoRA Model Merging</i>	
Chenyang Lu, Jiaru Li, Jinman Zhao, Xinran Chen, Yining Wang, Renyi Cai, Yuchen Li and Chao He	5289
<i>MCLE-Mol: Empowering LLM with Molecular Comprehension and Low-Cost Continual Evolution for Interpretable Property Prediction</i>	
Zhili Pu, Lantian Zhang, Hao Duan, Zhixing Zhang, Keyun Zhu, Yongqi Fan, Ruihui Hou, Tong Ruan and Yun Tang	5305
<i>MTR-DuplexBench: Towards a Comprehensive Evaluation of Multi-Round Conversations for Full-Duplex Speech Language Models</i>	
Zhang He, Wenqian Cui, Haoning Xu, Xiao-Hui Li, Lei Zhu, Haoli Bai, Ma Shaohua and Irwin King	5334
<i>Generating Effective CoT Traces for Mitigating Causal Hallucination</i>	
Yiheng Zhao and Jun Yan	5352
<i>SAE-FiRE: Enhancing Earnings Surprise Predictions Through Sparse Autoencoder Feature Selection</i>	
Huopu Zhang, Yanguang Liu, Miao Zhang, Zirui He and Mengnan Du	5363
<i>From Implicit Graph Encoding to Explicit Evidence: A Training-Free LLM Framework for Temporal Knowledge Graph Reasoning</i>	
Guo Tang, Ke Cheng, Huiming Fan, Heng Chang, Wenxiang Zheng, Xianhao Ou, Junjia Xiang, Ming Liu, Yujun Zhou, Li Lanyu and Bing Qin	5383
<i>LLM-Driven Multi-Perspective Location Completion for Next Location Prediction</i>	
Pengxiang Lan, Enneng Yang, Yuliang Liang, Jianzhe Zhao, Linying Jiang and Guibing Guo	5406
<i>WikiSeeker: Rethinking the Role of Vision-Language Models in Knowledge-Based Visual Question Answering</i>	
Yingjian Zhu, Xinming Wang, Kun Ding, Ying Wang, Bin Fan and Shiming Xiang	5429
<i>Understanding Conflicts in Multi-Objective Alignment through Reward Consistency</i>	
Zhihao Xu, Yongqi Tong, Xin Zhang, Jun Zhou and Xiting Wang	5450
<i>MetaMem: Evolving Meta-Memory for Knowledge Utilization through Self-Reflective Symbolic Optimization</i>	
Haidong Xin, Xinze Li, Zhenghao Liu, Yukun Yan, Shuo Wang, Cheng Yang, Yu Gu, Ge Yu and Maosong Sun	5473
<i>GeoLAN: Geometric Learning of Latent Explanatory Directions in Large Language Models</i>	
Tianyu Pan and Damon L. Woodard	5493
<i>SepSeq: A Training-Free Framework for Long Numerical Sequence Processing in LLMs</i>	
Jie Sun, Yu Liu, Lu Han, Qiwen Deng, Xiang Shu, Yang Xiao, Lintao Ma, Xingyu Lu, Jun Zhou, Pengfei Liu, Jiancan Wu and Xiang Wang	5522
<i>Ghost in the Shell: Synonym-Aware Logit Shaping Fingerprint for Copyright Protection of Large Vision-Language Models</i>	
Xiaofan Zheng, Xinghao Wang and Xiaojun Wan	5538
<i>Look Twice before You Leap: A Rational Framework for Localized Adversarial Text Anonymization</i>	
Donghang Duan, Xu Zheng, Yuefeng He, Chong Mu, Leyi Cai and Lizong Zhang	5556
<i>K-MetBench: A Multi-Dimensional Benchmark for Fine-Grained Evaluation of Expert Reasoning, Locality, and Multimodality in Meteorology</i>	
Soyeon Kim, Cheongwoong Kang, Myeongjin Lee, Eun-Chul Chang, Lee Jaedeok and Jaesik Choi	5574

<i>V-GameGym: Visual Game Generation for Code Large Language Models</i>	
Wei Zhang, Jian Yang, Renshuai Tao, Linzheng Chai, Shuyue Guo, Jiajun Wu, Xiaoming Chen, Ganqu Cui, Ning Ding, Xander Xu, HU Wei and Bowen Zhou	5613
<i>UCoder: Unsupervised Code Generation by Internal Probing of Large Language Models</i>	
Jiajun Wu, Jian Yang, Wei Zhang, Linzheng Chai, Yuchi Ma, Ensheng Shi, Yuqing Ma, Zhoujun Li and Xianglong Liu	5642
<i>Deciphering Cultural Representations in Large Language Models via Sparse Autoencoders</i>	
Chenye Zou, Difan Jiao and Lijie Hu	5656
<i>Fin-Bias: Comprehensive Evaluation for LLM Decision-Making under human bias in Finance Domain</i>	
Xiaoyu Hu and Jinman Zhao	5678
<i>Duplicate-Aware Controlled Code Generation: Enhancing Copyright Protection with Targeted Re-ordering Beam Search in LLMs</i>	
Junbo Fu, Guoshuai Zhao, Linkang Yang, Yunqi Mi and Xueming Qian	5695
<i>Graph-Based Chain-of-Thought Pruning for Reducing Redundant Reflections in Reasoning LLMs</i>	
Hongyuan Yuan, Xinran He, Run Shao, Bolei He, Xianwei Xue, Mengke Chen, Qitong Pan, Haiwei Wang and Haifeng Li	5708
<i>Too Good to be Bad: On the Failure of LLMs to Role-Play Villains</i>	
Zihao Yi, Qingxuan Jiang, Ruotian Ma, Xingyu Chen, Qu Yang, Mengru Wang, Fanghua Ye, Ying Shen, Zhaopeng Tu, Xiaolong Li and Liefeng Bo	5724
<i>AudioPrivacy: Parallel Audio Dataset for Speaker Profiling with Diverse Audio Types and Rich Attributes</i>	
Jiabei He, Yanzhe Zhang, Jiaming Zhou, Hui Wang, Haoqin Sun and Yong Qin	5735
<i>Federated LoRA Fine-Tuning with Pipelined Error-Mitigated Aggregation and Matrix-Wise Freezing</i>	
Haoran Wang, Xiong Wang, Yuqing Li, Jing Chen, Junyi Zhang, Nan Yan, Kun He and Wei Wang	5749
<i>RTCFake: Speech Deepfake Detection in Real-Time Communication</i>	
Jun Xue, Zhuolin Yi, Yihuan Huang, Yanzhen Ren, Yujie Chen, Cunhang Fan, Zicheng Su, Yongcheng Zhang and Bo Cai	5763
<i>Progress Ratio Embeddings: An Impatience Signal for Robust Length Control in Neural Text Generation</i>	
Ivanhoé Botcazou, Tassadit Amghar, Sylvain Lamprier and Frédéric Saubion	5776
<i>LeakDojo: Decoding the Leakage Threats of RAG Systems</i>	
Maosen Zhang, Jianshuo Dong, Lu Boting, Li Wenyue, Xiaoping Zhang, Tianwei Zhang and Han Qiu	5790
<i>CodeWiki: Evaluating AI's Ability to Generate Holistic Documentation for Large-Scale Codebases</i>	
Anh Nguyen Hoang, Minh Le-Anh, Bach Le and Nghi D. Q. Bui	5812
<i>ARISE: An Adaptive Resolution-Aware Metric for Test-Time Scaling Evaluation in Large Reasoning Models</i>	
Zhangyue Yin, Qiushi Sun, Zhiyuan Zeng, Zhiyuan Yu, Qipeng Guo, Xuanjing Huang and Xipeng Qiu	5828
<i>Breaking the Static Graph: Context-Aware Traversal for Graph-Based RAG</i>	
Kwun Hang Lau, Fangyuan Zhang, Boyu Ruan, Yingli Zhou, Qintian Guo, Ruiyuan Zhang and Xiaofang Zhou	5849

<i>POP: Prefill-Only Pruning for Efficient Large Model Inference</i> Junhui He, Zhihui Fu, Jun Wang and Qingan Li	5864
<i>Guardian-as-an-Advisor: Advancing Next-Generation Guardian Models for Trustworthy LLMs</i> Yue Huang, Haomin Zhuang, Jiayi Ye, Han Bao, Yanbo Wang, Hang Hua, Siyuan Wu, Pin-Yu Chen and Xiangliang Zhang	5878
<i>StructBreak: Structural Cognitive Overload-Induced Safety Failures in MLLMs</i> Yang Luo, Liu Xinran, TianTian Ji, Zhiyi Yin, Lingyun Peng and Shuyu Li	5901
<i>Role-Sensitive Neurons: A Neuron-Level Gain Control Mechanism for Confidence Steering</i> Peiwen Huang, Chih-Hao Hsu, Tzu-Hung Huang and Shou-De Lin	5924
<i>StreamingEval: A Unified Evaluation Framework towards Realistic Streaming Video Understanding</i> Tang Guowei, Tianwen Qian, Huanran Zheng, Wang Yifei and Xiaoling Wang	5945
<i>CogGen: A Cognitively Inspired Recursive Framework for Deep Research Report Generation</i> Kuo Tian, Pengfei Sun, Zhen Wu, Junran Ding and Xinyu Dai	5961
<i>Debate to Align: Reliable Entity Alignment through Two-Stage Multi-Agent Debate</i> Cunda Wang, Ziying Ma, Po Hu, Weihua Wang and Feilong Bao	5989
<i>Exploring Multilingual Pre-trained Language Model for Aspect-based Sentiment Analysis</i> Ye Wang, Ruijun Jiang, Zhongqing Wang and Guodong Zhou	6011
<i>Mid-Think: Training-Free Intermediate-Budget Reasoning via Token-Level Triggers</i> Van Yang, Shouren Wang, Debargha Ganguly, Xinpeng Li, Chaoda Song, Vikash Singh, Vipin Chaudhary and Xiaotian Han	6024
<i>SSR-Zero: Simple Self-Rewarding Reinforcement Learning for Machine Translation</i> Wenjie Yang, Mao Zheng, Mingyang Song, Zheng Li and Sitong Wang	6039
<i>HisDoc-OCR: Restoring Visual Grounding in MLLMs for Chinese Historical Document OCR</i> Jiahuan Cao, Yongxin Shi, Zeyu Shan, Zhengyang Lu and Lianwen Jin	6053
<i>One LLM Does Not Simulate All Students: Ability-Aware Student Simulation via Cognitive Diagnosis Guided LLM Assignment</i> Huixing Que, Qi Liu, Weibo Gao and Zhenya Huang	6067
<i>Interactive Learning for LLM Reasoning</i> Hehai Lin, Shilei Cao, Sudong Wang, Haotian Wu, Minzhi Li, Linyi Yang, Juepeng Zheng and Chengwei Qin	6085
<i>Style Amnesia: Investigating Speaking Style Degradation and Mitigation in Multi-Turn Spoken Language Models</i> Yu-Xiang Lin, Cheng-Han Chiang and Hung-yi Lee	6109
<i>NeuroSym-Cal: Bridging the Reasoning-Execution Gap in Code Generation via Hierarchical Calibration</i> Peiyang Liu, Yining Wang, Youru Li, Long Li, Zhi Cai and Wei Ye	6132
<i>GOBench: Stage-Wise Diagnostics and the Visual Paradox in Multimodal Graph Optimization</i> Yinghao Chen, Wantong Xie, Shuli Zeng, Sijia Zhang, Xiaotian Pan, Feng Wu and Xiangyang Li	6144

<i>Sparse Growing Transformer: Training-Time Sparse Depth Allocation via Progressive Attention Looping</i>	
Yao Chen, Yilong Chen, Yinqi Yang, Junyuan Shang, Zhenyu Zhang, Zefeng Zhang, Shuaiyi Nie, Shuohuan Wang, Yu Sun, Hua Wu, Haifeng Wang and Tingwen Liu	6168
<i>Computational Narrative Understanding for Expressive Text-to-Speech</i>	
Gaspard Michel, Elena V. Epure and Christophe Cerisara	6194
<i>DataSage: Multi-agent Collaboration for Insight Discovery with External Knowledge Retrieval, Multi-role Debating, and Multi-path Reasoning</i>	
Xiaochuan Liu, Yuanfeng Song, Xiaoming Yin and Xing Chen	6216
<i>Eliminating Out-of-Domain Recommendations in LLM-based Recommender Systems: A Unified View</i>	
Hao Liao, Jiwei Zhang, Jianxun Lian, Wensheng Lu, Mingqi Wu, Shuowangg, Yong Zhang, Yitian Huang, Mingyang Zhou and Rui Mao	6251
<i>Evian: Towards Explainable Visual Instruction-tuning Data Auditing</i>	
Zimu Jia, Mingjie Xu, Andrew Estornell and Jiaheng Wei	6272
<i>Iterative Self-Correction for Text-Driven Person Re-Identification with Large Vision-Language Models</i>	
Guijin Luo, Zequn Xie, Sihang Cai, Chuxin Wang, Zhou Zhao and Yixuan Tang	6292
<i>Breaking Contextual Inertia: Reinforcement Learning with Single-Turn Anchors for Stable Multi-Turn Interaction</i>	
Xingwu Chen, Zhanqiu Zhang, Steven Y. Guo and Difan Zou	6302
<i>Label Words as Local Task Vectors in In-Context Learning</i>	
Bowen Zheng, Ming Ma, Zhongqiao Lin and Tianming Yang	6319
<i>CoGR-MoE: Concept-Guided Expert Routing with Consistent Selection and Flexible Reasoning for Visual Question Answering</i>	
Xiyin Zeng, Yi Lu and Hao Wang	6333
<i>Which Pieces Does Unigram Tokenization Really Need?</i>	
Sander Land and Yuval Pinter	6351
<i>Step Potential Advantage Estimation: Harnessing Intermediate Confidence and Correctness for Efficient Mathematical Reasoning</i>	
Fei Wu, Zhenrong Zhang, Qikai Chang, Jianshu Zhang, Quan Liu and Jun Du	6361
<i>When depth is redundant: Efficient transformer-based speech anti-spoofing</i>	
Hoan My Tran, Damien Lolive, Aghilas Sini, Arnaud Delhay, Pierre-Francois Marteau and David Guennec	6380
<i>GCoT-Decoding: Unlocking Deep Reasoning Paths for Universal Question Answering</i>	
Guanran Luo, Wentao Qiu, Zhongquan Jian, Meihong Wang and Qingqiang Wu	6398
<i>PerMemSafe: Benchmarking Implicit Personalized Safety of Long Horizon Self-Evolving Agents</i>	
Hengyu An, Minxi Li, Naen Xu, Chunyi Zhou, Xiaogang Xu, Tianyu Du, Jinbao Li and Shouling Ji	6415
<i>TagRAG: Tag-guided Hierarchical Knowledge Graph Retrieval-Augmented Generation</i>	
Wenbiao Tao, Xinyuan Li, Yunshi Lan and Weining Qian	6434
<i>Danger Depends on the Mind: A Theory-of-Mind Grounded Dataset and Model for Context-Dependent Dangerous Speech</i>	
Yuanchen Shi, Longyin Zhang, Guodong Zhou and Fang Kong	6457

<i>One Token Is Enough: Improving Diffusion Language Models with a Sink Token</i> Zihou Zhang, Zheyong Xie, Li Zhong, Haifeng Liu and Shaosheng Cao	6479
<i>Thunder-KoNUBench: A Corpus-Aligned Benchmark for Korean Negation Understanding</i> Sungmok Jung, Yeonkyoung So, Joonhak Lee, Sangho Kim, Yelim Ahn and Jaejin Lee	6491
<i>Progra: Progress-Aware Reinforcement Learning for Multi-Turn Function Calling</i> Huacan Chai, Zijie Cao, Maolin Ran, Yingxuan Yang, Jianghao Lin, Xin Peng, Hairui Wang, Renjie Ding, Ziyu Wan, Muning Wen, Weiwen Liu, Weinan Zhang, Fei Huang and Ying Wen . . .	6519
<i>Towards Modern Topic Models: A Survey of Taxonomies and Paradigm Shifts from Algorithm-Centric to LLM-Centered Topic Analysis</i> Xuan Xu, Zhongliang Yang, Haolun Li, Rui Tian, Beilin Chu, J Song, Yu Li, Shaolin Tan and Linna Zhou	6536
<i>Beyond Templates: Dynamic Adaptation of Reasoning Demonstrations via Feasibility-Aware Exploration</i> Yong Wu, Weihang Pan, Ke Li, Chen Binhui, Ping Li and Binbin Lin	6562
<i>CFlowPsyD: An Analysis-Enhanced Dataset for Asynchronous Psychological Counseling through Self-Optimizing Multi-Agent Framework</i> Donghao Li, Yifan Deng, Jinta Weng, Xingsheng Zhang, Chenxu Niu, Jingyuan Tian, Heyan Huang and Yue Hu	6578
<i>The Thin Line Between Comprehension and Persuasion in LLMs</i> Adrian de Wynter and Tangming Yuan	6603
<i>Evaluating Machine Translation Datasets for Low-Web Data Languages: A Gendered Lens</i> Hellina Hailu Nigatu, Bethelhem Yemane Mamo, Bontu Fufa Balcha, Debora Taye Tesfaye, Elbethel Daniel Zewdie, Ikram Behiru Nesiru, Jitu Ewnetu Hailu and Senait Mengesha Yayo	6632
<i>Select2Reason: Efficient Instruction-Tuning Data Selection for Long-CoT Reasoning</i> Cehao Yang, Xueyuan Lin, Xiaojun Wu, Chengjin Xu, Xuhui Jiang, Honghao Liu, Hui Xiong and Jian Guo	6656
<i>Curriculum Learning based Hierarchical Scoring and Analysis Framework for Question Answering Task Evaluation</i> Qiong Wu, Tan Yue, Jianxin Liang, Zhen Li, Kai He, Shuai Zhao and Dongyan Zhao	6672
<i>TrigReason: Trigger-Based Collaboration between Small and Large Reasoning Models</i> Yi Zhao, Yajuan Peng, Cam-Tu Nguyen, Zuchao Li, Xiaoliang Wang, Xiaoming Fu and Hai Zhao	6700
<i>TURTLEAI: Benchmarking Multimodal Models for Visual Programming in Turtle Graphics</i> Chao Wen, Jacqueline Staub and Adish Singla	6719
<i>Calibrated Progressive Distillation: Co-Designing Curriculum and Target Mixing for Knowledge Distillation of Large Language Models</i> Mengxiang Zhang and Lingyuan Liu	6757
<i>Faithful-First Reasoning, Planning, and Acting for Multimodal LLMs</i> Junxian Li, Xinyue Xu, Sai Ma, Di Zhang and Sichao Li	6777
<i>SEAD: A Surrogate-free Label-only Membership Inference Attack against Pre-trained LLMs with Semantic-Aware Density</i> Biao Yi, Jiahao Li, Yiming Li, Yu He, Baolei Zhang, Zheli Liu and Dacheng Tao	6794

Faithfulness-Aware Uncertainty Quantification for Fact-Checking the Output of Retrieval-Augmented Generation

Ekaterina Fadeeva, Aleksandr Rubashevskii, Dzianis Piatrashyn, Roman Vashurin, Shehzaad Dhuliawala, Artem Shelmanov, Timothy Baldwin, Preslav Nakov, Mrinmaya Sachan and Maxim Panov
6814

AGGC: Adaptive Group Gradient Clipping for Stabilizing Large Language Model Training

Zhiyuan Li, Yuan Wu and Yi Chang 6837

MADRA: Multi-Agent Debate for Risk-Aware Embodied Planning

JunJian Wang, Lidan Zhao and Xi Sheryl Zhang 6852

Navigating Large-Scale Document Collections: MuDABench for Multi-Document Analytical QA

Zhanli Li, Yixuan Cao, Lvzhou Luo and Ping Luo 6877

CORBA: Contagious Recursive Blocking Attacks on Multi-Agent Systems Based on Large Language Models

Zhenhong Zhou, Zherui Li, Jie Zhang, Yuanhe Zhang, Kun Wang, Yang Liu and Qing Guo . 6899

Afrispeech Semantics: Evaluating Audio-Semantic Reasoning in Spoken Language Models Across Domains and Accents

Chibuzor Okocha and Christan Grant 6909

MoPrune: Scene-Guided Motion-Aware Token Pruning for Efficient Video Large Language Models

Wenhao Hong, Ziyang Wang, Yixin Zhang and Zilei Wang 6929

SCMAPR: Self-Correcting Multi-Agent Prompt Refinement for Complex-Scenario Text-to-Video Generation

Chengyi Yang, Pengzhen Li, Jiayin Qi, Aimin Zhou, Ji Wu and Ji Liu 6942

LoReFact: Bridging the Logic Gap in Fact-Checking

Qiming Xie, Wenjie Zheng, Xiangqing Shen and Rui Xia 6974

ODTQA-FoRe: An Open-Domain Tabular Question Answering Dataset for Future Data Forecasting and Reasoning

Zhensheng Wang, Xiaole Liu, Wenmian Yang, Kun Zhou, Yiquan Zhang and Weijia Jia . . . 6997

Rectified Sparse Attention for Efficient Long-Sequence Generation

Yutao Sun, Tianzhu Ye, Li Dong, Yuqing Xia, Jian Chen, Yizhao Gao, Shijie Cao, Jianyong Wang and Furu Wei 7023

Knowing-but-Doing: Diagnosing and Defending Role-Play-Driven LLMs Jailbreaks via Moral Disengagement

Haiming Qin, Jianxun Lian, Qimin Zhong, Mingyang Zhou, Hao Liao and Naipeng Chao . . 7035

Generative-to-Discriminative Test-Time Adaptation via Manifold-Aware Diffusion and Bayesian Distillation

Boyun Zhang, Zequn Xie, Fangming Feng, Zihan Zhang, Yongbo He, Chuxin Wang, Sihang Cai, Tao Jin and Qifei Zhang 7052

Evaluating Memory Capability in Continuous Lifelog Scenario

Jianjie Zheng, Zhichen Liu, Zhanyu Shen, Jingxiang Qu, Guanhua Chen, Yile Wang, Yang Xu, Yang Liu and Sijie Cheng 7063

GoT-RI: Internalizing Graph-of-Thought via Structural Reinforcement for High-Density Reasoning

Zuchao Li, Qiwei Li, Yao Yao, Hai Zhao, Lefei Zhang and Bo Du 7090

<i>Hey, That’s My Data! Token-Only Dataset Inference in Large Language Models</i> Chen Xiong, Zihao Wang, Rui Zhu, Tsung-Yi Ho, Pin-Yu Chen, Jingwei Xiong and Haixu Tang	7105
<i>ToxiTrace: Gradient-Aligned Training for Explainable Chinese Toxicity Detection</i> Boyang Li, Hongzhe Shou, Yuanyuan Liang, JingBin Zhang and Fang Zhou	7121
<i>Three Minds, One Legend: Jailbreak Large Reasoning Model with Adaptive Stacked Ciphers</i> Nguyen Viet Anh, Shiqian Zhao, Gia Dao, Runyi Hu, Yi Xie, Xiaobao Wu and Anh Tuan Luu	7139
<i>From Imitation to Discrimination: Progressive Curriculum Learning for Robust Web Navigation</i> Peng Chuang, Wei Zhang, Renshuai Tao, Xinhao Zhang and Jian Yang	7160
<i>Parallelism and Generation Order in Masked Diffusion Language Models: Limits Today, Potential Tomorrow</i> Yangyang Zhong, Yanmei Gu, Zhengqing Zang, Xiaomeng Li, Yuqi Ding, Xibei Jia, Yuting Shen, Zhenzhong Lan, Liwang Zhu, Weiping Liu, Junlin Zhou, Haisheng Liu, Zhong Xin Yu, Pengxin Luo, Donglian Qi, Yunfeng Yan and Junbo Zhao	7178
<i>Understanding New-Knowledge-Induced Factual Hallucinations in LLMs: Analysis and Interpretation</i> Renfei Dang, Peng Hu, Zhejian Lai, Changjiang Gao, Min Zhang and Shujian Huang	7219
<i>Attribution-Based Analysis and Optimization of Modular Agentic Workflows</i> Yingxuan Yang, Bo Huang, Siyuan Qi, Chao Feng, Haoyi Hu, Yuxuan Zhu, Jinbo Hu, Haoran Zhao, Ziyi He, Xiao Liu, ZongYu Wang, Muning Wen, Lin Qiu, Xuezhi Cao, Xunliang Cai, Yong Yu and Weinan Zhang	7247
<i>Synthetic Eggs in Many Baskets: The Impact of Synthetic Data Diversity on LLM Fine-Tuning</i> Max Schaffelder and Albert Gatt	7265
<i>SolidCoder: Bridging the Mental-Reality Gap in LLM Code Generation through Concrete Execution</i> Woojin Lee and Jin-Xia Huang	7294
<i>EgoMemory: Memory-Augmented Personalized Retrieval for Long-Context Egocentric Video</i> Yuanmin Tang, Jue Zhang, Xiaoting Qin, Jing Yu, Meikang Qiu, Gaopeng Gou, Gang Xiong, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang and Qi Wu	7317
<i>PERM: Psychology-grounded Empathetic Reward Modeling for Large Language Models</i> Chengbing Wang, Wuqiang Zheng, Yang Zhang, Fengbin Zhu, Junyi Cheng, Yi Xie, Wenjie Wang and Fuli Feng	7350
<i>Remedy-R: Generative Reasoning for Machine Translation Evaluation without Error Annotations</i> Shaomu Tan, Ryosuke Mitani, Ritvik Choudhary, Qiyu Wu, Toshiyuki Sekiya and Christof Monz	7374
<i>VizoMem: A Visual-Textual Memory Framework for Efficient Long-Horizon Reasoning</i> Weijie Liang, Yuanfeng Song, Xing Chen, Caleb Chen Cao, Sirui Han and Yike Guo	7399
<i>More Images, More Problems? A Controlled Analysis of VLM Failure Modes.</i> Anurag Das, Adrian Bulat, Alberto Baldrati, Ioannis Maniadis Metaxas, Bernt Schiele, Georgios Tzimiropoulos and Brais Martinez	7423
<i>EOP-LLM: Energy Oriented Pruning for Large Language Models</i> Guan Huang and Tao Shu	7443

<i>Sentipolis: Emotion-Aware Agents for Social Simulations</i> Chiyuan Fu, Lyuhao Chen, Yunze Xiao, Weihao Xuan, Carlos Busso and Mona T. Diab	7465
<i>Bidirectional Semantic Enhancement for Schema Routing Across Large-Scale Databases</i> Yuyang Wu, Xiaoliang Wang and Cam-Tu Nguyen	7495
<i>DeepPlanner: Scaling Planning Capability for Deep Research Agents via Advantage Shaping</i> Wei Fan, Wenlin Yao, Zheng Li, Feng Yao, Xin Liu, Liang Qiu, Qingyu Yin, Yangqiu Song and Bing Yin	7510
<i>SupChain-Bench: Benchmarking Large Language Models for Real-World Supply Chain Management</i> Shengyue Guan, Yihao Liu and Lang Cao	7526
<i>How Hypocritical Is Your LLM judge? Listener-Speaker Asymmetries in the Pragmatic Competence of Large Language Models</i> Judith Sieker and Sina Zarrieß	7551
<i>CoopQ: Cooperative Game Inspired Layerwise Mixed Precision Quantization for LLMs</i> Junchen Zhao, Ali Derakhshan, Jayden Hyman, Junhao Dong, Sangeetha Abdu Jyothi and Ian Harris	7566
<i>Breaking Language Preference in Multilingual RAG via Language-Controllable Retrieval and Language-Agnostic Reasoning</i> Wenshuai Huo, Xiaocheng Feng, Baohang Li, Chengpeng Fu, Yichong Huang, Hui Wang and Bing Qin	7579
<i>From Scoring to Explanations: Evaluating SHAP and LLM Rationales for Rubric-based Teaching Quality Assessment</i> Ivo Bueno, Babette Bühler, Philipp Stark, Tim Fütterer, Ulrich Trautwein, Dorottya Demszky, Heather Hill and Enkelejda Kasneci	7590
<i>LycheeCluster: Efficient Long-Context Inference with Structure-Aware Chunking and Hierarchical KV Indexing</i> Dongfang Li, Zixuan Liu, Gang Lin, Baotian Hu and Min Zhang	7607
<i>HiRAS: A Hierarchical Multi-Agent Framework for Paper-to-Code Generation and Execution</i> Hanhua Hong, Yizhi LI, Jiaoyan Chen, Sophia Ananiadou, Xiaoli Li, Jung-jae Kim and Chenghua Lin	7624
<i>RoadMapper: A Multi-Agent System for Roadmap Generation of Solving Complex Research Problems</i> Jiacheng Liu, Zichen Tang, Zhongjun Yang, Xinyi Hu, Xueyuan Lin, Linwei Jia, Ruofei Bai, Rongjin Li, Shiyao Peng, Haocheng Gao and Haihong E	7653
<i>On Temperature-Constrained Non-Deterministic Machine Translation: Potential and Evaluation</i> Weichuan Wang, Mingyang Liu, Chen Ma and Linqi Song	7677
<i>ExaGPT: Example-Based Machine-Generated Text Detection for Human Interpretability</i> Ryuto Koike, Masahiro Kaneko, Ayana Niwa, Preslav Nakov and Naoaki Okazaki	7702
<i>Zipage: Maintain High Request Concurrency for LLM Reasoning through Compressed PagedAttention</i> Mengqi Liao, Lu Wang, Chaoyun Zhang, Bo Qiao, Si Qin, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang and Huaiyu Wan	7716
<i>The Overlooked Role of Graded Relevance Thresholds in Multilingual Dense Retrieval</i> Tomer Wullach, Ori Shapira and Amir David Nissan Cohen	7738

<i>Test-Time Scaling in Multimodal Foundation Models: A Comprehensive Survey of Generation and Reasoning</i>	
Cong Wan, Ying He, Zhongzhan Huang and Hefeng Wu	7751
<i>PsyScore: A Psychometrically-Aware Framework for Trait-Adaptive Essay Scoring and ZPD-Scaffolded Feedback</i>	
Wei Xia, Jin Wu, Haoran Shi, Xiangyu Wang and Chanjin Zheng	7768
<i>FinMaster: A Holistic Benchmark for Full-Pipeline Financial Management with Large Language Models</i>	
Junzhe Jiang, Chang Yang, Aixin Cui, Sihan Jin, Yujing Zhang, Yilin Xiao, Ruiyu Wang, Bo Li, Xiao Huang, Danny Dongning Sun and Xinrun Wang	7787
<i>PersonaForge: Psychology-Grounded Dual-Process Architecture for Personality-Consistent Role-Playing Agents</i>	
Jizhou Tong and Sirui Zou	7845
<i>Graph Explorer: Training Faithful KG Agents with Visibility-Grounded Supervision</i>	
Yifeng Chen, Sicheng Wan, Tianyi Zhang and Xuezhou Zhang	7875
<i>Robust Uncertainty Quantification for Self-Evolving Large Language Models via Continual Domain Pretraining</i>	
Xiaofan Zhou and Lu Cheng	7891
<i>Exploring Graph Learning Tasks with Pure LLMs: A Comprehensive Benchmark and Investigation</i>	
Yuxiang Wang, Xinnan Dai, Wenqi Fan and Yao Ma	7909
<i>Beyond Rejection Sampling: Trajectory Fusion for Scaling Mathematical Reasoning</i>	
Jie Deng, Hanshuang Tong, Jun Li, Shining Liang, Ning Wu, Hongzhi Li and Yutao Xie ...	7943
<i>ReasonRec: A Reasoning-Augmented Multimodal Agent for Unified Recommendation</i>	
Yihua Zhang, Mingfu Liang, Jiyan Yang, Rong Jin, Wen-Yen Chen, Yiping Han, Huayu Li, Buyun Zhang, Liang Luo, Luke Simon, Sijia Liu, Tianlong Chen and Xi Liu	7960
<i>Does Chain-of-Thought Reasoning Help Mobile GUI Agents? An Empirical Study</i>	
Li Zhang, Longxi Gao and Mengwei Xu	7981
<i>Asymmetric Relational-Geometry Driven Universal Adversarial Perturbations for Vision-Language Models</i>	
Jiaxin Ye, Weihai Li, Ying Wang, Simeng Qin, Zhitao Zeng and Zikai Xu	7997
<i>PaperScope: A Multi-Modal Multi-Document Benchmark for Agentic Deep Research Across Massive Scientific Papers</i>	
Lei Xiong, Huaying Yuan, Zheng Liu, Zhao Cao and Zhicheng Dou	8015
<i>When Personalization Misleads: Understanding and Mitigating Hallucinations in Personalized LLMs</i>	
Zhongxiang Sun, Yi Zhan, Chenglei Shen, Weijie Yu, Xiao Zhang, Ming He and Jun Xu ...	8041
<i>Law in Silico: Simulating Legal Society with LLM-Based Agents</i>	
Yiding Wang, Yuxuan Chen, Fanxu Meng, Xifan Chen, Xiaolei Yang and Muhan Zhang ...	8061
<i>MT-Video-Bench: A Holistic Video Understanding Benchmark for Evaluating Multimodal LLMs in Multi-Turn Dialogues</i>	
Yaning Pan, Qianqian Xie, Guohui Zhang, Zekun Moore Wang, Yongqian Wen, Yuanxing Zhang, Haoxuan Hu, Zhiyu Pan, Yibing Huang, Zhidong Gan, Yonghong Lin, An Ping, Shihao Li, Yanghai Wang, Tianhao Peng and Jiaheng Liu	8105

<i>Query-Aware Graph Attention for Precise Subgraph Retrieval in Knowledge-Augmented Reasoning</i> Yuanye Xu, Linyi Guo, Yue Zhang and Fu Ning	8127
<i>LIBERTy: A Causal Framework for Benchmarking Concept-Based Explanations of LLMs with Structural Counterfactuals</i> Gilat Toker, Nitay Calderon, Ohad Amosy and Roi Reichart	8149
<i>Contract-Coding: Towards Repo-Level Generation via Structured Symbolic Paradigm</i> Yilin, Lujin Zhao and Yijie Shi	8187
<i>Adaptive Text Anonymization: Learning Privacy-Utility Trade-offs via Prompt Optimization</i> Gabriel Loiseau, Damien Sileo, Damien Riquet, Maxime Meyer and Marc Tommasi	8207
<i>SWE-QA: Can Language Models Answer Repository-level Code Questions?</i> Weihan Peng, Yuling Shi, Yuhang Wang, Xinyun Zhang, Beijun Shen and Xiaodong Gu ...	8230
<i>Efficient PRM Training Data Synthesis via Formal Verification</i> Ryo Kamoi, Yusen Zhang, Nan Zhang, Sarkar Snigdha Sarathi Das, Ranran Haoran Zhang, Wenpeng Yin and Rui Zhang	8246
<i>Revealing the Attention Floating Mechanism in Masked Diffusion Models</i> Xin Dai, Pengcheng Huang, Zhenghao Liu, Shuo Wang, Yukun Yan, Chaojun Xiao, Yu Gu, Ge Yu and Maosong Sun	8266
<i>Balancing Knowledge Breadth and Task Depth for Effective Domain Adaptation Fine-Tuning</i> Mu Zhang, Yuxiang Chu, Guangya Yu, Yongqi Fan, Weiyan Zhang, Hang Hu, Tong Ruan and Jingping Liu	8287
<i>Explain the Flag: Contextualizing Hate Speech Beyond Censorship</i> Jason Liartis, Eirini Kaldeli, Lamprini Gyftokosta, Eleftherios Chelioudakis and Orfeas Menis Mastromichalakis	8305
<i>EnvScaler: Scaling Tool-Interactive Environments for LLM Agent via Programmatic Synthesis</i> Xiaoshuai Song, Haofei Chang, Guanting Dong, Yutao Zhu, Ji-Rong Wen and Zhicheng Dou	8326
<i>Iterative Knowledge Graph Refinement and Integration for Medical Question Answering</i> Haochen Zou and Yongli Wang	8358
<i>Lang2Act: Fine-Grained Visual Reasoning through Self-Emergent Linguistic Toolchains</i> Yuqi Xiong, Chunyi Peng, Zhipeng Xu, Zhenghao Liu, Zulong Chen, Yukun Yan, Shuo Wang, Yu Gu and Ge Yu	8375
<i>Lost in Stories: Consistency Bugs in Long Story Generation by LLMs</i> Junjie Li, Xinrui Guo, Yuhao Wu, Roy Ka-Wei Lee, Hongzhi Li and Yutao Xie	8400
<i>Chinese Court Simulation with LLM-Based Agents System</i> Kaiyuan Zhang, Jiaqi Li, Yueyue Wu, Haitao Li, Cheng Luo, Shaokun Zou, Yujia Zhou, Weihang Su, Yiqun Liu and Qingyao Ai	8429
<i>SPASM: Stable Persona-driven Agent Simulation for Multi-turn Dialogue Generation</i> Han Luo and Guy Laban	8455
<i>Telling Speculative Stories to Help Humans Imagine the Harms of Healthcare AI</i> Xingmeng Zhao, Tongnian Wang, Dan Schumacher, Veronica Rammouz and Anthony Rios	8476
<i>Masking or Mitigating? Deconstructing the Impact of Query Rewriting on Retriever Biases in RAG</i> Agam Goyal, Koyel Mukherjee, Apoorv Saxena, Anirudh Phukan, Eshwar Chandrasekharan and Hari Sundaram	8517

<i>Unifying Inference-Time Planning Language Generation</i>	
Prabhu Prakash Kagitha, Bo Sun, Ishan Desai, Andrew Zhu, Cassie Huang, Manling Li, Ziyang Li and Li Zhang	8531
<i>EmoRes: Toward Adaptive Psychological Support via User-Agnostic Benchmark and Topic-Mining Agent</i>	
Zhengwei Zou, Xuanming Jiang, Baoyi An, Dingyu Nie, Zhengxing Fang, Qingyu Liu, Xueming Qian, Guoshuai Zhao and Zhongyu Yang	8575
<i>Budget-Aware Anytime Reasoning with LLM-Synthesized Preference Data</i>	
Xuanming Zhang, Shwan Ashrafi, Aziza Mirsaidova, Amir H. Rezaeian, Miguel Ballesteros, Lydia Chilton, Zhou Yu and Dan Roth	8587
<i>Mitigating Cultural Bias in LLMs via Multi-Agent Cultural Debate</i>	
Qian Tan, Lei Jiang, Yuting Zeng, Shuoyang Ding and Xiaohua Xu	8600
<i>ToolRM: Towards Agentic Tool-Use Reward Modeling</i>	
Renhao Li, Jianhong Tu, Yang Su, Yantao Liu, Fei Huang, Hamid Alinejad-Rokny, Derek F. Wong, Junyang Lin and Min Yang	8613
<i>MDocRAG-RL: Empowering Multi-Modal Document RAG via Complex Visual Reasoning with Reinforcement Learning</i>	
Zhongyu Wang	8641
<i>MARS: Unleashing the Power of Speculative Decoding via Margin-Aware Verification</i>	
Jingwei Song, Xinyu Wang, Hanbin Wang, Xiaoxuan Lei, Tianyu Shi, Shixin Han, Eric Yang, Xiao-Wen Chang and Lynn Ai	8652
<i>AI, Take the Wheel: What Drives Delegation and Trust in Human-Computer Cooperative Question Answering?</i>	
Maharshi Gor, Yoo Yeon Sung, Yu Hou, Eve Fleisig, Zhu Irene Ying, Tianyi Zhou and Jordan Lee Boyd-Graber	8664
<i>Self-Reinforcing Controllable Synthesis of Rare Relational Data via Bayesian Calibration</i>	
Chongsheng Zhang, Hao Wang, Zelong Yu, Esteban Garces Arias, Julian Rodemann, Zhanshuo Zhang, Qilong Li, Gaojuan Fan, Krikamol Muandet and Christian Heumann	8690
<i>WaveDetect: Robust Framework for Machine-Generated Text Detection via Wavelet Transform</i>	
Zhichen Liu, Kaitong Qin, Linhan He and Yang Xu	8712
<i>AHEAD: Attention Head Energy-Aware Dynamics for Hallucination Mitigation in MLLMs</i>	
Jiale Chang, Ying Li, Siliang Tang and Yueting Zhuang	8728
<i>TopoSHIELD: Reshaping the Flow of Malice via Spatio-Temporal Risk-Aware Topological Evolution in Multi-Agent Systems</i>	
Ruiyang Huang, Chenxi Wang, Tinghe Zhang, Fengrui Liu, Jiayan Sun, Haocheng Wang and Yifan Wu	8740
<i>Diagnosing and Mitigating Sycophancy and Skepticism in LLM Causal Judgment</i>	
Edward Y Chang	8769
<i>SuperWriter: Reflection-Driven Long-Form Generation with Large Language Models</i>	
Yuhao Wu, Yushi Bai, Zhiqiang Hu, Juanzi Li and Roy Ka-Wei Lee	8790
<i>ProxyPrompt: Securing System Prompts against Prompt Extraction Attacks</i>	
Zhixiong Zhuang, Maria-Irina Nicolae, Hui-Po Wang and Mario Fritz	8813

<i>Cogito: A Cognitive Agentic Framework Driven by Dynamic Graph of Thoughts for Financial Report Generation</i>	Chen Lifan, Wei Ding, Jingwen Yang, Xiuze Zhou, Jingan Chen and Fan Lin	8846
<i>From Coarse to Fine: A Multi-Granularity Multimodal Framework for Teacher Sentiment Analysis</i>	Zhiyi Duan, Xiangren Wang, Jiangshan Guan, Bing Jia and Qianli Xing	8861
<i>PBEBench: A Multi-Step Programming by Examples Reasoning Benchmark inspired by Historical Linguistics</i>	Atharva Naik, Prakam, Yash Mathur, Darsh Agrawal, Manav Nitin Kapadnis, Yuwei An, Clayton Marr, Carolyn Rose and David R. Mortensen	8877
<i>T2I-ReasonBench: Benchmarking Reasoning-Informed Text-to-Image Generation</i>	Kaiyue Sun, Rongyao Fang, Chengqi Duan, Xian Liu, Aoxue Li and Xihui Liu	8919
<i>CGBridge: Bridging Code Graphs and Large Language Models for Better Structure-Aware Code Understanding</i>	Zeqi Chen, Zhaoyang Chu, Yi Gui, Feng Guo, Yao Wan and Chuan Shi	8945
<i>Candidate-Aware Retrieval and Reranking for Multiple-Choice Question Answering: Arabic as a Case Study</i>	Yassine Bouziane, Youness Moukafih and Mounir Ghogho	8967
<i>Logical Structure as Knowledge: Enhancing LLM Reasoning via Structured Logical Knowledge Density Estimation</i>	Zhen Bi, Zhenlin Hu, Xueshu Chen, Mingyang Chen, Cheng Deng, Yida Xue, Zhen Wang, Qing Shen, Ningyu Zhang and Jungang Lou	8978
<i>ACSE: An Ancient Character Semantic-Aware Embedding for Large Language Models</i>	Zhihan Zhou, Daqian Shi, Lida Shi, Rui Song, Peiqiang Qiu, Xiaolei Diao and Hao Xu	9000
<i>Reasoning as Gradient: Scaling MLE Agents Beyond Tree Search</i>	Yifei Zhang, Xu Yang, Xiao Yang, Bowen Xian, Qizheng Li, Shikai Fang, Jingyuan Li, Jian Wang, Minrui Xu, Yuge Zhang, Weiqing Liu and Jiang Bian	9013
<i>Speak No Evil, Just Prompt: Low-resource Multilingual Toxic Speech Detection with Audio Language Model</i>	Mingzi Zuo, Lei Zhang, Hailiang Sun, Shengzhi Huo, Changyu Dong, Xin Wang, Bo Wang and Hao Liu	9039
<i>EmotionTalk: An Interactive Chinese Multimodal Emotion Dataset With Rich Annotations</i>	Haoqin Sun, Jinghua Zhao, Xuechen Wang, Shiwan Zhao, Jiaming Zhou, Hui Wang, Xi Yang, Yequan Wang and Yonghua Lin	9054
<i>What Prompts Don't Say: Understanding and Managing Underspecification in LLM Prompts</i>	Chenyang Yang, Yike Shi, Qianou Ma, Michael Xieyang Liu, Christian Kaestner and Tongshuang Wu	9072
<i>AffectCodec: Emotion-Preserving Neural Speech Codec for Expressive Speech Modeling</i>	Jiacheng Shi, Hongfei Du, Xinyuan Song, Y. Alicia Hong, Yanfu Zhang and Ashley Gao	9102
<i>Agentic Conversational Search with Contextualized Reasoning via Reinforcement Learning</i>	Fengran Mo, Yifan Gao, Sha Li, Hansi Zeng, Xin Liu, Zhaoxuan Tan, Xian Li, Jianshu Chen, Dakuo Wang and Meng Jiang	9125

<i>Do Multimodal RAG Systems Leak Data? A Comprehensive Evaluation of Membership Inference and Image Caption Retrieval Attacks</i>	Ali Al-Lawati and Suhang Wang	9139
<i>PED: Route-Decoupled Diagnostics for Persona Consistency in Spoken Agents</i>	Weihao Liu, Junrui Wei, Zhao Zhang and Ju Zhang	9155
<i>Interpretable Semantic Gradients in SSD: A PCA Sweep Approach and a Case Study on AI Discourse</i>	Hubert Plisiecki, Maria Leniarska, Jan Piotrowski and Marcin Zajenkowski	9169
<i>Self-Sum: Teaching an Agent to Decide Itself When and What to Summarize</i>	Hongru Wang, Rui Wang, Jushi Kai, Boyang Xue, Yongqi Li, Shijue Huang, Xiaoteng Ma, Jeff Z. Pan and Amos Storkey	9178
<i>GreekMMLU: A Native-Sourced Multitask Benchmark for Evaluating Language Models in Greek</i>	Yang Zhang, Mersin Konomi, Christos Xypolopoulos, Konstantinos Divriotis, Konstantinos Skianis, Giannis Nikolentzos, Giorgos Stamou, Guokan Shang and Michalis Vazirgiannis	9193
<i>User-Assistant Bias in LLMs</i>	Xu Pan, Jingxuan Fan, Zidi Xiong, Ely Hahami, Jorin Overwiening and Ziqian Xie	9218
<i>ICLAD: In-Context Learning with Comparison-Guidance for Audio Deepfake Detection</i>	Benjamin Shiue-Hal Chou, Yi Zhu and Surya Koppiseti	9242
<i>ReStyle-TTS: Relative and Continuous Style Control for Zero-Shot Speech Synthesis</i>	Haitao Li, Chunxiang Jin, Chenglin Li, Wenhao Guan, Zhengxing Huang and Xie Chen	9257
<i>Who Gets Which Message? Auditing Demographic Bias in LLM-Generated Targeted Text</i>	Tunazzina Islam	9270
<i>FACTS: Table Summarization via Offline Template Generation with Agentic Workflows</i>	Ye Yuan, Mohammad Amin Shabani and Siqi Liu	9298
<i>How Language Models Conflate Logical Validity with Plausibility: A Representational Analysis of Content Effects</i>	Leonardo Bertolazzi, Sandro Pezzelle and Raffaella Bernardi	9316
<i>LLMs Deceive Unintentionally: Emergent Misalignment in Dishonesty from Misaligned Samples to Biased Human-AI Interactions</i>	Xuhao Hu, Wang Peng, Xiaoya Lu, Dongrui Liu, Xuanjing Huang and Jing Shao	9348
<i>AlphaQuanter: An End-to-End Tool-Augmented Agentic Reinforcement Learning Framework for Stock Trading</i>	Zheyue Deng, Weixiang Yan, Changlong Yu and Jiashu Wang	9373
<i>Distribution Shift Alignment Helps LLMs Simulate Survey Response Distributions</i>	Ji Huang, Mengfei LI and Shuai Shao	9395
<i>PhysPRM: A Generative Process Reward Model with Fine-grained Diagnosis for Physics Problem Solving</i>	Yuxuan Dong, Xinyu Zhang, Lingling Zhang, Han Lai, Pengyu Li, Bifan Wei, Yaqiang Wu and Jun Liu	9410
<i>The Mark Fades: Adaptive Evolutionary Paraphrase-based Attack against LLM Watermarks</i>	Yusheng Zhao, Jian Zhao, Tianle Zhang, Feng Wei and Xuelong Li	9428

<i>CoRE: A Fine-Grained Code Reasoning Benchmark Beyond Output Prediction</i>	
Jun Gao, Yun Peng, Qian Qiao, Changhai Zhou, Yuhua Zhou, Shiyang Zhang, Shichao Weng, Zhenchang Xing and Xiaoxue Ren	9446
<i>Empowering Reliable Visual-Centric Instruction Following in MLLMs</i>	
Weilei He, Feng Ju, Zhiyuan Fan, Rui Min and Minhao Cheng	9460
<i>TT-SI: Self-Improving LLM Agents with Test-Time Training</i>	
Emre Can Acikgoz, Cheng Qian, Heng Ji, Dilek Hakkani-Tür and Gokhan Tur	9483
<i>Conflicts Make Large Reasoning Models Vulnerable to Attacks</i>	
Honghao Liu, Chengjin Xu, Xuhui Jiang, Cehao Yang, Shengming Yin, Zhengwu Ma, Lionel Ni and Jian Guo	9509
<i>Linear Script Representations in Speech Foundation Models Enable Zero-Shot Transliteration</i>	
Ryan Soh-Eun Shim, Kwanghee Choi, Calvin Chang, Ming-Hao Hsu, Florian Eichin, Zhizheng Wu, Alane Suhr, Michael A. Hedderich, David Harwath, David R. Mortensen and Barbara Plank	9532
<i>NSL-MT: Linguistically Informed Negative Samples for Efficient Machine Translation in African Low-Resource Languages</i>	
Mamadou K. Keita, Christopher M Homan and Huy Le	9545
<i>EduMARS: Can Vision-Language Models Grade Like Teachers? Benchmarking Multimodal, Rubric-Based Assessment on Chinese K-12 Answers</i>	
Xuan Zhao, Jiashun Chen, Wanting xu, Huiyuan Yan, Chaowei Fang and Xing Wei	9561
<i>MMR-GRPO: Accelerating GRPO-Style Training through Diversity-Aware Reward Reweighting</i>	
Kangda Wei and Ruihong Huang	9584
<i>StoryMI: Steerable Multi-Agent Therapeutic Dialogue Generation</i>	
Qingyu Meng, Min Chen, Dingming Liu, Yifan Mo, Yue Su, Xin Sun, Koen Hindriks and Jiahuan Pei	9606
<i>Disambiguation-Centric Finetuning Makes Enterprise Tool-Calling LLMs More Realistic and Less Risky</i>	
Ashutosh Hathidara, Julien Yu and Sebastian Schreiber	9624
<i>ToolGate: Contract-Grounded and Verified Tool Execution for LLMs</i>	
Yanming Liu, Xinyue Peng, Jiannan Cao, Xinyi Wang, Songhang Deng, Jintao Chen, Jianwei Yin and Xuhong Zhang	9653
<i>PersonaArena: Dynamic Simulation for Evaluating and Enhancing Persona-Level Role-Playing in Large Language Models</i>	
Wenlong Shi, Jianxun Lian, Mingqi Wu, Haiming Qin, Mingyang Zhou, Xing Xie, Naipeng Chao and Hao Liao	9685
<i>Poller: Are LLMs Suitable for Evaluating Poetry Understanding Task?</i>	
Shanshan Wang, Derek F. Wong, Jingming Yao and Lidia S. Chao	9720
<i>No-Worse Context-Aware Decoding: Preventing Neutral Regression in Context-Conditioned Generation</i>	
Yufei Tao and Ameeta Agrawal	9736
<i>Prompting Test-Time Scaling Is A Strong LLM Reasoning Data Augmentation</i>	
Sondos Mahmoud Bsharat and Zhiqiang Shen	9752

<i>GRASS: Gradient-based Adaptive Layer-wise Importance Sampling for Memory-efficient Large Language Model Fine-tuning</i>	
Kaiyuan Tian, Linbo Qiao, Yu Tang, Gongqingjian Jiang, Baihui Liu, Yifu Gao, Xialin Su and Dongsheng Li	9777
<i>ltzGLUE: Luxembourgish General Language Understanding Evaluation</i>	
Alistair Plum, Felicia Körner, Anne-Marie Lutgen, Laura Bernardy, Fred Philipp, Emilia Milano, Nils Rehlinger, Cedric Lothritz, Tharindu Ranasinghe, Barbara Plank and Christoph Purschke ...	9791
<i>AdapThink: Adaptive Thinking Preferences for Reasoning Language Models</i>	
Wenyue Xu, Xu Wan, Wei Wang, Wenqi Huang, Wotao Yin, Shengjie Zhao and Mingyang Sun	9808
<i>Improving the Throughput of Diffusion-based Large Language Models via a Training-Free Confidence-Aware Calibration</i>	
Jucheng Shen, Gaurav Sarkar, Yeonju Ro, Sharath Nittur Sridhar, Zhangyang Wang, Aditya Akella and Souvik Kundu	9826
<i>Same Claim, Different Judgment: Benchmarking Scenario-Induced Bias in Multilingual Financial Misinformation Detection</i>	
Zhiwei Liu, Yupeng Cao, Yuechen Jiang, Mohsinul Kabir, Polydoros Giannouris, Chen Xu, Ziyang Xu, Tianlei Zhu, Md. Tariquzzaman, Triantafillos Papadopoulos, Yan Wang, Lingfei Qian, Xueqing Peng, Zhuohan Xie, Ye Yuan, Saeed Almheiri, Abdulrazzaq Alnajjar, Ming-Bin Chen, Harry Stuart, Paul Thompson, Prayag Tiwari, Alejandro Lopez-Lira, Xue Liu, Jimin Huang and Sophia Ananiadou	9838
<i>MASS-RAG: Multi-Agent Synthesis Retrieval-Augmented Generation</i>	
Xingchen Xiao, Heyan Huang, Runheng Liu and Jincheng Xie	9865
<i>Bayesian Active Learning with Gaussian Processes Guided by LLM Relevance Scoring for Dense Passage Retrieval</i>	
Junyoung Kim, Anton Korikov, Jiazhou Liang, Justin Cui, Yifan Simon Liu, Qianfeng Wen, Mark Zhao and Scott Sanner	9884
<i>Reasoning-Based Refinement of Unsupervised Text Clusters with LLMs</i>	
Tunazzina Islam	9899
<i>Self-Evolving Multi-Agent Systems via Textual Backpropagation</i>	
Xiaowen Ma, Yunpu Ma, Chenyang Lin, Sikuan Yan, Jinhe Bi, Zixuan Cao, Yijun Tian, Volker Tresp and Hinrich Schuetze	9918
<i>Confidence over Time: Confidence Calibration with Temporal Logic for Large Language Model Reasoning</i>	
Zhenjiang Mao, Anirudh Venkat, Artem Bisliouk, Sindhura Kumbakonam Subramanian, Akshat Kothiyal, Saithej Singhu and Ivan Ruchkin	9952
<i>Automated Profile Inference with Language Model Agents</i>	
Yuntao Du, Zitao Li, Bolin Ding, Yaliang Li, Hanshen Xiao, Jingren Zhou and Ninghui Li .	9977
<i>When Debiasing Backfires: Counterintuitive Side Effects of Preprocessing-Based Stereotype Mitigation</i>	
Yahan Zheng, John J. Guerrerio, Soroush Vosoughi and Weicheng Ma	10001
<i>Scaling Laws for Code: Every Programming Language Matters</i>	
Jian Yang, Shuyue Guo, Linzheng Chai, Wei Zhang, Aishan Liu, Chuan Hao, Zhoujun Li, Xin Zhao, Xianglong Liu, Weifeng Lv and Bryan Dai	10022

<i>Beyond the Singular: Revealing the Value of Multiple Generations in Benchmark Evaluation</i> Wenbo Zhang, Hengrui Cai and Wenyu Chen	10033
<i>Safety Is Not Universal: The Selective Safety Trap in LLM Alignment</i> Iago Alves Brito, Walcy Rios, Julia Soares Dollis, Diogo Fernandes Costa Silva and Arlindo Rodrigues Galvão Filho	10044
<i>LLM Prompt Duel Optimizer: Efficient Label-Free Prompt Optimization</i> Yuanchen Wu, Saurabh Verma, Justin Lee, Fangzhou Xiong, Poppy Zhang, Amel Awadelkarim, Xu Chen, Yubai Yuan and Shawndra Hill	10066
<i>FrameNet-Cultures: A Benchmark for Evaluating LLMs via Cross-Cultural Frame Semantics</i> Neda Jamshidi, Anders Sjøgaard and Monica Bianchini	10090
<i>Meta-Tool: Efficient Few-Shot Tool Adaptation for Small Language Models</i> Sachin Kumar	10132
<i>Toward A Digital Twin of U.S. Congress</i> Hayden Helm, Tianyi Chen, Harvey McGuinness, Paige Lee, Brandon Duderstadt and Carey Priebe	10148
<i>OjaKV: Context-Aware Online Low-Rank KV Cache Compression</i> Yuxuan Zhu, David H. Yang, Mohammad Mohammadi Amiri, Keerthiram Murugesan, Tejaswini Pedapati and Pin-Yu Chen	10161
<i>AURORA: Neuro-Symbolic Continual Indexing for Evolving RAG Systems</i> Manoj Saravanan, Rohit Kumar Salla and Ramya Manasa Amancherla	10179
<i>Decide less, communicate more: On the construct validity of end-to-end fact-checking in medicine</i> Sebastian Antony Joseph, Lily Chen, Barry Wei, Michael Mackert, Iain James Marshall, Paul Pu Liang, Ramez Kouzy, Byron C Wallace and Junyi Jessy Li	10196
<i>K-GIP: Diagnosing Logical Fractures in Large Vision-Language Models via Verification Scene Graphs and Sequential Pruning</i> Yujun Hu, Xiaoyu Zhou, Changbo Wang and Gaoqi He	10222
<i>From Static Inference to Dynamic Interaction: A Survey of Streaming Large Language Models</i> Junlong Tong, Zilong Wang, YuJie Ren, Peiran Yin, Hao Wu, Wei Zhang and Xiaoyu Shen	10237
<i>How Adversarial Environments Mislead Agentic AI?</i> Zhonghao Zhan, Huichi Zhou, Zhenhao Li, Peiyuan Jing, Krinos Li and Hamed Haddadi .	10264
<i>GALA: Geometric Data Selection with Strategic Prospecting for Large Language Model Self-training</i> Zhongwei Xie, Ruihao Liao, Zimo Wang, Chong Chen, Xian-Sheng Hua and Xiao Luo ...	10281
<i>PII-VisBench: Evaluating Personally Identifiable Information Safety in Vision Language Models Along a Continuum of Visibility</i> G M Shahariar, Zabir Al Nazi, Md Olid Hasan Bhuiyan and Zhouxing Shi	10294
<i>Locate, Steer, and Improve: A Practical Survey of Actionable Mechanistic Interpretability in Large Language Models</i> Hengyuan Zhang, Zhihao Zhang, Ercong Nie, Mingyang Wang, Zunhai Su, Yiwei Wang, Qianli Wang, Shuzhou Yuan, Xufeng Duan, Qibo Xue, Zeping Yu, Chenming Shang, Xiao Liang, Jing Xiong, Hui Shen, Chaofan Tao, Zhengwu Liu, Senjie Jin, Zhiheng Xi, Dongdong Zhang, Sophia Ananiadou, Tao Gui, Ruobing Xie, Hayden Kwok-Hay So, Hinrich Schuetze, Xuanjing Huang, Qi Zhang and Ngai Wong	10317

<i>MTRAG-UN: A Benchmark for Open Challenges in Multi-Turn RAG Conversations</i>	
Sara Rosenthal, Yannis Katsis, Vraj Shah, Lihong He, Lucian Popa and Marina Danilevsky	10363
<i>Vision-Language Models Mistake Head Orientation for Gaze Direction: Nonverbal Conversation Cues</i>	
Zory Zhang, Pinyuan Feng, Bingyang Wang, Tianwei Zhao, Suyang Yu, Qingying Gao, Hokin Deng, Ziqiao Ma, Yijiang Li and Dezhi Luo	10370
<i>RepoShapley: Shapley-Enhanced Context Filtering for Repository-Level Code Completion</i>	
Yu Huo, Kun Zeng, Siyu Zhang, Yuquan LU, Cheng Yang, Yifu Guo and Xiaoying Tang	10390
<i>Red-Teaming NSFW Image Classifiers as Text-to-Image Safeguards</i>	
Tinghao Xie, Yueqi Xie, Alireza Zareian, Shuming Hu, Felix Juefei-Xu, Xiaowen Lin, Ankit Jain, Prateek Mittal and Li Chen	10413
<i>Model-Agnostic Meta Learning for Class Imbalance Adaptation</i>	
Hanshu Rao, Guangzeng Han and Xiaolei Huang	10442
<i>PAPERMIND: Benchmarking Agentic Reasoning and Critique over Scientific Papers in Multimodal LLMs</i>	
Yanjun Zhao, Tianxin Wei, Jiaru Zou, Xuying Ning, Yuanchen Bei, Lingjie Chen, Simmi Rana, Wendy H. Yang, Hanghang Tong and Jingrui He	10457
<i>From Scenes to Elements: Multi-Granularity Evidence Retrieval for Verifiable Multimodal RAG</i>	
Guanhua Chen, Chuyue Huang, Yutong Yao, Shudong Liu, Xueqing Song, Lidia S. Chao and Derek F. Wong	10475
<i>Region-RI: Reinforcing Query-Side Region Cropping for Multi-Modal Re-Ranking</i>	
Chan-Wei Hu and Zhengzhong Tu	10492
<i>A Reward-Guided Dual-Phase Framework for Adaptive Inference-Time Reasoning</i>	
Yingqian Cui, Zhenwei Dai, Pengfei He, Bing He, Hui Liu, Zhan Shi, Xianfeng Tang, Jingying Zeng, Suhang Wang, Yue Xing, Jiliang Tang and Benoit Dumoulin	10506
<i>LVLMM-Aware Multimodal Retrieval for RAG-Based Medical Diagnosis with General-Purpose Models</i>	
Nir Mazor and Tom Hope	10532
<i>Detecting Proxy Gaming in RL and LLM Alignment via Evaluator Stress Tests</i>	
Sanjeda Akter, Ibne Farabi Shihab and Anuj Sharma	10554
<i>Do Language Models Think Consistently? A Study of Value Preferences Across Varying Response Lengths</i>	
Inderjeet Jayakumar Nair and Lu Wang	10584
<i>LEDGER: Scaling Agentic Document Editing with Dependency-aware Graph Retrieval</i>	
Hang Wang, Utkarsh Garg, Reza Davari, Huitian Jiao, Hao Cheng, Baolin Peng, Si-Qing Chen and Tao Ge	10614
<i>NSF-CoT: Neuro-Symbolic Formal Verification of Chain-of-Thought Faithfulness in Contextual Question Answering</i>	
Vishal Pramanik, Maisha Maliha, Nathaniel D. Bastian, Alvaro Velasquez, Susmit Jha and Sumit Kumar Jha	10645
<i>Learning Flexible Large Multimodal Models with Arbitrary Modality Combinations</i>	
Xinyu Zhao, Kangqi Ni, Jie Peng, Ang Li and Tianlong Chen	10664
<i>1,729 vs. 1729: The Effect of Scripts and Formats on LLM Numeracy</i>	
Varshini Reddy, Craig W Schmidt, Seth Ebner, Adam Wiemerslage, Yuval Pinter and Chris Tanner	10679

<i>SimpleOCR: Rendering Visual Questions to Teach MLLMs to Read</i>	
Yibo Peng, Peng Xia, Ding Zhong, Kaide Zeng, Siwei Han, Yiyang Zhou, Jiaqi Liu, Ruiyi Zhang and Huaxiu Yao	10697
<i>Beyond Facts- Benchmarking Distributional Reading Comprehension in Large Language Models</i>	
Pei-Fu Guo, Ya An Tsai, Chun-Chia Hsu, Kai-Xin Chen, Yun-Da Tsai, Kai-Wei Chang, Nanyun Peng, Mi-Yen Yeh and Shou-De Lin	10711
<i>Preference Learning Unlocks LLMs’ Psycho-Counseling Skills</i>	
Mian Zhang, Shaun M. Eack and Zhiyu Chen	10729
<i>Large Language Models Can Help Mitigate Barren Plateaus in Quantum Neural Networks</i>	
Jun Zhuang and Chaowen Guan	10751
<i>Finch: Benchmarking Finance & Accounting across Spreadsheet-Centric Enterprise Workflows</i>	
Haoyu Dong, Pengkun Zhang, Yan Gao, Xuanyu Dong, Yilin Cheng, Mingzhe Lu, Adina Yakefu and Shuxin Zheng	10768
<i>Exploiting Tree Structure for Credit Assignment in Reinforcement Learning with Large Language Models</i>	
Hieu Tran, Zonghai Yao and Hong yu	10795
<i>Demystify the Role of Memory in Machine Learning Engineering Agents</i>	
Xinyu Zhao, Junpeng Wang, Yuzhong Chen, Menghai Pan, Chin-Chia Michael Yeh, Jiarui Sun, Yan Zheng, Mahashweta Das and Tianlong Chen	10811
<i>Characterizing Web Search in The Age of Generative AI</i>	
Elisabeth Kirsten, Jost Große Perdekamp, Qinyuan Wu, Mihir Upadhyay, Krishna P. Gummadi and Muhammad Bilal Zafar	10827
<i>HypEHR: Hyperbolic Modeling of Electronic Health Records for Efficient Question Answering</i>	
Yuyu Liu, Sarang Rajendra Patil, Mengjia Xu and Tengfei Ma	10849
<i>What Users Leave Unsaid: Under-Specified Queries Limit Vision-Language Models</i>	
Dasol Choi, Guijin Son, Hanwool Lee, Minhyuk Kim, Hyunwoo Ko, Teabin Lim, Eungyeol Ahn, Jungwhan Kim, Seunghyeok Hong and Youngsook Song	10863
<i>Selective Steering: Norm-Preserving Control Through Discriminative Layer Selection</i>	
Quy-Anh Dang and Chris Ngo	10887
<i>From Bytes to Subwords: Challenges of Input Representations in NLP</i>	
Rob Van Der Goot	10911
<i>The Learnability of Model-Theoretic Interpretation Functions in Artificial Neural Networks</i>	
Adrian Brasoveanu and Jakub Dotlacil	10920
<i>A Mechanistic Perspective and Difficulty Metric for Unlearning</i>	
Jiali Cheng, Ziheng Chen, Chirag Agarwal and Hadi Amiri	10950
<i>UCS: Estimating Unseen Coverage for Improved In-Context Learning</i>	
Jiayi Xin, Xiang Li, Evan Qiang, Weiqing He, Tianqi Shang, Weijie J Su and Qi Long	10965
<i>Talent or Luck? Evaluating Attribution Bias in Large Language Models</i>	
Chahat Raj, Mahika Banerjee, Jinhao Pan, Aylin Caliskan, Antonios Anastasopoulos and Ziwei Zhu	10982
<i>Uncertainty Quantification of Large Language Models through Multiple Uncertainty Sources</i>	
Tiejun Chen, Xiaoou Liu, Longchao Da, Jia Chen, Evangelos E. Papalexakis and Hua Wei .	11015

<i>Dual Hierarchical Dialogue Policy Learning for Legal Inquisitive Conversational Agents</i> Xubo Lin, Zezhi Deng, Shihao Wang, Grace Hui Yang and Yang Deng	11030
<i>[b] = [d] - [t] + [p]: Self-supervised Speech Models Discover Phonological Vector Arithmetic</i> Kwanghee Choi, Eunjung Yeo, Cheol Jun Cho, David Harwath and David R. Mortensen ..	11048
<i>Waking Up Blind: Cold-Start Optimization of Supervision-Free Agentic Trajectories for Grounded Visual Perception</i> Ashutosh Bajpai, Tamal Majumder, Akshay Nambi and Tanmoy Chakraborty	11070
<i>Spatial-RAG: Spatial Retrieval Augmented Generation for Real-World Geospatial Reasoning Questions</i> Dazhou Yu, Riyang Bao, Ruiyu Ning, Jinghong Peng, Gengchen Mai and Liang Zhao	11094
<i>Program Structure-aware Language Models: Targeted Software Testing beyond Textual Semantics</i> Khang Tran, Khoa Nguyen, Cristian Borcea and Hai Phan	11113
<i>Reasoning for Hierarchical Text Classification: The Case of Patents</i> Lekang Jiang, Wenjun Sun and Stefan Goetz	11127
<i>From Shijing to English and German: Resources and Evaluation for LLM Translation of Early Chinese Poetry</i> Ying Jiao and Meng Sun	11143
<i>One Mask to Rule Them All: On Hidden Facts after Editing and How to Find Them</i> Ali Kholmovaia, Paul Youssef, Nandi Schoots and Christin Seifert	11163
<i>Analysing the Safety Pitfalls of Steering Vectors</i> Yuxiao Li, Alina Fastowski, Efstratios Zaradoukas, Bardh Prenkaj and Gjergji Kasneci ...	11182
<i>Compartmentalised Agentic Reasoning for Clinical NLI</i> Mael Jullien, Andre Freitas, Marco Valentino and Lei Xu	11205
<i>RAG or Learning? Understanding the Limits of LLM Adaptation under Continuous Knowledge Drift in the Real World</i> Hanbing Liu, Lang Cao and Yang Li	11234
<i>Content Fuzzing for Escaping Information Cocoons on Digital Social Media</i> Yifeng He, Ziye Tang and Hao Chen	11253
<i>KnowledgeBerg: Evaluating Systematic Knowledge Coverage and Compositional Reasoning in Large Language Models</i> Xiao Zhang, Qianru Meng, Yongjian Chen, Yumeng Wang and Johan Bos	11272
<i>Hybrid Self-evolving Structured Memory for Computer-Use Agents</i> Sibo Zhu, Wenyi WU, Kun Zhou, Stephen Wang and Biwei Huang	11287
<i>Characterizing Selective Refusal Bias in Large Language Models</i> Adel Khorramrouz and Sharon Levy	11305
<i>Infusing Theory of Mind into Socially Intelligent LLM Agents</i> EunJeong Hwang, Yuwei Yin, Giuseppe Carenini, Peter West and Vered Shwartz	11327
<i>Do BabyLMs Wanna Learn Wanna Contraction? On the Learnability without Language-Specific Bias</i> Kangsan Noh and Sanghoun Song	11361
<i>OpenGenAlign: A Preference Dataset and Benchmark for Trustworthy Reward Modeling in Open-Ended, Long-Context Generation</i> Hanning Zhang, Juntong Song, Juno Zhu, Yuanhao Wu, Tong Zhang and Cheng Niu	11373

<i>Understanding LLMs' summarization capabilities: an analysis of biomedical abstract and lay summary generation</i>	Batuhan Nursal and Cassie S. Mitchell	11393
<i>FIFA: Unified Faithfulness Evaluation Framework for Text-to-Video and Video-to-Text Generation</i>	Liqiang Jing, Viet Dac Lai, Seunghyun Yoon, Trung Bui and Xinya Du	11418
<i>RLHS: Mitigating Misalignment in RLHF with Hindsight Simulation</i>	Kaiqu Liang, Haimin Hu, Ryan Liu, Thomas L. Griffiths and Jaime Fernández Fisac	11457
<i>Propaganda Signals in LLMs: Perspectival Divergence and Narrative Framing in the Russia-Ukraine War</i>	Ofir Shabat, Ido Guy and Kira Radinsky	11484
<i>DELTA: Dynamic Layer-Aware Token Attention for Efficient Long-Context Reasoning</i>	Hossein Entezari Zarch, Lei Gao, Chaoyi Jiang and Murali Annavaram	11502
<i>LANTERN in the Event Stream: Training-Free Temporal Knowledge Graph Forecasting by Balancing Inertia and Shifts</i>	Chengyuan Jin, Ao Chang, Daojian Zeng, Wenhao Teng, Xiangwen Liao, Kang Liu, Jun Zhao and Yubo Chen	11519
<i>Large Language Models for IT Automation Tasks: Are We There Yet?</i>	Md. Mahadi Hassan, John Salvador, Akond Ashfaqur Ur Rahman and Santu Karmaker	11534
<i>Follow the Path: Reasoning over Knowledge Graph Paths to Improve Large Language Model Factuality</i>	Mike Zhang, Johannes Bjerva and Russa Biswas	11574
<i>Time-RA: Towards Time Series Reasoning for Anomaly Diagnosis with LLM Feedback</i>	Yiyuan Yang, Zichuan Liu, Lei Song, Kai Ying, Stephen Wang, Joshua Thomas Bamford, Svitlana Vyetrenko, Jiang Bian and Qingsong Wen	11591
<i>Ground Then Rank: Revisiting Knowledge-Based VQA with Training-Free Entity Identification</i>	Qian Ma, Qiong Wu, Zhengyi Zhou and Yao Ma	11617
<i>Non-invasive electromyographic speech neuroprosthesis: a geometric perspective</i>	Harshavardhana T Gowda and Lee M. Miller	11636
<i>Context Attribution with Multi-Armed Bandit Optimization</i>	Deng Pan, Keerthiram Murugesan, Ting Hua, Nuno Moniz and Nitesh V. Chawla	11651
<i>Peering Behind the Shield: Guardrail Identification in Large Language Models</i>	Ziqing Yang, Yixin Wu, Rui Wen, Michael Backes and Yang Zhang	11663
<i>WildSci: Advancing Scientific Reasoning from In-the-Wild Literature</i>	Tengxiao Liu, Deepak Nathani, Zekun Li, Kevin Yang and William Yang Wang	11677
<i>WebDART: Dynamic Decomposition and Re-planning for Complex Web Tasks</i>	Jingbo Yang, Bairu Hou, Wei Wei, Shiyu Chang and Yujia Bao	11696
<i>Like a Therapist, But Not: Reddit Narratives of AI in Mental Health Contexts</i>	Elham Aghakhani and Rezvaneh Rezapour	11716
<i>VoiceStar: Robust Zero-Shot Autoregressive TTS with Duration Control and Extrapolation</i>	Puyuan Peng, Zhisheng Zheng, Shang-Wen Li, Abdelrahman Mohamed and David Harwath	11737
<i>Evaluating Implicit Biases in LLM Reasoning through Logic Grid Puzzles</i>	Fatima Jahara, Mark Dredze and Sharon Levy	11755

<i>InferPilot: Autonomous Inference Attacks Against ML Services With LLM-Based Agents</i> Yixin Wu, Rui Wen, Chi Cui, Michael Backes and Yang Zhang	11781
<i>MedErrBench: A Fine-Grained Multilingual Benchmark for Medical Error Detection and Correction with Clinical Expert Annotations</i> Congbo Ma, Yichun Zhang, Yousef Al-Jazzazi, Ahamed Foisal, Laasya Sharma, Youstra Sadqi, Khaled Saleh, Jihad Mallat and Farah E. Shamout	11802
<i>MetaScale: Test-Time Scaling with Evolving Meta-Thoughts</i> Qin Liu, Wenxuan Zhou, Nan Xu, James Y. Huang, Fei Wang, Sheng Zhang, Hoifung Poon and Muhao Chen	11828
<i>Arab Voices: Mapping Standard and Dialectal Arabic Speech Technology</i> Peter Sullivan, AbdelRahim A. Elmadany, Alcides Alcoba Inciarte and Muhammad Abdul-Mageed	11843
<i>CoDA: Restoring Contextual Dominance via Copy-Encouraged Attention Intervention for Mitigating RAG Hallucinations</i> JinWei Shi, Qizhuo Xie, Qianzi Hou, Zhipeng Wang, Wanting Su, Jianhua Zhao, Tao Zheng and Tieke He	11879
<i>CausalDetox: Causal Head Selection and Intervention for Language Model Detoxification</i> Yian Wang, Yuen Chen, Agam Goyal and Hari Sundaram	11893
<i>Commonsense Knowledge with Negation: A Resource to Enhance Negation Understanding</i> Zijie Wang, MohammadHossein Rezaei, Farzana Rashid and Eduardo Blanco	11915
<i>P-QuASAR: A Unified Probabilistic Framework for Holistic Patent Quality Assessment and Refinement</i> Xinyuan Song, Ziyi Ni, Fred Yang, Bo Zhang, Yijin Wang and Jane Zhang	11935
<i>From Heard to Lived Opinions: Simulating Opinion Dynamics with Grounded LLM Agents in Economic Environments</i> Ryuji Hashimoto, Masahiro Kaneko, Ryosuke Takata, Takehiro Takayanagi and Kiyoshi Izumi	11952
<i>HiSA: Hierarchical State Abstraction for Scalable GUI Agents</i> Weiming Li, Hye-young Paik and Yulei Sui	11965
<i>SecureWebArena: A Holistic Security Evaluation Benchmark for LVLM-based Web Agents</i> Zonghao Ying, Yangguang Shao, Jianle Gan, Gan Xu, Wenxin Zhang, Quanchen Zou, Junzheng Shi, Zhenfei Yin, Mingchuan Zhang, Aishan Liu and Xianglong Liu	11986
<i>The Mechanics of Interference: Defusing Distractors in RAG via Sparse Autoencoder Interventions</i> Christian Giannetti, Giovanni Trappolini, Nicola Tonello, Fabrizio Silvestri and Pietro Lio	11999
<i>Grounding Agent Memory in Contextual Intent</i> Ruozen Yang, Yucheng Jiang, Yueqi Jiang, Priyanka Kargupta, Yunyi Zhang and Jiawei Han	12008
<i>Persona-Assigned Large Language Models Exhibit Human-Like Motivated Reasoning</i> Saloni Dash, Amélie Reymond, Emma Spiro and Aylin Caliskan	12043
<i>Generalizable LLM Learning of Graph Synthetic Data with Post-training Alignment</i> Yizhuo Zhang, Heng Wang, Shangbin Feng, Zhaoxuan Tan, Xinyun Liu and Yulia Tsvetkov	12070

<i>Distilling the Essence: Efficient Reasoning Distillation via Sequence Truncation</i>	
Wei-Rui Chen, Vignesh Kothapalli, Ata Fatahibaarzi, Hejian Sang, Shao Tang, Qingquan Song, Zhipeng Wang and Muhammad Abdul-Mageed	12092
<i>Beyond Cross-Modal Alignment: Measuring and Leveraging Modality Gap in Vision-Language Models</i>	
Hanqi Yan, Xiangxiang Cui, Lu Yin, Jindong Gu, Paul Pu Liang, Yulan He and Yifei Wang	12123
<i>Data Swarms: Optimizable Generation of Synthetic Evaluation Data</i>	
Shangbin Feng, Yike Wang, Weijia Shi and Yulia Tsvetkov	12139
<i>Concept rather than Document: Context Compression via AMR-based Conceptual Entropy</i>	
Kaize Shi, Xueyao Sun, Xiaohui Tao, Lin Li, Qika Lin and Guandong Xu	12158
<i>Can LLMs Understand the Impact of Trauma? Costs and Benefits of LLMs Coding the Interviews of Firearm Violence Survivors</i>	
Jessica H Zhu, Shayla Stringfield, Vahe Zaprosyan, Michael Wagner, Michel Cukier and Joseph Richardson	12174
<i>"I Don't Know What to Say": A Fact-Filling Questionnaire Method to Help Non-Experts Talk to LegalAI Assistant</i>	
Yuting Huang, Yiquan Wu, Meitong Guo, Ang Li, Xiaozhong Liu, Keting Yin, Fei Wu and Kun Kuang	12193
<i>BnMMLU: Measuring Massive Multitask Language Understanding in Bengali</i>	
Saman Sarker Joy and Swakkhar Shatabda	12211
<i>Diagnosing LLMs via Information Spectrum Analysis: Tail Behavior and the Effects of Side Information</i>	
Yuuki Tachioka	12231
<i>Scaling Collaborative Effort with Agents</i>	
Shannon Zejiang Shen, Valerie Chen, Ken Gu, Alexis Ross, Zixian Ma, Jillian Ross, Alex Gu, Chenglei Si, Wayne Chi, Andi Peng, Jocelyn J Shen, Ameet Talwalkar, Tongshuang Wu and David Sontag	12254
<i>CriticSearch: Fine-Grained Credit Assignment for Search Agents via a Retrospective Critic</i>	
Yaocheng Zhang, Haohuan Huang, Zijun Song, Zijie Zhao, Qichao Zhang, Yuanheng Zhu and Dongbin Zhao	12272
<i>Double-Edged Sword or Sharp Tool? Designing and Evaluating Triadic LLM-Teacher Collaboration for K-12 Writing at Scale</i>	
Canran Wang, Yuwen Yang, Zhen Wang, Ming MA, Ding Yu, Chentai Wang, Keman Huang and Xiaoyong Du	12291
<i>IntroLM: Introspective Language Models via Prefilling-Time Self-Evaluation</i>	
Hossein Hosseini Kasnavieh, Gholamreza Haffari, Christopher Leckie and Adel N. Toosi .	12313
<i>A Herd of Language Models Makes a Better Zero-shot Annotator for Clinical Named Entity Recognition</i>	
Seiji Shimizu, Shoko Wakamiya and Eiji Aramaki	12327
<i>P2P: A Poison-to-Poison Remedy for Reliable Backdoor Defense in LLMs</i>	
Shuai Zhao, Xinyi Wu, Shiqian Zhao, Xiaobao Wu, Zhongliang Guo, Yanhao Jia and Anh Tuan Luu	12345
<i>Learning Temporally-Aware Sample Weights for Preference Optimization</i>	
Mengyang Li, Xudong Zhou and Pinlong Zhao	12361

<i>ToolPRMBench: Evaluating and Advancing Process Reward Models for Tool-using Agents</i>	
Dawei Li, Yuguang Yao, Zhen Tan, Huan Liu and Ruocheng Guo	12378
<i>Adapt to Thrive! Adaptive Power-Mean Policy Optimization for Improved LLM Reasoning</i>	
Yiming Huang, Zhenbo Shi, Shuzheng Gao, Cuiyun Gao, Peiyi Han and Chuanyi Liu	12392
<i>LEAF: Towards Lightweight Explainable Hateful Video Detection via Self-Grounding CoT Guided Stage-Wise Distillation</i>	
Jian Lang, Rongpei Hong, Meihui Zhong, Kaiju Li, Ting Zhong, Qiang Gao and Fan Zhou	12420
<i>DemMA: Dementia Multi-Turn Dialogue Agent with Expert-Guided Reasoning and Action Simulation</i>	
Yutong Song, Jiang Wu, Kazi Shaharair Sharif, Pengfei Zhang, Wenjun Huang, Honghui Xu, Nikil Dutt and Amir M. Rahmani	12439
<i>SplitThenMerge: Token-Level Skill-Compositional Sparse Mixture-of-Experts for Complex Domain-Specific Tasks</i>	
Yuting Huang, Jiawen Zhang, Yiquan Wu, Yinghao Hu, Fei Wu and Kun Kuang	12471
<i>Causal-Audit: Explicit and Auditable Graph-based Reasoning via Target-Aware Causal Chain Construction</i>	
Su Lan, Xuefei Yin, Yanming Zhu and Alan Wee-Chung Liew	12484
<i>RADS: Reinforcement Learning-Based Sample Selection Improves Transfer Learning in Low-resource and Imbalanced Clinical Settings</i>	
Wei Han, David Martinez Iraola, Anna Khanina, Lawrence Cavedon and Karin Verspoor	12501
<i>Watch Out Your Industrial Copilots: Stealthy Backdoor Attack Against LLM-Based PLC Code Generation</i>	
Xinyuan An, Liu Xiaoxia, Dongxia Wang, Zhanhang Xiong and Wenhai Wang	12517
<i>Balancing Classification and Calibration Performance in Decision-Making LLMs via Calibration Aware Reinforcement Learning</i>	
Duygu Nur Yaldiz, Evangelia Spiliopoulou, Zheng Qi, Siddharth Varia, Srikanth Doss and Nikolaos Pappas	12537
<i>Redefining Machine Simultaneous Interpretation: From Incremental Translation to Human-Like Strategies</i>	
Qianen Zhang, Zeyu Yang and Satoshi Nakamura	12554
<i>ChunQiuTR: Time-Keyed Temporal Retrieval in Classical Chinese Annals</i>	
Yihao Wang, Zijian He, Jie Ren and Keze Wang	12578
<i>OFFSIDE: Benchmarking Unlearning Misinformation in Multimodal Large Language Models</i>	
Hao Zheng, Zirui Pang, Ling Li, Zhijie Deng, Yuhan Pu, Zhaowei Zhu, Xiaobo Xia and Jiaheng Wei	12602
<i>PedagogyBench: A Cognitive-Driven Benchmark for Multimodal Instructional Video Understanding</i>	
Xiaokang Jin, Jia Zhu, Jingjiang Liu, Yabing Shi, Jueqi Guan, Hao Chen and Pasquale De Meo	12621
<i>FeedEval: Pedagogically Aligned Evaluation of LLM-Generated Essay Feedback</i>	
SeongYeub Chu, Jongwoo Kim and Mun Yong Yi	12648
<i>Probing Multimodal Large Language Models on Cognitive Biases in Chinese Short-Video Misinformation</i>	
Jen-tse Huang, Chang Chen, Shiyang Lai, Wenxuan Wang, Michelle R Kaufman and Mark Dredze	12675

<i>Emotion–Cause Pair Extraction in Conversations via Semantic Decoupling and Alignment</i>	
Tianxiang Ma, Weijie Feng, Xinyu Wang and Zhiyong Cheng	12697
<i>Failure makes the agent stronger: Enhancing Accuracy through Structured Reflection for Reliable Tool Interactions</i>	
Junhao Su, Yuanliang Wan, Junwei Yang, Hengyu Shi, Tianyang Han, Yurui Qiu and Junfeng Luo	12712
<i>Feedback to Reasoning: LLM-Assisted Molecular Optimization with Domain Feedback and Historical Reasoning</i>	
Wenhan Gao, Xiran Fan, Chin-Chia Michael Yeh, Jiarui Sun, Yuzhong Chen, Menghai Pan, Mahashweta Das and Yi Liu	12735
<i>Iterative Formalization and Planning in Partially Observable Environments</i>	
Liancheng Gong, Wang Bill Zhu, Jesse Thomason and Li Zhang	12755
<i>StructKV: Preserving the Structural Skeleton for Scalable Long-Context Inference</i>	
Zhirui Chen, Peiyang Liu and Ling Shao	12784
<i>Correct, Concise and Complete: Multi-stage Training For Adaptive Reasoning</i>	
Nathanaël Carraz Rakotonirina, Ren Pang, Neha Anna John, Michael Bohlke-Schneider and Momchil Hardalov	12798
<i>At Your Own PACE: A Causal Framework for Evaluating EQ in LLMs</i>	
Lei Lyu, Shengling Wang, Ke Chao and Yichao Wei	12811
<i>Identifying Collective Intelligence Factor in LLM Agent Groups for Generalizable Multi-Agent System Design</i>	
Zhilun Zhou, Zihan Liu, Jiahe Liu, Yihan Wang, Qingyu Shao, Fengli Xu, Depeng Jin and Yong Li	12827
<i>Rethinking Parameter Sharing for LLM Fine-Tuning with Multiple LoRAs</i>	
Hao Ban and Kaiyi Ji	12843
<i>Shorten After You’re Right: Lazy Length Penalties for Reasoning RL</i>	
Danlong Yuan, Tian Xie, Shaohan Huang, Huishuai Zhang, Zhuocheng Gong, Chong Luo, Furu Wei and Dongyan Zhao	12864
<i>MedDCR: Learning to Design Agentic Workflows for Medical Coding</i>	
Jiyang Zheng, Islam Nassar, Thanh Vu, Xu Zhong, Yang Lin, Tongliang Liu, Long Duong and Yuan-Fang Li	12878
<i>TokenPenalty: Alleviating Attention Sinks and Positional Decay in LVLMS</i>	
Xiaofeng Zhang, Yuanchao Zhu, Qiyao Zhao, Xiaosong Yuan, Jiawei Cao and Xuhang Chen	12894
<i>Beyond Screenshots: Evaluating VLMs’ Understanding of UI Animations</i>	
Chen Liang, Xirui Jiang, Naihao Deng, Eytan Adar and Anhong Guo	12907
<i>MemWeaver: Weaving Hybrid Memories for Traceable Long-Horizon Agentic Reasoning</i>	
Juexiang Ye, Xue Li, Yang Xinyu, Chengkai Huang, Lanshun Nie, Lina Yao and Dechen Zhan	12928
<i>From Mimesis to Metamorphosis: Evolving VLM Judges via In-Context Comparing and Knowledge Internalization</i>	
Juntuo Wang, Yuming Qiao, Yifan Yang, Lunxi Yuan, Liang Luo and Dan Meng	12957

<i>Generative Text-to-Image Retrieval via Hierarchical Identifiers and Semantic Internalization</i>	
Jie Huang, Junjie Wang, Xin Liao, Ziyong Jiang, Wenshuo Wang, Shoubin Li and Qing Wang	12972
<i>Bridging Reasoning and Action: Hybrid LLM–RL Framework for Efficient Cross-Domain Task-Oriented Dialogue</i>	
Yangyang Zhao, Linfan Dai, Li Cai, Bowen Xing and Libo Qin	12987
<i>JudgeAgent: Beyond Static Benchmarks for Knowledge-Driven and Dynamic LLM Evaluation</i>	
Zhichao Shi, Xuhui Jiang, Chengjin Xu, Cangli Yao, Shengjie Ma, Yinghan Shen, Zixuan Li, Jian Guo and Yuanzhuo Wang	13004
<i>Retrieve Only Relevant Tables Whether Few or Many: Adaptive Table Retrieval Method</i>	
Taehee Kim, Seungbin Yang, Jihwan Kim and Jaegul Choo	13031
<i>Code-Switching Information Retrieval: Benchmarks, Analysis, and the Limits of Current Retrievers</i>	
Qingcheng Zeng, Yuheng Lu, Zeqi Zhou, Heli Qi, Puxuan Yu, Fuheng Zhao, Hitomi Yanaka, Weihao Xuan and Naoto Yokoya	13055
<i>WebUncertainty: Dual-Level Uncertainty Driven Planning and Reasoning For Autonomous Web Agent</i>	
Lingfeng Zhang, Yongan Sun, Jinpeng Hu, Hui Ma, Ying Yang, Kuien Liu, Zenglin Shi and Meng Wang	13072
<i>Multi-Faceted Self-Consistent Preference Alignment for Query Rewriting in Conversational Search</i>	
Zhiyu Cao, Peifeng Li and Qiaoming Zhu	13083
<i>CogNet-KG: Empowering Tutoring Dialogues with a Cognitively-Structured Knowledge Graph for STEM Learning</i>	
Ding Yu, Yu Lu, Tengju Li, Shasha Xiong and Shengquan Yu	13101
<i>From 128K to 4M: Efficient Training of Ultra-Long Context Large Language Models</i>	
Chejian Xu, Wei Ping, Peng Xu, Zihan Liu, Boxin Wang, Mohammad Shoeybi, Bo Li and Bryan Catanzaro	13122
<i>A Few Bad Apples Spoil the Bunch: Preventing Global Entropy Collapse Driven by a Small Set of Tokens in LLM Reasoning</i>	
Jaeun Jang, Hansle Lee and Sangmin Kim	13134
<i>Efficient Training for Cross-lingual Speech Language Models</i>	
Yan Zhou, Qingkai Fang, Yun Hong and Yang Feng	13155
<i>Empirical Analysis of Task Mixture Effects in Small-scale Instruction Tuning: A Statistical Approach</i>	
Jeesu Jung and Sangkeun Jung	13168
<i>ResearchBench: Benchmarking LLMs in Scientific Discovery via Inspiration-Based Task Decomposition</i>	
Yujie Liu, Zonglin Yang, Tong Xie, Jinjie Ni, Ben Gao, Yuqiang Li, Shixiang Tang, Wanli Ouyang, Erik Cambria and Dongzhan Zhou	13187
<i>GraphLoRA: Structure-Aware Low-Rank Adaptation for Large Language Model Recommendation</i>	
Lin Mu, Guoji Wang, Li Ni, Lei Sang, Zhize Wu, Peiquan Jin and Yiwen Zhang	13208
<i>Spotlight and Shadow: Attention-Guided Dual-Anchor Introspective Decoding for MLLM Hallucination Mitigation</i>	
Yebo Wu, Han Jin, Zhijiang Guo and Li Li	13219

<i>StreamMeCo: Long-Term Agent Memory Compression for Efficient Streaming Video Understanding</i> Junxi Wang, Te Sun, Jiayi Zhu, Junxian Li, Haowen Xu, Zichen Wen, Xuming Hu, Zhiyu li and Linfeng Zhang	13234
<i>Reflective RAG: Self-Evaluation Driven Strategy Optimization in Agentic Retrieval-Augmented Generation</i> Haiyan Wu, Chenchen Wang, Chaoqun Sun, Chengxiong Lu, Yan-Hong Chen, Zhiqiang Zhang and Xiaoqing Feng	13252
<i>Meta-Reasoner: Dynamic Guidance for Optimized Inference-time Reasoning in Large Language Models</i> Yuan Sui, Yufei He, Tri Cao, Sophia Simeng Han, Yulin Chen and Bryan Hooi	13268
<i>Polynomial Expansion Rank Adaptation: Enhancing Low-Rank Fine-Tuning with High-Order Interactions</i> Wenhao Zhang, Lin Mu, Li Ni, Peiquan Jin and Yiwen Zhang	13287
<i>Efficient Test-Time Scaling via Temporal Reasoning Aggregation</i> Jiakun Li, Xingwei He, Kefan Li, Hongzheng Chai, Hongyue Yu and Yuan Yuan	13304
<i>A Survey on MLLM-based Visually Rich Document Understanding: Methods, Challenges, and Emerging Trends</i> Yihao Ding, Siwen Luo, Yue Dai, Yanbei Jiang, Zechuan Li, Qiang Sun, Geoffrey Martin, Wei Liu and Yifan Peng	13319
<i>Detecting Sensitive Personal Information in Japanese Pre-Training Corpora for Large Language Models</i> Rei Minamoto, Yusuke Oda and Daisuke Kawahara	13341
<i>Agentic Episodic Control</i> Xidong Yang, Wenhao Li, Junjie Sheng, Yun Hua, Haosheng Chen, Chuyun Shen and Xiangfeng Wang	13355
<i>TrajGuard: Streaming Hidden-state Trajectory Detection for Decoding-time Jailbreak Defense</i> Cheng Liu, Xiaolei Liu, Xingyu Li, Bangzhou Xin and Kangyi Ding	13371
<i>DeepPrune: Parallel Scaling without Inter-trace Redundancy</i> Shangqing Tu, Yaxuan Li, Yushi Bai, Lei Hou and Juanzi Li	13389
<i>Contrastive Decoding Mitigates Score Range Bias in LLM-as-a-Judge</i> Yoshinari Fujinuma	13404
<i>ContextCheck: Sentence-Level Faithfulness Verification with Context-Aware Disambiguation</i> Yueqin Yin, Yaxi Li, Xin Liu, Xun Wang, Kaiqiang Song, Simin Ma, Shujian Liu, Sathish Reddy Indurthi, Haoyun Deng, Pengcheng He, Mingyuan Zhou and Song Wang	13419
<i>Less Noise, More Voice: Reinforcement Learning for Reasoning via Instruction Purification</i> Yiju Guo, Tianyi Hu, Zexu Sun and Yankai Lin	13469
<i>MAST: A Multi-View Alignment Strategy for Optimal Transport-Based Contrastive Clustering of Short Text</i> Zijian Zheng and Yonghe Lu	13483
<i>AICA-Bench: Holistically Examining the Capabilities of VLMs in Affective Image Content Analysis</i> Dong She, Xianrong Yao, Liqun Chen, Jinghe Yu, Yang Gao and Zhanpeng Jin	13501

<i>You Only Need One Single Token to Refine Safety Alignment</i>	
Wenqian Yu, Shuo Chen, Zhijiang Li, Zhipeng Wang and Jindong Gu	13529
<i>CrisPrune: Combining Contextual Relevance and Intrinsic Saliency for Efficient Visual Token Pruning in MLLMs</i>	
Ziniu Liu, Shuheng Zhou, Mingqing Liu, Hao Deng and Huijia Zhu	13546
<i>Beyond Topology: Generative Node Importance Estimation via Structure-Guided Semantic Reasoning</i>	
Kuofei Fang, Siyan Wu, Yu Guo and Bin Wu	13565
<i>AGT^{AO}: Robust and Stabilized LLM Unlearning via Adversarial Gating Training with Adaptive Orthogonality</i>	
Pengyu Li, Lingling Zhang, Zhitao Gao, Yanrui Wu, Yuxuan Dong, Huan Liu, Bifan Wei and Jun Liu	13585
<i>LLMs Meet Isolation Kernel: Lightweight, Learning-free Binary Embeddings for Fast Retrieval</i>	
Zhibo Zhang, Yang Xu, Kai Ming Ting and Cam-Tu Nguyen	13601
<i>BOLT: Benchmarking Open-World Learning for Text Classification</i>	
Chuan Qin, Xi Chen, Jinpeng Li and Hengshu Zhu	13624
<i>MobileBench-OL: A Comprehensive Chinese Benchmark for Evaluating Mobile GUI Agents in Real-World Environment</i>	
Qinzhao Wu, Zhizhuo Yang, Hanhao Li, Pengzhi Gao, Wei Liu and Jian Luan	13659
<i>CondenseFlow: Scalable Latent Space Collaboration via Semantic Compression for Multi-Agent Systems</i>	
Xiaoyu Chen, Fengge Wu, Zhao Junsuo and Yun Fan	13694
<i>MAXS: Meta-Adaptive Exploration with LLM Agents</i>	
Jian Zhang, Zhiyuan Wang, Zhangqi Wang, Yu He, Haoran Luo, li Yuan, Lingling Zhang, Rui Mao, Qika Lin and Jun Liu	13713
<i>Leveraging Human and Machine Preferences for Zero-shot Detection of AI-Generated Text</i>	
Lei Jiang, Desheng Wu, Xiaolong Zheng and Cuicui Luo	13732
<i>MelTrim: Coarse-to-Fine Data Pruning for Speech Classification</i>	
Shaobo Wang, Tianle Niu, Xuan Ouyang, Xintong Li, Zhengkun Ge, Yue Min, Xiaoqian Liu, Hankun Wang and Linfeng Zhang	13751
<i>MINED: Probing and Updating with Multimodal Time-Sensitive Knowledge for Large Multimodal Models</i>	
Kailin Jiang, Ning Jiang, Yuntao Du, Yuchen Ren, Yuchen Li, Yifan Gao, Jinhe Bi, Yunpu Ma, Bin Li, Lei Liu and Qing Li	13766
<i>Beyond Output Confidence: Epistemic-Aware Hallucination Detection with Answer-Level Signals</i>	
Jieran Li, Xiuyuan Hu, Yang Zhao, Dongbiao Sun and Hao Zhang	13796
<i>PerfCoder: Large Language Models for Interpretable Code Performance Optimization</i>	
Jiuding Yang, Shengyao Lu, Hongxuan Liu, Shayan Shirahmad Gale Bagi, Zahra Fazel, Tomasz Czajkowski and Di Niu	13807
<i>MIPIC: Matryoshka Representation Learning via Self-Distilled Intra-Relational and Progressive Information Chaining</i>	
Phung Gia Huy, Hai An Vu, Minh-Phuc Truong, Thang Duc Tran, Linh Ngo Van, Thanh Hong Nguyen and Trung Le	13824

<i>Discovering the Gems in Early Layers: Accelerating Long-Context LLMs with 1000x Input Token Reduction</i>	
Zhenmei Shi, Yifei Ming, Xuan-Phi Nguyen, Yingyu Liang and Shafiq Joty	13839
<i>Self-Reflection Improves Safety of Large Reasoning Models</i>	
Qiang Huang, Wei Zhai, Feng Huang and Dejing Dou	13858
<i>WildGraphBench: Benchmarking GraphRAG with Wild-Source Corpora</i>	
Pengyu Wang, Benfeng Xu, Licheng Zhang, Shaohan Wang, Mingxuan Du, Chiwei Zhu and Zhendong Mao	13875
<i>CRPS: Curriculum Replay via Progressive Suffixes from Successful Trajectories for Long-Horizon LLM Agents</i>	
Zijing Zhang and Xiajie Yang	13891
<i>D-Artemis: A Deliberative Cognitive Framework for Mobile GUI Multi-Agents</i>	
Hongze Mi, Yibo Feng, WenJie Lu, Yuqi Wang, Jinyuan Li, Song Cao, He Cui, Tengfei Tian, Xuelin Zhang, Haotian Luo, Di Sun, Jun Fang, Hua Chai, Naiqiang Tan and Gang Pan	13905
<i>Temporal Flattening in LLM-Generated Text: Comparing Human and LLM Writing Trajectories</i>	
Zhanwei Cao, YeoJin Go, Yifan Hu and Shanu Sushmita	13932
<i>Beyond Superficial Tests: Adversarial Refinement for Reliable Property-Based Testing</i>	
Xiao Li, Runlin Liu, Zhe Zhang, Xiang Gao and Hailong Sun	13957
<i>FAITH: Factuality Alignment through Integrating Trustworthiness and Honesty</i>	
Xiaoning Dong, Chengyan Wu, Yajie Wen, Yu Chen, Yun Xue, Zhang Jing, Wei Xu and Bolei Ma	13975
<i>DRA-GRPO: Your GRPO Needs to Know Diverse Reasoning Paths for Mathematical Reasoning</i>	
Xiwen Chen, Wenhui Zhu, Peijie Qiu, Xuanzhao Dong, Hao Wang, Haiyu Wu, Huayu Li, Aris Sotiras, Yalin Wang and Abolfazl Razi	13995
<i>Web Fraud Attacks Against LLM-Driven Multi-Agent Systems</i>	
Dezhang Kong, Hujin Peng, Yilun Zhang, Lele Zhao, Zhenhua Xu, Shi Lin, Changting Lin and Meng Han	14020
<i>From Query to Logic: Ontology-Driven Multi-Hop Reasoning in LLMs</i>	
Haonan Bian, Yutao Qi, Rui Yang, Yuanxi Che, Jiaqian Wang, Heming Xia and Ranran Zhen	14035
<i>AudioStealer: Extracting Audio Prompts via Shapley Value-Guided Query Search</i>	
Yingbin Jin, Xingjian Du, Hanjun Luo, Zihao Wang, Haibo Hu, XiaoFeng Wang and Xinfeng Li	14052
<i>FaithLens: Detecting and Explaining Faithfulness Hallucination</i>	
Shuzheng Si, Qingyi Wang, Haozhe Zhao, Yuzhuo Bai, Guanqiao Chen, Kangyang Luo, Gang Chen, Fanchao Qi, Minjia Zhang, Baobao Chang and Maosong Sun	14068
<i>Escaping the Sisyphus Dilemma: Experience Replay for Robust Text-to-Optimization Modeling</i>	
Wantong Xie, Yinghao Chen, Yi-Xiang Hu, Feng Wu, Jieyang Xu, Sijia Zhang and Xiangyang Li	14100
<i>Toward Automated Robustness Evaluation of Mathematical Reasoning</i>	
Yutao Hou, Zeguan Xiao, Fei Yu, Yihan Jiang, Ma Shuguang, Zhaoqian Dai, Hailiang Huang, Yun Chen and Guanhua Chen	14117

<i>Simulating Crisis Cognition: A Computational Framework for Hypothesis Generation in Crisis Communication</i>	
Changsen Yuan, Yanghao Zhou, Chong Feng and Ge Shi	14140
<i>MedScore: Generalizable Factuality Evaluation of Open-ended Long-form Medical Answers by Domain-adapted Claim Decomposition and Verification</i>	
Heyuan Huang, Alexandra DeLucia, Vijay Murari Tiyyala and Mark Dredze	14149
<i>FinSafetyBench: Evaluating LLM Safety in Real-World Financial Scenarios</i>	
Yutao Hou, Yihan Jiang, Yuhan Xie, Jian Yang, Liwen Zhang, Hailiang Huang, Guanhua Chen and Yun Chen	14181
<i>Optimizing Native Sparse Attention with Latent Attention and Local Global Alternating Strategies</i>	
Yuxuan Hu, Jianchao Tan, Jiaqi Zhang, Wen Zan, Pingwei Sun, Yifan Lu, Xunliang Cai and Jing Zhang	14209
<i>BRIEF-Pro: Universal Context Compression with Short-to-Long Synthesis for Fast and Accurate Multi-Hop Reasoning</i>	
Jia-Chen Gu, Junyi Zhang, Di Wu, Yuankai Li, Kai-Wei Chang and Nanyun Peng	14221
<i>O1-Pruner: Length-Harmonizing Fine-Tuning for O1-Like Reasoning Pruning</i>	
Haotian Luo, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Xiaochun Cao, Dacheng Tao, Naiqiang Tan and Li Shen	14242
<i>One Agent to Serve All: a Lite-Adaptive Stylized AI Assistant for Millions of Multi-Style Official Accounts</i>	
Xingyu Fan, Li Feifei, Que Wenhui and Hailong Li	14258
<i>SAVOIR: Learning Social Savoir-Faire via Shapley-based Reward Attribution</i>	
Xiachong Feng, Yi Jiang, Xiaocheng Feng, Deyi Yin, Libo Qin, Yangfan Ye, Lei Huang, Weitao Ma, Yuxuan Gu, Chonghan Qin, Bing Qin and Lingpeng Kong	14276
<i>Language-Coupled Reinforcement Learning for Multilingual Retrieval-Augmented Generation</i>	
Rui Qi, Fengran Mo, Yufeng Chen, Xue Zhang, Shuo Wang, Hongliang Li, Xu Jinan, Meng Jiang, Jian-Yun Nie and Kaiyu Huang	14291
<i>Article and Comment Frames Shape the Quality of Online Comments</i>	
Matteo Guida, Yulia Otmakhova, Eduard Hovy and Lea Frermann	14311
<i>LAD: Learning Advantage Distribution for Reasoning</i>	
Wendi Li and Sharon Li	14327
<i>RealMem: Benchmarking LLMs in Real-World Memory-Driven Interaction</i>	
Haonan Bian, Zhiyuan Yao, Sen Hu, Zishan Xu, Shaolei Zhang, Yifu Guo, Ziliang Yang, Xueran Han, Huacan Wang and Ronghao Chen	14349
<i>AFMRL: Attribute-Enhanced Fine-Grained Multi-Modal Representation Learning in E-commerce</i>	
Biao Zhang, Lixin Chen, Bin Zhang, Wang Zongwei, Tong Liu and Bo Zheng	14366
<i>SAFER: A Controllable Safeguard for LLMs against Backdoor Attacks</i>	
Zirui Hu, Zheng Zhang, Yingjie Wang and Dacheng Tao	14380
<i>RIFT: Repurposing Negative Samples via Reward-Informed Fine-Tuning</i>	
Zehua Liu, Shuqi Liu, Tao Zhong and Mingxuan Yuan	14399
<i>LEPO: Latent Reasoning Policy Optimization for Large Language Models</i>	
Yuyan Zhou, Jiarui Yu, Hande Dong, Zhezhen Hao, Hong Wang, Jianqing Zhang and Qiang Lin	14416

<i>You Never Know a Person, You Only Know Their Defenses: Detecting Levels of Psychological Defense Mechanisms in Supportive Conversations</i>	
Hongbin Na, Zimu Wang, Zhaoming Chen, Peilin Zhou, Yining Hua, Grace Ziqi Zhou, Haiyang Zhang, Tao Shen, Wei Wang, John Torous, Shaoxiong Ji and Ling Chen	14428
<i>Beyond Length Scaling: Synergizing Breadth and Depth for Generative Reward Models</i>	
Qiyuan Zhang, Yufei Wang, Tianhe Wu, Can Xu, Qingfeng Sun, Kai Zheng, Xue Liu and Chen Ma	14449
<i>Pushing the Limits of LLM Tool Calling via Experiential Knowledge Integration and Activation</i>	
Yupu Hao, Zhuoran Jin, Huanxuan Liao, Kang Liu and Jun Zhao	14470
<i>Efficient Transformer Parameter Reuse via Zero-Token Mechanism</i>	
Guanghao Li, Wenhao Jiang, Li Shen, Ming Tang and Chun Yuan	14498
<i>x1: Learning to Think Adaptively Across Languages and Cultures</i>	
Yangfan Ye, Xiaocheng Feng, Xiachong Feng, Yichong Huang, Zekun Yuan, Lei Huang, Weitao Ma, Qichen Hong, Yunfei Lu, Dandan Tu and Bing Qin	14516
<i>Towards Bridging the Reward-Generation Gap in Direct Alignment Algorithms</i>	
Zeguan Xiao, Yun Chen, Jian Yang, Guanhua Chen and Ke Tang	14534
<i>FABLE: Fine-grained Fact Anchoring for Unstructured Model Editing</i>	
Peng Wang, Biyu Zhou, Xuehai Tang, Jizhong Han and Songlin Hu	14549
<i>SOP-Maze: Evaluating Large Language Models on Complicated Business Standard Operating Procedures</i>	
Jiaming Wang, Zhe Tang, Zehao Jin, Hefei Chen, Yilin Jin, Peng Ding, Xiaoyu Li and Xuezhi Cao	14568
<i>MalURLBench: A Benchmark Evaluating Agents' Vulnerabilities When Processing Web URLs</i>	
Dezhang Kong, Zhuxi Wu, Shiqi Liu, ZhiCheng Tan, Kuichen Lu, Minghao Li, Qichen Liu, Shengyu Chu, Zhenhua Xu, Xuan Liu and Meng Han	14589
<i>Representation-Guided Parameter-Efficient LLM Unlearning</i>	
Zeguan Xiao, Lang Mo, Yun Chen, Lei Yang, Jiehui Zhao, Lili Yang and Guanhua Chen ..	14602
<i>LPO: Towards Accurate GUI Agent Interaction via Location Preference Optimization</i>	
Jiaqi Tang, Yu Xia, Yi-Feng Wu, Yuwei Hu, Chen Yuhui, Qing-Guo Chen, Xiaogang Xu, Xiangyu Wu, Hao LU, Yanqing Ma, Shiyin Lu and Qifeng Chen	14617
<i>Mistake Notebook Learning: Batch-Clustered Failures for Training-Free Agent Adaptation</i>	
Xuanbo Su, Yingfang Zhang, Hao Luo, Xiaoteng Liu and Leo Huang	14629
<i>Large-Scale Multimodal Knowledge Graph about Classical Chinese Poetry: Fine-grained Method and Comprehensive Evaluation</i>	
Shuo Wang, Qing Zhu, Yang Xiao and Minglong Lei	14646
<i>MultiCodeAttack: Iterative Jailbreak Attacking on LLMs with Multi-Code Prompt Injection</i>	
Weifeng Sun, Meng Yan, Zhou Yang, Yuchen Chen, Song Sun and David Lo	14670
<i>SAMem: State-Aware Memory as a Fine-Grained Memory for LLM Agents in Decision-Making</i>	
Tong Wang, Pei Xu, Shiyue Cao, Likun Yang, Daipeng Li, Jianbin Jiao and Kaiqi Huang .	14691
<i>Learning to Translate by Translating: Stabilizing the Dual Loop via Semantic-Aware Self-Evolution</i>	
Kui Liu, Mingming Yin, Zuoli Tang, Zihao Li, Chilin Fu, Xiaolu Zhang, Jun Zhou, Lixin Zou and Chenliang Li	14711

<i>How to Allocate, How to Learn? Dynamic Rollout Allocation and Advantage Modulation for Policy Optimization</i>	
Yangyi Fang, Jiaye Lin, Xiaoliang Fu, Cong Qin, Haolin Shi, Chaowen Hu, Lu Pan, Ke Zeng and Xunliang Cai	14727
<i>Meeseeks: A Feedback-Driven, Iterative Self-Correction Benchmark evaluating LLMs' Instruction Following Capability</i>	
Jiaming Wang, Yunke Zhao, Peng Ding, Jun Kuang, Yibin Shen, Zhe Tang, Yilin Jin, ZongYu Wang, Xiaoyu Li and Xuezhi Cao	14745
<i>HSUGA: LLM-Enhanced Recommendation with Hierarchical Semantic Understanding and Group-Aware Alignment</i>	
Guorui Li, Dugang Liu, Lei Li, Xing Tang and Zhong Ming	14774
<i>AgenticEval: Toward Agentic and Self-Evolving Safety Evaluation of Large Language Models</i>	
Yixu Wang, Xin Wang, Yang Yao, Xinyuan Li, Xibang Yang, Yan Teng, Xingjun Ma and Yingchun Wang	14789
<i>AlphaEdit+: Model Editing in the Presence of Conflicting and Inconsistent Knowledge</i>	
Qing Liu, Jianhao Zhang, Ou Wu, Michael Ng and Yi Du	14809
<i>Spatiotemporal Sycophancy: Negation-Based Gaslighting in Video Large Language Models</i>	
Ziyao Tang, Pengkun Jiao, Bin Zhu, Huiyan Qi, Jingjing Chen and Yu-Gang Jiang	14836
<i>FreeChunker: A Cross-Granularity Chunking Framework</i>	
Zhang Wenxuan, Yuan-Hao Jiang, Yang Cao and Yonghe Wu	14853
<i>Invocation Refiner: A Plug-and-Play Module for Rectifying LLM Tool Invocations</i>	
Qirui Jiao, Dian Jiao, Nan Du, Ying Shen and Liang Lin	14866
<i>GoViG: Goal-Conditioned Visual Navigation Instruction Generation via Multimodal Reasoning</i>	
Fengyi Wu, Yifei Dong, Yilong Dai, Guangyu Chen, Qifeng Wu, Huiting Huang, Hang Wang, Qi Dai, Alexander G Hauptmann and Zhi-Qi Cheng	14899
<i>Temporal Contrastive Decoding: A Training-Free Method for Large Audio-Language Models</i>	
Yanda Li, Yuhao Liu, Zirui Song, Yunchao Wei, Martin Takáč and Salem Lahlou	14921
<i>When Internalization Fails: Finding Better Targets for Reasoning Compression</i>	
Mourad Heddaya, Manley Roberts, Rohan Wadhawan and Chenhao Tan	14935
<i>PatentMind: A Multi-Aspect Reasoning Graph for Patent Similarity Evaluation</i>	
Yongmin Yoo, Qionghai Xu and Longbing Cao	14947
<i>MemSearcher: Iterative Memory Integration for Search Agent via End-to-End Reinforcement Learning</i>	
Qianhao Yuan, Jie Lou, Zichao Li, Jiawei Chen, Yaojie Lu, Hongyu Lin, Le Sun, Debing Zhang and Xianpei Han	14965
<i>CoPA: Benchmarking Personalized Question Answering with Data-Informed Cognitive Factors</i>	
Hang Su, Zequn Liu, Chen Hu, Xuesong Lu, Yingce Xia and Liu Zhen	14978
<i>RAP-ID: Mechanistic Prompt Injection Detection via Impostor Behavior Analysis</i>	
Yuchen Yang, Lei Peng, Yujie He, Yang Yu, Zhongxin Wu and Yanlei Shi	15008
<i>When Safety Alignment Fails to Generalize: Probing with Language Game Jailbreaks</i>	
Zewen Long, Yu Peng, Fangming Dong, Congyi Li, Xingmao Guan, Shu Wu and Kai Chen	15020
<i>Self-EmoQ: Plutchik-Guided Value-based Planning to Drive Streaming Emotional TTS</i>	
Yue Zhao, Hongyan Li, Yong Chen and Luo Ji	15038

<i>Scaling is Not All You Need: Clinical-Oriented Reinforcement Learning Makes Parameter-Efficient Clinical Reasoning</i>	
Chi Liu, Yan Shu, Mengzhuo Chen, Hongming Piao, Zhijian Duan, Derek Li and Bryan Dai	15056
<i>Do We Need Distinct Representations for Every Speech Token? Unveiling and Exploiting Redundancy in Large Speech Language Models</i>	
Bajian Xiang, Tingwei Guo, Xuan Chen and Yang Han	15069
<i>Table-as-Search: Agentic Information Seeking is Table Completion</i>	
Tian Lan, Felix Henry, Bin Zhu, Qianghuai Jia, Junyang Ren, Qihang PU, Haijun Li, Longyue Wang, Zhao Xu and Weihua Luo	15088
<i>Minos: A Multimodal Evaluation Model for Bidirectional Generation Between Image and Text</i>	
Junzhe Zhang, Huixuan Zhang, Xinyu Hu, Li Lin, Mingqi Gao, Shi Qiu and Xiaojun Wan	15108
<i>Global Context or Local Detail? Adaptive Visual Grounding for Hallucination Mitigation</i>	
Yubo Jiang, Xin Yang, Abudukelimu Wuerkaixi, Zheming Yuan, Xuxin Cheng, Cao Liu, Ke Zeng, Fengying Xie, Zhiguo Jiang and Haopeng Zhang	15133
<i>FinToolSyn: A forward synthesis Framework for Financial Tool-Use Dialogue Data with Dynamic Tool Retrieval</i>	
Caishuang Huang, Yang Qiao, Rongyu Zhang, Junjie Ye, Pu Lu, Wuwenxi, Meng Zhou, Xiku Du, Qi Zhang, Tao Gui and Xuanjing Huang	15153
<i>From Fake to Real: Mitigating Out-of-Distribution Bias in In-Context Learning via Feedback Supervision from Large Language Models</i>	
Rui Song, Yingji Li, Jian Li, Fausto Giunchiglia and Hao Xu	15197
<i>Revealing and Enhancing Core Visual Regions: Harnessing Internal Attention Dynamics for Hallucination Mitigation in LLMs</i>	
Guangtao Lyu, Qi Liu, Chenghao Xu, Jiexi Yan, Muli Yang, Xueting Li, Fen Fang and Cheng Deng	15214
<i>Is Chain-of-Thought Reasoning of LLMs a Mirage? A Data Distribution Lens</i>	
Chengshuai Zhao, Zhen Tan, Pingchuan Ma, Dawei Li, Bohan Jiang, Yancheng Wang, Yingzhen Yang and Huan Liu	15231
<i>WindowsWorld: A Process-Centric Benchmark of Autonomous GUI Agents in Professional Cross-Application Environments</i>	
Jinchao Li, Yunxin li, Chenrui Zhao, Zhenran Xu, Baotian Hu and Min Zhang	15262
<i>EDU-CIRCUIT-HW: Evaluating Multimodal Large Language Models on Real-World University-Level STEM Student Handwritten Solutions</i>	
Weiyu Sun, Liangliang Chen, Yongnuo Cai, Huiru Xie, Yi Zeng and Ying Zhang	15281
<i>What Factors Affect LLMs and RLLMs in Financial Question Answering?</i>	
Peng Wang, Xuesi Hu, Jiageng Wu, Yuntao Zou, Qiancheng Zhang and Dagang Li	15315
<i>Superficial Success vs. Internal Breakdown: An Empirical Study of Generalization in Adaptive Multi-Agent Systems</i>	
Namyeong So, Seokgyu Jang and Taeuk Kim	15328
<i>D²-RAG: Dual-Decision Retrieval-Augmented Generation via Multi-Dimensional Uncertainty and Utility-Aware Decoding</i>	
Jinshuo Zhang, Xiaoding Zhou, Weiyu Zhang, Guoqiang Chen, Ying Lian, Xiaoyang Meng, Yonghe Chen, Hongjiao Guan, Jiasheng Si and Wenpeng Lu	15355

<i>ProFit: Leveraging High-Value Signals in SFT via Probability-Guided Token Selection</i>	
Tao Liu, Taiqiang Wu, Runming Yang, Shaoning Sun, Junjie Wang and Yujiu Yang	15383
<i>DisastQA: A Comprehensive Benchmark for Evaluating Question Answering in Disaster Management</i>	
Zhitong Chen, Kai Yin, Xiangjue Dong, Chengkai Liu, Xiangpeng Li, Bo Li, Junwei Ma, Yiming Xiao, Ali Mostafavi and James Caverlee	15402
<i>TAGS: A Test-Time Generalist–Specialist Framework with Retrieval-Augmented Reasoning and Verification</i>	
Jianghao Wu, Feilong Tang, Yulong Li, Ming Hu, Haochen Xue, Shoaib Jameel, Zongyuan Ge, Yutong Xie and Imran Razzak	15428
<i>How Retrieved Context Shapes Internal Representations in RAG</i>	
Samuel Yeh and Sharon Li	15446
<i>TabularMath: Understanding Math Reasoning over Tables with Large Language Models</i>	
Shi-Yu Tian, Zhi Zhou, Wei Dong, Kun-Yang Yu, Ming Yang, Zi-Jian Cheng, Lan-Zhe Guo and Yu-Feng Li	15469
<i>Seek-and-Solve: Benchmarking MLLMs for Visual Clue-Driven Reasoning in Daily Scenarios</i>	
Xiaomin Li, Tala Wang, Zichen Zhong, Ying Zhang, Zirui Zheng, Takashi Isobe, Dezhuang Li, Huchuan Lu, You He and Xu Jia	15499
<i>Structured and Abstractive Reasoning on Multi-modal Relational Knowledge Images</i>	
Yichi Zhang, Zhuo Chen, Lingbing Guo, Wen Zhang and Huajun Chen	15518
<i>Code Reffix: A Benchmark for Reflection-Guided Code Repair with Large Language Models</i>	
Zaiyuan Di, Jianting Chen, Yunxiao Yang, Xiaoying Gao, Li Yang, Zhihao Wang and Yang Xiang	15541
<i>ProtoCycle: Reflective Tool-Augmented Planning for Text-Guided Protein Design</i>	
Yutang Ge, Guojiang Zhao, Sihang Li, Zheng Cheng, Zifeng Zhao, Hanchen Xia, Guolin Ke, Linfeng Zhang, Zhifeng Gao and Yu Guang Wang	15562
<i>SEC-FinTables: Evaluating Large Language Models for Detecting Logical Inconsistencies on Tabular Data</i>	
Shuyan Ke, Qiong Wu, Hui Li and Liujuan Cao	15587
<i>QueryLink: Leveraging Query-Memory Alignment for Long-Term Reasoning in LLM Agents</i>	
Xuxian Hu, Zhu Teng, Wei Zhang, Ming He and Jianping Fan	15608
<i>AutoPKG: An Automated Framework for Dynamic E-commerce Product-Attribute Knowledge Graph Construction</i>	
Pollawat Hongwimol, Haoning Shang, Chutong Wang, Zhichao Wan, Yi Gao, Yuanming Li, Lin Gui, Wenhao Sun and Cheng Yu	15622
<i>Benchmarking LLMs on Authentic Cases from Medical Journals</i>	
Wanlong Liu, Junying Chen, Yunjin Yang, Prayag Tiwari, Wenyu Chen and Benyou Wang	15651
<i>Beyond Quantity: Trajectory Diversity Scaling for Code Agents</i>	
Guhong Chen, Chenghao Sun, Cheng Fu, Qiyao Wang, Zhihong Huang, ChaoPeng Wei, Guangxu Chen, Feiteng Fang, Ahmadreza Argha, Bing Zhao, Xander Xu, Qi Han, Hamid Alinejad-Rokny, Qiang Qu, Binhua Li, Shiwen Ni, Min Yang, HU Wei and Yongbin Li	15676
<i>Rethinking RL Evaluation: Can Benchmarks Truly Reveal Failures of RL Methods?</i>	
Zihan Chen, Yiming Zhang, Hengguang Zhou, Zenghui Ding, Yining Sun and Cho-Jui Hsieh	15692

<i>CaPEdit: Capability-Preserving Lifelong Knowledge Editing For Language Models</i> Song-Li Wu, Zhaocheng Du, Xianquan Wang and Jingyi Wang	15710
<i>Exploring Layer-wise Information Effectiveness for Post-Training Quantization in Small Language Models</i> He Xiao, Qingyao Yang, Dirui Xie, Wendong XU, Zunhai Su, Runming Yang, Haobo Liu, Wenyong Zhou, Zhengwu Liu and Ngai Wong	15753
<i>ELTLM: Evaluation of Longitudinal Temporal Large Multimodal Models in Clinical Scenarios</i> Gengyuan Hu, Haoxiang Liu, Chenhong Cao, Shilei Tan and Wei Gong	15765
<i>Easy Samples Are All You Need: Self-Evolving LLMs via Data-Efficient Reinforcement Learning</i> Zhiyin Yu, Bo Zhang, Qibin Hou, Zhonghai Wu, Xiao Luo and Lei Bai	15780
<i>Collaborative Multi-Agent Scripts Generation for Enhancing Imperfect-Information Reasoning in Murder Mystery Games</i> Keyang Zhong, Junlin Xie, Hefeng Wu, Haofeng Li and Guanbin Li	15796
<i>A Picture is Worth a Thousand Words? An Empirical Study of Aggregation Strategies for Visual Financial Document Retrieval</i> Ho Hung Lim and Yi Yang	15820
<i>ToolCPT: Improving Tool Utilization in LLM Agents via Continuous Pre-training</i> Yifan Yang, Jinghui Lu, Evadeng, Ao Yang, Peijie Yu, TingHao YU and Feng Zhang	15830
<i>A Universal Avoidance Method for Diverse Multi-branch Generation</i> Kyeongman Park, Minha Jhang and Kyomin Jung	15857
<i>CentaurTA: A Self-Improving Human-Agents Collaboration Framework for Thematic Analysis</i> Lei Wang, Min Huang and Eduard Dragut	15871
<i>Balancing Fidelity and Plasticity: Aligning Mixed-Precision Fine-Tuning with Linguistic Hierarchies</i> Changhai Zhou, Shiyang Zhang, Yuhua Zhou, Jun Gao, Qian Qiao, Shichao Weng, Weizhong Zhang and Cheng Jin	15885
<i>From Script to Stage: Automating Experimental Design for Social Simulations with LLMs</i> Yuwei Guo, Zihan Zhao, Xiaowei Liu, Xiangning Yu, Qun Ma, Deyu Zhou and Xiao Xue	15897
<i>Generative Giants, Retrieval Weaklings: Why do Multimodal Large Language Models Fail at Multimodal Retrieval?</i> Hengyi Feng, Zeang Sheng, Meiyi Qiang, Yang Li and Wentao Zhang	15917
<i>Learning to Extract Rational Evidence via Reinforcement Learning for Retrieval-Augmented Generation</i> Xinping Zhao, Shouzheng Huang, Yan Zhong, Xinshuo Hu, Meishan Zhang, Baotian Hu and Min Zhang	15934
<i>Tree-Notebook: A Context-Aware Agent with Tree Search and Entropy-Aware Data Shadow for Interactive Data Science</i> Junkun Qiu, Min Huang and Qinghai Miao	15957
<i>Resource-Limited Joint Multimodal Sentiment Reasoning and Classification via Chain-of-Thought Enhancement and Distillation</i> Haonan Shangguan, Xiaocui Yang, Shi Feng, Daling Wang, Yifei Zhang, Feiliang Ren and Ge Yu	15971

<i>Agentic-R: Learning to Retrieve for Agentic Search</i>	
Wenhan Liu, Xinyu Ma, Yutao Zhu, Yuchen Li, Daiting Shi, Dawei Yin and Zhicheng Dou	15987
<i>JailMeter: An Evidence-Based Evaluation Framework for Jailbreak Attacks on Large Language Models</i>	
Qingjia Huang, Jingyu Zhang, Jianguo Wu, Yakai Li, Weijuan Zhang, Yankai Rong, Junyi Yao, Shengzhi Zhang and Xiaoqi Jia	16006
<i>DataSeer: A Manager-Centric Collaborative Multi-Agent Framework with Multi-Branch Reasoning for Automated Insight Discovery</i>	
Suchen Liu, Yuanfeng Song, Jun Gao and Xing Chen	16030
<i>Penny Wise, Pixel Foolish": Bypassing Price Constraints in Multimodal Agents via Visual Adversarial Perturbations</i>	
Jiachen Qian and Zhaolu Kang	16059
<i>Beyond the Individual: Virtualizing Multi-Disciplinary Reasoning for Clinical Intake via Collaborative Agents</i>	
Huangwei Chen, Wu Li, Junhao Jia, Yining Chen, Xiaotao Pang, Ya-Long Chen, Li Gonghui, Haishuai Wang, Jiajun Bu and Lei Wu	16074
<i>SWE-Swiss: A Multi-Task Fine-Tuning and RL Recipe for High-Performance Issue Resolution</i>	
Zhenyu He, Qingping Yang, Wei Shen, Xiaojian Zhong, Kechi Zhang, Chenxin An, Wenlei Shi, Tianle Cai, Di He, Jiaze Chen and Jingjing Xu	16102
<i>BackdoorAgent: A Unified Framework for Backdoor Attacks on LLM-based Agents</i>	
Yunhao Feng, Yige Li, Yutao Wu, Yingshui Tan, Yanming Guo, Yifan Ding, Kun Zhai, Xingjun Ma and Yu-Gang Jiang	16115
<i>DKME: Rethinking Coupled Knowledge Memory for Lifelong Model Editing of Large Language Models</i>	
Guanyu Zheng, Wang Zhenyu, He Tingting, Xv Wang, Haochang Wang, Yaokai Huang and Tiejun Zhao	16128
<i>Self-Awareness before Action: Mitigating Logical Inertia via Proactive Cognitive Awareness</i>	
Fulong Fan, Peilin Liu, Liu Fengzhe, Shuyan Yang and Gang Yan	16151
<i>MuBench: Assessment of Multilingual Capabilities of Large Language Models Across 61 Languages</i>	
Wenhan Han, Yifan Zhang, Zhixun Chen, Binbinliu, Mykola Pechenizkiy, Meng Fang and Yin Zheng	16163
<i>PrefRAG: Correcting Semantic Errors in Auto-Formalization for Logical Reasoning with Program Preference RAG</i>	
Yuyin Zhou and Yongmei Liu	16193
<i>LoopCoder: Scaling Code Intelligence via Looped Language Models</i>	
Jian Yang, Wei Zhang, Shuyue Guo, Yizhi LI, Linzheng Chai, Zhengmao Ye, Shukai Liu, Yuyang Song, Jiajun Wu, Che Liu, Tianyu Zheng, Siwei Wu, Leo L, Xudong Ma, Chuan Hao, Ran Tao, Yan Xing, Jianzhou Wang, Mingjie Tang, Aishan Liu, Zhoujun Li, Xianglong Liu, Weifeng Lv and Bryan Dai	16209
<i>TempTool-RL: Tool-Augmented Reinforcement Learning for Temporal Knowledge Graph Question Answering</i>	
Zicheng Huang, Yajuan Tong, Xinhui Tu and Tingting He	16224
<i>Lightweight Haar Wavelet Subband Pruning for LLMs</i>	
Jiang Li, Pengfei Cao, Chenxi Zhou, Tian Lan, Xiangdong Su, Kang Liu, Jun Zhao and Guanglai Gao	16242

<i>Prompt-R1: Collaborative Automatic Prompting Framework via End-to-end Reinforcement Learning</i> Wenjin Liu, Haoran Luo, Xueyuan Lin, Haoming Liu, Tiesunlong Shen, Jiapu Wang, Rui Mao and Erik Cambria	16260
<i>ProCeedRL: Process Critic with Explorative Demonstration Reinforcement Learning for LLM Agentic Reasoning</i> Jingyue Gao, Yanjiang Guo, Chen Xiaoshuai and Jianyu Chen	16281
<i>Accelerating Prefilling via Decoding-time Contribution Sparsity</i> Zhiyuan He, Yike Zhang, Chengruidong Zhang, Huiqiang Jiang, Yuqing Yang and Lili Qiu	16296
<i>SpecBound: Adaptive Bounded Self-Speculation with Layer-wise Confidence Calibration</i> Zhuofan Wen and Yang Feng	16309
<i>Stable and Explainable Personality Trait Evaluation in Large Language Models with Internal Activations</i> Xiaoxu Ma, Xiangbo Zhang and Zhenyu Weng	16322
<i>Learning Continuous Temporal Dynamics on Symplectic Manifolds for Temporal Knowledge Graph Embedding</i> Jiang Li, Zehua Duo, Tian Lan, Feilong Bao, Guanglai Gao and Xiangdong Su	16341
<i>mPresenter: An Agentic Framework for Generating Multilingual Presentation Videos from Scientific Papers</i> Wenhan Han, Xiao Xiao, Mykola Pechenizkiy and Meng Fang	16358
<i>Tree-CoT-RT: An Explainable Multi-Path Tree-Guided Chain-of-Thought and Reinforcement Learning Framework for Aspect Sentiment Quad Prediction</i> Hao Zhang, Jiahao Wang, Zhenke Duan, Xin Yin, Haichuan Hu, Hualong Chen, Suyi, Congqing He, Yike Tan and Yu-N Cheah	16372
<i>PRISM-MCTS: Learning from Reasoning Trajectories with Metacognitive Reflection</i> Siyuan Cheng, Bozhong Tian, Yanchao Hao and Zheng Wei	16392
<i>CausalityCheck: A Framework for Evaluating Causal Reasoning in Large Language Models</i> Jiang Li, Zehua Duo, Guanglai Gao and Xiangdong Su	16403
<i>OmniDiagram: Advancing Unified Diagram Code Generation via Visual Interrogation Reward</i> Haoyue Yang, Xuanle Zhao, Xuexin Liu, Feibing Jiang and Yao Zhu	16430
<i>Why Agents Compromise Safety Under Pressure</i> Hengle Jiang and Ke Tang	16453
<i>CoMoL: Efficient Mixture of LoRA Experts via Dynamic Core Space Merging</i> Jie Cao, Zhenxuan Fan, Zhuonan Wang, Tianwei Lin, Ziyuan Zhao, Rolan Yan, Wenqiao Zhang, Feifei Shao, Hongwei Wang, Jun Xiao and Siliang Tang	16471
<i>QuCo-RAG: Quantifying Uncertainty from the Pre-training Corpus for Dynamic Retrieval-Augmented Generation</i> Dehai Min, Kailin Zhang, Tongtong Wu and Lu Cheng	16482
<i>CaM-HG: Causal-Enhanced MoE and Hypergraphs Network for Incomplete Multimodal Emotion Recognition in Conversations</i> Mingjian Yang, Yong Wang, Peng Liu and Wen Yin	16501
<i>Large-Scale Diverse Synthesis for Mid-Training</i> Xuemiao Zhang, Chengying Tu, Can Ren, Rongxiang Weng, Hongfei Yan, Jingang Wang and Xunliang Cai	16517

<i>DiSec: Mitigating Backdoors in Pre-trained Language Models via Disentanglement of Adversarial Weights for Secure Fine-Tuning</i>	
Sunanda Das and Qinghua Li	16540
<i>LUNE: Efficient LLM Unlearning via LoRA Fine-Tuning with Negative Examples</i>	
Yezi Liu, Hanning Chen, Wenjun Huang, Yang Ni and Mohsen Imani	16560
<i>Think Earlier, Not Longer: Prompt Optimization via Reducing Unhealthy Exploration</i>	
Ling-I Wu, Minyu Chen, Jingyang Li, Xi Chang and Guoqiang Li	16577
<i>Psychological Counseling Cannot Be Achieved Overnight: Automated Psychological Counseling Through Multi-Session Conversations</i>	
Bichen Wang, Junzhe Wang, Yixin Sun, Xing Fu, Yanyan Zhao and Bing Qin	16593
<i>Code over Words: Overcoming Semantic Inertia via Code-Grounded Reasoning</i>	
Manjie Xu, Isabella Yin, Xinyi Tu, Chi Zhang and Yixin Zhu	16610
<i>Bypassing Neural Evaluations for Fast Audio Editing via Adaptive Trajectory Extrapolation</i>	
Xiaoqian Liu, Zhengkun Ge, Jianjin Wang, Haoran Zhang, Yuan Ge, Kaiyan Chang, Chen Xu, Tong Xiao, Zhengtao Yu, Linfeng Zhang and JingBo Zhu	16633
<i>Agent-GWO: Collaborative Agents for Dynamic Prompt Optimization in Large Language Models</i>	
Xudong Wang, Chaoning Zhang, Chenghao Li, Shuxu Chen, Qigan Sun, Jiaquan Zhang, Fachrina Dewi Puspitasari, Tae-Ho Kim, Jiwei Wei, Malu Zhang, Guoqing Wang, Yang Yang and Heng Tao Shen	16648
<i>Structured Confidence-Guided Online Adaptation for LLM-based Multi-Label Classification</i>	
Pengyu Xu, JingRen Hou, Liping Jing and Jian Yu	16671
<i>SERM: Self-Evolving Relevance Model with Agent-Driven Learning from Massive Query Streams</i>	
Chenglong Wang, Canjia Li, Xingzhao Zhu, Yifu Huo, Huiyu Wang, Weixiong Lin, Yun Yang, Qiaozhi He, Tian Hua Zhou, Changxiaojia, JingBo Zhu and Tong Xiao	16687
<i>CLAG: Adaptive Memory Organization via Agent-Driven Clustering for Small Language Model Agents</i>	
Taeyun Roh, WonJune Jang, Junha Jung and Jaewoo Kang	16707
<i>ClinAlign: Scaling Healthcare Alignment from Clinician Preference</i>	
Shiwei Lyu, Xidong Wang, Hao Zhu, Lei Liu, Chaohe Zhang, Jian Wang, Jinjie Gu, Benyou Wang and Yue Shen	16727
<i>SessionIntentBench: A Multi-task Inter-session Intention-shift Modeling Benchmark for E-commerce Customer Behavior Understanding</i>	
Yuqi Yang, Weiqi Wang, Baixuan Xu, Wei Fan, Qing Zong, Chunkit Chan, Zheyue Deng, Xin Liu, Yifan Gao, Changlong Yu, Chen Luo, Yang Li, Zheng Li, Qingyu Yin, Bing Yin and Yangqiu Song	16748
<i>Adapting Where It Matters: Depth-Aware Adaptation for Efficient Multilingual Speech Recognition in Low-Resource Languages</i>	
Yang Xiao, Eun-Jung Holden and Ting Dang	16776
<i>Targeted Exploration via Unified Entropy Control for Reinforcement Learning</i>	
Chen Wang, Lai Wei, Yanzhi Zhang, Chenyang Shao, Zedong Dan, Weiran Huang, Ge Lan and Yue Wang	16789
<i>Remember Me, Refine Me: A Dynamic Procedural Memory Framework for Experience-Driven Agent Evolution</i>	
Zouying Cao, Jiaji Deng, LI YU, Weikang Zhou, Zhaoyang Liu, Bolin Ding and Hai Zhao	16803

<i>Knowing When to Quit: Diagnosing and Training LLMs to Abort Futile Reasoning</i> Xinyan Guan, Jiali Zeng, Chunlei Xin, Yaojie Lu, Hongyu Lin, Xianpei Han, Le Sun and Fandong Meng	16823
<i>AutoL2S: Auto Long-Short Reasoning for Efficient Large Language Models</i> Feng Luo, Yu-Neng Chuang, Guanchu Wang, Hoang Anh Duy Le, Shaochen Zhong, Hongyi Liu, Jiayi Yuan, Yang Sui, Vladimir Braverman, Vipin Chaudhary and Xia Hu	16836
<i>Doc-PP: Document Policy Preservation Benchmark for Large Vision-Language Models</i> Haeun Jang, Hwan Chang and Hwanhee Lee	16859
<i>RIPRAG: Hack a Black-box Retrieval-Augmented Generation Question-Answering System with Reinforcement Learning</i> Meng Xi, Sihan Lv, Yechen Jin, Guanjie Cheng, Naibo Wang, Ying Li and Jianwei Yin ...	16882
<i>CodeMEM: AST-Guided Adaptive Memory for Repository-Level Iterative Code Generation</i> Peiding Wang, Li Zhang, Fang Liu, Chongyang Tao and Yinghao Zhu	16903
<i>HalluGuard: Evidence-Grounded Small Reasoning Models to Mitigate Hallucinations in Retrieval-Augmented Generation</i> Loris Bergeron, Ioana Buhnila, Jerome Francois and Radu State	16918
<i>Capabilities and Evaluation Biases of Large Language Models in Classical Chinese Poetry Generation: A Case Study on Tang Poetry</i> Bolei Ma, Yina Yao and Anna-Carolina Haensch	16933
<i>SWE-QA-Pro: A Representative Benchmark and Scalable Training Recipe for Repository-Level Code Understanding</i> Songcheng Cai, Zhiheng Lyu, Yuansheng Ni, Xiangchao Chen, Baichuan Zhou, Shenzhe Zhu, Yi Lu, Haozhe Wang, Chi Ruan, Benjamin Schneider, Weixu Zhang, Xiang Li, Andy Zheng, Yuyu Zhang, Ping Nie and Wenhui Chen	16958
<i>Do MLLMs Understand Pointing? Benchmarking and Enhancing Referential Reasoning in Egocentric Vision</i> Chentao Li, Zirui Gao, Mingze Gao, Yinglian Ren, Jianjiang Feng and Jie Zhou	17000
<i>Can Post-Training Transform LLMs into Causal Reasoners?</i> Junqi Chen, Sirui Chen and Chaochao Lu	17020
<i>MoE Adapter for Large Audio Language Models: Sparsity, Disentanglement, and Gradient-Conflict-Free</i> Yishu Lei, Shuwei He, Hu Jing, Dan Zhang, Xianlong Luo, Danxiang Zhu, Shikun Feng, Rui Liu, Jingzhou HE, Yu Sun, Hua Wu and Haifeng Wang	17039
<i>QaRL: Rollout-Aligned Quantization-Aware RL for Fast and Stable Training under Training-Inference Mismatch</i> Hao Gu, Hao Wang, Jiacheng Liu, Lujun Li, Qiyuan Zhu, Bei Liu, Binxing Xu, Lei Wang, Xintong Yang, Sida Lin, Sirui Han and Yike Guo	17051
<i>Beyond Static Rules: Automated Discovery of Latent Vulnerabilities in Text-to-SQL</i> Hanqing Wang, Yongdong Chi, Jian Yang, Lei Yang, Jiehui Zhao, Yun Chen and Guanhua Chen	17065
<i>Can We Entrust Justice to AI?: How Persona Traps Contaminate Reasoning in Criminal Investigation</i> Jaewook Lee, Myeong-Cheol Kang and Jong-hun Shin	17083

<i>Beyond Semantic Similarity: Appraisal-Guided Chain-of-Thought Reasoning and Retrieval for Multi-modal Emotional Support Conversations</i>	
Yuqi Chu, Lizi Liao, Jinggui Liang, Boyang Li and Richang Hong	17104
<i>When and What to Ask: AskBench and Rubric-Guided RLVR for LLM Clarification</i>	
Jiale Zhao, Ke Fang and Lu Cheng	17120
<i>FreezeEmpath: Efficient Training for Empathetic Spoken Chatbots with Frozen LLMs</i>	
Yun Hong, Yan Zhou and Yang Feng	17141
<i>To Intervene or Not: Guiding Inference-time Alignment with Probabilistic Model Blending</i>	
Jin Gan, Xin Li and Jun Luo	17158
<i>FLEXITOKENS: Flexible Tokenization for Evolving Language Models</i>	
Abraham Toluwase Owodunni, Orevaoghene Ahia and Sachin Kumar	17170
<i>RAG in the Wild: On the (In)effectiveness of LLMs with Mixture-of-Knowledge Retrieval Augmentation</i>	
Ran Xu, Yuchen Zhuang, Yue Yu, Haoyu Wang, Wenqi Shi and Carl Yang	17191
<i>Chain-of-Procedure: Hierarchical Visual-Language Reasoning for Procedural QA</i>	
Guanhua Chen, Yutong Yao, Shenghe Sun, Ci-jun Gao, Shudong Liu, Lidia S. Chao, Feng Wan and Derek F. Wong	17207
<i>PEGRL: Improving Machine Translation by Post-Editing Guided Reinforcement Learning</i>	
Yunzhi Shen, Hao Zhou, Xin Huang, Xue Han, Junlan Feng and Shujian Huang	17225
<i>RanLoRA: Residual-aware Nonlinear Low-Rank Adaptation</i>	
Xu Luo, Yongbin Liu, Chunping Ouyang and Ying Yu	17243
<i>Multi-Hop Knowledge Editing via Critic-Guided Multi-Agent Reasoning</i>	
Xudong Li, Yuhang Tian, Dandan Song, Zhijing Wu, Shuhao Zhang, Jun Yang, Yongyu Huo, Changzhi Zhou, Xinyu Zhang, Chenhao Li, Huipeng Ma, Luan Zhang, Yan Xu and Qian Liu	17259
<i>Rethinking Post-Unlearning Behavior of Large Vision-Language Models</i>	
Minsung Kim, Nakyeong Yang and Kyomin Jung	17277
<i>Scattered Hypothesis Generation for Open-Ended Event Forecasting</i>	
He Chang, Zhulin Tao, Lifang Yang, Xianglin Huang and Yunshan Ma	17288
<i>Context-attended Adversarial Reinforcement Learning for Robust Multi-step Retrieval Augmented Generation</i>	
Yingtao Ren, Xiao Luo, Yu-Cheng Chang and Chin-teng Lin	17305
<i>From Pseudo-Balancing to True Specialization: Memory-Aware Routing for Mixture-of-Experts</i>	
Peixuan Hou, Yunbo Hou, Bin Chen, LI He, Jian Xu, Weiping Li, Bo Zheng and Guojie Song	17320
<i>Learning to Retrieve User History and Generate User Profiles for Personalized Persuasiveness Prediction</i>	
Sejun Park, Yoonah Park, Jongwon Lim and Yohan Jo	17338
<i>Understanding LLM Reasoning for Abstractive Summarization</i>	
Haohan Yuan and Haopeng Zhang	17360
<i>Feasible is Not Enough: Cost-Aware Optimal Tool-Chain Planning on Multi-Solution Tool Graphs</i>	
Liangliang Liu, Yanming Li, Yigang Liu, Jialong Han, Rujia Shen, Yi Guan, Yi Lin and Jingchi Jiang	17387

<i>DOS: Dependency-Oriented Sampler for Masked Diffusion Language Models</i> Xueyu Zhou, Yangrong Hu and Jian Huang	17404
<i>BIASEDTALES-ML: A Multilingual Dataset for Analyzing Narrative Attribute Distributions in LLM-Generated Stories</i> Yuxuan Ouyang, Yingfeng Luo, JingBo Zhu and Tong Xiao	17420
<i>Rethinking Reasoning: A Survey on Reasoning-based Backdoors in LLMs</i> Man Hu, Xinyi Wu, Zhufeng Suo, Jinbo Feng, Linghui Meng, Yanhao Jia, Anh Tuan Luu and Shuai Zhao	17437
<i>From Tokens to Steps: Verification-Aware Speculative Decoding for Efficient Multi-Step Reasoning</i> Kiran Purohit, Ramasuri Narayanam and Soumyabrata Pal	17457
<i>SPS: Steering Probability Squeezing for Better Exploration in Reinforcement Learning for Large Language Models</i> Yifu Huo, Chenglong Wang, Ziming Zhu, Shunjie Xing, Peinan Feng, Tongran Liu, Qiaozhi He, Tian Hua Zhou, Changxiaojia, JingBo Zhu, Zhengtao Yu and Tong Xiao	17472
<i>Memp: Exploring Agent Procedural Memory</i> Runnan Fang, Yuan Liang, Xiaobin Wang, Jialong Wu, Shuofei Qiao, Pengjun Xie, Fei Huang, Huajun Chen and Ningyu Zhang	17490
<i>ATLAS: Orchestrating Heterogeneous Models and Tools for Multi-Domain Complex Reasoning</i> Jinyang Wu, Guocheng Zhai, Ruihan Jin, Jiahao Yuan, Yuhao Shen, Shuai Zhang, Zhengqi Wen and Jianhua Tao	17503
<i>SWE-AGILE: A Software Agent Framework for Efficiently Managing Dynamic Reasoning Context</i> Shuquan Lian, Juncheng Liu, Yazhe Chen, Yuhong Chen and Hui Li	17536
<i>CAR: Empowering Agents with Dynamic Tool Synthesis and Global Trajectory Rectification</i> Kai Wang, Xu Wang, Zhiyuan Fu, Yudong Zhang and Yang Wang	17551
<i>FastDiSS: Few-step Match Many-step Diffusion Language Model on Sequence-to-Sequence Generation</i> Dat Nguyen Cong, Tung Kieu and Hoang Thanh-Tung	17572
<i>GCIG: GraphRAG-based Cross-document Instruction Generation for Boosting LLM Reasoning</i> Xiaoliang Xu, Huang Yuan, Junmei Wang and Can Xu	17593
<i>Towards General Agentic Intelligence via Environment Scaling</i> Runnan Fang, Shihao Cai, Baixuan Li, Jialong Wu, Guangyu Li, Wenbiao Yin, Xinyu Wang, Xiaobin Wang, Liangcai Su, Zhen Zhang, Shibin Wu, Zhengwei Tao, Yong Jiang, Pengjun Xie, Ningyu Zhang, Fei Huang, Wentao Zhang and Jingren Zhou	17610
<i>EAIR: Entity-aware Inference-Time Knowledge Routing for Multi-Hop Knowledge Editing</i> Jungyu Lee, Kunhui Lee, Gyun Lee and Seung-Hoon Na	17622
<i>Preventing Safety Drift in Large Language Models via Coupled Weight and Activation Constraints</i> Songping Peng, Zhiheng Zhang, Daojian Zeng, Lincheng Jiang and Xieping Gao	17646
<i>GAPO: Robust Advantage Estimation for Real-World Code LLMs</i> Jianqing Zhang, Zhezheng Hao, Wei Xia, Hande Dong, Hong Wang, Chenxing Wei, Yuyan Zhou, Yubin Qi, Qiang Lin and Jian Cao	17663
<i>What's Missing in Screen-to-Action? Towards a UI-in-the-Loop Paradigm for Multimodal GUI Reasoning</i> Songze Li, Xiaoke Guo, Tianqi Liu, Biao Yi, Zhaoyan Gong, Zhiqiang Liu, Huajun Chen and Wen Zhang	17674

<i>When to Trust Tools? Adaptive Tool Trust Calibration For Tool-Integrated Math Reasoning</i>	
Ruotao Xu, Yixin Ji, Yu Luo, Jinpeng Li, Dong Li, Peifeng Li, Juntao Li and Min Zhang . . .	17691
<i>V-MAGE: A Game Evaluation Framework for Assessing Vision-Centric Capabilities in Multimodal Large Language Models</i>	
Xiangxi Zheng, Linjie Li, Zhengyuan Yang, Ping Yu, Alex Jinpeng Wang, Rui Yan, Yuan Yao and Lijuan Wang	17707
<i>Understanding and Preventing Entropy Collapse in RLVR with On-Policy Entropy Flow Optimization</i>	
Huimin Xu, Shuai Zhao, Xiaobao Wu and Anh Tuan Luu	17759
<i>Reinforcement Learning with Semantic Rewards Enables Low-Resource Language Expansion without Alignment Tax</i>	
Zeli Su, Ziyin Zhang, Zhou Liu, Xuexian Song, Zhankai Xu, Longfei Zheng, Xiaolu Zhang, Rong Fu, Guixian Xu and Wentao Zhang	17772
<i>MEIC-DT: Memory-Efficient Incremental Clustering for Long-Text Coreference Resolution with Dual-Threshold Constraints</i>	
Kangyang Luo, Shuzheng Si, Yuzhuo Bai, Cheng Gao, Zhitong Wang, Cheng Huang, Yingli Shen, Yufeng Han, Wenhao Li, Cunliang Kong and Maosong Sun	17787
<i>Lost in Diffusion: Uncovering Hallucination Patterns and Failure Modes in Diffusion Large Language Models</i>	
Zhengnan Guo and Fei Tan	17805
<i>FARSS: Fisher-Optimized Adaptive Low-Rank and Singular-Vector Selection for Knowledge-Preserving Fine-Tuning</i>	
Renxing Chen, Ziwei Xiang, Peisong Wang, Hongjian Fang, Meng Li, Fanhu Zeng, Yanan Zhu, Peipei Yang, Xu-Yao Zhang and Jian Cheng	17819
<i>Align Documents to Questions: Question-Oriented Document Rewriting for Retrieval-Augmented Generation</i>	
Jiaang Li, Zhendong Mao, Quan Wang, Yuning Wan and Yongdong Zhang	17838
<i>GlimpRouter: Efficient Collaborative Inference by Glimpsing One Token of Thoughts</i>	
Wenhao Zeng, Xuteng Zhang, Yuling Shi, Chao Hu, Yuting Chen, Beijun Shen and Xiaodong Gu	17850
<i>Plug-and-Play Data Module for Code RL: Adaptive Ambiguity Replay</i>	
Jianqing Zhang, Wei Xia, Zhezheng Hao, Hong Wang, Hande Dong, Qiang Lin, Yang Liu, Jian Cao and Qiang Yang	17865
<i>Can AI Revise Research Papers with Human Review Feedback? An Empirical Study and Benchmark</i>	
Zihan Luo, Hong Huang, Jianxun Lian, Yu Chang, Xing Xie and Hai Jin	17876
<i>Dunhuang-Bench: How Well Do MLLMs Understand Cultural Heritage?</i>	
Junyi Yuan, Jian Zhang, Tianxiu Yu, Yanlin Zhou, Xiaobo Jin, Qiufeng Wang and Fangyu Wu	17894
<i>Multi-Docker-Eval: A ‘Shovel of the Gold Rush’ Benchmark on Automatic Environment Building for Software Engineering</i>	
Kelin Fu, Tianyu Liu, Zeyu Shang, Yingwei MA, Jiaheng Liu, Jian Yang and Kaigui Bian	17911
<i>Not All Modalities at Once: Dynamic Dropout and Bidirectional Fusion for Robust Multi-modal Knowledge Graph Completion</i>	
Jiashun Peng, Fu Zhang, Hongzhi Chen, Jingwei Cheng, Yingsong Ning and Xiaoke Wang	17928

<i>If Only My CGM Could Speak: A Privacy-Preserving Agent for Question Answering over Continuous Glucose Data</i>	
Yanjun Cui, Ali Emami, Temiloluwa Prioleau and Nikhil Singh	17941
<i>E-ABSA20K: A Dataset and Propose-and-Verify for Aspect-Based Sentiment Analysis in Long E-commerce Reviews</i>	
Tong Sun, Mingyang Ma and Cheng Yu	17959
<i>ECHA: Jailbreaking LVLMS via the Mismatch between Implicit Semantic Reconstruction and Explicit Safety Alignment</i>	
Chenxing Xu, Junyong Jiang, Zehu Zhang and Lu Dong	17974
<i>Thinking with Map: Reinforced Parallel Map-Augmented Agent for Geolocalization</i>	
Yuxiang Ji, Yong Wang, Ziyu Ma, Yiming Hu, Hailang Huang, Xuecai Hu, Guanhua Chen, Liaoni Wu and Xiangxiang Chu	17991
<i>Revisiting Entropy Regularization: Adaptive Coefficient Unlocks Its Potential for LLM Reinforcement Learning</i>	
Xiaoyun Zhang, Xiaojian Yuan, Di Huang, Wang You, Chen Hu, Jingqing Ruan, Kejiang Chen and Xing Hu	18005
<i>SLIP: Soft Label Mechanism and Key-Extraction-Guided CoT-based Defense Against Instruction Backdoor in APIs</i>	
Wu Zhengxian, Juan Wen, Wanli Peng, Haowei Chang, Yinghan Zhou and Xue Yiming ...	18021
<i>U-Fold: Dynamic Intent-Aware Context Folding for User-Centric Agents</i>	
Jin Su, Runnan Fang, Yeqiu Li, Xiaobin Wang, Shihao Cai, Pengjun Xie, Ningyu Zhang and Fajie Yuan	18043
<i>AdaTooler-V: Adaptive Tool-Use for Images and Videos</i>	
Chaoyang Wang, Kaituo Feng, Dongyang Chen, Zhongyu Wang, Zhixun Li, Sicheng Gao, Meng Meng, Xu Zhou, Manyuan Zhang, Yuzhang Shang and Xiangyu Yue	18058
<i>Exposing Privacy Risks in Graph Retrieval-Augmented Generation</i>	
Jiale Liu, Jiahao Zhang and Suhang Wang	18073
<i>BalanceSFT: Improving LLM Function Calling with Balanced Training Signals and Data Hardness</i>	
Bingguang Hao, Zengzhuang Xu, Maolin Wang, Yuntao Wen, Yicheng Chen, Cunyin Peng, Long Chen, Xiangyu Zhao, Jinjie Gu, Chenyi Zhuang and Ji Zhang	18094
<i>How Robust Are Large Language Models for Clinical Numeracy? An Empirical Study on Numerical Reasoning Abilities in Clinical Contexts</i>	
Minh-Vuong Nguyen, Fatemeh Shiri, Zhuang Li and Karin Verspoor	18113
<i>Safety-Aware Dialogue System for Postoperative Oral Cancer Care with Structured Clarification and a Clinically Curated Dataset</i>	
Tzu-Chi Liu, Hui-Ying Yang, Shiow-Ching Shun, Yu-Chi Chen, Lu-Yen Anny Chen and Yong-Sheng Chen	18130
<i>FTibSuite: A Comprehensive Resource Suite for Tibetan Vision–Language Modeling</i>	
Guixian Xu, Yide Liang, Zeli Su, Xuexian Song, Ziyin Zhang, Yushuang Dong, Ting Zhang and Xu Han	18143
<i>DeepMed: Building a Medical DeepResearch Agent via Multi-hop Med-Search Data and Turn-Controlled Agentic Training & Inference</i>	
Zihan Wang, Hao Wang, Shi Feng, Xiaocui Yang, Daling Wang, Yiqun Zhang, Jinghao Lin, Xiaozhong Ji and Haihua Yang	18160

<i>Softpick: No Attention Sink, No Massive Activations with Rectified Softmax</i>	
Zayd Muhammad Kawakibi Zuhri, Erland Hilman Fuadi and Alham Fikri Aji	18179
<i>UniCM: A Unified Consistency Model For Efficient Multimodal Generation and Understanding</i>	
Chenkai Xu, Xu Wang, Zhenyi Liao, Yishun Li, TianQi Hou and Zhijie Deng	18214
<i>Can Multi-agent Help Disambiguation in Multi-domain Translation?</i>	
Zhibo Man, Shaoyang Xu, Yujie Zhang, Yi Feng, Yuanmeng Chen, Yufeng Chen, Xu Jinan and Wenxuan Zhang	18234
<i>Agent-Dice: Disentangling Knowledge Updates via Geometric Consensus for Agent Continual Learning</i>	
Zheng Wu, Xingyu Lou, Xinbei Ma, Yansi Li, Weiwen Liu, Weinan Zhang, Jun Wang and Zhuosheng Zhang	18245
<i>Automatic Slide Updating with User-Defined Dynamic Templates and Natural Language Instructions</i>	
Kun Zhou, Jiakai He, Wenmian Yang, Zhensheng Wang, Yiquan Zhang and Weijia Jia	18263
<i>PretrainRL: Alleviating Factuality Hallucination of Large Language Models at the Beginning</i>	
Langming Liu, Kangtao Lv, Haibin Chen, Weidong Zhang, Yejing Wang, Shilei Liu, Xin Tong, Yujin Yuan, Yongwei Wang, Wenbo Su and Bo Zheng	18298
<i>OMHBench: Benchmarking Balanced and Grounded Omni-Modal Multi-Hop Reasoning</i>	
Seunghee Kim, Ingyu Bang, Seokgyu Jang, Changhyeon Kim, Sanghwan Bae, Jihun Choi, Richeng Xuan and Taeuk Kim	18311
<i>Beyond Uniform SVD: Dual-Level Optimization across Columns and Modules for LLM Compression</i>	
Lin Xv, Xian Gao, Ting Liu and Yuzhuo fu	18335
<i>RSA-Bench: Benchmarking Audio Large Models in Real-World Acoustic Scenarios</i>	
Yibo Zhang, Kaiwen Luo, Liang Lin, Shilinlu Yan, Jin Wang, Yaoqi Guo, Yitian Chen, Yalan Qin, Zhenhong Zhou, Kun Wang and Li Sun	18350
<i>MedLayBench-V: A Large-Scale Benchmark for Expert-Lay Semantic Alignment in Medical Vision Language Models</i>	
Han Jang, Junhyeok Lee, Heeseong Eum and Kyu Sung Choi	18375
<i>Divergent Thinking: Escape the Homogeneity Trap in Generative Commonsense Reasoning</i>	
Yiheng Tao, Kaiwen Cheng, Zhiwei Nie, Chang Liu and Jie Chen	18395
<i>SafeSteer: A Decoding-level Defense Mechanism for Multimodal Large Language Models</i>	
Xinyi Zeng, Xue Yang, Jingyuan Zhang, Huanqian Yan, Xiang Chen, Kaiwen Wei, Hankun Kang and Yu Tian	18411
<i>SEMA-RAG: A Self-Evolving Multi-Agent Retrieval-Augmented Generation Framework for Medical Reasoning</i>	
Yongfeng Huang, Ruiying Chen and James Cheng	18423
<i>On the Step Length Confounding in LLM Reasoning Data Selection</i>	
Bing Wang, Rui Miao, Chen Shen, Shaotian Yan, Kaiyuan Liu, Ximing Li, Xiaosong Yuan, Sinan Fan, Jun Zhang and Jieping Ye	18443
<i>SLIM: Stealthy Low-Coverage Black-Box Watermarking via Latent-Space Confusion Zones</i>	
Hengyu WU and Yang Cao	18458

<i>SQL-ASTRA: Alleviating Sparse Feedback in Agentic SQL via Column-Set Matching and Trajectory Aggregation</i>	
Long Li, Zhijian Zhou, Jiangxuan Long, Peiyang Liu, Weidi Xu, Zhe Wang, Shirui Pan and Chao Qu	18474
<i>Rotation Control Unlearning: Quantifying and Controlling Continuous Unlearning for LLM with The Cognitive Rotation Space</i>	
Xiang Zhang, Kun Wei, Xu Yang, Jiahua Li, Su Yan and Cheng Deng	18492
<i>SpecEdit: A Spectral Approach for Multi-Round Knowledge Editing</i>	
Junxian Liu, Zikun Deng, Xin Wu, Dazhen Deng and Yi Cai	18511
<i>SAGE: Sign-Adaptive Gradient for Memory-Efficient LLM Optimization</i>	
Woojin Lee and Hyun-Tae Kim	18525
<i>DP³: Differentially Private Prompt Perturbation for Multi-turn LLM Inference</i>	
Lele Zheng, Chao Zhang, Feiyang Yuan, Ke Cheng, Tao Zhang, Anxiao Song and Yulong Shen	18538
<i>RAMP: Risk-Aware Multi-Turn Planning for Jailbreak Red-Teaming</i>	
Yize Liu, Yunyun Hou and Aina Sui	18553
<i>LexGenius: An Expert-Level Benchmark for Large Language Models in Legal General Intelligence</i>	
Wenjin Liu, Haoran Luo, Xin Feng, Xiang Ji, Lijuan Zhou, Rui Mao, Jiapu Wang, Shirui Pan and Erik Cambria	18565
<i>From Attenuation to Attention: Variational Information Flow Manipulation for Fine-Grained Visual Perception</i>	
Jilong Zhu and Yang Feng	18586
<i>VulAgent: Hypothesis-Validation Driven Multi-Agent Architecture for Vulnerability Detection</i>	
Ziliang Wang, Ge Li, Jia Li, Hao Zhu and Zhi Jin	18598
<i>From Attack Surfaces to Actual Operations: A Survey of Modern LLM Jailbreaks</i>	
Ruikang Zhou, Changsheng Sun and Mark Huasong Meng	18617
<i>ARC: Active and Reflection-driven Context Management for Long-Horizon Information Seeking Agents</i>	
Yilun Yao, Shan Huang, Elsie Dai, Zhewen Tan, Zhenyu Duan, Shousheng Jia, Yanbing Jiang and Tong Yang	18644
<i>ToolDNA: Autonomous Evolution of Tool Metadata for Robust Dialogue Agents</i>	
Qiuyuan Ai, Cong Wang, Jiaqi Zhang, Zengxin Han and Jie Song	18660
<i>ExpertIVS: Sociological Expert Driven Individual Value Simulation in Large Language Models</i>	
Zhen Wang, Yuqi Ren, Yuehan Cui, Hongxiang Wang, Jianxiang Peng, Zhaoxia Zhang, Bingkun Zhu, Tongxuan Zhang, Dezhi Tong and Deyi Xiong	18679
<i>Be Your Own Red Teamer: Safety Alignment via Self-Play and Reflective Experience Replay</i>	
Hao Wang, Yanting Wang, Hao Li, Rui Li and Lei Sha	18700
<i>RetentiveKV: State-Space Memory for Uncertainty-Aware Multimodal KV Cache Eviction</i>	
Sihao Liu, YuFan Xiong, Zhonghua Jiang, Zhaode Wang, Chengfei lv and Shengyu Zhang	18717
<i>RCBSF: A Multi-Agent Framework for Automated Contract Revision via Stackelberg Game</i>	
Shijia Xu, Yu Wang, Xiaolong Jia, Zhou Wu, Kai Liu and April Xiaowen Dong	18728

<i>RAVR: Reference-Answer-guided Variational Reasoning for Large Language Models</i> Tianqianjin Lin, Xi Zhao, Xingyao Zhang, Rujiao Long, Yi Xu, Zhuoren Jiang, Wenbo Su and Bo Zheng	18757
<i>Data-Centric Continual Pre-training for 500+ Languages: A New Bilingual Translation Corpus and Multilingual Models</i> Shaoxiong Ji, Zihao Li, Jaakko Paavola, Hengyu Luo and Jörg Tiedemann	18776
<i>Failures are Treasures: Constructing a Pedagogical Bridge for Agentic Strategy Distillation</i> Jiaxin Guo, Hao Sun, Wenhao Zhang, Chunyu Yang and Yan Zhang	18808
<i>Privacy-Preserving Reasoning with Knowledge-Distilled Parametric Retrieval Augmented Generation</i> Jinwen Chen, Hainan Zhang, Liang Pang, Yongxin Tong, Haibo Zhou, Wei Lin and Zhiming Zheng	18824
<i>Beyond Sentence-level Labels: Integrating Conversational Context and Personal Experience for Natural Emotional Expression</i> Haiyang Sun, Chenyang Le, Wei Wang, Leying Zhang, Chuang Li, Bing Han, Chenda Li, Mengxiao Bi and Yanmin Qian	18839
<i>Structure-Aware Zero-Shot Relational Learning for Knowledge Graphs without External Knowledge</i> Kuan Xu, Baoxin Zhang, Shuyue Fan, Ming Chen, Zhipeng Ke, Jian Yu and Xuezhong Zhou	18855
<i>Continual Safety Alignment via Gradient-Based Sample Selection</i> Thong Bach, Dung Nguyen, Thao Minh Le and Truyen Tran	18870
<i>ClinicalMC: A Benchmark for Multi-Course Clinical Decision-Making with Large Language Models</i> Ruihui Hou, Siyi Zhu, Ziyue Huai, Guangya Yu, Yongqi Fan, ChunMing Wang and Tong Ruan	18888
<i>Safe-SAIL: Towards a Fine-grained Safety Landscape of Large Language Models via Sparse Autoencoder Interpretation Framework</i> Jiaqi Weng, Han Zheng, Hanyu Zhang, Ej Zhou, Qinqin He, Jialing Tao, Hui Xue, Zhixuan Chu and Xiting Wang	18916
<i>Graph-Assisted Large Language Models: A Perspective on Mitigating Intrinsic Limitations</i> Haitong Luo, Fali Wang, Weiyao Zhang, Xianren Zhang, Zhiwei Zhang, Tianxiang Zhao, Minhua Lin, Jiahao Zhang, Hui Liu, Xianfeng Tang, Qi He, Suhang Wang, Xuying Meng and Yujun Zhang	18936
<i>GUITester: Enabling GUI Agents for Exploratory Defect Discovery</i> Yifei Gao, Jiang Wu, Xiaoyi Chen, Yifan Yang, Zhe Cui, Tianyi Ma, Jiaming Zhang and Jitao Sang	18956
<i>BiCSRrouter: Bi-Level Cross-System Routing for Utility-Aware LLM Inference</i> Mao Keyu, Eiki Murata and Ukyo Honda	18979
<i>FAER: Benchmarking VLMs for Failure-Aware Embodied Reasoning</i> Hao Song, Kaifeng Liu, Yuanxing Liu, Xiang Tian, Xuesong Wang, Chen Yifan, Weinan Zhang and Ting Liu	18994
<i>Reverse Constitutional AI: A Framework for Controllable Toxic Data Generation via Probability-Clamped RLAI</i> Yuan Fang, Yiming Luo, Aimin Zhou and Fei Tan	19017

<i>ToolGrad: Efficient Tool-use Dataset Generation with Textual “Gradients”</i>	
Zhongyi Zhou, Kohei Uehara, Haoyu Zhang, Jingtao Zhou, Lin Gu, Ruofei Du, Zheng Xu and Tatsuya Harada	19040
<i>Unlocking Exploration in RLVR: Uncertainty-aware Advantage Shaping for Deeper Reasoning</i>	
Can Xie, Ruotong Pan, Xiangyu Wu, Zhang Yunfei, Jiayi Fu, Tingting Gao and Guorui Zhou	19057
<i>GlossaGen: Making Academic Translation Smarter with Glossing</i>	
Zixiao Wang, Duzhen Zhang, Juntian Zhang, Yuhan Liu, Guoming Li, Haolun Wu, Le Song and Xiuying Chen	19077
<i>DeCoVec: Building Decoding Space based Task Vector for Large Language Models via In-Context Learning</i>	
Feiyang Li and Yile Wang	19096
<i>Hebbian-Guided Bi-Directional Rank Adaptation for Parameter-Efficient Fine-Tuning</i>	
Xing Zhao, Liu Yang, Xi-Le Zhao, Yueqi Yang and Tianqi Jiang	19112
<i>StruNRAG: Evaluation of OCR-Induced Structural Noise on RAG Robustness</i>	
Mengna Gao, Dapeng Yin, Shuyue Zhu, Bingxuan Hou, Zhanpeng Ni and Junli Wang	19129
<i>Memory as Action: Autonomous Context Curation for Long-Horizon Agentic Tasks</i>	
Yuxiang Zhang, Jiangming Shu, Ye Ma, Xueyuan Lin, Shangxi Wu and Jitao Sang	19149
<i>Experience is the Teacher: Reusing Atomic Thoughts from LLMs to Improve Medical Dialogue</i>	
Guangya Yu, Hui Luo, Qi Ye, Ruihui Hou, Weiyang Zhang, Mingxi Shang, Xuanwu Li, ChunMing Wang and Tong Ruan	19165
<i>Beyond Static Personas: Situational Personality Steering for Large Language Models</i>	
Zesheng Wei, Mengxiang Li, Zilei Wang and Yang Deng	19185
<i>Born Pragmatic, Trained to Hallucinate? Quantifying the Origins of Contextual Bias in LLMs via the PaCE Benchmark</i>	
Ziming Li, Yu Tian, Tian Lan, Jiang Li, Zehua Duo, Guanglai Gao and Xiangdong Su	19211
<i>Domain-Specific Data Generation Framework for RAG Adaptation</i>	
Chris Xing Tian, Weihao Xie, Zhen Chen, Hui Liu, Zhengyuan Yi, Haoliang Li, Shiqi Wang and Siwei Ma	19236
<i>Tracing Mathematical Proficiency Through Problem-Solving Processes</i>	
Jungyang Park, Suho Kang, Jaewoo Park, Jae Hong Kim, Jaewoo Shin, Seonjoon Park and Young-jae Yu	19251
<i>Decision Biases and Intent-Irony Decoupling in Large Language Models</i>	
Kewei Guo, Lingyun Sun and Manhao Guan	19270
<i>Anonpsy: A Graph-Based Framework for Structure-Preserving De-identification of Psychiatric Narratives</i>	
Kyungho Lim and Byung-Hoon Kim	19288
<i>HearSay Benchmark: Do Audio LLMs Leak What They Hear?</i>	
Jin Wang, Kaiwen Luo, Liang Lin, Weiliu Wang, Yitian Chen, Moayad Aloqaily, Xuehai Tang, Zhenhong Zhou, Kun Wang, Li Sun and Qingsong Wen	19312
<i>LongInsightBench: A Comprehensive Benchmark for Evaluating Omni-Modal Models on Human-Centric Long-Video Understanding.</i>	
ZhaoYang Han, Qihan Lin, Hao Liang, Bowen Chen, Zhou Liu and Wentao Zhang	19332

<i>Fixing Semantic Blind Spots in Anchor Tokens of dMLLMs</i>	
Ruixuan Xu, Jiexi Xu, Qiyan Zhao and Xiaofeng Zhang	19359
<i>Beyond Transcription: Unified Audio Schema for Perception-Aware AudioLLMs</i>	
Linhao Zhang, Yuhan Song, Aiwei Liu, Chuhan Wu, Sijun Zhang, Wei Jia, Yuan Liu, Houfeng Wang and Zhou Xiao	19375
<i>From Style to Story: A Curriculum Learning Approach for Imitative Novel Generation</i>	
Xueran Han, Yuhan Liu, Mingzhe Li, Wei Liu, Sen Hu, Rui Yan, Zhiqiang xu and Xiuying Chen	19391
<i>ArkRepoBench: A Repository-Level Code Completion Benchmark for HarmonyOS Development</i>	
Yanlin Wang, Bowen Zhang, Yanli Wang, Daya Guo, Terry Yue Zhuo, Jiachi Chen, Mingwei Liu, Xingong Zhang and Zibin Zheng	19409
<i>e5-omni: Explicit Cross-modal Alignment for Omni-modal Embeddings</i>	
Haonan Chen, Sicheng Gao, Radu Timofte, Tetsuya Sakai and Zhicheng Dou	19430
<i>Large Language Models Are Still Misled by Simple Bias Ensembles</i>	
Zhouhao Sun, Zhiyuan Kan, Xiao Ding, Li Du, Bibo Cai, Yang Zhao, Bing Qin and Ting Liu	19444
<i>Text-to-TrajVis: Enabling Trajectory Data Visualizations from Natural Language Questions</i>	
Tian Bai, Huiyan Ying, Kailong Suo, Victor Junqiu Wei, Tao Fan and Yuanfeng Song	19456
<i>MemTR: Enhancing Tool-Calling Reliability via Uncertainty-Triggered FFN-Space Retracing</i>	
Hongtao Duan, Lu Jiang, Mingyong Zhang, Xiaobing Zhu, Tianpeng Bu, Hao Jiang, Xinyu Wei and Lulu hu	19476
<i>More Than Sum of Its Parts: Deciphering Intent Shifts in Multimodal Hate Speech Detection</i>	
Runze Sun, Yu Zheng, Zexuan Xiong, Zhongjin Qu, Lei Chen, Jie Zhou and Jiwen Lu	19494
<i>Specialization without Sparsity: Efficient and Expressive Split-Path Experts for LLM Fine-Tuning</i>	
Ganghao Liu, Qin Zhou, Zhe Wang, Xuehan Lu, Haihua Huang, Yunfei Tong and Heng Tian	19515
<i>Pardon? Evaluating Conversational Repair in Large Audio-Language Models</i>	
Shuanghong Huang, Jinlei Xu, Youchao Zhou, Yanghao Zhou, Xuan Zhao, Chong Feng and Wenxuan Zhang	19528
<i>ToxReason: A Benchmark for Mechanistic Chemical Toxicity Reasoning via Adverse Outcome Pathway</i>	
Jueon Park, WonJune Jang, Chanhwi Kim, Yein Park and Jaewoo Kang	19542
<i>CVRH: Cross-modal Variational Role Hypergraph Network via Semantic Enhancement for Multi-modal Event Argument Extraction</i>	
Bangze Pan, Yang Li, Ruili Pu, Suge Wang, Jian Liao, JianXing Zheng, Xiaoli Li and Deyu Li	19565
<i>DiffCL: Difference-Aware Contrastive Learning for Automatic Answer Grading with Multi-Level Semantic Modeling</i>	
Lei Chen, BoYu Gao, Zitao Liu, Tingjie Wan and Weiqi Luo	19576
<i>GaLa: Hypergraph-Guided Visual Language Models for Procedural Planning</i>	
Kun Wang, Li Yiming, Mingcheng Qu, Aqiang Zhang, Guang Yang and Tonghua Su	19590
<i>TwinVoice: A Multi-dimensional Benchmark Towards Digital Twins via LLM Persona Simulation</i>	
Bangde Du, Minghao Guo, Songming He, Ziyi Ye, Xi Zhu, Weihang Su, Shuqi Zhu, Yujia Zhou, Yongfeng Zhang, Qingyao Ai and Yiqun Liu	19604

<i>CE-RM: A Pointwise Generative Reward Model Optimized via Two-Stage Rollout and Unified Criteria</i> Xinyu Hu, Yancheng He, Weixun Wang, Tao Feng, Li Lin, Jiashun Liu, Wenbo Su, Bo Zheng and Xiaojun Wan	19629
<i>Hard to Read, Easy to Jailbreak: How Visual Degradation Bypasses MLLM Safety Alignment</i> Zhixue Song, Boyan Han, Yiwei Wang and Chi Zhang	19643
<i>Unleashing the Native Recommendation Potential: LLM-Based Generative Recommendation via Structured Term Identifiers</i> Zhiyang Zhang, Junda She, Kuo Cai, Bo Chen, Shiyao Wang, Xinchen Luo, Qiang Luo, Ruiming Tang, Han Li, Kun Gai and Guorui Zhou	19659
<i>Abstain-RI: Calibrated Abstention and Post-Refusal Clarification via Verifiable RL</i> Haotian Zhai, Jingcheng Liang and Dongyeop Kang	19674
<i>Learning Through Dialogue: Engagement and Efficacy Matter More Than Explanations</i> Shaz Furniturewala, Gerard Christopher Yeo and Kokil Jaidka	19696
<i>CodeBind: Decoupled Representation Learning for Multimodal Alignment with Unified Compositional Codebook</i> Zeyu Chen, Jie Li and Kai Han	19712
<i>MASS: Deep Research for Social Sciences with Memory-Augmented Social Simulation</i> Yongrui Liu and Deyi Xiong	19739
<i>LS-Guard: Adaptive Safety Guardrails Tailored to Individual LLMs</i> Jinggui Liang and Lizi Liao	19759
<i>Rethinking LLM Watermark Detection in Black-Box Settings: A Non-Intrusive Third-Party Framework</i> Zhuoshang Wang, Yubing Ren, Yanan Cao, Fang Fang, Xiaoxue Li and Li Guo	19773
<i>Deputy: Accelerating Large Language Model Inference with Dynamic Low-Rank Substitution</i> Yuhua Zhou, Shichao Weng, Changhai Zhou, Yuhan Wu, Qian Qiao, Jun Gao, Fei Yang and Aimin Pan	19791
<i>SGG-R³: From Next-Token Prediction to End-to-End Unbiased Scene Graph Generation</i> Jiaye Feng, Qixiang Yin, Yuankun Liu, Tong Mo and Weiping Li	19811
<i>Preserving Fairness and Safety in Quantized LLMs Through Critical Weight Protection</i> Muhammad Alif Al Hakim, Alfian Farizki Wicaksono and Fajri Koto	19831
<i>X-Router: Decoupling Knowledge and Reasoning for Cost-Effective LLM Inference</i> Zixuan Wang, Yinze Ding, Zihan Wang, Jinyu Guo, Zhenhong Zhou, Junhao Dong and Chaomeng Chen	19856
<i>SPD-Faith Bench: Diagnosing and Improving Faithfulness in Chain-of-Thought for Multimodal Large Language Models</i> Weijiang Lv, Yaoxuan Feng, Xiaobo Xia, Jiayu Wang, Yan Jing, Wenchao Chen and Bo Chen	19875
<i>HIPO: A Hierarchical Prompt Optimization Framework with Task Awareness and Fine-Grained Debugging</i> Lu Qi, Lei Chai, Hongrui Yu, Binhang Qi and Hailong Sun	19928
<i>MAPLE: A Meta-learning Framework for Cross-Prompt Essay Scoring</i> Salam Albatarni, May Bashendy, Sohaila Eltanbouly and Tamer Elsayed	19948

<i>RACER: Retrieval-Augmented Contextual Rapid Speculative Decoding</i>	
Zihong Zhang, Zuchao Li, Lefei Zhang, Ping Wang and Hai Zhao	19962
<i>PEC-Home: Interpretation of Progressively Elliptical Commands in Smart Homes</i>	
Yingyu Shan, Zeming Liu, Silin Li, Boao Qian, Jiashu Yao, Yuhang Guo and Haifeng Wang	19989
<i>Towards IP Intelligence: Benchmarking Large Language Models on Intellectual Property Knowledge and Practice</i>	
Qiyao Wang, Guhong Chen, Hongbo Wang, Huaren Liu, Minghui Zhu, Zhifei Qin, Li Linwei, Yilin Yue, Shiqiang Wang, Jiayan Li, Wu Yihang, Ziqiang Liu, Longze Chen, Run Luo, Liyang Fan, Jiaming Li, Lei Zhang, Kan Xu, Hamid Alinejad-Rokny, Chengming Li, Shiwen Ni, Yuan Lin and Min Yang	20007
<i>EchoMLLM: Incentivizing Echocardiographic Video Understanding with Keyframe Grounding and Report Generation</i>	
Heyu Huang, Wanran Sun, Chi Chen, Bo Chen, Zonghao Guo, Yuhua Li, Ruixuan Li, Kunlun He and Maosong Sun	20053
<i>HAUNTATTACK: When Attack Follows Reasoning as a Shadow</i>	
Jingyuan Ma, Rui Li, Zheng Li, Junfeng Liu, Heming Xia, Lei Sha and Zhifang Sui	20072
<i>Rethinking Multiple-Choice Questions for RLVR: Unlocking Potential via Distractor Design</i>	
Xu Guo, Qiming Ge, Jian Tong, Kedi Chen, Jin Zhang, Xiaogui Yang, Xuan Gao, Haijun Lv, Zhihui Lu, Yicheng Zou and Qipeng Guo	20092
<i>S2ST-Omni: Hierarchical Language-Aware SpeechLLM Adaptation for Multilingual Speech-to-Speech Translation</i>	
Yu Pan, Xiongfei Wu, Yang Yuguang, Jixun Yao, Maxime Cordy, Lei Ma and Jianjun Zhao	20114
<i>G-LoRA: Global-Local Decoupled Low-Rank Adaptation</i>	
Jiahao Xiong, Yihong Huang, Yihe Liu, Xianming Hu, Hongbo Zhao and Kai Zhang	20125
<i>Making MLLMs Blind: Adversarial Smuggling Attacks in MLLM Content Moderation</i>	
Zhiheng Li, Zongyang Ma, Yuntong Pan, Ziqi Zhang, Xiaolei Lv, Bo Li, Jun Gao, Jianing Zhang, Chunfeng Yuan, Bing Li and Weiming Hu	20142
<i>LatentRefusal: Latent-Signal Refusal for Unanswerable Text-to-SQL Queries</i>	
Xuan Cheng Ren, Shijing Hu, Zhihui Lu, Jiangqi Huang and Qiang Duan	20162
<i>Where meaning lives: Layer-wise accessibility of psycholinguistic features in encoder and decoder language models</i>	
Taisiia Tikhomirova and Dirk U. Wulff	20180
<i>DynamicFocalPO: Adaptive Focusing Strategy for Preference Optimization</i>	
Shu Zhou, Junan Chen, Rui Ling, Xin Wang, Tao Fan and Hao Wang	20206
<i>Graph-GRPO: Stabilizing Multi-Agent Topology Learning via Group Relative Policy Optimization</i>	
Yueyang Cang, Xiaoteng Zhang, Erlu Zhao, Zehua Ji, Yuhang Liu, Yuchen He, Zhiyuan Ning, Chen Yijun, Wenge Que and Li Shi	20222
<i>ReMedi: Reasoner for Medical Clinical Prediction</i>	
Yushi Cao, Yiming Chen, Hongchao Jiang, Hung-yi Lee and Robby T. Tan	20232
<i>JX4MEI: Multimodal Semantically-Enhanced LLM for Joint Multimodal Emotion-Intent Explanation and Classification</i>	
YiJie Huang, Xiaocui Yang, Shi Feng, Daling Wang, Yifei Zhang, Ning Yuan, Zhuoyue Jia and Wen Zhang	20242

<i>Breaking the Provable Security": Detecting Finite-Precision Artifacts in LLM-based Steganography via Low-Probability Vanishing</i>	20262
Wenzhao Cao, Yaofei Wang and Donghui Hu	
<i>FACT-E: Causality-Inspired Evaluation for Trustworthy Chain-of-Thought Reasoning</i>	20275
Yuxi Sun, Aoqi Zuo, Haotian Xie, Wei Gao, Mingming Gong and Jing Ma	
<i>DRIFT: Transferring Reasoning Priors for Efficient MLLM Fine-Tuning</i>	20294
Chao Huang, Zeliang Zhang, Jiang Liu, Ximeng Sun, Jialian Wu, Xiaodong Yu, Ze Wang, Chenliang Xu, Emad Barsoum and Zicheng Liu	
<i>Mitigating Hallucinations in VLMs: Enhancing Visual Attention via Head-Wise Perturbation</i>	20310
Zhenghua Wang, Yixin Wu, Feiran Zhang, Qi Qian, Changze Lv, Xuanjing Huang and Xiaoqing Zheng	
<i>Modeling Multi-Dimensional Cognitive States in Large Language Models under Cognitive Crowding</i>	20322
Lin Zhong, Siyu Zhu, Zizhen Yuan, Jinhao Cui, Xinyang Zhao, Lingzhi Wang, Hao Chen and Qing Liao	
<i>EmoMM: Benchmarking and Steering MLLM for Multimodal Emotion Recognition under Conflict and Missingness</i>	20351
Yueru Sun, Yimeng Zhang, Haoyu Gu, Nuo Chen, Dong She, Xianrong Yao, Yang Gao and Zhanpeng Jin	
<i>Veri-RI: Toward Precise and Faithful Claim Verification via Online Reinforcement Learning</i>	20372
Qi He, Cheng Qian, Xiusi Chen, Bingxiang He, Yi R. Fung and Heng Ji	
<i>MimicLM: Zero-Shot Voice Imitation through Autoregressive Modeling of Pseudo-Parallel Speech Corpora</i>	20392
Tao Feng, Yuxiang Wang, Yuancheng Wang, Xueyao Zhang, Dekun Chen, Chaoren Wang, Xun Guan and Zhizheng Wu	
<i>Fine-Grained Data Ordering Improves Fine-Tuning for Large Language Models</i>	20406
Xiaomeng Hu, Yixuan Tang, Haoze Li, Hao Chen, Qi Zhang, Zhanming Shen, Yiming Zhang, Haobo Wang and Junbo Zhao	
<i>LifeSim: Long-Horizon User Life Simulator for Personalized Assistant Evaluation</i>	20419
Feiyu Duan, Xuanjing Huang and Zhongyu Wei	
<i>Interactive Semantic Parsing with Reinforcement Learning for Knowledge Graph Reasoning</i>	20464
Yurun Song, Xiangqing Shen, Jianfei Yu and Rui Xia	
<i>CLewR: Curriculum Learning with Restarts for Machine Translation Preference Learning</i>	20485
Alexandra Dragomir, Florin Brad and Radu Tudor Ionescu	
<i>Targeting the Needle, Ignoring the Haystack: Anchoring Crucial Cues for Evolving Scam Call Detection via an LLM-Assisted Classifier</i>	20497
Tong Wu, Qinliang Su, Jianxing Yu, Bo Liang and Minhua Huang	
<i>Distributional Alignment for Large Language Models under Domain Shift</i>	20514
Viet Thanh Pham, Lizhen Qu, Zhuang Li and Gholamreza Haffari	
<i>Traces in the Brain: Neural Evidence for Syntactic Movement in English and Chinese</i>	20529
Yuhan Huang, Zhengwu Ma, Yuqi Jin, Beth Chan, Zheng Shen, Jackie Yan-Ki Lai, John T. Hale and Jixing Li	

<i>LinguaGame: A Linguistically Grounded Game-Theoretic Paradigm for Multi-Agent Dialogue Generation</i>	
Yuxiao Ye, Yiming Zhang, Yiran Rex Ma, Huiyuan Xie, Huining Zhu and Zhiyuan Liu . . .	20544
<i>A Dual-View Analysis of Multiple Languages in Colonial Newspapers</i>	
Zhan Su, Xiaoya Chen, Fengran Mo, Ida L. Vos, Prayag Tiwari, Yazhou Zhang, Qian Zheng and Natália da Silva Perez	20559
<i>CARE-CR: Context-Aware Routing and Expert Fusion for Multi-Preference Cognitive Restructuring</i>	
Hongzhi Qi, Liangcheng Wang, Yijing YU, Jianqiang Li, Bing Xiang Yang and Qing Zhao	20574
<i>EvoCoT: Overcoming the Exploration Bottleneck in Reinforcement Learning for LLMs</i>	
Huanyu Liu, Jia Li, Yihong Dong, Chang Yu, Taozhi Chen, Lecheng Wang, Yongding Tao, Bin Gu and Ge Li	20589
<i>Context as a Tool: Context Management for Long-Horizon SWE-Agents</i>	
Shukai Liu, Bo Jiang, Jian Yang, Yizhi LI, Jinyang Guo, Xianglong Liu and Bryan Dai . . .	20604
<i>Beyond Static Profiles: Capturing the Fluidity of User Preferences in Diverse Scenarios</i>	
Chunyang Gao, Yi Huang, Jingyu Yao, Xiaoting Wu and Junlan Feng	20618
<i>Mitigating Coordinate Prediction Bias from Positional Encoding Failures</i>	
Xingjian Tao, Yiwei Wang, Yujun Cai, Yihong Luo, Kai Han and Jing Tang	20635
<i>PerDucer: Keyphrase-Driven Personalization Inducer for Summarization from User Histories</i>	
Parthiv Chatterjee, Asish Joel Batha, Sourish Dasgupta and Tanmoy Chakraborty	20651
<i>Beyond A Fixed Seal: Adaptive Stealing Watermark in Large Language Models</i>	
Shuhao Zhang, Yuli Chen, Jiale Han, Bo Cheng and Jiabao Ma	20678
<i>UniGeM: Unifying Data Selection and Mixing via Geometric Exploration and Mining</i>	
Changhao Wang, Yunfeiyu, Xinhao Yao, Jiaolong Yang, Lu Yu, Junpeng Fang, Chaobo Li, Riccardo Cantoro, Qing Cui and Jun Zhou	20696
<i>Don't Just Listen, Try Planning: Graph-based Retrieval-Generation Agent for Long-form Audio Meeting Understanding</i>	
Quanwei Tang, Dong Zhang, Shoushan Li and Guodong Zhou	20715
<i>FLASH: Focused Layer Attention Sink Hijacking</i>	
Siyuan Deng, Yerong Li, Fanhua Shang and Hongying Liu	20743
<i>SPIDE: Serial and Parallel Intertwined Speculative Decoding</i>	
Wenru Xu, Peixuan Xu, Ziqi Yang, Ming Hu, Zihui Wang, Jianzhong Qi, Rongshan Yu, Xiaoliang Fan and Cheng Wang	20762
<i>MdEval: Massively Multilingual Code Debugging</i>	
Shukai Liu, Linzheng Chai, Jian Yang, Jiajun Shi, He Zhu, Liran Wang, Jin Ke, Wei Zhang, Hualei Zhu, Shuyue Guo, Tao Sun, Jiaheng Liu, Yunlong Duan, Yu Hao, Liqun Yang, Guanglin Niu, Ge Zhang and Zhoujun Li	20780
<i>Breaking Consensus Bias: Unsupervised Reinforcement Learning for Machine Translation</i>	
Shuting Jiang, Ran Song, Siqi Zhang, Yuxin Huang, Shengxiang Gao and Zhengtao Yu . . .	20798
<i>Reasoning-Aware AIGC Detection via Alignment and Reinforcement</i>	
Zhao Wang, Max Xiong, Jianxun Lian and Zhicheng Dou	20813

<i>Beyond Modality Collapse: Taming Guided Modality Entropy for Omni-modal Emotion Reasoning</i> Xian Zhao, Rui Hu, Yuxiang Zhang, Delai Qiu, Yining Wang, Shengping Liu, Jian Yu and Jitao Sang	20830
<i>ID10M-JAM: Stress-Testing Idiom Identification Under Challenging Context</i> Kai Golan Hashiloni, Lior Livyatan, Ofri Hefetz, Alon Mannor, Bar Cohen and Kfir Bar ..	20846
<i>MENTOR: Mitigating Identity Drift in Dynamic Role-Playing via Dual-Chain Structured Memory</i> Zhu Zhuoning, Xingyu Gao and Hailong Shi	20865
<i>Beyond Benchmarks: A Capability-Based Maturity Model for Systematic AI Integration in Hospitals</i> Rui Li, Xiaofen WU, Mingqian Liu, Xiaoxia Song, Xu Xiangjun, Jiacheng Qiao, Zhuang Boqin and Xu Chen	20883
<i>DiffuSpec: Unlocking Diffusion Language Models for Speculative Decoding</i> Guanghao Li, Zhihui Fu, Min Fang, Qibin Zhao, Ming Tang, Chun Yuan and Jun Wang ...	20896
<i>TENP: Trapezoidal Expert Neuron Pruning For Mixture-of-Experts</i> Jiangyang He, Shaolin Zhu and Deyi Xiong	20911
<i>FlowRAG: Synergizing Explicit Reasoning via Frequency-Aware Multi-Granularity Graph Flow</i> Bihao Zhan, Zongsheng Cao, Jie Zhou, Bo Zhang and Liang He	20926
<i>OntoGuard: Enforcing Action Admissibility for LLM Agents in Complex Interactive Environments</i> Pengxiang Liu, Tao Ren, Wei Xiong, Tingrui Yang, Junjie Wang and Jun HU	20937
<i>Self-Correcting RAG: Enhancing Faithfulness via MMKP Context Selection and NLI-Guided MCTS</i> Shijia Xu, Zhou Wu, Xiaolong Jia, Yu Wang, Kai Liu and April Xiaowen Dong	20953
<i>DiffER: Diffusion Entity-Relation Modeling for Reversal Curse in Diffusion Large Language Models</i> Shaokai He, Kaiwen Wei, Xinyi Zeng, Xiang Chen, Xue Yang, Zhenyang Li, Jiang Zhong and Yu Tian	20977
<i>Layer-aware Dual-directional Modulation for Low-resource Machine Translation</i> Siqi Zhang, Ran Song, Shuting Jiang, Yuxin Huang and Zhengtao Yu	20988
<i>MATCHA: Matching Text via Contrastive Semantic Alignment</i> Siran Li, Ece Sena Etoglu, Carsten Eickhoff and Seyed Ali Bahrainian	21001
<i>PlaM: Training-Free Plateau-Guided Model Merging for Better Visual Grounding in MLLMs</i> Zijing Wang, YongKang Liu, Mingyang Wang, Ercong Nie, Deyuan Chen, Zhengjie Zhao, Shi Feng, Daling Wang, Xiaocui Yang, Yifei Zhang and Hinrich Schuetze	21019
<i>Efficient Inference for Large Vision-Language Models: Bottlenecks, Techniques, and Prospects</i> Jun Zhang, Yicheng Ji, Feiyang Ren, Yihang Li, Bowen Zeng, Zonghao Chen, Ke Chen, Lidan Shou, Gang Chen and Huan Li	21036
<i>Adaptive Zooming via Relevance-Informed Positional Resource Allocation for Training-free LLM Context Extension</i> Hongbo Zhao, Huibin Wang, Bin Tang, Xianming Hu, Yihong Huang, Yijun Shen, Nuoyi Chen, Ping Li and Kai Zhang	21067
<i>Mitigating Position Bias in Transformers via Layer-Specific Positional Embedding Scaling</i> Zhenghua Wang, Yiran Ding, Changze Lv, Yixin Wu, Tianlong Li, Zhibo Xu, Muling Wu, Tianyuan Shi, Shizheng Li, Qi Qian, Xuanjing Huang and Xiaoqing Zheng	21084

<i>Semantic-pragmatic Annotations in the Prague Dependency Treebank</i>	
Marie Mikulová, Eva Hajicova, Jiří Mírovský, Anna Nedoluzhko, Michal Novák, Pavlína Synková, Jan Štěpánek, Barbora Štěpánková and Jan Hajič	21099
<i>Penetrating Linguistic Disguises: A Slang-aware Label-Aligned Framework for Fine-Grained Toxicity Extraction in Chinese Hate Speech Detection</i>	
Wei Liu, Xiaoliang Chen, Duoqian Miao, Xu Gu, Xianyong Li and Yajun Du	21111
<i>LightMoE: Task-Aware Expert Availability Management for Memory-Efficient MoE-LLM Inference</i>	
Puhan Luo, Yunhao Yao, Junyang Wang, Junyang Zhang and Xiangyang Li	21124
<i>SelfBudgeter: Adaptive Token Allocation for Efficient LLM Reasoning</i>	
Zheng Li, Qingxiu Dong, Jingyuan Ma, Di Zhang, Kai Jia and Zhifang Sui	21135
<i>Streaming Hallucination Detection in Long Chain-of-Thought Reasoning</i>	
Haolang Lu, Minghui Pan, Ripeng LI, Guoshun Nan, Jialin Zhuang, Zijie Zhao, Zhongxiang Sun, Kun Wang and Yang Liu	21157
<i>Training-Free Adaptive Speculative Decoding via Linguistic Priors</i>	
Jingyi Wang, Jiaqi Huang, Zunnan Xu, Jun Zhou, Kehong Yuan and Xiang Qian	21184
<i>Knowledge-Infused Multi-Bit Watermarking for RAG Knowledge Bases</i>	
Ke Yang, Shuguang Yuan, Jing Yu and Chi Chen	21195
<i>ClimAgent: LLM as Agents for Autonomous Open-ended Climate Science Analysis</i>	
Hao Wang, Jindong Han, Wei Fan and Hao Liu	21219
<i>CogBench: Benchmarking Cognitive Alignment of Large Language Models in Educational Question Answering</i>	
Tong Lu, Zhichun Wang, Yuanhao Sun, Yaoyu Zhou, Mingrui Li, Yiming Guan and Zhiyong Bai	21242
<i>CARL: Constraint-Aware Reinforcement Learning for Planning with LLMs</i>	
Qiuyi Qi, Jinjian Zhang, Mutian Bao, Tian Liang, Guocong Li, Dongnan Liu, Wei Zhou, Jie Liu, Ming Kong, Linjian Mo, Feng Zhang and Qiang Zhu	21257
<i>ChildEval: WHEN LARGE LANGUAGE MODELS MEET CHILDREN'S PERSONALITIES</i>	
Yanyan Luo, Xue Han, Chunxu Zhao, Ruiqiao Bai, Yaxing Zhang, Qian Hu, Lijun Mei and Junlan Feng	21282
<i>Belief in Authority: Impact of Authority in Multi-Agent Evaluation Framework</i>	
Junhyuk Choi, Jeongyoun Kwon, Heeju Kim, Haeun Cho, Hayeong Jung, Sehee Min and Bugeun Kim	21305
<i>ToMELP: A Theory-of-Mind Benchmark for Route-Controlled Persuasion under the Elaboration Likelihood Model</i>	
Ruirui Wang, Haoran Zhang, Tian Lan, Zehua Duo, Jiang Li, Guanglai Gao and Xiangdong Su	21320
<i>StateX: Enhancing RNN Recall via Post-training State Expansion</i>	
Xingyu Shen, Yingfa Chen, Zhen Leng Thai, Xu Han, Zhiyuan Liu and Maosong Sun	21339
<i>TOXIFRENCH: Benchmarking and Enhancing Language Models via CoT Fine-Tuning for French Toxicity Detection</i>	
Axel Delaval, Shujian Yang, Haicheng Wang, Han Qiu and Jialiang LU	21354

<i>Memory-Guided Hard Data Augmentation for Multimodal Named Entity Recognition</i>	
Xinyu Liu, Kai fu, Yinghan Shi, Quanyou Chu, Ming Du, Hongya Wang, Xiaojun Meng, Jiansheng Wei, Yanghua Xiao and Bo Xu	21376
<i>DUET: Joint Exploration of User–Item Profiles in Recommendation System</i>	
Yue Chen, Yifei Sun, Lu Wang, Fangkai Yang, Pu Zhao, Minjie Hong, Yifei Dong, Minghua He, Nan Hu, Jianjin Zhang, Zhiwei Dai, Yuefeng Zhan, Weihao Han, Hao Sun, Qingwei Lin, Weiwei Deng, Feng Sun, Qi Zhang, Saravan Rajmohan and Dongmei Zhang	21392
<i>ATAAT: Adaptive Threat-Aware Adversarial Tuning Framework against Backdoor Attacks on Vision-Language-Action Models</i>	
Kewei Chen, Yayu Long, Shuai Li and Mingsheng Shang	21407
<i>CoTrust: Privacy-Preserving Collaboration Between Large and Small Language Models in Trusted Execution Environments</i>	
Zhenya Ma, Tingyi Wang, Yongheng Deng, Ziqing Qiao, Yinggui Wang, Tao Wei, Lei Wang and Ju Ren	21423
<i>FAQ: Mitigating Quantization Error via Regenerating Calibration Data with Family-Aware Quantization</i>	
Haiyang Xiao, Weiqing Li, Jinyue Guo, Guochao Jiang, Guohua Liu and Yuewei Zhang ..	21441
<i>HACHIMI: Scalable and Controllable Student Persona Generation via Orchestrated Agents</i>	
Yilin Jiang, Fei Tan, Xuanyu Yin, Leng Jing and Aimin Zhou	21461
<i>D-QRELO: Training- and Data-Free Delta Compression for Large Language Models via Quantization and Residual Low-Rank Approximation</i>	
Junlin Li, Shuangyong Song, Guodong DU, Ngai Wong, Xuebo Liu, Yongxiang Li, Min Zhang, Jing Li and Xuelong Li	21507
<i>Beyond Static Toolsets: Self-Evolving LLM Tool Agents via Continual Documentation Adaptation</i>	
Bin Wu, Edgar Meij and Emine Yilmaz	21519
<i>ZoFia: Zero-Shot Fake News Detection with Entity-Guided Retrieval and Multi-LLM Interaction</i>	
Lvhua Wu, Xuefeng Jiang, Sheng Sun, Yan Lei, Tian Wen, Yuwei Wang and Min Liu	21540
<i>SenseJudge: Human-Centric Preference-Driven Judgment Framework</i>	
Rui Li, Junfeng Liu, Xiangwen Kong and Zhifang Sui	21557
<i>Reliability-Aware Adaptive Self-Consistency for Efficient Sampling in LLM Reasoning</i>	
Junseok Kim, Nakyeong Yang, Kyungmin Min and Kyomin Jung	21575
<i>CLaS-Bench: A Cross-Lingual Alignment and Steering Benchmark</i>	
Daniil Gurgurov, Yusser Al Ghussin, Tanja Baeumel, Cheng-Ting Chou, Patrick Schramowski, Marius Mosbach, Josef Van Genabith and Simon Ostermann	21591
<i>From Prediction to Intervention: Personalized Meal-Level Glucose Regulation via an LLM Agent</i>	
Mingyu Huang, Weiqing Min, Ying Jin, Yilin Wang and Shuqiang Jiang	21629
<i>PsyPath: Psychologically-guided Self-Exploration for Personality Detection</i>	
Zheng Li, Hongxin Ding, Chenyu Zhang, Weimin Xiong, Dawei Zhu and Sujian Li	21646
<i>CCTVBench: Contrastive Consistency Traffic VideoQA Benchmark for Multimodal LLMs</i>	
Xingcheng Zhou, Hao Guo, Rui Song, Walter Zimmer, Mingyu Liu, André Schamschurko, Hu Cao and Alois Knoll	21665

<i>PROGRAM: Programmatic Retrieval Optimization with Generative Reasoning and Augmented Multi-queries</i>	
Gun Il Kim, Jungkyu Shin, Jong Wook Kim and Beakcheol Jang	21685
<i>TiMem: Temporal-Hierarchical Memory Consolidation for Long-Horizon Conversational Agents</i>	
Kai Li, Xuanqing Yu, Ziyi Ni, Yi Zeng, Yao Xu, Zheqing Zhang, Xin Li, Jitao Sang, Xiaogang Duan, Xuele Wang, Chengbao Liu and Jie Tan	21700
<i>GraphDx: A Cost-Aware Knowledge-Enhanced Multi-Agent Framework for Sequential Diagnosis</i>	
Shaoting Tan, Ning Liu, Yuntao Du, Shuyue Wei, Wu Shuai, Qian Li, Yanyu Xu, Wei Zhang, Lizhen Cui and Haitao Yuan	21721
<i>ProUIE: A Macro-to-Micro Progressive Learning Method for LLM-based Universal Information Extraction</i>	
Wenda Liu, Song Zhigang, Shuai Nie, Guangyao Liu, Lisung Chen, Binyu Yang, Yaran Chen, Peng Zhou, Hongzhen Wang, Yuchen Liu, Wenyue Hu, Jiaming Xu, Runyu Shi and Ying Huang	21737
<i>MAVIS: Multi-Agent Video Retrieval via Structured Video Understanding</i>	
Jie Zhang, Qilang Ye, Hao Zhou, Haochen Liang and Fei Luo	21751
<i>ThinkBrake: Efficient Reasoning via Log-Probability Margin Guided Decoding</i>	
Sangjun Song, Minjae Oh, Seungkyu Lee, Sungmin Jo and Yohan Jo	21765
<i>MSIA: Adaptive Medical Multimodal Multi-turn Semantic Jailbreak</i>	
Zhiheng Han, Yao Zhang, Jun Wang and Zhenglu Yang	21791
<i>Why These Documents? Explainable Generative Retrieval with Hierarchical Category Paths</i>	
Sangam Lee, Ryang Heo, SeongKu Kang, Susik Yoon, Jinyoung Yeo and Dongha Lee	21807
<i>ETHICA-MT: Introducing a Framework and Dataset for Studying Ethical Orientations in LLM-based Machine Translation</i>	
Omri Asscher, Arif Ahmad, Ananya Agrawal and Monojit Choudhury	21825
<i>ArgBench: Benchmarking LLMs on Computational Argumentation Tasks</i>	
Yamen Ajjour, Carlotta Quensel, Nedim Lipka and Henning Wachsmuth	21846
<i>Generics are not quantificational: A new path from language models to semantic theory</i>	
Gustavo Cilleruelo Calderón, Mahrad Almotahari, Emily Allaway, Barry Haddow and Alexandra Birch	21875
<i>Beyond Neural Incompatibility: Cross-Scale Knowledge Transfer in Large Language Models through Latent Semantic Alignment</i>	
Jian Gu, Aldeida Aleti, Chunyang Chen and Hongyu Zhang	21893
<i>CURaTE: Continual Unlearning in Real Time with Ensured Preservation of LLM Knowledge</i>	
Seyun Bae, Seokhan Lee and Eunho Yang	21906
<i>Decoupling the Effect of Chain-of-Thought Reasoning: A Human Label Variation Perspective</i>	
Beiduo Chen, Tiancheng Hu, Caiqi Zhang, Robert Litschko, Anna Korhonen and Barbara Plank	21931
<i>CTRL: Control-Based Time Series Forecasting with LLM-Guided Residual Learning</i>	
Minkyung Kim, Daeun Ji, Yohan Lee, Beomsoo Kim and Beakcheol Jang	21952
<i>INDOTABVQA: A Benchmark for Cross-Lingual Table Understanding in Bahasa Indonesia Documents</i>	
Somraj Gautam, Anathapindika Dravichi and Gaurav Harit	21969

<i>Cross-Cultural Transfer of Emoji Semantics and Sentiment in Financial Social Media</i> Ahmed Mahrous and Roberto Di Pietro	21982
<i>EXPO-SQL: Execution-based Clause-level Policy Optimization for Text-to-SQL</i> Jaehoon Lee, CheolWon Na, Suyoung Bae, Jin-Seop Lee, Jihyung Lee, YunSeok Choi and Jee-Hyong Lee	22000
<i>Synapse: Empowering LLM Agents with Episodic-Semantic Memory via Spreading Activation</i> Hanqi Jiang, Junhao Chen, Yi Pan, Ling Chen, Weihang You, Yifan Zhou, Ruidong Zhang, Yohannes Abate and Tianming Liu	22020
<i>Similarity-Distance-Magnitude Activations</i> Allen Schmaltz	22037
<i>SDAR: A Synergistic Diffusion-AutoRegression Paradigm for Scalable Sequence Generation</i> Shuang Cheng, Yihan Bian, Dawei Liu, Yuhua Jiang, Yihao Liu, Linfeng Zhang, Qian Yao, Zhongbo Tian, Wenhai Wang, Qipeng Guo, Kai Chen, Biqing Qi and Bowen Zhou	22058
<i>Comparing Human and Large Language Model Interpretation of Implicit Information</i> Antonio De Santis, Tommaso Bonetti, Andrea Tocchetti and Marco Brambilla	22076
<i>Simplify-Pro: A Two-level and Progressive LLM-based Framework for Auto Long Text Simplification</i> Peng Zhou, Guangxin Li, Xiaoying Huang and Yiming Tang	22096
<i>Pattern Enhanced Multi-Turn Jailbreaking: Exploiting Structural Vulnerabilities in Large Language Models</i> Ragib Amin Nihal, Rui Wen, Kazuhiro Nakadai and Jun Sakuma	22123
<i>Data-Efficient Adaptation to Contextual Shifts in LLM-based Conversational Recommendation</i> Hyeongjun Yang, Donghyun Kim, Seokju Hwang, Midan Shim, KyuHwan Yeom, KaeHyun Um and Kyong-Ho Lee	22175
<i>PsychEval: A Multi-Session and Multi-Therapy Benchmark for High-Realism AI Psychological Counselor</i> Qianjun Pan, Junyi Wang, Jie Zhou, Yutao Yang, Junsong Li, Kaiyin Xu, Yougen Zhou, Yihan Li, JingYuan Zhao, Qin Chen, Ningning Zhou, Kai Chen and Liang He	22192
<i>GrandGuard: Taxonomy, Benchmark, and Safeguards for Elderly-Chatbot Interaction Safety</i> Changxuan Fan, Xi Yang, Yueyuan Zheng, Bin Zhou, Yuanping Wang, Wenbin Hu, Huihao Jing, Ki Sen Hung, Dazhao Du, Haoran Li, Janet Hui-wen Hsiao and Yangqiu Song	22213
<i>SciExplore: Evaluating Autonomous Agents from Scientific Navigation to Information Integration</i> Yinhao Tang, Youqing Fang, Yanan Sun, Wenran Liu, Weiming Zhang, Bin Liu, Kuikun Liu, Wenwei Zhang and Kai Chen	22249
<i>When Models Outthink Their Safety: Unveiling and Mitigating Self-Jailbreak in Large Reasoning Models</i> Yingzhi Mao, Chunkang Zhang, Junxiang Wang, Xinyan Guan, Boxi Cao, Yaojie Lu, Hongyu Lin, Xianpei Han and Le Sun	22274
<i>NASH: Numerically Aware Scoring Heuristic for Robust Semantic Similarity</i> Yu-Shiang Huang, Yun-Yu Lee, Tzu-Hsin Chou, Che Lin and Chuan-Ju Wang	22303
<i>Beyond Query Bias: Candidate-Aware Iterative Refinement for Zero-Shot Composed Image Retrieval</i> Nan Sun, Jing Tang, Lei Sun, Rui Chen, Yuxing Lu, Xiangxiang Chu, Hefei Ling and Yujun Cai	22318

<i>Large Reasoning Models Are (Not Yet) Multilingual Latent Reasoners</i>	Yihong Liu, Raoyuan Zhao, Hinrich Schuetze and Michael A. Hedderich	22330
<i>Towards Scalable Lightweight GUI Agents via Multi-role Orchestration</i>	Ziwei Wang, Junjie Zheng, Leyang Yang, Sheng Zhou, Xiaoxuan Tang, Fang Zhouhua, Zhiwei Liu, Dajun Chen, Yong Li and Jiajun Bu	22359
<i>Memory Matters More: Event-Centric Memory as a Logic Map for Agent Searching and Reasoning</i>	Yuyang Hu, Jiongnan Liu, Jiejun Tan, Yutao Zhu and Zhicheng Dou	22389
<i>Format-Adapter: Improving Reasoning Capability of LLMs by Adapting Suitable Format</i>	Dingzirui Wang, Xuanliang Zhang, Rongyu Cao, Longxu Dou, Xianzhen Luo, Yingwei MA, Qingfu Zhu, Binhua Li, Fei Huang and Yongbin Li	22408
<i>Behavior-Aware Item Modeling via Dynamic Procedural Solution Representations for Knowledge Tracing</i>	Jun Seo, Sangwon Ryu, Heejin Do, Hyoungun Kim and Gary Lee	22428
<i>Tracing the Light of Thought: A Probabilistic Self- and Cross-Consistency Verification Mechanism Improving Mathematical Reasoning in LLMs</i>	Xiaoyang Liu, Dawei Wang, Tian Li, Huizhi Liang, Gary Ushaw and Richard Davison	22444
<i>E2E-GMNER: End-to-End Generative Grounded Multimodal Named Entity Recognition</i>	Meng Zhang, Jinzhong Ning, Xiaolong Wu, Hongfei Lin and Yijia Zhang	22460
<i>Think Less, Know More: State-Aware Reasoning Compression with Knowledge Guidance for Efficient Reasoning</i>	Yi Sui, Chaozhuo Li and Dawei Song	22470
<i>Muse: Towards Reproducible Long-Form Song Generation with Fine-Grained Style Control</i>	Changhao Jiang, Jiahao Chen, Zhenghao Xiang, Zhixiong Yang, Hanchen Wang, Jiabao Zhuang, Xinmeng Che, Jiajun Sun, Hui Li, Yifei Cao, Shihan Dou, Ming Zhang, Junjie Ye, Tao Ji, Tao Gui, Qi Zhang and Xuanjing Huang	22492
<i>Distilling LLM Reasoning into Dense Encoders: Bridging the Accuracy-Efficiency Gap in Recommendation</i>	Donghee Han, Daeyoung Roh, A Young Kim, Hwanjun Song and Mun Yong Yi	22513
<i>Decoding the Multimodal Mind: Generalizable Brain-to-Text Translation via Multimodal Alignment and Adaptive Routing</i>	Chunyu Ye, Yunhao Zhang, Jingyuan Sun, Chong Li, Yang Zhao and Shaonan Wang	22532
<i>Which bird does not have wings: Negative-constrained KGQA with Schema-guided Semantic Matching and Self-directed Refinement</i>	Midan Shim, Seokju Hwang, KaeHyun Um and Kyong-Ho Lee	22547
<i>A Self-Evolving LLM Agent Framework for Role-Based Norm Compliance in Healthcare</i>	Haijie Ruan, Xiaowu Jiang, Zhanpeng LI, Wei Jia, Xuanwu Xu, Xiao-Fen Shan, Shujie Chen and Xindong Ye	22562
<i>M3TQA: Massively Multilingual Multitask Table Question Answering</i>	Daixin Shu, Jian Yang, Zhenhe Wu, Xianjie Wu, Xianfu Cheng, Guan Xiangyuan, Yanghai Wang, Pengfei Wu, Tingyang Yang, Hualei Zhu, Wei Zhang, Ge Zhang, Jiaheng Liu and Zhoujun Li	22578
<i>ShredBench: Evaluating the Semantic Reasoning Capabilities of Multimodal LLMs in Document Reconstruction</i>	Zichun Guo, Yuling Shi, Wenhao Zeng, Chao Hu, Haotian Lin, Terry Yue Zhuo, Jiawei Chen, Xiaodong Gu and Wenping Ma	22603

<i>TELL-TALE: Task Efficient LLMs with Task Aware Layer Elimination</i>	
Omar Naim, Krish Sharma, Niyar R Barman and Nicholas Asher	22616
<i>Reason-KE++: Aligning the Process, Not Just the Outcome, for Faithful LLM Knowledge Editing</i>	
Yuchen Wu, Liang Ding, Li Shen and Dacheng Tao	22639
<i>RACC: Regret-Aware Confidence Calibration for Consistent Masked Discrete Diffusion Decoding</i>	
Qinglin Zeng, Jusheng Zhang, Jing Yang, Ningyuan Liu and Keze Wang	22656
<i>The Impact of Off-Policy Training Data on Probe Generalisation</i>	
Adrians Skapars, Nathalie Maria Kirch, Samuel Dower, Ekdeep Singh Lubana and Dmitrii Krashenin- nikov	22673
<i>GRI: Reinforcement-Enhanced LLM for Geoscience Reasoning</i>	
Yule Xie, Jiaxin Ding, Cheng Deng, Shiqing Gao, Junran Zhang, Sibozhang, Zeyuan Wang, Ke Wu, Xin Ding, Luoyi Fu, Meng Jin and Xinbing Wang	22730
<i>FOCUS: A Fine-Grained Customer-Oriented Sentiment Dialogue Summarization Dataset for Chinese Customer Service</i>	
Qian Chen, Mengqiang Hu and Xin Guo	22744
<i>ABC-Bench: Benchmarking Agentic Backend Coding in Real-World Development</i>	
Jie Yang, Honglin Guo, Li Ji, Jiazheng Zhou, Rui Zheng, Zhikai Lei, Shuo Zhang, Zhiheng Xi, Shichun Liu, Yuxin Wang, Bo Wang, Yining Zheng, Tao Gui and Xipeng Qiu	22765
<i>LiveCANNBench: Benchmark SWE AI Coding for Ascend CANN</i>	
Sijie Wang, Kai Zhao, Wee Peng Tay, Shuo Zhang, Chengwen Liu, Quanjiang Guo, Ren Junhao, Xin Li, Heng Lian, Jingdi Lei, Rui She, Huacan Wang and Ronghao Chen	22788
<i>When Models Decide and When They Bind: A Two-Stage Computation for Multiple-Choice Question Answering</i>	
Hugh Mee Wong, Rick Nouwen and Albert Gatt	22804
<i>Who Watches the Watchmen? Humans Disagree With Translation Metrics on Unseen Domains</i>	
Finn Schmidt, Jan Philip Wahle, Terry Ruas and Bela Gipp	22822
<i>Stress-Testing Emotional Support Models: Moving from Homogeneous to Diverse Help Seekers</i>	
Chaewon Heo, Cheyon Jin and Yohan Jo	22842
<i>Progressive Planning and Reinforced Reasoning: Large Language Model-Guided Multi-hop Question Answering over Knowledge Graph</i>	
Xiang Li, Runhai Jiao, Ruifan Li, Dongnan Wu, Ruoqiao Qiao and Lei Liu	22870
<i>AED-RAG: Continuous Multi-Granular Context Fusion for Retrieval-Augmented Generation via Adap- tive Ensemble Decoding</i>	
Junzhe Zhou, Fulin Lin, Tairan Cheng, Shaowen Chen and Hongwei Wang	22885
<i>Neural Chain-of-Thought Search: Searching the Optimal Reasoning Path to Enhance Large Language Models</i>	
Guoming Ling, Zhongzhan Huang, Yupei Lin, Junxin Li, Shanshan Zhong, Hefeng Wu and Liang Lin	22900
<i>GOAT: A Training Framework for Goal-Oriented Agent with Tools</i>	
Hyunji Min, Sangwon Jung, Junyoung Sung, Dosung Lee, Leekyeung Han and Paul Hongsuck Seo	22934

<i>LLMs as ASP Programmers: Self-Correction Enables Task-Agnostic Nonmonotonic Reasoning</i> Adam Ishay and Joohyung Lee	22964
<i>DRIV-EX: Counterfactual Explanations for Driving LLMs</i> Amaia Cardiel, Eloi Zablocki, Elias Ramzi and Eric Gaussier	22994
<i>ROMA: Real-time Omni-Multimodal Assistant with Interactive Streaming Understanding</i> Tian Xueyun, Wei Li, Bingbing Xu, Heng Dong, Yuanzhuo Wang and Huawei Shen	23018
<i>Towards Provably Secure Generative AI: Reliable Consensus Sampling</i> Yu Cui, Hang Fu, Sicheng Pan, Zhuoyu Sun, Yifei Liu, Yuhong Nie, Bo Ran, Baohan Huang, Xufeng Zhang, Haibin Zhang, Cong Zuo and Licheng Wang	23040
<i>Pause or Fabricate? Training Language Models for Grounded Reasoning</i> Yiwen Qiu, Linjuan Wu, Yizhou Liu, Yuchen Yan, Jin Ma, Xu Tan, Yao Hu, Daoxin Zhang, Wenqi Zhang, Weiming Lu, Jun Xiao and Yongliang Shen	23056
<i>AgentV-RL: Scaling Reward Modeling with Agentic Verifier</i> Jiazheng Zhang, Ziche Fu, Zhiheng Xi, Wenqing Jing, Mingxu Chai, Wei He, Guoqiang Zhang, Chenghao Fan, Chenxin An, Wenxiang Chen, Zhicheng Liu, Haojie Pan, Dingwei Zhu, Tao Gui, Qi Zhang and Xuanjing Huang	23078
<i>DiTReducio: A Training-Free Acceleration for DiT-Based TTS via Progressive Calibration</i> Yanru Huo, Ziyue Jiang, Zuoli Tang, Qingyang Hong and Zhou Zhao	23101
<i>DcLM: Output Length Control of Large Language Models via Dynamic Length Markers</i> Zhe Chen, Jiaao Yu and Honglin Li	23117
<i>LongMP-Bench: A Benchmark for Multimodal Persona Understanding in Long-Term Dialogues</i> Zhuoqun Li, Zhaopei Huang, Wenxuan Wang and Qin Jin	23132
<i>TOOLCAD: Exploring Tool-Using Large Language Models in Text-to-CAD Generation with Reinforce- ment Learning</i> Yifei Gong, Xing Wu, Wenda Liu and Tukang	23161
<i>Beyond Overlap Metrics: Rewarding Reasoning and Preferences for Faithful Multi-Role Dialogue Sum- marization</i> Xiaoyong Mei, Tingting Zuo, Da Chen, Guangyu Hu, Xiangyu Wen, Chao Duan, Mingyan Zhang and Fudan Zheng	23189
<i>From Signal Degradation to Computation Collapse: Uncovering the Two Failure Modes of LLM Quan- tization</i> Chenxi Zhou, Pengfei Cao, Jiang Li, Bohan Yu, Jinyu Ye, Jun Zhao and Kang Liu	23204
<i>Establishing a Scale for Kullback-Leibler Divergence in Language Models Across Various Settings</i> Ryo Kishino, Yusuke Takase, Momose Oyama, Hiroaki Yamagiwa and Hidetoshi Shimodaira	23223
<i>Simple Role Assignment is Extraordinarily Effective for Safety Alignment</i> Zhou Ziheng, Jiakun Ding, Zhaowei Zhang, Ruosen Gao, Ying Nian Wu, Demetri Terzopoulos, Yipeng Kang, Fangwei Zhong and Junqi Wang	23249
<i>Task-Stratified Knowledge Scaling Laws for Post-Training Quantized Large Language Models</i> Chenxi Zhou, Pengfei Cao, Jiang Li, Bohan Yu, Jinyu Ye, Jun Zhao and Kang Liu	23268
<i>MemPO: Self-Memory Policy Optimization for Long-Horizon Agents</i> Ruoran Li, Xinghua Zhang, Haiyang Yu, Shitong Duan, Xiang Li, Wenxin Xiang, Chonghua Liao, Xudong Guo, Yongbin Li and Jinli Suo	23286

<i>Reason Only When Needed: Efficient Generative Reward Modeling via Model-Internal Uncertainty</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 . . .	23302
<i>Metacognitive Self-Correction for Multi-Agent System via Prototype-Guided Next-Execution Reconstruction</i>	
Xu Shen, Qi Zhang, Song Wang, Zhen Tan, Xinyu Zhao, Laura Yao, Vaishnav Tadiparthi, Hossein Nourkhiz Mahjoub, Ehsan Moradi Pari, Kwonjoon Lee and Tianlong Chen	23320
<i>DualGuard: Dual-stream Large Language Model Watermarking Defense against Paraphrase and Spoofing Attack</i>	
Hao Li, Yubing Ren, Yanan Cao, Yingjie Li, Fang Fang, Shi Wang and Li Guo	23338
<i>PeerCheck: Enhancing LLM-Generated Academic Reviews Towards Human-Level Quality</i>	
Zeyuan Chen, Ziqing Yang, Yihan Ma, Michael Backes and Yang Zhang	23362
<i>MorphoBench: A Benchmark with Difficulty Adaptive to Model Reasoning</i>	
Xukai Wang, Xuanbo Liu, Mingrui Chen, Haitian Zhong, Xuanlin Yang, Bohan Zeng, Jinbo Hu, Hao Liang, Junbo Niu, Xuchen Li, Ruitao Wu, Ruichuan An, Yang Shi, Liu Liu, Qiang Liu, Zhouchen Lin, Xu-Yao Zhang, Wentao Zhang and Bin Dong	23387
<i>Privacy Risks of Intermediate Representations: Attribute Inference in Distributed LLM Inference</i>	
Yang Lyu, Jin Cao, Yang Xiao, Zhe Sun, Ben Niu, Fenghua Li and Hui LI	23412
<i>ReCon: Active Defense against Large Vision-Language Model Jailbreaks via Reverse Safety Concept Injection</i>	
Zheng He, Yiwei Wang, Hongxing Wang and Yujun Cai	23427
<i>Verifying the Subjective: Structured Multilingual Rewards for Low-Resource Alignment</i>	
Jiu Sha and Mengxiao Zhu	23442
<i>Memory-Driven Role-Playing: Evaluation and Enhancement of Persona Knowledge Utilization in LLMs</i>	
Kai Wang, Haoyang You, Yang Zhang and Zhongjie Wang	23475
<i>Scripts Through Time: A Survey of the Evolving Role of Transliteration in NLP</i>	
Thanmay Jayakumar, Deepon Halder and Raj Dabre	23511
<i>Health-ORSC-Bench: A Benchmark for Measuring Over-Refusal and Safety Completion in Health Context</i>	
Zhihao Zhang, Liting Huang, Guanghao Wu, Preslav Nakov, Heng Ji and Usman Naseem	23525
<i>LLM Inductive Reasoning Through Multi-Agent Enhanced Monte Carlo Tree Search</i>	
Xiang Li, Yucheng Zhou, Xiangzhi Wei, Zesheng Shi, Haiyuan Wan, Gong Yifan, Fangming Liu and Jing Li	23548
<i>UniCreative: Unifying Long-form Logic and Short-form Sparkle via Reference-Free Reinforcement Learning</i>	
Xiaolong Wei, Zerun Zhu, Simin Niu, Xingyu Zhang, Peiying Yu, Changxuan Xiao, Yuchen Li, Jicheng Yang, Zhejun Zhao, Chong Meng, Long Xia and Daiting Shi	23563
<i>FFN Lens: How Transformers Divide Labor for Multilingual Tasks</i>	
Jiatong Li, Hailong Cao and Yang Liu	23584
<i>Seeing Beyond Words: MatVQA for Challenging Visual-Scientific Reasoning in Materials Science</i>	
Sifan Wu, Huan Zhang, Yizhan Li, Farshid Effaty, Hongyuan Mei, Amirreza Ataei and Bang Liu	23599

<i>RLShield: Dynamic Jailbreak Detection for LLMs via Reinforced Adaptive Learning</i>	
Zhao Tong, Pengfei Yang, Yimeng Gu, Haichao Shi, Qiang Liu, Xingcheng Xu, Shu Wu and Xiao-Yu Zhang	23615
<i>MEDSYN: Benchmarking Multi-Evidence SYNthesis in Complex Clinical Cases for Multimodal Large Language Models</i>	
Boqi Chen, Xudong Liu, Jiachuan Peng, Marianne Frey-Marti, Kyle Lam, Bang Zheng, Lin Li and Jianing Qiu	23631
<i>Sem-DPO: Mitigating Semantic Inconsistency in Preference Optimization for Prompt Engineering</i>	
Anas Mohamed, Azal Ahmad Khan, Xinran Wang, Ahmad Faraz Khan, Shuwen Ge, Saman Bahzad Khan, Ayaan Ahmad and Ali Anwar	23656
<i>BNLP: A Text Annotation Platform for Quality Control of LLM-Generated Annotations</i>	
Xinhao Zhuang, Qiongyu Tian, Yalin Chen, Tianle Xin, Yongyong Fu, Yunchao Ling and Guoqing Zhang	23675
<i>ElasticFlow: One-Step Physics-Consistent Policy with Elastic Time Horizons for Language-Guided Manipulation</i>	
Kewei Chen, Yayu Long, Shuai Li and Mingsheng Shang	23685
<i>AdaptiveK: Complexity-Driven Sparse Autoencoders for Interpretable Language Model Representations</i>	
Yifei Yao, Hanrong Zhang and Mengnan Du	23702
<i>Experience-Driven Multi-Agent Optimization for Black-Box Jailbreak Attacks on Large Language Models</i>	
Zhaoyang Han, Yihe Liu, Kai Zhang and Ping Li	23729
<i>CMIG: Conceptual Metaphor Theory-Inspired Framework for Metaphorical Image Generation</i>	
Qingbao Huang, Cheng Yang, Jiawei Yao, Zhiyue Liu, Yi Cai and Xingmao Zhang	23748
<i>Position: From Noise to Signal to Selbstzweck - Reframing Human Label Variation in the Era of Post-training in NLP</i>	
Shanshan Xu, Santosh T.y.s.s and Barbara Plank	23762
<i>Rethinking Assessments of Prompt Injection Attacks</i>	
Chi Cui, Yixin Wu, Michael Backes and Yang Zhang	23773
<i>SegDRE: A Salient Entity Guided Approach to Document-Level Relation Extraction</i>	
Xing Yang and Chengxiang Tan	23800
<i>Live-Aid: A Large-Scale Dialogue Dataset and Benchmark for Interleaved Multi-party Interactions in Live Streaming</i>	
Yiming Lei, Yize Fan, Zeming Liu, Jiayi Dong, Hui Qiu, Haitao Leng, Qingjie Liu, Kehai Chen, Tingting Gao and Yunhong Wang	23813
<i>Semi-Supervised Diseased Detection from Speech Dialogues with Multi-Level Data Modeling</i>	
Xingyuan Li and Mengyue Wu	23851
<i>From Isolated Scoring to Collaborative Ranking: A Comparison-Native Framework for LLM-Based Paper Evaluation</i>	
Pujun Zheng, Jiacheng Yao, Jinqian Zheng, Chenyang Gu, Guoxiu He, Jiawei Liu, Yong Huang, Tianrui Guo and Wei Lu	23863

<i>UnAC: Adaptive Visual Prompting with Abstraction and Stepwise Checking for Complex Multimodal Reasoning</i>	
Yifan Wang and Yun Fu	23884
<i>AnalystBench: Benchmarking professional long-form report generation with web-mined multimodal tasks</i>	
Chau Minh Pham, Zichao Wang, Puneet Mathur, Alexa Siu, Akriti Jain, Aparna Garimella, Ananya B. Sai, Nedim Lipka, Mohit Iyyer and Varun Manjunatha	23894
<i>BizCompass: Benchmarking the Reasoning Capabilities of LLMs in Business Knowledge and Applications</i>	
Jianing Hao, Yuhe Wu, Yuanjian Xu, Shichang Meng, Shuai Yuan, Wei Zeng, Zixuan Wang and Guang Zhang	23927
<i>When More Thinking Hurts: Overthinking in LLM Test-Time Compute Scaling</i>	
Shu Zhou, Rui Ling, Junan Chen, Xin Wang, Tao Fan and Hao Wang	23967
<i>Learning from Failures: Error Notebook-guided Secure Code Generation</i>	
Xinyu Zhong, Peng Lan and Zhifang Liao	23978
<i>Evaluating Reasoning Models for Queries with Presuppositions</i>	
Rose Sathyanathan, Kinshuk Vasisht and Danish Pruthi	23991
<i>GRE Score: Generative Risk Evaluation for Large Language Models</i>	
Zaitang LI, Pin-Yu Chen and Tsung-Yi Ho	24007
<i>MSVBench: Towards Human-Level Evaluation of Multi-Shot Video Generation</i>	
Haoyuan Shi, Yunxin li, Nanhao Deng, Zhenran Xu, Xinyu Chen, Longyue Wang, Baotian Hu and Min Zhang	24034
<i>M-TRACE: Detecting and Mitigating Time-Anchor Drift via Step-wise Conflict Checking in Temporal Reasoning</i>	
Danyu Huang, Jiayuan Jiang, Yao Zhang, Jun Wang, Huijia Li and Zhenglu Yang	24059
<i>Faithful Persona Steering under Incongruity via Dual-Stream Refinement</i>	
Yu-An Chu, Jen-Ren Pong, Chia-Yao Yeh and Meng-Fen Chiang	24078
<i>ChemVLR: Prioritizing Reasoning in Perception for Chemical Vision-Language Understanding</i>	
Xuanle Zhao, Xinyuan Cai, Xiang Cheng, Xiuyi Chen and Bo XU	24096
<i>SDC-LoRA: Singular-Subspace Drift Controlled LoRA to Mitigate Knowledge Forgetting</i>	
Geyuan Zhang, Xiaofei Zhou, Shihao Liu, Jingyuan Tian and Jizheng Ma	24116
<i>Regret-Now: A Physics-Inspired Regret Framework for Temporal Knowledge Graph Question Answering with LLMs</i>	
Danyu Huang, Yao Zhang, Jun Wang and Zhenglu Yang	24133
<i>Low-probability Tokens Sustain Exploration in Reinforcement Learning with Verifiable Reward</i>	
Guanhua Huang, Tingqiang Xu, Mingze Wang, Qi Yi, Xue Gong, Siheng Li, Ruibin Xiong, Kejiao Li, Yuhao Jiang and Bo Zhou	24158
<i>Preserving Language Capabilities in Vision-Language Models via Representation Regulation</i>	
ZiXuan Chen, Juncheng Tao and Ziqian Zeng	24189
<i>Me-Agent: A Personalized Mobile Agent with Two-Level User Habit Learning for Enhanced Interaction</i>	
Shuoxin Wang, Chang Liu, Gowen Loo, Lifan Zheng, Kaiwen Wei, Huanqian Yan, Xinyi Zeng, Jingyuan Zhang and Yu Tian	24206

<i>Dynamic PMI-Guided Contrastive Decoding Reduces Hallucination in Large Language Models: A Unified Framework of Fine-Grained Input Transformations</i>	
Dongsheng Chen, Yingqi Zhu, Xingyue Zhang, Wenqing Zhou and Lei Li	24223
<i>LSEG: A Fine-tuning Free Method for NL2FOL via Logic-Structure and Entropy Guided Inference Controlling</i>	
Yuwei Huang, Shi Wang, Kangli Zi, Tianyu Luo, Jiang Zhixiao, Ruiheng Wang and Yufei Wang	24236
<i>Compact Example-Based Explanations for Language Models</i>	
Loris Schoenegger and Benjamin Roth	24246
<i>MedCPI: A Construct–Personalize–Integrate Framework for KG-enhanced Clinical Prediction</i>	
Hang Wang, Hang Dong and Lu Liu	24266
<i>Across Programming Language Silos: A Study on Cross-Lingual Retrieval-Augmented Code Generation</i>	
Qiming Zhu, Jialun Cao, Xuanang Chen, Weili Zhang, Yaojie Lu, Hongyu Lin, Xianpei Han, Le Sun and Shing-Chi Cheung	24283
<i>ErrorRadar: Benchmarking Complex Mathematical Reasoning of Multimodal Large Language Models Via Error Detection</i>	
Yibo Yan, Shen Wang, Jiahao Huo, Hang Li, Boyan Li, Jiamin Su, Xiong Gao, YiFan Zhang, Tianlong Xu, Zhendong Chu, Aoxiao Zhong, Kun Wang, Hui Xiong, Philip S. Yu, Xuming Hu and Qingsong Wen	24297
<i>Disentangling Continued Pre-Training: Attention-Driven Routing and Semantic Hub Preservation in Language Adaptation</i>	
Khanh-Tung Tran, Vinh-Khanh Tran, Barry O’Sullivan and Hoang D. Nguyen	24335
<i>Do You Get the Hint? Benchmarking LLMs on the Board Game Concept</i>	
Ine Gevers and Walter Daelemans	24358
<i>Topic-Based Watermarks for Large Language Models</i>	
Alexander Nemecek, Yuzhou Jiang and Erman Ayday	24372
<i>SPaRk-Eval: Evaluating Structure-Aware Knowledge Acquisition in LLMs for Domain Adaptation to Industrial Records</i>	
Ekant Muljibhai Amin, Yuta Koreeda and Yasuhiro Sogawa	24403
<i>SOLAR-RL: Semi-Online Long-horizon Assignment Reinforcement Learning</i>	
Jichao Wang, Liuyang Bian, Yufeng Zhou, Han Xiao, Yue Pan, Guozhi Wang, Hao Wang, Zhaoxiong Wang, Yafei Wen, Xiaoxin Chen, Shuai Ren and Lingfang Zeng	24419
<i>Kumatigi: Quality-Driven Data Augmentation for Low-Resource Machine Translation</i>	
Cheick Tidiani Cissé	24433
<i>CAMO: An Agentic Framework for Automated Causal Discovery from Micro Behaviors to Macro Emergence in LLM Agent Simulations</i>	
Xiangning Yu, Yuwei Guo, Yuqi Hou, Xiao Xue and Qun Ma	24447
<i>MoLAN: A Unified Modality-Aware Noise Dynamic Editing Framework for Multimodal Sentiment Analysis</i>	
Xingle Xu, YongKang Liu, Dexian Cai, Shi Feng, Xiaocui Yang, Daling Wang and Yifei Zhang	24480

<i>C2PO: Diagnosing and Disentangling Bias Shortcuts in LLMs</i>	
Xuan Feng, Bo An, Tianlong Gu, Liang Chang, Fengrui Hao, Peipeng Yu and Shuai Zhao	24496
<i>Deconstruct, Diagnose, and Deliberate: A Protocol-Adaptive Role-Specific Multi-Agent Framework for Fake News Detection</i>	
Zhigan Zhang and Jie Sui	24516
<i>Position: Multimodal Large Language Models Can Significantly Advance Scientific Reasoning</i>	
Yibo Yan, Shen Wang, Jiahao Huo, Jingheng Ye, Zhendong Chu, Xuming Hu, Philip S. Yu, Carla P Gomes, Bart Selman and Qingsong Wen	24535
<i>E³-TIR: Enhanced Experience Exploitation for Tool-Integrated Reasoning</i>	
Weiyang Guo, Zesheng Shi, Liye Zhao, Jiayuan Ma, Zeen Zhu, Junxian He, Min Zhang and Jing Li	24575
<i>EVA: Evolving Semantic Adversaries for Red-Teaming GUI Agents Against Environmental Injection Attacks</i>	
Yijie Lu, Manman Zhao, Tianjie Ju, Zihe Yan, Xinbei Ma, Yuan Guo, Daizong Ding, Gongshen Liu and Zhuosheng Zhang	24597
<i>NEAT: Neuron-Based Early Exit for Large Reasoning Models</i>	
Kang Liu, YongKang Liu, Xiaocui Yang, Peidong Wang, Wen Zhang, Shi Feng, Yifei Zhang and Daling Wang	24616
<i>From “Thinking” to “Justifying”: Aligning High-Stakes Explainability with Professional Communication Standards</i>	
Chen Qian, Yimeng Wang, Yu Chen, Lingfei Wu and Andreas Stathopoulos	24628
<i>EHRAG: Bridging Semantic Gaps in Lightweight GraphRAG via Hybrid Hypergraph Construction and Retrieval</i>	
Yifan Song, Xingjian Tao, Zhicheng Yang, Yihong Luo and Jing Tang	24638
<i>PURE: Post-hoc Unlocking and REfinement for Discrete Diffusion Decoding</i>	
Yangryeol Park, Kunhui Lee, Hanback Choi, Cheoneum Park, Donghyeon Jeon, Inho Kang and Seung-Hoon Na	24653
<i>OSCR-Attack: One-Shot Character Level Attacks through Self-Optimizing Continuous Relaxation</i>	
Lingyi Kong, Zhuo Liu, Zhanghao Hu, Qilong Qiu, Yutao Yang, Jingjing Xue, Zheng Wang, Lin Gui and Feiping Nie	24668
<i>MTEB-NL and E5-NL: Embedding Benchmark and Models for Dutch</i>	
Nikolay Banar, Ehsan Lotfi, Jens Van Nooten, Cristina Arhiliuc, Marija Kliocaitė and Walter Daelemans	24684
<i>Perceptual Hallucination in Vision–Language Models: Definition, Analysis and Verification</i>	
Taewook Hwang, Inbum Heo, Sung Jun Lee and Sangkeun Jung	24710
<i>RATION: Entropy-Driven Task-Adaptive Visual Attention Allocation Framework for Multimodal Reasoning</i>	
Xingle Xu, Fanheng Kong, Dexian Cai, Shi Feng, Xiaocui Yang, Daling Wang and Yifei Zhang	24726
<i>Multimodal In-context Learning for ASR of Low-resource Languages</i>	
Zhaolin Li and Jan Niehues	24745

<i>a1: Steep Test-time Scaling Law via Environment Augmented Generation</i> Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Yuyao Ge, Jun Wan, Yurong Wu and Xueqi Cheng	24761
<i>Look and Think: Efficient Multimodal Reasoning via Modality-Decoupled Compression</i> Xidi Cai, Junhao Zheng, Jingye Li, Boyuan Li, Shaowei Zhang and Qianli Ma	24787
<i>Topology-Enhanced Alignment for Large Language Models: Trajectory Topology Loss and Topological Preference Optimization</i> Yurui Pan, Ke Xu and Bo Peng	24807
<i>Inference-Time Scaling of Verification: Self-Evolving Deep Research Agents via Test-Time Rubric-Guided Verification</i> Yuxuan Wan, Tianqing Fang, Zaitang LI, Yintong Huo, Wenxuan Wang, Haitao Mi, Dong Yu and Michael R. Lyu	24822
<i>FinCARDS: Card-Based Analyst Reranking for Financial Document Question Answering</i> Yixi Zhou, Fan Zhang, YU Chen, Haipeng Zhang, Preslav Nakov and Zhuohan Xie	24836
<i>A Unified Feature Mixture Framework for Joint Speech and Singing Deepfake Detection</i> Aastha Sharma and Guangjing Wang	24853
<i>From Atomic to Complex tasks: Cross-Tasking Improves Zero-Shot Argument Generation and Retrieval</i> Lynda Tamine, Ahmed Rayane Kebir, Merveille Dona Codjo, Enzo Pasquies and Jose G Moreno	24864
<i>Sculpting the Vector Space: Towards Efficient Multi-Vector Visual Document Retrieval via Prune-then-Merge Framework</i> Yibo Yan, Mingdong Ou, Yi Cao, Xin Zou, Jiahao Huo, Shuliang Liu, James Kwok and Xuming Hu	24883
<i>ThinkLinker: From Low-Rank Interaction to Knowledge-Aware Verification for Multimodal Entity Linking</i> Yingyao Ma, Yuanyuan Zhou, Congyu Zhang, Yi Yuan, Jiasong Wu, Lotfi Senhadji and Huazhong Shu	24926
<i>Logit Arithmetic Elicits Long Reasoning Capabilities Without Training</i> Yunxiang Zhang, Muhammad Khalifa, Lechen Zhang, Xin Liu, Ayoung Lee, Xinliang Frederick Zhang, Farima Fatahi Bayat and Lu Wang	24943
<i>Evolving Agentic Workflow Driven by Human-Agent Collaboration</i> Yuxin Liu, Jinxuan Zhang, Yuezhong Peng, Hefeng Zhou, Xiangfeng Wang, Jiong Lou, Chentao Wu, Jie LI, Jingjing Qu and Chaochao Lu	24960
<i>HintPilot: LLM-based Compiler Hint Synthesis for Code Optimization</i> Hanyun Jiang, Peisen Yao, Kaiyue Li, Tingting Lin, Chengpeng Wang and Kui Ren	24970
<i>What Makes an LLM a Good Optimizer? A Trajectory Analysis of LLM-Guided Evolutionary Search</i> Xinhao Zhang, Xi Chen, François Portet and Maxime Peyrard	24987
<i>Token-level Inference-Time Alignment for Vision-Language Models</i> Kejia Chen, Junjun Zheng, Jiawen Zhang, Manxi Lin, Xiao Pan, Jiacong Hu, Jian Lou, Zunlei Feng and Mingli Song	25012
<i>STELLA: A Multimodal LLM for Protein Functional Annotation via Unified Sequence-Structure Encoding</i> Hongwang Xiao, Wenjun Lin, Xi Chen, Hui Wang, Kai Chen, Jiashan Li, Yuancheng Sun, Sicheng Dai, Boya Wu and Qiwei Ye	25030

<i>RealFin: How Well Do LLMs Reason About Finance When Users Leave Things Unsaid?</i> Yuyang Dai, Yan Lin, Zhuohan Xie and Yuxia Wang	25050
<i>Preference Optimization for Review Question Generation Improves Writing Quality</i> Karun Sharma, Vidushee Vats, Shengzhi LI, Yuxiang Wang, Zhongtian Sun and Prayag Tiwari	25081
<i>Your LLM Agent Can Leak Your Data: Data Exfiltration via Backdoored Tool Use</i> Wuyang Zhang and Shichao Pei	25105
<i>EvoHyper: Evolving Hypergraph Topologies for Unified Collaboration in Multi-Agent Communication</i> Heng Zhang, Yihao Zhong, Lubin Gan, Zhihe Chen, Jiajun Wu, Yuling Shi, Xiaodong Gu, Hao Zhang, Haochen You and Jin Huang	25130
<i>Think Before you Write: QA-Guided Reasoning for Character Descriptions in Books</i> Argyrios Papoudakis, Mirella Lapata and Frank Keller	25144
<i>GroupRank: A Groupwise Paradigm for Effective and Efficient Passage Reranking with LLMs</i> Meixiu Long, Duolin Sun, Dan Yang, Yihan Jiao, Lei Liu, Jiahai Wang, Binbin Hu, Yue Shen, Jie Feng, Zhehao Tan, Junjie Wang, Lianzhen Zhong, Jian Wang, Peng Wei and Jinjie Gu	25165
<i>Mirage: A Diagnostic Framework for Evaluating the Realism of Synthetic Contact Center Dialogue Generation</i> Rishikesh Devanathan, Varun Nathan and Ayush Kumar	25187
<i>BloomEval: A Bloom’s Cognitive Taxonomy-Based Benchmark for Evaluating LRMs via Cognitive Hierarchy Trace</i> Zhiyi Duan, Lei Gao, Jiangshan Guan, Qi Wang and Rui Liu	25228
<i>CAIR: Causal Adaptive Information-based Reinforcement Learning for Multimodal Emotion Reasoning</i> Fengyu Zhang, Bin Liu, Jianhua Tao, Zhuofan Wen, Shun Chen, Hailiang Yao and Zhengqi Wen	25249
<i>A Multi-View Media Profiling Suite: Resources, Evaluation, and Analysis</i> Muhammad Arslan Manzoor, Dilshod Azizov, Daniil Orel, Umer Siddique, Zain Muhammad Mujahid, Yufang Hou and Preslav Nakov	25265
<i>Mechanistic Interpretability of Text-to-Image Diffusion Models via Cross-Attention Interventions</i> Maisha Maliha and Dean F. Hougen	25287
<i>Revisiting Entropy in Reinforcement Learning for Large Reasoning Models</i> Renren Jin, Pengzhi Gao, Yuqi Ren, Zhuowen Han, Tongxuan Zhang, Wuwei Huang, Wei Liu, Jian Luan and Deyi Xiong	25300
<i>How Large Language Models Balance Internal Knowledge with User and Document Assertions</i> Shuowei Li, Haoxin Li, Wenda Chu and Yi Fang	25323
<i>ComicVQA: A Benchmark for Visual Reasoning in Multimodal LLMs</i> Esther Gan, Hannah Brown, David Herel, Kenji Kawaguchi, Min-Yen Kan and Michael Qizhe Shieh	25347
<i>AJ-Bench: Benchmarking Agent-as-a-Judge for Environment-Aware Evaluation</i> Wentao Shi, Yu Wang, Yuyang Zhao, Yuxin Chen, Fuli Feng, Xueyuan Hao, Xi Su, Qi GU, Hui Su, Xunliang Cai and Xiangnan He	25371
<i>Can LLMs Estimate Student Struggles? Human-AI Difficulty Alignment with Proficiency Simulation for Item Difficulty Prediction</i> Ming Li, Han Chen, Yunze Xiao, Jian Chen, Hong Jiao and Tianyi Zhou	25414

<i>Transition-Matrix Regularization for Next Dialogue Act Prediction in Counselling Conversations</i>	
Eric Rudolph, Philipp Steigerwald and Jens Albrecht	25442
<i>ClimateCause: Complex and Implicit Causal Structures in Climate Reports</i>	
Liesbeth Allein, Nataly Pineda-Castañeda, Andrea Rocci and Marie-Francine Moens	25458
<i>Reference-Free Evaluation of Taxonomies</i>	
Pascal Wulschleger, Majid Zarharan, Donnacha Daly, Marc Pouly and Jennifer Foster	25489
<i>CSI: An Investigative Multi-Agent Framework for Explainable Short Video Fake News Detection</i>	
Yuxin Wang, Yang Yang and Huaiwen Zhang	25508
<i>Dial HEALTHDIAL for Advice: A Multilingual and Multi-Parallel Spoken Dialogue Dataset for Knowledge-Grounded Information Seeking</i>	
Songbo Hu, Yinhong Liu, Ej Zhou, Evgeniia Razumovskaia, Xiaobin Wang, Alexander Fraser, Ivan Vulić and Anna Korhonen	25529
<i>Fast and Effective On-Policy Distillation from Reasoning Prefixes</i>	
Dongxu Zhang, Zhichao Yang, Sepehr Janghorbani, Jun Han, Andrew Ressler II, Qian Qian, Gregory D Lyng, Sanjit Singh Batra and Robert E. Tillman	25553
<i>MaRF: Leveraging Representation-Level Fusion of Formula Semantics for Mathematical Information Retrieval</i>	
Suyuan Wang, Hongbo Zheng and Nickvash Kani	25570
<i>From Documents to Segments: A Contextual Reformulation for Topic Assignment</i>	
Hoonsang Yoon, Takyoun Kim, Wonkee Lee, Ilmin Cho, Dilek Hakkani-Tür and Stanley Jungkyu Choi	25586
<i>Can LLMs Compress (and Decompress)? Evaluating Code Understanding and Execution via Invertibility</i>	
Nickil Maveli, Antonio Vergari and Shay B Cohen	25625
<i>Confidence Estimation for LLMs in Multi-turn Interactions</i>	
Caiqi Zhang, Ruihan Yang, Xiaochen Zhu, Chengzu Li, Tiancheng Hu, Yijiang River Dong, Deqing Yang and Nigel Collier	25661
<i>Beyond Black-Box Labels: Interpretable Criteria for Diagnosing Subjective NLP Tasks</i>	
Nisrine Rair, Alban Goupil, Valeriu Vrabie and Emmanuel Chochoy	25677
<i>AutoRubric: Rubric-Based Generative Rewards for Faithful Multimodal Reasoning</i>	
Mengzhao Jia, Zhihan Zhang, Ignacio Cases, Zheyuan Liu, Meng Jiang and Peng Qi	25707
<i>HER: Human-like Reasoning and Reinforcement Learning for LLM Role-playing</i>	
Chengyu Du, Xintao Wang, Aili Chen, Weiyuan Li, Rui Xu, Junteng Liu, Zishan Huang, Rong Tian, Zijun Sun, Yuhao Li, Liheng Feng, Deming Ding, Pengyu Zhao and Yanghua Xiao	25725
<i>Do Generalisation Results Generalise?</i>	
Matteo Boglioni, Andrea Sgobbi, Gabriel Tavernini, Francesco Rita, Marius Mosbach and Tiago Pimentel	25763
<i>Do LLMs Overthink Basic Math Reasoning? Benchmarking the Accuracy-Efficiency Tradeoff in Language Models</i>	
Gaurav Srivastava, Aafiya Shamshad Hussain, Sriram Srinivasan and Xuan Wang	25784
<i>PV-SQL: Synergizing Database Probing and Rule-based Verification for Text-to-SQL Agents</i>	
Yuan Tian and Tianyi Zhang	25827

<i>Anatomy of Unlearning: The Dual Impact of Fact Saliency and Model Fine-Tuning</i>	
Anna Borisiuk, Andrey Savchenko, Alexander Panchenko and Elena Tutubalina	25846
<i>What Do Neural Speech Models Know About Phonology? Evidence from Structured Phoneme Confusions</i>	
Eli Stafford, Aimée Lahaussais and Guillaume Wisniewski	25860
<i>Rerank Before You Reason: Analyzing Reranking Tradeoffs through Effective Token Cost in Deep Search Agents</i>	
Sahel Sharifymoghaddam and Jimmy Lin	25874
<i>Position: LLM Watermarking Should Align Stakeholders' Incentives for Practical Adoption</i>	
Yepeng Liu, Xuandong Zhao, Dawn Song, Gregory W. Wornell and Yuheng Bu	25887
<i>Learning to Refine: Self-Refinement of Parallel Reasoning in LLMs</i>	
Qibin Wang, Pu Zhao, Shaohan Huang, Fangkai Yang, Lu Wang, Furu Wei, Qingwei Lin, Saravan Rajmohan and Dongmei Zhang	25904
<i>AUDITA: A New Dataset to Audit Humans vs. AI Skill at Audio QA</i>	
Tasnim Kabir, Dmytro Kurdydyk, Aadi Palnitkar, Liam Dorn, Ahmed Haj Ahmed and Jordan Lee Boyd-Graber	25922
<i>NeoAraBERT: A Modern Foundation Model for Arabic Embeddings with Diacritics-Aware Tokenization and POS-Targeted Masking</i>	
Chadi Abou Chakra, Hadi Khaled Hamoud, Osama Rakan Al Mraikhat, Qusai Abu Obaida, Mohammad Ballout and Fadi Zaraket	25952
<i>Steering Away from Refusal: A Black-box Jailbreak Method Based on First-Token Distribution</i>	
Shuangjie Fu, Du Su, Xin Chen, Fei Sun, Huawei Shen and Xueqi Cheng	25969
<i>Mina: A Multilingual LLM-Powered Legal Assistant Agent for Empowering Access to Justice in Bangladesh</i>	
Azmine Touseh Wasi, Wahid Faisal, Mst Rafia Islam and Md Rizwan Parvez	25980
<i>Zero-Shot Context-Aware ASR for Diverse Arabic Varieties</i>	
Bashar Talafha, Amin Abu Alhassan and Muhammad Abdul-Mageed	26029
<i>The Sonar Moment: An Audio Geo-Localization Benchmark for Audio-Language Models</i>	
Ruixing Zhang, Zihan Liu, Leilei Sun, Tongyu Zhu and Weifeng Lv	26045
<i>Towards Understanding the Robustness of Sparse Autoencoders</i>	
Ahson Saiyed, Sabrina Sadikeh and Chirag Agarwal	26066
<i>Beyond Local vs. External: A Game-Theoretic Framework for Trustworthy Knowledge Acquisition</i>	
Rujing Yao, Yufei Shi, Yang Wu, Ang Li, Zhuoren Jiang, XiaoFeng Wang, Haixu Tang and Xiaozhong Liu	26087
<i>False Sense of Security: Why Probing-based Malicious Input Detection Fails to Generalize</i>	
Cheng Wang, Zeming Wei, Qin Liu, Wenxuan Zhou and Muhao Chen	26100
<i>Do Transformers Grok Succinct Algorithms? Mechanistic Evidence for Counting Circuits</i>	
Yizheng Zhao	26114
<i>How Value Induction Reshapes LLM Behavior</i>	
Arnav Arora, Natalie Schluter, Katherine Metcalf and Maartje Ter Hoeve	26131
<i>Stable Evidence, Unstable Decisions: An Empirical Analysis of Model Decision Stability in Vision-Language Models</i>	
Ali Khoramfar, Mohammad Javad Dousti, Alireza Mohamadian and Hesham Faili	26153

<i>PlanRAG-Audio: Planning and Retrieval Augmented Generation for Long-form Audio Understanding</i> Masao Someki, Chien-yu Huang, Siddhant Arora, Samuele Cornell, Markus Müller, Nathan Susanj, Rupak Vignesh Swaminathan, Grant Strimel, Jing Liu and Shinji Watanabe	26167
<i>VecCISC: Improving Confidence-Informed Self-Consistency with Reasoning Trace Clustering and Candidate Answer Selection</i> James Petullo, Sonny George, Dylan Cashman and Nianwen Xue	26184
<i>Incentivizing In-depth Reasoning over Long Contexts with Process Advantage Shaping</i> Miao Peng, Weizhou Shen, Nuo Chen, Chenliang Li, Ming Yan and Jia Li	26201
<i>Dissecting Clinical Reasoning in Natural Language Inference for Large Language Models</i> Mael Jullien, Andre Freitas, Marco Valentino and Leonardo Ranaldi	26229
<i>LLM-Codec: Neural Audio Codec Meets Language Model Objectives</i> Ho-Lam Chung, Yiming Chen and Hung-yi Lee	26251
<i>Who is the richest club in the championship? Detecting and Rewriting Underspecified Questions Improve QA Performance</i> Yunchong Huang, Gianni Barlacchi and Sandro Pezzelle	26266
<i>Beyond Chunks and Graphs: Retrieval-Augmented Generation through Triplet-Driven Thinking</i> Shengbo Gong, Xianfeng Tang, Qi He, Carl Yang and Wei Jin	26282
<i>Where Reasoning Breaks: Logic-Aware Path Selection by Controlling Logical Connectives in LLMs Reasoning Chains</i> Seunghyun Park and Yuanyuan Lei	26309
<i>Reheat Nachos for Dinner? Evaluating AI Support for Cross-Cultural Communication of Neologisms</i> Dayeon Ki, Yu Hou, Rachel Rudinger, Hal Daumé Iii, Marine Carpuat and Fumeng Yang .	26325
<i>Learning Dynamic Representations and Policies from Multimodal Clinical Time-Series with Informative Missingness</i> Zihan Liang, Ziwen Pan and Ruoxuan Xiong	26363
<i>Quantize What Counts: More for Keys, Less for Values</i> Mohsen Hariri, Alan Luo, Weicong Chen, Tianyi Zhang, Qifan Wang, Xiaotian Han and Vipin Chaudhary	26393
<i>PersonaAgent: Bridging Memory and Action for Personalized LLM Agents</i> Weizhi Zhang, Xinyang Zhang, Chenwei Zhang, Liangwei Yang, Jingbo Shang, Zhepei Wei, Henry Peng Zou, Zijie Huang, Zhengyang Wang, Yifan Gao, Xiaoman Pan, Lian Xiong, Jingguo Liu, Philip S. Yu and Xian Li	26421
<i>Beyond Fixed Psychological Personas: State Beats Trait, but Language Models are State-Blind</i> Tamunotonye Harry, Ivoline C. Ngong, Chima Nweke, Yuanyuan Feng and Joseph Near . .	26440
<i>From RAG to Agentic RAG for Faithful Islamic Question Answering</i> Gagan Bhatia, Hamdy Mubarak, Mustafa Jarrar, George Mikros, Fadi Zaraket, Mahmoud Al-hirhani, Mutaz al-Khatib, Logan Cochrane, Kareem Mohamed Darwish, Rashid Yahiaoui and Firoj Alam	26469
<i>Losses that Cook: Topological Optimal Transport for Structured Recipe Generation</i> Mattia Ottoborgo, Daniele Rege Cambrin and Paolo Garza	26489
<i>Concept Tokens: Learning Behavioral Embeddings Through Concept Definitions</i> Ignacio Sastre and Aiala Rosá	26501

<i>Beyond Monolithic Rewards: Hybrid Multi-Aspect Reward Optimization</i>	
Radha Gulhane and Sathish Reddy Indurthi	26519
<i>VAUQ: Vision-Aware Uncertainty Quantification for LVLN Self-Evaluation</i>	
Seongheon Park, Changdae Oh, Hyeong Kyu Choi, Sean Du and Sharon Li	26534
<i>Will it Merge? On The Causes of Model Mergeability</i>	
Adir Rahamim, Asaf Yehudai, Boaz Carmeli, Leshem Choshen, Yosi Mass and Yonatan Belinkov	26551
<i>Carefully Considering Culture: Analyzing LLM Alignment in Single- and Multi-Cultural Settings using Cultural Consensus Theory</i>	
Krishna Pothugunta and John P. Lalor	26571
<i>Parallel-SFT: Improving Zero-Shot Cross-Programming-Language Transfer for Code RL</i>	
Zhaofeng Wu, Shiqi Wang, Boya Peng, Anuj Kumar Goyal, Melanie Kambadur, Sebastian Ruder, Yoon Kim and Chloe Bi	26583
<i>Distorted or Fabricated? A Survey on Hallucination in Video LLMs</i>	
Yiyang Huang, Yitian Zhang, Yizhou Wang, Mingyuan Zhang, Liang Shi, Huimin Zeng and Yun Fu	26599
<i>Generative Floor Plan Design with LLMs via Reinforcement Learning with Verifiable Rewards</i>	
Luis Lara, Aristides Milios, ZhiHao Luo, Aditya Sharma, Ge Ya Luo, Christopher Beckham, Florian Golemo and Christopher Pal	26612
<i>Creating Grammar Teaching Material for Endangered Languages with Hybrid Grammar Induction</i>	
Sebastien Christian	26628
<i>OpenExempt: A Diagnostic Benchmark for Legal Reasoning and a Framework for Creating Custom Benchmarks on Demand</i>	
Sergio Servantez, Sarah B. Lawsky, Rajiv Jain, Daniel W. Linna Jr. and Kristian J Hammond	26644
<i>Context-Conditioned Masked LoRA: Dynamic Rank Routing for Compute-Efficient Parameter-Efficient Fine-Tuning</i>	
Rifat Rafiuddin and Rafee Abdullah	26670
<i>A Survey on Evaluation of LLM-based Agents</i>	
Asaf Yehudai, Lilach Eden, Alan Li, Guy Uziel, Yilun Zhao, Roy Bar-Haim, Arman Cohan and Michal Shmueli-Scheuer	26690
<i>ProToM: Promoting Prosocial Behaviour via Theory of Mind-Informed Feedback</i>	
Matteo Bortoletto, Yichao Zhou, Lance Ying, Tianmin Shu and Andreas Bulling	26715
<i>JudgeMeNot: Personalizing Large Language Models to Emulate Judicial Reasoning in Hebrew</i>	
Itay Razumenko, Arnon Sturm and Nir Grinberg	26735
<i>Stabilizing Efficient Reasoning with Step-Level Advantage Selection</i>	
Han Wang, Xiaodong Yu, Jialian Wu, Jiang Liu, Ximeng Sun, Mohit Bansal and Zicheng Liu	26754
<i>RePromptT: Recurrent Prompt Tuning for Integrating Structured EHR Encoders with Large Language Models</i>	
Arya Hadizadeh Moghaddam, Drew Ross, Mohsen Nayebi Kerdabadi, Dongjie Wang and Zijun Yao	26766

<i>CLARO: Controlled Attribute-Driven Reasoning Optimization for Efficient Chain-of-Thought</i> Oded Schlesinger, Young Kyung Kim, J. Matias Di Martino and Guillermo Sapiro	26779
<i>Hidden States as Early Signals: Step-level Trace Evaluation and Pruning for Efficient Test-Time Scaling</i> Zhixiang Liang, Beichen Huang, Zheng Wang and Minjia Zhang	26800
<i>From Recall to Forgetting: Benchmarking Long-Term Memory for Personalized Agents</i> Md Nayem Uddin, Kumar Shubham, Eduardo Blanco, Chitta Baral and Gengyu Wang	26814
<i>Towards Inference-time Scaling for Continuous Space Reasoning</i> Minghan Wang, Thuy-Trang Vu, Ehsan Shareghi and Gholamreza Haffari	26842
<i>SynthFix: Adaptive Neuro-Symbolic Code Vulnerability Repair</i> Yifan Zhang, Jieyu Li, Kexin Pei, Yu Huang and Kevin Leach	26857
<i>Imperfectly Cooperative Human-AI Interactions: Comparing the Impacts of Human and AI Attributes in Simulated and User Studies</i> Myke C. Cohen, Mingqian Zheng, Neel Bhandari, Hsien-Te Kao, Xuhui Zhou, Daniel Nguyen, Laura Cassani, Maarten Sap and Svitlana Volkova	26871
<i>CheMM-R1: Enhancing Chemical Structure Recognition and Elucidation with Reasoning Multimodal Large Language Models</i> Liting Huang, Zhihao Zhang and Shoujin Wang	26902
<i>Agree, Disagree, Explain: Decomposing Human Label Variation in NLI through the Lens of Explanations</i> Pingjun Hong, Beiduo Chen, Siyao Peng, Marie-Catherine de Marneffe, Benjamin Roth and Barbara Plank	26922
<i>DeepSpecs: Expert-Level Question Answering in 5G</i> Aman Ganapathy Manvattira, Yifei Xu, Ziyue Dang and Songwu Lu	26935
<i>On-Policy Self-Distillation for Efficient Diffusion Language Models with Early-Stage Calibration</i> Huaisheng Zhu, MingYu Liu, Junze Liu, Zhen Ge, Tian Wang, Jiri Gesi, Dakuo Wang, Weiqi Zhang, Houyu Zhang, Yufan Guo, Xian Li, Bing Yin and Sujay Sanghavi	26954
<i>Arctic-Text2SQL-R1: Simple Rewards, Strong Reasoning in Text-to-SQL</i> Zhewei Yao, Guoheng Sun, Łukasz Borchmann, Zheyu Shen, Minghan Deng, Bohan Zhai, Hao Zhang, Ang Li and Yuxiong He	26966
<i>Do Emotions Influence Moral Judgment in Large Language Models?</i> Mohammad Saim and Tianyu Jiang	26996
<i>RShield: A User-level Traceable Backdoor Watermark for LLMs in Embedding-as-a-Service</i> Lingyun Xiang, Yufan Zhong, Chengfu Ou, Zhihua Xia, Chunfang Yang, Daojian Zeng and Zhangjie Fu	27014
<i>The Adaptive Interrogator: Detecting Trojan LLMs in Multi-Agent Systems via Evolved Conversational Strategies</i> Rana Muhammad Shahroz Khan, Ruichen Zhang, Zhen Tan, Charles Fleming and Tianlong Chen	27029
<i>SmartAD: Capacity-Aligned Agent Distillation for Small Language Models</i> Guokai Tang and Feng Zhao	27045
<i>TRM-Planner: Offline Target Planning and Distillation for Tiny Recursive Models</i> Euijin Baek, Housam Babiker, Mi-Young Kim and Randy Goebel	27058

<i>MAPLE: Multi-Aspect Panels of LLM Evaluators for Open-Ended Questions</i> Michinori Jinji, Kyohei Atarashi, Koh Takeuchi and Hisashi Kashima	27071
<i>MARCH: Evaluating the Intersection of Ambiguity Interpretation and Multi-hop Inference</i> Jeonghyun Park, Ingeol Baek, Seunghyun Yoon, Haeun Jang, Aparna Garimella, Akriti Jain, Nedim Lipka and Hwanhee Lee	27089
<i>Enhancing Multilingual RAG Systems with Debiased Language Preference-Guided Query Fusion</i> Jeonghyun Park, Byeongjeong Kim, Seojin Hwang and Hwanhee Lee	27116
<i>MT-OSC: Path for LLMs that Get Lost in Multi-Turn Conversation</i> Jyotika Singh, Fang Tu, Miguel Ballesteros, Weiyi Sun, Sandip Ghoshal, Michelle Yuan, Yassine Benajiba, Sujith Ravi and Dan Roth	27137
<i>Soft Head Selection for Injecting ICL-Derived Task Embeddings</i> Jungwon Park, Jimyeong Kim, Changin Choi and Wonjong Rhee	27161
<i>CORAL: Adaptive Retrieval Loop for Culturally-Aligned Multilingual RAG</i> Nayeon Lee, Jiwoo Song and Byeongcheol Kang	27215
<i>Safety Guardrails of Large Language Models Are Vulnerable to Value-Driven Adversarial Prompting</i> Xiaohao Luo, Ying Wei and Zhijun Li	27238
<i>Fair RAG: End-to-End Fairness Across Retrieval and Generation</i> Farsheed Haque, Zhe Fu, Ramit Aditya, Depeng Xu and Xi Niu	27256
<i>ManCC: A Task-Anchored Benchmark for Manchu–Classical Chinese Cross-Lingual Modeling</i> Meiqi Wang, Xiaoxin Sun, Dongjie Wang, Ruixin Yu, Xiantao Heng, Shuo Wang, Zhen Huang, Peng Zhao, Suhua Wang and Minghao Yin	27271
<i>Dynamic Graph Navigation via Triplet Chains for Structure-Aware Retrieval-Augmented Generation</i> Feng Zhao, Yufei Wu, Xianggan Liu and Ruilin Zhao	27293
<i>Style over Story: Measuring LLM Narrative Preferences via Structured Selection</i> Donghoon Jung, Jiwoo Choi, Songeun Chae and Seohyon Jung	27304
<i>Visual Interference in Speech Evaluation: Cultural Asymmetry and Cross-Modal Bias in MLLMs</i> Kyusik Kim, Hyunwoo Yoo, Jaehoon Choi, Gail Rosen and Bongwon Suh	27332
<i>DraDDP: A Multimodal Multi-Party Dialogue Discourse Parsing Dataset</i> Shannan Liu, Peifeng Li, Yaxin Fan and Qiaoming Zhu	27359
<i>Table-R1: Region-based Reinforcement Learning for Table Understanding</i> Zhenhe Wu, Jian Yang, Zhongjiang He, Changzai Pan, Jiaheng Liu, Xianjie Wu, Yu Zhao, Shuangyong Song, Yongxiang Li, Zhoujun Li and Xuelong Li	27374
<i>ARM2: Adaptive Reasoning Model with Vision Understanding and Executable Code</i> Jian Xie, Zhendong Chu, Aoxiao Zhong, Kai Zhang, Mingzhe Han, Xing Fan, Jialie Shen and Qingsong Wen	27392
<i>A Framework of Reflective Agents with Adaptive Collaboration for Attributed Summary Generation</i> Yu Chen, Peng Chen, Ziwei Zheng and Bang Wang	27406
<i>VLURes: Benchmarking Long-Text Grounding and Cross-Lingual Robustness in Vision Language Mod- els</i> Jesse Atuhurra, Iqra Ali, Tomoya Iwakura, Hidetaka Kamigaito and Tatsuya Hiraoka	27426

<i>Trustworthy and Explainable Causal Representation Learning in Transformers</i>	
Yang Liu, Yinghao Zhang, Lin Liu, Jiuyong Li, Debo Cheng and Zaiwen Feng	27482
<i>Affectron: Emotional Speech Synthesis with Affective and Contextually Aligned Nonverbal Vocalizations</i>	
Deok-Hyeon Cho, Hyung-Seok Oh, Seung-Bin Kim and Seong-Whan Lee	27502
<i>Reasoning-Guided Exploration for Online DPO</i>	
Zetian Hu, Shunyu Liu, Ting-En Lin, Fei Huang, Yongbin Li and Dacheng Tao	27526
<i>CAPA: Contribution-Aware Pruning and FFN Approximation for Efficient Large Vision-Language Models</i>	
Samyak Jha and Junho Kim	27543
<i>Strong Reasoning Isn't Enough: Evaluating Evidence Elicitation in Interactive Diagnosis</i>	
Zhuohan Long and Zhongyu Wei	27559
<i>IF-GEO: Conflict-Aware Instruction Fusion for Multi-Query Generative Engine Optimization</i>	
Heyang Zhou, Jiajia Chen, Xiaolu Chen, Jie Bao, Zhen Chen and Yong Liao	27576
<i>Investigating Human and LLMs' Decisions in Unverifiable Environments: A Case Study with GitHub Activity Overview</i>	
Zheng Jiang, Wei Wang, Gaowei Zhang, Yang Feng and Yi Wang	27591
<i>PlotGen-Bench: Evaluating VLMs on Generating Visualization Code from Diverse Plots across Multiple Libraries</i>	
Yi Zhao, Zhen Yang, Shuaiqi Duan, Wenmeng Yu, Zhe Su, Jibing Gong and Jie Tang	27619
<i>Chronos: Learning Temporal Dynamics of Reasoning Chains for Test-Time Scaling</i>	
Kai Zhang, Jiayi Liao, Chengpeng Li, Ziyuan Xie, Sihang Li and Xiang Wang	27651
<i>V-RoLoRA: RLVR-Driven MoE Routing for Steerable Pluralistic Alignment</i>	
Jing Wang, Yaomin Wu, Yinglin Wang and Yitong Yang	27665
<i>ParaCook: On Time-Efficient Planning for Multi-Agent Systems</i>	
Shiqi Zhang, Xinbei Ma, Yunqing Xu, Zouying Cao, Pengrui Lu, Haobo Yuan, Tiancheng Shen, Zhuosheng Zhang, Hai Zhao and Ming-Hsuan Yang	27683
<i>Multimodal Unlearning Across Vision, Language, Video, and Audio: Survey of Methods, Datasets, and Benchmarks</i>	
Nobin Sarwar, Shubhashis Roy Dipta, Zheyuan Liu and Vaidehi Patil	27702
<i>From Interpretability to Performance: Optimizing Retrieval Heads for Long-Context Language Models</i>	
Youmi Ma and Naoaki Okazaki	27731
<i>Bridging the Temporal Gap in Multimodal LLMs: Deeply Stacking Temporal Tokens for Audio-Visual Speech Recognition</i>	
Liyong Wang, Junliang Xing, Tianyu Hu, Jianfei Jiang, Jihuai Zhao and Huimin Ma	27748
<i>MHSafeEval: Role-Aware Interaction-Level Evaluation of Mental Health Safety in Large Language Models</i>	
Suhyun Lee, Palakorn Achananuparp, Neemesh Yadav, Ee-Peng Lim and Yang Deng	27760
<i>Enabling Stroke-Level Structural Analysis of Hieroglyphic Scripts without Language-Specific Priors</i>	
Fuwen Luo, Zihao Wan, Ziyue Wang, Yaluo Liu, Pau Tong Lin Xu, Xuanjia Qiao, Xiaolong Wang, Peng Li and Yang Liu	27794

<i>CODEPROMPTZIP: Code-specific Prompt Compression for Retrieval-Augmented Generation in Coding Tasks with LMs</i>	Pengfei He, Shaowei Wang and Tse-Hsun Chen	27811
<i>Detecting Hallucinations in Retrieval-Augmented Generation via Semantic-level Internal Reasoning Graph</i>	Jianpeng Hu, Yanzeng Li, Jialun Zhong, Lei Zou and Wenfa Qi	27826
<i>MeasHalu: Mitigation of Scientific Measurement Hallucinations for Large Language Models with Enhanced Reasoning</i>	Ruijun Huang, Zhiqiao Kang, Yuxuan Zhu, Junxiong Li, Jiahao Zhao, Minghuan Tan, Feng Jiang and Min Yang	27842
<i>CRAFTQA: A Code-Driven Adaptive Framework for Complex Structured Data Reasoning</i>	Chengtao Gan, Zhiqiang Liu, Long Jin, Yushan Zhu, Lei Liang and Wen Zhang	27860
<i>EMTIR-GRPO: Efficient Multi-Tool Augmented Large Language Models via Reinforcement Learning</i>	Shixin Jiang, Zhihao Zhu, Jiafeng Liang, Yang Wu, Ming Liu and Bing Qin	27877
<i>Multi-Persona Thinking for Bias Mitigation in Large Language Models</i>	Yuxing Chen, Guoqing Luo, Zijun Wu and Lili Mou	27895
<i>LCO: LLM-based Constraint Optimization for Safer Agentic LLMs in Real-world Tasks</i>	Jiayong Wan, Jiawei Chen, Zhaoxia Yin, Liu Shuyuan and Hang Su	27910
<i>MedKInstruct: A Multimodal Knowledge Graph Based Framework for Multi-Hop and Hard-Negative Instruction Data Synthesis in MedVQA</i>	Yinan Wu, Jihang Jin, Xuhao Bao, Weiyan Zhang, Hanjing Yan, Tong Ruan and ChunMing Wang	27935
<i>GeometryZero: Advancing Geometry Solving via Group Contrastive Policy Optimization</i>	Yikun Wang, Yibin Wang, Dianyi Wang, Zimian Peng, Qipeng Guo, Dacheng Tao and Jiaqi Wang	27948
<i>CollabCoder: Plan-Code Co-Evolution via Collaborative Decision-Making for Efficient Code Generation</i>	Duy Tung Doan, Quang Huy Phung, Dzung Nguyen and Khac-Hoai Nam Bui	27964
<i>Beyond the Safety Tax: Mitigating Unsafe Text-to-Image Generation via External Safety Rectification</i>	Xiangtao Meng, Yingkai Dong, Ning Yu, Li Wang, Zheng Li and Shanqing Guo	27986
<i>WenetSpeech-Wu: Datasets, Benchmarks, and Models for a Unified Chinese Wu Dialect Speech Processing Ecosystem</i>	Chengyou Wang, Mingchen Shao, Jingbin Hu, Zeyu Zhu, Hongfei Xue, Bingshen Mu, Xin Xu, Xingyi Duan, Binbin Zhang, Zhu Pengcheng, Chuang Ding, Xiaojun Zhang, Hui Bu and Lei Xie	27999
<i>ReaGeo: Reasoning-Enhanced End-to-End Geocoding with LLMs</i>	Jian Cui, Zhiyuan Ren, Desheng Weng, Yongqi Zhao, Gong Wenbin, Yu Lei and Zhenning Dong	28012
<i>EviReport: From Reasoned Outlines to Evidence Tracked Long-Form Reports</i>	Zihan Liu, Jianhui li, Zexin Wang, Fei Sun, Jingjing LI, Zheyuan Li, Ke Xiang, Hang Cui, Houhua Gong, Changhua Pei and Gaogang Xie	28024
<i>Preference Estimation via Opponent Modeling in Multi-Agent Negotiation</i>	Yuta Konishi, Kento Yamamoto, Eisuke Sonomoto, Rikuho Takeda, Ryo Furukawa, Yusuke Muraki, Takafumi Shimizu, Kazuma Fukumura, Yuya Kanemoto, Takayuki Ito and Shiyao Ding . . .	28049

<i>AutoSearch: Adaptive Search Depth for Efficient Agentic RAG via Reinforcement Learning</i>	
Jingbo Sun, Wenyue Chong, Songjun Tu, Qichao Zhang, Yaocheng Zhang, Jiajun Chai, Xiaohan Wang, Wei Lin, Guojun Yin and Dongbin Zhao	28059
<i>Beneficial Reasoning Behaviors in Agentic Search and Effective Training Methods to Obtain Them</i>	
Jiahe Jin, Abhijay Sai Paladugu and Chenyan Xiong	28080
<i>Beyond Prompt: Fine-grained Simulation of Cognitively Impaired Standardized Patients via Stochastic Steering</i>	
Weikang Zhang, Zimo Zhu, Zhichuan Yang, Chen Huang, Wenqiang Lei and See-Kiong Ng	28098
<i>Where Did It Go Wrong? Capability-Oriented Failure Attribution for Vision-and-Language Navigation Agents</i>	
Jianming Chen, Yawen Wang, Junjie Wang, Xiaofei Xie, Shoubin Li, Qing Wang and Fanjiang Xu	28132
<i>MultiFileTest: A Multi-File-Level LLM Unit Test Generation Benchmark and Impact of Error Fixing Mechanisms</i>	
Yibo Wang, Congying Xia, Wenting Zhao, Jiangshu Du, Chunyu Miao, Zhongfen Deng, Philip S. Yu and Chen Xing	28151
<i>SAMoRA: Semantic-Aware Mixture of LoRA Experts for Task-Adaptive Learning</i>	
Boyan Shi, Wei Chen, Shuyuan Zhao, Junfeng Shen, Shengnan Guo, Shaojiang Wang and Huaiyu Wan	28173
<i>DualAlign: Generating Clinically Grounded Synthetic Data</i>	
Rumeng Li, XWang and Hong yu	28189
<i>Breaking the Evaluation Paradox: Evaluating High-Entropy Search with Computationally Irreducible Constraints</i>	
Juntao Wu, Wei Wen, Xianting Huang, Shuai Pang, Ruizhi Qiao, Xing Sun and Ke Wang .	28209
<i>CoT-Edit: Reinforcement Learning of Chain-of-Thought Reasoning for Code Edit Suggestion</i>	
Wuya Chen, Yihao Yang and Yue Lin	28219
<i>MemORAI: Memory Organization and Retrieval via Adaptive Graph Intelligence for LLM Conversational Agents</i>	
Hung Pham Van, Nguyen Manh Hieu, Khang Pham Tran Tuan, Nam Le Hai, Linh Ngo Van, Nguyen Thi Ngoc Diep and Trung Le	28235
<i>MDC-Bench: A Multidisciplinary Causal Benchmark Based on Causal Structures for Evaluating Large Language Models</i>	
Peng Wang, Yuxiong Yan, Xiao Ding, Kai Xiong, Bibo Cai, Chao Peng, Yutai Hou, Dandan Tu, Bing Qin and Ting Liu	28254
<i>MathAgent: Adversarial Evolution of Constraint Graphs for Mathematical Reasoning Data Synthesis</i>	
Zixiong Yu, Jun Rao, Guhan Chen, Songtao Tian, Bohan Li, Jiansheng Wei, Min Zhang and Xiaojun Meng	28298
<i>Neural Induction of Finite-State Transducers</i>	
Michael Ginn, Alexis Palmer and Mans Hulden	28322
<i>Dynamic Sampling that Adapts: Self-Aware Iterative Data Persistent Optimization for Mathematical Reasoning</i>	
Jun Rao, Xuebo Liu, Hexuan Deng, Zepeng Lin, Zixiong Yu, Jiansheng Wei, Xiaojun Meng and Min Zhang	28337

<i>Forge: Quality-Aware Reinforcement Learning for NP-Hard Optimization in LLMs</i>	
Xiaozhe Li, Xinyu Fang, Shengyuan Ding, Yang Li, Linyang Li, Haodong Duan, Qingwen Liu and Kai Chen	28351
<i>Datasets for Scientific Literature Understanding: A Survey</i>	
Yuanzhe Zhang, Xun Zhao, Maodi Hu, Xi Sun, Donghuan Song and Zhixiong Zhang	28369
<i>ODASim: Ordered, Distinctive and Absolute Semantic Similarity for Code Explanation Evaluation</i>	
Prince Kumar, Vitobha Munigala, Jaydeep Sen, Ashish Mittal, Vishwajeet Kumar and Srikanth G. Tamilselvam	28390
<i>Almeyer-Oryx-BloomBench: A Bilingual Multimodal Benchmark for Cognitively Informed Evaluation of Vision-Language Models</i>	
Mohammad Mahdi Abootorabi, Omid Ghahroodi, Anas Madkoo, Marzia Nouri, Doratossadat Dastgheib and Ehsaneddin Asgari	28404
<i>OPT-BENCH: Evaluating the Iterative Self-Optimization of LLM Agents in Large-Scale Search Spaces</i>	
Xiaozhe Li, Jixuan Chen, Xinyu Fang, Shengyuan Ding, Haodong Duan, Qingwen Liu and Kai Chen	28437
<i>VocabTailor: Dynamic Vocabulary Selection for Downstream Tasks in Small Language Models</i>	
Hanling Zhang, Yayu Zhou, Tongcheng Fang, Zhihang Yuan, Guohao Dai, Wanli Ouyang and Yu Wang	28453
<i>Feedback Adaptation for Retrieval-Augmented Generation</i>	
Jihwan Bang, Seunghan Yang, Kyuhong Shim, Simyung Chang, Juntae Lee and Sungha Choi	28472
<i>From Implicit to Explicit: Token-Efficient Logical Supervision for Mathematical Reasoning in LLMs</i>	
Shaojie Wang and Liang Zhang	28486
<i>When Backdoors Go Beyond Triggers: Semantic Drift in Diffusion Models Under Encoder Attacks</i>	
Shenyang Chen and Liuwan Zhu	28501
<i>Reducing Peak Memory Usage for Modern Multimodal Large Language Model Pipelines</i>	
Junwan Kim and Hyunkyung Bae	28518
<i>MAKI: Multi-layer Aligned Knowledge Injection for Structure-aware Knowledge Graph Completion with Large Language Models</i>	
Zhiwen Xie, Xin Wang, Guangyou Zhou and Derek F. Wong	28527
<i>From Domains to Instances: Dual-Granularity Data Synthesis for LLM Unlearning</i>	
Xiaoyu Xu, Minxin Du, Zitong LI, Zi Liang, Zhibiao Guo, Zhang Shiyu, Peizhao Hu, Qingqing Ye and Haibo Hu	28540
<i>ERRV: Eliciting Efficient Reasoning through Reasoning Vectors for Policy Optimization in Large Language Models</i>	
Zhuowen Han, Lei Yang, Renren Jin, Dan Shi, Chenxi Sun and Deyi Xiong	28557
<i>LLM-induced Rationales for More Compact Explainable Style Classification Models</i>	
Ahmad Aljanaideh and Saeb Ganideh	28571
<i>MDTeamGPT: Mitigating Context Collapse and Enabling Self-Evolution in Medical Multi-Agent Reasoning</i>	
Kai Chen, Xinfeng Li, Tianpei Yang, Hewei Wang, Guang Yang, Jing Huo and Yang Gao .	28578

<i>TSPO: Breaking the Double Homogenization Dilemma in Multi-turn Search Policy Optimization</i>	
Shichao Ma, Zhiyuan Ma, Ming Yang, Xiaofan Li, Xing Wu, Jintao Du, Yu Cheng, Weiqiang Wang, Qiliang Liu, Zhengyang Zhou and Yang Wang	28607
<i>Towards semantic reliable clinical QA: Query pipeline optimization for cancer patient question answering systems</i>	
MaoLin He, Rena Wei Gao, Mike Conway and Brian E. Chapman	28624
<i>LayerNorm Induces Recency Bias in Transformer Decoders</i>	
Junu Kim, Xiao Liu, Zhenghao Lin, Lei Ji, Yeyun Gong and Edward Choi	28638
<i>Scaling Performance and Low-Resource Annotation with Many-Shot In-Context Learning for Named Entity Recognition</i>	
Qi Zhang, Fangping Lan, Cornelia Caragea, Longin Jan Latecki and Eduard Dragut	28653
<i>PhageBench: Can LLMs Understand Raw Bacteriophage Genomes?</i>	
Yusen Hou, Weicai Long, Haitao HU, Houcheng su, Junning Feng and Yanlin Zhang	28674
<i>POLAR: A Benchmark for Multilingual, Multicultural, and Multi-Event Online Polarization</i>	
Usman Naseem, Robert Geislinger, Juan Ren, Sarah Kohail, Rudy Alexandro Garrido Veliz, P Sam Sahil, Yiran Zhang, Idris Abdulmumin, Marco Antonio Stranisci, Özge Alacam, Cengiz Acarturk, Aisha Jabr, Saba Anwar, Abinew Ali Ayele, Simona Frenda, Alessandra Teresa Cignarella, Elena Tutubalina, Oleg Rogov, Aung Kyaw Htet, Xintong Wang, Surendrabikram Thapa, Kritesh Rauniyar, Tanmoy Chakraborty, MD Arfeen Zeeshan, Dheeraj Kodati, Satya Keerthi, Sahar Moradizyev, Firoj Alam, Md Arid Hasan, Syed Ishtiaque Ahmed, Ye Kyaw Thu, Shantipriya Parida, Ihsan Ayyub Qazi, Lilian Diana Awuor Wanzare, Nelson Odhiambo Onyango, Clemencia Siro, Jane Wanjiru Kimani, Ibrahim Said Ahmad, Adem Chanie Ali, Martin Semmann, Chris Biemann, Shamsuddeen Hassan Muhammad and Seid Muhie Yimam	28699
<i>VANE: Guiding High-Value Exploration in RLVR via Outcome-Process Novelty Shaping</i>	
Xu He, Jialiang Guo, Fucheng Xiong, Haodong Zhao, Xingyang li, Ke Zeng and Xunliang Cai	28721
<i>Exploring Coding Spot: Understanding Parametric Contributions to LLM Coding Performance</i>	
Dongjun Kim, Minhyuk Kim, Yongchan Chun, Chanjun Park and Heuseok Lim	28740
<i>FedGUI: Benchmarking Federated GUI Agents across Heterogeneous Platforms, Devices, and Operating Systems</i>	
WenHao Wang, Haoting Shi, Mengying Yuan, Yiquan Lin, Panrong Tong, Hanzhang Zhou, Guangyi Liu, Pengxiang Zhao, Yue Wang and Siheng Chen	28747
<i>Towards Proactive Information Probing: Customer Service Chatbots Harvesting Value from Conversation</i>	
Chen Huang, Zitan Jiang, Zou Changyi, Wenqiang Lei and See-Kiong Ng	28768
<i>Benchmarking Vision-Language Models on Chinese Ancient Documents: From OCR to Knowledge Reasoning</i>	
Haiyang Yu, Yuchuan Wu, Fan Shi, Jinghui Lu, Ke Niu, Xiaodong Ge, Minghan Zhuo, Jingqun Tang and Bin Li	28793
<i>DaMo: Data Mixing Optimizer in Fine-tuning Multimodal LLMs for Mobile Phone Agents</i>	
Kai Shi, Jun Yang, Ni Yang, Binqiang Pan, Qingsong Xie, Zhangchao, Zhenyu Yang, Tianhuang Su and Haonan Lu	28813
<i>From Outcome to Process: Optimizing MoE Load Balancing with MCTS</i>	
Wenjun Ke, Hengyuan Xu, Ziyu Shang, Yao He, Jiahao Wang, Zijie Xu, Peng Wang, Yuhang Lou and Jiajun Liu	28831

<i>Learning Optimal Message Representations for Agentic Communication</i>	
Shashwat Gupta, Anson Bastos, Mayukh Das, Supriyo Ghosh, Nagarajan Natarajan, Chetan Bansal and Saravan Rajmohan	28849
<i>CARO: Chain-of-Analogy Reasoning Optimization for Robust Content Moderation</i>	
Bingzhe Wu, Haotian Lu and Yuchen Mou	28880
<i>N-GRPO: Embedding-Level Neighbor Mixing for Enhanced Policy Optimization</i>	
Xukun Zhu, Hang Yu, Peng Di and Linchao Zhu	28893
<i>GFT: From Imitation to Reward Fine-Tuning with Unbiased Group Advantages and Dynamic Coefficient Rectification</i>	
Wangjie Gan, Miao Pan, Linbo Xi, Wenqi Zhang, Jintao Chen, Jianwei Yin and Xuhong Zhang	28909
<i>SciImpact: A Multi-Dimensional, Multi-Field Benchmark for Scientific Impact Prediction</i>	
Hangxiao Zhu, Yuyu Zhang, Ping Nie and Yu Zhang	28923
<i>When Helpers Become Hazards: A Benchmark for Analyzing Multimodal LLM-Powered Safety in Daily Life</i>	
Xinyue Lou, Xu Jinan, Jingyi Yin, Xiaolong Wang, Zhaolu Kang, Liaoyouwei, Yixuan Wang, Xiangyu Shi, Fengran Mo, SU Yao and Kaiyu Huang	28937
<i>MM-ShiftKV: Decode-Aware Prefill-Stage KV Selection for Multimodal Large Language Models</i>	
Jinsong Shu, Chenyang Wu, Zhongle Xie, Baokun Wang and Lidan Shou	28964
<i>Measuring What Matters Beyond Text: Evaluating Multimodal Summaries by Quality, Alignment, and Diversity</i>	
Abid Ali, Diego Molla and Usman Naseem	28983
<i>MonCulture-Eval: A Hierarchical Benchmark for Evaluating Mongolian Cultural Capabilities of Large Language Models across Scripts and Regions</i>	
Qulgan Minggad, Xiao Zinan and Yuan Sun	28997
<i>Less is More: Knowledge-Aware Compression for Long Legal Judgment Prediction</i>	
Fanghao Lou, Qiqi Wang, Guanyu Chen, Senbo Zhang, Kaiqi Zhao, Qian Liu and Huijia Li	29015
<i>Towards Visually Grounded Multimodal Summarization via Cross-Modal Transformer and Gated Attention</i>	
Abid Ali, Diego Molla and Usman Naseem	29032
<i>Fairness Evaluation and Inference Level Mitigation in LLMs</i>	
Afrozah Nadeem, Mark Dras and Usman Naseem	29048
<i>SCOPE: Preserving Modality-Specific Cues to Mitigate Modality Laziness in Multimodal Learning</i>	
Jingfan Yang, Rui Zhang, Liang Hong and Ke Yuan	29066
<i>Embedding-based In-Context Prompt Training for Enhancing LLMs as Text Encoders</i>	
Ailiang Lin, Zhuoyun Li, Keyu Mao, Kotaro Funakoshi and Manabu Okumura	29079
<i>Creating ConLangs to Probe the Metalinguistic Grammatical Knowledge of LLMs</i>	
Chihiro Taguchi and Richard Sproat	29096
<i>ReasonIF: Large Reasoning Models Fail to Follow Instructions During Reasoning</i>	
Yongchan Kwon, Shang Zhu, Federico Bianchi, Kaitlyn Zhou and James Zou	29149
<i>An LLM-Embedding Semantic Adaptation Network for Post-level Semantic Drift Evaluation</i>	
Ning Chen, Mingyu Kang, Jie Li and Linyuan Lü	29165

<i>CEAID: Benchmark of Multilingual Machine-Generated Text Detection Methods for Central European Languages</i>	
Dominik Macko and Jakub Kopál	29177
<i>Distilling the Essence, Discarding the Dross: Improving Fairness in Multimodal Large Language Models via Historical Reflection-Guided Prompt Optimization</i>	
Juncheng Hu, Jiming Yu, Rui Song, Kedi Lyu, Yingji Li and Zheli Liu	29191
<i>Tree-of-Evidence: Efficient System 2 Search for Faithful Multimodal Grounding</i>	
Micky C. Nnamdi, Benoit Louis Marteau, Yishan Zhong, J. Ben Tamo and May Dongmei Wang	29212
<i>Jailbreaking Large Language Models with Morality Attacks</i>	
Ying Su, Zheng Mingen, Weili Diao and Haoran Li	29228
<i>ExpSeek: Self-Triggered Experience Seeking for Web Agents</i>	
Wenyuan Zhang, Xinghua Zhang, Haiyang Yu, Shuaiyi Nie, Bingli Wu, Juwei Yue, Tingwen Liu and Yongbin Li	29255
<i>Do Personality Traits Interfere? Geometric Limitations of Steering in Large Language Models</i>	
Pranav Bhandari, Usman Naseem and Mehwish Nasim	29284
<i>AHA: Aligning Large Audio-Language Models for Reasoning Hallucinations via Counterfactual Hard Negatives</i>	
Yanxi Chen, Wenhui Zhu, Xiwen Chen, Zhipeng Wang, Xin Li, Peijie Qiu, Hao Wang, Xuanzhao Dong, Yujian Xiong, Anderson Schneider, Yuriy Nevmyvaka and Yalin Wang	29294
<i>Web Sitemap Knowledge Can Enhance Autonomous Browsing</i>	
Yuyao Zhang, Hongyu Lu, Jiajie Jin, Hongjin Qian, Shiyu Li, Zhao Yang, Yutao Zhu, Ji-Rong Wen and Zhicheng Dou	29307
<i>Coarse-to-Fine Multimodal Information Selection for Video Speaking Style Recognition with Large Language Models</i>	
Beibei Zhang, Yanan Lu, Lin Fen and Tongwei Ren	29322
<i>Large Language Models Require Curated Context for Reliable Political Fact-Checking—Even with Reasoning and Web Search</i>	
Matthew R. DeVerna, Kai-Cheng Yang, Harry Yaojun Yan and Filippo Menczer	29338
<i>IREASONER: Trajectory-Aware Intrinsic Reasoning Supervision for Self-Evolving Large Multimodal Models</i>	
Meghana Sunil, Manikandarajan Venmathimaran and Muthu Subash Kavitha	29361
<i>CLaRE-ty Amid Chaos: Quantifying Representational Entanglement to Predict Ripple Effects in LLM Editing</i>	
Manit Baser, Alperen Yildiz, Dinil Mon Divakaran and Mohan Gurusamy	29373
<i>STAR-Teaming: A Strategy-Response Multiplex Network Approach to Automated LLM Red Teaming</i>	
Min Jae Jung, YongTaek Lim, Chaeyun Kim, Junghwan Kim, Kihyun Kim and Minwoo Kim	29406
<i>Critic Rule Induction: Improving Temporal Knowledge Graph Forecasting with Generator-Critic Language Models</i>	
Yingsong Ning, Fu Zhang, Jingwei Cheng, Jiashun Peng and Xiaoke Wang	29436
<i>Context-Agent: Dynamic Discourse Trees for Non-Linear Dialogue</i>	
Junan Hu, Shudan Guo, Wenqi Liu, Jianhua Yin and Yinwei Wei	29449

<i>Can LLMs Really Judge? A Progressive Argumentation-Mining Framework for Distinguishing Understanding from Aggregation</i>	Fuyu Wang, Jiangtong Li, Kun Zhu and Changjun Jiang	29463
<i>SMART: Semantic Header Flattening and Pseudo-Code-Style Reasoning for LLM-based Complex Table Question Answering</i>	Yongshuo Zhang, Jixiong Chen, Kaihe xu, Wei Wei and Shihao Zou	29483
<i>ConMA : Confidence-Guided Kernel Sampling with Multi-Stage Aggregation for LLM Reasoning</i>	Yinuo Wang, Qingjie Li, Wenyao Cui, Qiuchi Li and Zhang Huaping	29500
<i>Human-Centered Supervision for Sentiment Analysis in Telugu: A Systematic Inquiry Beyond Accuracy</i>	Vallabhaneni Raj Kumar, Ashwin S, Supriya Manna, Niladri Sett, Cheedella V S N M S Hema Harshitha, Kurakula Harshitha, Basina Deepakraj, Anand Kumar Sharma, Tanuj Sarkar, Samanthapudi Shakeer and Bondada Navaneeth Krishna	29519
<i>Evidence-Aligned Entity Verification for Hallucination Detection in Retrieval-Augmented Generation</i>	Runsong Jia, Zhen Fang, Mengjia Wu, Jie Lu and Yi Zhang	29538
<i>From Personal to Collective: On the Role of Local and Global Knowledge in LLM Personalization</i>	Zehong Wang, Junlin Wu, Zhaoxuan Tan, Bolian Li, Xianrui Zhong, Zheli Liu and Qingkai Zeng	29555
<i>ProMediate: A Simulation Testbed for Evaluating Proactive Mediation in Multi-Party Negotiation</i>	Ziyi Liu, Bahareh Sarrafzadeh, Pei Zhou, Longqi Yang, Jieyu Zhao and Ashish Sharma . . .	29570
<i>InteracSPARQL : An Interactive System for SPARQL Query Refinement Using Natural Language Explanations</i>	Xiangru Jian, Zhengyuan Dong and M. Tamer Özsu	29599
<i>Are Large Language Models Reliable Reviewers? A Benchmark for Error Detection in Financial Documents</i>	Ying He, Zhouhong Gu, Zhecheng Hu, Yubo Zhou, Hao Shen, Jiaqing Liang, Zhaoqian Dai, Ma Shuguang, Fei Yu, Yanghua Xiao and Zhixu Li	29625
<i>Training-Free Test-Time Contrastive Learning for Large Language Models</i>	Kaiwen Zheng, Kai Zhou, Jinwu Hu, Te Gu, Mingkai Peng and Fei Liu	29644
<i>To Diff or Not to Diff? Structure-Aware and Adaptive Output Formats for Efficient LLM-based Code Editing</i>	Wei Cheng, Yongchang Cao, Chen Shen, Binhua Li, Jue Chen, Yongbin Li and Wei Hu . . .	29677
<i>CogEmp:A Cognitive Empathy-Oriented Dialogue System for Structured Psychological Counseling</i>	Yiguo Deng, Xia Lei, Yuan Zhang and Long Ye	29695
<i>LiveLongBench: Tackling Long-Context Understanding for Spoken Texts from Live Streams</i>	Yongxuan Wu, Runyu Chen, Peiyu Liu and Hongjin Qian	29713
<i>CPC-GRPO: Answer-Free Reinforcement Learning with Cross-Prompt Consensus Rewards</i>	Gyunyueop Kim and Sangwoo Kang	29733
<i>MASH: Evading Black-Box AI-Generated Text Detectors via Style Humanization</i>	Yongtong Gu, Songze Li and Xia Hu	29749
<i>MM-Doc-RI: Training Agents for Long Document Visual Question Answering through Multi-turn Reinforcement Learning</i>		

Jiahang Lin, Kai Hu, Binghai Wang, Yuhao Zhou, Zhiheng Xi, Honglin Guo, Shichun Liu, Junzhe Wang, Shihan Dou, Enyu Zhou, Hang Yan, Zhenhua Han, Tao Gui, Qi Zhang and Xuanjing Huang	29770
<i>Multimodal Chemical Structure-Text Coreference in Intellectual Property via Rule-guided Reinforcement Learning</i>	
Hanmeng Zhong, Wentao Wu, Linqing Chen and Peng Zhou	29784
<i>TheraAgent: Self-Improving Therapeutic Agent for Precise and Comprehensive Treatment Planning</i>	
Junkai Li, Yunghwei Lai, Tianyi Zhu, Zheng Long Lee, Weizhi Ma and Yang Liu	29797
<i>LearnAct: Few-Shot Mobile GUI Agent with a Unified Demonstration Benchmark</i>	
Guangyi Liu, Pengxiang Zhao, Liang Liu, Zhiming Chen, Yuxiang Chai, Yaozhen Liang, WenHao Wang, Siheng Chen, Zhengxi Lu, Shuai Ren, Hao Wang, Shibo He, Yong Liu and Wenchao Meng	29820
<i>From Local Perspective to Global Reasoning: A Neuro-Symbolic Framework for Zero-Shot Relation Extraction</i>	
Kailun Lyu, Fu Zhang, Zehan Li and Jingwei Cheng	29844
<i>DiFRa: A Unified Framework for Harmonizing Semantic Diversity and Factual Consistency in Question-Answer Generation</i>	
Zhenqin Li, ShengYong Ding and Shuangyin Li	29857
<i>Geoparsing: Diagram Parsing for Plane and Solid Geometry with a Unified Formal Language</i>	
Peijie Wang, Ming-Liang Zhang, Jun Cao, Chao Deng, Dekang Ran, Pi Bu, Hongda Sun, Xuan Zhang, Yingyao Wang, Jun Song, Bo Zheng, Fei Yin and Cheng-Lin Liu	29876
<i>Beyond Memorization: Testing LLM Reasoning on Unseen Theory of Computation Tasks</i>	
Shlok Shelat, Jay Raval, Souvik Roy and Manas Gaur	29904
<i>Beyond Dialogue Time: Temporal Semantic Memory for Personalized LLM Agents</i>	
Miao Su, Yucan Guo, Zhongni Hou, Long Bai, Zixuan Li, Yufei Zhang, Guojun Yin, Wei Lin, Xiaolong Jin, Jiafeng Guo and Xueqi Cheng	29935
<i>LongCLI-Bench: A Preliminary Benchmark and Study for Long-horizon Agentic Programming in Command-Line Interfaces</i>	
Yukang Feng, Jianwen Sun, Zelai Yang, Jiabin Ai, Chuanhao Li, Zizhen Li, Fanrui Zhang, Kang He, Rui Ma, Jifan Lin, Jie Sun, Yang Xiao, Sizhuo Zhou, Wenxiao Wu, Yiming Liu, Pengfei Liu, Shenglin Zhang and Kaipeng Zhang	29952
<i>Cognitive-Uncertainty Guided Knowledge Distillation for Accurate Classification of Student Misconceptions</i>	
Qirui Liu, Hao Chen, Weijie Shi, Jiajie Xu and Jia Zhu	29964
<i>Beyond Black-Box Interventions: Latent Probing for Faithful Retrieval-Augmented Generation</i>	
Linfeng Gao, Qinggang Zhang, Baolong Bi, Bo Zeng, Zheng Yuan, Zerui Chen, Zhimin Wei, Shenghua Liu, Linlong Xu, Longyue Wang, Weihua Luo and Jinsong Su	29981
<i>FalconCopilot: Empowering LLMs Towards Integrated Human-Machine Systems for Aviation Autonomy</i>	
Jingyuan Yan, Qingchen Liu, Qichao Ma and Jiahu Qin	30001
<i>Accelerating Training of Autoregressive Video Generation Models via Local Optimization with Representation Continuity</i>	
Yucheng Zhou and Jianbing Shen	30027

<i>RouteRAG: Efficient Retrieval-Augmented Generation from Text and Graph via Reinforcement Learning</i> Yucan Guo, Miao Su, Saiping Guan, Zihao Sun, Xiaolong Jin, Jiafeng Guo and Xueqi Cheng	30042
<i>HORIZON: A Benchmark for In-the-wild User Behaviour Modeling</i> Arnav Goel, Pranjal A Chitale, Bhawna Paliwal, Bishal Santra and Amit Sharma	30060
<i>Two-Stage Parameter Alignment for Multi-LoRA Merging in Large Language Models</i> Zijian Li, Xiachong Feng, Weitao Ma, Yichong Huang, Xiaocheng Feng and Bing Qin	30079
<i>The Agent’s First Day: Benchmarking Learning, Exploration, and Scheduling in the Workplace Scenarios</i> Daocheng Fu, Jianbiao Mei, Rong Wu, Xueming Yang, Jia Xu, Ding Wang, Pinlong Cai, Yong Liu, Licheng Wen and Botian Shi	30094
<i>Prefix: Understand and Adapt to User Preference in Human-Agent Interaction</i> Jialin Li, Zhenhao Chen, Hanjun Luo and Hanan Salam	30110
<i>How Do Answer Tokens Read Reasoning Traces? Self-Reading Patterns in Thinking LLMs for Quantitative Reasoning</i> Haoyang Chen, Yi Liu, Jianzhi Shao, Tao Zhang, Chengfu Huo and Wei Hu	30150
<i>MENTOR: Efficient Autoregressive Image Generation with Balanced Multimodal Control</i> Haozhe Zhao, Zefan Cai, Shuzheng Si, Liang Chen, Jiuxiang Gu, Wen Xiao, Minjia Zhang and Junjie Hu	30167
<i>Better and Worse with Scale: How Contextual Entrainment Diverges with Model Size</i> Dikshant Kukreja, Kshitij Sah, Gautam Gupta, Avinash Anand, Rajiv Ratn Shah, Zhengkui Wang, Aik Beng Ng and Erik Cambria	30194
<i>LLM-Guided Semantic Bootstrapping for Interpretable Text Classification with Tsetlin Machines</i> Jiechao Gao, Rohan Kumar Yadav, Yuangang Li, Yuandong Pan, Jie Wang, Ying Liu and Michael Lepech	30210
<i>Evaluating the Impact of Reviewer Guideline Design on LLM-Based Automated Peer Review</i> Haowen Li, Yoichi Ishibashi and Masafumi Oyamada	30223
<i>MI-CXR: A Benchmark for Longitudinal Reasoning over Multi-Interval Chest X-rays</i> Sunghwan Steve Cho, Yunseok Han and Jaeyoung Do	30241
<i>Relevance to Utility: Process-Supervised Rewrite for RAG</i> Jaeyoung Kim, Jongho Kim, Seung-won Hwang, Seoho Song and Young-In Song	30274
<i>On the Cultural Anachronism and Temporal Reasoning in Vision Language Models</i> Mukul Ranjan, Prince Jha, Khushboo Kumari and Zhiqiang Shen	30294
<i>GLA: Grounding Large Language Models in Molecular Hierarchy for Chemical Understanding</i> Yingxu Li, Jingjie Zeng, Zekun Wang, Hongfei Lin and Liang Yang	30311
<i>NeuRAG: End-to-End Neural Knowledge Augmentation via Hyper-Neurons</i> Liwei Zheng, Xuemin Liu and Jie Liu	30324
<i>DEAR: Distributional Error-Aware Reliability for Robust Multimodal Sentiment Analysis with Missing Modalities</i> Shihao Zou, Wei Wei and Yongshuo Zhang	30344
<i>OpenPhone: Mobile Agentic Foundation Models</i> Yangqin Jiang and Chao Huang	30362

<i>Fast Retrieval and Slow Reasoning for Explainable Multimodal Sentiment Analysis</i>	
Aoqiang Zhu, Min Hu and Yan Xing	30381
<i>Exploration-Exploitation Reshaping towards Efficient Reasoning for Large Language Models</i>	
Yufeng Shi, Weilin Luo, Yuxiang Zhang, Zongmeng Zhang, Haoyang Liu, Yubing Wang, Bin Wang, Wengang Zhou and Houqiang Li	30392
<i>COMPEL: Compensated Mixture-of-Experts Pruning with Expert-Layer distribution</i>	
Seohee Yoon and Yong Suk Choi	30408
<i>Towards Self-Evolving Agents: Enabling Autonomy through Interactive Experience Refinement</i>	
Cheng Yang, Xuemeng Yang, Licheng Wen, Daocheng Fu, Jianbiao Mei, Rong Wu, Pinlong Cai, Yufan Shen, Nianchen Deng, Jia Xu, Botian Shi, Yu Qiao and Haifeng Li	30424
<i>Good Reasoning Makes Good Demonstrations: Implicit Reasoning Quality Supervision via In-Context Reinforcement Learning</i>	
Tiehua Mei, Minxuan Lv, Leiyu Pan, Zhenpeng Su, Hongru Hou, Hengrui Chen, Ao Xu and Deqing Yang	30452
<i>Design First, Code Later: Aesthetically Pleasing Template-Free Slides Generation</i>	
Zhiyao Cui, Chenxu Wang, Shuyue Hu, Yiqun Zhang, Wenqi Shao, Qiaosheng Zhang and Zhen Wang	30470
<i>Benchmarking the Fine-Grained Discriminability in Image-Text Retrieval via Controlled Contrastive Differences</i>	
Zhen Wang, Xi Zhou, Yating Yang, Bo Ma, Lei Wang, Rui Dong, Azmat Anwar and Siru Miao	30491
<i>Learning on Imbalanced Noisy Data via Debaised Sample Selection and LLM-Driven Annotation</i>	
Bo Yuan, Yulin Chen and Yin Zhang	30504
<i>Beyond Task-Level Context: Class-Conditional Context Vectors for Implicit In-Context Learning</i>	
Jianxin Zhang, Yilu Hu, Teng Liu, Pei Guo and Juntao Li	30543
<i>DPN-LE: Dual Personality Neuron Localization and Editing for Large Language Models</i>	
Lifan Zheng, Xue Yang, Jiawei Chen, Chenyan WU, Jingyuan Zhang, Fanheng Kong, Xinyi Zeng, Xiang Chen and Yu Tian	30567
<i>Reasoning in a Combinatorial and Constrained World: Benchmarking LLMs on Natural-Language Combinatorial Optimization</i>	
Xia Jiang, Jing Chen, Cong Zhang, Jie Gao, Chengpeng Hu, Chenhao Zhang, Yaoxin Wu and Yingqian Zhang	30592
<i>DARL: Encouraging Diverse Answers for General Reasoning without Verifiers</i>	
Chongxuan Huang, Lei Lin, Xiaodong Shi, Wenping Hu and Ruiming Tang	30649
<i>Clustered Self-Assessment: A Simple yet Effective Method for Uncertainty Quantification in Large Language Models</i>	
Qi Cao, Takeshi Kojima, Andrew Gambardella, Helinyi Peng, Yutaka Matsuo and Yusuke Iwasawa	30666
<i>STEP: Success-Rate-Aware Trajectory-Efficient Policy Optimization</i>	
Yuhan Chen, Yuxuan Liu, Long Zhang, Pengzhi Gao, Jian Luan and Wei Liu	30681
<i>AscendKernelGen: LLM-Driven Kernel Generation for NPUs</i>	
Xinzi Cao, Jianyang Zhai, Pengfei Li, Zhiheng Hu, Cen Yan, Mubingxu, Guanghuan Fang, Bin She, Jiayu Li, Yihan Su, Dongyang Tao, Feidiao Yang, Chang-Dong Wang, Yutong Lu, Weicheng Xue, Bin Zhou and Yonghong Tian	30693

<i>Learning to Evolve: A Self-Improving Framework for Multi-Agent Systems via Textual Parameter Graph Optimization</i>	
Shan He, Runze Wang, Zhuoyun Du, Huiyu Bai, Zouying Cao, Yu Cheng and Bo Zheng . . .	30719
<i>Learning How to Remember: A Meta-Cognitive Management Method for Structured and Transferable Agent Memory</i>	
Sirui Liang, Pengfei Cao, Jian Zhao, Wenhao Teng, Xiangwen Liao, Jun Zhao and Kang Liu	30733
<i>LoRE: Enhancing Search Relevance with Progressive Chain-of-Thought and Preference Alignment</i>	
Chenji Lu, Zhuo Chen, Hui Zhao, Zhiyuan Zeng, Gang Zhao, Junjie Ren, Lihaoran, Songyan Liu, Pengjie Wang, Chuan Yu, Jian Xu and Bo Zheng	30754
<i>Multilingual Refusal Alignment for Safer Large Language Models</i>	
Aleksandra Krasnodębska, Wojciech Kusa and Aldo Lipani	30769
<i>View-RI: Asymmetric Policy Optimization for Difficulty-Aware Multimodal Reinforcement Learning</i>	
Minjie Hong, Zirun Guo, Jiabao Zhang, Zehan Wang, Ziang Zhang, Tao Jin and Zhou Zhao	30791
<i>PseudoGD: Enhancing Spatial Reasoning in Vision-Language Models through Pseudo Geometric Knowledge Distillation</i>	
Gwanghee Lee, Yeeun Choi and Kyoungson Jhang	30804
<i>DMSD: Dual-Modal Semantic Disentanglement for Compositional Zero-Shot Learning</i>	
Pan Yang, Jing Yang, Ruan Xiao li, Yuling Chen, Yuankai Wu, Quan Zhou and Xu Wang .	30815
<i>ReFreeKV: Towards Threshold-Free KV Cache Compression</i>	
Xuanfan Ni, Liyan Xu, Chenyang Lyu, Longyue Wang, Mo Yu, Lemao Liu, Fandong Meng, Jie Zhou and Piji Li	30827
<i>Safety Sidecar: Reflection-Driven Runtime Control for Safer Agents</i>	
Wang Bin, Quan Jiazheng, Xingrui Yu, Hu Hansen, Yu Hao, Anjun Gao, Zhenglin Wan, Hui LI and Ivor Tsang	30842
<i>MCGA: A Multi-task Classical Chinese Literary Genre Audio Corpus</i>	
Yexing Du, Kaiyuan Liu, Bihe Zhang, Youcheng Pan, Bo Yang, Liangyu Huo, Xiyuan Zhang, Jian Xie, Daojing He, Yang Xiang, Ming Liu and Bing Qin	30857
<i>VET: Verifiable Execution Tracing for Reliable Text-to-SQL Generation</i>	
Dongyu Wang, Jingyu Li, Lan Zhang, Ganggang.yu and Liang Huang	30867
<i>PACE: Prefix-Protected and Difficulty-Aware Compression for Efficient Reasoning</i>	
Ruixiang Feng, Yuntao Wen, Silin Zhou, Ke Shi, Yifan Wang, Ran Le, Zhenwei An, Zongchao Chen, Chen Yang, Guangyue Peng, Yiming Jia, Dongsheng Wang, Tao Zhang, Lisi Chen, Yang Song, Shen Gao and Shuo Shang	30884
<i>CreativeBench: Benchmarking and Enhancing Machine Creativity via Self-Evolving Challenges</i>	
Zihan Wang, Lam Nguyen, Zhengyang Zhao, Mengyue Yang, Chengwei Qin, Yujiu Yang and Linyi Yang	30904
<i>Machine Behavior in Relational Moral Dilemmas: Moral Rightness, Predicted Human Behavior, and Model Decisions</i>	
Jiseon Kim, Jea Kwon, Luiz Felipe Vecchietti, Wenchao Dong, Jaehong Kim and Meeyoung Cha	30938
<i>Prompting the Unknown: Understanding Response Uncertainty in Large Language Models</i>	
Ze Yu Zhang, Arun Verma, Finale Doshi-Velez and Bryan Kian Hsiang Low	30956

<i>Explainable Quantum Program Repair with Verifiable Proof Traces</i>	
Tingting Li, Ziming Zhao, Zhaoxuan Li, Jiongchi Yu, Xiaofei Yue and Jianwei Yin	30985
<i>MA²P: A Meta-Cognitive Autonomous Intelligent Agents Framework for Complex Persuasion</i>	
Dingyi Zhang, Ziqing Zhuang, Linhai Zhang, Ziyang Gao and Deyu Zhou	30996
<i>Synergizing Semantic Anchors and Ordinal Smoothed Cross-Entropy for Speech Fluency Classification</i>	
Mulati Kahaer, Sirajahmat Ruzmamat, XuDong Pang, Subinuer Maimaituexun, Zaokere Kadeer, Abudurexiti Rehemam, Wenwen Lu, Panpan Zheng and Aishan Wumaier	31018
<i>PAC-BENCH: Evaluating Multi-Agent Collaboration under Privacy Constraints</i>	
Minjun Park, Donghyun Kim, Hyeonjong Ju, Seungwon Lim, Dongwook Choi, Taeyoon Kwon, Minju Kim and Jinyoung Yeo	31030
<i>SciText2Eq: Assessing LLMs for Explainable Equation Generation for Scientific Creativity</i>	
Yifan Mo, Xiao Fu, Yue Su, Qingyu Meng, Koen Hindriks, Qingzhi Liu and Jiahuan Pei	31057
<i>Reward Yourself: Efficient Self Rewards for Trustworthy Sampling</i>	
Mingjie Li, Wai Man Si, Michael Backes and Yang Zhang	31070
<i>Analyze Like a Venture Capitalist: Information-Gain and Knowledge Enhanced Graph Reasoning for Startup Success Prediction</i>	
Haoyu Pei, Zhongyang Liu, Xiangyi Xiao, Xiaocong Du, Suting Hong, Kunpeng Zhang and Haipeng Zhang	31087
<i>Read the Room, Read the Image: Understanding Indirect Speech Acts in Multimodal Visual Contexts</i>	
Jaehee Kim, Ji Hoon Chung, Seoyoon Park, Unsol Kim, Kyungwon Park, JiHak Kim, Yi-Jun Chen and Hansaem Kim	31109
<i>Break Through the Compression Bottleneck: From Theory to Practice</i>	
Xiusheng Huang, Lu Wang, Yequan Wang, Jun Zhao and Kang Liu	31125
<i>Cognitive Analysis Graph-Guided Multi-Turn Safety Enhancement for Large Language Models</i>	
Lanxue Zhang, Yuqiang Xie, Fang Fang, Yubing Ren, Xuebin Wang and Yanan Cao	31143
<i>Entropy Ratio Clipping as a Soft Global Constraint for Stable Reinforcement Learning</i>	
Zhenpeng Su, Leiyu Pan, Minxuan Lv, Tiehua Mei, Zijia Lin, Yuntao Li, Wenping Hu, Ruiming Tang, Kun Gai and Guorui Zhou	31161
<i>BCL: Bayesian In-Context Learning Framework for Information Extraction</i>	
Haoliang Liu, Chengkun Cai, Xu Zhao, Han Zhu, Shizhou Huang, Xinglin Zhang, Tao Chen, Jenq-Neng Hwang, Zhang Huaping and Lei Li	31172
<i>G-HiRel: Enhancing the Adaption to Knowledge Updating for Large Language Model Reasoning</i>	
Yudai Pan, Jiajie Hong, Tianzhe Zhao, Lingyun Song, Lingling Zhang, Yixin Chen and Xuequn Shang	31190
<i>WoW-Bench: Evaluating Fine-Grained Acoustic Perception in Audio-Language Models via Marine Mammal Vocalizations</i>	
Jaeyeon Kim, Heeseung Yun, Tony Woo, Chao-Han Huck Yang and Gunhee Kim	31208
<i>AdaMARP: An Adaptive Multi-Agent Interaction Framework for General Immersive Role-Playing</i>	
Zhenhua Xu, Dongsheng Chen, Shuo Wang, Jian Li, Chengjie Wang, Meng Han and Yabiao Wang	31230
<i>Train in Vain: Functionality-Preserving Poisoning to Prevent Unauthorized Use of Code Datasets</i>	
Yuan Xiao, Jiaming Wang, Yuchen Chen, Wei Song, Jun Sun, Shiqing Ma, Yanzhou Mu, Juan Zhai, Chunrong Fang, Jin Song Dong and Zhenyu Chen	31284

<i>Revisiting the Uniform Information Density Hypothesis in LLM Reasoning</i> Minju Gwak, Guijin Son and Jaehyung Kim	31304
<i>pQuant: Towards Effective Low-Bit Language Models via Decoupled Linear Quantization-Aware Training</i> Wenzheng Zhang and Bingzheng Liu	31334
<i>Incomplete Prompt Jailbreaks in Large Language Models</i> Yeonjea Kim, Bumjin Park and Jaesik Choi	31352
<i>The Best of Both Worlds: Combining Parallel and Sequential Inference Scaling via Aggregation Fine-Tuning</i> Yafu Li, Zhilin Wang, Tingchen Fu, Ganqu Cui, Sen Yang and Yu Cheng	31369
<i>A.S.E: A Repository-Level Benchmark for Evaluating Security in AI-Generated Code</i> Keke Lian, Wang Bin, Lei Zhang, Libo Chen, Junjie Wang, Ziming Zhao, Yujiu Yang, Miaoqian Lin, Haotong Duan, Haoran Zhao, Shuang Liao, Mingda Guo, Quan Jiazheng, Yilu Zhong, Chenhao He, Chen Zichuan, Jie Wu, Haoling Li, Zhaoxuan Li, Jiongchi Yu, Hui LI and Dong Zhang	31390
<i>Large Language Models Are Overconfident in Their Own Responses</i> Mario Sanz-Guerrero, Manuel Mager and Katharina Von Der Wense	31406
<i>Structured Dialogue Refinement: Building Retrieval-Augmented Question Answering on Goal-Oriented Dialogues</i> Bin Wu, Sawan Kumar, Prasetya Ajie Utama and Mohamed Yahya	31419
<i>Thought-Action Graph Reasoning: Faithful and Efficient Reasoning of Large Language Models via Reusing Past Experience</i> Zhixiao Qi, Feng Huang, Yunqi Zhang, Shijie Zhang, Qingqing Sun, Yongfeng Huang, Minghu Jiang, Shuai Chen and Tianyi Zhang	31433
<i>XLSR-MamBo: Scaling the Hybrid Mamba-Attention Backbone for Audio Deepfake Detection</i> Kwok-Ho Ng, Tingting Song, Yongdong WU and Zhihua Xia	31450
<i>Anchored Cyclic Generation: A Novel Paradigm for Long-Sequence Symbolic Music Generation</i> Boyu Cao, Lekai Qian, Dehan Li, Haoyu Gu, Mingda Xu and Qi Liu	31463
<i>TRUST: Towards Robust Social Bot Detection via Uncertainty-Guided Pseudo-Labeling and Graph Structure Purification</i> Ruixuan Xu, Mengting Hu, Zhunheng Wang, Ming Jiang, Rui Ying, Zhen Zhang, Hang Gao, Shuaipeng Liu and Renhong Cheng	31477
<i>XGUARD: A Graded Benchmark for Evaluating Safety Failures of Large Language Models on Extremist Content</i> Vadivel Abishethvarman, Bhavik Chandna, Pratik Jalan and Usman Naseem	31492
<i>PiCSAR: Probabilistic Confidence Selection and Ranking for Reasoning Chains</i> Joshua Ong Jun Leang, Zheng Zhao, Aryo Pradipta Gema, Sohee Yang, Wai-Chung Kwan, Xuanli He, Wenda Li, Pasquale Minervini, Eleonora Giunchiglia and Shay B Cohen	31511
<i>DeepPresenter: Environment-Grounded Reflection for Agentic Presentation Generation</i> Hao Zheng, Guozhao Mo, Xinru Yan, Qianhao Yuan, Wenkai Zhang, Xuanang Chen, Yaojie Lu, Hongyu Lin, Xianpei Han and Le Sun	31545
<i>An Iterative Utility Judgment Framework Inspired by Philosophical Relevance via LLMs</i> Hengran Zhang, Keping Bi, Jiafeng Guo and Xueqi Cheng	31559

<i>FormulaReasoning: A Dataset for Formula-Based Numerical Reasoning</i>	
Xiao Li, Bolin Zhu, Kaiwen Shi, Sichen Liu, Yin Zhu, Yiwei Liu and Gong Cheng	31579
<i>CORD: Bridging the Audio–Text Reasoning Gap via Weighted On-policy Cross-modal Distillation</i>	
Hu Jing, Danxiang Zhu, Xianlong Luo, Dan Zhang, Shuwei He, Yishu Lei, Shikun Feng, Hai-Tao Zheng, Jingzhou HE, Yu Sun, Hua Wu and Haifeng Wang	31602
<i>PlanE: Meta Planning of Data, Tuning, and Inference for Extractive-based LLMs</i>	
Jiacheng Wang, Weiyan Zhang and Guangya Yu	31613
<i>Generalizable Prompt Tuning for Audio-Language Models via Semantic Expansion</i>	
Jaehyuk Jang, Wonjun Lee, Kangwook Ko and Changick Kim	31636
<i>Leveraging Label Semantics and Entity Description Generation for LLM-based Fine-grained Entity Typing</i>	
Yingying Ma, Hongliang Dai, Xia Zhang and Piji Li	31655
<i>Trident: Self-Supervised Preference Alignment via Triplet Regularization</i>	
Yingnan Guo, Kejia Chen, Xiaofeng Zhang, Zifei Wu and Yu Zhang	31668
<i>Why Do Multilingual Reasoning Gaps Emerge in Reasoning Language Models?</i>	
Deokhyung Kang, Seonjeong Hwang, Daehui Kim, Hyounghun Kim and Gary Lee	31684
<i>EvoMemKG: An Evolvable Memory Agent for Multi-hop Knowledge Graph Reasoning</i>	
Shiyu Tian, Shuyue Xing, Zhuoxin Han, Caixia Yuan and Xiaojie Wang	31717
<i>Scaling Laws or Threshold Effects: Exploring the Optimal Vocabulary Size for Balancing Performance and Efficiency in Low-Resource Languages</i>	
Ao Han, Andong Chen, Yuan Sun and Xiaobing Zhao	31741
<i>RouterHGC: Optimized Router for LLM-based Multi-Agent Systems via Heterogeneous Graph Contrastive Learning</i>	
Yitao Xiao, Shaoyong Guo, Guoming Yang, Qingnan Wang, Yinlin Ren, Xuesong Qiu and Qi Feng	31759
<i>Grouped Adaptive Weight Sharing (GAWS): An Inference-Efficient Adaptation Method for Large Language Models</i>	
Eman Alsuradi, Junhyun Lee, Kyenghun Lee, Hyeonmok Ko and Fahed Jubair	31790
<i>Lending Eyesight to Language Models: Modeling and Probing Human scanpath through Transformer Decoder</i>	
Junlin Li, David Robert Reich and Yu-Yin Hsu	31807
<i>Timesteps of Mamba Align with Human Reading Times</i>	
Yuji Yamamoto, Shinnosuke Isono, Yoshinobu Kawahara and Sho Yokoi	31820
<i>Astra: Activation-Space Tail-Eigenvector Low-Rank Adaptation of Large Language Models</i>	
Kainan Liu, Yong Zhang, Ning Cheng, Yun Zhu, Yanmeng Wang, Shaojun Wang and Jing Xiao	31833
<i>AROMA: Augmented Reasoning Over a Multimodal Architecture for Virtual Cell Genetic Perturbation Modeling</i>	
Wang Zhenyu, Geyan Ye, Wei Liu and Man Tat Alexander Ng	31855
<i>Exploring Two-Phase Continual Instruction Fine-tuning for Multilingual Adaptation in Large Language Models</i>	
Divyanshu Aggarwal, Sankarshan Damle, Navin Goyal, Satya Lokam and Sunayana Sitaram	31882

<i>SCALER: Synthetic Scalable Adaptive Learning Environment for Reasoning</i> Caijun Xu, Changyi Xiao, Zhongyuan Peng, Xinrun Wang and Yixin Cao	31905
<i>Probing Social Identity Bias in Chinese LLMs with Gendered Pronouns and Social Groups</i> Geng Liu, Li Feng, Junjie Mu, Mengxiao Zhu and Francesco Pierri	31924
<i>AutoVecCoder: Teaching LLMs to Generate Explicitly Vectorized Code</i> ShangZhan Li, Xinyu Yin, Xuanyu Jin, Ye He, Yuxin Zhou, Yuxuan Li, Xu Han, Wanxiang Che, Qi Shi, Ting Liu and Maosong Sun	31942
<i>CompassLLM: A Multi-Agent Approach toward Geo-Spatial Reasoning for Popular Path Query</i> Md. Nazmul Islam Ananto, Shamit Fatin, Mohammed Eunus Ali and Md Rizwan Parvez .	31960
<i>Free-MAD: Consensus-Free Multi-Agent Debate</i> Yu Cui, Hang Fu, Haibin Zhang, Licheng Wang and Cong Zuo	31977
<i>Dialectical Structured Reasoning for Explainable Multimodal Fake News Detection</i> Ruichao Yang, Yufan Bian, Wei Gao, Bo-Wen Zhang, Jing Ma, Hongzhan Lin, Ziyang Luo, Xiaobin Zhu and Xu-Cheng Yin	31998
<i>Unlocking Human-Like Visible Logic: How Logic Diagrams Boost Logic Reasoning in Large Language Models?</i> Xin Wu, Yuqi Bu, Mengchen Zhao, Qingbao Huang and Yi Cai	32014
<i>Let Retrievers Think Before Action: Thought-Augmented Embedding for Dense Retrieval</i> Ruiran Yan, Wen Xiong, Ze Liu, Chaozhuo Li, Hao Liao, Defu Lian and Zheng Liu	32032
<i>TravelBehaviorQA: A Benchmark Dataset for Behavioral Interpretation of GPS Trajectories</i> Dongyang Zhen, Niping Duan, Huan Zhou and Qingbin Cui	32053
<i>SpiralThinker: Latent Reasoning through an Iterative Process with Text–Latent Interleaving</i> Shengmin Piao and Sanghyun Park	32072
<i>Verifier-Free RL for LLMs via Intrinsic Gradient-Norm Reward</i> Xuexiang Wen, Hang Yu, Linchao Zhu and Gaoang Wang	32089
<i>NovBench: Evaluating Large Language Models on Academic Paper Novelty Assessment</i> Wenqing Wu, Yi Zhao, Yuzhuo Wang, Siyou Li, Juexi Shao, Yunfei Long and Chengzhi Zhang	32103
<i>LaMPE: Length-aware Multi-grained Positional Encoding for Adaptive Long-context Scaling Without Training</i> Sikui Zhang, Guangze Gao, Ziyun Gan, Chunfeng Yuan, Zefeng Lin, Houwen Peng, Bing Li and Weiming Hu	32134
<i>Can Small Vision–Language Models Perform Sign Language Translation?</i> Anal Roy Chowdhury and Debarshi Kumar Sanyal	32150
<i>FastKV: Decoupling of Context Reduction and KV Cache Compression for Prefill-Decoding Acceleration</i> Dongwon Jo, Jiwon Song, Yulhwa Kim and Jae-Joon Kim	32167
<i>Beyond Outcome Verification: Verifiable Process Reward Models for Structured Reasoning</i> Massimiliano Pronesti, Anya Belz and Yufang Hou	32187
<i>Auto-Stega: An Agent-Driven System for Lifelong Strategy Evolution in LLM-Based Text Steganography</i> Jiuan Zhou, Yu Cheng, Yuan Xie and Zhaoxia Yin	32203

<i>Detecting AI-Generated Video: A Vision–Language Dual-View Survey</i>	
Dylan Xinming Hou, Juntian Zhang, Xu Gu, Yichen Wu, Nils Lukas, Gus Xia, Xiuying Chen and Yuhan Liu	32221
<i>UnIte: Uncertainty-based Iterative Document Sampling for Domain Adaptation in Information Retrieval</i>	
Jongyoon Kim, Minseong Hwang and Seung-won Hwang	32256
<i>Improving LLM Code Reasoning via Semantic Equivalence Self-Play with Formal Verification</i>	
Poon Tsz Nok and Antonio Valerio Miceli Barone	32273
<i>Towards Identification and Intervention of Safety-Critical Parameters in Large Language Models</i>	
Weiwei Qi, Zefeng Wu, Tianhang Zheng, Zikang Zhang, Xiaojun Jia, Zhan Qin and Kui Ren	32293
<i>OPINE: A Prior-calibrated Scoring Framework for LLM-based Multi-label Scientific Opinion Classification</i>	
Mengting Zhang, Gaofeng Pan, Zhixiong Zhang, Yang Li and Guangyin Zhang	32313
<i>Vocab Diet: Reshaping the Vocabulary of LLMs via Vector Arithmetic</i>	
Yuval Reif, Guy Kaplan and Roy Schwartz	32334
<i>RecMem: Recurrence-based Memory Consolidation for Efficient and Effective Long-Running LLM Agents</i>	
Zijie Dai, Shiyuan Deng, Sheng Guan, Yizhou Tian, Xin Yao, Xiao Yan and James Cheng	32353
<i>Let the Comments Speak: A Multi-Agent Framework based on Large Language Model for Comment-Guided Code Refactoring</i>	
Zixuan Wang, Wutong Yu and Deyu Zhou	32377
<i>AdapTime: Enabling Adaptive Temporal Reasoning in Large Language Models</i>	
Yimin Deng, Yejing Wang, Zhenxi Lin, Zichuan Fu, Guoshuai Zhao, Derong Xu, Yefeng Zheng, Xiangyu Zhao, Xian Wu, Li Zhu and Xueming Qian	32394
<i>Hierarchical Representation Alignment Learning of Diffusion Transformers for Neural Audio Codec</i>	
Sang-Hoon Lee and Ha-Yeong Choi	32410
<i>CoV: Chain-of-View Prompting for Spatial Reasoning</i>	
Haoyu Zhao, Akide Liu, Zeyu Zhang, Weijie Wang, Feng Chen, Ruihan Zhu, Gholamreza Haffari and Bohan Zhuang	32427
<i>A Graph Talks, But Who’s Listening? Rethinking Evaluations for Graph-Language Models</i>	
Soham Petkar, Hari Aakash K, Anirudh Vempati, Akshit Sinha, Ponnurangam Kumaraguru and Chirag Agarwal	32441
<i>Rejection-to-Acceptance Transition: Model Editing-Based Jailbreak Backdoor Injection Not Limited to Few Output Tokens</i>	
Shiji Yang, Min Cai, Hao Xiong, Congyao Mei, Haodong Zou, Shicheng Tan, Jie Chen, Fulan Qian and Shu Zhao	32463
<i>Causal Evidence Extraction and Triangulation in Crisis Reports using Large Language Models: A ReliefWeb-based Study</i>	
Yuanjun Zhang and Mourad Oussalah	32478
<i>Sounding vs. Being an Expert: Disentangling Authority, Register and Cultural Impact in Sycophantic LLMs</i>	
Gabriele Maraia, Fabio Massimo Zanzotto and Leonardo Ranaldi	32492

<i>Datamart-Agent: LLM-Driven Game-Theoretic Agent for Data Marketplace Modeling</i> Pangjing Wu, Peter Q. Chen, Xiaodong Li, Wenqi Fan and Li Qing	32509
<i>Multi-Drafter Speculative Decoding with Alignment Feedback</i> Taehyeon Kim, Hojung Jung and Se-Young Yun	32532
<i>Probabilistic Depression Detection from Textual Time Series</i> Fabian Schmidt, Seyedehtoniba Ravan and Vladimir Vlassov	32574
<i>Distilling Examples into Task Instructions: Enhanced In-Context Learning for Real-World B2B Conversations</i> Guy Rotman, Adi Kopilov, Danit Berger Zalmanson and Omri Allouche	32590
<i>Multilingual Tokenization through the Lens of Indian Languages: Challenges and Insights</i> Maharaj Brahma, N J Karthika, Rajat Verma, Nagasai Saketh Naidu, Rohit Saluja, Maunendra Sankar Desarkar and Ganesh Ramakrishnan	32614
<i>TrimTokenator: Towards Adaptive Visual Token Pruning for Large Multimodal Models</i> Hao Zhang, Lyu Mengsi, Chenrui He, Yulong Ao and Yonghua Lin	32633
<i>TTVS: Boosting Self-Exploring Reinforcement Learning via Test-time Variational Synthesis</i> Sikai Bai, Haoxi Li, Jie Zhang, Yongjiang Liu and Song Guo	32651
<i>Parallel Context-of-Experts Decoding for Retrieval Augmented Generation</i> Giulio Corallo and Paolo Papotti	32666
<i>LLaTiSA: Towards Difficulty-Stratified Time Series Reasoning from Visual Perception to Semantics</i> Yueyang Ding, HaoPeng Zhang, Rui Dai, Yi Wang, Tianyu Zong, Kaikui Liu and Xiangxiang Chu	32677
<i>NeedleChain: Measuring Intact Context Comprehension Capability of Large Language Models</i> Hyeonseok Moon and Heuseok Lim	32718
<i>From Scores to Preferences: Redefining Evaluation Paradigm for Speech Quality Reward Modeling</i> Yifei Cao, Changhao Jiang, Jiabao Zhuang, Jiajun Sun, Ming Zhang, Zhiheng Xi, Hui Li, Shihan Dou, Yuran Wang, Yunke Zhang, Tao Ji, Tao Gui, Qi Zhang and Xuanjing Huang	32731
<i>FIND: Toward Multimodal Financial Reasoning and Question Answering for Indic Languages</i> Sarmistha Das, Vaibhav Vishal, Syed Ibrahim Ahmad, Sriparna Saha and Manish Gupta ..	32750
<i>Think Smart, Not Hard: Difficulty Adaptive Reasoning for Large Audio Language Models</i> Zhichao Sheng, Shilin Zhou, Chen Gong and Zhenghua Li	32771
<i>Can Large Language Models Infer Human Actions and Motives? Evaluation in Social Prediction and Inspection Games</i> Kaleen Shrestha, Abhinav Gupta, Harish Dukkupati, Zhonghao Shi and Maja Mataric	32791
<i>Self-Consistency from Only Two Samples: CoT-PoT Ensembling for Efficient LLM Reasoning</i> Raman Saparkhan, Majd Hawasly, Md Rizwan Parvez and Mohammad Raza	32804
<i>ZoomRAG: Hierarchical Random-walk Zooming across Multi-scale Information Graphs for Fast and Accurate RAG</i> Xianming Hu, Jingyang Chen, Bin Tang, Yihe Liu, Yihong Huang, Hongbo Zhao, Nuoyi Chen, Jie Zhang, Ping Li and Kai Zhang	32840
<i>Training Language Models to Use Prolog as a Tool</i> Niklas Mellgren, Peter Schneider-Kamp and Lukas Galke Poch	32860

<i>Frame of Reference: Addressing the Challenges of Common Ground Representation in Situational Dialogs</i>	
Biswesh Mohapatra, Théo Charlot, Giovanni Duca, Mayank Palan, Laurent Romary and Justine Cassell.....	32881
<i>MultiDx: A Multi-Source Knowledge Integration Framework towards Diagnostic Reasoning</i>	
Yimin Deng, Zhenxi Lin, Yejing Wang, Guoshuai Zhao, Pengyue Jia, Zichuan Fu, Derong Xu, Yefeng Zheng, Xiangyu Zhao, Li Zhu, Xian Wu and Xueming Qian.....	32904
<i>Training Verifier to Assessing Complex Real-World Tool-Use Trajectories</i>	
Linzhuang Sun, Mingyang Chen, Hao Liang, Tianpeng Li, Zhou Yijie, Chenzheng Zhu, Tianyu Guo, Huanyao Zhang, Jingxuan Wei, Bihui Yu, Fan Yang and Wentao Zhang.....	32922
<i>HyperAdaLoRA: Accelerating LoRA Rank Allocation During Training via Hypernetworks without Sacrificing Performance</i>	
Hao Zhang, Zhenjia Li, Yifan Gao, Xi Xiao, Heng Zhang, Shuyang Zhang, Xiaoxincc, Bo Huang, Yuhang Wu, Tianyang Wang and Hao Xu.....	32937
<i>ConSensus: Multi-Agent Collaboration for Multimodal Sensing</i>	
Hyungjun Yoon, Mohammad Malekzadeh, Sung-Ju Lee, Fahim Kawsar and Lorena Qendro	32950
<i>Beyond High-Entropy Exploration: Correctness-Aware Low-Entropy Segment-Based Advantage Shaping for Reasoning LLMs</i>	
Xinzhu Chen, Xuesheng Li, Zhongxiang Sun and Weijie Yu.....	32970
<i>GraphMind: LLMs as Dynamic Knowledge Builders for Sequential Decision-Making</i>	
Sunguk Shin, Hayeong Lee, Jun Ho Seo, Jinho Lee, Myunsoo Kim, Minsuk Chang and Byung-Jun Lee.....	32985
<i>EET: Experience-Driven Early Termination for Cost-Efficient Software Engineering Agents</i>	
Yaoqi Guo, Ying Xiao, Jie M. Zhang, Mark Harman, Yiling Lou, Yang Liu and Zhenpeng Chen	33008
<i>Fundamental Reasoning Paradigms Induce Out-of-Domain Generalization in Language Models</i>	
Mingzi Cao, Xingwei Tan, Mahmud Elahi Akhter, Marco Valentino, Maria Liakata, Xi Wang and Nikolaos Aletras.....	33025
<i>Rethinking Text-to-SQL: Dynamic Multi-turn SQL Interaction for Real-world Database Exploration</i>	
Linzhuang Sun, Tianyu Guo, Hao Liang, Ruitong Liu, Yuying Li, Qifeng Cai, Jingxuan Wei, Yuchen Wu, Bihui Yu, Xiangxiang Zhang, Wentao Zhang and Bin Cui.....	33047
<i>LogitSpec: Accelerating Retrieval-based Speculative Decoding via Next Next Token Speculation</i>	
Tianyu Liu, Qitan Lv, Hao Li, Xing Gao, Xiao Sun and Xiaoyan Sun.....	33070
<i>Answering Cross-Dimensional Geometric Visual Questions by Multi-constraint Spatial Reasoning</i>	
Dongling Li, Qi Chen, Jianxing Yu, Hanjiang Lai, Yanghui Rao, Wenqing Chen and Jian Yin	33093
<i>SudokuFill: A Multi-Agent Progressive Filling Framework for Document-Level Scientific Information Extraction</i>	
Yang Li, Yajiao Wang, Yu Zhang, Yuanzhe Zhang, Maodi Hu, Mengting Zhang, Xi Sun, Hua Yue and Zhixiong Zhang.....	33112
<i>Emergent Relational Order in LLM Agent Societies: From Collective Affect to Authority Stratification</i>	
Zhiyuan Ji, Xinyu Chen, Ziqi Dai, Shiyun Tang, Chunyu Wei and Yueguo Chen.....	33139

<i>VCB Bench: An Evaluation Benchmark for Audio-Grounded Large Language Model Conversational Agents</i>	
Jiliang Hu, Wenfu Wang, Zuchao Li, Chenxing Li, Yiyang Zhao, Hanzhao Li, Liqiang Zhang, Meng Yu and Dong Yu	33176
<i>LLM-FK: Multi-Agent LLM Reasoning for Foreign Key Detection in Large-Scale Complex Databases</i>	
Zijian Tang, Ying Zhang, Sibao Cai and Ruoxuan Wang	33201
<i>When Meaning Travels: A Granular Lens on Hybrid-MoE's Role in Idiomatic Understanding for Language Models</i>	
Sarmistha Das, Vaibhav Vishal, Shreyas Guha, Amaan Ali, Kitsuchart Pasupa and Sriparna Saha	33229
<i>IntentCoding: Amplifying User Intent in Code Generation</i>	
Zheng Fang, Yihong Dong, Lili Mou, Dongming Jin, Zhi Jin and Ge Li	33246
<i>UbuntuGuard: A Culturally-Grounded Policy Benchmark for Equitable AI Safety in African Languages.</i>	
Tassallah Abdullahi, Macton Mgonzo, Mardiyah Oduwole, Paul Okewunmi, Abraham Toluwase Owodunni, Ritambhara Singh and Carsten Eickhoff	33262
<i>microCLIP: Unsupervised CLIP Adaptation via Coarse-Fine Token Fusion for Fine-Grained Image Classification</i>	
Sathira Silva, Eman Ali, Chetan Arora and Muhammad Haris Khan	33277
<i>Prompting Across Time: Evaluating LLMs on Historical and Contemporary Offensive Language</i>	
Sanne Hoeken, Sophie Jasmin Spliethoff, Silke Schwandt, Özge Alacam and Sina Zarrieß	33295
<i>LDEDE: LRP-Driven Efficient Detection and Editing Framework for LLM Privacy Neurons</i>	
Zhao Zhengyuan, Cao Lifeng, Sunhaodong, Shi Haotian, Du Xuehui, Liu Aodi, Niu Lanjie and Yang Xiaocheng	33315
<i>Chimera: Compositional Jailbreak Attacks on LLMs via Judgment-Driven Search over Heterogeneous Strategies</i>	
Leo Hyun Park, Juwon Cho, Gyuhan Kim, YoonDong Yeo and Taekyoung Kwon.....	33330
<i>Bias Dynamics in BabyLMs: Towards a Compute-Efficient Sandbox for Democratising Pre-Training Debiasing</i>	
Filip Trhlík, Andrew Caines and Paula Buttery	33356
<i>PuzzleClone: A DSL-Powered Framework for Synthesizing Verifiable Data</i>	
Kai Xiong, Yanwei Huang, Rongjunchen Zhang, Kun Chen, Haipang WU and Yingcai Wu	33378
<i>Automated Knowledge Component Generation and Interpretable Knowledge Tracing in Coding Problems</i>	
Zhangqi Duan, Nigel Fernandez, Arun Balajee Lekshmi Narayanan, Mohammad Hassany, Rafaela Sampaio de Alencar, Peter Brusilovsky, Bitu Akram and Andrew Lan	33405
<i>Evaluation Pitfalls and Sparsity Limitations in LLM-based Confidence Estimates for Classification</i>	
Elena Merdjanovska, Omar Zaidan and Andreas Rücklé	33424
<i>MATA: Multi-Agent Framework for Reliable and Flexible Table Question Answering</i>	
Sieun Hyeon, Jusang Oh, Sunghwan Steve Cho and Jaeyoung Do.....	33436
<i>Hallucination Detection in Long-Form Text Generated by LLMs: A Benchmark and a Hyper-Relational Knowledge Graph Approach</i>	
Zituo Li, Guangzhou Chen, Jianbin Sun, Qi Song and Yang Kewei.....	33477

<i>Can Intelligent Agents Revolutionize Scale Generation?</i>	
Chenghao Jia, Zhitao Yuan, Zhaokang Zong, YiFei Yin, Zhe Chen, Man Lan and Shengjun Wu	33495
<i>On Finding Inconsistencies in Documents</i>	
Charles Lovering, Seth Ebner, Brandon Smock, Michael Krumdick, Muhammad Saad Rabbani, Ahmed Muhammad, Varshini Reddy and Chris Tanner	33523
<i>Towards Preference Following in Tool Calling Language Agents</i>	
Zhi-Yuan Chen, Siyu Lu, Qianlong Xie, Xingxing Wang and Yankai Lin	33565
<i>ReSQL: Self-Improving Framework for Reasoning-Aware Text-to-SQL Dataset Generation</i>	
Minjun Park, Yongju Seong, Myoseop Sim, Kyungkoo Min and Stanley Jungkyu Choi	33582
<i>Identifying and typifying demographic unfairness in phoneme-level embeddings of self-supervised speech recognition models</i>	
Felix Herron, Solange Rossato, Alexandre Allauzen and François Portet	33603
<i>MolSafeEval: A Benchmark for Uncovering Safety Risks in AI-Generated Molecules</i>	
Tong Xu, Xinzhe Cao, Zhihui Zhu, Keyan Ding and Huajun Chen	33621
<i>Dynamic Tool Dependency Retrieval for Lightweight Function Calling</i>	
Bhrij Patel, Davide Belli, Amir Jalalirad, Maximilian Arnold, Aleksandr Ermolov and Bence Major	33649
<i>Audit Me If You Can: Query-Efficient Active Fairness Auditing of Black-Box LLMs</i>	
David Hartmann, Lena Pohlmann, Lelia Hanslik, Noah Gießing, Bettina Berendt and Pieter De-lobelle	33673
<i>Mending the Holes: Mitigating Reward Hacking in Reinforcement Learning for Multilingual Translation</i>	
Yifeng Liu, Siqi Ouyang, Yatish H R and Lei Li	33699
<i>MedCoach: Enhancing Medical Reasoning in LLMs via Knowledge Graph-Augmented Chain-of-Thought Distillation</i>	
Chuan Li, Ye Lyu, Chengyu Wang, Mingyuan Fan and Cen Chen	33724
<i>Pru-CoT: Towards Efficient Reasoning Distillation via Pruning Chain-of-Thought</i>	
Han Liu, Shuotian Ma, Hui Li, Xiaotong Zhang, Fenglong Ma and Hong Yu	33744
<i>Curriculum-RLAIF: Curriculum Alignment with Reinforcement Learning from AI Feedback</i>	
Jiaye Lin, Mengdi Li, Xufeng Zhao, Wenhao Lu, Peilin Zhao, Stefan Wermter and Di Wang	33757
<i>Query-Focused Individual Simulation with Progressive Persona Completion</i>	
Weiwen SU, Naoki Yoshinaga and Masashi Toyoda	33778
<i>VIDA: A Visual Intent-driven Design Assistant for Proactive Multimodal Clarification</i>	
Yanshan Liu, Hongbo Zhang, Zhen Sun, Jiaheng Wei and Kaishun Wu	33793
<i>DeepSynth-Eval: Objectively Evaluating Information Consolidation in Deep Survey Writing</i>	
Hongzhi Zhang, Yuanze Hu, Tinghai Zhang, Jia Fu, Tao Wang, Junwei Jing, Zhaoxin Fan, Wei Bi, Ruiming Tang, Han Li, Guorui Zhou and Kun Gai	33805
<i>A Scalable Entity-Based Framework for Auditing Bias in Large Language Models</i>	
Akram Elbouanani, Aboubacar Tuo and Adrian Popescu	33823
<i>HiGMem: A Hierarchical and LLM-Guided Memory System for Long-Term Conversational Agents</i>	
Shuqi Cao, Jingyi He and Fei Tan	33853

<i>Self-Guided Plan Extraction for Instruction-Following Tasks with Goal-Conditional Reinforcement Learning</i>	
Zoya Volovikova, Nikita Sorokin, Dmitriy Lukashevskiy, Aleksandr Panov and Alexey Skrynnik	33863
<i>Adaptive Prompt Optimization for Open-Ended Tasks: Uncertainty Preference as a Secondary Signal</i>	
Shuyang Zhang, Zhixuan Liu, Zhichen Dong, Hao Zhang, Chaochao Lu and Chao Yang . .	33883
<i>DPLoRA: A Dual-Pruning Framework based on ILP Optimization and Progressive Pruning for Parameter-Efficient LoRA Fine-Tuning</i>	
Changjun Park, Sejong Yoon and Jaekwang Kim	33892
<i>Demystifying Multi-Agent Debate: The Role of Confidence and Diversity</i>	
Xiaochen Zhu, Caiqi Zhang, Yizhou Chi, Tom Stafford, Nigel Collier and Andreas Vlachos	33909
<i>RoleConflictBench: A Benchmark of Role Conflict Scenarios for Evaluating LLMs’ Contextual Sensitivity</i>	
Jisu Shin, Hoyun Song, Juhyun Oh, Changgeon Ko, Eunsu Kim, Chani Jung and Alice Oh	33931
<i>RbtAct: Rebuttal as Supervision for Actionable Review Feedback Generation</i>	
Sihong Wu, Yiling Ma, Yilun Zhao, Tiansheng Hu, Owen Jiang, Manasi Patwardhan and Arman Cohan	33965
<i>Is Grokking Worthwhile? Functional Analysis and Transferability of Generalization Circuits in Transformers</i>	
Kaiyu He, Mian Zhang, Peilin Wu, Xinya Du and Zhiyu Chen	33993
<i>When 20 Agents Fail to Sort: The Distributed Sorting Benchmark for Scalable Multi-Agent Systems</i>	
Xin Yang, Junhao Wang, Bintao Tang, Xuxin Cheng, Cao Liu, Ke Zeng and Wenyuan Jiang	34002
<i>RIGOURATE: Quantifying Scientific Exaggeration with Evidence-Aligned Claim Evaluation</i>	
Joseph James, Chenghao Xiao, Yucheng Li, Nafise Sadat Moosavi and Chenghua Lin	34022
<i>Glance-or-Gaze: Incentivizing LMMs to Adaptively Focus Search via Reinforcement Learning</i>	
Hongbo Bai, Yujin Zhou, Yile Wu, Chi-Min Chan, Pengcheng Wen, Kunhao Pan, Sirui Han and Yike Guo	34044
<i>Before Forgetting, Learn to Remember: Revisiting Foundational Learning Failures in LVLMM Unlearning Benchmarks</i>	
Junehyoung Kwon, MiHyeon Kim, Eunju Lee, JungMin Yun, Byeonggeuk Lim and YoungBin Kim	34063
<i>Investigating Links between Illicit Massage Businesses through Natural Language Processing and Graph Machine Learning</i>	
Vasuki Garg, Osman Ozaltin, Maria Mayorga and Sherrie Caltagirone	34080
<i>TopoRAG: Graph-based RAG via Topology-aware Approximate Nearest Neighbor Search</i>	
Tianhao Wu and Siqiang Luo	34097
<i>Self-Explaining Hate Speech Detection with Moral Rationales</i>	
Francielle Vargas, Jackson Trager, Diego Alves, Matteo Guida, Surendrabikram Thapa, Berk Atıl, Daryna Dementieva, Andrew J Smart and Ameeta Agrawal	34109
<i>Principled Detection of Hallucinations in Large Language Models via Multiple Testing</i>	
Jiawei Li, Akshayaa Magesh and Venugopal Veeravalli	34132

<i>What Do Vision–Language Models Encode for Personalized Image Aesthetics Assessment?</i> Koki Ryu and Hitomi Yanaka	34146
<i>A Diversity Diet for a Healthier Model: A Case Study of French ModernBERT</i> Louis Estève, Christophe Servan, Thomas Lavergne and Agata Savary	34168
<i>Is Your Language Model Ready for Monetization Decisions?</i> Jialu Gao, Hua Li, Tracy Ortman, Yeyun Gong and Jian Jiao	34182
<i>Activation Reward Models for Few-Shot Model Alignment</i> Tianning Chai, Chancharik Mitra, Brandon Huang, Gautam Rajendrakumar Gare, Zhiqiu Lin, Assaf Arbelle, Leonid Karlinsky, Rogerio Feris, Trevor Darrell, Deva Ramanan and Roei Herzig	34201
<i>Cross-lingual Matryoshka Representation Learning across Speech and Text</i> Yaya SY, Dioula Doucouré, Christophe Cerisara and Irina Illina	34218
<i>TRACE: Two-Phase RL for Causal Graph Exploration and Deeper Psychological Intervention in Dynamic Counseling Scenarios</i> Shilin Tang, Zunyi Yin, Xuefeng Liang, Guanghui Shi, Song Tong and Chen Guangyu	34229
<i>LLM Multi-Agent Systems for Long Triple Set Data-to-Text Generation</i> Chinonso Cynthia Osuji, Simon Mille, Mark Andrade, Jane Adkins, Ornait O’Connell, Elaine Uí Dhonnchadha, Bláithín Heffernan, Fírinne Nic an tSaoir, Anya Belz, Thiago Castro Ferreira and Brian Davis	34261
<i>Biomed-Enriched: Data-Efficient Biomedical Pretraining via Paragraph-Level Annotation</i> Rian Touchent, Nathan Godey and Éric Villemonte de la Clergerie	34276
<i>Identifying the Achilles’ Heel: An Iterative Method for Uncovering Factual Errors in Large Language Models</i> Wenxuan Wang, Yuk-Kit Chan, Zixuan Ling, Shi Juluan, Youliang Yuan, Jen-tse Huang, Yifei Zhang, Wenxiang Jiao, Zhaopeng Tu and Michael R. Lyu	34288
<i>Probing the Plasticity and Correlation of LLM Value Systems: LLM Value Rankings are Not Stable</i> Zhenheng Tang, Qihua Pan, Jingya Shen, Xiang Liu, Qian Wang, Bo Li and Xiaowen Chu	34310
<i>FAMA: Failure-Aware Meta-Agent Framework for Open-Source LLMs in Interactive Tool Use Environments</i> Amir Saeidi, Venkatesh Mishra, Souradeep Mukhopadhyay, Gaowen Liu, Ali Payani, Jayanth Srinivasa and Chitta Baral	34346
<i>Beyond Reasoning Gains: Mitigating General-Capability Forgetting in Large Reasoning Models</i> Hoang Phan, Xianjun Yang, Yuanshun Yao, Jingyu Zhang, Shengjie Bi, Xiaocheng Tang, Madian Khabsa, Lijuan Liu and Deren Lei	34368
<i>ActorMind: Emulating Human Actor Reasoning for Speech Role-Playing</i> Xi Chen, Wei Xue and Yike Guo	34399
<i>Neuronal Insights into LLM Attacks: Targeted Neuron Tuning for Precise and Robust Vulnerability Patching</i> Dan Shi, Renren Jin, Zhuowen Han, Yuqi Ren, Xinwei Wu, Zhigen Li and Deyi Xiong	34414
<i>Ro-SLM: Onboard Small Language Models for Robot Task Planning and Operation Code Generation</i> Wenhao Wang, Yanyan Li, Long Jiao and Jiawei Yuan	34436
<i>Just Use XML: Revisiting Joint Translation and Label Projection</i> Thennal D K, Chris Biemann and Hans Ole Hatzel	34461

<i>CaTS-Bench: Can Language Models Describe Time Series?</i>	
Luca Zhou, Pratham Yashwante, Marshall Fisher, Alessio Sampieri, Zihao Zhou, Fabio Galasso and Rose Yu	34479
<i>How do Visual Attributes Influence Web Agents? A Comprehensive Evaluation of User Interface Design Factors</i>	
Kuai Yu, Naicheng Yu, Han Wang, Rui Yang and Huan Zhang	34520
<i>Scaling Unverifiable Rewards: A Case Study on Visual Insights</i>	
Shuyu Gan, James Mooney, Pan Hao, Renxiang Wang, Mingyi Hong, Qianwen Wang and Dongyeop Kang	34537
<i>SilentDrift: Exploiting Action Chunking for Stealthy Backdoor Attacks on Vision-Language-Action Models</i>	
Bingxin Xu, Yuzhang Shang, Binghui Wang and Emilio Ferrara	34570
<i>X-ray Made Simple: Lay Radiology Report Generation and Robust Evaluation</i>	
Kun Zhao, Chenghao Xiao, Sixing Yan, Haoteng Tang, William K. Cheung, Noura Al Moubayed, Liang Zhan and Chenghua Lin	34583
<i>Not All Citations Are Equal: Entropy-Guided Citation Selection for Noise-Resistant Medical LLM</i>	
Minyu Gao, Hanlin Xiao, Ruoyu Wang, Shuai Yang, YeXuan Zhang, Xin Wu and Xingyu Liu	34599
<i>Decomposing Unitization and Typing for Efficient and Consistent Span-Bound Concept Annotation</i>	
Nupoor Gandhi, Michael Bada and Emma Strubell	34616
<i>InsightEval: An Expert-Curated Benchmark for Assessing Insight Discovery in LLM-Driven Data Agents</i>	
Zhenghao Zhu, Yuanfeng Song, Xing Chen, Chengzhong Liu, Cui Yakun, Caleb Chen Cao, Sirui Han and Yike Guo	34632
<i>ArrowGEV: Grounding Events in Video via Learning the Arrow of Time</i>	
Fangxu Yu, Ziyao Lu, Liqiang Niu, Fandong Meng and Jie Zhou	34657
<i>Multi-lingual Functional Evaluation for Large Language Models</i>	
Victor Ojewale, Inioluwa Deborah Raji and Suresh Venkatasubramanian	34672
<i>Why LoRA Fails to Forget: Regularized Low-Rank Adaptation Against Backdoors in Language Models</i>	
Hoang-Chau Luong and Lingwei Chen	34692
<i>BoundRL: Efficient Token-level Structured Text Segmentation through Reinforced Boundary Generation</i>	
Haoyuan Li, Zhengyuan Shen, Sullam Jeoung, Yueyan Chen, Jiayu Li, Qi Zhu, Shuai Wang, Vassilis N. Ioannidis and Huzefa Rangwala	34706
<i>AOT*: Efficient Synthesis Planning via LLM-Empowered AND-OR Tree Search</i>	
Xiaozhuang Song, Xuanhao Pan, Xinjian Zhao, Hangting Ye, Shufei Zhang, Jian Tang and Tian-shu Yu	34727
<i>AutoSUIT Bench - Automated Security UnIt Test Benchmark for LLM Coding</i>	
Samuel Osebe, Fan Yang, Junyi Li, Yue Gu, Yongxin Wang, Satyapriya Krishna, Kai-Wei Chang, Aram Galstyan, Rahul Gupta and Weitong Ruan	34759
<i>AnchorMem: Anchored Facts with Associative Contexts for Building Memory in Large Language Models</i>	
Zhanyu Shen, Sijie Cheng, Zhicheng Guo, Weiqin Wang, Yile Wang and Hui Huang	34784

<i>Filling in the Mechanisms: How do LMs Learn Filler-Gap Dependencies under Developmental Constraints?</i>	
Atrey Desai and Sathvik Nair	34799
<i>PictoEduca: Building a Dataset for Spanish Text-to-Pictogram Generation</i>	
Alfonso Manuel Paredes Umeres and Marco Antonio Sobrevilla Cabezedo	34816
<i>HCFD: A Benchmark for Audio Deepfake Detection in Healthcare</i>	
Mohd Mujtaba Akhtar, Girish and Muskaan Singh	34829
<i>Linear Semantic Segmentation for Low-Resource Spoken Dialects</i>	
Kirill Chirkunov, Younes Samih, Abed Alhakim Freihat and Hanan Aldarmaki	34844
<i>Thinking Like a Botanist: Challenging Multimodal Language Models with Intent Driven Chain-of-Inquiry</i>	
Syed Nazmus Sakib, Nafiul Haque, Shahrear Bin Amin, Hasan Muhammad Abdullah, Md Mehedi Hasan, Mohammad Zabed Hossain and Shifat E. Arman	34862
<i>Analysing Differences in Persuasive Language in LLM-Generated Text: Uncovering Stereotypical Gender Patterns</i>	
Amalie Brogaard Pauli, Maria Barrett, Max Müller-Eberstein, Isabelle Augenstein and Ira Assent	34893
<i>Robust In-Context Selection via Online Learned Position-Corrected Attention</i>	
Deeksha Koul, Gaurav Kumar, Yash Sabale and Sunita Sarawagi	34919
<i>Comprehensiveness Metrics for Automatic Evaluation of Factual Recall in Text Generation</i>	
Adam Dejl, James Barry, Alessandra Pascale and Javier Carnerero-Cano	34931
<i>TableVista: Benchmarking Multimodal Table Reasoning under Visual and Structural Complexity</i>	
Zheyuan Yang, Liqiang Shang, Junjie Chen, Xun Yang, Chenglong Xu, Bo Yuan, Chenyuan Jiao, Yaoru Sun and Yilun Zhao	34967
<i>Lost in Translation: Do LVLM Judges Generalize Across Languages?</i>	
Md Tahmid Rahman Laskar, Mohammed Saidul Islam, Mir Tafseer Nayeem, Amran Bhuiyan, Mizanur Rahman, Shafiq Joty, Enamul Hoque and Jimmy Huang	34986
<i>Making Revisions Understandable: A Survey of Edit Intentions, Methods, and Applications</i>	
Fangping Lan, Qi Zhang and Eduard Dragut	35003
<i>GSM-Noise: Exploring and Enhancing Large Language Models' Reasoning under Noisy Inputs</i>	
Zhengxin Zhang, Chengyu Huang, Xufu Liu, Dan Zhao, Jinyan Su and Claire Cardie	35020
<i>VALUE ALIGNMENT TAX: Measuring Value Trade-offs in LLM Alignment</i>	
Jiajun Chen and Hua Shen	35046
<i>Confidence-Aware Ranker Ensembles for Robust In-Context Knowledge Editing</i>	
Tejal Nair, Mahmud Wasif Nafee, Maiqi Jiang, Ashley Gao, Haipeng Chen and Yanfu Zhang	35070
<i>EMSDialog: Synthetic Multi-person Emergency Medical Service Dialogue Generation from Electronic Patient Care Reports via Multi-LLM Agents</i>	
Xueren Ge, Sahil Murtaza, Anthony Cortez and Homa Alemzadeh	35081
<i>Better LLM Reasoning via Dual-Play</i>	
Zhengxin Zhang, Chengyu Huang, Aochong Oliver Li and Claire Cardie	35111

<i>CobwebTM: Probabilistic Concept Formation for Lifelong and Hierarchical Topic Modeling</i> Karthik Singaravadelan, Anant Gupta, Zekun Wang and Christopher J. MacLellan	35140
<i>Adaptive Test-Time Compute Allocation with Evolving In-Context Demonstrations</i> Bowen Zuo, Dongruo Zhou and Yinglun Zhu	35156
<i>CCD: Mitigating Hallucinations in Radiology MLLMs via Clinical Contrastive Decoding</i> Xi Zhang, Zaiqiao Meng, Jake Lever and Edmond S. L. Ho	35174
<i>Can Large Language Models Adequately Perform Symbolic Reasoning Over Time Series?</i> Zewen Liu, Juntong Ni, Xianfeng Tang, Max SY Lau, Qi He, Wenpeng Yin and Wei Jin . . .	35201
<i>Reading Between the Lines: The One-Sided Conversation Problem</i> Victoria Ebert, Rishabh Singh, Tuochao Chen, Noah A. Smith and Shyamnath Gollakota .	35227
<i>Fico: Evaluating Vision-Language Models under Visual Fidelity and Compression at Scale</i> Jianhong Tu, Nicholas Crispino, Kyle Montgomery, Chenguang Wang and Dawn Song . . .	35261
<i>Sycobench-600: Measuring Sycophancy and Correction Selectivity in LLM Assistants</i> Debu Sinha	35278
<i>Calibrating Model-Based Evaluation Metrics for Summarization</i> Hongye Liu, Dhanajit Brahma and Ricardo Henao	35285
<i>SafeMERGE: Preserving Safety Alignment in Fine-Tuned Large Language Models via Selective Layer-Wise Model Merging</i> Aladin Djuhera, Swanand Ravindra Kadhe, Farhan Ahmed, Syed Zawad and Holger Boche	35316
<i>Learning to Control Summaries with Score Ranking</i> Hongye Liu, Liang Ding and Ricardo Henao	35336
<i>Raw Pointer Rewriting with LLMs for Translating C to Safer Rust</i> Yifei Gao, Chengpeng Wang, Pengxiang Huang, Xuwei Liu, Mingwei Zheng and Xiangyu Zhang	35360
<i>Beyond Single Plots: A Benchmark for Question Answering on Multi-Charts</i> Azher Ahmed Efat, Seok Hwan Song and Wallapak Tavanapong	35374
<i>Assessing Y-Axis Influence: Bias in Multimodal Language Models on Chart-to-Table Translation</i> Seok Hwan Song, Azher Ahmed Efat and Wallapak Tavanapong	35412
<i>FLOP-Efficient Training: Early Stopping Based on Test-Time Compute Awareness</i> Hossam Amer, Maryam Dialameh, Hossein Rajabzadeh, Walid Ahmed, Weiwei Zhang and Yang Liu	35431
<i>Save the Good Prefix: Precise Error Penalization via Process-Supervised RL to Enhance LLM Reasoning</i> Haolin Liu, Dian Yu, Sidi Lu, Yujun Zhou, Rui Liu, Zhenwen Liang, Haitao Mi, Chen-Yu Wei and Dong Yu	35450
<i>POSTCONDBENCH: Benchmarking Correctness and Completeness in Formal Postcondition Inference</i> Gehao Zhang and Juan Zhai	35478
<i>Who's Asking? Simulating Role-Based Questions for Conversational AI Evaluation</i> Navreet Kaur, Hoda Ayad, Hayoung Jung, Shravika Mittal, Munmun De Choudhury and Tanu Mitra	35508

<i>CircuitSynth: Reliable Synthetic Data Generation</i>	
Zehua Cheng, Wei Dai, Jiahao Sun and Thomas Lukasiewicz	35542
<i>MedRedFlag: Investigating how LLMs Redirect Misconceptions in Real-World Health Communication</i>	
Sraavya Sambara, Yuan Pu, Ayman Ali, Vishala Mishra, Lionel Wong and Monica Agrawal	35553
<i>Can VLMs Predict Future States? Bootstrapping World Models from Inverse Dynamics</i>	
Yifu Qiu, Yftah Ziser, Anna Korhonen, Shay B Cohen and Edoardo Ponti	35579
<i>PL-MTEB: Polish Massive Text Embedding Benchmark</i>	
Rafał Poświata, Sławomir Dadas and Michał Wiktor Perełkiewicz	35601
<i>RubRIX: Rubric-Driven Risk Mitigation in Caregiver-AI Interactions</i>	
Drishti Goel, Jeongah Lee, Qiuyue Zhong, Violeta J. Rodriguez, Daniel S. Brown, Ravi Karkar, Dong Whi Yoo and Koustuv Saha	35620
<i>Spectral Gravity Formant Estimation for Phonetic Segmentation</i>	
Michael S. Yantosca and Albert M. K. Cheng	35639
<i>Listen, Pause, and Reason: Toward Perception-Grounded Hybrid Reasoning for Audio Understanding</i>	
Jieyi Wang, Yazhe Niu, Dexuan Xu and Zhongyu Wei	35653
<i>DeFrame: Debiasing Large Language Models Against Framing Effects</i>	
Kahee Lim, Soyeon Kim and Steven Euijong Whang	35672
<i>Idea First, Code Later: Disentangling Problem Solving from Code Generation in Evaluating LLMs for Competitive Programming</i>	
Sama Hadhoud, Alaa Elsetohy, Frederikus Hudi, Jan Christian Blaise Cruz, Steven Halim and Alham Fikri Aji	35708
<i>LLMs in the Real World: Evaluating “AI” in Emergency Contexts</i>	
Sara Court, Lara Downing and Micha Elsner	35748
<i>A Survey of Toxicity Mitigation Strategies for Multilingual Language Models</i>	
Soham Dan, Himanshu Beniwal and Thomas Hartvigsen	35761
<i>LLM-KT: Enhancing Large Language Models with Knowledge Tracing via Multi-Level Plug-and-Play Alignment</i>	
Ziwei Wang, Jie Zhou, Qin Chen, Bo Jiang, Qingchun Bai, Liang Dou and Liang He	35775
<i>Fast-Decoding Diffusion Language Models via Progress-Aware Confidence Schedules</i>	
Amr Mohamed, Yang Zhang, Michalis Vazirgiannis and Guokan Shang	35793
<i>When Cultures Meet: Multicultural Text-to-Image Generation</i>	
Parth Bhalerao, Oana Ignat, Brian Trinh and Mounika Yalamarty	35808
<i>Aligning with Your Own Voice: Self-Corrected Preference Learning for Hallucination Mitigation in LVLMs</i>	
Byeonggeuk Lim, JungMin Yun, Junehyoung Kwon, Kyeonghyun Kim and YoungBin Kim	35829
<i>VocalRep: Structure-Aware Vocal Representations for Multimodal Generation</i>	
Da Shen, Zhenqiang Weng, Tianyu Liu, Gongyu Chen, Runhua Shi, Jiahui Chen, Chaofan Ding, Wei-Qiang Zhang and Zihao Chen	35850
<i>RealSec-bench: A Benchmark for Evaluating Secure Code Generation in Real-World Repositories</i>	
Yanlin Wang, Ziyao Zhang, Chong Wang, Xinyi Xu, Mingwei Liu, Yong Wang, Jiachi Chen and Zibin Zheng	35866

<i>InfiAgent: An Infinite-Horizon Framework for General-Purpose Autonomous Agents</i> Chenglin Yu, Yuchen Wang, Songmiao Wang, Hongxia Yang and Li Ming	35884
<i>Efficient and Effective Internal Memory Retrieval for LLM-Based Healthcare Prediction</i> Mingchen Li, Jiatan Huang, Zonghai Yao and Hong yu	35895
<i>GAVEL: Evidence-Contract Debate with Mechanized Scrutiny for Provenance-Grounded Fact-Checking</i> Ruoyu Xu, Gaoxiang Li and Victor S. Sheng	35907
<i>EMPATH: An Ensemble Method for Automatic Fine-Grained Turn-Level Dialogue Empathy Evaluation with a Novel Emotional Distance Metric</i> Dongning Rao, Zhihua Liang and Zhihua Jiang	35921
<i>Q2EI: Query-to-Entity Inference for Semantic Condensation in Domain-Specific Retrieval</i> Yixuan Sun, Zhenqin Xu, HanFeng Zhai, Zishu Yu and Xiaohui Peng	35943
<i>TVWorld: Foundations for Remote-Control TV Agents</i> Zhantao Ma, Quanfeng Lu, Shuai Zhong, Dahai Yu, Ping Luo and Michael Ng	35959
<i>RiTeK: A Dataset for Large Language Models Complex Reasoning over Textual Knowledge Graphs in Medicine</i> Jiatan Huang, Mingchen Li, Zonghai Yao, Dawei Li, Yuxin Zhang, Zhichao Yang, Yongkang Xiao, Feiyun Ouyang, Xiaohan Li, Shuo Han and Hong yu	35985
<i>PRA-RAG: Provably Robust Aggregation in Retrieval-Augmented Generation against Retrieval Corruption</i> Xue Tan, Yi Zheng, Chang Huo, Yunruo Zhang, Yu Liu, Hao Luan, Zhuyang Yu, Jun Dai, Xiaoyan Sun and Ping Chen	36002
<i>Evaluating Perspectival Biases in Cross-Modal Retrieval</i> Teerapol Saengsukhiran, Peerawat Chomphooyod, Narabodee Rodjananant, Chompakorn Chak-sangchaichot, Patawee Prakrankamanant, Witthawin Sripheanpol, Pak Lovichit, Sarana Nutanong and Ekapol Chuangsuwanich	36018
<i>Beyond Static Synthetic Noise: Assessing the Robustness of Large Language Models to Natural Context Variation in the Real World</i> Yulong Wu, Viktor Schlegel and Riza Batista-Navarro	36050
<i>Measuring Watermarking under Jailbreaking: ASR Inflation and Goal-Compliance Mismatch</i> Sungwoo Han, Sangjun Moon, Jinguin Kwon, Hidetaka Kamigaito and Manabu Okumura	36071
<i>Inventive Problem Solving with LLMs: A Benchmark for TRIZ Reasoning</i> Zhu Wang and Brian Uzzi	36084
<i>Do LLM Agents Really Mimic Humans? Diagnosing and Aligning Microeconomic Behaviors in Macro-ABMs</i> Guangya Liu, Cheng Wang, Jiangtong Li, Huafei Wu and Changjun Jiang	36102
<i>Explanation Quality Assessment as Ranking with Listwise Rewards</i> Thomas Bailleux, Tanmoy Mukherjee, Emmanuel Lonca, Pierre Marquis and Zied Bouraoui	36123
<i>ViLegalLM: Language Models for Vietnamese Legal Text</i> Truong-Phuc Nguyen, Quy-Nhan Nguyen and Minh-Tien Nguyen	36136
<i>From Heads to Neurons: Causal Attribution and Steering in Multi-Task Vision–Language Models</i> Qidong Wang, Junjie Hu and Ming Jiang	36151

<i>Syntax as a Rosetta Stone: Universal Dependencies for In-Context Coptic Translation</i> Abhishek Purushothama, Emma Thronson, Alexia Guo and Amir Zeldes	36176
<i>AI Agents for the Science of Science: A Survey of Tasks, Architectures, Evaluations, and Challenges</i> Yixuan Liu and Yicheng Zhang	36196
<i>A2O: LLM-based Agentic Learning of Action-to-Object Features for Video Action Recognition</i> Taiki Sekii and Fumiaki Sato	36212
<i>PrinciplismQA: A Philosophy-Grounded Approach to Assessing LLM-Human Clinical Medical Ethics Alignment</i> Chang Hong, Minghao Wu, Qingying Xiao, Yuchi Wang, Xiang Wan, Guangjun Yu, Benyou Wang and Yan Hu	36229
<i>Take Out Your Calculators: Estimating the Real Difficulty of Question Items with LLM Student Simulations</i> Christabel Acquaye, Yi Ting Huang, Marine Carpuat and Rachel Rudinger	36246
<i>SAFER: Advancing Safety Alignment via Efficient Ex-Ante Reasoning</i> Kehua Feng, Keyan Ding, Yuhao Wang, Menghan Li, Fanjunduo Wei, Xinda Wang and Huajun Chen	36268
<i>Speculative Verification: Exploiting Information Gain for Speculative Decoding</i> Sungkyun Kim, Jaemin Kim, Dogyeong Yun, Jiho Shin, Junyeol Lee and Jiwon Seo	36290
<i>EngiBench: A Benchmark for Evaluating Large Language Models on Engineering Problem Solving</i> Xiyuan Zhou, Xinlei Wang, Yirui He, Ruixi Zou, Yang Wu, Yuheng Cheng, Yulu Xie, Wenxuan Liu, Huan Zhao, Yan Xu, Jinjin Gu and Junhua Zhao	36308
<i>LLM-Based Human-Agent Collaboration and Interaction Systems: A Survey</i> Henry Peng Zou, Wei-Chieh Huang, Yaozu Wu, Jizhou Guo, Yankai Chen, Chunyu Miao, Hoang H Nguyen, Yue Zhou, Weizhi Zhang, Liancheng Fang, Hanrong Zhang, Fangxin Wang, Pengfei Zhang, Langzhou He, Yangning Li, Dongyuan Li, Renhe Jiang and Philip S. Yu	36335
<i>Thinking Twice Makes Large Language Models Safer and More Helpful</i> Yutao Mou, Yuxiao Luo, Shikun Zhang and Wei Ye	36365
<i>AEQ-Bench: Measuring Empathy of Omni-Modal Large Models</i> Xuan Luo, Lewei Yao, Lanqing Hong, Kai Chen, Dehua Tao, Daxin Tan, Yukun Deng, Ruifeng Xu and Jing Li	36390
<i>BAPO: Boundary-Aware Policy Optimization for Reliable Agentic Search</i> Shiyu Liu, Yongjing Yin, Jianhao Yan, Yunbo Tang, Qinggang Zhang, Bei Li, Xin Chen, Jingang Wang, Xunliang Cai and Jinsong Su	36408
<i>GRASP: Graph-Reasoning Aided Survey Planning for High-Fidelity Related Work Generation</i> Haoming Li and Jessica Ouyang	36427
<i>Benchmarking Agentic Newswriting via Journalistic Workflows</i> Yen-Che Chien, Kuang-Da Wang, Wei-Yao Wang and Wen-Chih Peng	36450
<i>Self-Correcting Text-to-Video Generation with Misalignment Detection and Localized Refinement</i> Daeun Lee, Jaehong Yoon, Jaemin Cho and Mohit Bansal	36464
<i>Switching Heads and Softening Tokens: Turnkey Solutions to Visually Grounded Document QA</i> Ximing Wen, Wenbo Li, Sudipta Paul, Yashas Malur Saidutta, Kalpa Gunaratna and Srinivas Chappidi	36490

<i>HTMuon: Improving Muon via Heavy-Tailed Spectral Correction</i>	
Tianyu Pang, Yujie Fang, Zihang Liu, Shenyang Deng, Lei Hsiung, Shuhua Yu and Yaoqing Yang	
	36504
<i>Can Large Language Models Effectively Support Decision-Making in Sudden Emergencies?</i>	
Mengna Zhu, Jibing Wu, Lihua Liu, Yuran Gong, Yang Hao, Fu Yachao, Mao Wang, Lei Hou and Juanzi Li	
	36536
<i>J-Shuwa: A Large-Scale Web-Collected Japanese Sign Language-Japanese Parallel Corpus</i>	
Junwen Mo, MinhDuc Vo, Noriki Nishida, Shin'ichi Satoh and Hideki Nakayama	
	36559
<i>Hallucinations as Orthogonal Noise: Inference-Time Manifold Alignment via Dynamic Contextual Orthogonalization</i>	
Mingkuan Zhao, Wentao Hu, Tianchen Huang, Yuheng Min, Suquan Chen, Yide Gao, Yanbo Zhai, Shuangyong Song and Xuelong Li	
	36575
<i>ConvX: A Lightweight Converter to Bridge Indexed Dense Representations and Large Language Models for Retrieval-Augmented Generation</i>	
Bonggeun Choi, Keunha Kim, Junho Han and Youngjoong Ko	
	36587
<i>Resonant Context Anchoring: Decoupling Attention Routing and Signal Gain at Inference Time</i>	
Mingkuan Zhao, Yide Gao, Wentao Hu, Suquan Chen, Tianchen Huang, Zhenhua An, Zetao Chang, Xiayu Sun and Yuheng Min	
	36600
<i>S2H-DPO: Hardness-Aware Preference Optimization for Vision–Language Models</i>	
Nitish Shukla, Sargan Jandial and Arun Ross	
	36612
<i>Dr. Assistant: Enhancing Clinical Diagnostic Inquiry via Structured Diagnostic Reasoning Data and Reinforcement Learning</i>	
Yue Guo, Fanfu Wang, Jianwei Lv, Xincheng Shi, Yuchen Li, Youya Wang, Yunsheng Zeng, Yujing Liu, Yunhao Qiao, Gen Li, Junfeng Wang and Bo Yuan	
	36624
<i>FigSIM: A Dataset for Fine-grained Suicide Severity and Figurative Language in Suicide Memes</i>	
Liuliu Chen, Elise Carrotte, Brian E. Chapman, Jo Robinson and Mike Conway	
	36659
<i>Activation Steering for Chain-of-Thought Compression</i>	
Seyedarmin Azizi, Erfan Baghaei Potraghloo, Souvik Kundu and Massoud Pedram	
	36676
<i>Decomposed Prompting Does Not Fix Knowledge Gaps, But Helps Models Say “I Don’t Know”</i>	
Dhruv Madhwal, Lyuxin David Zhang, Dan Roth, Tomer Wolfson and Vivek Gupta	
	36688
<i>VisTW: Benchmarking Vision-Language Models for Taiwanese Mandarin in Taiwan</i>	
Zhi Rui Tam, Yung-Yu Shih, Yen-Wei Lee, Ya-Ting Pai, Wen Yu Chang and Yun-Nung Chen	
	36711
<i>PsyChain: A Collaborative Chain-of-Agents Framework for Generating Personalized and Professional Counseling Dialogues</i>	
Yi Feng, Zijie Yang, Chen Zhang, Wenxuan Zhang, Dongming Zhang, Liping Jing and Jian Yu	
	36757
<i>Breaking the Autoregressive Chain: Hyper-Parallel Decoding for Efficient LLM-Based Attribute Value Extraction</i>	
Theodore Glavas, Nikhita Vedula, Dushyanta Dhyani, Yilun Zhu and Shervin Malmasi	
	36792
<i>On the Editability of Delta Parameters in Post-Trained Models</i>	
Qiaoyu Tang, Le Yu, Bowen Yu, Hongyu Lin, Keming Lu, Yaojie Lu, Xianpei Han and Le Sun	
	36809

<i>Reasoning with Memory: Adaptive Information Management for Retrieval-Augmented Generation</i> Hieu Man, Ro-ee Tal, Abhishek Kumar, Jaejin Cho and Benjamin Hsu	36825
<i>LiCoMemory: Lightweight and Cognitive Agentic Memory for Efficient Long-Term Reasoning</i> Zhengjun Huang, Zhoujin Tian, Qintian Guo, Fangyuan Zhang, Yingli Zhou, Di Jiang, Zeying Xie and Xiaofang Zhou	36842
<i>Multimodal Item Scoring for Natural Language Recommendation via Gaussian Process Regression with LLM Relevance Judgments</i> Yifan Simon Liu, Qianfeng Wen, Jiazhou Liang, Mark Zhao, Justin Cui, Anton Korikov, Armin Toroghi, Junyoung Kim and Scott Sanner	36859
<i>Dialectic-Med: Mitigating Diagnostic Hallucinations via Counterfactual Adversarial Multi-Agent Debate</i> Zhixiang Lu and Jionglong Su	36877
<i>EPIR: Capturing Promoting and Inhibiting Relationships between Events</i> Bowen Dong, Wenjun Wang, Xueli Liu and Quanlin Qiu	36895
<i>EULoInf: Efficient Hessian-Free Entropy Based Uncertainty-Aware Data Influence Approximation</i> Runxin Cai, Jingtian Wang and Bryan Kian Hsiang Low	36911
<i>CodeRise: Bootstrapping LLMs for Ultra Low-Resource Programming Languages via Progressive Self-Refinement Curriculum</i> Tengfei Wen, Xuanang Chen, Ben He, Xiaoliang Cong and Le Sun	36929
<i>What Tokens Truly Matter? The Logit Conflation Problem in LLM Sampling</i> Pinlong Zhao, Huijun Tang, Pengfei Jiao and Mengyang Li	36943
<i>Bears, all bears, and some bears. Language Constraints on Language Models' Inductive Inferences</i> Sriram Padmanabhan, Siyuan Song and Kanishka Misra	36962
<i>B-APO: Bias-Targeted Adversarial Preference Optimization for Debiasing Multimodal Large Language Models</i> Pinlong Zhao, Zike Ding, Zengshu Ye and Zhou Zhaoting	36979
<i>On the Fallacy of Global Token Perplexity in Spoken Language Model Evaluation</i> Chan-Jan Hsu, Liang-Hsuan Tseng, Yi-Cheng Lin, Yen-Chun Kuo, Ju-Chieh Chou, Kai-Wei Chang, Hung-yi Lee and Carlos Busso	37000
<i>Modeling and Solving Stable Matching under Probabilistic Preferences with Large Language Models</i> Yuqi Kong, Shiyu Liu, Jiaxu Li, Hongtao Liu, Qi Qi and Weiran Shen	37021
<i>XQ-MEval: A Dataset with Cross-lingual Parallel Quality for Benchmarking Translation Metrics</i> Jingxuan Liu, Zhi Qu, Jin Tei, Hidetaka Kamigaito, Lema Liu and Taro Watanabe	37034
<i>Faithfulness vs. Safety: Evaluating LLM Behavior Under Counterfactual Medical Evidence</i> Kaijie Mo, Siddhartha Venkatayogi, Chantal Shaib, Ramez Kouzy, Wei Xu, Byron C Wallace and Junyi Jessy Li	37053
<i>Your LLM Agents are Temporally Blind: The Misalignment Between Tool Use Decisions and Human Time Perception</i> Yize Cheng, Arshia Soltani Moakhar, Chenrui Fan, Parsa Hosseini, Kazem Faghieh, Zahra Soda-gar, Wenxiao Wang and Soheil Feizi	37082
<i>How Do LLMs Trust Unknown Knowledge? An Unknown Knowledge Based Jailbreak Attack</i> Yixiao Huang, Lan Zhang and Chaoran Wang	37105

<i>ToolSafe: Enhancing Tool Invocation Safety of LLM-based agents via Proactive Step-level Guardrail and Feedback</i>	
Yutao Mou, Zhangchi Xue, Lijun Li, Peiyang Liu, Shikun Zhang, Wei Ye and Jing Shao . .	37125
<i>A Survey of Large Models in Sports</i>	
Yichen Xu, Jianzhe Ma, Chuhan Wang, Zhonghao Cao, Liangyu Chen, Wenxuan Wang and Qin Jin	37154
<i>SafetyALFRED: Evaluating Safety-Conscious Planning of Vision Language Models</i>	
Josue Torres-Fonseca, Naihao Deng, Yinpei Dai, Shane Storks, Yichi Zhang, Rada Mihalcea, Casey Kennington and Joyce Chai	37190
<i>Towards Unified Multimodal Large Language Models: A survey</i>	
Xu Ma, Yitian Zhang and Yun Fu	37212
<i>♪ Something Just Like TRuST ♪ *: Toxicity Recognition of Span and Target</i>	
Berk Atıl, Namrata Suresddy and Rebecca J. Passonneau	37231
<i>Synergizing Stylometrics with Semantics: Dual-Path Framework for LLM Detection and Attribution</i>	
Xingyu Lu, Yumeng Ma, Xiang Zhou, Shengli Gan, Guiying Deng, Yang Wen and Yanbing Liu	37252
<i>Enhancing Multilingual Reasoning via Steerable Model Merging</i>	
Zhuoran Li, Rui Xu, Jian Yang, Junnan Liu, Zhijun Chen, Qianren Mao, Hongcheng Guo, Jiaheng Liu, Likang Xiao, Ming LI and Xiaojie Wang	37266
<i>Adaptive Backtracking for Privacy Protection in Large Language Models</i>	
Zhihao Yao, Yuxuan Gu, Xiachong Feng, Weitao Ma, Bo Li, Xiaocheng Feng and Bing Qin	37278
<i>SemPA: Improving Sentence Embeddings of Large Language Models through Semantic Preference Alignment</i>	
Ziyang Chen, Zhenxuan Huang, Yile Wang, Weiqin Wang, Lu Yin and Hui Huang	37299
<i>Interpretability from the Ground Up: Stakeholder-Centric Design of Automated Scoring in Educational Assessments</i>	
Yunsung Kim, Michael Hardy, Joseph Tey, Candace Thille and Christopher J Piech	37313
<i>DANCE: Diversity-attended Dynamic Caching with Asymmetric Quantization for Test-time Adaptation of Vision-Language Models</i>	
Shunge Zou, Changhu Wang, Wei Ju, Ziyue Qiao and Xiao Luo	37329
<i>How to Set the Learning Rate for Large-Scale Pre-training?</i>	
Yunhua Zhou, Shuhao Xing, Junhao Huang, Xipeng Qiu and Qipeng Guo	37344
<i>Evaluating Large Vision Language Models on Bangla Medical Visual Question Answering</i>	
Rafid Ahmed, Intesar Tahmid, Mir Sazzat Hossain, Tasnimul Hossain Tomal, Md Mahir Jawad, Anam Borhan Uddin, Md Fahim and Md Farhad Alam Bhuiyan	37362
<i>SIV-Bench: A Video Benchmark for Social Interaction Understanding and Reasoning</i>	
Fanqi Kong, Weiqin Zu, Xinyu Chen, Yaodong Yang, Song-Chun Zhu and Xue Feng	37379
<i>No One Fits All: From Fixed Prompting to Learned Routing in Multilingual LLMs</i>	
Wei-Chi Wu, Sheng-Lun Wei, Hen-Hsen Huang and Hsin-Hsi Chen	37404
<i>Understanding and Mitigating Spurious Signal Amplification in Test-Time Reinforcement Learning for Math Reasoning</i>	
Yongcan Yu, Lingxiao He, Jian Liang, Kuangpu Guo, Meng Wang, Qianlong Xie, Xingxing Wang and Ran He	37424

<i>Steering LLM Thinking with Budget Guidance</i>	
Junyan Li, Wenshuo Zhao, Yang Zhang and Chuang Gan	37437
<i>Distilling Long-CoT Reasoning through Collaborative Step-wise Multi-Teacher Decoding</i>	
Taewon Yun, Jisu Shin, Jeonghwan Choi, Seunghwan Bang and Hwanjun Song	37452
<i>Adversarial Yet Cooperative: Multi-Perspective Reasoning in Retrieved-Augmented Language Models</i>	
Can Xu, Lingyong Yan, Jiayi Wu, Haosen Wang, Shuaiqiang Wang, Yuchen Li, Jizhou Huang, Dawei Yin and Xiang Li	37469
<i>PyraMathBench: Evaluating and Improving Mathematical Capability in Large Language Models</i>	
Zetian Ouyang, Linlin Wang, Gerard de Melo and Liang He	37481
<i>D2PCM: A Multi-Turn Dialogue Dataset with Personalized Contextual Memory</i>	
Zhe Yang, Yi Huang, Yaqin Chen, Chunyang Gao, Jingyu Yao and Junlan Feng	37512
<i>MTP-RL: Acceleration of Reinforcement Learning Rollouts with Policy-Aligned Multi-Token Prediction</i>	
Ke Wang, Aohan Zeng, Zhengxiao Du, Yuxuan Hu, Bohan Zhang, Xinyi Wang, Jie Tang and Jing Zhang	37530
<i>ReflexiCoder: Teaching Large Language Models to Self-Reflect on Generated Code and Self-Correct It via Reinforcement Learning</i>	
Juyong Jiang, Jiasi Shen, Sunghun Kim, Kang Min Yoo, Jeonghoon Kim and Sungju Kim	37543
<i>TRAC: Token-level Reward Assignment for Coherent Abstractive Summarization</i>	
陈宣齐, 容梓莹, Xinfeng Liao, Lianxi Wang, Ying Gao and Shengyi Jiang	37563
<i>Planning Beyond Text: Graph-based Reasoning for Complex Narrative Generation</i>	
Hanwen Gu, Chao Guo, Junle Wang, Wenda Xie and Yisheng Lv	37579
<i>Enrich, Aggregate, and Generate: Three-stage Biomedical Data-to-Text Generation Using Large Language Models in Low-resource Scenarios</i>	
Yupian Lin, Guangya Yu, Yuang Bian, Cheng Yuan, Hui Luo and Tong Ruan	37611
<i>HiCoLoRA: Addressing Context-Prompt Misalignment via Hierarchical Collaborative LoRA for Zero-Shot DST</i>	
Shuyu Zhang, Yifan Wei, Xinru Wang, Yanmin Zhu, Yangfan He, Yixuan Weng, Yujie Liu and Bin Li	37623
<i>From Conversation to Evaluation: Benchmarking LLMs on Development Knowledge via SimpleDevQA</i>	
Jing Zhang, Lianghong Guo, Yanlin Wang, Terry Yue Zhuo, Yong Wang, Mingwei Liu, Jiachi Chen, Ensheng Shi, Yuchi Ma, Hongyu Zhang and Zibin Zheng	37649
<i>LLEOT: A Privacy-Enhancing Offsite Tuning Framework via Loss Landscape Elevation</i>	
Jin Zhong, Jinglin Liang, Tongtong Yang, Zijian Xie, Shuangping Huang and Hanlin Gu ..	37664
<i>Beyond Compromise: Pareto-Lenient Consensus for Efficient Multi-Preference LLM Alignment</i>	
Renxuan Tan, Rongpeng Li, Zhifeng Zhao and Honggang Zhang	37684
<i>HSS-Synth: Humanities and Social Sciences Data Synthesis for LLMs</i>	
Ru Peng, Tianyu Zhao, Xijun Gu, Zhiting Fan, Haokai Xu, Jinyang Zhang, Yawen Zeng, Yihong Zhuang, Kexin Yang, Junyang Lin, Dayiheng Liu and Junbo Zhao	37706
<i>LLMRouterBench: A Massive Benchmark and Unified Framework for LLM Routing</i>	
Hao Li, Yiqun Zhang, Zhaoyan Guo, Chenxu Wang, Shengji Tang, Qiaosheng Zhang, Yang Chen, Biqing Qi, Peng Ye, Lei Bai, Zhen Wang and Shuyue Hu	37733

<i>CEMT: Controllable Element-Oriented Machine Translation via Structured Linguistic Reasoning</i> Lingling Shi, Haoyu Jin, Ruiyu Fang, Shuangyong Song, Jinsong Su, Yongxiang Li and Xuelong Li	37755
<i>Orchestrating Tokens and Sequences: Dynamic Hybrid Policy Optimization for RLVR</i> Zijun Min, Bingshuai Liu, Ante Wang, Long Zhang, Anxiang Zeng, Haibo Zhang and Jinsong Su	37782
<i>Seeing Isn't Believing: Mitigating Belief Inertia via Active Intervention in Embodied Agents</i> Hanlin Wang, Chak Tou Leong, Jian Wang and Wenjie Li	37796
<i>AlgBench: To What Extent Do Large Reasoning Models Understand Algorithms?</i> Henan Sun, Kaichi Yu, Yuyao Wang, Bowen Liu, Xunkai Li, Rong-Hua Li, Nuo Chen and Jia Li	37816
<i>A Learnable Skill Combination Strategy for Multi-task Learning in Natural Language Understanding</i> Zhe Yang, Yi Huang, Yaqin Chen, Mengfei Guo, Xiaoting Wu and Junlan Feng	37839
<i>CoopValue: Revealing LLM Value Preferences Through Multi-Agent Cooperation</i> Zee Hen Tang and Mi-Yen Yeh	37846
<i>QUARTZ: Quantile-Aware Routing and Queueing for TTFT SLOs in LLM Serving</i> Zhipeng Liu, Yifan Zheng, Fanqi Kong and Ziming Zhao	37886
<i>MPTc-Bench: Measuring Cross-market Generative Ability of Vision-Language Models via Movie Poster Transcreation</i> Youyuan Lin, Yuan Li, Yahan Yu, Fei Cheng, Shin'ya Nishida and Chenhui Chu	37897
<i>Counterfactual Fairness Evaluation of LLM-Based Contact Center Agent Quality Assurance System</i> Kawin Mayilvaghanan, Siddhant Gupta and Ayush Kumar	37914
<i>Masked Text-to-Audio Flow-Matching and Reward Feedback Optimization</i> Rongjie Huang, Dongchao Yang, Wenxiang Guo, Huadai Liu, Xize Cheng, Zehan Wang, Zhou Zhao, Xixin Wu and Helen M. Meng	37953
<i>Orthogonal Representation Editing: Decoupling Semantic Entanglement in Batch Knowledge Editing of LLMs</i> Wenhao Yu, Zhicong Lu, Bo Lv, Fangyin Ma, Kaiwen Wei, Shihao Yang and Nayu Liu ...	37967
<i>Debiasing LLMs by Masking Unfairness-Driving Attention Heads</i> Tingxu Han, Wei Song, Ziqi Ding, Ziming Li, Chunrong Fang, Yuekang Li, Dongfang Liu, Zhenyu Chen and Zhenting Wang	37980
<i>Single-Agent Generation Surpasses Multi-Agent Systems in Semantic Diversity</i> Cui Encheng, Shaowen Peng, Kazuhiro Ito, XU Jinsha, Hisada Shohei, Shoko Wakamiya and Eiji Aramaki	37993
<i>Med-SRAF: A Multi-Agent Framework for Medical Reasoning via Semantic Routing and Agentic Fusion</i> Xiao Li, Zhuo Chen, Jun Xia, Hongxin Xiang, Chao Wang, Wenjie Du and Yang Wang ...	38009
<i>R-Search: Empowering LLM Reasoning with Search via Multi-Reward Reinforcement Learning</i> Qingfei Zhao, Ruobing Wang, Dingling Xu, Daren Zha, Ma Bowen, Zhichun Wang, Shijie Jia, Limin Liu and Xin Wang	38030
<i>Decoupled Reasoning with Implicit Fact Tokens (DRIFT): A Dual-Model Framework for Efficient Long-Context Inference</i> Wenxuan Xie and Xuhong Wang	38047

<i>AnchorAlign: A Novel Anchor Alignment-enhanced Generative Method for Joint Named Entity Recognition and Relation Extraction</i>	
Xiaolong Weng, Yuanyun Zhou, Boyu Qiu, Zehua Wang, Ying Xiong and Buzhou Tang . .	38068
<i>Why Can Distillation Work with Limited Resources? A Systematic Study</i>	
Xiao Hu, Xingyu Lu, Liyuan Mao, YiFan Zhang, Tianke Zhang, Bin Wen, Fan Yang, Tingting Gao and Guorui Zhou	38082
<i>Autoregressive Semantic Visual Reconstruction Helps VLMs Understand Better</i>	
Dianyi Wang, Wei Song, Yikun Wang, Siyuan Wang, Kaicheng Yu, Zhongyu Wei and Jiaqi Wang	38101
<i>MSCoDe: Advancing Human Motion-Language Understanding via Modality-Shared Codebook</i>	
Haoyu Shi and Huaiwen Zhang	38116
<i>RAG over Tables: Hierarchical Memory Index, Multi-Stage Retrieval, and Benchmarking</i>	
Jiaru Zou, Dongqi Fu, Sirui Chen, Xinrui He, Zihao Li, Yada Zhu, Jiawei Han and Jingrui He	38133
<i>CoTEvol: Self-Evolving Chain-of-Thoughts for Data Synthesis in Mathematical Reasoning</i>	
Zhuo Wang, Zhuo Zhang, Yafu Li, Yu Cheng, Lizhen Qu and Zenglin Xu	38153
<i>Feedback Is The Key for Automated Survey Generation</i>	
Tianyi Xu, Zhe Zhao, Tianshuo Wei, Yiqun Kou, Liuliu Han and Ye Wei	38174
<i>Sentient Agent as a Judge: Evaluating Higher-Order Social Cognition in Large Language Models</i>	
Bang Zhang, Ruotian Ma, Qingxuan Jiang, Peisong Wang, Jiaqi Chen, Zheng Xie, Xingyu Chen, Yue Wang, Fanghua Ye, Jian Li, Yifan Yang, Zhaopeng Tu and Xiaolong Li	38196
<i>Alignment Data Map for Efficient Preference Data Selection and Diagnosis</i>	
Seohyeong Lee, Eunwon Kim, Hwaran Lee and Buru Chang	38225
<i>Fine-Grained Detection of Context-Grounded Hallucinations Using LLMs</i>	
Yehonatan Peisakhovsky, Zorik Gekhman, Yosi Mass, Liat Ein-Dor and Roi Reichart	38242
<i>When Reviews Disagree: Fine-Grained Contradiction Analysis in Scientific Peer Reviews</i>	
Sandeep Kumar, Yash Kamdar, Abid Hossain, Bharti Kumari, Tanik Saikh and Asif Ekbal	38267
<i>One Script Instead of Hundreds? On Pretraining Romanized Encoder Language Models</i>	
Benedikt Ebing, Lennart Keller and Goran Glavaš	38291
<i>Relaxing the Constraints: A Dual-Importance Projection Mechanism for Lifelong Model Editing</i>	
Zhenghai Chen, Senbin Xu, Jiayi Tan, Xinhua Wu, Yan Zhang, Xiawu Zheng, Shengchuan Zhang, Ke Li, Sicheng Zhao, Liujuan Cao and Rongrong Ji	38308
<i>Not all ANIMALs are equal: metaphorical framing through source domains and semantic frames</i>	
Yulia Otmakhova, Matteo Guida and Lea Frermann	38334
<i>DualFact+: A Multimodal Fact Verification Framework for Procedural Video Captioning</i>	
Cennet Oguz, Yasser Hamidullah, Josef Van Genabith and Simon Ostermann	38356
<i>From Scaffolding to Assimilation: Progressive Structural Internalization for Format-Constrained Creative Text Generation</i>	
Wenhao Li, Yuwei Yang, Xiaoqing Wu, Yufeng Han, Cunliang Kong, Yuzhuo Bai, Xin Cong and Maosong Sun	38372

<i>Hierarchical Visual Agent: Managing Contexts in Joint Image-Text Space for Advanced Chart Reasoning</i>	
Qihua Dong, Ruozhen He, Junwen Chen, Yizhou Wang, Xu Ma, Songyao Jiang and Yun Fu	38390
<i>Semantic-Space Exploration and Exploitation in RLVR for LLM Reasoning</i>	
Fanding Huang, Guanbo Huang, Xiao Fan, Yi He, Xiao Liang, Xiao Chen, Qinting Jiang, Faisal Nadeem Khan, Jingyan Jiang and Zhi Wang	38402
<i>Towards Efficient Large Language Model Serving: A Survey on System-Aware KV Cache Optimization</i>	
Jiantong Jiang, Peiyu Yang, Rui Zhang and Feng Liu	38450
<i>A Counterfactual Explanation Framework for Retrieval Models</i>	
Bhavik Chandna and Procheta Sen	38477
<i>Teaching Language Models to Forecast Research Success Through Comparative Idea Evaluation</i>	
Srujan P Mule, Aniketh Garikaparthi and Manasi Patwardhan	38491
<i>Social Welfare Function Leaderboard: On the Emergence of LLM Agents as the Welfare Dictator</i>	
Zhengliang Shi, Ruotian Ma, Jen-tse Huang, Xinbei Ma, Xingyu Chen, Mengru Wang, Qu Yang, Yue Wang, Fanghua Ye, Ziyang Chen, Shanyi Wang, Cixing LI, Wenxuan Wang, Zhaopeng Tu, Xiaolong Li, Zhaochun Ren and Liefeng Bo	38530
<i>From Individual Excellence to Collective Sustainability: Seeking Strategic Equilibrium in Proactive Multi-Agent Teams</i>	
Tong Zhang, Yang Wu, Yufei Shi, Rujing Yao, Zhuoren Jiang and Xiaozhong Liu	38552
<i>Lost in the Prompt Order: Revealing the Limitations of Causal Attention in Language Models</i>	
Hyunjong Ok and Jaeho Lee	38566
<i>SR-RAG: Verifiable Multi-Hop Reasoning via On-the-fly Symbolic Graph Construction</i>	
Zehua Wang, Zhaojin Zhang, Boyu Qiu, Xiaolong Weng, Ying Xiong, Buzhou Tang and Min Zhang	38588
<i>KnowDR-REC: Auditing Knowledge-Conditioned Visual Grounding in Referring Expression Comprehension</i>	
Guanghao Jin, Jingpei Wu, Tianpei Guo, Yiyi Niu, Weidong Zhou, Linyi Yang and Guoyang Liu	38607
<i>Agent for Numerical Data Retrieval and Understanding by Code Generation and Multimodal Reasoning</i>	
Florian Baud, Feda Almuhsen and Dorian Midou	38630
<i>Cause-CSD: A Challenge Multimodal Conversational Stance Cause Detection Dataset and Effective Method</i>	
Fuqiang Niu, Bowen Zhang, Juntong Zhu, Qing Liao, Genan Dai and Hu Huang	38654
<i>The Sparse Frontier: Sparse Attention Trade-offs in Transformer LLMs</i>	
Piotr Nawrot, Jianing Li, Renjie Huang, Sebastian Ruder, Kelly Marchisio and Edoardo Ponti	38667
<i>Uncertainty-Aware Contrastive Sentence Embedding With Local Context Representation for Text Classification</i>	
Han Liu, Jiaqing Zhan, Zhichao Chen and Qin Zhang	38702
<i>ZipVoice-Dialog: Non-Autoregressive Spoken Dialogue Generation with Flow Matching</i>	
Han Zhu, Wei Kang, Liyong Guo, Zengwei Yao, Fangjun Kuang, Weiji Zhuang, Zhaoqing Li, Zhifeng Han, Dong Zhang, Xin Zhang, Xingchen Song, Lingxuan Ye, Long Lin and Daniel Povey	38717

<i>All Prompts Are Created Equal? Evaluating Robustness of LLM Judges Against Non-Adversarial Prompt Variations</i>	
Savita Bhat and Vasudeva Varma	38730
<i>Thinking Traps in Long Chain-of-Thought: A Measurable Study and Trap-Aware Adaptive Restart</i>	
Chenkang, Fan Yu, Junjie Nian, Sihan Zhao, Zhuoka Feng, Zijun Yao, Wang Heng, Yu Minshen and Yixin Cao	38746
<i>Instruction-Guided Poetry Generation in Arabic and Its Dialects</i>	
Abdelrahman Sadallah, Kareem Elozeiri, Mervat Abassy, Rania Elbadry, Mohamed Anwar, Abed Alhakim Freihat, Preslav Nakov and Fajri Koto	38761
<i>Agentic Verification for Ambiguous Query Disambiguation</i>	
Youngwon Lee, Seung-won Hwang, Ruofan Wu, Feng Yan, Danmei Xu, Moutasem Akkad, Zhewei Yao and Yuxiong He	38781
<i>DICA: Dual-Indicator Guided Contrastive Alignment in Multimodal Large Language Models</i>	
Hao Yang, Jin Wang and Xuejie Zhang	38797
<i>From Tasks to Teams: A Risk-First Evaluation Framework for Multi-Agent LLM Systems in Finance</i>	
Zichen Chen, Jianda Chen, Jiaao Chen and Misha Sra	38819
<i>Parametric Knowledge is Not All You Need: Toward Honest Large Language Models via Retrieval of Pretraining Data</i>	
Christopher Adrian Kusuma, Muhammad Reza Qorib and Hwee Tou Ng	38858
<i>TIME: Temporally Intelligent Meta-reasoning Engine for Context-Triggered Explicit Reasoning</i>	
Susmit Das	38872
<i>Omni-RI: Towards the Unified Generative Paradigm for Multimodal Reasoning</i>	
Dongjie Cheng, Yongqi Li, Zhixin Ma, Hongru Cai, Yupeng Hu, Wenjie Wang, Liqiang Nie and Wenjie Li	38907
<i>Time-for-Accuracy: Formalizing Chain-of-Thought as an Expansion of Logical Depth</i>	
Yue Wang, Zhi Zhang, Wang Xi, Chengjie Sun, Lili Shan and Bingquan Liu	38925
<i>DiffCoT: Diffusion-styled Chain-of-Thought Reasoning in LLMs</i>	
Shidong Cao, Hongzhan Lin, Yuxuan Gu, Ziyang Luo and Jing Ma	38942
<i>System-Mediated Attention Imbalances Make Vision-Language Models Say Yes</i>	
Tsan Tsai Chan, Varsha Suresh, Anisha Saha, Michael Hahn and Vera Demberg	38958
<i>SCURank: Ranking Multiple Candidate Summaries with Summary Content Units for Enhanced Summarization</i>	
Bo-Jyun Wang, Ying-Jia Lin and Hung-Yu Kao	38980
<i>Valid \neq Necessary: Diagnosing Latent Inefficiency in Chain-of-Thought</i>	
Daeyeop Lee and Hwanjo Yu	38998
<i>CausalGaze: Unveiling Hallucinations via Counterfactual Graph Intervention in Large Language Models</i>	
Linggang Kong, Lei Wu, Yunlong Zhang, Xiaofeng Zhong, Wang Zhen, Yongjie Wang and Yao Pan	39015
<i>Improved Policy Optimization for Mixture-of-Experts Models: Importance Sampling and Rewarding from an Expert-Centric Perspective</i>	
Yining Qian, Jinpeng Li, Fei Mi, Lifeng Shang and Xiang Zhang	39029

<i>From What Is Said to Why It Is Framed: Intent-Aware News Video Understanding</i>	
Xiangzheng Kong, Minnan Luo, Wenya Wang, Jiaying Wu, Zhi Zeng and Guang Dai	39039
<i>H-MAS: Hierarchical Multi-Agent Scheduling for Multi-Tenant LLM Serving</i>	
Yuhan Liu, Cong Xu, Qi Jia, Yihua Wang, Feiyu Chen, Liang Jin, Lu Liu, Yaqian Zhao, Yuting Ding and Xiang Li	39051
<i>EvoMD-LLM: Learning the Language of Species Evolution in Reactive Molecular Dynamics</i>	
Zhichen Tang, Zhengzheng Dang, Yulin Chen, Jixin Wu, Haiwen Li and Yanming Wang . .	39072
<i>Do Image-Text Metrics Respect Semantic Invariances?</i>	
Amit Agarwal, Hitesh Laxmichand Patel, Meizhu Liu, Jyotika Singh, Karan Dua, Hansa Megh-wani, Matthew Rowe, M. Avendi, Yassi Abbasi, Tao Sheng, Sujith Ravi and Dan Roth	39089
<i>Reasoning in the Dark: Interleaved Vision-Text Reasoning in Latent Space</i>	
Chao Chen, Zhixin Ma, Yongqi Li, Yupeng Hu, Yinwei Wei, Wenjie Li and Liqiang Nie . .	39117
<i>Hail to the Thief: Exploring Attacks and Defenses in Decentralised GRPO</i>	
Nikolay Blagoev, Oguzhan Ersoy and Lydia Chen	39130
<i>SocraticKG: Knowledge Graph Construction via QA-Driven Fact Extraction</i>	
Sanghyeok Choi, Woosang Jeon, Kyuseok Yang and Taehyeong Kim	39149
<i>GROLE: Instance-Level Group Relative Optimization for LoRA Experts in Incremental Learning</i>	
Yongyi Liao, Wencan Lai, Jun Fang, Jinjin Guo, Xiaohui Zhang, Zhiyuan Liu, Chao Liu, Pengzhang Liu and Qixia Jiang	39170
<i>SEARCH-R: Structured Entity-Aware Retrieval with Chain-of-Reasoning Navigator for Multi-hop Question Answering</i>	
FU Yuqing, Yimin Deng, Wanyu Wang, Yuhao Wang, Yejing Wang, Hongshi Liu, Yiqi Wang, Xiao Han, Maolin Wang, Guoshuai Zhao, Yi Chang and Xiangyu Zhao	39183
<i>Why Did Apple Fall: Evaluating Curiosity in Large Language Models</i>	
Haoyu Wang, Sihang Jiang, Yuyan Chen, Yitong Wang, Xiaojun Meng, Jiansheng Wei and Yanghua Xiao	39197
<i>Gaperon: A Peppered English-French Generative Language Model Suite</i>	
Nathan Godey, Wissam Antoun, Rian Touchent, Rachel Bawden, Éric Villemonte de la Clergerie, Benoît Sagot and Djamé Seddah	39216
<i>Rhombus: Incentivizing Coordination in Parallel Thinking through Reinforcement Learning</i>	
Ziyuan Nan, Qi Yi, Di Huang, Yutong Wu, Guanhua Huang, Xue Gong, Kejiao Li, Yuhao Jiang, Chenchen Zhang, Zenan Xu, Xing Hu and Bo Zhou	39258
<i>FLARE: Task-Agnostic Embedding Model Evaluation via Normalizing Flows</i>	
Jingzhou Jiang, Yixuan Tang, Yi Yang and Kar Yan Tam	39271
<i>RC TEA: Richness-guided Co-training for Temporal Entity Alignment</i>	
Jiayun Li, Wen Hua, Shiqi Fan, Fengmei Jin, Haiyang Jiang and Xue Li	39295
<i>SQLAgent: Learning to Explore Before Generating as a Data Engineer</i>	
Wenjia Jiang, Yiwei Wang, Boyan Han, Joey Tianyi Zhou and Chi Zhang	39311
<i>Reasoning Up the Instruction Ladder for Controllable Language Models</i>	
Zishuo Zheng, Vidhisha Balachandran, Chan Young Park, Faeze Brahman and Sachin Kumar	39332

<i>CRISP: Compressing Redundancy in Chain-of-Thought via Intrinsic Saliency Pruning</i> Yangsong Lan, Hongliang Dai and Piji Li	39355
<i>CascadeFix: Multi-Location Program Repair via Cascading Planning and Generation</i> Huan Zhang, Li Kuang, Yang Yang, Yilei Fang and Yingjie Xia	39374
<i>DGPO: Beyond Pairwise Preferences with Directional Consistent Groupwise Optimization</i> Mengyi Deng, Zhiwei Li, Xin Li, Tingyu Zhu, Yulan Yuan, Zhijiang Guo and Wei Wang ..	39386
<i>Emergence and Localisation of Semantic Role Circuits in LLMs</i> Nura Aljaafari, Danilo Carvalho and Andre Freitas	39402
<i>Thinking Economically: A Hierarchical Framework for Adaptive-Complexity Reasoning in LLMs</i> Yubo Gao, Haotian Wu, Hong Chen, Junquan Huang, Yibo Yan, Jungang Li, Zihao Dongfang, Sicheng Tao, PS Tan, Jie Zhang and Xuming Hu	39434
<i>Characterizing the Robustness of Black-Box LLM Planners Under Perturbed Observations with Adaptive Stress Testing</i> Neeloy Chakraborty, John Pohovey, Melkior Ornik and Katherine Rose Driggs-Campbell ..	39445
<i>ProMCP: Profiling Token Flows and Latency Costs in Model Context Protocol-Based LLM Agents</i> Sumera Anjum, Weijian Zheng, Rajkumar Kettimuthu, Heng Fan and Yunhe Feng	39476
<i>Reducing Hallucinations in LLMs via Factuality-Aware Preference Learning</i> Sindhuja Chaduvula, Ahmed Y. Radwan, Azib Farooq, Yani Ioannou and Shaina Raza	39488
<i>POLARIS: A Gödel Agent Framework for Small Language Models through Experience-Abstracted Policy Repair</i> Aditya Namdev Kakade, Vivek Srivastava and Shirish Karande	39505
<i>PhaseMI: A Motivational Interviewing Dataset for Enhancing Phase Progression in LLM-based Counseling</i> Jina Kim, Myeongho Jeon, Soohyun Cho, Chae-Gyun Lim, Jongmin Lim, Haewon Min and Eunho Yang	39538
<i>PsychEthicsBench: Evaluating Large Language Models Against Australian Mental Health Ethics</i> Yaling Shen, Stephanie Fong, Yiwen Jiang, Zimu Wang, Feilong Tang, Qingyang Xu, Xiangyu Zhao, Zhongxing Xu, Jiahe Liu, Jinpeng Hu, Dominic Dwyer and Zongyuan Ge	39571
<i>DynamixSFT: Dynamic Mixture Optimization of Instruction Tuning Collections</i> Haebin Shin, Lei Ji, Xiao Liu, Zhiwei Yu, Hyunwoo Yoo, Qi Chen and Yeyun Gong	39590
<i>BMAM: Brain-inspired Multi-Agent Memory Framework</i> Yang Li, Jiaxiang Liu, Yusong Wang, Yujie Wu and Mingkun Xu	39604
<i>Verified Critical Step Optimization for LLM Agents</i> Mukai Li, Qingcheng Zeng, Tianqing Fang, Zhenwen Liang, Linfeng Song, Qi Liu, Haitao Mi and Dong Yu	39627
<i>Nash-Pruned CredMAS: Dynamic Panel Pruning for VLM-MAS using Nash-based Selection and Doubly-Robust Credits</i> Yijia Fan, Mingyu Liu, Jing Yang, Jian Wang, Keze Wang and Jusheng Zhang	39640
<i>SWE-Mutation: Can LLMs Generate Reliable Test Suites in Software Engineering?</i> Yuxuan Sun, Yuze Zhao, Yufeng Wang, Yao Du, Zhiyuan Ma, Jinbo Wang, Mengdi Zhang, Kai Zhang and Zhenya Huang	39651

<i>STaD: Scaffolded Task Design for Identifying Compositional Skill Gaps in LLMs</i>	
Sungeun An, Swanand Ravindra Kadhe, Shailja Thakur, Chad DeLuca and Hima Patel . . .	39675
<i>PseudoSeer: a Search Engine for Pseudocode</i>	
Levent Toksoz, Mukund Srinath, Gang Tan and C. Lee Giles	39706
<i>Beyond Polarity: Continuous Affect-Enhanced Multimodal Aspect-Based Sentiment Classification</i>	
Ling-Ang Meng, Tianyu Zhao, Dawei Song, Jingxu Cao and Youhui Zuo	39717
<i>Topology Matters: Measuring Memory Leakage in Multi-Agent LLMs</i>	
Jinbo Liu, Defu Cao, Yifei Wei, Tianyao Su, Yuan Liang, Yushun Dong, Yan Liu, Yue Zhao and Xiyang Hu	39728
<i>Beyond 'Aha!': Toward Systematic Meta-Abilities Alignment in Large Reasoning Models</i>	
Zhiyuan Hu, Yibo Wang, Hanze Dong, Yuhui Xu, Amrita Saha, Caiming Xiong, Bryan Hooi and Junnan Li	39747
<i>Rewarding the Rare: Uniqueness-Aware RL for Creative Problem Solving in LLMs</i>	
Zhiyuan Hu, Yucheng Wang, Yufei He, Jiaying Wu, Yilun Zhao, See-Kiong Ng, Cynthia Breazeal, Anh Tuan Luu, Hae Won Park and Bryan Hooi	39765
<i>Empathy Applicability Modeling for General Health Queries</i>	
Shan Randhawa, Agha Ali Raza, Kentaro Toyama, Julie Hui and Mustafa Naseem	39791
<i>Worldwide LiveVQA: Real-Time Visual Knowledge Seeking and Updating Across Languages</i>	
Xuanao Huang, Xingjia Liu, Zetong Zhou, Yuyang Peng, Yao Wan and Dongping Chen . .	39819
<i>HCRE: LLM-based Hierarchical Classification for Cross-Document Relation Extraction with a Prediction-then-Verification Strategy</i>	
Guoqi Ma, Liang Zhang, Hongyao Tu, Hao Fu, Hui Li, Yujie Lin, Longyue Wang, Weihua Luo and Jinsong Su	39895
<i>VLMGuard-R1: Proactive Safety Alignment for VLMs via Reasoning-Driven Prompt Optimization</i>	
Menglan Chen, Xianghe Pang, Jingjing Dong, WenHao Wang, Yaxin Du and Siheng Chen	39914
<i>TinyAttack: Exploring Stylistic Vulnerabilities in Large Language Models</i>	
Mamta Mamta, Bogdan Grecu and Oana Cocarascu	39933
<i>Human-Agent Collaborative Paper-to-Page Crafting</i>	
Qianli Ma, Siyu Wang, Chen Yilin, Yin hao Tang, Yixiang Yang, Chang Guo, Bingjie Gao, Zhen- ing Xing, Yanan Sun and Zhipeng Zhang	39963
<i>Skill Weaving: Efficient LLM Improvement via Modular Skillpacks</i>	
Zhuo Li, Guodong DU, Zesheng Shi, Weiyang Guo, Weijun Yao, Yuan Zhou, Jiabo Zhang and Jing Li	40000
<i>Controlling What You Share: Assessing Language Model Adherence to Privacy Preferences</i>	
Guillem Ramírez, Alexandra Birch and Ivan Titov	40024
<i>GerAV: Towards New Heights in German Authorship Verification using Fine-Tuned LLMs on a New Benchmark</i>	
Lotta Kiefer, Christoph Leiter, Sotaro Takeshita, Elena Schmidt and Steffen Eger	40050
<i>LipoAgent: Coordinating Fine-Tuned LLM Agents for Safer Lipid Design</i>	
Leshu Li, An Lu, Haiyu Wang, Zhibin Feng, Conghui Duan, Qing Bao, Zongmin Zhao and Sai Qian Zhang	40070

<i>PsychePass: Calibrating LLM Therapeutic Competence via Trajectory-Anchored Tournaments</i> Zhuang Chen, Dazhen Wan, Zhangkai Zheng, Guanqun Bi, Xiyao Xiao, Binghang Li and Minlie Huang	40082
<i>Generalization or Memorization? Multi-Agent vs. Baseline LLMs and AutoML Models for Tabular Classification</i> Aida Sanatizadeh, Sorouralsadat Fatemi, Reza Mousavi and Ahmed Abbasi	40099
<i>GanitLLM: Difficulty-Aware Bengali Mathematical Reasoning through Curriculum-GRPO</i> Shubhashis Roy Dipta, Khairul Mahbub and Nadia Najjar	40133
<i>Can MLLMs Reason Beyond Language? VisReason: A Comprehensive Benchmark for Vision-Centric Reasoning</i> Longteng Guo, Yifan Wang, Pengkang Huo, Tailai Chen, Yuze Wu, Jing Liu and Xinxin Zhu	40149
<i>Multimodal Dual-Path Decoding for Medical Report Generation</i> Jinghan Sun, Dong Wei, Zhihong Zhu, Yuyang Xue, Steven McDonagh and Xian Wu	40193
<i>PanoramaRAG: Enabling Consistent Global Topic Awareness in Graph-Based RAG</i> Ding Deng, Xiang Li, Yaqing Zhang, Meng Li and Xiting Wang	40205
<i>Universally Empowering Zeroth-Order Optimization via Adaptive Layer-wise Sampling</i> Fei Wang, Li Shen, Liang Ding, Chao Xue, Ye Liu and Changxing Ding	40218
<i>Speculative Decoding with a Speculative Vocabulary</i> Miles Williams, Young D. Kwon, Rui Li, Alexandros Kouris and Stylianos I. Venieris	40240
<i>Entropy-Aware Reshaping of Reinforcement Signals for Multi-Answer Reasoning</i> Zhi Li, Huidan Xu, Zhen Hu, Yali Du and Ying Liu	40255
<i>TripTide: A Benchmark for Adaptive Travel Planning under Disruptions</i> Priyanshu Karmakar, Soumyabrata Chaudhuri, Shubhojit Mallick, Manish Gupta, Abhik Jana and Shreya Ghosh	40269
<i>Rational Synthesizers or Heuristic Followers? Analyzing LLMs in RAG-based Question-Answering</i> Atharv Naphade	40293
<i>IDEA: An Interpretable and Editable Decision-Making Framework for LLMs via Verbal-to-Numeric Calibration</i> Yanji He, Yuxin Jiang, Yiwen Wu, Bo Huang, Jiaheng Wei and Wei Wang	40312
<i>An Answer is just the Start: Related Insight Generation for Open-Ended Document-Grounded QA</i> Saransh Sharma, Pritika Ramu, Aparna Garimella and Koyel Mukherjee	40333
<i>Behavioral Consistency Validation for LLM Agents: An Analysis of Trading-Style Switching through Stock-Market Simulation</i> Zeping Li, Guancheng Wan, Keyang Chen, Yu Chen, Yiwen Zhao, Philip Torr, Guangnan Ye, Zhenfei Yin and Hongfeng Chai	40356
<i>NITI: Neural Plan Concretization for Incremental Execution, Bridging and Trigger Inference from Unspecified Human Policies</i> Ayan Banerjee, Shomrik Barua Banerjee and Sandeep Gupta	40371
<i>Automatic Combination of Sample Selection Strategies for Few-Shot Learning</i> Branislav Pecher, Ivan Srba, Maria Bielikova and Joaquin Vanschoren	40385

<i>LearnAlign: Data Selection for LLM Reinforcement Learning with Improved Gradient Alignment</i> Shipeng Li, Zhiqin Yang, Shikun Li, Xiaobo Xia, Hengyu Liu, Xinghua Zhang, Gaode Chen, Dong Fang, Ying Tai and Zhe Peng	40417
<i>Auditing LLM Responses to Harmful Stereotypes Targeting Mental Health Groups</i> Arka Dutta, Rijul Magu, Sean Kim, Seohee Yoon, Munmun De Choudhury and Ashiqur R. Khud- aBukhsh	40435
<i>ROSCO-Omni: Multimodal LLM-Based Communication Understanding for Non- and Minimally-Speaking Autistic Individuals</i> Siddhant Bikram Shah and Kristina T. Johnson	40453
<i>LQM: Linguistically Motivated Multidimensional Quality Metrics for Machine Translation</i> Samar M. Magdy, Fakhraddin Alwajih, Abdellah EL Mekki, Wesam El-Sayed and Muhammad Abdul-Mageed	40470
<i>Many-Shot Scaling of In-Context Learning with Self-Generated Demonstrations</i> Zhengyao Gu, Henry Peng Zou, Yankai Chen, Aiwei Liu, Weizhi Zhang and Philip S. Yu .	40494
<i>Small Data, Big Noise: Adversarial Training for Robust ParameterEfficient Fine-Tuning</i> Eitan Cohen, Idan Simai and Uri Shaham	40509
<i>LLMs as Lab Engineers: A Benchmark for Analytical Method Lifecycle Management</i> Xiaoyi Chen, Mahsa Monshizadeh, Chaoqi Zhang, Jianjun Lang, Yang Wu, Genevieve Mortensen, Xiaozhong Liu and Haixu Tang	40533
<i>Too Fast, Too Shallow – LLMs, Including Reasoning LLMs, Are Unreliable Constitutional Reasoners</i> Reto Gubelmann and Peter Hongler	40554
<i>JTPRO: A Joint Tool–Prompt Reflective Optimization Framework for Language Agents</i> Sandip Ghoshal, Anshul Mittal, Jyotika Singh, Miguel Ballesteros, Weiyi Sun, Fang Tu, Shailen- der Singh, Yassine Benajiba, Fahad Shah, Sujeeth Bharadwaj, Sujith Ravi and Dan Roth	40573
<i>CANDICE: Agentic Causal Disentanglement with Class Conditional Knowledge Integration for Long Tailed Domain Generalization</i> Midhat Urooj, Ayan Banerjee and Sandeep Gupta	40596
<i>AdaptEvolve: Improving Efficiency of Evolutionary AI Agents through Adaptive Model Selection</i> Pretam Ray, Pratik Prabhajan Brahma, Zicheng Liu and Emad Barsoum	40625
<i>OmniCode: A Benchmark for Evaluating Software Development Agents</i> Atharv Sonwane, Eng-Shen Tu, Wei-Chung Lu, Claas Beger, Carter Larsen, Debjit Dhar, Simon Alford, Rachel Chen, Ronit Pattanayak, Tuan Anh Dang, Guohao Chen, Gloria Geng, Kevin Ellis and Saikat Dutta	40634
<i>Supplement Generation Training for Enhancing Agentic Task Performance</i> Young Min Cho, Daniele Bonadiman, Divya Bhargavi, Tamer Alkhoul, Salvatore Romeo, Dong- wei Jiang, Khushbu Pahwa, Yubin Ge, Etsuko Ishii, Monica Sunkara and Yi Zhang	40662
<i>Automatic Paper Analysis and Categorisation for Systematic Reviews with Combined Reasoning-Augmented SFT and DAPO RL</i> Michela Lorandi, Anya Belz, Simon Mille and Craig Thomson	40676
<i>MotifAgent: Learning Molecular Assembly through Multi-Agent Collaboration for Chemical Language Understanding</i> Jinjia Feng, Wenda Wang and Zhewei Wei	40700

<i>Policy-Guided Stepwise Action Planning for Controllable LLM Reasoning</i> Jianpeng Zhou, Qisheng Hu, Jiahai Wang and Wenya Wang	40740
<i>Multiplication in Multimodal LLMs: Computation with Text, Image, and Audio Inputs</i> Samuel Gideon Balter, Ethan Jerzak and Connor Thomas Jerzak	40766
<i>Modeling Human Perspectives with Socio-Demographic Representations</i> Leixin Zhang and Cagri Coltekin	40781
<i>Weight Tying Biases Token Embeddings Towards the Output Space</i> Antonio Lopardo, Avyukth Harish, Catherine Arnett and Akshat Gupta	40795
<i>Structured Uncertainty guided Clarification for LLM Agents</i> Manan Suri, Puneet Mathur, Nedim Lipka, Franck Dernoncourt, Ryan A. Rossi and Dinesh Manocha	40811
<i>LRBench and Judge-RI: Principled Evaluation and Training of LLM-Based Judges for Long-Context Reasoning</i> Xinyi Zhao, Haoqi Hu, Ziyu Wang and Jinfeng Xiao	40839
<i>Beyond Reproduction: A Paired-Task Framework for Assessing LLM Comprehension and Creativity in Literary Translation</i> Ran Zhang, Steffen Eger, Arda Tezcan, Wei Zhao, Simone Paolo Ponzetto and Lieve Macken	40862
<i>Mechanistic Interpretability of Large-Scale Counting in LLMs through a System-2 Strategy</i> Hosein Hasani, Mohammadali Banayeeanzade, Ali Nafisi, Sadeh Mohammadian, Fatemeh Askari, Mobin Bagherian, Amirmohammad Izadi and Mahdieh Soleymani Baghshah	40885
<i>CodeScout: Contextual Problem Statement Enhancement for Software Agents</i> Manan Suri, Xiangci Li, Mehdi Shojaie, Songyang Han, Chao-Chun Hsu, Shweta Garg, Aniket Anand Deshmukh and Varun Kumar	40902
<i>Evaluating Scene-based In-Situ Item Labeling for Immersive Conversational Recommendation</i> Jiazhou Liang, Yifan Simon Liu, David Guo, Yilun Jiang, Minqi Sun and Scott Sanner ...	40932
<i>StanceAttack: Adversarial Attack for Stance Detection</i> Chenye Zhao and Cornelia Caragea	40954
<i>MEAV: Model Editing with Alignment Vectors for inference time LLM alignment in single and multidomain preference spectrum</i> Sadat Shahriar, Zheng Qi, Nikolaos Pappas, Srikanth Doss, Kishaloy Halder, Monica Sunkara, Manuel Mager and Yassine Benajiba	40972
<i>BioVLM: Routing Prompts, Not Parameters, for Cross-Modality Generalization in Biomedical VLMs</i> Mainak Singha, Tanisha Gupta, Ankit Jha, Muhammad Haris Khan, Sayantani Ghosh and Biplab Banerjee	40986
<i>SCALE: Upscaled Continual Learning of Large Language Models</i> Jin-woo Lee, Junhwa Choi, Bongkyu Hwang, Jinho Choo, Bogun Kim, Jeongseon Yi, Joonseok Lee, DongYoung Jung, Jaeseon Park, Kyoungwon Park and Suk-hoon Jung	41006
<i>Words that make SENSE: Sensorimotor Norms in Learned Lexical Token Representations</i> Abhinav Gupta, Toben Mintz and Jesse Thomason	41021
<i>Interpreting Style Representations via Style-Eliciting Prompts</i> Junghwan Kim and David Jurgens	41030

<i>Optimizing Conversational Quality in Spoken Dialogue Systems with Reinforcement Learning from AI Feedback</i>	
Siddhant Arora, Jinchuan Tian, Jiatong Shi, Hayato Futami, Yosuke Kashiwagi, Emiru Tsunoo and Shinji Watanabe	41049
<i>Conceptual Hierarchies within LLMs</i>	
Tiago Almeida, Zining Zhu and Yue Ning	41067
<i>Conjecture and Inquiry: Quantifying Software Performance Requirements via Interactive Retrieval-Augmented Preference Elicitation</i>	
Shihai Wang and Tao Chen	41080
<i>Measuring Large Language Models' Adversarial Behavior in Social Deduction Games</i>	
Marissa Zhao Li, Esha Shivakumar, Peiran Wang, Ying Li and Yuan Tian	41099
<i>ChangJuan: A Comprehensive Benchmark for Book-Length Chinese Story Evaluation</i>	
Dingyi Yang, Mingshuo Wang and Qin Jin	41116
<i>Progressive Re-ranking for Multimodal Retrieval-Augmented Generation via Curriculum Learning</i>	
Zhu Min, Yanchao Hao, Jian Liu, Shizhu He and Xi Chen	41135
<i>Confidence-Weighted Token Set Cover for Early Hypothesis Pruning in Self-Consistency</i>	
Md Arafat Sultan and Ramón Fernandez Astudillo	41148
<i>MMTabReal: Real-World Benchmark for Multimodal Table Understanding</i>	
Prasham Yatinkumar Titiya, Jainil Trivedi, Chitta Baral and Vivek Gupta	41156
<i>Inject to Heal: Alleviating hallucination in LVLMs via Context Embedding Injection</i>	
Mehrdad Fazli, Bowen Wei and Ziwei Zhu	41177
<i>Colorful Talks with Graphs: Human-Interpretable Graph Encodings for Large Language Models</i>	
Angelo Zangari, Peyman Bagher Shahi and Sourav Medya	41194
<i>SAM-NER: Semantic Archetype Mediation for Zero-Shot Named Entity Recognition</i>	
Ruichu Cai, Juntao Gan, Miao Mai, Zhifeng Hao and Boyan Xu	41216
<i>DRPG (Decompose, Retrieve, Plan, Generate): An Agentic Framework for Academic Rebuttal</i>	
Peixuan Han, Yingjie Yu, Jingjun Xu and Jiaxuan You	41232
<i>BiasGRPO: Stabilizing Bias Mitigation in High-Variance Reward Landscapes via Group-Relative Policy Optimization</i>	
Saket Reddy, Ke Yang and ChengXiang Zhai	41250
<i>MaskTab: Scalable Masked Tabular Pretraining with Scaling Laws and Distillation for Industrial Classification</i>	
Bo Zheng, Yudong Chen, Zihua Xiong, Shuai Fang, Peidong He, Yang Yang and Sheng Guo	41268
<i>Enhancing Zero-Shot Time Series Forecasting in Off-the-Shelf LLMs via Noise Injection Prompting</i>	
Xingyou Yin, Ceyao Zhang, Min Hu and Kai Chen	41281
<i>Text Embedding as Treatment: A Meta Causal Approach for Robust Sentiment Classification</i>	
Fengxiang Cheng, Chuan Zhou, Xiang Li, Haoxuan Li, Wen-li Wang, Jinkun Chen, Mingming Gong and Kun Zhang	41309
<i>CliniCAST: Benchmarking Acoustic Grounding and Text Dominance in Medical Triage</i>	
Kyusik Kim, Hyunwoo Yoo, Jaehoon Choi, Kitae Kim, Gail Rosen and Bongwon Suh	41321

<i>Whose Voice, Whose Avatar? Gender Matching Bias in Multimodal AI Teammates</i>	
Kyusik Kim, Jaehoon Choi, Hyunwoo Yoo and Bongwon Suh	41344
<i>WebAnchor: Anchoring Agent Planning to Stabilize Long-Horizon Web Reasoning</i>	
Xinmiao Yu, Liwen Zhang, Xiaocheng Feng, Pengjun Xie, Jingren Zhou, Bing Qin and Yong Jiang	41368
<i>Context-Driven and Reference-Guided Data Augmentation for Subtitle Translation</i>	
Hitoshi Ito, Naoto Shirai, Kazutaka Kinugawa, Hideya Mino, Rei Endo and Yoshihiko Kawai	41381
<i>Real or Robotic? Assessing Whether LLMs Accurately Simulate Qualities of Human Responses in Human-LLM Dialogue</i>	
Jonathan Ivey, Shivani Kumar, Jiayu Liu, Hua Shen, Sushrita Rakshit, Rohan Raju, Haotian Zhang, Aparna Ananthasubramaniam, Junghwan Kim, Bowen Yi, Dustin Wright, Abraham Israeli, Anders Giovanni Møller, Lechen Zhang and David Jurgens	41395
<i>Can Small LLMs Learn a Robust Theory of Mind via RLVR? Investigating Generalization through the False-Belief Task</i>	
Sneheel Sarangi and Hanan Salam	41433
<i>Rolling Out Data Quality Overnight, without losing the plot: A Multi-Agent System for Speech Data Quality Management</i>	
Rishabh Kumar, Abhinav Painuli, Chriss Philip Saji, Devesh Soni, Amrith Krishna and Ganesh Ramakrishnan	41449
<i>Schema-Guided Response Generation using Multi-Frame Dialogue State for Motivational Interviewing Systems</i>	
Jie Zeng and Yukiko Nakano	41493
<i>From Passive Metric to Active Signal: The Evolving Role of Uncertainty Quantification in Large Language Models</i>	
Jiaxin Zhang, Wendi Cui, Zhuohang Li, Lifu Huang, Bradley A. Malin, Caiming Xiong and Chien-Sheng Wu	41525
<i>Value–Action Alignment in Large Language Models under Privacy–Prosocial Conflict</i>	
Guanyu Chen, Chenxiao Yu and Xiyang Hu	41545
<i>See or Say Graphs: Agent-Driven Scalable Graph Understanding with Vision-Language Models</i>	
Shuo Han, Yukun Cao, Zezhong Ding, Zengyi Gao, S Kevin Zhou and Xike Xie	41565
<i>Recovered in Translation: Efficient Pipeline for Automated Translation of Benchmarks and Datasets</i>	
Hanna Yukhymenko, Anton Alexandrov and Martin Vechev	41590
<i>MorphBPE: Morphology-Aware Tokenization for Efficient LLM Training</i>	
Ehsaneddin Asgari, Yassine El Kheir, MohammadAli SadraeiJavaheri and Ali Nazari	41610
<i>From Storage to Experience: A Survey on the Evolution of LLM Agent Memory Mechanisms</i>	
Jinghao Luo, Yuchen Tian, Chuxue Cao, Ziyang Luo, Hongzhan Lin, Kaixin Li, Chuyi Kong, Ruichao Yang and Jing Ma	41622
<i>LLMs are Brittle to Simple Code Transformations: Introducing CETBench – A Benchmark for Code-Equivalence Checking</i>	
Neeva Oza, Ishaan Govil, Parul Gupta, Dinesh Khandelwal, Dinesh Garg and Parag Singla	41653

<i>Enhancing Factuality through Consensus and Consistency in Summarization Using Minimum Bayes Risk Decoding</i>	
Riza Setiawan Soetedjo, Yusuke Sakai, Hidetaka Kamigaito, Jingun Kwon, Manabu Okumura and Taro Watanabe	41686
<i>Moneyball with LLMs: Analyzing Tabular Summarization in Sports Narratives</i>	
Ritam Upadhyay, Naman Ahuja, Rishabh Baral, Aparna Garimella and Vivek Gupta	41714
<i>A Computational Method for Measuring Open Codes in Qualitative Analysis</i>	
John Chen, Alexandros Nikolaos Lotsos, Sihan Cheng, Lexie Zhao, Yanjia Zhang, Jessica Hullman, Bruce Sherin, Uri Wilensky and Michael Horn	41740
<i>Vulnerability of LLMs' Stated Belief? LLMs Belief Resistance Check Through Strategic Persuasive Conversation Interventions</i>	
Fan Huang, Haewoon Kwak and Jisun An	41759
<i>Choose Your Lens: Multi-Perspective Value Alignment of Chain-of-Thought Reasoning</i>	
Gejian Zhao, Hanzhou Wu and Xinpeng Zhang	41795
<i>Beyond Evidence: Belief-Chain Conditioning for Persuasive Misinformation Debunking Explanation</i>	
Yi-Li Hsu, Li-Wun Chang, Chun-Yu Hsu, Wei-Kuan Shih, Aiping Xiong and Lun-Wei Ku	41809
<i>CoTJuder: A Graph-Driven Framework for Automatic Evaluation of Chain-of-Thought Efficiency and Redundancy in LLMs</i>	
Siyi Li, Jiajun Shi, Shiwen Ni, Ge Zhang, Shuaimin Li, Shijian Wang, Zhoufutu Wen, Yizhi LI, Hamid Alinejad-Rokny, Jiaheng Liu, Min Yang and Wenhao Huang	41837
<i>The Inner Monologue of Language Models: When Reasoning Traces Reveal More Than They Hide</i>	
Pratham Singla, Shivank Garg, Ayush Singh, Ishan Garg and Ketan Suhaas Saichandran	41864
<i>Do VLMs Have a Moral Backbone? A Study on the Fragile Morality of Vision-Language Models</i>	
Zhining Liu, Tianyi Wang, Xiao Lin, Penghao Ouyang, Gaotang Li, Ze Yang, Hui Liu, Sumit Keswani, Vishwa Pardeshi, Huijun Zhao, Wei Fan and Hanghang Tong	41890
<i>ENPMR-Bench: Benchmarking Proactive Memory Retrieval for Emotional Support Agents</i>	
Xing Fu, Yulin Hu, Mengtong Ji, Haozhen Li, Yixin Sun, Weixiang Zhao, Yanyan Zhao and Bing Qin	41910
<i>Grounded Concreteness: Human-Like Concreteness Sensitivity in Vision–Language Models</i>	
Aryan Roy, Zekun Wang and Christopher J. MacLellan	41934
<i>PROPER Agents: Proactivity Driven Personalized Agents for Advancing Knowledge Gap Navigation</i>	
Kirandeep Kaur, Vinayak Gupta, Aditya Gupta and Chirag Shah	41951
<i>PO-KGQA: Preference Optimization for Low-Resource Complex Knowledge Graph Question Answering</i>	
Prerna Agarwal, Ayushman Kumar Singh and Srikanta Bedathur	41976
<i>Beyond Prompt Engineering: A Systematic Analysis of Prompt Lexical Sensitivity and Its Impacts on Quality</i>	
Qipeng Xie, Zi Liang, Jiafei Wu, Yufei Chen, Weizheng Wang, Wenao Ma, Zhong Ming, Haiqin Yang and Kaishun Wu	41998
<i>DUSK: Do Not Unlearn Shared Knowledge</i>	
Wonje Jeung, Sangyeon Yoon, Hyesoo Hong, Soeun Kim, Seungju Han, Youngjae Yu and Albert No	42013

<i>Towards Mitigating Hallucinations in Large Vision-Language Models by Refining Textual Embeddings</i> Aakriti Agrawal, Gouthaman KV, Rohith Aralikatti, Gauri Jagatap, Jiaxin Yuan, Sarvesh Baskar, Vijay Kamarshi, Andrea Fanelli and Furong Huang	42032
<i>Prior Beliefs Prejudice LLM-as-Judge: Evidence from Persuasion Evaluation</i> Pardis Sadat Zahraei, Xiaoning Wang, Nimet Beyza Bozdog, Gokhan Tur and Dilek Hakkani-Tür 42049	
<i>Direct Token Optimization: A Self-Contained Approach to Large Language Model Unlearning</i> Hong Kyu Lee, Ruixuan Liu and Li Xiong	42083
<i>EntroBench: Evaluating LLM Watermarking Under Multi-Entropy Scenarios and Practical User Operations</i> Pengyuan Qin, Linnan Tu, Yuhan Ke and Hefei Ling	42101
<i>Beyond Single-Shot: Multi-step Tool Retrieval via Query Planning</i> Wei Fang and James R. Glass	42119
<i>On Safety Risks in Experience-Driven Self-Evolving Agents</i> Weixiang Zhao, Yichen Zhang, Yingshuo Wang, Yang Deng, Yanyan Zhao, Xuda Zhi, Yongbo Huang, Hao He, Wanxiang Che, Bing Qin and Ting Liu	42145
<i>LLM Program Optimization via Retrieval Augmented Search</i> Sagnik Anupam, Alexander Shypula and Osbert Bastani	42170
<i>EnsemW2S: Enhancing Weak-to-Strong Generalization with Large Language Model Ensembles</i> Aakriti Agrawal, Mucong Ding, Chenghao Deng, Zora Che, Arjun Rajaram, Anirudh Satheesh, Bang An, C. Bayan Bruss, John Langford and Furong Huang	42187
<i>Stable On-Policy Distillation through Adaptive Target Reformulation</i> Ijun Jang, Jewon Yeom, Juan Yeo, Hyunggyu Lim and Taesup Kim	42217
<i>A Syntactic and Semantic Probe into Language Evolution based on Large Language Models</i> Hao Pang, Changcheng Li and Yingxue Liu	42228
<i>Assessing the Effect of Context in Multi-domain Acceptability Judgment</i> Eunike Andriani Kardinata, Yusuke Sakai and Taro Watanabe	42251
<i>Lost in Decomposition: Analyzing and Mitigating the Limitations of Long Context Methods via Context Dependency</i> Jiayuan Guo, Yueyang Su, Yuyao Ge, Saiping Guan, Lei Yu, Jiafeng Guo and Xueqi Cheng	42267
<i>Tandem: Riding Together with Large and Small Language Models for Efficient Reasoning</i> Zichuan Fu, Xian Wu, Guojing Li, Yejing Wang, Yijun Chen, Zhao Zihao, Luo Yixuan, Hanyu Yan, Yefeng Zheng and Xiangyu Zhao	42286
<i>PROBE: PROcess-Based BENCHMARK for Hallucination Detection</i> Yu Zhang, Peter Belcak, Shizhe Diao, Yonggan Fu, Shaona Ghosh, Morteza Mardani, Eileen Margaret Peters Long, Bei Yu and Pavlo Molchanov	42303
<i>Prompt Optimization for Relation Extraction using Reinforcement Learning</i> Ying Liu, Dong Shuai, Cui Zibo, TengQi Ye and Gang Wu	42321
<i>PolyAudio: Advancing Multi-Audio Reasoning in Large Audio Language Models with Interleaved Multi-Audio Contexts</i> Sonal Kumar, Sreyan Ghosh, Yueqian Lin, S Sakshi, Ashish Seth, Yiran Chen, Ramani Duraiswami and Dinesh Manocha	42335

<i>Scaling Evaluation-Time Compute with Reasoning Models as Evaluators</i>	
Seungone Kim, Ian Wu, Jinu Lee, Xiang Yue, Seongyun Lee, Minkyong Moon, Carolin Lawrence, Kiril Gashteovski, Julia Hockenmaier, Graham Neubig and Sean Welleck	42354
<i>Beyond Memorization: Extending Reasoning Depth with Recurrence, Memory and Test-Time Compute Scaling</i>	
Ivan Rodkin, Daniil Orel, Konstantin Smirnov, Arman Bolatov, Bilal Elbouardi, Beshar Hassan, Yuri Kuratov, Aydar Bulatov, Preslav Nakov, Timothy Baldwin, Artem Shelmanov and Mikhail Burtsev	42385
<i>Navigating the Alignment-Calibration Trade-off: A Pareto-Superior Frontier via Model Merging</i>	
Tiancheng Hu, Benjamin Minixhofer and Nigel Collier	42405
<i>From Representation to Choice: Tracing Decision Emergence Across Languages in LLMs</i>	
Divyanshu Bhatt, Abhinav Joshi, Ujjaval Patel and Ashutosh Modi	42423
<i>Unlocking the Edge deployment and ondevice acceleration of multi-LoRA enabled one-for-all foundational LLM</i>	
Sravanth Kodavanti, Sowmya Vajrala, Srinivas Soumitri Miriyala, Utsav Tiwari, Uttam Kumar, Utkarsh Kumar Mahawar, Achal Pratap Singh, Arya D, Narendra Mutyala, Vikram Nelvoy Rajendiran, Sharan Kumar Allur, Euntaik Lee, Dohyoung Kim, HyeonSu Lee, Gyusung Cho and JungBae Kim	42443
<i>SynthRL: Scaling Visual Reasoning with Verifiable Data Synthesis</i>	
Zijian Wu, Jinjie Ni, Xiangyan Liu, Zichen Liu, Hang Yan and Michael Qizhe Shieh	42456
<i>From Fragments to Facts: A Curriculum-Driven DPO Approach for Generating Hindi News Veracity Explanations</i>	
Pulkit Bansal, Raghvendra Kumar, Shakti Singh, Adam Jatowt and Sriparna Saha	42477
<i>Joint Optimization of Training Data and Policy in RLHF</i>	
Zhuohao Yu, Jiali Zeng, Weizheng Gu, Mengyuan Sun, Yidong Wang, Fandong Meng, Jie Zhou, Shikun Zhang and Wei Ye	42498
<i>Simul-COMET: A Quality Metric for Simultaneous Interpretation in Distant Language Pair Considering Word Order Difference</i>	
Kosuke Doi, Mana Makinae, Yusuke Sakai, Hidetaka Kamigaito and Taro Watanabe	42515
<i>Risk-Controlled Event-Driven Cascading Updates for Knowledge Graph Consistency Restoration</i>	
Bo Ni, Qinwen Ge, Haowei Fu, Ryan A. Rossi, Xiaorui Liu, Jiejun Xu and Tyler Derr	42534
<i>ContractEval: A Benchmark for Evaluating Contract-Satisfying Assertions in Code Generation</i>	
Soochan Lim, Joonghyuk Hahn, Hyunwoo Park, Sang-Ki Ko and Yo-Sub Han	42549
<i>ToM-Synth: Scaling Robust Theory of Mind in LLMs via 6,912 Structured Social Units</i>	
Guiyang Hou, Xiang Huang, Shangke Lyu, Yuchuan Wu, Weiyao Luo, Xinyu Mei, Yongliang Shen, Weiming Lu and Yongbin Li	42567
<i>Budget-Aware Routing for Long Clinical Text</i>	
Khizar Qureshi, Geoffrey Martin and Yifan Peng	42583
<i>Multi-LLM Collaborative Search for Complex Problem Solving</i>	
Sen Yang, Yafu Li, Wai Lam and Yu Cheng	42599
<i>SceneLM: 3D-Aware Language Models for Editable 3D Scene Synthesis</i>	
Xingbo Yao, Xiaoyu Chen, Doudou Zhang, Mingzhi Sheng, Boyuan Cao, Ying-Cong Chen and Hui Xiong	42615

<i>Leveraging Pretrained Language Models as Energy Functions for Glauber Dynamics Text Diffusion</i> Tarun Kathuria and Sachin Kumar	42625
<i>Uncovering Currency Bias and Syntax Gap in Text Embedding Models</i> Saurav Sudevan, Harsh Pal, Yatin Katyal and Sahil Manchanda	42644
<i>C-ReD: A Comprehensive Chinese Benchmark for AI-Generated Text Detection Derived from Real-World Prompts</i> Chenxi Qing, Junxi Wu, Zheng Liu, Yixiang Qiu, Hongyao Yu, Bin Chen, Hao Wu and Shu-Tao Xia	42703
<i>Seer Self-Consistency: Advance Budget Estimation for Adaptive Test-Time Scaling</i> Shiyu Ji, Yixuan Wang, Yijun Liu, Qingfu Zhu and Wanxiang Che	42734
<i>Context-Fidelity Boosting: Enhancing Faithful Generation through Watermark-Inspired Decoding</i> Weixu Zhang, Fanghua Ye, Qiang Gao, Jian Li, Haolun Wu, Yuxing Tian, Sijing Duan, Nan Du and Xiaolong Li	42748
<i>CaRVE: Critiquing and Refining Visual Elaborations for Figurative Language Illustrations</i> Manishit Kundu, Tejomay Kishor Padole, Sumit Shekhar, Biplab Banerjee and Pushpak Bhattacharyya	42760
<i>Defending LLMs against Jailbreak Attacks via Template-Based ICL with a Defensive Suffix</i> Ruiyang Ni, Changlong Li, Shuaibiao Han, Zhiyu Yi, Perley Xu and Wenjie Ruan	42778
<i>From Fallback to Frontline: When Can LLMs be Superior Annotators of Human Perspectives?</i> Hasan Amin, Harry Yizhou Tian, Xiaoni Duan, Chien-Ju Ho, Rajiv Khanna and Ming Yin	42798
<i>Does Theory of Mind Improvement Really Benefit Human-AI Interactions? Empirical Findings from Interactive Evaluations</i> Nanxu Gong, Zixin Chen, Haotian Li, Zishu Zhao, Jianxun Lian, Huamin Qu, Yanjie Fu and Xing Xie	42831
<i>SERE: Structural Example Retrieval for Enhancing LLMs in Event Causality Identification</i> Zhifeng Hao, Zhongjie Chen, Junhao Lu, Shengyin Yu, Guimin Hu, Keli Zhang, Ruichu Cai and Boyan Xu	42861
<i>Emotion Trajectory-aware Retrieval for Markov-driven Emotion Anticipation in LLM-based Emotional Support Conversation</i> Hongyan Wu, Zhiliang Tian, Zhen Huang, Tengyue Hu, Linbo Qiao, Yifu Gao, Feng Liu and Dongsheng Li	42887
<i>High-Throughput and Memory-Efficient Zeroth-Order Fine-tuning LLMs with Distributed Parallel Computing</i> Liangyu Wang, Huanyi Xie and Di Wang	42906
<i>AutoAgent: A Fully-Automated and Zero-Code Framework for LLM Agents</i> Jiabin Tang, Tianyu Fan and Chao Huang	42924
<i>DR-HM: Distill-then-Reinforce Training with Cognition-Aware Data Synthesis for Harmful Meme Detection</i> Zihan Cheng, Jianxiang Ma, Xiaocui Yang, Peidong Wang, Wen Zhang, Shi Feng, Daling Wang, Yifei Zhang and Mingfu Zhang	42975
<i>Long Story Short: Disentangling Compositionality and Long-Caption Understanding in Contrastive VLMs</i> Israfel Salazar, Desmond Elliott and Yova Kementchedjieva	42994

<i>Model in Distress: Sentiment Analysis on French Synthetic Social Media</i> Pierre-Carl Langlais, Pavel Chizhov, Yannick Detrois, Carlos Rosas Hinostroza, Ivan P. Yamshchikov and Bastien Perroy	43009
<i>Working Memory Constraints Scaffold Learning in Transformers under Data Scarcity</i> Pranava Madhyastha and Dagmar Adamcová	43022
<i>Internalizing Multi-Agent Reasoning for Accurate and Efficient LLM-based Recommendation</i> Yang Wu, Haoze Wang, Qian Li, Jun Zhang, Huan Yu and Jie Jiang	43039
<i>Do Language Models Use Logophoric Cues? Evidence from Mandarin Chinese Long-Distance Reflexive</i> Yunfang Dong	43061
<i>FineState-Bench: Benchmarking State-Conditioned Grounding for Fine-grained GUI State Setting</i> Fengxian Ji, Jingpu Yang, Zirui Song, Yuanxi Wang, Zhexuan Cui, Yuke Li, Qian Jiang and Xiuying Chen	43073
<i>Dictionary Guided Sparse Logit Editing for Reliable Jailbreak Attacks</i> Shuaibiao Han, Ruiyang Ni, Zhiyu Yi, Changlong Li, Perley Xu and Wenjie Ruan	43089
<i>Chain-of-Relations: Faithful and Efficient LLM Reasoning over Knowledge Graphs via Relation-Centric Exploration</i> Chenhui Liu, Jianpeng Zhou and Jiahai Wang	43108
<i>MedQPA-Gen: Medical Question Proposing and Answering for Report Generation</i> Weijie Liang, Xiyue Zhu, Ruike Zhu, Chenhao Li, Cheng Tang, Zhiyu Liu, Zhihua Gong, Shirui Luo, Yudu Li and Volodymyr Kindratenko	43120
<i>Faster MoE LLM Inference for Extremely Large Models</i> Haoqi Yang, Luohe Shi, Qiwei Li, Zuchao Li, Ping Wang, Hao Huang and Hai Zhao	43133
<i>Localized Cultural Knowledge is Conserved and Controllable in Large Language Models</i> Veniamin Veselovsky, Berke Argın, Benedikt Stroebel, Chris Wendler, Robert West, James Evans, Thomas L. Griffiths and Arvind Narayanan	43152
<i>The Confidence Paradox: Unveiling the Latent Discriminative Power of Diffusion Large Language Models in Mathematical Reasoning</i> Yansi Li, Gongshen Liu and Zhuosheng Zhang	43179
<i>Efficient Hallucination Detection in Automatic Code Generation</i> Georgii Andriushchenko, Roman Garaev, Lyudmila Rvanova, Artem Shelmanov and Vladimir V. Ivanov	43197
<i>DuET: Dual Execution for Test Output Prediction with Generated Code and Pseudocode</i> Hojae Han, Jaejin Kim, Seung-won Hwang, Yu Jin Kim and Moontae Lee	43221
<i>Attention Sinks in Diffusion Language Models</i> Maximo Eduardo Rulli, Simone Petrucci, Edoardo Michielon, Fabrizio Silvestri, Simone Scardapane and Alessio Devoto	43244
<i>R³-SQL: Ranking Reward and Resampling for Text-to-SQL</i> Hojae Han, Yeonseok Jeong, Seung-won Hwang, Zhewei Yao and Yuxiong He	43257
<i>Detecting Hallucinations in SpeechLLMs at Inference Time Using Attention Maps</i> Jonas Waldendorf, Bashar Awwad Shiekh Hasan and Evgenii Tsymbalov	43276

<i>Leveraging External Knowledge for Historical Document Restoration via Retrieval-Augmented Large Language Models</i>	
Gabeen Kim and Kyeongpil Kang	43290
<i>From Experience to Skill: Multi-Agent Generative Engine Optimization via Reusable Strategy Learning</i>	
Beining Wu, Fuyou Mao, Jiong Lin, Cheng Yang, Jiaxuan Lu, Yifu Guo, Siyu Zhang, Yifan Wu, Ying Huang and Fu Li	43305
<i>CodeRM-NT: Reward Model for Code RL without Unit Tests</i>	
Xiao Xia, Dan Zhang and Tianrui Sun	43316
<i>Fin-STAR: Structure-as-Semantics to Resolve Implicitness in Financial Retrieval</i>	
Yu Zou, Yan Chen, Lida He, Qi Zhou, Xiaorui Zhou, Aixi Zhong, Yi Wang, Wei Li, Qingyu Wang, Jiatao Li, Wei Gong, Jialei Zeng, Jingmei Zhao, Ke Jiang and Qing Li	43334
<i>When Rules Learn: A Self-Evolving Agent for Legal Case Retrieval</i>	
Mingxu Tao, Jiawei Hu, Xian Zhou, Wenpeng Hu, Jiajun Cheng, Yunbo Cao, Zhunchen Luo and Guotong Geng	43348
<i>SpecExtend: A Drop-in Enhancement for Speculative Decoding of Long Sequences</i>	
Jungyoub Cha, Hyunjong Kim and Sungzoon Cho	43366
<i>Task-Related In-Context Learning</i>	
Wenqiang Wang, Peng Chen, Yan Xiao, Yangshijie Zhang, Xiaoyue Lu, Jianjie Huang and Xiaochun Cao	43378
<i>Concise Math Reasoning via Difficulty-Aware Distillation</i>	
Yifan Wu, Jingze Shi, Bingheng Wu, Jiayi Zhang, Xiaotian Lin, Yizhang Zhu, Zhaoyang Yu, Bang Liu, Chenglin Wu, Nan Tang and Yuyu Luo	43401
<i>Corpora Generation for Urdu Grammatical Error Correction</i>	
Syed Ahad, Burhanuddin Aliasghar Ezzi, Muhammad Arsalan Hussain, Sandesh Kumar and Abdul Samad	43428
<i>Rethinking Retrieval-Augmented Generation as a Cooperative Decision-Making Problem</i>	
Lichang Song, Ting Long and Yi Chang	43445
<i>More Agents Improve Math Problem Solving but Adversarial Robustness Gap Persists</i>	
Khashayar Alavi, Zhastay Yeltay, Lucie Flek and Akbar Karimi	43457
<i>Indic-CodecFake meets SATYAM: Towards Detecting Neural Audio Codec Synthesized Speech Deep-fakes in Indic Languages</i>	
Girish, Mohd Mujtaba Akhtar, Orchid Chetia Phukan and Arun Balaji Buduru	43476
<i>Tracing Relational Knowledge Recall in Large Language Models</i>	
Nicholas Popovič and Michael Färber	43490
<i>LR-DWM: Efficient Watermarking for Diffusion Language Models</i>	
Ofek Raban, Gal Chechik and Ethan Fetaya	43510
<i>Beyond Unimodal Shortcuts: MLLMs as Cross-Modal Reasoners for Grounded Named Entity Recognition</i>	
Jinlong Ma, Yu Zhang, Xuefeng Bai, Kehai Chen, Yuwei Wang, Zeming Liu, Jun Yu and Min Zhang	43518
<i>BubbleRAG: Interactive Cognitive Offloading with Thought Bubble in Retrieval-Augmented Generation</i>	
Fuda Ye, Jiachuan Wang, Yongqi Zhang, Lei Chen and Shuangyin Li	43540