ACL 2025

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

Findings of the Association for Computational Linguistics: ACL 2025

July 27 - August 1, 2025

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

Diamond

Platinum

Gold

Silver

Bronze

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 979-8-89176-256-5

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 Mršic** 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 Blloshmi** 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 **Qiongkai 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 techical 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 Qiongkai 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 Storks 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 Storks, 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

Yunyao Li, Adobe
George Rehm, DFKI

Best Paper Committee Chairs

Rada Mihalcea, University of Michighan
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 Doğruöz, 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, Niranjan 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, Xinchi 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, Tosho 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, Jiaxin 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 Kamalloo, Hidetaka Kamigaito, Masahiro Kaneko, Dongyeop Kang, Diptesh Kanojia, Alina Karakanta, Börje F. Karlsson, Pei Ke, Yova Kementchedjhieva, 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 Krahmer, 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 Prasojo, Emily Prud'hommeaux, Danish Pruthi

Tao Qi, Chen Qian, Cheng Qian, Tieyun Qian, Libo Qin, Yanxia Qin, Chen Qiu, Han Qiu, Haoyi Qiu, Jielin 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 Rekabsaz, 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 Vasan 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, Qiongkai 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, Jiaxuan 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, Jiaxin 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 Adlakha, 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 Ahamed, 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 Akbartajari, Mohammad Akbari, Aditya Akella, Benjamin Akera, Mahmud Elahi Akhter, Berfin Aktas, Syeda Sabrina Akter, Hend Al-khalifa, Hadeel Al-negheimish, Nora Al-twairesh, Jesujoba Oluwadara 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, Felermino D. M. A. Ali, Iqra Ali, Sarwan Ali, Ilseyar Alimova, Pegah Alipoormolabashi, 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, Kaikai An, Na Min An, Shengnan An, Wenbin An, Zhenwei An, Aliki 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 Aßenmacher

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, Bobobolu, 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 Bouchekif, 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, Zheye 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 Bousso 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 Ducel, 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, Ghebriout 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 Eremeev, 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 Ferianc, 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 Subrahmanyan 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, Janghoon 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 Hasibi, Islam A. Hassan, Michael Hassid, Nabil Hathout, Hans Ole Hatzel, Shreya Havaldar, 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, Jiangen 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, Bhathiya 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 Moshiul 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 Hätty

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, Ziyan 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 Kargaran, Priyanka Kargupta, Akbar Karimi, Sanjeev Kumar Karn, Zohar Karnin, Marzena Karpinska, Fakhri Karray, Pradeep Karuturi, Pranav Kasela, Gautam Siddharth Kashyap, Gjergji Kasneci, Surya Kasturi, Tushar Kataria, Shunya Kato, Yoshihide Kato, Gregorios A Katsios, Ayush Kaushal, Muhammad Dehan Al Kautsar, Parsa Kavehzadeh, Pride Kavumba, 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, Baber Khalid, Muhammad Khalid, Osama Khalid, Mohammed Khalilia, Aleem Khan, Latifur Khan, Mohammed Safi Ur Rahman Khan, Salman Khan, Sujay Khandagale, Vivek Khetan, Zahra Khodagholi, Md Tawkat Islam Khondaker, Ashiqur R. Khudabukhsh, Urja Khurana, Varun Khurana, Dayeon Ki, Johannes Kiesel, Zlata Kikteva, Jihyung Kil, Halil Kilicoglu, Bugeun Kim, Byeonggeun Kim, Chae Won Kim, Chaeeun 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, Youna 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 Koupaee, 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, Keito Kudo, Sachit Kuhar, Adithya Kulkarni, Atharva Kulkarni, Atharva Kulkarni, Mayank Kulkarni, Anoop Kumar, Bhargava Kumar, Deepak Kumar, Karun Kumar, Nalin Kumar, Nitesh Kumar, Rishabh Kumar, Sachin Kumar, Sandeep Kumar, Shanu Kumar, Shashi Kumar, Sonal Kumar, Sujay S Kumar, Vishwajeet Kumar, Vivek Kumar, Yaman Kumar, Tharindu Sandaruwan Kumarage, Sadhana Kumaravel, Gitanjali Kumari, Tuhin Kundu, Po-nien Kung, Jenny Kunz, Olli Kuparinen, Yury Kuratov, Murathan Kurfalı, Kemal Kurniawan, Guy Kushilevitz, Mucahid Kutlu, Gleb Kuzmin, Andrey Kuznetsov, Ilia 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, Yuhang 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, Gyeongeun 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, Jiaxi Li, Jiaxuan 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, Heuiseok 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, Gangao 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, Jakub 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 Mahbub, 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, Ilia 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 Matusevych, 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 Miletić, 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 Nalisnick, 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, Yinhua Piao, Massimo Piccardi, Janet B. Pierrehumbert, Michał Pietruszka, Reshmi Gopalakrishna Pillai, Mathis Pink, Ramakrishna Pinninti, Bartosz Piotrowski, Andrew Piper, Hugo Pitorro, Priya Pitre, Lidia Pivovarova, 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, Mohammadreza Pourreza, Manuel Prada-corral, 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, Jiaxing 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, Shafin Rahman, Hossein A. Rahmani, Nishat Raihan, Tina Raissi, Lisa Raithel, Chahat Raj, Sahar Rajabi, Sara Rajaee, 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 Rücker

Ramaneswaran S, Jon Saad-falcon, Ashish Sabharwal, Ahmed Sabir, Sahand Sabour, Rachneet Singh Sachdeva, Kaushik Ram Sadagopan, Abdelrahman Boda Sadallah, Mobashir Sadat, Amir Saeidi, 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, Jivnesh 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 Saparina, 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 Schüz, 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, Jiaxin 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, Kiyoaki 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, Bokyung 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 Stammbach, 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 Uthayasooriyar, 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 Velldal, 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 Jiyoung 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, Xiongxiao 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, Jiaxi Yang, Jiaxi 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, Xinchen 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, Xinchen 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, Hyeongu 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 Zaporojets, Alessandra Zarcone, 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, Gaifan 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, Wenyuan 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, Xinping 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 Ziems, 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

Mirella 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. Mirella 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

Explicit vs. Implicit: Investigating Social Bias in Large Language Models through Self-Reflection
Yachao Zhao, Bo Wang, Yan Wang, Dongming Zhao, Ruifang He and Yuexian Hou 1

Beyond Perception: Evaluating Abstract Visual Reasoning through Multi-Stage Task
Yanbei Jiang, Yihao Ding, Chao Lei, Jiayang Ao, Jey Han Lau and Krista A. Ehinger 13

How Numerical Precision Affects Arithmetical Reasoning Capabilities of LLMs
Guhao Feng, Kai Yang, Yuntian Gu, Xinyue Ai, Shengjie Luo, Jiacheng Sun, Di He, Zhenguo Li and Liwei Wang 46

Diversifying the Expert Knowledge for Task-Agnostic Pruning in Sparse Mixture-of-Experts
Zeliang Zhang, Xiaodong Liu, Hao Cheng, Chenliang Xu and Jianfeng Gao 86

A Persona-Aware LLM-Enhanced Framework for Multi-Session Personalized Dialogue Generation
Dongshuo Liu, Zhijing Wu, Dandan Song and Heyan Huang 103

Exploring In-Image Machine Translation with Real-World Background
Yanzhi Tian, Zeming Liu, Zhengyang Liu and Yuhang Guo 124

BayesKD: Bayesian Knowledge Distillation for Compact LLMs in Constrained Fine-tuning Scenarios
Wei Li, Lujun Li, Mark G. Lee, Shengjie Sun, Lei Zhang, Wei Xue and Yike Guo 138

GOLFer: Smaller LMs-Generated Documents Hallucination Filter & Combiner for Query Expansion in Information Retrieval
Lingyuan Liu and Mengxiang Zhang 153

Exp4Fuse: A Rank Fusion Framework for Enhanced Sparse Retrieval using Large Language Model-based Query Expansion
Lingyuan Liu and Mengxiang Zhang 163

Emo Pillars: Knowledge Distillation to Support Fine-Grained Context-Aware and Context-Less Emotion Classification
Alexander Shvets 174

Multi-Prompting Decoder Helps Better Language Understanding
Zifeng Cheng, Zhaoling Chen, Zhiwei Jiang, Yafeng Yin, Cong Wang, Shiping Ge and Qing Gu 192

Visual Cues Enhance Predictive Turn-Taking for Two-Party Human Interaction
Sam O'Connor Russell and Naomi Harte 209

The Right Time Matters: Data Arrangement Affects Zero-Shot Generalization in Instruction Tuning
Bingxiang He, Ning Ding, Cheng Qian, Jia Deng, Ganqu Cui, Lifan Yuan, Haiwen Hong, Huang Gao, Longtao Huang, Hui Xue, Huimin Chen, Zhiyuan Liu and Maosong Sun 222

M^3FinMeeting: A Multilingual, Multi-Sector, and Multi-Task Financial Meeting Understanding Evaluation Dataset
Jie Zhu, Junhui Li, Yalong Wen, Xiandong Li, Lifan Guo and Feng Chen 244

ODDA: An OODA-Driven Diverse Data Augmentation Framework for Low-Resource Relation Extraction
Yijie Zhong, Yunfan Gao, Xiaolian Zhang and Haofen Wang 267

Detecting and Mitigating Challenges in Zero-Shot Video Summarization with Video LLMs
Luca Cagliero, Lorenzo Vaiani, Eliana Pastor, Alkis Koudounas, Elena Baralis, Vittorio Mazzia, Sandro Pollastrini, Thomas Gueudre, Manuel Giollo, Daniele Amberti and Yue Wu 286

Entity Framing and Role Portrayal in the News
Tarek Mahmoud, Zhuohan Xie, Dimitar Iliyanov Dimitrov, Nikolaos Nikolaidis, Purificação Silvano, Roman Yangarber, Shivam Sharma, Elisa Sartori, Nicolas Stefanovitch, Giovanni Da San Martino, Jakub Piskorski and Preslav Nakov 302

Derailer-Rerailer: Adaptive Verification for Efficient and Reliable Language Model Reasoning
Guangya Wan, Yuqi Wu, Hao Wang, Shengming Zhao, Jie Chen and Sheng Li 327

Leveraging Large Language Models for Conversational Multi-Doc Question Answering: The First Place of WSDM Cup 2024
Yiming Li and Zhao Zhang 349

TreeRAG: Unleashing the Power of Hierarchical Storage for Enhanced Knowledge Retrieval in Long Documents
Wenyu Tao, Xiaofen Xing, Yirong Chen, Linyi Huang and Xiangmin Xu 356

Attention with Dependency Parsing Augmentation for Fine-Grained Attribution
Qiang Ding, Lvzhou Luo, Yixuan Cao and Ping Luo 372

ASTRO: Automatic Strategy Optimization For Non-Cooperative Dialogues
Yikuan Hu, Chen Huang and Wenqiang Lei 388

Defensive Prompt Patch: A Robust and Generalizable Defense of Large Language Models against Jailbreak Attacks
Chen Xiong, Xiangyu Qi, Pin-Yu Chen and Tsung-Yi Ho 409

GUM-SAGE: A Novel Dataset and Approach for Graded Entity Salience Prediction
Jessica Lin and Amir Zeldes 438

Verifying the Steps of Deductive Reasoning Chains
Zacchary Sadeddine and Fabian M. Suchanek 456

Translate With Care: Addressing Gender Bias, Neutrality, and Reasoning in Large Language Model Translations
Pardis Sadat Zahraei and Ali Emami 476

Utilizing Semantic Textual Similarity for Clinical Survey Data Feature Selection
Benjamin C Warner, Ziqi Xu, Simon Haroutounian, Thomas Kannampallil and Chenyan Lu . 502

Distance between Relevant Information Pieces Causes Bias in Long-Context LLMs
Runchu Tian, Yanghao Li, Yuepeng Fu, Siyang Deng, Qinyu Luo, Cheng Qian, Shuo Wang, Xin Cong, Zhong Zhang, Yesai Wu, Yankai Lin, Huadong Wang and Xiaojiang Liu 521

Variable Layerwise Quantization: A Simple and Effective Approach to Quantize LLMs
Razvan-Gabriel Dumitru, Vikas Yadav, Rishabh Maheshwary, Paul Ioan Clotan, Sathwik Tejaswi Madhusudhan and Mihai Surdeanu 534

Why Are Positional Encodings Nonessential for Deep Autoregressive Transformers? A Petroglyph Revisited
Kazuki Irie 551

CRPO: Confidence-Reward Driven Preference Optimization for Machine Translation
Guofeng Cui, Pichao Wang, Yang Liu, Zemian Ke, Zhu Liu and Vimal Bhat 560

Talking Point based Ideological Discourse Analysis in News Events
Nishanth Sridhar Nakshatri, Nikhil Mehta, Siyi Liu, Sihao Chen, Daniel Hopkins, Dan Roth and Dan Goldwasser 575

FlashBack: Efficient Retrieval-Augmented Language Modeling for Fast Inference
Runheng Liu, Xingchen Xiao, Heyan Huang, Zewen Chi and Zhijing Wu 595

CMQCIC-Bench: A Chinese Benchmark for Evaluating Large Language Models in Medical Quality Control Indicator Calculation
Guangya Yu, Yanhao Li, Zongying Jiang, Yuxiong Jin, Li Dai, Yupian Lin, Ruihui Hou, Weiyan Zhang, Yongqi Fan, Qi Ye, Jingping Liu and Tong Ruan 609

ConKE: Conceptualization-Augmented Knowledge Editing in Large Language Models for Commonsense Reasoning
Liyu Zhang, Weiqi Wang, Tianqing Fang and Yangqiu Song 627

Exploring Multi-Modal Data with Tool-Augmented LLM Agents for Precise Causal Discovery
ChengAo Shen, Zhengzhang Chen, Dongsheng Luo, Dongkuan Xu, Haifeng Chen and Jingchao Ni 636

PARSQL: Enhancing Text-to-SQL through SQL Parsing and Reasoning
Yaxun Dai, Haiqin Yang, Mou Hao and Pingfu Chao 661

Probing the Geometry of Truth: Consistency and Generalization of Truth Directions in LLMs Across Logical Transformations and Question Answering Tasks
Yuntai Bao, Xuhong Zhang, Tianyu Du, Xinkui Zhao, Zhengwen Feng, Hao Peng and Jianwei Yin 682

Comparing Bad Apples to Good Oranges Aligning Large Language Models via Joint Preference Optimization
Hritik Bansal, Ashima Suvarna, Gantavya Bhatt, Nanyun Peng, Kai-Wei Chang and Aditya Grover 701

TestAgent: An Adaptive and Intelligent Expert for Human Assessment
Junhao Yu, Yan Zhuang, Yuxuan Sun, Weibo Gao, Qi Liu, Mingyue Cheng, Zhenya Huang and Enhong Chen 724

SPICA: Retrieving Scenarios for Pluralistic In-Context Alignment
Quan Ze Chen, Kevin Feng, Chan Young Park and Amy X Zhang 748

First-Step Advantage: Importance of Starting Right in Multi-Step Math Reasoning
Kushal Jain, Moritz Miller, Niket Tandon and Kumar Shridhar 766

Evaluating Instructively Generated Statement by Large Language Models for Directional Event Causality Identification
Wei Xiang, Chuanhong Zhan, Qing Zhang and Bang Wang 779

CoinMath: Harnessing the Power of Coding Instruction for Math LLM
Chengwei Wei, Bin Wang, Jung-jae Kim, Guimei Liu and Nancy F. Chen 786

Profiling News Media for Factuality and Bias Using LLMs and the Fact-Checking Methodology of Human Experts
Zain Muhammad Mujahid, Dilshod Azizov, Maha Tufail Agro and Preslav Nakov 798

Structured Discourse Representation for Factual Consistency Verification
Kun Zhang, Oana Balalau and Ioana Manolescu 820

SHARP: Unlocking Interactive Hallucination via Stance Transfer in Role-Playing LLMs
Chuyi Kong, Ziyang Luo, Hongzhan Lin, Zhiyuan Fan, Yaxin Fan, Yuxi Sun and Jing Ma . . . 839

Understanding the Gap: an Analysis of Research Collaborations in NLP and Language Documentation
Luke Gessler, Alexis Palmer and Katharina Von Der Wense 867

PersonaBench: Evaluating AI Models on Understanding Personal Information through Accessing (Synthetic) Private User Data
Juntao Tan, Liangwei Yang, Zuxin Liu, Zhiwei Liu, Rithesh R N, Tulika Manoj Awalgaonkar, Jianguo Zhang, Weiran Yao, Ming Zhu, Shirley Kokane, Silvio Savarese, Huan Wang, Caiming Xiong and Shelby Heinecke 878

Leveraging Variation Theory in Counterfactual Data Augmentation for Optimized Active Learning
Simret A Gebreegziabher, Kuangshi Ai, Zheng Zhang, Elena Glassman and Toby Jia-Jun Li . 894

ORBIT: Cost-Effective Dataset Curation for Large Language Model Domain Adaptation with an Astronomy Case Study
Eric Modesitt, Ke Yang, Spencer Hulsey, Xin Liu, ChengXiang Zhai and Volodymyr Kindratenko 907

Serial Position Effects of Large Language Models
Xiaobo Guo and Soroush Vosoughi 927

scRAG: Hybrid Retrieval-Augmented Generation for LLM-based Cross-Tissue Single-Cell Annotation
Zhiyin Yu, Chao Zheng, Chong Chen, Xian-Sheng Hua and Xiao Luo 954

Can Large Language Models Address Open-Target Stance Detection?
Abu Ubaida Akash, Ahmed Fahmy and Amine Trabelsi 971

Improve Language Model and Brain Alignment via Associative Memory
Congchi Yin, Yongpeng Zhang, Xuyun Wen and Piji Li 986

Towards Reliable Large Audio Language Model
Ziyang Ma, Xiquan Li, Yakun Song, Wenxi Chen, Chenpeng Du, Jian Wu, Yuanzhe Chen, Zhuo Chen, Yuping Wang, Yuxuan Wang and Xie Chen 1000

Large Vocabulary Size Improves Large Language Models
Sho Takase, Ryokan Ri, Shun Kiyono and Takuya Kato 1015

MUSE: A Multimodal Conversational Recommendation Dataset with Scenario-Grounded User Profiles
Zihan Wang, Xiaocui Yang, YongKang Liu, Shi Feng, Daling Wang and Yifei Zhang 1027

Machine Translation Models are Zero-Shot Detectors of Translation Direction
Michelle Wastl, Jannis Vamvas and Rico Sennrich 1054

Do Robot Snakes Dream like Electric Sheep? Investigating the Effects of Architectural Inductive Biases on Hallucination
Jerry Huang, Prasanna Parthasarathi, Mehdi Rezagholizadeh, Boxing Chen and Sarath Chandar 1075

GenTool: Enhancing Tool Generalization in Language Models through Zero-to-One and Weak-to-Strong Simulation
Jie He, Jennifer Neville, Mengting Wan, Longqi Yang, Hui Liu, Xiaofeng Xu, Xia Song, Jeff Z. Pan and Pei Zhou 1097

SWE-Fixer: Training Open-Source LLMs for Effective and Efficient GitHub Issue Resolution
Chengxing Xie, Bowen Li, Chang Gao, He Du, Wai Lam, Difan Zou and Kai Chen 1123

GlyphPattern: An Abstract Pattern Recognition for Vision-Language Models
Zixuan Wu, Yoolim Kim and Carolyn Jane Anderson 1140

FitCF: A Framework for Automatic Feature Importance-guided Counterfactual Example Generation
Qianli Wang, Nils Feldhus, Simon Ostermann, Luis Felipe Villa-Arenas, Sebastian Möller and Vera Schmitt 1176

From Misleading Queries to Accurate Answers: A Three-Stage Fine-Tuning Method for LLMs
Guocong Li, Weize Liu, Yihang Wu, Ping Wang, Shuaihan Huang, Hongxia Xu and Jian Wu 1192

Separate the Wheat from the Chaff: A Post-Hoc Approach to Safety Re-Alignment for Fine-Tuned Language Models
Di Wu, Xin Lu, Yanyan Zhao and Bing Qin 1210

Nuclear Deployed!: Analyzing Catastrophic Risks in Decision-making of Autonomous LLM Agents
Rongwu Xu, Xiaojian Li, Shuo Chen and Wei Xu 1226

MoRE: A Mixture of Low-Rank Experts for Adaptive Multi-Task Learning
Dacao Zhang, Kun Zhang, Shimao Chu, Le Wu, Xin Li and Si Wei 1311

Lunar Twins: We Choose to Go to the Moon with Large Language Models
Xin-Yu Xiao, Yalei Liu, Xiangyu Liu, Zengrui Li, Erwei Yin and Qianchen Xia 1325

SPHERE: An Evaluation Card for Human-AI Systems
Dora Zhao, Qianou Ma, Xinran Zhao, Chenglei Si, Chenyang Yang, Ryan Louie, Ehud Reiter, Diyi Yang and Tongshuang Wu 1340

Data-Centric Improvements for Enhancing Multi-Modal Understanding in Spoken Conversation Modeling
Maximillian Chen, Ruoxi Sun and Sercan O Arik 1366

Question-Aware Knowledge Graph Prompting for Enhancing Large Language Models
Haochen Liu, Song Wang, Chen Chen and Jundong Li 1388

UQ-Merge: Uncertainty Guided Multimodal Large Language Model Merging
Huaizhi Qu, Xinyu Zhao, Jie Peng, Kwonjoon Lee, Behzad Dariush and Tianlong Chen 1401

AQuAECHR: Attributed Question Answering for European Court of Human Rights
Korbinian Q. Weidinger, Santosh T.y.s.s, Oana Ichim and Matthias Grabmair 1418

Leveraging Unit Language Guidance to Advance Speech Modeling in Textless Speech-to-Speech Translation
Yuhao Zhang, Xiangnan Ma, Kaiqi Kou, Peizhuo Liu, Weiqiao Shan, Benyou Wang, Tong Xiao, Yuxin Huang, Zhengtao Yu and JingBo Zhu 1448

Ponder & Press: Advancing Visual GUI Agent towards General Computer Control
Yiqin Wang, Haoji Zhang, Jingqi Tian and Yansong Tang 1461

LogicGame: Benchmarking Rule-Based Reasoning Abilities of Large Language Models
Jiayi Gui, Yiming Liu, Jiale Cheng, Xiaotao Gu, Xiao Liu, Hongning Wang, Yuxiao Dong, Jie Tang and Minlie Huang 1474

LLM-Based Multi-Agent Systems are Scalable Graph Generative Models
Jiarui Ji, Runlin Lei, Jialing Bi, Zhewei Wei, Xu Chen, Yankai Lin, Xuchen Pan, Yaliang Li and Bolin Ding 1492

AD-LLM: Benchmarking Large Language Models for Anomaly Detection
Tiankai Yang, Yi Nian, Li Li, Ruiyao Xu, Yuangang Li, Jiaqi li, Zhuo Xiao, Xiyang Hu, Ryan A. Rossi, Kaize Ding, Xia Hu and Yue Zhao 1524

RTADev: Intention Aligned Multi-Agent Framework for Software Development
Jie Liu, Guohua Wang, Ronghui Yang, Jiajie Zeng, Mengchen Zhao and Yi Cai 1548

TACO-RL: Task Aware Prompt Compression Optimization with Reinforcement Learning
Shivam Shandilya, Menglin Xia, Supriyo Ghosh, Huiqiang Jiang, Jue Zhang, Qianhui Wu, Victor Rühle and Saravan Rajmohan 1582

A Character-Centric Creative Story Generation via Imagination
Kyeongman Park, Minbeom Kim and Kyomin Jung 1598

Proverbs Run in Pairs: Evaluating Proverb Translation Capability of Large Language Model
Minghan Wang, Viet Thanh Pham, Farhad Moghimifar and Thuy-Trang Vu 1646

Towards Efficient LLM Grounding for Embodied Multi-Agent Collaboration
Yang Zhang, Shixin Yang, Chenjia Bai, Fei Wu, Xiu Li, Zhen Wang and Xuelong Li 1663

UAQFact: Evaluating Factual Knowledge Utilization of LLMs on Unanswerable Questions
Chuanyuan Tan, Wenbiao Shao, Hao Xiong, Tong Zhu, Zhenhua Liu, Kai Shi and Wenliang Chen 1700

Exploring Knowledge Filtering for Retrieval-Augmented Discriminative Tasks
Minjie Qiang, Zhongqing Wang, Xiaoyi Bao, HaoYuan Ma, Shoushan Li and Guodong Zhou1716

Group then Scale: Dynamic Mixture-of-Experts Multilingual Language Model
Chong Li, Yingzhuo Deng, Jiajun Zhang and Chengqing Zong 1730

Beyond Verbal Cues: Emotional Contagion Graph Network for Causal Emotion Entailment
Fangxu Yu, Junjie Guo, Zhen Wu and Xinyu Dai 1755

Critic-CoT: Boosting the Reasoning Abilities of Large Language Model via Chain-of-Thought Critic
Xin Zheng, Jie Lou, Boxi Cao, Xueru Wen, Yuqiu Ji, Hongyu Lin, Yaojie Lu, Xianpei Han, Debing Zhang and Le Sun 1768

Systematic Generalization in Language Models Scales with Information Entropy
Sondre Wold, Lucas Georges Gabriel Charpentier and Étienne Simon 1807

The Inverse Scaling Effect of Pre-Trained Language Model Surprisal Is Not Due to Data Leakage
Byung-Doh Oh, Hongao Zhu and William Schuler 1820

Logical Consistency is Vital: Neural-Symbolic Information Retrieval for Negative-Constraint Queries
Ganlin Xu, Zhoujia Zhang, Wangyi Mei, Jiaqing Liang, Weijia Lu, Xiaodong Zhang, Zhifei Yang, Xiaofeng Ma, Yanghua Xiao and Deqing Yang 1828

'No' Matters: Out-of-Distribution Detection in Multimodality Multi-Turn Interactive Dialogue Download PDF
Rena Wei Gao, Xuetong Wu, Siwen Luo, Caren Han and Feng Liu 1848

Event Pattern-Instance Graph: A Multi-Round Role Representation Learning Strategy for Document-Level Event Argument Extraction
Qizhi Wan, LiuTao LiuTao, Changxuan Wan, Rong Hu, Keli Xiao and Yuxin Shuai 1865

EXECUTE: A Multilingual Benchmark for LLM Token Understanding
Lukas Edman, Helmut Schmid and Alexander Fraser 1878

Explainable Hallucination through Natural Language Inference Mapping
Wei-Fan Chen, Zhixue Zhao, Akbar Karimi and Lucie Flek 1888

HopRAG: Multi-Hop Reasoning for Logic-Aware Retrieval-Augmented Generation
Hao Liu, Zhengren Wang, Xi Chen, Zhiyu li, Feiyu Xiong, Qinhan Yu and Wentao Zhang . . 1897

Double Entendre: Robust Audio-Based AI-Generated Lyrics Detection via Multi-View Fusion
Markus Frohmann, Gabriel Meseguer-Brocal, Markus Schedl and Elena V. Epure 1914

Don't Miss the Forest for the Trees: Attentional Vision Calibration for Large Vision Language Models
Sangmin Woo, Donguk Kim, Jaehyuk Jang, Yubin Choi and Changick Kim 1927

SATA: A Paradigm for LLM Jailbreak via Simple Assistive Task Linkage
Xiaoning Dong, Wenbo Hu, Wei Xu and Tianxing He 1952

Chain-Talker: Chain Understanding and Rendering for Empathetic Conversational Speech Synthesis
Yifan Hu, Rui Liu, Yi Ren, Xiang Yin and Haizhou Li 1988

Parameter-Efficient Fine-Tuning via Circular Convolution
Aochuan Chen, Jiashun Cheng, Zijing Liu, Ziqi Gao, Fugee Tsung, Yu Li and Jia Li 2004

Alleviating Hallucinations in Large Language Models via Truthfulness-driven Rank-adaptive LoRA
Jiahao Li, Zhendong Mao and Quan Wang 2020

ScEdit: Script-based Assessment of Knowledge Editing
Xinye Li, Zunwen Zheng, Qian Zhang, Dekai Zhuang, Jiabao Kang, Liyan Xu, Qingbin Liu, Xi Chen, Zhiying Tu, Dianhui Chu and Dianbo Sui 2032

SafeRoute: Adaptive Model Selection for Efficient and Accurate Safety Guardrails in Large Language Models
Seanie Lee, Dong Bok Lee, Dominik Wagner, Minki Kang, Haebin Seong, Tobias Bocklet, Juho Lee and Sung Ju Hwang 2053

Moderation Matters: Measuring Conversational Moderation Impact in English as a Second Language Group Discussion
Rena Wei Gao, Ming-Bin Chen, Lea Frermann and Jey Han Lau 2070

Measuring Bias and Agreement in Large Language Model Presupposition Judgments
Katherine Atwell, Mandy Simons and Malihe Alikhani 2096

Harnessing PDF Data for Improving Japanese Large Multimodal Models
Jeonghun Baek, Akiko Aizawa and Kiyoharu Aizawa 2108

EnerGIZAr: Leveraging GIZA++ for Effective Tokenizer Initialization
Pranaydeep Singh, Eneko Agirre, Gorka Azkune, Orphee De Clercq and Els Lefever 2124

AMEX: Android Multi-annotation Expo Dataset for Mobile GUI Agents
Yuxiang Chai, Siyuan Huang, Yazhe Niu, Han Xiao, Liang Liu, Guozhi Wang, Dingyu Zhang, Shuai Ren and Hongsheng Li 2138

Drop Dropout on Single Epoch Language Model Pretraining
Houjun Liu, John Bauer and Christopher D Manning 2157

Robust and Minimally Invasive Watermarking for EaaS
Zongqi Wang, Baoyuan Wu, Jingyuan Deng and Yujiu Yang 2167

Task-Informed Anti-Curriculum by Masking Improves Downstream Performance on Text
Jarca Andrei, Florinel Alin Croitoru and Radu Tudor Ionescu 2192

CARMO: Dynamic Criteria Generation for Context Aware Reward Modelling
Taneesh Gupta, Shivam Shandilya, Xuchao Zhang, Rahul Madhavan, Supriyo Ghosh, Chetan Bansal, Huaxiu Yao and Saravan Rajmohan 2202

SLAM-Omni: Timbre-Controllable Voice Interaction System with Single-Stage Training
Wenxi Chen, Ziyang Ma, Ruiqi Yan, Yuzhe Liang, Xiquan Li, Ruiyang Xu, Zhikang Niu, Yanqiao Zhu, Yifan Yang, Zhanxun Liu, Kai Yu, Yuxuan Hu, Jinyu Li, Yan Lu, Shujie Liu and Xie Chen . 2262

C^2LEVA: Toward Comprehensive and Contamination-Free Language Model Evaluation
Yanyang Li, Wong Tin Long, Cheung To Hung, Jianqiao Zhao, Duo Zheng, Liu Ka Wai, Michael R. Lyu and Liwei Wang 2283

Texts or Images? A Fine-grained Analysis on the Effectiveness of Input Representations and Models for Table Question Answering
Wei Zhou, Mohsen Mesgar, Heike Adel and Annemarie Friedrich 2307

Adaptive-VP: A Framework for LLM-Based Virtual Patients that Adapts to Trainees' Dialogue to Facilitate Nurse Communication Training
Keyeun Lee, Seolhee Lee, Esther Hehsun Kim, Yena Ko, Jinsu Eun, Dahee Kim, Hyewon Cho, Haiyi Zhu, Robert E. Kraut, Eunyoung E. Suh, Eun-mee Kim and Hajin Lim 2319

Enhancing Multimodal Unified Representations for Cross Modal Generalization
Hai Huang, Yan Xia, Shengpeng Ji, Shulei Wang, Hanting Wang, Minghui Fang, Jieming Zhu, Zhenhua Dong, Sashuai Zhou and Zhou Zhao 2353

Domain Regeneration: How well do LLMs match syntactic properties of text domains?
Da JU, Hagen Blix and Adina Williams 2367

Structural Deep Encoding for Table Question Answering
Raphaël Mouravieff, Benjamin Piwowarski and Sylvain Lamprier 2389

MPL: Multiple Programming Languages with Large Language Models for Information Extraction
Bo Li, Gexiang Fang, Wei Ye, Zhenghua Xu, Jinglei Zhang, Hao Cheng and Shikun Zhang . 2403

Self-Critique Guided Iterative Reasoning for Multi-hop Question Answering
Zheng Chu, Huiming Fan, Jingchang Chen, Qianyu Wang, Mingda Yang, Jiafeng Liang, Zhongjie Wang, Hao Li, Guo Tang, Ming Liu and Bing Qin 2415

Anchored Answers: Unravelling Positional Bias in GPT-2's Multiple-Choice Questions
Ruizhe Li and Yanjun Gao 2439

Failing Forward: Improving Generative Error Correction for ASR with Synthetic Data and Retrieval Augmentation
Sreyan Ghosh, Mohammad Sadegh Rasooli, Michael Levit, Peidong Wang, Jian Xue, Dinesh Manocha and Jinyu Li 2466

LTRAG: Enhancing Autoformalization and Self-refinement for Logical Reasoning with Thought-Guided RAG
Ruikang Hu, Shaoyu Lin, Yeliang Xiu and Yongmei Liu 2483

Eta-WavLM: Efficient Speaker Identity Removal in Self-Supervised Speech Representations Using a Simple Linear Equation
Giuseppe Ruggiero, Matteo Testa, Jurgen Van De Walle and Luigi Di Caro 2494

MathCoder-VL: Bridging Vision and Code for Enhanced Multimodal Mathematical Reasoning
Ke Wang, Junting Pan, Linda Wei, Aojun Zhou, Weikang Shi, Zimu Lu, Han Xiao, Yunqiao Yang, Houxing Ren, Mingjie Zhan and Hongsheng Li 2505

MlingConf: A Comprehensive Study of Multilingual Confidence Estimation on Large Language Models
Boyang Xue, Hongru Wang, Rui Wang, Sheng Wang, Zezhong Wang, Yiming Du, Bin Liang, Wenxuan Zhang and Kam-Fai Wong 2535

COMPKE: Complex Question Answering under Knowledge Editing
Keyuan Cheng, Zijian Kan, Zhuoran Zhang, Muhammad Asif Ali, Lijie Hu and Di Wang . . 2557

RaaS: Reasoning-Aware Attention Sparsity for Efficient LLM Reasoning
Junhao Hu, Wenrui Huang, Weidong Wang, Zhenwen Li, Tiancheng hu, Zhixia Liu, Xusheng Chen, Tao Xie and Yizhou Shan 2577

One-for-All Pruning: A Universal Model for Customized Compression of Large Language Models
Rongguang Ye and Ming Tang 2591

CLaMP 3: Universal Music Information Retrieval Across Unaligned Modalities and Unseen Languages
Shangda Wu, Guo Zhancheng, Ruibin Yuan, Junyan Jiang, SeungHeon Doh, Gus Xia, Juhan Nam, Xiaobing Li, Feng Yu and Maosong Sun 2605

PFDial: A Structured Dialogue Instruction Fine-tuning Method Based on UML Flowcharts
Ming Zhang, Yuhui Wang, Yujiong Shen, Tingyi Yang, Changhao Jiang, Yilong Wu, Shihan Dou, Qinhao Chen, Zhiheng Xi, Zhihao Zhang, Yi Dong, Zhen Wang, Zhihui Fei, Mingyang Wan, Tao Liang, Guojun Ma, Qi Zhang, Tao Gui and Xuanjing Huang 2626

Listening to Patients: Detecting and Mitigating Patient Misreport in Medical Dialogue System
Lang Qin, Yao Zhang, Hongru Liang, Adam Jatowt and Zhenglu Yang 2650

Do Language Models Understand the Cognitive Tasks Given to Them? Investigations with the N-Back Paradigm
Xiaoyang Hu and Richard Lewis 2665

Graph-guided Cross-composition Feature Disentanglement for Compositional Zero-shot Learning
Yuxia Geng, Runkai Zhu, Jiaoyan Chen, Jintai Chen, Xiang Chen, Zhuo Chen, Shuofei Qiao, Yuxiang Wang, Xiaoliang Xu and Sheng-Jun Huang 2678

Training Long-Context LLMs Efficiently via Chunk-wise Optimization
Wenhao Li, Yuxin Zhang, Gen Luo, Daohai Yu and Rongrong Ji 2691

Revisiting LoRA through the Lens of Parameter Redundancy: Spectral Encoding Helps
Jiashun Cheng, Aochuan Chen, Nuo Chen, Ziqi Gao, Yuhan Li, Jia Li and Fugee Tsung . . . 2701

CODEMENV: Benchmarking Large Language Models on Code Migration
Keyuan Cheng, Xudong Shen, Yihao Yang, TengyueWang TengyueWang, Yang Cao, Muhammad Asif Ali, Hanbin Wang, Lijie Hu and Di Wang 2719

A Case Study of Cross-Lingual Zero-Shot Generalization for Classical Languages in LLMs
V.S.D.S.Mahesh Akavarapu, Hrishikesh Terdalkar, Pramit Bhattacharyya, Shubhangi Agarwal, Dr. Vishakha Deulgaonkar, Chaitali Dangarikar, Pralay Manna and Arnab Bhattacharya 2745

BrainECHO: Semantic Brain Signal Decoding through Vector-Quantized Spectrogram Reconstruction for Whisper-Enhanced Text Generation
Jilong Li, Zhenxi Song, Jiaqi Wang, Meishan Zhang, Honghai Liu, Min Zhang and Zhiguo Zhang 2762

Progressive LoRA for Multimodal Continual Instruction Tuning
Yahan Yu, Duzhen Zhang, Yong Ren, Xuanle Zhao, Xiuyi Chen and Chenhui Chu 2779

ARC 'Challenge' Is Not That Challenging
Łukasz Borchmann 2797

Cross-Lingual Transfer of Debiasing and Detoxification in Multilingual LLMs: An Extensive Investigation
Vera Neplenbroek, Arianna Bisazza and Raquel Fernández 2805

Tracr-Injection: Distilling Algorithms into Pre-trained Language Models
Tomás Vergara Browne and Alvaro Soto 2831

Model Performance-Guided Evaluation Data Selection for Effective Prompt Optimization
Ximing Dong, Shaowei Wang, Dayi Lin and Ahmed Hassan 2844

Revisiting Weak-to-Strong Generalization in Theory and Practice: Reverse KL vs. Forward KL
Wei Yao, Wenkai Yang, Ziqiao Wang, Yankai Lin and Yong Liu 2860

Stories that (are) Move(d by) Markets: A Causal Exploration of Market Shocks and Semantic Shifts across Different Partisan Groups
Felix Drinkall, Stefan Zohren, Michael McMahon and Janet B. Pierrehumbert 2889

NetSafe: Exploring the Topological Safety of Multi-agent System
Miao Yu, Shilong Wang, Guibin Zhang, Junyuan Mao, Chenlong Yin, Qijiong Liu, Kun Wang, Qingsong Wen and Yang Wang 2905

Reasoning is All You Need for Video Generalization: A Counterfactual Benchmark with Sub-question Evaluation
Qiji Zhou, YiFan Gong, Guangsheng Bao, Hongjie Qiu, Jinqiang Li, Xiangrong Zhu, Huajian Zhang and Yue Zhang 2939

Initializing and Retrofitting Key-Value Adaptors for Traceable Model Editing
Hanlun Zhu, Yunshi Lan, Xiang Li and Weining Qian 2958

Know the Unknown: An Uncertainty-Sensitive Method for LLM Instruction Tuning
Jiaqi Li, Yixuan Tang and Yi Yang 2972

Position-Aware Depth Decay Decoding (D^3): Boosting Large Language Model Inference Efficiency
Siqi Fan, Xuezhi Fang, Xingrun Xing, Peng Han, Shuo Shang and Yequan Wang 2990

Explaining Puzzle Solutions in Natural Language: An Exploratory Study on 6x6 Sudoku
Anirudh Maiya, Razan Alghamdi, Maria Leonor Pacheco, Ashutosh Trivedi and Fabio Somenzi 3002

Stress-testing Machine Generated Text Detection: Shifting Language Models Writing Style to Fool Detectors
Andrea Pedrotti, Michele Papucci, Cristiano Ciaccio, Alessio Miaschi, Giovanni Puccetti, Felice Dell'Orletta and Andrea Esuli 3010

InfiniSST: Simultaneous Translation of Unbounded Speech with Large Language Model
Siqi Ouyang, Xi Xu and Lei Li 3032

VSCBench: Bridging the Gap in Vision-Language Model Safety Calibration
Jiahui Geng, Qing Li, Zongxiong Chen, Yuxia Wang, Derui Zhu, Zhuohan Xie, Chenyang Lyu, Xiuying Chen, Preslav Nakov and Fakhri Karray 3047

To Code or not to Code? Adaptive Tool Integration for Math Language Models via Expectation-Maximization
Haozhe Wang, Long Li, Chao Qu, Weidi Xu, Fengming Zhu, Wei Chu and Fangzhen Lin . . 3060

GOODLIAR: A Reinforcement Learning-Based Deceptive Agent for Disrupting LLM Beliefs on Foundational Principles
Soo Kyung Kim and Hyunsoo Cho 3076

How Does Response Length Affect Long-Form Factuality
James Xu Zhao, Jimmy Z.j. Liu, Bryan Hooi and See-Kiong Ng 3102

Scaling LLMs' Social Reasoning: Sprinkle Cognitive Aha Momentinto Fundamental Long-thought Logical Capabilities
Guiyang Hou, Wenqi Zhang, Zhe Zheng, Yongliang Shen and Weiming Lu 3126

SimGRAG: Leveraging Similar Subgraphs for Knowledge Graphs Driven Retrieval-Augmented Generation
Yuzheng Cai, Zhenyue Guo, YiWen Pei, WanRui Bian and Weiguo Zheng 3139

RuleEdit: Towards Rule-Level Knowledge Generalization to Mitigate Over-Editing in Large Language Models
Bihan Zhou, HaoPeng Ren, li Yuan, Yi Cai, Liuwen Cao and Zikun Deng 3159

Eliciting In-context Retrieval and Reasoning for Long-context Large Language Models
Yifu Qiu, Varun R. Embar, Yizhe Zhang, Navdeep Jaitly, Shay B Cohen and Benjamin Han 3176

GeAR: Generation Augmented Retrieval
Haoyu Liu, Shaohan Huang, Jianfeng Liu, Yuefeng Zhan, Hao Sun, Weiwei Deng, Feng Sun, Furu Wei and Qi Zhang 3193

A Unified Taxonomy-Guided Instruction Tuning Framework for Entity Set Expansion and Taxonomy Expansion
Yanzhen Shen, Yu Zhang, Yunyi Zhang and Jiawei Han 3208

Zero-Shot Conversational Stance Detection: Dataset and Approaches
Yuzhe Ding, Kang He, Bobo Li, Li Zheng, Haijun He, Fei Li, Chong Teng and Donghong Ji3221

LongFaith: Enhancing Long-Context Reasoning in LLMs with Faithful Synthetic Data
Cehao Yang, Xueyuan Lin, Chengjin Xu, Xuhui Jiang, Shengjie Ma, Aofan Liu, Hui Xiong and Jian Guo 3236

SYNTHVERIFY: Enhancing Zero-Shot Claim Verification through Step-by-Step Synthetic Data Generation
Rongwen Zhao and Jeffrey Flanigan 3257

Domaino1s: Guiding LLM Reasoning for Explainable Answers in High-Stakes Domains
Xu Chu, Zhijie Tan, Hanlin Xue, Guanyu Wang, Tong Mo and Weiping Li 3275

Dynamic Prefix as Instructor for Incremental Named Entity Recognition: A Unified Seq2Seq Generation Framework
Zihao Wu, YongXiang Hua, Yongxin Zhu, Fang Zhang and Linli Xu 3294

Who Taught You That? Tracing Teachers in Model Distillation
Somin Wadhwa, Chantal Shaib, Silvio Amir and Byron C Wallace 3307

D-GEN: Automatic Distractor Generation and Evaluation for Reliable Assessment of Generative Models
Grace Byun and Jinho D. Choi 3316

HammerBench: Fine-Grained Function-Calling Evaluation in Real Mobile Assistant Scenarios
Jun Wang, Jiamu Zhou, Xihuai Wang, Xiaoyun Mo, Haoyu Zhang, Qiqiang Lin, Jincheng Jincheng, Muning Wen, Weinan Zhang, Qiuying Peng and Jun Wang 3350

Beyond In-Context Learning: Aligning Long-form Generation of Large Language Models via Task-Inherent Attribute Guidelines
Do Xuan Long, Duong Ngoc Yen, Do Xuan Trong, Anh Tuan Luu, Kenji Kawaguchi, Shafiq Joty, Min-Yen Kan and Nancy F. Chen 3377

GRAMMAR-LLM: Grammar-Constrained Natural Language Generation
Gabriele Tuccio, Luana Bulla, Maria Madonia, Aldo Gangemi and Misael Mongiovì 3412

MANBench: Is Your Multimodal Model Smarter than Human?
Han Zhou, Qitong Xu, Yiheng Dong and Xin Yang 3423

BanStereoSet: A Dataset to Measure Stereotypical Social Biases in LLMs for Bangla
Mahammed Kamruzzaman, Abdullah Al Monsur, Shrabon Kumar Das, Enamul Hassan and Gene Louis Kim 3450

mOSCAR: A Large-scale Multilingual and Multimodal Document-level Corpus
Matthieu Futeral, Armel Randy Zebaze, Pedro Ortiz Suarez, Julien Abadji, Rémi Lacroix, Cordelia Schmid, Rachel Bawden and Benoît Sagot 3461

NorEval: A Norwegian Language Understanding and Generation Evaluation Benchmark
Vladislav Mikhailov, Tita Enstad, David Samuel, Hans Christian Farsethås, Andrey Kutuzov, Erik Velldal and Lilja Øvrelid 3495

Massively Multilingual Instruction-Following Information Extraction
Thang Le, Huy Huu Nguyen, Anh Tuan Luu and Thien Huu Nguyen 3542

DALR: Dual-level Alignment Learning for Multimodal Sentence Representation Learning
Kang He, Yuzhe Ding, Haining Wang, Fei Li, Chong Teng and Donghong Ji 3586

Large Language Models in Bioinformatics: A Survey
Zhenyu Wang, Zikang Wang, Jiyue Jiang, Pengan Chen, Xiangyu Shi and Yu Li 3602

ChartEdit: How Far Are MLLMs From Automating Chart Analysis? Evaluating MLLMs' Capability via Chart Editing
Xuanle Zhao, Xuexin Liu, Yang Haoyue, Xianzhen Luo, Fanhu Zeng, Jianling li, Qi Shi and Chi Chen 3616

Unraveling and Mitigating Safety Alignment Degradation of Vision-Language Models
Qin Liu, Chao Shang, Ling Liu, Nikolaos Pappas, Jie Ma, Neha Anna John, Srikanth Doss, Lluis Marquez, Miguel Ballesteros and Yassine Benajiba 3631

Turbocharging Web Automation: The Impact of Compressed History States
Xiyue Zhu, Peng Tang, Haofu Liao and Srikar Appalaraju 3644

Making RALM Robust to Irrelevant Contexts via Layer Knowledge Guided Attention
Weijie Shi, Hao Chen, Jiaming Li, Yao Zhao, Yazhong Zhang, Qijin Chen, Jipeng Zhang, Ruiyuan Zhang, Jia Zhu, Jiajie Xu and Xiaofang Zhou 3652

Rewrite to Jailbreak: Discover Learnable and Transferable Implicit Harmfulness Instruction
Yuting Huang, Chengyuan Liu, Yifeng Feng, Yiquan Wu, Chao Wu, Fei Wu and Kun Kuang 3669

SignAlignLM: Integrating Multimodal Sign Language Processing into Large Language Models
Mert Inan, Anthony Sicilia and Malihe Alikhani 3691

NegVQA: Can Vision Language Models Understand Negation?
Yuhui Zhang, Yuchang Su, Yiming Liu and Serena Yeung-Levy 3707

Natural Language Reasoning in Large Language Models: Analysis and Evaluation
Debela Gemechu, Ramon Ruiz-Dolz, Henrike Beyer and Chris Reed 3717

SWE-Dev: Building Software Engineering Agents with Training and Inference Scaling
Haoran Wang, Zhenyu Hou, Yao Wei, Jie Tang and Yuxiao Dong 3742

The Two Paradigms of LLM Detection: Authorship Attribution vs Authorship Verification
Janek Bevendorff, Matti Wiegmann, Emmelie Richter, Martin Potthast and Benno Stein . . . 3762

Unveiling Confirmation Bias in Chain-of-Thought Reasoning
Yue Wan, Xiaowei Jia and Xiang Lorraine Li 3788

GRNFormer: A Biologically-Guided Framework for Integrating Gene Regulatory Networks into RNA Foundation Models
Mufan Qiu, Xinyu Hu, Fengwei Zhan, Sukwon Yun, Jie Peng, Ruichen Zhang, Bhavya Kailkhura, Jiekun Yang and Tianlong Chen 3805

RemoteRAG: A Privacy-Preserving LLM Cloud RAG Service
Yihang Cheng, Lan Zhang, Junyang Wang, Mu Yuan and Yunhao Yao 3820

My life is miserable, have to sign 500 autographs everyday": Exposing Humblebragging, the Brags in Disguise
Sharath Naganna, Saprativa Bhattacharjee, Biplab Banerjee and Pushpak Bhattacharyya . . . 3838

SCITAT: A Question Answering Benchmark for Scientific Tables and Text Covering Diverse Reasoning Types
Xuanliang Zhang, Dingzirui Wang, Baoxin Wang, Longxu Dou, Xinyuan Lu, Keyan Xu, Dayong Wu and Qingfu Zhu 3859

TokenShapley: Token Level Context Attribution with Shapley Value
Yingtai Xiao, Yuqing Zhu, Sirat Samyoun, Wanrong Zhang, Jiachen T. Wang and Jian Du . . 3882

Entropy-based Exploration Conduction for Multi-step Reasoning
Jinghan Zhang, Xiting Wang, Fengran Mo, Yeyang Zhou, Wanfu Gao and Kunpeng Liu . . . 3895

Taxonomizing Representational Harms using Speech Act Theory
Emily Corvi, Hannah Washington, Stefanie Reed, Chad Atalla, Alexandra Chouldechova, P. Alex Dow, Jean Garcia-Gathright, Nicholas J Pangakis, Emily Sheng, Dan Vann, Matthew Vogel and Hanna Wallach 3907

Turning Conversations into Workflows: A Framework to Extract and Evaluate Dialog Workflows for Service AI Agents
Prafulla Kumar Choubey, Xiangyu Peng, Shilpa Bhagavath, Caiming Xiong, Shiva Kumar Pentyala and Chien-Sheng Wu 3933

Statistical inference on black-box generative models in the data kernel perspective space
Hayden Helm, Aranyak Acharyya, Youngser Park, Brandon Duderstadt and Carey Priebe . . 3955

Do Large Language Models Perform Latent Multi-Hop Reasoning without Exploiting Shortcuts?
Sohee Yang, Nora Kassner, Elena Gribovskaya, Sebastian Riedel and Mor Geva 3971

AceMath: Advancing Frontier Math Reasoning with Post-Training and Reward Modeling
Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro and Wei Ping 3993

WXImpactBench: A Disruptive Weather Impact Understanding Benchmark for Evaluating Large Language Models
Yongan Yu, Qingchen Hu, Xianda Du, Jiayin Wang, Fengran Mo and Renée Sieber 4016

MeMoTune: A Measure and Moment-Driven Fine-Tuning Framework for Quantized Large Language Models
Yun Zhang, Xue Geng, Lizi Liao, Jintong Sun, Minghe Yu and Ge Yu 4036

MALAMUTE: A Multilingual, Highly-granular, Template-free, Education-based Probing Dataset
Sagi Shaier, George Arthur Baker, Chiranthan Sridhar, Lawrence Hunter and Katharina Von Der Wense 4051

Sentimental Image Generation for Aspect-based Sentiment Analysis
Xiaoyi Bao, Jinghang Gu, Zhongqing Wang and Chu-Ren Huang 4070

Long-form Hallucination Detection with Self-elicitation
Zihang Liu, Jiawei Guo, Hao Zhang, Hongyang Chen, Jiajun Bu and Haishuai Wang 4082

ComparisonQA: Evaluating Factuality Robustness of LLMs Through Knowledge Frequency Control and Uncertainty
Qing Zong, Zhaowei Wang, Tianshi Zheng, Xiyu Ren and Yangqiu Song 4101

One-Dimensional Object Detection for Streaming Text Segmentation of Meeting Dialogue
Rui He, Zhongqing Wang, Minjie Qiang, Hongling Wang, Yifan.zhang Yifan.zhang, Hua Xu, Shuai Fan and Guodong Zhou 4118

CodeTaxo: Enhancing Taxonomy Expansion with Limited Examples via Code Language Prompts
Qingkai Zeng, Yuyang Bai, Zhaoxuan Tan, Zhenyu Wu, Shangbin Feng and Meng Jiang . . . 4131

Predicate-Conditional Conformalized Answer Sets for Knowledge Graph Embeddings
Yuqicheng Zhu, Daniel Hernández, Yuan He, Zifeng Ding, Bo Xiong, Evgeny Kharlamov and Steffen Staab 4145

Autonomous Data Selection with Zero-shot Generative Classifiers for Mathematical Texts
Yifan Zhang, Yifan Luo, Yang Yuan and Andrew C Yao 4168

Learning from Committee: Reasoning Distillation from a Mixture of Teachers with Peer-Review
Zhuochun Li, Yuelyu Ji, Rui Meng and Daqing He 4190

Investigating Prosodic Signatures via Speech Pre-Trained Models for Audio Deepfake Source Attribution
Orchid Chetia Phukan, Drishti Singh, Swarup Ranjan Behera, Arun Balaji Buduru and Rajesh Sharma 4206

Multilingual Retrieval Augmented Generation for Culturally-Sensitive Tasks: A Benchmark for Cross-lingual Robustness
Bryan Li, Fiona Luo, Samar Haider, Adwait Agashe, Siyu Li, Runqi Liu, Miranda Muqing Miao, Shriya Ramakrishnan, Yuan Yuan and Chris Callison-Burch 4215

Bridging Relevance and Reasoning: Rationale Distillation in Retrieval-Augmented Generation
Pengyue Jia, Derong Xu, Xiaopeng Li, Zhaocheng Du, Xiangyang Li, Yichao Wang, Yuhao Wang, Qidong Liu, Maolin Wang, Huifeng Guo, Ruiming Tang and Xiangyu Zhao 4242

Scaling Laws for Multilingual Language Models
Yifei He, Alon Benhaim, Barun Patra, Praneetha Vaddamanu, Sanchit Ahuja, Parul Chopra, Vishrav Chaudhary, Han Zhao and Xia Song 4257

Corpus Poisoning via Approximate Greedy Gradient Descent
Jinyan Su, Preslav Nakov and Claire Cardie 4274

Taxonomy-Driven Knowledge Graph Construction for Domain-Specific Scientific Applications
Huitong Pan, Qi Zhang, Mustapha Adamu, Eduard Dragut and Longin Jan Latecki 4295

Wanda++: Pruning Large Language Models via Regional Gradients
Yifan Yang, Kai Zhen, Bhavana Ganesh, Aram Galstyan, Goeric Huybrechts, Markus Müller, Jonas M. Kübler, Rupak Vignesh Swaminathan, Athanasios Mouchtaris, Sravan Babu Bodapati, Nathan Susanj, Zheng Zhang, Jack FitzGerald and Abhishek Kumar 4321

MATCHED: Multimodal Authorship-Attribution To Combat Human Trafficking in Escort-Advertisement Data
Vageesh Kumar Saxena, Benjamin Ashpole, Gijs Van Dijck and Gerasimos Spanakis 4334

Fraud-R1 : A Multi-Round Benchmark for Assessing the Robustness of LLM Against Augmented Fraud and Phishing Inducements
Shu Yang, Shenzhe Zhu, Zeyu Wu, Keyu Wang, Junchi Yao, Junchao Wu, Lijie Hu, Mengdi Li, Derek F. Wong and Di Wang 4374

Mitigating Paraphrase Attacks on Machine-Text Detection via Paraphrase Inversion
Rafael Alberto Rivera Soto, Barry Y. Chen and Nicholas Andrews 4421

SANSKRITI: A Comprehensive Benchmark for Evaluating Language Models' Knowledge of Indian Culture
Arijit Maji, Raghvendra Kumar, Akash Ghosh, Anushka Anushka and Sriparna Saha 4434

System Prompt Hijacking via Permutation Triggers in LLM Supply Chains
Lu Yan, Siyuan Cheng, Xuan Chen, Kaiyuan Zhang, Guangyu Shen and Xiangyu Zhang . . . 4452

Frequency matters: Modeling irregular morphological patterns in Spanish with Transformers
Akhilesh Kakolu Ramarao, Kevin Tang and Dinah Baer-Henney 4474

From Heart to Words: Generating Empathetic Responses via Integrated Figurative Language and Semantic Context Signals
Gyeongeun Lee, Zhu Wang, Sathya N. Ravi and Natalie Parde 4490

There's No Such Thing as Simple Reasoning for LLMs
Nurul Fajrin Ariyani, Zied Bouraoui, Richard Booth and Steven Schockaert 4503

CLIX: Cross-Lingual Explanations of Idiomatic Expressions
Aaron Gluck, Katharina Von Der Wense and Maria Leonor Pacheco 4515

Beyond Semantic Entropy: Boosting LLM Uncertainty Quantification with Pairwise Semantic Similarity
Dang Nguyen, Ali Payani and Baharan Mirzasoleiman 4530

R^3Mem: Bridging Memory Retention and Retrieval via Reversible Compression
Xiaoqiang Wang, Suyuchen Wang, Yun Zhu and Bang Liu 4541

Vision Language Model Helps Private Information De-Identification in Vision Data
Tiejin Chen, Pingzhi Li, Kaixiong Zhou, Tianlong Chen and Hua Wei 4558

Unveiling Privacy Risks in Multi-modal Large Language Models: Task-specific Vulnerabilities and Mitigation Challenges
Tiejin Chen, Pingzhi Li, Kaixiong Zhou, Tianlong Chen and Hua Wei 4573

DeFine: Decision-Making with Analogical Reasoning over Factor Profiles
Yebowen Hu, Xiaoyang Wang, Wenlin Yao, Yiming Lu, Daoan Zhang, Hassan Foroosh, Dong Yu and Fei Liu 4587

SMART: Self-Aware Agent for Tool Overuse Mitigation
Cheng Qian, Emre Can Acikgoz, Hongru Wang, Xiusi Chen, Avirup Sil, Dilek Hakkani-Tür, Gokhan Tur and Heng Ji 4604

Continued Pretraining and Interpretability-Based Evaluation for Low-Resource Languages: A Galician Case Study
Pablo Rodríguez, Silvia Paniagua Suárez, Pablo Gamallo and Susana Sotelo Docio 4622

TC-Bench: Benchmarking Temporal Compositionality in Conditional Video Generation
Weixi Feng, Jiachen Li, Michael Saxon, Tsu-Jui Fu, Wenhu Chen and William Yang Wang . 4638

DAM: Dynamic Attention Mask for Long-Context Large Language Model Inference Acceleration
Hanzhi Zhang, Heng Fan, Kewei Sha, Yan Huang and Yunhe Feng 4663

Arbiters of Ambivalence: Challenges of using LLMs in No-Consensus tasks
Bhaktipriya Radharapu, Manon Revel, Megan Ung, Sebastian Ruder and Adina Williams . . 4677

Beyond Text: Characterizing Domain Expert Needs in Document Research
Sireesh Gururaja, Nupoor Gandhi, Jeremiah Milbauer and Emma Strubell 4732

Efficient but Vulnerable: Benchmarking and Defending LLM Batch Prompting Attack
Murong Yue and Ziyu Yao 4746

MM-R^3: On (In-)Consistency of Vision-Language Models (VLMs)
Shih-Han Chou, Shivam Chandhok, Jim Little and Leonid Sigal 4762

Investigating Context Faithfulness in Large Language Models: The Roles of Memory Strength and Evidence Style
Yuepei Li, Kang Zhou, Qiao Qiao, Bach Nguyen, Qing Wang and Qi Li 4789

Shadow-Activated Backdoor Attacks on Multimodal Large Language Models
Ziyi Yin, Muchao Ye, Yuanpu Cao, Jiaqi Wang, Aofei Chang, Han Liu, Jinghui Chen, Ting Wang and Fenglong Ma 4808

Why Vision Language Models Struggle with Visual Arithmetic? Towards Enhanced Chart and Geometry Understanding
Kung-Hsiang Huang, Can Qin, Haoyi Qiu, Philippe Laban, Shafiq Joty, Caiming Xiong and Chien-Sheng Wu 4830

K-order Ranking Preference Optimization for Large Language Models
Shihao Cai, Chongming Gao, Yang Zhang, Wentao Shi, Jizhi Zhang, Keqin Bao, Qifan Wang and Fuli Feng 4844

Spectral Insights into Data-Oblivious Critical Layers in Large Language Models
Xuyuan Liu, Lei Hsiung, Yaoqing Yang and Yujun Yan 4860

SynFix: Dependency-Aware Program Repair via RelationGraph Analysis
Xunzhu Tang, Jiechao Gao, Jin Xu, Tiezhu Sun, Yewei Song, Saad Ezzini, Wendkûuni C. Ouédraogo, Jacques Klein and Tegawendé F. Bissyandé 4878

EXIT: Context-Aware Extractive Compression for Enhancing Retrieval-Augmented Generation
Taeho Hwang, Sukmin Cho, Soyeong Jeong, Hoyun Song, SeungYoon Han and Jong C. Park 4895

Re-TASK: Revisiting LLM Tasks from Capability, Skill, and Knowledge Perspectives
Zhihu Wang, Shiwan Zhao, Yu Wang, Heyuan Huang, Sitao Xie, Yubo Zhang, Jiaxin Shi, Zhixing Wang, Hongyan Li and Junchi Yan 4925

Unlearning Backdoor Attacks for LLMs with Weak-to-Strong Knowledge Distillation
Shuai Zhao, Xiaobao Wu, Cong-Duy T Nguyen, Yanhao Jia, Meihuizi Jia, Feng Yichao and Anh Tuan Luu 4937

Packing Analysis: Packing Is More Appropriate for Large Models or Datasets in Supervised Fine-tuning
Shuhe Wang, Guoyin Wang, Yizhong Wang, Jiwei Li, Eduard Hovy and Chen Guo 4953

Better Red Teaming via Searching with Large Language Model
Yongkang Chen, Chongyang Zhao, Jianwentian Jianwentian, Guiling Cao, Hu LI and Xiaohui Kuang 4968

AdaV: Adaptive Text-visual Redirection for Vision-Language Models
Jiayi Han, Liang Du, Yiwen Wu, Guanming Liang, Xiangguo Zhou, Weibo Zheng, Donghong Han and Zixun Sun 4985

MegaAgent: A Large-Scale Autonomous LLM-based Multi-Agent System Without Predefined SOPs
Qian Wang, Tianyu Wang, Zhenheng Tang, Qinbin Li, Nuo Chen, Jingsheng Liang and Bingsheng He 4998

Persona-judge: Personalized Alignment of Large Language Models via Token-level Self-judgment
Xiaotian Zhang, Ruizhe Chen, Yang Feng and Zuozhu Liu 5037

A Self-Distillation Recipe for Neural Machine Translation
Hongfei Xu, Zhuofei Liang, Qiuhui Liu and Lingling Mu 5050

BlockPruner: Fine-grained Pruning for Large Language Models
Longguang Zhong, Fanqi Wan, Ruijun Chen, Xiaojun Quan and Liangzhi Li 5065

Evaluating Implicit Bias in Large Language Models by Attacking From a Psychometric Perspective
Yuchen Wen, Keping Bi, Wei Chen, Jiafeng Guo and Xueqi Cheng 5081

LongCite: Enabling LLMs to Generate Fine-grained Citations in Long-Context QA
Jiajie Zhang, Yushi Bai, Xin Lv, Wanjun Gu, Danqing Liu, Minhao Zou, Shulin Cao, Lei Hou, Yuxiao Dong, Ling Feng and Juanzi Li 5098

An Empirical Study of Group Conformity in Multi-Agent Systems
Min Choi, Keonwoo Kim, Sungwon Chae and Sangyeop Baek 5123

Combining the Best of Both Worlds: A Method for Hybrid NMT and LLM Translation
Zhanglin Wu, Daimeng Wei, Xiaoyu Chen, Hengchao Shang, Jiaxin Guo, Zongyao Li, Yuanchang Luo, Jinlong Yang, Zhiqiang Rao and Hao Yang 5140

ASPO: Adaptive Sentence-Level Preference Optimization for Fine-Grained Multimodal Reasoning
Yeyuan Wang, Dehong Gao, Rujiao Long, Lei Yi, Linbo Jin, Libin Yang and Xiaoyan Cai . . 5149

NovelCR: A Large-Scale Bilingual Dataset Tailored for Long-Span Coreference Resolution
MeiHan Tong and Shuai Wang 5161

Dynamic Attention-Guided Context Decoding for Mitigating Context Faithfulness Hallucinations in Large Language Models
Huangyw Huangyw, Yong Zhang, Ning Cheng, Zhitao Li, Shaojun Wang and Jing Xiao . . . 5174

Exploring the Choice Behavior of Large Language Models
Weidong Wu, Qinlin Zhao, Hao Chen, Lexin Zhou, Defu Lian and Hong Xie 5194

On-Policy Self-Alignment with Fine-grained Knowledge Feedback for Hallucination Mitigation
Xueru Wen, Jie Lou, Xinyu Lu, Yuqiu Ji, Xinyan Guan, Yaojie Lu, Hongyu Lin, Ben He, Xianpei Han, Debing Zhang and Le Sun 5215

From Phrases to Subgraphs: Fine-Grained Semantic Parsing for Knowledge Graph Question Answering
Yurun Song, Xiangqing Shen and Rui Xia 5232

StableToolBench-MirrorAPI: Modeling Tool Environments as Mirrors of 7,000+ Real-World APIs
Zhicheng Guo, Sijie Cheng, Yuchen Niu, Hao Wang, Sicheng Zhou, Wenbing Huang and Yang Liu 5247

ClaimPKG: Enhancing Claim Verification via Pseudo-Subgraph Generation with Lightweight Specialized LLM
Hoang Pham, Thanh-Do Nguyen and Khac-Hoai Nam Bui 5271

TriEmbed: Bridge the Gap between Text and Token Indices with Embedding Reparameterization
Baizhou Huang and Xiaojun Wan 5291

Chain of Methodologies: Scaling Test Time Computation without Training
Cong Liu, Jie Wu, Weigang Wu, Xu Chen, Liang Lin and Wei-Shi Zheng 5298

A Survey on Personalized Alignment—The Missing Piece for Large Language Models in Real-World Applications
Jian Guan, Junfei Wu, Jia-Nan Li, Chuanqi Cheng and Wei Wu 5313

SuLoRA: Subspace Low-Rank Adaptation for Parameter-Efficient Fine-Tuning
Chenhao Ding, Jiangyang Li, SongLin Dong, Xinyuan Gao, Yuhang He and Yihong Gong . 5334

MIRe: Enhancing Multimodal Queries Representation via Fusion-Free Modality Interaction for Multimodal Retrieval
Yeong-Joon Ju, Ho-Joong Kim and Seong-Whan Lee 5350

Correcting on Graph: Faithful Semantic Parsing over Knowledge Graphs with Large Language Models
Ruilin Zhao, Feng Zhao and Hong Zhang 5364

COPR: Continual Human Preference Learning via Optimal Policy Regularization
Han Zhang, Lin Gui, Yu Lei, Yuanzhao Zhai, Yehong Zhang, Zhuo Zhang, Yulan He, Hui Wang, Yue Yu, Kam-Fai Wong, Bin Liang and Ruifeng Xu 5377

Robust Preference Optimization via Dynamic Target Margins
Jie Sun, Junkang Wu, Jiancan Wu, Zhibo Zhu, Xingyu Lu, Jun Zhou, Lintao Ma and Xiang Wang 5399

AdaReTaKe: Adaptive Redundancy Reduction to Perceive Longer for Video-language Understanding
Xiao Wang, Qingyi Si, Shiyu Zhu, Jianlong Wu, Li Cao and Liqiang Nie 5417

Rethinking Stateful Tool Use in Multi-Turn Dialogues: Benchmarks and Challenges
Hongru Wang, Wenyu Huang, Yufei Wang, Yuanhao Xi, Jianqiao Lu, Huan Zhang, Nan Hu, Zeming Liu, Jeff Z. Pan and Kam-Fai Wong 5433

Open-Set Living Need Prediction with Large Language Models
Xiaochong Lan, Jie Feng, Yizhou Sun, Chen Gao, Jiahuan Lei, Xinleishi Xinleishi, Hengliang Luo and Yong Li 5454

Improve Rule Retrieval and Reasoning with Self-Induction and Relevance ReEstimate
Ziyang Huang, Wangtao Sun, Jun Zhao and Kang Liu 5473

Beyond Words: Integrating Theory of Mind into Conversational Agents for Human-Like Belief, Desire, and Intention Alignment
Mehdi Jafari, Yuncheng Hua, Hao Xue and Flora D. Salim 5489

Multimodal Causal Reasoning Benchmark: Challenging Multimodal Large Language Models to Discern Causal Links Across Modalities
Zhiyuan Li, Heng Wang, Dongnan Liu, Chaoyi Zhang, Ao Ma, Jieting Long and Weidong Cai 5509

Context-Aware Hierarchical Merging for Long Document Summarization
Litu Ou and Mirella Lapata 5534

VCD: A Dataset for Visual Commonsense Discovery in Images
Xiangqing Shen, Fanfan Wang, Siwei Wu and Rui Xia 5562

Self-Reasoning Language Models: Unfold Hidden Reasoning Chains with Few Reasoning Catalyst
Hongru Wang, Deng Cai, Wanjun Zhong, Shijue Huang, Jeff Z. Pan, Zeming Liu and Kam-Fai Wong 5578

HyperCRS: Hypergraph-Aware Multi-Grained Preference Learning to Burst Filter Bubbles in Conversational Recommendation System
Yongsen Zheng, Mingjie Qian, Guohua Wang, Yang Liu, Ziliang Chen, Mingzhi Mao, Liang Lin and Kwok-Yan Lam 5597

Is LLM an Overconfident Judge? Unveiling the Capabilities of LLMs in Detecting Offensive Language with Annotation Disagreement
Junyu Lu, Kai Ma, Kaichun Wang, Kelaiti Xiao, Roy Ka-Wei Lee, Bo Xu, Liang Yang and Hongfei Lin 5609

Language Repository for Long Video Understanding
Kumara Kahatapitiya, Kanchana Ranasinghe, Jongwoo Park and Michael S Ryoo 5627

Investigating Language Preference of Multilingual RAG Systems
Jeonghyun Park and Hwanhee Lee 5647

FGDGNN: Fine-Grained Dynamic Graph Neural Network for Rumor Detection on Social Media
Mei Guo, Chen Chen, Chunyan Hou, Yike Wu and Xiaojie Yuan 5676

Self-Tuning: Instructing LLMs to Effectively Acquire New Knowledge through Self-Teaching
Xiaoying Zhang, Baolin Peng, Ye Tian, Jingyan Zhou, Yipeng Zhang, Haitao Mi and Helen M. Meng 5688

QueryAttack: Jailbreaking Aligned Large Language Models Using Structured Non-natural Query Language
Qingsong Zou, Jingyu Xiao, Qing Li, Zhi Yan, Yuhang Wang, Li Xu, Wenxuan Wang, Kuofeng Gao, Ruoyu Li and Yong Jiang 5725

Memory or Reasoning? Explore How LLMs Compute Mixed Arithmetic Expressions
Chengzhi Li, Heyan Huang, Ping Jian, Zhen Yang, Chenxu Wang and Yifan Wang 5742

PersonaX: A Recommendation Agent-Oriented User Modeling Framework for Long Behavior Sequence
Yunxiao Shi, Wujiang Xu, Zhang Zeqi, Xing Zi, Qiang Wu and Min Xu 5764

Judge as A Judge: Improving the Evaluation of Retrieval-Augmented Generation through the Judge-Consistency of Large Language Models
Shuliang Liu, Xinze Li, Zhenghao Liu, Yukun Yan, Cheng Yang, Zheni Zeng, Zhiyuan Liu, Maosong Sun and Ge Yu 5788

Rationales Are Not Silver Bullets: Measuring the Impact of Rationales on Model Performance and Reliability
Chiwei Zhu, Benfeng Xu, An Yang, Junyang Lin, Quan Wang, Chang Zhou and Zhendong Mao 5808

CA-GAR: Context-Aware Alignment of LLM Generation for Document Retrieval
Heng Yu, Junfeng Kang, Rui Li, Qi Liu, Liyang He, Zhenya Huang, Shuanghong Shen and Junyu Lu 5836

AgentCourt: Simulating Court with Adversarial Evolvable Lawyer Agents
Guhong Chen, Liyang Fan, Zihan Gong, Nan Xie, Zixuan Li, Ziqiang Liu, Chengming Li, Qiang QU, Hamid Alinejad-Rokny, Shiwen Ni and Min Yang 5850

MLDebugging: Towards Benchmarking Code Debugging Across Multi-Library Scenarios
JinYang Huang, Xiachong Feng, Qiguang Chen, Hanjie Zhao, Zihui Cheng, Jiesong Bai, Jingxuan Zhou, Min Li and Libo Qin 5866

An Empirical Study of LLM-as-a-Judge for LLM Evaluation: Fine-tuned Judge Model is not a General Substitute for GPT-4
Hui Huang, Xingyuan Bu, Hongli Zhou, Yingqi Qu, Jing Liu, Muyun Yang, Bing Xu and Tiejun Zhao 5880

Expectation Confirmation Preference Optimization for Multi-Turn Conversational Recommendation Agent
Xueyang Feng, Jingsen Zhang, Jiakai Tang, Wei Li, Guohao Cai, Xu Chen, Quanyu Dai, Yue Zhu and Zhenhua Dong 5896

ProMedTS: A Self-Supervised, Prompt-Guided Multimodal Approach for Integrating Medical Text and Time Series
Shuai Niu, Jing Ma, Hongzhan Lin, Liang Bai, Zhihua Wang, V. W., Richard Yi Da Xu, Guo Li and Xian Yang 5915

CipherBank: Exploring the Boundary of LLM Reasoning Capabilities through Cryptography Challenge
Yu Li, Qizhi Pei, Mengyuan Sun, Honglin Lin, Chenlin Ming, Xin Gao, Jiang Wu, Conghui He and Lijun Wu 5929

Which Retain Set Matters for LLM Unlearning? A Case Study on Entity Unlearning
Hwan Chang and Hwanhee Lee 5966

Tell Me What You Don't Know: Enhancing Refusal Capabilities of Role-Playing Agents via Representation Space Analysis and Editing
Wenhao Liu, Siyu An, Junru Lu, Muling Wu, Tianlong Li, Xiaohua Wang, Changze Lv, Xiaoqing Zheng, di Yin, Xing Sun and Xuanjing Huang 5983

LR²Bench: Evaluating Long-chain Reflective Reasoning Capabilities of Large Language Models via Constraint Satisfaction Problems
Jianghao Chen, Zhenlin Wei, Zhenjiang Ren, Ziyong Li and Jiajun Zhang 6006

McBE: A Multi-task Chinese Bias Evaluation Benchmark for Large Language Models
Tian Lan, Xiangdong Su, Xu Liu, Ruirui Wang, Ke Chang, Jiang Li and Guanglai Gao 6033

MARK: Multi-agent Collaboration with Ranking Guidance for Text-attributed Graph Clustering
Yiwei Fu, Yuxing Zhang, Chunchun Chen, JianwenMa JianwenMa, Quan Yuan, Rong-Cheng Tu, Xinli Huang, Wei Ye, Xiao Luo and Minghua Deng 6057

Can Language Models Capture Human Writing Preferences for Domain-Specific Text Summarization?
Jingbao Luo, Ming Liu, Ran Liu, Yongpan Sheng, Xin Hu, Gang Li and WupengNjust Wupengnjust 6073

Mitigate Position Bias in LLMs via Scaling a Single Hidden States Channel
Yijiong Yu, Huiqiang Jiang, Xufang Luo, Qianhui Wu, Chin-Yew Lin, Dongsheng Li, Yuqing Yang, Yongfeng Huang and Lili Qiu 6092

Self-attention-based Graph-of-Thought for Math Problem Solving
Ruiqiao Bai, Xue Han, Shuo Lei, Junlan Feng, Yanyan Luo and Chao Deng 6112

BAR: A Backward Reasoning based Agent for Complex Minecraft Tasks
Weihong Du, Wenrui Liao, Binyu Yan, Hongru Liang, Anthony G Cohn and Wenqiang Lei . 6126

KAPA: A Deliberative Agent Framework with Tree-Structured Knowledge Base for Multi-Domain User Intent Understanding
Jiakai Tang, Shiqi Shen, ZhipengWang ZhipengWang, Gong Zhi, Xueyang Feng, Zexu Sun, Haoran Tan and Xu Chen 6150

RASD: Retrieval-Augmented Speculative Decoding
Guofeng Quan, Wenfeng Feng, Chuzhan Hao, Guochao Jiang, Yuewei Zhang and Hao Henry Wang 6167

FRAG: A Flexible Modular Framework for Retrieval-Augmented Generation based on Knowledge Graphs
Zengyi Gao, Yukun Cao, Hairu Wang, Ao Ke, Yuan Feng, S Kevin Zhou and Xike Xie 6178

Reefknot: A Comprehensive Benchmark for Relation Hallucination Evaluation, Analysis and Mitigation in Multimodal Large Language Models
Kening Zheng, Junkai Chen, Yibo Yan, Xin Zou, Huiyu Zhou and Xuming Hu 6193

Blessing of Multilinguality: A Systematic Analysis of Multilingual In-Context Learning
Yilei Tu, Andrew Xue and Freda Shi 6213

SEK: Self-Explained Keywords Empower Large Language Models for Code Generation
Lishui Fan, Mouxiang Chen and Zhongxin Liu 6249

Why Not Act on What You Know? Unleashing Safety Potential of LLMs via Self-Aware Guard Enhancement
Peng Ding, Jun Kuang, ZongYu Wang, Xuezhi Cao, Xunliang Cai, Jiajun Chen and Shujian Huang 6279

Explorer: Scaling Exploration-driven Web Trajectory Synthesis for Multimodal Web Agents
Vardaan Pahuja, Yadong Lu, Corby Rosset, Boyu Gou, Arindam Mitra, Spencer Whitehead, Yu Su and Ahmed Hassan Awadallah 6300

Advancing General Multimodal Capability of Vision-language Models with Pyramid-descent Visual Position Encoding
Zhanpeng Chen, Mingxiao Li, Ziyang Chen, Nan du, Xiaolong Li and Yuexian Zou 6324

P-React: Synthesizing Topic-Adaptive Reactions of Personality Traits via Mixture of Specialized LoRA Experts
Yuhao Dan, Jie Zhou, Qin Chen, Junfeng Tian and Liang He 6342

EssayJudge: A Multi-Granular Benchmark for Assessing Automated Essay Scoring Capabilities of Multimodal Large Language Models
Jiamin Su, Yibo Yan, Fangteng FU, Zhang Han, Jingheng Ye, Xiang Liu, Jiahao Huo, Huiyu Zhou and Xuming Hu 6363

Streamlining the Collaborative Chain of Models into A Single Forward Pass in Generation-Based Tasks
Yuanjie Lyu, Chao Zhang, Yuhao Chen, Yong Chen and Tong Xu 6390

Self-Correction is More than Refinement: A Learning Framework for Visual and Language Reasoning Tasks
Jiayi He, Hehai Lin, Qingyun Wang, Yi R. Fung and Heng Ji 6405

Beyond Reactive Safety: Risk-Aware LLM Alignment via Long-Horizon Simulation
Chenkai Sun, Denghui Zhang, ChengXiang Zhai and Heng Ji 6422

Probability-Consistent Preference Optimization for Enhanced LLM Reasoning
Yunqiao Yang, Houxing Ren, Zimu Lu, Ke Wang, Weikang Shi, Aojun Zhou, Junting Pan, Mingjie Zhan and Hongsheng Li 6435

IW-Bench: Evaluating Large Multimodal Models for Converting Image-to-Web
Hongcheng Guo, Wei Zhang, Junhao Chen, Yaonan Gu, Jian Yang, Junjia Du, Shaosheng Cao, Binyuan Hui, Tianyu Liu, Jianxin Ma, Chang Zhou and Zhoujun Li 6449

TDCSA: LLM-Guided Top-Down Approach for Robust Citation Sentiment Analysis
Fan Gao, Jieyang Peng, Xiaoming Tao and Wang Youzheng 6467

DeepRTL2: A Versatile Model for RTL-Related Tasks
Yi Liu, Hongji Zhang, Yunhao Zhou, Zhengyuan Shi, Changran XU and Qiang Xu 6485

The Self-Improvement Paradox: Can Language Models Bootstrap Reasoning Capabilities without External Scaffolding?
Yutao Sun, Mingshuai Chen, Tiancheng Zhao, Ruochen Xu, Zilun Zhang and Jianwei Yin . 6501

Cross-lingual Multimodal Sentiment Analysis for Low-Resource Languages via Language Family Disentanglement and Rethinking Transfer
Long Chen, Shuoyu Guan, Xiaohua Huang, Wen-Jing Wang, Cai Xu, Ziyu Guan and Wei Zhao 6513

Does Chain-of-Thought Reasoning Really Reduce Harmfulness from Jailbreaking?
Chengda Lu, Xiaoyu Fan, Yu Huang, Rongwu Xu, Jijie Li and Wei Xu 6523

InternLM-XComposer2.5-Reward: A Simple Yet Effective Multi-Modal Reward Model
Yuhang Zang, Xiaoyi Dong, Pan Zhang, Yuhang Cao, Ziyu Liu, Shengyuan Ding, Shenxi Wu, Yubo Ma, Haodong Duan, Wenwei Zhang, Kai Chen, Dahua Lin and Jiaqi Wang 6547

RATE-Nav: Region-Aware Termination Enhancement for Zero-shot Object Navigation with Vision-Language Models
Junjie Li, Nan Zhang, Xiaoyang Qu, Kai Lu, Guokuan Li, Jiguang Wan and Jianzong Wang 6564

RMoA: Optimizing Mixture-of-Agents through Diversity Maximization and Residual Compensation
Zhentao Xie, Chengcheng Han, Jinxin Shi, Wenjun Cui, Xin Zhao, Xingjiao Wu and Jiabao Zhao 6575

Instruction-Tuning Data Synthesis from Scratch via Web Reconstruction
Yuxin Jiang, Yufei Wang, Chuhan Wu, Xinyi Dai, Yan Xu, Weinan Gan, Yasheng Wang, Xin Jiang, Lifeng Shang, Ruiming Tang and Wei Wang 6603

RLKGF: Reinforcement Learning from Knowledge Graph Feedback Without Human Annotations
Lian Yan, Chen Tang, Yi Guan, Haotian Wang, Songyuan Wang, Haifeng Liu, Yang Yang and Jingchi Jiang 6619

Learning Task Representations from In-Context Learning
Baturay Saglam, Xinyang Hu, Zhuoran Yang, Dionysis Kalogerias and Amin Karbasi 6634

CAVGAN: Unifying Jailbreak and Defense of LLMs via Generative Adversarial Attacks on their Internal Representations
Xiaohu Li, Yunfeng Ning, Zepeng Bao, Mayi Xu, Jianhao Chen and Tieyun Qian 6664

Firm or Fickle? Evaluating Large Language Models Consistency in Sequential Interactions
Yubo Li, Yidi Miao, Xueying Ding, Ramayya Krishnan and Rema Padman 6679

OS-Kairos: Adaptive Interaction for MLLM-Powered GUI Agents
Pengzhou Cheng, Zheng Wu, Zongru Wu, Tianjie Ju, Aston Zhang, Zhuosheng Zhang and Gongshen Liu 6701

Red-Teaming LLM Multi-Agent Systems via Communication Attacks
Pengfei He, Yuping Lin, Shen Dong, Han Xu, Yue Xing and Hui Liu 6726

Can We Trust AI Doctors? A Survey of Medical Hallucination in Large Language and Large Vision-Language Models
Zhihong Zhu, Yunyan Zhang, Xianwei Zhuang, Fan Zhang, Zhongwei Wan, Yuyan Chen, QingqingLong QingqingLong, Yefeng Zheng and Xian Wu 6748

DRT: Deep Reasoning Translation via Long Chain-of-Thought
Jiaan Wang, Fandong Meng, Yunlong Liang and Jie Zhou 6770

CTPD: Cross-Modal Temporal Pattern Discovery for Enhanced Multimodal Electronic Health Records Analysis
Fuying Wang, Feng Wu, Yihan Tang and Lequan Yu 6783

Vision-aided Unsupervised Constituency Parsing with Multi-MLLM Debating
Dong Zhang, Haiyan Tian, Qingying Sun and Shoushan Li 6800

Inter-Passage Verification for Multi-evidence Multi-answer QA
Bingsen Chen, Shenji Wan, Xi Ye and Chen Zhao 6811

PROMTEC: Fast LLM Inference Decoding using Prompt Multi-Lookup with Template Database and Common Sequences
Alan Chi-Man Lee, Wing-Sun Cheng and Calvin Chun-Kit Chan 6830

Logical DA: Enhancing Data Augmentation for Logical Reasoning via a Multi-Agent System
Haoqi Zheng, DongWang DongWang, Silin Yang, Yunpeng Qi, Ruochun Jin and Liyang Xu6843

Adapting General-Purpose Embedding Models to Private Datasets Using Keyword-based Retrieval
Yubai Wei, Jiale Han and Yi Yang 6856

SQL Injection Jailbreak: A Structural Disaster of Large Language Models
Jiawei Zhao, Kejiang Chen, Weiming Zhang and Nenghai Yu 6871

TAMP: Token-Adaptive Layerwise Pruning in Multimodal Large Language Models
Jaewoo Lee, Keyang Xuan, Chanakya Ekbote, Sandeep Polisetty, Yi R. Fung and Paul Pu Liang 6892

Generative Music Models' Alignment with Professional and Amateur Users' Expectations
Zihao Wang, Jiaxing Yu, Haoxuan Liu, Zehui Zheng, Yuhang Jin, Shuyu Li, Shulei Ji and Kejun Zhang 6909

LLM-Forest: Ensemble Learning of LLMs with Graph-Augmented Prompts for Data Imputation
Xinrui He, Yikun Ban, Jiaru Zou, Tianxin Wei, Curtiss Cook and Jingrui He 6921

Task Calibration: Calibrating Large Language Models on Inference Tasks
Yingjie Li, Yun Luo, Xiaotian Xie and Yue Zhang 6937

MiniELM: A Lightweight and Adaptive Query Rewriting Framework for E-Commerce Search Optimization
Duy A. Nguyen, Rishi Kesav Mohan, Shimeng Yang, Pritom Saha Akash and Kevin Chen-Chuan Chang 6952

Visibility as Survival: Generalizing NLP for Native Alaskan Language Identification
Ivory Yang, Chunhui Zhang, Yuxin Wang, Zhongyu Ouyang and Soroush Vosoughi 6965

KodCode: A Diverse, Challenging, and Verifiable Synthetic Dataset for Coding
Zhangchen Xu, Yang Liu, Yueqin Yin, Mingyuan Zhou and Radha Poovendran 6980

Select, Read, and Write: A Multi-Agent Framework of Full-Text-based Related Work Generation
Xiaochuan Liu, Ruihua Song, Xiting Wang and Xu Chen 7009

Graph-Assisted Culturally Adaptable Idiomatic Translation for Indic languages
Pratik Rakesh Singh, Kritarth Prasad, Mohammadi Zaki and Pankaj Wasnik 7029

Question Answering in Climate Adaptation for Agriculture: Model Development and Evaluation with Expert Feedback
Vincent Nguyen, Sarvnaz Karimi, Willow Hallgren and Mahesh Prakash 7045

AGRec: Adapting Autoregressive Decoders with Graph Reasoning for LLM-based Sequential Recommendation
Xinfeng Wang, Jin Cui, Fumiyo Fukumoto and Yoshimi Suzuki 7076

Causal Denoising Prototypical Network for Few-Shot Multi-label Aspect Category Detection
Jin Cui, Xinfeng Wang, Yoshimi Suzuki and Fumiyo Fukumoto 7091

RealHiTBench: A Comprehensive Realistic Hierarchical Table Benchmark for Evaluating LLM-Based Table Analysis
Pengzuo Wu, Yuhang Yang, Guangcheng Zhu, Chao Ye, Hong Gu, Xu Lu, Ruixuan Xiao, Bowen Bao, Yijing He, Liangyu Zha, Wentao Ye, Junbo Zhao and Haobo Wang 7105

A Query-Response Framework for Whole-Page Complex-Layout Document Image Translation with Relevant Regional Concentration
Zhiyang Zhang, Yaping Zhang, Yupu Liang, Zhiyuan Chen, Lu Xiang, Yang Zhao, Yu Zhou and Chengqing Zong 7138

DependEval: Benchmarking LLMs for Repository Dependency Understanding
Junjia Du, Yadi Liu, Hongcheng Guo, Jiawei Wang, Haojian Huang, Yunyi Ni and Zhoujun Li 7150

A General Knowledge Injection Framework for ICD Coding
Xu Zhang, Kun Zhang, Wenxin ma, Rongsheng Wang, Chenxu Wu, Yingtai Li and S Kevin Zhou 7180

MMUnlearner: Reformulating Multimodal Machine Unlearning in the Era of Multimodal Large Language Models
Jiahao Huo, Yibo Yan, Xu Zheng, Yuanhuiyi Lyu, Xin Zou, Zhihua Wei and Xuming Hu . . . 7190

Generating Questions, Answers, and Distractors for Videos: Exploring Semantic Uncertainty of Object Motions
Wenjian Ding, Yao Zhang, Jun Wang, Adam Jatowt and Zhenglu Yang 7207

DiffSkip: Differential Layer Skipping in Large Language Models
Xuan Luo, Weizhi Wang and Xifeng Yan 7221

Towards Explainable Temporal Reasoning in Large Language Models: A Structure-Aware Generative Framework
Zihao Jiang, Ben Liu, Miao Peng, Wenjie Xu, Yao Xiao, Zhenyan Shan and Min Peng 7232

A Bounding Box is Worth One Token - Interleaving Layout and Text in a Large Language Model for Document Understanding
Jinghui Lu, Haiyang Yu, Yanjie Wang, Yongjie Ye, Jingqun Tang, Ziwei Yang, Binghong Wu, Qi Liu, Hao Feng, Han Wang, Hao Liu and Can Huang 7252

Self-Foveate: Enhancing Diversity and Difficulty of Synthesized Instructions from Unsupervised Text via Multi-Level Foveation
Mingzhe Li, Xin Lu and Yanyan Zhao 7274

TableDreamer: Progressive and Weakness-guided Data Synthesis from Scratch for Table Instruction Tuning
Mingyu Zheng, Zhifan Feng, Jia Wang, Lanrui Wang, Zheng Lin, Hao Yang and Weiping Wang 7290

Konooz: Multi-domain Multi-dialect Corpus for Named Entity Recognition
Nagham Hamad, Mohammed Khalilia and Mustafa Jarrar 7316

Self-Rewarding Large Vision-Language Models for Optimizing Prompts in Text-to-Image Generation
Hongji Yang, Yucheng Zhou, Wencheng Han and Jianbing Shen 7332

CodeV: Issue Resolving with Visual Data
Linhao Zhang, Daoguang Zan, Quanshun Yang, Zhirong Huang, Dong Chen, Bo Shen, Tianyu Liu, Yongshun Gong, Huang Pengjie, Xudong Lu, Guangtai Liang, Lizhen Cui and Qianxiang Wang 7350

A Survey of Large Language Models in Psychotherapy: Current Landscape and Future Directions
Hongbin Na, Yining Hua, Zimu Wang, Tao Shen, Beibei Yu, Lilin Wang, Wei Wang, John Torous and Ling Chen 7362

Breaking the Reasoning Barrier A Survey on LLM Complex Reasoning through the Lens of Self-Evolution
Tao He, Hao Li, Jingchang Chen, Runxuan Liu, Yixin Cao, Lizi Liao, Zihao Zheng, Zheng Chu, Jiafeng Liang, Ming Liu and Bing Qin 7377

SEE: Continual Fine-tuning with Sequential Ensemble of Experts
Zhilin Wang, Yafu Li, Xiaoye Qu and Yu Cheng 7418

Boosting Policy and Process Reward Models with Monte Carlo Tree Search in Open-Domain QA
Chi-Min Chan, Chunpu Xu, Junqi Zhu, Jiaming Ji, Donghai Hong, Pengcheng Wen, Chunyang Jiang, Zhen Ye, Yaodong Yang, Wei Xue, Sirui Han and Yike Guo 7433

Investigating and Enhancing Vision-Audio Capability in Omnimodal Large Language Models
Rui Hu, Delai Qiu, Shuyu Wei, Jiaming Zhang, Yining Wang, Shengping Liu and Jitao Sang7452

OpenHuEval: Evaluating Large Language Model on Hungarian Specifics
Haote Yang, Xingjian Wei, Jiang Wu, Noémi Ligeti-Nagy, Jiaxing Sun, Yinfan Wang, Győző Zijian Yang, Junyuan Gao, Jingchao Wang, Bowen Jiang, Shasha Wang, Nanjun Yu, Zihao Zhang, Shixin Hong, Hongwei Liu, Wei Li, Songyang Zhang, Dahua Lin, Lijun Wu, Gábor Prószéky and Conghui He 7464

StructFact: Reasoning Factual Knowledge from Structured Data with Large Language Models
Sirui Huang, Yanggan Gu, Zhonghao Li, Xuming Hu, Li Qing and Guandong Xu 7521

From Imitation to Introspection: Probing Self-Consciousness in Language Models
Sirui Chen, Shu Yu, Shengjie Zhao and Chaochao Lu 7553

DocFusion: A Unified Framework for Document Parsing Tasks
Mingxu Chai, Ziyu Shen, Chong Zhang, Yue Zhang, Xiao Wang, Shihan Dou, Jihua Kang, Jiazheng Zhang and Qi Zhang 7584

Hierarchical Safety Realignment: Lightweight Restoration of Safety in Pruned Large Vision-Language Models
Yue Li, Xin Yi, Dongsheng Shi, Gerard De Melo, Xiaoling Wang and Linlin Wang 7600

LongDPO: Unlock Better Long-form Generation Abilities for LLMs via Critique-augmented Stepwise Information
Bowen Ping, Jiali Zeng, Fandong Meng, Shuo Wang, Jie Zhou and Shanghang Zhang 7613

Reinforcing Compositional Retrieval: Retrieving Step-by-Step for Composing Informative Contexts
Quanyu Long, Jianda Chen, Zhengyuan Liu, Nancy F. Chen, Wenya Wang and Sinno Jialin Pan 7633

Towards A Better Initial Policy Model For Scalable Long-CoT Reinforcement Learning
Bofei Gao, Yejie Wang, Yibo Miao, Ruoyu Wu, Feifan Song, Longhui Yu, Tianyu Liu and Baobao Chang 7652

Topic Modeling for Short Texts via Optimal Transport-Based Clustering
Tu Vu, Manh Do, Tung Nguyen, Linh Ngo Van, Sang Dinh and Thien Huu Nguyen 7666

Lemmatisation & Morphological Analysis of Unedited Greek: Do Simple Tasks Need Complex Solutions?
Colin Swaelens, Ilse De Vos and Els Lefever 7681

FRAME: Feedback-Refined Agent Methodology for Enhancing Medical Research Insights
Chengzhang Yu, Yiming Zhang, Zhixin Liu, Zenghui Ding, Yining Sun and Zhanpeng Jin . . 7690

Chain-of-Scrutiny: Detecting Backdoor Attacks for Large Language Models
Xi Li, Ruofan Mao, Yusen Zhang, Renze Lou, Chen Wu and Jiaqi Wang 7705

Relevance Scores Calibration for Ranked List Truncation via TMP Adapter
Pavel Posokhov, Sergei Masliukhin, Skrylnikov Stepan, Danil Tirskikh and Olesia Makhnytkina 7728

Neuron Activation Modulation for Text Style Transfer: Guiding Large Language Models
Chaona Kong, Jianyi Liu, Yifan Tang and Ru Zhang 7735

MTVQA: Benchmarking Multilingual Text-Centric Visual Question Answering
Jingqun Tang, Qi Liu, Yongjie Ye, Jinghui Lu, Shu Wei, An-Lan Wang, Chunhui Lin, Hao Feng, Zhen Zhao, Yanjie Wang, Yuliang Liu, Hao Liu, Xiang Bai and Can Huang 7748

HICD: Hallucination-Inducing via Attention Dispersion for Contrastive Decoding to Mitigate Hallucinations in Large Language Models
Xinyan Jiang, Hang Ye, Yongxin Zhu, Xiaoying Zheng, Zikang Chen and Jun Gong 7764

Understanding the Repeat Curse in Large Language Models from a Feature Perspective
Junchi Yao, Shu Yang, Jianhua Xu, Lijie Hu, Mengdi Li and Di Wang 7787

Code-Switching Curriculum Learning for Multilingual Transfer in LLMs
Haneul Yoo, Cheonbok Park, Sangdoo Yun, Alice Oh and Hwaran Lee 7816

A Mousetrap: Fooling Large Reasoning Models for Jailbreak with Chain of Iterative Chaos
Yang Yao, Xuan Tong, Ruofan Wang, Yixu Wang, Lujundong Li, Liang Liu, Yan Teng and Yingchun Wang . . . 7837

Tag-Evol: Achieving Efficient Instruction Evolving via Tag Injection
Yixuan Wang, Shiqi Zhou, Chuanzhe Guo and Qingfu Zhu . . . 7856

Breaking the Ceiling: Exploring the Potential of Jailbreak Attacks through Expanding Strategy Space
Yao Huang, Yitong Sun, Shouwei Ruan, Yichi Zhang, Yinpeng Dong and Xingxing Wei . . . 7870

GeNRe: A French Gender-Neutral Rewriting System Using Collective Nouns
Enzo Doyen and Amalia Todirascu . . . 7889

LGAR: Zero-Shot LLM-Guided Neural Ranking for Abstract Screening in Systematic Literature Reviews
Christian Jaumann, Andreas Wiedholz and Annemarie Friedrich . . . 7910

LCHAIM - Investigating Long Context Reasoning in Hebrew
Ehud Malul, Oriel Perets, Ziv Mor, Yigal Kassel and Elior Sulem . . . 7928

CLeVeR: Multi-modal Contrastive Learning for Vulnerability Code Representation
Jiayuan Li, Lei Cui, Sen Zhao, Yun Yang, Lun Li and Hongsong Zhu . . . 7940

MEMIT-Merge: Addressing MEMIT's Key-Value Conflicts in Same-Subject Batch Editing for LLMs
Zilu Dong, Xiangqing Shen and Rui Xia . . . 7952

Large Language Models for Predictive Analysis: How Far Are They?
Qin Chen, Yuanyi Ren, Xiaojun Ma and Yuyang Shi . . . 7961

Think More, Hallucinate Less: Mitigating Hallucinations via Dual Process of Fast and Slow Thinking
Xiaoxue Cheng, Junyi Li, Xin Zhao and Ji-Rong Wen . . . 7979

Towards Adaptive Memory-Based Optimization for Enhanced Retrieval-Augmented Generation
Qitao Qin, Yucong Luo, Yihang Lu, Zhibo Chu, Xiaoman Liu and Xianwei Meng . . . 7991

Enhancing Cross-Tokenizer Knowledge Distillation with Contextual Dynamical Mapping
Yijie Chen, Yijin Liu, Fandong Meng, Yufeng Chen, Jinan Xu and Jie Zhou . . . 8005

A Semantic-Aware Layer-Freezing Approach to Computation-Efficient Fine-Tuning of Language Models
Jian Gu, Aldeida Aleti, Chunyang Chen and Hongyu Zhang . . . 8019

CNNSum: Exploring Long-Context Summarization with Large Language Models in Chinese Novels
Lingxiao Wei, He Yan, Lu Xiangju, Junmin Zhu, Jun Wang and Wei Zhang . . . 8034

Document Segmentation Matters for Retrieval-Augmented Generation
Zhitong Wang, Cheng Gao, Chaojun Xiao, Yufei Huang, Shuzheng Si, Kangyang Luo, Yuzhuo Bai, Wenhao Li, Tangjian Duan, Chuancheng Lv, Guoshan Lu, Gang Chen, Fanchao Qi and Maosong Sun . . . 8063

UBench: Benchmarking Uncertainty in Large Language Models with Multiple Choice Questions
Xunzhi Wang, Zhuowei Zhang, Gaonan Chen, Qiongyu Li, Bitong Luo, Zhixin Han, Haotian Wang, Zhiyu li, Hang Gao and Mengting Hu . . . 8076

Embracing Large Language Models in Traffic Flow Forecasting
Yusheng Zhao, Xiao Luo, Haomin Wen, Zhiping Xiao, Wei Ju and Ming Zhang . . . 8108

Flow2Code: Evaluating Large Language Models for Flowchart-based Code Generation Capability
Mengliang He, Jiayi Zeng, Yankai Jiang, Wei Zhang, Zeming Liu, Xiaoming Shi and Aimin Zhou 8124

Smarter, Not Harder: Training-Free Adaptive Computation for Transformers
Romain Storaï, Jaeseong Lee and Seung-won Hwang 8147

UCS-SQL: Uniting Content and Structure for Enhanced Semantic Bridging In Text-to-SQL
Zhenhe Wu, Zhongqiu Li, JieZhangChinaTele JieZhangChinaTele, Zhongjiang He, Jian Yang, Yu Zhao, Ruiyu Fang, Bing Wang, Hongyan Xie, Shuangyong Song and Zhoujun Li 8156

CodePRM: Execution Feedback-enhanced Process Reward Model for Code Generation
Qingyao Li, Xinyi Dai, Xiangyang Li, Weinan Zhang, Yasheng Wang, Ruiming Tang and Yong Yu 8169

STEM-POM: Evaluating Language Models Math-Symbol Reasoning in Document Parsing
Jiaru Zou, Qing Wang, Pratyush Thakur and Nickvash Kani 8183

Retrieval Visual Contrastive Decoding to Mitigate Object Hallucinations in Large Vision-Language Models
Jihoon Lee and Min Song 8200

Leveraging LLMs for Bangla Grammar Error Correction: Error Categorization, Synthetic Data, and Model Evaluation
Pramit Bhattacharyya and Arnab Bhattacharya 8220

Think Both Ways: Teacher-Student Bidirectional Reasoning Enhances MCQ Generation and Distractor Quality
Yimiao Qiu, Yang Deng, Quanming Yao, Zhimeng Zhang, Zhiang Dong, Chang Yao and Jingyuan Chen 8240

mmE5: Improving Multimodal Multilingual Embeddings via High-quality Synthetic Data
Haonan Chen, Liang Wang, Nan Yang, Yutao Zhu, Ziliang Zhao, Furu Wei and Zhicheng Dou 8254

Word2Passage: Word-level Importance Re-weighting for Query Expansion
Jeonghwan Choi, Minjeong Ban, Minseok Kim and Hwanjun Song 8276

MECoT: Markov Emotional Chain-of-Thought for Personality-Consistent Role-Playing
Yangbo Wei, Zhen Huang, Fangzhou Zhao, Qi Feng and Wei W. Xing 8297

FiDeLiS: Faithful Reasoning in Large Language Models for Knowledge Graph Question Answering
Yuan Sui, Yufei He, Nian Liu, Xiaoxin He, Kun Wang and Bryan Hooi 8315

REALM: A Dataset of Real-World LLM Use Cases
Jingwen Cheng, Kshitish Ghate, Wenyue Hua, William Yang Wang, Hong Shen and Fei Fang 8331

BABELEDITS: A Benchmark and a Modular Approach for Robust Cross-lingual Knowledge Editing of Large Language Models
Tommaso Green, Félix Gaschi, Fabian David Schmidt, Simone Paolo Ponzetto and Goran Glavaš 8342

CDS: Data Synthesis Method Guided by Cognitive Diagnosis Theory
Haokun Zhao, Jinyi Han, Jiaqing Liang, Yanghua Xiao, Xiaojun Meng and Jiansheng Wei . 8370

Problem-Solving Logic Guided Curriculum In-Context Learning for LLMs Complex Reasoning
Xuetao Ma, Wenbin Jiang and Hua Huang 8394

BESSTIE: A Benchmark for Sentiment and Sarcasm Classification for Varieties of English
Dipankar Srirag, Aditya Joshi, Jordan Painter and Diptesh Kanojia 8413

NavRAG: Generating User Demand Instructions for Embodied Navigation through Retrieval-Augmented LLM
Zihan Wang, Yaohui Zhu, Gim Hee Lee and Yachun Fan 8430

SQLForge: Synthesizing Reliable and Diverse Data to Enhance Text-to-SQL Reasoning in LLMs
Yu Guo, Dong Jin, Shenghao Ye, Shuangwu Chen, Jianyang Jianyang and Xiaobin Tan 8441

Retrieval-Augmented Process Reward Model for Generalizable Mathematical Reasoning
Jiachen Zhu, Congmin Zheng, Jianghao Lin, Kounianhua Du, Ying Wen, Yong Yu, Jun Wang and Weinan Zhang 8453

Contrastive Learning for Task-Independent SpeechLLM-Pretraining
Maike Züfle and Jan Niehues 8469

QiMeng-Attention: SOTA Attention Operator is generated by SOTA Attention Algorithm
Qirui Zhou, Shaohui Peng, Weiqiang Xiong, Haixin Chen, Yuanbo Wen, Haochen Li, Ling Li, Qi Guo, Yongwei Zhao, KE Gao, Ruizhi Chen, Yanjun Wu, Zhao Chen and Yunji Chen 8491

ALW: Adaptive Layer-Wise contrastive decoding enhancing reasoning ability in Large Language Models
Yuechi Zhou, Chuyue Zhou, Jianxin Zhang, Juntao Li and Min Zhang 8506

Mixture of Decoding: An Attention-Inspired Adaptive Decoding Strategy to Mitigate Hallucinations in Large Vision-Language Models
Xinlong Chen, Yuanxing Zhang, Qiang Liu, Junfei Wu, Fuzheng Zhang and Tieniu Tan 8525

VidCapBench: A Comprehensive Benchmark of Video Captioning for Controllable Text-to-Video Generation
Xinlong Chen, Yuanxing Zhang, Chongling Rao, Yushuo Guan, Jiaheng Liu, Fuzheng Zhang, Chengru Song, Qiang Liu, Di Zhang and Tieniu Tan 8543

Mitigating Demonstration Bias through Global Coevolutionary Reasoning
Chuan Gou, Bangwei Li, Jianhua Dai, Xiaoyang Han and Ming Cai 8564

A Representation Level Analysis of NMT Model Robustness to Grammatical Errors
Abderrahmane Issam, Yusuf Can Semerci, Jan Scholtes and Gerasimos Spanakis 8579

T^2DR: A Two-Tier Deficiency-Resistant Framework for Incomplete Multimodal Learning
Han Lin, Xiu Tang, Huan Li, Wenxue Cao, Sai Wu, Chang Yao, Lidan Shou and Gang Chen8602

From Specific-MLLMs to Omni-MLLMs: A Survey on MLLMs Aligned with Multi-modalities
Shixin Jiang, Jiafeng Liang, Jiyuan Wang, Xuan Dong, Heng Chang, Weijiang Yu, Jinhua Du, Ming Liu and Bing Qin 8617

Analyzing the Effect of Linguistic Similarity on Cross-Lingual Transfer: Tasks and Experimental Setups Matter
Verena Blaschke, Masha Fedzechkina and Maartje Ter Hoeve 8653

Agents generalize to novel levels of abstraction by using adaptive linguistic strategies
Kristina Kobrock, Xenia Ohmer, Elia Bruni and Nicole Gotzner 8685

The Linguistic Connectivities Within Large Language Models
Dan Wang, Boxi Cao, Ning Bian, Xuanang Chen, Yaojie Lu, Hongyu Lin, Jia Zheng, Le Sun, Shanshan Jiang, Bin Dong and Xianpei Han 8700

XFinBench: Benchmarking LLMs in Complex Financial Problem Solving and Reasoning
Zhihan Zhang, Yixin Cao and Lizi Liao 8715

Align^{2}LLaVA: Cascaded Human and Large Language Model Preference Alignment for Multi-modal Instruction Curation
Hongzhe Huang, Jiang Liu, Zhewen Yu, Li Cai, Dian Jiao, Wenqiao Zhang, Siliang Tang, Juncheng Li, Hao Jiang, Haoyuan Li and Yueting Zhuang 8759

Achieving binary weight and activation for LLMs using Post-Training Quantization
Siqing Song, Chuang Wang, Rui-Qi Wang, Yi Yang and Xu-Yao Zhang 8782

Mitigating Negative Interference in Multilingual Knowledge Editing through Null-Space Constraints
Wei Sun, Tingyu Qu, Mingxiao Li, Jesse Davis and Marie-Francine Moens 8796

From Awareness to Adaptability: Enhancing Tool Utilization for Scientific Reasoning
Wenjing Xie, Xiaobo Liang, Juntao Li, Wanfu Wang, Kehai Chen, Qiaoming Zhu and Min Zhang 8811

AMoPO: Adaptive Multi-objective Preference Optimization without Reward Models and Reference Models
Qi Liu, Jingqing Ruan, Hao Li, Haodong Zhao, Desheng Wang, Jiansong Chen, Wan Guanglu, Xunliang Cai, Zhi Zheng and Tong Xu 8832

Supervised Optimism Correction: Be Confident When LLMs Are Sure
Junjie Zhang, Rushuai Yang, Shunyu Liu, Ting-En Lin, Fei Huang, Yi Chen, Yongbin Li and Dacheng Tao 8867

Offline Reinforcement Learning for LLM Multi-step Reasoning
Huaijie Wang, Shibo Hao, Hanze Dong, Shenao Zhang, Yilin Bao, Ziran Yang and Yi Wu . . 8881

Sampling-based Pseudo-Likelihood for Membership Inference Attacks
Masahiro Kaneko, Youmi Ma, Yuki Wata and Naoaki Okazaki 8894

AgentStore: Scalable Integration of Heterogeneous Agents As Specialized Generalist Computer Assistant
Chengyou Jia, Minnan Luo, Zhuohang Dang, Qiushi Sun, Fangzhi Xu, Junlin Hu, Tianbao Xie and Zhiyong Wu 8908

Boosting Vulnerability Detection of LLMs via Curriculum Preference Optimization with Synthetic Reasoning Data
Xin-Cheng Wen, Yijun Yang, Cuiyun Gao, Yang Xiao and Deheng Ye 8935

$GA - S^3$: Comprehensive Social Network Simulation with Group Agents
Yunyao Zhang, Zikai Song, Hang Zhou, Wenfeng Ren, Yi-Ping Phoebe Chen, Junqing Yu and Wei Yang 8950

M-RangeDetector: Enhancing Generalization in Machine-Generated Text Detection through Multi-Range Attention Masks
Kaijie Jiao, Quan Wang, Licheng Zhang, Zikang Guo and Zhendong Mao 8971

Does Your Voice Assistant Remember? Analyzing Conversational Context Recall and Utilization in Voice Interaction Models
Heeseung Kim, Che Hyun Lee, Sangkwon Park, Jiheum Yeom, Nohil Park, Sangwon Yu and Sungroh Yoon 8984

NeuronMerge: Merging Models via Functional Neuron Groups
Wangyun Gu, Qianghua Gao, Zhang Li-Xin, Xu Shen and Jieping Ye 9015

HellaSwag-Pro: A Large-Scale Bilingual Benchmark for Evaluating the Robustness of LLMs in Commonsense Reasoning
Xiaoyuan Li, Moxin Li, Rui Men, Yichang Zhang, Keqin Bao, Wenjie Wang, Fuli Feng, Dayiheng Liu and Junyang Lin 9038

Self-Steering Optimization: Autonomous Preference Optimization for Large Language Models
Hao Xiang, Bowen Yu, Hongyu Lin, Keming Lu, Yaojie Lu, Xianpei Han, Ben He, Le Sun, Jingren Zhou and Junyang Lin 9073

LIME: Less Is More for MLLM Evaluation
King Zhu, Qianbo Zang, Shian Jia, Siwei Wu, Feiteng Fang, Yizhi LI, Shuyue Guo, Tianyu Zheng, Jiawei Guo, Bo Li, Haoning Wu, Xingwei Qu, Jian Yang, Ruibo Liu, Xiang Yue, Jiaheng Liu, Chenghua Lin, Hamid Alinejad-Rokny, Min Yang, Shiwen Ni, Wenhao Huang and Ge Zhang . . . 9086

Debate, Reflect, and Distill: Multi-Agent Feedback with Tree-Structured Preference Optimization for Efficient Language Model Enhancement
Xiaofeng Zhou, Heyan Huang and Lizi Liao 9122

CodeReviewQA: The Code Review Comprehension Assessment for Large Language Models
Hong Yi Lin, Chunhua Liu, Haoyu Gao, Patanamon Thongtanunam and Christoph Treude . 9138

Narrative Media Framing in Political Discourse
Yulia Otmakhova and Lea Frermann 9167

MHALO: Evaluating MLLMs as Fine-grained Hallucination Detectors
Yishuo Cai, Renjie Gu, Jiaxu Li, Xuancheng Huang, Junzhe Chen, Xiaotao Gu and Minlie Huang 9197

Semantic Topology: a New Perspective for Communication Style Characterization
Barbara Scalvini and Alireza Mashaghi 9223

Decoding LLM Personality Measurement: Forced-Choice vs. Likert
Xiaoyu Li, Haoran Shi, Zengyi Yu, Yukun Tu and Chanjin Zheng 9234

MultiMSD: A Corpus for Multilingual Medical Text Simplification from Online Medical References
Koki Horiguchi, Tomoyuki Kajiwara, Takashi Ninomiya, Shoko Wakamiya and Eiji Aramaki9248

BadWindtunnel: Defending Backdoor in High-noise Simulated Training with Confidence Variance
Ruyi Zhang, Songlei Jian, Yusong Tan, Heng Gao, Haifang Zhou and Kai Lu 9259

Multimodal Machine Translation with Text-Image In-depth Questioning
Yue Gao, Jing Zhao, Shiliang Sun, Xiaosong Qiao, Tengfei Song and Hao Yang 9274

ReKG-MCTS: Reinforcing LLM Reasoning on Knowledge Graphs via Training-Free Monte Carlo Tree Search
Xiaozhuang Song, Shufei Zhang and Tianshu Yu 9288

HTML: Hierarchical Topology Multi-task Learning for Semantic Parsing in Knowledge Base Question Answering
Aziguli Wulamu, Lyu Zhengyu, Kaiyuan Gong, Yu Han, Zewen Wang, Zhihong Zhu and Bowen Xing 9307

StructFlowBench: A Structured Flow Benchmark for Multi-turn Instruction Following
Jinnan Li, Jinzhe Li, Yue Wang, Yi Chang and Yuan Wu 9322

CMIE: Combining MLLM Insights with External Evidence for Explainable Out-of-Context Misinformation Detection
Fanxiao Li, Jiaying Wu, Canyuan He and Wei Zhou 9342

EtiCor++: Towards Understanding Etiquettical Bias in LLMs
Ashutosh Dwivedi, Siddhant Shivdutt Singh and Ashutosh Modi 9355

FinRipple: Aligning Large Language Models with Financial Market for Event Ripple Effect Awareness
Yuanjian Xu, Jianing Hao, Kunsheng Tang, Jingnan Chen, Anxian Liu, Peng Liu and Guang Zhang 9377

Beyond Decoder-only: Large Language Models Can be Good Encoders for Machine Translation
Yingfeng Luo, Tong Zheng, Yongyu Mu, Bei Li, Qinghong Zhang, Yongqi Gao, Ziqiang Xu, Peinan Feng, Xiaoqian Liu, Tong Xiao and JingBo Zhu 9399

EC-RAFT: Automated Generation of Clinical Trial Eligibility Criteria through Retrieval-Augmented Fine-Tuning
Nopporn Lekuthai, Nattawit Pewngam, Supitcha Sokrai and Titipat Achakulvisut 9432

Pitfalls of Scale: Investigating the Inverse Task of Redefinition in Large Language Models
Elena Stringli, Maria Lymperaiou, Giorgos Filandrianos, Athanasios Voulodimos and Giorgos Stamou 9445

Implicit Reasoning in Transformers is Reasoning through Shortcuts
Tianhe Lin, Jian Xie, Siyu Yuan and Deqing Yang 9470

Learning to Align Multi-Faceted Evaluation: A Unified and Robust Framework
Kaishuai Xu, Tiezheng YU, Yi Cheng, Wenjun Hou, Liangyou Li, Xin Jiang, Lifeng Shang, Qun Liu and Wenjie Li 9488

CortexDebate: Debating Sparsely and Equally for Multi-Agent Debate
Yiliu Sun, Zicheng Zhao, Sheng Wan and Chen Gong 9503

PAP2PAT: Benchmarking Outline-Guided Long-Text Patent Generation with Patent-Paper Pairs
Valentin Knappich, Anna Hätty, Simon Razniewski and Annemarie Friedrich 9524

Debt Collection Negotiations with Large Language Models: An Evaluation System and Optimizing Decision Making with Multi-Agent
Xiaofeng Wang, Zhixin Zhang, Jin Guang Zheng, Yiming Ai and Rui Wang 9555

Focused-DPO: Enhancing Code Generation Through Focused Preference Optimization on Error-Prone Points
Kechi Zhang, Ge Li, Jia Li, Yihong Dong, Jia Li and Zhi Jin 9578

Supervised and Unsupervised Probing of Shortcut Learning: Case Study on the Emergence and Evolution of Syntactic Heuristics in BERT
Elke Vandermeerschen and Miryam De Lhoneux 9592

GIMMICK: Globally Inclusive Multimodal Multitask Cultural Knowledge Benchmarking
Florian Schneider, Carolin Holtermann, Chris Biemann and Anne Lauscher 9605

R-VLM: Region-Aware Vision Language Model for Precise GUI Grounding
Joonhyung Park, Peng Tang, Sagnik Das, Srikar Appalaraju, Kunwar Yashraj Singh, R. Manmatha and Shabnam Ghadar 9669

Perspective Transition of Large Language Models for Solving Subjective Tasks
Xiaolong Wang, Yuanchi Zhang, Ziyue Wang, Yuzhuang Xu, Fuwen Luo, Yile Wang, Peng Li and Yang Liu 9686

TripTailor: A Real-World Benchmark for Personalized Travel Planning
Kaimin Wang, Yuanzhe Shen, Changze Lv, Xiaoqing Zheng and Xuanjing Huang 9705

Random Splitting Negatively Impacts NER Evaluation: Quantifying and Eliminating the Overestimation of NER Performance
Florian Babl, Moritz Hennen, Jakob Murauer and Michaela Geierhos 9724

Structure-adaptive Adversarial Contrastive Learning for Multi-Domain Fake News Detection
Lingwei Wei, Dou Hu, Wei Zhou, Philip S. Yu and Songlin Hu 9739

BiasGuard: A Reasoning-Enhanced Bias Detection Tool for Large Language Models
Zhiting Fan, Ruizhe Chen and Zuozhu Liu 9753

Qorǵau: Evaluating Safety in Kazakh-Russian Bilingual Contexts
Maiya Goloburda, Nurkhan Laiyk, Diana Turmakhan, Yuxia Wang, Mukhammed Togmanov, Jonibek Mansurov, Askhat Sametov, Nurdaulet Mukhituly, Minghan Wang, Daniil Orel, Zain Muhammad Mujahid, Fajri Koto, Timothy Baldwin and Preslav Nakov 9765

MMXU: A Multi-Modal and Multi-X-ray Understanding Dataset for Disease Progression
Linjie Mu, Zhongzhen Huang, Shengqian Qin, Yakun Zhu, Shaoting Zhang and Xiaofan Zhang 9785

Tree-of-Code: A Self-Growing Tree Framework for End-to-End Code Generation and Execution in Complex Tasks
Ziyi Ni, Yifan LI, Ning Yang, Dou Shen, Pin Lyu and Daxiang Dong 9804

Akan Cinematic Emotions (ACE): A Multimodal Multi-party Dataset for Emotion Recognition in Movie Dialogues
David Sasu, Zehui wu, Ziwei Gong, Run Chen, Pengyuan Shi, Lin Ai, Julia Hirschberg and Natalie Schluter 9820

A Cognitive Writing Perspective for Constrained Long-Form Text Generation
Kaiyang Wan, Honglin Mu, Rui Hao, Haoran Luo, Tianle Gu and Xiuying Chen 9832

Migician: Revealing the Magic of Free-Form Multi-Image Grounding in Multimodal Large Language Models
You Li, Heyu Huang, Chi Chen, Kaiyu Huang, Chao Huang, Zonghao Guo, Zhiyuan Liu, Jinan Xu, Yuhua Li, Ruixuan Li and Maosong Sun 9845

SIKeD: Self-guided Iterative Knowledge Distillation for Mathematical Reasoning
Shivam Adarsh, Kumar Shridhar, Caglar Gulcehre, Nicholas Monath and Mrinmaya Sachan9868

Chain of Attack: Hide Your Intention through Multi-Turn Interrogation
Xikang Yang, Biyu Zhou, Xuehai Tang, Jizhong Han and Songlin Hu 9881

MIG: Automatic Data Selection for Instruction Tuning by Maximizing Information Gain in Semantic Space
Yicheng Chen, Yining Li, Kai Hu, Ma Zerun, HaochenYe HaochenYe and Kai Chen 9902

Enhancing Automatic Term Extraction with Large Language Models via Syntactic Retrieval
Yongchan Chun, Minhyuk Kim, Dongjun Kim, Chanjun Park and Heuiseok Lim 9916

Explainable Depression Detection in Clinical Interviews with Personalized Retrieval-Augmented Generation
Linhai Zhang, Ziyang Gao, Deyu Zhou and Yulan He 9927

EMPEC: A Comprehensive Benchmark for Evaluating Large Language Models Across Diverse Healthcare Professions
Zheheng Luo, Chenhan Yuan, Qianqian Xie and Sophia Ananiadou 9945

Beyond Numeric Rewards: In-Context Dueling Bandits with LLM Agents
Fanzeng Xia, Hao Liu, Yisong Yue and Tongxin Li 9959

"Well, Keep Thinking": Enhancing LLM Reasoning with Adaptive Injection Decoding
Hyunbin Jin, Je Won Yeom, Seunghyun Bae and Taesup Kim 9989

SpeechT-RAG: Reliable Depression Detection in LLMs with Retrieval-Augmented Generation Using Speech Timing Information
Xiangyu Zhang, Hexin Liu, Qiquan Zhang, Beena Ahmed and Julien Epps 10019

Fine-grained Knowledge Enhancement for Retrieval-Augmented Generation
Jingxuan Han, Zhendong Mao, Yi Liu, Yexuan Che, Zheren Fu and Quan Wang 10031

Bayesian Optimization for Controlled Image Editing via LLMs
Chengkun Cai, Haoliang Liu, Xu Zhao, Zhongyu Jiang, Tianfang Zhang, Zongkai Wu, John Lee, Jenq-Neng Hwang and Lei Li 10045

SPOT: Zero-Shot Semantic Parsing Over Property Graphs
Francesco Cazzaro, Justin Kleindienst, Sofia Márquez Gomez and Ariadna Quattoni 10057

Reasoning Circuits in Language Models: A Mechanistic Interpretation of Syllogistic Inference
Geonhee Kim, Marco Valentino and Andre Freitas 10074

Multi-Hop Question Generation via Dual-Perspective Keyword Guidance
Maodong Li, Longyin Zhang and Fang Kong 10096

LoRMA: Low-Rank Multiplicative Adaptation for LLMs
Harsh Bihany, Shubham Patel and Ashutosh Modi 10113

DI-BENCH: Benchmarking Large Language Models on Dependency Inference with Testable Repositories at Scale
Linghao Zhang, Junhao Wang, Shilin He, Chaoyun Zhang, Yu Kang, Bowen Li, Jiaheng Wen, Chengxing Xie, Maoquan Wang, Yufan Huang, Elsie Nallipogu, Qingwei Lin, Yingnong Dang, Saravan Rajmohan, Dongmei Zhang and Qi Zhang 10134

Weak-to-Strong Honesty Alignment via Learning-to-Rank Supervision
YunfanXie YunfanXie, Lixin Zou, Dan Luo, Min Tang and Chenliang Li 10154

MultiHoax: A Dataset of Multi-hop False-premise questions
Mohammadamin Shafiei, Hamidreza Saffari and Nafise Sadat Moosavi 10169

Learning to Play Like Humans: A Framework for LLM Adaptation in Interactive Fiction Games
Jinming Zhang and Yunfei Long 10188

STATE ToxiCN: A Benchmark for Span-level Target-Aware Toxicity Extraction in Chinese Hate Speech Detection
Zewen Bai, Liang Yang, Shengdi Yin, Junyu Lu, Jingjie Zeng, Haohao Zhu, Yuanyuan Sun and Hongfei Lin 10206

RelEdit: Evaluating Conceptual Knowledge Editing in Language Models via Relational Reasoning
Yifan Niu, Miao Peng, Nuo Chen, Yatao Bian, Tingyang Xu and Jia Li 10220

Unlocking Speech Instruction Data Potential with Query Rewriting
Yonghua Hei, Yibo Yan, Shuliang Liu, Huiyu Zhou, Linfeng Zhang and Xuming Hu 10239

From Evasion to Concealment: Stealthy Knowledge Unlearning for LLMs
Tianle Gu, Kexin Huang, Ruilin Luo, Yuanqi Yao, Xiuying Chen, Yujiu Yang, Yan Teng and Yingchun Wang 10261

Context-DPO: Aligning Language Models for Context-Faithfulness
Baolong Bi, Shaohan Huang, Yiwei Wang, Tianchi Yang, Zihan Zhang, Haizhen Huang, Lingrui Mei, Junfeng Fang, Zehao Li, Furu Wei, Weiwei Deng, Feng Sun, Qi Zhang and Shenghua Liu . 10280

Reasoning Does Not Necessarily Improve Role-Playing Ability
Xiachong Feng, Longxu Dou and Lingpeng Kong 10301

TableLLM: Enabling Tabular Data Manipulation by LLMs in Real Office Usage Scenarios
Xiaokang Zhang, Sijia Luo, Bohan Zhang, Zeyao Ma, Jing Zhang, Yang Li, Guanlin Li, Zijun Yao, Kangli Xu, Jinchang Zhou, Daniel Zhang-Li, Jifan Yu, Shu Zhao, Juanzi Li and Jie Tang . . 10315

A Survey of LLM-based Agents in Medicine: How far are we from Baymax?
Wenxuan Wang, Zizhan Ma, Zheng Wang, Chenghan Wu, Jiaming Ji, Wenting Chen, Xiang Li and Yixuan Yuan 10345

Context-Robust Knowledge Editing for Language Models
Haewon Park, Gyubin Choi, Minjun Kim and Yohan Jo 10360

Multi-Agent Collaboration via Cross-Team Orchestration
Zhuoyun Du, Chen Qian, Wei Liu, Zihao Xie, YiFei Wang, Rennai Qiu, Yufan Dang, Weize Chen, Cheng Yang, Ye Tian, Xuantang Xiong and Lei Han 10386

Semantic Evaluation of Multilingual Data-to-Text Generation via NLI Fine-Tuning: Precision, Recall and F1 scores
William Soto Martinez, Yannick Parmentier and Claire Gardent 10407

Optimized Text Embedding Models and Benchmarks for Amharic Passage Retrieval
Kidist Amde Mekonnen, Yosef Worku Alemneh and Maarten de Rijke 10428

Enhancing Transformation from Natural Language to Signal Temporal Logic Using LLMs with Diverse External Knowledge
Yue Fang, Zhi Jin, Jie An, Hongshen Chen, Xiaohong Chen and Naijun Zhan 10446

DAGS: A Dependency-Based Dual-Attention and Global Semantic Improvement Framework for Metaphor Recognition
Puli Chen, Cheng Yang, Xingmao Zhang and Qingbao Huang 10459

ESF: Efficient Sensitive Fingerprinting for Black-Box Tamper Detection of Large Language Models
Xiaofan Bai, Pingyi Hu, Xiaojing Ma, Linchen Yu, Dongmei Zhang, Qi Zhang and Bin Benjamin Zhu 10477

The Lessons of Developing Process Reward Models in Mathematical Reasoning
Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou and Junyang Lin 10495

MinosEval: Distinguishing Factoid and Non-Factoid for Tailored Open-Ended QA Evaluation with LLMs
Yongqi Fan, Yating Wang, Guandong Wang, Zhai Jie, Jingping Liu, Qi Ye and Tong Ruan 10517

Towards Conditioning Clinical Text Generation for User Control
Osman Alperen Koraş, Rabi Bahnan, Jens Kleesiek and Amin Dada 10549

CoDet-M4: Detecting Machine-Generated Code in Multi-Lingual, Multi-Generator and Multi-Domain Settings
Daniil Orel, Dilshod Azizov and Preslav Nakov 10570

Q-Mamba: Towards more efficient Mamba models via post-training quantization
Chen Tianqi, Yuanteng Chen, Peisong Wang, Weixiang Xu, Zeyu Zhu and Jian Cheng 10594

P²Net: Parallel Pointer-based Network for Key Information Extraction with Complex Layouts
Kaiwen Wei, Jie Yao, Jiang Zhong, Yangyang Kang, Jingyuan Zhang, Changlong Sun, Xin Zhang, Fengmao Lv and Li Jin 10611

Refining Sentence Embedding Model through Ranking Sentences Generation with Large Language Models
Liyang He, Chenglong Liu, Rui Li, Zhenya Huang, Shulan Ruan, Jun Zhou and Enhong Chen 10627

RQT: Hierarchical Residual Quantization for Multi-Model Compression
Chen Tianqi, Peisong Wang, Weixiang Xu, Zeyu Zhu and Jian Cheng 10644

taz2024full: Analysing German Newspapers for Gender Bias and Discrimination across Decades
Stefanie Urchs, Veronika Thurner, Matthias Aßenmacher, Christian Heumann and Stephanie Thiemichen 10661

LCFO: Long Context and Long Form Output Dataset and Benchmarking
Marta R. Costa-jussà, Pierre Andrews, Mariano Coria Meglioli, Joy Chen, Joe Chuang, David Dale, Christophe Ropers, Alexandre Mourachko, Eduardo Sánchez, Holger Schwenk, Tuan A. Tran, Arina Turkatenko and Carleigh Wood 10672

Span-based Semantic Role Labeling as Lexicalized Constituency Tree Parsing
Yang Hou and Zhenghua Li 10701

Learning from Negative Samples in Biomedical Generative Entity Linking
Chanhwi Kim, Hyunjae Kim, Sihyeon Park, Jiwoo Lee, Mujeen Sung and Jaewoo Kang . . 10714

Self-play through Computational Runtimes improves Chart Reasoning
Tautvydas Misiūnas, Hassan Mansoor, Jasper Uijlings, Oriana Riva and Victor Carbune . . 10731

Towards Better Chain-of-Thought: A Reflection on Effectiveness and Faithfulness
Jiachun Li, Pengfei Cao, Yubo Chen, Jiexin Xu, Huaijun Li, Xiaojian Jiang, Kang Liu and Jun Zhao 10747

A Couch Potato is not a Potato on a Couch: Prompting Strategies, Image Generation, and Compositionality Prediction for Noun Compounds
Sinan Kurtyigit, Diego Frassinelli, Carina Silberer and Sabine Schulte Im Walde 10766

A Rose by Any Other Name: LLM-Generated Explanations Are Good Proxies for Human Explanations to Collect Label Distributions on NLI
Beiduo Chen, Siyao Peng, Anna Korhonen and Barbara Plank 10777

Measuring What Matters: Evaluating Ensemble LLMs with Label Refinement in Inductive Coding
Angelina Parfenova and Jürgen Pfeffer 10803

Dynamic Evil Score-Guided Decoding: An Efficient Decoding Framework For Red-Team Model
Cong Gao, Bo Zhang, Linkang Yang, Minghao Hu, Zhunchen Luo, Xiaoying Bai, Guotong Geng, Jun Zhang and Yunhua Xue 10817

CoMuMDR: Code-mixed Multi-modal Multi-domain corpus for Discourse paRsing in conversations
Divyaksh Shukla, Ritesh Baviskar, Dwijesh Gohil, Aniket Tiwari, Atul Shree and Ashutosh Modi 10834

Multi-word Measures: Modeling Semantic Change in Compound Nouns
Chris Jenkins, Filip Miletić and Sabine Schulte Im Walde 10850

Bridge-Coder: Transferring Model Capabilities from High-Resource to Low-Resource Programming Language
Jipeng Zhang, Jianshu Zhang, Yuanzhe LI, Renjie Pi, Rui Pan, Runtao Liu, Zheng Ziqiang and Tong Zhang 10865

ProBench: Judging Multimodal Foundation Models on Open-ended Multi-domain Expert Tasks
Yan Yang, Dongxu Li, Haoning Wu, Bei Chen, Liu Liu, Liyuan Pan and Junnan Li 10883

2M-BELEBELE: Highly Multilingual Speech and American Sign Language Comprehension Dataset Download PDF
Marta R. Costa-jussà, Bokai YU, Pierre Andrews, Belen Alastruey, Necati Cihan Camgoz, Joe Chuang, Jean Maillard, Christophe Ropers, Arina Turkatenko and Carleigh Wood 10893

LSC-Eval: A General Framework to Evaluate Methods for Assessing Dimensions of Lexical Semantic Change Using LLM-Generated Synthetic Data
Naomi Baes, Raphael Merx, Nick Haslam, Ekaterina Vylomova and Haim Dubossarsky . . 10905

Chain-of-Jailbreak Attack for Image Generation Models via Step by Step Editing
Wenxuan Wang, Kuiyi Gao, Youliang Yuan, Jen-tse Huang, Qiuzhi Liu, Shuai Wang, Wenxiang Jiao and Zhaopeng Tu 10940

Tokenization is Sensitive to Language Variation
Anna Wegmann, Dong Nguyen and David Jurgens 10958

WirelessMathBench: A Mathematical Modeling Benchmark for LLMs in Wireless Communications
Xin Li, Mengbing Liu, LI Wei, Jiancheng An, Merouane Abdelkader Debbah and Chau Yuen 10984

Self-Improvement Towards Pareto Optimality: Mitigating Preference Conflicts in Multi-Objective Alignment
Moxin Li, Yuantao Zhang, Wenjie Wang, Wentao Shi, Zhuo Liu, Fuli Feng and Tat-Seng Chua 11010

Investigating and Scaling up Code-Switching for Multilingual Language Model Pre-Training
Zhijun Wang, Jiahuan Li, Hao Zhou, Rongxiang Weng, Jingang Wang, Xin Huang, Xue Han, Junlan Feng, Chao Deng and Shujian Huang 11032

User Behavior Prediction as a Generic, Robust, Scalable, and Low-Cost Evaluation Strategy for Estimating Generalization in LLMs
Sougata Saha and Monojit Choudhury 11047

Beyond Browsing: API-Based Web Agents
Yueqi Song, Frank F. Xu, Shuyan Zhou and Graham Neubig 11066

MiLiC-Eval: Benchmarking Multilingual LLMs for China's Minority Languages
Chen Zhang, Mingxu Tao, Zhiyuan Liao and Yansong Feng 11086

ArgInstruct: Specialized Instruction Fine-Tuning for Computational Argumentation
Maja Stahl, Timon Ziegenbein, Joonsuk Park and Henning Wachsmuth 11103

Crabs: Consuming Resource via Auto-generation for LLM-DoS Attack under Black-box Settings
Yuanhe Zhang, Zhenhong Zhou, Wei Zhang, Xinyue Wang, Xiaojun Jia, Yang Liu and Sen Su 11128

Probabilistic Aggregation and Targeted Embedding Optimization for Collective Moral Reasoning in Large Language Models
Chenchen Yuan, Zheyu Zhang, Shuo Yang, Bardh Prenkaj and Gjergji Kasneci 11151

Unlocking Recursive Thinking of LLMs: Alignment via Refinement
Haoke Zhang, Xiaobo Liang, Cunxiang Wang, Juntao Li and Min Zhang 11169

CitaLaw: Enhancing LLM with Citations in Legal Domain
Kepu Zhang, Weijie Yu, Sunhao Dai and Jun Xu 11183

MEGen: Generative Backdoor into Large Language Models via Model Editing
Jiyang Qiu, Xinbei Ma, Zhuosheng Zhang, Hai Zhao, Yun Li and Qianren Wang 11197

Social Bias Benchmark for Generation: A Comparison of Generation and QA-Based Evaluations
Jiho Jin, Woosung Kang, Junho Myung and Alice Oh 11215

Generating Pedagogically Meaningful Visuals for Math Word Problems: A New Benchmark and Analysis of Text-to-Image Models
Junling Wang, Anna Rutkiewicz, April Wang and Mrinmaya Sachan 11229

RASPberry: Retrieval-Augmented Monte Carlo Tree Self-Play with Reasoning Consistency for Multi-Hop Question Answering
Baixuan Li, Yunlong Fan, Tianyi Ma, Miao Gao, Chuanqi Shi and Zhiqiang Gao 11258

All That Glitters is Not Gold: Improving Robust Retrieval-Augmented Language Models with Fact-Centric Preference Alignment
Jia Hao, Chunhong Zhang, Jiarun Liu, Haiyu Zhao, Zhiqiang Zhan and Zheng Hu 11277

FairSteer: Inference Time Debiasing for LLMs with Dynamic Activation Steering
Yichen Li, Zhiting Fan, Ruizhe Chen, Xiaotang Gai, Luqi Gong, Yan Zhang and Zuozhu Liu 11293

Listen, Watch, and Learn to Feel: Retrieval-Augmented Emotion Reasoning for Compound Emotion Generation
Zhuofan Wen, Zheng Lian, Shun Chen, Hailiang Yao, Longjiang Yang, Bin Liu and Jianhua Tao 11313

GLTW: Joint Improved Graph Transformer and LLM via Three-Word Language for Knowledge Graph Completion
Kangyang Luo, Yuzhuo Bai, Cheng Gao, Shuzheng Si, Zhu Liu, Yingli Shen, Zhitong Wang, Cunliang Kong, Wenhao Li, Yufei Huang, Ye Tian, Xuantang Xiong, Lei Han and Maosong Sun 11328

Learning to Select In-Context Demonstration Preferred by Large Language Model
Zheng Zhang, Shaocheng Lan, Lei Song, Jiang Bian, Yexin Li and Kan Ren 11345

Beyond the Spelling Miracle: Investigating Substring Awareness in Character-Blind Language Models
Cristiano Ciaccio, Marta Sartor, Alessio Miaschi and Felice Dell'Orletta 11361

DEMO: Reframing Dialogue Interaction with Fine-grained Element Modeling
Minzheng Wang, Xinghua Zhang, Kun Chen, Nan Xu, Haiyang Yu, Fei Huang, Wenji Mao and Yongbin Li 11373

InfiniteICL: Breaking the Limit of Context Window Size via Long Short-term Memory Transformation
Bowen Cao, Deng Cai and Wai Lam 11402

M3HG: Multimodal, Multi-scale, and Multi-type Node Heterogeneous Graph for Emotion Cause Triplet Extraction in Conversations
Qiao Liang, Ying Shen, Tiantian Chen and Lin Zhang 11416

Large Language Models Are Natural Video Popularity Predictors
Pratik Kayal, Pascal Mettes, Nima Dehmamy and Minsu Park 11432

DELMAN: Dynamic Defense Against Large Language Model Jailbreaking with Model Editing
Yi Wang, Fenghua Weng, Sibei Yang, Zhan Qin, Minlie Huang and Wenjie Wang 11465

You need to MIMIC to get FAME: Solving Meeting Transcript Scarcity with Multi-Agent Conversations
Frederic Kirstein, Muneeb Khan, Jan Philip Wahle, Terry Ruas and Bela Gipp 11482

Code-Switching and Syntax: A Large-Scale Experiment
Igor Sterner and Simone Teufel 11526

Optima: Optimizing Effectiveness and Efficiency for LLM-Based Multi-Agent System
Weize Chen, Jiarui Yuan, Chen Qian, Cheng Yang, Zhiyuan Liu and Maosong Sun 11534

Generating Domain-Specific Knowledge Graphs from Large Language Models
Marinela Parović, Ze Li and Jinhua Du 11558

Large Language Models are Miscalibrated In-Context Learners
Chengzu Li, Han Zhou, Goran Glavaš, Anna Korhonen and Ivan Vulić 11575

STeCa: Step-level Trajectory Calibration for LLM Agent Learning
Hanlin Wang, Jian Wang, Chak Tou Leong and Wenjie Li 11597

LEMMA: Learning from Errors for MatheMatical Advancement in LLMs
Zhuoshi Pan, Yu Li, Honglin Lin, Qizhi Pei, Zinan Tang, Wei Wu, Chenlin Ming, H. Vicky Zhao, Conghui He and Lijun Wu 11615

Voting or Consensus? Decision-Making in Multi-Agent Debate
Lars Benedikt Kaesberg, Jonas Becker, Jan Philip Wahle, Terry Ruas and Bela Gipp 11640

Rhetorical Device-Aware Sarcasm Detection with Counterfactual Data Augmentation
Qingqing Hong, Dongyu Zhang, Jiayi Lin, Dapeng Yin, Shuyue Zhu and Junli Wang 11672

Selecting Demonstrations for Many-Shot In-Context Learning via Gradient Matching
Jianfei Zhang, Bei Li, Jun Bai, Rumei Li, Yanmeng Wang, Chenghua Lin and Wenge Rong 11686

Cheap Character Noise for OCR-Robust Multilingual Embeddings
Andrianos Michail, Juri Opitz, Yining Wang, Robin Meister, Rico Sennrich and Simon Clematide 11705

Physics: Benchmarking Foundation Models on University-Level Physics Problem Solving
Kaiyue Feng, Yilun Zhao, Yixin Liu, Tianyu Yang, Chen Zhao, John Sous and Arman Cohan 11717

DOVE: A Large-Scale Multi-Dimensional Predictions Dataset Towards Meaningful LLM Evaluation
Eliya Habba, Ofir Arviv, Itay Itzhak, Yotam Perlitz, Elron Bandel, Leshem Choshen, Michal Shmueli-Scheuer and Gabriel Stanovsky 11744

ALPS: Attention Localization and Pruning Strategy for Efficient Adaptation of Large Language Models
Hao Chen, Haoze Li, Zhiqing Xiao, Lirong Gao, Qi Zhang, Xiaomeng Hu, Ningtao Wang, Xing Fu and Junbo Zhao 11764

DeTAM: Defending LLMs Against Jailbreak Attacks via Targeted Attention Modification
Yu Li, Han Jiang and Zhihua Wei 11781

A Survey of Mathematical Reasoning in the Era of Multimodal Large Language Model: Benchmark, Method & Challenges
Yibo Yan, Jiamin Su, Jianxiang He, Fangteng FU, Xu Zheng, Yuanhuiyi Lyu, Kun Wang, Shen Wang, Qingsong Wen and Xuming Hu 11798

Fast-and-Frugal Text-Graph Transformers are Effective Link Predictors
Andrei Catalin Coman, Christos Theodoropoulos, Marie-Francine Moens and James Henderson 11828

NeoQA: Evidence-based Question Answering with Generated News Events
Max Glockner, Xiang Jiang, Leonardo F. R. Ribeiro, Iryna Gurevych and Markus Dreyer . 11842

ChatMap: Mining Human Thought Processes for Customer Service Chatbots via Multi-Agent Collaboration
Xinyi Jiang, Tianyi Hu, Yuheng Qin, Guoming Wang, Zhou Huan, Kehan Chen, Gang Huang, Rongxing Lu and Siliang Tang 11927

P3: Prompts Promote Prompting
Xinyu Zhang, Yuanquan Hu, Fangchao Liu and Zhicheng Dou 11948

VAQUUM: Are Vague Quantifiers Grounded in Visual Data?
Hugh Mee Wong, Rick Nouwen and Albert Gatt 11966

Forgotten Polygons: Multimodal Large Language Models are Shape-Blind
William Rudman, Michal Golovanevsky, Amir Bar, Vedant Palit, Yann LeCun, Carsten Eickhoff and Ritambhara Singh 11983

MindBridge: Scalable and Cross-Model Knowledge Editing via Memory-Augmented Modality
Shuaike Li, Kai Zhang, Qi Liu and Enhong Chen 11999

FIHA: Automated Fine-grained Hallucinations Evaluations in Large Vision Language Models with Davidson Scene Graphs
Bowen Yan, Zhengsong Zhang, Liqiang Jing, Eftekhar Hossain and Xinya Du 12014

On the Role of Semantic Proto-roles in Semantic Analysis: What do LLMs know about agency?
Elizabeth Spaulding, Shafiuddin Rehan Ahmed and James Martin 12027

GeAR: Graph-enhanced Agent for Retrieval-augmented Generation
Zhili Shen, Chenxin Diao, Pavlos Vougiouklis, Pascual Merita, Shriram Piramanayagam, Enting Chen, Damien Graux, Andre Melo, Ruofei Lai, Zeren Jiang, Zhongyang Li, YE QI, Yang Ren, Dandan Tu and Jeff Z. Pan 12049

WebNLG-IT: Construction of an aligned RDF-Italian corpus through Machine Translation techniques
Michael Oliverio, Pier Felice Balestrucci, Alessandro Mazzei and Valerio Basile 12073

Towards Adapting Open-Source Large Language Models for Expert-Level Clinical Note Generation
Hanyin Wang, Chufan Gao, Bolun Liu, Qiping Xu, Guleid Hussein, Mohamad El Labban, Kingsley Iheasirim, Hariprasad Reddy Korsapati, Chuck Outcalt and Jimeng Sun 12084

Bridging Robustness and Generalization Against Word Substitution Attacks in NLP via the Growth Bound Matrix Approach
Mohammed Bouri and Adnane Saoud 12118

Neuro-Symbolic Query Compiler
Yuyao Zhang, Zhicheng Dou, Xiaoxi Li, Jiajie Jin, Yongkang Wu, Zhonghua Li, YE QI and Ji-Rong Wen 12138

Revealing and Mitigating the Local Pattern Shortcuts of Mamba
WangJie You, Zecheng Tang, Juntao Li, Lili Yao and Min Zhang 12156

Forget the Token and Pixel: Rethinking Gradient Ascent for Concept Unlearning in Multimodal Generative Models
Jiaqi Li, Chuanyi Zhang, Miaozeng Du, Hui Zhang, Yongrui Chen, Qianshan Wei, Junfeng Fang, Ruipeng Wang, Sheng Bi and Guilin Qi 12179

Slamming: Training a Speech Language Model on One GPU in a Day
Gallil Maimon, Avishai Elmakies and Yossi Adi 12201

Boosting LLM Translation Skills without General Ability Loss via Rationale Distillation
Junhong Wu, Yang Zhao, Yangyifan Xu, Bing Liu and Chengqing Zong 12217

Clarifying Underspecified Discourse Relations in Instructional Texts
Berfin Aktas and Michael Roth 12237

WMT24++: Expanding the Language Coverage of WMT24 to 55 Languages & Dialects
Daniel Deutsch, Eleftheria Briakou, Isaac Rayburn Caswell, Mara Finkelstein, Rebecca Galor, Juraj Juraska, Geza Kovacs, Alison Lui, Ricardo Rei, Jason Riesa, Shruti Rijhwani, Parker Riley, Elizabeth Salesky, Firas Trabelsi, Stephanie Winkler, Biao Zhang and Markus Freitag 12257

Exploring Graph Representations of Logical Forms for Language Modeling
Michael Sullivan 12285

SEA-HELM: Southeast Asian Holistic Evaluation of Language Models
Yosephine Susanto, Adithya Venkatadri Hulagadri, Jann Railey Montalan, Jian Gang Ngui, Xianbin Yong, Wei Qi Leong, Hamsawardhini Rengarajan, Peerat Limkonchotiwat, Yifan Mai and William Chandra Tjhi 12308

TRANS-ZERO: Self-Play Incentivizes Large Language Models for Multilingual Translation Without Parallel Data
Wei Zou, Sen Yang, Yu Bao, Shujian Huang, Jiajun Chen and Shanbo Cheng 12337

A Conformal Risk Control Framework for Granular Word Assessment and Uncertainty Calibration of CLIPScore Quality Estimates
Goncalo Emanuel Cavaco Gomes, Bruno Martins and Chrysoula Zerva 12348

SGDPO: Self-Guided Direct Preference Optimization for Language Model Alignment
Wenqiao Zhu, Ji Liu, Lulu Wang, Jun Wu and Yulun Zhang 12366

Socratic Style Chain-of-Thoughts Help LLMs to be a Better Reasoner
Jiangbo Pei, Peiyu Liu, Xin Zhao, Aidong Men and Yang Liu 12384

Quantile Regression with Large Language Models for Price Prediction
Nikhita Vedula, Dushyanta Dhyani, Laleh Jalali, Boris N. Oreshkin, Mohsen Bayati and Shervin Malmasi 12396

Training Turn-by-Turn Verifiers for Dialogue Tutoring Agents: The Curious Case of LLMs as Your Coding Tutors
Jian Wang, Yinpei Dai, Yichi Zhang, Ziqiao Ma, Wenjie Li and Joyce Chai 12416

AIGuard: A Benchmark and Lightweight Detection for E-commerce AIGC Risks
Wenhua Zhang, Weicheng Li, Xuanrong Rao, Lixin Zou, Xiangyang Luo, Chubin Zhuang, Yongjie Hong, Zhen Qin, Hengyu Chang, Chenliang Li and Bo Zheng 12437

A^2ATS: Retrieval-Based KV Cache Reduction via Windowed Rotary Position Embedding and Query-Aware Vector Quantization
Junhui He, Junna Xing, Nan Wang, Rui Xu, Shangyu Wu, Peng Zhou, Qiang Liu, Chun Jason Xue and Qingan Li 12451

TransBench: Breaking Barriers for Transferable Graphical User Interface Agents in Dynamic Digital Environments
Yuheng Lu, Qian Yu, Hongru Wang, Zeming Liu, Wei Su, Yanping Liu, Yuhang Guo, Maocheng Liang, Yunhong Wang and Haifeng Wang 12464

Order Matters: Investigate the Position Bias in Multi-constraint Instruction Following
Jie Zeng, Qianyu He, Qingyu Ren, Jiaqing Liang, Weikang Zhou, Zeye Sun, Fei Yu and Yanghua Xiao 12479

CoT-VTM: Visual-to-Music Generation with Chain-of-Thought Reasoning
Xikang Guan, Zheng Gu, Jing Huo, Tianyu Ding and Yang Gao 12493

A Tale of Evaluating Factual Consistency: Case Study on Long Document Summarization Evaluation
Yang Zhong and Diane Litman 12511

Evaluating Pretrained Causal Language Models for Synonymy
Ioana Ivan, Carlos Ramisch and Alexis Nasr 12533

MDIT-Bench: Evaluating the Dual-Implicit Toxicity in Large Multimodal Models
Bohan Jin, Shuhan Qi, Kehai Chen, Xinyi Guo and Xuan Wang 12552

CoVE: Compressed Vocabulary Expansion Makes Better LLM-based Recommender Systems
Haochen Zhang, Tianyi Zhang, Junze Yin, Oren Gal, Anshumali Shrivastava and Vladimir Braverman 12575

CtrlA: Adaptive Retrieval-Augmented Generation via Inherent Control
Liu Huanshuo, Hao Zhang, Zhijiang Guo, Jing Wang, Kuicai Dong, Xiangyang Li, Yi Quan Lee, Cong Zhang and Yong Liu 12592

Maximum Score Routing For Mixture-of-Experts
Bowen Dong, Yilong Fan, Yutao Sun, Zhenyu Li, Tengyu Pan, Zhou Xun and Jianyong Wang 12619

Time Course MechInterp: Analyzing the Evolution of Components and Knowledge in Large Language Models
Ahmad Dawar Hakimi, Ali Modarressi, Philipp Wicke and Hinrich Schuetze 12633

Well Begun is Half Done: Low-resource Preference Alignment by Weak-to-Strong Decoding
Feifan Song, Shaohang Wei, Wen Luo, Yuxuan Fan, Tianyu Liu, Guoyin Wang and Houfeng Wang 12654

Disentangling Text and Math in Word Problems: Evidence for the Bidimensional Structure of Large Language Models' Reasoning
Pedro Calais, Gabriel Franco, Zilu Tang, Themistoklis Nikas, Wagner Meira Jr., Evimaria Terzi and Mark Crovella 12671

Human-LLM Coevolution: Evidence from Academic Writing
Mingmeng Geng and Roberto Trotta 12689

Disentangled Multi-span Evolutionary Network against Temporal Knowledge Graph Reasoning
Hao Dong, Ziyue Qiao, Zhiyuan Ning, Qi Hao, Yi Du, Pengyang Wang and Yuanchun Zhou 12697

GRAF: Graph Retrieval Augmented by Facts for Romanian Legal Multi-Choice Question Answering
Cristian-George Craciun, Răzvan-Alexandru Smădu, Dumitru-Clementin Cercel and Mihaela-Claudia Cercel 12708

Express What You See: Can Multimodal LLMs Decode Visual Ciphers with Intuitive Semiosis Comprehension?
Jiayi Kuang, Yinghui Li, Chen Wang, Haohao Luo, Ying Shen and Wenhao Jiang 12743

ConFit v2: Improving Resume-Job Matching using Hypothetical Resume Embedding and Runner-Up Hard-Negative Mining
Xiao Yu, Ruize Xu, Chengyuan Xue, Jinzhong Zhang, Xu Ma and Zhou Yu 12775

Knowing Before Saying: LLM Representations Encode Information About Chain-of-Thought Success Before Completion
Anum Afzal, Florian Matthes, Gal Chechik and Yftah Ziser 12791

Grounding Task Assistance with Multimodal Cues from a Single Demonstration
Gabriel Herbert Sarch, Balasaravanan Thoravi Kumaravel, Sahithya Ravi, Vibhav Vineet and Andrew D Wilson 12807

Awes, Laws, and Flaws From Today's LLM Research
Adrian de Wynter 12834

Dual Debiasing for Noisy In-Context Learning for Text Generation
Siqi Liang, Sumyeong Ahn, Paramveer Dhillon and Jiayu Zhou 12855

DRS: Deep Question Reformulation With Structured Output
Zhecheng Li, Yiwei Wang, Bryan Hooi, Yujun Cai, Nanyun Peng and Kai-Wei Chang 12869

Towards Explainable Hate Speech Detection
Happy Khairunnisa Sariyanto, Diclehan Ulucan, Oguzhan Ulucan and Marc Ebner 12883

BioHopR: A Benchmark for Multi-Hop, Multi-Answer Reasoning in Biomedical Domain
Yunsoo Kim, Yusuf Abdulle and Honghan Wu 12894

PipeSpec: Breaking Stage Dependencies in Hierarchical LLM Decoding
Bradley McDanel, Sai Qian Zhang, Yunhai Hu and Zining Liu 12909

LAM SIMULATOR: Advancing Data Generation for Large Action Model Training via Online Exploration and Trajectory Feedback
Thai Quoc Hoang, Kung-Hsiang Huang, Shirley Kokane, Jianguo Zhang, Zuxin Liu, Ming Zhu, Jake Grigsby, Tian Lan, Michael S Ryoo, Chien-Sheng Wu, Shelby Heinecke, Huan Wang, Silvio Savarese, Caiming Xiong and Juan Carlos Niebles 12921

Rank, Chunk and Expand: Lineage-Oriented Reasoning for Taxonomy Expansion
Sahil Mishra, Kumar Arjun and Tanmoy Chakraborty 12935

Probing Subphonemes in Morphology Models
Gal Astrach and Yuval Pinter 12954

Exploiting Instruction-Following Retrievers for Malicious Information Retrieval
Parishad BehnamGhader, Nicholas Meade and Siva Reddy 12962

Improving Causal Interventions in Amnesic Probing with Mean Projection or LEACE
Alicja Dobrzeniecka, Antske Fokkens and Pia Sommerauer 12981

The Threat of PROMPTS in Large Language Models: A System and User Prompt Perspective
Zixuan Xia, Haifeng Sun, Jingyu Wang, Qi Qi, Huazheng Wang, Xiaoyuan Fu and Jianxin Liao 12994

RoseRAG: Robust Retrieval-augmented Generation with Small-scale LLMs via Margin-aware Preference Optimization
Tianci Liu, Haoxiang Jiang, Tianze Wang, Ran Xu, Yue Yu, Linjun Zhang, Tuo Zhao and Haoyu Wang 13036

Instruction-Tuning LLMs for Event Extraction with Annotation Guidelines
Saurabh Srivastava, Sweta Pati and Ziyu Yao 13055

mRAKL: Multilingual Retrieval-Augmented Knowledge Graph Construction for Low-Resourced Languages
Hellina Hailu Nigatu, Min Li, Maartje Ter Hoeve, Saloni Potdar and Sarah Chasins 13072

Mechanistic Interpretability of Emotion Inference in Large Language Models
Ala N. Tak, Amin Banayeeanzade, Anahita Bolourani, Mina Kian, Robin Jia and Jonathan Gratch 13090

RL-Guider: Leveraging Historical Decisions and Feedback for Drug Editing with Large Language Models
Xufeng Liu, Yixuan Ding, Jingxiang Qu, Yichi Zhang, Wenhan Gao and Yi Liu 13121

BriefMe: A Legal NLP Benchmark for Assisting with Legal Briefs
Jesse Woo, Fateme Hashemi Chaleshtori, Ana Marasovic and Kenneth Marino 13139

I see what you mean: Co-Speech Gestures for Reference Resolution in Multimodal Dialogue
Esam Ghaleb, Bulat Khaertdinov, Asli Ozyurek and Raquel Fernández 13191

World Knowledge Resolves Some Aspectual Ambiguity
Katarzyna Pruś, Mark Steedman and Adam Lopez 13207

ACCESS DENIED INC: The First Benchmark Environment for Sensitivity Awareness
Dren Fazlija, Arkadij Orlov and Sandipan Sikdar 13221

Spatial Coordinates as a Cell Language: A Multi-Sentence Framework for Imaging Mass Cytometry Analysis
Chi-Jane Chen, Yuhang Chen, Sukwon Yun, Natalie Stanley and Tianlong Chen 13241

HumanEval Pro and MBPP Pro: Evaluating Large Language Models on Self-invoking Code Generation Task
Zhaojian Yu, Yilun Zhao, Arman Cohan and Xiao-Ping Zhang 13253

TCSinger 2: Customizable Multilingual Zero-shot Singing Voice Synthesis
Yu Zhang, Wenxiang Guo, Changhao Pan, Dongyu Yao, Zhiyuan Zhu, Ziyue Jiang, Yuhan Wang, Tao Jin and Zhou Zhao 13280

Compute Optimal Scaling of Skills: Knowledge vs Reasoning
Nicholas Roberts, Niladri S. Chatterji, Sharan Narang, Mike Lewis and Dieuwke Hupkes . 13295

PECAN: LLM-Guided Dynamic Progress Control with Attention-Guided Hierarchical Weighted Graph for Long-Document QA
Xinyu Wang, Yanzheng Xiang, Lin Gui and Yulan He 13317

Lifelong Model Editing with Graph-Based External Memory
Yash Kumar Atri, Ahmed Alaa and Thomas Hartvigsen 13336

Multi-Sense Embeddings for Language Models and Knowledge Distillation
Qitong Wang, Mohammed J Zaki, Georgios Kollias and Vasileios Kalantzis 13353

CodeScientist: End-to-End Semi-Automated Scientific Discovery with Code-based Experimentation
Peter Jansen, Oyvind Tafjord, Marissa Radensky, Pao Siangliulue, Tom Hope, Bhavana Dalvi Mishra, Bodhisattwa Prasad Majumder, Daniel S Weld and Peter Clark 13370

Beyond Factual Accuracy: Evaluating Coverage of Diverse Factual Information in Long-form Text Generation
Chris Samarinas, Alexander Krubner, Alireza Salemi, Youngwoo Kim and Hamed Zamani 13468

Continual Quantization-Aware Pre-Training: When to transition from 16-bit to 1.58-bit pre-training for BitNet language models?
Jacob Nielsen, Peter Schneider-Kamp and Lukas Galke 13483

When Detection Fails: The Power of Fine-Tuned Models to Generate Human-Like Social Media Text
Hillary Dawkins, Kathleen C. Fraser and Svetlana Kiritchenko 13494

Not quite Sherlock Holmes: Language model predictions do not reliably differentiate impossible from improbable events
James A. Michaelov, Reeka Estacio, Zhien Zhang and Ben Bergen 13528

The Rotary Position Embedding May Cause Dimension Inefficiency in Attention Heads for Long-Distance Retrieval
Ting-Rui Chiang and Dani Yogatama 13552

IDEA: Enhancing the Rule Learning Ability of Large Language Model Agent through Induction, Deduction, and Abduction
Kaiyu He, Mian Zhang, Shuo Yan, Peilin Wu and Zhiyu Chen 13563

EnigmaToM: Improve LLMs' Theory-of-Mind Reasoning Capabilities with Neural Knowledge Base of Entity States
Hainiu Xu, Siya Qi, Jiazheng Li, Yuxiang Zhou, Jinhua Du, Caroline Catmur and Yulan He 13598

ReasonerRank: Redefining Language Model Evaluation with Ground-Truth-Free Ranking Frameworks
Jiamu Zhang, Jiayi Yuan, Andrew Wen, Hoang Anh Duy Le, Yu-Neng Chuang, Soo-Hyun Choi, Rui Chen and Xia Hu 13623

HyGenar: An LLM-Driven Hybrid Genetic Algorithm for Few-Shot Grammar Generation
Weizhi Tang, Yixuan Li, Chris Sypherd, Elizabeth Polgreen and Vaishak Belle 13640

Can Large Language Models Understand Argument Schemes?
Elfia Bezou-Vrakatseli, Oana Cocarascu and Sanjay Modgil 13666

MMInA: Benchmarking Multihop Multimodal Internet Agents
Shulin Tian, Ziniu Zhang, Liangyu Chen and Ziwei Liu 13682

ThinkGuard: Deliberative Slow Thinking Leads to Cautious Guardrails
Xiaofei Wen, Wenxuan Zhou, Wenjie Jacky Mo and Muhao Chen 13698

Neutralizing Bias in LLM Reasoning using Entailment Graphs
Liang Cheng, Tianyi Li, Zhaowei Wang, Tianyang Liu and Mark Steedman 13714

Dynamic Steering With Episodic Memory For Large Language Models
Van Dai Do, Quan Hung Tran, Svetha Venkatesh and Hung Le 13731

Eeyore: Realistic Depression Simulation via Expert-in-the-Loop Supervised and Preference Optimization
Siyang Liu, Bianca Brie, Wenda Li, Laura Biester, Andrew Lee, James Pennebaker and Rada Mihalcea 13750

Lost in Translation: Benchmarking Commercial Machine Translation Models for Dyslexic-Style Text
Gregory Price and Shaomei Wu 13771

Divide-Verify-Refine: Can LLMs Self-align with Complex Instructions?
Xianren Zhang, Xianfeng Tang, Hui Liu, Zongyu Wu, Qi He, Dongwon Lee and Suhang Wang 13783

LlamaPIE: Proactive In-Ear Conversation Assistants
Tuochao Chen, Nicholas Scott Batchelder, Alisa Liu, Noah A. Smith and Shyamnath Gollakota 13801

Task-Oriented Automatic Fact-Checking with Frame-Semantics
Jacob Devasier, Akshith Reddy Putta, Rishabh Mediratta and Chengkai Li 13825

Craw4LLM: Efficient Web Crawling for LLM Pretraining
Shi Yu, Zhiyuan Liu and Chenyan Xiong 13843

Be Cautious When Merging Unfamiliar LLMs: A Phishing Model Capable of Stealing Privacy
Guo Zhenyuan, Yi Shi, Wenlong Meng, Chen Gong, Chengkun Wei and Wenzhi Chen 13852

Understand User Opinions of Large Language Models via LLM-Powered In-the-Moment User Experience Interviews
Mengqiao Liu, Tevin Wang, Cassandra A. Cohen, Sarah Li and Chenyan Xiong 13872

HiCOT: Improving Neural Topic Models via Optimal Transport and Contrastive Learning
Hoang Tran Vuong, Tue Le, Tu Vu, Tung Nguyen, Linh Ngo Van, Sang Dinh and Thien Huu Nguyen 13894

FLAG-TRADER: Fusion LLM-Agent with Gradient-based Reinforcement Learning for Financial Trading
Guojun Xiong, Zhiyang Deng, Keyi Wang, Yupeng Cao, Haohang Li, Yangyang Yu, Xueqing Peng, Mingquan Lin, Kaleb E Smith, Xiao-Yang Liu, Jimin Huang, Sophia Ananiadou and Qianqian Xie 13921

The Silent Saboteur: Imperceptible Adversarial Attacks against Black-Box Retrieval-Augmented Generation Systems
Hongru Song, Yu-An Liu, Ruqing Zhang, Jiafeng Guo, Jianming Lv, Maarten de Rijke and Xueqi Cheng 13935

CROSSAGENTIE: Cross-Type and Cross-Task Multi-Agent LLM Collaboration for Zero-Shot Information Extraction
Meng Lu, Yuzhang Xie, Zhenyu Bi, Shuxiang Cao and Xuan Wang 13953

Decoupling Memories, Muting Neurons: Towards Practical Machine Unlearning for Large Language Models
Lishuai Hou, Zixiong Wang, Gaoyang Liu, Chen Wang, Wei Liu and Kai Peng 13978

Assimilation and Accommodation: Task-Adaptive Hierarchical Abstraction for Solving Web Tasks
Xinyu Pang, Ruixin Hong, Hongming Zhang and Changshui Zhang 14000

SafeLawBench: Towards Safe Alignment of Large Language Models
Chuxue Cao, Han Zhu, Jiaming Ji, Qichao Sun, Zhenghao Zhu, WU Yinyu, Josef Dai, Yaodong Yang, Sirui Han and Yike Guo 14015

3DM: Distill, Dynamic Drop, and Merge for Debiasing Multi-modal Large Language Models
Zhaoxi Zhang, Sanwoo Lee, Zhixiang Wang and Yunfang Wu 14049

CausalAbstain: Enhancing Multilingual LLMs with Causal Reasoning for Trustworthy Abstention
Yuxi Sun, Aoqi Zuo, Wei Gao and Jing Ma 14060

CapArena: Benchmarking and Analyzing Detailed Image Captioning in the LLM Era
Kanzhi Cheng, Wenpo Song, Jiaxin Fan, Zheng Ma, Qiushi Sun, Fangzhi Xu, Chenyang Yan, Nuo Chen, Jianbing Zhang and Jiajun Chen 14077

LLM-Empowered Class Imbalanced Graph Prompt Learning for Online Drug Trafficking Detection
Tianyi Ma, Yiyue Qian, Zehong Wang, Zheyuan Zhang, Chuxu Zhang and Yanfang Ye . . . 14095

CoLA: Collaborative Low-Rank Adaptation
Yiyun Zhou, Chang Yao and Jingyuan Chen 14115

GLiM: Integrating Graph Transformer and LLM for Document-Level Biomedical Relation Extraction with Incomplete Labeling
Hao Fang, Yuejie Zhang, Rui Feng, Yingwen Wang, Qing Wang, Wen He, Xiaobo Zhang, Tao Zhang and Shang Gao 14131

AnalyticKWS: Towards Exemplar-Free Analytic Class Incremental Learning for Small-footprint Keyword Spotting
Yang Xiao, Peng Tianyi, Rohan Kumar Das, Yuchen Hu and Huiping Zhuang 14147

Sleepless Nights, Sugary Days: Creating Synthetic Users with Health Conditions for Realistic Coaching Agent Interactions
Taedong Yun, Eric Yang, Mustafa Safdari, Jong Ha Lee, Vaishnavi Vinod Kumar, S. Sara Mahdavi, Jonathan Amar, Derek Peyton, Reut Aharony, Andreas Michaelides PhD, Logan Douglas Schneider, Isaac Galatzer-Levy, Yugang Jia, John Canny, Arthur Gretton and Maja Mataric 14159

Imagine to Hear: Auditory Knowledge Generation can be an Effective Assistant for Language Models
Suho Yoo, Hyunjong Ok and Jaeho Lee 14182

SafeEraser: Enhancing Safety in Multimodal Large Language Models through Multimodal Machine Unlearning
Junkai Chen, Zhijie Deng, Kening Zheng, Yibo Yan, Shuliang Liu, PeiJun WU, Peijie Jiang, Jia Liu and Xuming Hu 14194

Prediction-Augmented Generation for Automatic Diagnosis Tasks
Chan-Yang Ju and Dong-Ho Lee 14225

FedLEKE: Federated Locate-then-Edit Knowledge Editing for Multi-Client Collaboration
Zongkai Zhao, Guozeng Xu, Xiuhua Li, Kaiwen Wei and Jiang Zhong 14247

DiSCo: Device-Server Collaborative LLM-based Text Streaming Services
Ting Sun, Penghan Wang and Fan Lai 14259

Customizing In-context Learning for Dynamic Interest Adaption in LLM-based Recommendation
Keqin Bao, Ming Yan, Yang Zhang, Jizhi Zhang, Wenjie Wang, Fuli Feng and Xiangnan He 14278

Robust Data Watermarking in Language Models by Injecting Fictitious Knowledge
Xinyue Cui, Johnny Wei, Swabha Swayamdipta and Robin Jia 14292

LLM-Enhanced Query Generation and Retrieval Preservation for Task-Oriented Dialogue
Jiale Chen, Xuelian Dong, Wenxiu Xie, Ru Peng, Kun Zeng and Tianyong Hao 14307

ClozeMath: Improving Mathematical Reasoning in Language Models by Learning to Fill Equations
Quang Hieu Pham, Thuy Duong Nguyen, Tung Pham, Anh Tuan Luu and Dat Quoc Nguyen 14322

Low-Entropy Watermark Detection via Bayes' Rule Derived Detector
Beining Huang, Du Su, Fei Sun, Qi Cao, Huawei Shen and Xueqi Cheng 14330

CoD, Towards an Interpretable Medical Agent using Chain of Diagnosis
Junying Chen, Chi Gui, Anningzhe Gao, Ke Ji, Xidong Wang, Xiang Wan and Benyou Wang 14345

DaNet: Dual-Aware Enhanced Alignment Network for Multimodal Aspect-Based Sentiment Analysis
Aoqiang Zhu, Min Hu, Xiaohua Wang, Jiaoyun Yang, Yiming Tang and Ning An 14369

Exploring Multimodal Challenges in Toxic Chinese Detection: Taxonomy, Benchmark, and Findings
Shujian Yang, Shiyao Cui, Chuanrui Hu, Haicheng Wang, Tianwei Zhang, Minlie Huang, Jialiang LU and Han Qiu 14382

LDIR: Low-Dimensional Dense and Interpretable Text Embeddings with Relative Representations
Yile Wang, Zhanyu Shen and Hui Huang 14397

Ranked Voting based Self-Consistency of Large Language Models
Weiqin Wang, Yile Wang and Hui Huang 14410

SemanticCamo: Jailbreaking Large Language Models through Semantic Camouflage
Jihui Yan, Xiaocui Yang, Daling Wang, Shi Feng, Yifei Zhang and Yinzhi Zhao 14427

Assigning Distinct Roles to Quantized and Low-Rank Matrices Toward Optimal Weight Decomposition
Yoonjun Cho, Soeun Kim, Dongjae Jeon, Kyelim Lee, Beomsoo Lee and Albert No 14453

Better Process Supervision with Bi-directional Rewarding Signals
Wenxiang Chen, Wei He, Zhiheng Xi, Honglin Guo, Boyang Hong, Jiazheng Zhang, Nijun Li, Tao Gui, Yun Li, Qi Zhang and Xuanjing Huang 14471

KnowCoder-X: Boosting Multilingual Information Extraction via Code
Yuxin Zuo, Wenxuan Jiang, Wenxuan Liu, Zixuan Li, Long Bai, Hanbin Wang, Yutao Zeng, Xiaolong Jin, Jiafeng Guo and Xueqi Cheng 14486

MEIT: Multimodal Electrocardiogram Instruction Tuning on Large Language Models for Report Generation
Zhongwei Wan, Che Liu, Xin Wang, Chaofan Tao, Hui Shen, Jing Xiong, Rossella Arcucci, Huaxiu Yao and Mi Zhang 14510

Harnessing Large Language Models for Disaster Management: A Survey
Zhenyu Lei, Yushun Dong, Weiyu Li, Rong Ding, Qi R. Wang and Jundong Li 14528

Towards Medical Complex Reasoning with LLMs through Medical Verifiable Problems
Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang and Benyou Wang 14552

Monitoring Decoding: Mitigating Hallucination via Evaluating the Factuality of Partial Response during Generation
Yurui Chang, Bochuan Cao and Lu Lin 14574

LLM Critics Help Catch Bugs in Mathematics: Towards a Better Mathematical Verifier with Natural Language Feedback
Bofei Gao, Zefan Cai, Runxin Xu, Peiyi Wang, Ce Zheng, Runji Lin, Keming Lu, Dayiheng Liu, Chang Zhou, Wen Xiao, Tianyu Liu and Baobao Chang 14588

EvoBench: Towards Real-world LLM-Generated Text Detection Benchmarking for Evolving Large Language Models
Xiao Yu, Yi Yu, Dongrui Liu, Kejiang Chen, Weiming Zhang, Nenghai Yu and Jing Shao . 14605

MMSciBench: Benchmarking Language Models on Chinese Multimodal Scientific Problems
Xinwu Ye, Chengfan Li, Siming Chen, Wei Wei and Robert Tang 14621

Lightweight Query Checkpoint: Classifying Faulty User Queries to Mitigate Hallucinations in Large Language Model Question Answering
Minjoo Son, Jonghak Jang and Misuk Kim 14664

Exploring LLM Annotation for Adaptation of Clinical Information Extraction Models under Data-sharing Restrictions
Seiji Shimizu, Hisada Shohei, Yutaka Uno, Shuntaro Yada, Shoko Wakamiya and Eiji Aramaki 14678

Enhancing the Comprehensibility of Text Explanations via Unsupervised Concept Discovery
Yifan Sun, Danding Wang, Qiang Sheng, Juan Cao and Jintao Li 14695

RecordTwin: Towards Creating Safe Synthetic Clinical Corpora
Seiji Shimizu, Ibrahim Baroud, Lisa Raithel, Shuntaro Yada, Shoko Wakamiya and Eiji Aramaki 14714

Beyond Surface-Level Patterns: An Essence-Driven Defense Framework Against Jailbreak Attacks in LLMs
Shiyu Xiang, Ansen Zhang, Yanfei Cao, Fan Yang and Ronghao Chen 14727

Multimodal Invariant Sentiment Representation Learning
Aoqiang Zhu, Min Hu, Xiaohua Wang, Jiaoyun Yang, Yiming Tang and Ning An 14743

ChuLo: Chunk-Level Key Information Representation for Long Document Understanding
Yan Li, Caren Han, Yue Dai and Feiqi Cao 14756

REVS: Unlearning Sensitive Information in Language Models via Rank Editing in the Vocabulary Space
Tomer Ashuach, Martin Tutek and Yonatan Belinkov 14774

Is External Information Useful for Stance Detection with LLMs?
Quang Minh Nguyen and Taegyoon Kim 14798

Benchmarking Query-Conditioned Natural Language Inference
Marc E. Canby, Xinchi Chen, Xing Niu, Jifan Chen, Bonan Min, Sergul Aydore and Vittorio Castelli 14808

Flowchart-Based Decision Making with Large Language Models
Yuuki Yamanaka, Hiroshi Takahashi and Tomoya Yamashita 14836

NarGINA: Towards Accurate and Interpretable Children's Narrative Ability Assessment via Narrative Graphs
Jun Zhong, Longwei Xu, Li Kong, Xianzhuo Li, Dandan Liang and Junsheng Zhou 14843

Improving Efficiency in Large Language Models via Extendable Block Floating Point Representation
Dongyang Li, Zeyang Li, Bosheng Liu and Jigang Wu 14861

EpiCoDe: Boosting Model Performance Beyond Training with Extrapolation and Contrastive Decoding
Mingxu Tao, Jie Hu, Mingchuan Yang, Yunhuai Liu, Dongyan Zhao and Yansong Feng . . . 14874

NativQA: Multilingual Culturally-Aligned Natural Query for LLMs
Md. Arid Hasan, Maram Hasanain, Fatema Ahmad, Sahinur Rahman Laskar, Sunaya Upadhyay, Vrunda N Sukhadia, Mucahid Kutlu, Shammur Absar Chowdhury and Firoj Alam 14886

DoCIA: An Online Document-Level Context Incorporation Agent for Speech Translation
Xinglin Lyu, Wei Tang, Yuang Li, Xiaofeng Zhao, Ming Zhu, Junhui Li, Yunfei Lu, Min Zhang, Daimeng Wei, Hao Yang and Min Zhang 14910

RISE: Reasoning Enhancement via Iterative Self-Exploration in Multi-hop Question Answering
Bolei He, Xinran He, Mengke Chen, Xianwei Xue, Ying Zhu and Zhen-Hua Ling 14925

VADE: Visual Attention Guided Hallucination Detection and Elimination
Vishnu Prabhakaran, Purav Aggarwal, Vinay Kumar Verma, Gokul Swamy and Anoop Saladi 14949

PGPO: Enhancing Agent Reasoning via Pseudocode-style Planning Guided Preference Optimization
Zouying Cao, Runze Wang, Yifei Yang, Xinbei Ma, Xiaoyong Zhu, Bo Zheng and Hai Zhao14966

The Effectiveness of Uncased Tokeniziaion for Clinical Notes
Cory Paik and Katharina Von Der Wense 14986

AMXFP4: Taming Activation Outliers with Asymmetric Microscaling Floating-Point for 4-bit LLM Inference
Janghwan Lee, Jiwoong Park, Jinseok Kim, Yongjik Kim, Jungju Oh, Jinwook Oh and Jungwook Choi 14993

Improving Continual Pre-training Through Seamless Data Packing
Ruicheng Yin, Xuan Gao, Changze Lv, Xiaohua Wang, Xiaoqing Zheng and Xuanjing Huang 15014

The Impact of Name Age Perception on Job Recommendations in LLMs
Mahammed Kamruzzaman and Gene Louis Kim 15033

DAPI: Domain Adaptive Toxicity Probe Vector Intervention, for Fine-Grained Detoxification
Cho Hyeonsu, Dooyoung Kim and Youngjoong Ko 15059

Task Knowledge Injection via Interpolations and Reinstatement for Large Language Model Generalization
Yukun Zhao, Lingyong Yan, Zhenyang Li, Shuaiqiang Wang, Zhumin Chen, Zhaochun Ren and Dawei Yin 15070

STARS: A Unified Framework for Singing Transcription, Alignment, and Refined Style Annotation
Wenxiang Guo, Yu Zhang, Changhao Pan, Zhiyuan Zhu, Ruiqi Li, ZheTao Chen, Wenhao Xu, Fei Wu and Zhou Zhao 15081

Unveiling the Key Factors for Distilling Chain-of-Thought Reasoning
Xinghao Chen, Zhijing Sun, Guo Wenjin, Miaoran Zhang, Yanjun Chen, Yirong Sun, Hui Su, Yijie Pan, Dietrich Klakow, Wenjie Li and Xiaoyu Shen 15094

INT: Establishing Information Transfer for Multilingual Intent Detection and Slot Filling
Di Wu, Liting Jiang, Bohui Mao, Hongyan Xie, Haoxiang Su, Zhongjiang He, Ruiyu Fang, Shuangyong Song, Hao Huang and Xuelong Li 15120

Enhancing LLM Agent Safety via Causal Influence Prompting
Dongyoon Hahm, Woogyeol Jin, June Suk Choi, Sungsoo Ahn and Kimin Lee 15143

Position Paper: MeMo: Towards Language Models with Associative Memory Mechanisms
Fabio Massimo Zanzotto, Elena Sofia Ruzzetti, Giancarlo A. Xompero, Leonardo Ranaldi, Davide Venditti, Federico Ranaldi, Cristina Giannone, Andrea Favalli and Raniero Romagnoli 15169

DeRAGEC: Denoising Named Entity Candidates with Synthetic Rationale for ASR Error Correction
Solee Im, Wonjun Lee, JinMyeong AN, Yunsu Kim, Jungseul Ok and Gary Lee 15181

Rehearse With User: Personalized Opinion Summarization via Role-Playing based on Large Language Models
Yanyue Zhang, Yulan He and Deyu Zhou 15194

AdParaphrase v2.0: Generating Attractive Ad Texts Using a Preference-Annotated Paraphrase Dataset
Soichiro Murakami, Peinan Zhang, Hidetaka Kamigaito, Hiroya Takamura and Manabu Okumura 15212

Beyond the Average Reader: the Reader Embedding Approach
Calogero Jerik Scozzaro, Matteo Delsanto and Daniele P. Radicioni 15231

PredictaBoard: Benchmarking LLM Score Predictability
Lorenzo Pacchiardi, Konstantinos Voudouris, Ben Slater, Fernando Martínez-Plumed, Jose Hernandez-Orallo, Lexin Zhou and Wout Schellaert 15245

FedDQC: Data Quality Control in Federated Instruction-tuning of Large Language Models
Yaxin Du, Rui Ye, Fengting Yuchi, Wanru Zhao, Jingjing Qu, Yanfeng Wang and Siheng Chen 15267

Weed Out, Then Harvest: Dual Low-Rank Adaptation is an Effective Noisy Label Detector for Noise-Robust Learning
Bo Yuan, Yulin Chen and Yin Zhang 15292

"I understand your perspective": LLM Persuasion through the Lens of Communicative Action Theory
Esra Dönmez and Agnieszka Falenska 15312

Nunchi-Bench: Benchmarking Language Models on Cultural Reasoning with a Focus on Korean Superstition
Kyuhee Kim and Sangah Lee 15328

Let's Be Self-generated via Step by Step: A Curriculum Learning Approach to Automated Reasoning with Large Language Models
Kangyang Luo, Zichen Ding, Zhenmin Weng, Lingfeng Qiao, Meng Zhao, Xiang Li, di Yin and Jinlong Shu 15343

daDPO: Distribution-Aware DPO for Distilling Conversational Abilities
Zhengze Zhang, Shiqi Wang, Yiqun Shen, Simin Guo, Dahua Lin, Xiaoliang Wang, Nguyen Cam-Tu and Fei Tan 15421

Consultant Decoding: Yet Another Synergistic Mechanism
Chuanghao Ding, Jiaping Wang, Ziqing Yang, Xiaoliang Wang, Dahua Lin, Nguyen Cam-Tu and Fei Tan 15438

IntelliCockpitBench: A Comprehensive Benchmark to Evaluate VLMs for Intelligent Cockpit
Liang Lin, Siyuan Chai, Jiahao Wu, Hongbing Hu, Xiaotao Gu, Hao Hu, Fan Zhang, Wei Wang and Dan Zhang 15453

Analyzing Political Bias in LLMs via Target-Oriented Sentiment Classification
Akram Elbouanani, Evan Dufraisse and Adrian Popescu 15476

PISCO: Pretty Simple Compression for Retrieval-Augmented Generation
Maxime Louis, Hervé Déjean and Stéphane Clinchant 15506

AnchorCoT: Anchors Pave the Way for Multi-hop Reasoning
Tianshi Ming, Xian Wu, Yingying Zhang, Zichuan Fu and Dawei Cheng 15522

Token Pruning in Multimodal Large Language Models: Are We Solving the Right Problem?
Zichen Wen, Yifeng Gao, Weijia Li, Conghui He and Linfeng Zhang 15537

Federated Data-Efficient Instruction Tuning for Large Language Models
Zhen Qin, Zhaomin Wu, Bingsheng He and Shuiguang Deng 15550

They want to pretend not to understand: The Limits of Current LLMs in Interpreting Implicit Content of Political Discourse
Walter Paci, Alessandro Panunzi and Sandro Pezzelle 15569

ZeroNER: Fueling Zero-Shot Named Entity Recognition via Entity Type Descriptions
Alessio Cocchieri, Marcos Martínez Galindo, Giacomo Frisoni, Gianluca Moro, Claudio Sartori and Giuseppe Tagliavini 15594

Do Large Language Models Have "Emotion Neurons"? Investigating the Existence and Role
Jaewook Lee, Woojin Lee, Oh-Woog Kwon and Harksoo Kim 15617

Grammar-Based Code Representation: Is It a Worthy Pursuit for LLMs?
Qingyuan Liang, Zhao Zhang, Zeyu Sun, Zheng Lin, Qi Luo, Xiao Yueyi, Yizhou Chen, Yuqun Zhang, Haotian Zhang, Lu Zhang, Chenbin Chenbin and Yingfei Xiong 15640

Investigating Inference-time Scaling for Chain of Multi-modal Thought: A Preliminary Study
Yujie Lin, Ante Wang, Moye Chen, Jingyao Liu, Hao Liu, Jinsong Su and Xinyan Xiao . . . 15654

UI-E2I-Synth: Advancing GUI Grounding with Large-Scale Instruction Synthesis
Xinyi Liu, Xiaoyi Zhang, Ziyun Zhang and Yan Lu 15668

A Study into Investigating Temporal Robustness of LLMs
Jonas Wallat, Abdelrahman Abdallah, Adam Jatowt and Avishek Anand 15685

ToolExpNet: Optimizing Multi-Tool Selection in LLMs with Similarity and Dependency-Aware Experience Networks
Zijing Zhang, Zhanpeng Chen, He Zhu, Ziyang Chen, Nan du and Xiaolong Li 15706

SPILL: Domain-Adaptive Intent Clustering based on Selection and Pooling with Large Language Models
I-Fan Lin, Faegheh Hasibi and Suzan Verberne 15723

How Far are LLMs from Being Our Digital Twins? A Benchmark for Persona-Based Behavior Chain Simulation
Rui Li, Heming Xia, Xinfeng Yuan, Qingxiu Dong, Lei Sha, Wenjie Li and Zhifang Sui . . 15738

GRI-QA: a Comprehensive Benchmark for Table Question Answering over Environmental Data
Michele Luca Contalbo, Sara Pederzoli, Francesco Del Buono, Venturelli Valeria, Francesco Guerra and Matteo Paganelli 15764

WebUIBench: A Comprehensive Benchmark for Evaluating Multimodal Large Language Models in WebUI-to-Code
Zhiyu Lin, Zhengda Zhou, Zhiyuan Zhao, Tianrui Wan, Yilun Ma, Junyu Gao and Xuelong Li 15780

Optimizing Multi-Hop Document Retrieval Through Intermediate Representations
Linjiaen Linjiaen, Jingyu Liu and Yingbo Liu 15798

Towards Better Understanding of Program-of-Thought Reasoning in Cross-Lingual and Multilingual Environments
Patomporn Payoungkhamdee, Pume Tuchinda, Jinheon Baek, Samuel Cahyawijaya, Can Udomcharoenchaikit, Potsawee Manakul, Peerat Limkonchotiwat, Ekapol Chuangsuwanich and Sarana Nutanong 15810

A Fully Automated Pipeline for Conversational Discourse Annotation: Tree Scheme Generation and Labeling with Large Language Models
Kseniia Petukhova and Ekaterina Kochmar 15829

Can Language Models Serve as Analogy Annotators?
Xiaojing Zhang and Bochen Lyu 15853

Reward Generalization in RLHF: A Topological Perspective
Tianyi Alex Qiu, Fanzhi Zeng, Jiaming Ji, Dong Yan, Kaile Wang, Jiayi Zhou, Yang Han, Josef Dai, Xuehai Pan and Yaodong Yang 15884

Enhanced Data Synthesis for LLM through Reasoning Structures Generated by Hierarchical GFlowNet
Tianpeng Bu, Minying Zhang, Hongtao Duan, Shurui Li, Lulu hu and Yu Li 15931

Capturing Nuanced Preferences: Preference-Aligned Distillation for Small Language Models
Yanggan Gu, Junzhuo Li, Sirui Huang, Xin Zou, Zhenghua Li and Xuming Hu 15959

Token-level Preference Self-Alignment Optimization for Multi-style Outline Controllable Generation
Zihao Li, Xuekong Xu, Ziyao Chen, Lixin Zou, Ethanhjwu Ethanhjwu, Qiang Chen and Chenliang Li 15974

HatePRISM: Policies, Platforms, and Research Integration. Advancing NLP for Hate Speech Proactive Mitigation
Naquee Rizwan, Seid Muhie Yimam, Daryna Dementieva, Dr. Florian Skupin, Tim Fischer, Daniil Moskovskiy, Aarushi Ajay Borkar, Robert Geislinger, Punyajoy Saha, Sarthak Roy, Martin Semmann, Alexander Panchenko, Chris Biemann and Animesh Mukherjee 16008

Local Look-Ahead Guidance via Verifier-in-the-Loop for Automated Theorem Proving
Sara Rajaee, Kumar Pratik, Gabriele Cesa and Arash Behboodi 16023

Generalizable Cross-Lingual Cognitive Distortion Detection with Standardized Annotations and Multi-Task Learning
Hongzhi Qi, Nan Bai, Jianqiang Li, Wei Zhai, Qing Zhao, Qi Gao, Bing Xiang Yang and Guanghui FU 16041

How Do Multilingual Language Models Remember Facts?
Constanza Fierro, Negar Foroutan, Desmond Elliott and Anders Søgaard 16052

SeqPO-SiMT: Sequential Policy Optimization for Simultaneous Machine Translation
Ting Xu, Zhichao Huang, Jiankai Sun, Shanbo Cheng and Wai Lam 16107

Do Large Language Models Know Folktales? A Case Study of Yokai in Japanese Folktales
Ayuto Tsutsumi and Yuu Jinnai 16124

BOSE: A Systematic Evaluation Method Optimized for Base Models
Hongzhi Luan, Changxin Tian, Zhaoxin Huan, Xiaolu Zhang, Kunlong Chen, Zhiqiang Zhang and Jun Zhou 16147

DPGA-TextSyn: Differentially Private Genetic Algorithm for Synthetic Text Generation
Zhonghao Sun, Zhiliang Tian, Yiping Song, Yuyi Si, Juhua Zhang, Minlie Huang, Kai Lu, Zeyu Xiong, Xinwang Liu and Dongsheng Li 16159

Semantic Aware Linear Transfer by Recycling Pre-trained Language Models for Cross-lingual Transfer
Seungyoon Lee, Seongtae Hong, Hyeonseok Moon and Heuiseok Lim 16180

Boost, Disentangle, and Customize: A Robust System2-to-System1 Pipeline for Code Generation
Kounianhua Du, Hanjing Wang, Jianxing Liu, Jizheng Chen, Xinyi Dai, Yasheng Wang, Ruiming Tang, Yong Yu, Jun Wang and Weinan Zhang 16194

On the Consistency of Commonsense in Large Language Models
Guozheng Li, Peng Wang, Wenjun Ke, Zijie Xu, Jiajun Liu and Ziyu Shang 16205

Statement-Tuning Enables Efficient Cross-lingual Generalization in Encoder-only Models
Ahmed Elshabrawy, Thanh-Nhi Nguyen, Yeeun Kang, Lihan Feng, Annant Jain, Faadil Abdullah Shaikh, Jonibek Mansurov, Mohamed Fazli Mohamed Imam, Jesus-German Ortiz-Barajas, Rendi Chevi and Alham Fikri Aji 16226

Evaluating Large Language Models for Confidence-based Check Set Selection
Jane Arleth Dela Cruz, Iris Hendrickx and Martha Larson 16249

Training Multi-Modal LLMs through Dialogue Planning for HRI
Claudiu Daniel Hromei, Federico Borazio, Andrea Sensi, Elisa Passone, Danilo Croce and Roberto Basili 16266

MVL-SIB: A Massively Multilingual Vision-Language Benchmark for Cross-Modal Topical Matching
Fabian David Schmidt, Florian Schneider, Chris Biemann and Goran Glavaš 16285

The Rise of Darkness: Safety-Utility Trade-Offs in Role-Playing Dialogue Agents
Yihong Tang, Kehai Chen, Xuefeng Bai, Zheng-Yu Niu, Bo Wang, Jie Liu and Min Zhang 16313

SynGraph: A Dynamic Graph-LLM Synthesis Framework for Sparse Streaming User Sentiment Modeling
Xin Zhang, Qiyu Wei, Yingjie Zhu, Linhai Zhang, Deyu Zhou and Sophia Ananiadou 16338

Enhancing Tool Learning in Large Language Models with Hierarchical Error Checklists
Yue Cui, Liuyi Yao, Shuchang Tao, Weijie Shi, Yaliang Li, Bolin Ding and Xiaofang Zhou 16357

A Large and Balanced Corpus for Fine-grained Arabic Readability Assessment
Khalid Elmadani, Nizar Habash and Hanada Taha 16376

Can Medical Vision-Language Pre-training Succeed with Purely Synthetic Data?
Che Liu, Zhongwei Wan, Haozhe Wang, Yinda Chen, Talha Qaiser, Chen Jin, Nikolay Burlutskiy, Fariba Yousefi and Rossella Arcucci 16401

See the World, Discover Knowledge: A Chinese Factuality Evaluation for Large Vision Language Models
Jihao Gu, Yingyao Wang, Pi Bu, Chen Wang, Ziming Wang, Tengtao Song, Donglai Wei, Jiale Yuan, Yingxiu Zhao, Yancheng He, Shilong Li, Jiaheng Liu, Meng Cao, Jun Song, Yingshui Tan, Xiang Li, Wenbo Su, Xiaoyong Zhu and Bo Zheng 16422

Argus: Benchmarking and Enhancing Vision-Language Models for 3D Radiology Report Generation
Che Liu, Zhongwei Wan, Yuqi Wang, Hui Shen, Haozhe Wang, Kangyu Zheng, Mi Zhang and Rossella Arcucci 16448

Resource-Friendly Dynamic Enhancement Chain for Multi-Hop Question Answering
Binquan Ji, Haibo Luo, YifeiLu YifeiLu, Lei Hei, Jiaqi Wang, Tingjing Liao, Wang Lingyu, Shichao Wang and Feiliang Ren 16461

Evaluating LLMs' Assessment of Mixed-Context Hallucination Through the Lens of Summarization
Siya Qi, Rui Cao, Yulan He and Zheng Yuan 16480

TUBA: Cross-Lingual Transferability of Backdoor Attacks in LLMs with Instruction Tuning
Xuanli He, Jun Wang, Qiongkai Xu, Pasquale Minervini, Pontus Stenetorp, Benjamin I. P. Rubinstein and Trevor Cohn 16504

Eliciting Textual Descriptions from Representations of Continuous Prompts
Daniela Gottesman, Mor Geva and Dana Ramati 16545

Mitigating Hallucination in Multimodal Large Language Model via Hallucination-targeted Direct Preference Optimization
Yuhan Fu, Ruobing Xie, Xingwu Sun, Zhanhui Kang and Xirong Li 16563

Review-Instruct: A Review-Driven Multi-Turn Conversations Generation Method for Large Language Models
Jiangxu Wu, Cong Wang, Tianhuang Su, Lin Haozhi, JunYang JunYang, Zhangchao Zhangchao, Binqiang Pan, SongpanYang SongpanYang, Mingpeng Mingpeng, Kai Shi and Zixian LI 16578

Why Uncertainty Estimation Methods Fall Short in RAG: An Axiomatic Analysis
Heydar Soudani, Evangelos Kanoulas and Faegheh Hasibi 16596

EuroVerdict: A Multilingual Dataset for Verdict Generation Against Misinformation
Daniel Russo, Fariba Sadeghi, Stefano Menini and Marco Guerini 16617

LoFTI: Localization and Factuality Transfer to Indian Locales
Sona Elza Simon, Soumen Kumar Mondal, Abhishek Singhania, Sayambhu Sen and Preethi Jyothi 16635

Hierarchical Retrieval with Evidence Curation for Open-Domain Financial Question Answering on Standardized Documents
Jaeyoung Choe, Jihoon Kim and Woohwan Jung 16663

GNN-RAG: Graph Neural Retrieval for Efficient Large Language Model Reasoning on Knowledge Graphs
Costas Mavromatis and George Karypis 16682

ASTRID - An Automated and Scalable TRIaD for the Evaluation of RAG-based Clinical Question Answering Systems
Yajie Vera He, Mohita Chowdhury, Jared Joselowitz, Aisling Higham and Ernest Lim 16700

On Entity Identification in Language Models
Masaki Sakata, Benjamin Heinzerling, Sho Yokoi, Takumi Ito and Kentaro Inui 16717

RAPID: Efficient Retrieval-Augmented Long Text Generation with Writing Planning and Information Discovery
Hongchao Gu, Dexun Li, Kuicai Dong, Hao Zhang, Hang Lv, Hao Wang, Defu Lian, Yong Liu and Enhong Chen 16742

CHARPEVAL: Benchmarking Large Language Models' Contextual Reasoning in Knowledge-Grounded Dialogue
Abbas Ghaddar, David Alfonso-Hermelo, Philippe Langlais, Boxing Chen and Prasanna Parthasarathi 16764

Ask in Any Modality: A Comprehensive Survey on Multimodal Retrieval-Augmented Generation
Mohammad Mahdi Abootorabi, Amirhosein Zobeiri, Mahdi Dehghani, Mohammadali Mohammadkhani, Bardia Mohammadi, Omid Ghahroodi, Mahdieh Soleymani Baghshah and Ehsaneddin Asgari 16776

Debate4MATH: Multi-Agent Debate for Fine-Grained Reasoning in Math
Shaowei Zhang and Deyi Xiong 16810

Disambiguate First, Parse Later: Generating Interpretations for Ambiguity Resolution in Semantic Parsing
Irina Saparina and Mirella Lapata 16825

The Anatomy of Evidence: An Investigation Into Explainable ICD Coding
Katharina Beckh, Elisa Studeny, Sujan Sai Gannamaneni, Dario Antweiler and Stefan Rueping 16840

AVG-LLaVA: An Efficient Large Multimodal Model with Adaptive Visual Granularity
Zhibin Lan, Liqiang Niu, Fandong Meng, Wenbo Li, Jie Zhou and Jinsong Su 16852

Word Form Matters: LLMs' Semantic Reconstruction under Typoglycemia
Chenxi Wang, Tianle Gu, Zhongyu Wei, Lang Gao, Zirui Song and Xiuying Chen 16870

LLM-based Translation Inference with Iterative Bilingual Understanding
Andong Chen, Kehai Chen, Yang Xiang, Xuefeng Bai, Muyun Yang, Yang Feng, Tiejun Zhao and Min Zhang 16886

Vulnerability of Text-to-Image Models to Prompt Template Stealing: A Differential Evolution Approach
Yurong Wu, Fangwen Mu, Qiuhong Zhang, Jinjing Zhao, Xinrun Xu, Lingrui Mei, Yang Wu, Lin Shi, Junjie Wang, Zhiming Ding and Yiwei Wang 16903

mStyleDistance: Multilingual Style Embeddings and their Evaluation
Justin Qiu, Jiacheng Zhu, Ajay Patel, Marianna Apidianaki and Chris Callison-Burch 16917

SeqMMR: Sequential Model Merging and LLM Routing for Enhanced Batched Sequential Knowledge Editing
Shanbao Qiao, Xuebing Liu, Akshat Gupta and Seung-Hoon Na 16932

ReflectEvo: Improving Meta Introspection of Small LLMs by Learning Self-Reflection
Jiaqi Li, Xinyi Dong, Yang Liu, Zhizhuo Yang, Quansen Wang, Xiaobo Wang, Song-Chun Zhu, Zixia Jia and Zilong Zheng 16948

MAGIC-VQA: Multimodal And Grounded Inference with Commonsense Knowledge for Visual Question Answering
Shuo Yang, Caren Han, Siwen Luo and Eduard Hovy 16967

Automatic Transmission for LLM Tiers: Optimizing Cost and Accuracy in Large Language Models
Injae Na, Keonwoong Noh and Woohwan Jung 16987

Low-Rank Interconnected Adaptation across Layers
Yibo Zhong, Jinman Zhao and Yao Zhou 17005

GaRAGe: A Benchmark with Grounding Annotations for RAG Evaluation
Ionut Teodor Sorodoc, Leonardo F. R. Ribeiro, Rexhina Blloshmi, Christopher Davis and Adrià de Gispert 17030

Change Entity-guided Heterogeneous Representation Disentangling for Change Captioning
Yi Li, Yunbin Tu, Liang Li, Li Su and Qingming Huang 17050

RAG-RewardBench: Benchmarking Reward Models in Retrieval Augmented Generation for Preference Alignment
Zhuoran Jin, Hongbang Yuan, Tianyi Men, Pengfei Cao, Yubo Chen, Jiexin Xu, Huaijun Li, Xiaojian Jiang, Kang Liu and Jun Zhao 17061

Generate, Discriminate, Evolve: Enhancing Context Faithfulness via Fine-Grained Sentence-Level Self-Evolution
Kun LI, Tianhua Zhang, Yunxiang Li, Hongyin Luo, Abdalla Mohamed Salama Sayed Moustafa, Xixin Wu, James R. Glass and Helen M. Meng 17091

PAM: Paraphrase AMR-Centric Evaluation Metric
Afonso Sousa and Henrique Lopes Cardoso 17106

VP-MEL: Visual Prompts Guided Multimodal Entity Linking
Hongze Mi, Jinyuan Li, Zhangxuying Zhangxuying, Haoran Cheng, Jiahao Wang, Di Sun and Gang Pan 17122

FADE: Why Bad Descriptions Happen to Good Features
Bruno Puri, Aakriti Jain, Elena Golimblevskaia, Patrick Kahardipraja, Thomas Wiegand, Wojciech Samek and Sebastian Lapuschkin 17138

In the LLM era, Word Sense Induction remains unsolved
Anna Mosolova, Marie Candito and Carlos Ramisch 17161

Navigating the Political Compass: Evaluating Multilingual LLMs across Languages and Nationalities
Chadi Helwe, Oana Balalau and Davide Ceolin 17179

Who Can Withstand Chat-Audio Attacks? An Evaluation Benchmark for Large Audio-Language Models
Wanqi Yang, Yanda Li, Meng Fang, Yunchao Wei and Ling Chen 17205

Beyond the Tip of Efficiency: Uncovering the Submerged Threats of Jailbreak Attacks in Small Language Models
Sibo Yi, Tianshuo Cong, Xinlei He, Qi Li and Jiaxing Song 17221

EMRs2CSP : Mining Clinical Status Pathway from Electronic Medical Records
Yifei Chen, Ruihui Hou, Jingping Liu and Tong Ruan 17235

A Law Reasoning Benchmark for LLM with Tree-Organized Structures including Factum Probandum, Evidence and Experiences
Jiaxin Shen, Jinan Xu, Huiqi Hu, Luyi Lin, Guoyang Ma, Fei Zheng, Fandong Meng, Jie Zhou and Wenjuan Han 17252

Libra: Leveraging Temporal Images for Biomedical Radiology Analysis
Xi Zhang, Zaiqiao Meng, Jake Lever and Edmond S. L. Ho 17275

Stereotype Detection as a Catalyst for Enhanced Bias Detection: A Multi-Task Learning Approach
Aditya Tomar, Rudra Murthy and Pushpak Bhattacharyya 17304

Filling the Temporal Void: Recovering Missing Publication Years in the Project Gutenberg Corpus Using LLMs
Omar Momen, Manuel Schaaf and Alexander Mehler 17318

ExpliCa: Evaluating Explicit Causal Reasoning in Large Language Models
Martina Miliani, Serena Auriemma, Alessandro Bondielli, Emmanuele Chersoni, Lucia Passaro, Irene Sucameli and Alessandro Lenci 17335

Are Dialects Better Prompters? A Case Study on Arabic Subjective Text Classification
Leila Moudjari and Farah Benamara 17356

Natural Logic at the Core: Dynamic Rewards for Entailment Tree Generation
Jihao Shi, Xiao Ding, Kai Xiong, Hengwei Zhao, Bing Qin and Ting Liu 17372

R.R.: Unveiling LLM Training Privacy through Recollection and Ranking
Wenlong Meng, Guo Zhenyuan, Lenan Wu, Chen Gong, Wenyan Liu, Weixian Li, Chengkun Wei and Wenzhi Chen 17383

Nested-Refinement Metamorphosis: Reflective Evolution for Efficient Optimization of Networking Problems
Shuhan Guo, Nan Yin, James Kwok and Quanming Yao 17398

MC-MKE: A Fine-Grained Multimodal Knowledge Editing Benchmark Emphasizing Modality Consistency
Junzhe Zhang, Huixuan Zhang, Xunjian Yin, Baizhou Huang, Xu Zhang, Xinyu Hu and Xiaojun Wan 17430

Visualising Policy-Reward Interplay to Inform Zeroth-Order Preference Optimisation of Large Language Models
Alessio Galatolo, Zhenbang Dai, Katie Winkle and Meriem Beloucif 17446

Metaphor and Large Language Models: When Surface Features Matter More than Deep Understanding
Elisa Sanchez-Bayona and Rodrigo Agerri 17462

AskQE: Question Answering as Automatic Evaluation for Machine Translation
Dayeon Ki, Kevin Duh and Marine Carpuat 17478

ExPerT: Effective and Explainable Evaluation of Personalized Long-Form Text Generation
Alireza Salemi, Julian Killingback and Hamed Zamani 17516

Bridging Intuitive Associations and Deliberate Recall: Empowering LLM Personal Assistant with Graph-Structured Long-term Memory
Yujie Zhang, Weikang Yuan and Zhuoren Jiang 17533

Each graph is a new language: Graph Learning with LLMs
Huachi Zhou, Jiahe Du, Chuang Zhou, Chang Yang, Yilin Xiao, Yuxuan Xie and Xiao Huang 17548

100-LongBench: Are de facto Long-Context Benchmarks Literally Evaluating Long-Context Ability?
Van Yang, Hongye Jin, Shaochen Zhong, Song Jiang, Qifan Wang, Vipin Chaudhary and Xiaotian Han 17560

Multimodal Fusion and Coherence Modeling for Video Topic Segmentation
Hai Yu, Chong Deng, Qinglin Zhang, Jiaqing Liu, Qian Chen and Wen Wang 17577

Are Your LLMs Capable of Stable Reasoning?
Junnan Liu, Hongwei Liu, Linchen Xiao, Ziyi Wang, Kuikun Liu, Songyang Gao, Wenwei Zhang, Songyang Zhang and Kai Chen 17594

FANNO: Augmenting High-Quality Instruction Data with Open-Sourced LLMs Only
He Zhu, Yifan Ding, Yicheng Tao, Zhiwen Ruan, Yixia Li, Wenjia Zhang, Yun Chen and Guanhua Chen 17633

JEBS: A Fine-grained Biomedical Lexical Simplification Task
William Xia, Ishita Unde, Brian David Ondov and Dina Demner-Fushman 17654

Multi-Hop Reasoning for Question Answering with Hyperbolic Representations
Simon Welz, Lucie Flek and Akbar Karimi 17667

Look & Mark: Leveraging Radiologist Eye Fixations and Bounding boxes in Multimodal Large Language Models for Chest X-ray Report Generation
Yunsoo Kim, Jinge Wu, Su Hwan Kim, Pardeep Vasudev, Jiashu Shen and Honghan Wu . . 17680

Hatevolution: What Static Benchmarks Don't Tell Us
Chiara Di Bonaventura, Barbara McGillivray, Yulan He and Albert Meroño-Peñuela 17695

Tag-Instruct: Controlled Instruction Complexity Enhancement through Structure-based Augmentation
He Zhu, Zhiwen Ruan, Junyou Su, Xingwei He, Yun Chen, Wenjia Zhang and Guanhua Chen 17708

Code-SPA: Style Preference Alignment to Large Language Models for Effective and Robust Code Debugging
Tengfei Wen, Xuanang Chen, Ben He and Le Sun 17730

Open-World Authorship Attribution
Xinhao Tan, Songhua Liu, Xia Cong, Kunjun Li and Xinchao Wang 17744

What is in a name? Mitigating Name Bias in Text Embedding Similarity via Anonymization
Sahil Manchanda and Pannaga Shivaswamy 17759

BenNumEval: A Benchmark to Assess LLMs' Numerical Reasoning Capabilities in Bengali
Kawsar Ahmed, Md Osama, Omar Sharif, Eftekhar Hossain and Mohammed Moshiul Hoque 17782

LLM Agents for Coordinating Multi-User Information Gathering
Harsh Jhamtani, Jacob Andreas and Benjamin Van Durme 17800

C2KD: Cross-layer and Cross-head Knowledge Distillation for Small Language Model-based Recommendation
Xiao Chen, Changyi Ma, Wenqi Fan, Zhaoxiang Zhang and Li Qing 17827

Sign2Vis: Automated Data Visualization from Sign Language
Yao Wan, Yang Wu, Zhen Li, Guobiao Zhang, Hongyu Zhang, Zhou Zhao, Hai Jin and April Wang 17839

Transparentize the Internal and External Knowledge Utilization in LLMs with Trustworthy Citation
Jiajun Shen, Tong Zhou, Yubo Chen, Delai Qiu, Shengping Liu, Kang Liu and Jun Zhao . . 17858

JARVIS-VLA: Post-Training Large-Scale Vision Language Models to Play Visual Games with Keyboards and Mouse
Muyao Li, Zihao Wang, Kaichen He, Xiaojian Ma and Yitao Liang 17878

Generative Frame Sampler for Long Video Understanding
Linli Yao, Haoning Wu, Kun Ouyang, Yuanxing Zhang, Caiming Xiong, Bei Chen, Xu Sun and Junnan Li 17900

Annotating the Annotators: Analysis, Insights and Modelling from an Annotation Campaign on Persuasion Techniques Detection
Davide Bassi, Dimitar Iliyanov Dimitrov, Bernardo D'Auria, Firoj Alam, Maram Hasanain, Christian Moro, Luisa Orrù, Gian Piero Turchi, Preslav Nakov and Giovanni Da San Martino 17918

On the Generalization vs Fidelity Paradox in Knowledge Distillation
Suhas Kamasetty Ramesh, Ayan Sengupta and Tanmoy Chakraborty 17930

BEDAA: Bayesian Enhanced DeBERTa for Uncertainty-Aware Authorship Attribution
Iqra Zahid, Youcheng Sun and Riza Batista-Navarro 17952

Benchmarking the Benchmarks: Reproducing Climate-Related NLP Tasks
Tom Calamai, Oana Balalau and Fabian M. Suchanek 17967

Exploring Supervised Approaches to the Detection of Anthropomorphic Language in the Reporting of NLP Venues
Matthew Shardlow, Ashley Williams, Charlie Roadhouse, Filippos Ventirozos and Piotr Przybyła 18010

PersonaLens: A Benchmark for Personalization Evaluation in Conversational AI Assistants
Zheng Zhao, Clara Vania, Subhradeep Kayal, Naila Khan, Shay B Cohen and Emine Yilmaz 18023

iAgent: LLM Agent as a Shield between User and Recommender Systems
Wujiang Xu, Yunxiao Shi, Zujie Liang, Xuying Ning, Kai Mei, Kun Wang, Xi Zhu, Min Xu and Yongfeng Zhang 18056

FactLens: Benchmarking Fine-Grained Fact Verification
Kushan Mitra, Dan Zhang, Sajjadur Rahman and Estevam Hruschka 18085

Process-based Self-Rewarding Language Models
Shimao Zhang, Xiao Liu, Xin Zhang, Junxiao Liu, Zheheng Luo, Shujian Huang and Yeyun Gong 18097

The Devil Is in the Word Alignment Details: On Translation-Based Cross-Lingual Transfer for Token Classification Tasks
Benedikt Ebing and Goran Glavaš 18111

ShieldHead: Decoding-time Safeguard for Large Language Models
Zitao Xuan, Xiaofeng Mao, Da Chen, Xin Zhang, Yuhan Dong and Jun Zhou 18129

A Survey on Proactive Defense Strategies Against Misinformation in Large Language Models
Shuliang Liu, Hongyi Liu, Aiwei Liu, Duan Bingchen, Zheng Qi, Yibo Yan, He Geng, Peijie Jiang, Jia Liu and Xuming Hu 18144

Smotrom tvoja på ander drogoj verden! Resurrecting Dead Pidgin with Generative Models: Russenorsk Case Study
Alexey Tikhonov, Sergei Shteiner, Anna Bykova and Ivan P. Yamshchikov 18156

PromptCoT: Synthesizing Olympiad-level Problems for Mathematical Reasoning in Large Language Models
Xueliang Zhao, Wei Wu, Jian Guan and Lingpeng Kong 18167

Speculative Sampling via Exponential Races
Szymon Kobus and Deniz Gunduz 18189

Going Beyond Your Expectations in Latency Metrics for Simultaneous Speech Translation
Jorge Iranzo-Sánchez, Javier Iranzo-Sánchez, Adrià Giménez and Jorge Civera 18205

Towards a Design Guideline for RPA Evaluation: A Survey of Large Language Model-Based Role-Playing Agents
Chaoran Chen, Bingsheng Yao, Ruishi Zou, Wenyue Hua, Weimin Lyu, Toby Jia-Jun Li and Dakuo Wang 18229

Recursive Question Understanding for Complex Question Answering over Heterogeneous Personal Data
Philipp Christmann and Gerhard Weikum 18269

PreSumm: Predicting Summarization Performance Without Summarizing
Steven Koniaev, Ori Ernst and Jackie CK Cheung 18289

Mixture of Structural-and-Textual Retrieval over Text-rich Graph Knowledge Bases
Yongjia Lei, Haoyu Han, Ryan A. Rossi, Franck Dernoncourt, Nedim Lipka, Mahantesh M Halappanavar, Jiliang Tang and Yu Wang 18306

Fact Recall, Heuristics or Pure Guesswork? Precise Interpretations of Language Models for Fact Completion
Denitsa Saynova, Lovisa Hagström, Moa Johansson, Richard Johansson and Marco Kuhlmann 18322

FPE2M2: Approaching Lossless and Efficient Quantization with Native Floating Point
Ke Yi, Jianwei Zhang, Zhiying Xu, Xinlong Yang, Yang Zhou, Minmin Sun, Zengke Liu, Tong Zhang, Junyang Lin and Jingren Zhou 18350

Asymmetric Conflict and Synergy in Post-training for LLM-based Multilingual Machine Translation
Tong Zheng, Yan Wen, Huiwen Bao, Junfeng Guo and Heng Huang 18362

VISIAR: Empower MLLM for Visual Story Ideation
Zhaoyang Xia, Somdeb Sarkhel, Mehrab Tanjim, Stefano Petrangeli, Ishita Dasgupta, Yuxiao Chen, Jinxuan XU, Di Liu, Saayan Mitra and Dimitris N. Metaxas 18384

Same Company, Same Signal: The Role of Identity in Earnings Call Transcripts
Ding Yu, Zhuo Liu and Hangfeng He 18403

Understanding and Meeting Practitioner Needs When Measuring Representational Harms Caused by LLM-Based Systems
Emma Harvey, Emily Sheng, Su Lin Blodgett, Alexandra Chouldechova, Jean Garcia-Gathright, Alexandra Olteanu and Hanna Wallach 18423

Mind the (Belief) Gap: Group Identity in the World of LLMs
Angana Borah, Marwa Houalla and Rada Mihalcea 18441

A General Framework to Enhance Fine-tuning-based LLM Unlearning
Jie Ren, Zhenwei Dai, Xianfeng Tang, Hui Liu, Jingying Zeng, Zhen Li, Rahul Goutam, Suhang Wang, Yue Xing, Qi He and Hui Liu 18464

Right Answer, Wrong Score: Uncovering the Inconsistencies of LLM Evaluation in Multiple-Choice Question Answering
Francesco Maria Molfese, Luca Moroni, Luca Gioffré, Alessandro Scirè, Simone Conia and Roberto Navigli 18477

Machine Theory of Mind Needs Machine Validation
Adil Soubki and Owen Rambow 18495

MiniKV: Pushing the Limits of 2-Bit KV Cache via Compression and System Co-Design for Efficient Long Context Inference
Akshat Sharma, Hangliang Ding, Jianping Li, Neel Dani and Minjia Zhang 18506

Sci-LoRA: Mixture of Scientific LoRAs for Cross-Domain Lay Paraphrasing
Ming Cheng, Jiaying Gong and Hoda Eldardiry 18524

Trick or Neat: Adversarial Ambiguity and Language Model Evaluation
Antonia Karamolegkou, Oliver Eberle, Phillip Rust, Carina Kauf and Anders Søgaard 18542

Biases Propagate in Encoder-based Vision-Language Models: A Systematic Analysis From Intrinsic Measures to Zero-shot Retrieval Outcomes
Kshitish Ghate, Tessa Charlesworth, Mona T. Diab and Aylin Caliskan 18562

Stepwise Perplexity-Guided Refinement for Efficient Chain-of-Thought Reasoning in Large Language Models
Yingqian Cui, Pengfei He, Jingying Zeng, Hui Liu, Xianfeng Tang, Zhenwei Dai, Yan Han, Chen Luo, Jing Huang, Zhen Li, Suhang Wang, Yue Xing, Jiliang Tang and Qi He 18581

Can Multimodal Foundation Models Understand Schematic Diagrams? An Empirical Study on Information-Seeking QA over Scientific Papers
Yilun Zhao, Chengye Wang, Chuhan Li and Arman Cohan 18598

MultiChallenge: A Realistic Multi-Turn Conversation Evaluation Benchmark Challenging to Frontier LLMs
Kaustubh Deshpande, Ved Sirdeshmukh, Johannes Baptist Mols, Lifeng Jin, Ed-Yeremai Hernandez-Cardona, Dean Lee, Jeremy Kritz, Willow E. Primack, Summer Yue and Chen Xing 18632

Privacy Ripple Effects from Adding or Removing Personal Information in Language Model Training
Jaydeep Borkar, Matthew Jagielski, Katherine Lee, Niloofar Mireshghallah, David A. Smith and Christopher A. Choquette-Choo 18703

Safety is Not Only About Refusal: Reasoning-Enhanced Fine-tuning for Interpretable LLM Safety
Yuyou Zhang, Miao Li, William Han, Yihang Yao, Zhepeng Cen and Ding Zhao 18727

Is a cute puyfred cute? Context-dependent form-meaning systematicity in LLMs
Jaïr A. Waal and Giovanni Cassani 18747

MetaSynth: Meta-Prompting-Driven Agentic Scaffolds for Diverse Synthetic Data Generation
Haris Riaz, Sourav Sanjukta Bhabesh, Vinayak Arannil, Miguel Ballesteros and Graham Horwood 18770

MVTamperBench: Evaluating Robustness of Vision-Language Models
Amit Agarwal, Srikant Panda, Angeline Charles, Hitesh Laxmichand Patel, Bhargava Kumar, Priyaranjan Pattnayak, Taki Hasan Rafi, Tejaswini Kumar, Hansa Meghwani, Karan Gupta and Dong-Kyu Chae 18804

Multimodal Inconsistency Reasoning (MMIR): A New Benchmark for Multimodal Reasoning Models
Qianqi Yan, Yue Fan, Hongquan Li, Shan Jiang, Yang Zhao, Xinze Guan, Ching-Chen Kuo and Xin Eric Wang 18829

Vision-Language Models Struggle to Align Entities across Modalities
Iñigo Alonso, Gorka Azkune, Ander Salaberria, Jeremy Barnes and Oier Lopez De Lacalle 18846

A Multi-Labeled Dataset for Indonesian Discourse: Examining Toxicity, Polarization, and Demographics Information
Lucky Susanto, Musa Izzanardi Wijanarko, Prasetia Anugrah Pratama, Zilu Tang, Fariz Akyas, Traci Hong, Ika Karlina Idris, Alham Fikri Aji and Derry Tanti Wijaya 18863

MedCite: Can Language Models Generate Verifiable Text for Medicine?
Xiao Wang, Mengjue Tan, Qiao Jin, Guangzhi Xiong, Yu Hu, Aidong Zhang, Zhiyong Lu and Minjia Zhang 18891

Let The Jury Decide: Fair Demonstration Selection for In-Context Learning through Incremental Greedy Evaluation
Sadaf Md Halim, Chen Zhao, Xintao Wu, Latifur Khan, Christan Grant, Fariha Ishrat Rahman and Feng Chen 18914

The Lies Characters Tell: Utilizing Large Language Models to Normalize Adversarial Unicode Perturbations
Portia Cooper, Eduardo Blanco and Mihai Surdeanu 18932

Speech Act Patterns for Improving Generalizability of Explainable Politeness Detection Models
Ahmad Aljanaideh 18945

Systematic Evaluation of Auto-Encoding and Large Language Model Representations for Capturing Author States and Traits
Khushboo Singh, Vasudha Varadarajan, Adithya V Ganesan, August Håkan Nilsson, Nikita Soni, Syeda Mahwish, Pranav Chitale, Ryan L. Boyd, Lyle Ungar, Richard N Rosenthal and H. Schwartz 18955

TReMu: Towards Neuro-Symbolic Temporal Reasoning for LLM-Agents with Memory in Multi-Session Dialogues
Yubin Ge, Salvatore Romeo, Jason Cai, Raphael Shu, Yassine Benajiba, Monica Sunkara and Yi Zhang 18974

Conservative Bias in Large Language Models: Measuring Relation Predictions
Toyin Aguda, Erik Wilson, Allan Anzagira, Simerjot Kaur and Charese Smiley 18989

Mitigating Bias in RAG: Controlling the Embedder
Taeyoun Kim, Jacob Mitchell Springer, Aditi Raghunathan and Maarten Sap 18999

V-ALPHASOCIAL: Benchmark and Self-Reflective Chain-of-Thought Generation for Visual Social Commonsense Reasoning
Zongyu Lin, Zhikun Xu, Xiaohan Song, Yixin Wan, Xingcheng Yao, Tsung-Han Lin, Selina Song, Pranav Subbaraman, Ben Zhou, Kai-Wei Chang and Yizhou Sun 19025

AfroBench: How Good are Large Language Models on African Languages?
Jessica Ojo, Odunayo Ogundepo, Akintunde Oladipo, Kelechi Ogueji, Jimmy Lin, Pontus Stenetorp and David Ifeoluwa Adelani 19048

Training Bilingual LMs with Data Constraints in the Targeted Language
Skyler Seto, Maartje Ter Hoeve, Richard He Bai, Natalie Schluter and David Grangier 19096

ChartQAPro: A More Diverse and Challenging Benchmark for Chart Question Answering
Ahmed Masry, Mohammed Saidul Islam, Mahir Ahmed, Aayush Bajaj, Firoz Kabir, Aaryaman Kartha, Md Tahmid Rahman Laskar, Mizanur Rahman, Shadikur Rahman, Mehrad Shahmohammadi, Megh Thakkar, Md Rizwan Parvez, Enamul Hoque and Shafiq Joty 19123

From Observation to Understanding: Front-Door Adjustments with Uncertainty Calibration for Enhancing Egocentric Reasoning in LVLMs
Shenshen Li, Wenxin Meng, Lei Wang, Hao Yang, Chong Peng, Peng Yan, Fumin Shen, Jingkuan Song, Heng Tao Shen and Xing Xu 19152

Hypothetical Documents or Knowledge Leakage? Rethinking LLM-based Query Expansion
Yejun Yoon, Jaeyoon Jung, Seunghyun Yoon and Kunwoo Park 19170

Worse than Random? An Embarrassingly Simple Probing Evaluation of Large Multimodal Models in Medical VQA
Qianqi Yan, Xuehai He, Xiang Yue and Xin Eric Wang 19188

Optimizing Reasoning for Text-to-SQL with Execution Feedback
Bohan Zhai, Canwen Xu, Yuxiong He and Zhewei Yao 19206

Disentangling Logic: The Role of Context in Large Language Model Reasoning Capabilities
Wenyue Hua, Kaijie Zhu, Lingyao Li, Lizhou Fan, Mingyu Jin, Shuhang Lin, Haochen Xue, Zelong Li, Jindong Wang and Yongfeng Zhang 19219

Sens-Merging: Sensitivity-Guided Parameter Balancing for Merging Large Language Models
Shuqi Liu, Han Wu, Bowei He, Xiongwei Han, Mingxuan Yuan and Linqi Song 19243

EgoNormia: Benchmarking Physical-Social Norm Understanding
MohammadHossein Rezaei, Yicheng Fu, Phil Cuvin, Caleb Ziems, Yanzhe Zhang, Hao Zhu and Diyi Yang 19256

Large Language Models as Neurolinguistic Subjects: Discrepancy between Performance and Competence
Linyang He, Ercong Nie, Helmut Schmid, Hinrich Schuetze, Nima Mesgarani and Jonathan Brennan 19284

The Impact of Large Language Models in Academia: from Writing to Speaking
Mingmeng Geng, Caixi Chen, Yanru Wu, Yao Wan, Pan Zhou and Dongping Chen 19303

X-WebAgentBench: A Multilingual Interactive Web Benchmark for Evaluating Global Agentic System
Peng Wang, Ruihan Tao, Qiguang Chen, Mengkang Hu and Libo Qin 19320

MemBench: Towards More Comprehensive Evaluation on the Memory of LLM-based Agents
Haoran Tan, Zeyu Zhang, Chen Ma, Xu Chen, Quanyu Dai and Zhenhua Dong 19336

Adaptive LoRA Merge with Parameter Pruning for Low-Resource Generation
Ryota Miyano and Yuki Arase 19353

LongAttn: Selecting Long-context Training Data via Token-level Attention
Longyun Wu, Dawei Zhu, Guangxiang Zhao, Zhuocheng Yu, Junfeng Ran, Xiangyu Wong, Lin Sun and Sujian Li 19367

CoRE: Condition-based Reasoning for Identifying Outcome Variance in Complex Events
Sai P Vallurupalli and Francis Ferraro 19381

FaVe: Factored and Verified Search Rationale for Long-form Answer
Jihyuk Kim, Sungjin Lee, Seung-won Hwang and Yang Liu 19402

UnrealLLM: Towards Highly Controllable and Interactable 3D Scene Generation by LLM-powered Procedural Content Generation
SongTang SongTang, Kaiyong Zhao, Lei Wang, Yuliang Li, Xuebo Liu, Junyi Zou, Qiang Wang and Xiaowen Chu 19417

Tree-of-Prompts: Abstracting Control-Flow for Prompt Optimization
Jihyuk Kim, Shubham Garg, Lahari Poddar, Seung-won Hwang and Chris Hench 19436

Outlier-weighed Layerwise Sampling for LLM Fine-tuning
Pengxiang Li, Lu Yin, Xiaowei Gao and Shiwei Liu 19460

KVPR: Efficient LLM Inference with I/O-Aware KV Cache Partial Recomputation
Chaoyi Jiang, Lei Gao, Hossein Entezari Zarch and Murali Annavaram 19474

Direct Behavior Optimization: Unlocking the Potential of Lightweight LLMs
Hongming Yang, Shi Lin, Jun Shao, Changting Lin, Donghai Zhu, Meng Han and Qinglei Kong 19489

Whether LLMs Know If They Know: Identifying Knowledge Boundaries via Debiased Historical In-Context Learning
Bo Lv, Nayu Liu, Yang Shen, Xin Liu, Ping Luo and Yue Yu 19516

How do LLMs' Preferences Affect Event Argument Extraction? CAT: Addressing Preference Traps in Unsupervised EAE
Yunhao Wei, Kai Shuang, Zhiyi Li and Chenrui Mao 19529

Out-of-Distribution Detection via LLM-Guided Outlier Generation for Text-attributed Graph
Xiangwei Lv, Mengze Li, Jingyuan Chen, Zhiang Dong, Sirui Han and Beishui Liao 19544

Document-Level Relation Extraction with Global Relations and Entity Pair Reasoning
Fu Zhang, Yi Yan and Jingwei Cheng 19556

Towards Storage-Efficient Visual Document Retrieval: An Empirical Study on Reducing Patch-Level Embeddings
Yubo Ma, Jinsong Li, Yuhang Zang, Xiaobao Wu, Xiaoyi Dong, Pan Zhang, Yuhang Cao, Haodong Duan, Jiaqi Wang, Yixin Cao and Aixin Sun 19568

Step-by-Step Mastery: Enhancing Soft Constraint Following Ability of Large Language Models
Qingyu Ren, Jie Zeng, Qianyu He, Jiaqing Liang, Yanghua Xiao, Weikang Zhou, Zeye Sun and Fei Yu 19581

ZeroDL: Zero-shot Distribution Learning for Text Clustering via Large Language Models
Hwiyeol Jo, Hyunwoo Lee, Kang Min Yoo and Taiwoo Park 19597

Patterns Over Principles: The Fragility of Inductive Reasoning in LLMs under Noisy Observations
Chunyang Li, Weiqi Wang, Tianshi Zheng and Yangqiu Song 19608

LLMTaxo: Leveraging Large Language Models for Constructing Taxonomy of Factual Claims from Social Media
Haiqi Zhang, Zhengyuan Zhu, Zeyu Zhang and Chengkai Li 19627

AnCast++: Document-Level Evaluation of Graph-based Meaning Representations
Haibo Sun, Jayeol Chun and Nianwen Xue 19642

MMEvol: Empowering Multimodal Large Language Models with Evol-Instruct
Run Luo, Haonan Zhang, Longze Chen, Ting-En Lin, Xiong Liu, Yuchuan Wu, Min Yang, Yongbin Li, Minzheng Wang, Pengpeng Zeng, Lianli Gao, Heng Tao Shen, Yunshui Li, Hamid Alinejad-Rokny, Xiaobo Xia, Jingkuan Song and Fei Huang 19655

SciVerse: Unveiling the Knowledge Comprehension and Visual Reasoning of LMMs on Multi-modal Scientific Problems
Ziyu Guo, Renrui Zhang, Hao Chen, Jialin Gao, Dongzhi Jiang, Jiaze Wang and Pheng-Ann Heng 19683

Exploring Layer-wise Representations of English and Chinese Homonymy in Pre-trained Language Models
Matthew King-Hang Ma, Xie Chenwei, Wenbo Wang and William Shiyuan Wang 19705

DocMEdit: Towards Document-Level Model Editing
Li Zeng, Zeming Liu, Chong Feng, Heyan Huang and Yuhang Guo 19725

Adaptive Detoxification: Safeguarding General Capabilities of LLMs through Toxicity-Aware Knowledge Editing
Yifan Lu, Jing Li, Yigeng Zhou, Yihui Zhang, Wenya Wang, Xiucheng Li, Meishan Zhang, Fangming Liu, Jun Yu and Min Zhang 19744

Evaluating the Long-Term Memory of Large Language Models
Zixi Jia, Qinghua Liu, Hexiao Li, Yuyan Chen and Jiqiang Liu 19759

Explain-then-Process: Using Grammar Prompting to Enhance Grammatical Acceptability Judgments
Russell Scheinberg, Ameeta Agrawal, Amber Shore and So Young Lee 19778

Data Interpreter: An LLM Agent for Data Science
Sirui Hong, Yizhang Lin, Bang Liu, Bangbang Liu, Binhao Wu, Ceyao Zhang, Danyang Li, Jiaqi Chen, Jiayi Zhang, Jinlin Wang, Li Zhang, Lingyao Zhang, Min Yang, Mingchen Zhuge, Taicheng Guo, Tuo Zhou, Wei Tao, Robert Tang, Xiangtao Lu, Xiawu Zheng, Xinbing Liang, Yaying Fei, Yuheng Cheng, Yongxin Ni, Zhibin Gou, Zongze Xu, Yuyu Luo and Chenglin Wu 19796

DReSD: Dense Retrieval for Speculative Decoding
Milan Gritta, Huiyin Xue and Gerasimos Lampouras 19822

Core: Robust Factual Precision with Informative Sub-Claim Identification
Zhengping Jiang, Jingyu Zhang, Nathaniel Weir, Seth Ebner, Miriam Wanner, Kate Sanders, Daniel Khashabi, Anqi Liu and Benjamin Van Durme 19833

Rethinking Diverse Human Preference Learning through Principal Component Analysis
Feng Luo, Rui Yang, Hao Sun, Chunyuan Deng, Jiarui Yao, Jingyan Shen, Huan Zhang and Hanjie Chen 19857

Improving Word Alignment Using Semi-Supervised Learning
Zhongtao Miao, Qiyu Wu, Masaaki Nagata and Yoshimasa Tsuruoka 19871

How Do LLMs Acquire New Knowledge? A Knowledge Circuits Perspective on Continual Pre-Training
Yixin Ou, Yunzhi Yao, Ningyu Zhang, Hui Jin, Jiacheng Sun, Shumin Deng, Zhenguo Li and Huajun Chen 19889

LLM-Symbolic Integration for Robust Temporal Tabular Reasoning
Atharv Kulkarni, Kushagra Dixit, Vivek Srikumar, Dan Roth and Vivek Gupta 19914

Multimodal Large Language Models for Text-rich Image Understanding: A Comprehensive Review
Pei Fu, Tongkun Guan, Zining Wang, Zhentao Guo, Chen Duan, Hao Sun, Boming Chen, Qianyi Jiang, Jiayao Ma, Kai Zhou and Junfeng Luo 19941

PruneVid: Visual Token Pruning for Efficient Video Large Language Models
Xiaohu Huang, Hao Zhou and Kai Han 19959

PromptWizard: Optimizing Prompts via Task-Aware, Feedback-Driven Self-Evolution
Eshaan Agarwal, Raghav Magazine, Joykirat Singh, Vivek Dani, Tanuja Ganu and Akshay Nambi 19974

Exposing Numeracy Gaps: A Benchmark to Evaluate Fundamental Numerical Abilities in Large Language Models
Haoyang LI, Xuejia Chen, Zhanchao Xu, Darian Li, Nicole HU, Fei Teng, Yiming Li, Luyu Qiu, Chen Jason Zhang, Li Qing and Lei Chen 20004

TABGEN-ICL: Residual-Aware In-Context Example Selection for Tabular Data Generation
Liancheng Fang, Aiwei Liu, Hengrui Zhang, Henry Peng Zou, Weizhi Zhang and Philip S. Yu 20027

Benchmarking Multi-National Value Alignment for Large Language Models
Chengyi Ju, Weijie Shi, Chengzhong Liu, Jiaming Ji, Jipeng Zhang, Ruiyuan Zhang, Jiajie Xu, Yaodong Yang, Sirui Han and Yike Guo 20042

MotiveBench: How Far Are We From Human-Like Motivational Reasoning in Large Language Models?
Xixian Yong, Jianxun Lian, Xiaoyuan Yi, Xiao Zhou and Xing Xie 20059

Confidence Improves Self-Consistency in LLMs
Amir Taubenfeld, Tom Sheffer, Eran Ofek, Amir Feder, Ariel Goldstein, Zorik Gekhman and Gal Yona 20090

None of the Above, Less of the Right Parallel Patterns in Human and LLM Performance on Multi-Choice Questions Answering
Zhi Rui Tam, Cheng-Kuang Wu, Chieh-Yen Lin and Yun-Nung Chen 20112

In Search of the Lost Arch in Dialogue: A Dependency Dialogue Acts Corpus for Multi-Party Dialogues
Jon Cai, Brendan King, Peyton Cameron, Susan Windisch Brown, Miriam Eckert, Dananjay Srinivas, George Arthur Baker, V Kate Everson, Martha Palmer, James Martin and Jeffrey Flanigan 20135

ProMind-LLM: Proactive Mental Health Care via Causal Reasoning with Sensor Data
Xinzhe Zheng, Sijie JI, Jiawei Sun, Renqi Chen, Wei Gao and Mani Srivastava. 20150

Debiasing Online Preference Learning via Preference Feature Preservation
Dongyoung Kim, Jinsung Yoon, Jinwoo Shin and Jaehyung Kim 20172

ShortGPT: Layers in Large Language Models are More Redundant Than You Expect
Xin Men, Mingyu Xu, Qingyu Zhang, Qianhao Yuan, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han and Weipeng Chen 20192

ProjectEval: A Benchmark for Programming Agents Automated Evaluation on Project-Level Code Generation
Kaiyuan Liu, Youcheng Pan, Yang Xiang, Daojing He, Jing Li, Yexing Du and Tianrun Gao20205

Unveiling the Lack of LVLM Robustness to Fundamental Visual Variations: Why and Path Forward
Zhiyuan Fan, Yumeng Wang, Sandeep Polisetty and Yi R. Fung 20222

DYNTEXT: Semantic-Aware Dynamic Text Sanitization for Privacy-Preserving LLM Inference
Juhua Zhang, Zhiliang Tian, Minghang Zhu, Yiping Song, Taishu Sheng, Siyi Yang, Qiunan Du, Xinwang Liu, Minlie Huang and Dongsheng Li 20243

InImageTrans: Multimodal LLM-based Text Image Machine Translation
Fei Zuo, Kehai Chen, Yu Zhang, Zhengshan Xue and Min Zhang 20256

FRAME: Boosting LLMs with A Four-Quadrant Multi-Stage Pretraining Strategy
Xuemiao Zhang, Feiyu Duan, Xu Liangyu, Yongwei Zhou, Sirui Wang, Rongxiang Weng, Jingang Wang and Xunliang Cai 20278

When Large Language Models Meet Speech: A Survey on Integration Approaches
Zhengdong Yang, Shuichiro Shimizu, Yahan Yu and Chenhui Chu 20298

KE-MHISTO: Towards a Multilingual Historical Knowledge Extraction Benchmark for Addressing the Long-Tail Problem
Arianna Graciotti, Leonardo Piano, Nicolas Lazzari, Enrico Daga, Rocco Tripodi, Valentina Presutti and Livio Pompianu 20316

TailorKV: A Hybrid Framework for Long-Context Inference via Tailored KV Cache Optimization
Dingyu Yao, Bowen Shen, Zheng Lin, Wei Liu, Jian Luan, Bin Wang and Weiping Wang . 20340

The Elephant in the Room: Exploring the Role of Neutral Words in Language Model Group-Agnostic Debiasing
Xinwei Guo, Jiashi Gao, Junlei Zhou, Jiaxin Zhang, Guanhua Chen, Xiangyu Zhao, Quanying Liu, Haiyan Wu, Xin Yao and Xuetao Wei 20360

LLMs Can Achieve High-quality Simultaneous Machine Translation as Efficiently as Offline
Biao Fu, Minpeng Liao, Kai Fan, Chengxi Li, Liang Zhang, Yidong Chen and Xiaodong Shi 20372

Beyond Completion: A Foundation Model for General Knowledge Graph Reasoning
Yin Hua, Zhiqiang Liu, Mingyang Chen, Zheng Fang, Chi Man Wong, Lingxiao Li, Chi Man Vong, Huajun Chen and Wen Zhang 20396

Generative Error Correction for Emotion-aware Speech-to-text Translation
Zhengdong Yang, Sheng Li and Chenhui Chu 20413

SynapticRAG: Enhancing Temporal Memory Retrieval in Large Language Models through Synaptic Mechanisms
Yuki Hou, Haruki Tamoto, Qinghua Zhao and Homei Miyashita 20422

Localizing and Mitigating Errors in Long-form Question Answering
Rachneet Singh Sachdeva, Yixiao Song, Mohit Iyyer and Iryna Gurevych 20437

EMGLLM: Data-to-Text Alignment for Electromyogram Diagnosis Generation with Medical Numerical Data Encoding
Zefei Long, Zhenbiao Cao, Wei Chen and Zhongyu Wei 20470

LLMVoX: Autoregressive Streaming Text-to-Speech Model for Any LLM
Sambal Shikhar, Mohammed Irfan Kurpath, Sahal Shaji Mullappilly, Jean Lahoud, Fahad Shahbaz Khan, Rao Muhammad Anwer, Salman Khan and Hisham Cholakkal 20481

Act2P: LLM-Driven Online Dialogue Act Classification for Power Analysis
Zhangwenbo Zhangwenbo and Wang Yuhan 20494

MELABenchv1: Benchmarking Large Language Models against Smaller Fine-Tuned Models for Low-Resource Maltese NLP
Kurt Micallef and Claudia Borg 20505

TRATES: Trait-Specific Rubric-Assisted Cross-Prompt Essay Scoring
Sohaila Eltanbouly, Salam Albatarni and Tamer Elsayed 20528

DAST: Context-Aware Compression in LLMs via Dynamic Allocation of Soft Tokens
Shaoshen Chen, Yangning Li, Zishan Xu, Yongqin Zeng, Shunlong Wu, Xinshuo Hu, Zifei Shan, Xin Su, Jiwei Tang, Yinghui Li and Hai-Tao Zheng 20544

A Multi-Expert Structural-Semantic Hybrid Framework for Unveiling Historical Patterns in Temporal Knowledge Graphs
Yimin Deng, Yuxia Wu, Yejing Wang, Guoshuai Zhao, Li Zhu, Qidong Liu, Derong Xu, Zichuan Fu, Xian Wu, Yefeng Zheng, Xiangyu Zhao and Xueming Qian 20553

MWPO: Enhancing LLMs Performance through Multi-Weight Preference Strength and Length Optimization
Shiyue Xu, Fu Zhang, Jingwei Cheng and Linfeng Zhou 20566

CLEAR: Character Unlearning in Textual and Visual Modalities
Alexey Dontsov, Dmitrii Korzh, Alexey Zhavoronkin, Boris Mikheev, Denis Bobkov, Aibek Alanov, Oleg Rogov, Ivan Oseledets and Elena Tutubalina 20582

Assessing the Reasoning Capabilities of LLMs in the context of Evidence-based Claim Verification
John Dougrez-Lewis, Mahmud Elahi Akhter, Federico Ruggeri, Sebastian Löbbers, Yulan He and Maria Liakata 20604

Language Models Lack Temporal Generalization and Bigger is Not Better
Stella Verkijk, Piek Vossen and Pia Sommerauer 20629

DiffLM: Controllable Synthetic Data Generation via Diffusion Language Models
Ying Zhou, Xinyao Wang, Yulei Niu, Yaojie Shen, Lexin Tang, Fan Chen, Ben He, Le Sun and Longyin Wen 20638

Uncertainty Unveiled: Can Exposure to More In-context Examples Mitigate Uncertainty for Large Language Models?
Yifei Wang, Yu Sheng, Linjing Li and Daniel Dajun Zeng 20659

ToolSpectrum: Towards Personalized Tool Utilization for Large Language Models
Zihao Cheng, Hongru Wang, Zeming Liu, Yuhang Guo, Yuanfang Guo, Yunhong Wang and Haifeng Wang 20679

Reverse Preference Optimization for Complex Instruction Following
Xiang Huang, Ting-En Lin, Feiteng Fang, Yuchuan Wu, Hangyu Li, Yuzhong Qu, Fei Huang and Yongbin Li 20700

MMS-LLaMA: Efficient LLM-based Audio-Visual Speech Recognition with Minimal Multimodal Speech Tokens
Jeong Hun Yeo, Hyeongseop Rha, Se Jin Park and Yong Man Ro 20724

Def-DTS: Deductive Reasoning for Open-domain Dialogue Topic Segmentation
Seungmin Lee, Yongsang Yoo, Minhwa Jung and Min Song 20736

Exploring Jailbreak Attacks on LLMs through Intent Concealment and Diversion
Tiehan Cui, Yanxu Mao, Peipei Liu, Congying Liu and Datao You 20754

Verbosity-Aware Rationale Reduction: Sentence-Level Rationale Reduction for Efficient and Effective Reasoning
Joonwon Jang, Jaehee Kim, Wonbin Kweon, Seonghyeon Lee and Hwanjo Yu 20769

Exploring the Role of Mental Health Conversational Agents in Training Medical Students and Professionals: A Systematic Literature Review
Thushari Atapattu, Menasha Thilakaratne, Duc Nhan Do, Mahen Herath and Katrina E. Falkner 20785

Bandit-Based Prompt Design Strategy Selection Improves Prompt Optimizers
Rin Ashizawa, Yoichi Hirose, Nozomu Yoshinari, Kento Uchida and Shinichi Shirakawa . 20799

STORYTELLER: An Enhanced Plot-Planning Framework for Coherent and Cohesive Story Generation
Jiaming Li, Yukun Chen, Ziqiang Liu, Minghuan Tan, Lei Zhang, Yunshui Li, Run Luo, Longze Chen, Jing Luo, Ahmadreza Argha, Hamid Alinejad-Rokny, Wei Zhou and Min Yang 20818

SelectLLM: Query-Aware Efficient Selection Algorithm for Large Language Models
Kaushal Kumar Maurya, Kv Aditya Srivatsa and Ekaterina Kochmar 20847

SkyLLM: Cross-LLM-APIs Federation for Cost-effective Query Processing
Heng Zhao and Yifei Zhu 20864

Matina: A Culturally-Aligned Persian Language Model Using Multiple LoRA Experts
Sara Bourbour Hosseinbeigi, MohammadAli SeifKashani, Javad Seraj, Fatemeh Taherinezhad, Ali Nafisi, Fatemeh Nadi, Iman Barati, Hosein Hasani, Mostafa Amiri and Mostafa Masoudi . . . 20874

PM3-KIE: A Probabilistic Multi-Task Meta-Model for Document Key Information Extraction
Birgit Kirsch, Héctor Allende-Cid and Stefan Rueping 20890

TechniqueRAG: Retrieval Augmented Generation for Adversarial Technique Annotation in Cyber Threat Intelligence Text
Ahmed Lekssays, Utsav Shukla, Husrev Taha Sencar and Md Rizwan Parvez 20913

G2S: A General-to-Specific Learning Framework for Temporal Knowledge Graph Forecasting with Large Language Models
Long Bai, Zixuan Li, Xiaolong Jin, Jiafeng Guo, Xueqi Cheng and Tat-Seng Chua 20927

Disentangling Reasoning Tokens and Boilerplate Tokens For Language Model Fine-tuning
Ziang Ye, Zhenru Zhang, Yang Zhang, Jianxin Ma, Junyang Lin and Fuli Feng 20939

APT: Improving Specialist LLM Performance with Weakness Case Acquisition and Iterative Preference Training
Jun Rao, Zepeng Lin, Xuebo Liu, Xiaopeng Ke, Lian Lian, Dong Jin, Shengjun Cheng, Jun Yu and Min Zhang 20958

EasyEA: Large Language Model is All You Need in Entity Alignment Between Knowledge Graphs
Jingwei Cheng, Chenglong Lu, Linyan Yang, Guoqing Chen and Fu Zhang 20981

An Adaptive Multi-Threshold Loss and a General Framework for Collaborating Losses in Document-Level Relation Extraction
Huangming Xu, Fu Zhang and Jingwei Cheng 20996

RoleMRC: A Fine-Grained Composite Benchmark for Role-Playing and Instruction-Following
Junru Lu, Jiazheng Li, Guodong Shen, Lin Gui, Siyu An, Yulan He, di Yin and Xing Sun . 21008

C²RBench: A Chinese Complex Reasoning Benchmark for Large Language Models
Junru Wu, Tianhao Shen, Linxi Su and Deyi Xiong 21031

Unlocking LLMs' Self-Improvement Capacity with Autonomous Learning for Domain Adaptation
Ke Ji, Junying Chen, Anningzhe Gao, Wenya Xie, Xiang Wan and Benyou Wang 21051

How Personality Traits Shape LLM Risk-Taking Behaviour
John Hartley, Conor Brian Hamill, Dale Seddon, Devesh Batra, Ramin Okhrati and Raad Khraishi 21068

Word-Level Detection of Code-Mixed Hate Speech with Multilingual Domain Transfer
Karin Niederreiter and Dagmar Gromann 21093

Evaluation of Attribution Bias in Generator-Aware Retrieval-Augmented Large Language Models
Amin Abolghasemi, Leif Azzopardi, Seyyed Hadi Hashemi, Maarten de Rijke and Suzan Verberne 21105

Implicit Cross-Lingual Rewarding for Efficient Multilingual Preference Alignment
Wen Yang, Junhong Wu, Chen Wang, Chengqing Zong and Jiajun Zhang 21125

Diagnosing Failures in Large Language Models' Answers: Integrating Error Attribution into Evaluation Framework
Zishan Xu, Shuyi Xie, Qingsong Lv, Shupei Xiao, Linlin Song, Sui Wenjuan and Fan Lin . 21148

Encode Errors: Representational Retrieval of In-Context Demonstrations for Multilingual Grammatical Error Correction
Guangyue Peng, Wei Li, Wen Luo and Houfeng Wang 21166

Preference Curriculum: LLMs Should Always Be Pretrained on Their Preferred Data
Xuemiao Zhang, Xu Liangyu, Feiyu Duan, Yongwei Zhou, Sirui Wang, Rongxiang Weng, Jingang Wang and Xunliang Cai 21181

Can Input Attributions Explain Inductive Reasoning in In-Context Learning?
Mengyu Ye, Tatsuki Kuribayashi, Goro Kobayashi and Jun Suzuki 21199

Modal Dependency Parsing via Biaffine Attention with Self-Loop
Jayeol Chun and Nianwen Xue 21226

Beyond Profile: From Surface-Level Facts to Deep Persona Simulation in LLMs
Zixiao Wang, Duzhen Zhang, Ishita Agarwal, Shen Gao, Le Song and Xiuying Chen 21239

Measuring What Makes You Unique: Difference-Aware User Modeling for Enhancing LLM Personalization
Yilun Qiu, Xiaoyan Zhao, Yang Zhang, Yimeng Bai, Wenjie Wang, Hong Cheng, Fuli Feng and Tat-Seng Chua 21258

VideoRAG: Retrieval-Augmented Generation over Video Corpus
Soyeong Jeong, Kangsan Kim, Jinheon Baek and Sung Ju Hwang 21278

Synergistic Augmentation: Enhancing Cross-Domain Zero-Shot Slot Filling with Small Model-Assisted Large Language Models
Weizhen Li, Junbao Huang, Peijie Huang, Yuhong Xu and Jiekun Fan 21299

A Classifier of Word-Level Variants in Witnesses of Biblical Hebrew Manuscripts
Iglika Nikolova-Stoupak, Maxime Amblard, Sophie Robert-Hayek and Frédérique Rey . . . 21313

NOVA: An Iterative Planning Framework for Enhancing Scientific Innovation with Large Language Models
Xiang hu, Hongyu Fu, Jinge Wang, Yifeng Wang, Zhikun li, Renjun Xu, Yu Lu, Yaochu Jin, Lili Pan and Zhenzhong Lan 21330

Query-Driven Multimodal GraphRAG: Dynamic Local Knowledge Graph Construction for Online Reasoning
Chenyang Bu, Guojie Chang, Zihao Chen, CunYuan Dang, Zhize Wu, Yi He and Xindong Wu 21360

A Survey of Uncertainty Estimation Methods on Large Language Models
Zhiqiu Xia, Jinxuan XU, Yuqian Zhang and Hang Liu 21381

Beyond Single-Value Metrics: Evaluating and Enhancing LLM Unlearning with Cognitive Diagnosis
Yicheng Lang, Kehan Guo, Yue Huang, Yujun Zhou, Haomin Zhuang, Tianyu Yang, Yao Su and Xiangliang Zhang 21397

Natural Language Processing in Support of Evidence-based Medicine: A Scoping Review
Zihan Xu, Haotian Ma, Yihao Ding, Gongbo Zhang, Chunhua Weng and Yifan Peng 21421

How do Transformer Embeddings Represent Compositions? A Functional Analysis
Aishik Nagar, Ishaan Singh Rawal, Mansi Dhanania and Cheston Tan 21444

Entriever: Energy-based Retriever for Knowledge-Grounded Dialog Systems
Yucheng Cai, Ke Li, Yi Huang, Junlan Feng and Zhijian Ou 21462

MONTROSE: LLM-driven Monte Carlo Tree Search Self-Refinement for Cross-Domain Rumor Detection
Shanshan Liu, Menglong Lu, Zhen Huang, Zejiang He, Liu Liu, Zhigang Sun and Dongsheng Li 21475

PEToolLLM: Towards Personalized Tool Learning in Large Language Models
Qiancheng Xu, Yongqi Li, Heming Xia, Fan Liu, Min Yang and Wenjie Li 21488

A Comprehensive Graph Framework for Question Answering with Mode-Seeking Preference Alignment
Quanwei Tang, Sophia Yat Mei Lee, Junshuang Wu, Dong Zhang, Shoushan Li, Erik Cambria and Guodong Zhou 21504

A MISMATCHED Benchmark for Scientific Natural Language Inference
Firoz Shaik, Mobashir Sadat, Nikita Gautam, Doina Caragea and Cornelia Caragea 21524

TagRouter: Learning Route to LLMs through Tags for Open-Domain Text Generation Tasks
Zhou Chen, Zhiqiang Wei, Yuqi Bai, Xue Xiong and Jianmin WU 21539

The Reasoning-Memorization Interplay in Language Models Is Mediated by a Single Direction
Yihuai Hong, Meng Cao, Dian Zhou, Lei Yu and Zhijing Jin 21565

MPBench: A Comprehensive Multimodal Reasoning Benchmark for Process Errors Identification
xu Zhao Pan, Pengfei Zhou, Jiaxin Ai, Wangbo Zhao, Kai Wang, Xiaojiang Peng, Wenqi Shao, Hongxun Yao and Kaipeng Zhang 21586

CRAB: Cross-environment Agent Benchmark for Multimodal Language Model Agents
Tianqi Xu, Linyao Chen, Dai-Jie Wu, Yanjun Chen, Zecheng Zhang, Xiang Yao, Zhiqiang Xie, Yongchao Chen, Shilong Liu, Bochen Qian, Anjie Yang, Zhaoxuan Jin, Jianbo Deng, Philip Torr, Bernard Ghanem and Guohao Li 21607

Towards A "Novel" Benchmark: Evaluating Literary Fiction with Large Language Models
Wenqing Wang, Mingqi Gao, Xinyu Hu and Xiaojun Wan 21648

A Reinforcement Learning Framework for Cross-Lingual Stance Detection Using Chain-of-Thought Alignment
Binghui Li, Minghui Zou, Xiaowang Zhang, Shizhan Chen and Zhiyong Feng 21674

CARE-STaR: Constraint-aware Self-taught Reasoner
Zhiliang Li, Bo Tang, Yijun Niu, Beihong Jin, Qiwen Shi, Yuchen Feng, Zhiyu li, Jie Hu, Mingchuan Yang and Feiyu Xiong 21689

Is It JUST Semantics? A Case Study of Discourse Particle Understanding in LLMs
William Berkeley Sheffield, Kanishka Misra, Valentina Pyatkin, Ashwini Deo, Kyle Mahowald and Junyi Jessy Li 21704

War of Thoughts: Competition Stimulates Stronger Reasoning in Large Language Models
Yibin Chen, Jinyi Liu, Yan Zheng, Yifu Yuan and Jianye Hao 21716

Does Rationale Quality Matter? Enhancing Mental Disorder Detection via Selective Reasoning Distillation
Hoyun Song, Huije Lee, Jisu Shin, Sukmin Cho, Changgeon Ko and Jong C. Park 21738

Rethinking Table Instruction Tuning
Naihao Deng and Rada Mihalcea 21757

CliniDial: A Naturally Occurring Multimodal Dialogue Dataset for Team Reflection in Action During Clinical Operation
Naihao Deng, Kapotaksha Das, Rada Mihalcea, Vitaliy Popov and Mohamed Abouelenien 21781

Chumor 2.0: Towards Better Benchmarking Chinese Humor Understanding from (Ruo Zhi Ba)
Ruiqi He, Yushu He, Longju Bai, Jiarui Liu, Zhenjie Sun, Zenghao Tang, He Wang, Hanchen Xia, Rada Mihalcea and Naihao Deng 21799

Explicit Bayesian Inference to Uncover the Latent Themes of Large Language Models
Raymond Li, Chuyuan Li, Gabriel Murray and Giuseppe Carenini 21819

Improving Occupational ISCO Classification of Multilingual Swiss Job Postings with LLM-Refined Training Data
Ann-Sophie Gnehm and Simon Clematide 21834

Brevity is the soul of sustainability: Characterizing LLM response lengths
Soham Poddar, Paramita Koley, Janardan Misra, Niloy Ganguly and Saptarshi Ghosh 21848

Adversarial Preference Learning for Robust LLM Alignment
Yuanfu Wang, Pengyu Wang, Chenyang Xi, Bo Tang, Junyi Zhu, Wenqiang Wei, Chen Chen, Chao Yang, Jingfeng Zhang, Chaochao Lu, Yijun Niu, Keming Mao, Zhiyu li, Feiyu Xiong, Jie Hu and Mingchuan Yang 21865

gMBA: Expression Semantic Guided Mixed Boolean-Arithmetic Deobfuscation Using Transformer Architectures
Youjeong Roh, Joon-Young Paik, Jingun Kwon and Eun-Sun Cho 21882

READoc: A Unified Benchmark for Realistic Document Structured Extraction
Zichao Li, Aizier Abulaiti, Yaojie Lu, Xuanang Chen, Jia Zheng, Hongyu Lin, Xianpei Han, Shanshan Jiang, Bin Dong and Le Sun 21889

TicTac: Time-aware Supervised Fine-tuning for Automatic Text Dating
Han Ren and Minna Peng 21906

Dolphin: Document Image Parsing via Heterogeneous Anchor Prompting
Hao Feng, Shu Wei, Xiang Fei, Wei Shi, Yingdong Han, Lei Liao, Jinghui Lu, Binghong Wu, Qi Liu, Chunhui Lin, Jingqun Tang, Hao Liu and Can Huang 21919

FanChuan: A Multilingual and Graph-Structured Benchmark For Parody Detection and Analysis
Yilun Zheng, Sha Li, Fangkun Wu, Yang Ziyi, Lin Hongchao, Zhichao Hu, Cai Xinjun, Ziming Wang, Jinxuan Chen, Sitao Luan, Jiahao Xu and Lihui Chen 21937

P-CoT: A Pedagogically-motivated Participatory Chain-of-Thought Prompting for Phonological Reasoning in LLMs
Dongjun Jang, Youngchae Ahn and Hyopil Shin 21958

DynaCode: A Dynamic Complexity-Aware Code Benchmark for Evaluating Large Language Models in Code Generation
Wenhao Hu, Jinhao Duan, Chunchen Wei, Li Zhang, Yue Zhang and Kaidi Xu 21980

Small Encoders Can Rival Large Decoders in Detecting Groundedness
Istabrak Abbes, Gabriele Prato, Quentin Fournier, Fernando Rodriguez, Alaa Boukhary, Adam Elwood and Sarath Chandar 21998

KITAB-Bench: A Comprehensive Multi-Domain Benchmark for Arabic OCR and Document Understanding
Ahmed Heakl, Muhammad Abdullah Sohail, Mukul Ranjan, Rania Elbadry, Ghazi Shazan Ahmad, Mohamed El-Geish, Omar Maher, Zhiqiang Shen, Fahad Shahbaz Khan and Salman Khan 22006

Robustness and Confounders in the Demographic Alignment of LLMs with Human Perceptions of Offensiveness
Shayan Alipour, Indira Sen, Mattia Samory and Tanu Mitra 22025

AL-QASIDA: Analyzing LLM Quality and Accuracy Systematically in Dialectal Arabic
Nathaniel Romney Robinson, Shahd Abdelmoneim, Kelly Marchisio and Sebastian Ruder 22048

Is Large Language Model Performance on Reasoning Tasks Impacted by Different Ways Questions Are Asked?
Seok Hwan Song, Mohna Chakraborty, Qi Li and Wallapak Tavanapong 22066

MutantPrompt: Prompt Optimization via Mutation Under a Budget on Modest-sized LMs
Arijit Nag, Animesh Mukherjee, Niloy Ganguly and Soumen Chakrabarti 22082

Heuristic-based Search Algorithm in Automatic Instruction-focused Prompt Optimization: A Survey
Wendi Cui, Jiaxin Zhang, Zhuohang Li, Hao Sun, Damien Lopez, Kamalika Das, Bradley A. Malin and Sricharan Kumar 22093

CONSENSAGENT: Towards Efficient and Effective Consensus in Multi-Agent LLM Interactions Through Sycophancy Mitigation
Priya Pitre, Naren Ramakrishnan and Xuan Wang 22112

The Structural Safety Generalization Problem
Julius Broomfield, Tom Gibbs, George Ingebretsen, Ethan Kosak-Hine, Tia Nasir, Jason Zhang, Reihaneh Iranmanesh, Sara Pieri, Reihaneh Rabbany and Kellin Pelrine 22134

DPO Kernels: A Semantically-Aware, Kernel-Enhanced, and Divergence-Rich Paradigm for Direct Preference Optimization
Amitava Das, Suranjana Trivedy, Danush Khanna, Yaswanth Narsupalli, Basab Ghosh, Rajarshi Roy, Gurpreet Singh, Vinija Jain, Vasu Sharma, Aishwarya Naresh Reganti and Aman Chadha . 22174

Model-Dependent Moderation: Inconsistencies in Hate Speech Detection Across LLM-based Systems
Neil Fasching and Yphtach Lelkes 22271

Label-semantics Aware Generative Approach for Domain-Agnostic Multilabel Classification
Subhendu Khatuya, Shashwat Naidu, Saptarshi Ghosh, Pawan Goyal and Niloy Ganguly . 22286

Unsupervised Morphological Tree Tokenizer
Qingyang Zhu, Xiang Hu, Pengyu Ji, Wei Wu and Kewei Tu 22299

CausalLink: An Interactive Evaluation Framework for Causal Reasoning
Jinyue Feng and Frank Rudzicz 22313

Toward Global AI Inclusivity: A Large-Scale Multilingual Terminology Dataset (GIST)
Jiarui Liu, Iman Ouzzani, Wenkai Li, Lechen Zhang, Tianyue Ou, Houda Bouamor, Zhijing Jin and Mona T. Diab 22327

A Joint Optimization Framework for Enhancing Efficiency of Tool Utilization in LLM Agents
Bin Wu, Edgar Meij and Emine Yilmaz 22361

When Claims Evolve: Evaluating and Enhancing the Robustness of Embedding Models Against Misinformation Edits
Jabez Magomere, Emanuele La Malfa, Manuel Tonneau, Ashkan Kazemi and Scott A. Hale22374

Splintering Nonconcatenative Languages for Better Tokenization
Bar Gazit, Shaltiel Shmidman, Avi Shmidman and Yuval Pinter 22405

Aria-UI: Visual Grounding for GUI Instructions
Yuhao Yang, Yue Wang, Dongxu Li, Ziyang Luo, Bei Chen, Chao Huang and Junnan Li . . 22418

Revealing Hidden Mechanisms of Cross-Country Content Moderation with Natural Language Processing
Neemesh Yadav, Jiarui Liu, Francesco Ortu, Roya Ensafi, Zhijing Jin and Rada Mihalcea . 22434

Unilogit: Robust Machine Unlearning for LLMs Using Uniform-Target Self-Distillation
Stefan Vasilev, Christian Herold, Baohao Liao, Seyyed Hadi Hashemi, Shahram Khadivi and Christof Monz 22453

Creating a Lens of Chinese Culture: A Multimodal Dataset for Chinese Pun Rebus Art Understanding
Tuo Zhang, Tiantian Feng, Yibin Ni, Mengqin Cao, Ruying Liu, Kiana Avestimehr, Katharine Butler, Yanjun Weng, Mi Zhang, Shrikanth Narayanan and Salman Avestimehr 22473

FastDraft: How to Train Your Draft
Ofir Zafrir, Igor Margulis, Dorin Shteyman, Shira Guskin and Guy Boudoukh 22488

SignMusketeers: An Efficient Multi-Stream Approach for Sign Language Translation at Scale
Shester Gueuwou, Xiaodan Du, Greg Shakhnarovich and Karen Livescu 22506

GUI Agents: A Survey
Dang Nguyen, Jian Chen, Yu Wang, Gang Wu, Namyong Park, Zhengmian Hu, Hanjia Lyu, Junda Wu, Ryan Aponte, Yu Xia, Xintong Li, Jing Shi, Hongjie Chen, Viet Dac Lai, Zhouhang Xie, Sungchul Kim, Ruiyi Zhang, Tong Yu, Mehrab Tanjim, Nesreen K. Ahmed, Puneet Mathur, Seunghyun Yoon, Lina Yao, Branislav Kveton, Jihyung Kil, Thien Huu Nguyen, Trung Bui, Tianyi Zhou, Ryan A. Rossi and Franck Dernoncourt 22522

MEDEC: A Benchmark for Medical Error Detection and Correction in Clinical Notes
Asma Ben Abacha, Wen-wai Yim, Yujuan Fu, Zhaoyi Sun, Meliha Yetisgen, Fei Xia and Thomas Lin 22539

Understanding the Influence of Synthetic Data for Text Embedders
Jacob Mitchell Springer, Vaibhav Adlakha, Siva Reddy, Aditi Raghunathan and Marius Mosbach 22551

Dynamic Knowledge Integration for Evidence-Driven Counter-Argument Generation with Large Language Models
Anar Yeginbergen, Maite Oronoz and Rodrigo Agerri 22568

Tell, Don't Show: Leveraging Language Models' Abstractive Retellings to Model Literary Themes
Li Lucy, Camilla Griffiths, Sarah Levine, Jennifer L Eberhardt, Dorottya Demszky and David Bamman 22585

BottleHumor: Self-Informed Humor Explanation using the Information Bottleneck Principle
EunJeong Hwang, Peter West and Vered Shwartz 22611

Financial Language Model Evaluation (FLaME)
Glenn Matlin, Mika Okamoto, Huzaifa Pardawala, Yang Yang and Sudheer Chava 22633

CausalRAG: Integrating Causal Graphs into Retrieval-Augmented Generation
Nengbo Wang, Xiaotian Han, Jagdip Singh, Jing Ma and Vipin Chaudhary 22680

Towards Safety Reasoning in LLMs: AI-agentic Deliberation for Policy-embedded CoT Data Creation
Tharindu Kumarage, Ninareh Mehrabi, Anil Ramakrishna, Xinyan Zhao, Richard Zemel, Kai-Wei Chang, Aram Galstyan, Rahul Gupta and Charith Peris 22694

Explain then Rank: Scale Calibration of Neural Rankers Using Natural Language Explanations from LLMs
Puxuan Yu, Daniel Cohen, Hemank Lamba, Joel R. Tetreault and Alejandro Jaimes 22716

Beyond instruction-conditioning, MoTE: Mixture of Task Experts for Multi-task Embedding Models
Miguel Romero Calvo, Shuoyang Ding, Corey D Barrett, Georgiana Dinu and George Karypis 22731

Metagent-P: A Neuro-Symbolic Planning Agent with Metacognition for Open Worlds
YanfangZhou YanfangZhou, Yuntao Liu, Xiaodong Li, Yongqiang Zhao, Xintong Wang, Jinlong Tian, Zhenyu Li and Xinhai Xu 22747

Q-STRUM Debate: Query-Driven Contrastive Summarization for Recommendation Comparison
George-Kirollos Saad and Scott Sanner 22765

Inductive Linguistic Reasoning with Large Language Models
Raghav Ramji and Keshav Ramji 22783

Evaluating LLMs' Mathematical and Coding Competency through Ontology-guided Interventions
Pengfei Hong, Navonil Majumder, Deepanway Ghosal, Somak Aditya, Rada Mihalcea and Soujanya Poria 22811

Exploiting Phonetics and Glyph Representation at Radical-level for Classical Chinese Understanding
Junyi Xiang and Maofu Liu 22850

Tokens for Learning, Tokens for Unlearning: Mitigating Membership Inference Attacks in Large Language Models via Dual-Purpose Training
Toan Tran, Ruixuan Liu and Li Xiong 22872

Verify with Caution: The Pitfalls of Relying on Imperfect Factuality Metrics
Ameya Godbole and Robin Jia 22889

TabXEval: Why this is a Bad Table? An eXhaustive Rubric for Table Evaluation
Vihang Pancholi, Jainit Sushil Bafna, Tejas Anvekar, Manish Shrivastava and Vivek Gupta 22913

LADDER: Language-Driven Slice Discovery and Error Rectification in Vision Classifiers
Shantanu Ghosh, Rayan Syed, Chenyu Wang, Vaibhav Choudhary, Binxu Li, Clare B Poynton, Shyam Visweswaran and Kayhan Batmanghelich 22935

GSQ-Tuning: Group-Shared Exponents Integer in Fully Quantized Training for LLMs On-Device Fine-tuning
Sifan Zhou, Shuo Wang, Zhihang Yuan, Mingjia Shi, Yuzhang Shang and Dawei Yang . . . 22971

Evaluation of LLMs in Medical Text Summarization: The Role of Vocabulary Adaptation in High OOV Settings
Gunjan Balde, Soumyadeep Roy, Mainack Mondal and Niloy Ganguly 22989

UniT: One Document, Many Revisions, Too Many Edit Intention Taxonomies
Fangping Lan, Abdullah Aljebreen and Eduard Dragut 23005

Predicting Depression in Screening Interviews from Interactive Multi-Theme Collaboration
Xianbing Zhao, Yiqing Lyu, Di Wang and Buzhou Tang 23025

Your Language Model May Think Too Rigidly: Achieving Reasoning Consistency with Symmetry-Enhanced Training
Yihang Yao, Zhepeng Cen, Miao Li, William Han, Yuyou Zhang, Emerson Liu, Zuxin Liu, Chuang Gan and Ding Zhao 23036

TritonBench: Benchmarking Large Language Model Capabilities for Generating Triton Operators
Jianling li, ShangZhan Li, Zhenye Gao, Qi Shi, Yuxuan Li, Zefan Wang, Jiacheng Huang, WangHaojie WangHaojie, Jianrong Wang, Xu Han, Zhiyuan Liu and Maosong Sun 23053

Just KIDDIN' : Knowledge Infusion and Distillation for Detection of INdecent Memes
Rahul Garg, Trilok Padhi, Hemang Jain, Ugur Kursuncu and Ponnurangam Kumaraguru . . 23067

Dynamic Personality in LLM Agents: A Framework for Evolutionary Modeling and Behavioral Analysis in the Prisoner's Dilemma
Weiqi Zeng, Bo Wang, Dongming Zhao, Zongfeng Qu, Ruifang He, Yuexian Hou and Qinghua Hu 23087

Building A Proof-Oriented Programmer That Is 64% Better Than GPT-4o Under Data Scarcity
Dylan Zhang, Justin Wang and Tianran Sun 23101

On the Robust Approximation of ASR Metrics
Abdul Waheed, Hanin Atwany, Rita Singh and Bhiksha Raj 23119

Are the Values of LLMs Structurally Aligned with Humans? A Causal Perspective
Yipeng Kang, Junqi Wang, Yexin Li, Mengmeng Wang, Wenming Tu, Quansen Wang, Hengli Li, Tingjun Wu, Xue Feng, Fangwei Zhong and Zilong Zheng 23147

LLMs Can Also Do Well! Breaking Barriers in Semantic Role Labeling via Large Language Models
Xinxin Li, Huiyao Chen, Chengjun Liu, Jing Li, Meishan Zhang, Jun Yu and Min Zhang . 23162

Lost in Transcription, Found in Distribution Shift: Demystifying Hallucination in Speech Foundation Models
Hanin Atwany, Abdul Waheed, Rita Singh, Monojit Choudhury and Bhiksha Raj 23181

M2PA: A Multi-Memory Planning Agent for Open Worlds Inspired by Cognitive Theory
YanfangZhou YanfangZhou, Xiaodong Li, Yuntao Liu, Yongqiang Zhao, Xintong Wang, Zhenyu Li, Jinlong Tian and Xinhai Xu 23204

AnnaAgent: Dynamic Evolution Agent System with Multi-Session Memory for Realistic Seeker Simulation
Ming Wang, Peidong Wang, Lin Wu, Xiaocui Yang, Daling Wang, Shi Feng, Yuxin Chen, Bixuan Wang and Yifei Zhang 23221

Diversification Catalyzes Language Models' Instruction Generalization To Unseen Semantics
Dylan Zhang, Justin Wang and Francois Charton 23236

DecompileBench: A Comprehensive Benchmark for Evaluating Decompilers in Real-World Scenarios
Zeyu Gao, Yuxin Cui, Hao Wang, Siliang Qin, Yuanda Wang, Zhang Bolun and Chao Zhang 23250

Thinking Before Running! Efficient Code Generation with Thorough Exploration and Optimal Refinement
Xiaoqing Zhang, Yuhan Liu, Flood Sung, Xiuying Chen, Shuo Shang and Rui Yan 23268

Edit Once, Update Everywhere: A Simple Framework for Cross-Lingual Knowledge Synchronization in LLMs
Yuchen Wu, Liang Ding, Li Shen and Dacheng Tao 23282

SafeChain: Safety of Language Models with Long Chain-of-Thought Reasoning Capabilities
Fengqing Jiang, Zhangchen Xu, Yuetai Li, Luyao Niu, Zhen Xiang, Bo Li, Bill Yuchen Lin and Radha Poovendran 23303

ETRQA: A Comprehensive Benchmark for Evaluating Event Temporal Reasoning Abilities of Large Language Models
Sigang Luo, Yinan Liu, Dongying Lin, Yingying Zhai, Bin Wang, Xiaochun Yang and Junpeng Liu 23321

The Law of Knowledge Overshadowing: Towards Understanding, Predicting and Preventing LLM Hallucination
Yuji Zhang, Sha Li, Cheng Qian, Jiateng Liu, Pengfei Yu, Chi Han, Yi R. Fung, Kathleen McKeown, ChengXiang Zhai, Manling Li and Heng Ji 23340

LegoMT2: Selective Asynchronous Sharded Data Parallel Training for Massive Neural Machine Translation
Fei Yuan, Yinquan Lu, Lei Li and Jingjing Xu 23359

Pruning General Large Language Models into Customized Expert Models
Yiran Zhao, Guizhen Chen, Kenji Kawaguchi, Lidong Bing and Wenxuan Zhang 23377

Enhance Multimodal Consistency and Coherence for Text-Image Plan Generation
Xiaoxin Lu, Ranran Haoran Zhang, Yusen Zhang and Rui Zhang 23392

Un-considering Contextual Information: Assessing LLMs' Understanding of Indexical Elements
Metehan Oğuz, Yavuz Faruk Bakman and Duygu Nur Yaldiz 23410

Behavioral Analysis of Information Salience in Large Language Models
Jan Trienes, Jörg Schlötterer, Junyi Jessy Li and Christin Seifert 23428

The Behavior Gap: Evaluating Zero-shot LLM Agents in Complex Task-Oriented Dialogs
Avinash Baidya, Kamalika Das and Xiang Gao 23455

Task Facet Learning: A Structured Approach To Prompt Optimization
Gurusha Juneja, Gautam Jajoo, Hua Li, Jian Jiao, Nagarajan Natarajan and Amit Sharma . 23473

LLM as Effective Streaming Processor: Bridging Streaming-Batch Mismatches with Group Position Encoding
Junlong Tong, Jinlan Fu, Zixuan Lin, Yingqi Fan, Anhao Zhao, Hui Su and Xiaoyu Shen . 23497

YinYang-Align: A new Benchmark for Competing Objectives and Introducing Multi-Objective Preference based Text-to-Image Alignment
Amitava Das, Yaswanth Narsupalli, Gurpreet Singh, Vinija Jain, Vasu Sharma, Suranjana Trivedy, Aman Chadha and Amit Sheth 23518

FREE: Fast and Robust Vision Language Models with Early Exits
Divya Jyoti Bajpai and Manjesh Kumar Hanawal 23599

REPRO-Bench: Can Agentic AI Systems Assess the Reproducibility of Social Science Research?
Chuxuan Hu, Liyun Zhang, Yeji Lim, Aum Wadhwani, Austin Peters and Daniel Kang . . . 23616

Time Travel: A Comprehensive Benchmark to Evaluate LMMs on Historical and Cultural Artifacts
Sara Ghaboura, Ketan Pravin More, Ritesh Thawkar, Wafa Al Ghallabi, Omkar Thawakar, Fahad Shahbaz Khan, Hisham Cholakkal, Salman Khan and Rao Muhammad Anwer 23627

Unveiling and Addressing Pseudo Forgetting in Large Language Models
Huashan Sun, Yizhe Yang, Yinghao Li, Jiawei Li and Yang Gao 23642

Improving MLLM's Document Image Machine Translation via Synchronously Self-reviewing Its OCR Proficiency
Yupu Liang, Yaping Zhang, Zhiyang Zhang, Zhiyuan Chen, Yang Zhao, Lu Xiang, Chengqing Zong and Yu Zhou 23659

HG-InsightLog: Context Prioritization and Reduction for Question Answering with Non-Natural Language Construct Log Data
Supriya Bajpai, Athira Gopal, Chandrakant Harjpal and Niraj Kumar 23679

Dialect Normalization using Large Language Models and Morphological Rules
Antonios Dimakis, John Pavlopoulos and Antonios Anastasopoulos 23696

USDC: A Dataset of User Stance and Dogmatism in Long Conversations
Mounika Marreddy, Subba Reddy Oota, Venkata Charan Chinni, Manish Gupta and Lucie Flek 23715

Learning to Insert [PAUSE] Tokens for Better Reasoning
Eunki Kim, Sangryul Kim and James Thorne 23760

Understand the Implication: Learning to Think for Pragmatic Understanding
Settaluri Lakshmi Sravanthi, Kishan Maharaj, Sravani Gunnu, Abhijit Mishra and Pushpak Bhattacharyya 23778

WASA: WAtermark-based Source Attribution for Large Language Model-Generated Data
Xinyang Lu, Jingtan Wang, Zitong Zhao, Zhongxiang Dai, Chuan-Sheng Foo, See-Kiong Ng and Bryan Kian Hsiang Low 23791

Dense Retrieval with Quantity Comparison Intent
Prayas Agrawal, Nandeesh Kumar K M, Muthusamy Chelliah, Surender Kumar and Soumen Chakrabarti 23825

Reflection on Knowledge Graph for Large Language Models Reasoning
Yigeng Zhou, Wu Li, Yifan Lu, Jing Li, Fangming Liu, Meishan Zhang, Yequan Wang, Daojing He, Honghai Liu and Min Zhang 23840

Revisiting 3D LLM Benchmarks: Are We Really Testing 3D Capabilities?
Jiahe Jin, Yanheng He and Mingyan Yang 23858

DIESEL: A Lightweight Inference-Time Safety Enhancement for Language Models
Ben Ganon, Alon Zolfi, Omer Hofman, Inderjeet Singh, Hisashi Kojima, Yuval Elovici and Asaf Shabtai 23870

Toward Structured Knowledge Reasoning: Contrastive Retrieval-Augmented Generation on Experience
Jiawei Gu, Ziting Xian, Yuanzhen Xie, Ye Liu, Enjie Liu, Ruichao Zhong, Mochi Gao, Yunzhi Tan, Bo Hu and Zang Li 23891

Structured Pruning for Diverse Best-of-N Reasoning Optimization
Hieu Trung Nguyen, Bao Nguyen and Viet Anh Nguyen 23911

PodAgent: A Comprehensive Framework for Podcast Generation
Yujia Xiao, Lei He, Haohan Guo, Feng-Long Xie and Tan Lee 23923

STORM-BORN: A Challenging Mathematical Derivations Dataset Curated via a Human-in-the-Loop Multi-Agent Framework
Wenhao Liu, Zhenyi Lu, Xinyu Hu, Jerry Zhang, Dailin Li, Jiacheng Cen, Huilin Cao, Haiteng Wang, Yuhan Li, Xie Kun, Dandan Li, Pei Zhang, Chengbo Zhang, Yuxiang Ren, Xiaohong Huang and Yan Ma 23938

iMOVE : Instance-Motion-Aware Video Understanding
Jiaze Li, Yaya Shi, Zongyang Ma, Haoran Xu, Yandong.bai Yandong.bai, Huihui Xiao, Ruiwen Kang, Fan Yang, Tingting Gao and Di Zhang 23959

SceneGram: Conceptualizing and Describing Tangrams in Scene Context
Simeon Junker and Sina Zarrieß 23976

Relevant or Random: Can LLMs Truly Perform Analogical Reasoning?
Chengwei Qin, Wenhan Xia, Tan Wang, Fangkai Jiao, Yuchen Hu, Bosheng Ding, Ruirui Chen and Shafiq Joty 23993

MERIT: Multi-Agent Collaboration for Unsupervised Time Series Representation Learning
Shu Zhou, Yunyang Xuan, Yuxuan Ao, Xin Wang, Tao Fan and Hao Wang 24011

JsonTuning: Towards Generalizable, Robust, and Controllable Instruction Tuning
Chang Gao, Wenxuan Zhang, Guizhen Chen and Wai Lam 24029

RedundancyLens: Revealing and Exploiting Visual Token Processing Redundancy for Efficient Decoder-Only MLLMs
Hongliang Li, Jiaxin Zhang, Wenhui Liao, Dezhi Peng, Kai Ding and Lianwen Jin 24056

Memory-augmented Query Reconstruction for LLM-based Knowledge Graph Reasoning
Mufan Xu, Gewen Liang, Kehai Chen, Wei Wang, Xun Zhou, Muyun Yang, Tiejun Zhao and Min Zhang 24068

KaFT: Knowledge-aware Fine-tuning for Boosting LLMs' Domain-specific Question-Answering Performance
Qihuang Zhong, Liang Ding, Xiantao Cai, Juhua Liu, Bo Du and Dacheng Tao 24085

Are Multimodal Large Language Models Pragmatically Competent Listeners in Simple Reference Resolution Tasks?
Simeon Junker, Manar Ali, Larissa Koch, Sina Zarrieß and Hendrik Buschmeier 24101

Removing Prompt-template Bias in Reinforcement Learning from Human Feedback
Chaojie Wang, Haonan Shi, Long Tian, Bo An and Shuicheng Yan 24110

Latent Distribution Decouple for Uncertain-Aware Multimodal Multi-label Emotion Recognition
Jingwang Huang, Jiang Zhong, Qin Lei, Gaojinpeng Gaojinpeng, Ymyang Ymyang, Sirui Wang, PeiguangLi PeiguangLi and Kaiwen Wei 24123

Are LLMs Rational Investors? A Study on the Financial Bias in LLMs
Yuhang Zhou, Yuchen Ni, Zhiheng Xi, Zhangyue Yin, Yu He, Gan Yunhui, Xiang Liu, Zhang Jian, Sen Liu, Xipeng Qiu, Yixin Cao, Guangnan Ye and Hongfeng Chai 24139

Seeing What Tastes Good: Revisiting Multimodal Distributional Semantics in the Billion Parameter Era
Dan Oneata, Desmond Elliott and Stella Frank 24174

Communication-Efficient and Tensorized Federated Fine-Tuning of Large Language Models
Sajjad Ghiasvand, Yifan Yang, Zhiyu Xue, Mahnoosh Alizadeh, Zheng Zhang and Ramtin Pedarsani 24192

A rebuttal of two common deflationary stances against LLM cognition
Zak Hussain, Rui Mata and Dirk U. Wulff 24208

COVER: Context-Driven Over-Refusal Verification in LLMs
Giovanni Sullutrone, Riccardo A. Vigliermo, Sonia Bergamaschi and Luca Sala 24214

MOSAIC: Multiple Observers Spotting AI Content
Matthieu Dubois, François Yvon and Pablo Piantanida 24230

GUIDEX: Guided Synthetic Data Generation for Zero-Shot Information Extraction
Neil De La Fuente, Oscar Sainz, Iker García-Ferrero and Eneko Agirre 24248

Missing the Margins: A Systematic Literature Review on the Demographic Representativeness of LLMs
Indira Sen, Marlene Lutz, Elisa Rogers, David Garcia and Markus Strohmaier 24263

LlamaV-o1: Rethinking Step-by-step Visual Reasoning in LLMs
Omkar Thawakar, Dinura Dissanayake, Ketan Pravin More, Ritesh Thawkar, Ahmed Heakl, Noor Ahsan, Yuhao Li, Ilmuz Zaman Mohammed Zumri, Jean Lahoud, Rao Muhammad Anwer, Hisham Cholakkal, Ivan Laptev, Mubarak Shah, Fahad Shahbaz Khan and Salman Khan 24290

Burn After Reading: Do Multimodal Large Language Models Truly Capture Order of Events in Image Sequences?
Yingjin Song, Yupei Du, Denis Paperno and Albert Gatt 24316

Full-Step-DPO: Self-Supervised Preference Optimization with Step-wise Rewards for Mathematical Reasoning
Huimin Xu, Xin Mao, Feng-Lin Li, Xiaobao Wu, Wang Chen, Wei Zhang and Anh Tuan Luu 24343

Do Emotions Really Affect Argument Convincingness? A Dynamic Approach with LLM-based Manipulation Checks
Yanran Chen and Steffen Eger 24357

SCOPE: Compress Mathematical Reasoning Steps for Efficient Automated Process Annotation
Huimin Xu, Xin Mao, Feng-Lin Li, Xiaobao Wu, Wang Chen, Wei Zhang and Anh Tuan Luu 24382

Compositional Syntactico-SemBanking for English as a Second or Foreign Language
Wenxi Li, Xihao Wang and Weiwei Sun 24395

Semantics-aware prompting for translating NOtices To AirMen
Minal Nitin Dani, Aishwarya Maheswaran and Maunendra Sankar Desarkar 24407

Stereotype or Personalization? User Identity Biases Chatbot Recommendations
Anjali Kantharuban, Jeremiah Milbauer, Maarten Sap, Emma Strubell and Graham Neubig24418

Automated main concept generation for narrative discourse assessment in aphasia
Ankita Gupta, Marisa Hudspeth, Polly Stokes, Jacquie Kurland and Brendan O'Connor . . . 24437

Can VLMs Actually See and Read? A Survey on Modality Collapse in Vision-Language Models
Mong Yuan Sim, Wei Emma Zhang, Xiang Dai and Biaoyan Fang 24452

"You are Beautiful, Body Image Stereotypes are Ugly!" BIStereo: A Benchmark to Measure Body Image Stereotypes in Language Models
Narjis Asad, Nihar Ranjan Sahoo, Rudra Murthy, Swaprava Nath and Pushpak Bhattacharyya 24471

Retrieval Models Aren't Tool-Savvy: Benchmarking Tool Retrieval for Large Language Models
Zhengliang Shi, Yuhan Wang, Lingyong Yan, Pengjie Ren, Shuaiqiang Wang, Dawei Yin and Zhaochun Ren 24497

FineCite: A Novel Approach For Fine-Grained Citation Context Analysis
Lasse M. Jantsch, Dong-Jae Koh, Seonghwan Yoon, Jisu Lee, Anne Lauscher and Young-Kyoon Suh 24525

Decoupling Reasoning and Knowledge Injection for In-Context Knowledge Editing
Changyue Wang, Weihang Su, Qingyao Ai, Yujia Zhou and Yiqun Liu 24543

Entrospect: Information-Theoretic Self-Reflection Elicits Better Response Refinement of Small Language Models
Tianqiang Yan, Ziqiao Lin, Lin Zhang, Zhenglong Sun and Yuan Gao 24563

Iterative Repair with Weak Verifiers for Few-shot Transfer in KBQA with Unanswerability
Riya Sawhney, Samrat Yadav, Indrajit Bhattacharya and Mausam Mausam 24578

Safeguarding RAG Pipelines with GMTP: A Gradient-based Masked Token Probability Method for Poisoned Document Detection
San Kim, Jonghwi Kim, Yejin Jeon and Gary Lee 24597

EnSToM: Enhancing Dialogue Systems with Entropy-Scaled Steering Vectors for Topic Maintenance
Heejae Suh, Yejin Jeon, Deokhyung Kang, Taehee Park, Yejin Min and Gary Lee 24615

MultiTEND: A Multilingual Benchmark for Natural Language to NoSQL Query Translation
Zhiqian Qin, Yuanfeng Song, Jinwei Lu, Yuanwei Song, Shuaimin Li and Chen Jason Zhang 24632

Tool learning via Inference-time Scaling and Cycle Verifier
Xiaobo Liang, Wenjin Xie, Juntao Li, Wanfu Wang, Yibin Chen, Kehai Chen and Min Zhang 24658

When Benchmarks Talk: Re-Evaluating Code LLMs with Interactive Feedback
Jane Pan, Ryan Shar, Jacob Pfau, Ameet Talwalkar, He He and Valerie Chen 24672

Reranking-based Generation for Unbiased Perspective Summarization
Narutatsu Ri, Nicholas Deas and Kathleen McKeown 24701

KARPA: A Training-free Method of Adapting Knowledge Graph as References for Large Language Model's Reasoning Path Aggregation
Siyuan Fang, Kaijing Ma, Tianyu Zheng, Xeron Du, Ningxuan Lu, Ge Zhang and Qingkun Tang 24724

Enhancing LLM-based Hatred and Toxicity Detection with Meta-Toxic Knowledge Graph
Yibo Zhao, Jiapeng Zhu, Can Xu, Yao Liu and Xiang Li 24747

Mixture-of-Personas Language Models for Population Simulation
Ngoc Bui, Hieu Trung Nguyen, Shantanu Kumar, Julian Theodore, Weikang Qiu, Viet Anh Nguyen and Rex Ying 24761

ClusComp: A Simple Paradigm for Model Compression and Efficient Finetuning
Baohao Liao, Christian Herold, Seyyed Hadi Hashemi, Stefan Vasilev, Shahram Khadivi and Christof Monz 24779

Decomposed Opinion Summarization with Verified Aspect-Aware Modules
Miao Li, Jey Han Lau, Eduard Hovy and Mirella Lapata 24805

Token-Budget-Aware LLM Reasoning
Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma and Zhenyu Chen . 24842

HATA: Trainable and Hardware-Efficient Hash-Aware Top-k Attention for Scalable Large Model Inference
Ping Gong, Jiawei Yi, Shengnan Wang, Juncheng Zhang, Zewen Jin, Ouxiang Zhou, Ruibo Liu, Guanbin Xu, Youhui Bai, Bowen Ye, Kun Yuan, Tong Yang, Gong Zhang, Renhai Chen, Feng Wu and Cheng Li 24856

Answer When Needed, Forget When Not: Language Models Pretend to Forget via In-Context Knowledge Unlearning
Shota Takashiro, Takeshi Kojima, Andrew Gambardella, Qi Cao, Yusuke Iwasawa and Yutaka Matsuo 24872

LIST: Linearly Incremental SQL Translator for Single-Hop Reasoning, Generation and Verification
Kaiyuan Guan, Ruoxin Li, Xudong Guo, Zhenning Huang, Xudong Weng, Hehuan Liu, Zheng Wei and Zang Li 24886

MAGI: Multi-Agent Guided Interview for Psychiatric Assessment
Guanqun Bi, Zhuang Chen, Zhoufu Liu, Hongkai Wang, Xiyao Xiao, Yuqiang Xie, Wen Zhang, Yongkang Huang, Yuxuan Chen, Libiao Peng and Minlie Huang 24898

TituLLMs: A Family of Bangla LLMs with Comprehensive Benchmarking
Shahriar Kabir Nahin, Rabindra Nath Nandi, Sagor Sarker, Quazi Sarwar Muhtaseem, Md Kowsher, Apu Chandraw Shill, Md Ibrahim, Mehadi Hasan Menon, Tareq Al Muntasir and Firoj Alam 24922

WikiMixQA: A Multimodal Benchmark for Question Answering over Tables and Charts
Negar Foroutan, Angelika Romanou, Matin Ansaripour, Julian Martin Eisenschlos, Karl Aberer and Rémi Lebret 24941

Let's Fuse Step by Step: A Generative Fusion Decoding Algorithm with LLMs for Robust and Instruction-Aware ASR and OCR
Chan-Jan Hsu, Yi-Chang Chen, Feng-Ting Liao, Pei-Chen Ho, Yu-Hsiang Wang, Po-Chun Hsu and Da-shan Shiu 24959

HPSS: Heuristic Prompting Strategy Search for LLM Evaluators
Bosi Wen, Pei Ke, Yufei Sun, Cunxiang Wang, Xiaotao Gu, Jinfeng Zhou, Jie Tang, Hongning Wang and Minlie Huang 24974

A Fully Generative Motivational Interviewing Counsellor Chatbot for Moving Smokers Towards the Decision to Quit
Zafarullah Mahmood, Soliman Ali, Jiading Zhu, Mohamed Abdelwahab, Michelle Yu Collins, Sihan Chen, Yi Cheng Zhao, Jodi Wolff, Osnat C. Melamed, Nadia Minian, Marta Maslej, Carolynne Cooper, Matt Ratto, Peter Selby and Jonathan Rose 25008

LegalCore: A Dataset for Event Coreference Resolution in Legal Documents
Kangda Wei, Xi Shi, Jonathan Tong, Sai Ramana Reddy , Anandhavelu Natarajan, Rajiv Jain, Aparna Garimella and Ruihong Huang 25044

Rectifying Belief Space via Unlearning to Harness LLMs' Reasoning
Ayana Niwa, Masahiro Kaneko and Kentaro Inui 25060

MemeDetoxNet: Balancing Toxicity Reduction and Context Preservation
Gitanjali Kumari, Jitendra Solanki and Asif Ekbal 25076

Should I Trust You? Detecting Deception in Negotiations using Counterfactual RL
Wichayaporn Wongkamjan, Yanze Wang, Feng Gu, Denis Peskoff, Jonathan K. Kummerfeld, Jonathan May and Jordan Lee Boyd-Graber 25099

Multi-matrix Factorization Attention
Jingcheng Hu, Houyi Li, Yinmin Zhang, Zili Wang, Shuigeng Zhou, Xiangyu Zhang and Heung-Yeung Shum 25114

Self-Training Elicits Concise Reasoning in Large Language Models
Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin Yang, Yujin Kim and Se-Young Yun 25127

Reason from Future: Reverse Thought Chain Enhances LLM Reasoning
Yinlong Xu, Yanzhao Zheng, Shuoshuo Sun, Shuaihan Huang, Baohua Dong, Zhu Hangcheng, Ruohui Huang, Gang Yu, Hongxia Xu and Jian Wu 25153

LLMs as Planning Formalizers: A Survey for Leveraging Large Language Models to Construct Automated Planning Models
Marcus Tantakoun, Christian Muise and Xiaodan Zhu 25167

From Conversation to Automation: Leveraging LLMs for Problem-Solving Therapy Analysis
Elham Aghakhani, Lu Wang, Karla T. Washington, George Demiris, Jina Huh-Yoo and Rezvaneh Rezapour 25189

Revisiting Self-Consistency from Dynamic Distributional Alignment Perspective on Answer Aggregation
Yiwei Li, Ji Zhang, Shaoxiong Feng, Peiwen Yuan, Xinglin Wang, Jiayi Shi, Yueqi Zhang, Chuyi Tan, Boyuan Pan, Yao Hu and Kan Li 25208

Don't Say No: Jailbreaking LLM by Suppressing Refusal
Yukai Zhou, Jian Lou, Zhijie Huang, Zhan Qin, Sibei Yang and Wenjie Wang 25224

From Perception to Reasoning: Enhancing Vision-Language Models for Mobile UI Understanding
Settaluri Lakshmi Sravanthi, Ankit Mishra, Debjyoti Mondal, Subhadarshi Panda, Rituraj Singh and Pushpak Bhattacharyya 25250

Lemmas Matter, But Not Like That: Predictors of Lemma-Based Generalization in Morphological Inflection
Sarah Ruth Brogden Payne and Jordan Kodner 25270

Mosaic-IT: Cost-Free Compositional Data Synthesis for Instruction Tuning
Ming Li, Pei Chen, Chenguang Wang, Hongyu Zhao, Yijun Liang, YuPeng Hou, Fuxiao Liu and Tianyi Zhou 25287

MAM: Modular Multi-Agent Framework for Multi-Modal Medical Diagnosis via Role-Specialized Collaboration
Yucheng Zhou, Lingran Song and Jianbing Shen 25319

ATLAS: Agent Tuning via Learning Critical Steps
Zhixun Chen, Ming Li, Yuxuan Huang, Yali Du, Meng Fang and Tianyi Zhou 25334

Syntactic Control of Language Models by Posterior Inference
Vicky Xefteri, Tim Vieira, Ryan Cotterell and Afra Amini 25350

Small Models Struggle to Learn from Strong Reasoners
Yuetai Li, Xiang Yue, Zhangchen Xu, Fengqing Jiang, Luyao Niu, Bill Yuchen Lin, Bhaskar Ramasubramanian and Radha Poovendran 25366

Sparse Rewards Can Self-Train Dialogue Agents
Barrett Martin Lattimer, Varun Prashant Gangal, Ryan McDonald and Yi Yang 25395

Almost AI, Almost Human: The Challenge of Detecting AI-Polished Writing
Shoumik Saha and Soheil Feizi 25414

The Reader is the Metric: How Textual Features and Reader Profiles Explain Conflicting Evaluations of AI Creative Writing
Guillermo Marco, Julio Gonzalo and Víctor Fresno 25432

Summary Factual Inconsistency Detection Based on LLMs Enhanced by Universal Information Extraction
Anguo Li and Lei Yu 25450

ELI-Why: Evaluating the Pedagogical Utility of Language Model Explanations
Brihi Joshi, Keyu He, Sahana Ramnath, Sadra Sabouri, Kaitlyn Zhou, Souti Chattopadhyay, Swabha Swayamdipta and Xiang Ren 25466

Beyond Generation: Leveraging LLM Creativity to Overcome Label Bias in Classification
Xiaoyue Wang and Xin Liu 25500

CogSteer: Cognition-Inspired Selective Layer Intervention for Efficiently Steering Large Language Models
Xintong Wang, Jingheng Pan, Liang Ding, Longyue Wang, Longqin Jiang, Xingshan Li and Chris Biemann 25507

PASTEL : Polarity-Aware Sentiment Triplet Extraction with LLM-as-a-Judge
Aaditya Bodke, Avinoor Singh Kohli, Hemant Subhash Pardeshi and Prathamesh Bhosale 25523

COSMIC: Generalized Refusal Direction Identification in LLM Activations
Vincent Siu, Nicholas Crispino, Zihao Yu, Sam Pan, Zhun Wang, Yang Liu, Dawn Song and Chenguang Wang 25534

Red Queen: Exposing Latent Multi-Turn Risks in Large Language Models
Yifan Jiang, Kriti Aggarwal, Tanmay Laud, Kashif Munir, Jay Pujara and Subhabrata Mukherjee 25554

MDBench: A Synthetic Multi-Document Reasoning Benchmark Generated with Knowledge Guidance
Joseph J Peper, Wenzhao Qiu, Ali Payani and Lu Wang 25592

DiaLLMs: EHR-Enhanced Clinical Conversational System for Clinical Test Recommendation and Diagnosis Prediction
Weijieying Ren, Tianxiang Zhao, Lei Wang, Tianchun Wang and Vasant G Honavar 25622

Can Hallucination Correction Improve Video-Language Alignment?
Lingjun Zhao, Mingyang Xie, Paola Cascante-Bonilla, Hal Daumé Iii and Kwonjoon Lee . 25636

IMPARA-GED: Grammatical Error Detection is Boosting Reference-free Grammatical Error Quality Estimator
Yusuke Sakai, Takumi Goto and Taro Watanabe 25647

Do Language Models Mirror Human Confidence? Exploring Psychological Insights to Address Overconfidence in LLMs
Chenjun Xu, Bingbing Wen, Bin Han, Robert Wolfe, Lucy Lu Wang and Bill Howe 25655

Why Multi-Interest Fairness Matters: Hypergraph Contrastive Multi-Interest Learning for Fair Conversational Recommender System
Yongsen Zheng, Zongxuan Xie, Guohua Wang, Ziyao Liu, Liang Lin and Kwok-Yan Lam 25673

Cautious Next Token Prediction
Yizhou Wang, Lingzhi Zhang, Yue Bai, Mang Tik Chiu, Zhengmian Hu, Mingyuan Zhang, Qihua Dong, Yu Yin, Sohrab Amirghodsi and Yun Fu 25685

Reasoning with Graphs: Structuring Implicit Knowledge to Enhance LLMs Reasoning
Haoyu Han, Yaochen Xie, Hui Liu, Xianfeng Tang, Sreyashi Nag, William Headden, Yang Li, Chen Luo, Shuiwang Ji, Qi He and Jiliang Tang 25698

Enhancing Medical Dialogue Generation through Knowledge Refinement and Dynamic Prompt Adjustment
Hongda Sun, Jiaren Peng, Wenzhong Yang, Liang He, Bo Du and Rui Yan 25715

Feature-Level Insights into Artificial Text Detection with Sparse Autoencoders
Kristian Kuznetsov, Laida Kushnareva, Anton Razzhigaev, Polina Druzhinina, Anastasia Voznyuk, Irina Piontkovskaya, Evgeny Burnaev and Serguei Barannikov 25727

Low-Resource Grammatical Error Correction: Selective Data Augmentation with Round-Trip Machine Translation
Frank Palma Gomez and Alla Rozovskaya 25749

Just Put a Human in the Loop? Investigating LLM-Assisted Annotation for Subjective Tasks
Hope Schroeder, Deb Roy and Jad Kabbara 25771

Research Community Perspectives on Intelligenceand Large Language Models
Bertram Højer, Terne Sasha Thorn Jakobsen, Anna Rogers and Stefan Heinrich 25796

LEMONADE: A Large Multilingual Expert-Annotated Abstractive Event Dataset for the Real World
Sina Semnani, Pingyue Zhang, Wanyue Zhai, Haozhuo Li, Ryan Beauchamp, Trey Billing, Katayoun Kishi, Manling Li and Monica Lam 25813

Memorization vs. Reasoning: Updating LLMs with New Knowledge
Aochong Oliver Li and Tanya Goyal 25853

CourtEval: A Courtroom-Based Multi-Agent Evaluation Framework
Sandeep Kumar, Abhijit A Nargund and Vivek Sridhar 25875

Multilingual Definition Modeling
Edison Marrese-Taylor, Erica K. Shimomoto, A. Solano and Enrique Reid 25888

Human Bias in the Face of AI: Examining Human Judgment Against Text Labeled as AI Generated
Tiffany Zhu, Iain Weissburg, Kexun Zhang and William Yang Wang 25907

Redundancy, Isotropy, and Intrinsic Dimensionality of Prompt-based Text Embeddings
Hayato Tsukagoshi and Ryohei Sasano 25915

Harnessing Whisper for Prosodic Stress Analysis
Samuel S. Sohn, Sten Knutsen and Karin Stromswold 25931

Can You Share Your Story? Modeling Clients' Metacognition and Openness for LLM Therapist Evaluation
Minju Kim, Dongje Yoo, Yeonjun Hwang, Minseok Kang, Namyoung Kim, Minju Gwak, Beongwoo Kwak, Hyungjoo Chae, Harim Kim, Yunjoong Lee, Min Hee Kim, Dayi Jung, Kyong-Mee Chung and Jinyoung Yeo 25943

Dictionaries to the Rescue: Cross-Lingual Vocabulary Transfer for Low-Resource Languages Using Bilingual Dictionaries
Haruki Sakajo, Yusuke Ide, Justin Vasselli, Yusuke Sakai, Yingtao Tian, Hidetaka Kamigaito and Taro Watanabe 25963

When Should Dense Retrievers Be Updated in Evolving Corpora? Detecting Out-of-Distribution Corpora Using GradNormIR
Dayoon Ko, Jinyoung Kim, Sohyeon Kim, Jinhyuk Kim, Jaehoon Lee, Seonghak Song, Minyoung Lee and Gunhee Kim 25977

The Million Authors Corpus: A Cross-Lingual and Cross-Domain Wikipedia Dataset for Authorship Verification
Abraham Israeli, Shuai Liu, Jonathan May and David Jurgens 25997

BridG MT: Enhancing LLMs' Machine Translation Capabilities with Sentence Bridging and Gradual MT
Seungwoo Choi, Gahyun Yoo and Jay-Yoon Lee 26018

Text2World: Benchmarking Large Language Models for Symbolic World Model Generation
Mengkang Hu, Tianxing Chen, Yude Zou, Yuheng Lei, Qiguang Chen, Ming Li, Yao Mu, Hongyuan Zhang, Wenqi Shao and Ping Luo 26043

Blinded by Context: Unveiling the Halo Effect of MLLM in AI Hiring
Kyusik Kim, Jeongwoo Ryu, Hyeonseok Jeon and Bongwon Suh 26067

CoT-UQ: Improving Response-wise Uncertainty Quantification in LLMs with Chain-of-Thought
Boxuan Zhang and Ruqi Zhang 26114

ADO: Automatic Data Optimization for Inputs in LLM Prompts
Sam Lin, Wenyue Hua, Lingyao Li, Zhenting Wang and Yongfeng Zhang 26134

Large Language Models Still Exhibit Bias in Long Text
Wonje Jeung, Dongjae Jeon, Ashkan Yousefpour and Jonghyun Choi 26147

Do Vision-Language Models Have Internal World Models? Towards an Atomic Evaluation
Qiyue Gao, Xinyu Pi, Kevin Liu, Junrong Chen, Ruolan Yang, Xinqi Huang, Xinyu Fang, Lu Sun, Gautham Kishore, Bo Ai, Stone Tao, Mengyang Liu, Jiaxi Yang, Chao-Jung Lai, Chuanyang Jin, Jiannan Xiang, Benhao Huang, Zeming Chen, David Danks, Hao Su, Tianmin Shu, Ziqiao Ma, Lianhui Qin and Zhiting Hu 26170

Protecting Users From Themselves: Safeguarding Contextual Privacy in Interactions with Conversational Agents
Ivoline C. Ngong, Swanand Ravindra Kadhe, Hao Wang, Keerthiram Murugesan, Justin D. Weisz, Amit Dhurandhar and Karthikeyan Natesan Ramamurthy 26196

Enhancing Persona Consistency for LLMs' Role-Playing using Persona-Aware Contrastive Learning
Ke Ji, Yixin Lian, Linxu Li, Jingsheng Gao, Weiyuan Li and Bin Dai 26221

M^2-TabFact: Multi-Document Multi-Modal Fact Verification with Visual and Textual Representations of Tabular Data
Mingyang Zhou, Lingyu Zhang, Sophia Horng, Maximillian Chen, Kung-Hsiang Huang and Shih-Fu Chang 26239

Fuzzy Speculative Decoding for a Tunable Accuracy-Runtime Tradeoff
Maximilian Holsman, Yukun Huang and Bhuwan Dhingra 26257

PLAY2PROMPT: Zero-shot Tool Instruction Optimization for LLM Agents via Tool Play
Wei Fang, Yang Zhang, Kaizhi Qian, James R. Glass and Yada Zhu 26274

Towards the Pedagogical Steering of Large Language Models for Tutoring: A Case Study with Modeling Productive Failure
Romain Puech, Jakub Macina, Julia Chatain, Mrinmaya Sachan and Manu Kapur 26291

Spotting Out-of-Character Behavior: Atomic-Level Evaluation of Persona Fidelity in Open-Ended Generation
Jisu Shin, Juhyun Oh, Eunsu Kim, Hoyun Song and Alice Oh 26312

What Language Do Non-English-Centric Large Language Models Think in?
Chengzhi Zhong, Qianying Liu, Fei Cheng, Junfeng Jiang, Zhen Wan, Chenhui Chu, Yugo Murawaki and Sadao Kurohashi 26333

T^5Score: A Methodology for Automatically Assessing the Quality of LLM Generated Multi-Document Topic Sets
Itamar Trainin and Omri Abend 26347

Uncertainty-Aware Contrastive Decoding
Hakyung Lee, Subeen Park, Joowang Kim, Sungjun Lim and Kyungwoo Song 26376

GEMS: Generation-Based Event Argument Extraction via Multi-perspective Prompts and Ontology Steering
Run Lin, Yao Liu, Yanglei Gan, Yuxiang Cai, Tian Lan and Qiao Liu 26392

RomanLens: The Role Of Latent Romanization In Multilinguality In LLMs
Alan Saji, Jaavid Aktar Husain, Thanmay Jayakumar, Raj Dabre, Anoop Kunchukuttan and Ratish Puduppully 26410

7 Points to Tsinghua but 10 Points to ? Assessing Large Language Models in Agentic Multilingual National Bias
Qianying Liu, Katrina Qiyao Wang, Fei Cheng and Sadao Kurohashi 26430

Search-in-Context: Efficient Multi-Hop QA over Long Contexts via Monte Carlo Tree Search with Dynamic KV Retrieval
Jiabei Chen, Guang Liu, Shizhu He, Kun Luo, Yao Xu, Jun Zhao and Kang Liu 26443

LLM-as-an-Interviewer: Beyond Static Testing Through Dynamic LLM Evaluation
Eunsu Kim, Juyoung Suk, Seungone Kim, Niklas Muennighoff, Dongkwan Kim and Alice Oh 26456

IntentionESC: An Intention-Centered Framework for Enhancing Emotional Support in Dialogue Systems
Xinjie Zhang, Wenxuan Wang and Qin Jin 26494

Beyond Context to Cognitive Appraisal: Emotion Reasoning as a Theory of Mind Benchmark for Large Language Models
Gerard Christopher Yeo and Kokil Jaidka 26517

CSTRL: Context-Driven Sequential Transfer Learning for Abstractive Radiology Report Summarization
Mst. Fahmida Sultana Naznin, Adnan Ibney Faruq, Mostafa Rifat Tazwar, Md Jobayer, Md. Mehedi Hasan Shawon and Md Rakibul Hasan 26526

Rethinking Prompt-based Debiasing in Large Language Model
Xinyi Yang, Runzhe Zhan, Shu Yang, Junchao Wu, Lidia S. Chao and Derek F. Wong 26538

Exploring In-context Example Generation for Machine Translation
Dohyun Lee, Seungil Chad Lee, Chanwoo Yang, Yujin Baek and Jaegul Choo 26554

Knowledge Base Construction for Knowledge-Augmented Text-to-SQL
Jinheon Baek, Horst Samulowitz, Oktie Hassanzadeh, Dharmashankar Subramanian, Sola Shirai, Alfio Gliozzo and Debarun Bhattacharjya 26569

NBDESCRIB: A Dataset for Text Description Generation from Tables and Code in Jupyter Notebooks with Guidelines
Xuye Liu, Tengfei Ma, Yimu Wang, Fengjie Wang and Jian Zhao 26584

ECoRAG: Evidentiality-guided Compression for Long Context RAG
Yeonseok Jeong, Jinsu Kim, Dohyeon Lee and Seung-won Hwang 26607

From Complexity to Clarity: AI/NLP's Role in Regulatory Compliance
Jivitesh Jain, Nivedhitha Dhanasekaran and Mona T. Diab 26629

EXPERT: An Explainable Image Captioning Evaluation Metric with Structured Explanations
Hyunjong Kim, Sangyeop Kim, Jongheon Jeong, Yeongjae Cho and Sungzoon Cho 26642

Mind Your Theory: Theory of Mind Goes Deeper Than Reasoning
Eitan Wagner, Nitay Alon, Joseph M Barnby and Omri Abend 26658

LLMs are Biased Evaluators But Not Biased for Fact-Centric Retrieval Augmented Generation
Yen-Shan Chen, Jing Jin, Peng-Ting Kuo, Chao-Wei Huang and Yun-Nung Chen 26669

Standard Quality Criteria Derived from Current NLP Evaluations for Guiding Evaluation Design and Grounding Comparability and AI Compliance Assessments
Anya Belz, Simon Mille and Craig Thomson 26685

skLEP: A Slovak General Language Understanding Benchmark
Marek Suppa, Andrej Ridzik, Daniel Hládek, Tomáš Javůrek, Viktória Ondrejová, Kristína Sásiková, Martin Tamajka and Marian Simko 26716

Can Vision Language Models Understand Mimed Actions?
Hyundong Justin Cho, Spencer Lin, Tejas Srinivasan, Michael Saxon, Deuksin Kwon, Natali T. Chavez and Jonathan May 26744

Training Language Model to Critique for Better Refinement
Tianshu Yu, Chao Xiang, Mingchuan Yang, Pei Ke, Bosi Wen, Cunxiang Wang, Jiale Cheng, Li Zhang, Xinyu Mu, Chuxiong Sun and Minlie Huang 26760

Dynamic Task Vector Grouping for Efficient Multi-Task Prompt Tuning
Peiyi Zhang, Richong Zhang, Zhijie Nie and Ziqiao Wang 26805

DICE-BENCH: Evaluating the Tool-Use Capabilities of Large Language Models in Multi-Round, Multi-Party Dialogues
Kyochul Jang, Donghyeon Lee, Kyusik Kim, Dongseok Heo, Taewhoo Lee, Woojeong Kim and Bongwon Suh 26822

HASH-RAG: Bridging Deep Hashing with Retriever for Efficient, Fine Retrieval and Augmented Generation
Jinyu Guo, Xunlei Chen, Qiyang Xia, Zhaokun Wang, Jie Ou, Libo Qin, Shunyu Yao and Wenhong Tian 26847

A Constrained Text Revision Agent via Iterative Planning and Searching
Hannan Cao and Hwee Tou Ng 26859

MMRefine: Unveiling the Obstacles to Robust Refinement in Multimodal Large Language Models
Gio Paik, Geewook Kim and Jinbae Im 26883

How Programming Concepts and Neurons Are Shared in Code Language Models
Amir Hossein Kargaran, Yihong Liu, François Yvon and Hinrich Schuetze 26905

DynaQuest: A Dynamic Question Answering Dataset Reflecting Real-World Knowledge Updates
Qian Lin, Junyi Li and Hwee Tou Ng 26918

ProcrustesGPT: Compressing LLMs with Structured Matrices and Orthogonal Transformations
Ekaterina Grishina, Mikhail Gorbunov and Maxim Rakhuba 26937

Revisiting In-Context Learning with Long Context Language Models
Jinheon Baek, Sun Jae Lee, Prakhar Gupta, Geunseob Oh, Siddharth Dalmia and Prateek Kolhar 26950

Rationalize and Align: Enhancing Writing Assistance with Rationale via Self-Training for Improved Alignment
Hannan Cao, Hai Ye and Hwee Tou Ng 26967

Accelerating Adaptive Retrieval Augmented Generation via Instruction-Driven Representation Reduction of Retrieval Overlaps
Jie Ou, Jinyu Guo, Shuaihong Jiang, Zhaokun Wang, Libo Qin, Shunyu Yao and Wenhong Tian 26983

MEXA: Multilingual Evaluation of English-Centric LLMs via Cross-Lingual Alignment
Amir Hossein Kargaran, Ali Modarressi, Nafiseh Nikeghbal, Jana Diesner, François Yvon and Hinrich Schuetze 27001

Automated Fine-Grained Mixture-of-Experts Quantization
Zhanhao Xie, Yuexiao Ma, Xiawu Zheng, Fei Chao, Wanchen Sui, Yong Li, Shen Li and Rongrong Ji 27024

Enhancing Complex Reasoning in Knowledge Graph Question Answering through Query Graph Approximation
Hongjun Jeong, Minji Kim, Heesoo Jung, Ko Keun Kim and Hogun Park 27038