

ACL 2024

**62nd Annual Meeting of the Association for Computational
Linguistics (ACL 2024)**

Proceedings of the Conference

August 11-16, 2024

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

Diamond



Platinum



Megagon Labs



Meituan



Gold



Silver



Bronze



Diversity Inclusion Ally



Local Sponsors





©2024 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 979-8-89176-094-3

Message from the General Chair

Originally named the Association for Machine Translation and Computational Linguistics (AMTCL), the Association for Computational Linguistics was founded in 1962 and renamed the ACL in 1968.

The ACL is run by some 20 volunteers overseeing the administration of the Association (organising elections, deciding on new actions, adapting to the fast changing trends of our fields), the publication of two journals (Computational Linguistics and the Transaction of the ACL) and the organisation of the ACL, EACL, NAACL and AACL conferences. The ACL executive board is regularly renewed based on elections from the membership. To ensure continuity, some of the volunteers (e.g., secretary, treasurer and anthology editor) serve longer term mandates and a professional Business Manager gives much needed support to the daily management of the Association.

While ACL runs multiple conferences and workshops each year, the ACL conference is the flagship conference of the Association and like the field in general, it has seen drastic changes over the years both in terms of approaches (from symbolic to statistical to neural) and in terms of popularity (from a few academics to a large population of several thousands academics and industrials with widespread geographical coverage).

It is a testimony to the strength of the Association that the ACL meeting has been held annually since 1962. Also remarkable is that, thanks to Steven Bird's initiative, the proceedings are available online in the ACL anthology since ACL 1979 (which hosted a whopping 23 papers!).

In this note, I would like to welcome all participants to ACL 2024, the 62nd Annual Meeting of the Association for Computational Linguistics (held in Bangkok, Thailand, August 11-16, 2024), and to express my gratitude to another large set of volunteers who made ACL 2024 possible.

First and foremost, I would like to thank the three Program Chairs, Lun-Wei Ku, Andre Martins and Vivek Srikumar who oversaw the reviewing process, selected the keynote speakers and created the conference program. 2024 was the first year where all *CL conferences moved entirely to the ACL Rolling Review model. This meant a novel process working in tight interaction with the ARR team while handling a huge number of submissions. Thanks to Lun-Wei, Andre and Vivek, to the ARR Editors in Chief (Mausam, Viviane Moreira, Vincent Ng, Lilja Øvrelid, Tamar Solorio, Jun Suzuki), to the Senior Area Chairs, Area Chairs, reviewers, and to the Best Paper Committee, who worked together to select the ACL 2024 program.

The workshop and the tutorial program was created by the Workshop (Eunsol Choi, Xipeng Qiu) and the Tutorial (Luis Chiruzzo, Hung-yi Lee, Leonardo Ribeiro) Chairs, who collaborated with EACL, NAACL and EMNLP 2024 to select 32 workshops and 6 tutorials that cover both technical and societal areas of Natural Language Processing.

The program also includes demonstrations selected by the Demonstration Chairs (Yixin Cao, Yang Feng, Deyi Xiong), as well as the traditional Student Research Workshop, which was put together by the SRW Chairs (Eve Fleisig, Xiyan Fu), with the guidance and support of the faculty advisors (Ekapol Chuangsuwanich, Yuval Pinter). Thank you all!

Preparing the proceedings is another large, time consuming task. Thanks to the Publication Chairs (Miruna Clinciu, Zhiyu Zoey Chen, Chen Liang, Bing Liu) for coordinating the preparation of all proceedings, including the main conference proceedings, findings, demonstration, SRW and workshop proceedings.

The Website Chairs (Yun-Nung Chen, Vipas Sutantayawalee) created the website and were quick in responding to our queries for updates; thank you Vipas and Yun-Nun. Thanks also to the Publicity and Social Media Chairs (Yuki Arase, Dimitra Gkatzia, Jing Jiang) who communicated and publicized the conference through various social media channels, enhancing the visibility and reach of the conference. Thanks also to the Handbook Chairs (Loic Barrault, Pierre Colombo) for creating the conference handbook that will guide you through the conference program.

Making ACL accessible to a wider community was the task of the Diversity and Inclusion Chairs (Aparna Garimella, Lin Gui, Jing Li, Steven Wilson). Thank you all for helping in fostering a diverse and inclusive environment.

Thanks to the Ethics Chairs (Aurélie Névéol, Alice Oh), who checked papers flagged with ethics issues. They had to process many more papers than expected but handled the overflow brilliantly.

The Technical Open Review Chairs (Thiago Castro-Ferreira, Taro Watanabe) helped out with Open Review related requests and the Virtual Infrastructure Chairs (Gaël Guibon, Gözde Gül, Rachada Kongkrachantra), made various enhancements to the virtual platform to ensure an engaging conference experience. Thank you!

The conference is also a successful recipient for thousands of emails. Many thanks go to the Internal Communication Chairs (Claudia Borg, Yannick Parmentier, Valentina Pyatkin), for their efficient and much needed processing of the multiple emails sent to ACL 2024.

The registration fees would be considerably higher without our sponsors generous contributions. Sincere thanks to them and to Chris Callison-Burch, the ACL sponsorship Director who, together with the Sponsorship Chairs (Lluís Marquez, Kobbrit Viriyayudhakorn) succeeded in securing the sponsorships that are crucial in helping keep registration fees down.

Sol Rosenberg and his team provide the Underline virtual platform - thank for their support in collaborating with us to meet our needs and accommodate ACL feature requests.

The local team (Thepchai Supnithi, Prachya Bookwan, Thanaruk Theeramunkong) did a wonderful job locating the venue, providing help with local and visa information, booking hotels for the participants and organizing a social program.

During the conference, student volunteers help make the conference run smoothly. Many thanks to them and to the Student Volunteer Chairs (Hao Fei, Margot Mieskes, Liangming Pan) who reviewed applications, selected the student volunteers, and assigned them their tasks.

Much of the know-how for the various chair positions is based on insights gleaned from earlier events. Thanks to previous ACL conference Chairs for sharing their experience, and to the ACL Exec for their support.

Special thanks go to Jennifer Rachford, ACL Business Manager, whose remarkable ability to juggle multiple conferences and interact with overloaded scientists make the impossible possible. Her knowledge of the ACL conferences and of their multiple aspects is essential for the success of our conferences.

Finally, let me thank you all, authors, reviewers, presenters, workshop organizers and participants of the conference. Thank you for choosing to be part of ACL 2024, I wish you a very enjoyable conference!

ACL 2024 General Chair

Claire Gardent

CNRS, France

Message from the Program Chairs

Welcome to the 62nd Annual Meeting of the Association for Computational Linguistics! ACL 2024 will feature a hybrid format, allowing attendees to join us in person in Bangkok, Thailand, or to participate remotely from anywhere in the world. We are pleased to be hosting the conference in Bangkok, which was the original planned venue for ACL 2021 before the COVID-19 pandemic forced a change. Organizing ACL 2024 has been a team effort involving thousands of people. We would like to thank the support and contributions of the following people:

- The General Chair, Claire Gardent;
- The ARR Editors-in-Chief of the Feb 2025 cycle (Viviane Moreira, Jun Suzuki) and the entire team (Lilja Øvrelid, Mausam, Thamar Solorio, Vincent Ng, Jonathan Kummerfield, Sudipta Kar);
- The OpenReview team;
- The 72 Senior Area Chairs;
- The 718 Area Chairs and the 4209 reviewers;
- The awards committee chairs, Hal Daumé, Raquel Fernández and Yuji Matsumoto, and the 20 awards committee members;
- The ethics committee led by Alice Oh and Aurélie Névéol, along with Malihe Alikhani and Vinodkumar Prabhakaran from ARR;
- The website chairs, Yun-Nung (Vivian) Chen and Vipas Sutantayawalee;
- The publication chairs, Miruna Clinciu, Bing Liu, Zhiyu Zoey Chen and Chen Liang;
- The handbook chairs Pierre Colombo and Loic Barrault
- The local organization chairs, Thepchai Supnithi, Prachya Bookwan, Thanaruk Theeramunkong, and their team;
- The publicity and social media chairs, Yuki Arase, Jing Jiang, and Dimitra Gkatzia;
- The student volunteer chairs, Margot Mieskes, Hao Fei and Liangming Pan;
- 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 NAACL 2024 Program Chairs (Kevin Duh, Helena Gomez, and Steven Bethard) and the ACL 2023 Program Chairs (Anna Rogers, Jordan Boyd-Graber and Naoki Okazaki);
- Damira Mršić and Underline Team;
- Jennifer Rachford and entire conference support staff;
- All the authors of papers submitted for review and committed to the conference.

Review process All ACL 2024 submissions were channeled through a two-stage review process: Submissions were first sent to ACL Rolling Review (ARR) for reviews (by reviewers) and meta-reviews (by area chairs). Then, authors could choose to commit their reviewed papers to ACL via a separate ACL 2024 commitment site for recommendations by senior area chairs and final acceptance decisions by the program chairs. In this, ACL 2024 follows EACL 2024 and NAACL 2024.

We worked closely with the ARR team, especially the February 2024 Editors-in-Chief, and served as guest Editors-in-Chief for this cycle. We helped recruit new reviewers and area chairs to ARR, resulting in 4209 reviewers and 718 ACs in the 2024 February ARR cycle to which most ACL 2024 papers were submitted. The 72 senior area chairs recruited by ACL helped oversee the review and meta-review process during this phase. Overall, the ARR process went mostly smoothly, successfully delivering at least three reviews and a meta-review for all papers submitted.

For the ACL commitment part of the process, the senior area chairs made acceptance recommendations for 2931 committed papers based on the papers, reviews, and meta-reviews, and program chairs finalized the recommendations into acceptance decisions.

Acceptance rate The acceptance rate calculation follows precedent set by previous conferences that also go through ARR, e.g. EACL 2024, NAACL 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. The denominator includes:

- Papers in the ARR February 2024 cycle that selected ACL as a preferred venue.
- Papers in the ARR February 2024 cycle that did not select any conference as a preferred venue.
- Papers in the ARR February 2024 cycle that selected another conference, but then committed to ACL 2024.
- Papers in the ARR cycles before February 2024 that committed to ACL 2024.

In total, we had 4,835 submissions in the ARR February 2024 cycle. Among these, 276 were withdrawn before reviews were released and 169 were desk rejected. Among the remaining, 4244 had either an unspecified venue or included ACL as a desired venue. Among the submissions that selected other venues, three papers were committed to ACL. Finally, an additional 160 papers from other cycles were committed to ACL. In total, the denominator for the acceptance rate calculation is $4244 + 3 + 160 = 4407$. Among these, 2931 were committed to ACL.

Among the committed papers, 940 were accepted to the Main Conference. The acceptance rate for Main Conference papers is therefore $940 / 4407 = 21.3\%$. A further 975 papers were accepted to Findings of ACL, representing solid work that is has been judged worthy of publication with sufficient substance, quality and novelty. The acceptance rate for Findings of ACL is $975 / 4407 = 22.1\%$.

Special Theme: Open science, open data, and open models for reproducible NLP research The rise of large language models as a general purpose tool for NLP has opened up exciting possibilities for NLP. But their widespread adoption via closed APIs has also raised concerns about transparency and reproducibility. When we do not have access to information about how these models were trained or the data they learned from, it becomes challenging to build upon existing research and compare new approaches fairly. This lack of openness poses a risk to progress in our field. With this perspective in mind, for ACL 2024, we invited submissions to a special theme titled “Open science, open data, and open models for reproducible NLP research”.

With this theme track, we sought a discussion on increased transparency in the field by promoting the use of open models and open-source initiatives in NLP as an alternative to closed approaches. We encouraged contributions related to the release of high quality datasets, novel ideas for evaluation, non-trivial algorithm and toolbox implementations, and models which are properly documented (e.g. via model cards).

We received 55 submissions to the theme track during the review phase. Among these, 22 papers were accepted to the main conference and a further 16 to Findings of ACL.

Paper Awards ACL 2024 implemented the updated ACL award policy that seeks to expand the pool of work recognized as outstanding. In total 102 papers were nominated by the reviewers, area chairs and senior area chairs for consideration. The Awards Committee assessed these papers to select the best papers (featuring $\leq 0.25\%$ of accepted papers), outstanding papers (featuring $\leq 2.5\%$ of accepted papers), and special awards for social impact and best resource. Separately, the senior area chairs of individual tracks also selected papers in their track for the area chair awards. Finally, the program chairs selected one paper from the papers submitted to the conference theme track as the best theme paper.

The recent change in ACL policy allows papers to be non-anonymous during the review process via public preprints. To recognize submissions that remained anonymous, we followed the policy recommendation to have separate best and outstanding paper awards for submissions that remained anonymous to the public during the whole process.

All the awards will be announced in a dedicated plenary session on the last day of the conference.

Program composition & presentation modes Based on feedback from the conference support staff and the Underline team after NAACL 2024, we decided to hold the virtual poster session separately, and *after* the in-person conference ended. The post-conference virtual sessions were scheduled to avoid conflicts with the in-person attendees of the conference. The goal was to encourage all attendees — both virtual and in-person — to join the virtual conference. All in-person papers accepted to the main conference were given a poster slot. In addition, 102 papers were assigned oral presentations. These papers were selected by the program chairs and the decision was motivated by the goal of having a well-rounded program with a diverse set of topics. Additionally, all Findings papers were also assigned a poster presentation in separate Findings posters sessions in the conference.

The ACL program features three exciting keynote speakers: Sunita Sarawagi, Subbarao Kambhampati and Barbara Plank. Further, to celebrate the venue of ACL, the program also including a panel discussion focusing on Southeast Asian languages featuring panelists Ayu Purwarianti, William Tjhi and Sarana Nutanong.

The ACL program also includes 6 papers accepted by the Computational Linguistics journal and 31 papers accepted by the Transactions of ACL (TACL). All journal papers whose authors were attending the conference in person were given oral presentation slots and thematically distributed in the conference in appropriate sessions. The program is rounded out with dedicated sessions during the main conference for the demonstrations track and student research workshop.

We hope that you will enjoy this year's program and conference!

Lun-Wei Ku (Academia Sinica, Taiwan)

André F. T. Martins (Instituto Superior Técnico, Instituto de Telecomunicações, Unbabel, Portugal)

Vivek Srikumar (University of Utah, USA)

ACL 2024 Program Committee Co-Chairs

Organizing Committee

General Chair

Claire Gardent, CNRS and Université de Lorraine

Program Chairs

Lun-Wei Ku, Academia Sinica

Andre Martins, Instituto Superior Técnico / Instituto de Telecomunicações / Unbabel

Vivek Srikumar, University of Utah

Local Organization

Thepchai Supnithi, NECTEC and AIAT

Prachya Bookwan, NECTEC and AIAT

Thanaruk Theeramunkong, SIIT and AIAT

Workshop Chairs

Xipeng Qiu, Fudan University

Eunsol Choi, The University of Texas at Austin

Tutorial Chairs

Luis Chiruzzo, Universidad de la República

Hung-yi Lee, National Taiwan University

Leonardo Ribeiro, Amazon Alexa Seattle

Demonstration Chairs

Yixin Cao, Singapore Management University

Yang Feng, Chinese Academy of Science

Deyi Xiong, Tianjin University

Student Research

Xiyan Fu, Heidelberg University

Eve Fleisig, UC Berkeley

Student Research: Faculty Advisor

Ekapol Chuangsuwanich, Chulalongkorn University

Yuval Pinter, Ben Gurion University

Publicity and Social Media Chairs

Yuki Arase, Osaka University

Jing Jiang, Singapore Management University

Dimitra Gkatzia, Napier University

Publication Chairs

Miruna Clinciu, University of Edinburgh
Bing Liu, Meta AI
Zhiyu Zoey Chen, University of Texas at Dallas
Chen Liang, Google DeepMind

Handbook Chairs

Pierre Colombo, Université Paris Saclay
Loic Barrault, Meta AI

Technical Open Review

Taro Watanabe, Nara Institute of Science and Technology (NAIST)
Thiago Castro, Federal University of Minas Gerais

Diversity and Inclusion

Jing Li, Hong Kong Polytechnic University
Aparna Garimella, Adobe Research
Steven Wilson, Oakland University
Lin Gui, King's College London

Ethic Committee

Alice Oh, KAIST
Aurélie Névéol, CNRS and Université Paris Saclay

Internal Communications

Claudia Borg, University of Malta
Valentina Pyatkin, Allen Institute for AI
Yannick Parmentier, Université de Lorraine

Student Volunteer

Margot Mieskes, University of Applied Science, Darmstadt
Hao Fei, National University of Singapore
Liangming Pan, University of California, Santa Barbara

Virtual Infrastructure

Gözde Gül, Koç University
Gael Guibon, University of Lorraine
Rachada Kongkrachantra, Thammasat University

Website and Conference App

Yun-Nung (Vivian) Chen, National Taiwan University
Vipas Sutantayawalee, Artificial Intelligence Entrepreneur Association of Thailand

Sponsorship Chairs

Lluis Marquez, Amazon
Kobkrit Viriyayudhakorn, iAPP Co., Ltd

Program Committee

SRW Program Chairs

Miruna Clinciu
Eve Fleisig
Xiyan Fu

SRW Reviewers

Piush Aggarwal, Chandan Akiti, Abeer Aldayel, Miguel A. Alonso, Begoa Altuna, Sergi Alvarez-Vidal, Bharat Ram Ambati, Evelin Amorim, Tatiana Anikina, Raghav Arora

Parsa Bagherzadeh, Ahsaas Bajaj, Valerio Basile, Anil Batra, Tim Baumgrtner, Rachel Bawden, Shabnam Behzad, Gabriel Bernier-Colborne, Anne Beyer, Isil Poyraz Bilgin, Shikha Bordia

Meng Cao, Arlene Casey, Silvia Casola, Canyu Chen, William Chen, Hyundong Justin Cho, Sukmin Cho, Juhwan Choi, Arijit Ghosh Chowdhury, Ekapol Chuangsuwanich, Rodolfo Corona, Amanda Cercas Curry

Aryia Dattamajumdar, Brian Davis, Alok Debnath, Ayse Ayyuce Demirbas, Nesat Dereli, Chris Develder, Noshin Tabassum Dina, Bonaventure F. P. Dossou, Zi-Yi Dou, Ritam Dutt, Aditi Dutta

Micha Elsner

Neele Falk, Antnio Farinhas, Fauzan Farooqui, Eve Fleisig, Francesca Franzon, Dayne Freitag, Xiyan Fu

Diana Galvan, Shubham Garg, Aina Gar Soler, Ratul Ghosh, Kristina Gligoric, Anmol Goel, Satya Kesav Gundabathula, Ankita Gupta

Ivan Habernal, Youssef Al Hariri, Yifan Hou, Phu Mon Htut, Tiancheng Hu, Zhongzhan Huang

Robert L. Logan IV

Glorianna Jagfeld, Labiba Jahan, Thanmay Jayakumar, Jyoti Jha, Kenan Jiang, Kristen Johnson

Jad Kabbara, Patrick Kahardipraja, Mihir Kale, Amr Keleg, Yova Kementchedjhieva, Mohammad Aflah Khan, Harshita Khandelwal, Anirudh Khattri, Elena Kochkina, Prashant Kodali, Mamoru Komachi, Fajri Koto, Pranoy Kovuri, Devang Kulshreshtha, Kemal Kurniawan, Mascha Kurpicz-Briki

Alexandra Lavrentovich, Amirhossein Layegh, Messi H.j. Lee, Yoonsang Lee, Jia Li, Manling Li, Sha Li, Siyan Li, Baohao Liao, Jasy Suet Yan Liew, Zheng Wei Lim, Peerat Limkonchotiwat, Lucy H. Lin, Tatiana Litvinova, Wei Liu, Li Lucy, Bohan Lyu

Bolei Ma, Valentin Malykh, Lara J. Martin, Sandeep Mathias, Shubhanshu Mishra, Sai Ajay Mody, Moku

Hongbin Na, Murali Nandan Nagarapu, Sachin Narayan Nagargoje, Masaaki Nagata, Pakawat

Nakwijit, Kota Shamanth Ramanath Nayak, Nihal V. Nayak, Mariana Neves, Kiet Van Nguyen, Ratna Nirupama

Naoki Otani, Kun Ouyang

Maria Leonor Pacheco, Jiayi Pan, Artemis Panagopoulou, Sara Papi, Tanmay Parekh, Ajay Patel, Naiara Perez, Gabriele Pergola, Ben Peters, Shubham Milind Phal, Van-Thuy Phi, Dina Pisarevskaya, Dongqi Pu, Rajkumar Pujari

Yusu Qian, Yusu Qian

Sunny Rai, Surangika Ranathunga, Pretam Ray, Frederick Riemenschneider, Shruti Rijhwani, Hossein Rouhizadeh, Markian Rybchuk, Sangwon Ryu

Farig Sadeque, Philipp Sadler, Oscar Sainz, Nicholas Sanders, Sashank Santhanam, Abulhair Saparov, Ryohei Sasano, Tatjana Scheffler, Fabian Schmidt, Sebastian Schuster, Indira Sen, Sheikh Shafayat, Aditya Shah, Roshan Sharma, Sina Sheikholeslami, Kejian Shi, Manish Shrivastava, Chenglei Si, Pushpdeep Singh, Kevin Small, Cesare Spinoso-Di Piano, Richard Sproat, Marija Stanojevic, Mihai Surdeanu, Dage Srg, Dage Srg

Shabnam Tafreshi, Zeerak Talat, Alberto Testoni, Nicholas Tomlin, Son Quoc Tran

Can Udomcharoenchaikit

Sowmya Vajjala, Andrea Varga, Francielle Vargas, Deeksha Varshney, David Vilar

Hao Wang, Mingyang Wang, Bonnie Webber, Alexander Wettig, Chenglin Wu, Xinyi Wu

Qi Xin, Rui Xing, Fuzhao Xue

Dixi Yao, Michael Miller Yoder

Olga Zamaraeva, Zenan Zhai, Chenhao Zhang, Haoran Zhang, Mike Zhang, Yangjun Zhang, Yao Zhang, Yiqun Zhang, Zhehao Zhang, Siheng Zhao, Wenjie Zhao, Xin Zhao, Yu Zhao, Xinyi Zheng, Shanshan Zhong, Naitian Zhou, Rongxin Zhu, Elena Zotova

Omri Abend, Hebrew University of Jerusalem

Fernando Alva-Manchego, Cardiff University

Maria Antoniak

Marianna Apidianaki, University of Pennsylvania, University of Pennsylvania

Esma Balkir, Wisetech Global

Maria Barrett, IT University of Copenhagen

Tilman Beck

Terra Blevins, University of Washington

Su Lin Blodgett, Microsoft

Elia Bruni, Universitt Osnabrck

Agostina Calabrese, University of Edinburgh

Ronald Cardenas

Asli Celikyilmaz, FAIR

Ming-Wei Chang, Google
Yen-Chun Chen, Microsoft
Pinzhen Chen, University of Edinburgh
Mingda Chen, Meta AI
Danqi Chen, Department of Computer Science, Princeton University
Christos Christodoulopoulos, Amazon
Alessandra Teresa Cignarella, aequa-tech srl
Volkan Cirik, ASAPP
Xiang Dai, CSIRO
Marco Del Tredici, Amazon
David Demeter
Laura Dietz, University of New Hampshire
Steffen Eger, Universitt Mannheim
Yanai Elazar, Allen Institute for Artificial Intelligence and Department of Computer Science
Amir Feder, Columbia University and Google
Cesar Ferri
David Gaddy, Google
Kuzman Ganchev, Google and School of Engineering and Applied Science
Tianyu Gao
Marcos Garcia, Universidade de Santiago de Compostela
Iker Garca-Ferrero
Jonas Groschwitz
Vivek Gupta, University of Pennsylvania, United States
Barry Haddow, University of Edinburgh
Andreas Hanselowski
Junxian He, Hong Kong University of Science and Technology
Yulan He, King's College London, University of London
Dieuwke Hupkes, Facebook
Mimansa Jaiswal, Norm AI
Zhijing Jin
Tomoyuki Kajiwara, Ehime University
Ehsan Kamalloo, ServiceNow Inc
Sudipta Kar, Amazon
Alina Karakanta, Leiden University
Shima Khanehzar, CSIRO
Daniel Khashabi, Johns Hopkins University
Gunhee Kim, Seoul National University
Mandy Barrett Korpusik, Loyola Marymount University
Julia Kreutzer, Cohere for AI
Kalpesh Krishna, Google
Ilia Kuznetsov, TU Darmstadt
Gorka Labaka, Universidad del Pas Vasco
Gabriella Lapesa, GESIS Leibniz Institute for the Social Sciences and Heinrich-Heine University
Dsseldorf
Ji-Ung Lee, TU Darmstadt
Lei Li, University of Hong Kong
Tianyi Li
None

Keynote

Does In-Context-Learning Offer the Best Tradeoff in Accuracy, Robustness, and Efficiency for Model Adaptation?

Sunita Sarawagi

Indian Institute of Technology Bombay, India



08/12/2024 – Time: 09:30 - 10:30 – Room: Convention Center B1

Abstract: Adapting a model trained on vast amounts of data to new tasks with limited labeled data has long been a challenging problem, and over the years, a diverse range of techniques have been explored. Effective model adaptation requires achieving high accuracy through task-specific specialization without forgetting previous learnings, robustly handling the high variance from limited task-relevant supervision, and doing so efficiently with minimal compute and memory overheads. Recently, large language models (LLMs) have demonstrated remarkable ease of adaptation to new tasks with just a few examples provided in context, without any explicit training for such a capability. Puzzled by this apparent success, many researchers have sought to explain why in-context learning (ICL) works, but we still have only an incomplete understanding. In this talk, we examine this emerging phenomenon and assess its potential to meet our longstanding model adaptation goals in terms of accuracy, robustness, and efficiency.

Bio: Sunita Sarawagi researches in the fields of databases, machine learning, and applied NLP. She got her PhD in databases from the University of California at Berkeley and a bachelors degree from IIT Kharagpur. She has also worked at Google Research, CMU, and IBM Almaden Research Center. She is an ACM fellow, was awarded the Infosys Prize in 2019 for Engineering and Computer Science, and the distinguished Alumnus award from IIT Kharagpur. She has several publications in database, machine learning, and NLP conferences including notable paper awards at ACM SIGMOD, ICDM, and NeurIPS conferences.

Keynote

Can LLMs Reason and Plan?

Subbarao Kambhampati
Arizona State University, USA



08/13/2024 – Time: 09:00 - 10:00 – Room: Convention Center B1

Abstract: Large Language Models (LLMs) are on track to reverse what seemed like an inexorable shift of AI from explicit to tacit knowledge tasks. Trained as they are on everything ever written on the web, LLMs exhibit “approximate omniscience”—they can provide answers to all sorts of queries, but with nary a guarantee. This could herald a new era for knowledge-based AI systems—with LLMs taking the role of (blowhard?) experts. But first, we have to stop confusing the impressive style/form of the generated knowledge for correct/factual content, and resist the temptation to ascribe reasoning, planning, self-critiquing etc. powers to approximate retrieval by these n-gram models on steroids. We have to focus instead on LLM-Modulo techniques that complement the unfettered idea generation of LLMs with careful vetting by model-based verifiers (the models underlying which themselves can be teased out from LLMs in semi-automated fashion). In this talk, I will reify this vision and attendant caveats in the context of our ongoing work on understanding the role of LLMs in planning tasks.

Bio: Subbarao Kambhampati is a professor of computer science at Arizona State University. Kambhampati studies fundamental problems in planning and decision making, motivated in particular by the challenges of human-aware AI systems. He is a fellow of Association for the Advancement of Artificial Intelligence, American Association for the Advancement of Science, and Association for Computing machinery. He served as the president of the Association for the Advancement of Artificial Intelligence, a trustee of the International Joint Conference on Artificial Intelligence, the chair of AAAS Section T (Information, Communication and Computation), and a founding board member of Partnership on AI. Kambhampati’s research as well as his views on the progress and societal impacts of AI have been featured in multiple national and international media outlets. He can be followed on Twitter @rao2z.

Keynote
**Are LLMs Narrowing Our Horizon? Let's Embrace
Variation in NLP!**

Barbara Plank
Ludwig Maximilian University of Munich, Germany



08/14/2024 – Time: 09:00 - 10:00 – Room: Convention Center B1

Abstract: NLP research has made significant progress, and our community's achievements are becoming deeply integrated in society. The recent paradigm shift due to rapid advances in Large Language Models (LLMs) offers immense potential, but also led NLP to become more homogeneous. In this talk, I will argue for the importance of embracing variation in research, which will lead to more innovation, and in turn, trust. I will give an overview of current challenges and show how they led to the loss of trust in our models. To counter this, I propose to embrace variation in three key areas: inputs to models, outputs of models and research itself. Embracing variation holistically will be crucial to move our field towards more trustworthy human-facing NLP.

Bio: Barbara Plank is Professor and co-director of the Center for Information and Language Processing at LMU Munich. She holds the Chair for AI and Computational Linguistics at LMU Munich and is an affiliated Professor at the Computer Science department at the IT University of Copenhagen. Her MaiNLP research lab (Munich AI and NLP lab, pronounced “my NLP”) focuses on robust machine learning for Natural Language Processing with an emphasis on human-inspired and data-centric approaches. Her research has been funded by distinguished grants, including an Amazon Research Award (2018), the Danish Research Council (Sapere Aude Research Leader Grant, 2020-2024), and the European Research Council (ERC Consolidator Grant, 2022-2027). Barbara is a Scholar of ELLIS (the European Laboratory for Learning and Intelligent Systems) and regularly serves on international committees, including the Association for Computational Linguistics (ACL), the European Chapter of the ACL, and the Northern European Association for Language Technology (NEALT).

Table of Contents

<i>Quantized Side Tuning: Fast and Memory-Efficient Tuning of Quantized Large Language Models</i> Zhengxin Zhang, Dan Zhao, Xupeng Miao, Gabriele Oliaro, Zhihao Zhang, Qing Li, Yong Jiang and Zhihao Jia	1
<i>Unsupervised Multimodal Clustering for Semantics Discovery in Multimodal Utterances</i> Hanlei Zhang, Hua Xu, Fei Long, Xin Wang and Kai Gao	18
<i>MAGE: Machine-generated Text Detection in the Wild</i> Yafu Li, Qintong Li, Leyang Cui, Wei Bi, Zhilin Wang, Longyue Wang, Linyi Yang, Shuming Shi and Yue Zhang	36
<i>PrivLM-Bench: A Multi-level Privacy Evaluation Benchmark for Language Models</i> Haoran Li, Dadi Guo, Donghao Li, Wei Fan, Qi Hu, Xin Liu, Chunkit Chan, Duanyi Yao, Yuan Yao and Yangqiu Song	54
<i>GenTranslate: Large Language Models are Generative Multilingual Speech and Machine Translators</i> Yuchen Hu, Chen Chen, Chao-Han Huck Yang, Ruizhe Li, Dong Zhang, Zhehuai Chen and EngSiong Chng	74
<i>Exploring Chain-of-Thought for Multi-modal Metaphor Detection</i> Yanzhi Xu, Yueying Hua, Shichen Li and Zhongqing Wang	91
<i>BitDistiller: Unleashing the Potential of Sub-4-Bit LLMs via Self-Distillation</i> DaYou Du, Yijia Zhang, Shijie Cao, Jiaqi Guo, Ting Cao, Xiaowen Chu and Ningyi Xu	102
<i>A Unified Temporal Knowledge Graph Reasoning Model Towards Interpolation and Extrapolation</i> Kai Chen, Ye Wang, Yitong Li, Aiping Li, Han Yu and Xin Song	117
<i>Unsupervised Information Refinement Training of Large Language Models for Retrieval-Augmented Generation</i> Shicheng Xu, Liang Pang, Mo Yu, Fandong Meng, Huawei Shen, Xueqi Cheng and Jie Zhou	133
<i>CSCD-NS: a Chinese Spelling Check Dataset for Native Speakers</i> Yong Hu, Fandong Meng and Jie Zhou	146
<i>Evaluating Dynamic Topic Models</i> Charu Karakkaparambil James, Mayank Nagda, Nooshin Haji Ghassemi, Marius Kloft and Sophie Fellenz	160
<i>How Abilities in Large Language Models are Affected by Supervised Fine-tuning Data Composition</i> Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou and Jingren Zhou	177
<i>Through the Lens of Split Vote: Exploring Disagreement, Difficulty and Calibration in Legal Case Outcome Classification</i> Shanshan Xu, Santosh T.y.s.s, Oana Ichim, Barbara Plank and Matthias Grabmair	199
<i>Inference to the Best Explanation in Large Language Models</i> Dhairya Dalal, Marco Valentino, Andre Freitas and Paul Buitelaar	217
<i>A Novel Cartography-Based Curriculum Learning Method Applied on RoNLI: The First Romanian Natural Language Inference Corpus</i> Eduard Poesina, Cornelia Caragea and Radu Tudor Ionescu	236

<i>MinPrompt: Graph-based Minimal Prompt Data Augmentation for Few-shot Question Answering</i> Xiusi Chen, Jyun-Yu Jiang, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu and Wei Wang	254
<i>SportsMetrics: Blending Text and Numerical Data to Understand Information Fusion in LLMs</i> Yebowen Hu, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Hassan Foroosh, Dong Yu and Fei Liu	267
<i>SciMON: Scientific Inspiration Machines Optimized for Novelty</i> Qingyun Wang, Doug Downey, Heng Ji and Tom Hope	279
<i>Expedited Training of Visual Conditioned Language Generation via Redundancy Reduction</i> Yiren Jian, Tingkai Liu, Yunzhe Tao, Chunhui Zhang, Soroush Vosoughi and Hongxia Yang	300
<i>Confidence Under the Hood: An Investigation into the Confidence-Probability Alignment in Large Language Models</i> Abhishek Kumar, Robert Morabito, Sanzhar Umbet, Jad Kabbara and Ali Emami	315
<i>Retrieval-Augmented Multilingual Knowledge Editing</i> Weixuan Wang, Barry Haddow and Alexandra Birch	335
<i>Picturing Ambiguity: A Visual Twist on the Winograd Schema Challenge</i> Brendan Park, Madeline Janecek, Naser Ezzati-Jivan, Yifeng Li and Ali Emami	355
<i>Subtle Biases Need Subtler Measures: Dual Metrics for Evaluating Representative and Affinity Bias in Large Language Models</i> Abhishek Kumar, Sarfaroz Yunusov and Ali Emami	375
<i>Framing in the Presence of Supporting Data: A Case Study in U.S. Economic News</i> Alexandria Leto, Elliot E. Pickens, Coen D. Needell, David Rothschild and Maria Leonor Pacheco	393
<i>Mementos: A Comprehensive Benchmark for Multimodal Large Language Model Reasoning over Image Sequences</i> Xiyao Wang, Yuhang Zhou, Xiaoyu Liu, Hongjin Lu, Yuancheng Xu, Feihong He, Jaehong Yoon, Taixi Lu, Fuxiao Liu, Gedas Bertasius, Mohit Bansal, Huaxiu Yao and Furong Huang	416
<i>TTM-RE: Memory-Augmented Document-Level Relation Extraction</i> Chufan Gao, Xuan Wang and Jimeng Sun	443
<i>Answer is All You Need: Instruction-following Text Embedding via Answering the Question</i> Letian Peng, Yuwei Zhang, Zilong Wang, Jayanth Srinivasa, Gaowen Liu, Zihan Wang and Jingbo Shang	459
<i>Explore Spurious Correlations at the Concept Level in Language Models for Text Classification</i> Yuhang Zhou, Paiheng Xu, Xiaoyu Liu, Bang An, Wei Ai and Furong Huang	478
<i>Every Answer Matters: Evaluating Commonsense with Probabilistic Measures</i> Qi Cheng, Michael Boratko, Pranay Kumar Yelugam, Tim O’Gorman, Nalini Singh, Andrew McCallum and Xiang Lorraine Li	493
<i>GradSafe: Detecting Jailbreak Prompts for LLMs via Safety-Critical Gradient Analysis</i> Yueqi Xie, Minghong Fang, Renjie Pi and Neil Zhenqiang Gong	507
<i>Pouring Your Heart Out: Investigating the Role of Figurative Language in Online Expressions of Empathy</i> Gyeongun Lee, Christina Wong, Meghan Guo and Natalie Parde	519

<i>An Information-Theoretic Approach to Analyze NLP Classification Tasks</i>	
Luran Wang, Mark Gales and Vatsal Raina	530
<i>Can Your Model Tell a Negation from an Implicature? Unravelling Challenges With Intent Encoders</i>	
Yuwei Zhang, Siffi Singh, Sailik Sengupta, Igor Shalymov, Hang Su, Hwanjun Song and Saab Mansour	552
<i>Wav2Gloss: Generating Interlinear Glossed Text from Speech</i>	
Taiqi He, Kwanghee Choi, Lindia Tjuatja, Nathaniel Romney Robinson, Jiatong Shi, Shinji Watanabe, Graham Neubig, David R Mortensen and Lori Levin	568
<i>Leveraging Codebook Knowledge with NLI and ChatGPT for Zero-Shot Political Relation Classification</i>	
Yibo Hu, Erick Skorupa Parolin, Latifur Khan, Patrick Brandt, Javier Osorio and Vito D’Orazio	583
<i>SPOR: A Comprehensive and Practical Evaluation Method for Compositional Generalization in Data-to-Text Generation</i>	
Ziyao Xu and Houfeng Wang	604
<i>OPEX: A Component-Wise Analysis of LLM-Centric Agents in Embodied Instruction Following</i>	
Haochen Shi, Zhiyuan Sun, Xingdi Yuan, Marc-Alexandre Côté and Bang Liu	622
<i>Multimodal Instruction Tuning with Conditional Mixture of LoRA</i>	
Ying Shen, Zhiyang Xu, Qifan Wang, Yu Cheng, Wenpeng Yin and Lifu Huang	637
<i>DocLens: Multi-aspect Fine-grained Medical Text Evaluation</i>	
Yiqing Xie, Sheng Zhang, Hao Cheng, Pengfei Liu, Zelalem Gero, Cliff Wong, Tristan Naumann, Hoifung Poon and Carolyn Rose	649
<i>FOFO: A Benchmark to Evaluate LLMs’ Format-Following Capability</i>	
Congying Xia, Chen Xing, Jiangshu Du, Xinyi Yang, Yihao Feng, Ran Xu, Wenpeng Yin and Caiming Xiong	680
<i>Hyper-CL: Conditioning Sentence Representations with Hypernetworks</i>	
Young Hyun Yoo, Jii Cha, Changhyeon Kim and Taeuk Kim	700
<i>Analysis of Multi-Source Language Training in Cross-Lingual Transfer</i>	
Seonghoon Lim, Taejun Yun, Jinhyeon Kim, Jihun Choi and Taeuk Kim	712
<i>ABEX: Data Augmentation for Low-Resource NLU via Expanding Abstract Descriptions</i>	
Sreyan Ghosh, Utkarsh Tyagi, Sonal Kumar, Chandra Kiran Reddy Evuru, Ramaneswaran S, S Sakshi and Dinesh Manocha	726
<i>The Bebebe Benchmark: a Parallel Reading Comprehension Dataset in 122 Language Variants</i>	
Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer and Madian Khabsa	749
<i>Learn from Failure: Fine-tuning LLMs with Trial-and-Error Data for Intuitionistic Propositional Logic Proving</i>	
Chenyang An, Zhibo Chen, Qihao Ye, Emily First, Letian Peng, Jiayun Zhang, Zihan Wang, Sorin Lerner and Jingbo Shang	776
<i>Interactive Text-to-Image Retrieval with Large Language Models: A Plug-and-Play Approach</i>	
Saehyung Lee, Sangwon Yu, Junsung Park, Jihun Yi and Sungroh Yoon	791

<i>IMBUE: Improving Interpersonal Effectiveness through Simulation and Just-in-time Feedback with Human-Language Model Interaction</i>	
Inna Wanyin Lin, Ashish Sharma, Christopher Michael Rytting, Adam S Miner, Jina Suh and Tim Althoff	810
<i>Token-wise Influential Training Data Retrieval for Large Language Models</i>	
Huawei Lin, Jikai Long, Zhaozhuo Xu and Weijie Zhao	841
<i>Tree-of-Counterfactual Prompting for Zero-Shot Stance Detection</i>	
Maxwell Weinzierl and Sanda Harabagiu	861
<i>VisualWebArena: Evaluating Multimodal Agents on Realistic Visual Web Tasks</i>	
Jing Yu Koh, Robert Lo, Lawrence Keunho Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Russ Salakhutdinov and Daniel Fried	881
<i>FineSurE: Fine-grained Summarization Evaluation using LLMs</i>	
Hwanjun Song, Hang Su, Igor Shalyminov, Jason Cai and Saab Mansour	906
<i>Tuning Large Multimodal Models for Videos using Reinforcement Learning from AI Feedback</i>	
Daechul Ahn, Yura Choi, Youngjae Yu, Dongyeop Kang and Jonghyun Choi	923
<i>Prompt Refinement with Image Pivot for Text-to-Image Generation</i>	
Jingtao Zhan, Qingyao Ai, Yiqun Liu, Yingwei Pan, Ting Yao, Jiabin Mao, Shaoping Ma and Tao Mei	941
<i>Striking Gold in Advertising: Standardization and Exploration of Ad Text Generation</i>	
Masato Mita, Soichiro Murakami, Akihiko Kato and Peinan Zhang	955
<i>AbsInstruct: Eliciting Abstraction Ability from LLMs through Explanation Tuning with Plausibility Estimation</i>	
Zhaowei Wang, Wei Fan, Qing Zong, Hongming Zhang, Sehyun Choi, Tianqing Fang, Xin Liu, Yangqiu Song, Ginny Wong and Simon See	973
<i>Reflect-RL: Two-Player Online RL Fine-Tuning for LMs</i>	
Runlong Zhou, Simon Shaolei Du and Beibin Li	995
<i>Can ChatGPT's Performance be Improved on Verb Metaphor Detection Tasks? Bootstrapping and Combining Tacit Knowledge</i>	
Cheng Yang, Puli Chen and Qingbao Huang	1016
<i>Self-Distillation Bridges Distribution Gap in Language Model Fine-Tuning</i>	
Zhaorui Yang, Tianyu Pang, Haozhe Feng, Han Wang, Wei Chen, Minfeng Zhu and Qian Liu	1028
<i>An Information Bottleneck Perspective for Effective Noise Filtering on Retrieval-Augmented Generation</i>	
Kun Zhu, Xiaocheng Feng, Xiyuan Du, Yuxuan Gu, Weijiang Yu, Haotian Wang, Qianglong Chen, Zheng Chu, Jingchang Chen and Bing Qin	1044
<i>RORA: Robust Free-Text Rationale Evaluation</i>	
Zhengping Jiang, Yining Lu, Hanjie Chen, Daniel Khashabi, Benjamin Van Durme and Anqi Liu	1070
<i>Tell Me More! Towards Implicit User Intention Understanding of Language Model Driven Agents</i>	
Cheng Qian, Bingxiang He, Zhong Zhuang, Jia Deng, Yujia Qin, Xin Cong, Zhong Zhang, Jie Zhou, Yankai Lin, Zhiyuan Liu and Maosong Sun	1088
<i>InstructProtein: Aligning Human and Protein Language via Knowledge Instruction</i>	
Zeyuan Wang, Qiang Zhang, Keyan Ding, Ming Qin, Xiang Zhuang, Xiaotong Li and Huajun Chen	1114

<i>ConSiDERS-The-Human Evaluation Framework: Rethinking Human Evaluation for Generative Large Language Models</i>	
Aparna Elangovan, Ling Liu, Lei Xu, Sravan Babu Bodapati and Dan Roth	1137
<i>Linguistically Conditioned Semantic Textual Similarity</i>	
Jingxuan Tu, Keer Xu, Liulu Yue, Bingyang Ye, Kyeongmin Rim and James Pustejovsky . .	1161
<i>Navigate through Enigmatic Labyrinth A Survey of Chain of Thought Reasoning: Advances, Frontiers and Future</i>	
Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin and Ting Liu	1173
<i>TimeBench: A Comprehensive Evaluation of Temporal Reasoning Abilities in Large Language Models</i>	
Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Haotian Wang, Ming Liu and Bing Qin	1204
<i>BeamAggR: Beam Aggregation Reasoning over Multi-source Knowledge for Multi-hop Question Answering</i>	
Zheng Chu, Jingchang Chen, Qianglong Chen, Haotian Wang, Kun Zhu, Xiyuan Du, Weijiang Yu, Ming Liu and Bing Qin	1229
<i>ANALOGYKB: Unlocking Analogical Reasoning of Language Models with A Million-scale Knowledge Base</i>	
Siyu Yuan, Jiangjie Chen, Changzhi Sun, Jiaqing Liang, Yanghua Xiao and Deqing Yang . .	1249
<i>TaSL: Continual Dialog State Tracking via Task Skill Localization and Consolidation</i>	
Yujie Feng, Xu Chu, Yongxin Xu, Guangyuan Shi, Bo Liu and Xiao-Ming Wu	1266
<i>DeepSeekMoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models</i>	
Damai Dai, Chengqi Deng, Chenggang Zhao, R.x. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y.k. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui and Wenfeng Liang	1280
<i>Grounding Language Model with Chunking-Free In-Context Retrieval</i>	
Hongjin Qian, Zheng Liu, Kelong Mao, Yujia Zhou and Zhicheng Dou	1298
<i>Advancing Abductive Reasoning in Knowledge Graphs through Complex Logical Hypothesis Generation</i>	
Jiaxin Bai, Yicheng Wang, Tianshi Zheng, Yue Guo, Xin Liu and Yangqiu Song	1312
<i>Active Prompting with Chain-of-Thought for Large Language Models</i>	
Shizhe Diao, Pengcheng Wang, Yong Lin, Rui Pan, Xiang Liu and Tong Zhang	1330
<i>EasyGen: Easing Multimodal Generation with BiDiffuser and LLMs</i>	
Xiangyu Zhao, Bo Liu, Qijiong Liu, Guangyuan Shi and Xiao-Ming Wu	1351
<i>Rewriting the Code: A Simple Method for Large Language Model Augmented Code Search</i>	
Haochen Li, Xin Zhou and Zhiqi Shen	1371
<i>A Multidimensional Framework for Evaluating Lexical Semantic Change with Social Science Applications</i>	
Naomi Baes, Nick Haslam and Ekaterina Vylomova	1390
<i>Mitigating Catastrophic Forgetting in Large Language Models with Self-Synthesized Rehearsal</i>	
Jianheng Huang, Leyang Cui, Ante Wang, Chengyi Yang, Xinting Liao, Linfeng Song, Junfeng Yao and Jinsong Su	1416

<i>Enhancing Large Language Models in Coding Through Multi-Perspective Self-Consistency</i> Baizhou Huang, Shuai Lu, Xiaojun Wan and Nan Duan	1429
<i>Citation-Enhanced Generation for LLM-based Chatbots</i> Weitao Li, Junkai Li, Weizhi Ma and Yang Liu	1451
<i>Transitive Consistency Constrained Learning for Entity-to-Entity Stance Detection</i> Haoyang Wen, Eduard Hovy and Alexander G Hauptmann	1467
<i>Feature-Adaptive and Data-Scalable In-Context Learning</i> Jiahao Li, Quan Wang, Licheng Zhang, Guoqing Jin and Zhendong Mao	1481
<i>Probing the Multi-turn Planning Capabilities of LLMs via 20 Question Games</i> Yizhe Zhang, Jiarui Lu and Navdeep Jaitly	1495
<i>WaterBench: Towards Holistic Evaluation of Watermarks for Large Language Models</i> Shangqing Tu, Yuliang Sun, Yushi Bai, Jifan Yu, Lei Hou and Juanzi Li	1517
<i>Dependency Transformer Grammars: Integrating Dependency Structures into Transformer Language Models</i> Yida Zhao, Chao Lou and Kewei Tu	1543
<i>A Non-autoregressive Generation Framework for End-to-End Simultaneous Speech-to-Any Translation</i> Zhengrui Ma, Qingkai Fang, Shaolei Zhang, Shoutao Guo, Yang Feng and Min Zhang	1557
<i>Probing Language Models for Pre-training Data Detection</i> Zhenhua Liu, Tong Zhu, Chuanyuan Tan, Bing Liu, Haonan Lu and Wenliang Chen	1576
<i>Analyzing Temporal Complex Events with Large Language Models? A Benchmark towards Temporal, Long Context Understanding</i> Zhihan Zhang, Yixin Cao, Chenchen Ye, Yunshan Ma, Lizi Liao and Tat-Seng Chua	1588
<i>IBSEN: Director-Actor Agent Collaboration for Controllable and Interactive Drama Script Generation</i> Senyu Han, Lu Chen, Li-Min Lin, Zhengshan Xu and Kai Yu	1607
<i>Language Model Adaption for Reinforcement Learning with Natural Language Action Space</i> Jiangxing Wang, Jiachen Li, Xiao Han, Deheng Ye and Zongqing Lu	1620
<i>Evaluating Intention Detection Capability of Large Language Models in Persuasive Dialogues</i> Hiromasa Sakurai and Yusuke Miyao	1635
<i>LongLLMLingua: Accelerating and Enhancing LLMs in Long Context Scenarios via Prompt Compression</i> Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang and Lili Qiu	1658
<i>Persuading across Diverse Domains: a Dataset and Persuasion Large Language Model</i> Chuhao Jin, Kening Ren, Lingzhen Kong, Xiting Wang, Ruihua Song and Huan Chen	1678
<i>HealMe: Harnessing Cognitive Reframing in Large Language Models for Psychotherapy</i> Mengxi Xiao, Qianqian Xie, Ziyang Kuang, Zhicheng Liu, Kailai Yang, Min Peng, Weiguang Han and Jimin Huang	1707
<i>Multimodal Prompt Learning with Missing Modalities for Sentiment Analysis and Emotion Recognition</i> Zirun Guo, Tao Jin and Zhou Zhao	1726

<i>An Effective Pronunciation Assessment Approach Leveraging Hierarchical Transformers and Pre-training Strategies</i>	
Bi-Cheng Yan, Jiun-Ting Li, Yi-Cheng Wang, Hsin Wei Wang, Tien-Hong Lo, Yung-Chang Hsu, Wei-Cheng Chao and Berlin Chen	1737
<i>Detection-Correction Structure via General Language Model for Grammatical Error Correction</i>	
Wei Li and Houfeng Wang	1748
<i>Generative Pre-trained Speech Language Model with Efficient Hierarchical Transformer</i>	
Yongxin Zhu, Dan Su, Liqiang He, Linli Xu and Dong Yu	1764
<i>Selene: Pioneering Automated Proof in Software Verification</i>	
Lichen Zhang, Shuai Lu and Nan Duan	1776
<i>Dissecting Human and LLM Preferences</i>	
Junlong Li, Fan Zhou, Shichao Sun, Yikai Zhang, Hai Zhao and Pengfei Liu	1790
<i>UniCoder: Scaling Code Large Language Model via Universal Code</i>	
Tao Sun, Linzheng Chai, Jian Yang, Yuwei Yin, Hongcheng Guo, Jiaheng Liu, Bing Wang, Liquan Yang and Zhoujun Li	1812
<i>AoE: Angle-optimized Embeddings for Semantic Textual Similarity</i>	
Xianming LI and Jing Li	1825
<i>InCharacter: Evaluating Personality Fidelity in Role-Playing Agents through Psychological Interviews</i>	
Xintao Wang, Yunze Xiao, Jen-tse Huang, Siyu Yuan, Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang Leng, Wei Wang, Jiangjie Chen, Cheng Li and Yanghua Xiao	1840
<i>Does DetectGPT Fully Utilize Perturbation? Bridging Selective Perturbation to Fine-tuned Contrastive Learning Detector would be Better</i>	
Shengchao Liu, Xiaoming Liu, Yichen Wang, Zehua Cheng, Chengzhengxu Li, Zhaohan Zhang, Yu Lan and Chao Shen	1874
<i>AFaCTA: Assisting the Annotation of Factual Claim Detection with Reliable LLM Annotators</i>	
Jingwei Ni, Minjing Shi, Dominik Stammach, Mrinmaya Sachan, Elliott Ash and Markus Leippold	1890
<i>Towards Faithful and Robust LLM Specialists for Evidence-Based Question-Answering</i>	
Tobias Schimanski, Jingwei Ni, Mathias Kraus, Elliott Ash and Markus Leippold	1913
<i>LoRAMoE: Alleviating World Knowledge Forgetting in Large Language Models via MoE-Style Plugin</i>	
Shihan Dou, Enyu Zhou, Yan Liu, Songyang Gao, Wei Shen, Limao Xiong, Yuhao Zhou, Xiao Wang, Zhiheng Xi, Xiaoran Fan, Shiliang Pu, Jiang Zhu, Rui Zheng, Tao Gui, Qi Zhang and Xuanjing Huang	1932
<i>Self-Alignment for Factuality: Mitigating Hallucinations in LLMs via Self-Evaluation</i>	
Xiaoying Zhang, Baolin Peng, Ye Tian, Jingyan Zhou, Lifeng Jin, Linfeng Song, Haitao Mi and Helen M. Meng	1946
<i>M-RAG: Reinforcing Large Language Model Performance through Retrieval-Augmented Generation with Multiple Partitions</i>	
Zheng Wang, Shu Xian Teo, Jieer Ouyang, Yongjun xu and Wei Shi	1966
<i>AIR-Bench: Benchmarking Large Audio-Language Models via Generative Comprehension</i>	
Qian Yang, Jin Xu, Wenrui Liu, Yunfei Chu, Ziyue Jiang, Xiaohuan Zhou, Yichong Leng, Yuanjun Lv, Zhou Zhao, Chang Zhou and Jingren Zhou	1979

<i>Navigating the Metrics Maze: Reconciling Score Magnitudes and Accuracies</i> Tom Kocmi, Vilém Zouhar, Christian Federmann and Matt Post	1999
<i>ValueBench: Towards Comprehensively Evaluating Value Orientations and Understanding of Large Language Models</i> Yuanyi Ren, Haoran Ye, Hanjun Fang, Xin Zhang and Guojie Song	2015
<i>DM-BLI: Dynamic Multiple Subspaces Alignment for Unsupervised Bilingual Lexicon Induction</i> Ling Hu and Yuemei Xu	2041
<i>SparseFit: Few-shot Prompting with Sparse Fine-tuning for Jointly Generating Predictions and Natural Language Explanations</i> Jesus Solano, Mardhiyah Sanni, Oana-Maria Camburu and Pasquale Minervini	2053
<i>Handling Ambiguity in Emotion: From Out-of-Domain Detection to Distribution Estimation</i> Wen Wu, Bo Li, Chao Zhang, Chung-Cheng Chiu, Qiuqia Li, Junwen Bai, Tara N Sainath and Phil Woodland	2078
<i>REANO: Optimising Retrieval-Augmented Reader Models through Knowledge Graph Generation</i> Jinyuan Fang, Zaiqiao Meng and Craig MacDonald	2094
<i>Learning Disentangled Semantic Spaces of Explanations via Invertible Neural Networks</i> Yingji Zhang, Danilo Carvalho and Andre Freitas	2113
<i>MoPS: Modular Story Premise Synthesis for Open-Ended Automatic Story Generation</i> Yan Ma, Yu Qiao and Pengfei Liu	2135
<i>Open-Set Semi-Supervised Text Classification via Adversarial Disagreement Maximization</i> Junfan Chen, Richong Zhang, Junchi Chen and Chunming Hu	2170
<i>ToolSword: Unveiling Safety Issues of Large Language Models in Tool Learning Across Three Stages</i> Junjie Ye, Sixian Li, Guanyu Li, Caishuang Huang, Songyang Gao, Yilong Wu, Qi Zhang, Tao Gui and Xuanjing Huang	2181
<i>A synthetic data approach for domain generalization of NLI models</i> Mohammad Javad Hosseini, Andrey Petrov, Alex Fabrikant and Annie Louis	2212
<i>Enhancing Contrastive Learning with Noise-Guided Attack: Towards Continual Relation Extraction in the Wild</i> Ting Wu, Jingyi Liu, Rui Zheng, Tao Gui, Qi Zhang and Xuanjing Huang	2227
<i>LRQuant: Learnable and Robust Post-Training Quantization for Large Language Models</i> Jiaqi Zhao, Miao Zhang, Chao Zeng, Ming Wang, Xuebo Liu and Liqiang Nie	2240
<i>VariErr NLI: Separating Annotation Error from Human Label Variation</i> Leon Weber-Genzel, Siyao Peng, Marie-Catherine De Marneffe and Barbara Plank	2256
<i>Benchmarking Knowledge Boundary for Large Language Models: A Different Perspective on Model Evaluation</i> Xunjian Yin, Xu Zhang, Jie Ruan and Xiaojun Wan	2270
<i>ListT5: Listwise Reranking with Fusion-in-Decoder Improves Zero-shot Retrieval</i> Soyoung Yoon, Eunbi Choi, Jiyeon Kim, Hyeonju Yun, Yireun Kim and Seung-won Hwang	2287
<i>Exploring the Potential of Large Language Models in Computational Argumentation</i> Guizhen Chen, Liying Cheng, Anh Tuan Luu and Lidong Bing	2309

<i>TaxoLLaMA: WordNet-based Model for Solving Multiple Lexical Semantic Tasks</i>	
Viktor Moskvoretskii, Ekaterina Neminova, Alina Lobanova, Alexander Panchenko and Irina Nikishina	2331
<i>CANDLE: Iterative Conceptualization and Instantiation Distillation from Large Language Models for Commonsense Reasoning</i>	
Weiqi Wang, Tianqing Fang, Chunyang Li, Haochen Shi, Wenxuan Ding, Baixuan Xu, Zhaowei Wang, Jiaxin Bai, Xin Liu, Cheng Jiayang, Chunkit Chan and Yangqiu Song	2351
<i>MEFT: Memory-Efficient Fine-Tuning through Sparse Adapter</i>	
Jitai Hao, Weiwei Sun, Xin Xin, Qi Meng, Zhumin Chen, Pengjie Ren and Zhaochun Ren ..	2375
<i>Surgical Feature-Space Decomposition of LLMs: Why, When and How?</i>	
Arnav Chavan, Nahush Lele and Deepak Gupta	2389
<i>Reasoning in Flux: Enhancing Large Language Models Reasoning through Uncertainty-aware Adaptive Guidance</i>	
Zhangyue Yin, Qiushi Sun, Qipeng Guo, Zhiyuan Zeng, Xiaonan Li, Junqi Dai, Qinyuan Cheng, Xuanjing Huang and Xipeng Qiu	2401
<i>Modality-Aware Integration with Large Language Models for Knowledge-Based Visual Question Answering</i>	
Junnan Dong, Qinggang Zhang, Huachi Zhou, Daochen Zha, Pai Zheng and Xiao Huang ..	2417
<i>Unlocking Data-free Low-bit Quantization with Matrix Decomposition for KV Cache Compression</i>	
Peiyu Liu, Ze-Feng Gao, Xin Zhao, Yipeng Ma, Tao Wang and Ji-Rong Wen	2430
<i>VerifiNER: Verification-augmented NER via Knowledge-grounded Reasoning with Large Language Models</i>	
Seoyeon Kim, Kwangwook Seo, Hyungjoo Chae, Jinyoung Yeo and Dongha Lee	2441
<i>Making Long-Context Language Models Better Multi-Hop Reasoners</i>	
Yanyang Li, Shuo Liang, Michael Lyu and Liwei Wang	2462
<i>TransliCo: A Contrastive Learning Framework to Address the Script Barrier in Multilingual Pretrained Language Models</i>	
Yihong Liu, Chunlan Ma, Haotian Ye and Hinrich Schuetze	2476
<i>Extreme Miscalibration and the Illusion of Adversarial Robustness</i>	
Vyas Raina, Samson Tan, Volkan Cevher, Aditya Rawal, Sheng Zha and George Karypis ..	2500
<i>HyCoRec: Hypergraph-Enhanced Multi-Preference Learning for Alleviating Matthew Effect in Conversational Recommendation</i>	
Yongsen Zheng, Ruilin Xu, Ziliang Chen, Guohua Wang, Mingjie Qian, Jinghui Qin and Liang Lin	2526
<i>Co-training for Low Resource Scientific Natural Language Inference</i>	
Mobashir Sadat and Cornelia Caragea	2538
<i>RLHFPoison: Reward Poisoning Attack for Reinforcement Learning with Human Feedback in Large Language Models</i>	
Jiong Xiao Wang, Junlin Wu, Muhao Chen, Yevgeniy Vorobeychik and Chaowei Xiao	2551
<i>Time is Encoded in the Weights of Finetuned Language Models</i>	
Kai Nylund, Suchin Gururangan and Noah A. Smith	2571

<i>Long-Context Language Modeling with Parallel Context Encoding</i> Howard Yen, Tianyu Gao and Danqi Chen	2588
<i>SirLLM: Streaming Infinite Retentive LLM</i> Yao Yao, Zuchao Li and Hai Zhao	2611
<i>IMO: Greedy Layer-Wise Sparse Representation Learning for Out-of-Distribution Text Classification with Pre-trained Models</i> Tao Feng, Lizhen Qu, Zhuang Li, Haolan Zhan, Yuncheng Hua and Reza Haf	2625
<i>Generative Pretrained Structured Transformers: Unsupervised Syntactic Language Models at Scale</i> Xiang Hu, Pengyu Ji, Qingyang Zhu, Wei Wu and Kewei Tu	2640
<i>MELA: Multilingual Evaluation of Linguistic Acceptability</i> Ziyin Zhang, Yikang Liu, Weifang Huang, Junyu Mao, Rui Wang and Hai Hu	2658
<i>CopyNE: Better Contextual ASR by Copying Named Entities</i> Shilin Zhou, Zhenghua Li, Yu Hong, Min Zhang, Zhefeng Wang and Baoxing Huai	2675
<i>Is Table Retrieval a Solved Problem? Exploring Join-Aware Multi-Table Retrieval</i> Peter Baile Chen, Yi Zhang and Dan Roth	2687
<i>Generalizing Conversational Dense Retrieval via LLM-Cognition Data Augmentation</i> Haonan Chen, Zhicheng Dou, Kelong Mao, Jiongnan Liu and Ziliang Zhao	2700
<i>ItD: Large Language Models Can Teach Themselves Induction through Deduction</i> Wangtao Sun, Haotian Xu, Xuanqing Yu, Pei Chen, Shizhu He, Jun Zhao and Kang Liu	2719
<i>MathGenie: Generating Synthetic Data with Question Back-translation for Enhancing Mathematical Reasoning of LLMs</i> Zimu Lu, Aojun Zhou, Houxing Ren, Ke Wang, Weikang Shi, Junting Pan, Mingjie Zhan and Hongsheng Li	2732
<i>Rethinking Task-Oriented Dialogue Systems: From Complex Modularity to Zero-Shot Autonomous Agent</i> Heng-Da Xu, Xian-Ling Mao, Puhai Yang, Fanshu Sun and Heyan Huang	2748
<i>On Context Utilization in Summarization with Large Language Models</i> Mathieu Ravaut, Aixin Sun, Nancy F. Chen and Shafiq Joty	2764
<i>INTERS: Unlocking the Power of Large Language Models in Search with Instruction Tuning</i> Yutao Zhu, Peitian Zhang, Chenghao Zhang, Yifei Chen, Binyu Xie, Zheng Liu, Ji-Rong Wen and Zhicheng Dou	2782
<i>Enhancing In-Context Learning via Implicit Demonstration Augmentation</i> Xiaoling Zhou, Wei Ye, Yidong Wang, Chaoya Jiang, Zhemg Lee, Rui Xie and Shikun Zhang	2810
<i>PRoLoRA: Partial Rotation Empowers More Parameter-Efficient LoRA</i> Sheng Wang, Boyang Xue, Jiacheng Ye, Jiyue Jiang, Liheng Chen, Lingpeng Kong and Chuan Wu	2829
<i>Improving Event Definition Following For Zero-Shot Event Detection</i> Zefan Cai, Po-Nien Kung, Ashima Suvarna, Mingyu Derek Ma, Hritik Bansal, Baobao Chang, P. Jeffrey Brantingham, Wei Wang and Nanyun Peng	2842
<i>Through the MUD: A Multi-Defendant Charge Prediction Benchmark with Linked Crime Elements</i> Xiao Wei, Qi Xu, Hang Yu, Qian Liu and Erik Cambria	2864

<i>Interpreting Conversational Dense Retrieval by Rewriting-Enhanced Inversion of Session Embedding</i> Yiruo Cheng, Kelong Mao and Zhicheng Dou	2879
<i>Stumbling Blocks: Stress Testing the Robustness of Machine-Generated Text Detectors Under Attacks</i> Yichen Wang, Shangbin Feng, Abe Bohan Hou, Xiao Pu, Chao Shen, Xiaoming Liu, Yulia Tsve- tkov and Tianxing He	2894
<i>Training Language Models to Generate Text with Citations via Fine-grained Rewards</i> Chengyu Huang, Zeqiu Wu, Yushi Hu and Wenya Wang	2926
<i>Hypergraph based Understanding for Document Semantic Entity Recognition</i> Qiwei Li, Zuchao Li, Ping Wang, Haojun Ai and Hai Zhao	2950
<i>GSM-Plus: A Comprehensive Benchmark for Evaluating the Robustness of LLMs as Mathematical Problem Solvers</i> Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong and Wei Bi	2961
<i>Synergetic Event Understanding: A Collaborative Approach to Cross-Document Event Coreference Resolution with Large Language Models</i> Qingkai Min, Qipeng Guo, Xiangkun Hu, Songfang Huang, Zheng Zhang and Yue Zhang ..	2985
<i>AutoAct: Automatic Agent Learning from Scratch for QA via Self-Planning</i> Shuofei Qiao, Ningyu Zhang, Runnan Fang, Yujie Luo, Wangchunshu Zhou, Yuchen Eleanor Jiang, Chengfei lv and Huajun Chen	3003
<i>ChronosLex: Time-aware Incremental Training for Temporal Generalization of Legal Classification Tasks</i> Santosh T.y.s.s, Tuan-Quang Vuong and Matthias Grabmair	3022
<i>Virtual Compiler Is All You Need For Assembly Code Search</i> Zeyu Gao, Hao Wang, Yuanda Wang and Chao Zhang	3040
<i>MELoRA: Mini-Ensemble Low-Rank Adapters for Parameter-Efficient Fine-Tuning</i> Pengjie Ren, Chengshun Shi, Shiguang Wu, Mengqi Zhang, Zhaochun Ren, Maarten de Rijke, Zhumin Chen and Jiahuan Pei	3052
<i>Can LLMs Learn from Previous Mistakes? Investigating LLMs' Errors to Boost for Reasoning</i> Yongqi Tong, Dawei Li, Sizhe Wang, Yujia Wang, Fei Teng and Jingbo Shang	3065
<i>An Iterative Associative Memory Model for Empathetic Response Generation</i> Zhou Yang, Zhaochun Ren, Wang Yufeng, Haizhou Sun, Chao Chen, Xiaofei Zhu and Xiangwen Liao	3081
<i>Detoxifying Large Language Models via Knowledge Editing</i> Mengru Wang, Ningyu Zhang, Ziwen Xu, Zekun Xi, Shumin Deng, Yunzhi Yao, Qishen Zhang, Linyi Yang, Jindong Wang and Huajun Chen	3093
<i>LongBench: A Bilingual, Multitask Benchmark for Long Context Understanding</i> Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang and Juanzi Li	3119
<i>Dr.Academy: A Benchmark for Evaluating Questioning Capability in Education for Large Language Models</i> Yuyan Chen, Songzhou Yan, Panjun Liu and Yanghua Xiao	3138
<i>UniBridge: A Unified Approach to Cross-Lingual Transfer Learning for Low-Resource Languages</i> Trinh Pham, Khoi M. Le and Anh Tuan Luu	3168

<i>VISTA: Visualized Text Embedding For Universal Multi-Modal Retrieval</i> Junjie Zhou, Zheng Liu, Shitao Xiao, Bo Zhao and Yongping Xiong	3185
<i>Black-Box Prompt Optimization: Aligning Large Language Models without Model Training</i> Jiale Cheng, Xiao Liu, Kehan Zheng, Pei Ke, Hongning Wang, Yuxiao Dong, Jie Tang and Minlie Huang	3201
<i>Open Ko-LLM Leaderboard: Evaluating Large Language Models in Korean with Ko-H5 Benchmark</i> Chanjun Park, Hyeonwoo Kim, Dahyun Kim, SeongHwan Cho, Sanghoon Kim, Sukyung Lee, Yungi Kim and Hwalsuk Lee	3220
<i>Unified Hallucination Detection for Multimodal Large Language Models</i> Xiang Chen, Chenxi Wang, Yida Xue, Ningyu Zhang, Xiaoyan Yang, Qiang Li, Yue Shen, Lei Liang, Jinjie GU and Huajun Chen	3235
<i>Empowering Character-level Text Infilling by Eliminating Sub-Tokens</i> Houxing Ren, Mingjie Zhan, Zhongyuan Wu and Hongsheng Li	3253
<i>Landmark Embedding: A Chunking-Free Embedding Method For Retrieval Augmented Long-Context Large Language Models</i> Kun Luo, Zheng Liu, Shitao Xiao, Tong Zhou, Yubo Chen, Jun Zhao and Kang Liu	3268
<i>GrowOVER: How Can LLMs Adapt to Growing Real-World Knowledge?</i> Dayoon Ko, Jinyoung Kim, Hahyeon Choi and Gunhee Kim	3282
<i>Attribute First, then Generate: Locally-attributable Grounded Text Generation</i> Aviv Slobodkin, Eran Hirsch, Arie Cattan, Tal Schuster and Ido Dagan	3309
<i>T2S-GPT: Dynamic Vector Quantization for Autoregressive Sign Language Production from Text</i> Aoxiong Yin, Haoyuan Li, Kai Shen, Siliang Tang and Yueting Zhuang	3345
<i>OceanGPT: A Large Language Model for Ocean Science Tasks</i> Zhen Bi, Ningyu Zhang, Yida Xue, Yixin Ou, Daxiong Ji, Guozhou Zheng and Huajun Chen	3357
<i>Beyond Memorization: The Challenge of Random Memory Access in Language Models</i> Tongyao Zhu, Qian Liu, Liang Pang, Zhengbao Jiang, Min-Yen Kan and Min Lin	3373
<i>BIPED: Pedagogically Informed Tutoring System for ESL Education</i> Soonwoo Kwon, Sojung Kim, Minju Park, Seunghyun Lee and Kyuseok Kim	3389
<i>Timeline-based Sentence Decomposition with In Context Learning for Temporal Fact Extraction</i> Jianhao Chen, Haoyuan Ouyang, Junyang Ren, Wentao Ding, Wei Hu and Yuzhong Qu	3415
<i>Collaboration or Corporate Capture? Quantifying NLP's Reliance on Industry Artifacts and Contributions</i> Will Aitken, Mohamed Abdalla, Karen Rudie and Catherine Stinson	3433
<i>Prompt Expansion for Adaptive Text-to-Image Generation</i> Siddhartha Datta, Alexander Ku, Deepak Ramachandran and Peter Anderson	3449
<i>Progressively Modality Freezing for Multi-Modal Entity Alignment</i> Yani Huang, Xuefeng Zhang, Richong Zhang, Junfan Chen and Jaein Kim	3477
<i>Llama2Vec: Unsupervised Adaptation of Large Language Models for Dense Retrieval</i> Chaofan Li, Zheng Liu, Shitao Xiao, Yingxia Shao and Defu Lian	3490

<i>Democratizing LLMs for Low-Resource Languages by Leveraging their English Dominant Abilities with Linguistically-Diverse Prompts</i>	
Xuan-Phi Nguyen, Mahani Aljunied, Shafiq Joty and Lidong Bing	3501
<i>Metaphor Understanding Challenge Dataset for LLMs</i>	
Xiaoyu Tong, Rochelle Choenni, Martha Lewis and Ekaterina Shutova	3517
<i>A Multi-Task Embedder For Retrieval Augmented LLMs</i>	
Peitian Zhang, Zheng Liu, Shitao Xiao, Zhicheng Dou and Jian-Yun Nie	3537
<i>Language Models Don't Learn the Physical Manifestation of Language</i>	
Bruce W. Lee and Jaehyuk Lim	3554
<i>What Does the Bot Say? Opportunities and Risks of Large Language Models in Social Media Bot Detection</i>	
Shangbin Feng, Herun Wan, Ningnan Wang, Zhaoxuan Tan, Minnan Luo and Yulia Tsvetkov	3580
<i>Self-Contrast: Better Reflection Through Inconsistent Solving Perspectives</i>	
Wenqi Zhang, Yongliang Shen, Linjuan Wu, Qiuying Peng, Jun Wang, Yueting Zhuang and Weiming Lu	3602
<i>Relying on the Unreliable: The Impact of Language Models' Reluctance to Express Uncertainty</i>	
Kaitlyn Zhou, Jena D. Hwang, Xiang Ren and Maarten Sap	3623
<i>Unity in Diversity: Collaborative Pre-training Across Multimodal Medical Sources</i>	
Xiaochen Wang, Junyu Luo, Jiaqi Wang, Yuan Zhong, Xiaokun Zhang, Yaqing Wang, Parminder Bhatia, Cao Xiao and Fenglong Ma	3644
<i>When Good and Reproducible Results are a Giant with Feet of Clay: The Importance of Software Quality in NLP</i>	
Sara Papi, Marco Gaido, Andrea Pilzer and Matteo Negri	3657
<i>SBAAM! Eliminating Transcript Dependency in Automatic Subtitling</i>	
Marco Gaido, Sara Papi, Matteo Negri, Mauro Cettolo and Luisa Bentivogli	3673
<i>StreamAtt: Direct Streaming Speech-to-Text Translation with Attention-based Audio History Selection</i>	
Sara Papi, Marco Gaido, Matteo Negri and Luisa Bentivogli	3692
<i>ARL2: Aligning Retrievers with Black-box Large Language Models via Self-guided Adaptive Relevance Labeling</i>	
LingXi Zhang, Yue Yu, Kuan Wang and Chao Zhang	3708
<i>Crayon: Customized On-Device LLM via Instant Adapter Blending and Edge-Server Hybrid Inference</i>	
Jihwan Bang, Juntae Lee, Kyuhong Shim, Seunghan Yang and Simyung Chang	3720
<i>FLEUR: An Explainable Reference-Free Evaluation Metric for Image Captioning Using a Large Multimodal Model</i>	
Yebin Lee, Imseong Park and Myungjoo Kang	3732
<i>MentalManip: A Dataset For Fine-grained Analysis of Mental Manipulation in Conversations</i>	
Yuxin Wang, Ivory Yang, Saeed Hassanpour and Soroush Vosoughi	3747
<i>MPCoder: Multi-user Personalized Code Generator with Explicit and Implicit Style Representation Learning</i>	
Zhenlong Dai, Chang Yao, WenKang Han, Yuanying Yuanying, Zhipeng Gao and Jingyuan Chen	3765

<i>DataDreamer: A Tool for Synthetic Data Generation and Reproducible LLM Workflows</i> Ajay Patel, Colin Raffel and Chris Callison-Burch.....	3781
<i>Understanding and Addressing the Under-Translation Problem from the Perspective of Decoding Objective</i> Chenze Shao, Fandong Meng, Jiali Zeng and Jie Zhou.....	3800
<i>Identifying while Learning for Document Event Causality Identification</i> Cheng Liu, Wei Xiang and Bang Wang.....	3815
<i>OlympiadBench: A Challenging Benchmark for Promoting AGI with Olympiad-Level Bilingual Multimodal Scientific Problems</i> Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu and Maosong Sun.....	3828
<i>Insert or Attach: Taxonomy Completion via Box Embedding</i> Wei Xue, Yongliang Shen, Wenqi Ren, Jietian Guo, Shiliang Pu and Weiming Lu.....	3851
<i>Semiparametric Token-Sequence Co-Supervision</i> Hyunji Lee, Doyoung Kim, Jihoon Jun, Se June Joo, Joel Jang, Kyoung-Woon On and Minjoon Seo.....	3864
<i>Instruction Fusion: Advancing Prompt Evolution through Hybridization</i> Weidong Guo, Jiuding Yang, Kaitong Yang, Xiangyang Li, Zhuwei Rao, Yu Xu and Di Niu.....	3883
<i>TimeArena: Shaping Efficient Multitasking Language Agents in a Time-Aware Simulation</i> Yikai Zhang, Siyu Yuan, Caiyu Hu, Kyle Richardson, Yanghua Xiao and Jiangjie Chen....	3894
<i>Exploring Memorization in Fine-tuned Language Models</i> Shenglai Zeng, Yaxin Li, Jie Ren, Yiding Liu, Han Xu, Pengfei He, Yue Xing, Shuaiqiang Wang, Jiliang Tang and Dawei Yin.....	3917
<i>Towards Real-world Scenario: Imbalanced New Intent Discovery</i> Shun Zhang, Yan Chaoran, Jian Yang, Jiaheng Liu, Ying Mo, Jiaqi Bai, Tongliang Li and Zhoujun Li.....	3949
<i>M4GT-Bench: Evaluation Benchmark for Black-Box Machine-Generated Text Detection</i> Yuxia Wang, Jonibek Mansurov, Petar Ivanov, Jinyan Su, Artem Shelmanov, Akim Tsvigun, Osama Mohammed Afzal, Tarek Mahmoud, Giovanni Puccetti, Thomas Arnold, Alham Fikri Aji, Nizar Habash, Iryna Gurevych and Preslav Nakov.....	3964
<i>Instruct Once, Chat Consistently in Multiple Rounds: An Efficient Tuning Framework for Dialogue</i> Jian Wang, Chak Tou Leong, Jiashuo Wang, Dongding Lin, Wenjie Li and Xiaoyong Wei..	3993
<i>SoftDedup: an Efficient Data Reweighting Method for Speeding Up Language Model Pre-training</i> Nan He, Weichen Xiong, Hanwen Liu, Yi Liao, Lei Ding, Kai Zhang, Guohua Tang, Xiao Han and Yang Wei.....	4011
<i>Rule or Story, Which is a Better Commonsense Expression for Talking with Large Language Models?</i> Ning Bian, Xianpei Han, Hongyu Lin, Yaojie Lu, Ben He and Le Sun.....	4023
<i>Learning Global Controller in Latent Space for Parameter-Efficient Fine-Tuning</i> Zeqi Tan, Yongliang Shen, Xiaoxia Cheng, Chang Zong, Wenqi Zhang, Jian Shao, Weiming Lu and Yueting Zhuang.....	4044
<i>CaMML: Context-Aware Multimodal Learner for Large Models</i> Yixin Chen, Shuai Zhang, Boran Han, Tong He and Bo Li.....	4056

<i>MAVEN-ARG: Completing the Puzzle of All-in-One Event Understanding Dataset with Event Argument Annotation</i>	
Xiaozhi Wang, Hao Peng, Yong Guan, Kaisheng Zeng, Jianhui Chen, Lei Hou, Xu Han, Yankai Lin, Zhiyuan Liu, Ruobing Xie, Jie Zhou and Juanzi Li	4072
<i>NPHardEval: Dynamic Benchmark on Reasoning Ability of Large Language Models via Complexity Classes</i>	
Lizhou Fan, Wenyue Hua, Lingyao Li, Haoyang Ling and Yongfeng Zhang	4092
<i>Can Watermarks Survive Translation? On the Cross-lingual Consistency of Text Watermark for Large Language Models</i>	
Zhiwei He, Binglin Zhou, Hongkun Hao, Aiwei Liu, Xing Wang, Zhaopeng Tu, Zhuosheng Zhang and Rui Wang	4115
<i>Multi-Level Feedback Generation with Large Language Models for Empowering Novice Peer Counselors</i>	
Alicja Chaszczewicz, Raj Sanjay Shah, Ryan Louie, Bruce A Arnow, Robert Kraut and Diyi Yang	4130
<i>In-context Mixing (ICM): Code-mixed Prompts for Multilingual LLMs</i>	
Bhavani Shankar, Preethi Jyothi and Pushpak Bhattacharyya	4162
<i>Respond in my Language: Mitigating Language Inconsistency in Response Generation based on Large Language Models</i>	
Liang Zhang, Qin Jin, Haoyang Huang, Dongdong Zhang and Furu Wei	4177
<i>Transferable Embedding Inversion Attack: Uncovering Privacy Risks in Text Embeddings without Model Queries</i>	
Yu-Hsiang Huang, Yuche Tsai, Hsiang Hsiao, Hong-Yi Lin and Shou-De Lin	4193
<i>Enhancing Reinforcement Learning with Label-Sensitive Reward for Natural Language Understanding</i>	
Kuo Liao, Shuang Li, Meng Zhao, Liqun Liu, Mengge Xue, Zhenyu hu, Honglin Han and Chengguo Yin	4206
<i>Intuitive or Dependent? Investigating LLMs' Behavior Style to Conflicting Prompts</i>	
Jiahao Ying, Yixin Cao, Kai Xiong, Long Cui, Yidong He and Yongbin Liu	4221
<i>CoCA: Fusing Position Embedding with Collinear Constrained Attention in Transformers for Long Context Window Extending</i>	
Shiyi Zhu, Jing Ye, Wei Jiang, Siqiao Xue, Qi Zhang, Yifan Wu and Jianguo Li	4247
<i>InfoLossQA: Characterizing and Recovering Information Loss in Text Simplification</i>	
Jan Trienes, Sebastian Antony Joseph, Jörg Schlötterer, Christin Seifert, Kyle Lo, Wei Xu, Byron C Wallace and Junyi Jessy Li	4263
<i>CoGenesis: A Framework Collaborating Large and Small Language Models for Secure Context-Aware Instruction Following</i>	
Kaiyan Zhang, Jianyu Wang, Ermo Hua, Biqing Qi, Ning Ding and Bowen Zhou	4295
<i>DAPR: A Benchmark on Document-Aware Passage Retrieval</i>	
Kexin Wang, Nils Reimers and Iryna Gurevych	4313
<i>Strengthened Symbol Binding Makes Large Language Models Reliable Multiple-Choice Selectors</i>	
Mengge Xue, Zhenyu hu, Liqun Liu, Kuo Liao, Shuang Li, Honglin Han, Meng Zhao and Chengguo Yin	4331

<i>SAC-KG: Exploiting Large Language Models as Skilled Automatic Constructors for Domain Knowledge Graph</i>	Hanzhu Chen, Xu Shen, Qitan Lv, Jie Wang, Xiaoqi Ni and Jieping Ye	4345
<i>Uncertainty-Guided Modal Rebalance for Hateful Memes Detection</i>	Chuanpeng Yang, Yaxin Liu, Fuqing Zhu, Jizhong Han and Songlin Hu	4361
<i>Missci: Reconstructing Fallacies in Misrepresented Science</i>	Max Glockner, Yufang Hou, Preslav Nakov and Iryna Gurevych	4372
<i>Uncovering the Full Potential of Visual Grounding Methods in VQA</i>	Daniel Reich and Tanja Schultz	4406
<i>Small Models, Big Insights: Leveraging Slim Proxy Models To Decide When and What to Retrieve for LLMs</i>	Jiejun Tan, Zhicheng Dou, Yutao Zhu, Peidong Guo, Kun Fang and Ji-Rong Wen	4420
<i>Favi-Score: A Measure for Favoritism in Automated Preference Ratings for Generative AI Evaluation</i>	Pius Von Däniken, Jan Milan Deriu, Don Tuggener and Mark Cieliebak	4437
<i>LLM-based Rewriting of Inappropriate Argumentation using Reinforcement Learning from Machine Feedback</i>	Timon Ziegenbein, Gabriella Skitalinskaya, Alireza Bayat Makou and Henning Wachsmuth	4455
<i>Graph Language Models</i>	Moritz Plenz and Anette Frank	4477
<i>Analyzing Semantic Change through Lexical Replacements</i>	Francesco Periti, Pierluigi Cassotti, Haim Dubossarsky and Nina Tahmasebi	4495
<i>Exploiting Intrinsic Multilateral Logical Rules for Weakly Supervised Natural Language Video Localization</i>	Zhe Xu, Kun Wei, Xu Yang and Cheng Deng	4511
<i>Interpretability of Language Models via Task Spaces</i>	Lucas Weber, Jaap Jumelet, Elia Bruni and Dieuwke Hupkes	4522
<i>Using Synchronic Definitions and Semantic Relations to Classify Semantic Change Types</i>	Pierluigi Cassotti, Stefano De Pascale and Nina Tahmasebi	4539
<i>Factual Confidence of LLMs: on Reliability and Robustness of Current Estimators</i>	Matéo Mahaut, Laura Aina, Paula Czarnowska, Momchil Hardalov, Thomas Müller and Lluís Marquez	4554
<i>StepCoder: Improving Code Generation with Reinforcement Learning from Compiler Feedback</i>	Shihan Dou, Yan Liu, Haoxiang Jia, Enyu Zhou, Limao Xiong, Junjie Shan, Caishuang Huang, Xiao Wang, Xiaoran Fan, Zhiheng Xi, Yuhao Zhou, Tao Ji, Rui Zheng, Qi Zhang, Tao Gui and Xuanjing Huang	4571
<i>One-Shot Learning as Instruction Data Prospector for Large Language Models</i>	Yunshui Li, Binyuan Hui, Xiaobo Xia, Jiayi Yang, Min Yang, Lei Zhang, Shuzheng Si, Ling-Hao Chen, Junhao Liu, Tongliang Liu, Fei Huang and Yongbin Li	4586
<i>Navigating the OverKill in Large Language Models</i>	Chenyu Shi, Xiao Wang, Qiming Ge, Songyang Gao, Xianjun Yang, Tao Gui, Qi Zhang, Xuanjing Huang, Xun Zhao and Dahua Lin	4602

<i>A Chain-of-Thought Is as Strong as Its Weakest Link: A Benchmark for Verifiers of Reasoning Chains</i> Alon Jacovi, Yonatan Bitton, Bernd Bohnet, Jonathan Herzig, Or Honovich, Michael Tseng, Michael Collins, Roei Aharoni and Mor Geva	4615
<i>Re3: A Holistic Framework and Dataset for Modeling Collaborative Document Revision</i> Qian Ruan, Ilia Kuznetsov and Iryna Gurevych	4635
<i>NextLevelBERT: Masked Language Modeling with Higher-Level Representations for Long Documents</i> Tamara Czinczoll, Christoph Hönes, Maximilian Schall and Gerard De Melo	4656
<i>FollowBench: A Multi-level Fine-grained Constraints Following Benchmark for Large Language Models</i> Yuxin Jiang, Yufei Wang, Xingshan Zeng, Wanjun Zhong, Liangyou Li, Fei Mi, Lifeng Shang, Xin Jiang, Qun Liu and Wei Wang	4667
<i>Learning to Edit: Aligning LLMs with Knowledge Editing</i> Yuxin Jiang, Yufei Wang, Chuhan Wu, Wanjun Zhong, Xingshan Zeng, Jiahui Gao, Liangyou Li, Xin Jiang, Lifeng Shang, Ruiming Tang, Qun Liu and Wei Wang	4689
<i>DolphCoder: Echo-Locating Code Large Language Models with Diverse and Multi-Objective Instruction Tuning</i> Yejie Wang, Keqing He, Guanting Dong, Pei Wang, Weihao Zeng, Muxi Diao, Weiran Xu, Jintang Wang, Mengdi Zhang and Xunliang Cai	4706
<i>When Only Time Will Tell: Interpreting How Transformers Process Local Ambiguities Through the Lens of Restart-Incrementality</i> Brielen Madureira, Patrick Kahardipraja and David Schlangen	4722
<i>SpaRC and SpaRP: Spatial Reasoning Characterization and Path Generation for Understanding Spatial Reasoning Capability of Large Language Models</i> Md Imbesat Hassan Rizvi, Xiaodan Zhu and Iryna Gurevych	4750
<i>Planning Like Human: A Dual-process Framework for Dialogue Planning</i> Tao He, Lizi Liao, Yixin Cao, Yuanxing Liu, Ming Liu, Zerui Chen and Bing Qin	4768
<i>Spectral Filters, Dark Signals, and Attention Sinks</i> Nicola Cancedda	4792
<i>DiffuCOMET: Contextual Commonsense Knowledge Diffusion</i> Silin Gao, Mete Ismayilzada, Mengjie Zhao, Hiromi Wakaki, Yuki Mitsufuji and Antoine Bosselut	4809
<i>Systematic Task Exploration with LLMs: A Study in Citation Text Generation</i> Furkan Şahinoç, Ilia Kuznetsov, Yufang Hou and Iryna Gurevych	4832
<i>Limits of Theory of Mind Modelling in Dialogue-Based Collaborative Plan Acquisition</i> Matteo Bortoletto, Constantin Ruhdorfer, Adnen Abdessaied, Lei Shi and Andreas Bulling	4856
<i>Temporal Knowledge Question Answering via Abstract Reasoning Induction</i> Ziyang Chen, Dongfang Li, Xiang Zhao, Baotian Hu and Min Zhang	4872
<i>Who Wrote this Code? Watermarking for Code Generation</i> Taehyun Lee, Seokhee Hong, Jaewoo Ahn, Ilgee Hong, Hwaran Lee, Sangdoon Yun, Jamin Shin and Gunhee Kim	4890
<i>MapCoder: Multi-Agent Code Generation for Competitive Problem Solving</i> Md. Ashraful Islam, Mohammed Eunus Ali and Md Rizwan Parvez	4912

<i>RelayAttention for Efficient Large Language Model Serving with Long System Prompts</i>	
Lei Zhu, Xinjiang Wang, Wayne Zhang and Rynson W. H. Lau	4945
<i>Boosting Language Models Reasoning with Chain-of-Knowledge Prompting</i>	
Jianing Wang, Qiushi Sun, Xiang Li and Ming Gao	4958
<i>Open Grounded Planning: Challenges and Benchmark Construction</i>	
Shiguang Guo, Ziliang Deng, Hongyu Lin, Yaojie Lu, Xianpei Han and Le Sun	4982
<i>LLM Knows Body Language, Too: Translating Speech Voices into Human Gestures</i>	
Chenghao Xu, Guangtao Lyu, Jiexi Yan, Muli Yang and Cheng Deng	5004
<i>QueryAgent: A Reliable and Efficient Reasoning Framework with Environmental Feedback based Self-Correction</i>	
Xiang Huang, Sitao Cheng, Shanshan Huang, Jiayu Shen, Yong Xu, Chaoyun Zhang and Yuzhong Qu	5014
<i>PITA: Prompting Task Interaction for Argumentation Mining</i>	
Yang Sun, Muyi Wang, Jianzhu Bao, Bin Liang, Xiaoyan Zhao, Caihua Yang, Min Yang and Ruifeng Xu	5036
<i>Shifting Attention to Relevance: Towards the Predictive Uncertainty Quantification of Free-Form Large Language Models</i>	
Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura and Kaidi Xu	5050
<i>Babel-ImageNet: Massively Multilingual Evaluation of Vision-and-Language Representations</i>	
Gregor Geige, Radu Timofte and Goran Glavaš	5064
<i>Estimating Agreement by Chance for Sequence Annotation</i>	
Diya Li, Carolyn Rose, Ao Yuan and Chunxiao Zhou	5085
<i>Are Emergent Abilities in Large Language Models just In-Context Learning?</i>	
Sheng Lu, Irina Bigoulaeva, Rachneet Singh Sachdeva, Harish Tayyar Madabushi and Iryna Gurevych	5098
<i>WaveCoder: Widespread And Versatile Enhancement For Code Large Language Models By Instruction Tuning</i>	
Zhaojian Yu, Xin Zhang, Ning Shang, Yangyu Huang, Can Xu, Yishujie Zhao, Wenxiang Hu and Qiufeng Yin	5140
<i>Eliciting Better Multilingual Structured Reasoning from LLMs through Code</i>	
Bryan Li, Tamer Alkhouli, Daniele Bonadiman, Nikolaos Pappas and Saab Mansour	5154
<i>OLIVE: Object Level In-Context Visual Embeddings</i>	
Timothy Ossowski and Junjie Hu	5170
<i>Quantifying Uncertainty in Answers from any Language Model and Enhancing their Trustworthiness</i>	
Jiuhai Chen and Jonas Mueller	5186
<i>Marathon: A Race Through the Realm of Long Context with Large Language Models</i>	
Lei Zhang, Yunshui Li, Ziqiang Liu, Jiayi Yang, Junhao Liu, Longze Chen, Run Luo and Min Yang	5201
<i>Beyond Scaling: Predicting Patent Approval with Domain-specific Fine-grained Claim Dependency Graph</i>	
Xiaochen Kev Gao, Feng Yao, Kewen Zhao, Beilei He, Animesh Kumar, Vish Krishnan and Jingbo Shang	5218

<i>PCAD: Towards ASR-Robust Spoken Language Understanding via Prototype Calibration and Asymmetric Decoupling</i>	
Xianwei Zhuang, Xuxin Cheng, Liming Liang, Yuxin Xie, Zhichang Wang, Zhiqi Huang and Yuexian Zou	5235
<i>Rethinking the Multimodal Correlation of Multimodal Sequential Learning via Generalizable Attentional Results Alignment</i>	
Tao Jin, Wang Lin, Ye Wang, Linjun Li, Xize Cheng and Zhou Zhao	5247
<i>UHGEval: Benchmarking the Hallucination of Chinese Large Language Models via Unconstrained Generation</i>	
Xun Liang, Shichao Song, Simin Niu, Zhiyu li, Feiyu Xiong, Bo Tang, Yezhaohui Wang, Dawei He, Cheng Peng, Zhonghao Wang and Haiying Deng	5266
<i>PreFLMR: Scaling Up Fine-Grained Late-Interaction Multi-modal Retrievers</i>	
Weizhe Lin, Jingbiao Mei, Jinghong Chen and Bill Byrne	5294
<i>Triple-Encoders: Representations That Fire Together, Wire Together</i>	
Justus-Jonas Erker, Florian Mai, Nils Reimers, Gerasimos Spanakis and Iryna Gurevych . . .	5317
<i>Improving Hateful Meme Detection through Retrieval-Guided Contrastive Learning</i>	
Jingbiao Mei, Jinghong Chen, Weizhe Lin, Bill Byrne and Marcus Tomalin	5333
<i>Agent-Pro: Learning to Evolve via Policy-Level Reflection and Optimization</i>	
Wenqi Zhang, Ke Tang, Hai Wu, Mengna Wang, Yongliang Shen, Guiyang Hou, Zeqi Tan, Peng Li, Yueting Zhuang and Weiming Lu	5348
<i>Your Transformer is Secretly Linear</i>	
Anton Razzhigaev, Matvey Mikhailchuk, Elizaveta Goncharova, Nikolai Gerasimenko, Ivan Oseledets, Denis Dimitrov and Andrey Kuznetsov	5376
<i>Noise Correction on Subjective Datasets</i>	
Uthman Jinadu and Yi Ding	5385
<i>Generative Explore-Exploit: Training-free Optimization of Generative Recommender Systems using LLM Optimizers</i>	
Lütfi Kerem Senel, Besnik Fetahu, Davis Yoshida, Zhiyu Chen, Giuseppe Castellucci, Nikhita Vedula, Jason Ingyu Choi and Shervin Malmasi	5396
<i>Instruction-tuned Language Models are Better Knowledge Learners</i>	
Zhengbao Jiang, Zhiqing Sun, Weijia Shi, Pedro Rodriguez, Chunting Zhou, Graham Neubig, Xi Victoria Lin, Wen-tau Yih and Srini Iyer	5421
<i>What Do Language Models Hear? Probing for Auditory Representations in Language Models</i>	
Jerry Ngo and Yoon Kim	5435
<i>Threads of Subtlety: Detecting Machine-Generated Texts Through Discourse Motifs</i>	
Zae Myung Kim, Kwang Hee Lee, Preston Zhu, Vipul Raheja and Dongyeop Kang	5449
<i>Jailbreak Open-Sourced Large Language Models via Enforced Decoding</i>	
Hangfan Zhang, Zhimeng Guo, Huaisheng Zhu, Bochuan Cao, Lu Lin, Jinyuan Jia, Jinghui Chen and Dinghao Wu	5475
<i>NICE: To Optimize In-Context Examples or Not?</i>	
Pragya Srivastava, Satvik Golechha, Amit Deshpande and Amit Sharma	5494

<i>CodeScope: An Execution-based Multilingual Multitask Multidimensional Benchmark for Evaluating LLMs on Code Understanding and Generation</i>	
Weixiang Yan, Haitian Liu, Yunkun Wang, Yunzhe Li, Qian Chen, Wen Wang, Tingyu Lin, Weishan Zhao, Li Zhu, Hari Sundaram and Shuiguang Deng	5511
<i>Digital Socrates: Evaluating LLMs through Explanation Critiques</i>	
Yuling Gu, Oyvind Tafjord and Peter Clark	5559
<i>SafeDecoding: Defending against Jailbreak Attacks via Safety-Aware Decoding</i>	
Zhangchen Xu, Fengqing Jiang, Luyao Niu, Jinyuan Jia, Bill Yuchen Lin and Radha Poovendran	5587
<i>Multi-Task Inference: Can Large Language Models Follow Multiple Instructions at Once?</i>	
Guijin Son, SangWon Baek, Sangdae Nam, Ilgyun Jeong and Seungone Kim	5606
<i>Experiential Co-Learning of Software-Developing Agents</i>	
Chen Qian, Yufan Dang, Jiahao Li, Wei Liu, Zihao Xie, YiFei Wang, Weize Chen, Cheng Yang, Xin Cong, Xiaoyin Che, Zhiyuan Liu and Maosong Sun	5628
<i>Learning Geometry-Aware Representations for New Intent Discovery</i>	
Kai Tang, Junbo Zhao, Xiao Ding, Runze Wu, Lei Feng, Gang Chen and Haobo Wang	5641
<i>Speaker Verification in Agent-generated Conversations</i>	
Yizhe Yang, Palakorn Achananuparp, Heyan Huang, Jing Jiang and Ee-Peng Lim	5655
<i>Benchmarking Data Science Agents</i>	
Yuge Zhang, Qiyang Jiang, XingyuHan XingyuHan, Nan Chen, Yuqing Yang and Kan Ren	5677
<i>Language-Specific Neurons: The Key to Multilingual Capabilities in Large Language Models</i>	
Tianyi Tang, Wenyang Luo, Haoyang Huang, Dongdong Zhang, Xiaolei Wang, Xin Zhao, Furu Wei and Ji-Rong Wen	5701
<i>Forgetting before Learning: Utilizing Parametric Arithmetic for Knowledge Updating in Large Language Models</i>	
Shiwen Ni, Dingwei Chen, Chengming Li, Xiping Hu, Ruifeng Xu and Min Yang	5716
<i>A Deep Dive into the Trade-Offs of Parameter-Efficient Preference Alignment Techniques</i>	
Megh Thakkar, Quentin Fournier, Matthew D Riemer, Pin-Yu Chen, Amal Zouaq, Payel Das and Sarath Chandar	5732
<i>Zero-Shot Cross-Domain Dialogue State Tracking via Dual Low-Rank Adaptation</i>	
Xiang Luo, Zhiwen Tang, Jin Wang and Xuejie Zhang	5746
<i>PRP-Graph: Pairwise Ranking Prompting to LLMs with Graph Aggregation for Effective Text Re-ranking</i>	
Jian Luo, Xuanang Chen, Ben He and Le Sun	5766
<i>RepCodec: A Speech Representation Codec for Speech Tokenization</i>	
Zhichao Huang, Chutong Meng and Tom Ko	5777
<i>GumbelSoft: Diversified Language Model Watermarking via the GumbelMax-trick</i>	
Jiayi Fu, Xuandong Zhao, Ruihan Yang, Yuansen Zhang, Jiangjie Chen and Yanghua Xiao .	5791
<i>Event-Radar: Event-driven Multi-View Learning for Multimodal Fake News Detection</i>	
Zihan Ma, Minnan Luo, Hao Guo, Zhi Zeng, Yiran Hao and Xiang Zhao	5809
<i>Fine-Grained Modeling of Narrative Context: A Coherence Perspective via Retrospective Questions</i>	
Liyang Xu, Jiangnan Li, Mo Yu and Jie Zhou	5822

<i>Stealthy Attack on Large Language Model based Recommendation</i>	
Jinghao Zhang, Yuting Liu, Qiang Liu, Shu Wu, Guibing Guo and Liang Wang	5839
<i>Multi-Dimensional Optimization for Text Summarization via Reinforcement Learning</i>	
Sangwon Ryu, Heejin Do, Yunsu Kim, Gary Lee and Jungseul Ok	5858
<i>Masked Thought: Simply Masking Partial Reasoning Steps Can Improve Mathematical Reasoning Learning of Language Models</i>	
Changyu Chen, Xiting Wang, Ting-En Lin, Ang Lv, Yuchuan Wu, Xin Gao, Ji-Rong Wen, Rui Yan and Yongbin Li	5872
<i>SEER: Facilitating Structured Reasoning and Explanation via Reinforcement Learning</i>	
Guoxin Chen, Kexin Tang, Chao Yang, Fuying Ye, Yu Qiao and Yiming Qian	5901
<i>Towards Robust and Generalized Parameter-Efficient Fine-Tuning for Noisy Label Learning</i>	
Yeanchan Kim, Junho Kim and SangKeun Lee	5922
<i>SparseFlow: Accelerating Transformers by Sparsifying Information Flows</i>	
Yeanchan Kim and SangKeun Lee	5937
<i>ProtT3: Protein-to-Text Generation for Text-based Protein Understanding</i>	
Zhiyuan Liu, An Zhang, Hao Fei, Enzhi Zhang, Xiang Wang, Kenji Kawaguchi and Tat-Seng Chua	5949
<i>KIEval: A Knowledge-grounded Interactive Evaluation Framework for Large Language Models</i>	
Zhuohao Yu, Chang Gao, Wenjin Yao, Yidong Wang, Wei Ye, Jindong Wang, Xing Xie, Yue Zhang and Shikun Zhang	5967
<i>EmoBench: Evaluating the Emotional Intelligence of Large Language Models</i>	
Sahand Sabour, Siyang Liu, Zheyuan Zhang, June M. Liu, Jinfeng Zhou, Alviaona Shiergetya Sunaryo, Tatia M.c. Lee, Rada Mihalcea and Minlie Huang	5986
<i>Are AI-Generated Text Detectors Robust to Adversarial Perturbations?</i>	
Guanhua Huang, Yuchen Zhang, Zhe Li, Yongjian You, Mingze Wang and Zhouwang Yang	6005
<i>FinTextQA: A Dataset for Long-form Financial Question Answering</i>	
Jian Chen, Peilin Zhou, Yining Hua, Loh Ying Xin, Kehui Chen, Ziyuan Li, Bing Zhu and Junwei Liang	6025
<i>On Measuring Faithfulness or Self-consistency of Natural Language Explanations</i>	
Letitia Parcalabescu and Anette Frank	6048
<i>Learning or Self-aligning? Rethinking Instruction Fine-tuning</i>	
Mengjie Ren, Boxi Cao, Hongyu Lin, Cao Liu, Xianpei Han, Ke Zeng, Wan Guanglu, Xunliang Cai and Le Sun	6090
<i>Rethinking the Bounds of LLM Reasoning: Are Multi-Agent Discussions the Key?</i>	
Qineng Wang, Zihao Wang, Ying Su, Hanghang Tong and Yangqiu Song	6106
<i>Soft Knowledge Prompt: Help External Knowledge Become a Better Teacher to Instruct LLM in Knowledge-based VQA</i>	
Qunbo Wang, Ruyi Ji, Tianhao Peng, Wenjun Wu, Zechao Li and Jing Liu	6132
<i>TasTe: Teaching Large Language Models to Translate through Self-Reflection</i>	
Yutong Wang, Jiali Zeng, Xuebo Liu, Fandong Meng, Jie Zhou and Min Zhang	6144

<i>Not All Experts are Equal: Efficient Expert Pruning and Skipping for Mixture-of-Experts Large Language Models</i>	
Xudong Lu, Qi Liu, Yuhui Xu, Aojun Zhou, Siyuan Huang, Bo Zhang, Junchi Yan and Hongsheng Li	6159
<i>UNIMO-G: Unified Image Generation through Multimodal Conditional Diffusion</i>	
Wei Li, Xue Xu, Jiachen Liu and Xinyan Xiao	6173
<i>The Fine-Tuning Paradox: Boosting Translation Quality Without Sacrificing LLM Abilities</i>	
David Stap, Eva Hasler, Bill Byrne, Christof Monz and Ke Tran	6189
<i>Blinded by Generated Contexts: How Language Models Merge Generated and Retrieved Contexts When Knowledge Conflicts?</i>	
Hexiang Tan, Fei Sun, Wanli Yang, Yuanzhuo Wang, Qi Cao and Xueqi Cheng	6207
<i>Unveiling Linguistic Regions in Large Language Models</i>	
Zhihao Zhang, Jun Zhao, Qi Zhang, Tao Gui and Xuanjing Huang	6228
<i>Text-to-Song: Towards Controllable Music Generation Incorporating Vocal and Accompaniment</i>	
Zhiqing Hong, Rongjie Huang, Xize Cheng, Yongqi Wang, Ruiqi Li, Fuming You, Zhou Zhao and Zhimeng Zhang	6248
<i>FastFiD: Improve Inference Efficiency of Open Domain Question Answering via Sentence Selection</i>	
Yufei Huang, Xu Han and Maosong Sun	6262
<i>Discursive Socratic Questioning: Evaluating the Faithfulness of Language Models' Understanding of Discourse Relations</i>	
Yisong Miao, Hongfu Liu, Wenqiang Lei, Nancy F. Chen and Min-Yen Kan	6277
<i>An Open Multilingual System for Scoring Readability of Wikipedia</i>	
Mykola Trokhymovych, Indira Sen and Martin Gerlach	6296
<i>Unlearning Traces the Influential Training Data of Language Models</i>	
Masaru Isonuma and Ivan Titov	6312
<i>Exploring Alignment in Shared Cross-lingual Spaces</i>	
Basel Mousi, Nadir Durrani, Fahim Dalvi, Majd Hawasly and Ahmed Abdelali	6326
<i>Not All Countries Celebrate Thanksgiving: On the Cultural Dominance in Large Language Models</i>	
Wenxuan Wang, Wenxiang Jiao, Jingyuan Huang, Ruyi Dai, Jen-tse Huang, Zhaopeng Tu and Michael Lyu	6349
<i>Self-Evolving GPT: A Lifelong Autonomous Experiential Learner</i>	
Jinglong Gao, Xiao Ding, Yiming Cui, Jianbai Zhao, Hepeng Wang, Ting Liu and Bing Qin	6385
<i>WRP: Weight Recover Prune for Structured Sparsity</i>	
Zhendong Tan, Xingjun Zhang and Zheng Wei	6433
<i>Error-preserving Automatic Speech Recognition of Young English Learners' Language</i>	
Janick Michot, Manuela Hürlimann, Jan Milan Deriu, Luzia Sauer, Katsiaryna Mlynchuk and Mark Cieliebak	6444
<i>DiFiNet: Boundary-Aware Semantic Differentiation and Filtration Network for Nested Named Entity Recognition</i>	
Yuxiang Cai, Qiao Liu, Yanglei Gan, Run Lin, Changlin Li, Xueyi Liu, Da Luo and JiayeYang	6455

<i>Legal Case Retrieval: A Survey of the State of the Art</i> Yi Feng, Chuanyi Li and Vincent Ng	6472
<i>Benchmarking and Improving Compositional Generalization of Multi-aspect Controllable Text Generation</i> Tianqi Zhong, Zhaoyi Li, Quan Wang, Linqi Song, Ying Wei, Defu Lian and Zhendong Mao	6486
<i>LLaMA Pro: Progressive LLaMA with Block Expansion</i> Chengyue Wu, Yukang Gan, Yixiao Ge, Zeyu Lu, Jiahao Wang, Ye Feng, Ying Shan and Ping Luo	6518
<i>Generating Contrastive Narratives Using the Brownian Bridge Process for Narrative Coherence Learning</i> Feiteng Mu and Wenjie Li	6538
<i>A Causal Approach for Counterfactual Reasoning in Narratives</i> Feiteng Mu and Wenjie Li	6556
<i>SIP: Injecting a Structural Inductive Bias into a Seq2Seq Model by Simulation</i> Matthias Lindemann, Alexander Koller and Ivan Titov	6570
<i>The Hidden Space of Transformer Language Adapters</i> Jesujoba Oluwadara Alabi, Marius Mosbach, Matan Eyal, Dietrich Klakow and Mor Geva .	6588
<i>A Ship of Theseus: Curious Cases of Paraphrasing in LLM-Generated Texts</i> Nafis Irtiza Tripto, Saranya Venkatraman, Dominik Macko, Robert Moro, Ivan Srba, Adaku Uchendu, Thai Le and Dongwon Lee	6608
<i>Advancing Large Language Models to Capture Varied Speaking Styles and Respond Properly in Spoken Conversations</i> Guan-Ting Lin, Cheng-Han Chiang and Hung-yi Lee	6626
<i>RetinaQA: A Robust Knowledge Base Question Answering Model for both Answerable and Unanswerable Questions</i> Prayushi Faldu, Indrajit Bhattacharya and Mausam	6643
<i>GroundingGPT: Language Enhanced Multi-modal Grounding Model</i> Zhaowei Li, Qi Xu, Dong Zhang, Hang Song, YiQing Cai, Qi Qi, Ran Zhou, Junting Pan, Zefeng Li, Vu Van Tu, Zhida Huang and Tao Wang	6657
<i>Automated Justification Production for Claim Veracity in Fact Checking: A Survey on Architectures and Approaches</i> Islam Eldifrawi, Shengrui Wang and Amine Trabelsi	6679
<i>Decoupled Vocabulary Learning Enables Zero-Shot Translation from Unseen Languages</i> Carlos Mullov, Quan Pham and Alexander Waibel	6693
<i>SwapMoE: Serving Off-the-shelf MoE-based Large Language Models with Tunable Memory Budget</i> Rui Kong, Yuan Chun Li, Qingtian Feng, Weijun Wang, Xiaozhou Ye, Ye Ouyang, Linghe Kong and Yunxin Liu	6710
<i>PixT3: Pixel-based Table-To-Text Generation</i> Iñigo Alonso, Eneko Agirre and Mirella Lapata	6721
<i>Narrowing the Knowledge Evaluation Gap: Open-Domain Question Answering with Multi-Granularity Answers</i> Gal Yona, Roei Aharoni and Mor Geva	6737

<i>TAMS: Translation-Assisted Morphological Segmentation</i>	
Enora Rice, Ali Marashian, Luke Gessler, Alexis Palmer and Katharina Von Der Wense	6752
<i>XCodeEval: An Execution-based Large Scale Multilingual Multitask Benchmark for Code Understanding, Generation, Translation and Retrieval</i>	
Mohammad Abdullah Matin Khan, M Saiful Bari, Do Xuan Long, Weishi Wang, Md Rizwan Parvez and Shafiq Joty	6766
<i>ProxyQA: An Alternative Framework for Evaluating Long-Form Text Generation with Large Language Models</i>	
Haochen Tan, Zhijiang Guo, Zhan Shi, Lu Xu, Zhili Liu, Yunlong Feng, Xiaoguang Li, Yasheng Wang, Lifeng Shang, Qun Liu and Linqi Song	6806
<i>A Glitch in the Matrix? Locating and Detecting Language Model Grounding with Fakepedia</i>	
Giovanni Monea, Maxime Peyrard, Martin Josifoski, Vishrav Chaudhary, Jason Eisner, Emre Kiciman, Hamid Palangi, Barun Patra and Robert West	6828
<i>Muffin or Chihuahua? Challenging Multimodal Large Language Models with Multipanel VQA</i>	
Yue Fan, Jing Gu, Kaiwen Zhou, Qianqi Yan, Shan Jiang, Ching-Chen Kuo, Yang Zhao, Xinze Guan and Xin Eric Wang	6845
<i>WebVoyager: Building an End-to-End Web Agent with Large Multimodal Models</i>	
Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan and Dong Yu	6864
<i>Translation-based Lexicalization Generation and Lexical Gap Detection: Application to Kinship Terms</i>	
Senyu Li, Bradley Hauer, Ning Shi and Grzegorz Kondrak	6891
<i>Leveraging Machine-Generated Rationales to Facilitate Social Meaning Detection in Conversations</i>	
Ritam Dutt, Zhen Wu, Jiaxin Shi, Divyanshu Sheth, Prakhar Gupta and Carolyn Rose	6901
<i>Robust Frame-Semantic Models with Lexical Unit Trees and Negative Samples</i>	
Jacob Devasier, Yogesh Gurjar and Chengkai Li	6930
<i>Harnessing the Power of Large Language Models for Natural Language to First-Order Logic Translation</i>	
Yuan Yang, Siheng Xiong, Ali Payani, Ehsan Shareghi and Faramarz Fekri	6942
<i>Lightweight reranking for language model generations</i>	
Siddhartha Jain, Xiaofei Ma, Anoop Deoras and Bing Xiang	6960
<i>ARIES: A Corpus of Scientific Paper Edits Made in Response to Peer Reviews</i>	
Mike D’Arcy, Alexis Ross, Erin Bransom, Bailey Kuehl, Jonathan Bragg, Tom Hope and Doug Downey	6985
<i>The Unreasonable Effectiveness of Easy Training Data for Hard Tasks</i>	
Peter Hase, Mohit Bansal, Peter Clark and Sarah Wiegrefe	7002
<i>PLUG: Leveraging Pivot Language in Cross-Lingual Instruction Tuning</i>	
Zhihan Zhang, Dong-Ho Lee, Yuwei Fang, Wenhao Yu, Mengzhao Jia, Meng Jiang and Francesco Barbieri	7025
<i>MIDGARD: Self-Consistency Using Minimum Description Length for Structured Commonsense Reasoning</i>	
Inderjeet Jayakumar Nair and Lu Wang	7047

<i>ReConcile: Round-Table Conference Improves Reasoning via Consensus among Diverse LLMs</i> Justin Chen, Swarnadeep Saha and Mohit Bansal	7066
<i>Mirror: Multiple-perspective Self-Reflection Method for Knowledge-rich Reasoning</i> Hanqi Yan, Qinglin Zhu, Xinyu Wang, Lin Gui and Yulan He	7086
<i>Where Do People Tell Stories Online? Story Detection Across Online Communities</i> Maria Antoniak, Joel Mire, Maarten Sap, Elliott Ash and Andrew Piper	7104
<i>Large Language Models Are No Longer Shallow Parsers</i> Yuanhe Tian, Fei Xia and Yan Song	7131
<i>Dialogue Summarization with Mixture of Experts based on Large Language Models</i> Yuanhe Tian, Fei Xia and Yan Song	7143
<i>ChiMed-GPT: A Chinese Medical Large Language Model with Full Training Regime and Better Alignment to Human Preferences</i> Yuanhe Tian, Ruyi Gan, Yan Song, Jiaying Zhang and Yongdong Zhang	7156
<i>An Investigation of Neuron Activation as a Unified Lens to Explain Chain-of-Thought Eliciting Arithmetic Reasoning of LLMs</i> Daking Rai and Ziyu Yao	7174
<i>Leveraging Large Language Models for Learning Complex Legal Concepts through Storytelling</i> Hang Jiang, Xiajie Zhang, Robert Mahari, Daniel Kessler, Eric Ma, Tal August, Irene Li, Alex Pentland, Yoon Kim, Deb Roy and Jad Kabbara	7194
<i>Intrinsic Task-based Evaluation for Referring Expression Generation</i> Guanyi Chen, Fahime Same and Kees Van Deemter	7220
<i>From Moments to Milestones: Incremental Timeline Summarization Leveraging Large Language Models</i> Qisheng Hu, Geonsik Moon and Hwee Tou Ng	7232
<i>End-to-end Learning of Logical Rules for Enhancing Document-level Relation Extraction</i> Kunxun Qi, Jianfeng Du and Hai Wan	7247
<i>Can We Achieve High-quality Direct Speech-to-Speech Translation without Parallel Speech Data?</i> Qingkai Fang, Shaolei Zhang, Zhengrui Ma, Min Zhang and Yang Feng	7264
<i>Enhancing EEG-to-Text Decoding through Transferable Representations from Pre-trained Contrastive EEG-Text Masked Autoencoder</i> Jiaqi Wang, Zhenxi Song, Zhengyu Ma, Xipeng Qiu, Min Zhang and Zhiguo Zhang	7278
<i>CQIL: Inference Latency Optimization with Concurrent Computation of Quasi-Independent Layers</i> Longwei Zou, Qingyang Wang, Han Zhao, Jiangangkong Jiangangkong, YI Yang and Yangdong Deng	7293
<i>Prompt Optimization via Adversarial In-Context Learning</i> Do Xuan Long, Yiran Zhao, Hannah Brown, Yuxi Xie, James Xu Zhao, Nancy F. Chen, Kenji Kawaguchi, Michael Shieh and Junxian He	7308
<i>StreamVoice: Streamable Context-Aware Language Modeling for Real-time Zero-Shot Voice Conversion</i> Zhichao Wang, Yuanzhe Chen, Xinsheng Wang, Lei Xie and Yuping Wang	7328

<i>Generate-then-Ground in Retrieval-Augmented Generation for Multi-hop Question Answering</i> Zhengliang Shi, Shuo Zhang, Weiwei Sun, Shen Gao, Pengjie Ren, Zhumin Chen and Zhaochun Ren.....	7339
<i>Multimodal Contextualized Semantic Parsing from Speech</i> Jordan Voas, David Harwath and Ray Mooney	7354
<i>LaMP: When Large Language Models Meet Personalization</i> Alireza Salemi, Sheshera Mysore, Michael Bendersky and Hamed Zamani	7370
<i>AboutMe: Using Self-Descriptions in Webpages to Document the Effects of English Pretraining Data Filters</i> Li Lucy, Suchin Gururangan, Luca Soldaini, Emma Strubell, David Bamman, Lauren Klein and Jesse Dodge	7393
<i>MT-Bench-101: A Fine-Grained Benchmark for Evaluating Large Language Models in Multi-Turn Dialogues</i> Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo Su, Tiezheng Ge, Bo Zheng and Wanli Ouyang	7421
<i>EFSA: Towards Event-Level Financial Sentiment Analysis</i> Tianyu Chen, Yiming Zhang, Guoxin Yu, Dapeng Zhang, Li Zeng, Qing He and Xiang Ao ..	7455
<i>What Evidence Do Language Models Find Convincing?</i> Alexander Wan, Eric Wallace and Dan Klein.....	7468
<i>Advancement in Graph Understanding: A Multimodal Benchmark and Fine-Tuning of Vision-Language Models</i> Qihang Ai, Jiafan Li, Jincheng Dai, Jianwu Zhou, Lema Liu, Haiyun Jiang and Shuming Shi	7485
<i>LangBridge: Multilingual Reasoning Without Multilingual Supervision</i> Dongkeun Yoon, Joel Jang, Sungdong Kim, Seungone Kim, Sheikh Shafayat and Minjoon Seo	7502
<i>Can LLMs Reason with Rules? Logic Scaffolding for Stress-Testing and Improving LLMs</i> Siyuan Wang, Zhongyu Wei, Yejin Choi and Xiang Ren	7523
<i>SEGO: Sequential Subgoal Optimization for Mathematical Problem-Solving</i> Xueliang Zhao, Xinting Huang, Wei Bi and Lingpeng Kong	7544
<i>Unlocking the Power of Large Language Models for Entity Alignment</i> Xuhui Jiang, Yinghan Shen, Zhichao Shi, Chengjin Xu, Wei Li, Zixuan Li, Jian Guo, Huawei Shen and Yuanzhuo Wang	7566
<i>Trial and Error: Exploration-Based Trajectory Optimization of LLM Agents</i> Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li and Bill Yuchen Lin	7584
<i>ReFT: Reasoning with Reinforced Fine-Tuning</i> Luong Quoc Trung, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin and Hang Li.....	7601
<i>Cognitive Visual-Language Mapper: Advancing Multimodal Comprehension with Enhanced Visual Knowledge Alignment</i> Yunxin li, Xinyu Chen, Baotian Hu, Haoyuan Shi and Min Zhang.....	7615
<i>FreeCtrl: Constructing Control Centers with Feedforward Layers for Learning-Free Controllable Text Generation</i> Zijian Feng, Hanzhang Zhou, Kezhi Mao and Zixiao Zhu	7627

<i>HD-Eval: Aligning Large Language Model Evaluators Through Hierarchical Criteria Decomposition</i> Yuxuan Liu, Tianchi Yang, Shaohan Huang, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng, Feng Sun and Qi Zhang	7641
<i>Conundrums in Cross-Prompt Automated Essay Scoring: Making Sense of the State of the Art</i> Shengjie Li and Vincent Ng	7661
<i>Angry Men, Sad Women: Large Language Models Reflect Gendered Stereotypes in Emotion Attribution</i> Flor Miriam Plaza-del-Arco, Amanda Cercas Curry, Alba Cercas Curry, Gavin Abercrombie and Dirk Hovy	7682
<i>Label Augmentation for Zero-Shot Hierarchical Text Classification</i> Lorenzo Paletto, Valerio Basile and Roberto Esposito	7697
<i>STICKERCONV: Generating Multimodal Empathetic Responses from Scratch</i> Yiqun Zhang, Fanheng Kong, Peidong Wang, Shuang Sun, SWangLing SWangLing, Shi Feng, Daling Wang, Yifei Zhang and Kaisong Song	7707
<i>EIT: Enhanced Interactive Transformer</i> Tong Zheng, Bei Li, Huiwen Bao, Tong Xiao and JingBo Zhu	7734
<i>MARS: Meaning-Aware Response Scoring for Uncertainty Estimation in Generative LLMs</i> Yavuz Faruk Bakman, Duygu Nur Yaldiz, Baturalp Buyukates, Chenyang Tao, Dimitrios Dimitriadis and Salman Avestimehr	7752
<i>EXAMS-V: A Multi-Discipline Multilingual Multimodal Exam Benchmark for Evaluating Vision Language Models</i> Rocktim Jyoti Das, Simeon Emilov Hristov, Haonan Li, Dimitar Iliyanov Dimitrov, Ivan Koychev and Preslav Nakov	7768
<i>Order-Agnostic Data Augmentation for Few-Shot Named Entity Recognition</i> Huiming Wang, Liying Cheng, Wenxuan Zhang, De Wen Soh and Lidong Bing	7792
<i>Text Embedding Inversion Security for Multilingual Language Models</i> Yiyi Chen, Heather Lent and Johannes Bjerva	7808
<i>Large Language Models are Superpositions of All Characters: Attaining Arbitrary Role-play via Self-Alignment</i> Keming Lu, Bowen Yu, Chang Zhou and Jingren Zhou	7828
<i>PlatoLM: Teaching LLMs in Multi-Round Dialogue via a User Simulator</i> Chuyi Kong, Yaxin Fan, Xiang Wan, Feng Jiang and Benyou Wang	7841
<i>Synthesizing Text-to-SQL Data from Weak and Strong LLMs</i> Jiaxi Yang, Binyuan Hui, Min Yang, Jian Yang, Junyang Lin and Chang Zhou	7864
<i>STRUCTSUM Generation for Faster Text Comprehension</i> Parag Jain, Andreea Marzoca and Francesco Piccinno	7876
<i>Analysing The Impact of Sequence Composition on Language Model Pre-Training</i> Yu Zhao, Yuanbin Qu, Konrad Staniszewski, Szymon Tworowski, Wei Liu, Piotr Miłoś, Yuxiang Wu and Pasquale Minervini	7897
<i>NACL: A General and Effective KV Cache Eviction Framework for LLM at Inference Time</i> Yilong Chen, Guoxia Wang, Junyuan Shang, Shiyao Cui, Zhenyu Zhang, Tingwen Liu, Shuohuan Wang, Yu Sun, Dianhai Yu and Hua Wu	7913

<i>SpikeVoice: High-Quality Text-to-Speech Via Efficient Spiking Neural Network</i>	
Kexin Wang, Jiahong Zhang, Yong Ren, Man Yao, Di Shang, Bo XU and Guoqi Li	7927
<i>Context-aware Difference Distilling for Multi-change Captioning</i>	
Yunbin Tu, Liang Li, Li Su, Zheng-Jun Zha, Chenggang Yan and Qingming Huang	7941
<i>Dataflow-Guided Retrieval Augmentation for Repository-Level Code Completion</i>	
Wei Cheng, Yuhan Wu and Wei Hu	7957
<i>Chain-of-Exemplar: Enhancing Distractor Generation for Multimodal Educational Question Generation</i>	
Haohao Luo, Yang Deng, Ying Shen, See-Kiong Ng and Tat-Seng Chua	7978
<i>LLMEmbed: Rethinking Lightweight LLM’s Genuine Function in Text Classification</i>	
ChunLiu ChunLiu, Hongguang Zhang, Kainan Zhao, Xinghai Ju and Lin Yang	7994
<i>LEMON: Reviving Stronger and Smaller LMs from Larger LMs with Linear Parameter Fusion</i>	
Yilong Chen, Junyuan Shang, Zhenyu Zhang, Shiyao Cui, Tingwen Liu, Shuohuan Wang, Yu Sun and Hua Wu	8005
<i>Speech Sense Disambiguation: Tackling Homophone Ambiguity in End-to-End Speech Translation</i>	
Tengfei Yu, Xuebo Liu, Liang Ding, Kehai Chen, Dacheng Tao and Min Zhang	8020
<i>To be Continuous, or to be Discrete, Those are Bits of Questions</i>	
Yiran Wang and Masao Utiyama	8036
<i>Moûsai: Efficient Text-to-Music Diffusion Models</i>	
Flavio Schneider, Ojasv Kamal, Zhijing Jin and Bernhard Schölkopf	8050
<i>PokeMQA: Programmable knowledge editing for Multi-hop Question Answering</i>	
Hengrui Gu, Kaixiong Zhou, Xiaotian Han, Ninghao Liu, Ruobing Wang and Xin Wang	8069
<i>MemeGuard: An LLM and VLM-based Framework for Advancing Content Moderation via Meme Intervention</i>	
Prince Jha, Raghav Jain, Konika Mandal, Aman Chadha, Sriparna Saha and Pushpak Bhattacharyya	8084
<i>Efficient OCR for Building a Diverse Digital History</i>	
Jacob Carlson, Tom Bryan and Melissa Dell	8105
<i>Acquiring Clean Language Models from Backdoor Poisoned Datasets by Downscaling Frequency Space</i>	
Zongru Wu, Zhuosheng Zhang, Pengzhou Cheng and Gongshen Liu	8116
<i>ANAH: Analytical Annotation of Hallucinations in Large Language Models</i>	
Ziwei Ji, Yuzhe Gu, Wenwei Zhang, Chengqi Lyu, Dahua Lin and Kai Chen	8135
<i>Aligning Large Language Models for Controllable Recommendations</i>	
Wensheng Lu, Jianxun Lian, Wei Zhang, Guanghua Li, Mingyang Zhou, Hao Liao and Xing Xie	8159
<i>Revealing the Parametric Knowledge of Language Models: A Unified Framework for Attribution Methods</i>	
Haeun Yu, Pepa Atanasova and Isabelle Augenstein	8173
<i>Full Parameter Fine-tuning for Large Language Models with Limited Resources</i>	
Kai Lv, Yuqing Yang, Tengxiao Liu, Qipeng Guo and Xipeng Qiu	8187

<i>M³ CoT: A Novel Benchmark for Multi-Domain Multi-step Multi-modal Chain-of-Thought</i> Qiguang Chen, Libo Qin, Jin Zhang, Zhi Chen, Xiao Xu and Wanxiang Che	8199
<i>Long Context is Not Long at All: A Prospector of Long-Dependency Data for Large Language Models</i> Longze Chen, Ziqiang Liu, Wanwei He, Yinhe Zheng, Hao Sun, Yunshui Li, Run Luo and Min Yang	8222
<i>Label-Synchronous Neural Transducer for E2E Simultaneous Speech Translation</i> Keqi Deng and Phil Woodland	8235
<i>Hard Prompts Made Interpretable: Sparse Entropy Regularization for Prompt Tuning with RL</i> Yunseon Choi, Sangmin Bae, Seonghyun Ban, Minchan Jeong, Chuheng Zhang, Lei Song, Li Zhao, Jiang Bian and Kee-Eung Kim	8252
<i>A Modular Approach for Multimodal Summarization of TV Shows</i> Louis Mahon and Mirella Lapata	8272
<i>Think Twice: Perspective-Taking Improves Large Language Models' Theory-of-Mind Capabilities</i> Alex Wilf, Sihyun Shawn Lee, Paul Pu Liang and Louis-Philippe Morency	8292
<i>BizBench: A Quantitative Reasoning Benchmark for Business and Finance</i> Michael Krumdick, Rik Koncel-Kedziorski, Viet Dac Lai, Varshini Reddy, Charles Lovering and Chris Tanner	8309
<i>Direct Metric Optimization for Image Captioning through Reward-Weighted Augmented Data Utilization</i> Takumi Takada, Yuma Suzuki, Hiroki Takushima, Hayato Tanoue, Haruki Sato, Aiswariya Manoj Kumar, Hiroki Nishihara, Takayuki Hori and Kazuya Ueki	8333
<i>Deciphering Hate: Identifying Hateful Memes and Their Targets</i> Eftekhar Hossain, Omar Sharif, Mohammed Moshui Hoque and Sarah Masud Preum	8347
<i>Inducing Systematicity in Transformers by Attending to Structurally Quantized Embeddings</i> Yichen Jiang, Xiang Zhou and Mohit Bansal	8360
<i>Label-Efficient Model Selection for Text Generation</i> Shir Ashury Tahan, Ariel Gera, Benjamin Sznajder, Leshem Choshen, Liat Ein-Dor and Eyal Shnarch	8384
<i>Machine Unlearning of Pre-trained Large Language Models</i> Jin Yao, Eli Chien, Minxin Du, Xinyao Niu, Tianhao Wang, Zezhou Cheng and Xiang Yue	8403
<i>Competition of Mechanisms: Tracing How Language Models Handle Facts and Counterfactuals</i> Francesco Ortu, Zhijing Jin, Diego Doimo, Mrinmaya Sachan, Alberto Cazzaniga and Bernhard Schölkopf	8420
<i>FactPICO: Factuality Evaluation for Plain Language Summarization of Medical Evidence</i> Sebastian Antony Joseph, Lily Chen, Jan Trienes, Hannah Louisa Göke, Monika Coers, Wei Xu, Byron C Wallace and Junyi Jessy Li	8437
<i>BvSP: Broad-view Soft Prompting for Few-Shot Aspect Sentiment Quad Prediction</i> Yinhao Bai, Yalan Xie, Xiaoyi Liu, Yuhua Zhao, Zhixin Han, Mengting Hu, Hang Gao and Renhong Cheng	8465
<i>Safety Alignment in NLP Tasks: Weakly Aligned Summarization as an In-Context Attack</i> Yu Fu, Yufei Li, Wen Xiao, Cong Liu and Yue Dong	8483

<i>Speech language models lack important brain-relevant semantics</i> Subba Reddy Oota, Emin Çelik, Fatma Deniz and Mariya Toneva	8503
<i>DocLLM: A Layout-Aware Generative Language Model for Multimodal Document Understanding</i> Dongsheng Wang, Natraj Raman, Mathieu Sibue, Zhiqiang Ma, Petr Babkin, Simerjot Kaur, Yulong Pei, Armineh Nourbakhsh and Xiaomo Liu	8529
<i>Bypassing LLM Watermarks with Color-Aware Substitutions</i> Qilong Wu and Varun Chandrasekaran	8549
<i>Parallel Structures in Pre-training Data Yield In-Context Learning</i> Yanda Chen, Chen Zhao, Zhou Yu, Kathleen McKeown and He He	8582
<i>OpenToM: A Comprehensive Benchmark for Evaluating Theory-of-Mind Reasoning Capabilities of Large Language Models</i> Hainiu Xu, Runcong Zhao, Lixing Zhu, Jinhua Du and Yulan He	8593
<i>Towards Privacy-Aware Sign Language Translation at Scale</i> Phillip Rust, Bowen Shi, Skyler Wang, Necati Cihan Camgoz and Jean Maillard	8624
<i>Arithmetic Control of LLMs for Diverse User Preferences: Directional Preference Alignment with Multi-Objective Rewards</i> Haoxiang Wang, Yong Lin, Wei Xiong, Rui Yang, Shizhe Diao, Shuang Qiu, Han Zhao and Tong Zhang	8642
<i>Towards Real-World Writing Assistance: A Chinese Character Checking Benchmark with Faked and Misspelled Characters</i> Yinghui Li, Zishan Xu, Shaoshen Chen, Haojing Huang, Yangning Li, Shirong Ma, Yong Jiang, Zhongli Li, Qingyu Zhou, Hai-Tao Zheng and Ying Shen	8656
<i>RAVEL: Evaluating Interpretability Methods on Disentangling Language Model Representations</i> Jing Huang, Zhengxuan Wu, Christopher Potts, Mor Geva and Atticus Geiger	8669
<i>Large Language Models as Zero-shot Dialogue State Tracker through Function Calling</i> Zekun Li, Zhiyu Chen, Mike Ross, Patrick Huber, Seungwhan Moon, Zhaojiang Lin, Xin Luna Dong, Adithya Sagar, Xifeng Yan and Paul A. Crook	8688
<i>Faithful Chart Summarization with ChaTS-Pi</i> Syrine Krichene, Francesco Piccinno, Fangyu Liu and Julian Martin Eisenschlos	8705
<i>Enhancing Dialogue State Tracking Models through LLM-backed User-Agents Simulation</i> Cheng Niu, Xingguang Wang, Xuxin Cheng, Juntong Song and Tong Zhang	8724
<i>MetaSumPerceiver: Multimodal Multi-Document Evidence Summarization for Fact-Checking</i> Ting-Chih Chen, Chia-Wei Tang and Chris Thomas	8742
<i>KnowCoder: Coding Structured Knowledge into LLMs for Universal Information Extraction</i> Zixuan Li, Yutao Zeng, Yuxin Zuo, Weicheng Ren, Wenxuan Liu, Miao Su, Yucan Guo, Yantao Liu, Lixiang Lixiang, Zhilei Hu, Long Bai, Wei Li, Yidan Liu, Pan Yang, Xiaolong Jin, Jiafeng Guo and Xueqi Cheng	8758
<i>ERA-CoT: Improving Chain-of-Thought through Entity Relationship Analysis</i> Yanming Liu, Xinyue Peng, Tianyu Du, Jianwei Yin, Weihao Liu and Xuhong Zhang	8780
<i>On the Multi-turn Instruction Following for Conversational Web Agents</i> Yang Deng, Xuan Zhang, Wenxuan Zhang, Yifei Yuan, See-Kiong Ng and Tat-Seng Chua	8795

<i>Mobile-Bench: An Evaluation Benchmark for LLM-based Mobile Agents</i>	
Shihan Deng, Weikai Xu, Hongda Sun, Wei Liu, Tao Tan, Liujianfeng Liu, Ang Li, Jian Luan, Bin Wang, Rui Yan and Shuo Shang	8813
<i>MC²: Towards Transparent and Culturally-Aware NLP for Minority Languages in China</i>	
Chen Zhang, Mingxu Tao, Quzhe Huang, Jiuheg Lin, Zhibin Chen and Yansong Feng	8832
<i>Decoder-only Streaming Transformer for Simultaneous Translation</i>	
Shoutao Guo, Shaolei Zhang and Yang Feng	8851
<i>Defending Large Language Models Against Jailbreaking Attacks Through Goal Prioritization</i>	
Zhexin Zhang, Junxiao Yang, Pei Ke, Fei Mi, Hongning Wang and Minlie Huang	8865
<i>I am a Strange Dataset: Metalinguistic Tests for Language Models</i>	
Tristan Thrush, Jared Moore, Miguel Monares, Christopher Potts and Douwe Kiela	8888
<i>TruthX: Alleviating Hallucinations by Editing Large Language Models in Truthful Space</i>	
Shaolei Zhang, Tian Yu and Yang Feng	8908
<i>ProtLLM: An Interleaved Protein-Language LLM with Protein-as-Word Pre-Training</i>	
Le Zhuo, Zewen Chi, Minghao Xu, Heyan Huang, Jianan Zhao, Heqi Zheng, Conghui He, Xian-Ling Mao and Wentao Zhang	8950
<i>StreamSpeech: Simultaneous Speech-to-Speech Translation with Multi-task Learning</i>	
Shaolei Zhang, Qingkai Fang, Shoutao Guo, Zhengrui Ma, Min Zhang and Yang Feng	8964
<i>Investigating Multi-Hop Factual Shortcuts in Knowledge Editing of Large Language Models</i>	
Tianjie Ju, Yijin Chen, Xinwei Yuan, Zhuosheng Zhang, Wei Du, Yubin Zheng and Gongshen Liu	8987
<i>Why Don't Prompt-Based Fairness Metrics Correlate?</i>	
Abdelrahman Zayed, Goncalo Mordido, Ioana Baldini and Sarath Chandar	9002
<i>NaijaHate: Evaluating Hate Speech Detection on Nigerian Twitter Using Representative Data</i>	
Manuel Tonneau, Pedro Vitor Quinta De Castro, Karim Lasri, Ibrahim Sambo Farouq, Lakshmi Subramanian, Victor Orozco-Olvera and Samuel Fraiberger	9020
<i>M³AV: A Multimodal, Multigenre, and Multipurpose Audio-Visual Academic Lecture Dataset</i>	
Zhe Chen, Heyang Liu, Wenyi Yu, Guangzhi Sun, Hongcheng Liu, Ji Wu, Chao Zhang, Yu Wang and Yanfeng Wang	9041
<i>Mitigating Biases for Instruction-following Language Models via Bias Neurons Elimination</i>	
Nakyeong Yang, Taegwan Kang, Stanley Jungkyu Choi, Honglak Lee and Kyomin Jung	9061
<i>Domain Adaptation for Subjective Induction Questions Answering on Products by Adversarial Disentangled Learning</i>	
Yufeng Zhang, Jianxing Yu, Yanghui Rao, Libin Zheng, Qinliang Su, Huaijie Zhu and Jian Yin	9074
<i>Revisiting Demonstration Selection Strategies in In-Context Learning</i>	
Keqin Peng, Liang Ding, Yancheng Yuan, Xuebo Liu, Min Zhang, Yuanxin Ouyang and Dacheng Tao	9090
<i>Multimodal Table Understanding</i>	
Mingyu Zheng, Xinwei Feng, Qingyi Si, Qiaoqiao She, Zheng Lin, Wenbin Jiang and Weiping Wang	9102

<i>Ex³: Automatic Novel Writing by Extracting, Excelsior and Expanding</i>	
Huang Lei, Jiaming Guo, Guanhua He, Xishan Zhang, Rui Zhang, Shaohui Peng, Shaoli Liu and Tianshi Chen	9125
<i>Few-shot Transfer Learning for Knowledge Base Question Answering: Fusing Supervised Models with In-Context Learning</i>	
Mayur Patidar, Riya Sawhney, Avinash Kumar Singh, Biswajit Chatterjee, Mausam . and Indrajit Bhattacharya	9147
<i>WatME: Towards Lossless Watermarking Through Lexical Redundancy</i>	
Liang Chen, Yatao Bian, Yang Deng, Deng Cai, Shuaiyi Li, Peilin Zhao and Kam-Fai Wong	9166
<i>Text-like Encoding of Collaborative Information in Large Language Models for Recommendation</i>	
Yang Zhang, Keqin Bao, Ming Yan, Wenjie Wang, Fuli Feng and Xiangnan He	9181
<i>MM-SAP: A Comprehensive Benchmark for Assessing Self-Awareness of Multimodal Large Language Models in Perception</i>	
Yuhao Wang, Yusheng Liao, Heyang Liu, Hongcheng Liu, Yanfeng Wang and Yu Wang ...	9192
<i>Focus on Your Question! Interpreting and Mitigating Toxic CoT Problems in Commonsense Reasoning</i>	
Jiachun Li, Pengfei Cao, Chenhao Wang, Zhuoran Jin, Yubo Chen, Daojian Zeng, Kang Liu and Jun Zhao	9206
<i>Multi-Aspect Controllable Text Generation with Disentangled Counterfactual Augmentation</i>	
Yi Liu, Xiangyu Liu, Xiangrong Zhu and Wei Hu	9231
<i>Reward-based Input Construction for Cross-document Relation Extraction</i>	
Byeonghu Na, Suhyeon Jo, Yeongmin Kim and Il-chul Moon	9254
<i>Hyperspherical Multi-Prototype with Optimal Transport for Event Argument Extraction</i>	
Guangjun Zhang, Hu Zhang, YuJie Wang, Ru Li, Hongye Tan and Jiye Liang	9271
<i>Understanding Retrieval Robustness for Retrieval-augmented Image Captioning</i>	
Wenyan Li, Jiaang Li, Rita Ramos, Raphael Tang and Desmond Elliott	9285
<i>Semi-Supervised Spoken Language Glossification</i>	
Huijie Yao, Wengang Zhou, Hao Zhou and Houqiang Li	9300
<i>SeeClick: Harnessing GUI Grounding for Advanced Visual GUI Agents</i>	
Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Li YanTao, Jianbing Zhang and Zhiyong Wu	9313
<i>InterrogateLLM: Zero-Resource Hallucination Detection in LLM-Generated Answers</i>	
Yakir Yehuda, Itzik Malkiel, Oren Barkan, Jonathan Weill, Royi Ronen and Noam Koenigstein	9333
<i>F-Eval: Assessing Fundamental Abilities with Refined Evaluation Methods</i>	
Yu Sun, Keyuchen Keyuchen, Shujie Wang, Peiji Li, Qipeng Guo, Hang Yan, Xipeng Qiu, Xuanjing Huang and Dahua Lin	9348
<i>Comparing Inferential Strategies of Humans and Large Language Models in Deductive Reasoning</i>	
Philipp Mondorf and Barbara Plank	9370
<i>Whose Preferences? Differences in Fairness Preferences and Their Impact on the Fairness of AI Utilizing Human Feedback</i>	
Maria Emilia Agis Lerner, Florian E. Dorner, Elliott Ash and Naman Goel	9403

<i>Math-Shepherd: Verify and Reinforce LLMs Step-by-step without Human Annotations</i>	
Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu and Zhifang Sui	9426
<i>Large Language Models are not Fair Evaluators</i>	
Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu and Zhifang Sui	9440
<i>Improving Large Language Models in Event Relation Logical Prediction</i>	
Meiqi Chen, Yubo Ma, Kaitao Song, Yixin Cao, Yan Zhang and Dongsheng Li	9451
<i>Synchronized Video Storytelling: Generating Video Narrations with Structured Storyline</i>	
Dingyi Yang, Chunru Zhan, Ziheng Wang, Biao Wang, Tiezheng Ge, Bo Zheng and Qin Jin	9479
<i>Fine-Grained Image-Text Alignment in Medical Imaging Enables Explainable Cyclic Image-Report Generation</i>	
Wenting Chen, Linlin Shen, Jingyang Lin, Jiebo Luo, Xiang Li and Yixuan Yuan	9494
<i>T-Eval: Evaluating the Tool Utilization Capability of Large Language Models Step by Step</i>	
Zehui Chen, Weihua Du, Wenwei Zhang, Kuikun Liu, Jiangning Liu, Miao Zheng, Jingming Zhuo, Songyang Zhang, Dahua Lin, Kai Chen and Feng Zhao	9510
<i>Are LLM-based Evaluators Confusing NLG Quality Criteria?</i>	
Xinyu Hu, Mingqi Gao, Sen Hu, Yang Zhang, Yicheng Chen, Teng XU and Xiaojun Wan . .	9530
<i>Synergistic Interplay between Search and Large Language Models for Information Retrieval</i>	
Jiazhan Feng, Chongyang Tao, Xiubo Geng, Tao Shen, Can Xu, Guodong Long, Dongyan Zhao and Daxin Jiang	9571
<i>Linear Transformers with Learnable Kernel Functions are Better In-Context Models</i>	
Yaroslav Aksenov, Nikita Balagansky, Sofia Maria Lo Cicero Vaina, Boris Shaposhnikov, Alexey Gorbатовski and Daniil Gavrilov	9584
<i>Temperature-scaling surprisal estimates improve fit to human reading times – but does it do so for the “right reasons”?</i>	
Tong Liu, Iza Škrjanec and Vera Demberg	9598
<i>Beyond Recognising Entailment: Formalising Natural Language Inference from an Argumentative Perspective</i>	
Ameer Saadat-Yazdi and Nadin Kökciyan	9620
<i>AnyGPT: Unified Multimodal LLM with Discrete Sequence Modeling</i>	
Jun Zhan, Junqi Dai, Jiasheng Ye, Yunhua Zhou, Dong Zhang, Zhigeng Liu, Xin Zhang, Ruibin Yuan, Ge Zhang, Linyang Li, Hang Yan, Jie Fu, Tao Gui, Tianxiang Sun, Yu-Gang Jiang and Xipeng Qiu	9637
<i>CofiPara: A Coarse-to-fine Paradigm for Multimodal Sarcasm Target Identification with Large Multimodal Models</i>	
Zixin Chen, Hongzhan Lin, Ziyang Luo, Mingfei Cheng, Jing Ma and Guang Chen	9663
<i>Direct Large Language Model Alignment Through Self-Rewarding Contrastive Prompt Distillation</i>	
Aiwei Liu, Haoping Bai, Zhiyun Lu, Xiang Kong, Xiaoming Simon Wang, Jiulong Shan, Meng Cao and Lijie Wen	9688
<i>Diffusion Lens: Interpreting Text Encoders in Text-to-Image Pipelines</i>	
Michael Toker, Hadas Orgad, Mor Ventura, Dana Arad and Yonatan Belinkov	9713

<i>Parrot: Enhancing Multi-Turn Instruction Following for Large Language Models</i>	
Yuchong Sun, Che Liu, Kun Zhou, Jinwen Huang, Ruihua Song, Xin Zhao, Fuzheng Zhang, Di Zhang and Kun Gai	9729
<i>Robust Singing Voice Transcription Serves Synthesis</i>	
Ruiqi Li, Yu Zhang, Yongqi Wang, Zhiqing Hong, Rongjie Huang and Zhou Zhao	9751
<i>VulLibGen: Generating Names of Vulnerability-Affected Packages via a Large Language Model</i>	
Tianyu Chen, Lin Li, ZhuLiuchuan ZhuLiuchuan, Zongyang Li, Xueqing Liu, Guangtai Liang, Qianxiang Wang and Tao Xie	9767
<i>Self-Modifying State Modeling for Simultaneous Machine Translation</i>	
Donglei Yu, Xiaomian Kang, Yuchen Liu, Yu Zhou and Chengqing Zong	9781
<i>MapGPT: Map-Guided Prompting with Adaptive Path Planning for Vision-and-Language Navigation</i>	
Jiaqi Chen, Bingqian Lin, Ran Xu, Zhenhua Chai, Xiaodan Liang and Kwan-Yee K. Wong ..	9796
<i>BadAgent: Inserting and Activating Backdoor Attacks in LLM Agents</i>	
Yifei Wang, Dizhan Xue, Shengjie Zhang and Shengsheng Qian	9811
<i>DetermLR: Augmenting LLM-based Logical Reasoning from Indeterminacy to Determinacy</i>	
Hongda Sun, Weikai Xu, Wei Liu, Jian Luan, Bin Wang, Shuo Shang, Ji-Rong Wen and Rui Yan	9828
<i>LePaRD: A Large-Scale Dataset of Judicial Citations to Precedent</i>	
Robert Mahari, Dominik Stammach, Elliott Ash and Alex Pentland	9863
<i>To Generate or to Retrieve? On the Effectiveness of Artificial Contexts for Medical Open-Domain Question Answering</i>	
Giacomo Frisoni, Alessio Cocchieri, Alex Presepi, Gianluca Moro and Zaiqiao Meng	9878
<i>MERA: A Comprehensive LLM Evaluation in Russian</i>	
Alena Fenogenova, Artem Chervyakov, Nikita Martynov, Anastasia Kozlova, Maria Tikhonova, Albina Akhmetgareeva, Anton Emelyanov, Denis Shevelev, Pavel Lebedev, Leonid S Sinev, Ulyana Isaeva, Katerina Kolomeytseva, Daniil Moskovskiy, Elizaveta Goncharova, Nikita Savushkin, Polina Mikhailova, Anastasia Minaeva, Denis Dimitrov, Alexander Panchenko and Sergey Markov	9920
<i>SC2: Towards Enhancing Content Preservation and Style Consistency in Long Text Style Transfer</i>	
Jie Zhao, Ziyu Guan, Cai Xu, Wei Zhao and Yue Jiang	9949
<i>Dodo: Dynamic Contextual Compression for Decoder-only LMs</i>	
Guanghui Qin, Corby Rosset, Ethan C. Chau, Nikhil Rao and Benjamin Van Durme	9961
<i>POMP: Probability-driven Meta-graph Prompter for LLMs in Low-resource Unsupervised Neural Machine Translation</i>	
Shilong Pan, Zhiliang Tian, Liang Ding, Haoqi Zheng, Zhen Huang, Zhihua Wen and Dongsheng Li	9976
<i>NewsBench: A Systematic Evaluation Framework for Assessing Editorial Capabilities of Large Language Models in Chinese Journalism</i>	
Miao Li, Ming-Bin Chen, Bo Tang, ShengbinHou ShengbinHou, Pengyu Wang, Haiying Deng, Zhiyu li, Feiyu Xiong, Keming Mao, Cheng Peng and Yi Luo	9993
<i>MAPO: Advancing Multilingual Reasoning through Multilingual-Alignment-as-Preference Optimization</i>	
Shuaijie She, Wei Zou, Shujian Huang, Wenhao Zhu, Xiang Liu, Xiang Geng and Jiajun Chen	10015

<i>Enhancing Noise Robustness of Retrieval-Augmented Language Models with Adaptive Adversarial Training</i>	
Feiteng Fang, Yuelin Bai, Shiwen Ni, Min Yang, Xiaojun Chen and Ruifeng Xu	10028
<i>Predicting Text Preference Via Structured Comparative Reasoning</i>	
Jing Nathan Yan, Tianqi Liu, Justin T Chiu, Jiaming Shen, Zhen Qin, Yue Yu, Charumathi Lakshmanan, Yair Kurzion, Alexander M Rush, Jialu Liu and Michael Bendersky	10040
<i>CoELM: Construction-Enhanced Language Modeling</i>	
Lvxiaowei Xu, Zhilin Gong, Jianhua Dai, Tianxiang Wang, Ming Cai and Jiawei Peng	10061
<i>Uni-Dubbing: Zero-Shot Speech Synthesis from Visual Articulation</i>	
Songju Lei, Xize Cheng, Mengjiao Lyu, Jianqiao Hu, Jintao Tan, Runlin Liu, Lingyu Xiong, Tao Jin, Xiandong Li and Zhou Zhao	10082
<i>On the Impact of Calibration Data in Post-training Quantization and Pruning</i>	
Miles Williams and Nikolaos Aletras	10100
<i>SymKGQA: Few-Shot Knowledge Graph Question Answering via Symbolic Program Generation and Execution</i>	
Prerna Agarwal, Nishant Kumar and Srikanta J. Bedathur	10119
<i>Meta-Task Prompting Elicits Embeddings from Large Language Models</i>	
Yibin Lei, Di Wu, Tianyi Zhou, Tao Shen, Yu Cao, Chongyang Tao and Andrew Yates	10141
<i>A Sentiment Consolidation Framework for Meta-Review Generation</i>	
Miao Li, Jey Han Lau and Eduard Hovy	10158
<i>Revisiting Structured Sentiment Analysis as Latent Dependency Graph Parsing</i>	
Chengjie Zhou, Bobo Li, Hao Fei, Fei Li, Chong Teng and Donghong Ji	10178
<i>OWSM-CTC: An Open Encoder-Only Speech Foundation Model for Speech Recognition, Translation, and Language Identification</i>	
Yifan Peng, Yui Sudo, Muhammad Shakeel and Shinji Watanabe	10192
<i>Do Large Language Models Latently Perform Multi-Hop Reasoning?</i>	
Sohee Yang, Elena Gribovskaya, Nora Kassner, Mor Geva and Sebastian Riedel	10210
<i>MuggleMath: Assessing the Impact of Query and Response Augmentation on Math Reasoning</i>	
Chengpeng Li, Zheng Yuan, Hongyi Yuan, Guanting Dong, Keming Lu, Jiancan Wu, Chuanqi Tan, Xiang Wang and Chang Zhou	10230
<i>Harnessing Toulmin’s theory for zero-shot argument explication</i>	
Ankita Gupta, Ethan Zuckerman and Brendan O’Connor	10259
<i>BinaryAlign: Word Alignment as Binary Sequence Labeling</i>	
Gaetan Lopez Latouche, Marc-André Carbonneau and Benjamin Swanson	10277
<i>Quantifying the Persona Effect in LLM Simulations</i>	
Tiancheng Hu and Nigel Collier	10289
<i>Artifacts or Abduction: How Do LLMs Answer Multiple-Choice Questions Without the Question?</i>	
Nishant Balepur, Abhilasha Ravichander and Rachel Rudinger	10308
<i>Retrieval Augmented Fact Verification by Synthesizing Contrastive Arguments</i>	
Zhenrui Yue, Huimin Zeng, Lanyu Shang, Yifan Liu, Yang Zhang and Dong Wang	10331

<i>SyllabusQA: A Course Logistics Question Answering Dataset</i>	
Nigel Fernandez, Alexander Scarlatos and Andrew Lan	10344
<i>MindMap: Knowledge Graph Prompting Sparks Graph of Thoughts in Large Language Models</i>	
Yilin Wen, Zifeng Wang and Jimeng Sun	10370
<i>AGB-DE: A Corpus for the Automated Legal Assessment of Clauses in German Consumer Contracts</i>	
Daniel Braun and Florian Matthes	10389
<i>Examining the robustness of LLM evaluation to the distributional assumptions of benchmarks</i>	
Charlotte Siska, Katerina Marazopoulou, Melissa Ailem and James Bono	10406
<i>Re-Tuning: Overcoming the Compositionality Limits of Large Language Models with Recursive Tuning</i>	
Eric Pasewark, Kyle Montgomery, Kefei Duan, Dawn Song and Chenguang Wang	10422
<i>Bridging the Preference Gap between Retrievers and LLMs</i>	
Zixuan Ke, Weize Kong, Cheng Li, Mingyang Zhang, Qiaozhu Mei and Michael Bendersky	10438
<i>Large Language Models Can Learn Temporal Reasoning</i>	
Siheng Xiong, Ali Payani, Ramana Rao Kompella and Faramarz Fekri	10452
<i>Learning Relational Decomposition of Queries for Question Answering from Tables</i>	
Raphaël Mouravieff, Benjamin Piwowarski and Sylvain Lamprier	10471
<i>Characterizing Similarities and Divergences in Conversational Tones in Humans and LLMs by Sampling with People</i>	
Dun-Ming Huang, Pol Van Rijn, Ilia Sucholutsky, Raja Marjeh and Nori Jacoby	10486
<i>Pareto Optimal Learning for Estimating Large Language Model Errors</i>	
Theodore Zhao, Mu Wei, J. Samuel Preston and Hoifung Poon	10513
<i>Simul-LLM: A Framework for Exploring High-Quality Simultaneous Translation with Large Language Models</i>	
Victor Agostinelli, Max Wild, Matthew Raffel, Kazi Ahmed Asif Fuad and Lizhong Chen	10530
<i>Defending Against Alignment-Breaking Attacks via Robustly Aligned LLM</i>	
Bochuan Cao, Yuanpu Cao, Lu Lin and Jinghui Chen	10542
<i>Interactive-KBQA: Multi-Turn Interactions for Knowledge Base Question Answering with Large Language Models</i>	
Guanming Xiong, Junwei Bao and Wen Zhao	10561
<i>LLMs in the Imaginarium: Tool Learning through Simulated Trial and Error</i>	
Boshi Wang, Hao Fang, Jason Eisner, Benjamin Van Durme and Yu Su	10583
<i>HyperMoE: Towards Better Mixture of Experts via Transferring Among Experts</i>	
Hao Zhao, Zihan Qiu, Huijia Wu, Zili Wang, Zhaofeng He and Jie Fu	10605
<i>Aligning Large Language Models with Human Preferences through Representation Engineering</i>	
Wenhao Liu, Xiaohua Wang, Muling Wu, Tianlong Li, Changze Lv, Zixuan Ling, Zhu JianHao, Cenyuan Zhang, Xiaoqing Zheng and Xuanjing Huang	10619
<i>CODIS: Benchmarking Context-dependent Visual Comprehension for Multimodal Large Language Models</i>	
Fuwen Luo, Chi Chen, Zihao Wan, Zhaolu Kang, Qidong Yan, Yingjie Li, Xiaolong Wang, Siyu Wang, Ziyue Wang, Xiaoyue Mi, Peng Li, Ning Ma, Maosong Sun and Yang Liu	10639

<i>ARAIDA: Analogical Reasoning-Augmented Interactive Data Annotation</i>	
Chen Huang, Yiping Jin, Ilija Iliovski, Wenqiang Lei and Jiancheng Lv	10660
<i>PolCLIP: A Unified Image-Text Word Sense Disambiguation Model via Generating Multimodal Complementary Representations</i>	
Qihao Yang, Yong Li, Xuelin Wang, Fu Lee Wang and Tianyong Hao	10676
<i>Prompted Aspect Key Point Analysis for Quantitative Review Summarization</i>	
An Quang Tang, Xiuzhen Zhang, Minh Ngoc Dinh and Erik Cambria	10691
<i>Ask Again, Then Fail: Large Language Models' Vacillations in Judgment</i>	
Qiming Xie, Zengzhi Wang, Yi Feng and Rui Xia	10709
<i>CLAMBER: A Benchmark of Identifying and Clarifying Ambiguous Information Needs in Large Language Models</i>	
Tong Zhang, Peixin Qin, Yang Deng, Chen Huang, Wenqiang Lei, Junhong Liu, Dingnan Jin, Hongru Liang and Tat-Seng Chua	10746
<i>Multimodal Reasoning with Multimodal Knowledge Graph</i>	
Junlin Lee, Yequan Wang, Jing Li and Min Zhang	10767
<i>Confidence is not Timeless: Modeling Temporal Validity for Rule-based Temporal Knowledge Graph Forecasting</i>	
Rikui Huang, Wei Wei, Xiaoye Qu, Shengzhe Zhang, Danyang Chen and Yu Cheng	10783
<i>CARE: A Clue-guided Assistant for CSRs to Read User Manuals</i>	
Weihong Du, Jia Liu, Zujie Wen, Dingnan Jin, Hongru Liang and Wenqiang Lei	10795
<i>Enhancing Numerical Reasoning with the Guidance of Reliable Reasoning Processes</i>	
Dingzirui Wang, Longxu Dou, Xuanliang Zhang, Qingfu Zhu and Wanxiang Che	10812
<i>PAGED: A Benchmark for Procedural Graphs Extraction from Documents</i>	
Weihong Du, Wenrui Liao, Hongru Liang and Wenqiang Lei	10829
<i>Navigating the Shadows: Unveiling Effective Disturbances for Modern AI Content Detectors</i>	
Ying Zhou, Ben He and Le Sun	10847
<i>RAGTruth: A Hallucination Corpus for Developing Trustworthy Retrieval-Augmented Language Models</i>	
Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, KaShun Shum, Randy Zhong, Juntong Song and Tong Zhang	10862
<i>The Dawn After the Dark: An Empirical Study on Factuality Hallucination in Large Language Models</i>	
Junyi Li, Jie Chen, Ruiyang Ren, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie and Ji-Rong Wen	10879
<i>Revisiting Knowledge Distillation for Autoregressive Language Models</i>	
Qihuang Zhong, Liang Ding, Li Shen, Juhua Liu, Bo Du and Dacheng Tao	10900
<i>Continual Learning with Semi-supervised Contrastive Distillation for Incremental Neural Machine Translation</i>	
Yunlong Liang, Fandong Meng, Jiaan Wang, Jinan Xu, Yufeng Chen and Jie Zhou	10914
<i>Make-A-Voice: Revisiting Voice Large Language Models as Scalable Multilingual and Multitask Learners</i>	
Rongjie Huang, Chunlei Zhang, Yongqi Wang, Dongchao Yang, Jinchuan Tian, Zhenhui Ye, Luping Liu, Zehan Wang, Ziyue Jiang, Xuankai Chang, Jiatong Shi, Chao Weng, Zhou Zhao and Dong Yu	10929

<i>Chat Vector: A Simple Approach to Equip LLMs with Instruction Following and Model Alignment in New Languages</i>	
Shih-Cheng Huang, Pin-Zu Li, Yu-chi Hsu, Kuang-Ming Chen, Yu Tung Lin, Shih-Kai Hsiao, Richard Tzong-Han Tsai and Hung-yi Lee	10943
<i>PRP: Propagating Universal Perturbations to Attack Large Language Model Guard-Rails</i>	
Neal Mangaokar, Ashish Hooda, Jihye Choi, Shreyas Chandrashekar, Kassem Fawaz, Somesh Jha and Atul Prakash	10960
<i>Hide and Seek in Noise Labels: Noise-Robust Collaborative Active Learning with LLMs-Powered Assistance</i>	
Bo Yuan, Yulin Chen, Yin Zhang and Wei Jiang	10977
<i>CLOMO: Counterfactual Logical Modification with Large Language Models</i>	
Yinya Huang, Ruixin Hong, Hongming Zhang, Wei Shao, Zhicheng Yang, Dong Yu, Changshui Zhang, Xiaodan Liang and Linqi Song	11012
<i>Exploring Hybrid Question Answering via Program-based Prompting</i>	
Qi Shi, Han Cui, Haofeng Wang, Qingfu Zhu, Wanxiang Che and Ting Liu	11035
<i>IndicGenBench: A Multilingual Benchmark to Evaluate Generation Capabilities of LLMs on Indic Languages</i>	
Harman Singh, Nitish Gupta, Shikhar Bharadwaj, Dinesh Tewari and Partha Talukdar	11047
<i>Simple but Effective Compound Geometric Operations for Temporal Knowledge Graph Completion</i>	
Rui Ying, Mengting Hu, Jianfeng Wu, Yalan Xie, Xiaoyi Liu, Zhunheng Wang, Ming Jiang, Hang Gao, Linlin Zhang and Renhong Cheng	11074
<i>Uncertainty Aware Learning for Language Model Alignment</i>	
Yikun Wang, Rui Zheng, Liang Ding, Qi Zhang, Dahua Lin and Dacheng Tao	11087
<i>Interpretable User Satisfaction Estimation for Conversational Systems with Large Language Models</i>	
Ying-Chun Lin, Jennifer Neville, Jack W Stokes, Longqi Yang, Tara Safavi, Mengting Wan, Scott Counts, Siddharth Suri, Reid Andersen, Xiaofeng Xu, Deepak Gupta, Sujay Kumar Jauhar, Xia Song, Georg Buscher, Saurabh Tiwary, Brent Hecht and Jaime Teevan	11100
<i>Fundamental Capabilities of Large Language Models and their Applications in Domain Scenarios: A Survey</i>	
Jiawei Li, Yizhe Yang, Yu Bai, Xiaofeng Zhou, Yinghao Li, Huashan Sun, Yuhang Liu, Xingpeng Si, Yuhao Ye, Yixiao Wu, , Bin Xu, Ren Bowen, Chong Feng, Yang Gao and Heyan Huang	11116
<i>Measuring Political Bias in Large Language Models: What Is Said and How It Is Said</i>	
Yejin Bang, Delong Chen, Nayeon Lee and Pascale Fung	11142
<i>Fortify the Shortest Stave in Attention: Enhancing Context Awareness of Large Language Models for Effective Tool Use</i>	
Yuhan Chen, Ang Lv, Ting-En Lin, Changyu Chen, Yuchuan Wu, Fei Huang, Yongbin Li and Rui Yan	11160
<i>Layer-Condensed KV Cache for Efficient Inference of Large Language Models</i>	
Haoyi Wu and Kewei Tu	11175
<i>Enhancing Multilingual Capabilities of Large Language Models through Self-Distillation from Resource-Rich Languages</i>	
Yuanchi Zhang, Yile Wang, Zijun Liu, Shuo Wang, Xiaolong Wang, Peng Li, Maosong Sun and Yang Liu	11189

<i>Benchmarking Chinese Commonsense Reasoning of LLMs: From Chinese-Specifics to Reasoning-Memorization Correlations</i>	
Jiaxing Sun, Weiquan Huang, Jiang Wu, Chenya Gu, Wei Li, Songyang Zhang, Hang Yan and Conghui He	11205
<i>Browse and Concentrate: Comprehending Multimodal Content via Prior-LLM Context Fusion</i>	
Ziyue Wang, Chi Chen, Yiqi Zhu, Fuwen Luo, Peng Li, Ming Yan, Ji Zhang, Fei Huang, Maosong Sun and Yang Liu	11229
<i>Model Composition for Multimodal Large Language Models</i>	
Chi Chen, Yiyang Du, Zheng Fang, Ziyue Wang, Fuwen Luo, Peng Li, Ming Yan, Ji Zhang, Fei Huang, Maosong Sun and Yang Liu	11246
<i>Draft</i>	
<i>Verify: Lossless Large Language Model Acceleration via Self-Speculative Decoding</i>	
Jun Zhang, Jue Wang, Huan Li, Lidan Shou, Ke Chen, Gang Chen and Sharad Mehrotra ..	11263
<i>Soul-Mix: Enhancing Multimodal Machine Translation with Manifold Mixup</i>	
Xuxin Cheng, Ziyu Yao, Yifei Xin, Hao An, Hongxiang Li, Yaowei Li and Yuexian Zou ..	11283
<i>Measuring Meaning Composition in the Human Brain with Composition Scores from Large Language Models</i>	
Changjiang Gao, Jixing Li, Jiajun Chen and Shujian Huang	11295
<i>MIST: Mutual Information Maximization for Short Text Clustering</i>	
Krissanee Kamthawee, Can Udomcharoenchaikit and Sarana Nutanong	11309
<i>Self-chats from Large Language Models Make Small Emotional Support Chatbot Better</i>	
Zhonghua Zheng, Lizi Liao, Yang Deng, Libo Qin and Liqiang Nie	11325
<i>Improving Conversational Abilities of Quantized Large Language Models via Direct Preference Alignment</i>	
Janghwan Lee, Seongmin Park, Sukjin Hong, Minsoo Kim, Du-Seong Chang and Jungwook Choi	11346
<i>Complex Reasoning over Logical Queries on Commonsense Knowledge Graphs</i>	
Tianqing Fang, Zeming Chen, Yangqiu Song and Antoine Bosselut	11365
<i>An Expert is Worth One Token: Synergizing Multiple Expert LLMs as Generalist via Expert Token Routing</i>	
Ziwei Chai, Guoyin Wang, Jing Su, Tianjie Zhang, Xuanwen Huang, Xuwu Wang, Jingjing Xu, Jianbo Yuan, Hongxia Yang, Fei Wu and Yang Yang	11385
<i>Learning to Plan and Generate Text with Citations</i>	
Constanza Fierro, Reinald Kim Amplayo, Fantine Huot, Nicola De Cao, Joshua Maynez, Shashi Narayan and Mirella Lapata	11397
<i>Exploring Precision and Recall to assess the quality and diversity of LLMs</i>	
Florian Le Bronnec, Alexandre Verine, Benjamin Negrevergne, Yann Chevaleyre and Alexandre Allauzen	11418
<i>Aligning Large Language Models by On-Policy Self-Judgment</i>	
Sangkyu Lee, Sungdong Kim, Ashkan Yousefpour, Minjoon Seo, Kang Min Yoo and Youngjae Yu	11442

<i>IL-TUR: Benchmark for Indian Legal Text Understanding and Reasoning</i>	
Abhinav Joshi, Shounak Paul, Akshat Sharma, Pawan Goyal, Saptarshi Ghosh and Ashutosh Modi	
11460	
<i>JumpCoder: Go Beyond Autoregressive Coder via Online Modification</i>	
Mouxian Chen, Hao Tian, Zhongxin Liu, Xiaoxue Ren and Jianling Sun	11500
<i>Aya Dataset: An Open-Access Collection for Multilingual Instruction Tuning</i>	
Shivalika Singh, Freddie Vargus, Daniel D'souza, Börje F. Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura O'Mahony, Mike Zhang, Ramith Hettiarachchi, Joseph Wilson, Marina Machado, Luisa Souza Moura, Dominik Krzemiński, Hakimeh Fadaei, Irem Ergun, Ifeoma Okoh, Aisha Alaagib, Oshan Ivantha Mudannayake, Zaid Alyafeai, Vu Minh Chien, Sebastian Ruder, Surya Guthikonda, Emad A. Alghamdi, Sebastian Gehrmann, Niklas Muennighoff, Max Bartolo, Julia Kreutzer, Ahmet Üstün, Marzieh Fadaee and Sara Hooker	11521
<i>Language Models can Exploit Cross-Task In-context Learning for Data-Scarce Novel Tasks</i>	
Anwoy Chatterjee, Eshaan Tanwar, Subhabrata Dutta and Tanmoy Chakraborty	11568
<i>Split and Rephrase with Large Language Models</i>	
Antonio David Ponce Martínez, Thierry Etchegoyhen, Jesus Javier Calleja Perez and Harritxu Gete	11588
<i>ChunkAttention: Efficient Self-Attention with Prefix-Aware KV Cache and Two-Phase Partition</i>	
Lu Ye, Ze Tao, Yong Huang and Yang Li	11608
<i>AlignBench: Benchmarking Chinese Alignment of Large Language Models</i>	
Xiao Liu, Xuanyu Lei, Shengyuan Wang, Yue Huang, Andrew Zhuoer Feng, Bosi Wen, Jiale Cheng, Pei Ke, Yifan Xu, Weng Lam Tam, Xiaohan Zhang, Lichao Sun, Xiaotao Gu, Hongning Wang, Jing Zhang, Minlie Huang, Yuxiao Dong and Jie Tang	11621
<i>SAPT: A Shared Attention Framework for Parameter-Efficient Continual Learning of Large Language Models</i>	
Weixiang Zhao, Shilong Wang, Yulin Hu, Yanyan Zhao, Bing Qin, Xuanyu Zhang, Qing Yang, Dongliang Xu and Wanxiang Che	11641
<i>DoRA: Enhancing Parameter-Efficient Fine-Tuning with Dynamic Rank Distribution</i>	
Yulong Mao, Kaiyu Huang, Changhao Guan, Ganglin Bao, Fengran Mo and Jinan Xu	11662
<i>Cross-Lingual Knowledge Editing in Large Language Models</i>	
Jiaan Wang, Yunlong Liang, Zengkui Sun, Yuxuan Cao, Jiarong Xu and Fandong Meng	11676
<i>Argument Mining in Data Scarce Settings: Cross-lingual Transfer and Few-shot Techniques</i>	
Anar Yeginbergen, Maite Oronoz and Rodrigo Agerri	11687
<i>Learning Task Decomposition to Assist Humans in Competitive Programming</i>	
Jiaxin Wen, Ruiqi Zhong, Pei Ke, Zhihong Shao, Hongning Wang and Minlie Huang	11700
<i>An Entropy-based Text Watermarking Detection Method</i>	
Yijian LU, Aiwei Liu, Dianzhi Yu, Jingjing Li and Irwin King	11724
<i>Enhancing Explainable Rating Prediction through Annotated Macro Concepts</i>	
Huachi Zhou, Shuang Zhou, Hao Chen, Ninghao Liu, Fan Yang and Xiao Huang	11736
<i>How to Engage your Readers? Generating Guiding Questions to Promote Active Reading</i>	
Peng Cui, Vilém Zouhar, Xiaoyu Zhang and Mrinmaya Sachan	11749

<i>Less is More: Mitigating Multimodal Hallucination from an EOS Decision Perspective</i> Zihao Yue, Liang Zhang and Qin Jin	11766
<i>Integrate the Essence and Eliminate the Dross: Fine-Grained Self-Consistency for Free-Form Language Generation</i> Xinglin Wang, Yiwei Li, Shaoxiong Feng, Peiwen Yuan, Boyuan Pan, Heda Wang, Yao Hu and Kan Li	11782
<i>More frequent verbs are associated with more diverse valency frames: Efficient principles at the lexicon-grammar interface</i> Siyu Tao, Lucia Donatelli and Michael Hahn	11795
<i>Quantifying Generalizations: Exploring the Divide Between Human and LLMs' Sensitivity to Quantification</i> Claudia Collacciani, Giulia Rambelli and Marianna Bolognesi	11811
<i>Can Large Language Models Interpret Noun-Noun Compounds? A Linguistically-Motivated Study on Lexicalized and Novel Compounds</i> Giulia Rambelli, Emmanuele Chersoni, Claudia Collacciani and Marianna Bolognesi	11823
<i>CharacterEval: A Chinese Benchmark for Role-Playing Conversational Agent Evaluation</i> Quan Tu, Shilong Fan, Zihang Tian, Tianhao Shen, Shuo Shang, Xin Gao and Rui Yan ...	11836
<i>Generative Cross-Modal Retrieval: Memorizing Images in Multimodal Language Models for Retrieval and Beyond</i> Yongqi Li, Wenjie Wang, Leigang Qu, Liqiang Nie, Wenjie Li and Tat-Seng Chua	11851
<i>Self-Training with Pseudo-Label Scorer for Aspect Sentiment Quad Prediction</i> Yice Zhang, Jie Zeng, Weiming Hu, Ziyi Wang, Shiwei Chen and Ruifeng Xu	11862
<i>Learning to Generate Answers with Citations via Factual Consistency Models</i> Rami Aly, Zhiqiang Tang, Samson Tan and George Karypis	11876
<i>Improving Text Embeddings with Large Language Models</i> Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder and Furu Wei .	11897
<i>Self-Training with Direct Preference Optimization Improves Chain-of-Thought Reasoning</i> Tianduo Wang, Shichen Li and Wei Lu	11917
<i>UltraLink: An Open-Source Knowledge-Enhanced Multilingual Supervised Fine-tuning Dataset</i> Haoyu Wang, Shuo Wang, Yukun Yan, Xujia Wang, Zhiyu Yang, Yuzhuang Xu, Zhenghao Liu, Liner Yang, Ning Ding, Xu Han, Zhiyuan Liu and Maosong Sun	11929
<i>Document-level Claim Extraction and Decontextualisation for Fact-Checking</i> Zhenyun Deng, Michael Sejr Schlichtkrull and Andreas Vlachos	11943
<i>PairCFR: Enhancing Model Training on Paired Counterfactually Augmented Data through Contrastive Learning</i> Xiaoqi Qiu, Yongjie Wang, Xu Guo, Zhiwei Zeng, Yu Yue, Yuhong Feng and Chunyan Miao	11955
<i>LLMs Learn Task Heuristics from Demonstrations: A Heuristic-Driven Prompting Strategy for Document-Level Event Argument Extraction</i> Hanzhang Zhou, Junlang Qian, Zijian Feng, LU Hui, Zixiao Zhu and Kezhi Mao	11972

<i>Investigating and Mitigating the Multimodal Hallucination Snowballing in Large Vision-Language Models</i>	
Weihong Zhong, Xiaocheng Feng, Liang Zhao, Qiming Li, Lei Huang, Yuxuan Gu, Weitao Ma, Yuan Xu and Bing Qin	11991
<i>mCoT: Multilingual Instruction Tuning for Reasoning Consistency in Language Models</i>	
Huiyuan Lai and Malvina Nissim	12012
<i>GunStance: Stance Detection for Gun Control and Gun Regulation</i>	
Nikesh Gyawali, Justin Sirbu, Tiberiu Sosea, Sarthak Khanal, Doina Caragea, Traian Rebedea and Cornelia Caragea	12027
<i>Beyond Traditional Benchmarks: Analyzing Behaviors of Open LLMs on Data-to-Text Generation</i>	
Zdeněk Kasner and Ondrej Dusek	12045
<i>Don't Go To Extremes: Revealing the Excessive Sensitivity and Calibration Limitations of LLMs in Implicit Hate Speech Detection</i>	
Min Zhang, Jianfeng He, Taoran Ji and Chang-Tien Lu	12073
<i>Don't Rank, Combine! Combining Machine Translation Hypotheses Using Quality Estimation</i>	
Giorgos Vernikos and Andrei Popescu-Belis	12087
<i>Generating and Evaluating Plausible Explanations for Knowledge Graph Completion</i>	
Antonio Di Mauro, Zhao Xu, Wiem Ben Rim, Timo Szttyler and Carolin Lawrence	12106
<i>One Prompt To Rule Them All: LLMs for Opinion Summary Evaluation</i>	
Tejpal Singh Siledar, Swaroop Nath, Sankara Sri Raghava Ravindra Muddu, Rupasai Rangaraju, Swaprava Nath, Pushpak Bhattacharyya, Suman Banerjee, Amey Patil, Sudhanshu Shekhar Singh, Muthusamy Chelliah and Nikesh Garera	12119
<i>LANDeRMT: Detecting and Routing Language-Aware Neurons for Selectively Finetuning LLMs to Machine Translation</i>	
Shaolin Zhu, Leiyu Pan, Bo Li and Deyi Xiong	12135
<i>A Joint Coreference-Aware Approach to Document-Level Target Sentiment Analysis</i>	
Hongjie Cai, Heqing Ma, Jianfei Yu and Rui Xia	12149
<i>VisDiaHalBench: A Visual Dialogue Benchmark For Diagnosing Hallucination in Large Vision-Language Models</i>	
Qingxing Cao, Junhao Cheng, Xiaodan Liang and Liang Lin	12161
<i>AutoDSL: Automated domain-specific language design for structural representation of procedures with constraints</i>	
Yu-Zhe Shi, Haofei Hou, Zhangqian Bi, Fanxu Meng, Xiang Wei, Lecheng Ruan and Qining Wang	12177
<i>Multipath parsing in the brain</i>	
Berta Franzluebbers, Donald Dunagan, Miloš Stanojević, Jan Buys and John T. Hale	12215
<i>Search-Adaptor: Embedding Customization for Information Retrieval</i>	
Jinsung Yoon, Yanfei Chen, Sercan O Arik and Tomas Pfister	12230
<i>Back to Basics: Revisiting REINFORCE-Style Optimization for Learning from Human Feedback in LLMs</i>	
Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün and Sara Hooker	12248

<i>VIEScore: Towards Explainable Metrics for Conditional Image Synthesis Evaluation</i>	
Max Ku, Dongfu Jiang, Cong Wei, Xiang Yue and Wenhui Chen	12268
<i>Tree Transformer’s Disambiguation Ability of Prepositional Phrase Attachment and Garden Path Effects</i>	
Lingling Zhou, Suzan Verberne and Gijs Wijnholds	12291
<i>Tree-of-Traversals: A Zero-Shot Reasoning Algorithm for Augmenting Black-box Language Models with Knowledge Graphs</i>	
Elan Sopher Markowitz, Anil Ramakrishna, Jwala Dhamala, Ninareh Mehrabi, Charith Peris, Rahul Gupta, Kai-Wei Chang and Aram Galstyan	12302
<i>Structured Tree Alignment for Evaluation of (Speech) Constituency Parsing</i>	
Freda Shi, Kevin Gimpel and Karen Livescu	12320
<i>ViSAGE: A Global-Scale Analysis of Visual Stereotypes in Text-to-Image Generation</i>	
Akshita Jha, Vinodkumar Prabhakaran, Remi Denton, Sarah Laszlo, Shachi Dave, Rida Qadri, Chandan K. Reddy and Sunipa Dev	12333
<i>Transferable and Efficient Non-Factual Content Detection via Probe Training with Offline Consistency Checking</i>	
Xiaokang Zhang, Zijun Yao, Jing Zhang, Kaifeng Yun, Jifan Yu, Juanzi Li and Jie Tang ...	12348
<i>What Do Language Models Learn in Context? The Structured Task Hypothesis.</i>	
Jiaoda Li, Yifan Hou, Mrinmaya Sachan and Ryan Cotterell	12365
<i>Agent Lumos: Unified and Modular Training for Open-Source Language Agents</i>	
Da Yin, Faeze Brahman, Abhilasha Ravichander, Khyathi Chandu, Kai-Wei Chang, Yejin Choi and Bill Yuchen Lin	12380
<i>Investigating Cultural Alignment of Large Language Models</i>	
Badr AlKhamissi, Muhammad ElNokrashy, Mai Alkhamissi and Mona T. Diab	12404
<i>More Victories, Less Cooperation: Assessing Cicero’s Diplomacy Play</i>	
Wichayaporn Wongkamjan, Feng Gu, Yanze Wang, Ulf Hermjakob, Jonathan May, Brandon M. Stewart, Jonathan K. Kummerfeld, Denis Peskoff and Jordan Lee Boyd-Graber	12423
<i>VoiceCraft: Zero-Shot Speech Editing and Text-to-Speech in the Wild</i>	
Puyuan Peng, Po-Yao Huang, Shang-Wen Li, Abdelrahman Mohamed and David Harwath	12442
<i>RAID: A Shared Benchmark for Robust Evaluation of Machine-Generated Text Detectors</i>	
Liam Dugan, Alyssa Hwang, Filip Trhlík, Andrew Zhu, Josh Magnus Ludan, Hainiu Xu, Daphne Ippolito and Chris Callison-Burch	12463
<i>Silent Signals, Loud Impact: LLMs for Word-Sense Disambiguation of Coded Dog Whistles</i>	
Julia Kruk, Michela Marchini, Rijul Magu, Caleb Ziems, David Muchlinski and Diyi Yang	12493
<i>On the Representational Capacity of Neural Language Models with Chain-of-Thought Reasoning</i>	
Franz Nowak, Anej Svete, Alexandra Butoi and Ryan Cotterell	12510
<i>Analyzing LLM Behavior in Dialogue Summarization: Unveiling Circumstantial Hallucination Trends</i>	
Sanjana Ramprasad, Elisa Ferracane and Zachary Chase Lipton	12549
<i>LLM in a flash: Efficient Large Language Model Inference with Limited Memory</i>	
Keivan Alizadeh, Seyed Iman Mirzadeh, Dmitry Belenko, S. Karen Khatamifard, Minsik Cho, Carlo C Del Mundo, Mohammad Rastegari and Mehrdad Farajtabar	12562

<i>Video-ChatGPT: Towards Detailed Video Understanding via Large Vision and Language Models</i> Muhammad Maaz, Hanoona Abdul Rasheed, Salman Khan and Fahad Shahbaz Khan	12585
<i>To Distill or Not to Distill? On the Robustness of Robust Knowledge Distillation</i> Abdul Waheed, Karima Kadaoui and Muhammad Abdul-Mageed	12603
<i>LayerSkip: Enabling Early Exit Inference and Self-Speculative Decoding</i> Mostafa Elhoushi, Akshat Shrivastava, Diana Liskovich, Basil Hosmer, Bram Wasti, Liangzhen Lai, Anas Mahmoud, Bilge Acun, Saurabh Agarwal, Ahmed Roman, Ahmed A Aly, Beidi Chen and Carole-Jean Wu	12622
<i>Classist Tools: Social Class Correlates with Performance in NLP</i> Amanda Cercas Curry, Giuseppe Attanasio, Zeerak Talat and Dirk Hovy	12643
<i>ActionIE: Action Extraction from Scientific Literature with Programming Languages</i> Xianrui Zhong, Yufeng Du, Siru Ouyang, Ming Zhong, Tingfeng Luo, Qirong Ho, Hao Peng, Heng Ji and Jiawei Han	12656
<i>A Community-Centric Perspective for Characterizing and Detecting Anti-Asian Violence-Provoking Speech</i> Gaurav Verma, Rynaa Grover, Jiawei Zhou, Binny Mathew, Jordan Kraemer, Munmun De Choudhury and Srijan Kumar	12672
<i>Retaining Key Information under High Compression Ratios: Query-Guided Compressor for LLMs</i> Zhiwei Cao, Qian Cao, Yu Lu, Ningxin Peng, Luyang Huang, Shanbo Cheng and Jinsong Su	12685
<i>COSMIC: Mutual Information for Task-Agnostic Summarization Evaluation</i> Maxime Darrin, Philippe Formont, Jackie CK Cheung and Pablo Piantanida	12696
<i>EUROPA: A Legal Multilingual Keyphrase Generation Dataset</i> Olivier Salaün, Frédéric Piedboeuf, Guillaume Le Berre, David Alfonso-Hermelo and Philippe Langlais	12718
<i>GLIMPSE: Pragmatically Informative Multi-Document Summarization for Scholarly Reviews</i> Maxime Darrin, Ines Arous, Pablo Piantanida and Jackie CK Cheung	12737
<i>Peacock: A Family of Arabic Multimodal Large Language Models and Benchmarks</i> Fakhraddin Alwajih, El Moatez Billah Nagoudi, Gagan Bhatia, Abdelrahman Mohamed and Muhammad Abdul-Mageed	12753
<i>Generating Coherent Sequences of Visual Illustrations for Real-World Manual Tasks</i> João Bordalo, Vasco Ramos, Rodrigo Valério, Diogo Glória-Silva, Yonatan Bitton, Michal Yarom, Idan Szpektor and Joao Magalhaes	12777
<i>Cheetah: Natural Language Generation for 517 African Languages</i> Ife Adebara, AbdelRahim A. Elmadany and Muhammad Abdul-Mageed	12798
<i>TaPERA: Enhancing Faithfulness and Interpretability in Long-Form Table QA by Content Planning and Execution-based Reasoning</i> Yilun Zhao, Lyuhao Chen, Arman Cohan and Chen Zhao	12824
<i>KnowledgeFMATH: A Knowledge-Intensive Math Reasoning Dataset in Finance Domains</i> Yilun Zhao, Hongjun Liu, Yitao Long, Rui Zhang, Chen Zhao and Arman Cohan	12841

<i>API-BLEND: A Comprehensive Corpora for Training and Benchmarking API LLMs</i>	
Kinjal Basu, Ibrahim Abdelaziz, Subhajit Chaudhury, Soham Dan, Maxwell Crouse, Asim Munawar, Vernon Austel, Sadhana Kumaravel, Vinod Muthusamy, Pavan Kapanipathi and Luis A. Lastras	12859
<i>LoRA-Flow: Dynamic LoRA Fusion for Large Language Models in Generative Tasks</i>	
Hanqing Wang, Bowen Ping, Shuo Wang, Xu Han, Yun Chen, Zhiyuan Liu and Maosong Sun	12871
<i>Harder Task Needs More Experts: Dynamic Routing in MoE Models</i>	
Quzhe Huang, Zhenwei An, Nan Zhuang, Mingxu Tao, Chen Zhang, Yang Jin, Kun Xu, Kun Xu, Liwei Chen, Songfang Huang and Yansong Feng	12883
<i>XLAVS-R: Cross-Lingual Audio-Visual Speech Representation Learning for Noise-Robust Speech Perception</i>	
HyoJung Han, Mohamed Anwar, Juan Pino, Wei-Ning Hsu, Marine Carpuat, Bowen Shi and Changhan Wang	12896
<i>SOTOPIA-: Interactive Learning of Socially Intelligent Language Agents</i>	
Ruiyi Wang, Haofei Yu, Wenxin Sharon Zhang, Zhengyang Qi, Maarten Sap, Yonatan Bisk, Graham Neubig and Hao Zhu	12912
<i>XFT: Unlocking the Power of Code Instruction Tuning by Simply Merging Upcycled Mixture-of-Experts</i>	
Yifeng Ding, Jiawei Liu, Yuxiang Wei and Lingming Zhang	12941
<i>Generalizability of Mixture of Domain-Specific Adapters from the Lens of Signed Weight Directions and its Application to Effective Model Pruning</i>	
Tuc Van Nguyen and Thai Le	12956
<i>Learning to Decode Collaboratively with Multiple Language Models</i>	
Zejiang Shen, Hunter Lang, Bailin Wang, Yoon Kim and David Sontag	12974
<i>DRAGIN: Dynamic Retrieval Augmented Generation based on the Real-time Information Needs of Large Language Models</i>	
Weihang Su, Yichen Tang, Qingyao Ai, Zhijing Wu and Yiqun Liu	12991
<i>Living in the Moment: Can Large Language Models Grasp Co-Temporal Reasoning?</i>	
Zhaochen Su, Juntao Li, Jun Zhang, Tong Zhu, Xiaoye Qu, Pan Zhou, Yan Bowen, Yu Cheng and Min Zhang	13014
<i>CritiqueLLM: Towards an Informative Critique Generation Model for Evaluation of Large Language Model Generation</i>	
Pei Ke, Bosi Wen, Andrew Zhuoer Feng, Xiao Liu, Xuanyu Lei, Jiale Cheng, Shengyuan Wang, Aohan Zeng, Yuxiao Dong, Hongning Wang, Jie Tang and Minlie Huang	13034
<i>LLMArena: Assessing Capabilities of Large Language Models in Dynamic Multi-Agent Environments</i>	
Junzhe Chen, Xuming Hu, Shuodi Liu, Shiyu Huang, Wei-Wei Tu, Zhaofeng He and Lijie Wen	13055
<i>Small But Funny: A Feedback-Driven Approach to Humor Distillation</i>	
Sahithya Ravi, Patrick Huber, Akshat Shrivastava, Vered Shwartz and Arash Einolghozati	13078
<i>Symbol-LLM: Towards Foundational Symbol-centric Interface For Large Language Models</i>	
Fangzhi Xu, Zhiyong Wu, Qiushi Sun, Siyu Ren, Fei Yuan, Shuai Yuan, Qika Lin, Yu Qiao and Jun Liu	13091

<i>From Sights to Insights: Towards Summarization of Multimodal Clinical Documents</i> Akash Ghosh, Mohit Singh Tomar, Abhisek Tiwari, Sriparna Saha, Jatin Avinash Salve and Setu Sinha	13117
<i>When Phrases Meet Probabilities: Enabling Open Relation Extraction with Cooperating Large Language Models</i> Jiaxin Wang, Lingling Zhang, Wee Sun Lee, Yujie Zhong, Liwei Kang and Jun Liu	13130
<i>Effects of diversity incentives on sample diversity and downstream model performance in LLM-based text augmentation</i> Jan Cegin, Branislav Pecher, Jakub Simko, Ivan Srba, Maria Bielikova and Peter Brusilovsky	13148
<i>Beyond Orthography: Automatic Recovery of Short Vowels and Dialectal Sounds in Arabic</i> Yassine El Kheir, Hamdy Mubarak, Ahmed Ali and Shammur Absar Chowdhury	13172
<i>Document-Level Machine Translation with Large-Scale Public Parallel Corpora</i> Proyag Pal, Alexandra Birch and Kenneth Heafield	13185
<i>Bridging the Empirical-Theoretical Gap in Neural Network Formal Language Learning Using Minimum Description Length</i> Nur Lan, Emmanuel Chemla and Roni Katzir	13198
<i>Context versus Prior Knowledge in Language Models</i> Kevin Du, Vésteinn Snæbjarnarson, Niklas Stoehr, Jennifer C. White, Aaron Schein and Ryan Cotterell	13211
<i>Word Matters: What Influences Domain Adaptation in Summarization?</i> Yinghao Li, Siyu Miao, Heyan Huang and Yang Gao	13236
<i>Visualization Recommendation with Prompt-based Reprogramming of Large Language Models</i> Xinhang Li, Jingbo Zhou, Wei Chen, Derong Xu, Tong Xu and Enhong Chen	13250
<i>HOLMES: Hyper-Relational Knowledge Graphs for Multi-hop Question Answering using LLMs</i> Pranoy Panda, Ankush Agarwal, Chaitanya Devaguptapu, Manohar Kaul and Prathosh AP	13263
<i>Toward In-Context Teaching: Adapting Examples to Students' Misconceptions</i> Alexis Ross and Jacob Andreas	13283
<i>Bridging Word-Pair and Token-Level Metaphor Detection with Explainable Domain Mining</i> Yuan Tian, Ruike Zhang, Nan Xu and Wenji Mao	13311
<i>Faithful Logical Reasoning via Symbolic Chain-of-Thought</i> Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee and Wynne Hsu	13326
<i>S²GSL: Incorporating Segment to Syntactic Enhanced Graph Structure Learning for Aspect-based Sentiment Analysis</i> Bingfeng Chen, Qihan Ouyang, Yongqi Luo, Boyan Xu, Ruichu Cai and Zhifeng Hao	13366
<i>Maverick: Efficient and Accurate Coreference Resolution Defying Recent Trends</i> Giuliano Martinelli, Edoardo Barba and Roberto Navigli	13380
<i>ESCoT: Towards Interpretable Emotional Support Dialogue Systems</i> Tenggan Zhang, Xinjie Zhang, Jinming Zhao, Li Zhou and Qin Jin	13395
<i>PathReasoner: Modeling Reasoning Path with Equivalent Extension for Logical Question Answering</i> Fangzhi Xu, Qika Lin, Tianzhe Zhao, JiaweiHan JiaweiHan and Jun Liu	13413

<i>WARDEN: Multi-Directional Backdoor Watermarks for Embedding-as-a-Service Copyright Protection</i> Anudeex Shetty, Yue Teng, Ke He and Qiongkai Xu	13430
<i>Advancing Parameter Efficiency in Fine-tuning via Representation Editing</i> Muling Wu, Wenhao Liu, Xiaohua Wang, Tianlong Li, Changze Lv, Zixuan Ling, Zhu JianHao, Cenyuan Zhang, Xiaoqing Zheng and Xuanjing Huang	13445
<i>Context Consistency between Training and Inference in Simultaneous Machine Translation</i> Meizhi Zhong, Lemao Liu, Kehai Chen, Mingming Yang and Min Zhang	13465
<i>Using Natural Language Explanations to Improve Robustness of In-context Learning</i> Xuanli He, Yuxiang Wu, Oana-Maria Camburu, Pasquale Minervini and Pontus Stenetorp	13477
<i>Chunk, Align, Select: A Simple Long-sequence Processing Method for Transformers</i> Jiawen Xie, Pengyu Cheng, Xiao Liang, Yong Dai and Nan du	13500
<i>ArchCode: Incorporating Software Requirements in Code Generation with Large Language Models</i> Hojae Han, Jaejin Kim, Jaeseok Yoo, Youngwon Lee and Seung-won Hwang	13520
<i>Combining Supervised Learning and Reinforcement Learning for Multi-Label Classification Tasks with Partial Labels</i> Zixia Jia, Junpeng Li, Shichuan Zhang, Anji Liu and Zilong Zheng	13553
<i>MULFE: A Multi-Level Benchmark for Free Text Model Editing</i> Chenhao Wang, Pengfei Cao, Zhuoran Jin, Yubo Chen, Daojian Zeng, Kang Liu and Jun Zhao	13570
<i>MobileSpeech: A Fast and High-Fidelity Framework for Mobile Zero-Shot Text-to-Speech</i> Shengpeng Ji, Ziyue Jiang, Hanting Wang, Jialong Zuo and Zhou Zhao	13588
<i>Spatially-Aware Speaker for Vision-and-Language Navigation Instruction Generation</i> Muraleekrishna Gopinathan, Martin Masek, Jumana Abu-Khalaf and David Suter	13601
<i>HiRoPE: Length Extrapolation for Code Models Using Hierarchical Position</i> Kechi Zhang, Ge Li, Huangzhao Zhang and Zhi Jin	13615
<i>Never Lost in the Middle: Mastering Long-Context Question Answering with Position-Agnostic Decompositional Training</i> Junqing He, Kunhao Pan, Xiaoqun Dong, Zhuoyang Song, LiuYiBo LiuYiBo, Qiangsun Qiangsun, Yuxin Liang, Hao Wang, Enming Zhang and Jiaying Zhang	13628
<i>CodeAgent: Enhancing Code Generation with Tool-Integrated Agent Systems for Real-World Repo-level Coding Challenges</i> Kechi Zhang, Jia Li, Ge Li, Xianjie Shi and Zhi Jin	13643
<i>When is Tree Search Useful for LLM Planning? It Depends on the Discriminator</i> Ziru Chen, Michael White, Ray Mooney, Ali Payani, Yu Su and Huan Sun	13659
<i>LogicBench: Towards Systematic Evaluation of Logical Reasoning Ability of Large Language Models</i> Mihir Parmar, Nisarg Patel, Neeraj Varshney, Mutsumi Nakamura, Man Luo, Santosh Mashetty, Arindam Mitra and Chitta Baral	13679
<i>Meta-Tuning LLMs to Leverage Lexical Knowledge for Generalizable Language Style Understanding</i> Ruohao Guo, Wei Xu and Alan Ritter	13708
<i>Reducing Privacy Risks in Online Self-Disclosures with Language Models</i> Yao Dou, Isadora Krsek, Tarek Naous, Anubha Kabra, Sauvik Das, Alan Ritter and Wei Xu	13732

<i>Navigating the Dual Facets: A Comprehensive Evaluation of Sequential Memory Editing in Large Language Models</i>	
Zihao Lin, Mohammad Beigi, Hongxuan Li, Yufan Zhou, Yuxiang Zhang, Qifan Wang, Wenpeng Yin and Lifu Huang	13755
<i>REFINESUMM: Self-Refining MLLM for Generating a Multimodal Summarization Dataset</i>	
Vaidehi Patil, Leonardo F. R. Ribeiro, Mengwen Liu, Mohit Bansal and Markus Dreyer . . .	13773
<i>When Benchmarks are Targets: Revealing the Sensitivity of Large Language Model Leaderboards</i>	
Norah A. Alzahrani, Hisham Abdullah Alyahya, Yazeed Alnumay, Sultan AlRashed, Shaykhah Z. Alsubaie, Yousef Almushayqih, Faisal Abdulrahman Mirza, Nouf M. Alotaibi, Nora Al-Twairesh, Areeb Alowisheq, M Saiful Bari and Haidar Khan	13787
<i>LLM-Rubric: A Multidimensional, Calibrated Approach to Automated Evaluation of Natural Language Texts</i>	
Helia Hashemi, Jason Eisner, Corby Rosset, Benjamin Van Durme and Chris Kedzie	13806
<i>LIEDER: Linguistically-Informed Evaluation for Discourse Entity Recognition</i>	
Xiaomeng Zhu and Robert Frank	13835
<i>Evaluating Very Long-Term Conversational Memory of LLM Agents</i>	
Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri and Yuwei Fang	13851
<i>Prototypical Reward Network for Data-Efficient Model Alignment</i>	
Jinghan Zhang, Xiting Wang, Yiqiao Jin, Changyu Chen, Xinhao Zhang and Kunpeng Liu	13871
<i>NEO-BENCH: Evaluating Robustness of Large Language Models with Neologisms</i>	
Jonathan Zheng, Alan Ritter and Wei Xu	13885
<i>Impacts of Misspelled Queries on Translation and Product Search</i>	
Greg Hanneman, Natawut Monaikul and Taichi Nakatani	13907
<i>Skin-in-the-Game: Decision Making via Multi-Stakeholder Alignment in LLMs</i>	
Bilgehan Sel, Priya Shanmugasundaram, Mohammad Kachuee, Kun Zhou, Ruoxi Jia and Ming Jin	13921
<i>The MERSA Dataset and a Transformer-Based Approach for Speech Emotion Recognition</i>	
Enshi Zhang, Rafael Trujillo and Christian Poellabauer	13960
<i>Transparent and Scrutable Recommendations Using Natural Language User Profiles</i>	
Jerome Ramos, Hossein A. Rahmani, Xi Wang, Xiao Fu and Aldo Lipani	13971
<i>Fora: A corpus and framework for the study of facilitated dialogue</i>	
Hope Schroeder, Deb Roy and Jad Kabbara	13985
<i>Explanation-aware Soft Ensemble Empowers Large Language Model In-context Learning</i>	
Yue Yu, Jiaming Shen, Tianqi Liu, Zhen Qin, Jing Nathan Yan, Jialu Liu, Chao Zhang and Michael Bendersky	14002
<i>What is the Best Way for ChatGPT to Translate Poetry?</i>	
Shanshan Wang, Derek F. Wong, Jingming Yao and Lidia S. Chao	14025
<i>Rephrasing the Web: A Recipe for Compute and Data-Efficient Language Modeling</i>	
Pratyush Maini, Skyler Seto, Richard He Bai, David Grangier, Yizhe Zhang and Navdeep Jaitly	14044

<i>DeCoT: Debiasing Chain-of-Thought for Knowledge-Intensive Tasks in Large Language Models via Causal Intervention</i>	
Junda Wu, Tong Yu, Xiang Chen, Haoliang Wang, Ryan A. Rossi, Sungchul Kim, Anup Rao and Julian McAuley	14073
<i>Representation Learning with Conditional Information Flow Maximization</i>	
Dou Hu, Lingwei Wei, Wei Zhou and Songlin Hu	14088
<i>GPT is Not an Annotator: The Necessity of Human Annotation in Fairness Benchmark Construction</i>	
Virginia K. Felkner, Jennifer A. Thompson and Jonathan May	14104
<i>Quantifying Contamination in Evaluating Code Generation Capabilities of Language Models</i>	
Martin Riddell, Ansong Ni and Arman Cohan	14116
<i>Language Models are Homer Simpson! Safety Re-Alignment of Fine-tuned Language Models through Task Arithmetic</i>	
Rishabh Bhardwaj, Duc Anh Do and Soujanya Poria	14138
<i>Tracking the Newsworthiness of Public Documents</i>	
Alexander Spangher, Serdar Tumgoren, Ben Welsh, Nanyun Peng, Emilio Ferrara and Jonathan May	14150
<i>EWEK-QA : Enhanced Web and Efficient Knowledge Graph Retrieval for Citation-based Question Answering Systems</i>	
Mohammad Dehghan, Mohammad Ali Alomrani, Sunyam Bagga, David Alfonso-Hermelo, Khalil Bibi, Abbas Ghaddar, Yingxue Zhang, Xiaoguang Li, Jianye Hao, Qun Liu, Jimmy Lin, Boxing Chen, Prasanna Parthasarathi, Mahdi Biparva and Mehdi Rezagholizadeh	14169
<i>Multi-modal Preference Alignment Remedies Degradation of Visual Instruction Tuning on Language Models</i>	
Shengzhi LI, Rongyu Lin and Shichao Pei	14188
<i>Multistage Collaborative Knowledge Distillation from a Large Language Model for Semi-Supervised Sequence Generation</i>	
Jiachen Zhao, Wenlong Zhao, Andrew Drozdov, Benjamin Rozonoyer, Md Arafat Sultan, Jay-Yoon Lee, Mohit Iyyer and Andrew McCallum	14201
<i>Controlled Text Generation for Black-box Language Models via Score-based Progressive Editor</i>	
Sangwon Yu, Changmin Lee, Hojin Lee and Sungroh Yoon	14215
<i>LogogramNLP: Comparing Visual and Textual Representations of Ancient Logographic Writing Systems for NLP</i>	
Danlu Chen, Freda Shi, Aditi Agarwal, Jacobo Myerston and Taylor Berg-Kirkpatrick	14238
<i>Superfiltering: Weak-to-Strong Data Filtering for Fast Instruction-Tuning</i>	
Ming Li, Yong Zhang, Shwai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng and Tianyi Zhou	14255
<i>Confabulation: The Surprising Value of Large Language Model Hallucinations</i>	
Peiqi Sui, Eamon Duede, Sophie Wu and Richard Jean So	14274
<i>IAPT: Instance-Aware Prompt Tuning for Large Language Models</i>	
Wei Zhu, Aaron Xuxiang Tian, Congrui Yin, Yuan Ni, Xiaoling Wang and Guotong Xie ..	14285

Program

Saturday, August 10, 2024

14:00 - 19:00 *Registration*

Sunday, August 11, 2024

07:30 - 20:30 *Registration Check in*

10:30 - 11:00 *Break*

12:30 - 14:00 *Box Lunch Provided*

15:30 - 16:00 *Break*

18:30 - 21:00 *Welcome Reception*

Monday, August 12, 2024

07:30 - 16:30 *Registration*

09:00 - 10:30 *Session 1: Plenary - Opening Session & Keynote TBA*

10:30 - 11:00 *Break*

11:00 - 12:30 *Session 2: Oral's/Posters/ Demo Presentations A*

12:45 - 13:45 *Session 3: Finding Presentations 1*

14:00 - 15:30 *Session 4: Oral's/Posters/ Demo Presentations B*

15:30 - 16:00 *Break*

16:00 - 17:30 *Session 5: Oral's/Posters/ Demo Presentations C*

17:45 - 18:45 *Session 6: Findings Presentations 2*

Tuesday, August 13, 2024

08:30 - 16:30 *Registration*

09:00 - 10:00 *Session 7: Plenary - Keynote TBA*

10:00 - 10:30 *Break*

10:30 - 12:00 *Session 8: Oral's/Posters/ Demo Presentations D*

12:15 - 13:15 *Session 9: Finding Presentations 3*

13:00 - 14:00 *Session 10: Plenary - Business Meeting (all attendees welcome)*

14:30 - 15:30 *Session 11: Plenary - Panel*

15:30 - 16:00 *Break*

16:00 - 17:30 *Session 12: Oral's/Posters/ Demo Presentations E*

19:00 - 22:00 *Social Event Dinner*

Wednesday, August 14, 2024

08:30 - 16:30 *Registration*

09:00 - 10:00 *Session 13: Plenary - Keynote TBA*

10:00 - 10:30 *Break*

10:30 - 12:00 *Session 14: Oral's/Posters/ Demo Presentations F*

12:15 - 13:15 *Session 15: Finding Presentations 4*

13:30 - 15:00 *Session 16: Plenary - Lifetime Achievement & ToT Awards*

15:00 - 15:30 *Break*

15:30 - 17:00 *Session 17: Plenary - Paper Awards*

17:15 - 18:00 *Session 18: Plenary - Closing Session*