

ACL 2025

**The 63rd Annual Meeting of the Association for
Computational Linguistics (ACL 2025)**

Proceedings of the Conference - Volume 1: Long Papers

July 27 - August 1, 2025

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

Diamond



Platinum



Tencent

 **Toloka**

Gold

CHEN TIANQIAO
& CHRISSY
INSTITUTE

 **KUAISHOU**


**Thomson
Reuters™**

Silver

Appen

Adobe

 **nexdata**

Bronze

 **abelscape®**

DataoceanAI

DEFINE.D.AI

JPMorganChase



**NatWest
Group**

NEC

 **translated.**

Diversity and Inclusion



©2025 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 979-8-89176-251-0

Message from the General Chair

It is my great pleasure and honor to welcome you to the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025), held in beautiful Vienna, Austria, from July 27 to August 1, 2025. ACL 2025 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.

I want to express my deepest thanks to our Program Chairs — **Wanxiang Che, Joyce Nabende, Mohammad Taher Pilehvar** and **Ekaterina Shutova** — who have overseen the reviewing and selection process, and shaped a compelling and diverse scientific program. Overall, we received more than 8300 submissions and accepted 1700 at the main conference and 1392 as findings. This was made possible through close coordination with the ACL Rolling Review (ARR) team, led by the ARR Editors-in-Chief — **Jun Suzuki, Jing Jiang, and Xiaodan Zhu** — and I sincerely thank them. Although some improvements are still possible, the ARR workflow was already well integrated with the conference review and decision-making process. We are also deeply grateful to the (many!) Senior Area Chairs, Area Chairs, reviewers, and the Best Paper Committee (led by **Rada Mihalcea** and **Roi Reichart**), whose dedication ensured the high quality of our program. I am also grateful to our Technical Open Review Chairs, **Niket Tandon** and **Lizhen Qu**, for their behind-the-scenes work. For those joining virtually, our Virtual Infrastructure Chairs — **Manling Li, Yang Liu, and Avi Sil** — have worked to make the hybrid conference experience inclusive and engaging. I would also like to recognize the efforts of our Ethics Chairs — **Karën Fort** and **Bjorn Ross** — who have led the work aimed at ensuring that submissions meet the ethical standards of our field.

ACL 2025 hosts a rich set of 28 workshops and 8 tutorials, thanks to the dedicated efforts of our Workshop Chairs — **Terra Blevins** and **Christophe Gravier** — and our Tutorial Chairs — **Yuki Arase, David Jurgens, and Fei Xia**. Our Demonstration Track Chairs — **Pushkar Mishra, Smaranda Muresan, and Tao Yu** — 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** and **Georg Rehm** who created a “conference in the conference”. Given the growing interaction between academia and industry, I hope this tradition, which stems from NAACL and which I reinstated at ACL, will be continued in the following editions.

As is customary, TACL and CL papers are presented at the conference. Thanks go to the TACL Editors-in-Chief (**Asli Celikyilmaz, Roi Reichart, and Dilek Hakkani-Tur**) and the CL Editor-in-Chief, **Wei Lu**, for their coordination efforts. I am also proud to announce that this year we are introducing the new CL Doctoral Dissertation Award at the conference. Other notable innovations include virtual presentation sessions held in parallel with the in-person poster sessions, a reception event dedicated to Findings presentations, and the introduction of combined oral and panel sessions designed to foster more thorough and interactive research discussions.

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 — in particular, **Megan Haddad**, our ACL Administrator and Registration Manager — especially given the largest number of participants ever at ACL, estimated at the time of writing to be around 5000! Thanks also go to **Damira Mršić** and the Underline team for managing all content for the hybrid experience. The Local Organization Chairs, **Benjamin Roth** and **Dagmar Gromann**, worked hard on the bidding process, providing huge support with all local arrangements, the social event, and the visa process, eventually ensuring an outstanding experience for our in-person attendees. Our Visa Chairs — **Rexhina Biloshmi** and **Eleni Ilkou** — have helped numerous participants with their visa needs and any logistical difficulties in attending the conference. I also thank our Student Volunteer Chairs — **Pedro Henrique Luz de Araujo** and **Eleonora Mancini** — for their behind-the-scenes work, which will prove indispensable during the conference.

Additional thanks go to **Zhu Liu, Mingyang Wang, and Jin Zhao**, our Student Research Workshop Chairs, who, with the support of faculty advisors **Lea Frermann, Daniel Hershcovich, and Tristan Miller**, 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 — **Pierpaolo Basile, Libo Qin, and Zhenghao Liu** — have ensured timely and high-quality conference proceedings. I also thank the Handbook Chairs, **Els Lefever and Qionгкаi Xu**, 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 — **Raffaella Bernardi and Thomas Scialom** — and to **Chris Callison-Burch**, 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 — **Senja Pollak, Maria Ryskina, Shane Storks, and Hwaran Lee** — 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 — **Anette Frank, Shruti Rijhwani, and Horacio Saggion** — 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 — **Xudong Han and Alessandro Raganato** — 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, **Sara Tonelli and Yiquan Wu**. We are very grateful!

I am particularly proud of the creation of a new role, the Documentation Chair, taken by **Chenghua Lin**, who did strenuous work to assemble a new handbook with up-to-date information coming from all chairs — whom I also thank for their valuable contributions — and maintain a clear and consistent record of all our processes. Given the speed at which all conference roles evolve in their duties and commitments, I hope this role will become standard in upcoming conferences to ensure a smooth transition.

Last but not least, I wish to thank the ACL Executive Committee for their guidance, and my fellow past General and Program Chairs, and especially **Claire Gardent and Colin Cherry**, for generously sharing their experience and insights.

ACL 2025 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 Vienna. Welcome to ACL 2025!

Roberto Navigli

Sapienza University of Rome & Babelscape

General Chair, ACL 2025

Message from the Program Chairs

Welcome to the 63rd Annual Meeting of the Association for Computational Linguistics!

ACL 2025 will be held in a hybrid format, offering attendees the option to join us in person in vibrant Vienna, Austria, or to participate remotely from anywhere in the world. Organizing ACL 2025 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, Roberto Navigli;
- The ARR Editors-in-Chief of the February 2025 cycle (Jun Suzuki, Jing Jiang, and Xiaodan Zhu) and the entire team (Mausam, Viviane Moreira, Vincent Ng, Lilja Øvrelid, Anna Rogers, Michael White, Margot Mieskes, Sarvnaz Karimi);
- The technical OpenReview chairs, Niket Tandon, Lizhen Qu, and the OpenReview support team, in particular Rachel, for multiple rounds of technical help in setting up ACL 2025 on the Open Review platform;
- The 169 Senior Area Chairs;
- The 1,937 Area Chairs and the 11,720 reviewers;
- The best paper committee chairs, Rada Mihalcea and Roi Reichart, and the best paper committee members;
- The ethics chairs, Karën Fort and Bjorn Ross;
- The workshop chairs, Terra Blevins and Christophe Gravier;
- The tutorial chairs, Yuki Arase, David Jurgens and Fei Xia;
- The industry track chairs, Yunyao Li and Georg Rehm;
- The demonstration chairs, Pushkar Mishra, Smaranda Muresan, and Tao Yu;
- The internal communications chairs, Sara Tonelli and Yiquan Wu;
- The website and conference app chairs, Xudong Han and Alessandro Raganato;
- The publication chairs, Pierpaolo Basile, Libo Qin, and Zhenghao Liu;
- The handbook chairs, Els Lefever and Qionghai Xu;
- The local organization chairs, Benjamin Roth and Dagmar Gromann, and their team;
- The visa chairs, Rexhina Blloshmi and Eleni Ilkou;
- The publicity and social media chairs, Anette Frank, Shruti Rijhwani and Horacio Saggion;
- The documentation chair, Chenghua Lin;
- The student research workshop chairs, Zhu Liu, Mingyang Wang and Jin Zhao;
- The student research workshop chairs faculty advisors, Lea Frermann, Daniel Hershcovich and Tristan Miller;
- The student volunteer chairs, Pedro Henrique Luz de Araujo and Eleonora Mancini;

- The diversity and inclusion chairs, Senja Pollak, Maria Ryskina, Shane Storcks and Hwaran Lee;
- The sponsorship chairs, Raffaella Bernardi and Thomas Scialom;
- The virtual infrastructure chairs, Manling Li, Yang Liu and Avi Sil;
- The ACL Anthology Director Matt Post and his team;
- The TACL editors-in-chief (Asli Celikyilmaz, Roi Reichart, Dilek Hakkani-Tur) and CL Editor in-Chief Wei Lu for coordinating TACL and CL presentations with us;
- The ACL 2024 Program Chairs, Lun-Wei Ku, André F. T. Martins, Vivek Srikumar, for information and support;
- Damira Mrsic and Underline Team;
- Jennifer Rachford and entire conference support staff;
- All the authors of papers who submitted their papers for review in the ARR 2025 February cycle and those who committed to the ACL 2025 conference.

Review Process

All submissions to ACL 2025 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 2025 commitment site. At this stage, Senior Area Chairs provided recommendations, and final acceptance decisions were made by the Program Chairs. This process is consistent with previous conferences, including ACL 2024, EACL 2024, and NAACL 2025.

We worked closely with the ARR team, particularly the ARR February 2025 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 11,720 reviewers and 1,942 Area Chairs in the February 2025 ARR cycle to which most ACL 2025 papers were submitted. ACL also recruited 169 Senior Area Chairs to oversee the review and meta-review process. Overall, the ARR process ran smoothly, ensuring that all submitted papers received at least three reviews and a meta-review. For the ACL commitment phase, Senior Area Chairs made recommendations for 5,356 committed papers based on the reviews, meta-reviews, and the papers themselves, with final acceptance decisions made by the Program Chairs.

Acceptance Rate

In total, there are 1,699 papers accepted to the Main Conference and 1,392 papers accepted to Findings. The acceptance rate calculation follows precedent set by previous conferences that go through ACL Rolling Review (ARR), e.g. ACL 2024. The calculation takes into account the multi-stage process of ARR where a paper may get revised in ARR and then later committed to the conference. In total, we had 8,360 unique submissions across the December 2024 and February 2025 ARR cycles of which 5,501 papers were committed to ACL. The acceptance rate is 20.3% for the Main Conference papers and a further 16.7% for Findings papers.

Special Theme: Generalization of NLP Models

ACL 2025’s special theme is generalization of NLP models. Generalization is crucial for ensuring that models behave robustly, reliably, and fairly when making predictions on data different from their training data. Achieving good generalization is critically important for models used in real-world applications, as they should emulate human-like behavior. Humans are known for their ability to generalize well, and models should aspire to this standard. The theme track invites empirical and theoretical research

and position and survey papers reflecting on the Generalization of NLP Models. The possible topics of discussion include (but are not limited to) the following:

1. How can we enhance the generalization of NLP models across various dimensions—compositional, structural, cross-task, cross-lingual, cross-domain, and robustness?
2. What factors affect the generalization of NLP models?
3. What are the most effective methods for evaluating the generalization capabilities of NLP models?
4. While Large Language Models (LLMs) significantly enhance the generalization of NLP models, what are the key limitations of LLMs in this regard?

We received 128 submissions to the theme track during the review phase. Among these, 41 papers were accepted to the main conference and a further 33 to Findings of ACL 2025. The conference will also feature a panel discussion on the theme of Generalization, with the participation of leading experts in this area.

Best Paper Selection

ACL 2025 implemented the updated ACL award policy that seeks to expand the pool of work recognized as outstanding. In total 117 papers were nominated by the reviewers, area chairs and senior area chairs for best paper consideration. The best paper committee assessed these papers to select the best papers (featuring $(\leq)0.6\%$ of accepted papers), outstanding papers (featuring $(\leq)2.5\%$ of accepted papers), and special awards for social impact and best resource. Based on the review by the Best Paper committee, 39 papers have been selected for awards in the above categories. Separately, the senior area chairs also nominated their favourite papers as SAC Highlights.

In addition, we have awards for test of time award for a paper published in TACL in 2013 or 2014 and the best paper award for a paper published in TACL in 2024. The final selection was made by the best paper committee, and the winners will be announced during the closing ceremony. The ACL 2025 Best Papers will also be given an opportunity to present their work in the closing ceremony.

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 presentations sessions during the main conference. This enables us to align the virtual sessions with time slots when in-person participants are available. 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.

This year, 218 main conference papers were selected for oral presentations by the program chairs, with the goal of creating a well-rounded program featuring a diverse set of topics instead of selecting papers based on their review scores. This year we have introduced a panel section of the conference. Out of the selected oral presenters, 25 will have an opportunity to not only present their work but also participate in a panel discussion. We have lined up five panels on the theme of: Generalisation in NLP, LLM alignment, Human-centred NLP, Interpretability and model analysis, Multilinguality and language diversity.

In addition to the main conference papers, the ACL program also includes 18 papers accepted by Computational Linguistics and 42 papers accepted by Transactions of the ACL (TACL). Among these, 11 journal papers will be presented in-person as oral presentations, thematically distributed across appropriate sessions.

This year the conference will also for the first time feature a dedicated Findings poster session with reception. All Findings presenters have been assigned poster presentations in this session or alongside other posters for main conference papers in the same track. Rounding out the program are dedicated sessions for the demonstrations track and the student research workshop.

Keynotes and Panel This year’s program features an impressive lineup of two keynote presentations:

- **Prof. Luke Zettlemoyer** from Paul G. Allen School of Computer Science & Engineering at the University of Washington, and a Senior Research Director at Meta, will share his insights on “*Rethinking Pretraining: Data and Architecture.*”
- **Prof. Verena Rieser** from Google DeepMind will present on “*Whose Gold? Re-imagining Alignment for Truly Beneficial AI.*”

Alongside these keynotes, we are thrilled to host a panel discussion aiming to answer the question *Can large language models (LLMs) generalize?*. Our esteemed panelists include:

- **Prof. Eduard Hovy**, University of Melbourne (who will also act as the panel chair).
- **Prof. Mirella Lapata**, University of Edinburgh
- **Prof. Yue Zhang**, Westlake University
- **Prof. Dan Roth**, University of Pennsylvania and Oracle

This diverse group of panelists will provide a comprehensive view of the latest trends and challenges in Generalisation of NLP in the Large Language Models era. We hope you enjoy this year’s diverse and engaging program!

Ekaterina Shutova (University of Amsterdam and Stanford University)

Mohammad Taher Pilehvar (Cardiff University and Tehran Institute for Advanced Studies)

Joyce Nakatumba-Nabende (Makerere University)

Wanxiang Che (Harbin Institute of Technology)

ACL 2025 Program Co-Chairs

Organizing Committee

General Chair

Roberto Navigli, Sapienza University Rome & Babelscape

Program Chairs

Wanxiang Che, Harbin Institute of Technology

Joyce Nabende, Makerere University

Ekaterina Shutova, University of Amsterdam

Mohammad Taher Pilehvar, Cardiff University / Tehran Institute for Advanced Studies

Local Organization

Benjamin Roth, University of Vienna

Dagmar Gromann, University of Vienna

Workshop Chairs

Terra Blevins, University of Vienna

Christophe Gravier, Université Jean Monnet

Tutorial Chairs

Yuki Arase, Tokyo Institute of Technology

David Jurgens, University of Michigan

Fei Xia, University of Washington

Demonstration Chairs

Pushkar Mishra, Google DeepMind / Cambridge

Smaranda Muresan, Barnard College, Columbia University

Tao Yu, The University of Hong Kong

Student Research Workshop Chairs

Zhu Liu, Tsinghua University

Mingyang Wang, LMU

Jin Zhao, Brandeis University

Student Research Workshop Chairs: Faculty Advisors

Lea Frermann, University of Melbourne

Daniel Hershcovich, University of Copenhagen

Tristan Miller, University of Manitoba

Publication Chairs

Pierpaolo Basile, University of Bari
Libo Qin, Central South University
Zhenghao Liu, Northeastern University

Handbook Chair

Els Lefever, Ghent University
Qiongkai Xu, Macquarie University

Sponsorship Chairs

Raffaella Bernardi, University of Trento
Thomas Scialom, Meta AI

Diversity and Inclusion Chairs

Senja Pollak, Jožef Stefan Institute
Maria Ryskina, MIT
Shane Storcks, University of Michigan
Hwaran Lee, NAVER

Publicity and Social Media Chairs

Anette Frank, Heidelberg University
Shruti Rijhwani, Google DeepMind
Horacio Saggion, Pompeu Fabra University

Website and Conference App Chairs

Xudong Han, LibrAI / MBZUAI
Alessandro Raganato, University of Milano-Bicocca

Ethics Chairs

Karën Fort, Université de Lorraine / LORIA
Bjorn Ross, University of Edinburgh

Student Volunteer Chairs

Pedro Henrique Luz de Araujo, University of Vienna
Eleonora Mancini, University of Bologna

Technical Open Review Chairs

Niket Tandon, Allen AI
Lizhen Qu, Monash University

Virtual Infrastructure Chair

Manling Li, Northwestern University

Yang Liu, Microsoft
Avil Sil, IBM Research

Internal Communication Chairs

Sara Tonelli, FBK
Yiquan Wu, Zhejiang University

Industry Track Chairs

Yun Yao Li, Adobe
George Rehm, DFKI

Best Paper Committee Chairs

Rada Mihalcea, University of Michigan
Roi Reichart, Technion - Israel Institute of Technology

Visa Chairs

Rexhina Blloshmi, Amazon AGI
Eleni Ilkou, University of Hannover

Documentation Chairs

Chenghua Lin, University of Manchester

ARR Editors in Chief

Jun Suzuki, Tohoku University
Jing Jiang, Australian National University
Xiaodan Zhu, Queen's University

Program Committee

Senior Area Chairs

David Ifeoluwa Adelani, McGill University
Alan Akbik, Humboldt Universität Berlin
Raquel Alhama, University of Amsterdam
Afra Alishahi, Tilburg University
Jacob Andreas, Massachusetts Institute of Technology
Maria Antoniak, Allen Institute for Artificial Intelligence
Marianna Apidianaki, University of Pennsylvania
Jun Araki, Bosch Research
Ehsaneddin Asgari, Qatar Computing Research Institute
Edoardo Barba, Sapienza University of Rome
Jasmijn Bastings, Google DeepMind
Beata Beigman Klebanov, Educational Testing Service
Ahmad Beirami, Google DeepMind
Yonatan Belinkov, Technion
Meriem Beloucif, Uppsala University
Jonathan Berant, Tel Aviv University
Steven Bethard, University of Arizona
Alexandra Birch-Mayne, University of Edinburgh
Jelke Bloem, University of Amsterdam
Danushka Bollegala, University of Liverpool
Houda Bouamor, Carnegie Mellon University
Jordan Boyd-Graber, University of Maryland
Emanuele Bugliarello, Google DeepMind
Jan Buys, University of Cape Town
Aoife Cahill, Dataminr
Jose Camacho-Collados, Cardiff University
Hanjie Chen, Rice University
Mingda Chen, Meta FAIR
Danqi Chen, Princeton University
Christos Christodoulopoulos, Amazon
Arman Cohan, Yale University
Simone Conia, Sapienza University of Rome
Ryan Cotterell, Swiss Federal Institute of Technology
Yiming Cui, iFLYTEK Research
Raj Dabre, National Institute of Information and Communications Technology (NICT)
Tejaswini Deoskar, Utrecht University
Leon Derczynski, University of Copenhagen
Yuxiao Dong, Tsinghua University
Li Dong, Microsoft Research
Yue Dong, University of California, Riverside
A. Seza Dođruöz, Ghent University
Greg Durrett, University of Texas at Austin
Ondřej Dušek, Charles University, Prague
Nouha Dziri, Allen Institute for Artificial Intelligence
Desmond Elliott, University of Copenhagen
Guy Emerson, University of Cambridge

Alexander Fabbri, Scale AI
Anna Feldman, Montclair State University
Fuli Feng, University of Science and Technology of China
Xiaocheng Feng, Harbin Institute of Technology
Anjalie Field, Johns Hopkins University
Michael Flor, Educational Testing Service
Antske Fokkens, VU University Amsterdam
Zhe Gan, Apple
Albert Gatt, Utrecht University
Mor Geva, Tel Aviv University
Marjan Ghazvininejad, Facebook
Goran Glavaš, Julius-Maximilians-Universität Würzburg
Yvette Graham, Trinity College Dublin
Jiafeng Guo, Chinese Academy of Sciences
Xianpei Han, Institute of Software, CAS
Peter Hase, Schmidt Sciences
Devamanyu Hazarika, Facebook
Yulan He, King's College London
Mohammad Javad Hosseini, Google
Neil Houlsby, Google
Dieuwke Hupkes, Facebook
Ignacio Iacobacci, Elm Europe
Mohit Iyyer, University of Massachusetts
Zhijing Jin, Max Planck Institute and University of Toronto
Shafiq Joty, Nanyang Technological University and Salesforce.com
Rabeeh Karimi Mahabadi, NVIDIA
Frank Keller, University of Edinburgh
Daniel Khashabi, Johns Hopkins University
Najoung Kim, Boston University
Svetlana Kiritchenko, National Research Council Canada
Ekaterina Kochmar, Mohamed bin Zayed University of Artificial Intelligence
Julia Kreutzer, Cohere for AI
Kalpesh Krishna, Google DeepMind
Ranjay Krishna, University of Washington
Matthieu Labeau, Télécom ParisTech
Anne Lauscher, Universität Hamburg
Carolin Lawrence, NEC Laboratories Europe
Wenqiang Lei, Sichuan University
Martha Lewis, University of Amsterdam
Junyi Jessy Li, University of Texas at Austin
Zhenghua Li, Soochow University
Dongxu Li, Salesforce.com
Piji Li, Nanjing University of Aeronautics and Astronautics, China
Yiqun Liu, Tsinghua University
Fangyu Liu, Google DeepMind
Zhiyuan Liu, Tsinghua University
Kang Liu, Institute of Automation, Chinese Academy of Sciences
Edison Marrese-Taylor, National Institute of Advanced Industrial Science and Technology (AIST)
Joshua Maynez, Google
Dipendra Misra, Mosaic Research, Databricks
Saif Mohammad, National Research Council of Canada

Nafise Sadat Moosavi, University of Sheffield
Aaron Mueller, Northeastern University
Philippe Muller, IRIT, University of Toulouse
Preslav Nakov, Mohamed bin Zayed University of Artificial Intelligence
Vivi Nastase, University of Geneva
Wolfgang Nejdl, Leibniz University Hannover
Xuefei Ning, Tsinghua University
Debora Nozza, Bocconi University
Giorgio Maria Di Nunzio, University of Padua
Nedjma Ousidhoum, Cardiff University
Alexandros Papangelis, Apple
Gabriella Pasi, University of Milano-Bicocca
Rebecca Passonneau, Pennsylvania State University
Viviana Patti, University of Turin
Hao Peng, University of Illinois Urbana-Champaign
Carla Perez-Almendros, Cardiff University
Tiago Pimentel, ETH Zurich
Edoardo Ponti, University of Edinburgh
Simone Paolo Ponzetto, Universität Mannheim
Daniel Preotiuc-Pietro, Bloomberg
Emily Prud'hommeaux, Boston College
Peng Qi, Orby AI
Marek Rei, Imperial College London
Verena Rieser, Google DeepMind
Roland Roller, German Research Center for AI
Alla Rozovskaya, Queens College, City University of New York
Michael Sejr Schlichtkrull, Queen Mary, University of London
Steven Schockaert, Cardiff University
Rico Sennrich, University of Zurich
Ehsan Shareghi, Monash University
Qinlan Shen, Oracle
Kevin Small, Amazon
Mihai Surdeanu, University of Arizona
Anders Søgaard, Copenhagen University
Chenhao Tan, University of Chicago
Hao Tan, Adobe Systems
Harish Tayyar Madabushi, University of Bath
Alberto Testoni, Amsterdam UMC
Trang Tran, USC Institute for Creative Technologies
Marcos Vinicius Treviso, Instituto de Telecomunicações, Lisbon
Yulia Tsvetkov, Carnegie Mellon University
Daniel Varab, German Research Center for AI
Suzan Verberne, Universiteit Leiden
Karin Verspoor, Royal Melbourne Institute of Technology
Ivan Vulić, University of Cambridge
Henning Wachsmuth, Leibniz Universität Hannover
Byron Wallace, Northeastern University
Qifan Wang, Meta AI
Tong Xiao, Northeastern University
Deyi Xiong, Tianjin University
Ruifeng Xu, Harbin Institute of Technology

Yadollah Yaghoobzadeh, University of Tehran
Diyi Yang, Stanford University
Min Yang, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences
Helen Yannakoudakis, King's College London
Huaxiu Yao, University of North Carolina at Chapel Hill
Andrew Yates, Johns Hopkins University
Kai Yu, Shanghai Jiao Tong University
Luke Zettlemoyer, University of Washington and Meta
Weinan Zhang, Harbin Institute of Technology
Meishan Zhang, Harbin Institute of Technology, Shenzhen
Yue Zhang, Westlake University
Qi Zhang, Fudan University
Xin Zhao, Renmin University of China
Yi Zhou, Cardiff University
Arkaitz Zubiaga, Queen Mary University of London

Area Chairs

Omri Abend, Gavin Abercrombie, Heike Adel, Rodrigo Agerri, Ameeta Agrawal, Sweta Agrawal, Ibrahim Said Ahmad, Xi Ai, Alham Fikri Aji, Yamen Ajjour, Md Shad Akhtar, Firoj Alam, Mohammad Aliannejadi, Emily Allaway, Milad Alshomary, Hadi Amiri, Antonios Anastasopoulos, Xiang Ao, Catherine Arnett, Shane Arora, Giuseppe Attanasio, Katherine Atwell

Fan Bai, Jiaxin Bai, Xuefeng Bai, JinYeong Bak, Niranjana Balasubramanian, Timothy Baldwin, Miguel Ballesteros, David Bamman, Sivaji Bandyopadhyay, Xiaoyi Bao, Yu Bao, Roy Bar-Haim, Chitta Baral, Fazl Barez, Jeremy Barnes, Valentin Barriere, Joe Barrow, Alberto Barrón-Cedeño, Max Bartolo, Amrit Singh Bedi, Shabnam Behzad, Adrian Benton, Yevgeni Berzak, Federico Bianchi, Yonatan Bitton, Johannes Bjerva, Terra Blevins, Su Lin Blodgett, Nadav Borenstein, Florian Boudin, Zied Bouraoui, Kevin Bowden, Stephanie Brandl, Pavel Braslavski, Chloé Braud, Davide Buscaldi

Sky CH-Wang, Elena Cabrio, Deng Cai, Xiangrui Cai, Zhaowei Cai, Andrew Caines, Ruken Cakici, Nitay Calderon, Jie Cao, Meng Cao, Meng Cao, Pengfei Cao, Qingqing Cao, Yang Trista Cao, Yixin Cao, Yong Cao, Ziqiang Cao, Doina Caragea, Dallas Card, Dumitru-Clementin Cercel, Soumen Chakrabarti, Saikat Chakraborty, Bharathi Raja Chakravarthi, Chunkit Chan, Angel X Chang, Ernie Chang, Haw-Shiuan Chang, Yi Chang, Snigdha Chaturvedi, Aditi Chaudhary, Chung-Chi Chen, Francine Chen, Guanhua Chen, Hui Chen, Jiangjie Chen, Jifan Chen, Junfan Chen, Kehai Chen, Lu Chen, Muhao Chen, Pin-Yu Chen, Qian Chen, Qianglong Chen, Wei Chen, Wenliang Chen, Xilun Chen, Xinchu Chen, Xinyun Chen, Xiuying Chen, Yangyi Chen, Yidong Chen, Yiming Chen, Yulong Chen, Yun-Nung Chen, Yunmo Chen, Zhiyu Chen, Zhiyu Chen, Zhuang Chen, Fei Cheng, Gong Cheng, Lu Cheng, Pengyu Cheng, Wei Cheng, Xueqi Cheng, Yu Cheng, Jianfeng Chi, Nadezhda Chirkova, Hyunsoo Cho, Sangwoo Cho, Eleanor Chodroff, Edward Choi, Eunsol Choi, Jinho D. Choi, YunSeok Choi, Jaegul Choo, Prafulla Kumar Choubey, Alexandra Chronopoulou, Chenhui Chu, Yi-Ling Chung, Simon Clematide, Shay B Cohen, John M Conroy, Caio Corro, Benoit Crabbé, Heriberto Cuayahuitl, Ruixiang Cui, Aron Culotta

Giovanni Da San Martino, Damai Dai, Bhavana Dalvi Mishra, Rossana Damiano, Sandipan Dandapat, Pradeep Dasigi, Orphee De Clercq, Budhaditya Deb, Pieter Delobelle, Vera Demberg, Daryna Dementieva, Chenlong Deng, Shumin Deng, Yang Deng, Yuntian Deng, Daniel Deutsch, Chris Develder, Kuntal Dey, Jwala Dhamala, Shizhe Diao, Haibo Ding, Kaize Ding, Liang Ding, Liviu P Dinu, Lucia Donatelli, Hanze Dong, Ruihai Dong, Zhicheng Dou, Markus Dreyer, Li Du,

Mengnan Du, Xinya Du, Brian DuSell, Xiangyu Duan, Marie-Catherine de Marneffe

Lukas Edman, Asif Ekbal, Aparna Elangovan, Maha Elbayad, Ahmed Elgohary, Ali Emami, Cristina España-Bonet, Ralph Ewerth

Agnieszka Falenska, Kai Fan, Yaxin Fan, Meng Fang, Marcello Federico, Hao Fei, Naomi Feldman, Jiazhan Feng, Shangbin Feng, Shaoxiong Feng, Shi Feng, Xiachong Feng, Yi Feng, Javier Ferrando, Anjalie Field, Simone Filice, Mark Fishel, Jeffrey Flanigan, Lucie Flek, Radu Florian, Mary Ellen Foster, Andre Freitas, Lea Frermann, Annemarie Friedrich, Yi R. Fung

Saadia Gabriel, Marco Gaido, Matthias Gallé, Niloy Ganguly, Chongming Gao, Yifan Gao, Ze-Feng Gao, Jon Gauthier, Jiahui Geng, Luke Gessler, Abbas Ghaddar, Debanjan Ghosh, Saptarshi Ghosh, Salvatore Giorgi, Mario Giulianelli, Nathan Godey, Seraphina Goldfarb-Tarrant, Omer Goldman, Dan Goldwasser, Sharon Goldwater, Sian Gooding, Sruthi Gorantla, Kyle Gorman, Matthew R. Gormley, Kartik Goyal, Loïc Grobol, Jiuxiang Gu, Qing Gu, Lin Gui, Junliang Guo, Yue Guo, Raghav Gupta, Vivek Gupta, Izzeddin Gur, Bernal Jiménez Gutiérrez, Carlos Gómez-Rodríguez

Nizar Habash, Ivan Habernal, Barry Haddow, Michael Hahn, Dilek Hakkani-tür, Rujun Han, Jie Hao, Mohammed Hasanuzzaman, Shirley Anugrah Hayati, Ben He, Bin He, Bobby He, He He, Jie He, Junxian He, Keqing He, Shizhu He, Tianxing He, Wei He, Daniel Hershcovich, Toshio Hirasawa, Julia Hockenmaier, Valentin Hofmann, Christopher M Homan, Lingzi Hong, Enamul Hoque, Lei Hou, Yufang Hou, Alexander Miserlis Hoyle, Dou Hu, Joy Hu, Joy Hu, Junjie Hu, Linmei Hu, Mengting Hu, Wei Hu, Wenpeng Hu, Xuming Hu, Chen Huang, Jiabin Huang, Kuan-Hao Huang, Kung-Hsiang Huang, Ruihong Huang, Tenghao Huang, Xinting Huang, Xuanjing Huang, Yi Huang, Yichong Huang, Yinya Huang, Zhenya Huang, Pere-Lluís Huguet Cabot, Chia-Chien Hung, Wonseok Hwang, seung-won hwang

Timour Igamberdiev, Filip Ilievski, Sathish Reddy Indurthi, Naoya Inoue, Radu Tudor Ionescu, Daphne Ippolito, Etsuko Ishii, Hayate Iso, Masaru Isonuma

Matthew Jagielski, Kokil Jaidka, Masoud Jalili Sabet, Shoaib Jameel, Hyeju Jang, Adam Jatowt, Ganesh Jawahar, Akshita Jha, Harsh Jhamtani, Shaoxiong Ji, Tao Ji, Yangfeng Ji, Ziwei Ji, Robin Jia, Zixia Jia, Chao Jiang, Huiqiang Jiang, Zhiying Jiang, Zhuoxuan Jiang, Zhanming Jie, Lifeng Jin, Qiao Jin, Qin Jin, Yohan Jo, Richard Johansson, Wei Ju, Jaap Jumelet, David Jurgens, Preethi Jyothi

Jad Kabbara, Hour Kaing, Anup K. Kalia, Ehsan Kamaloo, Hidetaka Kamigaito, Masahiro Kaneko, Dongyeop Kang, Diptesh Kanojia, Alina Karakanta, Börje F. Karlsson, Pei Ke, Yova Kementchedjheva, Ilker Kesen, Salam Khalifa, Khalid Al Khatib, Omar Khattab, Hyunjae Kim, Hyunwoo Kim, Jin-Dong Kim, Joo-Kyung Kim, Taeuk Kim, Yeachan Kim, Yunsu Kim, Christo Kirov, Hirokazu Kiyomaru, Dietrich Klakow, Roman Klinger, Christine De Kock, Christine De Kock, Philipp Koehn, Konstantinos Kogkalidis, Keshav Kolluru, Mamoru Komachi, Rik Koncel-Kedziorski, Lingkai Kong, Ioannis Konstas, Michalis Korakakis, Parisa Kordjamshidi, Raphael Koster, Fajri Koto, Emiel Kraemer, Amrith Krishna, Kundan Krishna, Shivani Kumar, Anoop Kunchukuttan, Souvik Kundu, Tatsuki Kuribayashi, Shuhei Kurita

Wai Lam, Wuwei Lan, Yunshi Lan, Md Tahmid Rahman Laskar, Hung Le, Thai Le, Eric Le Ferrand, Andrew Lee, Chia-Hsuan Lee, Dong-Ho Lee, Hung-yi Lee, Hwanhee Lee, Hwaran Lee, Jaeho Lee, Jay-Yoon Lee, Roy Ka-Wei Lee, SangKeun Lee, Seolhwa Lee, Oliver Lemon, Heather Lent, Michael A. Lepori, Kristina Lerman, Piyawat Lertvittayakumjorn, Chenliang Li, Dongfang

Li, Fei Li, Jing Li, Jingjing Li, Jiyi Li, Junhui Li, Junyi Li, Lei Li, Lei Li, Linyang Li, Manling Li, Mingxiao Li, Peifeng Li, Peng Li, Qian Li, Ru Li, Ruizhe Li, Shuai Li, Tao Li, Wei Li, Wenjie Li, Xiang Lorraine Li, Xiang Li, Xiaoli Li, Xin Li, Xintong Li, Yaliang Li, Yanzeng Li, Yaoyiran Li, Yinghui Li, Yingjie Li, Yu Li, Yuan-Fang Li, Zhuang Li, Zixuan Li, Zixuan Li, Bin Liang, Paul Pu Liang, Jindřich Libovický, Constantine Lignos, Chenghua Lin, Hongfei Lin, Hsien-chin Lin, Nankai Lin, Shou-De Lin, Yankai Lin, Yong Lin, Zhaojiang Lin, Pierre Lison, Bang Liu, Bingquan Liu, Dongqi Liu, Fenglin Liu, Guangliang Liu, Hao Liu, Hao Liu, Hui Liu, Jiahao Liu, Ming Liu, Qi Liu, Qianying Liu, Qun Liu, Tianyu Liu, Tingwen Liu, Wei Liu, Xiao Liu, Xiao Liu, Xuebo Liu, Yang Liu, Ye Liu, Yixin Liu, Zhenghao Liu, Zhiyuan Liu, Zuozhu Liu, Karen Livescu, Shayne Longpre, Adam Lopez, Hongyuan Lu, Wei Lu, Weiming Lu, Yaojie Lu, Yaojie Lu, Li Lucy, Thomas Lukasiewicz, Michal Lukasik, Junyu Luo, Xiao Luo, Anh Tuan Luu, Ang Lv

Fenglong Ma, Jing Ma, Mingyu Derek Ma, Ruotian Ma, Xiao Ma, Yubo Ma, Manuel Mager, Kyle Mahowald, Valentin Malykh, Potsawee Manakul, Xin Mao, Zhendong Mao, David Mareček, Katerina Margatina, Sérgio Matos, Jonathan May, Sahisnu Mazumder, Ninareh Mehrabi, Clara Meister, Dheeraj Mekala, Zaiqiao Meng, Sabrina J Mielke, Simon Mille, Kanishka Misra, Fengran Mo, Daichi Mochihashi, Ashutosh Modi, Hosein Mohebbi, Tasnim Mohiuddin, Diego Molla, Masoud Monajatipoor, Syrielle Montariol, Manuel Montes, Roser Morante, Makoto Morishita, David R. Mortensen, Marius Mosbach, Lili Mou, Youssef Mroueh, Rudra Murthy

Masaaki Nagata, Linyong Nan, Jason Naradowsky, Pranav Narayanan Venkit, Usman Naseem, Mir Tafseer Nayeem, Isar Nejadgholi, Hwee Tou Ng, Dat Quoc Nguyen, Dong Nguyen, Hoang H Nguyen, Minh-Tien Nguyen, Ansong Ni, Shiwen Ni, Garrett Nicolai, Jan Niehues, Joel Niklaus, Mitja Nikolaus, Qiang Ning, Takashi Ninomiya, Tong Niu, Farhad Nooralahzadeh, Pierre Nugues

Maciej Ogrodniczuk, Alice Oh, Naoaki Okazaki, Manabu Okumura, Alexandra Olteanu, Arturo Oncevay, Matan Orbach, Yohei Oseki, Simon Ostermann, Jessica Ouyang, Mete Ozay

Maria Leonor Pacheco, Aishwarya Padmakumar, Sebastian Padó, Serguei V. S. Pakhomov, Liangming Pan, Ashwinee Panda, Liang Pang, Isabel Papadimitriou, Sara Papi, Letitia Parcalabescu, Cecile Paris, Chanjun Park, Joonsuk Park, Kunwoo Park, Roma Patel, Siddharth Patwardhan, Debjit Paul, Ellie Pavlick, Hao Peng, Hao Peng, Minlong Peng, Siyao Peng, Veronica Perez-Rosas, Sandro Pezzelle, Chau Thi Minh Pham, Chau Thi Minh Minh Pham, Yuval Pinter, Irina Piontkovskaya, Benjamin Piwowarski, Bryan A. Plummer, Soujanya Poria, Shrimai Prabhumoye, Animesh Prasad, Radityo Eko Prasajo, Emily Prud'hommeaux, Danish Pruthi

Tao Qi, Chen Qian, Cheng Qian, Tiejun Qian, Libo Qin, Yanxia Qin, Chen Qiu, Han Qiu, Haoyi Qiu, Jieli Qiu, Xipeng Qiu, Lizhen Qu, Tingyu Qu, Xiaojun Quan

Dinesh Raghu, Vyas Raina, Hossein Rajaby Faghihi, Dhananjay Ram, Deepak Ramachandran, Carlos Ramisch, Alan Ramponi, Leonardo Ranaldi, Tharindu Ranasinghe, Surangika Ranathunga, Yanghui Rao, Hannah Rashkin, Charvi Rastogi, Shauli Ravfogel, Jishnu Ray Chowdhury, Anton Razzhigaev, Traian Rebedea, Ehud Reiter, Navid Rekasaz, Haopeng Ren, Ruiyang Ren, Yi Ren, Zhaochun Ren, Leonardo F. R. Ribeiro, Anthony Rios, Julia Romberg, Salvatore Romeo, Michael Roth, Joseph Le Roux, Frank Rudzicz, Ramon Ruiz-Dolz, Paul Röttger

Mrinmaya Sachan, Mehrnoosh Sadrzadeh, Kenji Sagae, Horacio Saggion, Gaurav Sahu, Hassan Sajjad, Elizabeth Salesky, Shailaja Keyur Sampat, Maarten Sap, Abulhair Saparov, Gabriele Sarti, Asad B. Sayeed, Yves Scherrer, Viktor Schlegel, Hinrich Schuetze, Hendrik Schuff, Raphael Schumann, João Sedoc, Indira Sen, Shubhashis Sengupta, Maureen De Seyssel, Lei Sha, Guokan Shang, Jingbo Shang, Ori Shapira, Pratyusha Sharma, Vasu Sharma, Hassan Shavarani, Ravi

Shekhar, Jiaming Shen, Xiaoyu Shen, Yongliang Shen, Peng Shi, Tianze Shi, Zhengliang Shi, Jiasheng Si, Qingyi Si, Carina Silberer, Edwin Simpson, Maneesh Kumar Singh, Koustuv Sinha, Sunayana Sitaram, Konstantinos Skianis, Noah A. Smith, Wei Song, Xingyi Song, Yangqiu Song, Yixiao Song, Sandeep Soni, Aitor Soroa, Gerasimos Spanakis, Balaji Vasani Srinivasan, Shashank Srivastava, Jacopo Staiano, Gabriel Stanovsky, Elias Stengel-Eskin, Kevin Stowe, Marco Antonio Stranisci, Jinsong Su, Qi Su, Qinliang Su, Yu Su, Yusheng Su, Saku Sugawara, Alessandro Suglia, Dianbo Sui, Elior Sulem, Aixin Sun, Fei Sun, Hao Sun, Jiashuo Sun, Kai Sun, Lichao Sun, Mingming Sun, Ruoxi Sun, Yuan Sun, Ziteng Sun, Shahbaz Syed, Michal Štefánik

Santosh T.Y.S.S, Hiroya Takamura, Zeerak Talat, Estefania Talavera, Minghuan Tan, Zhaoxuan Tan, Zhixing Tan, Gongbo Tang, Hao Tang, Hongyin Tang, Siliang Tang, Tianyi Tang, Xiangru Tang, Zhiyang Teng, Chris Thomas, Jesse Thomason, Zhiliang Tian, Maria Tikhonova, Ivan Titov, Amalia Todirascu, Antonio Toral, Kristina Toutanova, Amine Trabelsi, Quan Hung Tran, Chen-Tse Tsai, Richard Tzong-Han Tsai, Reut Tsarfaty, Yu-Hsiang Tseng, Yoshimasa Tsuruoka, Jingxuan Tu, Kewei Tu, Yunbin Tu, Zhaopeng Tu, Yi-Lin Tuan, Gokhan Tur, Ferhan Ture, Martin Tutek

Ana Sabina Uban, Adaku Uchendu

Esaú VILLATORO-TELLO, Neeraj Varshney, Alakananda Vempala, Pat Verga, Gaurav Verma, Serena Villata, Juraj Vladika, Rob Voigt, Thang Vu, Thuy Vu, Tu Vu, Yogarshi Vyas

V. W., Xiaojun Wan, Xingchen Wan, Benyou Wang, Bin Wang, Boshi Wang, Boxin Wang, Chengyu Wang, Fei Wang, Guoyin Wang, Hao Wang, Haoyu Wang, Jiaan Wang, Jin Wang, Jingjing Wang, Keze Wang, Liang Wang, Lingzhi Wang, Liwei Wang, Lucy Lu Wang, Peifeng Wang, Renzhi Wang, Rui Wang, Shijin Wang, Shuohang Wang, Siyuan Wang, Weiqi Wang, Wenjie Wang, Wenjie Wang, Wenya Wang, Xiaolei Wang, Xiaozhi Wang, Xin Eric Wang, Xin Wang, Xindi Wang, Xinyu Wang, Xuan Wang, Yan Wang, Yile Wang, Yu Wang, Zhaowei Wang, Zhongqing Wang, Zifeng Wang, Zihao Wang, Alex Warstadt, Shinji Watanabe, Taro Watanabe, Kellie Webster, Wei Wei, Shira Wein, Nathaniel Weir, Leonie Weissweiler, Haoyang Wen, Derry Tanti Wijaya, Thilini Wijesiriwardene, Genta Indra Winata, Dustin Wright, Bingzhe Wu, Chien-Sheng Wu, Chuhan Wu, Fangzhao Wu, Hongqiu Wu, Lijun Wu, Qiyu Wu, Xian Wu, Xiao-Ming Wu, Xiaobao Wu, Yiquan Wu, Yuanbin Wu, Yunfang Wu, Zhen Wu, Zhuofeng Wu

Heming Xia, Qianqian Xie, Tengyang Xie, Chenyan Xiong, Bo Xu, Chunpu Xu, Han Xu, Hongfei Xu, Jiacheng Xu, Jinan Xu, Qionгкаi Xu, Silei Xu, Tong Xu, Xinnuo Xu, Yumo Xu, Zhaozhuo Xu, Nianwen Xue

Shuntaro Yada, Shweta Yadav, Ikuya Yamada, Ikuya Yamada, An Yan, Hanqi Yan, Jianhao Yan, Jun Yan, Lingyong Yan, Rui Yan, Hitomi Yanaka, Chao-Han Huck Yang, Deqing Yang, Hao Yang, Jian Yang, Jie Yang, Jingkang Yang, Mingming Yang, Qian Yang, Shaohua Yang, Songlin Yang, Xianjun Yang, Yezhou Yang, Zhenglu Yang, Zhiwei Yang, Zonglin Yang, Bingsheng Yao, Ziyu Yao, Zonghai Yao, Gregory Yauney, Fanghua Ye, Hai Ye, Jinyoung Yeo, Seid Muhie Yimam, Da Yin, Jian Yin, Yongjing Yin, Kang Min Yoo, Naoki Yoshinaga, Jiakuan You, Jianfei Yu, Jianxing Yu, Juntao Yu, Lei Yu, Liang-Chih Yu, Tong Yu, Wenhao Yu, Youngjae Yu, Yue Yu, Zheng Yuan, Zheng Yuan, Xiang Yue, Zhenrui Yue

Fengbin ZHU, Fabio Massimo Zanzotto, Xingshan Zeng, Torsten Zesch, Chen Zhang, Chen Zhang, Hao Zhang, Haopeng Zhang, Hongyang Zhang, Jiajun Zhang, Jiabin Zhang, Jing Zhang, Kai Zhang, Lei Zhang, Lu Zhang, Ningyu Zhang, Rongzhi Zhang, Rui Zhang, Ruiyi Zhang, Shaolei Zhang, Shichang Zhang, Shikun Zhang, Shiyue Zhang, Shuai Zhang, Shuai Zhang, Taolin

Zhang, Tong Zhang, Wei Zhang, Weinan Zhang, Wenqi Zhang, Wenxuan Zhang, Xinsong Zhang, Yan Zhang, Yang Zhang, Yang Zhang, Yi Zhang, Yi Zhang, Yiming Zhang, Ying Zhang, Yong Zhang, Yu Zhang, Yu Zhang, Yuhui Zhang, Zeyu Zhang, Zhe Zhang, Zheng Zhang, Zhengyan Zhang, Zhenyu Zhang, Zhuosheng Zhang, Chao Zhao, Chen Zhao, Chenye Zhao, Dongyan Zhao, Jieyu Zhao, Jinming Zhao, Jun Zhao, Runcong Zhao, Tiejun Zhao, Xinran Zhao, Xuandong Zhao, Yang Zhao, Yanpeng Zhao, Zhixue Zhao, Chujie Zheng, Zaixiang Zheng, Zilong Zheng, Ben Zhou, Guangyou Zhou, Hanzhang Zhou, Jiawei Zhou, Jie Zhou, Kaitlyn Zhou, Qingyu Zhou, Tianyi Zhou, Wenxuan Zhou, Xiang Zhou, Yichao Zhou, Yucheng Zhou, Hao Zhu, Wanrong Zhu, Wenhao Zhu, Yilun Zhu, Zining Zhu, Terry Yue Zhuo, Shi Zong, Yicheng Zou, hai zhao

Reviewers

Aakanksha, Asma Ben Abacha, Micheal Abaho, Mahdi Abavisani, Ahmed Abbasi, Mohamed Abdalla, Abdelrahman Abdallah, Mohamed Abdellatif, Muhammad Abdul-mageed, Amir Abdullah, Badr M Abdullah, Idris Abdulmumin, Giuseppe Abrami, Mitchell Abrams, Abdalghani Abujabal, Artem Abzaliev, Lasha Abzianidze, Arkadeep Acharya, Sharada Acharya, Emre Can Acikgoz, Gowthamkrishna Addluri, Tosin Adewumi, Suman Adhya, Somak Aditya, Vaibhav Adlakhia, Muhammad Abdullah Adnan, Helena Gomez Adorno, Estelle Aflalo, Haithem Afli, Aishwarya Agarwal, Amit Agarwal, Ankush Agarwal, Saurabh Agarwal, Siddhant Agarwal, Ameya Agaskar, Arshiya Aggarwal, Divyanshu Aggarwal, Karan Aggarwal, Milan Aggarwal, Mohamed Aghzal, Eneko Agirre, Ayush Agrawal, Samarth Agrawal, Carlos Alejandro Aguirre, N.nasurudeen Ahmed, Sina Ahmadi, Arash Ahmadian, Murtadha Ahmed, Shafiuddin Rehan Ahmed, Zeeshan Ahmed, Jaewoo Ahn, Sumyeong Ahn, Kian Ahrabian, Sanchit Ahuja, Lin Ai, Qingyao Ai, Wei Ai, Ankit Aich, Taichi Aida, Laura Aina, Akiko Aizawa, Pritom Saha Akash, Mohammad Akbar-tajari, Mohammad Akbari, Aditya Akella, Benjamin Akera, Mahmud Elahi Akhter, Berfin Aktas, Syeda Sabrina Akter, Hend Al-khalifa, Hadeel Al-negheimish, Nora Al-twairesh, Jesujoba Oluwada-dara Alabi, Ozge Alacam, Mehwish Alam, Lisa Alazraki, Hanan Aldarmaki, Anton Mikhailovich Alekseev, Laura Alonso Alemany, Georgios Alexandridis, David Alfter, Israa Alghanmi, Bashar Alhafni, Alnur Ali, Felermimo D. M. A. Ali, Iqra Ali, Sarwan Ali, Ilseyar Alimova, Pegah Alipoor-molabashi, Dan Alistarh, Badr Alkhamissi, Hani Alomari, Iñigo Alonso, Kenneth Alperin, Sawsan Alqahtani, Hilal Alquabeh, Ana Alves, Duarte Miguel Alves, Fakhraddin Alwajih, Zaid Alyafeai, Mahsa Amani, Alfonso Amayuelas, Maxime Amblard, Enrique Amigo, Muhammad Saad Amin, Saadullah Amin, Silvio Amir, Samuel Joseph Amouyal, Bang An, Chenxin An, Haozhe An, Kai-kai An, Na Min An, Shengnan An, Wenbin An, Zhenwei An, Aliko Anagnostopoulou, Ashish Anand, Amith Ananthram, Judith Jeyafreeda Andrew, Nicholas Andrews, Ion Androutsopoulos, Piotr Andruszkiewicz, Rico Angell, Krasimir Angelov, Tatiana Anikina, Sifatul Anindho, Murali Annavaram, Miriam Anschütz, Wang Ante, Jacy Reese Anthis, Diego Matteo Antognini, Dimosthenis Antypas, David Anugraha, Abrar Anwar, Saba Anwar, Junyi Ao, Tatsuya Aoyama, Srikar Appalaraju, Mario Ezra Aragón, Eiji Aramaki, Oscar Araque, Pedro Henrique Luz De Araujo, Jerónimo Arenas-garcía, Arturo Argueta, Esteban Garces Arias, Samee Arif, Nurul Fajrin Ariyani, Konstantine Arkoudas, Jordi Armengol-estapé, Akshatha Arodi, Akhil Arora, Arnav Arora, Gaurav Arora, Sanchit Arora, Siddhant Arora, Ines Arous, David Arps, Mousa Arraf, Katya Artemova, Mikel Artetxe, Keren Artiaga, Mohammad Arvan, Bibek Aryal, Daiki Asami, Yuki Asano, Nicholas Asher, Saleh Ashkboos, Joshua Ashkinaze, Salena Torres Ashton, Dennis Assenmacher, Md. Atabuzzaman, Duygu Ataman, Thushari Atapattu, Yash Kumar Atri, Omar Attia, Joseph Attieh, Berk Atıl, P. V. S. Avinesh, Abinew Ali Ayele, Ansar Aynetdinov, Hammad Ayyubi, Amos Azaria, Israel Abebe Azime, Abdul Aziz, Dilshod Azizov, Gorka Azkune, Daichi Azuma, Matthias Aßemacher

Premjith B, Yukino Baba, Nikolay Babakov, Susanna Maria Baby, Stephen Bach, Sher Badshah,

Ingeol Baek, Jinheon Baek, Roxanne El Baff, Niyati Bafna, Mahdieh Baghshah, Seyed Ali Bahrainian, Ge Bai, Haoli Bai, He Bai, Jiaqi Bai, Jun Bai, Long Bai, Xiang Bai, Xiaopeng Bai, Yimeng Bai, Yu Bai, Yu Bai, Yuelin Bai, Yuyang Bai, Yuzhuo Bai, Divya Jyoti Bajpai, Pradeep Bajracharya, Gurkirat Singh Bajwa, George Arthur Baker, Yavuz Faruk Bakman, Oana Denisa Balalau, Nihal Balani, Ananth Balashankar, Gargi Balasubramaniam, Nishant Balepur, Pier Felice Balestrucci, Wolf-tilo Balke, Dibyanayan Bandyopadhyay, Rabiraj Bandyopadhyay, Sambaran Bandyopadhyay, Somnath Banerjee, Yejin Bang, Naman Bansal, Srijan Bansal, Guangsheng Bao, Jianzhu Bao, Keqin Bao, Rong Bao, Runxue Bao, Siqi Bao, Yuwei Bao, Kfir Bar, Libby Barak, Serguei Barannikov, Alexander Baranov, David La Barbera, M Saiful Bari, Guntis Barzdins, Arslan Basharat, Pierpaolo Basile, Elias Bassani, Elisa Bassignana, Abhipsa Basu, Kinjal Basu, Samyadeep Basu, Riza Batista-navarro, Kayhan Batmanghelich, Anil Batra, Soumya Batra, John Bauer, Tim Baumgärtner, Rachel Bawden, Deniz Bayazit, Adrián Bazaga, Omid Bazgir, Adrien Bazoge, Klaudia Bałazy, Nathanaël Beau, Helard Becerra, Tilman Beck, Karin Becker, Srikanta J. Bedathur, Peter Anthony Beerel, Parishad Behnamghader, Andrei Stefan Bejgu, Gemma Belenguix, Tadesse Destaw Belay, Peter Belcak, Peter Bell, Jonas Belouadi, Tony Belpaeme, Luisa Bentivogli, Uri Berger, Benjamin Bergner, Jordi Bernad, Guillaume Bernard, Timothée Bernard, Raffaella Bernardi, Gabriel Bernier-colborne, Dario Bertero, Amanda Bertsch, Katrien Beuls, Michele Bevilacqua, Marie Bexte, Sourav Sanjukta Bhabesh, Shilpa Bhagavath, Milan Bhan, Manasa Bharadwaj, Rohit K Bharadwaj, Suma Bhat, Hanoz Bhathena, Gagan Bhatia, Rohan Bhatia, Sumit Bhatia, Shaily Bhatt, Amrita Bhattacharjee, Arnab Bhattacharya, Indrajit Bhattacharya, Sourangshu Bhattacharya, Pushpak Bhattacharyya, Bhavya Bhavya, Nikhil Bhendawade, Rajarshi Bhowmik, Amran Bhuiyan, Baolong Bi, Bin Bi, Guanqun Bi, Keping Bi, Ziqian Bi, Laura Biester, Shankar Biradar, Steven Bird, Tegawendé F. Bissyandé, Johanna Björklund, Arno Blaas, Verena Blaschke, Rexhina Blloshmi, Frederic Blum, Boboboliu, Joanne Boisson, Marco Bombieri, Vamshi Krishna Bonagiri, Francis Bond, Claire Bonial, Philipp Borchert, Łukasz Borchmann, Claudia Borg, Hsuvas Borkakoty, Emanuela Boros, Daniel Borrajo, Matteo Bortoletto, Antal Van Den Bosch, Cristina Bosco, Digbalay Bose, Sombit Bose, Abdesselam Boucekif, Tom Bourgeade, Andrey Bout, Johan Boye, Leonid Boytsov, Murad Bozik, Marco Braga, Maharaj Brahma, David Adam Braude, Daniel Braun, Jannik Brinkmann, Tomás Vergara Browne, Dominique Brunato, Elia Bruni, Mikael Brunila, Alessio Brutti, Christopher Bryant, Xingyuan Bu, Yuqi Bu, Ana-maria Bucur, Kyle Robert Buettner, Nam Khac-hoai Bui, Trung Bui, Andreas Bulling, Laurie Burchell, Sergio Gastón Burdisso, Sophie Burkhardt, Andrea Burns, Mikhail Burtsev, Hendrik Buschmeier, Yash Parag Butala, Miriam Butt, Happy Buzaaba, Bill Byrne, Hoyoon Byun, Sungjoo Byun, Necva Bölücü

Isabel Cachola, Luca Cagliero, Erica Cai, Fengyu Cai, Guohao Cai, Hengyi Cai, Jinglun Cai, Kunlin Cai, Li Cai, Ming Cai, Pengshan Cai, Ruichu Cai, Shu-bin Cai, Weidong Cai, Yi Cai, Yiqing Cai, Yujun Cai, Yuxiang Cai, Zefan Cai, Cristian Buc Calderon, Aylin Caliskan, Iacer Calixto, Jesus Calvillo, Nguyen Cam-tu, Erik Cambria, Necati Cihan Camgoz, Stefano Campese, Duy-cat Can, Marc Canby, Nicola Cancedda, Marie Candito, Biwei Cao, Bochuan Cao, Bowen Cao, Boxi Cao, Hailong Cao, Hannan Cao, Hongye Cao, Jiuxin Cao, Juan Cao, Lang Cao, Liuwen Cao, Minxuan Cao, Qian Cao, Qingxing Cao, Rui Cao, Ruisheng Cao, Shitong Cao, Shuyang Cao, Yanan Cao, Yixuan Cao, Yu Cao, Yu Cao, Yuan Cao, Yunbo Cao, Yupeng Cao, Yuxuan Cao, Zhao Cao, Zhe Cao, Zhiyu Cao, Zouying Cao, Victor Carbune, Ronald Cardenas, Rémi Cardon, Henrique Lopes Cardoso, Luigi Di Caro, Lucien Carroll, Danilo Carvalho, Tommaso Caselli, Silvia Casola, Roman Castagné, Giuseppe Castellucci, Camilla Casula, Paulo Rodrigo Cavalin, Francesco Cazzaro, Jan Cegin, Yukuo Cen, Dong-kyu Chae, Hyungjoo Chae, Joyce Chai, Linzheng Chai, Wenhao Chai, Yekun Chai, Ziwei Chai, Megha Chakraborty, Souradip Chakraborty, Gaël De Chalendar, Alvin Chan, Chimin Chan, David Chan, Hou Pong Chan, Sophia Chan, Christina A Chance, Navin Chandak, Mohit Chandra, Aofei Chang, Buru Chang, Heng Chang, Ma Chang, Maria Chang, Ruidi Chang, Serina Chang, Ting-yun Chang, Tyler A Chang, Yapei Chang,

Yin-wen Chang, Zhiyuan Chang, Wenhan Chao, Yang Chao, Thierry Charnois, Lucas Georges Gabriel Charpentier, Anwoy Chatterjee, Usashi Chatterjee, Akshay Chaturvedi, Sagar Chaturvedi, Duen Chau, Anuj Chauhan, Arnav Santosh Chavan, Kushal Chawla, Muthusamy Chelliah, Ivan Chelombiev, Emmanuel Chemla, Andong Chen, Anthony Chen, Beiduo Chen, Ben Chen, Berlin Chen, Bin Chen, Bingsen Chen, Bo Chen, Bo Chen, Boxing Chen, C. L. Philip Chen, Canyu Chen, Cen Chen, Chacha Chen, Changyu Chen, Chen Chen, Chenhua Chen, Chi Chen, Chien Hung Chen, Congliang Chen, Danlu Chen, Delong Chen, Dingfan Chen, Dongping Chen, Dubing Chen, Enhong Chen, Feng Chen, Gang Chen, Guanjie Chen, Guanyi Chen, Guanzheng Chen, Guiming Hardy Chen, Guizhen Chen, Haibin Chen, Haihua Chen, Hang Chen, Hannah Chen, Hanzhu Chen, Hao Chen, Hao Chen, Hao Chen, Haonan Chen, Hong Chen, Hongyang Chen, Hongyu Chen, Hongyu Chen, Hongzhan Chen, Huiyao Chen, Hung-ting Chen, Jian Chen, Jianda Chen, Jiaqi Chen, Jiawei Chen, Jingchang Chen, Jinghui Chen, Jingyi Chen, Jiuhai Chen, Jixuan Chen, Jiyu Chen, John Chen, Junhao Chen, Junjie Chen, Junying Chen, Junzhe Chen, Kai Chen, Kai Chen, Kai Chen, Kaiyuan Chen, Kang Chen, Kangjie Chen, Kedi Chen, Kejia Chen, Kejiang Chen, Kuang-ming Chen, Lei Chen, Leshang Chen, Liang Chen, Lihu Chen, Lin Chen, Lizhi Chen, Lizhong Chen, Long Chen, Longze Chen, Luoxin Chen, Meiqi Chen, Meng Chen, Mouxiang Chen, Nuo Chen, Nuo Chen, Nuo Chen, Patrick Chen, Pinzhen Chen, Po-chun Chen, Qian Chen, Qiguang Chen, Qin Chen, Qingcai Chen, Qizhou Chen, Qun Chen, Ruijun Chen, Ruirui Chen, Ruizhe Chen, Shan Chen, Shijie Chen, Shisong Chen, Shiwei Chen, Shuo Chen, Simin Chen, Sishuo Chen, Tianlong Chen, Tianxiang Chen, Tong Chen, Tong Chen, Tongfei Chen, Valerie Chen, Wei Chen, Wei Chen, Wei-fan Chen, Weidong Chen, Weize Chen, Wenqing Chen, Wenting Chen, Wentong Chen, Xi Chen, Xi Chen, Xiang Chen, Xiaojun Chen, Xie Chen, Xinyue Chen, Xinyue Chen, Xinyue Chen, Xu Chen, Xuanang Chen, Yang Chen, Yang Chen, Yankai Chen, Yanquan Chen, Yanran Chen, Yi Chen, Yi-pei Chen, Yi-syuan Chen, Yifan Chen, Yige Chen, Yijun Chen, Yilong Chen, Yin-jen Chen, Yingfa Chen, Yiqun Chen, Yiran Chen, Yiwen Chen, Yixin Chen, Yiyi Chen, Yongrui Chen, Yuhan Chen, Yuheng Chen, Yulin Chen, Yun Chen, Yuxuan Chen, Yuyan Chen, Zhanpeng Chen, Zhengyu Chen, Zhi-yuan Chen, Zhihong Chen, Zhipeng Chen, Zhongwu Chen, Zhumin Chen, Zhuowei Chen, Ziru Chen, Zixin Chen, Ziyang Chen, Zulong Chen, Daixuan Cheng, Dawei Cheng, Fan Cheng, Gang Cheng, Hong Cheng, Jiale Cheng, Jian Cheng, Jianpeng Cheng, Jiayang Cheng, Jingwei Cheng, Julius Cheng, Kewei Cheng, Liang Cheng, Luyao Cheng, Minhao Cheng, Myra Cheng, Pengzhou Cheng, Qinyuan Cheng, Reynold Cheng, Sijie Cheng, Sitao Cheng, Siyuan Cheng, Wei Cheng, Xianfu Cheng, Xiaoxia Cheng, Xiaoxue Cheng, Xin Cheng, Xinle Cheng, Xize Cheng, Yao Cheng, Yi Cheng, Zehua Cheng, Zifeng Cheng, Liu Chengwu, Lin Lee Cheong, Yun-gyung Cheong, Emmanuele Chersoni, Rendi Chevi, Oscar Chew, Anshuman Chhabra, Pohan Chi, Yew Ken Chia, Cheng-han Chiang, Ting-rui Chiang, Si-chi Chin, Abhinav Chinta, Patricia Chiril, Luis Chiruzzo, Elena Chistova, Pranjal A Chitale, Nima Chitsazan, Ming-chang Chiu, Yu Ying Chiu, Hyundong Justin Cho, Hyunsouk Cho, Ikhyun Cho, Ji Yong Cho, Kyunghyun Cho, Myeongah Cho, Sukmin Cho, Won Ik Cho, Young Min Cho, Rochelle Choenni, Donghee Choi, Eunkyung Choi, Hongseok Choi, Jaesik Choi, Jason Ingyu Choi, Joon-young Choi, Juhwan Choi, Jungwook Choi, Minseok Choi, Sehyun Choi, Seungtaek Choi, Sungchul Choi, Yongsuk Choi, Yubin Choi, Mukund Choudhary, Monojit Choudhury, Sagnik Ray Choudhury, Katsuki Chousa, Yinlam Chow, Arijit Ghosh Chowdhury, Somnath Basu Roy Chowdhury, Tanya Chowdhury, George Chrysostomou, Minh Duc Chu, Xiaowen Chu, Zheng Chu, Tat-seng Chua, Yu-neng Chuang, Yun-shiuan Chuang, Yung-sung Chuang, Ekapol Chuangsuwanich, Chanyoung Chung, Hyunseung Chung, John Joon Young Chung, Abu Nowshed Chy, Philipp Cimiano, Jorge Civera, Christian Clark, Christopher Clark, Elizabeth Clark, Christopher Clarke, Chloé Clavel, Stéphane Clinchant, Miruna Adriana Clinciu, Maximin Coavoux, Alexandru Coca, Oana Cocarascu, Amir David Nissan Cohen, Daniel Cohen, Nachshon Cohen, Vanya Cohen, Vincent Cohen-addad, Luisa Coheur, Trevor Cohn, Eliana Colunga, Prakash Mandayam Comar, Tianshuo Cong, Xin Cong, Mathieu Constant, Fierro Constanza, Danish Contractor, Owen Cook, Rodolfo Corona, Luciano Del Corro, Adrian Cosma, Luis Morgado Da Costa, Marta R. Costa-

jussà, Adriana Valentina Costache, Ryan Cotterell, Mickael Coustaty, Alan Cowap, Florinel Alin Croitoru, Huimin Cui, Jiajia Cui, Jin Cui, Qu Cui, Shaobo Cui, Shiyao Cui, Wanyun Cui, Wendi Cui, Yuanning Cui, Washington Cunha, Anna Currey, Alba Cercas Curry, Amanda Cercas Curry, Tamara Czinczoll, Erion Çano

Mathieu D'aquin, Jennifer D'souza, Gautier Dagan, Ido Dagan, Nico Daheim, Deborah A. Dahl, Daniel Dahlmeier, Bin Dai, Hong-jie Dai, Hongliang Dai, Huangyu Dai, Quanyu Dai, Shih-chieh Dai, Sunhao Dai, Xiang Dai, Yong Dai, Parag Pravin Dakle, Dwip Dalal, David Dale, Hercules Dalianis, Majid Daliri, Elizabeth M. Daly, Greta Damo, Soham Dan, Yuhao Dan, Ankit Dangi, Mai Hoang Dao, Zhang Daoan, Guy Dar, Anirban Das, Dipankar Das, Sarkar Snigdha Sarathi Das, Souvik Das, Mihai Dascalu, Debajyoti Dasgupta, Sourish Dasgupta, Tirthankar Dasgupta, Saloni Dash, Do Huu Dat, Debtanu Datta, Sam Davidson, Brian Davis, Christopher Davis, Hillary Dawkins, Min-yuh Day, Jiang Dazhi, Nicholas Deas, Alok Debnath, Hiroyuki Deguchi, Derrick Goh Xin Deik, Jean-benoit Delbrouck, Gianluca Demartini, Çağatay Demiralp, Dina Demner-fushman, Steve Deneefe, Ailin Deng, Cheng Deng, Chunyuan Deng, Guanzhi Deng, Hexuan Deng, Jiawen Deng, Jingcheng Deng, Naihao Deng, Shuiguang Deng, Shuwen Deng, Yue Deng, Zhenyun Deng, Zheyue Deng, Zhijie Deng, Pavel Denisov, Omer Faruk Deniz, Sourabh Dattatray Deoghare, Al Depope, Jan Milan Deriu, Franck Dernoncourt, Aniket Deroy, Akshar Prabhu Desai, Meera Desai, Tejash Desai, Maunendra Desarkar, Vijeta Deshpande, Nina Dethlefs, Chaitanya Devaguptapu, Joseph Dexter, Suvodip Dey, Tejas Dhamecha, Prajit Dhar, Ruchira Dhar, Paramveer Dhillon, Aissatou Diallo, Xiaolei Diao, Xingjian Diao, Adji Bouso Dieng, Stefan Dietze, Barbara Dieugenio, Dimitar Iliyanov Dimitrov, Vania Dimitrova, Jacob Dineen, Bosheng Ding, Chuanghao Ding, Dujian Ding, Guiguang Ding, Hanxing Ding, Peng Ding, Shuangrui Ding, Shuoyang Ding, Wenjian Ding, Wenxuan Ding, Xiao Ding, Yifan Ding, Yifeng Ding, Yifu Ding, Yihao Ding, Ying Ding, Ying Ding, Yuyang Ding, Zifeng Ding, Zixiang Ding, Song Dingjie, Saket Dingliwal, Tu Anh Dinh, Tanvi Dinkar, Shubhashis Roy Dipta, Ajay Divakaran, Anuj Diwan, Amirbek Djanibekov, Heejin Do, Jaeyoung Do, Phong Nguyen-thuan Do, Susana Sotelo Docio, Sumanth Doddapaneni, Rama Doddipatla, Pierre Dognin, Kosuke Doi, Tobias Domhan, Gregor Donabauer, Bin Dong, Daize Dong, Guanting Dong, Hang Dong, Haoyu Dong, Jianfeng Dong, Junnan Dong, Kuicai Dong, Lu Dong, Ming Dong, Mingwen Dong, Peijie Dong, Rui Dong, Xiangjue Dong, Yijiang River Dong, Yuhao Dong, Yushun Dong, Zican Dong, Bonaventure F. P. Dossou, Longxu Dou, Qingyun Dou, Yao Dou, Antoine Doucet, Mohammad Javad Dousti, Mauro Dragoni, Mark Dras, Mark Dredze, Felix Drinkall, Rotem Dror, Dayou Du, Gaoyuan Du, Haowei Du, Jusen Du, Minxin Du, Weihong Du, Wenjie Du, Wenyu Du, Xiaoyong Du, Yali Du, Yanrui Du, Yifan Du, Yiming Du, Yupei Du, Haodong Duan, Jiaxin Duan, Jinhao Duan, Jiuding Duan, Juanyong Duan, Junwen Duan, Lixin Duan, Shaohua Duan, Sufeng Duan, Zenghao Duan, Zhichao Duan, Haim Dubossarsky, Fanny Duce, George Dueñas, Richard Dufour, Lavinia Dunagan, Ewan Dunbar, Benjamin Van Durme, Ritam Dutt, Subhabrata Dutta, Jennifer Dy, Nils Dycke, Esra Dönmez

Haihong E, Oliver Eberle, Seth Ebner, Jessica Maria Echterhoff, Ali Edalat, Ali Edalati, Ghebruiot Mohamed Imed Eddine, Lilach Eden, Carl Edwards, Paul Van Eecke, Steffen Eger, Koji Eguchi, Yo Ehara, Maud Ehrmann, Bryan Eikema, Liat Ein-dor, Alon Eirew, Roald Eiselen, Mohamed Elaraby, Passant Elchafei, Hoda Eldardiry, Mohamed Elgaar, T. Mark Ellison, Muhammad Nael Elnokrashy, Tamer Elsayed, Mai Elsherief, Micha Elsner, Yahya Emara, Matthew M. Engelhard, Sugyeong Eo, Elena V. Epure, Pierre Erbacher, Maksim Ereemeev, Justus-jonas Christoph Erker, Brendan Raffel Evers, Chandra Kiran Reddy Evuru, Ana Ezquerro

Marzieh Fadaee, Fartash Faghri, Md Fahim, Neele Falk, Boi Faltings, Jiabin Fan, Rui Fan, Shengda Fan, Siqi Fan, Wenqi Fan, Xiran Fan, Yao-chung Fan, Yongqi Fan, Yue Fan, Zhaoxin Fan, Zhiyuan Fan, Chunrong Fang, Feiteng Fang, Guian Fang, Haishuo Fang, Jiabao Fang, Jinyuan Fang, Ming Fang, Minghong Fang, Minghui Fang, Qingkai Fang, Qixiang Fang, Tao Fang, Tianqing

Fang, Wei Fang, Xinyu Fang, Yi Fang, Yin Fang, Yuan Fang, Yuejian Fang, Zhaohao Fang, Zhixuan Fang, Youmna Farag, Reza Farahbakhsh, Gonçalo Rui Faria, Delia Irazu Hernandez Farias, António Farinhas, Richard Farkas, Sakina Fatima, Pedro Faustini, Mohsen Fayyaz, Arnisa Fazla, Pooyan Fazli, Amir Feder, Fabio James Fehr, Yu Fei, Zhaoye Fei, Nils Feldhus, Sergey Feldman, Andrew Zhuoer Feng, Dongji Feng, Guhao Feng, Huawen Feng, Lei Feng, Tao Feng, Weixi Feng, Xiaoyue Feng, Xincan Feng, Yang Feng, Yansong Feng, Yi Feng, Yu Feng, Yubo Feng, Yujie Feng, Zihao Feng, Zijian Feng, Pan Fengjun, Alena Fenogenova, Teddy Ferdinan, Martin Ferienc, Libby Ferland, Jared Fernandez, Nigel Fernandez, Aloka Fernando, Thomas Palmeira Ferraz, Sheridan Feucht, Julian Fierrez, Elena Filatova, James Darrell Finch, Sarah E. Finch, Matthew Finlayson, Hamed Firooz, Asja Fischer, Jillian Fisher, Jack Fitzgerald, Margaret Fleck, Eve Fleisig, James Flemings, Lorenzo Jaime Yu Flores, José A.r. Fonollosa, Marcio Fonseca, Negar Foroutan, Karën Fort, Robert Frank, Stella Frank, Alexander M. Fraser, Diego Frassinelli, Dayne Freitag, Simon Frieder, Giacomo Frisoni, Markus Frohmann, Arvid Frydenlund, Biao Fu, Dayuan Fu, Harvey Yiyun Fu, Jinlan Fu, Peng Fu, Tingchen Fu, Xiyan Fu, Yicheng Fu, Yun Fu, Zichuan Fu, Zuohui Fu, Dennis Fucci, Yoshinari Fujinuma, Fumiyo Fukumoto, Kotaro Funakoshi, Hiroaki Funayama, Shaz Furniturewala, Hiroki Furuta

Matteo Gabburo, Guillaume Gadek, Gaelguibon, Baban Gain, Robert J. Gaizauskas, Ya'akov Gal, Jay Gala, Alessio Galatolo, Marcos Martínez Galindo, Lukas Paul Achatius Galke, Martina Galletti, Jack Gallifant, Aram Galstyan, Pablo Gamallo, Andrew Gambardella, Yanglei Gan, Yujian Gan, Sriram Ganapathy, Mona Gandhi, Adithya V Ganesan, Balaji Ganesan, Debasis Ganguly, William Gantt, Tanuja Ganu, Bofei Gao, Changjiang Gao, Chao Gao, Chen Gao, Chufan Gao, Cuiyun Gao, Dehong Gao, Difei Gao, Ge Gao, Hongcheng Gao, Jiahui Gao, Jing Gao, Jinglong Gao, Jingsheng Gao, Jinhua Gao, Jun Gao, Jun Gao, Lang Gao, Mingqi Gao, Muhan Gao, Neng Gao, Pengzhi Gao, Qiang Gao, Qiyue Gao, Rena Wei Gao, Rujun Gao, Shangqian Gao, Shen Gao, Shengxiang Gao, Songyang Gao, Wei Gao, Xiang Gao, Xiaoxue Gao, Yingbo Gao, Yisen Gao, Zhujin Gao, Claire Gardent, Ankur Garg, Shweta Garg, Joseph Gatto, Himaanshu Gauba, Eric Gaussier, Vagrant Gautam, Jiaxin Ge, Mengshi Ge, Shiping Ge, Xiou Ge, Yubin Ge, Tom Gedeon, Michaela Geierhos, Robert Geislinger, Aryo Pradipta Gema, Debela Tesfaye Gemechu, Josef Van Genabith, Barbara Gendron, Ruiying Geng, Saibo Geng, Xiang Geng, Yuxia Geng, Enfa Rose George, Ariel Gera, Walter Gerych, Sadaf Ghaffari, Abdul Ghafoor, Omid Ghahroodi, Karim Ghonim, Kripabandhu Ghosh, Sayan Ghosh, Shaona Ghosh, Soumitra Ghosh, Sreyan Ghosh, Michael Ginn, Luca Gioffrè, Leander Girrbach, Raja Giryes, Adrià De Gispert, Fausto Giunchiglia, Michael Glass, Anna Glazkova, Kristina Gligoric, Max Glockner, Ameya Godbole, Anmol Goel, Arnav Goel, Evangelia Gogoulou, Sim Kuan Goh, Usman Gohar, Marcel Gohsen, Yoav Goldberg, Jonas Golde, Tomas Goldsack, Shyamnath Gollakota, Sujatha Das Gollapalli, Michal Golovanevsky, Juan Pablo Zuluaga Gomez, Marcos A Goncalves, Elizaveta Goncharova, Chen Gong, Chen Gong, Cheng Gong, Jiaying Gong, Kaixiong Gong, Neil Gong, Shansan Gong, Yihong Gong, Zhiguo Gong, Ziwei Gong, Aitor Gonzalez-agirre, Julio Gonzalo, Deepak Gopinath, Dhiman Goswami, Koustava Goswami, Sujata Goswami, Boyu Gou, Yunhao Gou, Antoine Gourru, Venkata Subrahmanyam Govindarajan, Pawan Goyal, Jorge Gracia, Carine Graff, Michael Granitzer, Christan Grant, Candida Maria Greco, Tommaso Green, Derek Greene, Rachel Greenstadt, Matt Grenander, Thomas L. Griffiths, Charlie Grimshaw, Milan Gritta, Georg Groh, Dagmar Gromann, Paul Groth, Bowen Gu, Hengrui Gu, Jia-chen Gu, Jinghang Gu, Naibin Gu, Shuhao Gu, Tiancheng Gu, Tianle Gu, Tianlong Gu, Xiangming Gu, Yi Gu, Yu Gu, Yuxian Gu, Yuxuan Gu, Zhouhong Gu, Zijin Gu, Eleonora Gualdoni, Jian Guan, Nan Guan, Quanlong Guan, Shuhao Guan, Xin Guan, Xinyan Guan, Yong Guan, Li Hong Guang, Reto Gubelmann, Bhanu Prakash Reddy Guda, Imane Guellil, Marco Guerini, Michael Guerzhoy, Dipanwita Guhathakurta, Mustafa Omer Gul, Varun Gumma, Tom Gunter, Boyang Guo, Fenfei Guo, Han Guo, Hangyu Guo, Hongcheng Guo, Jia Guo, Jiaxin Guo, Jinyu Guo, Kan Guo, Kehan Guo, Lingbing Guo, Meiqi Guo, Qingpei Guo, Qipeng Guo, Ruohao Guo, Shangwei Guo, Shaoru Guo, Shasha Guo,

Shoutao Guo, Siyi Guo, Siyuan Guo, Taicheng Guo, Wenya Guo, Xiaobo Guo, Xinnan Guo, Xu Guo, Yanzhu Guo, Yike Guo, Yue Guo, Yuhang Guo, Zhen Guo, Zhicheng Guo, Zhiyu Guo, Ziyu Guo, Akash Gupta, Ankita Gupta, Ashim Gupta, Deepak Gupta, Gopal Gupta, Jai Gupta, Lavanya Gupta, Manish Gupta, Sameer Gupta, Sonam Gupta, Umang Gupta, Vipul Gupta, Timon Gurcke, Iryna Gurevych, V Gurucharan, Sireesh Gururaja, Yoan Gutierrez, Ximena Gutierrez-vasques, Daehoon Gwak, Jeremy Gwinnup

Thanh-le Ha, Shohreh Haddadan, Nikita Haduong, Samar Haider, Mao Haitao, Jan Hajič, Huda Hakami, Ahmad Dawar Hakimi, Ido Hakimi, Sherzod Hakimov, Scott A. Hale, Karina H Halevy, Coleman Haley, Patrick Haller, Skyler Hallinan, Hassan Hamad, Karen Hambardzumyan, Ayoub Hammal, Boran Han, Caren Han, Cheng Han, Chengcheng Han, Donghoon Han, Dou Shi Han, Haoyu Han, Hojae Han, Janghoo Han, Jiale Han, Jiuzhou Han, Kanyao Han, Lifeng Han, Meng Han, Namgi Han, Seungju Han, Shijie Han, Simeng Han, Wei Han, William Han, Xiaoqi Han, Xu Han, Xu Han, Xudong Han, Xue Han, Yang Han, Yo-sub Han, Ziwen Han, Adeep Hande, Zhang Hanlin, Michael Hanna, Chuzhan Hao, Hailin Hao, Jianye Hao, Qirui Hao, Tianyong Hao, Xiaoyang Hao, Yaru Hao, Yongchang Hao, Wang Haotian, Rejwanul Haque, Romain Harang, David Harbecke, Karin Harbusch, Momchil Hardalov, Fabrice Y Harel-canada, Stefan Harmeling, Ian Harris, Thomas Hartvigsen, Md Rakibul Hasan, Md. Arid Hasan, Mohammad Saqib Hasan, Maram Hasanain, Taku Hasegawa, Mark A. Hasegawa-johnson, Kazuma Hashimoto, Faegheh Hasiibi, Islam A. Hassan, Michael Hassid, Nabil Hathout, Hans Ole Hatzel, Shreya Havaladar, Akio Hayakawa, Katsuhiko Hayashi, Amir Hazem, Rima Hazra, Bailan He, Bingsheng He, Bingxiang He, Bolei He, Bowei He, Chaoqun He, Conghui He, Estrid He, Guoxiu He, Hangfeng He, Jianfeng He, Jianguan He, Junda He, Junqing He, Kai He, Kun He, Liqi He, Liyang He, Pengfei He, Pengfei He, Pinjia He, Qianyu He, Qing He, Ruifang He, Shwai He, Tao He, Tingting He, Xiangjian He, Xiangnan He, Xingwei He, Xinyi He, Xuehai He, Xuming He, Yifei He, Zexue He, Zhaofeng He, Zhitao He, Zihao He, Jennifer Healey, Larry Heck, Tobias Hecking, Johannes Heinecke, Maximilian Heinrich, Benjamin Heinzerling, Jindřich Helcl, William Barr Held, Oliver Hellwig, Chadi Helwe, Bhatiya Hemanthage, James Henderson, Iris Hendrickx, Amey Hengle, Aron Henriksson, Christian Herold, Jack Hessel, Anthony Hills, Tatsuya Hiraoka, Eran Hirsch, Julia Hirschberg, Youssef Hmamouche, Namgyu Ho, Xanh Ho, Nhat Minh Hoang, Thai Quoc Hoang, Thanh Tung Hoang, Vu Cong Duy Hoang, Sanne Hoeken, David Crossland Hogg, Carolin Holtermann, Ukyo Honda, Vlad Hondru, Giwon Hong, Jiwoo Hong, Jixiang Hong, Pengfei Hong, Ruixin Hong, Seokhee Hong, Yihuai Hong, Yu Hong, Yuan Hong, Zijin Hong, Hongbinzhanghitsz, Liu Hongjun, Jia Hongrui, Or Honovich, Anthony Hoogs, Zachary William Hopton, Mohammed Moshul Hoque, Andrea Horbach, Zachary Horvitz, Sho Hoshino, Faezeh Hosseini, Peyman Hosseini, Veronique Hoste, Abe Bohan Hou, Bairu Hou, Benjamin Hou, Guiyang Hou, Lu Hou, Rui Hou, Wenjun Hou, Yifan Hou, Yingyan Hou, Yu Hou, Zejiang Hou, Zhenyu Hou, Phillip Howard, David M Howcroft, Nick Howell, Chi-yuan Hsiao, I-hung Hsu, Wynne Hsu, Yili Hsu, Phu-mon Htut, Baotian Hu, Bodun Hu, Brian H Hu, Chenxi Hu, Guimin Hu, Hai Hu, Haifeng Hu, Hanxu Hu, He Hu, Hongxin Hu, Jinpeng Hu, Jinyi Hu, Jun Hu, Lianyu Hu, Lijie Hu, Mengkang Hu, Mengzhou Hu, Minda Hu, Minghao Hu, Nan Hu, Qi Hu, Qisheng Hu, Renfen Hu, Renjun Hu, Songbo Hu, Tiancheng Hu, Vincent Tao Hu, Wenbin Hu, Wenbo Hu, Wenxiang Hu, Xiao Hu, Xiaomeng Hu, Xiaoyu Hu, Xinshuo Hu, Xinyu Hu, Xiyang Hu, Yebowen Hu, Yiran Hu, Yuchen Hu, Yutong Hu, Zhe Hu, Zhengmian Hu, Zhengyu Hu, Zhiyuan Hu, Zi-yuan Hu, Hang Hua, Ting Hua, Wenyue Hua, Yuncheng Hua, Baizhou Huang, Chao-wei Huang, Chengkai Huang, Chengsong Huang, Chengyu Huang, Chengyue Huang, Chenyang Huang, Chu-ren Huang, Chun-ying Huang, Dong Huang, Fan Huang, Feiran Huang, Hanxian Huang, Hao Huang, Haojing Huang, Haoyu Huang, Hen-hsen Huang, Heng Huang, Hongtao Huang, Hsiu-yuan Huang, Hui Huang, Hujiang Huang, Jen-tse Huang, Jerry Huang, Jianheng Huang, Jimin Huang, Jin-xia Huang, Jinfa Huang, Jingyuan Huang, Junwei Huang, Kaiyu Huang, Kaizhu Huang, Kun Huang, Langlin Huang, Lei Huang, Li Huang, Liusheng Huang, Minghui Huang, Minlie Huang, Peijie Huang, Pengcheng Huang,

Qingbao Huang, Qiushi Huang, Quzhe Huang, Rikui Huang, Ruixuan Huang, Ruizhang Huang, Shijia Huang, Shizhou Huang, Shujian Huang, Shuo Huang, Siyuan Huang, Wei Huang, Weijing Huang, Wenhao Huang, Wenxin Huang, Xiang Huang, Xiaolei Huang, Yan Huang, Yi Huang, Yin Jou Huang, Yinxuan Huang, Yongfeng Huang, Youcheng Huang, Yu-shiang Huang, Yue Huang, Yuhao Huang, Yukun Huang, Zanming Huang, Zeyu Huang, Zhen Huang, Zhichao Huang, Zhongzhan Huang, Zixian Huang, Yuanxiang Huangfu, Zhang Huaping, Wang Huazheng, Patrick Huber, Frederikus Hudi, Mathew John Huerta-enochian, Bo Hui, Zheng Hui, Rudali Huidrom, Huiyingcao, Huiyingcao, Lawrence Hunter, Shuning Huo, Maeve Hutchinson, Nguyen Thanh Huy, Dae Yon Hwang, Eunjeong Hwang, Jena D. Hwang, Seonjeong Hwang, Truong Son Hy, Katharina Hämmerl, Mika Hämäläinen, Anna Häty

Mahardika Krisna Ihsani, Hal Daumé Iii, Taichi Iki, Nikolai Ilinykh, Dmitry Ilvovsky, Solee Im, Ayyoob Imani, Joseph Marvin Imperial, Muhammad Imran, Michimasa Inaba, Mert Inan, Go Inoue, Ioan-bogdan Iordache, David Martinez Iraola, Ozan Irsoy, Yoichi Ishibashi, Tatsuya Ishigaki, Alvi Md Ishmam, Md Asiful Islam, Md Rakibul Islam, Md Saiful Islam, Mohammed Saidul Islam, Mete Ismayilzada, Takumi Ito, Alexei V. Ivanov, Vladimir Ivanov, Hamish Ivison, Yusuke Iwasawa, Bhavani Iyer, Rishabh Iyer, Gautier Izacard

Karthika N J, Manoj Balaji J, Musashi Jacobs-harukawa, Aditya Ajay Jadhav, Aishwarya Jadhav, Aref Jafari, Akshay Jagatap, Arushi Jain, Rishabh Jain, Sanil Jain, Sparsh Jain, Yash Jain, Mimansa Jaiswal, Terne Sasha Thorn Jakobsen, Shruti Jalan, Surgan Jandial, Dongjun Jang, Eugene Jang, Jihyoung Jang, Joonwon Jang, Myeongjun Jang, Seongbo Jang, Young Kyun Jang, Yunah Jang, Peter Jansen, Arkadiusz Janz, Amit Jaspal, Sujay Kumar Jauhar, Dávid Javorský, Sai Muralidhar Jayanthi, Nicolaas Paul Jedema, Christopher William Jenkins, Eojin Jeon, Hyewon Jeon, Yejin Jeon, Minbyul Jeong, Soyeong Jeong, Sullam Jeoung, Bin Ji, Donghong Ji, Heng Ji, Jiayi Ji, Ke Ji, Shengpeng Ji, Tianbo Ji, Wei Ji, Yixin Ji, Yuelyu Ji, Zhenzhou Ji, Ziwei Ji, Zongcheng Ji, Furong Jia, Jinghan Jia, Pengyue Jia, Qi Jia, Ruipeng Jia, Xueying Jia, Yuxiang Jia, Liu Jiaheng, Ping Jian, Yiren Jian, Zhongquan Jian, Chao Ya Jiang, Chunyang Jiang, Dongfu Jiang, Dongwei Jiang, Fan Jiang, Fengqing Jiang, Gongyao Jiang, Guochao Jiang, Haiyun Jiang, Jinhao Jiang, Jiyue Jiang, Junfeng Jiang, Junzhe Jiang, Jyun-yu Jiang, Lai Jiang, Lan Jiang, Liwei Jiang, Meng Jiang, Ming Jiang, Nan Jiang, Nan Jiang, Ning Jiang, Qian Jiang, Shan Jiang, Tianyu Jiang, Tintin Jiang, Wenyuan Jiang, Xiaotong Jiang, Yichen Jiang, Yidi Jiang, Yuxin Jiang, Zhiwei Jiang, Zhongtao Jiang, Zhuoren Jiang, Ziyang Jiang, Ziyue Jiang, Zui Jiang, Jianghoucheng, Fangkai Jiao, Jian Jiao, Wenpin Jiao, Wenxiang Jiao, Yizhu Jiao, Bowen Jin, Chuanyang Jin, Guoqing Jin, Hai Jin, Han Jin, Hongye Jin, Jiho Jin, Junfeng Jin, Kyohoon Jin, Li Jin, Lianwen Jin, Lisa Jin, Mali Jin, Mingyu Jin, Peiquan Jin, Renren Jin, Rize Jin, Ruochen Jin, Ruochun Jin, Souyoung Jin, Tao Jin, Xiaolong Jin, Xiaomeng Jin, Xisen Jin, Yaohui Jin, Yeying Jin, Yiqiao Jin, Zhe Jin, Zhi Jin, Zhuoran Jin, Disha Jindal, Ishan Jindal, Liping Jing, Liqiang Jing, Wu Jing, Yuu Jinnai, Hwiyeol Jo, Yongrae Jo, Moa Johansson, Neha Anna John, Kristen Johnson, Jiahuei Joo, Alipio Jorge, Joemon M. Jose, Abhinav Joshi, Aditya Joshi, Anant Joshi, Brihi Joshi, Harshit Joshi, Rishabh Joshi, Mingxuan Ju, Tianjie Ju, Xincheng Ju, Yiming Ju, Gurusha Juneja, Changsoo Jung, Dahyun Jung, Dongwon Jung, Haeji Jung, Sangkeun Jung, Woohwan Jung, Yeonjoon Jung, He Junwei, Lu Junyu, Juraj Juraska, Nika Jurov, Tom S. Juzek

Karthikeyan K, Mohsinul Kabir, Anubha Kabra, David Kaczér, Ram Mohan Rao Kadiyala, Marek Kadlčík, Kazuma Kadowaki, Patrick Kahardipraja, Tomasz Jan Kajdanowicz, Abisek Rajakumar Kalarani, Jan-christoph Kalo, Nishant Kambhatla, Hirotaka Kameko, Sabyasachi Kamila, Ryo Kamoi, Jonathan Kamp, Onno P. Kampman, Jaap Kamps, Mahammed Kamruzzaman, Hiroshi Kanayama, Daniel Kang, Deokhyung Kang, Dongjin Kang, Gi-cheon Kang, Jaewoo Kang, Jaewook Kang, Junmo Kang, Liu Ze Kang, Minki Kang, Pilsung Kang, Xiaoxi Kang, Zhao Kang, Yoshinobu Kano, Hung-yu Kao, Rajvi Kapadia, Pinar Karagoz, Antonia Karamolegkou, Mladen Karan,

Hima Karanam, Amir Hossein Kargar, Priyanka Kargupta, Akbar Karimi, Sanjeev Kumar Karn, Zohar Karnin, Marzena Karpinska, Fakhri Karray, Pradeep Karuturi, Pranav Kasela, Gautam Sid-dharth Kashyap, Gjergji Kasneci, Surya Kasturi, Tushar Kataria, Shunya Kato, Yoshihide Kato, Gregorios A Katsios, Ayush Kaushal, Muhammad Dehan Al Kautsar, Parsa Kavehzadeh, Pride Ka-vumba, Daisuke Kawahara, Masayuki Kawarada, Yoshifumi Kawasaki, Roman Kazakov, Hideto Kazawa, Ashkan Kazemi, Seyed Mehran Kazemi, Wenjun Ke, Xiaolan Ke, Sedrick Scott Keh, Mamadou K. Keita, Katherine A. Keith, Simon Keizer, Amr Keleg, Mikaela Keller, Orgad Keller, Casey Kennington, Roman Kern, Natthawut Kertkeidkachorn, Vaishakh Keshava, Lee Kezar, Ba-ber Khalid, Muhammad Khalid, Osama Khalid, Mohammed Khalilia, Aleem Khan, Latifur Khan, Mohammed Safi Ur Rahman Khan, Salman Khan, Sujay Khandagale, Vivek Khetan, Zahra Kho-daghli, Md Tawkat Islam Khondaker, Ashiqur R. Khudabukhsh, Urja Khurana, Varun Khurana, Dayeon Ki, Johannes Kiesel, Zlata Kikteva, Jihyung Kil, Halil Kilicoglu, Bugeun Kim, Byeong-geun Kim, Chae Won Kim, Chaeun Kim, Changhun Kim, Dahyun Kim, Dongwoo Kim, Eunsu Kim, Geewook Kim, Gene Louis Kim, Gunhee Kim, Gyeongman Kim, Gyuwan Kim, Hansaem Kim, Harksoo Kim, Hazel H. Kim, Heeseung Kim, Hongjin Kim, Hwichan Kim, Hyounghun Kim, Hyuhng Joon Kim, Jaehyung Kim, Jaein Kim, Jaekyeom Kim, Jaeyoung Kim, Jeonghoon Kim, Jeonghwan Kim, Jieyong Kim, Jiho Kim, Jihyuk Kim, Jinsung Kim, Jiseon Kim, Jongho Kim, Jongyoon Kim, Jung-jae Kim, Junho Kim, Junyeob Kim, Junyeong Kim, Kang-min Kim, Kangil Kim, Keonwoo Kim, Kon Woo Kim, Michelle Kim, Minbeom Kim, Minju Kim, Minseok Kim, Minsoo Kim, Minsoo Kim, Minsu Kim, Misoo Kim, Misuk Kim, Seungone Kim, Sun Kim, Sungchul Kim, Sungdong Kim, Taehwan Kim, Takyoung Kim, Yeeun Kim, Yongil Kim, You-na Kim, Youngbin Kim, Youngkyoung Kim, Yujin Kim, Yumin Kim, Zae Myung Kim, Brendan King, Tracy Holloway King, Frederic Kirstein, Chunyu Kit, Diego Klabjan, Christopher Klamm, Jens Kleesiek, Đorđe Klisura, Marius Kloft, Marianne De Heer Kloots, Kate Knill, Dayoon Ko, Jongwoo Ko, Sang-ki Ko, Youngjoong Ko, Goro Kobayashi, Ichiro Kobayashi, Sosuke Kobayashi, Elena Kochkina, Prashant Kodali, Jordan Kodner, Hyukhun Koh, Ryosuke Kohita, Guneet Singh Kohli, Ryuto Koike, Stanley Kok, Jachym Kolar, Olga Kolesnikova, Prateek Kolhar, Alexander Koller, Sai Koneru, Cunliang Kong, Fang Kong, Fanshuang Kong, Jun Kong, Keyi Kong, Linghe Kong, Lingxing Kong, Tao Kong, Myoung-wan Koo, Seonmin Koo, Maarit Koponen, Yeganeh Kordi, Katerina Korre, Roman Koshkin, Katsunori Kotani, Neema Kotonya, Ziyi Kou, Mahnaz Koupae, George Kour, Punit Singh Koura, Geza Kovacs, Adriana Kovashka, Lea Krause, Robert Kraut, Markus Kreuzthaler, Satyapriya Krishna, Adit Krishnan, Nikhil Krishnaswamy, Canasai Kruengkrai, Joni Kruijsbergen, Lun-wei Ku, Max Ku, Jiayi Kuang, Jun Kuang, Marek Kubis, Kei-to Kudo, Sachit Kuhar, Adithya Kulkarni, Atharva Kulkarni, Atharva Kulkarni, Mayank Kulkarni, Anoop Kumar, Bhargava Kumar, Deepak Kumar, Karun Kumar, Nalin Kumar, Nitesh Kumar, Ri-shabh Kumar, Sachin Kumar, Sandeep Kumar, Shanu Kumar, Shashi Kumar, Sonal Kumar, Sujay S Kumar, Vishwajeet Kumar, Vivek Kumar, Yaman Kumar, Tharindu Sandaruwan Kumarage, Sa-dhana Kumaravel, Gitanjali Kumari, Tuhin Kundu, Po-nien Kung, Jenny Kunz, Olli Kuparinen, Yury Kuratov, Murathan Kurfali, Kemal Kurniawan, Guy Kushilevitz, Mucahid Kutlu, Gleb Ku-zmin, Andrey Kuznetsov, Iliia Kuznetsov, Kristian Kuznetsov, Alice Saebom Kwak, Beongwoo Kwak, Sunjun Kweon, Wonbin Kweon, Deuksin Kwon, Jingun Kwon, Se Jung Kwon, Taeyoon Kwon, Daeun Kyung, Abdullatif Köksal

Jorma Tapio Laaksonen, Gorka Labaka, Philippe Laban, Vincent Labatut, Cyril Labbe, Yanis Labrak, Tiziano Labruna, Oier Lopez De Lacalle, Avishek Lahiri, Allison Claire Lahnala, Eason Lai, Hanjiang Lai, Huiyuan Lai, Peichao Lai, Tuan Lai, Viet Dac Lai, Wen Lai, Yilong Lai, Yu-hang Lai, Yuxuan Lai, Kushnareva Laida, Francesco Laiti, Rice Lake, Yair Lakretz, John P. Lalor, Albert Y.s. Lam, Jessica Lam, Hemank Lamba, Vasileios Lampos, Gerasimos Lampouras, Andre Lamurias, Disen Lan, Tian Lan, Zhenzhong Lan, Paul Landes, Hao Lang, Lukas Lange, Loic De Langhe, Karin De Langis, Philippe Langlais, David Langlois, Mateusz Lango, Mirella Lapata, Gabriella Lapesa, Martha Larson, Stefan Larson, Chaak-ming Lau, Jey Han Lau, Leo Laugier,

Ivano Lauriola, Leonard Lausen, Jack Laviolette, John Lawrence, Dung Duy Le, Duy Le, Hang Le, Hoang-quynh Le, Nghia T. Le, Nham Le, Thang Le, Kevin Leach, Joshua Ong Jun Leang, Gwénolé Lecorvé, Beomseok Lee, Bruce W Lee, Byung-jun Lee, Changhun Lee, Changki Lee, Changmin Lee, Cheryl Lee, Dohyeon Lee, Dong-ho Lee, Dongha Lee, Dongwon Lee, Fei-tzin Lee, Gary Lee, Gyeongun Lee, Gyubok Lee, Hodong Lee, Hojae Lee, Hyunju Lee, Jaejin Lee, Jaeseong Lee, Janghwan Lee, Ji-ung Lee, Jimin Lee, Jinu Lee, Jiyoung Lee, John Sie Yuen Lee, Jongwuk Lee, Jung Hyun Lee, Jungseob Lee, Kangwook Lee, Katherine Lee, Kwonjoon Lee, Kyong-ho Lee, Lillian Lee, Mingyu Lee, Minhwa Lee, Nayeon Lee, Patrick Lee, Saehyung Lee, Sang-gil Lee, Sang-goo Lee, Sang-woo Lee, Sanwoo Lee, Seanie Lee, Seonghyeon Lee, Seungeon Lee, Seungyoon Lee, Sunbowen Lee, Sunkyung Lee, Suyeon Lee, Young-jun Lee, Younghun Lee, Youngwon Lee, Yukyung Lee, Fabrice Lefèvre, Joël Legrand, Fangyu Lei, Yibin Lei, Yuanyuan Lei, Zhenyu Lei, Markus Leippold, Christoph Leiter, Maxyn Leitner, Gaël Lejeune, Chak Tou Leong, Bruno Lepri, Pietro Lesci, Lauren Elizabeth Levine, Sergey Levine, Sharon Levy, William D. Lewis, Ang Li, Anqi Li, Bangzheng Li, Bing Li, Bo Li, Bobo Li, Bocheng Li, Bryan Li, Changhao Li, Changmao Li, Changqun Li, Changye Li, Chaozhuo Li, Chengkai Li, Chengming Li, Chengpeng Li, Chengxi Li, Chengzhengxu Li, Chengzu Li, Chong Li, Chuhan Li, Chunyang Li, Chuyuan Li, Cuiping Li, Daifeng Li, Dexun Li, Dingcheng Li, Diya Li, Dong Li, Dongsheng Li, Dongsheng Li, Dongyang Li, Dongyuan Li, Feng-lin Li, Ge Yi Li, Guozheng Li, Hang Li, Hao Li, Hao Li, Hao Li, Haochen Li, Haohang Li, Haonan Li, Haoran Li, Haoran Li, Haoxin Li, Haoyuan Li, Hongyu Li, Huan Li, Huayang Li, Hui Li, Hui Li, Huihan Li, Huiying Li, Jia Li, Jia Li, Jiaang Li, Jiahao Li, Jiahuan Li, Jialu Li, Jiang Li, Jianjun Li, Jianxin Li, Jiaoda Li, Jiaqi Li, Jiayi Li, Jiakuan Li, Jing Li, Jingxuan Li, Jinning Li, Jinyang Li, Jinyuan Li, Jixing Li, Juanzi Li, Juntao Li, Junxian Li, Junyi Li, Junzhuo Li, Kaixin Li, Kan Li, Ke Li, Kun Li, Kun Li, Kun Li, Lei Li, Li Li, Liang Li, Liang Li, Liangze Li, Linjing Li, Linjun Li, Linsen Li, Lishuang Li, Liyao Li, Lujun Li, Meng Li, Mengze Li, Miao Li, Min Li, Ming Li, Mingchen Li, Mingda Li, Mingjie Li, Mingming Li, Mingxiao Li, Mingzhe Li, Minzhi Li, Mukai Li, Na Li, Nian Li, Ning Li, Oliver Li, Peng Li, Ping Li, Qi Li, Qian Li, Qing Li, Qing Li, Qing Li, Qing Li, Qintong Li, Qiuchi Li, Ran Li, Raymond Li, Renhao Li, Rong Li, Rongsheng Li, Rui Li, Rui Li, Rui Li, Rui Li, Rui Li, Ruifan Li, Ruihao Li, Ruiqi Li, Ruixuan Li, Rumei Li, Ruosen Li, Sha Li, Shaobo Li, Shasha Li, Sheng Li, Sheng Li, Shengjie Li, Shichen Li, Shimin Li, Shiyao Li, Shuaiyi Li, Shuangyin Li, Shuo Li, Siyan Li, Sujian Li, Tianhao Li, Tianlin Li, Tianrui Li, Wei Li, Weijie Li, Wenda Li, Wenhao Li, Xia Li, Xiang Li, Xiang Li, Xiangci Li, Xiangtai Li, Xiangyang Li, Xianming Li, Xiao Li, Xiaochuan Li, Xiaomin Li, Xiaonan Li, Xiaopeng Li, Xiaopeng Li, Xiaoxuan Li, Xingxuan Li, Xinjin Li, Xintong Li, Xinze Li, Xinze Li, Xiujun Li, Xiuxing Li, Xuchen Li, Xuelong Li, Ya Li, Yafu Li, Yang Li, Yangning Li, Yanyang Li, Yawei Li, Yi Li, Yifei Li, Ying Li, Yinghao Li, Yingji Li, Yingya Li, Yinqiao Li, Yiqi Li, Yitong Li, Yitong Li, Yiwei Li, Yixia Li, Yixing Li, Yizhi Li, Yongqi Li, Yu Li, Yucheng Li, Yue Li, Yuepei Li, Yufei Li, Yuhan Li, Yunshui Li, Yunxin Li, Yunxuan Li, Yuqing Li, Zaijing Li, Zang Li, Zechen Li, Zehan Li, Zejun Li, Zekun Li, Zexin Li, Zhaonan Li, Zhaowei Li, Zhen Li, Zheng Li, Zheng Li, Zhenhao Li, Zhenwen Li, Zhixu Li, Zhiyu Li, Zhuohang Li, Zhuohao Li, Zichao Li, Zihan Li, Zihao Li, Zihao Li, Ziniu Li, Zizhong Li, Zongxi Li, Zongxia Li, Zongyao Li, Zuchao Li, Defu Lian, Jiachen Lian, Jianxun Lian, Zheng Lian, Chao-chun Liang, Chen Liang, Jenny Liang, Jiafeng Liang, Jiaqing Liang, Jinggui Liang, Junbang Liang, Sheng Liang, Shuailong Liang, Shuo Liang, Tian Liang, Tingting Liang, Wenxin Liang, Xiaobo Liang, Xiwen Liang, Xun Liang, Yingyu Liang, Yitao Liang, Yunlong Liang, Yupu Liang, Zhenwen Liang, Zihan Liang, Baohao Liao, Hao Liao, Huanxuan Liao, Jian Liao, Lizi Liao, Ming Liao, Minpeng Liao, Pei-jun Liao, Ruotong Liao, Weiheng Liao, Wenhui Liao, Xingming Liao, Yong Liao, Yusheng Liao, Zihan Liao, John Licato, Luca Di Liello, Wencke Liermann, Heuseok Lim, Kyungtae Lim, Ser-nam Lim, Woohyung Lim, Woosang Lim, Tomasz Limisiewicz, Peerat Limkonchotiwat, Antonie Lin, Chi-heng Lin, Chu-cheng Lin, Chun Yi Lin, Guan-ting Lin, Haokun Lin, Haonan Lin, Hongzhan Lin, Huawei Lin, I-fan Lin, Inna Wanyin Lin, Jiayu Lin, Jieyu Lin, Lu Lin, Mingbao Lin, Minhua Lin, Pin-jie Lin, Qika Lin, Ruixi Lin, Run Lin, Tony Lin,

Weizhe Lin, Xueyuan Lin, Yen-ting Lin, Yi-cheng Lin, Ying-jia Lin, Youyuan Lin, Yueqian Lin, Yupian Lin, Zeqi Lin, Zhe Lin, Zhen Lin, Zheng Lin, Zhong Lin, Zhouhan Lin, Zihao Lin, Zongyu Lin, Matthias Lindemann, Peter Lindes, Chen Ling, Yuan Ling, Zhen-hua Ling, Zijian Ling, Tal Linzen, Gili Lior, Enrico Liscio, Johann-mattis List, Aiwei Liu, Alisa Liu, An Liu, Anthony Zhe Liu, Ben Liu, Bing Liu, Bingyan Liu, Bo Liu, Bo Liu, Bo Liu, Chao-lin Liu, Chaoqun Liu, Chen Cecilia Liu, Chengyuan Liu, Chuanyi Liu, Chunhua Liu, Cong Liu, Cong Liu, Daizong Liu, Danni Liu, Dongfang Liu, Dugang Liu, Emmy Liu, Enshu Liu, Fangchao Liu, Fei Liu, Fuxiao Liu, Ganga Liu, Gongshen Liu, Guangyi Liu, Guisheng Liu, Han Liu, Hang Liu, Hao Liu, Haotian Liu, Hongfu Liu, Hongyi Liu, Huadai Liu, Hui Liu, Hui Liu, Huijun Liu, Ji Liu, Jiacheng Liu, Jiamou Liu, Jianfeng Liu, Jiang Liu, Jiangming Liu, Jianxing Liu, Jiarui Liu, Jiateng Liu, Jiawei Liu, Jiejyun Liu, Jin Liu, Jin Liu, Jinghui Liu, Jingping Liu, Jingyao Liu, Juhua Liu, Jun Liu, Junlong Liu, Junpeng Liu, Junpeng Liu, Kang Liu, Kunpeng Liu, Lan Liu, Lihui Liu, Linfeng Liu, Ling Liu, Meng Liu, Ming Liu, Ming Liu, Minqian Liu, Nayu Liu, Nelson F. Liu, Ninghao Liu, Peipei Liu, Peiyu Liu, Pengyuan Liu, Ping Liu, Qiang Liu, Qiao Liu, Qijiong Liu, Qin Liu, Rui Liu, Rui Liu, Ruikang Liu, Ruixuan Liu, Ruiyang Liu, Ruoqi Liu, Sheng Liu, Shengjie Liu, Shengzhong Liu, Shicheng Liu, Shichun Liu, Shuai Liu, Shuheng Liu, Shujie Liu, Shukai Liu, Shuliang Liu, Shuqi Liu, Sichen Liu, Sijia Liu, Siyang Liu, Siyi Liu, Siyi Liu, Tengxiao Liu, Tianci Liu, Tianming Liu, Tianyang Liu, Tianyang Liu, Tianyu Liu, Ting Liu, Tingkai Liu, Tong Liu, Wei Liu, Wei Liu, Wei Liu, Weifeng Liu, Weihao Liu, Weize Liu, Wenhe Liu, Xiang Liu, Xiangrui Liu, Xiao Liu, Xiaoming Liu, Xiaoze Liu, Xin Liu, Xin Liu, Xinyue Liu, Xiping Liu, Xixian Liu, Xuebing Liu, Xueqing Liu, Xuye Liu, Yahui Liu, Yan Liu, Yanchen Liu, Yanchi Liu, Yang Liu, Yang Liu, Yang Aron Liu, Yanming Liu, Yantao Liu, Yanting Liu, Ye Liu, Ye Liu, Yi Liu, Yicun Liu, Yidi Liu, Yiding Liu, Yihong Liu, Ying Liu, Yinhong Liu, Yong Liu, Yong Liu, Yonghao Liu, Yongkang Liu, Yongmei Liu, Yu Lu Liu, Yuanchao Liu, Yuanxing Liu, Yuchen Liu, Yue Liu, Yuhan Liu, Yuhan Liu, Yuqing Liu, Yuxiang Liu, Yuxuan Liu, Zeming Liu, Zhengyuan Liu, Zhexiong Liu, Zheyuan Liu, Zhili Liu, Zhining Liu, Zhiwei Liu, Zhiyue Liu, Zhuo Liu, Zihan Liu, Zihang Liu, Zijie Liu, Zijing Liu, Zijun Liu, Zikang Liu, Ziqi Liu, Ziqiang Liu, Zitao Liu, Ziyi Liu, Zizhou Liu, Elena Lloret, Chun Hei Lo, Soda Marem Lo, Tien-hong Lo, Tyler Loakman, Ekaterina Lobacheva, Anamika Lochab, Andrea Loddo, Mengsay Loem, Rahul Lokesh, Erika Lombart, Dingkun Long, Do Xuan Long, Quanyu Long, Teng Long, Wanqiu Long, Xinwei Long, Yunfei Long, Michela Lorandi, Cedric Lothritz, Jian Lou, Qian Lou, Maxime Louis, Natalia Loukachevitch, Justin Lovelace, Bin Lu, Bo-ru Lu, Chaochao Lu, Heng-yang Lu, Hengtong Lu, Jianqiao Lu, Jinghui Lu, Jinliang Lu, Junru Lu, Lei Lu, Ning Lu, Peng Lu, Qingyu Lu, Qixi Lu, Sidi Lu, Wenpeng Lu, Xiaolei Lu, Xiaoxin Lu, Xin Lu, Xinyuan Lu, Yichao Lu, Yida Lu, Yining Lu, Yu Lu, Yu Lu, Yunfan Lu, Yuxing Lu, Yuxuan Lu, Zexin Lu, Zhiyong Lu, Nurul Lubis, Evan Lucas, Jason S Lucas, Bernd Ludwig, Riccardo Lunardi, Chu Fei Luo, Dongsheng Luo, Feng Luo, Fuwen Luo, Ge Luo, Guoqing Luo, Haoran Luo, Haozheng Luo, Hongyin Luo, Jiaming Luo, Jinghao Luo, Jinqi Luo, Junyu Luo, Man Luo, Minnan Luo, Ruilin Luo, Run Luo, Sichun Luo, Tianze Luo, Tongxu Luo, Weidi Luo, Xiang Luo, Xianzhen Luo, Xufang Luo, Yanchen Luo, Yang Luo, Ye Luo, Yingfeng Luo, Yiran Luo, Yong Luo, Yuanchang Luo, Yun Luo, Yuyu Luo, Zheheng Luo, Zhuang Luo, Ziyang Luo, Thang Luong, Mohammad Luqman, Lin Luyang, Bo Lv, Changze Lv, Fengmao Lv, Kai Lv, Olga Lyashevskaya, Chris Lydia, Maria Lymperaiou, Bohan Lyu, Chenyang Lyu, Chunchuan Lyu, Hanjia Lyu, Michael Lyu, Weimin Lyu, Xinglin Lyu, Xinxi Lyu, Yiwei Lyu, Yougang Lyu, Alexander Lyzhov

Bolei Ma, Chen Ma, Chenhao Ma, Fukun Ma, Guangyuan Ma, Haokai Ma, Ji Ma, Jing Ma, Longxuan Ma, Minghua Ma, Nianzu Ma, Pingchuan Ma, Qianli Ma, Qiyao Ma, Rao Ma, Shuo Ma, Weicheng Ma, Weitao Ma, Weizhi Ma, Xiaomeng Ma, Xinbei Ma, Xinge Ma, Xinyin Ma, Xueguang Ma, Youmi Ma, Yunpu Ma, Yuqing Ma, Zexiong Ma, Zhengrui Ma, Zhiyuan Ma, Zihan Ma, Ziyang Ma, Sashank Macha, Jakob Macina, Mounica Maddela, Marco Madeddu, Rahul Madhavan, Manas Madine, Andreas Madsen, Tharindu Madusanka, Joao Magalhaes, Walid Mag-

dy, Pierre Magistry, Bernardo Magnini, Khyati Mahajan, Shweti Mahajan, Kishan Maharaj, Sainik Kumar Mahata, Ridwan Mahub, A V S D S Mahesh, Ayush Maheshwari, Gaurav Maheshwari, Rishabh Maheshwary, Aishwarya Maheswaran, Syed Hasan Amin Mahmood, Mohammad Mahoor, Sijie Mai, Arindam Majee, Navonil Majumder, Márton Makrai, Emanuele La Malfa, Manuj Malik, Advaith Malladi, Christopher Malon, Mamta Mamta, Zhibo Man, Sunny Manchanda, Avshalom Manevich, Varun Manjunatha, Christopher D. Manning, Ioana Manolescu, Jonibek Mansurov, Hangyu Mao, Kelong Mao, Qianren Mao, Rui Mao, Xiaoxi Mao, Yu Mao, Zhiming Mao, Zhuoyuan Mao, Guillermo Marco, Marion Di Marco, Alda Mari, Vukosi N. Marivate, Katja Markert, Iliia Markov, Shaul Markovitch, Elan Sopher Markowitz, Mounika Marreddy, Jason Mars, Federico Martelli, Anna Martin, James Martin, Lara J. Martin, Paul Martin, Matej Martinc, Richard Diehl Martinez, Eugenio Martinez-camara, Bruno Martins, Antonio David Ponce Martínez, Reem Masoud, Shahed Masoudian, Ahmed Masry, Sarah Masud, Ryo Masumura, Ved Mathai, Sandeep Mathias, Leena Mathur, Puneet Mathur, Yuji Matsumoto, Florian Matthes, Yevgen Matusovich, Evgeny Matusov, Kaushal Kumar Maurya, Costas Mavromatis, Diana Maynard, Philipp Mayr, Alessandro Mazzei, Sabrina Mccallum, R. Thomas Mccoy, John Philip Mccrae, Daniel McDuff, Euan McGill, Nick Mckenna, Kathleen Mckeown, Nicholas Meade, Sourav Medya, Jannat Ara Meem, Michiel Van Der Meer, Alexander Mehler, Houman Mehrafarin, Kartik Mehta, Maitrey Mehta, Sanket Vaibhav Mehta, Jingbiao Mei, Lingrui Mei, Sen Mei, Stephen Meisenbacher, Abdellah El Mekki, Jesus Lovón Melgarejo, Elena Alvarez Mellado, Gerard Melo, Velayuthan Menan, Julia Mendelsohn, Ethan Adrian Mendes, Marcelo Mendoza, Arul Menezes, João Menezes, Chuan Meng, Lingyuan Meng, Xuying Meng, Yuanliang Meng, Stefano Menini, Rakesh R Menon, Wolfgang Menzel, Fabio Mercorio, Paola Merlo, Mike A Merrill, William Merrill, Rin Metcalf, Guillaume Metzler, Francois Meyer, Selina Meyer, Stefano Mezza, Md Messal Monem Miah, Yibo Miao, Yisong Miao, Alessio Miaschi, Kurt Micallef, James Michaelov, Zulfat Miftahutdinov, Rada Mihalcea, Tsvetomila Mihaylova, Vladislav Mikhailov, Jeremiah Milbauer, Filip Miletic, Aleksandra Miletic, Timothy A. Miller, David Mimno, Dehai Min, Do June Min, Qingkai Min, Mario Mina, Pasquale Minervini, Koji Mineshima, Tianshi Ming, Yifei Ming, Benjamin Minixhofer, Santiago Miret, Seyedabolghasem Mirroshandel, Aziza Mirsaidova, Mehrnoosh Mirtaheri, Abhijit Mishra, Aditya M. Mishra, Kshitij Mishra, Pruthwik Mishra, Amita Misra, Masato Mita, Cassie S. Mitchell, Avijit Mitra, Dr. Satanik Mitra, Ashish Mittal, Happy Mittal, Makoto Miwa, Taiki Miyanishi, Yusuke Miyao, Taro Miyazaki, Wenjie Mo, Ali Modarressi, Sandip Modha, Arkadiusz Modzelewski, Lucas Moeller, Abdelrahman Mohamed, Amr Mohamed, Youssef Sherif Mansour Mohamed, Manas Ranjan Mohanty, Jisoo Mok, Laxmikant Kishor Mokadam, Francesco Maria Molfese, Nicholas Monath, Erfan Moosavi Monazzah, Ishani Mondal, Mainack Mondal, Philipp Mondorf, Helena Silva Moniz, Adrian Pastor Lopez Monroy, Syed Mostofa Monsur, Stefano Montanelli, Arturo Montejo-ráez, Kyle Montgomery, Hyeonseok Moon, Suhong Moon, Ray Mooney, Guy Mor, Raha Moraffah, Gianmarco De Francisci Morales, Tetsuro Morimura, Luca Moroni, Jacob Daniel Morrison, Isaac Thani Moses, Viktor Moskvoretskii, Pablo Mosteiro, Josiane Mothe, Yutao Mou, Ali Mousavi, Seyed Mahed Mousavi, Emily Mower, Li Moxin, Maximilian Mozes, Honglin Mu, Yida Mu, Yongyu Mu, Hamdy Mubarak, Shamsuddeen Hassan Muhammad, Rohit Mujumdar, Animesh Mukherjee, Anjishnu Mukherjee, Sagnik Mukherjee, Jonathan Mukiibi, Matthew Mulholland, Carlos Mullov, Hyeongdon Mun, Rajesh Kumar Mundotiya, Vishvak Saivenkat Murahari, Koji Murakami, Soichiro Murakami, Masayasu Muraoka, Yugo Murawaki, Andrei Ioan Muresanu, Keerthiram Murugesan, Pradeep Kumar Murukannaiah, John Murzaku, Elio Musacchio, Tomáš Musil, Skatje Myers, Agnieszka Mykowiecka

Byeonghu Na, Cheolwon Na, Hongbin Na, Seung-hoon Na, Nona Naderi, Mohamed Nadif, Hubert Naets, Arijit Nag, Sayan Nag, Ryo Nagata, Md Mahadi Hasan Nahid, Aakanksha Naik, Atharva Naik, Inderjeet Jayakumar Nair, Rahul Nair, Tetsuji Nakagawa, Itay Nakash, Yuta Nakashima, Ndapa Nakashole, Makoto Nakatsuji, Eric Nalnick, Akshay Nambi, Guoshun Nan, Federico Nanni, Tarek Naous, Diane Napolitano, Sharan Narasimhan, Shashi Narayan, Daniele Nardi, Ne-

ha Narwal, Jamal Nasir, Anandhavelu Natarajan, Abhijnan Nath, Shah Nawaz, Shravan Nayak, Tapas Nayak, Hamada Nayel, Renato Negrinho, Anil Kumar Nelakanti, Sree Harsha Nelaturu, Poli Nemkova, Vera Neplenbroek, Graham Neubig, Günter Neumann, Mariana Neves, Benjamin Newman, Jun-ping Ng, Raymond T. Ng, See-kiong Ng, Youyang Ng, Henry Nghiem, Nghia Trung Ngo, Anh Nguyen, Bang Nguyen, Chien Van Nguyen, Duc-vu Nguyen, Duy Nguyen, Giang Nguyen, Huy Nguyen, Khanh Xuan Nguyen, Kiem-hieu Nguyen, Long Hb Nguyen, Luan Thanh Nguyen, Minh Van Nguyen, Ngan Luu-thuy Nguyen, Thao Nguyen, Thong Nguyen, Tin Nguyen, Tin Trung Nguyen, Truc-vien T. Nguyen, Tuc Van Nguyen, Tung Nguyen, Van Bach Nguyen, Vincent Nguyen, Jingwei Ni, Junrui Ni, Xuanfan Ni, Yuansheng Ni, Zhaoheng Ni, Dong Nie, Ercong Nie, Jian-yun Nie, Pengyu Nie, Ping Nie, Yuxiang Nie, Zhijie Nie, Elizabeth Kaye Nielsen, Malvina Nikandrou, Anna Nikiforovskaya, Irina Nikishina, Christina Niklaus, Sergey I. Nikolenko, Jinzhong Ning, Kun-peng Ning, Ruoxi Ning, Noriki Nishida, Sergiu Nisioi, Malvina Nissim, Vikram Nitin, Chaoyue Niu, Guanglin Niu, Hao Niu, Jianwei Niu, Jingcheng Niu, Luyao Niu, Tianhao Niu, Tianyi Niu, Xueyan Niu, Joakim Nivre, Ayana Niwa, Bill Noble, Rodrigo Frassetto Nogueira, Keonwoong Noh, Adam Nohejl, Tadashi Nomoto, Rik Van Noord, Kai North, Marzia Nouri, Damien Nouvel, Scott Novotney, Sarana Nutanong, Ifeoma Nwogu

Alexander O’connor, Nicolás Benjamín Ocampo, Daniela Occhipinti, Brendan Oconnor, Yusuke Oda, Stephan Oepen, Owen Oertell, Alex Oesterling, Bahadorreza Ofoghi, Perez Atieno Ogayo, Michael Ogezi, Byung-doh Oh, Hanseok Oh, Hongseok Oh, Jaehoon Oh, Jungseul Ok, Tatsushi Oka, Yui Oka, Shu Okabe, Tsuyoshi Okita, Eda Okur, Tobi Olatunji, Hugo Gonçalo Oliveira, Adjali Omar, Ethel Chua Joy Ong, Subba Reddy Oota, Andreas Opedal, Juri Opitz, Alina Oprea, Maite Oronoz, John Ortega, Naoki Otani, Yulia Otmakhova, Jiao Ou, Jiefu Ou, Yixin Ou, Zhijian Ou, Siqi Ouyang, Siru Ouyang, Yang Ouyang, Zhongyu Ouyang, Risako Owan, Momose Oyama, Masafumi Oyamada, Savas Ozkan, Arzucan Özgür

Inkit Padhi, Trilok Padhi, Tejomay Kishor Padole, Aline Paes, Artidoro Pagnoni, Vardaan Pahuja, Partha Pakray, Christopher J Pal, Kuntal Kumar Pal, Proyag Pal, Chester Palen-michel, Matteo Palmonari, Flavio Di Palo, John Palowitch, Shramay Palta, Bo Pan, Deng Pan, Huitong Pan, Jane Pan, Junting Pan, Rong Pan, Rui Pan, Ruihao Pan, Shirui Pan, Xiang Pan, Xiao Pan, Yanzhou Pan, Youcheng Pan, Zhenyu Pan, Artemis Panagopoulou, Ruchi Panchanadikar, Alexander Panchenko, Pranoy Panda, Subhadarshi Panda, Nilay Pande, Punya Syon Pandey, Saurabh Kumar Pandey, Shrey Pandit, Chenxi Pang, Tianyu Pang, Sandeep Kumar Pani, Abhishek Panigrahi, Panlei, Aleksandr Panov, Evangelos E. Papalexakis, Sean Papay, Paolo Papotti, Duccio Pappadopulo, Ashraful Islam Paran, Natalie Parde, Amit Parekh, Tanmay Parekh, Angelina Parfenova, Priyank Parikh, Chaehun Park, Chan Young Park, Cheonbok Park, Cheoneum Park, Choonghyun Park, Eunhwan Park, Eunhyeok Park, Gyeong-moon Park, Hogun Park, Jinyoung Park, Jun-hyung Park, Jungsoo Park, Jungyeul Park, Seongmin Park, Shinwoo Park, Sungwon Park, Mihir Parmar, Marinela Parovic, Pranjali Ajay Parse, Md Rizwan Parvez, Ankita Pasad, Nivranshu Pasricha, Valeria Pastorino, Ajay Patel, Hitesh Laxmichand Patel, Maitreya Patel, Nilay Patel, Sarang Rajendra Patil, Vaidehi Patil, Jasabanta Patro, Priyaranjan Pattnayak, Parth Patwa, Manasi Patwardhan, Indraneil Arun Paul, Mithun Paul, Shounak Paul, Amalie Brogaard Pauli, John Pavlopoulos, Vera Pavlova, Sachin Pawar, Siddhesh Pawar, Ali Payani, Patomporn Payoungkhamdee, Tzuf Paz-argaman, Branislav Pecher, Pavel Pecina, Changhua Pei, Qizhi Pei, Georgios Peikos, Catherine Pelachaud, Lotem Peled-cohen, Kellin Pelrine, Abhirama Subramanyam Penamakuri, Baolin Peng, Bo Peng, Dezhi Peng, Dezhong Peng, Jiawei Peng, Jing Peng, Keqin Peng, Letian Peng, Lim Jia Peng, Miao Peng, Min Peng, Qiming Peng, Qiwei Peng, Qiyao Peng, Shaohui Peng, Tianhao Peng, Tianshuo Peng, Wei Peng, Wen-chih Peng, Xi Peng, Xiangyu Peng, Yifan Peng, Yifan Peng, Gerald Penn, Valerio Pepe, Joseph Joshua Peper, Gal Peretz, Jesus Javier Calleja Perez, Laura Perez-beltrachini, Gabriele Pergola, Nathanaël Perraudin, Stefano Perrella, Ali Pesaranghader, Stanislav Peshterliev, Denis Peskov, Jan-thorsten Peter, Ben Peters, Pavel Petrushkov, Pouya Pezeshkpour, Hanspeter Pfister,

Quang Hieu Pham, Fred Philippy, Renjie Pi, Yinhuo Piao, Massimo Piccardi, Janet B. Pierrehumbert, Michał Pietruszka, Reshmi Gopalakrishna Pillai, Mathis Pink, Ramakrishna Pinninti, Bartosz Piotrowski, Andrew Piper, Hugo Pitorro, Priya Pitre, Lidia Pivovarov, Paul Piwek, Barbara Plank, Moritz Plenz, Max Ploner, Soham Poddar, Christian Poelitz, Thierry Poibeau, Rhitabrat Pokharel, Sandeep Polisetty, Ramesh Poluru, Josiah Poon, Adrian Popescu, Nicholas Popovic, Ian Porada, Beatrice Portelli, François Portet, Matt Post, Martin Potthast, Amrit Poudel, Mohammedreza Pourreza, Manuel Prada-corrall, Ritu Prajapati, Aniket Pramanick, A Pranav, Jakob Prange, Shivika Prasanna, Anurag Pratik, Gabriele Prato, Ian Pratt-hartmann, Frederic Precioso, Damith Premasiri, Valentina Presutti, Sarah Masud Preum, Priyanshu Priya, Lorenzo Proietti, Victor Prokhorov, Ruili Pu, Shi Pu, Yewen Pu, Yifan Pu, Giovanni Puccetti, Giulia Pucci, Ratish Puduppully, Saurabh Pujar, Rajkumar Pujari, Sukannya Purkayastha, Alberto Purpura, Matthew Purver, Ayu Purwarianti, Rifki Afina Putri, Meghana Puvvadi, Juan Antonio Pérez-ortiz

Ayesha Qamar, Jianzhong Qi, Jirui Qi, Peng Qi, Weijian Qi, Xuan Qi, Yunjia Qi, Dong Qian, Hongjin Qian, Kun Qian, Xueming Qian, Yining Qian, Yusu Qian, Zhaozhi Qian, Zhenxing Qian, Zhong Qian, Jipeng Qiang, Minjie Qiang, Zewen Qiang, Shanbao Qiao, Shuofei Qiao, Bowen Qin, Chengwei Qin, Haotong Qin, Jianbin Qin, Jinghui Qin, Tian Qin, Wei Qin, Yuehan Qin, Zhanyue Qin, Zhenyue Qin, Peijun Qing, Huachuan Qiu, Lin Qiu, Linlu Qiu, Mengyang Qiu, Xihe Qiu, Xin Ying Qiu, Yifu Qiu, Zhaopeng Qiu, Zihan Qiu, Muhammad Reza Qorib, Jiaying Qu, Xiaoye Qu, Xingwei Qu, Yanru Qu, Yincen Qu, Yuzhong Qu, Xin Quan, Ariadna Quattoni, Haoran Que, Solen Quiniou, Mohammed Rameez Qureshi, Rizwan Qureshi

Gowtham R, Narendhiran R, Reihaneh Rabbany, Hamid R. Rabiee, Ella Rabinovich, Karthik Radhakrishnan, Bhaktipriya Radharapu, Daniele P. Radicioni, Taki Hasan Rafi, Davood Rafiei, Alessandro Raganato, Luca Ragazzi, Vipul Raheja, Vahid Rahimzadeh, Mizanur Rahman, Sha-fin Rahman, Hossein A. Rahmani, Nishat Raihan, Tina Raissi, Lisa Raithel, Chahat Raj, Sahar Rajabi, Sara Rajae, Kanagasabai Rajaraman, Rishanth Rajendhran, Shahab Raji, Sarah Michele Rajtmajer, Nathanaël Carraz Rakotonirina, Anil Ramakrishna, Bhaskar Ramasubramanian, Giulia Rambelli, Owen Rambow, Aida Ramezani, Sahana Ramnath, Sanjana Ramprasad, Shivaen Ramshetty, Pritika Ramu, Guillem Ramírez, Shihao Ran, Yide Ran, Ashish Rana, Federico Ranaldi, Anku Rani, Dongning Rao, Gengyu Rao, Jun Rao, Rajath Rao, Zhi Rao, Ahmad Rashid, Vipul Kumar Rathore, Shivam Raval, Federico Ravenda, Sahithya Ravi, Ajay Madhavan Ravichandran, Renjith P Ravindran, Yogesh S Rawat, Sourjyadip Ray, Simon Razniewski, Olesya Razuvayevskaya, Imran Razzak, Varshini Reddy, Chris Reed, Michaela Regneri, Abudurexiti Reheman, Nils Reiter, Sravan Reddy Rekula, Houxing Ren, Jie Ren, Jie Ren, Kan Ren, Kui Ren, Qingyang Ren, Runtao Ren, Shuo Ren, Weijieying Ren, Xuanfei Ren, Xubin Ren, Yubing Ren, Yujie Ren, Yarik Menchaca Resendiz, Achim Rettinger, Anja Reusch, Artem Revenko, Mohammadhossein Rezaei, Rezvaneh Rezapour, Kathryn Ricci, Kyle Richardson, Stephen D. Richardson, Korbinian Riedhammer, Maarten De Rijke, Laura Rimell, Md Imbesat Hassan Rizvi, Jesse Roberts, Gabriel Roccabruna, Gil Rocha, Julian Martin Rodemann, Christophe Rodrigues, Pedro Rodriguez, Melissa Roemmele, Lina Maria Rojas-barahona, Nathan Roll, Wenge Rong, Jonathan Rose, Mohammad Rostami, Aiala Rosá, Guy Rotman, Kay Rottmann, Dmitri Roussinov, Bryan R. Routledge, Amartya Roy, Amit Roy, Aniruddha Roy, Shamik Roy, Soumyadeep Roy, Sudipta Singha Roy, Nikolai Rozanov, Chong Ruan, Jiacheng Ruan, Jie Ruan, Jingqing Ruan, Tong Ruan, Terry Ruas, Patrick Ruch, William Rudman, Daniel Ruffinelli, Federico Ruggeri, Ruibochen, Mukund Rungta, Hu Ruofan, Jonathan Rusert, Stefan Ruseti, Jenna Russell, Daniel Russo, Irene Russo, Derek Ruths, Elena Sofia Ruzzetti, Pavel Rychlý, Sangwon Ryu, Rafal Rzepka, Susanna Rucker

Ramaneswaran S, Jon Saad-falcon, Ashish Sabharwal, Ahmed Sabir, Sahand Sabour, Rachneet Singh Sachdeva, Kaushik Ram Sadagopan, Abdelrahman Boda Sadallah, Mobashir Sadat, Amir Saedi, Takaaki Saeki, Bardia Safaei, Mustafa Safdari, Benoît Sagot, Amit Kumar Sah, Avirup

Saha, Punyajoy Saha, Rohan Saha, Sougata Saha, Soumya Saha, Sriparna Saha, Tulika Saha, Sattvik Sahai, Vaibhav Sahai, Magnus Sahlgren, Nihar Ranjan Sahoo, Oscar Sainz, Yusuke Sakai, Andrey Sakhovskiy, Alireza Salemi, Jonne Saleva, Vishal Vivek Saley, Suchir Salhan, Flora Salim, Avneesh Saluja, Tanja Samardzic, Ranbir Singh Sanasam, Javier Iranzo Sanchez, Jivne-sh Sandhan, Raghav Sandhu, Erik F. Tjong Kim Sang, Samin Mahdizadeh Sani, Maziar Sanjabi, Ashwin Sankar, Chinnadhurai Sankar, Aruna Sankaranarayanan, Scott Sanner, Leonid Sanochkin, Selene Baez Santamaria, Sashank Santhanam, Andrea Santilli, Bishal Santra, Debarshi Kumar Sanyal, Irina Sapparina, Xabier Saralegi, Vasanth Sarathy, Gabriel Herbert Sarch, Arindam Sarkar, Rupak Sarkar, Souvika Sarkar, Sujoy Sarkar, Somdeb Sarkhel, Shota Sasaki, David Sasu, Ranjan Satapathy, Prasanna Sattigeri, Agata Savary, Andrey Savchenko, Beatrice Savoldi, Guergana K Savova, Manolis Savva, Apoorv Umang Saxena, Rohit Saxena, Michael Saxon, Salim Sazzed, Daniel Scalena, Carolina Scarton, Elizabeth J. Schaefer, Markus Schedl, Tatjana Scheffler, Marten Van Schijndel, David Schlangen, Helmut Schmid, Fabian David Schmidt, Patrícia Schmidtová, Vera Schmitt, Florian Schneider, Gerold Schneider, Julian Moreno Schneider, Nathan Schneider, Patrick Schramowski, Christopher Schröder, Dan Schumacher, Elliot Schumacher, Jakob Schuster, Stefan Schweter, Ulrich Schäfer, Simeon Schütz, Alessandro Scirè, Amit Seker, Bilgehan Sel, David Semedo, Sina Semnani, Jaydeep Sen, Meghdut Sengupta, Jaehyung Seo, Jeongyeon Seo, Sofia Serrano, Sergio Servantez, Agrima Seth, Agam Shah, Alay Dilipbhai Shah, Mubarak Shah, Shalin Shah, Syed Afaq Ali Shah, Haz Sameen Shahgir, Rifat Shahriyar, Rana Shahroz, Chantal Shaib, Sagi Shaier, Omar Shaikh, Cory Shain, Gregory Shakhnarovich, Valerie Shalin, Hengchao Shang, Lanyu Shang, Ziyao Shangguan, Chen Shani, Abhilash Shankarampeta, Chenze Shao, Jing Shao, Wei Shao, Wenqi Shao, Yanjun Shao, Yihua Shao, Yijia Shao, Yingxia Shao, Yunfan Shao, Nada Sharaf, Omar Sharif, Abhishek Sharma, Divya V Sharma, Kartik Sharma, Mohit Sharma, Prashant Kumar Sharma, Raksha Sharma, Roshan Sharma, Behzad Shayegh, Artem Shelmanov, Bowen Shen, Huanming Shen, Huawei Shen, Jiabin Shen, Jiayi Shen, Jocelyn J Shen, Li Shen, Linlin Shen, Ming Shen, Siqi Shen, Tao Shen, Tianhao Shen, Xiangqing Shen, Xingyu Shen, Xinyue Shen, Yatian Shen, Ying Shen, Ying Shen, Zejiang Shen, Zhan Shen, Emily Sheng, Jiawei Sheng, Qiang Sheng, Quan Z. Sheng, Zhang Jia Sheng, Shengbinyue, Ivaxi Sheth, Anudeex Shetty, Pooja Shetty, Prakriti Shankar Shetty, Chufan Shi, Dachuan Shi, Daqian Shi, Haochen Shi, Haoxiang Shi, Jiatong Shi, Jiatong Shi, Jie Shi, Kejian Shi, Lida Shi, Qi Shi, Taiwei Shi, Tianyuan Shi, Weijia Shi, Wenqi Shi, Xiao Shi, Xiaoming Shi, Xiaoshuang Shi, Yaorui Shi, Yaya Shi, Yucheng Shi, Zesheng Shi, Zhan Shi, Zhengxiang Shi, Zijing Shi, Tomohide Shibata, Min-han Shih, Li Shilong, Hidetoshi Shimodaira, Erica Kido Shimomoto, Anastasia Shimorina, Haebin Shin, Hyopil Shin, Hyunseo Shin, Jisu Shin, Seungwon Shin, Kazutoshi Shinoda, Kiyooki Shirai, Yow-ting Shiue, Pannaga Shivaswamy, Michal M Shlapentokh-rothman, Vitaly Shmatikov, Michal Shmueli-scheuer, Lidan Shou, Kumar Shridhar, Kai Shu, Tianmin Shu, Yao Shu, Yiheng Shu, Hong-han Shuai, Wang Shuai, Kai Shuang, Liu Shudong, Ruihao Shui, Zeren Shui, Shreya Shukla, Utsav Shukla, Kashun Shum, Ammon C Shurtz, Shuwen, Chenglei Si, Shijing Si, Shuzheng Si, Anthony Sicilia, Lucia Siciliani, A.b. Siddique, Grigori Sidorov, Noah Yamamoto Siegel, Zhou Sijia, Sandipan Sikdar, Damien Sileo, Eliezer De Souza Da Silva, Nisansa De Silva, Purificação Silvano, Fabrizio Silvestri, Patrick Simianer, Antoine Simoulin, Elena Simperl, Amrita Singh, Apoorva Singh, Avyav Kumar Singh, Ayushman Singh, Bhavyajeet Singh, Chandan Singh, Janvijay Singh, Joykirat Singh, Mantek Singh, Mukul Singh, Pranaydeep Singh, Prateek Singh, Rituraj Singh, Sonali Singh, Mohit Singhal, Iustin Sirbu, Sonish Sivarajkumar, Johan Sjons, Steven Skiena, Gabriella Skitalinskaya, Paweł Marek Skórzewski, Charese Smiley, David A. Smith, Răzvan-alexandru Smădu, Vésteinn Snæbjarnarson, Daria Soboleva, Jy-yong Sohn, Felix Soldner, Taiga Someya, Pia Sommerauer, Bokyoung Son, Junyoung Son, Seungwoo Son, Chan Hee Song, Chunyao Song, Dandan Song, Dawei Song, Dawn Song, Feifan Song, Guojie Song, Haiyue Song, Hoyun Song, Hwanjun Song, Hyun-je Song, Jongyoon Song, Kaisong Song, Kaitao Song, Kehui Song, Kyungwoo Song, Lingran Song, Lingyun Song, Linqi Song, Linxin Song, Mingyang Song, Mingyang Song, Peiyang Song, Ran Song, Rui Song, Ruihua Song, Xiaoshuai Song, Xin Song,

Yingjin Song, Yiping Song, Yuanfeng Song, Yueqi Song, Yun-zhu Song, Yurun Song, Zhao Song, Zhenxi Song, Zifan Song, Zirui Song, Atin Sood, Lay-ki Soon, Alessandro Sordoni, Jeffrey S. Sorensen, Taylor Sorensen, Ionut Sorodoc, Alexey Sorokin, Tiberiu Sosea, Anna Sotnikova, Alvaro Soto, Rafael Alberto Rivera Soto, Sajad Sotudeh, Zhivar Sourati, Dilara Soylu, Alexander Spangher, Lucia Specia, Alessandro Sperduti, Michal Spiegel, Andreas Spitz, Maximilian Spliethöver, Zayne Rea Sprague, Ivan Srba, Makesh Narsimhan Sreedhar, Neha Srikanth, Mukund Srinath, Dananjay Srinivas, Jayanth Srinivasa, Tejas Srinivasan, Vijay Srinivasan, Ankit Kumar Srivastava, Aseem Srivastava, Dhruv Srivastava, Mani Srivastava, Saurabh Srivastava, Kv Aditya Srivatsa, Joe Stacey, Ieva Staliunaite, Dominik Stambach, Giorgos B. Stamou, Karolina Stanczak, Katherine Stasaski, Mark Steedman, Julius Steen, Nicolas Stefanovitch, Sebastian Steindl, Andreas Stephan, Igor Sterner, Hannah Sterz, Ian Stewart, Asa Cooper Stickland, Shane Storks, Svetlana Stoyanchev, Ivelina Stoyanova, Sebastian Stüker, Guixin Su, Hang Su, Hao Su, Hongjin Su, Hung-ting Su, Puzhen Su, Sen Su, Weihang Su, Xiangdong Su, Xin Su, Yi Su, Yongye Su, Zhaochen Su, Zhe Su, Zhenpeng Su, D. Subhalingam, Malliga Subramanian, Prasanna Lakkur Subramanyam, Peiqi Sui, Xuhui Sui, Yuan Sui, Bin Sun, Changzhi Sun, Chengjie Sun, Chenkai Sun, Dachun Sun, Guangzhi Sun, Guangzhong Sun, Hanshi Sun, Hao Sun, Haotian Sun, Huaman Sun, Jingchen Sun, Jingyi Sun, Jun Sun, Kai Sun, Kaiser Sun, Ke Sun, Kexuan Sun, Kun Sun, Ling Sun, Linzhuang Sun, Manxi Sun, Mengshu Sun, Penglei Sun, Qiao Sun, Qiushi Sun, Renliang Sun, Rui Sun, Ruoyu Sun, Shiliang Sun, Si Sun, Siqi Sun, Tiening Sun, Wei Sun, Xiaobing Sun, Xiaoshuai Sun, Yang Sun, Yuan Sun, Yuanyuan Sun, Yuchong Sun, Yuqing Sun, Yuxi Sun, Zechen Sun, Zengkui Sun, Zequn Sun, Zexu Sun, Zhaoyue Sun, Zhewei Sun, Zhongxiang Sun, Zhouhao Sun, Anirudh S Sundar, Jayaprakash Sundararaj, Megha Sundriyal, Hakyung Sung, Mujeen Sung, Yoo Yeon Sung, Sunhao, Monica Sunkara, Marek Suppa, Siddharth Suresh, Tarun Suresh, Manan Suri, Aditya Kaushik Surikuchi, Simon Suster, Lintang Sutawika, Yoshimi Suzuki, Pedro Javier Ortiz Suárez, Anej Svete, Colin Swaelens, Sarathkrishna Swaminathan, Swabha Swayamdipta, Zsolt Szanto, Andreas Säuberli, Mihailo Škorić

Jaesung Tae, Shabnam Tafreshi, Andrea Tagarelli, Marzieh S. Tahaei, Francesco Taioli, Yu Takagi, Takehiro Takayanagi, Koichi Takeuchi, Tetsuya Takiguchi, Ece Takmaz, Ekaterina Taktasheva, Zhi Rui Tam, Ales Tamchyna, Chao-hong Tan, Cheng Tan, Cheston Tan, Chuanyuan Tan, Fei Tan, Hanzhuo Tan, Haochen Tan, Hongye Tan, Qiaoyu Tan, Shaomu Tan, Shiwei Tan, Weiting Tan, Xiaoyu Tan, Xin Tan, Yusong Tan, Zeqi Tan, Zhen Tan, Zhentao Tan, Shohei Tanaka, Bo Tang, Buzhou Tang, Chen Tang, Chenming Tang, Haoran Tang, Jing Tang, Jing Tang, Kevin Tang, Kewen Tang, Liyan Tang, Lumingyuan Tang, Nan Tang, Ningzhi Tang, Raphael Tang, Ruixiang Tang, Wei Tang, Xiaohang Tang, Xuemei Tang, Xunzhu Tang, Yihong Tang, Yixuan Tang, Yixuan Tang, Yu-chien Tang, Yubao Tang, Yun Tang, Zecheng Tang, Zheng Tang, Zhenheng Tang, Zhiwen Tang, Zichen Tang, Zilu Tang, Kumar Tanmay, Chris Tanner, Farhan Tanvir, Eshaan Tanwar, Chaofan Tao, Chongyang Tao, Guanhong Tao, Jianhua Tao, Mingxu Tao, Wenbiao Tao, Xijia Tao, Zhang Wen Tao, Zhengwei Tao, Zhiqiang Tao, Zhuo Tao, Maksym Taranukhin, Stephen Eugene Taylor, Diane Tchuindjo, Matthieu Tehenan, Selma Tekir, Atharva Tendle, Damien Teney, Yan Teng, Simone Teufel, Katherine Thai, Patricia Thaine, Megh Thakkar, Surendrabikram Thapa, Ashish V Thapliyal, Ramya Keerthy Thatikonda, Anirudh Thatipelli, Barry-john Theobald, Sarubi Thillainathan, Craig Alexander Thomson, James Thorne, Veronika Thost, Poojitha Thota, David Thulke, Bozhong Tian, Chang Tian, Jiahao Tian, Jidong Tian, Jinchuan Tian, Junfeng Tian, Qi Tian, Shiyu Tian, Shulin Tian, Yapeng Tian, Yu Tian, Yuan Tian, Yuchen Tian, Roelien C Timmer, Abhisek Tiwari, Lindia Tjuatja, Antonela Tommasel, Hanghang Tong, Hao Tong, Meihan Tong, Xiaoyu Tong, Jonathan Tonglet, Atnafu Lambebo Tonja, Manuel Tonneau, Marwan Torki, Rishit Toteja, Samia Touileb, Julien Tourille, Itamar Trainin, Hieu Tran, Khiem Vinh Tran, Khoi-nguyen Tran, Quyen Tran, Son Quoc Tran, Marco Del Tredici, Volker Tresp, Ilias Triantafyllopoulos, Jan Trienes, Remi Van Trijp, Rocco Tripodi, Ashutosh Trivedi, Harsh Trivedi, Shubhendu Trivedi, Yun Da Tsai, Yu Tsao, Dimitrios Tsarapatsanis, Bo-hsiang Tseng, Shao-yen Tseng, Talia Tseriotou,

Ioannis Tsiamas, Eleftheria Tsipidi, Hayato Tsukagoshi, Akim Tsvigun, Lifu Tu, Mei Tu, Quan Tu, Rong-cheng Tu, Shangqing Tu, Elena Tutubalina, Aman Tyagi, Utkarsh Tyagi

Takuma Udagawa, Akhil Udathu, Can Udomcharoenchaikit, Ryo Ueda, Gökçe Uludoğan, Alexandra Uma, Inigo Jauregi Unanue, Rheeya Uppaal, Ashok Urlana, Nicolas Usunier, Benno Uthaya-sooriyar, Takehito Utsuro, Ozlem Uzuner

Robert Vacareanu, Saujas Vaduguru, Teemu Vahtola, Ashwini Vaidya, Sowmya Vajjala, Nidhi Vakil, Francisco Valentini, Marco Valentino, Andric Valdez Valenzuela, Estrella Vallecillo-rodríguez, Sai P Vallurupalli, Hoang Van, Linh Ngo Van, Vincent Vandeghinste, Natalia Vanetik, Vasudha Varadarajan, Siddharth Varia, Shiva Maleki Varnosfaderani, Deeksha Varshney, Prasoon Varshney, Rajashekar Vasantha, Aniket Vashishtha, Oleg Vasilyev, Justin Vasselli, Pavlo Vasylenko, Artem Vazhentsev, Eva Maria Vecchi, Nikhita Vedula, Hariram Veeramani, Dan John Velasco, Erik Vellidal, Raviteja Vemulapalli, Suresh Venkatasubramanian, Ganesh Venkatesh, Kavana Venkatesh, Praveen Venkateswaran, Apurv Verma, Arun Verma, Neha Verma, Rajeev Verma, Yash Verma, Keith Vertanen, Gianluca Vico, Renata Vieira, Lovekesh Vig, Anvesh Rao Vijjini, David Vilares, Dan Vilenchik, Danae Sanchez Villegas, Vibhav Vineet, Krishnapriya Vishnubhotla, Amar Viswanathan, Vijay Viswanathan, Aditya Visweswaran, Nikolas Vitsakis, Jean-noël Vittaut, Andreas Vlachos, Kiran Vodrahalli, Ekaterina Voloshina, Pavlos Vougiouklis, Duy-khanh Vu, Huy-the Vu, Trang Vu, Renato Vukovic, Ugurcan Vurgun, Ivan Vykopal, Ekaterina Vylomova, Horacio Jarquín Vásquez, Juan Vásquez, Raúl Vázquez

Takashi Wada, Sahil Wadhwa, Somin Wadhwa, Claudia Wagner, Eitan Wagner, Joachim Wagner, Abdul Waheed, Jan Philip Wahle, Eric Peter Wairagala, Hiromi Wakaki, Sabine Schulte Im Walde, Andreas Waldis, Jonas Wallat, Changxuan Wan, Chenwei Wan, Fanqi Wan, Hai Wan, Herun Wan, Huaiyu Wan, Kaiyang Wan, Qizhi Wan, Stephen Wan, Yao Wan, Yingjia Alisa Wan, Yixin Wan, Zhen Wan, Zhongwei Wan, An Wang, Bang Wang, Baoxin Wang, Bichen Wang, Bing Wang, Bingbing Wang, Bingjie Wang, Bingqing Wang, Bo Wang, Bo Wang, Boxiang Wang, Chaojie Wang, Chaokun Wang, Chen Wang, Chen Wang, Cheng Wang, Chenglin Wang, Chenglong Wang, Chenxu Wang, Cong Wang, Cunxiang Wang, Daisy Zhe Wang, Daling Wang, Danqing Wang, Darren Yow-bang Wang, Di Wang, Di Wang, Dingmin Wang, Dingzirui Wang, Dongxia Wang, Fali Wang, Fan Wang, Gaoang Wang, Hai Wang, Haishuai Wang, Han Wang, Hao Wang, Hao Wang, Hao Wang, Haobo Wang, Haochun Wang, Haohan Wang, Haoran Wang, Haoran Wang, Haozhao Wang, Hebin Wang, Heng Wang, Heng Wang, Hengyi Wang, Hongning Wang, Hongru Wang, Hongya Wang, Houfeng Wang, Huadong Wang, Huazheng Wang, Huimin Wang, Huimu Wang, Ji Wang, Jia-ning Wang, Jialu Wang, Jian Wang, Jian Wang, Jiankun Wang, Jianyong Wang, Jianzong Wang, Jiapeng Wang, Jiaqi Wang, Jiashuo Wang, Jiawei Wang, Jiawei Wang, Jiayi Wang, Jie Wang, Jingwei Wang, Jingyu Wang, Jinqiao Wang, Jun Wang, Jun Wang, Junda Wang, Junlin Wang, Junsheng Wang, Junzhe Wang, Kai Wang, Kangrui Wang, Ke Wang, Kexiang Wang, Kexin Wang, Lanjun Wang, Lanrui Wang, Lei Wang, Li Wang, Liang Wang, Lianxi Wang, Lin Wang, Linlin Wang, Longbiao Wang, Lu Wang, Mengna Wang, Mengru Wang, Mengyu Wang, Ming Wang, Mingyang Wang, Pan Wang, Peiqi Wang, Peng Wang, Pengcheng Wang, Ping Wang, Qiang Wang, Qiang Wang, Qianli Wang, Qianlong Wang, Qianning Wang, Qing Wang, Qingni Wang, Qiqi Wang, Qiufeng Wang, Quan Wang, Renxi Wang, Rui Wang, Rui Wang, Ruibo Wang, Ruida Wang, Runhui Wang, Runze Wang, Ruoyao Wang, Shangfei Wang, Sheng Wang, Shengyuan Wang, Shida Wang, Shihan Wang, Shiqi Wang, Shuang Wang, Shuo Wang, Shuting Wang, Siqi Wang, Siqi Wang, Sirui Wang, Siyin Wang, Song Wang, Song Wang, Suge Wang, Suyuchen Wang, Tianchun Wang, Tianduo Wang, Tianshu Wang, Wei Wang, Wei-yao Wang, Weichuan Wang, Weixuan Wang, Weiyun Wang, Wen Wang, Wenjin Wang, Wenxuan Wang, Xi Wang, Xiang Wang, Xiangfeng Wang, Xianquan Wang, Xiao Wang, Xiaochen Wang, Xiaodong Wang, Xiaohan Wang, Xiaohua Wang, Xiaojie Wang, Xiaoling Wang, Xiaolong Wang, Xiaoqiang Wang, Xiawei

Wang, Xiaxia Wang, Xin Wang, Xin Wang, Xinfeng Wang, Xinglin Wang, Xinpeng Wang, Xintong Wang, Xinyi Wang, Xinyou Wang, Xinyu Wang, Xinyu Wang, Xinze Wang, Xiting Wang, Xu Wang, Xueqian Wang, Xunzhi Wang, Xuwu Wang, Yanbo Wang, Yanfeng Wang, Yang Wang, Yanhao Wang, Yanling Wang, Yaoming Wang, Yaqing Wang, Yasheng Wang, Ye Wang, Ye Wang, Ye Wang, Yibo Wang, Yichao Wang, Yichen Wang, Yifan Wang, Yifan Wang, Yihan Wang, Yijue Wang, Yike Wang, Yikun Wang, Yiming Wang, Yiming Wang, Yimu Wang, Yinggui Wang, Yining Wang, Yiran Wang, Yixu Wang, Yixuan Wang, Yixuan Wang, Yongjie Wang, Yu Wang, Yuchi Wang, Yue Wang, Yue Wang, Yue Wang, Yueqian Wang, Yuhang Wang, Yuhao Wang, Yuhao Wang, Yujie Wang, Yujie Wang, Yun Cheng Wang, Yuqi Wang, Yusong Wang, Yutong Wang, Yuxia Wang, Yuxin Wang, Yuxuan Wang, Zecheng Wang, Zekun Wang, Zekun Wang, Zekun Wang, Zezhong Wang, Zhaoyang Wang, Zhen Wang, Zheng Wang, Zheng Wang, Zhenglin Wang, Zhenhailong Wang, Zhenting Wang, Zhenyi Wang, Zhichun Wang, Zhilin Wang, Zhiruo Wang, Zhixiang Wang, Zhiyuan Wang, Zhizheng Wang, Zhizhi Wang, Zhongjie Wang, Zhuoer Wang, Zichao Wang, Zifu Wang, Zihan Wang, Zihan Wang, Zihan Wang, Zihao Wang, Zihao Wang, Ziheng Wang, Zijie Wang, Zili Wang, Zilong Wang, Zimu Wang, Ziqi Wang, Zirui Wang, Ziyue Wang, Wanli.li, Lilian Diana Awuor Wanzare, Nigel G. Ward, Greta Warren, Christian Wartena, Ishaan Watts, Ruvan Weerasinghe, Tharindu Cyril Weerasooriya, Anna Wegmann, Bowen Wei, Daimeng Wei, Fangyun Wei, Jerry Weng Wei, Jingxuan Wei, Johnny Wei, Kaiwen Wei, Kangda Wei, Lingwei Wei, Penghui Wei, Sheng-lun Wei, Xiao Wei, Xiaoyong Wei, Ying Wei, Yuancheng Wei, Yuchen Wei, Zhihua Wei, Zhongyu Wei, Zihao Wei, Gerhard Weikum, Maxwell Weinzierl, Fei Weizhi, Ji-rong Wen, Kaiyue Wen, Liang Wen, Xiaofei Wen, Xie Jia Wen, Ximing Wen, Yanlong Wen, Yule Wen, Yuqiao Wen, Zhiyuan Wen, Zhiyuan Wen, Wenbosi, Yixuan Weng, Peter West, Joyce Jiyong Whang, Steven Euijong Whang, Haryo Akbarianto Wibowo, Philipp Wicke, Rachel Wicks, Michael Wiegand, Matti Wiegmann, Adam Wiemerslage, Ethan Wilcox, Bryan Wilie, Miles Williams, Steven R. Wilson, Guillaume Wisniewski, Sondre Wold, Tomer Wolfson, Ngai Wong, Wichayaporn Wongkamjan, Honguk Woo, Zach Wood-doughty, Phil Woodland, Anna Wróblewska, Alexander Wu, Anne Wu, Ben Peng Wu, Bin Wu, Bohao Wu, Bowen Wu, Changli Wu, Cheng-kuang Wu, Di Wu, Di Wu, Fang Wu, Feijie Wu, Han Wu, Haolun Wu, Haoyi Wu, Haoyuan Wu, Honghan Wu, Hongyan Wu, Hongyi Wu, Hui Wu, Ji Wu, Jiageng Wu, Jialong Wu, Jian Wu, Jiancan Wu, Jiawei Wu, Jiaying Wu, Jie Wu, Jincenzi Wu, Jinge Wu, Junchao Wu, Junhong Wu, Junjie Wu, Junkang Wu, Kaixin Wu, Lianwei Wu, Likang Wu, Linjuan Wu, Mengyue Wu, Minghao Wu, Mingyan Wu, Muling Wu, Peilin Wu, Peng Wu, Qianhui Wu, Qingqiang Wu, Qinzhuo Wu, Sai Wu, Shangyu Wu, Shengguang Wu, Shih-hung Wu, Shu Wu, Shujin Wu, Sifan Wu, Sixing Wu, Taiqiang Wu, Ting-wei Wu, Tongtong Wu, Tsung Han Wu, Wei Wu, Weiqi Wu, Weiyi Wu, Winston Wu, Xianchao Wu, Xin Wu, Xindong Wu, Xiyang Wu, Xuansheng Wu, Xueqing Wu, Xuyang Wu, Yang Wu, Yang Wu, Yangjun Wu, Yating Wu, Yawen Wu, Yifan Wu, Yihan Wu, Yimeng Wu, Yin Wu, Yingnian Wu, Yixin Wu, Yuan Wu, Yuanhao Wu, Yurong Wu, Yuwei Wu, Yuxia Wu, Zekun Wu, Zengqing Wu, Zhaoxuan Wu, Zhijing Wu, Zhonghai Wu, Zhuoyang Wu, Zihui Wu, Zijun Wu, Zirui Wu, Zongru Wu, Zongyu Wu, Jörn Wübker

Wang Xi, Zhaohan Xi, Haotian Xia, Jingbo Xia, Lei Xia, Lianghao Xia, Peng Xia, Qingrong Xia, Rui Xia, Tian Xia, Tingyu Xia, Wu Xia, Xiao Xia, Yan Xia, Yu Xia, Hu Xiang, Jiannan Xiang, Lu Xiang, Wei Xiang, Yue Xianghu, Xi Xiangyu, Bei Xiao, Chaojun Xiao, Chenghao Xiao, Chuan Xiao, Chunyang Xiao, Jin Xiao, Junbin Xiao, Qingfa Xiao, Rong Xiao, Shitao Xiao, Tingsong Xiao, Tong Xiao, Yang Xiao, Yao Xiao, Yijia Xiao, Yisheng Xiao, Yongjie Xiao, Yunxuan Xiao, Yunze Xiao, Yuxin Xiao, Zhiqing Xiao, Ziang Xiao, Zilin Xiao, Chen Xiaoshu, Dong Xiaoyu, Alex Xie, Binzhu Xie, Hong Xie, Hongyan Xie, Kaige Xie, Lei Xie, Rui Xie, Ruobing Xie, Sean Xie, Sihong Xie, Tao Xie, Tianbao Xie, Wenya Xie, Xuan Xie, Yaqi Xie, Yiqing Xie, Yong Xie, Yuanzhen Xie, Yubo Xie, Yuexiang Xie, Yuqiang Xie, Yuxi Xie, Zhiwen Xie, Zhixin Xie, Zhongbin Xie, Zhongwei Xie, Zhouhang Xie, Zhuohan Xie, Amy Xin, Bowen Xing, Chen Xing, Linzi Xing, Rui Xing, Xingrun Xing, Xinyu Xing, Yue Xing, Bo Xiong, Hui Xiong, Jing

Xiong, Jinjun Xiong, Kai Xiong, Siheng Xiong, Wei Xiong, Weimin Xiong, Yun Xiong, Wei Wang Xjtlu, Baixuan Xu, Benfeng Xu, Bin Xu, Bo Xu, Boyan Xu, Cai Xu, Canwen Xu, Chejian Xu, Chen Xu, Chenchen Xu, Cong Xu, Dehong Xu, Depeng Xu, Derong Xu, Dongkuan Xu, Fangyuan Xu, Fangzhi Xu, Frank F. Xu, Hainiu Xu, Haiyang Xu, Hanzi Xu, Hao Xu, Haotian Xu, Heng-da Xu, Hong Xu, Hongshen Xu, Hongxia Xu, Hongyuan Xu, Hua Xu, Huimin Xu, Jia Xu, Jiahao Xu, Jialiang Xu, Jiarong Xu, Jin Xu, Jing Xu, Jitao Xu, Jun Xu, Justin Xu, Kaidi Xu, Kaishuai Xu, Ke Xu, Lei Xu, Lei Xu, Liang Xu, Linli Xu, Liyan Xu, Lu Xu, Lu Xu, Lvxiaowei Xu, Mayi Xu, Meilong Xu, Nan Xu, Paiheng Xu, Pengyu Xu, Qi Xu, Ran Xu, Renjing Xu, Rongwu Xu, Ruochen Xu, Shanshan Xu, Shicheng Xu, Tianyang Xu, Ting Xu, Wang Xu, Wanshi Xu, Wei Xu, Weijia Xu, Weijie Xu, Weikai Xu, Weiwen Xu, Wujiang Xu, Xi Xu, Xiangmin Xu, Xiangzhe Xu, Xiao Xu, Xin Xu, Xiong Xiao Xu, Yan Xu, Yang Xu, Yang Xu, Yang Xu, Yao Xu, Yifan Xu, Yige Xu, Yongxin Xu, Yuemei Xu, Yunqiu Xu, Yuzhuang Xu, Zhangchen Xu, Zhao Xu, Zhe Xu, Zhen Xu, Zhenran Xu, Zhichao Xu, Zhijian Xu, Zhikun Xu, Zhipeng Xu, Zishan Xu, Zunnan Xu, Qi Xuan, Boyang Xue, Chun Jason Xue, Dongyu Xue, Huiyin Xue, Mingfeng Xue, Shangzi Xue, Wei Xue, Yun Xue, Zhiyu Xue, Zihan Xue, Christos Xypolopoulos

Naganand Yadati, Ajay Yadav, Prateek Yadav, Srishti Yadav, Aditya Yadavalli, Kosuke Yamada, Junichi Yamagishi, Hiroaki Yamagiwa, Atsuki Yamaguchi, Ivan P. Yamshchikov, Bing Yan, Cheng Yan, Dong Yan, Jing Nathan Yan, Junchi Yan, Junyi Yan, Ming Yan, Ruiyi Yan, Shi-qi Yan, Tianyi Yan, Weixiang Yan, Xifeng Yan, Yang Yan, Yibo Yan, Yuchen Yan, Zhaohui Yan, Zihe Yan, Bang Yang, Bohao Yang, Carl Yang, Changbing Yang, Cheng Yang, Cheng Yang, Chenxu Yang, Chih-kai Yang, Diji Yang, Dingkang Yang, Dingyi Yang, Fan Yang, Fangkai Yang, Fengyi Yang, Grace Hui Yang, Guang Yang, Guang Yang, Guanqun Yang, Haoran Yang, Hongji Yang, Ivory Yang, Jack Yang, Jiayi Yang, Jiayi Yang, Jingbo Yang, Jinlong Yang, Jiuding Yang, June Yong Yang, Kaicheng Yang, Kailai Yang, Kejuan Yang, Li Yang, Liang Yang, Lingyu Yang, Longfei Yang, Ming-hsuan Yang, Mu Yang, Muyun Yang, Nan Yang, Pei Yang, Qikai Yang, Qiyuan Yang, Rui Yang, Ruosong Yang, Sen Yang, Shu Yang, Shuo Yang, Sohee Yang, Soyoung Yang, Tao Yang, Tiankai Yang, Tianyi Yang, Tianyu Yang, Tianyu Yang, Wen Yang, Wenkai Yang, Wenmian Yang, Xiao-wen Yang, Xiaocheng Yang, Xiaocui Yang, Xincheng Yang, Xinyi Yang, Xinyu Yang, Xiulin Yang, Xu Yang, Yahan Yang, Yajing Yang, Yi Yang, Yi Yang, Yifan Yang, Yifan Yang, Yifan Yang, Yifei Yang, Yingxiang Yang, Yiyuan Yang, Yizhe Yang, Yongjin Yang, Yuanhang Yang, Yue Yang, Yujiu Yang, Yuming Yang, Yuqing Yang, Ze Yang, Zeyuan Yang, Zhao Yang, Zhe Yang, Zhengdong Yang, Zhenhua Yang, Zhichao Yang, Zhiyu Yang, Zhou Yang, Zi Yang, Ziqing Yang, Ziyi Yang, Roman Yangarber, Ken Yano, Binwei Yao, Chang Yao, Dingyu Yao, Duanyi Yao, Fangzhou Yao, Feng Yao, Jiashu Yao, Jihan Yao, Jing Yao, Kaisheng Yao, Liang Yao, Lili Yao, Lina Yao, Man Yao, Maxwell Yao, Peiran Yao, Quanming Yao, Wei Yao, Xingcheng Yao, Yixiang Yao, Yuekun Yao, Yuhang Yao, Yunzhi Yao, Zijun Yao, Mahsa Yarmohammadi, Michihiro Yasunaga, Mert Yazan, Bingyang Ye, Chenchen Ye, Haoran Ye, Haotian Ye, Jiabo Ye, Jiasheng Ye, Jieping Ye, Jingheng Ye, Junjie Ye, Seonghyeon Ye, Tong Ye, Wei Ye, Wentao Ye, Xiao Ye, Yang Ye, Yuyang Ye, Min-hsuan Yeh, Sung-lin Yeh, An-zi Yen, Howard Yen, Seren Yenikent, Eunjung Yeo, Gerard Christopher Yeo, Jeonghun Yeo, Antonio Jimeno Yepes, Yfguo, Euiin Yi, Jingru Yi, Jingwei Yi, Mun Yong Yi, Xiaoyuan Yi, Du Yichao, Emine Yilmaz, Jinyeong Yim, Aoxiong Yin, Congchi Yin, Dawei Yin, Hang Yin, Heqi Yin, Hongzhi Yin, Jiong Yin, Kayo Yin, Pengcheng Yin, Weichong Yin, Wenjie Yin, Xunjian Yin, Yu Yin, Yuwei Yin, Zhangyue Yin, Zhenfei Yin, Ziyi Yin, Huaiyuan Ying, Jiahao Ying, Hong Yung Yip, Siu Ming Yiu, Gal Yona, Haneul Yoo, Hee Suk Yoon, Jinsung Yoon, Seunghyun Yoon, Soyoung Yoon, Sungroh Yoon, Sunjae Yoon, Susik Yoon, Wonjin Yoon, Ori Yoran, Issei Yoshida, Ryo Yoshida, Tomoya Yoshida, Fuming You, Wangjie You, Zhiwen You, Bei Yu, Botao Yu, Dezhi Yu, Dingyao Yu, Erxin Yu, Fangxu Yu, Guoxin Yu, Hang Yu, Hang Yu, Hao Yu, Hao Yu, Haofei Yu, Hwanjo Yu, Jifan Yu, Jun Yu, Le Yu, Lequan Yu, Mengxia Yu, Nenghai Yu, Peilin Yu, Pengfei Yu, Sangwon Yu, Seunguk Yu, Shi Yu, Shoubin Yu, Sicheng Yu, Sungduk Yu, Tan Yu, Tian Yu, Tianshu Yu, Tianshu Yu, Tiezheng Yu, Weijiang

Yu, Weijie Yu, Xiao Yu, Xiaobing Yu, Xiaodong Yu, Xiaoyan Yu, Xin Yu, Xinchun Yu, Xinyan Yu, Yahan Yu, Yangyang Yu, Zeping Yu, Zhang Yu, Zhi Yu, Zhuohao Yu, Zihan Yu, Bo Yuan, Caixia Yuan, Chenhan Yuan, Chunyuan Yuan, Fei Yuan, Hongbang Yuan, Hongyi Yuan, Jiahao Yuan, Jiaqing Yuan, Jiayi Yuan, Jingyang Yuan, Li Yuan, Peiwen Yuan, Shaozu Yuan, Shuhan Yuan, Shuzhou Yuan, Siyu Yuan, Yancheng Yuan, Ye Yuan, Ye Yuan, Yige Yuan, Youliang Yuan, Zhangdie Yuan, Zheng Yuan, Zhihang Yuan, Yuanrao, Linan Yue, Murong Yue, Xiangyu Yue, Yuanhao Yue, Hyeongun Yun, Jungmin Yun, Se-young Yun, Seongjun Yun, Sukwon Yun, Taewon Yun, Tian Yun, Yunfanxie, Frances Yung, François Yvon

Joe Z, Muhammad Bilal Zafar, Hamada Zahera, Iqra Zahid, Pardis Sadat Zahraei, Mohd Abbas Zaidi, Mohammadi Zaki, Mohammed J Zaki, Mahdi Zakizadeh, Kerem Zaman, Hamed Zamani, Olga Zamaraeva, Changtong Zan, Daoguang Zan, Hongying Zan, Yuhang Zang, Andrea Zaninello, Klim Zaporozjets, Alessandra Zarccone, Sina Zarrieß, Niamat Zawad, Piotr Zelasko, Aohan Zeng, Daojian Zeng, Fengzhu Zeng, Gaby Zeng, Guangtao Zeng, Hongchuan Zeng, Huimin Zeng, Jayden Zeng, Jiali Zeng, Jichuan Zeng, Jingjie Zeng, Jinshan Zeng, Min Zeng, Peijian Zeng, Qi Zeng, Qingcheng Zeng, Qingkai Zeng, Shenglai Zeng, Weixin Zeng, Xin Zeng, Yutao Zeng, Zhiyuan Zeng, Ding Zepeng, Yang Zerui, Chrysoula Zerva, George Zerveas, Albert Zeyer, Deniz Zeyrek, Sheng Zha, Zheng-jun Zha, Chengxiang Zhai, Weihe Zhai, Haolan Zhan, Hongli Zhan, Huixin Zhan, Li-ming Zhan, Pengwei Zhan, Qiusi Zhan, Runzhe Zhan, Tianyu Zhan, Zheng Zhan, Andrew Zhang, Bing Zhang, Binyan Zhang, Bo Zhang, Bohan Zhang, Bokai Zhang, Bowen Zhang, Chao Zhang, Chao Zhang, Chao Zhang, Chaoran Zhang, Chaoyun Zhang, Chen Zhang, Chenhao Zhang, Chenyang Zhang, Chi Zhang, Chong Zhang, Chong Zhang, Chunhui Zhang, Chunxia Zhang, Chuxu Zhang, Congzhi Zhang, Dong Zhang, Dongyu Zhang, Duzhen Zhang, Fan Zhang, Feng Zhang, Fengji Zhang, Fu Zhang, Fujun Zhang, Fuwei Zhang, Fuzheng Zhang, Gai-fan Zhang, Ge Zhang, Gengyuan Zhang, Guchun Zhang, Haichao Zhang, Hainan Zhang, Hanlei Zhang, Hanqing Zhang, Hao Zhang, Haoran Ranran Zhang, Haowei Zhang, Haoyu Zhang, Hongbo Zhang, Hongyu Zhang, Hongzhi Zhang, Hu Zhang, Huajian Zhang, Huishuai Zhang, Huixuan Zhang, Jiajie Zhang, Jianjin Zhang, Jianshu Zhang, Jiarui Zhang, Jie Zhang, Jinchuan Zhang, Jinghao Zhang, Jinghuai Zhang, Jinghui Zhang, Jingyu Zhang, Jingyuan Zhang, Jinpeng Zhang, Jintian Zhang, Jinyi Zhang, Jipeng Zhang, Jiwen Zhang, Jizhi Zhang, Jue Zhang, Jun Zhang, Juntian Zhang, Kai Zhang, Kai Zhang, Kai Zhang, Kai Zhang, Kaike Zhang, Kaipeng Zhang, Kaiyan Zhang, Kaiyi Zhang, Ke Zhang, Ke Zhang, Kechi Zhang, Kejun Zhang, Kexun Zhang, Kun Zhang, Kun Zhang, Lan Zhang, Le Zhang, Lei Zhang, Lei Zhang, Lei Zhang, Liang Zhang, Liang Zhang, Licheng Zhang, Lin Zhang, Lingling Zhang, Lingming Zhang, Linhai Zhang, Linhao Zhang, Litian Zhang, Longhui Zhang, Longyin Zhang, Meiru Zhang, Mengqi Zhang, Mi Zhang, Mian Zhang, Michael J Zhang, Mike Zhang, Min Zhang, Mingyang Zhang, Mozhi Zhang, Muru Zhang, Peinan Zhang, Peitian Zhang, Peiyuan Zhang, Peng Zhang, Qi Zhang, Qi Zhang, Qiang Zhang, Qiang Zhang, Qihui Zhang, Qing Zhang, Qinggang Zhang, Qingjie Zhang, Qingru Zhang, Qiyuan Zhang, Richong Zhang, Rui Zhang, Ruiquan Zhang, Ruochen Zhang, Ruoyu Zhang, Sai Qian Zhang, Shao Zhang, Sheng Zhang, Shengyu Zhang, Shujian Zhang, Shuo Zhang, Sicheng Zhang, Songming Zhang, Songyang Zhang, Tao Zhang, Tianhui Zhang, Tianlin Zhang, Tianwei Zhang, Tianyang Zhang, Tianyi Zhang, Tianyi Zhang, Tianyi Zhang, Wei Zhang, Wei Emma Zhang, Weichen Zhang, Weizhi Zhang, Wen Zhang, Wen Zhang, Wenbin Zhang, Wenjun Zhang, Wenqian Zhang, Wentao Zhang, Wenyan Zhang, Wenzheng Zhang, Xi Zhang, Xiang Zhang, Xiang Zhang, Xiangyu Zhang, Xiangyu Zhang, Xiantao Zhang, Xiao Zhang, Xiaofan Zhang, Xiaofeng Zhang, Xiaokun Zhang, Xiaoming Zhang, Xiaotong Zhang, Xiaowang Zhang, Xiaoying Zhang, Xiaoying Zhang, Xin Zhang, Xin Zhang, Xinghua Zhang, Xinjie Zhang, Xinliang Frederick Zhang, Xinlu Zhang, Xitong Zhang, Xu Zhang, Xu Zhang, Xuan Zhang, Xuan Zhang, Xuanming Zhang, Xueyao Zhang, Xuhong Zhang, Ya Zhang, Yan Zhang, Yancheng Zhang, Yanyong Zhang, Yanyue Zhang, Yanzhe Zhang, Yao Zhang, Yaping Zhang, Yazhou Zhang, Ye Zhang, Yehong Zhang, Yi Zhang, Yice Zhang, Yichao Zhang, Yichi Zhang, Yichi Zhang, Yichi Zhang, Yidan Zhang, Yifan Zhang,

Yifan Zhang, Yigeng Zhang, Yihe Zhang, Yijia Zhang, Yikai Zhang, Yiming Zhang, Yin Zhang, Ying Zhang, Yingji Zhang, Yinqi Zhang, Yiqun Zhang, Yizhuo Zhang, Yongheng Zhang, Yongqi Zhang, You Zhang, Yu Zhang, Yuan Zhang, Yuansen Zhang, Yuanxing Zhang, Yuanzhe Zhang, Yuchen Zhang, Yue Zhang, Yue Zhang, Yuejie Zhang, Yuhao Zhang, Yuhao Zhang, Yuji Zhang, Yunfan Zhang, Yunjian Zhang, Yunxiang Zhang, Yunyi Zhang, Yuqi Zhang, Yusen Zhang, Yuwei Zhang, Yuxiang Zhang, Zecheng Zhang, Zekai Zhang, Zeliang Zhang, Zequn Zhang, Zhanqiu Zhang, Zhao Zhang, Zhao Zhang, Zhehao Zhang, Zhen Zhang, Zhen Zhang, Zhen Zhang, Zhen-ru Zhang, Zheng Zhang, Zhengkun Zhang, Zhenhao Zhang, Zhenyu Zhang, Zhexin Zhang, Zheyu Zhang, Zheyuan Zhang, Zhi Zhang, Zhihan Zhang, Zhikang Zhang, Zhiyang Zhang, Zhuo Zhang, Zhuowei Zhang, Zihan Zhang, Zijian Zhang, Ziyin Zhang, Baining Zhao, Bowen Zhao, Chenxi Zhao, Fei Zhao, Feng Zhao, Feng Zhao, Guangxiang Zhao, Guangzhen Zhao, Guoshuai Zhao, Haiquan Zhao, Haiteng Zhao, Haiyan Zhao, Hengyuan Zhao, Hui Zhao, Jiahao Zhao, Jianqiao Zhao, Jianyu Zhao, Jiaxu Zhao, Jie Zhao, Jin Zhao, Jinman Zhao, Junbo Zhao, Kai Zhao, Kaiyan Zhao, Kun Zhao, Libo Zhao, Lili Zhao, Lingchen Zhao, Lingjun Zhao, Lulu Zhao, Mengjie Zhao, Mengyu Zhao, Minyi Zhao, Pengyu Zhao, Pu Zhao, Qinghua Zhao, Rui Zhao, Ruining Zhao, Runhao Zhao, Ruochen Zhao, Sendong Zhao, Shiwan Zhao, Shu Zhao, Shuai Zhao, Wei Zhao, Weilin Zhao, Weixiang Zhao, Wenlong Zhao, Wenting Zhao, Xiangyu Zhao, Xiangyu Zhao, Xiaoyan Zhao, Xingmeng Zhao, Xingyi Zhao, Xinpeng Zhao, Xinyu Zhao, Xiutian Zhao, Xuechen Zhao, Xueliang Zhao, Xujiang Zhao, Yang Zhao, Yang Yang Zhao, Yanyan Zhao, Yi Zhao, Yibo Zhao, Yiran Zhao, Yiren Zhao, Yizhou Zhao, Yongwei Zhao, Yu Zhao, Yu Zhao, Yukun Zhao, Yunxiao Zhao, Yutian Zhao, Yuze Zhao, Zhe Zhao, Zheng Zhao, Zhengyu Zhao, Zhenjie Zhao, Zhou Zhao, Zihao Zhao, Zijian Zhao, Zirui Zhao, Zhaolingchen, Bi Zhen, Ranran Zhen, Changmeng Zheng, Duo Zheng, Guangmin Zheng, Huan Zheng, Huanran Zheng, Jingnan Zheng, Jonathan Zheng, Kai Zheng, Mengxin Zheng, Mingyu Zheng, Shen Zheng, Tianshi Zheng, Tianyu Zheng, Tong Zheng, Weiguo Zheng, Weixiong Zheng, Xiaolin Zheng, Xiaoqing Zheng, Xiawu Zheng, Yefeng Zheng, Yi Zheng, Yilun Zheng, Yinhe Zheng, Yongsen Zheng, Zhe Zheng, Zhi Zheng, Zi'ou Zheng, Zibin Zheng, Zihao Zheng, Shi Zhi, Zhibinlan, Jialun Zhong, Jianyuan Zhong, Liangjun Zhong, Meizhi Zhong, Mian Zhong, Qihuang Zhong, Shanshan Zhong, Shaochen Zhong, Xian Zhong, Yang Zhong, Yujie Zhong, Zhi Zhong, Baohang Zhou, Bin Zhou, Changzhi Zhou, Cheng-jie Zhou, Chuang Zhou, Dongzhan Zhou, Fan Zhou, Gengze Zhou, Hao Zhou, Hao Zhou, Houquan Zhou, Huachi Zhou, Huichi Zhou, Huiwei Zhou, Jiaming Zhou, Jiayu Zhou, Jinfeng Zhou, Junsheng Zhou, Kang Zhou, Kun Zhou, Lexin Zhou, Li Zhou, Linqi Zhou, Luowei Zhou, Meng Zhou, Mengyu Zhou, Mingyang Zhou, Mingyao Zhou, Naitian Zhou, Peilin Zhou, Qiji Zhou, Shijia Zhou, Shilin Zhou, Shuigeng Zhou, Tianshuo Zhou, Tianyi Zhou, Tong Zhou, Wei Zhou, Wei Zhou, Weixiao Zhou, Wengang Zhou, Xiabing Zhou, Xiangyang Zhou, Xiaobing Zhou, Xiaoling Zhou, Xin Zhou, Xin Zhou, Xinyu Zhou, Yan Zhou, Yang Zhou, Yang Zhou, Yaqian Zhou, Yijie Zhou, Yin Zhou, Ying Zhou, Yongxin Zhou, Yougen Zhou, Yue Zhou, Yufa Zhou, Yufan Zhou, Yuhang Zhou, Yujun Zhou, Yukai Zhou, Yunhua Zhou, Yuqi Zhou, Yuxiang Zhou, Zhanke Zhou, Zhejian Zhou, Zhenhong Zhou, Zhoujunhao, Bolin Zhu, Conghui Zhu, Dawei Zhu, Derui Zhu, Fangqi Zhu, Fangwei Zhu, Hanlun Zhu, Henghui Zhu, Hongsong Zhu, Jia Zhu, Jian Zhu, Jianke Zhu, Jing Zhu, Junda Zhu, Junnan Zhu, Junyi Zhu, Kaijie Zhu, Kenny Q. Zhu, Kun Zhu, Linchao Zhu, Luyao Zhu, Mengna Zhu, Minfeng Zhu, Ming Zhu, Minjun Zhu, Muhua Zhu, Peide Zhu, Qi Zhu, Qingfu Zhu, Rui-jie Zhu, Shanfeng Zhu, Shaolin Zhu, Song-chun Zhu, Tinghui Zhu, Tong Zhu, Wang Zhu, Wei Zhu, Xiangrong Zhu, Xiaofeng Zhu, Xiaofeng Zhu, Xuekai Zhu, Yaochen Zhu, Yeshuang Zhu, Yi Zhu, Yi Zhu, Yingjie Zhu, Yinglun Zhu, Yonghua Zhu, Yongxin Zhu, Yutao Zhu, Yuying Zhu, Zhihong Zhu, Ziwei Zhu, Zixiao Zhu, Fuzhen Zhuang, Haojie Zhuang, Haomin Zhuang, Huiping Zhuang, Shengyao Zhuang, Xiang Zhuang, Xianwei Zhuang, Xinlin Zhuang, Yuan Zhuang, Yuchen Zhuang, Ziyu Zhuang, Jingming Zhuo, Alireza Salkhordeh Ziabari, Caleb Ziem, Brian Douglas Zimmerman, Roger Zimmermann, Yftah Ziser, Xu Zixiang, Zixuanwang, Stefan Zohren, Chang Zong, Linlin Zong, Qing Zong, Zongke, Bocheng Zou, Henry Peng Zou, Heqing Zou, Jiaru Zou, Jie Zou, Lixin Zou, Yanyan Zou, Zhuoyang Zou, Amal Zouaq, Vilém

Zouhar, Andrea Zugarini, Liang Zujie, Jialung Zuo, Kaiwen Zuo, Max Zuo, Yuxin Zuo

Keynote Talk

Rethinking Pretraining: Data and Architecture

Luke Zettlemoyer

Paul G. Allen School of Computer Science and Engineering, University of Washington



Mon, July 28th, 2025 – Time: 09:30 – 10:30 – Room: Austria Center Vienna

Abstract: Large language model training follows a standard pipeline: tokenization, pretraining, possibly mid-training, and post training or alignment. Despite its wild success, we understand relatively little about this recipe and are almost certainly missing many opportunities to improve it. In this talk, I will focus on three such cases. I’ll describe our work on data efficient post training (e.g. LIMA, ALMA, and s1) where we argue that nearly all advanced model capabilities ultimately come from the pretraining data, even if effective alignment is still essential for controlling model behavior. I will also describe new methods for extracting more signal from the pretraining data, including new hierarchical architectures for byte-level language models (e.g. BLT) that are both tokenizer-free and scale better than traditional BPE-based methods, especially in the long tail. Finally, I will discuss decentralized, modular training algorithms (e.g. BTM) that better isolate and control the influence of specific data on specific model components and behaviors. Together, these methods promise to simplify training and improve scaling, by centering and amplifying the influence of data in architecture design.

Bio: Luke Zettlemoyer is a Professor in the Paul G. Allen School of Computer Science & Engineering at the University of Washington, and a Senior Research Director at Meta. His research focuses on empirical methods for natural language semantics, and involves designing machine learning algorithms, introducing new tasks and datasets, and, most recently, studying how to best develop new architectures and self-supervision signals for pre-training. His honors include being elected ACL President, named an ACL Fellow, winning a PECASE award, an Allen Distinguished Investigator award, and multiple best paper awards. Luke was an undergrad at NC State, received his PhD from MIT and was a postdoc at the University of Edinburgh.

Keynote Talk

Whose Gold? Re-imagining Alignment for Truly Beneficial AI

Verena Rieser
Google DeepMind



Tue, July 29th, 2025 – Time: 09:00 – 10:00 – Room: Austria Center Vienna

Abstract: Human feedback is often the “gold standard” for AI alignment, but what if this “gold” reflects diverse, even contradictory human values? This keynote explores the technical and ethical challenges of building beneficial AI when values conflict – not just between individuals, but also within them. My talk advocates for a dual expansion of the AI alignment framework: moving beyond a single, monolithic viewpoint to a plurality of perspectives, and transcending narrow safety and engagement metrics to promote comprehensive human well-being.

Bio: Verena Rieser is a Senior Staff Research Scientist at Google DeepMind, where she founded the VOICES team (Voices-of-all in alignment). Her team is a core contributor to Gemini with a mission to enhance model safety and usability for diverse communities. Verena has pioneered work in data-driven multimodal Dialogue Systems and Natural Language Generation, encompassing conversational RL agents, faithful data-to-text generation, spoken language understanding, evaluation methodologies, and applications of AI for societal good. Verena previously directed the NLP lab as a full professor at Heriot-Watt University, Edinburgh, and held a Royal Society Leverhulme Senior Research Fellowship. She earned her PhD from Saarland University.

Panel

Generalization of NLP Models

Mon, July 28th, 2025 – Time: 16:30 – 17:30 – Room: Austria Center Vienna - Level 2 Hall A - C

Moderator: Eduard Hovy, Melbourne Connect University of Melbourne



Eduard Hovy is the Executive Director of Melbourne Connect (a research and tech transfer centre at the University of Melbourne), a professor at the University of Melbourne’s School of Computing and Information Systems, and an adjunct professor at the Language Technologies Institute in the School of Computer Science at Carnegie Mellon University. In 2020–21 he served as Program Manager in DARPA’s Information Innovation Office (I2O), where he managed programs in Natural Language Technology and Data Analytics. Dr. Hovy completed a Ph.D. in Computer Science (Artificial Intelligence) at Yale University in 1987 and was awarded honorary doctorates from the National Distance Education University (UNED) in Madrid in 2013 and the University of Antwerp in 2015. He is one of the initial 17 Fellows of the Association for Computational Linguistics (ACL) and is also a Fellow of the Association for the Advancement of Artificial Intelligence (AAAI). Dr. Hovy’s research focuses on computational semantics of language and addresses various areas in Natural Language Processing and Data Analytics, including in-depth machine reading of text, information extraction, automated text summarization, question answering, the semi-automated construction of large lexicons and ontologies, and machine translation. In early 2025 his Google h-index was 109, with about 68,000 citations. Dr. Hovy is the author or co-editor of eight books and over 400 technical articles and is a popular invited speaker. He regularly co-taught Ph.D.-level courses and has served on Advisory and Review Boards for both research institutes and funding organizations in Germany, Italy, Netherlands, Ireland, Singapore, and the USA. From 2003 to 2015 he was co-Director of Research for the Department of Homeland Security’s Center of Excellence for Command, Control, and Interoperability Data Analytics, a distributed cooperation of 17 universities. In 2001 Dr. Hovy served as President of the international Association of Computational Linguistics (ACL), in 2001–03 as President of the International Association of Machine Translation (IAMT), and in 2010–11 as President of the Digital Government Society (DGS).

Panelists:

Mirella Lapata, University of Edinburgh



Mirëlla Lapata is professor of natural language processing in the School of Informatics at the University of Edinburgh. Her research focuses on getting computers to understand, reason with, and generate natural language. She is the first recipient (2009) of the British Computer Society and Information Retrieval Specialist Group (BCS/IRSG) Karen Sparck Jones award and a Fellow of the Royal Society of Edinburgh, the ACL, and Academia Europaea. Mirëlla has also received best paper awards in leading NLP conferences and has served on the editorial boards of the Journal of Artificial Intelligence Research, the Transactions of the ACL, and Computational Linguistics. She was president of SIGDAT (the group that organizes EMNLP) in 2018. She has been awarded an ERC consolidator grant, a Royal Society Wolfson Research Merit Award, and a UKRI Turing AI World-Leading Researcher Fellowship.

Dan Roth, University of Pennsylvania and Oracle



Dan Roth is the Eduardo D. Glandt Distinguished Professor at the University of Pennsylvania and Chief AI Scientist at Oracle. Until June 2024 Dan was a VP/Distinguished Scientist at AWS AI where he led the scientific effort behind Amazon’s first-generation GenAI products, including Titan Models, Amazon Q, and Amazon Bedrock. Dan is a Fellow of the AAAS, ACM, AAAI, and ACL, and a recipient of the IJCAI John McCarthy Award “for major conceptual and theoretical advances in the modeling of natural language understanding, machine learning, and reasoning.” He has published broadly in natural language processing, machine learning, knowledge representation and reasoning, and learning theory, was the Editor-in-Chief of the Journal of Artificial Intelligence Research (JAIR) and has served as a Program Chair and Conference Chair for the major conferences in his research areas. Roth has been involved in several ML/NLP/GenAI startups in domains that range from legal and compliance to health care. Dan received his B.A Summa cum laude in Mathematics from the Technion, Israel and his Ph.D. in Computer Science from Harvard University in 1995.

Yue Zhang, Westlake University



Yue Zhang is a tenured Professor at Westlake University (<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” and served as a PC co-chair for CCL 2020 and EMNLP 2022, test-of-time award committee co-chairs for ACL 2024 and 2025, action editor for TACL, and associate editor for TASLP, TALLIP, TBD, and CSL. He won the best paper awards of IALP 2017 and COLING 2018, best paper honorable mention of SemEval 2020, and best paper nomination for ACL 2018 and ACL 2023.

Table of Contents

<i>EcomScriptBench: A Multi-task Benchmark for E-commerce Script Planning via Step-wise Intention-Driven Product Association</i>	
Weiqi Wang, Limeng Cui, Xin Liu, Sreyashi Nag, Wenju Xu, Chen Luo, Sheikh Muhammad Sarwar, Yang Li, Hansu Gu, Hui Liu, Changlong Yu, Jiaxin Bai, Yifan Gao, Haiyang Zhang, Qi He, Shuiwang Ji and Yangqiu Song	1
<i>GraphNarrator: Generating Textual Explanations for Graph Neural Networks</i>	
Bo Pan, Zhen Xiong, Guanchen Wu, Zheng Zhang, Yifei Zhang, Yuntong Hu and Liang Zhao	23
<i>M-RewardBench: Evaluating Reward Models in Multilingual Settings</i>	
Srishti Gureja, Lester James Validad Miranda, Shayekh Bin Islam, Rishabh Maheshwary, Drishti Sharma, Gusti Triandi Winata, Nathan Lambert, Sebastian Ruder, Sara Hooker and Marzieh Fadaee	43
<i>ELABORATION: A Comprehensive Benchmark on Human-LLM Competitive Programming</i>	
Xinwei Yang, Zhaofeng Liu, Chen Huang, Jiashuai Zhang, Tong Zhang, Yifan Zhang and Wenqiang Lei	59
<i>The Impossibility of Fair LLMs</i>	
Jacy Reese Anthis, Kristian Lum, Michael Ekstrand, Avi Feller and Chenhao Tan	105
<i>Intuitive Fine-Tuning: Towards Simplifying Alignment into a Single Process</i>	
Ermo Hua, Biqing Qi, Kaiyan Zhang, Kai Tian, Xingtai Lv, Ning Ding and Bowen Zhou	121
<i>Bias in Language Models: Beyond Trick Tests and Towards RUTEd Evaluation</i>	
Kristian Lum, Jacy Reese Anthis, Kevin Robinson, Chirag Nagpal and Alexander Nicholas D’Amour	137
<i>Sliding Windows Are Not the End: Exploring Full Ranking with Long-Context Large Language Models</i>	
Wenhan Liu, Xinyu Ma, Yutao Zhu, Ziliang Zhao, Shuaiqiang Wang, Dawei Yin and Zhicheng Dou	162
<i>The Impact of Auxiliary Patient Data on Automated Chest X-Ray Report Generation and How to Incorporate It</i>	
Aaron Nicolson, Shengyao Zhuang, Jason Dowling and Bevan Koopman	177
<i>CLEME2.0: Towards Interpretable Evaluation by Disentangling Edits for Grammatical Error Correction</i>	
Jingheng Ye, Zishan Xu, Yinghui Li, Linlin Song, Qingyu Zhou, Hai-Tao Zheng, Ying Shen, Wenhao Jiang, Hong-Gee Kim, Ruitong Liu, Xin Su and Zifei Shan	204
<i>StrucText-Eval: Evaluating Large Language Model’s Reasoning Ability in Structure-Rich Text</i>	
Zhouhong Gu, Haoning Ye, Xingzhou Chen, Zeyang Zhou, Hongwei Feng and Yanghua Xiao	223
<i>Literature Meets Data: A Synergistic Approach to Hypothesis Generation</i>	
Haokun Liu, Yangqiaoyu Zhou, Mingxuan Li, Chenfei Yuan and Chenhao Tan	245
<i>GAPO: Learning Preferential Prompt through Generative Adversarial Policy Optimization</i>	
Zhouhong Gu, Xingzhou Chen, Xiaoran Shi, Tao Wang, Suhang Zheng, Tianyu Li, Hongwei Feng and Yanghua Xiao	282
<i>Tree-of-Evolution: Tree-Structured Instruction Evolution for Code Generation in Large Language Models</i>	
Ziyang Luo, Kaixin Li, Hongzhan Lin, Yuchen Tian, Mohan Kankanhalli and Jing Ma	297

<i>Delving into Multilingual Ethical Bias: The MSQAD with Statistical Hypothesis Tests for Large Language Models</i>	
Seunguk Yu, Juhwan Choi and YoungBin Kim	317
<i>ReSCORE: Label-free Iterative Retriever Training for Multi-hop Question Answering with Relevance-Consistency Supervision</i>	
Dosung Lee, Wonjun Oh, Boyoung Kim, Minyoung Kim, Joonsuk Park and Paul Hongsuck Seo	341
<i>FACT-AUDIT: An Adaptive Multi-Agent Framework for Dynamic Fact-Checking Evaluation of Large Language Models</i>	
Hongzhan Lin, Yang Deng, Yuxuan Gu, Wenxuan Zhang, Jing Ma, See-Kiong Ng and Tat-Seng Chua	360
<i>Statistical Deficiency for Task Inclusion Estimation</i>	
Loïc Fosse, Frederic Bechet, Benoit Favre, Géraldine Damnati, Gwénoél Lecorvé, Maxime Darin, Philippe Formont and Pablo Piantanida	382
<i>Towards Robust and Efficient Federated Low-Rank Adaptation with Heterogeneous Clients</i>	
Jabin Koo, Minwoo Jang and Jungseul Ok	416
<i>LLM-Powered Test Case Generation for Detecting Bugs in Plausible Programs</i>	
Kaibo Liu, Zhenpeng Chen, Yiyang Liu, Jie M. Zhang, Mark Harman, Yudong Han, Yun Ma, Yihong Dong, Ge Li and Gang Huang	430
<i>Capture the Key in Reasoning to Enhance CoT Distillation Generalization</i>	
Chengwei Dai, Kun Li, Wei Zhou and Songlin Hu	441
<i>How to Enable Effective Cooperation Between Humans and NLP Models: A Survey of Principles, Formalizations, and Beyond</i>	
Chen Huang, Yang Deng, Wenqiang Lei, Jiancheng Lv, Tat-Seng Chua and Jimmy Huang . .	466
<i>Enhancing Hyperbole and Metaphor Detection with Their Bidirectional Dynamic Interaction and Emotion Knowledge</i>	
Li Zheng, Sihang Wang, Hao Fei, Zuquan Peng, Fei Li, Jianming Fu, Chong Teng and Donghong Ji	489
<i>UniICL: An Efficient ICL Framework Unifying Compression, Selection, and Generation</i>	
Jun Gao, Qi Lv, Zili Wang, Tianxiang Wu, Ziqiang Cao and Wenjie Li	500
<i>BelarusianGLUE: Towards a Natural Language Understanding Benchmark for Belarusian</i>	
Maksim Aparovich, Volha Harytskaya, Vladislav Poritski, Oksana Volchek and Pavel Smrz .	511
<i>A Survey on Foundation Language Models for Single-cell Biology</i>	
Fan Zhang, Hao Chen, Zhihong Zhu, Ziheng Zhang, Zhenxi Lin, Ziyue Qiao, Yefeng Zheng and Xian Wu	528
<i>RuleArena: A Benchmark for Rule-Guided Reasoning with LLMs in Real-World Scenarios</i>	
Ruiwen Zhou, Wenyue Hua, Liangming Pan, Sitao Cheng, Xiaobao Wu, En Yu and William Yang Wang	550
<i>Extending LLM Context Window with Adaptive Grouped Positional Encoding: A Training-Free Method</i>	
Xinhao Xu, Jiabin Li, Hui Chen, Zijia Lin, Jungong Han and Guiguang Ding	573
<i>Semantic Exploration with Adaptive Gating for Efficient Problem Solving with Language Models</i>	
Sungjae Lee, Hyejin Park, Jaechang Kim and Jungseul Ok	588

<i>HotelMatch-LLM: Joint Multi-Task Training of Small and Large Language Models for Efficient Multimodal Hotel Retrieval</i>	
Arian Askari, Emmanouil Stergiadis, Ilya Gusev and Moran Beladev	607
<i>Can Multimodal Large Language Models Understand Spatial Relations?</i>	
Jingping Liu, Ziyang Liu, Zhedong Cen, Yan Zhou, Yinan Zou, Weiyan Zhang, Haiyun Jiang and Tong Ruan	620
<i>S³ - Semantic Signal Separation</i>	
Márton Kardos, Jan Kostkan, Kenneth Enevoldsen, Arnault-Quentin Vermillet, Kristoffer Nielbo and Roberta Rocca	633
<i>TrimLLM: Progressive Layer Dropping for Domain-Specific LLMs</i>	
Lanxiang Hu, Tajana Rosing and Hao Zhang	667
<i>JuStRank: Benchmarking LLM Judges for System Ranking</i>	
Ariel Gera, Odellia Boni, Yotam Perlitz, Roy Bar-Haim, Lilach Eden and Asaf Yehudai	682
<i>Generating Diverse Training Samples for Relation Extraction with Large Language Models</i>	
Zexuan Li, Hongliang Dai and Piji Li	713
<i>MultiSocial: Multilingual Benchmark of Machine-Generated Text Detection of Social-Media Texts</i>	
Dominik Macko, Jakub Kopál, Robert Moro and Ivan Srba	727
<i>Efficient and Accurate Prompt Optimization: the Benefit of Memory in Exemplar-Guided Reflection</i>	
Cilin Yan, Jingyun Wang, Lin Zhang, Ruihui Zhao, Xiaopu Wu, Kai Xiong, Qingsong Liu, Guoliang Kang and Yangyang Kang	753
<i>Evaluation of LLM Vulnerabilities to Being Misused for Personalized Disinformation Generation</i>	
Aneta Zugecova, Dominik Macko, Ivan Srba, Robert Moro, Jakub Kopál, Katarína Marcinčinová and Matúš Mesarčík	780
<i>EscapeBench: Towards Advancing Creative Intelligence of Language Model Agents</i>	
Cheng Qian, Peixuan Han, Qinyu Luo, Bingxiang He, Xiushi Chen, Yuji Zhang, Hongyi Du, Jiarui Yao, Xiaocheng Yang, Denghui Zhang, Yunzhu Li and Heng Ji	798
<i>BPP-Search: Enhancing Tree of Thought Reasoning for Mathematical Modeling Problem Solving</i>	
Teng Wang, Wing Yin YU, Zhenqi He, Zehua Liu, HaileiGong HaileiGong, Han Wu, Xiongwei Han, Wei Shi, Ruifeng She, Fangzhou Zhu and Tao Zhong	821
<i>LACA: Improving Cross-lingual Aspect-Based Sentiment Analysis with LLM Data Augmentation</i>	
Jakub Šmíd, Pavel Priban and Pavel Kral	839
<i>Fusing Highly Specialized Language Models for Comprehensive Expertise</i>	
Ning Ding, Yulin Chen, Ganqu Cui, Xingtai Lv, Weilin Zhao, Kaiyan Zhang, Ruobing Xie, Bowen Zhou, Zhiyuan Liu and Maosong Sun	854
<i>HybGRAG: Hybrid Retrieval-Augmented Generation on Textual and Relational Knowledge Bases</i>	
Meng-Chieh Lee, Qi Zhu, Costas Mavromatis, Zhen Han, Soji Adeshina, Vassilis N. Ioannidis, Huzefa Rangwala and Christos Faloutsos	879
<i>Re-ranking Using Large Language Models for Mitigating Exposure to Harmful Content on Social Media Platforms</i>	
Rajvardhan Oak, Muhammad Haroon, Claire Wonjeong jo, Magdalena Wojcieszak and Anshuman Chhabra	894

<i>Aligning AI Research with the Needs of Clinical Coding Workflows: Eight Recommendations Based on US Data Analysis and Critical Review</i>	
Yidong Gan, Maciej Rybinski, Ben Hachey and Jonathan K. Kummerfeld	909
<i>MIND: A Multi-agent Framework for Zero-shot Harmful Meme Detection</i>	
Ziyan Liu, Chunxiao Fan, Haoran Lou, Yuexin Wu and Kaiwei Deng	923
<i>EvoWiki: Evaluating LLMs on Evolving Knowledge</i>	
Wei Tang, Yixin Cao, Yang Deng, Jiahao Ying, Bo Wang, Yizhe Yang, Yuyue Zhao, Qi Zhang, Xuanjing Huang, Yu-Gang Jiang and Yong Liao	948
<i>Rethinking Repetition Problems of LLMs in Code Generation</i>	
Yihong Dong, Yuchen Liu, Xue Jiang, Bin Gu, Zhi Jin and Ge Li	965
<i>PunchBench: Benchmarking MLLMs in Multimodal Punchline Comprehension</i>	
Kun Ouyang, Yuanxin Liu, Shicheng Li, Yi Liu, Hao Zhou, Fandong Meng, Jie Zhou and Xu Sun	986
<i>ProcessBench: Identifying Process Errors in Mathematical Reasoning</i>	
Chujie Zheng, Zhenru Zhang, Beichen Zhang, Runji Lin, Keming Lu, Bowen Yu, Dayiheng Liu, Jingren Zhou and Junyang Lin	1009
<i>Model Extrapolation Expedites Alignment</i>	
Chujie Zheng, Ziqi Wang, Heng Ji, Minlie Huang and Nanyun Peng	1025
<i>ATLANTIS: Weak-to-Strong Learning via Importance Sampling</i>	
Yi Liu, Guoyin Wang, Shicheng Li, Feifan Song and Xu Sun	1042
<i>MPVStance: Mitigating Hallucinations in Stance Detection with Multi-Perspective Verification</i>	
ZhaoDan Zhang, Zhao Zhang, Jin Zhang, Hui Xu and Xueqi Cheng	1053
<i>Personality-Guided Code Generation Using Large Language Models</i>	
Yaoqi Guo, Zhenpeng Chen, Jie M. Zhang, Yang Liu and Yun Ma	1068
<i>PsyDT: Using LLMs to Construct the Digital Twin of Psychological Counselor with Personalized Counseling Style for Psychological Counseling</i>	
Haojie Xie, Yirong Chen, Xiaofen Xing, Jingkai Lin and Xiangmin Xu	1081
<i>BIPro: Zero-shot Chinese Poem Generation via Block Inverse Prompting Constrained Generation Framework</i>	
Xu Zou	1116
<i>LongDocURL: a Comprehensive Multimodal Long Document Benchmark Integrating Understanding, Reasoning, and Locating</i>	
Chao Deng, Jiale Yuan, Pi Bu, Peijie Wang, Zhong-Zhi Li, Jian Xu, Xiao-Hui Li, Yuan Gao, Jun Song, Bo Zheng and Cheng-Lin Liu	1135
<i>ObfusLM: Privacy-preserving Language Model Service against Embedding Inversion Attacks</i>	
Yu Lin, Ruining Yang, Yunlong Mao, Qizhi Zhang, Jue Hong, Quanwei Cai, Ye Wu, Huiqi Liu, Zhiyu Chen, Bing Duan and Sheng Zhong	1160
<i>Interlocking-free Selective Rationalization Through Genetic-based Learning</i>	
Federico Ruggeri and Gaetano Signorelli	1175
<i>Re-identification of De-identified Documents with Autoregressive Infilling</i>	
Lucas Georges Gabriel Charpentier and Pierre Lison	1192

<i>Modeling Uncertainty in Composed Image Retrieval via Probabilistic Embeddings</i>	
Haomiao Tang, Jinpeng Wang, Yuang Peng, GuangHao Meng, Ruisheng Luo, Bin Chen, Long Chen, Yaowei Wang and Shu-Tao Xia	1210
<i>Untie the Knots: An Efficient Data Augmentation Strategy for Long-Context Pre-Training in Language Models</i>	
Junfeng Tian, Da Zheng, Yang Chen, Rui Wang, Colin Zhang and Debing Zhang	1223
<i>APPL: A Prompt Programming Language for Harmonious Integration of Programs and Large Language Model Prompts</i>	
Honghua Dong, Qidong Su, Yubo Gao, Zhaoyu Li, Yangjun Ruan, Gennady Pekhimenko, Chris J. Maddison and Xujie Si	1243
<i>Evaluating Lexical Proficiency in Neural Language Models</i>	
Cristiano Ciaccio, Alessio Miaschi and Felice Dell’Orletta	1267
<i>Autoregressive Speech Synthesis without Vector Quantization</i>	
Lingwei Meng, Long Zhou, Shujie Liu, Sanyuan Chen, Bing Han, Shujie HU, Yanqing Liu, Jinyu Li, Sheng Zhao, Xixin Wu, Helen M. Meng and Furu Wei	1287
<i>Cuckoo: An IE Free Rider Hatched by Massive Nutrition in LLM’s Nest</i>	
Letian Peng, Zilong Wang, Feng Yao and Jingbo Shang	1301
<i>FedEx-LoRA: Exact Aggregation for Federated and Efficient Fine-Tuning of Large Language Models</i>	
Raghav Singhal, Kaustubh Ponshe and Praneeth Vepakomma	1316
<i>Measuring Social Biases in Masked Language Models by Proxy of Prediction Quality</i>	
Rahul Zalkikar and Kanchan Chandra	1337
<i>Capturing Author Self Beliefs in Social Media Language</i>	
Siddharth Mangalik, Adithya V Ganesan, Abigail B. Wheeler, Nicholas Kerry, Jeremy D. W. Clifton, H. Schwartz and Ryan L. Boyd	1362
<i>Neural Topic Modeling with Large Language Models in the Loop</i>	
Xiaohao Yang, He Zhao, Weijie Xu, Yuanyuan QI, Jueqing Lu, Dinh Phung and Lan Du . . .	1377
<i>HALoGEN: Fantastic LLM Hallucinations and Where to Find Them</i>	
Abhilasha Ravichander, Shruti Ghela, David Wadden and Yejin Choi	1402
<i>Synergizing LLMs with Global Label Propagation for Multimodal Fake News Detection</i>	
Shuguo Hu, Jun Hu and Huaiwen Zhang	1426
<i>Yes, My LoRD. Guiding Language Model Extraction with Locality Reinforced Distillation</i>	
Zi Liang, Qingqing Ye, Yanyun Wang, Sen Zhang, Yaxin Xiao, RongHua Li, Jianliang Xu and Haibo Hu	1441
<i>Jailbreak Large Vision-Language Models Through Multi-Modal Linkage</i>	
Yu Wang, Xiaofei Zhou, Yichen Wang, Geyuan Zhang and Tianxing He	1466
<i>Wait, that’s not an option: LLMs Robustness with Incorrect Multiple-Choice Options</i>	
Gracjan Góral, Emilia Wiśnios, Piotr Sankowski and Paweł Budzianowski	1495
<i>The Hidden Attention of Mamba Models</i>	
Ameen Ali Ali, Itamar Zimerman and Lior Wolf	1516
<i>KV-Latent: Dimensional-level KV Cache Reduction with Frequency-aware Rotary Positional Embedding</i>	
Shi Luohe, Zuchao Li, Lefei Zhang, Baoyuan Qi, Liu Guoming and Hai Zhao	1535

<i>LEANCODE: Understanding Models Better for Code Simplification of Pre-trained Large Language Models</i>	
Yan Wang, Ling Ding, Tien N Nguyen, Shaohua Wang and Yanan Zheng	1551
<i>MARS: Benchmarking the Metaphysical Reasoning Abilities of Language Models with a Multi-task Evaluation Dataset</i>	
Wei qi Wang and Yang qiu Song	1568
<i>Ask-Before-Detection: Identifying and Mitigating Conformity Bias in LLM-Powered Error Detector for Math Word Problem Solutions</i>	
Hang Li, Tianlong Xu, Kaiqi Yang, Yucheng Chu, Yanling Chen, Yichi Song, Qingsong Wen and Hui Liu	1597
<i>Real-time Factuality Assessment from Adversarial Feedback</i>	
Sanxing Chen, Yukun Huang and Bhuwan Dhingra	1610
<i>Improve Vision Language Model Chain-of-thought Reasoning</i>	
Ruohong Zhang, Bowen Zhang, Yanghao Li, Haotian Zhang, Zhiqing Sun, Zhe Gan, Yinfei Yang, Ruoming Pang and Yiming Yang	1631
<i>On the Mutual Influence of Gender and Occupation in LLM Representations</i>	
Haozhe An, Connor Baumler, Abhilasha Sancheti and Rachel Rudinger	1663
<i>Disentangling Memory and Reasoning Ability in Large Language Models</i>	
Mingyu Jin, Weidi Luo, Sitao Cheng, Xinyi Wang, Wenyue Hua, Ruixiang Tang, William Yang Wang and Yongfeng Zhang	1681
<i>Open-World Attribute Mining for E-Commerce Products with Multimodal Self-Correction Instruction Tuning</i>	
Jiaqi Li, Yanming Li, Xiaoli Shen, Chuanyi Zhang, Guilin Qi and Sheng Bi	1702
<i>Normalized AOPC: Fixing Misleading Faithfulness Metrics for Feature Attributions Explainability</i>	
Joakim Edin, Andreas Geert Motzfeldt, Casper L. Christensen, Tuukka Ruotsalo, Lars Maaløe and Maria Maistro	1715
<i>Takin-VC: Expressive Zero-Shot Voice Conversion via Adaptive Hybrid Content Encoding and Enhanced Timbre Modeling</i>	
Yang Yuguang, Yu Pan, Jixun Yao, Xiang Zhang, Jianhao Ye, Hongbin Zhou, Lei Xie, Lei Ma and Jianjun Zhao	1731
<i>LangSAMP: Language-Script Aware Multilingual Pretraining</i>	
Yihong Liu, Haotian Ye, Chunlan Ma, Mingyang Wang and Hinrich Schuetze	1743
<i>RelationalCoder: Rethinking Complex Tables via Programmatic Relational Transformation</i>	
Haoyu Dong, Yue Hu, Huailiang Peng and Yanan Cao	1771
<i>Algorithmic Fidelity of Large Language Models in Generating Synthetic German Public Opinions: A Case Study</i>	
Bolei Ma, Berk Yoztyurk, Anna-Carolina Haensch, Xinpeng Wang, Markus Herklotz, Frauke Kreuter, Barbara Plank and Matthias Aßenmacher	1785
<i>TUNA: Comprehensive Fine-grained Temporal Understanding Evaluation on Dense Dynamic Videos</i>	
Fanheng Kong, Jingyuan Zhang, Hongzhi Zhang, Shi Feng, Daling Wang, Linhao Yu, Xingguang Ji, Yu Tian, V. W. and Fuzheng Zhang	1810

<i>Self-Instructed Derived Prompt Generation Meets In-Context Learning: Unlocking New Potential of Black-Box LLMs</i>	
Zhuo Li, Yuhao Du, Jinpeng Hu, Xiang Wan and Anningzhe Gao	1840
<i>Binary Classifier Optimization for Large Language Model Alignment</i>	
Seungjae Jung, Gunsoo Han, Daniel Wontae Nam and Kyoung-Woon On	1858
<i>UnSeenTimeQA: Time-Sensitive Question-Answering Beyond LLMs’ Memorization</i>	
Md Nayem Uddin, Amir Saeidi, Divij Handa, Agastya Seth, Tran Cao Son, Eduardo Blanco, Steven Corman and Chitta Baral	1873
<i>From Information to Insight: Leveraging LLMs for Open Aspect-Based Educational Summarization</i>	
Yang Zhong and Diane Litman	1914
<i>AfriMed-QA: A Pan-African, Multi-Specialty, Medical Question-Answering Benchmark Dataset</i>	
Charles Nimo, Tobi Olatunji, Abraham Toluwase Owodunni, Tassallah Abdullahi, Emmanuel Ayodele, Mardhiyah Sanni, Ezinwanne C. Aka, Folafunmi Omofoye, Foutse Yuehgoh, Timothy Fani-ran, Bonaventure F. P. Dossou, Moshood O. Yekini, Jonas Kemp, Katherine A Heller, Jude Chidubem Omeke, Chidi Asuzu MD, Naome A Etori, Aimérou Ndiaye, Ifeoma Okoh, Evans Doe Ocansey, Wendy Kinara, Michael L. Best, Irfan Essa, Stephen Edward Moore, Chris Fourie and Mercy Nyamewaa Asiedu	1948
<i>Root Defense Strategies: Ensuring Safety of LLM at the Decoding Level</i>	
Xinyi Zeng, Yuying Shang, Jiawei Chen, Jingyuan Zhang and Yu Tian	1974
<i>In-the-wild Audio Spatialization with Flexible Text-guided Localization</i>	
Tianrui Pan, Jie Liu, Zewen Huang, Jie Tang and Gangshan Wu	1989
<i>L4Q: Parameter Efficient Quantization-Aware Fine-Tuning on Large Language Models</i>	
Hyesung Jeon, Yulhwa Kim and Jae-Joon Kim	2002
<i>Second Language (Arabic) Acquisition of LLMs via Progressive Vocabulary Expansion</i>	
Jianqing Zhu, Huang Huang, Zhihang Lin, Juhao Liang, Zhengyang Tang, Khalid Almubarak, Mosen Alharthi, Bang An, Juncal He, Xiangbo Wu, Fei Yu, Junying Chen, MA Zhuoheng, Yuhao Du, He Zhang, Saied Alshahrani, Emad A. Alghamdi, Lian Zhang, Ruoyu Sun, Haizhou Li, Benyou Wang and Jinchao Xu	2025
<i>What Really Matters in Many-Shot Attacks? An Empirical Study of Long-Context Vulnerabilities in LLMs</i>	
Sangyeop Kim, Yohan Lee, Yongwoo Song and Kimin Lee	2043
<i>ECERC: Evidence-Cause Attention Network for Multi-Modal Emotion Recognition in Conversation</i>	
Tao Zhang and Zhenhua Tan	2064
<i>CompileAgent: Automated Real-World Repo-Level Compilation with Tool-Integrated LLM-based Agent System</i>	
Li Hu, Guoqiang Chen, Xiuwei Shang, Shaoyin Cheng, Benlong Wu, LiGangyang LiGangyang, Xu Zhu, Weiming Zhang and Nenghai Yu	2078
<i>Beyond Demographics: Fine-tuning Large Language Models to Predict Individuals’ Subjective Text Perceptions</i>	
Matthias Orlikowski, Jiaxin Pei, Paul Röttger, Philipp Cimiano, David Jurgens and Dirk Hovy	2092
<i>Exploring Forgetting in Large Language Model Pre-Training</i>	
Chonghua Liao, Ruobing Xie, Xingwu Sun, Haowen Sun and Zhanhui Kang	2112

<i>Bias in the Mirror : Are LLMs opinions robust to their own adversarial attacks</i> Virgile Rennard, Christos Xypolopoulos and Michalis Vazirgiannis	2128
<i>AndroidLab: Training and Systematic Benchmarking of Android Autonomous Agents</i> Yifan Xu, Xiao Liu, Xueqiao Sun, Siyi Cheng, Hao Yu, Hanyu Lai, Shudan Zhang, Dan Zhang, Jie Tang and Yuxiao Dong	2144
<i>Modular Sentence Encoders: Separating Language Specialization from Cross-Lingual Alignment</i> Yongxin Huang, Kexin Wang, Goran Glavaš and Iryna Gurevych	2167
<i>Multimodal Transformers are Hierarchical Modal-wise Heterogeneous Graphs</i> Yijie Jin, Junjie Peng, Xuanchao Lin, Haochen Yuan, Lan Wang and Cangzhi Zheng	2188
<i>Have We Designed Generalizable Structural Knowledge Promptings? Systematic Evaluation and Re- thinking</i> Yichi Zhang, Zhuo Chen, Lingbing Guo, Yajing Xu, Shaokai Chen, Mengshu Sun, Binbin Hu, Zhiqiang Zhang, Lei Liang, Wen Zhang and Huajun Chen	2210
<i>LLäMmlein: Transparent, Compact and Competitive German-Only Language Models from Scratch</i> Jan Pfister, Julia Wunderle and Andreas Hotho	2227
<i>Speaking Beyond Language: A Large-Scale Multimodal Dataset for Learning Nonverbal Cues from Video-Grounded Dialogues</i> Youngmin Kim, Jiwan Chung, Jisoo Kim, Sunghyun Lee, Sangkyu Lee, Junhyeok Kim, Cheol- jong Yang and Youngjae Yu	2247
<i>How Much Do Encoder Models Know About Word Senses?</i> Simone Teglia, Simone Tedeschi and Roberto Navigli	2266
<i>When Backdoors Speak: Understanding LLM Backdoor Attacks Through Model-Generated Explana- tions</i> Huaizhi Ge, Yiming Li, Qifan Wang, Yongfeng Zhang and Ruixiang Tang	2278
<i>HateDay: Insights from a Global Hate Speech Dataset Representative of a Day on Twitter</i> Manuel Tonneau, Diyi Liu, Niyati Malhotra, Scott A. Hale, Samuel Fraiberger, Victor Orozco- Olvera and Paul Röttger	2297
<i>LegalAgentBench: Evaluating LLM Agents in Legal Domain</i> Haitao Li, Junjie Chen, Jingli Yang, Qingyao Ai, Wei Jia, Youfeng Liu, Kai Lin, Yueyue WU, Guozhi Yuan, Yiran HU, Wuyue Wang, Yiqun Liu and Minlie Huang	2322
<i>Inference Compute-Optimal Video Vision Language Models</i> Peiqi Wang, ShengYun Peng, Xuewen Zhang, Hanchao Yu, Yibo Yang, Lifu Huang, Fujun Liu and Qifan Wang	2345
<i>Steering into New Embedding Spaces: Analyzing Cross-Lingual Alignment Induced by Model Interven- tions in Multilingual Language Models</i> Anirudh Sundar, Sinead Williamson, Katherine Metcalf, Barry-John Theobald, Skyler Seto and Masha Fedzechkina	2375
<i>Digital Gatekeepers: Google’s Role in Curating Hashtags and Subreddits</i> Amrit Poudel, Yifan Ding, Tim Weninger and Jürgen Pfeffer	2402
<i>Behind Closed Words: Creating and Investigating the forePLay Annotated Dataset for Polish Erotic Discourse</i> Anna Kołos, Katarzyna Lorenc, Emilia Wiśnios and Agnieszka Karlińska	2416

<i>Assessment and manipulation of latent constructs in pre-trained language models using psychometric scales</i>	
Maor Reuben, Ortal Slobodin, Idan-Chaim Cohen, Aviad Elyashar, Orna Braun-Lewensohn, Odeya Cohen and Rami Puzis	2433
<i>Did Translation Models Get More Robust Without Anyone Even Noticing?</i>	
Ben Peters and Andre Martins	2445
<i>Nemotron-CC: Transforming Common Crawl into a Refined Long-Horizon Pretraining Dataset</i>	
Dan SU, Kezhi Kong, Ying Lin, Joseph Jennings, Brandon Norick, Markus Kliegl, Mostofa Patwary, Mohammad Shoeybi and Bryan Catanzaro	2459
<i>Hierarchical Level-Wise News Article Clustering via Multilingual Matryoshka Embeddings</i>	
Hans William Alexander Hanley and Zakir Durumeric	2476
<i>Contrastive Perplexity for Controlled Generation: An Application in Detoxifying Large Language Models</i>	
Tassilo Klein and Moin Nabi	2493
<i>INVESTORBENCH: A Benchmark for Financial Decision-Making Tasks with LLM-based Agent</i>	
Haohang Li, Yupeng Cao, Yangyang Yu, Shashidhar Reddy Javaji, Zhiyang Deng, Yueru He, Yuechen Jiang, Zining Zhu, K.p. Subbalakshmi, Jimin Huang, Lingfei Qian, Xueqing Peng, Jordan W. Suchow and Qianqian Xie	2509
<i>Smarter, Better, Faster, Longer: A Modern Bidirectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference</i>	
Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Griffin Thomas Adams, Jeremy Howard and Iacopo Poli	2526
<i>Gender Inclusivity Fairness Index (GIFI): A Multilevel Framework for Evaluating Gender Diversity in Large Language Models</i>	
Zhengyang Shan, Emily Diana and Jiawei Zhou	2548
<i>D.Va: Validate Your Demonstration First Before You Use It</i>	
Qi Zhang, Zhiqing Xiao, Ruixuan Xiao, Lirong Gao and Junbo Zhao	2580
<i>Are Any-to-Any Models More Consistent Across Modality Transfers Than Specialists?</i>	
Jiwan Chung, Janghan Yoon, Junhyeong Park, Sangeyl Lee, Joowon Yang, Sooyeon Park and Youngjae Yu	2595
<i>MAIN-RAG: Multi-Agent Filtering Retrieval-Augmented Generation</i>	
Chia-Yuan Chang, Zhimeng Jiang, Vineeth Rakesh, Menghai Pan, Chin-Chia Michael Yeh, Guanchu Wang, Mingzhi Hu, Zhichao Xu, Yan Zheng, Mahashweta Das and Na Zou	2607
<i>Unraveling the Mechanics of Learning-Based Demonstration Selection for In-Context Learning</i>	
Hui Liu, Wenya Wang, Hao Sun, Chris Xing Tian, Chenqi Kong, Xin Dong and Haoliang Li	2623
<i>Direct Prompt Optimization with Continuous Representations</i>	
Yangkun Wang, Zihan Wang and Jingbo Shang	2642
<i>uMedSum: A Unified Framework for Clinical Abstractive Summarization</i>	
Aishik Nagar, Yutong Liu, Andy T. Liu, Viktor Schlegel, Vijay Prakash Dwivedi, Arun-Kumar Kaliya-Perumal, Guna Pratheep Kalanchiam, Yili Tang and Robby T. Tan	2653

<i>GigaSpeech 2: An Evolving, Large-Scale and Multi-domain ASR Corpus for Low-Resource Languages with Automated Crawling, Transcription and Refinement</i>	
Yifan Yang, Zheshu Song, Jianheng Zhuo, Mingyu Cui, Jinpeng Li, Bo Yang, Yexing Du, Ziyang Ma, Xunying Liu, Ziyuan Wang, Ke Li, Shuai Fan, Kai Yu, Wei-Qiang Zhang, Guoguo Chen and Xie Chen.....	2673
<i>Context-Aware Sentiment Forecasting via LLM-based Multi-Perspective Role-Playing Agents</i>	
Fanhang Man, Huandong Wang, Jianjie Fang, Zhaoyi Deng, Baining Zhao, Xinlei Chen and Yong Li.....	2687
<i>TARGA: Targeted Synthetic Data Generation for Practical Reasoning over Structured Data</i>	
Xiang Huang, Jiayu Shen, Shanshan Huang, Sitao Cheng, Xiaxia Wang and Yuzhong Qu ..	2704
<i>AndroidGen: Building an Android Language Agent under Data Scarcity</i>	
Hanyu Lai, Junjie Gao, Xiao Liu, Yifan Xu, Shudan Zhang, Yuxiao Dong and Jie Tang	2727
<i>Prompt Candidates, then Distill: A Teacher-Student Framework for LLM-driven Data Annotation</i>	
Mingxuan Xia, Haobo Wang, Yixuan Li, Zewei Yu, Jindong Wang, Junbo Zhao and Runze Wu	2750
<i>A Survey of Post-Training Scaling in Large Language Models</i>	
Hanyu Lai, Xiao Liu, Junjie Gao, Jiale Cheng, Zehan Qi, Yifan Xu, Shuntian Yao, Dan Zhang, Jinhua Du, Zhenyu Hou, Xin Lv, Minlie Huang, Yuxiao Dong and Jie Tang	2771
<i>Position-aware Automatic Circuit Discovery</i>	
Tal Haklay, Hadas Orgad, David Bau, Aaron Mueller and Yonatan Belinkov	2792
<i>HyperFM: Fact-Centric Multimodal Fusion for Link Prediction over Hyper-Relational Knowledge Graphs</i>	
Yuhuan Lu, Weijian Yu, Xin Jing and Dingqi Yang	2818
<i>Centurio: On Drivers of Multilingual Ability of Large Vision-Language Model</i>	
Gregor Geigle, Florian Schneider, Carolin Holtermann, Chris Biemann, Radu Timofte, Anne Lauscher and Goran Glavaš	2831
<i>Less for More: Enhanced Feedback-aligned Mixed LLMs for Molecule Caption Generation and Fine-Grained NLI Evaluation</i>	
Dimitris Gkoumas and Maria Liakata	2882
<i>Ensemble Watermarks for Large Language Models</i>	
Georg Niess and Roman Kern	2903
<i>ConInstruction: Universal Jailbreaking of Multimodal Large Language Models via Non-Textual Modalities</i>	
Jiahui Geng, Thy Thy Tran, Preslav Nakov and Iryna Gurevych	2917
<i>TRACT: Regression-Aware Fine-tuning Meets Chain-of-Thought Reasoning for LLM-as-a-Judge</i>	
Cheng-Han Chiang, Hung-yi Lee and Michal Lukasik	2934
<i>DioR: Adaptive Cognitive Detection and Contextual Retrieval Optimization for Dynamic Retrieval-Augmented Generation</i>	
Hanghai Guo, Jia Zhu, Shimin Di, Weijie Shi, Zhangze Chen and Jiajie Xu	2953
<i>Unveiling the Power of Source: Source-based Minimum Bayes Risk Decoding for Neural Machine Translation</i>	
Boxuan Lyu, Hidetaka Kamigaito, Kotaro Funakoshi and Manabu Okumura	2976

<i>ToolHop: A Query-Driven Benchmark for Evaluating Large Language Models in Multi-Hop Tool Use</i>	
Junjie Ye, Zhengyin Du, Xuesong Yao, Weijian Lin, Yufei Xu, Zehui Chen, Zaiyuan Wang, Sining Zhu, Zhiheng Xi, Siyu Yuan, Tao Gui, Qi Zhang, Xuanjing Huang and Jiecao Chen	2995
<i>Mixture of insightful Experts (MoTE): The Synergy of Reasoning Chains and Expert Mixtures in Self-Alignment</i>	
Zhili Liu, Yunhao Gou, Kai Chen, Lanqing Hong, Jiahui Gao, Fei Mi, Yu Zhang, Zhenguo Li, Xin Jiang, Qun Liu and James Kwok	3022
<i>MAPS: Motivation-Aware Personalized Search via LLM-Driven Consultation Alignment</i>	
Weicong Qin, Yi Xu, Weijie Yu, Chenglei Shen, Ming He, Jianping Fan, Xiao Zhang and Jun Xu	3039
<i>Aristotle: Mastering Logical Reasoning with A Logic-Complete Decompose-Search-Resolve Framework</i>	
Jundong Xu, Hao Fei, Meng Luo, Qian Liu, Liangming Pan, William Yang Wang, Preslav Nakov, Mong-Li Lee and Wynne Hsu	3052
<i>LADM: Long-context Training Data Selection with Attention-based Dependency Measurement for LLMs</i>	
Jianghao Chen, Junhong Wu, Yangyifan Xu and Jiajun Zhang	3076
<i>Iron Sharpens Iron: Defending Against Attacks in Machine-Generated Text Detection with Adversarial Training</i>	
Yuanfan Li, Zhaohan Zhang, Chengzhengxu Li, Chao Shen and Xiaoming Liu	3091
<i>Cultural Learning-Based Culture Adaptation of Language Models</i>	
Chen Cecilia Liu, Anna Korhonen and Iryna Gurevych	3114
<i>A-TASC: Asian TED-Based Automatic Subtitling Corpus</i>	
Yuhan Zhou and Naoki Yoshinaga	3135
<i>Refuse Whenever You Feel Unsafe: Improving Safety in LLMs via Decoupled Refusal Training</i>	
Youliang Yuan, Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Jiahao Xu, Tian Liang, Pinjia He and Zhaopeng Tu	3149
<i>Token Prepending: A Training-Free Approach for Eliciting Better Sentence Embeddings from LLMs</i>	
Yuchen Fu, Zifeng Cheng, Zhiwei Jiang, Zhonghui Wang, Yafeng Yin, Zhengliang Li and Qing Gu	3168
<i>No Questions are Stupid, but some are Poorly Posed: Understanding Poorly-Posed Information-Seeking Questions</i>	
Neha Srikanth, Rachel Rudinger and Jordan Lee Boyd-Graber	3182
<i>Understanding Common Ground Misalignment in Goal-Oriented Dialog: A Case-Study with Ubuntu Chat Logs</i>	
Rupak Sarkar, Neha Srikanth, Taylor Pellegrin, Rachel Rudinger, Claire Bonial and Philip Resnik	3200
<i>Addressing Blind Guessing: Calibration of Selection Bias in Multiple-Choice Question Answering by Video Language Models</i>	
Olga Loginova, Oleksandr Bezrukov, Ravi Shekhar and Alexey Kravets	3216
<i>Towards Reward Fairness in RLHF: From a Resource Allocation Perspective</i>	
Sheng Ouyang, Yulan Hu, Ge Chen, Qingyang Li, Fuzheng Zhang and Yong Liu	3247

<i>Taming LLMs with Gradient Grouping</i>	
Siyuan Li, Juanxi Tian, Zedong Wang, Xin Jin, Zicheng Liu, Wentao Zhang and Dan Xu . . .	3260
<i>LazyReview: A Dataset for Uncovering Lazy Thinking in NLP Peer Reviews</i>	
Sukannya Purkayastha, Zhuang Li, Anne Lauscher, Lizhen Qu and Iryna Gurevych	3280
<i>Revisiting Common Assumptions about Arabic Dialects in NLP</i>	
Amr Keleg, Sharon Goldwater and Walid Magdy	3309
<i>Retrieve to Explain: Evidence-driven Predictions for Explainable Drug Target Identification</i>	
Ravi Patel, Angus Brayne, Rogier Hintzen, Daniel Jaroslawicz, Georgiana Neculae and Dane S. Corneil	3328
<i>Whose Boat Does it Float? Improving Personalization in Preference Tuning via Inferred User Personas</i>	
Nishant Balepur, Vishakh Padmakumar, Fumeng Yang, Shi Feng, Rachel Rudinger and Jordan Lee Boyd-Graber	3371
<i>Which of These Best Describes Multiple Choice Evaluation with LLMs? A) Forced B) Flawed C) Fixable D) All of the Above</i>	
Nishant Balepur, Rachel Rudinger and Jordan Lee Boyd-Graber	3394
<i>Detection of Human and Machine-Authored Fake News in Urdu</i>	
Muhammad Zain Ali, Yuxia Wang, Bernhard Pfahringer and Tony C Smith	3419
<i>An Efficient Task-Oriented Dialogue Policy: Evolutionary Reinforcement Learning Injected by Elite Individuals</i>	
Yangyang Zhao, Ben Niu, Libo Qin and Shihan Wang	3429
<i>SR-LLM: Rethinking the Structured Representation in Large Language Model</i>	
Jiahuan Zhang, Tianheng Wang, Ziyi Huang, Yulong Wu, Hanqing WU, DongbaiChen Dongbai-Chen, Linfeng Song, Yue Zhang, Guozheng Rao and Kaicheng Yu	3443
<i>Taming Language Models for Text-attributed Graph Learning with Decoupled Aggregation</i>	
Chuang Zhou, Zhu Wang, Shengyuan Chen, Jiahe Du, Qiyuan Zheng, Zhaozhuo Xu and Xiao Huang	3463
<i>Contrastive Prompting Enhances Sentence Embeddings in LLMs through Inference-Time Steering</i>	
Zifeng Cheng, Zhonghui Wang, Yuchen Fu, Zhiwei Jiang, Yafeng Yin, Cong Wang and Qing Gu	3475
<i>Cracking the Code of Hallucination in LVLMS with Vision-aware Head Divergence</i>	
Jinghan He, Kuan Zhu, Haiyun Guo, Junfeng Fang, Zhenglin Hua, Yuheng Jia, Ming Tang, Tat-Seng Chua and Jinqiao Wang	3488
<i>Hierarchical Document Refinement for Long-context Retrieval-augmented Generation</i>	
Jiajie Jin, Xiaoxi Li, Guanting Dong, Yuyao Zhang, Yutao Zhu, Yongkang Wu, Zhonghua Li, YE QI and Zhicheng Dou	3502
<i>Comparing Moral Values in Western English-speaking societies and LLMs with Word Associations</i>	
Chaoyi Xiang, Chunhua Liu, Simon De Deyne and Lea Frermann	3521
<i>TEACH: A Contrastive Knowledge Adaptive Distillation Framework for Classical Chinese Understanding</i>	
Yuting Wei, Qi Meng, Yuanxing Xu and Bin Wu	3537
<i>RAG-Critic: Leveraging Automated Critic-Guided Agentic Workflow for Retrieval Augmented Generation</i>	
Guanting Dong, Jiajie Jin, Xiaoxi Li, Yutao Zhu, Zhicheng Dou and Ji-Rong Wen	3551

<i>Progressive Multimodal Reasoning via Active Retrieval</i>	
Guanting Dong, Chenghao Zhang, Mengjie Deng, Yutao Zhu, Zhicheng Dou and Ji-Rong Wen	
	3579
<i>Pre-training Distillation for Large Language Models: A Design Space Exploration</i>	
Hao Peng, Xin Lv, Yushi Bai, Zijun Yao, Jiajie Zhang, Lei Hou and Juanzi Li	3603
<i>Teaching Vision-Language Models to Ask: Resolving Ambiguity in Visual Questions</i>	
Pu Jian, Donglei Yu, Wen Yang, Shuo Ren and Jiajun Zhang	3619
<i>LongBench v2: Towards Deeper Understanding and Reasoning on Realistic Long-context Multitasks</i>	
Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, Jie Tang and Juanzi Li	3639
<i>Battling against Tough Resister: Strategy Planning with Adversarial Game for Non-collaborative Dialogues</i>	
Haiyang Wang, Zhiliang Tian, Yuchen Pan, Xin Song, Xin Niu, Minlie Huang and Bin Zhou	3665
<i>Cross-model Transferability among Large Language Models on the Platonic Representations of Concepts</i>	
Youcheng Huang, Chen Huang, Duanyu Feng, Wenqiang Lei and Jiancheng Lv	3686
<i>FoldMoE: Efficient Long Sequence MoE Training via Attention-MoE Pipelining</i>	
Guichao Zhu, Lintian Lei, Yuhao Qing, Yichao Fu, Fanxin Li, Dong Huang, Zekai Sun and Heming Cui	3705
<i>LongReward: Improving Long-context Large Language Models with AI Feedback</i>	
Jiajie Zhang, Zhongni Hou, Xin Lv, Shulin Cao, Zhenyu Hou, Yilin Niu, Lei Hou, Yuxiao Dong, Ling Feng and Juanzi Li	3718
<i>Influences on LLM Calibration: A Study of Response Agreement, Loss Functions, and Prompt Styles</i>	
Yuxi Xia, Pedro Henrique Luz De Araujo, Klim Zaporozhets and Benjamin Roth	3740
<i>UTBoost: Rigorous Evaluation of Coding Agents on SWE-Bench</i>	
Boxi Yu, Yuxuan Zhu, Pinjia He and Daniel Kang	3762
<i>Towards Better Evaluation for Generated Patent Claims</i>	
Lekang Jiang, Pascal A. Scherz and Stefan Goetz	3775
<i>Fine-Tuning on Diverse Reasoning Chains Drives Within-Inference CoT Refinement in LLMs</i>	
Haritz Puerto, Tilek Chubakov, Xiaodan Zhu, Harish Tayyar Madabushi and Iryna Gurevych	3789
<i>Establishing Trustworthy LLM Evaluation via Shortcut Neuron Analysis</i>	
Kejian Zhu, Shangqing Tu, Zhuoran Jin, Lei Hou, Juanzi Li and Jun Zhao	3809
<i>Do Large Language Models have an English Accent? Evaluating and Improving the Naturalness of Multilingual LLMs</i>	
Yanzhu Guo, Simone Conia, Zelin Zhou, Min Li, Saloni Potdar and Henry Xiao	3823
<i>Enhancing Character-Level Understanding in LLMs through Token Internal Structure Learning</i>	
Zhu Xu, Zhiqiang Zhao, Zihan Zhang, Yuchi Liu, Quanwei Shen, Fei Liu, Yu Kuang, Jian He and Conglin Liu	3839
<i>Conformity in Large Language Models</i>	
Xiaochen Zhu, Caiqi Zhang, Tom Stafford, Nigel Collier and Andreas Vlachos	3854

<i>Interpret and Improve In-Context Learning via the Lens of Input-Label Mappings</i>	
Chenghao Sun, Zhen Huang, Yonggang Zhang, Le Lu, Houqiang Li, Xinmei Tian, Xu Shen and Jieping Ye	3873
<i>Positional Overload: Positional Debiasing and Context Window Extension for Large Language Models using Set Encoding</i>	
Lukas Kinder, Lukas Edman, Alexander Fraser and Tobias Käfer	3896
<i>FR-Spec: Accelerating Large-Vocabulary Language Models via Frequency-Ranked Speculative Sampling</i>	
Weilin Zhao, Tengyu Pan, Xu Han, Yudi Zhang, Sun Ao, Yuxiang Huang, Kaihuo Zhang, Weilun Zhao, Yuxuan Li, Jie Zhou, Hao Zhou, Jianyong Wang, Maosong Sun and Zhiyuan Liu	3909
<i>VReST: Enhancing Reasoning in Large Vision-Language Models through Tree Search and Self-Reward Mechanism</i>	
Congzhi Zhang, Jiawei Peng, Zhenglin Wang, Yilong Lai, Haowen Sun, Heng Chang, Fei Ma and Weijiang Yu	3922
<i>Past Meets Present: Creating Historical Analogy with Large Language Models</i>	
Nianqi Li, Siyu Yuan, Jiangjie Chen, Jiaqing Liang, Feng Wei, Zujie Liang, Deqing Yang and Yanghua Xiao	3942
<i>Meta-Reflection: A Feedback-Free Reflection Learning Framework</i>	
Yaoke Wang, Yun Zhu, Xintong Bao, Xintong Bao, Wenqiao Zhang, Suyang Dai, Kehan Chen, Wenqiang Li, Gang Huang, Siliang Tang and Yueting Zhuang	3958
<i>Read it in Two Steps: Translating Extremely Low-Resource Languages with Code-Augmented Grammar Books</i>	
Chen Zhang, Jiuheng Lin, Xiao Liu, Zekai Zhang and Yansong Feng	3977
<i>Confidence v.s. Critique: A Decomposition of Self-Correction Capability for LLMs</i>	
Zhe Yang, Yichang Zhang, Yudong Wang, Ziyao Xu, Junyang Lin and Zhifang Sui	3998
<i>Automating Legal Interpretation with LLMs: Retrieval, Generation, and Evaluation</i>	
Kangcheng Luo, Quzhe Huang, Cong Jiang and Yansong Feng	4015
<i>Visual Evidence Prompting Mitigates Hallucinations in Large Vision-Language Models</i>	
Wei Li, Zhen Huang, Houqiang Li, Le Lu, Yang Lu, Xinmei Tian, Xu Shen and Jieping Ye ..	4048
<i>Leveraging Dual Process Theory in Language Agent Framework for Real-time Simultaneous Human-AI Collaboration</i>	
Shao Zhang, Xihuai Wang, Wenhao Zhang, Chaoran Li, Junru Song, Tingyu Li, Lin Qiu, Xuezhi Cao, Xunliang Cai, Wen Yao, Weinan Zhang, Xinbing Wang and Ying Wen	4081
<i>TokAlign: Efficient Vocabulary Adaptation via Token Alignment</i>	
Chong Li, Jiajun Zhang and Chengqing Zong	4109
<i>AdaEdit: Advancing Continuous Knowledge Editing For Large Language Models</i>	
Qi Li and Xiaowen Chu	4127
<i>The Impact of Token Granularity on the Predictive Power of Language Model Surprisal</i>	
Byung-Doh Oh and William Schuler	4150
<i>Segment-Level Diffusion: A Framework for Controllable Long-Form Generation with Diffusion Language Models</i>	
Xiaochen Zhu, Georgi Karadzhov, Chenxi Whitehouse and Andreas Vlachos	4163

<i>BELLE: A Bi-Level Multi-Agent Reasoning Framework for Multi-Hop Question Answering</i> Taolin Zhang, Dongyang Li, Qizhou Chen, Chengyu Wang and Xiaofeng He	4184
<i>Dynamic and Generalizable Process Reward Modeling</i> Zhangyue Yin, Qiushi Sun, Zhiyuan Zeng, Qinyuan Cheng, Xipeng Qiu and Xuanjing Huang	4203
<i>AdamMeme: Adaptively Probe the Reasoning Capacity of Multimodal Large Language Models on Harmfulness</i> Zixin Chen, Hongzhan Lin, Kaixin Li, Ziyang Luo, Zhen Ye, Guang Chen, Zhiyong Huang and Jing Ma	4234
<i>Towards Text-Image Interleaved Retrieval</i> Xin Zhang, Ziqi Dai, Yongqi Li, Yanzhao Zhang, Dingkun Long, Pengjun Xie, Meishan Zhang, Jun Yu, Wenjie Li and Min Zhang	4254
<i>Large Margin Representation Learning for Robust Cross-lingual Named Entity Recognition</i> Guangcheng Zhu, Ruixuan Xiao, Haobo Wang, Zhen Zhu, Gengyu Lyu and Junbo Zhao ...	4270
<i>An Efficient and Precise Training Data Construction Framework for Process-supervised Reward Model in Mathematical Reasoning</i> Wei Sun, Qianlong Du, Fuwei Cui and Jiajun Zhang	4292
<i>QAEncoder: Towards Aligned Representation Learning in Question Answering Systems</i> Zhengren Wang, Qinhan Yu, Shida Wei, Zhiyu li, Feiyu Xiong, Xiaoxing Wang, Simin Niu, Hao Liang and Wentao Zhang	4306
<i>Game Development as Human-LLM Interaction</i> Jiale Hong, Hongqiu Wu and Hai Zhao	4333
<i>Can LLMs Simulate L2-English Dialogue? An Information-Theoretic Analysis of L1-Dependent Biases</i> Rena Wei Gao, Xuotong Wu, Tatsuki Kuribayashi, Mingrui Ye, Siya Qi, Carsten Roever, Yuanxing Liu, Zheng Yuan and Jey Han Lau	4355
<i>DeepSolution: Boosting Complex Engineering Solution Design via Tree-based Exploration and Bi-point Thinking</i> Zhuoqun Li, Haiyang Yu, Xuanang Chen, Hongyu Lin, Yaojie Lu, Fei Huang, Xianpei Han, Yongbin Li and Le Sun	4380
<i>SurveyPilot: an Agentic Framework for Automated Human Opinion Collection from Social Media</i> Viet Thanh Pham, Lizhen Qu, Zhuang Li, Suraj Sharma and Gholamreza Haffari	4397
<i>Sharper and Faster mean Better: Towards More Efficient Vision-Language Model for Hour-scale Long Video Understanding</i> Daoze Zhang, Yuze Zhao, Jintao Huang and Yingda Chen	4423
<i>Auto-Arena: Automating LLM Evaluations with Agent Peer Battles and Committee Discussions</i> Ruo Chen Zhao, Wenxuan Zhang, Yew Ken Chia, Weiwen Xu, Deli Zhao and Lidong Bing .	4440
<i>How Humans and LLMs Organize Conceptual Knowledge: Exploring Subordinate Categories in Italian</i> Andrea Pedrotti, Giulia Rambelli, Caterina Villani and Marianna Bolognesi	4464
<i>PTQ1.61: Push the Real Limit of Extremely Low-Bit Post-Training Quantization Methods for Large Language Models</i> Jiaqi Zhao, Miao Zhang, Ming Wang, Yuzhang Shang, Kaihao Zhang, Weili Guan, Yaowei Wang and Min Zhang	4483

<i>ProtoLens: Advancing Prototype Learning for Fine-Grained Interpretability in Text Classification</i> Bowen Wei and Ziwei Zhu	4503
<i>Fine-grained Video Dubbing Duration Alignment with Segment Supervised Preference Optimization</i> Chaoqun Cui, Liangbin Huang, Shijing Wang, Zhe Tong, Zhaolong Huang, Xiao Zeng and Xiaofeng Liu	4524
<i>Sparse Latents Steer Retrieval-Augmented Generation</i> Chunlei Xin, Shuheng Zhou, Huijia Zhu, Weiqiang Wang, Xuanang Chen, Xinyan Guan, Yaojie Lu, Hongyu Lin, Xianpei Han and Le Sun	4547
<i>Unveiling Language-Specific Features in Large Language Models via Sparse Autoencoders</i> Boyi Deng, Yu Wan, Baosong Yang, Yidan Zhang and Fuli Feng	4563
<i>SafeRAG: Benchmarking Security in Retrieval-Augmented Generation of Large Language Model</i> Xun Liang, Simin Niu, Zhiyu li, Sensen Zhang, Hanyu Wang, Feiyu Xiong, Zhaoxin Fan, Bo Tang, Jihao Zhao, Jiawei Yang, Shichao Song and Mengwei Wang	4609
<i>AnRe: Analogical Replay for Temporal Knowledge Graph Forecasting</i> Guo Tang, Zheng Chu, Wenxiang Zheng, Junjia Xiang, Yizhuo Li, Weihao Zhang, Ming Liu and Bing Qin	4632
<i>Revisiting the Test-Time Scaling of o1-like Models: Do they Truly Possess Test-Time Scaling Capabilities?</i> Zhiyuan Zeng, Qinyuan Cheng, Zhangyue Yin, Yunhua Zhou and Xipeng Qiu	4651
<i>Text is All You Need: LLM-enhanced Incremental Social Event Detection</i> Zitai Qiu, Congbo Ma, Jia Wu and Jian Yang	4666
<i>Multimodal Pragmatic Jailbreak on Text-to-image Models</i> Tong Liu, Zhixin Lai, Jiawen Wang, Gengyuan Zhang, Shuo Chen, Philip Torr, Vera Demberg, Volker Tresp and Jindong Gu	4681
<i>Principled Understanding of Generalization for Generative Transformer Models in Arithmetic Reasoning Tasks</i> Xingcheng Xu, Zibo Zhao, Haipeng Zhang and Yanqing Yang	4721
<i>Discourse Relation-Enhanced Neural Coherence Modeling</i> Wei Liu and Michael Strube	4748
<i>Benchmarking Open-ended Audio Dialogue Understanding for Large Audio-Language Models</i> Kuofeng Gao, Shu-Tao Xia, Ke Xu, Philip Torr and Jindong Gu	4763
<i>from Benign import Toxic: Jailbreaking the Language Model via Adversarial Metaphors</i> Yu Yan, Sheng Sun, Zenghao Duan, Teli Liu, Min Liu, Zhiyi Yin, LeiJingyu LeiJingyu and Qi Li	4785
<i>ShifCon: Enhancing Non-Dominant Language Capabilities with a Shift-based Multilingual Contrastive Framework</i> Hengyuan Zhang, Chenming Shang, Sizhe Wang, Dongdong Zhang, Yiyao Yu, Feng Yao, Renliang Sun, Yujiu Yang and Furu Wei	4818
<i>MorphMark: Flexible Adaptive Watermarking for Large Language Models</i> Zongqi Wang, Tianle Gu, Baoyuan Wu and Yujiu Yang	4842

<i>A Silver Bullet or a Compromise for Full Attention? A Comprehensive Study of Gist Token-based Context Compression</i>	
Chenlong Deng, Zhisong Zhang, Kelong Mao, Shuaiyi Li, Xinting Huang, Dong Yu and Zhicheng Dou	4861
<i>On the Limit of Language Models as Planning Formalizers</i>	
Cassie Huang and Li Zhang	4880
<i>Learning to Generate Structured Output with Schema Reinforcement Learning</i>	
Yaxi Lu, Haolun Li, Xin Cong, Zhong Zhang, Yesai Wu, Yankai Lin, Zhiyuan Liu, Fangming Liu and Maosong Sun	4905
<i>Enhancing Unsupervised Sentence Embeddings via Knowledge-Driven Data Augmentation and Gaussian-Decayed Contrastive Learning</i>	
Peichao Lai, Zhengfeng Zhang, Wentao Zhang, Fangcheng Fu and Bin Cui	4919
<i>Improve Safety Training of Large Language Models with Safety-Critical Singular Vectors Localization</i>	
Peijian Gu, Quan Wang and Zhendong Mao	4941
<i>WarriorCoder: Learning from Expert Battles to Augment Code Large Language Models</i>	
Huawen Feng, Pu Zhao, Qingfeng Sun, Can Xu, Fangkai Yang, Lu Wang, Qianli Ma, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang and Qi Zhang	4955
<i>A Triple-View Framework for Fine-Grained Emotion Classification with Clustering-Guided Contrastive Learning</i>	
Junqing Gong, Binhan Yang and Wei Shen	4970
<i>Quantification of Large Language Model Distillation</i>	
Sunbowen Lee, Junting Zhou, Chang Ao, Kaige Li, Xeron Du, Sirui He, Haihong Wu, Tianci Liu, Jiaheng Liu, Hamid Alinejad-Rokny, Min Yang, Yitao Liang, Zhoufutu Wen and Shiwen Ni	4985
<i>Demons in the Detail: On Implementing Load Balancing Loss for Training Specialized Mixture-of-Expert Models</i>	
Zihan Qiu, Zeyu Huang, Bo Zheng, Kaiyue Wen, Zekun Wang, Rui Men, Ivan Titov, Dayiheng Liu, Jingren Zhou and Junyang Lin	5005
<i>Pandora’s Box or Aladdin’s Lamp: A Comprehensive Analysis Revealing the Role of RAG Noise in Large Language Models</i>	
Jinyang Wu, Shuai Zhang, Feihu Che, Mingkuan Feng, Pengpeng Shao and Jianhua Tao ...	5019
<i>Stepwise Reasoning Disruption Attack of LLMs</i>	
Jingyu Peng, Maolin Wang, Xiangyu Zhao, Kai Zhang, Wanyu Wang, Pengyue Jia, Qidong Liu, Ruo Cheng Guo and Qi Liu	5040
<i>Crowd Comparative Reasoning: Unlocking Comprehensive Evaluations for LLM-as-a-Judge</i>	
Qiyuan Zhang, Yufei Wang, Yuxin Jiang, Liangyou Li, Chuhan Wu, Yasheng Wang, Xin Jiang, Lifeng Shang, Ruiming Tang, Fuyuan Lyu and Chen Ma	5059
<i>Lost in Multilinguality: Dissecting Cross-lingual Factual Inconsistency in Transformer Language Models</i>	
Mingyang Wang, Heike Adel, Lukas Lange, Yihong Liu, Ercong Nie, Jannik Strötgen and Hinrich Schuetze	5075
<i>Optimizing Decomposition for Optimal Claim Verification</i>	
Yining Lu, Noah Ziemis, Hy Dang and Meng Jiang	5095

<i>GradOT: Training-free Gradient-preserving Offsite-tuning for Large Language Models</i> Kai Yao, Zhaorui Tan, Penglei Gao, Lichun Li, Kaixin Wu, Yinggui Wang, Yuan Zhao, Yixin Ji, Jianke Zhu and Wei Wang	5115
<i>Knowledge Boundary of Large Language Models: A Survey</i> Moxin Li, Yong Zhao, Wenxuan Zhang, Shuaiyi Li, Wenya Xie, See-Kiong Ng, Tat-Seng Chua and Yang Deng	5131
<i>Mitigating Visual Forgetting via Take-along Visual Conditioning for Multi-modal Long CoT Reasoning</i> Hai-Long Sun, Zhun Sun, Houwen Peng and Han-Jia Ye	5158
<i>MoC: Mixtures of Text Chunking Learners for Retrieval-Augmented Generation System</i> Jihao Zhao, Zhiyuan Ji, Zhaoxin Fan, Hanyu Wang, Simin Niu, Bo Tang, Feiyu Xiong and Zhiyu li	5172
<i>Mitigating Selection Bias with Node Pruning and Auxiliary Options</i> Hyeong Kyu Choi, Weijie Xu, Chi Xue, Stephanie Eckman and Chandan K. Reddy	5190
<i>Dually Self-Improved Counterfactual Data Augmentation Using Large Language Model</i> Luhao Zhang, Xinyu Zhang, Linmei Hu, Dandan Song and Liqiang Nie	5216
<i>RPO: Retrieval Preference Optimization for Robust Retrieval-Augmented Generation</i> Shi-Qi Yan, Quan Liu and Zhen-Hua Ling	5228
<i>Learning to Reason from Feedback at Test-Time</i> Yanyang Li, Michael R. Lyu and Liwei Wang	5241
<i>L-CiteEval: A Suite for Evaluating Fidelity of Long-context Models</i> Zecheng Tang, Keyan Zhou, Juntao Li, Baibei Ji, Jianye Hou and Min Zhang	5254
<i>SECRET: Semi-supervised Clinical Trial Document Similarity Search</i> Trisha Das, Afrah Shafquat, Mandis Beigi, Jacob Aptekar and Jimeng Sun	5278
<i>Geometric Signatures of Compositionality Across a Language Model’s Lifetime</i> Jin Hwa Lee, Thomas Jiralerspong, Lei Yu, Yoshua Bengio and Emily Cheng	5292
<i>Pattern Recognition or Medical Knowledge? The Problem with Multiple-Choice Questions in Medicine</i> Maxime Griot, Jean Vanderdonckt, Demet Yuksel and Coralie Hemptinne	5321
<i>People who frequently use ChatGPT for writing tasks are accurate and robust detectors of AI-generated text</i> Jenna Russell, Marzena Karpinska and Mohit Iyyer	5342
<i>YuLan-Mini: Pushing the Limits of Open Data-efficient Language Model</i> Hu Yiwen, Huatong Song, Jie Chen, Jia Deng, Jiapeng Wang, Kun Zhou, Yutao Zhu, Jinhao Jiang, Zican Dong, Yang LU, Xu Miao, Xin Zhao and Ji-Rong Wen	5374
<i>Your Model is Overconfident, and Other Lies We Tell Ourselves</i> Timothee Mickus, Aman Sinha and Raúl Vázquez	5401
<i>Bridging the Language Gaps in Large Language Models with Inference-Time Cross-Lingual Intervention</i> Weixuan Wang, Minghao Wu, Barry Haddow and Alexandra Birch	5418
<i>Plug-in and Fine-tuning: Bridging the Gap between Small Language Models and Large Language Models</i> Kyeonghyun Kim, Jinhee Jang, Juhwan Choi, Yoonji Lee, Kyohoon Jin and YoungBin Kim	5434

<i>What is Stigma Attributed to? A Theory-Grounded, Expert-Annotated Interview Corpus for Demystifying Mental-Health Stigma</i>	
Han Meng, Yancan Chen, Yunan Li, Yitian Yang, Jungup Lee, Renwen Zhang and Yi-Chieh Lee	
	5453
<i>ATRI: Mitigating Multilingual Audio Text Retrieval Inconsistencies by Reducing Data Distribution Errors</i>	
Yuguo Yin, Yuxin Xie, Wenyuan Yang, Dongchao Yang, Jinghan Ru, Xianwei Zhuang, Liming Liang and Yuexian Zou	5491
<i>Enhancing Transformers for Generalizable First-Order Logical Entailment</i>	
Tianshi Zheng, Jiazheng Wang, Zihao Wang, Jiabin Bai, Hang Yin, Zheyang Deng, Yangqiu Song and Jianxin Li	5505
<i>Self-Taught Agentic Long Context Understanding</i>	
Yufan Zhuang, Xiaodong Yu, Jialian Wu, Ximeng Sun, Ze Wang, Jiang Liu, Yusheng Su, Jingbo Shang, Zicheng Liu and Emad Barsoum	5525
<i>Hallucination Detox: Sensitivity Dropout (SenD) for Large Language Model Training</i>	
Shahrad Mohammadzadeh, Juan David Guerra, Marco Bonizzato, Reihaneh Rabbany and Golnoosh Farnadi	5538
<i>OS-Genesis: Automating GUI Agent Trajectory Construction via Reverse Task Synthesis</i>	
Qiusi Sun, Kanzhi Cheng, Zichen Ding, Chuanyang Jin, Yian Wang, Fangzhi Xu, Zhenyu Wu, Chengyou Jia, Liheng Chen, Zhoumianze Liu, Ben Kao, Guohao Li, Junxian He, Yu Qiao and Zhiyong Wu	5555
<i>CORAL: Learning Consistent Representations across Multi-step Training with Lighter Speculative Drafter</i>	
Yepeng Weng, Dianwen Mei, Huishi Qiu, Xujie Chen, Li Liu, Jiang Tian and Zhongchao Shi	5580
<i>ConSim: Measuring Concept-Based Explanations' Effectiveness with Automated Simulatability</i>	
Antonin Poché, Alon Jacovi, Agustin Martin Picard, Victor Boutin and Fanny Jourdan	5594
<i>Decoding Reading Goals from Eye Movements</i>	
Omer Shubi, Cfir Avraham Hadar and Yevgeni Berzak	5616
<i>Uncovering Visual-Semantic Psycholinguistic Properties from the Distributional Structure of Text Embedding Space</i>	
Si Wu and Sebastian Bruch	5638
<i>GUI-explorer: Autonomous Exploration and Mining of Transition-aware Knowledge for GUI Agent</i>	
Bin Xie, Rui Shao, Gongwei Chen, Kaiwen Zhou, Yinchuan Li, Jie Liu, Min Zhang and Liqiang Nie	5650
<i>P² Law: Scaling Law for Post-Training After Model Pruning</i>	
Xiaodong Chen, Yuxuan Hu, Xiaokang Zhang, Yanling Wang, Cuiping Li, Hong Chen and Jing Zhang	5668
<i>Making FETCH! Happen: Finding Emergent Dog Whistles Through Common Habitats</i>	
Kuleen Sasse, Carlos Alejandro Aguirre, Isabel Cachola, Sharon Levy and Mark Dredze	5687
<i>Lost in the Context: Insufficient and Distracted Attention to Contexts in Preference Modeling</i>	
Shihan Dou, Jiayi Chen, Chenhao Huang, Feng Chen, Wei Chengzhi, Huiyuan Zheng, Shichun Liu, Yan Liu, Chenxiao Liu, Chao Xin, Lin Yan, Zongzhang Zhang, Tao Gui, Qi Zhang and Xuanjing Huang	5710

<i>Entailment-Preserving First-order Logic Representations in Natural Language Entailment</i>	
Jinu Lee, Qi Liu, Runzhi Ma, Vincent Han, Ziqi Wang, Heng Ji and Julia Hockenmaier	5729
<i>Enhancing Multimodal Continual Instruction Tuning with BranchLoRA</i>	
Duzhen Zhang, Yong Ren, Zhong-Zhi Li, Yahan Yu, Jiahua Dong, Chenxing Li, Zhilong Ji and Jinfeng Bai	5743
<i>Enhancing Automated Interpretability with Output-Centric Feature Descriptions</i>	
Yoav Gur-Arieh, Roy Mayan, Chen Agassy, Atticus Geiger and Mor Geva	5757
<i>Towards Effective and Efficient Continual Pre-training of Large Language Models</i>	
Jie Chen, Zhipeng Chen, Jiapeng Wang, Kun Zhou, Yutao Zhu, Jinhao Jiang, Yingqian Min, Xin Zhao, Zhicheng Dou, Jiaxin Mao, Yankai Lin, Ruihua Song, Jun Xu, Xu Chen, Rui Yan, Zhewei Wei, Di Hu, Wenbing Huang and Ji-Rong Wen	5779
<i>Efficient Universal Goal Hijacking with Semantics-guided Prompt Organization</i>	
Yihao Huang, Chong Wang, Xiaojun Jia, Qing Guo, Felix Juefei-Xu, Jian Zhang, Yang Liu and Geguang Pu	5796
<i>mPLUG-DocOwl2: High-resolution Compressing for OCR-free Multi-page Document Understanding</i>	
Anwen Hu, Haiyang Xu, Liang Zhang, Jiabo Ye, Ming Yan, Ji Zhang, Qin Jin, Fei Huang and Jingren Zhou	5817
<i>What Makes a Good Natural Language Prompt?</i>	
Do Xuan Long, Duy Dinh, Ngoc-Hai Nguyen, Kenji Kawaguchi, Nancy F. Chen, Shafiq Joty and Min-Yen Kan	5835
<i>X-TURING: Towards an Enhanced and Efficient Turing Test for Long-Term Dialogue Agents</i>	
Weiqi Wu, Hongqiu Wu and Hai Zhao	5874
<i>Are Rules Meant to be Broken? Understanding Multilingual Moral Reasoning as a Computational Pipeline with UniMoral</i>	
Shivani Kumar and David Jurgens	5890
<i>Modality-Aware Neuron Pruning for Unlearning in Multimodal Large Language Models</i>	
Zheyuan Liu, Guangyao Dou, Xiangchi Yuan, Chunhui Zhang, Zhaoxuan Tan and Meng Jiang	5913
<i>NGQA: A Nutritional Graph Question Answering Benchmark for Personalized Health-aware Nutritional Reasoning</i>	
Zheyuan Zhang, Yiyang Li, Nhi Ha Lan Le, Zehong Wang, Tianyi Ma, Vincent Galassi, Keerthiram Murugesan, Nuno Moniz, Werner Geyer, Nitesh V Chawla, Chuxu Zhang and Yanfang Ye . .	5934
<i>ReLearn: Unlearning via Learning for Large Language Models</i>	
Haoming Xu, Ningyuan Zhao, Liming Yang, Sendong Zhao, Shumin Deng, Mengru Wang, Bryan Hooi, Nay Oo, Huajun Chen and Ningyu Zhang	5967
<i>Understanding Cross-Domain Adaptation in Low-Resource Topic Modeling</i>	
Pritom Saha Akash and Kevin Chen-Chuan Chang	5988
<i>UAlign: Leveraging Uncertainty Estimations for Factuality Alignment on Large Language Models</i>	
Boyang Xue, Fei Mi, Qi Zhu, Hongru Wang, Rui Wang, Sheng Wang, Erxin Yu, Xuming Hu and Kam-Fai Wong	6002
<i>CoT-Valve: Length-Compressible Chain-of-Thought Tuning</i>	
Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang and Xinchao Wang	6025

HoH: A Dynamic Benchmark for Evaluating the Impact of Outdated Information on Retrieval-Augmented Generation

Jie Ouyang, Tingyue Pan, Mingyue Cheng, Ruiran Yan, Yucong Luo, Jiaying Lin and Qi Liu 6036

Uncertainty Propagation on LLM Agent

Qiwei Zhao, Dong Li, Yanchi Liu, Wei Cheng, Yiyu Sun, Mika Oishi, Takao Osaki, Katsushi Matsuda, Huaxiu Yao, Chen Zhao, Haifeng Chen and Xujiang Zhao 6064

Beyond Position: the emergence of wavelet-like properties in Transformers

Valeria Ruscio, Umberto Nanni and Fabrizio Silvestri 6074

Are the Hidden States Hiding Something? Testing the Limits of Factuality-Encoding Capabilities in LLMs

Giovanni Servedio, Alessandro De Bellis, Dario Di Palma, Vito Walter Anelli and Tommaso Di Noia 6089

Disentangling Biased Knowledge from Reasoning in Large Language Models via Machine Unlearning

Zheyuan Liu, Suraj Maharjan, Fanyou Wu, Rahil Parikh, Belhassen Bayar, Srinivasan H. Sengamedu and Meng Jiang 6105

LLaMAs Have Feelings Too: Unveiling Sentiment and Emotion Representations in LLaMA Models Through Probing

Dario Di Palma, Alessandro De Bellis, Giovanni Servedio, Vito Walter Anelli, Fedelucio Narducci and Tommaso Di Noia 6124

CxGGEC: Construction-Guided Grammatical Error Correction

Yayu Cao, Tianxiang Wang, Lvxiaowei Xu, Zhenyao Wang and Ming Cai 6143

Beyond Sequences: Two-dimensional Representation and Dependency Encoding for Code Generation

Xiangyu Zhang, Yu Zhou, Guang Yang, Wei Cheng and Taolue Chen 6157

HD-NDEs: Neural Differential Equations for Hallucination Detection in LLMs

Qing Li, Jiahui Geng, Zongxiong Chen, Derui Zhu, Yuxia Wang, Congbo Ma, Chenyang Lyu and Fakhri Karray 6173

What Is That Talk About? A Video-to-Text Summarization Dataset for Scientific Presentations

Dongqi Liu, Chenxi Whitehouse, Xi Yu, Louis Mahon, Rohit Saxena, Zheng Zhao, Yifu Qiu, Mirella Lapata and Vera Demberg 6187

NeuSym-RAG: Hybrid Neural Symbolic Retrieval with Multiview Structuring for PDF Question Answering

Ruisheng Cao, Hanchong Zhang, Tiancheng Huang, Zhangyi Kang, Yuxin Zhang, Liangtai Sun, Hanqi Li, Yuxun Miao, Shuai Fan, Lu Chen and Kai Yu 6211

ProvBench: A Benchmark of Legal Provision Recommendation for Contract Auto-Reviewing

Xiuxuan Shen, Zhongyuan Jiang, Junsan Zhang, Junxiao Han, Yao Wan, Chengjie Guo, Bingcheng Liu, Jie Wu, Renxiang Li and Philip S. Yu 6240

F5-TTS: A Fairytaler that Fakes Fluent and Faithful Speech with Flow Matching

Yushen Chen, Zhikang Niu, Ziyang Ma, Keqi Deng, Chunhui Wang, JianZhao JianZhao, Kai Yu and Xie Chen 6255

AutoMedEval: Harnessing Language Models for Automatic Medical Capability Evaluation

Xiechi Zhang, Zetian Ouyang, Linlin Wang, Gerard De Melo, Zhu Cao, Xiaoling Wang, Ya Zhang, Yanfeng Wang and Liang He 6272

<i>CoT-based Synthesizer: Enhancing LLM Performance through Answer Synthesis</i> Bohan Zhang, Xiaokang Zhang, Jing Zhang, Jifan Yu, Sijia Luo and Jie Tang	6286
<i>Efficiently Identifying Watermarked Segments in Mixed-Source Texts</i> Xuandong Zhao, Chenwen Liao, Yu-Xiang Wang and Lei Li	6304
<i>Assessing Dialect Fairness and Robustness of Large Language Models in Reasoning Tasks</i> Fangru Lin, Shaoguang Mao, Emanuele La Malfa, Valentin Hofmann, Adrian de Wynter, Xun Wang, Si-Qing Chen, Michael J. Wooldridge, Janet B. Pierrehumbert and Furu Wei	6317
<i>Towards a More Generalized Approach in Open Relation Extraction</i> Qing Wang, Yuepei Li, Qiao Qiao, Kang Zhou and Qi Li	6343
<i>Adaptive Retrieval Without Self-Knowledge? Bringing Uncertainty Back Home</i> Viktor Moskvoretiskii, Maria Marina, Mikhail Salnikov, Nikolay Ivanov, Sergey Pletenev, Daria Galimzianova, Nikita Krayko, Vasily Konovalov, Irina Nikishina and Alexander Panchenko	6355
<i>Evaluating Language Models as Synthetic Data Generators</i> Seungone Kim, Juyoung Suk, Xiang Yue, Vijay Viswanathan, Seongyun Lee, Yizhong Wang, Kiril Gashteovski, Carolin Lawrence, Sean Welleck and Graham Neubig	6385
<i>Can Graph Descriptive Order Affect Solving Graph Problems with LLMs?</i> Yuyao Ge, Shenghua Liu, Baolong Bi, Yiwei Wang, Lingrui Mei, Wenjie Feng, Lizhe Chen and Xueqi Cheng	6404
<i>Learning to Rewrite: Generalized LLM-Generated Text Detection</i> Wei Hao, Ran Li, Weiliang Zhao, Junfeng Yang and Chengzhi Mao	6421
<i>Evaluating Multimodal Large Language Models on Video Captioning via Monte Carlo Tree Search</i> Linhao Yu, Xingguang Ji, Yahui Liu, Fanheng Kong, Chenxi Sun, Jingyuan Zhang, Hongzhi Zhang, V. W., Fuzheng Zhang and Deyi Xiong	6435
<i>GIFT-SW: Gaussian noise Injected Fine-Tuning of Salient Weights for LLMs</i> Maxim Zhelnin, Viktor Moskvoretiskii, Egor Shvetsov, Maria Krylova, Venediktov Egor, Zuev Aleksandr and Evgeny Burnaev	6463
<i>Quaff: Quantized Parameter-Efficient Fine-Tuning under Outlier Spatial Stability Hypothesis</i> Hong Huang and Dapeng Wu	6481
<i>Unsolvable Problem Detection: Robust Understanding Evaluation for Large Multimodal Models</i> Atsuyuki Miyai, Jingkang Yang, Jingyang Zhang, Yifei Ming, Qing Yu, Go Irie, Yixuan Li, Hai Helen Li, Ziwei Liu and Kiyoharu Aizawa	6497
<i>AlignMMBench: Evaluating Chinese Multimodal Alignment in Large Vision-Language Models</i> Yuhang Wu, Wenmeng Yu, Yean Cheng, Yan Wang, Xiaohan Zhang, Jiazheng Xu, Ming Ding and Yuxiao Dong	6541
<i>Biased LLMs can Influence Political Decision-Making</i> Jillian Fisher, Shangbin Feng, Robert Aron, Thomas Richardson, Yejin Choi, Daniel W Fisher, Jennifer Pan, Yulia Tsvetkov and Katharina Reinecke	6559
<i>LexTempus: Enhancing Temporal Generalizability of Legal Language Models Through Dynamic Mixture of Experts</i> Santosh T.y.s.s and Tuan-Quang Vuong	6608
<i>That is Unacceptable: the Moral Foundations of Canceling</i> Soda Marem Lo, Oscar Araque, Rajesh Sharma and Marco Antonio Stranisci	6625

<i>FloorPlan-LLaMa: Aligning Architects’ Feedback and Domain Knowledge in Architectural Floor Plan Generation</i>	
Jun Yin, Pengyu Zeng, Haoyuan Sun, Yuqin Dai, Han Zheng, Miao Zhang, Yachao Zhang and Shuai Lu	6640
<i>TheoremExplainAgent: Towards Video-based Multimodal Explanations for LLM Theorem Understanding</i>	
Max Ku, Cheuk Hei Chong, Jonathan Leung, Krish Shah, Alvin Yu and Wenhui Chen	6663
<i>FineReason: Evaluating and Improving LLMs’ Deliberate Reasoning through Reflective Puzzle Solving</i>	
Guizhen Chen, Weiwen Xu, Hao Zhang, Hou Pong Chan, Chaoqun Liu, Lidong Bing, Deli Zhao, Anh Tuan Luu and Yu Rong	6685
<i>The TIP of the Iceberg: Revealing a Hidden Class of Task-in-Prompt Adversarial Attacks on LLMs</i>	
Sergey Berezin, Reza Farahbakhsh and Noel Crespi	6716
<i>Identifying Reliable Evaluation Metrics for Scientific Text Revision</i>	
Leane Jourdan, Nicolas Hernandez, Florian Boudin and Richard Dufour	6731
<i>Can Language Models Reason about Individualistic Human Values and Preferences?</i>	
Liwei Jiang, Taylor Sorensen, Sydney Levine and Yejin Choi	6757
<i>BERT-like Models for Slavic Morpheme Segmentation</i>	
Dmitry Morozov, Lizaveta Astapenka, Anna Glazkova, Timur Garipov and Olga Lyashevskaya	6795
<i>Turning Trash into Treasure: Accelerating Inference of Large Language Models with Token Recycling</i>	
Xianzhen Luo, Yixuan Wang, Qingfu Zhu, Zhiming Zhang, Xuanyu Zhang, Qing Yang and Dongliang Xu	6816
<i>Unlocking General Long Chain-of-Thought Reasoning Capabilities of Large Language Models via Representation Engineering</i>	
Xinyu Tang, Xiaolei Wang, Zhihao Lv, Yingqian Min, Xin Zhao, Binbin Hu, Ziqi Liu and Zhiqiang Zhang	6832
<i>Drift: Enhancing LLM Faithfulness in Rationale Generation via Dual-Reward Probabilistic Inference</i>	
Jiazheng Li, Hanqi Yan and Yulan He	6850
<i>Fairness through Difference Awareness: Measuring Desired Group Discrimination in LLMs</i>	
Angelina Wang, Michelle Phan, Daniel E. Ho and Sanmi Koyejo	6867
<i>MergePrint: Merge-Resistant Fingerprints for Robust Black-box Ownership Verification of Large Language Models</i>	
Shojiro Yamabe, Futa Kai Waseda, Tsubasa Takahashi and Koki Wataoka	6894
<i>Dynamic Scaling of Unit Tests for Code Reward Modeling</i>	
Zeyao Ma, Xiaokang Zhang, Jing Zhang, Jifan Yu, Sijia Luo and Jie Tang	6917
<i>UniConv: Unifying Retrieval and Response Generation for Large Language Models in Conversations</i>	
Fengran Mo, Yifan Gao, Chuan Meng, Xin Liu, Zhuofeng Wu, Kelong Mao, Zhengyang Wang, Pei Chen, Zheng Li, Xian Li, Bing Yin and Meng Jiang	6936
<i>Tracking Life’s Ups and Downs: Mining Life Events from Social Media Posts for Mental Health Analysis</i>	
Minghao Lv, Siyuan Chen, Haoan Jin, Minghao Yuan, Qianqian Ju, Yujia Peng, Kenny Q. Zhu and Mengyue Wu	6950

<i>ControlSpeech: Towards Simultaneous and Independent Zero-shot Speaker Cloning and Zero-shot Language Style Control</i>	
Shengpeng Ji, Qian Chen, Wen Wang, Jialong Zuo, Minghui Fang, Ziyue Jiang, Hai Huang, Zehan Wang, Xize Cheng, Siqi Zheng and Zhou Zhao	6966
<i>PIC: Unlocking Long-Form Text Generation Capabilities of Large Language Models via Position ID Compression</i>	
Haoran Que and Wenge Rong	6982
<i>Towards Effective Extraction and Evaluation of Factual Claims</i>	
Dasha Metropolitansky and Jonathan Larson	6996
<i>Beyond Facts: Evaluating Intent Hallucination in Large Language Models</i>	
Yijie Hao, Haofei Yu and Jiaxuan You	7046
<i>A Systematic Study of Compositional Syntactic Transformer Language Models</i>	
Yida Zhao, Hao Xve, Xiang Hu and Kewei Tu	7070
<i>M-MAD: Multidimensional Multi-Agent Debate for Advanced Machine Translation Evaluation</i>	
Zhaopeng Feng, Jiayuan Su, Jiamei Zheng, Jiahan Ren, Yan Zhang, Jian Wu, Hongwei Wang and Zuozhu Liu	7084
<i>SongComposer: A Large Language Model for Lyric and Melody Generation in Song Composition</i>	
Shuangrui Ding, Zihan Liu, Xiaoyi Dong, Pan Zhang, Rui Qian, Junhao Huang, Conghui He, Dahua Lin and Jiaqi Wang	7108
<i>Personalized Text Generation with Contrastive Activation Steering</i>	
Jinghao Zhang, Yuting Liu, Wenjie Wang, Qiang Liu, Shu Wu, Liang Wang and Tat-Seng Chua	7128
<i>Gumbel Reranking: Differentiable End-to-End Reranker Optimization</i>	
Siyuan Huang, Zhiyuan Ma, Jintao Du, Changhua Meng, Weiqiang Wang, Jingwen Leng, Minyi Guo and Zhouhan Lin	7142
<i>Hybrid Preferences: Learning to Route Instances for Human vs. AI Feedback</i>	
Lester James Validad Miranda, Yizhong Wang, Yanai Elazar, Sachin Kumar, Valentina Pyatkin, Faeze Brahman, Noah A. Smith, Hannaneh Hajishirzi and Pradeep Dasigi	7162
<i>SEOE: A Scalable and Reliable Semantic Evaluation Framework for Open Domain Event Detection</i>	
Yi-Fan Lu, Xian-Ling Mao, Tian Lan, Tong Zhang, Yu-Shi Zhu and Heyan Huang	7201
<i>The UD-NewsCrawl Treebank: Reflections and Challenges from a Large-scale Tagalog Syntactic Annotation Project</i>	
Angelina Aspra Aquino, Lester James Validad Miranda and Elsie Marie T. Or	7219
<i>DRAG: Distilling RAG for SLMs from LLMs to Transfer Knowledge and Mitigate Hallucination via Evidence and Graph-based Distillation</i>	
Jennifer Chen, Aidar Myrzakhan, Yaxin Luo, Hassaan Muhammad Khan, Sondos Mahmoud Bsharat and Zhiqiang Shen	7240
<i>G-Safeguard: A Topology-Guided Security Lens and Treatment on LLM-based Multi-agent Systems</i>	
Shilong Wang, Guibin Zhang, Miao Yu, Guancheng Wan, Fanci Meng, Chongye Guo, Kun Wang and Yang Wang	7261
<i>Deontological Keyword Bias: The Impact of Modal Expressions on Normative Judgments of Language Models</i>	
Bumjin Park, Leejinsil Leejinsil and Jaesik Choi	7277

<i>LegalReasoner: Step-wised Verification-Correction for Legal Judgment Reasoning</i>	
Weijie Shi, Han Zhu, Jiaming Ji, Mengze Li, Jipeng Zhang, Ruiyuan Zhang, Jia Zhu, Jiajie Xu, Sirui Han and Yike Guo	7297
<i>Rolling the DICE on Idiomaticity: How LLMs Fail to Grasp Context</i>	
Maggie Mi, Aline Villavicencio and Nafise Sadat Moosavi	7314
<i>ChartCoder: Advancing Multimodal Large Language Model for Chart-to-Code Generation</i>	
Xuanle Zhao, Xianzhen Luo, Qi Shi, Chi Chen, Shuo Wang, Zhiyuan Liu and Maosong Sun	7333
<i>The Cross-linguistic Role of Animacy in Grammar Structures</i>	
Nina Gregorio, Matteo Gay, Sharon Goldwater and Edoardo Ponti	7349
<i>LexGen: Domain-aware Multilingual Lexicon Generation</i>	
Ayush Maheshwari, Atul Kumar Singh, N J Karthika, Krishnakant Bhatt, Preethi Jyothi and Ganesh Ramakrishnan	7364
<i>How to Train Long-Context Language Models (Effectively)</i>	
Tianyu Gao, Alexander Wettig, Howard Yen and Danqi Chen	7376
<i>MathFusion: Enhancing Mathematical Problem-solving of LLM through Instruction Fusion</i>	
Qizhi Pei, Lijun Wu, Zhuoshi Pan, Yu Li, Honglin Lin, Chenlin Ming, Xin Gao, Conghui He and Rui Yan	7400
<i>Mining Complex Patterns of Argumentative Reasoning in Natural Language Dialogue</i>	
Ramon Ruiz-Dolz, Zlata Kikteva and John Lawrence	7421
<i>OS Agents: A Survey on MLLM-based Agents for Computer, Phone and Browser Use</i>	
Xueyu Hu, Tao Xiong, Biao Yi, Zishu Wei, Ruixuan Xiao, Yurun Chen, Jiasheng Ye, Meiling Tao, Xiangxin Zhou, Ziyu Zhao, Yuhuai Li, Shengze Xu, Shenzi Wang, Xinchun Xu, Shuofei Qiao, Zhaokai Wang, Kun Kuang, Tiejong Zeng, Liang Wang, Jiwei Li, Yuchen Eleanor Jiang, Wangchunshu Zhou, Guoyin Wang, Keting Yin, Zhou Zhao, Hongxia Yang, Fan Wu, Shengyu Zhang and Fei Wu	7436
<i>Data Quality Issues in Multilingual Speech Datasets: The Need for Sociolinguistic Awareness and Proactive Language Planning</i>	
Mingfei Lau, Qian Chen, Yeming Fang, Tingting Xu, Tongzhou Chen and Pavel Golik	7466
<i>LLM as a Broken Telephone: Iterative Generation Distorts Information</i>	
Amr Mohamed, Mingmeng Geng, Michalis Vazirgiannis and Guokan Shang	7493
<i>VLM2-Bench: A Closer Look at How Well VLMs Implicitly Link Explicit Matching Visual Cues</i>	
Jianshu Zhang, Dongyu Yao, Renjie Pi, Paul Pu Liang and Yi R. Fung	7510
<i>Alleviating Distribution Shift in Synthetic Data for Machine Translation Quality Estimation</i>	
Xiang Geng, Zhejian Lai, Jiajun Chen, Hao Yang and Shujian Huang	7546
<i>Evaluation Agent: Efficient and Promptable Evaluation Framework for Visual Generative Models</i>	
Fan Zhang, Shulin Tian, Ziqi Huang, Yu Qiao and Ziwei Liu	7561
<i>Large Language Models Struggle to Describe the Haystack without Human Help: A Social Science-Inspired Evaluation of Topic Models</i>	
Zongxia Li, Lorena Calvo-Bartolomé, Alexander Miserlis Hoyle, Paiheng Xu, Daniel Kofi Stephens, Juan Francisco Fung, Alden Dima and Jordan Lee Boyd-Graber	7583
<i>ActiView: Evaluating Active Perception Ability for Multimodal Large Language Models</i>	
Ziyue Wang, Chi Chen, Fuwen Luo, Yurui Dong, Yuanchi Zhang, Yuzhuang Xu, Xiaolong Wang, Peng Li and Yang Liu	7605

<i>Enough Coin Flips Can Make LLMs Act Bayesian</i>	
Ritwik Gupta, Rodolfo Corona, Jiaxin Ge, Eric Wang, Dan Klein, Trevor Darrell and David M. Chan	7634
<i>GAMEBot: Transparent Assessment of LLM Reasoning in Games</i>	
Wenye Lin, Jonathan Roberts, Yunhan Yang, Samuel Albanie, Zongqing Lu and Kai Han . .	7656
<i>A Text is Worth Several Tokens: Text Embedding from LLMs Secretly Aligns Well with The Key Tokens</i>	
Zhijie Nie, Richong Zhang and Zhanyu Wu	7683
<i>Commonsense Reasoning in Arab Culture</i>	
Abdelrahman Sadallah, Junior Cedric Tonga, Khalid Almubarak, Saeed Almheiri, Farah Atif, Chatrine Qwaider, Karima Kadaoui, Sara Shatnawi, Yaser Alesh and Fajri Koto	7695
<i>AXIS: Efficient Human-Agent-Computer Interaction with API-First LLM-Based Agents</i>	
Junting Lu, Zhiyang Zhang, Fangkai Yang, Jue Zhang, Lu Wang, Chao Du, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang and Qi Zhang	7711
<i>Translation and Fusion Improves Cross-lingual Information Extraction</i>	
Yang Chen, Vedaant Shah and Alan Ritter	7744
<i>Conditional Dichotomy Quantification via Geometric Embedding</i>	
Shaobo Cui, Wenqing Liu, Yiyang Feng, Jiawei Zhou and Boi Faltings	7765
<i>Aligning Large Language Models with Implicit Preferences from User-Generated Content</i>	
Zhaoxuan Tan, Zheng Li, Tianyi Liu, Haodong Wang, Hyokun Yun, Ming Zeng, Pei Chen, Zhihan Zhang, Yifan Gao, Ruijie Wang, Priyanka Nigam, Bing Yin and Meng Jiang	7792
<i>VQAGuider: Guiding Multimodal Large Language Models to Answer Complex Video Questions</i>	
Yuyan Chen, Jiyuan Jia, Jiaxin Lu, Siyue Li, Yu Guan, Ming Yang and Qingpei Guo	7821
<i>Large Language Models are Good Relational Learners</i>	
Fang Wu, Vijay Prakash Dwivedi and Jure Leskovec	7835
<i>SpaRE: Enhancing Spatial Reasoning in Vision-Language Models with Synthetic Data</i>	
Michael Ogezi and Freda Shi	7855
<i>Distilling an End-to-End Voice Assistant Without Instruction Training Data</i>	
William Barr Held, Yanzhe Zhang, Weiyan Shi, Minzhi Li, Michael J Ryan and Diyi Yang .	7876
<i>CoMet: Metaphor-Driven Covert Communication for Multi-Agent Language Games</i>	
Shuhang Xu and Fangwei Zhong	7892
<i>CER: Confidence Enhanced Reasoning in LLMs</i>	
Ali Razghandi, Seyed Mohammad Hadi Hosseini and Mahdieh Soleymani Baghshah	7918
<i>Watermarking Large Language Models: An Unbiased and Low-risk Method</i>	
Minjia Mao, Dongjun Wei, Zeyu Chen, Xiao Fang and Michael Chau	7939
<i>On Synthetic Data Strategies for Domain-Specific Generative Retrieval</i>	
Haoyang Wen, Jiang Guo, Yi Zhang, Jiarong Jiang and Zhiguo Wang	7961
<i>LLM Braces: Straightening Out LLM Predictions with Relevant Sub-Updates</i>	
Ying Shen and Lifu Huang	7977
<i>CONFETTI: Conversational Function-Calling Evaluation Through Turn-Level Interactions</i>	
Tamer Alkhouli, Katerina Margatina, James Gung, Raphael Shu, Claudia Zaghi, Monica Sunkara and Yi Zhang	7993

<i>Evaluating Theory of (an uncertain) Mind: Predicting the Uncertain Beliefs of Others from Conversational Cues</i>	
Anthony Sicilia and Malihe Alikhani	8007
<i>Uncertainty in Causality: A New Frontier</i>	
Shaobo Cui, Luca Mouchel and Boi Faltings	8022
<i>SynthesizeMe! Inducing Persona-Guided Prompts for Personalized Reward Models in LLMs</i>	
Michael J Ryan, Omar Shaikh, Aditri Bhagirath, Daniel Frees, William Barr Held and Diyi Yang	8045
<i>When People are Floods: Analyzing Dehumanizing Metaphors in Immigration Discourse with Large Language Models</i>	
Julia Mendelsohn and Ceren Budak	8079
<i>AGrail: A Lifelong Agent Guardrail with Effective and Adaptive Safety Detection</i>	
Weidi Luo, Shenghong Dai, Xiaogeng Liu, Suman Banerjee, Huan Sun, Muhao Chen and Chaowei Xiao	8104
<i>Improving Model Factuality with Fine-grained Critique-based Evaluator</i>	
Yiqing Xie, Wenxuan Zhou, Pradyot Prakash, Di Jin, Yuning Mao, Quintin Fettes, Arya Talebzadeh, Sinong Wang, Han Fang, Carolyn Rose, Daniel Fried and Hejia Zhang	8140
<i>Building a Long Text Privacy Policy Corpus with Multi-Class Labels</i>	
Florecia Marotta-Wurgler and David Stein	8156
<i>R2-MultiOmnia: Leading Multilingual Multimodal Reasoning via Self-Training</i>	
Leonardo Ranaldi, Federico Ranaldi and Giulia Pucci	8220
<i>When the LM misunderstood the human chuckled: Analyzing garden path effects in humans and language models</i>	
Samuel Joseph Amouyal, Aya Meltzer-Asscher and Jonathan Berant	8235
<i>Cross-Lingual Pitfalls: Automatic Probing Cross-Lingual Weakness of Multilingual Large Language Models</i>	
Zixiang Xu, Yanbo Wang, Yue Huang, Xiuying Chen, Jieyu Zhao, Meng Jiang and Xiangliang Zhang	8254
<i>VLSBench: Unveiling Visual Leakage in Multimodal Safety</i>	
Xuhao Hu, Dongrui Liu, Hao Li, Xuanjing Huang and Jing Shao	8285
<i>Browsing Lost Unformed Recollections: A Benchmark for Tip-of-the-Tongue Search and Reasoning</i>	
Sky CH-Wang, Darshan Girish Deshpande, Smaranda Muresan, Anand Kannappan and Rebecca Qian	8317
<i>Data Laundering: Artificially Boosting Benchmark Results through Knowledge Distillation</i>	
Jonibek Mansurov, Akhmed Sakip and Alham Fikri Aji	8332
<i>Conspiracy Theories and Where to Find Them on TikTok</i>	
Francesco Corso, Francesco Pierri and Gianmarco De Francisci Morales	8346
<i>Growing Through Experience: Scaling Episodic Grounding in Language Models</i>	
Chunhui Zhang, Sirui Wang, Zhongyu Ouyang, Xiangchi Yuan and Soroush Vosoughi	8363
<i>Exploiting the Shadows: Unveiling Privacy Leaks through Lower-Ranked Tokens in Large Language Models</i>	
Yuan Zhou, Zhuo Zhang and Xiangyu Zhang	8376

<i>Attacking Vision-Language Computer Agents via Pop-ups</i> Yanzhe Zhang, Tao Yu and Diyi Yang	8387
<i>Explicit and Implicit Data Augmentation for Social Event Detection</i> Congbo Ma, Yuxia Wang, Jia Wu, Jian Yang, Jing Du, Zitai Qiu, Qing Li, Hu Wang and Preslav Nakov	8402
<i>In Prospect and Retrospect: Reflective Memory Management for Long-term Personalized Dialogue Agents</i> Zhen Tan, Jun Yan, I-Hung Hsu, Rujun Han, Zifeng Wang, Long Le, Yiwen Song, Yanfei Chen, Hamid Palangi, George Lee, Anand Rajan Iyer, Tianlong Chen, Huan Liu, Chen-Yu Lee and Tomas Pfister	8416
<i>Revisiting Classical Chinese Event Extraction with Ancient Literature Information</i> Xiaoyi Bao, Zhongqing Wang, Jinghang Gu and Chu-Ren Huang	8440
<i>Unanswerability Evaluation for Retrieval Augmented Generation</i> Xiangyu Peng, Prafulla Kumar Choubey, Caiming Xiong and Chien-Sheng Wu	8452
<i>SCALE: Towards Collaborative Content Analysis in Social Science with Large Language Model Agents and Human Intervention</i> Chengshuai Zhao, Zhen Tan, Chau-Wai Wong, Xinyan Zhao, Tianlong Chen and Huan Liu	8473
<i>Self-Error-Instruct: Generalizing from Errors for LLMs Mathematical Reasoning</i> Erxin Yu, Jing Li, Ming Liao, Qi Zhu, Boyang Xue, Minghui Xu, Baojun Wang, Lanqing Hong, Fei Mi and Lifeng Shang	8504
<i>RAGEval: Scenario Specific RAG Evaluation Dataset Generation Framework</i> Kunlun Zhu, Yifan Luo, Dingling Xu, Yukun Yan, Zhenghao Liu, Shi Yu, Ruobing Wang, Shuo Wang, Yishan Li, Nan Zhang, Xu Han, Zhiyuan Liu and Maosong Sun	8520
<i>A Survey on Patent Analysis: From NLP to Multimodal AI</i> Homaira Huda Shomee, Zhu Wang, Sathya N. Ravi and Sourav Medya	8545
<i>SciVer: Evaluating Foundation Models for Multimodal Scientific Claim Verification</i> Chengye Wang, Yifei Shen, Zexi Kuang, Arman Cohan and Yilun Zhao	8562
<i>MultiAgentBench : Evaluating the Collaboration and Competition of LLM agents</i> Kunlun Zhu, Hongyi Du, Zhaochen Hong, Xiaocheng Yang, Shuyi Guo, Zhe Wang, Zhenhailong Wang, Cheng Qian, Robert Tang, Heng Ji and Jiakuan You	8580
<i>Sinhala Encoder-only Language Models and Evaluation</i> Tharindu Ranasinghe, Hansi Hettiarachchi, Nadeesha Chathurangi Naradde Vidana Pathirana, Damith Premasiri, Lasitha Uyangodage, Isuri Nanomi Arachchige, Alistair Plum, Paul Rayson and Ruslan Mitkov	8623
<i>LLMs can Perform Multi-Dimensional Analytic Writing Assessments: A Case Study of L2 Graduate-Level Academic English Writing</i> Zhengxiang Wang, Veronika Makarova, Zhi Li, Jordan Kodner and Owen Rambow	8637
<i>SEUF: Is Unlearning One Expert Enough for Mixture-of-Experts LLMs?</i> Haomin Zhuang, Yihua Zhang, Kehan Guo, Jinghan Jia, Gaowen Liu, Sijia Liu and Xiangliang Zhang	8664
<i>Pragmatics in the Era of Large Language Models: A Survey on Datasets, Evaluation, Opportunities and Challenges</i> Bolei Ma, Yuting Li, Wei Zhou, Ziwei Gong, Yang Janet Liu, Katja Jasinskaja, Annemarie Friedrich, Julia Hirschberg, Frauke Kreuter and Barbara Plank	8679

<i>LocAgent: Graph-Guided LLM Agents for Code Localization</i>	
Zhaoling Chen, Robert Tang, Gangda Deng, Fang Wu, Jialong Wu, Zhiwei Jiang, Viktor Prasanna, Arman Cohan and Xingyao Wang	8697
<i>COSMMIC: Comment-Sensitive Multimodal Multilingual Indian Corpus for Summarization and Headline Generation</i>	
Raghvendra Kumar, Mohammed Salman S A, Aryan Sahu, Tridib Nandi, Pragathi Y P, Sriparna Saha and Jose G Moreno	8728
<i>Mind the Gap: Static and Interactive Evaluations of Large Audio Models</i>	
Minzhi Li, William Barr Held, Michael J Ryan, Kunat Pipatanakul, Potsawee Manakul, Hao Zhu and Diyi Yang	8749
<i>Understanding In-Context Machine Translation for Low-Resource Languages: A Case Study on Manchu</i>	
Renhao Pei, Yihong Liu, Peiqin Lin, François Yvon and Hinrich Schuetze	8767
<i>CKnowEdit: A New Chinese Knowledge Editing Dataset for Linguistics, Facts, and Logic Error Correction in LLMs</i>	
Jizhan Fang, Tianhe Lu, Yunzhi Yao, Ziyang Jiang, Xin Xu, Huajun Chen and Ningyu Zhang	8789
<i>TripleFact: Defending Data Contamination in the Evaluation of LLM-driven Fake News Detection</i>	
Cheng Xu and Nan Yan	8808
<i>Meaning Beyond Truth Conditions: Evaluating Discourse Level Understanding via Anaphora Accessibility</i>	
Xiaomeng Zhu, Zhenghao Zhou, Simon Charlow and Robert Frank	8824
<i>Large Language and Reasoning Models are Shallow Disjunctive Reasoners</i>	
Irtaza Khalid, Amir Masoud Nourollah and Steven Schockaert	8843
<i>Warmup Generations: A Task-Agnostic Approach for Guiding Sequence-to-Sequence Learning with Unsupervised Initial State Generation</i>	
Senyu Li, Zipeng Sun, Jiayi Wang, Xue Liu, Pontus Stenetorp, Siva Reddy and David Ifeoluwa Adelani	8870
<i>Building Better: Avoiding Pitfalls in Developing Language Resources when Data is Scarce</i>	
Nedjma Ousidhoum, Meriem Beloucif and Saif M. Mohammad	8881
<i>BRIGHTER: BRIdging the Gap in Human-Annotated Textual Emotion Recognition Datasets for 28 Languages</i>	
Shamsuddeen Hassan Muhammad, Nedjma Ousidhoum, Idris Abdulmumin, Jan Philip Wahle, Terry Ruas, Meriem Beloucif, Christine de Kock, Nirmal Surange, Daniela Teodorescu, Ibrahim Said Ahmad, David Ifeoluwa Adelani, Alham Fikri Aji, Felermino D. M. A. Ali, Ilseyar Alimova, Vladimir Araujo, Nikolay Babakov, Naomi Baes, Ana-Maria Bucur, Andiswa Bukula, Guanqun Cao, Rodrigo Tufiño, Rendi Chevi, Chiamaka Ijeoma Chukwunke, Alexandra Ciobotaru, Daryna Dementieva, Murja Sani Gadanya, Robert Geislinger, Bela Gipp, Oumaima Hourrane, Oana Ignat, Falalu Ibrahim Lawan, Rooweither Mabuya, Rahmad Mahendra, Vukosi Marivate, Alexander Panchenko, Andrew Piper, Charles Henrique Porto Ferreira, Vitaly Protasov, Samuel Rutunda, Manish Shrivastava, Aura Cristina Udrea, Lilian Diana Awuor Wanzare, Sophie Wu, Florian Valentin Wunderlich, Hanif Muhammad Zhafran, Tianhui Zhang, Yi Zhou and Saif M. Mohammad	8895
<i>SkillVerse : Assessing and Enhancing LLMs with Tree Evaluation</i>	
Yufei Tian, Jiao Sun, Nanyun Peng and Zizhao Zhang	8917

<i>CypherBench: Towards Precise Retrieval over Full-scale Modern Knowledge Graphs in the LLM Era</i> Yanlin Feng, Simone Papicchio and Sajjadur Rahman	8934
<i>Empathy Prediction from Diverse Perspectives</i> Francine Chen, Scott Carter, Tatiana Lau, Nayeli Suseth Bravo, Sumanta Bhattacharyya, Kate Sieck and Charlene C. Wu	8959
<i>Are LLMs effective psychological assessors? Leveraging adaptive RAG for interpretable mental health screening through psychometric practice</i> Federico Ravenda, Seyed Ali Bahrainian, Andrea Raballo, Antonietta Mira and Noriko Kando	8975
<i>INTERACT: Enabling Interactive, Question-Driven Learning in Large Language Models</i> Aum Kendapadi, Kerem Zaman, Rakesh R Menon and Shashank Srivastava	8992
<i>Circuit Stability Characterizes Language Model Generalization</i> Alan Sun.....	9025
<i>Comparing LLM-generated and human-authored news text using formal syntactic theory</i> Olga Zamaraeva, Dan Flickinger, Francis Bond and Carlos Gómez-Rodríguez	9041
<i>Improving Preference Extraction In LLMs By Identifying Latent Knowledge Through Classifying Probes</i> Sharan Maiya, Yinhong Liu, Ramit Debnath and Anna Korhonen	9061
<i>White Men Lead, Black Women Help? Benchmarking and Mitigating Language Agency Social Biases in LLMs</i> Yixin Wan and Kai-Wei Chang	9082
<i>AIMSCheck: Leveraging LLMs for AI-Assisted Review of Modern Slavery Statements Across Jurisdictions</i> Adriana Eufrosina Bora, Akshatha Arodi, Duoyi Zhang, Jordan Bannister, Mirko Bronzi, Arsene Fansi Tchango, Md Abul Bashar, Richi Nayak and Kerrie Mengersen	9109
<i>Collapse of Dense Retrievers: Short, Early, and Literal Biases Outranking Factual Evidence</i> Mohsen Fayyaz, Ali Modarressi, Hinrich Schuetze and Nanyun Peng	9136
<i>SelfElicit: Your Language Model Secretly Knows Where is the Relevant Evidence</i> Zhining Liu, Rana Ali Amjad, Ravinarayana Adkathimar, Tianxin Wei and Hanghang Tong	9153
<i>The Male CEO and the Female Assistant: Evaluation and Mitigation of Gender Biases in Text-To-Image Generation of Dual Subjects</i> Yixin Wan and Kai-Wei Chang	9174
<i>Mitigating Shortcut Learning with InterpoLated Learning</i> Michalis Korakakis, Andreas Vlachos and Adrian Weller	9191
<i>Toward Automatic Discovery of a Canine Phonetic Alphabet</i> Theron S. Wang, Xingyuan Li, Hridayesh Lekhak, Tuan Minh Dang, Mengyue Wu and Kenny Q. Zhu.....	9207
<i>DavIR: Data Selection via Implicit Reward for Large Language Models</i> Haotian Zhou, Tingkai Liu, Qianli Ma, Yufeng Zhang, Jianbo Yuan, Pengfei Liu, Yang You and Hongxia Yang	9220
<i>Byte Latent Transformer: Patches Scale Better Than Tokens</i> Artidoro Pagnoni, Ramakanth Pasunuru, Pedro Rodriguez, John Nguyen, Benjamin Muller, Margaret Li, Chunting Zhou, Lili Yu, Jason E Weston, Luke Zettlemoyer, Gargi Ghosh, Mike Lewis, Ari Holtzman and Srini Iyer	9238

<i>DiffuseDef: Improved Robustness to Adversarial Attacks via Iterative Denoising</i> Zhenhao Li, Huichi Zhou, Marek Rei and Lucia Specia	9259
<i>Identifying Cellular Niches in Spatial Transcriptomics: An Investigation into the Capabilities of Large Language Models</i> Huanhuan Wei, Xiao Luo, Hongyi Yu, Jinping Liang, Luning Yang, Lixing Lin, Alexandra Popa and Xiting Yan	9275
<i>Culture Matters in Toxic Language Detection in Persian</i> Zahra Bokaei, Walid Magdy and Bonnie Webber	9290
<i>Bitnet.cpp: Efficient Edge Inference for Ternary LLMs</i> Jinheng Wang, Hansong Zhou, Ting Song, Shijie Cao, Yan Xia, Ting Cao, Jianyu Wei, Shuming Ma, Hongyu Wang and Furu Wei	9305
<i>Instance-Selection-Inspired Undersampling Strategies for Bias Reduction in Small and Large Language Models for Binary Text Classification</i> Guilherme Fonseca, Washington Cunha, Gabriel Prenassi, Marcos André Gonçalves and Leonardo Chaves Dutra Da Rocha	9323
<i>Forward Knows Efficient Backward Path: Saliency-Guided Memory-Efficient Fine-tuning of Large Language Models</i> Yeachan Kim and SangKeun Lee	9341
<i>Focus on What Matters: Enhancing Medical Vision-Language Models with Automatic Attention Alignment Tuning</i> Aofei Chang, Le Huang, Alex James Boyd, Parminder Bhatia, Taha Kass-Hout, Cao Xiao and Fenglong Ma	9357
<i>LLMs + Persona-Plug = Personalized LLMs</i> Jiongnan Liu, Yutao Zhu, Shuting Wang, Xiaochi Wei, Erxue Min, Yu Lu, Shuaiqiang Wang, Dawei Yin and Zhicheng Dou	9373
<i>Developmentally-plausible Working Memory Shapes a Critical Period for Language Acquisition</i> Masato Mita, Ryo Yoshida and Yohei Oseki	9386
<i>IRIS: An Iterative and Integrated Framework for Verifiable Causal Discovery in the Absence of Tabular Data</i> Tao Feng, Lizhen Qu, Niket Tandon and Gholamreza Haffari	9400
<i>INJONGO: A Multicultural Intent Detection and Slot-filling Dataset for 16 African Languages</i> Hao Yu, Jesujoba Oluwadara Alabi, Andiswa Bukula, Jian Yun Zhuang, En-Shiun Annie Lee, Tadesse Kebede Guge, Israel Abebe Azime, Happy Buzaaba, Blessing Kudzaishe Sibanda, Godson Koffi Kalipe, Jonathan Mukiibi, Salomon Kabongo Kabenamualu, Mmasibidi Setaka, Lolwethu Ndolela, Nkiruka Odu, Rooweither Mabuya, Shamsuddeen Hassan Muhammad, Salomey Osei, Sokhar Samb, Dietrich Klakow and David Ifeoluwa Adelani	9429
<i>Boosting Long-Context Information Seeking via Query-Guided Activation Refilling</i> Hongjin Qian, Zheng Liu, Peitian Zhang, Zhicheng Dou and Defu Lian	9453
<i>Efficient Pretraining Data Selection for Language Models via Multi-Actor Collaboration</i> Tianyi Bai, Ling Yang, Zhen Hao Wong, Fupeng Sun, Xinlin Zhuang, Jiahui Peng, Chi Zhang, Lijun Wu, Qiu Jiantao, Wentao Zhang, Binhang Yuan and Conghui He	9465
<i>AdaDHP: Fine-Grained Fine-Tuning via Dual Hadamard Product and Adaptive Parameter Selection</i> Han Liu, Changya Li, Xiaotong Zhang, Feng Zhang, Fenglong Ma, Wei Wang and Hong Yu	9492

<i>KG-Agent: An Efficient Autonomous Agent Framework for Complex Reasoning over Knowledge Graph</i> Jinhao Jiang, Kun Zhou, Xin Zhao, Yang Song, Chen Zhu, Hengshu Zhu and Ji-Rong Wen	9505
<i>Curriculum Debiasing: Toward Robust Parameter-Efficient Fine-Tuning Against Dataset Biases</i> Mingyu Lee, Yeachan Kim, Wing-Lam Mok and SangKeun Lee	9524
<i>Does Context Matter? ContextualJudgeBench for Evaluating LLM-based Judges in Contextual Settings</i> Austin Xu, Srijan Bansal, Yifei Ming, Semih Yavuz and Shafiq Joty	9541
<i>On the Reliability of Large Language Models for Causal Discovery</i> Tao Feng, Lizhen Qu, Niket Tandon, Zhuang Li, Xiaoxi Kang and Gholamreza Haffari	9565
<i>Value-Spectrum: Quantifying Preferences of Vision-Language Models via Value Decomposition in Social Media Contexts</i> Jingxuan Li, Yuning Yang, Shengqi Yang, Linfan Zhang and Ying Nian Wu	9591
<i>TeRDy: Temporal Relation Dynamics through Frequency Decomposition for Temporal Knowledge Graph Completion</i> Ziyang Liu and Chaokun Wang	9611
<i>Incorporating Domain Knowledge into Materials Tokenization</i> Yerim Oh, Jun-Hyung Park, Junho Kim, SungHo Kim and SangKeun Lee	9623
<i>PIG: Privacy Jailbreak Attack on LLMs via Gradient-based Iterative In-Context Optimization</i> Yidan Wang, Yanan Cao, Yubing Ren, Fang Fang, Zheng Lin and Binxing Fang	9645
<i>Agents Under Siege: Breaking Pragmatic Multi-Agent LLM Systems with Optimized Prompt Attacks</i> Rana Shahroz, Zhen Tan, Sukwon Yun, Charles Fleming and Tianlong Chen	9661
<i>Semantic-Eval : A Semantic Comprehension Evaluation Framework for Large Language Models Generation without Training</i> Shusheng Li, Jiale Li, Yifei Qu, Xinwei Shi, Yanliang Guo, Ziyi He, Yubo Wang and Wenjun Tan	9675
<i>Between Circuits and Chomsky: Pre-pretraining on Formal Languages Imparts Linguistic Biases</i> Michael Y. Hu, Jackson Petty, Chuan Shi, William Merrill and Tal Linzen	9691
<i>When to Speak, When to Abstain: Contrastive Decoding with Abstention</i> Hyuhng Joon Kim, Youna Kim, Sang-goo Lee and Taeuk Kim	9710
<i>On the Risk of Evidence Pollution for Malicious Social Text Detection in the Era of LLMs</i> Herun Wan, Minnan Luo, Zhixiong Su, Guang Dai and Xiang Zhao	9731
<i>Investigating and Extending Homans' Social Exchange Theory with Large Language Model based Agents</i> Lei Wang, Zheqing Zhang and Xu Chen	9762
<i>A Drop-In Solution for On-the-Fly Adaptation of Speculative Decoding in Large Language Models</i> Jiesong Liu, Brian Park and Xipeng Shen	9778
<i>If Attention Serves as a Cognitive Model of Human Memory Retrieval, What is the Plausible Memory Representation?</i> Ryo Yoshida, Shinnosuke Isono, Kohei Kajikawa, Taiga Someya, Yushi Sugimoto and Yohei Oseki	9795
<i>Aligning VLM Assistants with Personalized Situated Cognition</i> Yongqi Li, Shen Zhou, Xiaohu Li, Xin Miao, Jintao Wen, Mayi Xu, Jianhao Chen, Birong Pan, Hankun Kang, Yuanyuan Zhu, Ming Zhong and Tiejun Qian	9813

<i>Attention Entropy is a Key Factor: An Analysis of Parallel Context Encoding with Full-attention-based Pre-trained Language Models</i>	
Zhisong Zhang, Yan Wang, Xinting Huang, Tianqing Fang, Hongming Zhang, Chenlong Deng, Shuaiyi Li and Dong Yu	9840
<i>Faster Speculative Decoding via Effective Draft Decoder with Pruned Candidate Tree</i>	
Huanran Zheng and Xiaoling Wang	9856
<i>Selecting and Merging: Towards Adaptable and Scalable Named Entity Recognition with Large Language Models</i>	
Zhuojun Ding, Wei Wei and Chenghao Fan	9869
<i>Embracing Imperfection: Simulating Students with Diverse Cognitive Levels Using LLM-based Agents</i>	
Tao Wu, Jingyuan Chen, Wang Lin, Mengze Li, Yumeng Zhu, Ang Li, Kun Kuang and Fei Wu	9887
<i>CADReview: Automatically Reviewing CAD Programs with Error Detection and Correction</i>	
Jiali Chen, Xusen Hei, HongFei Liu, Yuancheng Wei, Zikun Deng, Jiayuan Xie, Yi Cai and Li Qing	9909
<i>Think&Cite: Improving Attributed Text Generation with Self-Guided Tree Search and Progress Reward Modeling</i>	
Junyi Li and Hwee Tou Ng	9928
<i>The Lawyer That Never Thinks: Consistency and Fairness as Keys to Reliable AI</i>	
Dana R Alsagheer, Abdulrahman Kamal, Mohammad Kamal, Cosmo Yang Wu and Weidong Shi	9943
<i>Polishing Every Facet of the GEM: Testing Linguistic Competence of LLMs and Humans in Korean</i>	
SungHo Kim, Nayeon Kim, Taehee Jeon and SangKeun Lee	9955
<i>SpeechFake: A Large-Scale Multilingual Speech Deepfake Dataset Incorporating Cutting-Edge Generation Methods</i>	
Wen Huang, Yanmei Gu, Zhiming Wang, Huijia Zhu and Yanmin Qian	9985
<i>ReflectionCoder: Learning from Reflection Sequence for Enhanced One-off Code Generation</i>	
Houxing Ren, Mingjie Zhan, Zhongyuan Wu, Aojun Zhou, Junting Pan and Hongsheng Li .	9999
<i>InvestAlign: Overcoming Data Scarcity in Aligning Large Language Models with Investor Decision-Making Processes Under Herd Behavior</i>	
Huisheng Wang, Zhuoshi Pan, Hangjing Zhang, Mingxiao Liu, Hanqing Gao and H. Vicky Zhao	10021
<i>Enhancing Neural Machine Translation Through Target Language Data: A kNN-LM Approach for Domain Adaptation</i>	
Abudurexiti Rehemani, Hongyu Liu, Junhao Ruan, Abudukeyumu Abudula, Yingfeng Luo, Tong Xiao and JingBo Zhu	10053
<i>Multi-level Relevance Document Identifier Learning for Generative Retrieval</i>	
Fuwei Zhang, Xiaoyu Liu, Xinyu Jia, Yingfei Zhang, Shuai Zhang, Xiang Li, Fuzhen Zhuang, Wei Lin and Zhao Zhang	10066
<i>EfficientQAT: Efficient Quantization-Aware Training for Large Language Models</i>	
Mengzhao Chen, Wenqi Shao, Peng Xu, Jiahao Wang, Peng Gao, Kaipeng Zhang and Ping Luo	10081

<i>Exploring How Generative MLLMs Perceive More Than CLIP with the Same Vision Encoder</i> Siting Li, Pang Wei Koh and Simon Shaolei Du	10101
<i>NexusSum: Hierarchical LLM Agents for Long-Form Narrative Summarization</i> Hyuntak Kim and Byung-Hak Kim	10120
<i>HAIC: Improving Human Action Understanding and Generation with Better Captions for Multi-modal Large Language Models</i> Xiao Wang, Jingyun Hua, Weihong Lin, Yuanxing Zhang, Fuzheng Zhang, Jianlong Wu, Di Zhang and Liqiang Nie	10158
<i>Uni-Retrieval: A Multi-Style Retrieval Framework for STEM’s Education</i> Yanhao Jia, Xinyi Wu, Li Hao, QinglinZhang QinglinZhang, Yuxiao Hu, Shuai Zhao and Wenqi Fan	10182
<i>DenseLoRA: Dense Low-Rank Adaptation of Large Language Models</i> Lin Mu, Xiaoyu Wang, Li Ni, Yang Li, Zhize Wu, Peiquan Jin and Yiwen Zhang	10198
<i>Exploring the Potential of LLMs as Personalized Assistants: Dataset, Evaluation, and Analysis</i> Jisoo Mok, Ik-hwan Kim, Sangkwon Park and Sungroh Yoon	10212
<i>Cracking Factual Knowledge: A Comprehensive Analysis of Degenerate Knowledge Neurons in Large Language Models</i> Yuheng Chen, Pengfei Cao, Yubo Chen, Yining Wang, Shengping Liu, Kang Liu and Jun Zhao	10240
<i>Towards Context-Robust LLMs: A Gated Representation Fine-tuning Approach</i> Shenglai Zeng, Pengfei He, Kai Guo, Tianqi Zheng, Hanqing Lu, Yue Xing and Hui Liu ..	10262
<i>On Support Samples of Next Word Prediction</i> Yuqian Li, Yupei Du, Yufang Liu, Feifei Feng, Mou Xiao Feng and Yuanbin Wu	10277
<i>WebWalker: Benchmarking LLMs in Web Traversal</i> Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie and Fei Huang	10290
<i>From Trade-off to Synergy: A Versatile Symbiotic Watermarking Framework for Large Language Models</i> Yidan Wang, Yubing Ren, Yanan Cao and Binxing Fang	10306
<i>AutoGUI: Scaling GUI Grounding with Automatic Functionality Annotations from LLMs</i> Hongxin Li, Jingfan Chen, Jingran Su, Yuntao Chen, Li Qing and Zhaoxiang Zhang	10323
<i>Introducing Graph Context into Language Models through Parameter-Efficient Fine-Tuning for Lexical Relation Mining</i> Jingwen Sun, Zhiyi Tian, Yu He, Jingwei Sun and Guangzhong Sun	10359
<i>S-RAG: A Novel Audit Framework for Detecting Unauthorized Use of Personal Data in RAG Systems</i> Zhirui Zeng, Jiamou Liu, Meng-Fen Chiang, Jialing He and Zijian Zhang	10375
<i>Praetor: A Fine-Grained Generative LLM Evaluator with Instance-Level Customizable Evaluation Criteria</i> Yongqi Leng, Renren Jin, Yue Chen, Zhuowen Han, Ling Shi, Jianxiang Peng, Lei Yang, Juesi Xiao and Deyi Xiong	10386
<i>Mitigating Confounding in Speech-Based Dementia Detection through Weight Masking</i> Zhecheng Sheng, Xiruo Ding, Brian Hur, Changye Li, Trevor Cohen and Serguei V. S. Pakhomov	10419

<i>MCS-Bench: A Comprehensive Benchmark for Evaluating Multimodal Large Language Models in Chinese Classical Studies</i>	Yang Liu, Jiahuan Cao, Hiuyi Cheng, Yongxin Shi, Kai Ding and Lianwen Jin	10435
<i>The Knowledge Microscope: Features as Better Analytical Lenses than Neurons</i>	Yuheng Chen, Pengfei Cao, Kang Liu and Jun Zhao	10493
<i>From Real to Synthetic: Synthesizing Millions of Diversified and Complicated User Instructions with Attributed Grounding</i>	Chiwei Zhu, Benfeng Xu, Xiaorui Wang and Zhendong Mao	10516
<i>PrivaCI-Bench: Evaluating Privacy with Contextual Integrity and Legal Compliance</i>	Haoran Li, Wenbin Hu, Huihao Jing, Yulin Chen, Qi Hu, Sirui Han, Tianshu Chu, Peizhao Hu and Yangqiu Song	10544
<i>Unveiling Environmental Impacts of Large Language Model Serving: A Functional Unit View</i>	Yanran Wu, Inez Hua and Yi Ding	10560
<i>ExpeTrans: LLMs Are Experiential Transfer Learners</i>	Jinglong Gao, Xiao Ding, Lingxiao Zou, Bibo Cai, Bing Qin and Ting Liu	10577
<i>Cool-Fusion: Fuse Large Language Models without Training</i>	Cong Liu, Xiaojun Quan, Yan Pan, Weigang Wu, Xu Chen and Liang Lin	10617
<i>DAPE V2: Process Attention Score as Feature Map for Length Extrapolation</i>	Chuanyang Zheng, Yihang Gao, Han Shi, Jing Xiong, Jiankai Sun, Jingyao Li, Minbin Huang, Xiaozhe Ren, Michael Ng, Xin Jiang, Zhenguo Li and Yu Li	10628
<i>MuSC: Improving Complex Instruction Following with Multi-granularity Self-Contrastive Training</i>	Hui Huang, Jiaheng Liu, Yancheng He, Shilong Li, Bing Xu, Conghui Zhu, Muyun Yang and Tiejun Zhao	10667
<i>LongReD: Mitigating Short-Text Degradation of Long-Context Large Language Models via Restoration Distillation</i>	Zican Dong, Junyi Li, Jinhao Jiang, Mingyu Xu, Xin Zhao, Bingning Wang and Weipeng Chen	10687
<i>APB: Accelerating Distributed Long-Context Inference by Passing Compressed Context Blocks across GPUs</i>	Yuxiang Huang, Mingye Li, Xu Han, Chaojun Xiao, Weilin Zhao, Sun Ao, Hao Zhou, Jie Zhou, Zhiyuan Liu and Maosong Sun	10708
<i>PPT: A Minor Language News Recommendation Model via Cross-Lingual Preference Pattern Transfer</i>	Yiyang Zhang and Nan Chen	10728
<i>GainRAG: Preference Alignment in Retrieval-Augmented Generation through Gain Signal Synthesis</i>	Yi Jiang, Sendong Zhao, Jianbo Li, Haochun Wang and Bing Qin	10746
<i>Top-$n\sigma$: Eliminating Noise in Logit Space for Robust Token Sampling of LLM</i>	Chenxia Tang, Jianchun Liu, Hongli Xu and Liusheng Huang	10758
<i>SCOPE: Optimizing Key-Value Cache Compression in Long-context Generation</i>	Jialong Wu, Zhenglin Wang, Linhai Zhang, Yilong Lai, Yulan He and Deyu Zhou	10775
<i>Mitigating Non-Representative Prototypes and Representation Bias in Few-Shot Continual Relation Extraction</i>	Thanh Duc Pham, Nam Le Hai, Linh Ngo Van, Nguyen Thi Ngoc Diep, Sang Dinh and Thien Huu Nguyen	10791

<i>MoQAE: Mixed-Precision Quantization for Long-Context LLM Inference via Mixture of Quantization-Aware Experts</i>	
Wei Tao, Haocheng Lu, Xiaoyang Qu, Bin Zhang, Kai Lu, Jiguang Wan and Jianzong Wang	10810
<i>PrivacyRestore: Privacy-Preserving Inference in Large Language Models via Privacy Removal and Restoration</i>	
Ziqian Zeng, Jianwei Wang, Junyao Yang, Zhengdong Lu, Haoran Li, Huiping Zhuang and Cen Chen	10821
<i>Meta-rater: A Multi-dimensional Data Selection Method for Pre-training Language Models</i>	
Xinlin Zhuang, Jiahui Peng, Ren Ma, Yinfan Wang, Tianyi Bai, Xingjian Wei, Qiu Jiantao, Chi Zhang, Ying Qian and Conghui He	10856
<i>GuessArena: Guess Who I Am? A Self-Adaptive Framework for Evaluating LLMs in Domain-Specific Knowledge and Reasoning</i>	
Qingchen Yu, Zifan Zheng, Ding Chen, Simin Niu, Bo Tang, Feiyu Xiong and Zhiyu li	10897
<i>Sample-Efficient Human Evaluation of Large Language Models via Maximum Discrepancy Competition</i>	
Kehua Feng, Keyan Ding, Tan Hongzhi, Kede Ma, Zhihua Wang, Shuangquan Guo, Cheng Yuzhou, Ge Sun, Guozhou Zheng, Qiang Zhang and Huajun Chen	10913
<i>DTCRS: Dynamic Tree Construction for Recursive Summarization</i>	
Guanran Luo, Zhongquan Jian, Wentao Qiu, Meihong Wang and Qingqiang Wu	10948
<i>A Generative Adaptive Replay Continual Learning Model for Temporal Knowledge Graph Reasoning</i>	
Zhiyu Zhang, Wei Chen, Youfang Lin and Huaiyu Wan	10964
<i>ARise: Towards Knowledge-Augmented Reasoning via Risk-Adaptive Search</i>	
Yize Zhang, Tianshu Wang, Sirui Chen, Kun Wang, Xingyu Zeng, Hongyu Lin, Xianpei Han, Le Sun and Chaochao Lu	10978
<i>PKAG-DDI: Pairwise Knowledge-Augmented Language Model for Drug-Drug Interaction Event Text Generation</i>	
Ziyan Wang, Zhankun Xiong, Feng Huang and Wen Zhang	10996
<i>Knowledge-Augmented Multimodal Clinical Rationale Generation for Disease Diagnosis with Small Language Models</i>	
Shuai Niu, Jing Ma, Hongzhan Lin, Liang Bai, Zhihua Wang, Richard Yi Da Xu, Yunya Song and Xian Yang	11011
<i>TWIST: Text-encoder Weight-editing for Inserting Secret Trojans in Text-to-Image Models</i>	
Xindi Li, Zhe Liu, Tong Zhang, Jiahao Chen, Qingming Li, Jinbao Li and Shouling Ji	11025
<i>Frictional Agent Alignment Framework: Slow Down and Don't Break Things</i>	
Abhijnan Nath, Carine Graff, Andrei Bachinin and Nikhil Krishnaswamy	11042
<i>Powerformer: Efficient and High-Accuracy Privacy-Preserving Language Model with Homomorphic Encryption</i>	
Dongjin Park, Eunsang Lee and Joon-Woo Lee	11090
<i>Beware of Your Po! Measuring and Mitigating AI Safety Risks in Role-Play Fine-Tuning of LLMs</i>	
Weixiang Zhao, Yulin Hu, Yang Deng, Jiahe Guo, Xingyu Sui, Xinyang Han, An Zhang, Yanyan Zhao, Bing Qin, Tat-Seng Chua and Ting Liu	11112

<i>Can Graph Neural Networks Learn Language with Extremely Weak Text Supervision?</i> Zihao Li, Lecheng Zheng, Bowen Jin, Dongqi Fu, Baoyu Jing, Yikun Ban, Jingrui He and Jiawei Han	11138
<i>Towards Enhanced Immersion and Agency for LLM-based Interactive Drama</i> Hongqiu Wu, Weiqi Wu, Tianyang Xu, Jiameng Zhang and Hai Zhao	11166
<i>Disambiguating Reference in Visually Grounded Dialogues through Joint Modeling of Textual and Multimodal Semantic Structures</i> Shun Inadumi, Nobuhiro Ueda and Koichiro Yoshino	11183
<i>Improving Factuality with Explicit Working Memory</i> Mingda Chen, Yang Li, Karthik Padthe, Rulin Shao, Alicia Yi Sun, Luke Zettlemoyer, Gargi Ghosh and Wen-tau Yih	11199
<i>Gradient-Adaptive Policy Optimization: Towards Multi-Objective Alignment of Large Language Models</i> Chengao Li, Hanyu Zhang, Yunkun Xu, Hongyan Xue, Xiang Ao and Qing He	11214
<i>Dynamic Parallel Tree Search for Efficient LLM Reasoning</i> Yifu Ding, Wentao Jiang, Shunyu Liu, Yongcheng Jing, Jinyang Guo, Yingjie Wang, Jing Zhang, Zengmao Wang, Ziwei Liu, Bo Du, Xianglong Liu and Dacheng Tao	11233
<i>Pre³: Enabling Deterministic Pushdown Automata for Faster Structured LLM Generation</i> Junyi Chen, Shihao Bai, Zaijun Wang, Siyu Wu, Chuheng Du, Hailong Yang, Ruihao Gong, Shengzhong Liu, Fan Wu and Guihai Chen	11253
<i>SHARE: An SLM-based Hierarchical Action CorREction Assistant for Text-to-SQL</i> Ge Qu, Jinyang Li, Bowen Qin, Xiaolong Li, Nan Huo, Chenhao Ma and Reynold Cheng ..	11268
<i>GenderAlign: An Alignment Dataset for Mitigating Gender Bias in Large Language Models</i> Tao Zhang, Ziqian Zeng, YuxiangXiao YuxiangXiao, Huiping Zhuang, Cen Chen, James R. Foulds and Shimei Pan	11293
<i>Large Language and Protein Assistant for Protein-Protein Interactions Prediction</i> Peng Zhou, Pengsen Ma, Jianmin Wang, Xibao Cai, Haitao Huang, Wei Liu, Longyue Wang, Lai Hou Tim and Xiangxiang Zeng	11312
<i>An Empirical Study of Many-to-Many Summarization with Large Language Models</i> Jiaan Wang, Fandong Meng, Zengkui Sun, Yunlong Liang, Yuxuan Cao, Jiarong Xu, Haoxiang Shi and Jie Zhou	11328
<i>Locate-and-Focus: Enhancing Terminology Translation in Speech Language Models</i> Suhang Wu, Jialong Tang, Chengyi Yang, Pei Zhang, Baosong Yang, Junhui Li, Junfeng Yao, Min Zhang and Jinsong Su	11345
<i>GuideBench: Benchmarking Domain-Oriented Guideline Following for LLM Agents</i> Lingxiao Diao, Xinyue Xu, Wanxuan Sun, Cheng Yang and Zhuosheng Zhang	11361
<i>TC-RAG: Turing-Complete RAG's Case study on Medical LLM Systems</i> Xinke Jiang, Yue Fang, Rihong Qiu, Haoyu Zhang, Yongxin Xu, Hao Chen, Wentao Zhang, Ruizhe Zhang, Yuchen Fang, Xinyu Ma, Xu Chu, Junfeng Zhao and Yasha Wang	11400
<i>SoRFT: Issue Resolving with Subtask-oriented Reinforced Fine-Tuning</i> Zexiong Ma, Chao Peng, Pengfei Gao, Xiangxin Meng, Yanzhen Zou and Bing Xie	11427

<i>MiniLongBench: The Low-cost Long Context Understanding Benchmark for Large Language Models</i> Zhongzhan Huang, Guoming Ling, Shanshan Zhong, Hefeng Wu and Liang Lin	11442
<i>Divide-Then-Align: Honest Alignment based on the Knowledge Boundary of RAG</i> Xin Sun, Jianan Xie, Zhongqi Chen, Qiang Liu, Shu Wu, Yuehe Chen, Bowen Song, Zilei Wang, Weiqiang Wang and Liang Wang	11461
<i>PwnGPT: Automatic Exploit Generation Based on Large Language Models</i> Wanzong Peng, Lin Ye, Xuetao Du, Hongli Zhang, Dongyang Zhan, Yunting Zhang, Yicheng Guo and Chen Zhang	11481
<i>VMLU Benchmarks: A comprehensive benchmark toolkit for Vietnamese LLMs</i> Cuc Thi Bui, Nguyen Truong Son, Truong Van Trang, Lam Viet Phung, Pham Nhut Huy, Hoang Anh Le, Quoc Huu Van, Phong Nguyen-Thuan Do, Van Le Tran Truc, Duc Thanh Chau and Le-Minh Nguyen	11495
<i>Scaling up the State Size of RNN LLMs for Long-Context Scenarios</i> Kai Liu, Jianfei Gao and Kai Chen	11516
<i>Unifying Continuous and Discrete Text Diffusion with Non-simultaneous Diffusion Processes</i> Bocheng Li, Zhujin Gao and Linli Xu	11530
<i>A Strategic Coordination Framework of Small LMs Matches Large LMs in Data Synthesis</i> Xin Gao, Qizhi Pei, Zinan Tang, Yu Li, Honglin Lin, Jiang Wu, Lijun Wu and Conghui He	11552
<i>Defining and Evaluating Visual Language Models' Basic Spatial Abilities: A Perspective from Psycho- metrics</i> Wenrui Xu, Dalin Lyu, Weihang Wang, Jie Feng, Chen Gao and Yong Li	11571
<i>SPHERE: Unveiling Spatial Blind Spots in Vision-Language Models Through Hierarchical Evaluation</i> Wenyu Zhang, Wei En Ng, Lixin Ma, Yuwen Wang, Junqi Zhao, Allison Koenecke, Boyang Li and Wanglu Wanglu	11591
<i>User-side Model Consistency Monitoring for Open Source Large Language Models Inference Services</i> Qijun Miao and Zhixuan Fang	11610
<i>Jailbreaking? One Step Is Enough!</i> Weixiong Zheng, Peijian Zeng, YiWei Li, Hongyan Wu, Nankai Lin, Junhao Chen, Aimin Yang and Yongmei Zhou	11623
<i>Parenting: Optimizing Knowledge Selection of Retrieval-Augmented Language Models with Parameter Decoupling and Tailored Tuning</i> Yongxin Xu, Ruizhe Zhang, Xinke Jiang, Yujie Feng, Yuzhen Xiao, Xinyu Ma, Runchuan Zhu, Xu Chu, Junfeng Zhao and Yasha Wang	11643
<i>PaSa: An LLM Agent for Comprehensive Academic Paper Search</i> Yichen He, Guanhua Huang, Peiyuan Feng, Yuan Lin, Yuchen Zhang, Hang Li and Weinan E	11663
<i>Less Mature is More Adaptable for Sentence-level Language Modeling</i> Abhilasha Sancheti, David Dale, Artyom Kozhevnikov and Maha Elbayad	11680
<i>EpMAN: Episodic Memory Attention for Generalizing to Longer Contexts</i> Subhajit Chaudhury, Payel Das, Sarathkrishna Swaminathan, Georgios Kollias, Elliot Nelson, Khushbu Pahwa, Tejaswini Pedapati, Igor Melnyk and Matthew Riemer	11696

<i>UORA: Uniform Orthogonal Reinitialization Adaptation in Parameter Efficient Fine-Tuning of Large Models</i>	
Xueyan Zhang, Jinman Zhao, Zhifei Yang, Yibo Zhong, Shuhao Guan, Linbo Cao and Yining Wang	11709
<i>Agri-CM³: A Chinese Massive Multi-modal, Multi-level Benchmark for Agricultural Understanding and Reasoning</i>	
Haotian Wang, Yi Guan, Fanshu Meng, Chao Zhao, Lian Yan, Yang Yang and Jingchi Jiang	11729
<i>TROVE: A Challenge for Fine-Grained Text Provenance via Source Sentence Tracing and Relationship Classification</i>	
Junnan Zhu, Min Xiao, Yining Wang, Feifei Zhai, Yu Zhou and Chengqing Zong	11755
<i>CaLMQA: Exploring culturally specific long-form question answering across 23 languages</i>	
Shane Arora, Marzena Karpinska, Hung-Ting Chen, Ipsita Bhattacharjee, Mohit Iyyer and Eunsol Choi	11772
<i>Croppable Knowledge Graph Embedding</i>	
Yushan Zhu, Wen Zhang, Zhiqiang Liu, Mingyang Chen, Lei Liang and Huajun Chen ...	11818
<i>HyKGE: A Hypothesis Knowledge Graph Enhanced RAG Framework for Accurate and Reliable Medical LLMs Responses</i>	
Xinke Jiang, Ruizhe Zhang, Yongxin Xu, Rihong Qiu, Yue Fang, Zhiyuan Wang, Jinyi Tang, Hongxin Ding, Xu Chu, Junfeng Zhao and Yasha Wang	11836
<i>LongRecipe: Recipe for Efficient Long Context Generalization in Large Language Models</i>	
Zhiyuan Hu, Yuliang Liu, Jinman Zhao, Suyuchen Wang, WangYan WangYan, Wei Shen, Qing Gu, Anh Tuan Luu, See-Kiong Ng, Zhiwei Jiang and Bryan Hooi	11857
<i>BeamLoRA: Beam-Constraint Low-Rank Adaptation</i>	
Naibin Gu, Zhenyu Zhang, Xiyu Liu, Peng Fu, Zheng Lin, Shuohuan Wang, Yu Sun, Hua Wu, Weiping Wang and Haifeng Wang	11871
<i>GODBench: A Benchmark for Multimodal Large Language Models in Video Comment Art</i>	
Yiming Lei, Chenkai Zhang, Zeming Liu, Haitao Leng, ShaoGuo Liu, Tingting Gao, Qingjie Liu and Yunhong Wang	11884
<i>UniLR: Unleashing the Power of LLMs on Multiple Legal Tasks with a Unified Legal Retriever</i>	
Ang Li, Yiquan Wu, Yifei Liu, Ming Cai, Lizhi Qing, Shihang Wang, Yangyang Kang, Chengyuan Liu, Fei Wu and Kun Kuang	11953
<i>Generative Psycho-Lexical Approach for Constructing Value Systems in Large Language Models</i>	
Haoran Ye, TianZe Zhang, Yuhang Xie, Liyuan Zhang, Yuanyi Ren, Xin Zhang and Guojie Song	11968
<i>Beyond Dialogue: A Profile-Dialogue Alignment Framework Towards General Role-Playing Language Model</i>	
Yeyong Yu, Runsheng Yu, Haojie Wei, Zhanqiu Zhang and Quan Qian	11992
<i>ACECODER: Acing Coder RL via Automated Test-Case Synthesis</i>	
Huaye Zeng, Dongfu Jiang, Haozhe Wang, Ping Nie, Xiaotong Chen and Wenhua Chen ...	12023
<i>Quantifying Semantic Emergence in Language Models</i>	
Hang Chen, Xinyu Yang, Jiaying Zhu and Wenya Wang	12041

<i>DebateCoder: Towards Collective Intelligence of LLMs via Test Case Driven LLM Debate for Code Generation</i>	
Jizheng Chen, Kounianhua Du, Xinyi Dai, Weiming Zhang, Xihuai Wang, Yasheng Wang, Ruiming Tang, Weinan Zhang and Yong Yu	12055
<i>The Tug of War Within: Mitigating the Fairness-Privacy Conflicts in Large Language Models</i>	
Chen Qian, Dongrui Liu, Jie Zhang, Yong Liu and Jing Shao	12066
<i>GraphInsight: Unlocking Insights in Large Language Models for Graph Structure Understanding</i>	
Yukun Cao, Shuo Han, Zengyi Gao, Zezhong Ding, Xike Xie and S Kevin Zhou	12096
<i>Phonotomizer: A Compact, Unsupervised, Online Training Approach to Real-Time, Multilingual Phonetic Segmentation</i>	
Michael S. Yantosca and Albert M. K. Cheng	12135
<i>A Multi-persona Framework for Argument Quality Assessment</i>	
Bojun Jin, Jianzhu Bao, Yufang Hou, Yang Sun, Yice Zhang, Huajie Wang, Bin Liang and Ruifeng Xu	12148
<i>Safe: Enhancing Mathematical Reasoning in Large Language Models via Retrospective Step-aware Formal Verification</i>	
Chengwu Liu, Ye Yuan, Yichun Yin, Yan Xu, Xin Xu, Zaoyu Chen, Yasheng Wang, Lifeng Shang, Qun Liu and Ming Zhang	12171
<i>SAM Decoding: Speculative Decoding via Suffix Automaton</i>	
Yuxuan Hu, Ke Wang, Xiaokang Zhang, Fanjin Zhang, Cuiping Li, Hong Chen and Jing Zhang	12187
<i>PsyAdvisor: A Plug-and-Play Strategy Advice Planner with Proactive Questioning in Psychological Conversations</i>	
Yuxin Hu, Danni Liu, Bo Liu, Yida Chen, Jiuxin Cao and Yan Liu	12205
<i>HomeBench: Evaluating LLMs in Smart Homes with Valid and Invalid Instructions Across Single and Multiple Devices</i>	
Silin Li, Yuhang Guo, Jiashu Yao, Zeming Liu and Haifeng Wang	12230
<i>Advancing Zero-shot Text-to-Speech Intelligibility across Diverse Domains via Preference Alignment</i>	
Xueyao Zhang, Yuancheng Wang, Chaoren Wang, Ziniu Li, Zhuo Chen and Zhizheng Wu	12251
<i>GiFT: Gibbs Fine-Tuning for Code Generation</i>	
Haochen Li, Wanjin Feng, Xin Zhou and Zhiqi Shen	12271
<i>Enhancing Interpretable Image Classification Through LLM Agents and Conditional Concept Bottleneck Models</i>	
Yiwen Jiang, Deval Mehta, Wei Feng and Zongyuan Ge	12285
<i>Reliably Bounding False Positives: A Zero-Shot Machine-Generated Text Detection Framework via Multiscaled Conformal Prediction</i>	
Xiaowei Zhu, Yubing Ren, Yanan Cao, Xixun Lin, Fang Fang and Yangxi Li	12298
<i>RSCF: Relation-Semantics Consistent Filter for Entity Embedding of Knowledge Graph</i>	
Junsik Kim, Jinwook Park and Kangil Kim	12320
<i>RolePlot: A Systematic Framework for Evaluating and Enhancing the Plot-Progression Capabilities of Role-Playing Agents</i>	
Pinyi Zhang, Siyu An, Lingfeng Qiao, Yifei Yu, Jingyang Chen, Jie Wang, di Yin, Xing Sun and Kai Zhang	12337

<i>TreeRL: LLM Reinforcement Learning with On-Policy Tree Search</i>	
Zhenyu Hou, Ziniu Hu, Yujiang Li, Rui Lu, Jie Tang and Yuxiao Dong	12355
<i>Can a Single Model Master Both Multi-turn Conversations and Tool Use? CoALM: A Unified Conversational Agentic Language Model</i>	
Emre Can Acikgoz, Jeremiah Greer, Akul Datta, Ze Yang, William Zeng, Oussama Elachqar, Emmanouil Koukoumidis, Dilek Hakkani-Tür and Gokhan Tur	12370
<i>Single-to-mix Modality Alignment with Multimodal Large Language Model for Document Image Machine Translation</i>	
Yupu Liang, Yaping Zhang, Zhiyang Zhang, Yang Zhao, Lu Xiang, Chengqing Zong and Yu Zhou	12391
<i>SDPO: Segment-Level Direct Preference Optimization for Social Agents</i>	
Aobo Kong, Wentao Ma, Shiwan Zhao, Yongbin Li, Yuchuan Wu, Ke Wang, Xiaoqian Liu, Qicheng Li, Yong Qin and Fei Huang	12409
<i>KokoroChat: A Japanese Psychological Counseling Dialogue Dataset Collected via Role-Playing by Trained Counselors</i>	
Zhiyang Qi, Takumasa Kaneko, Keiko Takamizo, Mariko Ukiyo and Michimasa Inaba	12424
<i>SURVEYFORGE : On the Outline Heuristics, Memory-Driven Generation, and Multi-dimensional Evaluation for Automated Survey Writing</i>	
Xiangchao Yan, Shiyang Feng, Jiakang Yuan, Renqiu Xia, Bin Wang, Lei Bai and Bo Zhang	12444
<i>Making LLMs Better Many-to-Many Speech-to-Text Translators with Curriculum Learning</i>	
Yexing Du, Youcheng Pan, Ziyang Ma, Bo Yang, Yifan Yang, Keqi Deng, Xie Chen, Yang Xiang, Ming Liu and Bing Qin	12466
<i>AbGen: Evaluating Large Language Models in Ablation Study Design and Evaluation for Scientific Research</i>	
Yilun Zhao, Weiyuan Chen, Zhijian Xu, Manasi Patwardhan, Chengye Wang, Yixin Liu, Lovekesh Vig and Arman Cohan	12479
<i>Redundancy Principles for MLLMs Benchmarks</i>	
Zicheng Zhang, Xiangyu Zhao, Xinyu Fang, Chunyi Li, Xiaohong Liu, Xiongkuo Min, Haodong Duan, Kai Chen and Guangtao Zhai	12492
<i>WavRAG: Audio-Integrated Retrieval Augmented Generation for Spoken Dialogue Models</i>	
Yifu Chen, Shengpeng Ji, Haoxiao Wang, Ziqing Wang, Siyu Chen, Jinzheng He, Jin Xu and Zhou Zhao	12505
<i>ChildMandarin: A Comprehensive Mandarin Speech Dataset for Young Children Aged 3-5</i>	
Jiaming Zhou, Shiyao Wang, Shiwan Zhao, Jiabei He, Haoqin Sun, Hui Wang, Cheng Liu, Aobo Kong, Yujie Guo, Xi Yang, Yequan Wang, Yonghua Lin and Yong Qin	12524
<i>Finding the Sweet Spot: Preference Data Construction for Scaling Preference Optimization</i>	
Yao Xiao, Hai Ye, Linyao Chen, Hwee Tou Ng, Lidong Bing, Xiaoli Li and Roy Ka-Wei Lee	12538
<i>Enhancing Safe and Controllable Protein Generation via Knowledge Preference Optimization</i>	
Yuhao Wang, Keyan Ding, Kehua Feng, Zeyuan Wang, Ming Qin, Xiaotong Li, Qiang Zhang and Huajun Chen	12553
<i>SINCon: Mitigate LLM-Generated Malicious Message Injection Attack for Rumor Detection</i>	
Mingqing Zhang, Qiang Liu, Xiang Tao, Shu Wu and Liang Wang	12570

<i>Outlier-Safe Pre-Training for Robust 4-Bit Quantization of Large Language Models</i>	
Jungwoo Park, Taewhoo Lee, Chanwoong Yoon, Hyeon Hwang and Jaewoo Kang	12582
<i>Agentic Knowledgeable Self-awareness</i>	
Shuofei Qiao, Zhisong Qiu, Baochang Ren, Xiaobin Wang, Xiangyuan Ru, Ningyu Zhang, Xiang Chen, Yong Jiang, Pengjun Xie, Fei Huang and Huajun Chen	12601
<i>A Unified Agentic Framework for Evaluating Conditional Image Generation</i>	
Jifang Wang, Yangxue Yangxue, Longyue Wang, Zhenran Xu, Yiyu Wang, Yaowei Wang, Weihua Luo, Kaifu Zhang, Baotian Hu and Min Zhang	12626
<i>Planning-Driven Programming: A Large Language Model Programming Workflow</i>	
Chao Lei, Yanchuan Chang, Nir Lipovetzky and Krista A. Ehinger	12647
<i>Can Knowledge Graphs Make Large Language Models More Trustworthy? An Empirical Study Over Open-ended Question Answering</i>	
Yuan Sui, Yufei He, Zifeng Ding and Bryan Hooi	12685
<i>Nudging: Inference-time Alignment of LLMs via Guided Decoding</i>	
Yu Fei, Yasaman Razeghi and Sameer Singh	12702
<i>Unveiling Attractor Cycles in Large Language Models: A Dynamical Systems View of Successive Paraphrasing</i>	
Zhilin Wang, Yafu Li, Jianhao Yan, Yu Cheng and Yue Zhang	12740
<i>SCAR: Data Selection via Style Consistency-Aware Response Ranking for Efficient Instruction-Tuning of Large Language Models</i>	
Zhuang Li, Yuncheng Hua, Thuy-Trang Vu, Haolan Zhan, Lizhen Qu and Gholamreza Haffari	12756
<i>HFT: Half Fine-Tuning for Large Language Models</i>	
Tingfeng Hui, Zhenyu Zhang, Shuohuan Wang, Weiran Xu, Yu Sun and Hua Wu	12791
<i>Beyond Surface Simplicity: Revealing Hidden Reasoning Attributes for Precise Commonsense Diagnosis</i>	
Huijun Lian, Zekai Sun, Keqi Chen, Yingming Gao and Ya Li	12820
<i>From Objectives to Questions: A Planning-based Framework for Educational Mathematical Question Generation</i>	
Cheng Cheng, Zhenya Huang, GuanHao Zhao, Yuxiang Guo, Xin Lin, Jinze Wu, Xin Li and Shijin Wang	12836
<i>RankCoT: Refining Knowledge for Retrieval-Augmented Generation through Ranking Chain-of-Thoughts</i>	
Mingyan Wu, Zhenghao Liu, Yukun Yan, Xinze Li, Shi Yu, Zheni Zeng, Yu Gu and Ge Yu	12857
<i>Lost in Literalism: How Supervised Training Shapes Translationese in LLMs</i>	
Yafu Li, Ronghao Zhang, Zhilin Wang, Huajian Zhang, Leyang Cui, Yongjing Yin, Tong Xiao and Yue Zhang	12875
<i>Accurate KV Cache Quantization with Outlier Tokens Tracing</i>	
Yi Su, Yuechi Zhou, Quantong Qiu, Juntao Li, Qingrong Xia, Ping Li, Xinyu Duan, Zhefeng Wang and Min Zhang	12895
<i>Can Large Language Models Understand Internet Buzzwords Through User-Generated Content</i>	
Chen Huang, Junkai Luo, Xinzuo Wang, Wenqiang Lei and Jiancheng Lv	12916

<i>EAC-MoE: Expert-Selection Aware Compressor for Mixture-of-Experts Large Language Models</i> Yuanteng Chen, Yuantian Shao, Peisong Wang and Jian Cheng	12942
<i>Activation Steering Decoding: Mitigating Hallucination in Large Vision-Language Models through Bidirectional Hidden State Intervention</i> Jingran Su, Jingfan Chen, Hongxin Li, Yuntao Chen, Li Qing and Zhaoxiang Zhang	12964
<i>Interactive Evolution: A Neural-Symbolic Self-Training Framework For Large Language Models</i> Fangzhi Xu, Qiushi Sun, Kanzhi Cheng, Jun Liu, Yu Qiao and Zhiyong Wu	12975
<i>Improving Medical Large Vision-Language Models with Abnormal-Aware Feedback</i> Yucheng Zhou, Lingran Song and Jianbing Shen	12994
<i>Upcycling Instruction Tuning from Dense to Mixture-of-Experts via Parameter Merging</i> Tingfeng Hui, Zhenyu Zhang, Shuohuan Wang, Yu Sun, Hua Wu and Sen Su	13012
<i>MapNav: A Novel Memory Representation via Annotated Semantic Maps for VLM-based Vision-and-Language Navigation</i> Lingfeng Zhang, Xiaoshuai Hao, Qinwen Xu, Qiang Zhang, Xinyao Zhang, Pengwei Wang, Jing Zhang, Zhongyuan Wang, Shanghang Zhang and Renjing Xu	13032
<i>Exploring Compositional Generalization of Multimodal LLMs for Medical Imaging</i> Zhenyang Cai, Junying Chen, Rongsheng Wang, Weihong Wang, Yonglin Deng, Dingjie Song, Yize Chen, Zixu Zhang and Benyou Wang	13057
<i>CLAIM: Mitigating Multilingual Object Hallucination in Large Vision-Language Models with Cross-Lingual Attention Intervention</i> Zekai Ye, Qiming Li, Xiaocheng Feng, Libo Qin, Yichong Huang, Baohang Li, Kui Jiang, Yang Xiang, Zhirui Zhang, Yunfei Lu, Duyu Tang, Dandan Tu and Bing Qin	13080
<i>Wizard of Shopping: Target-Oriented E-commerce Dialogue Generation with Decision Tree Branching</i> Xiangci Li, Zhiyu Chen, Jason Ingyu Choi, Nikhita Vedula, Besnik Fetahu, Oleg Rokhlenko and Shervin Malmasi	13095
<i>Qwen2.5-xCoder: Multi-Agent Collaboration for Multilingual Code Instruction Tuning</i> Jian Yang, Wei Zhang, Yibo Miao, Shanghaoran Quan, Zhenhe Wu, Qiyao Peng, Liqun Yang, Tianyu Liu, Zeyu Cui, Binyuan Hui and Junyang Lin	13121
<i>Cultivating Gaming Sense for Yourself: Making VLMs Gaming Experts</i> Wenxuan lu, Jiangyang He, Zhanqiu Zhang, Steven Y. Guo and Tianning Zang	13132
<i>Genius: A Generalizable and Purely Unsupervised Self-Training Framework For Advanced Reasoning</i> Fangzhi Xu, Hang Yan, Chang Ma, Haiteng Zhao, Qiushi Sun, Kanzhi Cheng, Junxian He, Jun Liu and Zhiyong Wu	13153
<i>Extending Complex Logical Queries on Uncertain Knowledge Graphs</i> Weizhi Fei, Zihao Wang, Hang Yin, Yang Duan and Yangqiu Song	13168
<i>Knowledge Decoupling via Orthogonal Projection for Lifelong Editing of Large Language Models</i> Haoyu Xu, Pengxiang Lan, Enneng Yang, Guibing Guo, Jianzhe Zhao, Linying Jiang and Xingwei Wang	13194
<i>ϕ-Decoding: Adaptive Foresight Sampling for Balanced Inference-Time Exploration and Exploitation</i> Fangzhi Xu, Hang Yan, Chang Ma, Haiteng Zhao, Jun Liu, Qika Lin and Zhiyong Wu	13214
<i>Can LLM Watermarks Robustly Prevent Unauthorized Knowledge Distillation?</i> Leyi Pan, Aiwei Liu, Shiyu Huang, Yijian LU, Xuming Hu, Lijie Wen, Irwin King and Philip S. Yu	13228

<i>Rethinking Reward Model Evaluation Through the Lens of Reward Overoptimization</i> Sunghwan Kim, Dongjin Kang, Taeyoon Kwon, Hyungjoo Chae, Dongha Lee and Jinyoung Yeo	13252
<i>Inducing lexicons of in-group language with socio-temporal context</i> Christine de Kock	13281
<i>LLaSE-G1: Incentivizing Generalization Capability for LLaMA-based Speech Enhancement</i> Boyi Kang, Xinfu Zhu, Zihan Zhang, Zhen Ye, Mingshuai Liu, Ziqian Wang, Yike Zhu, Guobin Ma, Jun Chen, Longshuai Xiao, Chao Weng, Wei Xue and Lei Xie	13292
<i>MadaKV: Adaptive Modality-Perception KV Cache Eviction for Efficient Multimodal Long-Context Inference</i> Kunxi Li, Zhonghua Jiang, Zhouzhou Shen, ZhaodeWang ZhaodeWang, Chengfei lv, Shengyu Zhang, Fan Wu and Fei Wu	13306
<i>Efficient OpAmp Adaptation for Zoom Attention to Golden Contexts</i> Haoyuan Wu, Rui Ming, Haisheng Zheng, Zhuolun He and Bei Yu	13319
<i>Language-Codec: Bridging Discrete Codec Representations and Speech Language Models</i> Shengpeng Ji, Minghui Fang, Jialong Zuo, Ziyue Jiang, Dingdong Wang, Hanting Wang, Hai Huang and Zhou Zhao	13332
<i>Adaptive Tool Use in Large Language Models with Meta-Cognition Trigger</i> Wenjun Li, Dexun Li, Kuicai Dong, Cong Zhang, Hao Zhang, Weiwen Liu, Yasheng Wang, Ruiming Tang and Yong Liu	13346
<i>MMLU-CF: A Contamination-free Multi-task Language Understanding Benchmark</i> Qihao Zhao, Yangyu Huang, Tengchao Lv, Lei Cui, Qinzheng Sun, Shaoguang Mao, Xin Zhang, Ying Xin, Qiufeng Yin, Scarlett Li and Furu Wei	13371
<i>Code-Switching Red-Teaming: LLM Evaluation for Safety and Multilingual Understanding</i> Haneul Yoo, Yongjin Yang and Hwaran Lee	13392
<i>Unleashing LLM Reasoning Capability via Scalable Question Synthesis from Scratch</i> Yuyang Ding, Xinyu Shi, Xiaobo Liang, Juntao Li, Zhaopeng Tu, Qiaoming Zhu and Min Zhang	13414
<i>DREsS: Dataset for Rubric-based Essay Scoring on EFL Writing</i> Haneul Yoo, Jieun Han, So-Yeon Ahn and Alice Oh	13439
<i>PQR: Improving Dense Retrieval via Potential Query Modeling</i> Junfeng Kang, Rui Li, Qi Liu, Yanjiang Chen, Zheng Zhang, Junzhe Jiang, Heng Yu and Yu Su	13455
<i>Cross-Lingual Generalization and Compression: From Language-Specific to Shared Neurons</i> Frederick Riemenschneider and Anette Frank	13470
<i>SDBench: A Survey-based Domain-specific LLM Benchmarking and Optimization Framework</i> Cheng Guo, Hu Kai, Shuxian Liang, Yiyang Jiang, Yi Gao, Xian-Sheng Hua and Wei Dong	13492
<i>ReflecTool: Towards Reflection-Aware Tool-Augmented Clinical Agents</i> Yusheng Liao, Shuyang Jiang, Yanfeng Wang and Yu Wang	13507
<i>Lexical Recall or Logical Reasoning: Probing the Limits of Reasoning Abilities in Large Language Models</i> Henrike Beyer and Chris Reed	13532

<i>ChainEdit: Propagating Ripple Effects in LLM Knowledge Editing through Logical Rule-Guided Chains</i> Zilu Dong, Xiangqing Shen, Zinong Yang and Rui Xia	13558
<i>HiDe-LLaVA: Hierarchical Decoupling for Continual Instruction Tuning of Multimodal Large Language Model</i> Haiyang Guo, Fanhu Zeng, Ziwei Xiang, Fei Zhu, Da-Han Wang, Xu-Yao Zhang and Cheng-Lin Liu	13572
<i>Self-supervised Quantized Representation for Seamlessly Integrating Knowledge Graphs with Large Language Models</i> Qika Lin, Tianzhe Zhao, Kai He, Zhen Peng, Fangzhi Xu, Ling Huang, Jingying Ma and Mengling Feng	13587
<i>Finite State Automata Inside Transformers with Chain-of-Thought: A Mechanistic Study on State Tracking</i> Yifan Zhang, Wenyu Du, Dongming Jin, Jie Fu and Zhi Jin	13603
<i>TeamLoRA: Boosting Low-Rank Adaptation with Expert Collaboration and Competition</i> Tianwei Lin, Jiang Liu, Wenqiao Zhang, Yang Dai, Haoyuan Li, Zhelun Yu, Wangui He, Juncheng Li, Jiannan Guo, Hao Jiang, Siliang Tang and Yueting Zhuang	13622
<i>CRiskEval: A Chinese Multi-Level Risk Evaluation Benchmark Dataset for Large Language Models</i> Ling Shi and Deyi Xiong	13638
<i>STUN: Structured-Then-Unstructured Pruning for Scalable MoE Pruning</i> Jaeseong Lee, Seung-won Hwang, Aurick Qiao, Daniel F Campos, Zhewei Yao and Yuxiong He	13660
<i>Mimicking the Familiar: Dynamic Command Generation for Information Theft Attacks in LLM Tool-Learning System</i> Ziyong Jiang, Mingyang Li, Guowei Yang, Junjie Wang, Yuekai Huang, Zhiyuan Chang and Qing Wang	13677
<i>FlashAudio: Rectified Flow for Fast and High-Fidelity Text-to-Audio Generation</i> Huadai Liu, Jialei Wang, Rongjie Huang, Yang Liu, Heng Lu, Zhou Zhao and Wei Xue ...	13694
<i>How does Misinformation Affect Large Language Model Behaviors and Preferences?</i> Miao Peng, Nuo Chen, Jianheng Tang and Jia Li	13711
<i>YESciEval: Robust LLM-as-a-Judge for Scientific Question Answering</i> Jennifer D’Souza, Hamed Babaei Giglou and Quentin Münch	13749
<i>GALLa: Graph Aligned Large Language Models for Improved Source Code Understanding</i> Ziyin Zhang, Hang Yu, Sage Lee, Peng Di, Jianguo Li and Rui Wang	13784
<i>MEDDxAgent: A Unified Modular Agent Framework for Explainable Automatic Differential Diagnosis</i> Daniel Philip Rose, Chia-Chien Hung, Marco Lepri, Israa Alqassem, Kiril Gashteovski and Carolin Lawrence	13803
<i>A Training-free LLM-based Approach to General Chinese Character Error Correction</i> Houquan Zhou, Bo Zhang, Zhenghua Li, Ming Yan and Min Zhang	13827
<i>HSCR: Hierarchical Self-Contrastive Rewarding for Aligning Medical Vision Language Models</i> Songtao Jiang, Yan Zhang, Yeying Jin, Zhihang Tang, Yangyang Wu, Yang Feng, Jian Wu and Zuozhu Liu	13853

<i>MAMMO-TH-VL: Eliciting Multimodal Reasoning with Instruction Tuning at Scale</i>	
Jiawei Guo, Tianyu Zheng, Yizhi LI, Yuelin Bai, Bo Li, Yubo Wang, King Zhu, Graham Neubig, Wenhu Chen and Xiang Yue	13869
<i>SIFT-50M: A Large-Scale Multilingual Dataset for Speech Instruction Fine-Tuning</i>	
Prabhat Pandey, Rupak Vignesh Swaminathan, K V Vijay Girish, Arunasish Sen, Jian. Xie, Grant Strimel and Andreas Schwarz	13921
<i>Recent Advances in Speech Language Models: A Survey</i>	
Wenqian Cui, Dianzhi Yu, Xiaoqi Jiao, Ziqiao Meng, Guangyan Zhang, Qichao Wang, Steven Y. Guo and Irwin King	13943
<i>LexCLiPR: Cross-Lingual Paragraph Retrieval from Legal Judgments</i>	
Rohit Upadhyaya and Santosh T.y.s.s	13971
<i>Multi-task Adversarial Attacks against Black-box Model with Few-shot Queries</i>	
Wenqiang Wang, Yan Xiao, Hao Lin, Yangshijie Zhang and Xiaochun Cao	13994
<i>SPECTRA: Faster Large Language Model Inference with Optimized Internal and External Speculation</i>	
Nguyen-Khang Le, Truong Dinh Do and Le-Minh Nguyen	14015
<i>Multi-level Association Refinement Network for Dialogue Aspect-based Sentiment Quadruple Analysis</i>	
Zeliang Tong, Wei Wei, Xiaoye Qu, Rikui Huang, Zhixin Chen and Xingyu Yan	14035
<i>Innovative Image Fraud Detection with Cross-Sample Anomaly Analysis: The Power of LLMs</i>	
QiWen Wang, Junqi Yang, Zhenghao Lin, Zhenzhe Ying, Weiqiang Wang and Chen Lin ..	14058
<i>Cooperative or Competitive? Understanding the Interaction between Attention Heads From A Game Theory Perspective</i>	
Xiaoye Qu, Zengqi Yu, Dongrui Liu, Wei Wei, Daizong Liu, Jianfeng Dong and Yu Cheng	14079
<i>MM-Verify: Enhancing Multimodal Reasoning with Chain-of-Thought Verification</i>	
Lin Zhuang Sun, Hao Liang, Jingxuan Wei, Bihui Yu, Tianpeng Li, Fan Yang, Zenan Zhou and Wentao Zhang	14100
<i>Graph-Structured Trajectory Extraction from Travelogues</i>	
Aitaro Yamamoto, Hiroyuki Otomo, Hiroki Ouchi, Shohei Higashiyama, Hiroki Teranishi, Hiroyuki Shindo and Taro Watanabe	14116
<i>Learning First-Order Logic Rules for Argumentation Mining</i>	
Yang Sun, Guanrong Chen, Hamid Alinejad-Rokny, Jianzhu Bao, Yuqi Huang, Bin Liang, Kam-Fai Wong, Min Yang and Ruifeng Xu	14133
<i>Investigating and Enhancing the Robustness of Large Multimodal Models Against Temporal Inconsistency</i>	
Jiafeng Liang, Shixin Jiang, Xuan Dong, Ning Wang, Zheng Chu, Hui Su, Jinlan Fu, Ming Liu, See-Kiong Ng and Bing Qin	14149
<i>UniRAG: Unified Query Understanding Method for Retrieval Augmented Generation</i>	
Rui Li, Liyang He, Qi Liu, Zheng Zhang, Heng Yu, Yuyang Ye, Linbo Zhu and Yu Su	14163
<i>Contextual Experience Replay for Self-Improvement of Language Agents</i>	
Yitao Liu, Chenglei Si, Karthik R Narasimhan and Shunyu Yao	14179
<i>Emma-X: An Embodied Multimodal Action Model with Grounded Chain of Thought and Look-ahead Spatial Reasoning</i>	
Qi Sun, Pengfei Hong, Tej Deep Pala, Vernon Toh, U-Xuan Tan, Deepanway Ghosal and Soujanya Poria	14199

<i>Towards Comprehensive Argument Analysis in Education: Dataset, Tasks, and Method</i>	
Yupei Ren, Xinyi Zhou, Ning Zhang, Shangqing Zhao, Man Lan and Xiaopeng Bai	14215
<i>Browsing Like Human: A Multimodal Web Agent with Experiential Fast-and-Slow Thinking</i>	
Haohao Luo, Jiayi Kuang, Wei Liu, Ying Shen, Jian Luan and Yang Deng	14232
<i>MaXIFE: Multilingual and Cross-lingual Instruction Following Evaluation</i>	
Yile Liu, Ziwei Ma, Xiu Jiang, Jinglu HU, ChangJing ChangJing and Liang Li	14252
<i>Linguistic Generalizability of Test-Time Scaling in Mathematical Reasoning</i>	
Guijin Son, Jiwoo Hong, Hyunwoo Ko and James Thorne	14333
<i>Can MLLMs Understand the Deep Implication Behind Chinese Images?</i>	
Chenhao Zhang, Xi Feng, Yuelin Bai, Xeron Du, Jinchang Hou, Kaixin Deng, Guangzeng Han, Qinrui Li, Bingli Wang, Jiaheng Liu, Xingwei Qu, Yifei Zhang, Qixuan Zhao, Yiming Liang, Ziqiang Liu, Feiteng Fang, Min Yang, Wenhao Huang, Chenghua Lin, Ge Zhang and Shiwen Ni	14369
<i>KazMMLU: Evaluating Language Models on Kazakh, Russian, and Regional Knowledge of Kazakhstan</i>	
Mukhammed Togmanov, Nurdaulet Mukhituly, Diana Turmakhan, Jonibek Mansurov, Maiya Goloburda, Akhmed Sakip, Zhuohan Xie, Yuxia Wang, Bekassyl Syzdykov, Nurkhan Laiyk, Alham Fikri Aji, Ekaterina Kochmar, Preslav Nakov and Fajri Koto	14403
<i>Towards Multi-dimensional Evaluation of LLM Summarization across Domains and Languages</i>	
Hyangsuk Min, Yuho Lee, Minjeong Ban, Jiaqi Deng, Nicole Hee-Yeon Kim, Taewon Yun, Hang Su, Jason Cai and Hwanjun Song	14417
<i>ClusterAttn: KV Cache Compression under Intrinsic Attention Clustering</i>	
Minwei Zhang, Haifeng Sun, Jingyu Wang, Shaolong Li, Wanyi Ning, Qi Qi, Zirui Zhuang and Jianxin Liao	14451
<i>SHARE: Shared Memory-Aware Open-Domain Long-Term Dialogue Dataset Constructed from Movie Script</i>	
Eunwon Kim, Chanho Park and Buru Chang	14474
<i>Incongruity-aware Tension Field Network for Multi-modal Sarcasm Detection</i>	
Jiecheng Zhang, C.L.Philip Chen, Shuzhen Li and Tong Zhang	14499
<i>Instruction Tuning on Public Government and Cultural Data for Low-Resource Language: a Case Study in Kazakh</i>	
Nurkhan Laiyk, Daniil Orel, Rituraj Joshi, Maiya Goloburda, Yuxia Wang, Preslav Nakov and Fajri Koto	14509
<i>Stealing Training Data from Large Language Models in Decentralized Training through Activation Inversion Attack</i>	
Chenxi Dai, Lin Lu and Pan Zhou	14539
<i>From Selection to Generation: A Survey of LLM-based Active Learning</i>	
Yu Xia, Subhojyoti Mukherjee, Zhouhang Xie, Junda Wu, Xintong Li, Ryan Aponte, Hanjia Lyu, Joe Barrow, Hongjie Chen, Franck Dernoncourt, Branislav Kveton, Tong Yu, Ruiyi Zhang, Jiuxiang Gu, Nesreen K. Ahmed, Yu Wang, Xiang Chen, Hanieh Deilamsalehy, Sungchul Kim, Zhengmian Hu, Yue Zhao, Nedim Lipka, Seunghyun Yoon, Ting-Hao Kenneth Huang, Zichao Wang, Puneet Mathur, Soumyabrata Pal, Koyel Mukherjee, Zehao Zhang, Namyong Park, Thien Huu Nguyen, Jiebo Luo, Ryan A. Rossi and Julian McAuley	14552
<i>OmniFlatten: An End-to-end GPT Model for Seamless Voice Conversation</i>	
Qinglin Zhang, Luyao Cheng, Chong Deng, Qian Chen, Wen Wang, Siqi Zheng, Jiaqing Liu, Hai Yu, Chao-Hong Tan, Zhihao Du and ShiLiang Zhang	14570

<i>DoMIX: An Efficient Framework for Exploiting Domain Knowledge in Fine-Tuning</i>	
Dohoon Kim, Donghun Kang and Taesup Moon	14581
<i>EAGLE: Expert-Guided Self-Enhancement for Preference Alignment in Pathology Large Vision-Language Model</i>	
Meidan Ding, Jipeng Zhang, Wenxuan Wang, Haiqin Zhong, Xiaoqin Wang, Xinheng Lyu, Wenting Chen and Linlin Shen	14603
<i>CoT-ICL Lab: A Synthetic Framework for Studying Chain-of-Thought Learning from In-Context Demonstrations</i>	
Vignesh Kothapalli, Hamed Firooz and Maziar Sanjabi	14620
<i>Flexora: Flexible Low-Rank Adaptation for Large Language Models</i>	
Chenxing Wei, Yao Shu, Ying Tiffany He and Fei Yu	14643
<i>QDTSynth: Quality-Driven Formal Theorem Synthesis for Enhancing Proving Performance of LLMs</i>	
Lei Wang, Ruobing Zuo, Gaolei He, Jianlin Wang and Zhengfeng Yang	14683
<i>RSVP: Reasoning Segmentation via Visual Prompting and Multi-modal Chain-of-Thought</i>	
Yi Lu, Jiawang Cao, Yongliang Wu, Bozheng Li, Licheng Tang, Yangguang Ji, Chong Wu, Jay Wu and Wenbo Zhu	14699
<i>QAEval: Mixture of Evaluators for Question-Answering Task Evaluation</i>	
Tan Yue, Rui Mao, Xuzhao Shi, Shuo Zhan, Zuhao Yang and Dongyan Zhao	14717
<i>Debiasing the Fine-Grained Classification Task in LLMs with Bias-Aware PEFT</i>	
Daiying Zhao, Xinyu Yang and Hang Chen	14731
<i>Demystifying Small Language Models for Edge Deployment</i>	
Zhenyan Lu, Xiang Li, Dongqi Cai, Rongjie Yi, Fangming Liu, Wei Liu, Jian Luan, Xiwen Zhang, Nicholas D. Lane and Mengwei Xu	14747
<i>Adapt Once, Thrive with Updates: Transferable Parameter-Efficient Fine-Tuning on Evolving Base Models</i>	
Naibin Gu, Peng Fu, Xiyu Liu, Ke Ma, Zheng Lin and Weiping Wang	14765
<i>Can Vision-Language Models Evaluate Handwritten Math?</i>	
Oikantik Nath, Hanani Bathina, Mohammed Safi Ur Rahman Khan and Mitesh M Khapra	14784
<i>Continual Gradient Low-Rank Projection Fine-Tuning for LLMs</i>	
Chenxu Wang, Yilin Lyu, Zicheng Sun and Liping Jing	14815
<i>Towards Objective Fine-tuning: How LLMs' Prior Knowledge Causes Potential Poor Calibration?</i>	
Ziming Wang, Zeyu Shi, Haoyi Zhou, Shiqi Gao, Qingyun Sun and Jianxin Li	14830
<i>Towards Robust ESG Analysis Against Greenwashing Risks: Aspect-Action Analysis with Cross-Category Generalization</i>	
Keane Ong, Rui Mao, Deeksha Varshney, Erik Cambria and Gianmarco Mengaldo	14854
<i>HiddenDetect: Detecting Jailbreak Attacks against Multimodal Large Language Models via Monitoring Hidden States</i>	
Yilei Jiang, Xinyan Gao, Tianshuo Peng, Yingshui Tan, Xiaoyong Zhu, Bo Zheng and Xiangyu Yue	14880
<i>SwiLTr-Bench: The Swiss Legal Translation Benchmark</i>	
Joel Niklaus, Jakob Merane, Luka Nenadic, Sina Ahmadi, Yingqiang Gao, Cyril A. H. Chevalley, Claude Humbel, Christophe Gösken, Lorenzo Tanzi, Thomas Lüthi, Stefan Palombo, Spencer Poff,	

Boling Yang, Nan Wu, Matthew Guilloid, Robin Mamié, Daniel Brunner, Julio Pereyra and Niko Grupen	14894
<i>Two Intermediate Translations Are Better Than One: Fine-tuning LLMs for Document-level Translation Refinement</i>	
Yichen Dong, Xinglin Lyu, Junhui Li, Daimeng Wei, Min Zhang, Shimin Tao and Hao Yang	14917
<i>Circuit Compositions: Exploring Modular Structures in Transformer-Based Language Models</i>	
Philipp Mondorf, Sondre Wold and Barbara Plank	14934
<i>Can LLMs Ground when they (Don't) Know: A Study on Direct and Loaded Political Questions</i>	
Clara Lachenmaier, Judith Sieker and Sina Zarrieß	14956
<i>GraphCheck: Breaking Long-Term Text Barriers with Extracted Knowledge Graph-Powered Fact-Checking</i>	
Yingjian Chen, Haoran Liu, Yinhong Liu, Jinxiang Xie, Rui Yang, Han Yuan, Yanran Fu, Peng Yuan Zhou, Qingyu Chen, James Caverlee and Irene Li	14976
<i>SCULPT: Systematic Tuning of Long Prompts</i>	
Shanu Kumar, Akhila Yesantarao Venkata, Shubhanshu Khandelwal, Bishal Santra, Parag Agrawal and Manish Gupta	14996
<i>Crab: A Novel Configurable Role-Playing LLM with Assessing Benchmark</i>	
Kai He, Yucheng Huang, Wenqing Wang, Delong Ran, Dongming Sheng, Junxuan Huang, Qika Lin, Jiaying Xu, Wenqiang Liu and Mengling Feng	15030
<i>Chinese SafetyQA: A Safety Short-form Factuality Benchmark for Large Language Models</i>	
Yingshui Tan, Boren Zheng, Baihui Zheng, Kerui Cao, Huiyun Jing, Jincheng Wei, Jiaheng Liu, Yancheng He, Wenbo Su, Xiaoyong Zhu, Bo Zheng and Kaifu Zhang	15053
<i>TRIDENT: Enhancing Large Language Model Safety with Tri-Dimensional Diversified Red-Teaming Data Synthesis</i>	
Xiaorui Wu, Xiaofeng Mao, Fei Li, Xin Zhang, Xuanhong Li, Chong Teng, Donghong Ji and Zhuang Li	15077
<i>Cross-Lingual Optimization for Language Transfer in Large Language Models</i>	
Jungseob Lee, Seongtae Hong, Hyeonseok Moon and Heuseok Lim	15100
<i>CART: A Generative Cross-Modal Retrieval Framework With Coarse-To-Fine Semantic Modeling</i>	
Minghui Fang, Shengpeng Ji, Jialong Zuo, Hai Huang, Yan Xia, Jieming Zhu, Xize Cheng, Xiaoda Yang, Wenrui Liu, Gang Wang, Zhenhua Dong and Zhou Zhao	15120
<i>MMMU-Pro: A More Robust Multi-discipline Multimodal Understanding Benchmark</i>	
Xiang Yue, Tianyu Zheng, Yuansheng Ni, Yubo Wang, Kai Zhang, Shengbang Tong, Yuxuan Sun, Botao Yu, Ge Zhang, Huan Sun, Yu Su, Wenhua Chen and Graham Neubig	15134
<i>Cheems: A Practical Guidance for Building and Evaluating Chinese Reward Models from Scratch</i>	
Xueru Wen, Jie Lou, Zichao Li, Yaojie Lu, Xingyu Xingyu, Yuqiu Ji, Guohai Xu, Hongyu Lin, Ben He, Xianpei Han, Le Sun and Debing Zhang	15187
<i>Why Safeguarded Ships Run Aground? Aligned Large Language Models' Safety Mechanisms Tend to Be Anchored in The Template Region</i>	
Chak Tou Leong, Qingyu Yin, Jian Wang and Wenjie Li	15212

<i>LLaVA Steering: Visual Instruction Tuning with 500x Fewer Parameters through Modality Linear Representation-Steering</i>	
Jinhe Bi, Yujun Wang, Haokun Chen, Xun Xiao, Artur Hecker, Volker Tresp and Yunpu Ma	15230
<i>Efficient Long Context Language Model Retrieval with Compression</i>	
Minju Seo, Jinheon Baek, Seongyun Lee and Sung Ju Hwang	15251
<i>Ontology-Guided Reverse Thinking Makes Large Language Models Stronger on Knowledge Graph Question Answering</i>	
Runxuan Liu, Luobei Luobei, Jiaqi Li, Baoxin Wang, Ming Liu, Dayong Wu, Shijin Wang and Bing Qin	15269
<i>Towards Omni-RAG: Comprehensive Retrieval-Augmented Generation for Large Language Models in Medical Applications</i>	
Zhe Chen, Yusheng Liao, Shuyang Jiang, Pingjie Wang, YiQiu Guo, Yanfeng Wang and Yu Wang	15285
<i>Predicting Turn-Taking and Backchannel in Human-Machine Conversations Using Linguistic, Acoustic, and Visual Signals</i>	
Yuxin Lin, Yinglin Zheng, Ming Zeng and Wangzheng Shi	15310
<i>A New Formulation of Zipf’s Meaning-Frequency Law through Contextual Diversity</i>	
Ryo Nagata and Kumiko Tanaka-Ishii	15323
<i>The Mirage of Model Editing: Revisiting Evaluation in the Wild</i>	
Wanli Yang, Fei Sun, Jiajun Tan, Xinyu Ma, Qi Cao, Dawei Yin, Huawei Shen and Xueqi Cheng	15336
<i>LAQuer: Localized Attribution Queries in Content-grounded Generation</i>	
Eran Hirsch, Aviv Slobodkin, David Wan, Elias Stengel-Eskin, Mohit Bansal and Ido Dagan	15355
<i>EPO: Explicit Policy Optimization for Strategic Reasoning in LLMs via Reinforcement Learning</i>	
Xiaoqian Liu, Ke Wang, Yongbin Li, Yuchuan Wu, Wentao Ma, Aobo Kong, Fei Huang, Jianbin Jiao and Junge Zhang	15371
<i>DCG-SQL: Enhancing In-Context Learning for Text-to-SQL with Deep Contextual Schema Link Graph</i>	
Jihyung Lee, Jin-Seop Lee, Jaehoon Lee, YunSeok Choi and Jee-Hyong Lee	15397
<i>PreP-OCR: A Complete Pipeline for Document Image Restoration and Enhanced OCR Accuracy</i>	
Shuhao Guan, Moule Lin, Cheng Xu, Xinyi Liu, Jinman Zhao, Jiexin Fan, Qi Xu and Derek Greene	15413
<i>Digest the Knowledge: Large Language Models empowered Message Passing for Knowledge Graph Question Answering</i>	
Junhong Wan, Tao Yu, Kunyu Jiang, Yao Fu, Weihao Jiang and Jiang Zhu	15426
<i>RecLM: Recommendation Instruction Tuning</i>	
Yangqin Jiang, Yuhao Yang, Lianghao Xia, Da Luo, Kangyi Lin and Chao Huang	15443
<i>DS²-ABSA: Dual-Stream Data Synthesis with Label Refinement for Few-Shot Aspect-Based Sentiment Analysis</i>	
Hongling Xu, Yice Zhang, Qianlong Wang and Ruifeng Xu	15460
<i>MISP-Meeting: A Real-World Dataset with Multimodal Cues for Long-form Meeting Transcription and Summarization</i>	
HangChen HangChen, Chao-Han Huck Yang, Jia-Chen Gu, Sabato Marco Siniscalchi and Jun Du	15479

<i>Learning Together to Perform Better: Teaching Small-Scale LLMs to Collaborate via Preferential Rationale Tuning</i>	
Sohan Patnaik, Milan Aggarwal, Sumit Bhatia and Balaji Krishnamurthy	15493
<i>MolRAG: Unlocking the Power of Large Language Models for Molecular Property Prediction</i>	
Ziting Xian, Jiawei Gu, Lingbo Li and Shangsong Liang	15513
<i>SkillAggregation: Reference-free LLM-Dependent Aggregation</i>	
Guangzhi Sun, Anmol Kagrecha, Potsawee Manakul, Phil Woodland and Mark Gales	15532
<i>MasRouter: Learning to Route LLMs for Multi-Agent Systems</i>	
Yanwei Yue, Guibin Zhang, Boyang Liu, Guancheng Wan, Kun Wang, Dawei Cheng and Yiyang Qi	15549
<i>Beyond Single Labels: Improving Conversational Recommendation through LLM-Powered Data Augmentation</i>	
Haozhe Xu, Xiaohua Wang, Changze Lv and Xiaoqing Zheng	15573
<i>Beyond One-Size-Fits-All: Tailored Benchmarks for Efficient Evaluation</i>	
Peiwen Yuan, Yueqi Zhang, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Jiayi Shi, Chuyi Tan, Boyuan Pan, Yao Hu and Kan Li	15591
<i>iQUEST: An Iterative Question-Guided Framework for Knowledge Base Question Answering</i>	
Shuai Wang and Yinan Yu	15616
<i>IRT-Router: Effective and Interpretable Multi-LLM Routing via Item Response Theory</i>	
Wei Song, Zhenya Huang, Cheng Cheng, Weibo Gao, Bihan Xu, GuanHao Zhao, Fei Wang and Runze Wu	15629
<i>MLAS-LoRA: Language-Aware Parameters Detection and LoRA-Based Knowledge Transfer for Multilingual Machine Translation</i>	
Tianyu Dong, Bo Li, Jinsong Liu, Shaolin Zhu and Deyi Xiong	15645
<i>M2RC-EVAL: Massively Multilingual Repository-level Code Completion Evaluation</i>	
Jiaheng Liu, Ken Deng, Congnan Liu, Jian Yang, Shukai Liu, He Zhu, Peng Zhao, Linzheng Chai, Yanan Wu, JinKe JinKe, Ge Zhang, Zekun Moore Wang, Guoan Zhang, Yingshui Tan, Bangyu Xiang, Zhaoxiang Zhang, Wenbo Su and Bo Zheng	15661
<i>Evaluating Design Decisions for Dual Encoder-based Entity Disambiguation</i>	
Susanna Rücker and Alan Akbik	15685
<i>How to Compare Things Properly? A Study of Argument Relevance in Comparative Question Answering</i>	
Irina Nikishina, Saba Anwar, Nikolay Dolgov, Maria Manina, Daria Ignatenko, Artem Shelmanov and Chris Biemann	15702
<i>FinanceReasoning: Benchmarking Financial Numerical Reasoning More Credible, Comprehensive and Challenging</i>	
Zichen Tang, Haihong E, Ziyang Ma, Haoyang He, Jiacheng Liu, Zhongjun Yang, Zihua Rong, Rongjin Li, Kun Ji, Qing Huang, Xinyang Hu, Yang Liu and Qianhe Zheng	15721
<i>Controllable Style Arithmetic with Language Models</i>	
Weiqi Wang, Wengang Zhou, Zongmeng Zhang, Jie Zhao and Houqiang Li	15750
<i>Masks Can be Learned as an Alternative to Experts</i>	
Peiyu Liu, Tianwen Wei, Bo Zhu, Xin Zhao and Shuicheng Yan	15800

<i>Program Synthesis Benchmark for Visual Programming in XLogoOnline Environment</i> Chao Wen, Jacqueline Staub and Adish Singla	15812
<i>Removal of Hallucination on Hallucination: Debate-Augmented RAG</i> Wentao Hu, Wengyu Zhang, Yiyang Jiang, Chen Jason Zhang, Xiaoyong Wei and Li Qing	15839
<i>CodeDPO: Aligning Code Models with Self Generated and Verified Source Code</i> Kechi Zhang, Ge Li, Yihong Dong, Jingjing Xu, Jun Zhang, Jing Su, Yongfei Liu and Zhi Jin	15854
<i>ProxAnn: Use-Oriented Evaluations of Topic Models and Document Clustering</i> Alexander Miserlis Hoyle, Lorena Calvo-Bartolomé, Jordan Lee Boyd-Graber and Philip Resnik	15872
<i>BOOKWORLD: From Novels to Interactive Agent Societies for Story Creation</i> Yiting Ran, Xintao Wang, Tian Qiu, Jiaqing Liang, Yanghua Xiao and Deqing Yang	15898
<i>Quantifying Lexical Semantic Shift via Unbalanced Optimal Transport</i> Ryo Kishino, Hiroaki Yamagiwa, Ryo Nagata, Sho Yokoi and Hidetoshi Shimodaira	15913
<i>Agentic Reward Modeling: Integrating Human Preferences with Verifiable Correctness Signals for Reliable Reward Systems</i> Hao Peng, Yunjia Qi, Xiaozhi Wang, Zijun Yao, Bin Xu, Lei Hou and Juanzi Li	15934
<i>Adaptive and Robust Translation from Natural Language to Multi-model Query Languages</i> Gengyuan Shi, Chaokun Wang, Liu Yabin and Jiawei Ren	15950
<i>SAKE: Steering Activations for Knowledge Editing</i> Marco Scialanga, Thibault Laugel, Vincent Grari and Marcin Detyniecki	15966
<i>Middle-Layer Representation Alignment for Cross-Lingual Transfer in Fine-Tuned LLMs</i> Danni Liu and Jan Niehues	15979
<i>Can External Validation Tools Improve Annotation Quality for LLM-as-a-Judge?</i> Arduin Findeis, Floris Weers, Guoli Yin, Ke Ye, Ruoming Pang and Tom Gunter	15997
<i>One for All: Update Parameterized Knowledge Across Multiple Models with Once Edit</i> Weitao Ma, Xiyuan Du, Xiaocheng Feng, Lei Huang, Yichong Huang, Huiyi Zhang, Xiaoliang Yang, Baohang Li, Xiachong Feng, Ting Liu and Bing Qin	16021
<i>VLMInferSlow: Evaluating the Efficiency Robustness of Large Vision-Language Models as a Service</i> Xiasi Wang, Tianliang Yao, Simin Chen, Runqi Wang, Lei YE, Kuofeng Gao, Yi Huang and Yuan Yao	16035
<i>The Alternative Annotator Test for LLM-as-a-Judge: How to Statistically Justify Replacing Human Annotators with LLMs</i> Nitay Calderon, Roi Reichart and Rotem Dror	16051
<i>CrisisTS: Coupling Social Media Textual Data and Meteorological Time Series for Urgency Classification</i> Romain Meunier, Farah Benamara, Véronique Moriceau, Zhongzheng Qiao and Savitha Ramasamy	16082
<i>How to Mitigate Overfitting in Weak-to-strong Generalization?</i> Junhao Shi, Qinyuan Cheng, Zhaoye Fei, Yining Zheng, Qipeng Guo and Xipeng Qiu	16100

<i>Com² : A Causal-Guided Benchmark for Exploring Complex Commonsense Reasoning in Large Language Models</i>	
Kai Xiong, Xiao Ding, Yixin Cao, Yuxiong Yan, Li Du, Yufei Zhang, Jinglong Gao, Jiaqian Liu, Bing Qin and Ting Liu	16119
<i>Dynamic Head Selection for Neural Lexicalized Constituency Parsing</i>	
Yang Hou and Zhenghua Li	16141
<i>My Words Imply Your Opinion: Reader Agent-Based Propagation Enhancement for Personalized Implicit Emotion Analysis</i>	
Jian Liao, Yu Feng, Yujin Zheng, Jun Zhao, Suge Wang and JianXing Zheng	16156
<i>EvolveBench: A Comprehensive Benchmark for Assessing Temporal Awareness in LLMs on Evolving Knowledge</i>	
Zhiyuan Zhu, Yusheng Liao, Zhe Chen, Yuhao Wang, Yunfeng Guan, Yanfeng Wang and Yu Wang	16173
<i>Enabling LLM Knowledge Analysis via Extensive Materialization</i>	
Yujia Hu, Tuan-Phong Nguyen, Shrestha Ghosh and Simon Razniewski	16189
<i>Rhythm Controllable and Efficient Zero-Shot Voice Conversion via Shortcut Flow Matching</i>	
Jialong Zuo, Shengpeng Ji, Minghui Fang, Mingze Li, Ziyue Jiang, Xize Cheng, Xiaoda Yang, Chen Feiyang, Xinyu Duan and Zhou Zhao	16203
<i>Llama See, Llama Do: A Mechanistic Perspective on Contextual Entrainment and Distraction in LLMs</i>	
Jingcheng Niu, Xingdi Yuan, Tong Wang, Hamidreza Saghir and Amir H. Abdi	16218
<i>CritiQ: Mining Data Quality Criteria from Human Preferences</i>	
Honglin Guo, Kai Lv, Qipeng Guo, Tianyi Liang, Zhiheng Xi, Demin Song, Qiuyinzhe Zhang, Yu Sun, Kai Chen, Xipeng Qiu and Tao Gui	16240
<i>Theoretical Guarantees for Minimum Bayes Risk Decoding</i>	
Yuki Ichihara, Yuu Jinnai, Kaito Ariu, Tetsuro Morimura and Eiji Uchibe	16262
<i>Mutual-Taught for Co-adapting Policy and Reward Models</i>	
Tianyuan Shi, Canbin Huang, Fanqi Wan, Longguang Zhong, Ziyi Yang, Weizhou Shen, Xiaojun Quan and Ming Yan	16285
<i>Enhancing Cross-Lingual Transfer through Reversible Transliteration: A Huffman-Based Approach for Low-Resource Languages</i>	
Wenhao Zhuang, Yuan Sun and Xiaobing Zhao	16299
<i>Unmasking Style Sensitivity: A Causal Analysis of Bias Evaluation Instability in Large Language Models</i>	
Jiaxu Zhao, Meng Fang, Kun Zhang and Mykola Pechenizkiy	16314
<i>MockConf: A Student Interpretation Dataset: Analysis, Word- and Span-level Alignment and Baselines</i>	
Dávid Javorský, Ondřej Bojar and François Yvon	16339
<i>BMIKE-53: Investigating Cross-Lingual Knowledge Editing with In-Context Learning</i>	
Ercong Nie, Bo Shao, Mingyang Wang, Zifeng Ding, Helmut Schmid and Hinrich Schuetze	16357
<i>What Matters in Evaluating Book-Length Stories? A Systematic Study of Long Story Evaluation</i>	
Dingyi Yang and Qin Jin	16375
<i>PROPER: A Progressive Learning Framework for Personalized Large Language Models with Group-Level Adaptation</i>	
Linhai Zhang, Jialong Wu, Deyu Zhou and Yulan He	16399

<i>Enhancing Event-centric News Cluster Summarization via Data Sharpening and Localization Insights</i> Longyin Zhang, Bowei Zou and AiTi Aw	16412
<i>MMBoundary: Advancing MLLM Knowledge Boundary Awareness through Reasoning Step Confidence Calibration</i> Zhitao He, Sandeep Polisetty, Zhiyuan Fan, Yuchen Huang, Shujin Wu and Yi R. Fung ...	16427
<i>LIFBench: Evaluating the Instruction Following Performance and Stability of Large Language Models in Long-Context Scenarios</i> Xiaodong Wu, Minhao Wang, Yichen Liu, Xiaoming Shi, He Yan, Lu Xiangju, Junmin Zhu and Wei Zhang	16445
<i>Aligning Large Language Models to Follow Instructions and Hallucinate Less via Effective Data Filtering</i> Shuzheng Si, Haozhe Zhao, Gang Chen, Cheng Gao, Yuzhuo Bai, Zhitong Wang, Kaikai An, Kangyang Luo, Chen Qian, Fanchao Qi, Baobao Chang and Maosong Sun	16469
<i>One-Shot is Enough: Consolidating Multi-Turn Attacks into Efficient Single-Turn Prompts for LLMs</i> Junwoo Ha, Hyunjun Kim, Sangyoon Yu, Haon Park, Ashkan Yousefpour, Yuna Park and Suhyun Kim	16489
<i>RAEmoLLM: Retrieval Augmented LLMs for Cross-Domain Misinformation Detection Using In-Context Learning Based on Emotional Information</i> Zhiwei Liu, Kailai Yang, Qianqian Xie, Christine de Kock, Sophia Ananiadou and Eduard Hovy	16508
<i>Task-Specific Information Decomposition for End-to-End Dense Video Captioning</i> Zhiyue Liu, Xinru Zhang and Jinyuan Liu	16524
<i>CalibraEval: Calibrating Prediction Distribution to Mitigate Selection Bias in LLMs-as-Judges</i> Haitao Li, Junjie Chen, Qingyao Ai, Zhumin Chu, Yujia Zhou, Qian Dong and Yiqun Liu .	16537
<i>Explaining Matters: Leveraging Definitions and Semantic Expansion for Sexism Detection</i> Sahrish Khan, Arshad Jhumka and Gabriele Pergola	16553
<i>Private Memorization Editing: Turning Memorization into a Defense to Strengthen Data Privacy in Large Language Models</i> Elena Sofia Ruzzetti, Giancarlo A. Xompero, Davide Venditti and Fabio Massimo Zanzotto	16572
<i>PhysReason: A Comprehensive Benchmark towards Physics-Based Reasoning</i> Xinyu Zhang, Yuxuan Dong, Yanrui Wu, Jiaying Huang, Chengyou Jia, Basura Fernando, Mike Zheng Shou, Lingling Zhang and Jun Liu	16593
<i>Does Time Have Its Place? Temporal Heads: Where Language Models Recall Time-specific Information</i> Yein Park, Chanwoong Yoon, Jungwoo Park, Minbyul Jeong and Jaewoo Kang	16616
<i>Velocitune: A Velocity-based Dynamic Domain Reweighting Method for Continual Pre-training</i> Zheheng Luo, Xin Zhang, Xiao Liu, Haoling Li, Yeyun Gong, Qi Chen and Peng Cheng ..	16644
<i>Sheep's Skin, Wolf's Deeds: Are LLMs Ready for Metaphorical Implicit Hate Speech?</i> Jingjie Zeng, Liang Yang, Zekun Wang, Yuanyuan Sun and Hongfei Lin	16657
<i>Neuron-Level Sequential Editing for Large Language Models</i> Houcheng Jiang, Junfeng Fang, Tianyu Zhang, Baolong Bi, An Zhang, Ruipeng Wang, Tao Liang and Xiang Wang	16678

<i>Automatic Expert Discovery in LLM Upcycling via Sparse Interpolated Mixture-of-Experts</i>	Shengzhuang Chen, Ying Wei and Jonathan Richard Schwarz	16703
<i>SimulS2S-LLM: Unlocking Simultaneous Inference of Speech LLMs for Speech-to-Speech Translation</i>	Keqi Deng, Wenxi Chen, Xie Chen and Phil Woodland	16718
<i>VoxEval: Benchmarking the Knowledge Understanding Capabilities of End-to-End Spoken Language Models</i>	Wenqian Cui, Xiaoqi Jiao, Ziqiao Meng and Irwin King	16735
<i>RetroLLM: Empowering Large Language Models to Retrieve Fine-grained Evidence within Generation</i>	Xiaoxi Li, Jiajie Jin, Yujia Zhou, Yongkang Wu, Zhonghua Li, YE QI and Zhicheng Dou .	16754
<i>The Role of Deductive and Inductive Reasoning in Large Language Models</i>	Chengkun Cai, Xu Zhao, Haoliang Liu, Zhongyu Jiang, Tianfang Zhang, Zongkai Wu, Jenq-Neng Hwang and Lei Li	16780
<i>Disentangling the Roles of Representation and Selection in Data Pruning</i>	Yupei Du, Yingjin Song, Hugh Mee Wong, Daniil Ignatev, Albert Gatt and Dong Nguyen .	16791
<i>FRACTAL: Fine-Grained Scoring from Aggregate Text Labels</i>	Yukti Makhija, Priyanka Agrawal, Rishi Saket and Aravindan Raghuv eer	16810
<i>ACT: Knowledgeable Agents to Design and Perform Complex Tasks</i>	Makoto Nakatsuji, Shuhei Tateishi, Yasuhiro Fujiwara, Ayaka Matsumoto, Narichika Nomoto and Yoshihide Sato	16831
<i>Logical forms complement probability in understanding language model (and human) performance</i>	Yixuan Wang and Freda Shi	16862
<i>Length Controlled Generation for Black-box LLMs</i>	Yuxuan Gu, Wenjie Wang, Xiaocheng Feng, Weihong Zhong, Kun Zhu, Lei Huang, Ting Liu, Bing Qin and Tat-Seng Chua	16878
<i>Improving Contextual Faithfulness of Large Language Models via Retrieval Heads-Induced Optimization</i>	Lei Huang, Xiaocheng Feng, Weitao Ma, Yuchun Fan, Xiachong Feng, Yangfan Ye, Weihong Zhong, Yuxuan Gu, Baoxin Wang, Dayong Wu, Guoping Hu and Bing Qin	16896
<i>Global Eye: Breaking the Fixed Thinking Pattern during the Instruction Expansion Process</i>	Wenxuan lu, Wei Liu, Jian Luan, Bin Wang, Songhao Jiang and Tianning Zang	16914
<i>On Synthesizing Data for Context Attribution in Question Answering</i>	Gorjan Radevski, Kiril Gashteovski, Shahbaz Syed, Christopher Malon, Sebastien Nicolas, Chia-Chien Hung, Timo Sztyler, Verena Heußer, Wiem Ben Rim, Masafumi Enomoto, Kunihiro Takeoka, Masafumi Oyamada, Goran Glavaš and Carolin Lawrence	16929
<i>TST: A Schema-Based Top-Down and Dynamic-Aware Agent of Text-to-Table Tasks</i>	Peiwen Jiang, Haitong Jiang, Ruhui Ma, Yvonne Jie Chen and Jinhua Cheng	16951
<i>EventRAG: Enhancing LLM Generation with Event Knowledge Graphs</i>	Zairun Yang, Yilin Wang, Zhengyan Shi, Yuan Yao, Lei Liang, Keyan Ding, Emine Yilmaz, Huajun Chen and Qiang Zhang	16967
<i>Analyzing the Rapid Generalization of SFT via the Perspective of Attention Head Activation Patterns</i>	Yang Zhao, Li Du, Xiao Ding, Kai Xiong, Ting Liu and Bing Qin	16980

<i>Can't See the Forest for the Trees: Benchmarking Multimodal Safety Awareness for Multimodal LLMs</i> Wenxuan Wang, Xiaoyuan Liu, Kuiyi Gao, Jen-tse Huang, Youliang Yuan, Pinjia He, Shuai Wang and Zhaopeng Tu	16993
<i>Mis-prompt: Benchmarking Large Language Models for Proactive Error Handling</i> Jiayi Zeng, Yizhe Feng, Mengliang He, Wenhui Lei, Wei Zhang, Zeming Liu, Xiaoming Shi and Aimin Zhou	17007
<i>TripCraft: A Benchmark for Spatio-Temporally Fine Grained Travel Planning</i> Soumyabrata Chaudhuri, Pranav Purkar, Ritwik Raghav, Shubhojit Mallick, Manish Gupta, Abhik Jana and Shreya Ghosh	17035
<i>DualGuard: A Parameter Space Transformation Approach for Bidirectional Defense in Split-Based LLM Fine-Tuning</i> Zihan Liu, Yizhen Wang, Rui Wang and Sai Wu	17065
<i>Movie101v2: Improved Movie Narration Benchmark</i> Zihao Yue, Yepeng Zhang, Ziheng Wang and Qin Jin	17081
<i>Can LLMs Evaluate Complex Attribution in QA? Automatic Benchmarking using Knowledge Graphs</i> Nan Hu, Jiaoyan Chen, Yike Wu, Guilin Qi, Hongru Wang, Sheng Bi, Yongrui Chen, Tongtong Wu and Jeff Z. Pan	17096
<i>Value Portrait: Assessing Language Models' Values through Psychometrically and Ecologically Valid Items</i> Jongwook Han, Dongmin Choi, Woojung Song, Eun-Ju Lee and Yohan Jo	17119
<i>FEA-Bench: A Benchmark for Evaluating Repository-Level Code Generation for Feature Implementa- tion</i> Wei Li, Xin Zhang, Zhongxin Guo, Shaoguang Mao, Wen Luo, Guangyue Peng, Yangyu Huang, Houfeng Wang and Scarlett Li	17160
<i>Do not Abstain! Identify and Solve the Uncertainty</i> Jingyu Liu, JingquanPeng JingquanPeng, Xiaopeng Wu, Xubin Li, Tiezheng Ge, Bo Zheng and Yong Liu	17177
<i>Decoding by Contrasting Knowledge: Enhancing Large Language Model Confidence on Edited Facts</i> Baolong Bi, Shenghua Liu, Lingrui Mei, Yiwei Wang, Junfeng Fang, Pengliang Ji and Xueqi Cheng	17198
<i>ImpliHateVid: A Benchmark Dataset and Two-stage Contrastive Learning Framework for Implicit Hate Speech Detection in Videos</i> Mohammad Zia Ur Rehman, Anukriti Bhatnagar, Omkar Kabde, Shubhi Bansal and Dr. Nagendra Kumar	17209
<i>Improving Chain-of-Thought Reasoning via Quasi-Symbolic Abstractions</i> Leonardo Ranaldi, Marco Valentino and Andre Freitas	17222
<i>Information Extraction from Visually Rich Documents using LLM-based Organization of Documents into Independent Textual Segments</i> Aniket Bhattacharyya, Anurag Tripathi, Ujjal Das, Archan Karmakar, Amit Pathak and Maneesh Gupta	17241
<i>Enhancing Open-Domain Task-Solving Capability of LLMs via Autonomous Tool Integration from Gi- tHub</i> Bohan Lyu, Xin Cong, Heyang Yu, Pan Yang, Cheng Qian, Zihe Wang, Yujia Qin, Yining Ye, Yaxi Lu, Chen Qian, Zhong Zhang, Yukun Yan, Yankai Lin, Zhiyuan Liu and Maosong Sun	17257

<i>LLMs Can Simulate Standardized Patients via Agent Coevolution</i>	
Zhuoyun Du, LujieZheng LujieZheng, Renjun Hu, Yuyang Xu, Xiawei Li, Ying Sun, Wei Chen, Jian Wu, HaoLei Cai and HaoChao Ying	17278
<i>Donate or Create? Comparing Data Collection Strategies for Emotion-labeled Multimodal Social Media Posts</i>	
Christopher Bagdon, Aidan Combs, Carina Silberer and Roman Klingler	17307
<i>Which Demographics do LLMs Default to During Annotation?</i>	
Johannes Schäfer, Aidan Combs, Christopher Bagdon, Jiahui Li, Nadine Probol, Lynn Greschner, Sean Papay, Yarik Menchaca Resendiz, Aswathy Velutharambath, Amelie Wuehrl, Sabine Weber and Roman Klingler	17331
<i>Can You Really Trust Code Copilot? Evaluating Large Language Models from a Code Security Perspective</i>	
Yutao Mou, Xiao Deng, Yuxiao Luo, Shikun Zhang and Wei Ye	17349
<i>From Sub-Ability Diagnosis to Human-Aligned Generation: Bridging the Gap for Text Length Control via MarkerGen</i>	
Peiwen Yuan, Chuyi Tan, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Yueqi Zhang, Jiayi Shi, Boyuan Pan, Yao Hu and Kan Li	17370
<i>AGD: Adversarial Game Defense Against Jailbreak Attacks in Large Language Models</i>	
Shilong Pan, Zhiliang Tian, Zhen Huang, Wanlong Yu, Zhihua Wen, Xinwang Liu, Kai Lu, Minlie Huang and Dongsheng Li	17391
<i>SCOP: Evaluating the Comprehension Process of Large Language Models from a Cognitive View</i>	
Yongjie Xiao, Hongru Liang, Peixin Qin, Yao Zhang and Wenqiang Lei	17407
<i>Table-Critic: A Multi-Agent Framework for Collaborative Criticism and Refinement in Table Reasoning</i>	
Peiyong Yu, Guoxin Chen and Jingjing Wang	17432
<i>An Expanded Massive Multilingual Dataset for High-Performance Language Technologies (HPLT)</i>	
Laurie Burchell, Ona De Gibert Bonet, Nikolay Arefyev, Mikko Aulamo, Marta Bañón, Pinzhen Chen, Mariia Fedorova, Liane Guillou, Barry Haddow, Jan Hajič, Jindřich Helcl, Erik Henriksson, Mateusz Klimaszewski, Ville Komulainen, Andrey Kutuzov, Joonas Kytöniemi, Veronika Laippala, Petter Mæhlum, Bhavitvya Malik, Farrokh Mehryary, Vladislav Mikhailov, Nikita Moghe, Amanda Myntti, Dayyán O’Brien, Stephan Open, Proyag Pal, Jousia Piha, Sampo Pyysalo, Gema Ramírez-Sánchez, David Samuel, Pavel Stepachev, Jörg Tiedemann, Dušan Variš, Tereza Vojtěchová and Jaume Zaragoza-Bernabeu	17452
<i>Scaling Text-Rich Image Understanding via Code-Guided Synthetic Multimodal Data Generation</i>	
Yue Yang, Ajay Patel, Matt Deitke, Tanmay Gupta, Luca Weihs, Andrew Head, Mark Yatskar, Chris Callison-Burch, Ranjay Krishna, Aniruddha Kembhavi and Christopher Clark	17486
<i>Hierarchical Attention Generates Better Proofs</i>	
Jianlong Chen, Chao Li, Yang Yuan and Andrew C Yao	17506
<i>Agent-RewardBench: Towards a Unified Benchmark for Reward Modeling across Perception, Planning, and Safety in Real-World Multimodal Agents</i>	
Tianyi Men, Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu and Jun Zhao	17521
<i>It’s Not Bragging If You Can Back It Up: Can LLMs Understand Braggings?</i>	
Jingjie Zeng, Huayang Li, Liang Yang, Yuanyuan Sun and Hongfei Lin	17542
<i>A Troublemaker with Contagious Jailbreak Makes Chaos in Honest Towns</i>	
Tianyi Men, Pengfei Cao, Zhuoran Jin, Yubo Chen, Kang Liu and Jun Zhao	17561

<i>Meta-Learning Neural Mechanisms rather than Bayesian Priors</i>	
Michael Eric Goodale, Salvador Mascarenhas and Yair Lakretz	17588
<i>Shifting from Ranking to Set Selection for Retrieval Augmented Generation</i>	
Dahyun Lee, Yongrae Jo, Haeju Park and Moontae Lee	17606
<i>Understanding Large Language Model Vulnerabilities to Social Bias Attacks</i>	
Jiaxu Zhao, Meng Fang, Fanghua Ye, Ke Xu, Qin Zhang, Joey Tianyi Zhou and Mykola Pechenizkiy	17620
<i>ChatSOP: An SOP-Guided MCTS Planning Framework for Controllable LLM Dialogue Agents</i>	
Zhigen Li, Jianxiang Peng, Yanmeng Wang, Yong Cao, Tianhao Shen, Minghui Zhang, Linxi Su, Shang Wu, Yihang Wu, YuQian Wang, Ye Wang, Wei Hu, Jianfeng Li, Shaojun Wang, Jing Xiao and Deyi Xiong	17637
<i>Pixel-Level Reasoning Segmentation via Multi-turn Conversations</i>	
Dexian Cai, Xiaocui Yang, YongKang Liu, Daling Wang, Shi Feng, Yifei Zhang and Soujanya Poria	17660
<i>Fixing Distribution Shifts of LLM Self-Critique via On-Policy Self-Play Training</i>	
Rong Bao, Donglei Yu, Kai Fan and Minpeng Liao	17680
<i>Inferring Functionality of Attention Heads from their Parameters</i>	
Amit Elhelo and Mor Geva	17701
<i>Faithful and Robust LLM-Driven Theorem Proving for NLI Explanations</i>	
Xin Quan, Marco Valentino, Louise A. Dennis and Andre Freitas	17734
<i>Revealing the Deceptiveness of Knowledge Editing: A Mechanistic Analysis of Superficial Editing</i>	
Jiakuan Xie, Pengfei Cao, Yubo Chen, Kang Liu and Jun Zhao	17756
<i>Masking in Multi-hop QA: An Analysis of How Language Models Perform with Context Permutation</i>	
Wenyu Huang, Pavlos Vougiouklis, Mirella Lapata and Jeff Z. Pan	17781
<i>From Human Reading to NLM Understanding: Evaluating the Role of Eye-Tracking Data in Encoder-Based Models</i>	
Luca Dini, Lucia Domenichelli, Dominique Brunato and Felice Dell’Orletta	17796
<i>Optimizing Question Semantic Space for Dynamic Retrieval-Augmented Multi-hop Question Answering</i>	
Linhao YE, Lang Yu, Zhikai Lei, Qin Chen, Jie Zhou and Liang He	17814
<i>Insight Over Sight: Exploring the Vision-Knowledge Conflicts in Multimodal LLMs</i>	
Xiaoyuan Liu, Wenxuan Wang, Youliang Yuan, Jen-tse Huang, Qiuzhi Liu, Pinjia He and Zhao-peng Tu	17825
<i>SceneGenAgent: Precise Industrial Scene Generation with Coding Agent</i>	
Xiao Xia, Dan Zhang, Zibo Liao, Zhenyu Hou, Tianrui Sun, Jing Li, Ling Fu and Yuxiao Dong	17847
<i>ToolCoder: A Systematic Code-Empowered Tool Learning Framework for Large Language Models</i>	
Hanxing Ding, Shuchang Tao, Liang Pang, Zihao Wei, Jinyang Gao, Bolin Ding, Huawei Shen and Xueqi Cheng	17876
<i>Enhancing Text Editing for Grammatical Error Correction: Arabic as a Case Study</i>	
Bashar Alhafni and Nizar Habash	17892
<i>From Isolates to Families: Using Neural Networks for Automated Language Affiliation</i>	
Frederic Blum, Steffen Herbold and Johann-Mattis List	17915

<i>ELBA-Bench: An Efficient Learning Backdoor Attacks Benchmark for Large Language Models</i>	
Xuxu Liu, Siyuan Liang, Mengya Han, Yong Luo, Aishan Liu, Xiantao Cai, Zheng He and Da-cheng Tao	17928
<i>Less, but Better: Efficient Multilingual Expansion for LLMs via Layer-wise Mixture-of-Experts</i>	
Xue Zhang, Yunlong Liang, Fandong Meng, Songming Zhang, Yufeng Chen, Jinan Xu and Jie Zhou	17948
<i>When Harry Meets Superman: The Role of The Interlocutor in Persona-Based Dialogue Generation</i>	
Daniela Occhipinti, Marco Guerini and Malvina Nissim	17964
<i>ICR Probe: Tracking Hidden State Dynamics for Reliable Hallucination Detection in LLMs</i>	
Zhenliang Zhang, Xinyu Hu, Huixuan Zhang, Junzhe Zhang and Xiaojun Wan	17986
<i>Revisit Self-Debugging with Self-Generated Tests for Code Generation</i>	
Xiancai Chen, Zhengwei Tao, Kechi Zhang, Changzhi Zhou, Xinyu Zhang, Wanli Gu, Yuanpeng He, Mengdi Zhang, Xunliang Cai, Haiyan Zhao and Zhi Jin	18003
<i>InSertor: Speech Instruction Following with Unsupervised Interleaved Pre-training</i>	
Dingdong Wang, Jin Xu, Ruihang Chu, Zhifang Guo, Xiong Wang, Jincenzi Wu, Dongchao Yang, Shengpeng Ji and Junyang Lin	18024
<i>Exploring LLMs' Ability to Spontaneously and Conditionally Modify Moral Expressions through Text Manipulation</i>	
Candida Maria Greco, Lucio La Cava, Lorenzo Zangari and Andrea Tagarelli	18047
<i>Mixture of Ordered Scoring Experts for Cross-prompt Essay Trait Scoring</i>	
Po-Kai Chen, Bo-Wei Tsai, Shao Kuan Wei, Chien-Yao Wang, Jia-Ching Wang and Yi-Ting Huang	18071
<i>Sparse Logit Sampling: Accelerating Knowledge Distillation in LLMs</i>	
Anshumann Anshumann, Mohd Abbas Zaidi, Akhil Kedia, Jinwoo Ahn, Taehwak Kwon, Kangwook Lee, Haejun Lee and Joohyung Lee	18085
<i>Enhancing Spoken Discourse Modeling in Language Models Using Gestural Cues</i>	
Varsha Suresh, M. Hamza Mughal, Christian Theobalt and Vera Demberg	18109
<i>ExploraCoder: Advancing Code Generation for Multiple Unseen APIs via Planning and Chained Exploration</i>	
Yunkun Wang, Yue Zhang, Zhen Qin, Chen Zhi, Binhua Li, Fei Huang, Yongbin Li and Shuiguang Deng	18124
<i>Segment First or Comprehend First? Explore the Limit of Unsupervised Word Segmentation with Large Language Models</i>	
Zihong Zhang, Liqi He, Zuchao Li, Lefei Zhang, Hai Zhao and Bo Du	18146
<i>RUBY: An Effective Framework for Multi-Constraint Multi-Hop Question Generation</i>	
Wenzhuo Zhao and Shuangyin Li	18164
<i>Can Indirect Prompt Injection Attacks Be Detected and Removed?</i>	
Yulin Chen, Haoran Li, Yuan Sui, Yufei He, Yue Liu, Yangqiu Song and Bryan Hooi	18189
<i>Identifying Open Challenges in Language Identification</i>	
Rob Van Der Goot	18207
<i>The Distracting Effect: Understanding Irrelevant Passages in RAG</i>	
Chen Amiraz, Florin Cuconasu, Simone Filice and Zohar Karnin	18228

<i>Multilingual Encoder Knows more than You Realize: Shared Weights Pretraining for Extremely Low-Resource Languages</i>	
Zeli Su, Ziyin Zhang, Guixian Xu, Jianing Liu, Xu Han, Ting Zhang and Yushuang Dong .	18259
<i>Graphically Speaking: Unmasking Abuse in Social Media with Conversation Insights</i>	
Célia Nouri, Chloé Clavel and Jean-Philippe Cointet	18271
<i>CodeTool: Enhancing Programmatic Tool Invocation of LLMs via Process Supervision</i>	
YifeiLu YifeiLu, Fanghua Ye, Jian Li, Qiang Gao, Cheng Liu, Haibo Luo, Nan du, Xiaolong Li and Feiliang Ren	18287
<i>RARE: Retrieval-Augmented Reasoning Enhancement for Large Language Models</i>	
Hieu Tran, Zonghai Yao, Zhichao Yang, Junda Wang, Yifan Zhang, Shuo Han, Feiyun Ouyang and Hong yu	18305
<i>Defense Against Prompt Injection Attack by Leveraging Attack Techniques</i>	
Yulin Chen, Haoran Li, Zihao Zheng, Dekai Wu, Yangqiu Song and Bryan Hooi	18331
<i>Acquisition and Application of Novel Knowledge in Large Language Models</i>	
Ziyu Shang, Jianghan Liu, Zhizhao Luo, Peng Wang, Wenjun Ke, Jiajun Liu, Zijie Xu and Guozheng Li	18348
<i>DNCASR: End-to-End Training for Speaker-Attributed ASR</i>	
Xianrui Zheng, Chao Zhang and Phil Woodland	18369
<i>Exploring Persona Sentiment Sensitivity in Personalized Dialogue Generation</i>	
Yonghyun Jun and Hwanhee Lee	18384
<i>AntiLeakBench: Preventing Data Contamination by Automatically Constructing Benchmarks with Updated Real-World Knowledge</i>	
Xiaobao Wu, Liangming Pan, Yuxi Xie, Ruiwen Zhou, Shuai Zhao, Yubo Ma, Mingzhe Du, Rui Mao, Anh Tuan Luu and William Yang Wang	18403
<i>LLM-Guided Semantic-Aware Clustering for Topic Modeling</i>	
Jianghan Liu, Ziyu Shang, Wenjun Ke, Peng Wang, Zhizhao Luo, Jiajun Liu, Guozheng Li and Yining Li	18420
<i>Hierarchical Bracketing Encodings for Dependency Parsing as Tagging</i>	
Ana Ezquerro, David Vilares, Anssi Yli-Jyrä and Carlos Gómez-Rodríguez	18436
<i>OASIS: Order-Augmented Strategy for Improved Code Search</i>	
Gao Zuchen, Zizheng Zhan, Xianming LI, Erxin Yu, Haotian Zhang, Chenbin Chenbin, Yuqun Zhang and Jing Li	18451
<i>Can Large Language Models Detect Errors in Long Chain-of-Thought Reasoning?</i>	
Yancheng He, Shilong Li, Jiaheng Liu, Weixun Wang, Xingyuan Bu, Ge Zhang, Z.y. Peng, Zhaoxiang Zhang, Zhicheng Zheng, Wenbo Su and Bo Zheng	18468
<i>OmniAlign-V: Towards Enhanced Alignment of MLLMs with Human Preference</i>	
Xiangyu Zhao, Shengyuan Ding, Zicheng Zhang, Haiyan Huang, Maosongcao Maosongcao, Jiaqi Wang, Weiyun Wang, Xinyu Fang, Wenhai Wang, Guangtao Zhai, Hua Yang, Haodong Duan and Kai Chen	18490
<i>Tree-KG: An Expandable Knowledge Graph Construction Framework for Knowledge-intensive Domains</i>	
Songjie Niu, Kaisen Yang, Rui Zhao, Yichao Liu, Zonglin Li, Hongning Wang and Wenguang Chen	18516

<i>Measuring Data Diversity for Instruction Tuning: A Systematic Analysis and A Reliable Metric</i>	
Yuming Yang, Yang Nan, Junjie Ye, Shihan Dou, Xiao Wang, Shuo Li, Huijie Lv, Tao Gui, Qi Zhang and Xuanjing Huang	18530
<i>Micro-Act: Mitigate Knowledge Conflict in Question Answering via Actionable Self-Reasoning</i>	
Nan Huo, Jinyang Li, Bowen Qin, Ge Qu, Xiaolong Li, Xiaodong Li, Chenhao Ma and Reynold Cheng	18550
<i>Minimal Pair-Based Evaluation of Code-Switching</i>	
Igor Sterner and Simone Teufel	18575
<i>DNASpeech: A Contextualized and Situated Text-to-Speech Dataset with Dialogues, Narratives and Actions</i>	
Chuanqi Cheng, Hongda Sun, Bo Du, Shuo Shang, Xinrong Hu and Rui Yan	18599
<i>LLaMA-Omni 2: LLM-based Real-time Spoken Chatbot with Autoregressive Streaming Speech Synthesis</i>	
Qingkai Fang, Yan Zhou, Shoutao Guo, Shaolei Zhang and Yang Feng	18617
<i>Error Comparison Optimization for Large Language Models on Aspect-Based Sentiment Analysis</i>	
Qianlong Wang, Keyang Ding, Hengxin Gao, Hui Wang and Ruifeng Xu	18630
<i>The AI Gap: How Socioeconomic Status Affects Language Technology Interactions</i>	
Elisa Bassignana, Amanda Cercas Curry and Dirk Hovy	18647
<i>Probing LLMs for Multilingual Discourse Generalization Through a Unified Label Set</i>	
Florian Eichin, Yang Janet Liu, Barbara Plank and Michael A. Hedderich	18665
<i>Crowdsource, Crawl, or Generate? Creating SEA-VL, a Multicultural Vision-Language Dataset for Southeast Asia</i>	
Samuel Cahyawijaya, Holy Lovenia, Joel Ruben Antony Moniz, Tack Hwa Wong, Mohammad Rifqi Farhansyah, Thant Thiri Maung, Frederikus Hudi, David Anugraha, Muhammad Ravi Shulthan Habibi, Muhammad Reza Qorib, Amit Agarwal, Joseph Marvin Imperial, Hitesh Laxmichand Patel, Vicky Feliren, Bahrul Ilmi Nasution, Manuel Antonio Rufino, Genta Indra Winata, Rian Adam Rajagede, Carlos Rafael Catalan, Mohamed Fazli Mohamed Imam, Priyaranjan Pattnayak, Salsabila Zahirah Pranida, Kevin Pratama, Yeshil Bangera, Adisai Na-Thalang, Patricia Nicole Monderin, Yueqi Song, Christian Simon, Lynnette Hui Xian Ng, Richardy Lobo Sapan, Taki Hasan Rafi, Bin Wang, Supryadi, Kanyakorn Veerakanjana, Piyalitt Ittichaiwong, Matthew Theodore Roque, Karissa Vincentio, Takdanai Kreangphet, Phakphum Arkaew, Kadek Hendrawan Palgunadi, Yanzhi Yu, Rochana Prih Hastuti, William Nixon, Mithil Bangera, Adrian Xuan Wei Lim, Aye Hninn Khine, Hanif Muhammad Zhafran, Teddy Ferdinan, Audra Aurora Izzani, Ayushman Singh, Evan Evan, Jauza Akbar Krito, Michael Anugraha, Fenal Ashokbhai Ilasariya, Haochen Li, John Amadeo Daniswara, Filbert Aurelian Tjiaranata, Eryawan Presma Yulianrifat, Can Udomcharoenchaikit, Fadil Risdian Ansori, Mahardika Krisna Ihsani, Giang Nguyen, Anab Maulana Barik, Dan John Velasco, Rifo Ahmad Genadi, Saptarshi Saha, Chengwei Wei, Isaiah Edri W. Flores, Kenneth Chen Ko Han, Anjela Gail D. Santos, Wan Shen Lim, Kaung Si Phyo, Tim Santos, Meisyyarah Dwiastuti, Jiayun Luo, Jan Christian Blaise Cruz, Ming Shan Hee, Ikhlusal Akmal Hanif, M.Alif Al Hakim, Muhammad Rizky Sya'ban, Kun Kerdtthaisong, Lester James Validad Miranda, Fajri Koto, Tirana Noor Fatyanosa, Alham Fikri Aji, Jostin Jerico Rosal, Jun Kevin, Robert Wijaya, Onno P. Kampman, Ruo Chen Zhang, Börje F. Karlsson and Peerat Limkonchotiwat	18685
<i>Soundwave: Less is More for Speech-Text Alignment in LLMs</i>	
Yuhao Zhang, Zhiheng Liu, Fan Bu, Ruiyu Zhang, Benyou Wang and Haizhou Li	18718

<i>RoToR: Towards More Reliable Responses for Order-Invariant Inputs</i>	
Soyoung Yoon, Dongha Ahn, Youngwon Lee, Minkyu Jung, HyungJoo Jang and Seung-won Hwang	18739
<i>Global MMLU: Understanding and Addressing Cultural and Linguistic Biases in Multilingual Evaluation</i>	
Shivalika Singh, Angelika Romanou, Clémentine Fourier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiawat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Sebastian Ruder, Wei-Yin Ko, Antoine Bosselut, Alice Oh, Andre Martins, Leshem Choshen, Daphne Ippolito, Enzo Ferrante, Marzieh Fadaee, Beyza Ermis and Sara Hooker	18761
<i>Improving Dialogue Discourse Parsing through Discourse-aware Utterance Clarification</i>	
Yaxin Fan, Peifeng Li and Qiaoming Zhu	18800
<i>ImPart: Importance-Aware Delta-Sparsification for Improved Model Compression and Merging in LLMs</i>	
Yan Yang, Yixia Li, Hongru Wang, Xuetao Wei, James Jianqiao Yu, Yun Chen and Guanhua Chen	18817
<i>Words of Warmth: Trust and Sociability Norms for over 26k English Words</i>	
Saif M. Mohammad	18830
<i>BehaviorBox: Automated Discovery of Fine-Grained Performance Differences Between Language Models</i>	
Lindia Tjauatja and Graham Neubig	18851
<i>HAF-RM: A Hybrid Alignment Framework for Reward Model Training</i>	
Shujun Liu, Xiaoyu Shen, Yuhang Lai, Siyuan Wang, Shengbin Yue, Zengfeng Huang, Xuanjing Huang and Zhongyu Wei	18874
<i>CULEMO: Cultural Lenses on Emotion - Benchmarking LLMs for Cross-Cultural Emotion Understanding</i>	
Tadesse Destaw Belay, Ahmed Haj Ahmed, Alvin C Grissom II, Iqra Ameer, Grigori Sidorov, Olga Kolesnikova and Seid Muhie Yimam	18894
<i>DiffPO: Diffusion-styled Preference Optimization for Inference Time Alignment of Large Language Models</i>	
Ruizhe Chen, Wenhao Chai, Zhifei Yang, Xiaotian Zhang, Ziyang Wang, Tony Quek, Joey Tianyi Zhou, Soujanya Poria and Zuozhu Liu	18910
<i>MemeQA: Holistic Evaluation for Meme Understanding</i>	
Khoi P. N. Nguyen, Terrence Li, Derek Lou Zhou, Gabriel Xiong, Pranav Balu, Nandhan Alahari, Alan Huang, Tanush Chauhan, Harshavardhan Bala, Emre Guzelordu, Affan Kashfi, Aaron Xu, Suyesh Shrestha, Megan Vu, Jerry Wang and Vincent Ng	18926
<i>LoGU: Long-form Generation with Uncertainty Expressions</i>	
Ruihan Yang, Caiqi Zhang, Zhisong Zhang, Xinting Huang, Sen Yang, Nigel Collier, Dong Yu and Deqing Yang	18947
<i>KiRAG: Knowledge-Driven Iterative Retriever for Enhancing Retrieval-Augmented Generation</i>	
Jinyuan Fang, Zaiqiao Meng and Craig MacDonald	18969
<i>Enhancing Lexicon-Based Text Embeddings with Large Language Models</i>	
Yibin Lei, Tao Shen, Yu Cao and Andrew Yates	18986

<i>CoCoLex: Confidence-guided Copy-based Decoding for Grounded Legal Text Generation</i>	
Santosh T.y.s.s, Youssef Tarek Elkhayat, Oana Ichim, Pranav Shetty, Dongsheng Wang, Zhiqiang Ma, Armineh Nourbakhsh and Xiaomo Liu	19002
<i>Beyond N-Grams: Rethinking Evaluation Metrics and Strategies for Multilingual Abstractive Summarization</i>	
Itai Mondshine, Tzuf Paz-Argaman and Reut Tsarfaty	19019
<i>CC-Tuning: A Cross-Lingual Connection Mechanism for Improving Joint Multilingual Supervised Fine-Tuning</i>	
Yangfan Ye, Xiaocheng Feng, Zekun Yuan, Xiachong Feng, Libo Qin, Lei Huang, Weitao Ma, Yichong Huang, Zhirui Zhang, Yunfei Lu, Xiaohui Yan, Duyu Tang, Dandan Tu and Bing Qin. .	19036
<i>SConU: Selective Conformal Uncertainty in Large Language Models</i>	
Zhiyuan Wang, Qingni Wang, Yue Zhang, Tianlong Chen, Xiaofeng Zhu, Xiaoshuang Shi and Kaidi Xu	19052
<i>MegaPairs: Massive Data Synthesis for Universal Multimodal Retrieval</i>	
Junjie Zhou, Yongping Xiong, Zheng Liu, Ze Liu, Shitao Xiao, Yuezhe Wang, Bo Zhao, Chen Jason Zhang and Defu Lian	19076
<i>When GPT Spills the Tea: Comprehensive Assessment of Knowledge File Leakage in GPTs</i>	
Xinyue Shen, Yun Shen, Michael Backes and Yang Zhang	19096
<i>UniCodec: Unified Audio Codec with Single Domain-Adaptive Codebook</i>	
Yidi Jiang, Qian Chen, Shengpeng Ji, Yu Xi, Wen Wang, Chong Zhang, Xianghu Yue, ShiLiang Zhang and Haizhou Li	19112
<i>KERL: Knowledge-Enhanced Personalized Recipe Recommendation using Large Language Models</i>	
Fnu Mohbat and Mohammed J Zaki	19125
<i>Multilingual Arbitration: Optimizing Data Pools to Accelerate Multilingual Progress</i>	
Ayomide Odumakinde, Daniel D’souza, Pat Verga, Beyza Ermis and Sara Hooker	19142
<i>Controlled Low-Rank Adaptation with Subspace Regularization for Continued Training on Large Language Models</i>	
Yuheng Lu, Bingshuo Qian, Caixia Yuan, Huixing Jiang and Xiaojie Wang	19165
<i>Chinese SimpleQA: A Chinese Factuality Evaluation for Large Language Models</i>	
Yancheng He, Shilong Li, Jiaheng Liu, Yingshui Tan, Weixun Wang, Hui Huang, Xingyuan Bu, Hangyu Guo, Chengwei Hu, Boren Zheng, Zhuoran Lin, Dekai Sun, Zhicheng Zheng, Wenbo Su and Bo Zheng	19182
<i>PVP: An Image Dataset for Personalized Visual Persuasion with Persuasion Strategies, Viewer Characteristics, and Persuasiveness Ratings</i>	
Junseo Kim, Jongwook Han, Dongmin Choi, Jongwook Yoon, Eun-Ju Lee and Yohan Jo. .	19209
<i>Any Information Is Just Worth One Single Screenshot: Unifying Search With Visualized Information Retrieval</i>	
Zheng Liu, Ze Liu, Zhengyang Liang, Junjie Zhou, Shitao Xiao, Chao Gao, Chen Jason Zhang and Defu Lian	19238
<i>Tunable LLM-based Proactive Recommendation Agent</i>	
Mingze Wang, Chongming Gao, Wenjie Wang, Yangyang Li and Fuli Feng	19262

<i>AgentRM: Enhancing Agent Generalization with Reward Modeling</i>	
Yu Xia, Jingru Fan, Weize Chen, Siyu Yan, Xin Cong, Zhong Zhang, Yaxi Lu, Yankai Lin, Zhiyuan Liu and Maosong Sun	19277
<i>From Outcomes to Processes: Guiding PRM Learning from ORM for Inference-Time Alignment</i>	
Bin Xie, Bingbing Xu, Yige Yuan, Shengmao Zhu and Huawei Shen	19291
<i>Segment-Based Attention Masking for GPTs</i>	
Shahar Katz, Liran Ringel, Yaniv Romano and Lior Wolf	19308
<i>Cramming 1568 Tokens into a Single Vector and Back Again: Exploring the Limits of Embedding Space Capacity</i>	
Yuri Kuratov, Mikhail Arkhipov, Aydar Bulatov and Mikhail Burtsev	19323
<i>Bi-Tuning with Collaborative Information for Controllable LLM-based Sequential Recommendation</i>	
Xinyu Zhang, Linmei Hu, Luhao Zhang, Wentao Cheng, Yashen Wang, Ge Shi, Chong Feng and Liqiang Nie	19340
<i>A Modular Approach for Clinical SLMs Driven by Synthetic Data with Pre-Instruction Tuning, Model Merging, and Clinical-Tasks Alignment</i>	
Jean-Philippe Corbeil, Amin Dada, Jean-Michel Attendu, Asma Ben Abacha, Alessandro Sordoni, Lucas Caccia, Francois Beaulieu, Thomas Lin, Jens Kleesiek and Paul Vozila	19352
<i>DIVE into MoE: Diversity-Enhanced Reconstruction of Large Language Models from Dense into Mixture-of-Experts</i>	
Yuchen Feng, Bowen Shen, Naibin Gu, Jiaxuan Zhao, Peng Fu, Zheng Lin and Weiping Wang	19375
<i>DAC: A Dynamic Attention-aware Approach for Task-Agnostic Prompt Compression</i>	
Yi Zhao, Zuchao Li, Hai Zhao, Baoyuan Qi and Liu Guoming	19395
<i>Computation Mechanism Behind LLM Position Generalization</i>	
Chi Han and Heng Ji	19408
<i>IPO: Your Language Model is Secretly a Preference Classifier</i>	
Shivank Garg, Ayush Singh, Shweta Singh and Paras Chopra	19425
<i>Reversal of Thought: Enhancing Large Language Models with Preference-Guided Reverse Reasoning Warm-up</i>	
Jiahao Yuan, Dehui du, Hao Zhang, Zixiang Di and Usman Naseem	19442
<i>Déjà Vu? Decoding Repeated Reading from Eye Movements</i>	
Yoav Meiri, Omer Shubi, Cfir Avraham Hadar, Ariel Kreisberg Nitzav and Yevgeni Berzak	19460
<i>LLMs can be easily Confused by Instructional Distractions</i>	
Yerin Hwang, Yongil Kim, Jahyun Koo, Taegwan Kang, Hyunkyung Bae and Kyomin Jung	19483
<i>PlanGenLLMs: A Modern Survey of LLM Planning Capabilities</i>	
Hui Wei, Zihao Zhang, Shenghua He, Tian Xia, Shijia Pan and Fei Liu	19497
<i>IAM: Efficient Inference through Attention Mapping between Different-scale LLMs</i>	
Yi Zhao, Zuchao Li and Hai Zhao	19522
<i>nvAgent: Automated Data Visualization from Natural Language via Collaborative Agent Workflow</i>	
Geliang Ouyang, Jingyao Chen, Zhihe Nie, Yi Gui, Yao Wan, Hongyu Zhang and Dongping Chen	19534

<i>ZIPA: A family of efficient models for multilingual phone recognition</i> Jian Zhu, Farhan Samir, Eleanor Chodroff and David R. Mortensen	19568
<i>GRACE: A Granular Benchmark for Evaluating Model Calibration against Human Calibration</i> Yoo Yeon Sung, Eve Fleisig, Yu Hou, Ishan Upadhyay and Jordan Lee Boyd-Graber	19586
<i>Dynamic Evaluation with Cognitive Reasoning for Multi-turn Safety of Large Language Models</i> Lanxue Zhang, Yanan Cao, Yuqiang Xie, Fang Fang and Yangxi Li	19588
<i>From Tools to Teammates: Evaluating LLMs in Multi-Session Coding Interactions</i> Nathanaël Carraz Rakotonirina, Mohammed Hamdy, Jon Ander Campos, Lucas Weber, Alberto Testoni, Marzieh Fadaee, Sandro Pezzelle and Marco Del Tredici	19609
<i>Guiding not Forcing: Enhancing the Transferability of Jailbreaking Attacks on LLMs via Removing Superfluous Constraints</i> Junxiao Yang, Zhexin Zhang, Shiyao Cui, Hongning Wang and Minlie Huang	19643
<i>Multilingual Text-to-Image Generation Magnifies Gender Stereotypes</i> Felix Friedrich, Katharina Hämmerl, Patrick Schramowski, Manuel Brack, Jindřich Libovický, Alexander Fraser and Kristian Kersting	19656
<i>Adversarial Alignment with Anchor Dragging Drift (A^3D^2): Multimodal Domain Adaptation with Partially Shifted Modalities</i> Jun Sun, Xinxin Zhang, Simin Hong, Jian Zhu and Lingfang Zeng	19680
<i>A Reality Check on Context Utilisation for Retrieval-Augmented Generation</i> Lovisa Hagström, Sara Vera Marjanovic, Haeun Yu, Arnav Arora, Christina Lioma, Maria Maestro, Pepa Atanasova and Isabelle Augenstein	19691
<i>CU-MAM: Coherence-Driven Unified Macro-Structures for Argument Mining</i> Debela Gemechu and Chris Reed	19731
<i>Safer or Luckier? LLMs as Safety Evaluators Are Not Robust to Artifacts</i> Hongyu Chen and Seraphina Goldfarb-Tarrant	19750
<i>Text-to-ES Bench: A Comprehensive Benchmark for Converting Natural Language to Elasticsearch Query</i> DonggeXue DonggeXue, Zhili Pu, Zhentao Xia, Hongli Sun, Ruihui Hou, Guangya Yu, Yupian Lin, Yongqi Fan, Jingping Liu and Tong Ruan	19767
<i>AlignDistil: Token-Level Language Model Alignment as Adaptive Policy Distillation</i> Songming Zhang, Xue Zhang, Tong Zhang, Bojie Hu, Yufeng Chen and Jinan Xu	19791
<i>DARS: Dynamic Action Re-Sampling to Enhance Coding Agent Performance by Adaptive Tree Traversal</i> Vaibhav Aggarwal, Ojasv Kamal, Abhinav Japesh, Zhijing Jin and Bernhard Schölkopf	19808
<i>Steering off Course: Reliability Challenges in Steering Language Models</i> Patrick Queiroz Da Silva, Hari Sethuraman, Dheeraj Rajagopal, Hannaneh Hajishirzi and Sachin Kumar	19856
<i>Impartial Multi-task Representation Learning via Variance-invariant Probabilistic Decoding</i> Dou Hu, Lingwei Wei, Wei Zhou and Songlin Hu	19883
<i>If Eleanor Rigby Had Met ChatGPT: A Study on Loneliness in a Post-LLM World</i> Adrian de Wynter	19898

<i>Integrating Audio, Visual, and Semantic Information for Enhanced Multimodal Speaker Diarization on Multi-party Conversation</i>	
Luyao Cheng, Hui Wang, Chong Deng, Siqi Zheng, Yafeng Chen, Rongjie Huang, Qinglin Zhang, Qian Chen, Xihao Li and Wen Wang	19914
<i>Vulnerability of LLMs to Vertically Aligned Text Manipulations</i>	
Zhecheng Li, Yiwei Wang, Bryan Hooi, Yujun Cai, Zhen Xiong, Nanyun Peng and Kai-Wei Chang	19929
<i>AutoMixer: Checkpoint Artifacts as Automatic Data Mixers</i>	
Ernie Chang, Yang Li, Patrick Huber, Vish Vogeti, David Kant, Yangyang Shi and Vikas Chandra	19942
<i>Generalized Attention Flow: Feature Attribution for Transformer Models via Maximum Flow</i>	
Behrooz Azarkhalili and Maxwell W. Libbrecht	19954
<i>Beyond Prompting: An Efficient Embedding Framework for Open-Domain Question Answering</i>	
Zhanghao Hu, Hanqi Yan, Qinglin Zhu, Zhenyi Shen, Yulan He and Lin Gui	19975
<i>AIR-Bench: Automated Heterogeneous Information Retrieval Benchmark</i>	
Jianlyu Chen, Nan Wang, Chaofan Li, Bo Wang, Shitao Xiao, Han Xiao, Hao Liao, Defu Lian and Zheng Liu	19991
<i>We-Math: Does Your Large Multimodal Model Achieve Human-like Mathematical Reasoning?</i>	
Runqi Qiao, Qiuna Tan, Guanting Dong, MinhuiWu MinhuiWu, Chong Sun, Xiaoshuai Song, Jiapeng Wang, Zhuoma GongQue, Shanglin Lei, YiFan Zhang, Zhe Wei, Miaoxuan Zhang, Runfeng Qiao, Xiao Zong, Yida Xu, Peiqing Yang, Zhimin Bao, Muxi Diao, Chen Li and Honggang Zhang	20023
<i>Modeling the Evolution of English Noun Compounds with Feature-Rich Diachronic Compositionality Prediction</i>	
Filip Miletic and Sabine Schulte Im Walde	20071
<i>What's the Difference? Supporting Users in Identifying the Effects of Prompt and Model Changes Through Token Patterns</i>	
Michael A. Hedderich, Anyi Wang, Raoyuan Zhao, Florian Eichin, Jonas Fischer and Barbara Plank	20093
<i>V-Oracle: Making Progressive Reasoning in Deciphering Oracle Bones for You and Me</i>	
Runqi Qiao, Qiuna Tan, Guanting Dong, MinhuiWu MinhuiWu, Jiapeng Wang, YiFan Zhang, Zhuoma GongQue, Chong Sun, Yida Xu, Yadong Xue, Ye Tian, Zhimin Bao, Lan Yang, Chen Li and Honggang Zhang	20124
<i>Unveiling Cultural Blind Spots: Analyzing the Limitations of mLLMs in Procedural Text Comprehension</i>	
Amir Hossein Yari and Fajri Koto	20151
<i>Improving Language and Modality Transfer in Translation by Character-level Modeling</i>	
Ioannis Tsiamas, David Dale and Marta R. Costa-jussà	20171
<i>DialUp! Modeling the Language Continuum by Adapting Models to Dialects and Dialects to Models</i>	
Niyati Bafna, Emily Chang, Nathaniel Romney Robinson, David R. Mortensen, Kenton Murray, David Yarowsky and Hale Sirin	20188
<i>AutoMixAlign: Adaptive Data Mixing for Multi-Task Preference Optimization in LLMs</i>	
Nicholas E. Corrado, Julian Katz-Samuels, Adithya M Devraj, Hyokun Yun, Chao Zhang, Yi Xu, Yi Pan, Bing Yin and Trishul Chilimbi	20234

<i>Modeling Complex Semantics Relation with Contrastively Fine-Tuned Relational Encoders</i> Naïm Es-sebbani, Esteban Marquer and Zied Bouraoui	20259
<i>Error-driven Data-efficient Large Multimodal Model Tuning</i> Barry Menglong Yao, Qifan Wang and Lifu Huang	20289
<i>Planning with Diffusion Models for Target-Oriented Dialogue Systems</i> Hanwen Du, Bo Peng and Xia Ning	20307
<i>Interactive and Expressive Code-Augmented Planning with Large Language Models</i> Anthony Zhe Liu, Xinhe Wang, Jacob Sansom, Yao Fu, Jongwook Choi, Sungryull Sohn, Jae- kyeom Kim and Honglak Lee	20330
<i>Synergistic Weak-Strong Collaboration by Aligning Preferences</i> Yizhu Jiao, Xuchao Zhang, Zhaoyang Wang, Yubo Ma, Zhun Deng, Rujia Wang, Chetan Bansal, Saravan Rajmohan, Jiawei Han and Huaxiu Yao	20355
<i>Understanding Silent Data Corruption in LLM Training</i> Jeffrey Jian Ma, Hengzhi Pei, Leonard Lausen and George Karypis	20372
<i>Align-SLM: Textless Spoken Language Models with Reinforcement Learning from AI Feedback</i> Guan-Ting Lin, Prashanth Gurunath Shivakumar, Aditya Gourav, Yile Gu, Ankur Gandhe, Hung- yi Lee and Ivan Bulyko	20395
<i>Can LLMs Help Uncover Insights about LLMs? A Large-Scale, Evolving Literature Analysis of Frontier LLMs</i> Jungsoo Park, Junmo Kang, Gabriel Stanovsky and Alan Ritter	20412
<i>BIG5-CHAT: Shaping LLM Personalities Through Training on Human-Grounded Data</i> Wenkai Li, Jiarui Liu, Andy Liu, Xuhui Zhou, Mona T. Diab and Maarten Sap	20434
<i>Deep Temporal Reasoning in Video Language Models: A Cross-Linguistic Evaluation of Action Dura- tion and Completion through Perfect Times</i> Olga Loginova and Sofia Ortega Loguinova	20472
<i>Amplifying Trans and Nonbinary Voices: A Community-Centred Harm Taxonomy for LLMs</i> Eddie L. Ungless, Sunipa Dev, Cynthia L. Bennett, Rebecca Gulotta, Jasmijn Bastings and Remi Denton	20503
<i>Enhancing Human Evaluation in Machine Translation with Comparative Judgement</i> Yixiao Song, Parker Riley, Daniel Deutsch and Markus Freitag	20536
<i>Infogen: Generating Complex Statistical Infographics from Documents</i> Akash Ghosh, Aparna Garimella, Pritika Ramu, Sambaran Bandyopadhyay and Sriparna Saha 20552	
<i>Partial Colexifications Improve Concept Embeddings</i> Arne Rubehn and Johann-Mattis List	20571
<i>Improved Unbiased Watermark for Large Language Models</i> Ruibo Chen, Yihan Wu, Junfeng Guo and Heng Huang	20587
<i>MaCP: Minimal yet Mighty Adaptation via Hierarchical Cosine Projection</i> Yixian Shen, Qi Bi, Jia-hong Huang, Hongyi Zhu, Andy D. Pimentel and Anuj Pathania ..	20602
<i>Multi-Attribute Steering of Language Models via Targeted Intervention</i> Duy Nguyen, Archiki Prasad, Elias Stengel-Eskin and Mohit Bansal	20619

<i>AdaptAgent: Adapting Multimodal Web Agents with Few-Shot Learning from Human Demonstrations</i> Gaurav Verma, Rachneet Kaur, Nishan Srishankar, Zhen Zeng, Tucker Balch and Manuela Veloso	20635
<i>Can LLMs Identify Critical Limitations within Scientific Research? A Systematic Evaluation on AI Research Papers</i> Zhijian Xu, Yilun Zhao, Manasi Patwardhan, Lovekesh Vig and Arman Cohan	20652
<i>On the Acquisition of Shared Grammatical Representations in Bilingual Language Models</i> Catherine Arnett, Tyler A. Chang, James A. Michaelov and Ben Bergen	20707
<i>Using Shapley interactions to understand how models use structure</i> Divyansh Singhvi, Diganta Misra, Andrej Erkelens, Raghav Jain, Isabel Papadimitriou and Naomi Saphra	20727
<i>Adversarial Tokenization</i> Renato Geh, Zilei Shao and Guy Van Den Broeck	20738
<i>Classifying Unreliable Narrators with Large Language Models</i> Anneliese Brei, Katharine Henry, Abhishek Sharma, Shashank Srivastava and Snigdha Chaturvedi	20766
<i>ConceptCarve: Dynamic Realization of Evidence</i> Eylon Caplan and Dan Goldwasser	20792
<i>QQSUM: A Novel Task and Model of Quantitative Query-Focused Summarization for Review-based Product Question Answering</i> An Quang Tang, Xiuzhen Zhang, Minh Ngoc Dinh and Zhuang Li	20810
<i>Navigating Rifts in Human-LLM Grounding: Study and Benchmark</i> Omar Shaikh, Hussein Mozannar, Gagan Bansal, Adam Fourney and Eric Horvitz	20832
<i>Substance over Style: Evaluating Proactive Conversational Coaching Agents</i> Vidya Srinivas, Xuhai Xu, Xin Liu, Kumar Ayush, Isaac Galatzer-Levy, Shwetak Patel, Daniel McDuff and Tim Althoff	20848
<i>Open-World Planning via Lifted Regression with LLM-Inferred Affordances for Embodied Agents</i> Xiaotian Liu, Ali Pesaranghader, Hanze LI, Punyaphat Sukcharoenchaikul, Jaehong Kim, Tanmana Sadhu, Hyejeong Jeon and Scott Sanner	20881
<i>(RSA)²: A Rhetorical-Strategy-Aware Rational Speech Act Framework for Figurative Language Understanding</i> Cesare Spinoso-Di Piano, David Eric Austin, Pablo Piantanida and Jackie CK Cheung	20898
<i>SYNTHIA: Novel Concept Design with Affordance Composition</i> Hyeonjeong Ha, Xiaomeng Jin, Jeonghwan Kim, Jiateng Liu, Zhenhailong Wang, Khanh Duy Nguyen, Ansel Blume, Nanyun Peng, Kai-Wei Chang and Heng Ji	20939
<i>Consistent Client Simulation for Motivational Interviewing-based Counseling</i> Yizhe Yang, Palakorn Achananuparp, Heyan Huang, Jing Jiang, Nicholas Gabriel Lim, Cameron Tan Shi Ern, Phey Ling Kit, Jenny Giam Xiuhui, John Pinto and Ee-Peng Lim	20959
<i>AUTALIC: A Dataset for Anti-AUTistic Ableist Language In Context</i> Naba Rizvi, Harper Strickland, Daniel Gitelman, Alexis Morales Flores, Tristan Cooper, Aekta Kallepalli, Akshat Alurkar, Haaset Owens, Saleha Ahmedi, Isha Khirwadkar, Imani N. S. Munyaka and Nedjma Ousidhoum	20999

<i>Structural Reasoning Improves Molecular Understanding of LLM</i> Yunhui Jang, Jaehyung Kim and Sungsoo Ahn	21016
<i>CAMI: A Counselor Agent Supporting Motivational Interviewing through State Inference and Topic Exploration</i> Yizhe Yang, Palakorn Achananuparp, Heyan Huang, Jing Jiang, Phey Ling Kit, Nicholas Gabriel Lim, Cameron Tan Shi Ern and Ee-Peng Lim	21037
<i>Know You First and Be You Better: Modeling Human-Like User Simulators via Implicit Profiles</i> Kuang Wang, Xianfei Li, Shenghao Yang, Li Zhou, Feng Jiang and Haizhou Li	21082
<i>Targeted Syntactic Evaluation for Grammatical Error Correction</i> Aomi Koyama, Masato Mita, Su-Youn Yoon, Yasufumi Takama and Mamoru Komachi ...	21108
<i>VF-Eval: Evaluating Multimodal LLMs for Generating Feedback on AIGC Videos</i> Tingyu Song, Tongyan Hu, Guo Gan and Yilun Zhao	21126
<i>Language Model Fine-Tuning on Scaled Survey Data for Predicting Distributions of Public Opinions</i> Joseph Suh, Erfan Jahanparast, Suhong Moon, Minwoo Kang and Serina Chang	21147
<i>TESS 2: A Large-Scale Generalist Diffusion Language Model</i> Jaesung Tae, Hamish Ivison, Sachin Kumar and Arman Cohan	21171
<i>KatFishNet: Detecting LLM-Generated Korean Text through Linguistic Feature Analysis</i> Shinwoo Park, Shubin Kim, Do-Kyung Kim and Yo-Sub Han	21189
<i>Uncovering the Impact of Chain-of-Thought Reasoning for Direct Preference Optimization: Lessons from Text-to-SQL</i> Hanbing Liu, Haoyang Li, Xiaokang Zhang, Ruotong Chen, Haiyong Xu, Tian Tian, Qi Qi and Jing Zhang	21223
<i>On Generalization across Measurement Systems: LLMs Entail More Test-Time Compute for Underrepresented Cultures</i> Minh Duc Bui, Kyung Eun Park, Goran Glavaš, Fabian David Schmidt and Katharina Von Der Wense	21262
<i>CORDIAL: Can Multimodal Large Language Models Effectively Understand Coherence Relationships?</i> Aashish Anantha Ramakrishnan, Aadarsh Anantha Ramakrishnan and Dongwon Lee	21277
<i>Veracity Bias and Beyond: Uncovering LLMs' Hidden Beliefs in Problem-Solving Reasoning</i> Yue Zhou and Barbara Di Eugenio	21298
<i>Optimal Transport-Based Token Weighting scheme for Enhanced Preference Optimization</i> Meng Li, Guangda Huzhang, Haibo Zhang, Xiting Wang and Anxiang Zeng	21311
<i>LLM Meets Scene Graph: Can Large Language Models Understand and Generate Scene Graphs? A Benchmark and Empirical Study</i> Dongil Yang, Minjin Kim, Sunghwan Kim, Beong-woo Kwak, Minjun Park, Jinseok Hong, Woontack Woo and Jinyoung Yeo	21335
<i>Beyond Frameworks: Unpacking Collaboration Strategies in Multi-Agent Systems</i> Haochun Wang, Sendong Zhao, Jingbo Wang, Zewen Qiang, Bing Qin and Ting Liu	21361
<i>The Invisible Hand: Unveiling Provider Bias in Large Language Models for Code Generation</i> Xiaoyu Zhang, Juan Zhai, Shiqing Ma, Qingshuang Bao, Weipeng Jiang, Qian Wang, Chao Shen and Yang Liu	21376

<i>K/DA: Automated Data Generation Pipeline for Detoxifying Implicitly Offensive Language in Korean</i> Minkyong Jeon, Hyemin Jeong, Yerang Kim, Jiyoung Kim, Jae Hyeon Cho and Byung-Jun Lee	21404
<i>THOR-MoE: Hierarchical Task-Guided and Context-Responsive Routing for Neural Machine Translation</i> Yunlong Liang, Fandong Meng and Jie Zhou	21433
<i>Neuron Empirical Gradient: Discovering and Quantifying Neurons' Global Linear Controllability</i> Xin Zhao, Zehui Jiang and Naoki Yoshinaga	21446
<i>Can Third Parties Read Our Emotions?</i> Jiayi Li, Yingfan Zhou, Pranav Narayanan Venkit, Halima Binte Islam, Sneha Arya, Shomir Wilson and Sarah Rajtmajer	21478
<i>OZSpeech: One-step Zero-shot Speech Synthesis with Learned-Prior-Conditioned Flow Matching</i> Nghia Huynh Nguyen Hieu, Ngoc Son Nguyen, Huynh Nguyen Dang, Thieu Vo, Truong-Son Hy and Van Nguyen	21500
<i>World Modeling Makes a Better Planner: Dual Preference Optimization for Embodied Task Planning</i> Siyin Wang, Zhaoye Fei, Qinyuan Cheng, Shiduo Zhang, Panpan Cai, Jinlan Fu and Xipeng Qiu	21518
<i>JailbreakRadar: Comprehensive Assessment of Jailbreak Attacks Against LLMs</i> Junjie Chu, Yugeng Liu, Ziqing Yang, Xinyue Shen, Michael Backes and Yang Zhang	21538
<i>CogniBench: A Legal-inspired Framework and Dataset for Assessing Cognitive Faithfulness of Large Language Models</i> Xiaqiang Tang, Jian Li, Keyu Hu, Nan du, Xiaolong Li, Xi Zhang, Weigao Sun and Sihong Xie	21567
<i>Neural Incompatibility: The Unbridgeable Gap of Cross-Scale Parametric Knowledge Transfer in Large Language Models</i> Yuqiao Tan, Shizhu He, Kang Liu and Jun Zhao	21586
<i>Enhancing Mathematical Reasoning in LLMs by Stepwise Correction</i> Zhenyu Wu, Qingkai Zeng, Zhihan Zhang, Zhaoxuan Tan, Chao Shen and Meng Jiang	21602
<i>PsyDial: A Large-scale Long-term Conversational Dataset for Mental Health Support</i> Huachuan Qiu and Zhenzhong Lan	21624
<i>Enhancing Goal-oriented Proactive Dialogue Systems via Consistency Reflection and Correction</i> Didi Zhang, Yaxin Fan, Peifeng Li and Qiaoming Zhu	21656
<i>Exclusion of Thought: Mitigating Cognitive Load in Large Language Models for Enhanced Reasoning in Multiple-Choice Tasks</i> Qihang Fu, Yongbin Qin, Ruizhang Huang, Yanping Chen, Yulin Zhou and Lintao Long	21673
<i>Registering Source Tokens to Target Language Spaces in Multilingual Neural Machine Translation</i> Zhi Qu, Yiran Wang, Jiannan Mao, Chenchen Ding, Hideki Tanaka, Masao Utiyama and Taro Watanabe	21687
<i>VisuoThink: Empowering LVLM Reasoning with Multimodal Tree Search</i> Yikun Wang, Siyin Wang, Qinyuan Cheng, Zhaoye Fei, Liang Ding, Qipeng Guo, Dacheng Tao and Xipeng Qiu	21707

<i>Automated CAD Modeling Sequence Generation from Text Descriptions via Transformer-Based Large Language Models</i>	
JianXing Liao, Junyan Xu, Yatao Sun, Maowen Tang, Sicheng He, Jingxian Liao, Shui Yu, Yun Li and Xiaohong Guan	21720
<i>LED-Merging: Mitigating Safety-Utility Conflicts in Model Merging with Location-Election-Disjoint</i>	
Qianli Ma, Dongrui Liu, Qian Chen, Linfeng Zhang and Jing Shao	21749
<i>Dolphin: Moving Towards Closed-loop Auto-research through Thinking, Practice, and Feedback</i>	
Jiakang Yuan, Xiangchao Yan, Bo Zhang, Tao Chen, Botian Shi, Wanli Ouyang, Yu Qiao, Lei Bai and Bowen Zhou	21768
<i>PerSphere: A Comprehensive Framework for Multi-Faceted Perspective Retrieval and Summarization</i>	
Yun Luo, Yingjie Li, Xiangkun Hu, Qinglin Qi, Fang Guo, Qipeng Guo, Zheng Zhang and Yue Zhang	21790
<i>Prompt-Guided Internal States for Hallucination Detection of Large Language Models</i>	
Fujie Zhang, Peiqi Yu, Biao Yi, Baolei Zhang, Tong Li and Zheli Liu	21806
<i>Typology-Guided Adaptation in Multilingual Models</i>	
Ndapa Nakashole	21819
<i>Don't Erase, Inform! Detecting and Contextualizing Harmful Language in Cultural Heritage Collections</i>	
Orfeas Menis Mastromichalakis, Jason Liartis, Kristina Rose, Antoine Isaac and Giorgos Stamou	21836
<i>ECLM: Entity Level Language Model for Spoken Language Understanding with Chain of Intent</i>	
Shangjian Yin, Peijie Huang, JiaTian Chen, Haojing Huang and Yuhong Xu	21851
<i>FaithfulRAG: Fact-Level Conflict Modeling for Context-Faithful Retrieval-Augmented Generation</i>	
Qinggong Zhang, Zhishang Xiang, Yilin Xiao, Le Wang, Junhui Li, Xinrun Wang and Jinsong Su	21863
<i>Knowledge Image Matters: Improving Knowledge-Based Visual Reasoning with Multi-Image Large Language Models</i>	
Guanghui Ye, Huan Zhao, Zhixue Zhao, Xupeng Zha, Yang Liu and Zhihua Jiang	21883
<i>Evaluating Personalized Tool-Augmented LLMs from the Perspectives of Personalization and Proactivity</i>	
Yupu Hao, Pengfei Cao, Zhuoran Jin, Huanxuan Liao, Yubo Chen, Kang Liu and Jun Zhao	21897
<i>GUICourse: From General Vision Language Model to Versatile GUI Agent</i>	
Wentong Chen, Junbo Cui, Jinyi Hu, Yujia Qin, Junjie Fang, Yue Zhao, Chongyi Wang, Jun Liu, Guirong Chen, Yupeng Huo, Yuan Yao, Yankai Lin, Zhiyuan Liu and Maosong Sun	21936
<i>Evaluating Visual and Cultural Interpretation: The K-Viscuit Benchmark with Human-VLM Collaboration</i>	
ChaeHun Park, Yujin Baek, Jaeseok Kim, Yu-Jung Heo, Du-Seong Chang and Jaegul Choo	21960
<i>Maximizing the Effectiveness of Larger BERT Models for Compression</i>	
Wen-Shu Fan, Su Lu, Shangyu Xing, Xin-Chun Li and De-Chuan Zhan	21975
<i>Can LLMs Reason About Program Semantics? A Comprehensive Evaluation of LLMs on Formal Specification Inference</i>	
Thanh Le-Cong, Bach Le and Toby Murray	21991

<i>HACo-Det: A Study Towards Fine-Grained Machine-Generated Text Detection under Human-AI Coauthoring</i>	
Zhixiong Su, Yichen Wang, Herun Wan, Zhaohan Zhang and Minnan Luo	22015
<i>IndicSynth: A Large-Scale Multilingual Synthetic Speech Dataset for Low-Resource Indian Languages</i>	
Divya V Sharma, Vijval Ekbote and Anubha Gupta	22037
<i>Reinforced IR: A Self-Boosting Framework For Domain-Adapted Information Retrieval</i>	
Chaofan Li, Jianlyu Chen, Yingxia Shao, Chaozhuo Li, Quanqing Xu, Defu Lian and Zheng Liu	22061
<i>CoIR: A Comprehensive Benchmark for Code Information Retrieval Models</i>	
Xiangyang Li, Kuicai Dong, Yi Quan Lee, Wei Xia, Hao Zhang, Xinyi Dai, Yasheng Wang and Ruiming Tang	22074
<i>Enhancing Multimodal Retrieval via Complementary Information Extraction and Alignment</i>	
Delong Zeng, Yuexiang Xie, Yaliang Li and Ying Shen	22092
<i>JoPA: Explaining Large Language Model’s Generation via Joint Prompt Attribution</i>	
Yurui Chang, Bochuan Cao, Yujia Wang, Jinghui Chen and Lu Lin	22106
<i>Proxy-Driven Robust Multimodal Sentiment Analysis with Incomplete Data</i>	
Aoqiang Zhu, Min Hu, Xiaohua Wang, Jiaoyun Yang, Yiming Tang and Ning An	22123
<i>Not All Terms Matter: Recall-Oriented Adaptive Learning for PLM-aided Query Expansion in Open-Domain Question Answering</i>	
Xinran Chen, Ben He, Xuanang Chen and Le Sun	22139
<i>A Mutual Information Perspective on Knowledge Graph Embedding</i>	
Jiang Li, Xiangdong Su, Zehua Duo, Tian Lan, Xiaotao Guo and Guanglai Gao	22152
<i>Aligned but Blind: Alignment Increases Implicit Bias by Reducing Awareness of Race</i>	
Lihao Sun, Chengzhi Mao, Valentin Hofmann and Xuechunzi Bai	22167
<i>IOPO: Empowering LLMs with Complex Instruction Following via Input-Output Preference Optimization</i>	
Xinghua Zhang, Haiyang Yu, Cheng Fu, Fei Huang and Yongbin Li	22185
<i>ProMALex: Progressive Modular Adapters for Multi-Jurisdictional Legal Language Modeling</i>	
Santosh T.y.s.s and Mohamed Hesham Elganayni	22201
<i>Flipping Knowledge Distillation: Leveraging Small Models’ Expertise to Enhance LLMs in Text Matching</i>	
Mingzhe Li, Jing Xiang, Qishen Zhang, Kaiyang Wan and Xiuying Chen	22218
<i>Disentangling Language and Culture for Evaluating Multilingual Large Language Models</i>	
Jiahao Ying, Wei Tang, Yiran Zhao, Yixin Cao, Yu Rong and Wenxuan Zhang	22230
<i>Detecting Sockpuppetry on Wikipedia Using Meta-Learning</i>	
Luc Raszewski and Christine de Kock	22252
<i>Diversity-oriented Data Augmentation with Large Language Models</i>	
Zaitian Wang, Jinghan Zhang, Xinhao Zhang, Kunpeng Liu, Pengfei Wang and Yuanchun Zhou	22265

<i>CoreEval: Automatically Building Contamination-Resilient Datasets with Real-World Knowledge toward Reliable LLM Evaluation</i>	
Jingqian Zhao, Bingbing Wang, Geng Tu, Yice Zhang, Qianlong Wang, Bin Liang, Jing Li and Ruifeng Xu	22284
<i>RiOT: Efficient Prompt Refinement with Residual Optimization Tree</i>	
Chenyi Zhou, Zhengyan Shi, Yuan Yao, Lei Liang, Huajun Chen and Qiang Zhang	22307
<i>Caution for the Environment: Multimodal LLM Agents are Susceptible to Environmental Distractions</i>	
Xinbei Ma, Yiting Wang, Yao Yao, Tongxin Yuan, Aston Zhang, Zhuosheng Zhang and Hai Zhao	22324
<i>Automatic Evaluation for Text-to-image Generation: Task-decomposed Framework, Distilled Training, and Meta-evaluation Benchmark</i>	
Rong-Cheng Tu, Zi-Ao Ma, Tian Lan, Yuehao Zhao, Heyan Huang and Xian-Ling Mao . .	22340
<i>Mitigating Lost-in-Retrieval Problems in Retrieval Augmented Multi-Hop Question Answering</i>	
Rongzhi Zhu, Xiangyu Liu, Zequn Sun, Yiwei Wang and Wei Hu	22362
<i>TableLoRA: Low-rank Adaptation on Table Structure Understanding for Large Language Models</i>	
Xinyi He, Yihao Liu, Mengyu Zhou, Yeye He, Haoyu Dong, Shi Han, Zejian Yuan and Dongmei Zhang	22376
<i>Condor: Enhance LLM Alignment with Knowledge-Driven Data Synthesis and Refinement</i>	
Maosongcao Maosongcao, Taolin Zhang, Mo Li, Chuyu Zhang, Yunxin Liu, Conghui He, Haodong Duan, Songyang Zhang and Kai Chen	22392
<i>CulFiT: A Fine-grained Cultural-aware LLM Training Paradigm via Multilingual Critique Data Synthesis</i>	
Ruixiang Feng, Shen Gao, Xiuying Chen, Lisi Chen and Shuo Shang	22413
<i>Decoding Knowledge Attribution in Mixture-of-Experts: A Framework of Basic-Refinement Collaboration and Efficiency Analysis</i>	
Junzhuo Li, Bo Wang, Xiuze Zhou, Peijie Jiang, Jia Liu and Xuming Hu	22431
<i>ChartLens: Fine-grained Visual Attribution in Charts</i>	
Manan Suri, Puneet Mathur, Nedim Lipka, Franck Dernoncourt, Ryan A. Rossi and Dinesh Manocha	22447
<i>LESA: Learnable LLM Layer Scaling-Up</i>	
Yifei Yang, Zouying Cao, Xinbei Ma, Yao Yao, Zhi Chen, Libo Qin and Hai Zhao	22463
<i>MMRC: A Large-Scale Benchmark for Understanding Multimodal Large Language Model in Real-World Conversation</i>	
Haochen Xue, Feilong Tang, Ming Hu, Yexin Liu, Qidong Huang, Yulong Li, Chengzhi Liu, Zhongxing Xu, Chong Zhang, Chun-Mei Feng, Yutong Xie, Imran Razzak, Zongyuan Ge, Jionglong Su, Junjun He and Yu Qiao	22477
<i>Towards the Law of Capacity Gap in Distilling Language Models</i>	
Chen Zhang, Qiuchi Li, Dawei Song, Zheyu Ye, Yan Gao and Yao Hu	22504
<i>WhiSPA: Semantically and Psychologically Aligned Whisper with Self-Supervised Contrastive and Student-Teacher Learning</i>	
Rajath Rao, Adithya V Ganesan, Oscar Kjell, Jonah Luby, Akshay Raghavan, Scott M. Feltman, Whitney Ringwald, Ryan L. Boyd, Benjamin J. Luft, Camilo J. Ruggero, Neville Ryant, Roman Kotov and H. Schwartz	22529

<i>Keys to Robust Edits: From Theoretical Insights to Practical Advances</i> Jianhao Yan, Futing Wang, Yun Luo, Yafu Li and Yue Zhang	22545
<i>Boosting LLM’s Molecular Structure Elucidation with Knowledge Enhanced Tree Search Reasoning</i> Xiang Zhuang, Bin Wu, Jiyu Cui, Kehua Feng, Xiaotong Li, Huabin Xing, Keyan Ding, Qiang Zhang and Huajun Chen	22561
<i>MEMERAG: A Multilingual End-to-End Meta-Evaluation Benchmark for Retrieval Augmented Generation</i> María Andrea Cruz Blandón, Jayasimha Talur, Bruno Charron, Dong Liu, Saab Mansour and Marcello Federico	22577
<i>The Role of Visual Modality in Multimodal Mathematical Reasoning: Challenges and Insights</i> Yufang Liu, Yao Du, Tao Ji, Jianing Wang, Yang Liu, Yuanbin Wu, Aimin Zhou, Mengdi Zhang and Xunliang Cai	22596
<i>The Essence of Contextual Understanding in Theory of Mind: A Study on Question Answering with Story Characters</i> Chulun Zhou, Qiuqing Wang, Mo Yu, Xiaoqian Yue, Rui Lu, Jiangnan Li, Yifan Zhou, Shunchi Zhang, Jie Zhou and Wai Lam	22612
<i>S²R: Teaching LLMs to Self-verify and Self-correct via Reinforcement Learning</i> Ruotian Ma, Peisong Wang, Cheng Liu, Xingyan Liu, Jiaqi Chen, Bang Zhang, Xin Zhou, Nan du and Jia Li	22632
<i>Advancing Collaborative Debates with Role Differentiation through Multi-Agent Reinforcement Learning</i> Haoran Li, Ziyi Su, Yun Xue, Zhiliang Tian, Yiping Song and Minlie Huang	22655
<i>Retrieval-Augmented Fine-Tuning With Preference Optimization For Visual Program Generation</i> Deokhyung Kang, Jeonghun Cho, Yejin Jeon, Sunbin Jang, Minsub Lee, Jawoon Cho and Gary Lee	22667
<i>STRICTA: Structured Reasoning in Critical Text Assessment for Peer Review and Beyond</i> Nils Dycke, Matej Zečević, Ilia Kuznetsov, Beatrix Suess, Kristian Kersting and Iryna Gurevych	22687
<i>XDAC: XAI-Driven Detection and Attribution of LLM-Generated News Comments in Korean</i> Wooyoung Go, Hyoungshick Kim, Alice Oh and Yongdae Kim	22728
<i>CENTAUR: Bridging the Impossible Trinity of Privacy, Efficiency, and Performance in Privacy-Preserving Transformer Inference</i> Jinglong Luo, Guanzhong Chen, Yehong Zhang, Shiyu Liu, Hui Wang, Yue Yu, Xun Zhou, Yuan Qi and Zenglin Xu	22751
<i>Silencing Empowerment, Allowing Bigotry: Auditing the Moderation of Hate Speech on Twitch</i> Prarabdh Shukla, Wei Yin Chong, Yash Patel, Brennan Schaffner, Danish Pruthi and Arjun Bha- goji	22771
<i>EdiText: Controllable Coarse-to-Fine Text Editing with Diffusion Language Models</i> Che Hyun Lee, Heeseung Kim, Jiheum Yeom and Sungroh Yoon	22798
<i>TUMLU: A Unified and Native Language Understanding Benchmark for Turkic Languages</i> Jafar Isbarov, Arofat Akhundjanova, Mammad Hajili, Kavsar Huseynova, Dmitry Gaynullin, Anar Rzaev, Osman Tursun, Aizirek Turdubaeva, Ilshat Saetov, Rinat Kharisov, Saule Belginova, Ariana Kenbayeva, Amina Alisheva, Abdullatif Köksal, Samir Rustamov and Duygu Ataman	22816

<i>Look Both Ways and No Sink: Converting LLMs into Text Encoders without Training</i> Ziyong Lin, Haoyi Wu, Shu Wang, Kewei Tu, Zilong Zheng and Zixia Jia	22839
<i>A Statistical and Multi-Perspective Revisiting of the Membership Inference Attack in Large Language Models</i> Bowen Chen, Namgi Han and Yusuke Miyao	22854
<i>Around the World in 24 Hours: Probing LLM Knowledge of Time and Place</i> Carolin Holtermann, Paul Röttger and Anne Lauscher	22875
<i>Mining the uncertainty patterns of humans and models in the annotation of moral foundations and human values</i> Neele Falk and Gabriella Lapesa	22898
<i>What do you call a dog that is incontrovertibly true? Dogma”: Testing LLM Generalization through Humor</i> Alessio Cocchieri, Luca Ragazzi, Paolo Italiani, Giuseppe Tagliavini and Gianluca Moro .	22922
<i>Towards Harmonized Uncertainty Estimation for Large Language Models</i> Rui Li, Jing Long, Muge Qi, Heming Xia, Lei Sha, Peiyi Wang and Zhifang Sui	22938
<i>VITAL: A New Dataset for Benchmarking Pluralistic Alignment in Healthcare</i> Anudeex Shetty, Amin Beheshti, Mark Dras and Usman Naseem	22954
<i>Are We in the AI-Generated Text World Already? Quantifying and Monitoring AIGT on Social Media</i> Zhen Sun, Zongmin Zhang, Xinyue Shen, Ziyi Zhang, Yule Liu, Michael Backes, Yang Zhang and Xinlei He	22975
<i>From English to Second Language Mastery: Enhancing LLMs with Cross-Lingual Continued Instruction Tuning</i> Linjuan Wu, Hao-Ran Wei, Baosong Yang and Weiming Lu.....	23006
<i>WET: Overcoming Paraphrasing Vulnerabilities in Embeddings-as-a-Service with Linear Transformation Watermarks</i> Anudeex Shetty, Qiongkai Xu and Jey Han Lau.....	23024
<i>HoPE: A Novel Positional Encoding Without Long-Term Decay for Enhanced Context Awareness and Extrapolation</i> Yuhan Chen, Ang Lv, Jian Luan, Bin Wang and Wei Liu	23044
<i>One QuantLLM for ALL: Fine-tuning Quantized LLMs Once for Efficient Deployments</i> Ke Yi, Yuhui Xu, Heng Chang, Yuan Meng, Tong Zhang and Jia Li	23057
<i>Beyond Logits: Aligning Feature Dynamics for Effective Knowledge Distillation</i> Guoqiang Gong, Jiaying Wang, Jin Xu, Deping Xiang, Zicheng Zhang, Leqi Shen, Yifeng Zhang, JunhuaShu JunhuaShu, ZhaolongXing ZhaolongXing, Zhen Chen, Pengzhang Liu and Ke Zhang	23067
<i>Native Sparse Attention: Hardware-Aligned and Natively Trainable Sparse Attention</i> Jingyang Yuan, Huazuo Gao, Damai Dai, Junyu Luo, Liang Zhao, Zhengyan Zhang, Zhenda Xie, Yuxing Wei, Lean Wang, Zhiping Xiao, Yuqing Wang, Chong Ruan, Ming Zhang, Wenfeng Liang and Wangding Zeng.....	23078
<i>DRAE: Dynamic Retrieval-Augmented Expert Networks for Lifelong Learning and Task Adaptation in Robotics</i> Yayu Long, Kewei Chen, Long Jin and Mingsheng Shang.....	23098

<i>MT-RAIG: Novel Benchmark and Evaluation Framework for Retrieval-Augmented Insight Generation over Multiple Tables</i>	
Kwangwook Seo, Donguk Kwon and Dongha Lee	23142
<i>Enhancing Chain-of-Thought Reasoning with Critical Representation Fine-tuning</i>	
Chenxi Huang, Shaotian Yan, Liang Xie, Binbin Lin, Sinan Fan, Yue Xin, Deng Cai, Chen Shen and Jieping Ye	23173
<i>Does the Emotional Understanding of LLMs Vary Under High-Stress Environments and Across Different Demographic Attributes?</i>	
Jaewook Lee, Yeajin Jang, Oh-Woong Kwon and Harksoo Kim	23196
<i>S2WTM: Spherical Sliced-Wasserstein Autoencoder for Topic Modeling</i>	
Suman Adhya and Debarshi Kumar Sanyal	23211
<i>Learning to Look at the Other Side: A Semantic Probing Study of Word Embeddings in LLMs with Enabled Bidirectional Attention</i>	
Zhaoxin Feng, Jianfei Ma, Emmanuele Chersoni, Xiaojing Zhao and Xiaoyi Bao	23226
<i>Tracing and Dissecting How LLMs Recall Factual Knowledge for Real World Questions</i>	
Yiqun Wang, Chaoqun Wan, Sile Hu, Yonggang Zhang, Xiang Tian, Yaowu Chen, Xu Shen and Jieping Ye	23246
<i>Employing Discourse Coherence Enhancement to Improve Cross-Document Event and Entity Coreference Resolution</i>	
Xinyu Chen, Peifeng Li and Qiaoming Zhu	23272
<i>Data Whisperer: Efficient Data Selection for Task-Specific LLM Fine-Tuning via Few-Shot In-Context Learning</i>	
Shaobo Wang, Xiangqi Jin, Ziming Wang, Jize Wang, Jiajun Zhang, Kaixin Li, Zichen Wen, Zhong Li, Conghui He, Xuming Hu and Linfeng Zhang	23287
<i>Synthesizing Post-Training Data for LLMs through Multi-Agent Simulation</i>	
Shuo Tang, Xianghe Pang, Zexi Liu, Bohan Tang, Rui Ye, Tian Jin, Xiaowen Dong, Yanfeng Wang and Siheng Chen	23306
<i>SoftCoT: Soft Chain-of-Thought for Efficient Reasoning with LLMs</i>	
Yige Xu, Xu Guo, Zhiwei Zeng and Chunyan Miao	23336
<i>FCMR: Robust Evaluation of Financial Cross-Modal Multi-Hop Reasoning</i>	
Seunghee Kim, Changhyeon Kim and Taeuk Kim	23352
<i>Beyond Prompt Engineering: Robust Behavior Control in LLMs via Steering Target Atoms</i>	
Mengru Wang, Ziwen Xu, Shengyu Mao, Shumin Deng, Zhaopeng Tu, Huajun Chen and Ningyu Zhang	23381
<i>MobiLoRA: Accelerating LoRA-based LLM Inference on Mobile Devices via Context-aware KV Cache Optimization</i>	
Borui Li, Yitao Wang, Haoran Ma, Ligeng Chen, Jun Xiao and Shuai Wang	23400
<i>Language Models Resist Alignment: Evidence From Data Compression</i>	
Jiaming Ji, Kaile Wang, Tianyi Alex Qiu, Boyuan Chen, Jiayi Zhou, Changye Li, Hantao Lou, Josef Dai, Yunhuai Liu and Yaodong Yang	23411
<i>Beyond the Answer: Advancing Multi-Hop QA with Fine-Grained Graph Reasoning and Evaluation</i>	
Qichuan Liu, Chentao Zhang, Chenfeng Zheng, Guosheng Hu, Xiaodong Li and Zhihong Zhang	23433

<i>Mamba Knockout for Unraveling Factual Information Flow</i>	
Nir Endy, Idan Daniel Grosbard, Yuval Ran-Milo, Yonatan Slutzky, Itay Tshuva and Raja Giryes	
	23457
<i>Small Changes, Big Impact: How Manipulating a Few Neurons Can Drastically Alter LLM Aggression</i>	
Jaewook Lee, Junseo Jang, Oh-Woog Kwon and Harksoo Kim	23478
<i>Marco-o1 v2: Towards Widening The Distillation Bottleneck for Reasoning Models</i>	
Huifeng Yin, Yu Zhao, Minghao Wu, Xuanfan Ni, Bo Zeng, Huaiyu.wh Huaiyu.wh, Tianqi Shi, Liangying Shao, Chenyang Lyu, Longyue Wang, Weihua Luo and Kaifu Zhang	23506
<i>Curiosity-Driven Reinforcement Learning from Human Feedback</i>	
Haoran Sun, Yekun Chai, Shuohuan Wang, Yu Sun, Hua Wu and Haifeng Wang	23517
<i>T2A-Feedback: Improving Basic Capabilities of Text-to-Audio Generation via Fine-grained AI Feedback</i>	
Zehan Wang, Ke Lei, Chen Zhu, Jiawei Huang, Sashuai Zhou, Luping Liu, Xize Cheng, Sheng-peng Ji, Zhenhui Ye, Tao Jin and Zhou Zhao	23535
<i>CoE: A Clue of Emotion Framework for Emotion Recognition in Conversations</i>	
Zhiyu Shen, Yunhe Pang, Yanghui Rao and Jianxing Yu	23548
<i>MPO: Multilingual Safety Alignment via Reward Gap Optimization</i>	
Weixiang Zhao, Yulin Hu, Yang Deng, Tongtong Wu, Wenxuan Zhang, Jiahe Guo, An Zhang, Yanyan Zhao, Bing Qin, Tat-Seng Chua and Ting Liu	23564
<i>QualiSpeech: A Speech Quality Assessment Dataset with Natural Language Reasoning and Descriptions</i>	
Siyin Wang, Wenyi Yu, Xianzhao Chen, Xiaohai Tian, Jun Zhang, Lu Lu, Yu Tsao, Junichi Yamagishi, Yuxuan Wang and Chao Zhang	23588
<i>On the Relation Between Fine-Tuning, Topological Properties, and Task Performance in Sense-Enhanced Embeddings</i>	
Deniz Ekin Yavas, Timothée Bernard, Benoit Crabbé and Laura Kallmeyer	23610
<i>Finding Needles in Images: Can Multi-modal LLMs Locate Fine Details?</i>	
Parth Thakkar, Ankush Agarwal, Prasad Kasu, Pulkit Bansal and Chaitanya Devaguptapu	23626
<i>Don't Half-listen: Capturing Key-part Information in Continual Instruction Tuning</i>	
Yongquan He, Wenyuan Zhang, Xuancheng Huang, Peng Zhang, Lingxun Meng, Xiang Zhou, Ke Zeng and Xunliang Cai	23649
<i>Generating Plausible Distractors for Multiple-Choice Questions via Student Choice Prediction</i>	
Yooseop Lee, Suin Kim and Yohan Jo	23669
<i>Exploring Explanations Improves the Robustness of In-Context Learning</i>	
Ukyo Honda and Tatsushi Oka	23693
<i>Prediction Hubs are Context-Informed Frequent Tokens in LLMs</i>	
Beatrix Miranda Ginn Nielsen, Iuri Macocco and Marco Baroni	23715
<i>Capability Salience Vector: Fine-grained Alignment of Loss and Capabilities for Downstream Task Scaling Law</i>	
Qiming Ge, Shuhao Xing, Songyang Gao, Yunhua Zhou, Yicheng Zou, Songyang Zhang, Zhi Chen, Hang Yan, Qi Zhang, Qipeng Guo and Kai Chen	23746

<i>CRUXEVAL-X: A Benchmark for Multilingual Code Reasoning, Understanding and Execution</i>	
Ruiyang Xu, Jialun Cao, Yaojie Lu, Ming Wen, Hongyu Lin, Xianpei Han, Ben He, Shing-Chi Cheung and Le Sun	23762
<i>Graph of Records: Boosting Retrieval Augmented Generation for Long-context Summarization with Graphs</i>	
Haozhen Zhang, Tao Feng and Jiaxuan You	23780
<i>Rubrik’s Cube: Testing a New Rubric for Evaluating Explanations on the CUBE dataset</i>	
Diana Galvan-Sosa, Gabrielle Gaudeau, Pride Kavumba, Yunmeng Li, Hongyi gu, Zheng Yuan, Keisuke Sakaguchi and Paula Buttery	23800
<i>A Dual-Mind Framework for Strategic and Expressive Negotiation Agent</i>	
Yutong Liu, Lida Shi, Rui Song and Hao Xu	23840
<i>Ref-Long: Benchmarking the Long-context Referencing Capability of Long-context Language Models</i>	
Junjie Wu, Gefei Gu, Yanan Zheng, Dit-Yan Yeung and Arman Cohan	23861
<i>Revisiting Scaling Laws for Language Models: The Role of Data Quality and Training Strategies</i>	
Zhengyu Chen, Siqi Wang, Teng Xiao, Yudong Wang, Shiqi Chen, Xunliang Cai, Junxian He and Jingang Wang	23881
<i>Limited Generalizability in Argument Mining: State-Of-The-Art Models Learn Datasets, Not Arguments</i>	
Marc Feger, Katarina Boland and Stefan Dietze	23900
<i>Enhancing Machine Translation with Self-Supervised Preference Data</i>	
Haoxiang Sun, Ruize Gao, Pei Zhang, Baosong Yang and Rui Wang	23916
<i>Unveil: Unified Visual-Textual Integration and Distillation for Multi-modal Document Retrieval</i>	
Hao Sun, Yingyan Hou, Jiayan Guo, Bo Wang, Chunyu Yang, Jinsong Ni and Yan Zhang ..	23935
<i>Don’t Get Lost in the Trees: Streamlining LLM Reasoning by Overcoming Tree Search Exploration Pitfalls</i>	
Ante Wang, Linfeng Song, Ye Tian, Dian Yu, Haitao Mi, Xiangyu Duan, Zhaopeng Tu, Jinsong Su and Dong Yu	23946
<i>MEXMA: Token-level objectives improve sentence representations</i>	
João Maria Janeiro, Benjamin Piwowarski, Patrick Gallinari and Loic Barrault	23960
<i>Uncertainty-Aware Iterative Preference Optimization for Enhanced LLM Reasoning</i>	
Lei Li, Hehuan Liu, Yaxin Zhou, ZhaoYang Gui, Xudong Weng, Yi Yuan, Zheng Wei and Zang Li	23996
<i>AgentDropout: Dynamic Agent Elimination for Token-Efficient and High-Performance LLM-Based Multi-Agent Collaboration</i>	
Zhexuan Wang, Yutong Wang, Xuebo Liu, Liang Ding, Miao Zhang, Jie Liu and Min Zhang	24013
<i>Towards Dynamic Theory of Mind: Evaluating LLM Adaptation to Temporal Evolution of Human States</i>	
Yang Xiao, Jiashuo Wang, Qiancheng Xu, Changhe Song, Chunpu Xu, Yi Cheng, Wenjie Li and Pengfei Liu	24036
<i>Marco-Bench-MIF: On Multilingual Instruction-Following Capability of Large Language</i>	
Bo Zeng, Chenyang Lyu, Sinuo Liu, Mingyan Zeng, Minghao Wu, Xuanfan Ni, Tianqi Shi, Yu Zhao, Yefeng Liu, Chenyu Zhu, Ruizhe Li, Jiahui Geng, Qing Li, Yu Tong, Longyue Wang, Weihua Luo and Kaifu Zhang	24058

<i>Representation Bending for Large Language Model Safety</i>	
Ashkan Yousefpour, Taeheon Kim, Ryan Sungmo Kwon, Seungbeen Lee, Wonje Jeung, Seungju Han, Alvin Wan, Harrison Ngan, Youngjae Yu and Jonghyun Choi	24073
<i>Analyzing LLMs’ Knowledge Boundary Cognition Across Languages Through the Lens of Internal Representations</i>	
Chenghao Xiao, Hou Pong Chan, Hao Zhang, Mahani Aljunied, Lidong Bing, Noura Al Mou-bayed and Yu Rong	24099
<i>Enhancing Retrieval-Augmented Generation via Evidence Tree Search</i>	
Hao Sun, Hengyi Cai, Yuchen Li, Xuanbo Fan, Xiaochi Wei, Shuaiqiang Wang, Yan Zhang and Dawei Yin	24116
<i>HalluLens: LLM Hallucination Benchmark</i>	
Yejin Bang, Ziwei Ji, Alan Schelten, Anthony Hartshorn, Tara Fowler, Cheng Zhang, Nicola Cancedda and Pascale Fung	24128
<i>DEEPER Insight into Your User: Directed Persona Refinement for Dynamic Persona Modeling</i>	
Aili Chen, Chengyu Du, Jiangjie Chen, Jinghan Xu, Yikai Zhang, Siyu Yuan, Zulong Chen, Liangyue Li and Yanghua Xiao	24157
<i>Asclepius: A Spectrum Evaluation Benchmark for Medical Multi-Modal Large Language Models</i>	
Jie Liu, Wenxuan Wang, SU Yihang, Jingyuan Huang, Yudi Zhang, Cheng-Yi Li, Wenting Chen, Xiaohan Xing, Kao-Jung Chang, Linlin Shen and Michael R. Lyu	24181
<i>InstructPart: Task-Oriented Part Segmentation with Instruction Reasoning</i>	
Zifu Wan, Yaqi Xie, Ce Zhang, Zhiqiu Lin, Zihan Wang, Simon Stepputtis, Deva Ramanan and Katia P. Sycara	24202
<i>GRaMPa: Subword Regularisation by Skewing Uniform Segmentation Distributions with an Efficient Path-counting Markov Model</i>	
Thomas Bauwens, David Kaczér and Miryam De Lhoneux	24228
<i>Evaluating the Evaluation of Diversity in Commonsense Generation</i>	
Tianhui Zhang, Bei Peng and Danushka Bollegala	24258
<i>Generate First, Then Sample: Enhancing Fake News Detection with LLM-Augmented Reinforced Sampling</i>	
Zhao Tong, Yimeng Gu, Huidong Liu, Qiang Liu, Shu Wu, Haichao Shi and Xiao-Yu Zhang	24276
<i>ChemActor: Enhancing Automated Extraction of Chemical Synthesis Actions with LLM-Generated Data</i>	
Yu Zhang, Ruijie Yu, Jidong Tian, Feng Zhu, Jiapeng Liu, Xiaokang Yang, Yaohui Jin and Yanyan Xu	24291
<i>Towards Fully Exploiting LLM Internal States to Enhance Knowledge Boundary Perception</i>	
Shiyu Ni, Keping Bi, Jiafeng Guo, Lulu Yu, Baolong Bi and Xueqi Cheng	24315
<i>ALGEN: Few-shot Inversion Attacks on Textual Embeddings via Cross-Model Alignment and Generation</i>	
Yiyi Chen, Qiongkai Xu and Johannes Bjerva	24330
<i>Decoding on Graphs: Faithful and Sound Reasoning on Knowledge Graphs through Generation of Well-Formed Chains</i>	
Kun LI, Tianhua Zhang, Xixin Wu, Hongyin Luo, James R. Glass and Helen M. Meng . . .	24349

<i>STaR-SQL: Self-Taught Reasoner for Text-to-SQL</i>	
Mingqian He, Yongliang Shen, Wenqi Zhang, Qiuying Peng, Jun Wang and Weiming Lu .	24365
<i>Fairness Beyond Performance: Revealing Reliability Disparities Across Groups in Legal NLP</i>	
Santosh T.y.s.s and Irtiza Chowdhury	24376
<i>Beyond Similarity: A Gradient-based Graph Method for Instruction Tuning Data Selection</i>	
Yang Zhao, Li Du, Xiao Ding, Yangou Ouyang, Hepeng Wang, Kai Xiong, Jinglong Gao, Zhouhao Sun, Dongliang Xu, Qing Yang, Dongchen Li, Bing Qin and Ting Liu	24391
<i>FastMCTS: A Simple Sampling Strategy for Data Synthesis</i>	
Peiji Li, Kai Lv, Yunfan Shao, Yichuan Ma, Linyang Li, Xiaoqing Zheng, Xipeng Qiu and Qipeng Guo	24405
<i>Dialogue-RAG: Enhancing Retrieval for LLMs via Node-Linking Utterance Rewriting</i>	
Qiwei Li, Teng Xiao, Zuchao Li, Ping Wang, Mengjia Shen and Hai Zhao	24423
<i>Using Information Theory to Characterize Prosodic Typology: The Case of Tone, Pitch-Accent and Stress-Accent</i>	
Ethan Wilcox, Cui Ding, Giovanni Acampa, Tiago Pimentel, Alex Warstadt and Tamar I Regev	24439
<i>Evaluating LLMs for Portuguese Sentence Simplification with Linguistic Insights</i>	
Arthur Mariano Rocha De Azevedo Scalercio, Elvis A. De Souza, Maria José Bocorny Finatto and Aline Paes	24452
<i>LaTIM: Measuring Latent Token-to-Token Interactions in Mamba Models</i>	
Hugo Pitorro and Marcos Vinicius Treviso	24478
<i>Improving Low-Resource Morphological Inflection via Self-Supervised Objectives</i>	
Adam Wiemerslage and Katharina Von Der Wense	24494
<i>Don't Reinvent the Wheel: Efficient Instruction-Following Text Embedding based on Guided Space Transformation</i>	
Yingchaojie Feng, Yiqun Sun, Yandong Sun, Minfeng Zhu, Qiang Huang, Anthony Kum Hoe Tung and Wei Chen	24511
<i>BOOKCOREF: Coreference Resolution at Book Scale</i>	
Giuliano Martinelli, Tommaso Bonomo, Pere-Lluís Huguet Cabot and Roberto Navigli ...	24526
<i>OMGM: Orchestrate Multiple Granularities and Modalities for Efficient Multimodal Retrieval</i>	
Wei Yang, Jingjing Fu, Rui Wang, Jinyu Wang, Lei Song and Jiang Bian	24545
<i>Alleviating Hallucinations from Knowledge Misalignment in Large Language Models via Selective Ab-stention Learning</i>	
Lei Huang, Xiaocheng Feng, Weitao Ma, Yuchun Fan, Xiachong Feng, Yuxuan Gu, Yangfan Ye, Liang Zhao, Weihong Zhong, Baoxin Wang, Dayong Wu, Guoping Hu, Lingpeng Kong, Tong Xiao, Ting Liu and Bing Qin	24564
<i>Retrospective Learning from Interactions</i>	
Zizhao Chen, Mustafa Omer Gul, Yiwei Chen, Gloria Geng, Anne Wu and Yoav Artzi ...	24580
<i>Personalized Generation In Large Model Era: A Survey</i>	
Yiyang Xu, Jinghao Zhang, Alireza Salemi, Xinting Hu, Wenjie Wang, Fuli Feng, Hamed Zamani, Xiangnan He and Tat-Seng Chua	24607

<i>Graph Counselor: Adaptive Graph Exploration via Multi-Agent Synergy to Enhance LLM Reasoning</i> Junqi Gao, Xiang Zou, Ying Ai, Dong Li, Yichen Niu, Biqing Qi and Jianxing Liu	24650
<i>SOTOPIA-: Dynamic Strategy Injection Learning and Social Instruction Following Evaluation for Social Agents</i> Wenyuan Zhang, Tianyun Liu, Mengxiao Song, Xiaodong Li and Tingwen Liu	24669
<i>Can Language Models Replace Programmers for Coding? REPOCOD Says ‘Not Yet’</i> Shanchao Liang, Nan Jiang, Yiran Hu and Lin Tan	24698
<i>Leveraging In-Context Learning for Political Bias Testing of LLMs</i> Patrick Haller, Jannis Vamvas, Rico Sennrich and Lena Ann Jäger	24718
<i>ACORD: An Expert-Annotated Retrieval Dataset for Legal Contract Drafting</i> Steven H Wang, Maksim Zubkov, Kexin Fan, Sarah Harrell, Yuyang Sun, Wei Chen, Andreas Plesner and Roger Wattenhofer	24739
<i>LLMs know their vulnerabilities: Uncover Safety Gaps through Natural Distribution Shifts</i> Qibing Ren, Hao Li, Dongrui Liu, Zhanxu Xie, Xiaoya Lu, Yu Qiao, Lei Sha, Junchi Yan, Lizhuang Ma and Jing Shao	24763
<i>WAFFLE: Fine-tuning Multi-Modal Model for Automated Front-End Development</i> Shanchao Liang, Nan Jiang, Shangshu Qian and Lin Tan	24786
<i>Math Neurosurgery: Isolating Language Models’ Math Reasoning Abilities Using Only Forward Passes</i> Bryan R Christ, Zachary Gottesman, Jonathan Kropko and Thomas Hartvigsen	24803
<i>Multiple LLM Agents Debate for Equitable Cultural Alignment</i> Dayeon Ki, Rachel Rudinger, Tianyi Zhou and Marine Carpuat	24841
<i>RefreshKV: Updating Small KV Cache During Long-form Generation</i> Fangyuan Xu, Tanya Goyal and Eunsol Choi	24878
<i>SEA: Low-Resource Safety Alignment for Multimodal Large Language Models via Synthetic Embeddings</i> Weikai Lu, Hao Peng, Huiping Zhuang, Cen Chen and Ziqian Zeng	24894
<i>Chain-of-Reasoning: Towards Unified Mathematical Reasoning in Large Language Models via a Multi-Paradigm Perspective</i> Yiyao Yu, Yuxiang Zhang, Dongdong Zhang, Xiao Liang, Hengyuan Zhang, Xingxing Zhang, Mahmoud Khademi, Hany Hassan Awadalla, Junjie Wang, Yujiu Yang and Furu Wei	24914
<i>Language Models Grow Less Humanlike beyond Phase Transition</i> Tatsuya Aoyama and Ethan Wilcox	24938
<i>PCoT: Persuasion-Augmented Chain of Thought for Detecting Fake News and Social Media Disinformation</i> Arkadiusz Modzelewski, Witold Sosnowski, Tiziano Labruna, Adam Wierzbicki and Giovanni Da San Martino	24959
<i>Coordinating Chaos: A Structured Review of Linguistic Coordination Methodologies</i> Benjamin Roger Litterer, David Jurgens and Dallas Card	24984
<i>iNews: A Multimodal Dataset for Modeling Personalized Affective Responses to News</i> Tiancheng Hu and Nigel Collier	25000
<i>Mind the Gesture: Evaluating AI Sensitivity to Culturally Offensive Non-Verbal Gestures</i> Akhila Yerukola, Saadia Gabriel, Nanyun Peng and Maarten Sap	25041

<i>500xCompressor: Generalized Prompt Compression for Large Language Models</i> Zongqian Li, Yixuan Su and Nigel Collier	25081
<i>Estimating Privacy Leakage of Augmented Contextual Knowledge in Language Models</i> James Flemings, Bo Jiang, Wanrong Zhang, Zafar Takhirov and Murali Annavaram	25092
<i>Document-Level Event-Argument Data Augmentation for Challenging Role Types</i> Joseph Gatto, Omar Sharif, Parker Seegmiller and Sarah Masud Preum	25109
<i>Mapping the Podcast Ecosystem with the Structured Podcast Research Corpus</i> Benjamin Roger Litterer, David Jurgens and Dallas Card	25132
<i>Unravelling the Logic: Investigating the Generalisation of Transformers in Numerical Satisfiability Problems</i> Tharindu Madusanka, Marco Valentino, Iqra Zahid, Ian Pratt-Hartmann and Riza Batista-Navarro	25155
<i>The Nature of NLP: Analyzing Contributions in NLP Papers</i> Aniket Pramanick, Yufang Hou, Saif M. Mohammad and Iryna Gurevych	25169
<i>GeLLM³O: Generalizing Large Language Models for Multi-property Molecule Optimization</i> Vishal Dey, Xiao Hu and Xia Ning	25192
<i>Follow-up Question Generation For Enhanced Patient-Provider Conversations</i> Joseph Gatto, Parker Seegmiller, Timothy E. Burdick, Inas S. Khayal, Sarah DeLozier and Sarah Masud Preum	25222
<i>Unveiling Privacy Risks in LLM Agent Memory</i> Bo Wang, Weiyi He, Shenglai Zeng, Zhen Xiang, Yue Xing, Jiliang Tang and Pengfei He ..	25241
<i>Watching the Watchers: Exposing Gender Disparities in Machine Translation Quality Estimation</i> Emmanouil Zaranis, Giuseppe Attanasio, Sweta Agrawal and Andre Martins	25261
<i>Language Constrained Multimodal Hyper Adapter For Many-to-Many Multimodal Summarization</i> Nayu Liu, Fanglong Yao, Haoran Luo, Yong Yang, Chen Tang and Bo Lv	25285
<i>PRMBench: A Fine-grained and Challenging Benchmark for Process-Level Reward Models</i> Mingyang Song, Zhaochen Su, Xiaoye Qu, Jiawei Zhou and Yu Cheng	25299
<i>Efficient Ensemble for Fine-tuning Language Models on Multiple Datasets</i> Dongyue Li, Ziniu Zhang, Lu Wang and Hongyang R. Zhang	25347
<i>Library-Like Behavior In Language Models is Enhanced by Self-Referencing Causal Cycles</i> Munachiso S Nwadike, Zangir Iklassov, Toluwani Aremu, Tatsuya Hiraoka, Benjamin Heinzerling, Velibor Bojkovic, Hilal AlQuabeh, Martin Takáč and Kentaro Inui	25365
<i>Shaping the Safety Boundaries: Understanding and Defending Against Jailbreaks in Large Language Models</i> Lang Gao, Jiahui Geng, Xiangliang Zhang, Preslav Nakov and Xiuying Chen	25378
<i>ASPERA: A Simulated Environment to Evaluate Planning for Complex Action Execution</i> Alexandru Coca, Mark Gaynor, Zhenxing Zhang, Jianpeng Cheng, Bo-Hsiang Tseng, Peter Boothroyd, Hector Martinez Alonso, Diarmuid O Seaghdha and Anders Johannsen	25399
<i>ReflectDiffu: Reflect between Emotion-intent Contagion and Mimicry for Empathetic Response Generation via a RL-Diffusion Framework</i> Jiahao Yuan, Zixiang Di, Zhiqing Cui, Guisong Yang and Usman Naseem	25435

<i>SARA: Saliency-Aware Reinforced Adaptive Decoding for Large Language Models in Abstractive Summarization</i>	
Nayu Liu, Junnan Zhu, Yiming Ma, Zhicong Lu, Wenlei Xu, Yong Yang, Jiang Zhong and Kaiwen Wei	25450
<i>Embedding-Converter: A Unified Framework for Cross-Model Embedding Transformation</i>	
Jinsung Yoon and Sercan O Arik	25464
<i>Improving Automatic Evaluation of Large Language Models (LLMs) in Biomedical Relation Extraction via LLMs-as-the-Judge</i>	
Md Tahmid Rahman Laskar, Israt Jahan, Elham Dolatabadi, Chun Peng, Enamul Hoque and Jimmy Huang	25483
<i>Answering Complex Geographic Questions by Adaptive Reasoning with Visual Context and External Commonsense Knowledge</i>	
Fan Li, Jianxing Yu, Jielong Tang, Wenqing Chen, Hanjiang Lai, Yanghui Rao and Jian Yin	25498
<i>Safety Alignment via Constrained Knowledge Unlearning</i>	
Zesheng Shi, Yucheng Zhou, Jing Li, Yuxin Jin, YU LI, Daojing He, Fangming Liu, Saleh Alharbi, Jun Yu and Min Zhang	25515
<i>Response Wide Shut? Surprising Observations in Basic Vision Language Model Capabilities</i>	
Shivam Chandhok, Wan-Cyuan Fan, Vered Shwartz, Vineeth N. Balasubramanian and Leonid Sigal	25530
<i>EffiVLM-BENCH: A Comprehensive Benchmark for Evaluating Training-Free Acceleration in Large Vision-Language Models</i>	
Zekun Wang, Minghua Ma, Zexin Wang, Rongchuan Mu, Liping Shan, Ming Liu and Bing Qin	25546
<i>Pre-Training Curriculum for Multi-Token Prediction in Language Models</i>	
Ansar Aynedinov and Alan Akbik	25573
<i>Can We Further Elicit Reasoning in LLMs? Critic-Guided Planning with Retrieval-Augmentation for Solving Challenging Tasks</i>	
Xingxuan Li, Weiwen Xu, Ruochen Zhao, Fangkai Jiao, Shafiq Joty and Lidong Bing	25589
<i>On Many-Shot In-Context Learning for Long-Context Evaluation</i>	
Kaijian Zou, Muhammad Khalifa and Lu Wang	25605
<i>HelpSteer3: Human-Annotated Feedback and Edit Data to Empower Inference-Time Scaling in Open-Ended General-Domain Tasks</i>	
Zhilin Wang, Jiaqi Zeng, Olivier Delalleau, Daniel Egert, Ellie Evans, Hoo-Chang Shin, Felipe Soares, Yi Dong and Oleksii Kuchaiev	25640
<i>CulturalBench: A Robust, Diverse and Challenging Benchmark for Measuring LMs' Cultural Knowledge Through Human-AI Red-Teaming</i>	
Yu Ying Chiu, Liwei Jiang, Bill Yuchen Lin, Chan Young Park, Shuyue Stella Li, Sahithya Ravi, Mehar Bhatia, Maria Antoniak, Yulia Tsvetkov, Vered Shwartz and Yejin Choi	25663
<i>Balancing the Budget: Understanding Trade-offs Between Supervised and Preference-Based Finetuning</i>	
Mohit Raghavendra, Junmo Kang and Alan Ritter	25702
<i>All That Glitters is Not Novel: Plagiarism in AI Generated Research</i>	
Tarun Gupta and Danish Pruthi	25721

<i>Writing Like the Best: Exemplar-Based Expository Text Generation</i> Yuxiang Liu and Kevin Chen-Chuan Chang	25739
<i>Temporal Relation Extraction in Clinical Texts: A Span-based Graph Transformer Approach</i> Rochana Chaturvedi, Peyman Baghersahi, Sourav Medya and Barbara Di Eugenio	25765
<i>Finding A Voice: Exploring the Potential of African American Dialect and Voice Generation for Chatbots</i> Sarah E. Finch, Ellie S. Paek, Ikseon Choi and Jinho D. Choi	25789
<i>Delta-KNN: Improving Demonstration Selection in In-Context Learning for Alzheimer’s Disease Detection</i> Chuyuan Li, Raymond Li, Thalia S. Field and Giuseppe Carenini	25807
<i>Help Me Write a Story: Evaluating LLMs’ Ability to Generate Writing Feedback</i> Hannah Rashkin, Elizabeth Clark, Fantine Huot and Mirella Lapata	25827
<i>Language Fusion for Parameter-Efficient Cross-lingual Transfer</i> Philipp Borchert, Ivan Vulić, Marie-Francine Moens and Jochen De Weerd	25848
<i>Culture is Not Trivia: Sociocultural Theory for Cultural NLP</i> Naitian Zhou, David Bamman and Isaac L. Bleaman	25869
<i>AAD-LLM: Neural Attention-Driven Auditory Scene Understanding</i> Xilin Jiang, Sukru Samet Dindar, Vishal Choudhari, Stephan Bickel, Ashesh Mehta, Guy M McKhann, Daniel Friedman, Adeen Flinker and Nima Mesgarani	25887
<i>Do Language Models Have Semantics? On the Five Standard Positions</i> Anders Søgaard	25910
<i>Dehumanizing Machines: Mitigating Anthropomorphic Behaviors in Text Generation Systems</i> Myra Cheng, Su Lin Blodgett, Alicia DeVrio, Lisa Egede and Alexandra Olteanu	25923
<i>Evaluating Multimodal Language Models as Visual Assistants for Visually Impaired Users</i> Antonia Karamolegkou, Malvina Nikandrou, Georgios Pantazopoulos, Danae Sanchez Villegas, Phillip Rust, Ruchira Dhar, Daniel Hershcovich and Anders Søgaard	25949
<i>HumT DumT: Measuring and controlling human-like language in LLMs</i> Myra Cheng, Sunny Yu and Dan Jurafsky	25983
<i>ChatBench: From Static Benchmarks to Human-AI Evaluation</i> Serina Chang, Ashton Anderson and Jake M. Hofman	26009
<i>Teaching an Old LLM Secure Coding: Localized Preference Optimization on Distilled Preferences</i> Mohammad Saqib Hasan, Saikat Chakraborty, Santu Karmaker and Niranjana Balasubramanian	26039
<i>Anything Goes? A Crosslinguistic Study of (Im)possible Language Learning in LMs</i> Xiulin Yang, Tatsuya Aoyama, Yuekun Yao and Ethan Wilcox	26058
<i>Ranking Unraveled: Recipes for LLM Rankings in Head-to-Head AI Combat</i> Roland Daynauth, Christopher Clarke, Krisztian Flautner, Lingjia Tang and Jason Mars	26078
<i>LLM Agents Making Agent Tools</i> Georg Wölflein, Dyke Ferber, Daniel Truhn, Ognjen Arandjelovic and Jakob Nikolas Kather	26092

<i>CraText Benchmark: Advancing Instruction Following in Complex Multimodal Open-Ended World</i> Zoya Volovikova, Gregory Gorbov, Petr Kuderov, Aleksandr Panov and Alexey Skrynnik	26131
<i>QG-SMS: Enhancing Test Item Analysis via Student Modeling and Simulation</i> Bang Nguyen, Tingting Du, Mengxia Yu, Lawrence Angrave and Meng Jiang	26152
<i>Causal Graph based Event Reasoning using Semantic Relation Experts</i> Mahnaz Koupae, Xueying Bai, Mudan Chen, Greg Durrett, Nathanael Chambers and Niranjana Balasubramanian	26169
<i>LogicPro: Improving Complex Logical Reasoning via Program-Guided Learning</i> Jin Jiang, Yuchen Yan, Yang Liu, Jianing Wang, Shuai Peng, Xunliang Cai, Yixin Cao, Mengdi Zhang and Liangcai Gao	26200
<i>Do LLMs Understand Dialogues? A Case Study on Dialogue Acts</i> Ayesha Qamar, Jonathan Tong and Ruihong Huang	26219
<i>Research Borderlands: Analysing Writing Across Research Cultures</i> Shaily Bhatt, Tal August and Maria Antoniak	26238
<i>CEAES: Bidirectional Reinforcement Learning Optimization for Consistent and Explainable Essay Assessment</i> Xia Li and Wenjing Pan	26267
<i>DeAL: Decoding-time Alignment for Large Language Models</i> James Y. Huang, Sailik Sengupta, Daniele Bonadiman, Yi-An Lai, Arshit Gupta, Nikolaos Pappas, Saab Mansour, Katrin Kirchhoff and Dan Roth	26280
<i>Cultural Bias Matters: A Cross-Cultural Benchmark Dataset and Sentiment-Enriched Model for Understanding Multimodal Metaphors</i> Senqi Yang, Dongyu Zhang, Jing Ren, Ziqi Xu, Xiuzhen Zhang, Yiliao Song, Hongfei Lin and Feng Xia	26301
<i>OmniCharacter: Towards Immersive Role-Playing Agents with Seamless Speech-Language Personality Interaction</i> Haonan Zhang, Run Luo, Xiong Liu, Yuchuan Wu, Ting-En Lin, Pengpeng Zeng, Qiang QU, Feiteng Fang, Min Yang, Lianli Gao, Jingkuan Song, Fei Huang and Yongbin Li	26318
<i>Mixtures of In-Context Learners</i> Giwon Hong, Emile Van Krieken, Edoardo Ponti, Nikolay Malkin and Pasquale Minervini	26332
<i>Balancing Diversity and Risk in LLM Sampling: How to Select Your Method and Parameter for Open-Ended Text Generation</i> Yuxuan Zhou, Margret Keuper and Mario Fritz	26352
<i>RADAR: Enhancing Radiology Report Generation with Supplementary Knowledge Injection</i> Wenjun Hou, Yi Cheng, Kaishuai Xu, Heng Li, Yan Hu, Wenjie Li and Jiang Liu	26366
<i>Can LLMs Deceive CLIP? Benchmarking Adversarial Compositionality of Pre-trained Multimodal Representation via Text Updates</i> Jaewoo Ahn, Heeseung Yun, Dayoon Ko and Gunhee Kim	26382
<i>Attention Speaks Volumes: Localizing and Mitigating Bias in Language Models</i> Rishabh Adiga, Besmira Nushi and Varun Chandrasekaran	26403
<i>MTSA: Multi-turn Safety Alignment for LLMs through Multi-round Red-teaming</i> Weiyang Guo, Jing Li, Wenya Wang, YU LI, Daojing He, Jun Yu and Min Zhang	26424

<i>The Efficiency vs. Accuracy Trade-off: Optimizing RAG-Enhanced LLM Recommender Systems Using Multi-Head Early Exit</i>	
Huixue Zhou, Hengrui Gu, Zaifu Zhan, Xi Liu, Kaixiong Zhou, Yongkang Xiao, Mingfu Liang, Srinivas Prasad Govindan, Piyush Chawla, Jiyan Yang, Xiangfei Meng, Huayu Li, Buyun Zhang, Liang Luo, Wen-Yen Chen, Yiping Han, Bo Long, Rui Zhang and Tianlong Chen	26443
<i>Unraveling LoRA Interference: Orthogonal Subspaces for Robust Model Merging</i>	
Haobo Zhang and Jiayu Zhou	26459
<i>BIG-Bench Extra Hard</i>	
Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John Palowitch, Chrysovalantis Anastasiou, San- ket Vaibhav Mehta, Lalit K Jain, Virginia Aglietti, Disha Jindal, Peter Chen, Nishanth Dikkala, Gladys Tyen, Xin Liu, Uri Shalit, Silvia Chiappa, Kate Olszewska, Yi Tay, Vinh Q. Tran, Quoc V Le and Orhan Firat	26473
<i>CSTree-SRI: Introspection-Driven Cognitive Semantic Tree for Multi-Turn Question Answering over Extra-Long Contexts</i>	
Zhaowen Wang, Xiang Wei, Kangshao Du, Yiting Zhang, Libo Qin, Yingjie Xia and Li Kuang	26502
<i>InductionBench: LLMs Fail in the Simplest Complexity Class</i>	
Wenyue Hua, Tyler Wong, Fei Sun, Liangming Pan, Adam Jardine and William Yang Wang	26526
<i>RATIONALYST: Pre-training Process-Supervision for Improving Reasoning</i>	
Dongwei Jiang, Guoxuan Wang, Yining Lu, Andrew Wang, Jingyu Zhang, Chuyu Liu, Benjamin Van Durme and Daniel Khashabi	26547
<i>Make Imagination Clearer! Stable Diffusion-based Visual Imagination for Multimodal Machine Tran- slation</i>	
Andong Chen, Yuchen Song, Kehai Chen, Xuefeng Bai, Muyun Yang, Liqiang Nie, Jie Liu, Tiejun Zhao and Min Zhang	26567
<i>Advancing SMoE for Continuous Domain Adaptation of MLLMs: Adaptive Router and Domain-Specific Loss</i>	
Liang Zhang, Ziyao Lu, Fandong Meng, Hui Li, Jie Zhou and Jinsong Su	26584
<i>Multi-document Summarization through Multi-document Event Relation Graph Reasoning in LLMs: a case study in Framing Bias Mitigation</i>	
Yuanyuan Lei and Ruihong Huang	26603
<i>Who Writes What: Unveiling the Impact of Author Roles on AI-generated Text Detection</i>	
Jiatao Li and Xiaojuan Wan	26620
<i>RoCoFT: Efficient Finetuning of Large Language Models with Row-Column Updates</i>	
Md Kowsher, Tara Esmaeilbeig, Chun-Nam Yu, Chen Chen, Mojtaba Soltanian and Niloofar Yousefi	26659
<i>Scaling Laws and Efficient Inference for Ternary Language Models</i>	
Tejas Vaidhya, Ayush Kaushal, Vineet Jain, Francis Couture-Harpin, Prashant Shishodia, Majid Behbahani, Yuriy Nevmyvaka and Irina Rish	26679
<i>Exploring the Impact of Instruction-Tuning on LLM’s Susceptibility to Misinformation</i>	
Kyubeen Han, Junseo Jang, Hongjin Kim, Geunyeong Jeong and Harksoo Kim	26711
<i>Do Language Models Understand Honorific Systems in Javanese?</i>	
Mohammad Rifqi Farhansyah, Iwan Darmawan, Adryan Kusumawardhana, Genta Indra Winata, Alham Fikri Aji and Derry Tanti Wijaya	26732

<i>Generative Reward Modeling via Synthetic Criteria Preference Learning</i>	
Xiaobo Liang, Haoke Zhang, Juntao Li, Kehai Chen, Qiaoming Zhu and Min Zhang	26755
<i>Exploring Multimodal Relation Extraction of Hierarchical Tabular Data with Multi-task Learning</i>	
Xinyu Zhang, Aibo Song, Jingyi Qiu, Jiahui Jin, Tianbo Zhang and Xiaolin Fang	26770
<i>A Self-Denoising Model for Robust Few-Shot Relation Extraction</i>	
Liang Zhang, Yang Zhang, Ziyao Lu, Fandong Meng, Jie Zhou and Jinsong Su	26782
<i>QuASAR: A Question-Driven Structure-Aware Approach for Table-to-Text Generation</i>	
WeiJie Liu, Yibin Zheng and Fang Kong	26798
<i>Automated Structured Radiology Report Generation</i>	
Jean-Benoit Delbrouck, Justin Xu, Johannes Moll, Alois Thomas, Zhihong Chen, Sophie Ostmeier, Asfandyar Azhar, Kelvin Zhenghao Li, Andrew Johnston, Christian Bluethgen, Eduardo Pontes Reis, Mohamed S Muneer, Maya Varma and Curtis Langlotz	26813
<i>LPOI: Listwise Preference Optimization for Vision Language Models</i>	
Fatemeh Pesaran Zadeh, Yoojin Oh and Gunhee Kim	26830
<i>Predicting Through Generation: Why Generation Is Better for Prediction</i>	
Md Kowsher, Nusrat Jahan Prottasha, Prakash Bhat, Chun-Nam Yu, Mojtaba Soltanalian, Ivan Garibay, Ozlem Garibay, Chen Chen and Niloofar Yousefi	26845
<i>Give Me BF16 or Give Me Death"? Accuracy-Performance Trade-Offs in LLM Quantization</i>	
Eldar Kurtic, Alexandre Noll Marques, Shubhra Pandit, Mark Kurtz and Dan Alistarh	26872
<i>StitchLLM: Serving LLMs, One Block at a Time</i>	
Bodun Hu, Shuozhe Li, Saurabh Agarwal, Myungjin Lee, Akshay Jajoo, Jiamin Li, Le Xu, Geon-Woo Kim, Donghyun Kim, Hong Xu, Amy Zhang and Aditya Akella	26887
<i>Walk in Others' Shoes with a Single Glance: Human-Centric Visual Grounding with Top-View Perspective Transformation</i>	
Yuqi Bu, Xin Wu, Zirui Zhao, Yi Cai, David Hsu and Qiong Liu	26904
<i>Is linguistically-motivated data augmentation worth it?</i>	
Ray Groshan, Michael Ginn and Alexis Palmer	26924
<i>From Lists to Emojis: How Format Bias Affects Model Alignment</i>	
Xuanchang Zhang, Wei Xiong, Lichang Chen, Tianyi Zhou, Heng Huang and Tong Zhang	26940
<i>Colloquial Singaporean English Style Transfer with Fine-Grained Explainable Control</i>	
Jingui Liang, Dung Vo, Yap Hong Xian, Hai Leong Chieu, Kian Ming A. Chai, Jing Jiang and Lizi Liao	26962
<i>From Informal to Formal – Incorporating and Evaluating LLMs on Natural Language Requirements to Verifiable Formal Proofs</i>	
Jialun Cao, Yaojie Lu, Meiziniu Li, Haoyang Ma, Haokun Li, Mengda He, Cheng Wen, Le Sun, Hongyu Zhang, Shengchao Qin, Shing-Chi Cheung and Cong Tian	26984
<i>CoAM: Corpus of All-Type Multiword Expressions</i>	
Yusuke Ide, Joshua Tanner, Adam Nohejl, Jacob Hoffman, Justin Vasselli, Hidetaka Kamigaito and Taro Watanabe	27004
<i>SeaKR: Self-aware Knowledge Retrieval for Adaptive Retrieval Augmented Generation</i>	
Zijun Yao, Weijian Qi, Liangming Pan, Shulin Cao, Linmei Hu, Liu Weichuan, Lei Hou and Juanzi Li	27022

<i>Exposing the Achilles' Heel: Evaluating LLMs Ability to Handle Mistakes in Mathematical Reasoning</i> Joykirat Singh, Akshay Nambi and Vibhav Vineet.....	27044
<i>Understanding the Dark Side of LLMs' Intrinsic Self-Correction</i> Qingjie Zhang, Di Wang, Haoting Qian, Yiming Li, Tianwei Zhang, Minlie Huang, Ke Xu, Hewu Li, Liu Yan and Han Qiu.....	27066
<i>VideoVista-CulturalLingo: 360° Horizons-Bridging Cultures, Languages, and Domains in Video Comprehension</i> Xinyu Chen, Yunxin li, Haoyuan Shi, Baotian Hu, Wenhan Luo, Yaowei Wang and Min Zhang	27102
<i>What are the Essential Factors in Crafting Effective Long Context Multi-Hop Instruction Datasets? Insights and Best Practices</i> Zhi Chen, Qiguang Chen, Libo Qin, Qipeng Guo, Haijun Lv, Yicheng Zou, Hang Yan, Kai Chen and Dahua Lin.....	27129
<i>Knowledge Graph Retrieval-Augmented Generation for LLM-based Recommendation</i> Shijie Wang, Wenqi Fan, Yue Feng, Lin Shanru, Xinyu Ma, Shuaiqiang Wang and Dawei Yin	27152
<i>SudoLM: Learning Access Control of Parametric Knowledge with Authorization Alignment</i> Qin Liu, Fei Wang, Chaowei Xiao and Muhao Chen.....	27169
<i>IOT: Embedding Standardization Method Towards Zero Modality Gap</i> Na Min An, Euniki Kim, James Thorne and Hyunjung Shim.....	27182
<i>Odysseus Navigates the Sirens' Song: Dynamic Focus Decoding for Factual and Diverse Open-Ended Text Generation</i> Wen Luo, Feifan Song, Wei Li, Guangyue Peng, Shaohang Wei and Houfeng Wang.....	27200
<i>Better Embeddings with Coupled Adam</i> Felix Stollenwerk and Tobias Stollenwerk.....	27219
<i>Bone Soups: A Seek-and-Soup Model Merging Approach for Controllable Multi-Objective Generation</i> Guofu Xie, Xiao Zhang, Ting Yao and Yunsheng Shi.....	27237
<i>Controllable and Reliable Knowledge-Intensive Task-Oriented Conversational Agents with Declarative Genie Worksheets</i> Harshit Joshi, Shicheng Liu, James Chen, Larsen Weigle and Monica Lam.....	27264
<i>Benchmarking Long-Context Language Models on Long Code Understanding</i> Jia Li, Xuyuan Guo, Lei Li, Kechi Zhang, Ge Li, Jia Li, Zhengwei Tao, Fang Liu, Chongyang Tao, Yuqi Zhu and Zhi Jin.....	27309
<i>MAGNET: Augmenting Generative Decoders with Representation Learning and Infilling Capabilities</i> Savya Khosla, Aditi Tiwari, Kushal Kafle, Simon Jenni, Handong Zhao, John Collomosse and Jing Shi.....	27328
<i>Internal Value Alignment in Large Language Models through Controlled Value Vector Activation</i> Haoran Jin, Meng Li, Xiting Wang, Zhihao Xu, Minlie Huang, Yantao Jia and Defu Lian .	27347
<i>A Dual-Perspective NLG Meta-Evaluation Framework with Automatic Benchmark and Better Interpretability</i> Xinyu Hu, Mingqi Gao, Li Lin, Zhenghan Yu and Xiaojun Wan.....	27372

<i>Recurrent Knowledge Identification and Fusion for Language Model Continual Learning</i> Yujie Feng, Xujia Wang, Zexin LU, Shenghong Fu, Guangyuan Shi, Yongxin Xu, Yasha Wang, Philip S. Yu, Xu Chu and Xiao-Ming Wu	27396
<i>Data-Constrained Synthesis of Training Data for De-Identification</i> Thomas Vakili, Aron Henriksson and Hercules Dalianis	27414
<i>Just a Scratch: Enhancing LLM Capabilities for Self-harm Detection through Intent Differentiation and Emoji Interpretation</i> Soumitra Ghosh, Gopendra Vikram Singh, Shambhavi Shambhavi, Sabarna Choudhury and Asif Ekbal	27428
<i>Contrastive Learning on LLM Back Generation Treebank for Cross-domain Constituency Parsing</i> Peiming Guo, Meishan Zhang, Jianling li, Min Zhang and Yue Zhang	27446
<i>MMDEND: Dendrite-Inspired Multi-Branch Multi-Compartment Parallel Spiking Neuron for Sequence Modeling</i> Kexin Wang, Yuhong Chou, Di Shang, Shijie Mei, Jiahong Zhang, Yanbin Huang, Man Yao, Bo XU and Guoqi Li	27459
<i>Understanding Impact of Human Feedback via Influence Functions</i> Taywon Min, Haeone Lee, Yongchan Kwon and Kimin Lee	27471
<i>T2I-FactualBench: Benchmarking the Factuality of Text-to-Image Models with Knowledge-Intensive Concepts</i> Ziwei Huang, Wanggui He, Quanyu Long, Yandi Wang, Haoyuan Li, Zhelun Yu, Fangxun Shu, Weilong Dai, Hao Jiang, Fei Wu and Leilei Gan	27501
<i>InspireDebate: Multi-Dimensional Subjective-Objective Evaluation-Guided Reasoning and Optimiza- tion for Debating</i> Fuyu Wang, Jiangtong Li, Kun Zhu and Changjun Jiang	27525
<i>OpenWebVoyager: Building Multimodal Web Agents via Iterative Real-World Exploration, Feedback and Optimization</i> Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Hongming Zhang, Tianqing Fang, Zhenz- hong Lan and Dong Yu	27545
<i>FOCUS: Evaluating Pre-trained Vision-Language Models on Underspecification Reasoning</i> Kankan Zhou, Eason Lai, Kyriakos Mouratidis and Jing Jiang	27565
<i>Sightation Counts: Leveraging Sighted User Feedback in Building a BLV-aligned Dataset of Diagram Descriptions</i> Wan Ju Kang, Eunki Kim, Na Min An, Sangryul Kim, Haemin Choi, Ki Hoon Kwak and James Thorne	27585
<i>Personal Travel Solver: A Preference-Driven LLM-Solver System for Travel Planning</i> Zijian Shao, Jiancan Wu, Weijian Chen and Xiang Wang	27622
<i>Counterspeech the ultimate shield! Multi-Conditioned Counterspeech Generation through Attributed Prefix Learning</i> Aswini Kumar Padhi, Anil Bandhakavi and Tanmoy Chakraborty	27643
<i>LLM×MapReduce: Simplified Long-Sequence Processing using Large Language Models</i> Zihan Zhou, Chong Li, Xinyi Chen, Shuo Wang, Yu Chao, Zhili Li, Haoyu Wang, Qi Shi, Zhixing Tan, Xu Han, Xiaodong Shi, Zhiyuan Liu and Maosong Sun	27664

<i>CheXalign: Preference fine-tuning in chest X-ray interpretation models without human feedback</i>	
Dennis Hein, Zhihong Chen, Sophie Ostmeier, Justin Xu, Maya Varma, Eduardo Pontes Reis, Arne Edward Michalson MD, Christian Bluethgen, Hyun Joo Shin, Curtis Langlotz and Akshay S Chaudhari	27679
<i>Knowledge Tracing in Programming Education Integrating Students' Questions</i>	
Doyoun Kim, Suin Kim and Yohan Jo	27703
<i>PRISM: A Framework for Producing Interpretable Political Bias Embeddings with Political-Aware Cross-Encoder</i>	
Yiqun Sun, Qiang Huang, Anthony Kum Hoe Tung and Jun Yu	27719
<i>Representations of Fact, Fiction and Forecast in Large Language Models: Epistemics and Attitudes</i>	
Meng Li, Michael Vrazitulis and David Schlangen	27734
<i>Lexical Diversity-aware Relevance Assessment for Retrieval-Augmented Generation</i>	
Zhange Zhang, Yuqing Ma, Yulong Wang, Shan He, Tianbo Wang, Siqi He, Jiakai Wang and Xianglong Liu	27758
<i>Weaving Context Across Images: Improving Vision-Language Models through Focus-Centric Visual Chains</i>	
Juntian Zhang, Chuanqi Cheng, Yuhan Liu, Wei Liu, Jian Luan and Rui Yan	27782
<i>Online Iterative Self-Alignment for Radiology Report Generation</i>	
Ting Xiao, Lei Shi, Yang Zhang, HaoFeng Yang, Zhe Wang and Chenjia Bai	27799
<i>Chinese Inertial GAN for Handwriting Signal Generation and Recognition</i>	
Yifeng Wang and Yi Zhao	27815
<i>LLMs Caught in the Crossfire: Malware Requests and Jailbreak Challenges</i>	
Haoyang Li, Huan Gao, Zhiyuan Zhao, Zhiyu Lin, Junyu Gao and Xuelong Li	27833
<i>Evaluating Sequence Labeling on the basis of Information Theory</i>	
Enrique Amigo, Elena Álvarez-Mellado, Julio Gonzalo and Jorge Carrillo-de-Albornoz ...	27849
<i>GRAT: Guiding Retrieval-Augmented Reasoning through Process Rewards Tree Search</i>	
Xianshu Peng and Wei Wei	27861
<i>T-REG: Preference Optimization with Token-Level Reward Regularization</i>	
Wenxuan Zhou, Shujian Zhang, Lingxiao Zhao and Tao Meng	27876
<i>Gödel Agent: A Self-Referential Agent Framework for Recursively Self-Improvement</i>	
Xunjian Yin, Xinyi Wang, Liangming Pan, Li Lin, Xiaojun Wan and William Yang Wang .	27890
<i>AgentGym: Evaluating and Training Large Language Model-based Agents across Diverse Environments</i>	
Zhiheng Xi, Yiwen Ding, Wenxiang Chen, Boyang Hong, Honglin Guo, Junzhe Wang, Xin Guo, Dingwen Yang, Chenyang Liao, Wei He, Songyang Gao, Lu Chen, Rui Zheng, Yicheng Zou, Tao Gui, Qi Zhang, Xipeng Qiu, Xuanjing Huang, Zuxuan Wu and Yu-Gang Jiang	27914
<i>Rethinking the Role of Prompting Strategies in LLM Test-Time Scaling: A Perspective of Probability Theory</i>	
Yexiang Liu, Zekun Li, Zhi Fang, Nan Xu, Ran He and Tieniu Tan	27962
<i>Information Locality as an Inductive Bias for Neural Language Models</i>	
Taiga Someya, Anej Svete, Brian DuSell, Timothy J. O'Donnell, Mario Giulianelli and Ryan Cotterell	27995

<i>Learning to Reason Over Time: Timeline Self-Reflection for Improved Temporal Reasoning in Language Models</i>	
Adrián Bazaga, Rexhina Blloshmi, Bill Byrne and Adrià de Gispert	28014
<i>Query-driven Document-level Scientific Evidence Extraction from Biomedical Studies</i>	
Massimiliano Pronesti, Joao H Bettencourt-Silva, Paul Flanagan, Alessandra Pascale, Oisín Redmond, Anya Belz and Yufang Hou	28034
<i>Towards Robust Universal Information Extraction: Dataset, Evaluation, and Solution</i>	
Jizhao Zhu, Akang Shi, Zixuan Li, Long Bai, Xiaolong Jin, Jiafeng Guo and Xueqi Cheng	28052
<i>Multi-perspective Alignment for Increasing Naturalness in Neural Machine Translation</i>	
Huiyuan Lai, Esther Ploeger, Rik Van Noord and Antonio Toral	28071
<i>Temporal reasoning for timeline summarisation in social media</i>	
Jiayu Song, Mahmud Elahi Akhter, Dana Atzil-Slonim and Maria Liakata	28085
<i>Beyond Negative Stereotypes – Non-Negative Abusive Utterances about Identity Groups and Their Semantic Variants</i>	
Tina Lommel, Elisabeth Eder, Josef Ruppenhofer and Michael Wiegand	28102
<i>Persistent Homology of Topic Networks for the Prediction of Reader Curiosity</i>	
Manuel D.s. Hopp, Vincent Labatut, Arthur Amalvy, Richard Dufour, Hannah Stone, Hayley K Jach and Kou Murayama	28121
<i>Tokenisation is NP-Complete</i>	
Philip Whittington, Gregor Bachmann and Tiago Pimentel	28133
<i>Training Dynamics Underlying Language Model Scaling Laws: Loss Deceleration and Zero-Sum Learning</i>	
Andrei Mircea, Supriyo Chakraborty, Nima Chitsazan, Irina Rish and Ekaterina Lobacheva	28154
<i>Parameter-Aware Contrastive Knowledge Editing: Tracing and Rectifying based on Critical Transmission Paths</i>	
Songlin Zhai, Yuan Meng, Yuxin Zhang and Guilin Qi	28189
<i>Many Heads Are Better Than One: Improved Scientific Idea Generation by A LLM-Based Multi-Agent System</i>	
Haoyang Su, Renqi Chen, Shixiang Tang, Zhenfei Yin, Xinzhe Zheng, Jinzhe Li, Biqing Qi, Qi Wu, Hui Li, Wanli Ouyang, Philip Torr, Bowen Zhou and Nanqing Dong	28201
<i>Inner Thinking Transformer: Leveraging Dynamic Depth Scaling to Foster Adaptive Internal Thinking</i>	
Yilong Chen, Junyuan Shang, Zhenyu Zhang, Yanxi Xie, Jiawei Sheng, Tingwen Liu, Shuohuan Wang, Yu Sun, Hua Wu and Haifeng Wang	28241
<i>Document-Level Text Generation with Minimum Bayes Risk Decoding using Optimal Transport</i>	
Yuu Jinnai	28260
<i>Opt-Out: Investigating Entity-Level Unlearning for Large Language Models via Optimal Transport</i>	
Minseok Choi, Daniel Rim, Dohyun Lee and Jaegul Choo	28280
<i>Mixture of Small and Large Models for Chinese Spelling Check</i>	
Ziheng Qiao, Houquan Zhou and Zhenghua Li	28298
<i>DISC: Plug-and-Play Decoding Intervention with Similarity of Characters for Chinese Spelling Check</i>	
Ziheng Qiao, Houquan Zhou, Yumeng Liu, Zhenghua Li, Min Zhang, Bo Zhang, Chen Li, Ji Zhang and Fei Huang	28312

<i>Causal Estimation of Tokenisation Bias</i>	
Pietro Lesci, Clara Meister, Thomas Hofmann, Andreas Vlachos and Tiago Pimentel	28325
<i>Value Residual Learning</i>	
Zhanchao Zhou, Tianyi Wu, Zhiyun Jiang, Fares Obeid and Zhenzhong Lan	28341
<i>SGIC: A Self-Guided Iterative Calibration Framework for RAG</i>	
Guanhua Chen, Yutong Yao, Lidia S. Chao, Xuebo Liu and Derek F. Wong	28357
<i>NusaAksara: A Multimodal and Multilingual Benchmark for Preserving Indonesian Indigenous Scripts</i>	
Muhammad Farid Adilazuarda, Musa Izzanardi Wijanarko, Lucky Susanto, Khumaisa Nur'aini, Derry Tanti Wijaya and Alham Fikri Aji	28371
<i>LLM-based Rumor Detection via Influence Guided Sample Selection and Game-based Perspective Analysis</i>	
Zhiliang Tian, Jingyuan Huang, Zejiang He, Zhen Huang, Menglong Lu, Linbo Qiao, Songzhu Mei, Yijie Wang and Dongsheng Li	28402
<i>Hierarchical-Task-Aware Multi-modal Mixture of Incremental LoRA Experts for Embodied Continual Learning</i>	
Ziqi Jia, Anmin Wang, Xiaoyang Qu, Xiaowen Yang and Jianzong Wang	28415
<i>SpindleKV: A Novel KV Cache Reduction Method Balancing Both Shallow and Deep Layers</i>	
Zicong Tang, Shi Luohe, Zuchao Li, Baoyuan Qi, Liu Guoming, Lefei Zhang and Ping Wang	28428
<i>Medical Graph RAG: Evidence-based Medical Large Language Model via Graph Retrieval-Augmented Generation</i>	
Junde Wu, Jiayuan Zhu, Yunli Qi, Jingkun Chen, Min Xu, Filippo Menolascina, Yueming Jin and Vicente Grau	28443
<i>Unifying Uniform and Binary-coding Quantization for Accurate Compression of Large Language Models</i>	
Seungcheol Park, Jeongin Bae, Beomseok Kwon, Minjun Kim, Byeongwook Kim, Se Jung Kwon, U Kang and Dongsoo Lee	28468
<i>Agentic Reasoning: A Streamlined Framework for Enhancing LLM Reasoning with Agentic Tools</i>	
Junde Wu, Jiayuan Zhu, Yuyuan Liu, Min Xu and Yueming Jin	28489
<i>Probing Relative Interaction and Dynamic Calibration in Multi-modal Entity Alignment</i>	
Chenxiao Li, Jingwei Cheng, Qiang Tong, Fu Zhang and Cairui Wang	28504
<i>Learn to Memorize: Scalable Continual Learning in Semiparametric Models with Mixture-of-Neighbors Induction Memory</i>	
Guangyue Peng, Tao Ge, Wen Luo, Wei Li and Houfeng Wang	28517
<i>Adverse Event Extraction from Discharge Summaries: A New Dataset, Annotation Scheme, and Initial Findings</i>	
Imane Guellil, Salomé Andres, Atul Anand, Bruce Guthrie, Huayu Zhang, Abul Hasan, Honghan Wu and Beatrice Alex	28532
<i>Speed Up Your Code: Progressive Code Acceleration Through Bidirectional Tree Editing</i>	
Longhui Zhang, Jiahao Wang, Meishan Zhang, GaoXiong Cao, Ensheng Shi, Mayuchi Mayuchi, Jun Yu, Honghai Liu, Jing Li and Min Zhang	28563
<i>Multi-Facet Blending for Faceted Query-by-Example Retrieval</i>	
Heejin Do, Sangwon Ryu, Jonghwi Kim and Gary Lee	28577

<i>PIPER: Benchmarking and Prompting Event Reasoning Boundary of LLMs via Debiasing-Distillation Enhanced Tuning</i>	
Zhicong Lu, Changyuan Tian, PeiguangLi PeiguangLi, Li Jin, Sirui Wang, Wei Jia, Ying Shen and Guangluan Xu	28591
<i>MIR: Methodology Inspiration Retrieval for Scientific Research Problems</i>	
Aniketh Garikaparthi, Manasi Patwardhan, Aditya Sanjiv Kanade, Aman Hassan, Lovekesh Vig and Arman Cohan	28614
<i>Sticking to the Mean: Detecting Sticky Tokens in Text Embedding Models</i>	
Kexin Chen, Dongxia Wang, Yi Liu, Haonan Zhang and Wenhai Wang	28660
<i>Memorizing is Not Enough: Deep Knowledge Injection Through Reasoning</i>	
Ruoxi Xu, Yunjie Ji, Boxi Cao, Yaojie Lu, Hongyu Lin, Xianpei Han, Ben He, Yingfei Sun, Xiangang Li and Le Sun	28682
<i>Improving Dialogue State Tracking through Combinatorial Search for In-Context Examples</i>	
Haesung Pyun, Yoonah Park and Yohan Jo	28694
<i>Pretraining Context Compressor for Large Language Models with Embedding-Based Memory</i>	
Yuhong Dai, Jianxun Lian, Yitian Huang, Wei Zhang, Mingyang Zhou, Mingqi Wu, Xing Xie and Hao Liao	28715
<i>Dialogue Systems for Emotional Support via Value Reinforcement</i>	
Juhee Kim, Chung-hu Mok, Jisun Lee, Hyang Sook Kim and Yohan Jo	28733
<i>Length-Induced Embedding Collapse in PLM-based Models</i>	
Yuqi Zhou, Sunhao Dai, Zhanshuo Cao, Xiao Zhang and Jun Xu	28767
<i>SHuBERT: Self-Supervised Sign Language Representation Learning via Multi-Stream Cluster Prediction</i>	
Shester Gueuwou, Xiaodan Du, Greg Shakhnarovich, Karen Livescu and Alexander H. Liu	28792
<i>ERU-KG: Efficient Reference-aligned Unsupervised Keyphrase Generation</i>	
Lam Thanh Do, Aaditya Bodke, Pritom Saha Akash and Kevin Chen-Chuan Chang	28811
<i>Know Your Mistakes: Towards Preventing Overreliance on Task-Oriented Conversational AI Through Accountability Modeling</i>	
Suvodip Dey, Yi-Jyun Sun, Gokhan Tur and Dilek Hakkani-Tür	28830
<i>LLMs Trust Humans More, That’s a Problem! Unveiling and Mitigating the Authority Bias in Retrieval-Augmented Generation</i>	
Yuxuan Li, Xinwei Guo, Jiashi Gao, Guanhua Chen, Xiangyu Zhao, Jiabin Zhang, Quanying Liu, Haiyan Wu, Xin Yao and Xuetao Wei	28844
<i>Divide-Then-Aggregate: An Efficient Tool Learning Method via Parallel Tool Invocation</i>	
Dongsheng Zhu, Weixian Shi, Zhengliang Shi, Zhaochun Ren, Shuaiqiang Wang, Lingyong Yan and Dawei Yin	28859
<i>Reviving Cultural Heritage: A Novel Approach for Comprehensive Historical Document Restoration</i>	
Yuyi Zhang, Peirong Zhang, Zhenhua Yang, Pengyu Yan, Yongxin Shi, Pengwei Liu, Fengjun Guo and Lianwen Jin	28876
<i>PopAlign: Diversifying Contrasting Patterns for a More Comprehensive Alignment</i>	
Zekun Moore Wang, Shenzi Wang, King Zhu, Jiaheng Liu, Ke Xu, Jie Fu, Wangchunshu Zhou and Wenhao Huang	28893

<i>Robust Utility-Preserving Text Anonymization Based on Large Language Models</i> Tianyu Yang, Xiaodan Zhu and Iryna Gurevych.....	28922
<i>SEAL: Scaling to Emphasize Attention for Long-Context Retrieval</i> Changhun Lee, Minsang Seok, Jun-gyu Jin, YoungHyun Cho and Eunhyeok Park.....	28942
<i>From Neurons to Semantics: Evaluating Cross-Linguistic Alignment Capabilities of Large Language Models via Neurons Alignment</i> Chongxuan Huang, Yongshi Ye, Biao Fu, Qifeng Su and Xiaodong Shi.....	28956
<i>A³: Automatic Alignment Framework for Attributed Text Generation</i> Yue Wang, Haoke Zhang, Juntao Li, Jinxiong Chang and Min Zhang.....	28975
<i>Towards Better Value Principles for Large Language Model Alignment: A Systematic Evaluation and Enhancement</i> Bingbing Xu, Jing Yao, Xiaoyuan Yi, Aishan Maolinyazi, Xing Xie and Xiaofeng Meng .	28991
<i>Language Models, Graph Searching, and Supervision Adulteration: When More Supervision is Less and How to Make More More</i> Arvid Frydenlund.....	29011
<i>Diversity Explains Inference Scaling Laws: Through a Case Study of Minimum Bayes Risk Decoding</i> Hidetaka Kamigaito, Hiroyuki Deguchi, Yusuke Sakai, Katsuhiko Hayashi and Taro Watanabe	29060
<i>Performance Gap in Entity Knowledge Extraction Across Modalities in Vision Language Models</i> Ido Cohen, Daniela Gottesman, Mor Geva and Raja Giryes	29095
<i>SDD: Self-Degraded Defense against Malicious Fine-tuning</i> ZiXuan Chen, Weikai Lu, Xin Lin and Ziqian Zeng	29109
<i>CoachMe: Decoding Sport Elements with a Reference-Based Coaching Instruction Generation Model</i> Wei-Hsin Yeh, Yu-An Su, Chih-Ning Chen, Yi-Hsueh Lin, Calvin Ku, Wenhsin Chiu, Min-Chun Hu and Lun-Wei Ku.....	29126
<i>DRPruning: Efficient Large Language Model Pruning through Distributionally Robust Optimization</i> Hexuan Deng, Wenxiang Jiao, Xuebo Liu, Jing Li, Min Zhang and Zhaopeng Tu.....	29152
<i>How LLMs Comprehend Temporal Meaning in Narratives: A Case Study in Cognitive Evaluation of LLMs</i> Karin De Langis, Jong Inn Park, Andreas Schramm, Bin Hu, Khanh Chi Le and Dongyeop Kang	29174
<i>Data Caricatures: On the Representation of African American Language in Pretraining Corpora</i> Nicholas Deas, Blake Vente, Amith Ananthram, Jessica A Grieser, Desmond U. Patton, Shana Kleiner, James R. Shepard Iii and Kathleen McKeown.....	29192
<i>Language Model Probabilities are Not Calibrated in Numeric Contexts</i> Charles Lovering, Michael Krumdick, Viet Dac Lai, Varshini Reddy, Seth Ebner, Nilesh Kumar, Rik Koncel-Kedziorski and Chris Tanner	29218
<i>MDCure: A Scalable Pipeline for Multi-Document Instruction-Following</i> Gabrielle Kaili-May Liu, Bowen Shi, Avi Caciularu, Idan Szpektor and Arman Cohan....	29258
<i>Cross-Lingual Auto Evaluation for Assessing Multilingual LLMs</i> Sumanth Doddapaneni, Mohammed Safi Ur Rahman Khan, Dilip Venkatesh, Raj Dabre, Anoop Kunchukuttan and Mitesh M Khapra.....	29297

<i>DeepReview: Improving LLM-based Paper Review with Human-like Deep Thinking Process</i> Minjun Zhu, Yixuan Weng, Linyi Yang and Yue Zhang	29330
<i>Bypass Back-propagation: Optimization-based Structural Pruning for Large Language Models via Policy Gradient</i> Yuan Gao, Zujing Liu, Weizhong Zhang, Bo Du and Gui-Song Xia	29356
<i>Tree-of-Debate: Multi-Persona Debate Trees Elicit Critical Thinking for Scientific Comparative Analysis</i> Priyanka Kargupta, Ishika Agarwal, Tal August and Jiawei Han	29378
<i>Hierarchical Memory Organization for Wikipedia Generation</i> Eugene J. Yu, Dawei Zhu, Yifan Song, Xiangyu Wong, Jiebin Zhang, Wenxuan Shi, Xiaoguang Li, Qun Liu and Sujian Li	29404
<i>Class Distillation with Mahalanobis Contrast: An Efficient Training Paradigm for Pragmatic Language Understanding Tasks</i> Chenlu Wang, Weimin Lyu and Ritwik Banerjee	29428
<i>Structure-aware Domain Knowledge Injection for Large Language Models</i> Kai Liu, Ze Chen, Zhihang Fu, Wei Zhang, Rongxin Jiang, Fan Zhou, Yaowu Chen, Yue Wu and Jieping Ye	29443
<i>FinMME: Benchmark Dataset for Financial Multi-Modal Reasoning Evaluation</i> Junyu Luo, Zhizhuo Kou, Liming Yang, Xiao Luo, Jinsheng Huang, Zhiping Xiao, Jingshu Peng, Chengzhong Liu, Jiaming Ji, Xuanzhe Liu, Sirui Han, Ming Zhang and Yike Guo	29465
<i>Dialectal Coverage And Generalization in Arabic Speech Recognition</i> Amirbek Djanibekov, Hawau Olamide Toyin, Raghad Alshalan, Abdullah Alatir and Hanan Al-darmaki	29490
<i>EditInspector: A Benchmark for Evaluation of Text-Guided Image Edits</i> Ron Yosef, Yonatan Bitton, Dani Lischinski and Moran Yanuka	29503
<i>Reconsidering LLM Uncertainty Estimation Methods in the Wild</i> Yavuz Faruk Bakman, Duygu Nur Yaldiz, Sungmin Kang, Tuo Zhang, Baturalp Buyukates, Salman Avestimehr and Sai Praneeth Karimireddy	29531
<i>Bregman Conditional Random Fields: Sequence Labeling with Parallelizable Inference Algorithms</i> Caio Corro, Mathieu Lacroix and Joseph Le Roux	29557
<i>SEE: Strategic Exploration and Exploitation for Cohesive In-Context Prompt Optimization</i> Wendi Cui, Jiaxin Zhang, Zhuohang Li, Hao Sun, Damien Lopez, Kamalika Das, Bradley A. Malin and Sricharan Kumar	29575
<i>Programming by Example meets Historical Linguistics: A Large Language Model Based Approach to Sound Law Induction</i> Atharva Naik, Darsh Agrawal, Hong Sng, Clayton Marr, Kexun Zhang, Nathaniel Romney Robinson, Calvin Chang, Rebecca Byrnes, Aravind Mysore, Carolyn Rose and David R. Mortensen	29628
<i>Synergizing Unsupervised Episode Detection with LLMs for Large-Scale News Events</i> Priyanka Kargupta, Yunyi Zhang, Yizhu Jiao, Siru Ouyang and Jiawei Han	29648
<i>Beyond True or False: Retrieval-Augmented Hierarchical Analysis of Nuanced Claims</i> Priyanka Kargupta, Runchu Tian and Jiawei Han	29664

<i>The Task Shield: Enforcing Task Alignment to Defend Against Indirect Prompt Injection in LLM Agents</i> Feiran Jia, Tong Wu, Xin Qin and Anna Squicciarini	29680
<i>Sandcastles in the Storm: Revisiting the (Im)possibility of Strong Watermarking</i> Fabrice Y Harel-Canada, Boran Erol, Connor Choi, Jason Liu, Gary Jiarui Song, Nanyun Peng and Amit Sahai	29698
<i>Time-MQA: Time Series Multi-Task Question Answering with Context Enhancement</i> Yaxuan Kong, Yiyuan Yang, Yoontae Hwang, Wenjie Du, Stefan Zohren, Zhangyang Wang, Ming Jin and Qingsong Wen.....	29736
<i>From Perceptions to Decisions: Wildfire Evacuation Decision Prediction with Behavioral Theory- informed LLMs</i> Ruxiao Chen, Chenguang Wang, Yuran Sun, Xilei Zhao and Susu Xu	29754
<i>GETReason: Enhancing Image Context Extraction through Hierarchical Multi-Agent Reasoning</i> Shikhar Singh, Abhinav Rawat, Chitta Baral and Vivek Gupta	29779
<i>Hanging in the Balance: Pivotal Moments in Crisis Counseling Conversations</i> Vivian Nguyen, Lillian Lee and Cristian Danescu-Niculescu-Mizil.....	29801
<i>Unveiling the Potential of BERT-family: A New Recipe for Building Scalable, General and Competitive Large Language Models</i> Yisheng Xiao, Juntao Li, Wenpeng Hu, Zhunchen Luo and Min Zhang	29818
<i>TaxoAdapt: Aligning LLM-Based Multidimensional Taxonomy Construction to Evolving Research Cor- pora</i> Priyanka Kargupta, Nan Zhang, Yunyi Zhang, Rui Zhang, Prasenjit Mitra and Jiawei Han .	29834
<i>An Empirical Study of Iterative Refinements for Non-autoregressive Translation</i> Yisheng Xiao, Pei Guo, Zechen Sun, Juntao Li, Kai Song and Min Zhang	29851
<i>Retrofitting Large Language Models with Dynamic Tokenization</i> Darius Feher, Ivan Vulić and Benjamin Minixhofer	29866
<i>Principled Content Selection to Generate Diverse and Personalized Multi-Document Summaries</i> Vishakh Padmakumar, Zichao Wang, David Arbour and Jennifer Healey	29884
<i>Bilingual Zero-Shot Stance Detection</i> Chenye Zhao and Cornelia Caragea.....	29900
<i>GrammarMT: Improving Machine Translation with Grammar-Informed In-Context Learning</i> Rita Ramos, Evelyn Asiko Chimoto, Maartje Ter Hoeve and Natalie Schluter	29920
<i>Theorem Prover as a Judge for Synthetic Data Generation</i> Joshua Ong Jun Leang, Giwon Hong, Wenda Li and Shay B Cohen	29941
<i>Measuring the Effect of Transcription Noise on Downstream Language Understanding Tasks</i> Ori Shapira, Shlomo Chazan and Amir David Nissan Cohen	29978
<i>Assessing Reliability and Political Bias In LLMs' Judgements of Formal and Material Inferences With Partisan Conclusions</i> Reto Gubelmann and Ghassen Karray.....	30005
<i>PARME: Parallel Corpora for Low-Resourced Middle Eastern Languages</i> Sina Ahmadi, Rico Sennrich, Erfan Karami, Ako Marani, Parviz Fekrazad, Gholamreza Akbarza- deh Baghban, Hanah Hadi, Semko Heidari, Mahîr Dogan, Pedram Asadi, Dashne Bashir, Mohammad	

Amin Ghodrati, Kourosh Amini, Zeynab Ashourinezhad, Mana Baladi, Farshid Ezzati, Alireza Ghasemifar, Daryoush Hosseinpour, Behrooz Abbaszadeh, Amin Hassanpour, Bahaddin Jalal Hamaamin, Saya Kamal Hama, Ardeshir Mousavi, Sarko Nazir Hussein, Isar Nejadgholi, Mehmet Ölmez, Horam Osmanpour, Rashid Roshan Ramezani, Aryan Sediq Aziz, Ali Salehi, Mohammadreza Yadegari, Kewyar Yadegari and Sedighe Zamani Roodsari	30032
<i>METAL: A Multi-Agent Framework for Chart Generation with Test-Time Scaling</i> Bingxuan Li, Yiwei Wang, Jiuxiang Gu, Kai-Wei Chang and Nanyun Peng	30054
<i>ConLoan: A Contrastive Multilingual Dataset for Evaluating Loanwords</i> Sina Ahmadi, Micha David Hess, Elena Álvarez-Mellado, Alessia Battisti, Cui Ding, Anne Göhring, Yingqiang Gao, Zifan Jiang, Andrianos Michail, Peshmerge Morad, Joel Niklaus, Maria Christina Panagiotopoulou, Stefano Perrella, Juri Opitz, Anastassia Shaitarova and Rico Sennrich	30070
<i>A Theory of Response Sampling in LLMs: Part Descriptive and Part Prescriptive</i> Sarath Sivaprasad, Pramod Kaushik, Sahar Abdelnabi and Mario Fritz	30091
<i>MEraser: An Effective Fingerprint Erasure Approach for Large Language Models</i> Jingxuan Zhang, Zhenhua Xu, Rui Hu, Wenpeng Xing, Xuhong Zhang and Meng Han	30136
<i>VISA: Retrieval Augmented Generation with Visual Source Attribution</i> Xueguang Ma, Shengyao Zhuang, Bevan Koopman, Guido Zuccon, Wenhua Chen and Jimmy Lin	30154
<i>DRAMA: Diverse Augmentation from Large Language Models to Smaller Dense Retrievers</i> Xueguang Ma, Xi Victoria Lin, Barlas Oğuz, Jimmy Lin, Wen-tau Yih and Xilun Chen	30170
<i>Stochastic Chameleons: Irrelevant Context Hallucinations Reveal Class-Based (Mis)Generalization in LLMs</i> Ziling Cheng, Meng Cao, Marc-Antoine Rondeau and Jackie CK Cheung	30187
<i>MAPoRL: Multi-Agent Post-Co-Training for Collaborative Large Language Models with Reinforcement Learning</i> Chanwoo Park, Seungju Han, Xingzhi Guo, Asuman E. Ozdaglar, Kaiqing Zhang and Joo-Kyung Kim	30215
<i>Map&Make: Schema Guided Text to Table Generation</i> Naman Ahuja, Fenil Bardoliya, Chitta Baral and Vivek Gupta	30249
<i>IRIS: Interpretable Retrieval-Augmented Classification for Long Interspersed Document Sequences</i> Fengnan Li, Elliot D. Hill, Jiang Shu, Jiaxin Gao and Matthew M. Engelhard	30263
<i>Symmetrical Visual Contrastive Optimization: Aligning Vision-Language Models with Minimal Contrastive Images</i> Shengguang Wu, Fan-Yun Sun, Kaiyue Wen and Nick Haber	30284
<i>Can we Retrieve Everything All at Once? ARM: An Alignment-Oriented LLM-based Retrieval Method</i> Peter Baile Chen, Yi Zhang, Mike Cafarella and Dan Roth	30298
<i>R2D2: Remembering, Replaying and Dynamic Decision Making with a Reflective Agentic Memory</i> Tenghao Huang, Kinjal Basu, Ibrahim Abdelaziz, Pavan Kapanipathi, Jonathan May and Muhao Chen	30318
<i>FairI Tales: Evaluation of Fairness in Indian Contexts with a Focus on Bias and Stereotypes</i> Janki Atul Nawale, Mohammed Safi Ur Rahman Khan, Janani D, Mansi Gupta, Danish Pruthi and Mitesh M Khapra	30331

<i>SpeechIQ: Speech-Agentic Intelligence Quotient Across Cognitive Levels in Voice Understanding by Large Language Models</i>	
Zhen Wan, Chao-Han Huck Yang, Yahan Yu, Jinchuan Tian, Sheng Li, Ke Hu, Zhehuai Chen, Shinji Watanabe, Fei Cheng, Chenhui Chu and Sadao Kurohashi	30381
<i>Predicting Implicit Arguments in Procedural Video Instructions</i>	
Anil Batra, Laura Sevilla-Lara, Marcus Rohrbach and Frank Keller	30399
<i>PIGuard: Prompt Injection Guardrail via Mitigating Overdefense for Free</i>	
Hao Li, Xiaogeng Liu, Ning Zhang and Chaowei Xiao	30420
<i>CLIPerase: Efficient Unlearning of Visual-Textual Associations in CLIP</i>	
Tianyu Yang, Lisen Dai, Xiangqi Wang, Minhao Cheng, Yapeng Tian and Xiangliang Zhang	30438
<i>ViGiL3D: A Linguistically Diverse Dataset for 3D Visual Grounding</i>	
Austin Wang, ZeMing Gong and Angel X Chang	30453
<i>The time scale of redundancy between prosody and linguistic context</i>	
Tamar I Regev, Chiebuka Ohams, Shaylee Xie, Lukas Wolf, Evelina Fedorenko, Alex Warstadt, Ethan Wilcox and Tiago Pimentel	30476
<i>Basic Reading Distillation</i>	
Zhi Zhou, Sirui Miao, Xiangyu Duan, Hao Yang and Min Zhang	30489
<i>Quantized Can Still Be Calibrated: A Unified Framework to Calibration in Quantized Large Language Models</i>	
Mingyu Zhong, Guanchu Wang, Yu-Neng Chuang and Na Zou	30503
<i>A Spatio-Temporal Point Process for Fine-Grained Modeling of Reading Behavior</i>	
Francesco Ignazio Re, Andreas Opedal, Glib Manaiev, Mario Giulianelli and Ryan Cotterell	30518
<i>More is not always better? Enhancing Many-Shot In-Context Learning with Differentiated and Reweighting Objectives</i>	
Xiaoqing Zhang, Ang Lv, Yuhan Liu, Flood Sung, Wei Liu, Jian Luan, Shuo Shang, Xiuying Chen and Rui Yan	30539
<i>Astute RAG: Overcoming Imperfect Retrieval Augmentation and Knowledge Conflicts for Large Language Models</i>	
Fei Wang, Xingchen Wan, Ruoxi Sun, Jiefeng Chen and Sercan O Arik	30553
<i>SubLIME: Subset Selection via Rank Correlation Prediction for Data-Efficient LLM Evaluation</i>	
Gayathri Saranathan, Cong Xu, Mahammad Parwez Alam, Tarun Kumar, Martin Foltin, Soon Yee Wong and Suparna Bhattacharya	30572
<i>M³GQA: A Multi-Entity Multi-Hop Multi-Setting Graph Question Answering Benchmark</i>	
Boci Peng, Yongchao Liu, Xiaohe Bo, Jiaxin Guo, Yun Zhu, Xuanbo Fan, Chuntao Hong and Yan Zhang	30594
<i>LSSF: Safety Alignment for Large Language Models through Low-Rank Safety Subspace Fusion</i>	
Guanghao Zhou, Panjia Qiu, Cen Chen, Hongyu Li, Jason Chu, Xin Zhang and Jun Zhou .	30621
<i>ETF: An Entity Tracing Framework for Hallucination Detection in Code Summaries</i>	
Kishan Maharaj, Vitobha Munigala, Srikanth G. Tamilselvam, Prince Kumar, Sayandeep Sen, Palani Kodeswaran, Abhijit Mishra and Pushpak Bhattacharyya	30639

<i>Meta-Tool: Unleash Open-World Function Calling Capabilities of General-Purpose Large Language Models</i>	
Shengqian Qin, Yakun Zhu, Linjie Mu, Shaoting Zhang and Xiaofan Zhang	30653
<i>Benchmarking and Improving Large Vision-Language Models for Fundamental Visual Graph Understanding and Reasoning</i>	
Yingjie Zhu, Xuefeng Bai, Kehai Chen, Yang Xiang, Jun Yu and Min Zhang	30678
<i>ISR: Self-Refining Referring Expressions for Entity Grounding</i>	
Zhuocheng Yu, Bingchan Zhao, Yifan Song, Sujian Li and Zhonghui HE	30702
<i>Activating Distributed Visual Region within LLMs for Efficient and Effective Vision-Language Training and Inference</i>	
Siyuan Wang, Dianyi Wang, Chengxing Zhou, Zejun Li, Zhihao Fan, Xuanjing Huang and Zhongyu Wei	30715
<i>CCHall: A Novel Benchmark for Joint Cross-Lingual and Cross-Modal Hallucinations Detection in Large Language Models</i>	
Yongheng Zhang, Xu Liu, Ruoxi Zhou, Qiguang Chen, Hao Fei, Wenpeng Lu and Libo Qin	30728
<i>TestNUC: Enhancing Test-Time Computing Approaches and Scaling through Neighboring Unlabeled Data Consistency</i>	
Henry Peng Zou, Zhengyao Gu, Yue Zhou, Yankai Chen, Weizhi Zhang, Liancheng Fang, Yibo Wang, Yangning Li, Kay Liu and Philip S. Yu	30750
<i>The Esethu Framework: Reimagining Sustainable Dataset Governance and Curation for Low-Resource Languages</i>	
Jenalea Rajab, Anuoluwapo Aremu, Everlyn Asiko Chimoto, Dale Dunbar, Graham Morrissey, Fadel Thior, Luandrie Potgieter, Jessica Ojo, Atnafu Lambebo Tonja, Wilhelmina NdapewaOnyothi Nekoto, Pelonomi Moilola, Jade Abbott, Vukosi Marivate and Benjamin Rosman	30763
<i>Theoretical Analysis of Hierarchical Language Recognition and Generation by Transformers without Positional Encoding</i>	
Daichi Hayakawa and Issei Sato	30777
<i>Less is More: Explainable and Efficient ICD Code Prediction with Clinical Entities</i>	
James C. Douglas, Yidong Gan, Ben Hachey and Jonathan K. Kummerfeld	30835
<i>Benchmarking LLMs and LLM-based Agents in Practical Vulnerability Detection for Code Repositories</i>	
Alperen Yildiz, Sin G Teo, Yiling Lou, Yebo Feng, Chong Wang and Dinil Mon Divakaran	30848
<i>Multi-Modality Expansion and Retention for LLMs through Parameter Merging and Decoupling</i>	
Junlin Li, Guodong DU, Jing Li, Sim Kuan Goh, Wenya Wang, Yequan Wang, Fangming Liu, Ho-Kin Tang, Saleh Alharbi, Daojing He and Min Zhang	30866
<i>Serial Lifelong Editing via Mixture of Knowledge Experts</i>	
YuJu Cheng, Yu-Chu Yu, Kai-Po Chang and Yu-Chiang Frank Wang	30888
<i>A Survey on Efficient Large Language Model Training: From Data-centric Perspectives</i>	
Junyu Luo, Bohan Wu, Xiao Luo, Zhiping Xiao, Yiqiao Jin, Rong-Cheng Tu, Nan Yin, Yifan Wang, Jingyang Yuan, Wei Ju and Ming Zhang	30904
<i>IMOL: Incomplete-Modality-Tolerant Learning for Multi-Domain Fake News Video Detection</i>	
Zhi Zeng, Jiaying Wu, Minnan Luo, Herun Wan, Xiangzheng Kong, Zihan Ma, Guang Dai and Qinghua Zheng	30921

<i>DDxTutor: Clinical Reasoning Tutoring System with Differential Diagnosis-Based Structured Reasoning</i>	
Qian Wu, Zheyao Gao, Longfei Gou and Qi Dou	30934
<i>SocialEval: Evaluating Social Intelligence of Large Language Models</i>	
Jinfeng Zhou, Yuxuan Chen, Yihan Shi, Xuanming Zhang, Leqi Lei, Yi Feng, Zexuan Xiong, Miao Yan, Xunzhi Wang, Yaru Cao, Jianing Yin, Shuai Wang, Quanyu Dai, Zhenhua Dong, Hongning Wang and Minlie Huang	30958
<i>Hidden in Plain Sight: Evaluation of the Deception Detection Capabilities of LLMs in Multimodal Settings</i>	
Md Messal Monem Miah, Adrita Anika, Xi Shi and Ruihong Huang	31013
<i>Analyzing and Mitigating Inconsistency in Discrete Speech Tokens for Neural Codec Language Models</i>	
Wenrui Liu, Zhifang Guo, Jin Xu, Yuanjun Lv, Yunfei Chu, Zemin Liu and Junyang Lin ..	31035
<i>PlanningArena: A Modular Benchmark for Multidimensional Evaluation of Planning and Tool Learning</i>	
Zihan Zheng, Tianle Cui, Chuwen Xie, Jiahui Pan, Qianglong Chen and Lewei He	31047
<i>FocusLLM: Precise Understanding of Long Context by Dynamic Condensing</i>	
Zhenyu Li, Yike Zhang, Tengyu Pan, Yutao Sun, Zhichao Duan, Junjie Fang, Rong Han, Zixuan Wang and Jianyong Wang	31087
<i>Negative Matters: Multi-Granularity Hard-Negative Synthesis and Anchor-Token-Aware Pooling for Enhanced Text Embeddings</i>	
Tengyu Pan, Zhichao Duan, Zhenyu Li, Bowen Dong, Ning Liu, Xiuxing Li and Jianyong Wang	31102
<i>GPT-4 as a Homework Tutor Can Improve Student Engagement and Learning Outcomes</i>	
Alessandro Vanzo, Sankalan Pal Chowdhury and Mrinmaya Sachan	31119
<i>Diffusion Models Through a Global Lens: Are They Culturally Inclusive?</i>	
Zahra Bayramli, Ayhan Suleymanzade, Na Min An, Huzama Ahmad, Eunsu Kim, Junyeong Park, James Thorne and Alice Oh	31137
<i>Efficient Safety Alignment of Large Language Models via Preference Re-ranking and Representation-based Reward Modeling</i>	
Deng Qiyuan, Xuefeng Bai, Kehai Chen, Yaowei Wang, Liqiang Nie and Min Zhang	31156
<i>English-based acoustic models perform well in the forced alignment of two English-based Pacific Creoles</i>	
Sam Passmore, Lila San Roque, Kirsty Gillespie, Saurabh Nath, Kira Davey, Keira Mullan, Tim Cawley, Jennifer Biggs, Rosey Billington, Bethwyn Evans, Nick Thieberger and Danielle Barth	31172
<i>Subtle Errors in Reasoning: Preference Learning via Error-injected Self-editing</i>	
Kaishuai Xu, Tiezheng YU, Wenjun Hou, Yi Cheng, Chak Tou Leong, Liangyou Li, Xin Jiang, Lifeng Shang, Qun Liu and Wenjie Li	31184
<i>Truth Knows No Language: Evaluating Truthfulness Beyond English</i>	
Blanca Calvo Figueras, Eneko Sagarzazu, Julen Etxaniz, Jeremy Barnes, Pablo Gamallo, Iria de-Dios-Flores and Rodrigo Agerri	31204
<i>Revisiting Compositional Generalization Capability of Large Language Models Considering Instruction Following Ability</i>	
Yusuke Sakai, Hidetaka Kamigaito and Taro Watanabe	31219

<i>Batayan: A Filipino NLP benchmark for evaluating Large Language Models</i>	
Jann Railey Montalan, Jimson Paulo Layacan, David Demitri Africa, Richell Isaiah S. Flores, Michael T. Lopez II, Theresa Denise Magsajo, Anjanette Cayabyab and William Chandra Tjhi .	31239
<i>HintsOfTruth: A Multimodal Checkworthiness Detection Dataset with Real and Synthetic Claims</i>	
Michiel Van Der Meer, Pavel Korshunov, Sébastien Marcel and Lonneke Van Der Plas . . .	31274
<i>CityNavAgent: Aerial Vision-and-Language Navigation with Hierarchical Semantic Planning and Global Memory</i>	
Weichen Zhang, Chen Gao, Shiquan Yu, Ruiying Peng, Baining Zhao, Qian Zhang, Jinqiang Cui, Xinlei Chen and Yong Li	31292
<i>It's Not a Walk in the Park! Challenges of Idiom Translation in Speech-to-text Systems</i>	
Iuliia Zaitova, Badr M. Abdullah, Wei Xue, Dietrich Klakow, Bernd Möbius and Tania Avgustina	31310
<i>PolyNarrative: A Multilingual, Multilabel, Multi-domain Dataset for Narrative Extraction from News Articles</i>	
Nikolaos Nikolaidis, Nicolas Stefanovitch, Purificação Silvano, Dimitar Iliyanov Dimitrov, Roman Yangarber, Nuno Guimarães, Elisa Sartori, Ion Androutsopoulos, Preslav Nakov, Giovanni Da San Martino and Jakub Piskorski	31323
<i>A Parameter-Efficient and Fine-Grained Prompt Learning for Vision-Language Models</i>	
Yongbin Guo, Shuzhen Li, Zhulin Liu, Tong Zhang and C.L.Philip Chen	31346
<i>Persona Dynamics: Unveiling the Impact of Persona Traits on Agents in Text-Based Games</i>	
Seungwon Lim, Seungbeen Lee, Dongjun Min and Youngjae Yu	31360
<i>SeedBench: A Multi-task Benchmark for Evaluating Large Language Models in Seed Science</i>	
Jie Ying, Zihong Chen, Zhefan Wang, Wanli Jiang, Chenyang Wang, Zhonghang Yuan, Haoyang Su, Huanjun Kong, Fan Yang and Nanqing Dong	31395
<i>-Stance: A Large-Scale Real World Dataset of Stances in Legal Argumentation</i>	
Ankita Gupta, Douglas Rice and Brendan O'Connor	31450
<i>Re³Syn: A Dependency-Based Data Synthesis Framework for Long-Context Post-training</i>	
Zhiyang Zhang, Ziqiang Liu, Huiming Wang, Renke Shan, Li Kuang, Lu Wang and De Wen Soh	31468
<i>Enabling Chatbots with Eyes and Ears: An Immersive Multimodal Conversation System for Dynamic Interactions</i>	
Jihyoung Jang, Minwook Bae, Minji Kim, Dilek Hakkani-Tür and Hyounghun Kim	31481
<i>Multimodal Coreference Resolution for Chinese Social Media Dialogues: Dataset and Benchmark Approach</i>	
Xingyu Li, Chen Gong and Guohong Fu	31513
<i>TACLR: A Scalable and Efficient Retrieval-based Method for Industrial Product Attribute Value Identification</i>	
Yindu Su, Huike Zou, Lin Sun, Ting Zhang, Haiyang Yang, Chen Li Yu, David Lo, Qingheng Zhang, Shuguang Han and Jufeng Chen	31526
<i>Theory of Mind in Large Language Models: Assessment and Enhancement</i>	
Ruirui Chen, Weifeng Jiang, Chengwei Qin and Cheston Tan	31539
<i>Completing A Systematic Review in Hours instead of Months with Interactive AI Agents</i>	
Rui Qiu, Shijie Chen, Yu Su, Po-Yin Yen and Han Wei Shen	31559

<i>CMHKF: Cross-Modality Heterogeneous Knowledge Fusion for Weakly Supervised Video Anomaly Detection</i>	
Guohua Wang, Shengping Song, Wuchun He and Yongsen Zheng	31594
<i>CLaSp: In-Context Layer Skip for Self-Speculative Decoding</i>	
Longze Chen, Renke Shan, Huiming Wang, Lu Wang, Ziqiang Liu, Run Luo, Jiawei Wang, Hamid Alinejad-Rokny and Min Yang	31608
<i>Teaching Text Agents to Learn Sequential Decision Making from Failure</i>	
Canasai Kruengkrai and Koichiro Yoshino.....	31619
<i>The Harmonic Structure of Information Contours</i>	
Eleftheria Tsipidi, Samuel Kiegeand, Franz Nowak, Tianyang Xu, Ethan Wilcox, Alex Warstadt, Ryan Cotterell and Mario Giulianelli.....	31636
<i>REAL-MM-RAG: A Real-World Multi-Modal Retrieval Benchmark</i>	
Navve Wasserman, Roi Pony, Oshri Naparstek, Adi Raz Goldfarb, Eli Schwartz, Udi Barzelay and Leonid Karlinsky	31660
<i>Only a Little to the Left: A Theory-grounded Measure of Political Bias in Large Language Models</i>	
Mats Faulborn, Indira Sen, Max Pellert, Andreas Spitz and David Garcia	31684
<i>LongSafety: Evaluating Long-Context Safety of Large Language Models</i>	
Yida Lu, Jiale Cheng, Zhexin Zhang, Shiyao Cui, Cunxiang Wang, Xiaotao Gu, Yuxiao Dong, Jie Tang, Hongning Wang and Minlie Huang.....	31705
<i>Exploiting Contextual Knowledge in LLMs through \mathcal{V}-usable Information based Layer Enhancement</i>	
Xiaowei Yuan, Zhao Yang, Ziyang Huang, Yequan Wang, Siqi Fan, Yiming Ju, Jun Zhao and Kang Liu.....	31726
<i>Unintended Harms of Value-Aligned LLMs: Psychological and Empirical Insights</i>	
Sooyung Choi, Jaehyeok Lee, Xiaoyuan Yi, Jing Yao, Xing Xie and JinYeong Bak	31742
<i>Maximal Matching Matters: Preventing Representation Collapse for Robust Cross-Modal Retrieval</i>	
Hani Alomari, Anushka Sivakumar, Andrew Zhang and Chris Thomas	31769
<i>The Noisy Path from Source to Citation: Measuring How Scholars Engage with Past Research</i>	
Hong Chen, Misha Teplitskiy and David Jurgens.....	31786
<i>MAPLE: Enhancing Review Generation with Multi-Aspect Prompt LEarning in Explainable Recommendation</i>	
Ching-Wen Yang, Zhi-Quan Feng, Ying-Jia Lin, Che Wei Chen, Kun-da Wu, Hao Xu, Yao Jui-Feng and Hung-Yu Kao.....	31803
<i>Separating Tongue from Thought: Activation Patching Reveals Language-Agnostic Concept Representations in Transformers</i>	
Clément Dumas, Chris Wendler, Veniamin Veselovsky, Giovanni Monea and Robert West.....	31822
<i>Behavioural vs. Representational Systematicity in End-to-End Models: An Opinionated Survey</i>	
Ivan Vegner, Sydelle De Souza, Valentin Forch, Martha Lewis and Leonidas A. A. Doumas.....	31842
<i>Dynamic Chunking and Selection for Reading Comprehension of Ultra-Long Context in Large Language Models</i>	
Boheng Sheng, Jiacheng Yao, Meicong Zhang and Guoxiu He	31857

<i>DualRAG: A Dual-Process Approach to Integrate Reasoning and Retrieval for Multi-Hop Question Answering</i>	Rong Cheng, Jinyi Liu, Yan Zheng, Fei Ni, Jiazhen Du, Hangyu Mao, Fuzheng Zhang, Bo Wang and Jianye Hao	31877
<i>Deliberate Reasoning in Language Models as Structure-Aware Planning with an Accurate World Model</i>	Siheng Xiong, Ali Payani, Yuan Yang and Faramarz Fekri	31900
<i>Refining Saliency-Aware Sparse Fine-Tuning Strategies for Language Models</i>	Xinxin Liu, Aaron Thomas, Cheng Zhang, Jianyi Cheng, Yiren Zhao and Xitong Gao	31932
<i>Efficient Many-Shot In-Context Learning with Dynamic Block-Sparse Attention</i>	Emily Xiao, Chin-Jou Li, Yilin Zhang, Graham Neubig and Amanda Bertsch	31946
<i>ScaleBiO: Scalable Bilevel Optimization for LLM Data Reweighting</i>	Rui Pan, Dylan Zhang, Hanning Zhang, Xingyuan Pan, Minrui Xu, Jipeng Zhang, Renjie Pi, Xiaoyu Wang and Tong Zhang	31959
<i>PKU-SafeRLHF: Towards Multi-Level Safety Alignment for LLMs with Human Preference</i>	Jiaming Ji, Donghai Hong, Borong Zhang, Boyuan Chen, Josef Dai, Boren Zheng, Tianyi Alex Qiu, Jiayi Zhou, Kaile Wang, Boxun Li, Sirui Han, Yike Guo and Yaodong Yang	31983
<i>What Happened in LLMs Layers when Trained for Fast vs. Slow Thinking: A Gradient Perspective</i>	Ming Li, Yanhong Li and Tianyi Zhou	32017
<i>Beyond Text Compression: Evaluating Tokenizers Across Scales</i>	Jonas F. Lotz, António V. Lopes, Stephan Peitz, Hendra Setiawan and Leonardo Emili	32155
<i>Emergent Abilities of Large Language Models under Continued Pre-training for Language Adaptation</i>	Ahmed Elhady, Eneko Agirre and Mikel Artetxe	32174
<i>R-Fairness: Assessing Fairness of Ranking in Subjective Data</i>	Lorenzo Balzotti, Donatella Firmani, Jerin George Mathew, Riccardo Torlone and Sihem Amer-Yahia	32187
<i>RePanda: Pandas-powered Tabular Verification and Reasoning</i>	Atoosa Chegini, Keivan Rezaei, Hamid Eghbalzadeh and Soheil Feizi	32200
<i>Towards Style Alignment in Cross-Cultural Translation</i>	Shreya Havaldar, Adam Stein, Eric Wong and Lyle Ungar	32213
<i>TiC-LM: A Web-Scale Benchmark for Time-Continual LLM Pretraining</i>	Jeffrey Li, Mohammadreza Armandpour, Seyed Iman Mirzadeh, Sachin Mehta, Vaishaal Shankar, Raviteja Vemulapalli, Samy Bengio, Oncel Tuzel, Mehrdad Farajtabar, Hadi Pouransari and Fartash Faghri	32231
<i>Entailed Between the Lines: Incorporating Implication into NLI</i>	Shreya Havaldar, Hamidreza Alvari, John Palowitch, Mohammad Javad Hosseini, Senaka Buthpitiya and Alex Fabrikant	32274
<i>Multi-Level Explanations for Generative Language Models</i>	Lucas Monteiro Paes, Dennis Wei, Hyo Jin Do, Hendrik Strobelt, Ronny Luss, Amit Dhurandhar, Manish Nagireddy, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Werner Geyer and Soumya Ghosh	32291
<i>A Multi-Agent Framework for Mitigating Dialect Biases in Privacy Policy Question-Answering Systems</i>	Dorđe Klisura, Astrid R Bernaga Torres, Anna Karen Gárate-Escamilla, Rajesh Roshan Biswal, Ke Yang, Hilal Pataci and Anthony Rios	32318

<i>Low-Bit Quantization Favors Undertrained LLMs</i>	
Xu Ouyang, Tao Ge, Thomas Hartvigsen, Zhisong Zhang, Haitao Mi and Dong Yu	32338
<i>Enhancing User-Controlled Text-to-Image Generation with Layout-Aware Personalization</i>	
Hongliang Luo and Wei Xi	32349
<i>LETS-C: Leveraging Text Embedding for Time Series Classification</i>	
Rachneet Kaur, Zhen Zeng, Tucker Balch and Manuela Veloso	32365
<i>UrbanVideo-Bench: Benchmarking Vision-Language Models on Embodied Intelligence with Video Data in Urban Spaces</i>	
Baining Zhao, Jianjie Fang, Zichao Dai, Ziyou Wang, Jirong Zha, Weichen Zhang, Chen Gao, Yue Wang, Jinqiang Cui, Xinlei Chen and Yong Li	32400
<i>HELIOS: Harmonizing Early Fusion, Late Fusion, and LLM Reasoning for Multi-Granular Table-Text Retrieval</i>	
Sungho Park, Joohyung Yun, Jongwuk Lee and Wook-Shin Han	32424
<i>ONEBench to Test Them All: Sample-Level Benchmarking Over Open-Ended Capabilities</i>	
Adhiraj Ghosh, Sebastian Dziadzio, Ameya Prabhu, Vishaal Udandara, Samuel Albanie and Matthias Bethge	32445
<i>La Leaderboard: A Large Language Model Leaderboard for Spanish Varieties and Languages of Spain and Latin America</i>	
María Grandury, Javier Aula-Blasco, Júlia Falcão, Clémentine Fourier, Miguel González Saiz, Gonzalo Martínez, Gonzalo Santamaria Gomez, Rodrigo Agerri, Nuria Aldama García, Luis Chiruzzo, Javier Conde, Helena Gomez Adorno, Marta Guerrero Nieto, Guido Ivetta, Natàlia López Fuertes, Flor Miriam Plaza-del-Arco, María-Teresa Martín-Valdivia, Helena Montoro Zamorano, Carmen Muñoz Sanz, Pedro Reviriego, Leire Rosado Plaza, Alejandro Vaca Serrano, Estrella Vallecillo-Rodríguez, Jorge Vallego and Irune Zubiaga	32482
<i>Why Prompt Design Matters and Works: A Complexity Analysis of Prompt Search Space in LLMs</i>	
Xiang Zhang, Juntao Cao, Chenyu You and Dujian Ding	32525
<i>Energy Considerations of Large Language Model Inference and Efficiency Optimizations</i>	
Jared Fernandez, Clara Na, Vashisth Tiwari, Yonatan Bisk, Sasha Luccioni and Emma Strubell	32556
<i>Optimizing Pre-Training Data Mixtures with Mixtures of Data Expert Models</i>	
Lior Belenki, Alekh Agarwal, Tianze Shi and Kristina Toutanova	32570
<i>BFS-Prover: Scalable Best-First Tree Search for LLM-based Automatic Theorem Proving</i>	
Ran Xin, Chenguang Xi, Jie Yang, Feng Chen, Hang Wu, Xia Xiao, Yifan Sun, Shen Zheng and Ming Ding	32588
<i>Magnet: Multi-turn Tool-use Data Synthesis and Distillation via Graph Translation</i>	
Fan Yin, Zifeng Wang, I-Hung Hsu, Jun Yan, Ke Jiang, Yanfei Chen, Jindong Gu, Long Le, Kai-Wei Chang, Chen-Yu Lee, Hamid Palangi and Tomas Pfister	32600
<i>Logic-Regularized Verifier Elicits Reasoning from LLMs</i>	
Xinyu Wang, Changzhi Sun, Lian Cheng, Yuanbin Wu, Dell Zhang, Xiaoling Wang and Xuelong Li	32617
<i>Squeezed Attention: Accelerating Long Context Length LLM Inference</i>	
Coleman Richard Charles Hooper, Sehoon Kim, Hiva Mohammadzadeh, Monishwaran Mahe- swaran, Sebastian Zhao, June Paik, Michael W. Mahoney, Kurt Keutzer and Amir Gholami	32631

<i>LangMark: A Multilingual Dataset for Automatic Post-Editing</i>	
Diego Velazquez, Mikaela Grace, Konstantinos Karageorgos, Lawrence Carin, Aaron Schliem, Dimitrios Zaikis and Roger Wechsler	32653
<i>Neural Parameter Search for Slimmer Fine-Tuned Models and Better Transfer</i>	
Guodong DU, Zitao Fang, Jing Li, Junlin Li, Runhua Jiang, Shuyang Yu, Yifei Guo, Yangneng Chen, Sim Kuan Goh, Ho-Kin Tang, Daojing He, Honghai Liu and Min Zhang	32668
<i>Merge Hijacking: Backdoor Attacks to Model Merging of Large Language Models</i>	
Zenghui Yuan, Yangming Xu, Jiawen Shi, Pan Zhou and Lichao Sun	32688
<i>Where Are We? Evaluating LLM Performance on African Languages</i>	
Ife Adebara, Hawau Olamide Toyin, Nahom Tesfu Ghebremichael, AbdelRahim A. Elmadany and Muhammad Abdul-Mageed	32704
<i>Beyond Output Matching: Bidirectional Alignment for Enhanced In-Context Learning</i>	
Chengwei Qin, Wenhan Xia, Fangkai Jiao, Chen Chen, Yuchen Hu, Bosheng Ding, Ruirui Chen and Shafiq Joty	32732
<i>CiteEval: Principle-Driven Citation Evaluation for Source Attribution</i>	
Yumo Xu, Peng Qi, Jifan Chen, Kunlun Liu, Rujun Han, Lan Liu, Bonan Min, Vittorio Castelli, Arshit Gupta and Zhiguo Wang	32759
<i>HiAgent: Hierarchical Working Memory Management for Solving Long-Horizon Agent Tasks with Large Language Model</i>	
Mengkang Hu, Tianxing Chen, Qiguang Chen, Yao Mu, Wenqi Shao and Ping Luo	32779
<i>EducationQ: Evaluating LLMs' Teaching Capabilities Through Multi-Agent Dialogue Framework</i>	
Yao Shi, Rongkeng Liang and Yong Xu	32799
<i>KRISTEVA: Close Reading as a Novel Task for Benchmarking Interpretive Reasoning</i>	
Peiqi Sui, Juan Diego Rodriguez, Philippe Laban, J. Dean Murphy, Joseph P. Dexter, Richard Jean So, Samuel Baker and Primit Chaudhuri	32829
<i>Efficient Domain Continual pretraining by Mitigating the Stability Gap</i>	
Yiduo Guo, Jie Fu, Huishuai Zhang and Dongyan Zhao	32850
<i>Palm: A Culturally Inclusive and Linguistically Diverse Dataset for Arabic LLMs</i>	
Fakhraddin Alwajih, Abdellah El Mekki, Samar Mohamed Magdy, AbdelRahim A. Elmadany, Omer Nacar, El Moatez Billah Nagoudi, Reem Abdel-Salam, Hanin Atwany, Youssef Nafea, Abdulfattah Mohammed Yahya, Rahaf Alhamouri, Hamzah A. Alsayadi, Hiba Zayed, Sara Shatnawi, Serry Sibae, Yasir Ech-chammakhy, Walid Al-Dhabyani, Marwa Mohamed Ali, Imen Jarraya, Ahmed Oumar El-Shangiti, Aisha Alraeesi, Mohammed Anwar AL-Ghrawi, Abdulrahman S. Al-Batati, Elgizouli Mohamed, Noha Taha Elgindi, Muhammed Saeed, Houdaifa Atou, Issam Ait Yahia, Abdelhak Bouayad, Mohammed Machrouh, Amal Makouar, Dania Alkawi, Mukhtar Mohamed, Safaa Taher Abdelfadil, Amine Ziad Ounnoughene, Anfel Rouabhia, Rwa Assi, Ahmed Sorkatti, Mohamedou Cheikh Tourad, Anis Koubaa, Ismail Berrada, Mustafa Jarrar, Shady Shehata and Muhammad Abdul-Mageed ..	32871
<i>NewsInterview: a Dataset and a Playground to Evaluate LLMs' Grounding Gap via Informational Interviews</i>	
Alexander Spangher, Michael Lu, Sriya Kalyan, Hyundong Justin Cho, Tenghao Huang, Weiyan Shi and Jonathan May	32895
<i>CFBench: A Comprehensive Constraints-Following Benchmark for LLMs</i>	
Tao Zhang, ChengLin Zhu, Yanjun Shen, Wenjing Luo, Yan Zhang, Hao Liang, Tao Zhang, Fan Yang, Mingan Lin, Yujing Qiao, Weipeng Chen, Bin Cui, Wentao Zhang and Zenan Zhou	32926

<i>Towards Building Large Scale Datasets and State-of-the-Art Automatic Speech Translation Systems for 14 Indian Languages</i>	
Ashwin Sankar, Sparsh Jain, Nikhil Narasimhan, Devilal Choudhary, Dhairya Suman, Mohammed Safi Ur Rahman Khan, Anoop Kunchukuttan, Mitesh M Khapra and Raj Dabre	32945
<i>CoRe-MMRAG: Cross-Source Knowledge Reconciliation for Multimodal RAG</i>	
Yang Tian, Fan Liu, Jingyuan Zhang, V. W., Yupeng Hu and Liqiang Nie	32967
<i>Mapping 1,000+ Language Models via the Log-Likelihood Vector</i>	
Momose Oyama, Hiroaki Yamagiwa, Yusuke Takase and Hidetoshi Shimodaira	32983
<i>ConsistencyChecker: Tree-based Evaluation of LLM Generalization Capabilities</i>	
Zhaochen Hong, Haofei Yu and Jiaxuan You	33039
<i>Robust Estimation of Population-Level Effects in Repeated-Measures NLP Experimental Designs</i>	
Alejandro Benito-Santos, Adrian Ghajari and Víctor Fresno	33076
<i>FactBench: A Dynamic Benchmark for In-the-Wild Language Model Factuality Evaluation</i>	
Farima Fatahi Bayat, Lechen Zhang, Sheza Munir and Lu Wang	33090
<i>Training-free LLM Merging for Multi-task Learning</i>	
Zichuan Fu, Xian Wu, Yejing Wang, Wanyu Wang, Shanshan Ye, Hongzhi Yin, Yi Chang, Yefeng Zheng and Xiangyu Zhao	33111
<i>Inferring from Logits: Exploring Best Practices for Decoding-Free Generative Candidate Selection</i>	
Mingyu Derek Ma, Yanna Ding, Zijie Huang, Jianxi Gao, Yizhou Sun and Wei Wang	33125
<i>Comparison-based Active Preference Learning for Multi-dimensional Personalization</i>	
Minhyeon Oh, Seungjoon Lee and Jungseul Ok	33145
<i>OpenCoder: The Open Cookbook for Top-Tier Code Large Language Models</i>	
Siming Huang, Tianhao Cheng, Jason Klein Liu, Weidi Xu, Jiaran Hao, Liuyihan Song, Yang Xu, Jian Yang, Jiaheng Liu, Chenchen Zhang, Linzheng Chai, Ruifeng Yuan, Xianzhen Luo, Qiufeng Wang, YuanTao Fan, Qingfu Zhu, Zhaoxiang Zhang, Yang Gao, Jie Fu, Qian Liu, Houyi Li, Ge Zhang, Yuan Qi, Xu Yinghui, Wei Chu and Zili Wang	33167
<i>LlamaDuo: LLMops Pipeline for Seamless Migration from Service LLMs to Small-Scale Local LLMs</i>	
Chansung Park, Juyong Jiang, Fan Wang, Sayak Paul and Jing Tang	33194
<i>AmbiK: Dataset of Ambiguous Tasks in Kitchen Environment</i>	
Anastasia Ivanova, Bakaeva Eva, Zoya Volovikova, Alexey Kovalev and Aleksandr Panov	33216
<i>SocialCC: Interactive Evaluation for Cultural Competence in Language Agents</i>	
Jincenzi Wu, Jianxun Lian, Dingdong Wang and Helen M. Meng	33242
<i>Scalable Vision Language Model Training via High Quality Data Curation</i>	
Hongyuan Dong, Zijian Kang, Weijie Yin, LiangXiao LiangXiao, ChaoFeng ChaoFeng and Ran Jiao	33272
<i>GRAM: Generative Recommendation via Semantic-aware Multi-granular Late Fusion</i>	
Sunkyoung Lee, Minjin Choi, Eunseong Choi, Hye-young Kim and Jongwuk Lee	33294
<i>Towards Economical Inference: Enabling DeepSeek’s Multi-Head Latent Attention in Any Transformer-based LLMs</i>	
Tao Ji, Bin Guo, Yuanbin Wu, Qipeng Guo, Shenlixing Shenlixing, Chenzhan Chenzhan, Xipeng Qiu, Qi Zhang and Tao Gui	33313

<i>TETRIS: Optimal Draft Token Selection for Batch Speculative Decoding</i> Zhaoxuan Wu, Zijian Zhou, Arun Verma, Alok Prakash, Daniela Rus and Bryan Kian Hsiang Low	33329
<i>Introducing Verification Task of Set Consistency with Set-Consistency Energy Networks</i> Mooho Song, Hye Ryung Son and Jay-Yoon Lee	33346
<i>Language Models can Subtly Deceive Without Lying: A Case Study on Strategic Phrasing in Legislation</i> Atharvan Dogra, Krishna Pillutla, Ameet Deshpande, Ananya B. Sai, John J Nay, Tanmay Rajpurohit, Ashwin Kalyan and Balaraman Ravindran	33367
<i>AfroCS-xs: Creating a Compact, High-Quality, Human-Validated Code-Switched Dataset for African Languages</i> Kayode Olaleye, Arturo Oncevay, Mathieu Sibue, Nombuyiselo Zondi, Michelle Terblanche, Sibongile Mapikitla, Richard Lastrucci, Charese Smiley and Vukosi Marivate	33391
<i>Just Go Parallel: Improving the Multilingual Capabilities of Large Language Models</i> Muhammad Reza Qorib, Junyi Li and Hwee Tou Ng	33411
<i>Design Choices for Extending the Context Length of Visual Language Models</i> Mukai Li, Lei Li, Shansan Gong and Qi Liu	33425