

EMNLP 2024

**The 2024 Conference on Empirical Methods in Natural  
Language Processing**

**Proceedings of the Conference**

November 12-16, 2024

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

## Diamond



CITADEL | CITADEL Securities

## Platinum



Megagon Labs



## Gold



## Silver



## Bronze

JPMorganChase



## D&I Champion



©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](mailto:acl@aclweb.org)

ISBN 979-8-89176-164-3

## Message from the General Chair

It's a great honor to welcome you to the 2024 edition of the Conference on Empirical Methods in Natural Language Processing! This year marks the 29th edition of EMNLP, at least according to ACL Anthology proceedings. I counted 14 papers in that first edition; how times have changed and how much our community has grown!

Similar to recent editions, we continue to see growth in the number of papers submitted and the number of registered participants. At the time of this writing, we have reached over 3,500 registrations, with a little over 3,200 attending in-person. This is, without a doubt, a strong signal that we have plenty of interesting research to discuss, and exciting, outstanding challenges continue to motivate us and keep us engaged in the field.

As much as we like to see growth, the organization's logistics have become much more complex with a growing number of submissions and attendees. The effort and time from the many volunteers poured into making this meeting happen are tremendous and worthy of recognition. My first acknowledgments go to the entire organization committee. One of the perks of being a GC is that you get to pick the team that will help in the different aspects of the organization (an exception to this is the PCs that SIGDAT appoints). Suffice it to say that I did have a heavy involvement in inviting people to the team, and this made it extra fun because the final team composition included strong, committed, and just easy-to-work people. My heartfelt thank you goes to all of them:

- Program Chairs: Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung (Vivian) Chen
- ARR Guest Program Chair: Vincent Ng
- Local Chairs: Mark Finlayson, and Zoey Liu
- Industry Track Chairs: Franck Dernoncourt, Daniel Preotiuc-Pietro, and Anastasia Shimorina
- Workshop Chairs: David Vilar, Xiaodan Zhu, and Marta R. Costa-Jussa
- Tutorial Chairs: Jessy Li, and Fei Liu
- Ethics Chairs: Luciana Benotti, Snigdha Chaturvedi, and Sunipa Dev
- Internal Communications Chairs: Jing Li, Yixin Cao
- Demonstration Chairs: Delia Irazu Hernandez Farias, Tom Hope, and Manling Li
- Publication Chairs: Milad Alshomary, Danilo Croce, and Gözde Gül Şahin
- Handbook Chairs: Marco Polignano
- Publicity Chairs: Shruti Rijhwani, and Elias Stengel-Eskin
- Student Volunteer Chairs: Shubhra Kanti (Santu) Karmaker, Nafise Sadat Moosavi, and Emily Prud'hommeaux
- Diversity & Inclusion Chairs: Christos Christodoulopoulos, Veronica Perez-Rosas, and Danish Pruthi
- Sponsorship Chairs: Heba Elfardy, and Leonardo Neves
- Website Chairs: Raj Dabre, and Tiago Torrent
- Virtual Infrastructure Chair: Lianhui Qin, and Vladimir Araujo

- Past Program Chair Advisors: Juan Pino, and Kevin Duh

I would like to give special kudos to some of these names for going well beyond compromising their research, and often, personal time, to get things done. Here's a list of particular thank yous.

I'm thankful for the fantastic effort of the PCs: their eagerness to shape a strong technical program and their constant dedication to the process made my job much easier; thank you Mohit, Vivane, and Yaser!

Many thanks to Vincent, the ARR guest PC! For circumstances beyond the ARR control, he ended up carrying an extra heavy load of EiC duties, and he did so with such a great attitude and commitment, thank you Vincent!

A special thanks go to local arrangement chairs, Mark and Zoey, because they weathered with admirable patience the many demands to support visa letters, scouting for venues for the social event, as well as help address concerns regarding the venue from different constituents; thank you both!

I also want to give special thanks to diversity and inclusion co-chairs. They implemented new ideas to make the event more inclusive, and they handled the challenging task of selecting recipients of D&I awards when the budget was tight and the needs were large. One of the new ideas they launched was the ACL Fellows Lunch, where first-time EMNLP attendees from underrepresented groups are paired with highly regarded members of the ACL (our fellows) to go for lunch together and discuss research, mentorship, and any other topic of relevance to them. Imagine being your first time at an ACL event and having a chance to meet and talk to an ACL fellow. What a lovely first experience at a conference!

I also want to highlight the super responsive website co-chairs, Raj and Tiago. We sent them hundreds of requests to update the website, and every time, we heard back almost immediately and were pleasantly surprised to see the updates right away, thank you!

One committee that is not on the spotlight that much, unless things don't go as planned, is the publications committee. They need to coordinate all the different teams that will generate proceedings or that will contribute content to the proceedings. We were lucky this year to have Danilo Croce leading these efforts. His patience and attention to detail kept us moving towards meeting the deadlines.

Lastly, a big thanks to our sponsorship co-chairs, Leonardo and Heba. They went above and beyond the line of service to bring sponsors to EMNLP. Thanks to their efforts, we landed robust engagements with private and non-profit organizations to support the event. Many thanks for being so proactive about this!

Besides the organization committee, I also want to express my gratitude to Jenn Rachford, ACL business manager, to underline (in particular to Damira Masic), for all her work handling the online component of the conference. To the SIGDAT board, Isabelle Augenstein, Kai-Wei Chang, Alice Oh, and Juan Pino, I want to also express gratitude for all their support and flexibility, thanks to which we were able to grant the equivalent of a virtual registration as a thank you discount to all OC members.

My last special thank you goes to the super large number of unpaid volunteers (reviewers, area chairs, and senior area chairs) who have dedicated their time and energy to making this meeting possible.

I hope this conference brings you inspiration, motivation, energy, and stronger connections; enjoy your time in Miami!

Thamar Solorio

Mohamed bin Zayed University of Artificial Intelligence (MBZUAI) & University of Houston

EMNLP General Chair

## Message from the Program Chairs

Welcome to the 2024 Conference on Empirical Methods in Natural Language Processing!

EMNLP 2024 will be held in a hybrid format, offering attendees the option to join us in person in vibrant Miami, Florida, or to participate remotely from anywhere in the world. We are excited to host this year's conference in Miami. Organizing EMNLP 2024 has been a collaborative effort, made possible by the dedication and hard work of thousands of people. We would like to thank the support and contributions of the following people:

- The General Chair, Thamar Solorio;
- The ARR Editors-in-Chief of the June 2024 cycle (Vincent Ng), Technical Staff (Jonathan K. Kummerfeld), and the entire team (Mausam, Viviane Moreira, Lilja Øvreliid, Anna Rogers, Jun Suzuki, Jing Jiang, Michael White);
- The OpenReview team, especially Celeste Martinez for multiple rounds of technical help in setting up EMNLP 2024 on the OR platform;
- The 99 Senior Area Chairs;
- The 1,458 Area Chairs and the 10,309 reviewers;
- The awards committee chairs, Luke Zettlemoyer, Ivan Titov, and Claire T. Cardie, and the awards committee members;
- The ethics chairs, Luciana Benotti, Snigdha Chaturvedi, and Sunipa Dev;
- The industry track chairs, Franck Dernoncourt, Daniel Preoțiuc-Pietro, Daniel Preoțiuc-Pietro, and Anastasia Shimorina;
- The demonstration chairs, Delia Irazu Hernandez Farias, Tom Hope, and Manling Li;
- The internal communications chairs, Jing Li and Yixin Cao;
- The website chairs, Raj Dabre and Tiago Torrent;
- The publication chairs, Milad Alshomary, Danilo Croce, and Gözde Gül Şahin;
- The handbook chair, Marco Polignano;
- The local organization chairs, Mark Finlayson and Zoey Liu, and their team;
- The publicity chairs, Shruti Rijhwani and Elias Stengel-Eskin;
- The student volunteer chairs, Shubhra Kanti (Santu) Karmaker, Nafise Sadat Moosavi, and Emily Prud'hommeaux;
- The diversity/inclusion chairs, Christos Christodoulopoulos, Veronica Perez-Rosas, and Danish Pruthi;
- The virtual infrastructure chairs, Lianhui Qin and Vladimir Araujo;
- 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 2024 Program Chairs (Lun-Wei Ku, André F. T. Martins, Vivek Srikumar);
- Damira Mrsic 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 submissions to EMNLP 2024 went through a two-stage review process. First, papers were submitted to the ACL Rolling Review (ARR), where they were reviewed by reviewers and received meta-reviews from area chairs. Authors then had the option to commit their reviewed papers to EMNLP via a separate EMNLP 2024 commitment site. At this stage, senior area chairs provided recommendations, and final acceptance decisions were made by the program chairs. This process is consistent with previous conferences, including EACL 2024, NAACL 2024, and ACL 2024.

We worked closely with the ARR team, particularly the Editors-in-Chief for the June 2024 cycle, and served as guest Editors-in-Chief for this round. We also recruited additional reviewers and area chairs for ARR, bringing the total to 10,309 reviewers and 1,458 area chairs in the June 2024 ARR cycle, which handled the majority of EMNLP 2024 submissions. EMNLP also recruited 99 senior area chairs to oversee the review and meta-review process.

Overall, the ARR process ran smoothly, ensuring that all submitted papers received at least three reviews and a meta-review. For the EMNLP commitment phase, senior area chairs made recommendations for 6,108 committed papers based on the reviews, meta-reviews, and the papers themselves, with final acceptance decisions made by the program chairs.

## Acceptance Rate

In total, there are 1271 papers accepted to the Main Conference and 1029 papers accepted to Findings. The acceptance rate for Main Conference papers is 20.8% and one for Findings papers is further 16.9%. The acceptance rate calculation follows precedent set by previous conferences that also go through ACL Rolling Review (ARR), e.g. NAACL 2024, ACL 2024. The calculation takes into account the multi-stage process of ARR where a paper may get revised in ARR and then later committed to the conference. The denominator includes:

- Papers in the ARR April/June 2024 cycle that selected EMNLP (or including EMNLP): 526
- Papers in the ARR April/June 2024 cycle that did not select a preferred venue: 5797
- Papers in the ARR April/June 2024 cycle that selected another conference, but then committed to EMNLP 2024: 7
- Papers in the ARR cycles before April 2024 that committed to EMNLP 202: 65

Among these submissions, 70 were withdrawn before they received 3 reviews and 220 were desk-rejected. In total, the denominator for acceptance rate calculation is  $6395 - 70 - 220 = 6105$ .

The acceptance rate for Main Conference papers is therefore:  $1271/6105 = 20.8\%$ . The final Main Conference proceedings will include 1125 long papers and 143 short papers (3 papers withdrawn after the acceptance decision).

Findings papers are those which are not accepted at the Main Conference, but nevertheless have been judged worthy of publication as “solid work with sufficient substance, quality and novelty”. The next  $1029/6105 = 16.9\%$  of papers were accepted to EMNLP Findings. The final Findings proceedings will include 874 long papers and 129 short papers (26 papers withdrawn after the acceptance decision).



## Special Theme: Efficiency in Model Algorithms, Training, and Inference

This track provides a platform for researchers to explore key aspects of making model algorithms, training, and inference more efficient, e.g., quantization, data requirements, and model size. We welcome submissions that propose innovative approaches, methodologies, and techniques to streamline the training and inference process for language models while optimizing resource utilization and reducing model size. Authors are encouraged to explore various ways to enhance efficiency, including parameter-efficient tuning and methods for learning with less data and smaller model sizes, ultimately leading to more scalable, practical, and resource-efficient NLP systems.

We received 138 submissions to the theme track during the review phase. Among these, 54 papers were accepted to the main conference and a further 42 to Findings of EMNLP.

### Best Paper Selection

This year, we included best paper awards to recognize a broader range of exceptional work as in previous events:

- **Best Papers** (featuring  $\leq 0.25\%$  of accepted papers) and **Outstanding Papers** (featuring  $\leq 2.5\%$  of accepted papers) present fascinating, controversial, surprising, impressive, and/or potentially field-changing ideas.
- **Senior Area Chair’s awards** are similar to Outstanding papers, but specific to this research track.
- **Best Theme Papers** (= Senior Area Chair’s awards for the special theme track) make significant new contributions to efficiency in model algorithms, training, and inference.
- **Social Impact Papers** have the potential for significant positive societal impact.
- **Resource Papers** announce, describe, and share a fascinating, valuable, or potentially field-changing new resource.

Based on nominations from SACs and ACs, 114 candidates have been shortlisted for consideration for the above awards. The final selection is made by the Best Paper Award Committee, and the winners will be announced and will present their work during the closing ceremony.

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 closing ceremony.

### Program Composition & Presentation Modes

Based on feedback from the conference support staff and the Underline team after ACL 2024, we decided to hold the virtual poster sessions during the main conference, aligning them with time slots when in-person participants are also available. This approach allows virtual attendees to participate concurrently with the physical event, avoiding the need for organizers and attendees to engage with the conference twice.

This year, 168 main conference papers were selected for oral presentations by the program chairs, with the goal of creating a well-rounded program featuring a diverse set of topics instead of selecting papers based on their review scores. In addition to the main conference papers, the EMNLP program also includes 13 papers accepted by Computational Linguistics and 32 papers accepted by Transactions of the ACL (TACL). Among these, 30 journal papers will be presented in-person as oral presentations, thematically distributed across appropriate sessions, with formats chosen based on author preferences. Additionally, all Findings papers have been assigned poster presentations alongside other posters for main conference papers in the same track. Rounding out the program are dedicated sessions for the demonstrations track and the student research workshop.

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

- **Prof. Percy Liang** from Stanford University will present on “*Open-Source and Science in the Era of Foundation Models.*”
- **Prof. Anca Dragan** from the University of California Berkeley and Google Deepmind, will share her insights on “*My Journey in AI Safety and Alignment.*”
- **Prof. Tom Griffiths** from Princeton University will explore “*Bayes in the Age of Intelligent Machines.*”

Alongside these keynotes, we are thrilled to host a panel discussion on the *significance of NLP in the age of large language models (LLMs)*. Our esteemed panelists include:

- Prof. Monojit Choudhury, MBZUAI
- Prof. Heng Ji, University of Illinois Urbana-Champaign; Amazon
- Prof. Rada Mihalcea, University of Michigan
- Prof. Alice Oh, Korea Advanced Institute of Science and Technology
- Prof. Sasha Rush, Cornell University; HuggingFace

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

Yaser Al-Onaizan, Saudi Data and AI Authority, National Center for AI  
Mohit Bansal, University of North Carolina at Chapel Hill  
Yun-Nung (Vivian) Chen, National Taiwan University

EMNLP 2024 Program Co-Chairs

## Message from the Local Chair

Dear EMNLP 2024 Participants,

It is our great pleasure to welcome you to EMNLP 2024, held in the lovely tropical city of Miami, Florida, which is North America's gateway to South America and the Caribbean.

Miami's unique language characteristics are at the heart of its identity. Spanish is the dominant language here, but not central American Spanish, as is often found elsewhere in the United States. The dominant dialect is Cuban Spanish, with extensive local enclaves of Venezuelans, Colombians, and Argentinians. Indeed, every nationality and cultural group from South American and Caribbean is well represented here. You will find neighborhoods which speak primarily Haitian Creole, as well as Brazilian Portuguese. When you add in the large populations of Europeans from France, Spain, and the Balkans, the local language picture becomes quite rich indeed. Miami even has its own dialect of English, which was identified by Florida International University socio-linguist Phillip Carter: this dialect overlays standard English with Spanish syntax and Cuban vocabulary and slang, and comes with its own distinctive accent. So Miami has truly a distinctive language mix!

Given Miami's differences from much of the rest of the United States—in language, population, and climate—locals joke that we are not really in the United States at all, but rather in Latin America. The joke continues with the observation that one of Miami's most convenient attributes, as a supposedly independent Latin American nation, is its proximity to the United States and the fact that we share an open border and a currency. While tongue-in-cheek, if you have traveled elsewhere in the United States you will see a grain of truth in this, and we hope you enjoy and appreciate Miami's unusual character.

While you are here we hope you take full advantage of the cultural richness Miami has to offer, and the many fun things to do. This includes our vibrant, burgeoning, world-class food scene (with many Michelin starred restaurants), our world-renowned nightlife, and our beautiful beaches. Try the widely available Latin American food, such as *pastelitos* or *croquetas*, or our many varieties of speciality coffee, such as *cafecitas* (cuban coffee) or *cortaditos*. Enjoy a night out salsa dancing on *Calle Ocho* (8th Street) in Little Havana, or catch a Latin music concert at the many concert halls in Downtown, Brickell, or Wynwood. Dance the night away at our dance clubs that feature world-class electronic music DJs, or visit some of our fantastic museums, such as the Perez Art Museum Miami (the PAMM), the Frost Museum of Science (where the social event will be held), or the Viscaya Museum and Gardens. Indulge your dark desire for conspicuous consumption of luxury goods in our high-end stores in Brickell City Center or the Miami Design District.

In the past Miami has been famous for its nightlife and Latin flavors, as a place to visit, relax, and have fun. This is still true, but Miami has grown tremendously as a city in recent years in many other ways. For example, we now boast two major research universities: Florida International University, a public research university, is home to over 50,000 students and has recently been ranked as a top-50 public university in the United States and a top-100 university overall. The University of Miami, also ranked in the top 100, boasts a beautiful campus in Coral Gables, nearly 20,000 students, and a major research hospital. Miami is also home to a fast-growing tech startup and cryptocurrency scene, with a variety of startup accelerators, incubators, funders, and networking organizations, including Endeavor Miami, The Knight Foundation, The Lab, Rokk3r Labs, eMerge Americas, the Miami Angels and Wyncode. Finally, Miami continues to solidify its standing as a major hub of international finance for Latin America, with many banks and other financial firms opening major branches here or even moving their headquarters to Miami.

Returning to EMNLP, we would like to extend our thanks to Jennifer Rachford and Megs Haddad, both of the ACL business office, who provided quick, gracious, and ever-informative help in the quite laborious

process of issuing many hundreds of visa invitation letters for those coming from abroad. If you were one of those who needed a visa letter, and see them at the conference, please take a moment to thank them for their hard work.

In closing, we hope that you will thoroughly enjoy your stay in Miami, exploring its rich culture, taking advantage of the many opportunities for fun, all the while getting the most out the extensive technical program of EMNLP.

*¡Bienvenido a Miami!*

Mark Finlayson  
Florida International University, Miami, FL

Zoey Liu  
University of Florida, Gainesville, FL

Local Chairs, EMNLP 2024

# Organizing Committee

## General Chair

Thamar Solorio, Mohamed bin Zayed University of Artificial Intelligence  
and University of Houston

## Program Chairs

Yaser Al-Onaizan, Saudi Data and AI Authority, National Center for AI  
Mohit Bansal, University of North Carolina at Chapel Hill  
Yun-Nung (Vivian) Chen, National Taiwan University

## Local Chair

Mark Finlayson, Florida International University  
Zoey Liu, University of Florida

## Industry Track Chairs

Franck Dernoncourt, Adobe Research  
Daniel Preotiu-Pietro, Bloomberg  
Anastasia Shimorina, Orange Innovation

## Workshop Chairs

David Vilar, Google Inc.  
Xiaodan Zhu, Queen's University  
Marta R. Costa-Jussa, Meta AI

## Tutorial Chairs

Jessy Li, The University of Texas at Austin  
Fei Liu, Emory University

## Ethics Chairs

Luciana Benotti, Facultad de Matemática, Astronomía, Física y Computación  
Snigdha Chaturvedi, University of North Carolina at Chapel Hill  
Sunipa Dev, Google Research

## Demonstration Chairs

Delia Irazu Hernandez Farias, Instituto Nacional de Astrofísica, Óptica y Electrónica  
Tom Hope, AI2  
Manling Li, Northwestern University

### **Publication Chairs**

Milad Alshomary, Columbia University  
Danilo Croce, University of Rome Tor Vergata  
Gözde Gül Şahin, Koç University

### **Handbook Chair**

Marco Polignano, University of Bari Aldo Moro

### **Publicity Chairs**

Shruti Rijhwani, Google DeepMind  
Elias Stengel-Eskin, University of North Carolina

### **Student Volunteer Chairs**

Shubhra Kanti (Santu) Karmaker, University of Central Florida  
Nafise Sadat Moosavi, University of Sheffield  
Emily Prud'hommeaux, Boston College

### **Diversity and Inclusion Chairs**

Christos Christodoulopoulos, Amazon  
Veronica Perez-Rosas, University of Michigan  
Danish Pruthi, Indian Institute of Science (IISc), Bangalore

### **Sponsorship Chairs**

Heba Elfardy, Amazon  
Leonardo Neves, Snap Inc.

### **Website Chairs**

Raj Dabre, National Institute of Information and Communications Technology (NICT), Japan  
Tiago Torrent, Federal University of Juiz de Fora

### **Virtual Infrastructure Chair**

Lianhui Qin, AI2  
Vladimir Araujo, KU Leuven

### **Past Program Chair Advisors**

Juan Pino, Meta  
Kevin Duh, Johns Hopkins University

# Program Committee

## Computational Social Science and Cultural Analytics

Chenhao Tan, University of Chicago  
David Bamman, School of Information at UC Berkeley  
Svitlana Volkova, Aptima Inc.  
Tanmoy Chakraborty, Indian Institute of Technology Delhi

## Dialogue and Interactive Systems

Larry P. Heck, Tech AI  
Luis Fernando D'Haro, Universidad Politécnica de Madrid  
Minlie Huang, Tsinghua University  
Rebecca Passonneau, Indian Institute of Technology Delhi

## Discourse and Pragmatics

Malihe Alikhani, Northeastern University  
Vered Shwartz, University of British Columbia

## Ethics, Bias, and Fairness

Malvina Nissim, University of Groningen  
Monojit Choudhury, Mohamed bin Zayed University of Artificial Intelligence  
Natalie Schluter, Apple  
Kai-Wei Chang, University of California, Los Angeles

## Generation

Mirella Lapata, Edinburgh University  
Naoaki Okazaki, Tokyo Institute of Technology  
Sebastian Gehrmann, Bloomberg LP  
Yangfeng Ji, University of Virginia

## Human-Centered NLP

David Mimno, Cornell University  
Jeff Bigham, Carnegie Mellon University  
Marine Carpuat, University of Maryland

## Information Extraction

Derry Tanti Wijaya, Boston University  
Lifu Huang, University of California  
Ndapa Nakashole, University of California  
Ruihong Huang, Texas A&M University  
Scott Yih, Facebook AI Research

## **Information Retrieval and Text Mining**

Luca Soldaini, Allen Institute for AI  
Pawan Goyal, Indian Institute of Technology  
Wenhu Chen, University of Waterloo

## **Interpretability, Interactivity and Analysis of Models for NLP**

Grzegorz Chrupala, Tilburg University  
Lingpeng Kong, University of Hong Kong  
Nadir Durrani, Arabic Language Technologies (ALT)  
Tal Linzen, New York University  
Ren Chen, University of Southern California  
Xuanjing Huang, Fudan University

## **Language Modeling**

Anna Rumshisky, UMass Lowell  
Nanyun Peng, University of California, Los Angeles  
Swabha Swayamdipta, University of Southern California  
Tatsunori Hashimoto, Stanford University

## **Linguistic Theories, Cognitive Modeling and Psycholinguistics**

Frank Keller, University of Edinburgh  
Najoung Kim, Boston University

## **Low-resource Methods for NLP**

Kareem Darwish, Qatar Computing Research Institute  
Miryam De Lhoneux, KU Leuven in Belgium  
Shafiq Joty, Salesforce  
Wenpeng Yin, Penn State University  
Yue Dong, University of California

## **Machine Learning for NLP**

Gunhee Kim, Seoul National University  
Partha Talukdar, Google Research and IISc Bangalore  
Shashank Srivastava, UNC Chapel Hill  
Taylor Berg-Kirkpatrick, University of California

## **Machine Translation**

Alexander Fraser, Technical University of Munich  
Lei Li, Carnegie Mellon University  
Paco Guzmán, Meta AI

## **Multilinguality and Language Diversity**

Antonis Anastasopoulos, George Mason Computer Science  
Manaal Faruqui, Google



Steven Bird, Charles Darwin University  
Zornitsa Kozareva, SliceX AI

### **Multimodality and Language Grounding to Vision, Robotics and Beyond**

Gabriel Stanovsky, Hebrew University of Jerusalem  
Hao Tan, Nottingham University Business School China  
Jack Hessel, Samaya.ai  
Jesse Thomason, University of Southern California  
Roma Patel, Brown University  
Zhe Gan, Apple

### **NLP Applications**

Avirup Sil, IBM Research AI  
Gholamreza Haffari, Monash University  
Gokhan Tur, University of Illinois Urbana-Champaign  
Joel Tetreault, University of Rochester  
Kevin Small, Amazon  
Makoto Miwa, Toyota Technological Institute  
Parisa Kordjamshidi, Michigan State University  
Roman Klinger, University of Bamberg  
Sudha Rao, Microsoft Research  
Wei Lu, University of Michigan

### **Phonology, Morphology and Word Segmentation**

Brian Roark, Google

### **Question Answering**

Eunsol Choi, New York University  
Huan Sun, The Ohio State University  
Mrinmaya Sachan, ETH Zürich  
Siva Reddy, McGill University

### **Resources and Evaluation**

Adina Williams, Facebook AI Research  
Alane Suhr, UC Berkeley  
Eduardo Blanco, University of Arizona  
Jimmy Lin, University of Waterloo  
Masayuki Asahara, National Institute for Japanese Language and Linguistics  
Sujian Li, Peking University  
Wei Xu, Georgia Institute of Technology  
Yue Zhang, Westlake University  
Yufang Hou, IBM Research

### **Semantics: Lexical, Sentence level, Document Level, Textual Inference, etc.**

Marianna Apidianaki, University of Pennsylvania

Sujith Ravi, Temple University

### **Sentiment Analysis, Stylistic Analysis, and Argument Mining**

Saif Mohammad, National Research Council Canada

Veronique Hoste, Ghent University

Zhongyu Wei, Fudan University

### **Special Theme: Efficiency in Model Algorithms, Training, and Inference**

Emma Strubell, Carnegie Mellon University

Nafise Sadat Moosavi, University of Sheffield

Sara Hooker, Cohere For AI

### **Speech Processing and Spoken Language Understanding**

Julia Hirschberg, Columbia University

Preethi Jyothi, IIT Bombay

### **Summarization**

Ramakanth Pasunuru, FAIR at Meta

Xiaojun Wan, Peking University

### **Syntax, Parsing and their Applications**

Lingpeng Kong, University of Hong Kong

Najoung Kim, Boston University

### **Area Chairs**

Mausam ., Mohamed Abdalla, Tarek F. Abdelzaher, Omri Abend, Gavin Abercrombie, Heike Adel, David Ifeoluwa Adelani, Somak Aditya, Vaibhav Adlakha, Ameeta Agrawal, Priyanka Agrawal, Sweta Agrawal, Roei Aharoni, Alham Fikri Aji, Alan Akbik, Arjun Reddy Akula, James Allan, Fernando Alva-Manchego, Prithviraj Ammanabrolu, Aijun An, Jisun An, Antonios Anastasopoulos, Antonios Anastasopoulos, Carolyn Jane Anderson, Jacob Andreas, Maria Antoniak, Xiang Ao, Alessio Palmero Arosio, Jun Araki, Yuki Arase, Akari Asai, Ehsaneddin Asgari, Pepa Atanasova, Giuseppe Attardi, Fan Bai, Richard He Bai, Xuefeng Bai, JinYeong Bak, Timothy Baldwin, David Bamman, Sivaji Bandyopadhyay, Chitta Baral, M Saiful Bari, Jeremy Barnes, Valentin Barriere, Pierpaolo Basile, Ali Basirat, Jasmijn Bastings, Timo Baumann, Eyal Ben-David, Yassine Benajiba, Taylor Berg-Kirkpatrick, Leon Bergen, Raffaella Bernardi, Steven Bethard, Wei Bi, Federico Bianchi, Chris Biemann, Lidong Bing, Johannes Bjerva, Eduardo Blanco, Terra Blevins, Gemma Boleda, Danushka Bollegala, Marcel Bollmann, Aurélien Bossard, Florian Boudin, Jordan Lee Boyd-Graber, Leonid Boytsov, Bridger Bridger, Chris Brockett, Emanuele Bugliarello, Wray Buntine, Elena Cabrio, Avi Caciularu, Deng Cai, Zhaowei Cai, Ruken Cakici, Pengfei Cao, Qingqing Cao, Yixin Cao, Ziqiang Cao, Doina Caragea, Dallas Card, Giuseppe Carenini, Samuel Carton, Paula Carvalho, Tommaso Caselli, Dumitru-Clementin Cercel, Tuhin Chakrabarty, Tanmoy Chakraborty, Ilias Chalkidis, Sarath Chandar, Angel X Chang, Ernie Chang, Shiyu Chang, Yung-Chun Chang, Snigdha Chaturvedi, Aditi Chaudhary, Chen Chen, Chung-Chi Chen, Hanjie Chen, Hui Chen, Kehai Chen, Long Chen, Lu Chen, Muhao Chen, Pinzhen Chen, Qian Chen, Tianlong Chen, Wei Chen, Wenhui Chen, Wenliang Chen, Xilun Chen, Xinchu Chen, Xiuying

Chen, Yangyi Chen, Yidong Chen, Yulong Chen, Yun-Nung Chen, Zhiyu Chen, Fei Cheng, Hao Cheng, Lu Cheng, Wei Cheng, Yu Cheng, Ta-Chung Chi, Hai Leong Chieu, Eleanor Chodroff, Edward Choi, Eunsol Choi, Jinho D. Choi, Jaegul Choo, Leshem Choshen, Prafulla Kumar Choubey, Christos Christodouloupoulos, Chenhui Chu, Yun-Wei Chu, Kenneth Church, Stephen Clark, Simon Clematide, Oana Cocarascu, Arman Cohan, Simone Conia, John M Conroy, Caio Filippo Corro, Benoit Crabbé, Leyang Cui, Yiming Cui, Amanda Cercas Curry, Alexander Nicholas D'Amour, Luis Fernando D'Haro, Giovanni Da San Martino, Raj Dabre, Daniel Dakota, Fahim Dalvi, Bhavana Dalvi Mishra, Sandipan Dandapat, Pradeep Dasigi, Marie-Catherine De Marneffe, José G. C. De Souza, Vera Demberg, Daryna Dementieva, Yang Deng, Yuntian Deng, Zhenyun Deng, Leon Derczynski, Maunendra Sankar Desarkar, Barry Devereux, Shizhe Diao, Liang Ding, Liviu P Dinu, Li Dong, Yushun Dong, Bonnie J Dorr, Zhicheng Dou, Zi-Yi Dou, Antoine Doucet, Eduard Dragut, Markus Dreyer, Andrew Drozdov, Xinya Du, Dheeru Dua, Kevin Duh, Nadir Durani, Ondrej Dusek, Lukas Edman, Vladimir Eidelman, Bryan Eikema, Julian Martin Eisenschlos, Asif Ekbal, Yanai Elazar, Maha Elbayad, Ali Emami, Cristina España-Bonet, Luis Espinosa-Anke, Allyson Ettinger, Alexander Fabbri, Agnieszka Falenska, Zhihao Fan, Meng Fang, Yi Fang, Amir Feder, Hao Fei, Nils Feldhus, Naomi Feldman, Fuli Feng, Shangbin Feng, Shi Feng, Xiachong Feng, Xiaocheng Feng, Yansong Feng, Fernando Fernández-Martínez, Anjalie Field, Mark Fishel, Jeffrey Flanigan, Jennifer Foster, Anette Frank, Markus Freitag, Lea Frermann, Daniel Fried, Jintan Fu, Peng Fu, Yi Fung, Saadia Gabriel, Matthias Gallé, Zhe Gan, Jiahui Gao, Tianyu Gao, Yang Gao, Yifan Gao, Judith Gaspers, Jon Gauthier, Tao Ge, Sebastian Gehrmann, Alborz Geramifard, Luke Gessler, Mor Geva, Abbas Ghaddar, Deepanway Ghosal, Debanjan Ghosh, Mario Giulianelli, Goran Glavaš, Tejas Gokhale, Seraphina Goldfarb-Tarrant, Omer Goldman, Sharon Goldwater, Helena Gomez Adorno, Hila Gonen, Joseph E. Gonzalez, Matthew R. Gormley, Cyril Goutte, Kartik Goyal, Jia-Chen Gu, Jiatao Gu, Qing Gu, Frank Guerin, Lin Gui, Tao Gui, Honglei Guo, Zhijiang Guo, Deepak Gupta, Manish Gupta, Vivek Gupta, Sairam Gurajada, Suchin Gururangan, Francisco Guzmán, Carlos Gómez-Rodríguez, Le An Ha, Nizar Habash, Ivan Habernal, Barry Haddow, Gholamreza Haffari, Michael Hahn, Keith Hall, Xianpei Han, Viktor Hangya, Daniel Hardt, Peter Hase, Tatsunori Hashimoto, Di He, Jie He, Junxian He, Keqing He, Shizhu He, Daniel Hershcovich, Jack Hessel, Ryuichiro Higashinaka, Julia Hockenmaier, Valentin Hofmann, Christopher M Homan, Yicong Hong, Enamul Hoque, Sameera Horawalavithana, Md Mosharaf Hossain, Andreas Hotho, Cho-Jui Hsieh, I-Hung Hsu, Baotian Hu, Hai Hu, Jennifer Hu, Junjie Hu, Mengting Hu, Xuming Hu, Fei Huang, Jiaxin Huang, Jie Huang, Kuan-Hao Huang, Kung-Hsiang Huang, Lifu Huang, Peijie Huang, Shujian Huang, Ting-Hao Kenneth Huang, Xuanjing Huang, Binyuan Hui, Wonseok Hwang, Diana Inkpen, Naoya Inoue, Kentaro Inui, Radu Tudor Ionescu, Daphne Ippolito, Mohit Iyyer, Kokil Jaidka, Abhik Jana, Hyeju Jang, Harsh Jhamtani, Yangfeng Ji, Robin Jia, Yichen Jiang, Yong Jiang, Lifeng Jin, Mali Jin, Qin Jin, Zhijing Jin, Preethi Jyothi, Jad Kabbara, Ehsan Kamaloo, Hidetaka Kamigaito, Masahiro Kaneko, Diptesh Kanojia, Giannis Karamanolakis, Sarvnaz Karimi, Santu Karmaker, Marzena Karpinska, Divyansh Kaushik, Pei Ke, Mitesh M Khapra, Daniel Khashabi, Khalid Al Khatib, Tushar Khot, Hyounghun Kim, Jin-Hwa Kim, Joo-Kyung Kim, Taeuk Kim, Young Jin Kim, Christo Kirov, Ekaterina Kochmar, Mamoru Komachi, Rik Koncel-Kedziorski, Ioannis Konstas, Parisa Kordjamshidi, Zornitsa Kozaeva, Amrith Krishna, Kalpesh Krishna, Nikhil Krishnaswamy, Udo Kruschwitz, Sachin Kumar, Shankar Kumar, Sumeet Kumar, Vishwajeet Kumar, Anoop Kunchukuttan, Tatsuki Kuribayashi, Haewoon Kwak, Philippe Laban, Yuxuan Lai, Wai Lam, Gerasimos Lampouras, Lukas Lange, Philippe Langlais, Mirella Lapata, Mark Last, Ivano Lauriola, Anne Lauscher, Hady W. Lauw, Eric Le Ferrand, Dong-Ho Lee, Hwanhee Lee, Hwaran Lee, Mark G. Lee, Moontae Lee, Roy Ka-Wei Lee, SangKeun Lee, Wenqiang Lei, Yuanyuan Lei, Oliver Lemon, Piyawat Lertvittayakumjorn, Gina-Anne Levow, Roger P. Levy, Bing Li, Chenliang Li, Chuyuan Li, Dongfang Li, Dongxu Li, Fei Li, Huayang Li, Jing Li, Jiyi Li, Junhui Li, Juntao Li, Lei Li, Lei Li, Linjie Li, Linyang Li, Manling Li, Peng Li, Piji Li, Ru Li, Tao Li, Wei Li, Wenjie Li, Xian Li, Xiang Lorraine Li, Xiaoli Li, Xin Li, Xiujun Li, Yaliang Li, Yanghao Li, Yanyang Li, Yongbin Li, Yu Li, Yuan-Fang Li,

Zejun Li, Zhenghua Li, Zhifei Li, Maria Liakata, Bin Liang, Paul Pu Liang, Xinnian Liang, Lizi Liao, Jindřich Libovický, Bill Yuchen Lin, Hongfei Lin, Hongyu Lin, Junyang Lin, Shou-De Lin, Yankai Lin, Zhaojiang Lin, Nedim Lipka, Robert Litschko, Bang Liu, Bing Liu, Bingquan Liu, Hao Liu, Hao Liu, Jiahao Liu, Kang Liu, Lemao Liu, Ming Liu, Pengfei Liu, Qi Liu, Qian Liu, Tingwen Liu, Xiao Liu, Xiaodong Liu, Xuebo Liu, Yang Liu, Yang Liu, Yi Liu, Yuanxin Liu, Zhenghao Liu, Zhiyuan Liu, Zoey Liu, Kyle Lo, Wei Lu, Ximing Lu, Yaojie Lu, Michal Lukasik, Anh Tuan Luu, Hanjia Lyu, Qing Lyu, Jing Ma, Ruotian Ma, Tengfei Ma, Wei-Yun Ma, Xuezhe Ma, Sean MacAvaney, Andrea Madotto, Kyle Mahowald, Jean Maillard, Bodhisattwa Prasad Majumder, Navonil Majumder, Prodromos Malakasiotis, Chaitanya Malaviya, Saab Mansour, Xiaoxi Mao, Yuning Mao, Kelly Marchisio, Vukosi Marivate, Lara J. Martin, Sérgio Matos, Yuichiroh Matsubayashi, Yuji Matsumoto, Jonathan May, Julian McAuley, Kathleen McKeown, Clara Meister, Dheeraj Mekala, Julia Mendelsohn, Fandong Meng, Yu Meng, Yuxian Meng, Florian Metze, Antonio Valerio Miceli Barone, Timothee Mickus, Margot Mieskes, Timothy A Miller, Sewon Min, Niloofar Mireshghallah, Pushkar Mishra, Kanishka Misra, Koh Mitsuda, Yasuhide Miura, Daichi Mochihashi, Ashutosh Modi, Marie-Francine Moens, Saif M. Mohammad, Syrielle Montariol, Terufumi Morishita, David R Mortensen, Marius Mosbach, Lili Mou, Hussein Mozannar, Khalil Mrini, Aaron Mueller, Shamsuddeen Hassan Muhammad, Benjamin Muller, Tsendsuren Munkhdalai, Smaranda Muresan, Rudra Murthy, Nona Naderi, Masaaki Nagata, Ndapa Nakashole, Preslav Nakov, Linyong Nan, Jason Naradowsky, Shrikanth Narayanan, Isar Nejadgholi, Mariana Neves, Hwee Tou Ng, Patrick Ng, Le-Minh Nguyen, Thien Huu Nguyen, Thong Thanh Nguyen, Vahid Partovi Nia, Jan Niehues, Takashi Ninomiya, Tong Niu, Farhad Nooralahzadeh, Naoaki Okazaki, Manabu Okumura, Alexandra Olteanu, Matan Orbach, Vicente Ordonez, Nedjma Ousidhoum, Maria Leonor Pacheco, Serguei V. S. Pakhomov, Alexis Palmer, Liangming Pan, Richard Yuanzhe Pang, Pinelopi Papalampidi, Nikolaos Pappas, Natalie Parde, Joonsuk Park, Alicia Parrish, Gabriella Pasi, Rebecca J. Passonneau, Siddharth Patwardhan, Debjit Paul, Ellie Pavlick, Hao Peng, Nanyun Peng, Lis Pereira, Maxime Peyrard, Jonas Pfeiffer, Mohammad Taher Pilehvar, Tiago Pimentel, Irina Piontkovskaya, Bryan A. Plummer, Simone Paolo Ponzetto, Maja Popovic, Alexandros Potamianos, Saloni Potdar, Christopher Potts, Emily Prud'hommeaux, Danish Pruthi, Peng Qi, Tao Qi, Peng Qian, Lianhui Qin, Yujia Qin, Li Qing, Xiaojun Quan, Deepak Ramachandran, Leonardo Ranaldi, Sudha Rao, Yanghui Rao, Hannah Rashkin, Traian Rebedea, Marek Rei, Machel Reid, Ehud Reiter, Xiang Ren, Zhaochun Ren, Leonardo F. R. Ribeiro, Kyle Richardson, German Rigau, Brian Roark, Oleg Rokhlenko, Subendhu Rongali, Joseph Le Roux, Alla Rozovskaya, Terry Ruas, Attapol Rutherford, Paul Röttger, Kenji Sagae, Swarnadeep Saha, Hassan Sajjad, Keisuke Sakaguchi, Tetsuya Sakai, Hiroki Sakaji, Bidisha Samanta, Maarten Sap, Naomi Saphra, Sudeshna Sarkar, Ryohei Sasano, Asad B. Sayeed, Dominik Schlechtweg, Viktor Schlegel, Michael Sejr Schlichtkrull, William Schuler, H. Schwartz, João Sedoc, Shubhashis Sengupta, Rico Sennrich, Kevin Seppi, Lei Sha, Jingbo Shang, Yutong Shao, Ori Shapira, Vasu Sharma, Hassan Shavarani, Jiaming Shen, Yongliang Shen, Tom Sherborne, Peng Shi, Shuming Shi, Weiyang Shi, Manish Shrivastava, Lei Shu, Avirup Sil, Carina Silberer, Maneesh Kumar Singh, Mayank Singh, Pranaydeep Singh, Koustuv Sinha, Gabriel Skantze, Kevin Small, Noah A. Smith, Luca Soldaini, Pia Sommerauer, Linfeng Song, Wei Song, Xingyi Song, Yan Song, Sandeep Soni, Jeffrey S. Sorensen, Aitor Soroa, Shashank Srivastava, Felix Stahlberg, Gabriel Stanovsky, Manfred Stede, Shane Steinert-Threlkeld, Karl Stratos, Jannik Strötgen, Jinsong Su, Qi Su, Qinliang Su, Yu Su, Fabian M. Suchanek, Saku Sugawara, Alessandro Suglia, Dianbo Sui, Aixin Sun, Hao Sun, Jiao Sun, Kai Sun, Lin Sun, Shao-Hua Sun, Mihai Surdeanu, Jun Suzuki, György Szarvas, Santosh T.y.s.s, Hiroya Takamura, Zeerak Talat, Chenhao Tan, Xu Tan, Hao Tang, Raphael Tang, Siliang Tang, Xiangru Tang, Zhiyang Teng, Maartje Ter Hoeve, Alberto Testoni, Laure Thompson, Craig Thomson, Yuanhe Tian, Zhiliang Tian, Amalia Todirascu, Nadi Tomeh, Sara Tonelli, Mariya Toneva, Antonio Toral, Paolo Torrioni, Kristina Toutanova, Amine Trabelsi, Trang Tran, Marcos V Treviso, Enrica Troiano, Richard Tzong-Han Tsai, Yu-Hsiang Tseng, Yoshimasa Tsuruoka, Zhao-peng Tu, Ferhan Ture, Martin Tutek, Ana Sabina Uban, Shyam Upadhyay, Ahmet Üstün, Clara

Vania, Shikhar Vashishth, Pat Verga, Yannick Versley, Esaú Villatoro-tello, Soroush Vosoughi, Thang Vu, Thuy Vu, Tu Vu, Byron C Wallace, Yao Wan, Bo Wang, Chengyu Wang, Fei Wang, Hao Wang, Jingjing Wang, Longyue Wang, Lucy Lu Wang, Shijin Wang, Shuohang Wang, Sirui Wang, Siyuan Wang, Tao Wang, Tianlu Wang, Wei Wang, Weiqi Wang, Wenya Wang, Xiaozhi Wang, Xin Wang, Xin Eric Wang, Xing Wang, Xinyu Wang, Yan Wang, Yizhong Wang, Zhaowei Wang, Zhongqing Wang, Zihao Wang, Alex Warstadt, Ingmar Weber, Daimeng Wei, Leonie Weisweiler, Haoyang Wen, Tim Weninger, Michael White, Sarah Wiegrefe, John Frederick Wieting, Derry Tanti Wijaya, Bryan Wilie, Genta Indra Winata, Kam-Fai Wong, Chien-Sheng Wu, Chuan Wu, Lijun Wu, Yuanbin Wu, Yunfang Wu, Zhiyong Wu, Fei Xia, Yang Xiang, Jing Xiao, Tong Xiao, Wen Xiao, Ruobing Xie, Chenyan Xiong, Bo Xu, Haiyang Xu, Hongfei Xu, Jinan Xu, Qionikai Xu, Ruifeng Xu, Wei Xu, Xinnuo Xu, Zenglin Xu, Zhaozhuo Xu, Nianwen Xue, Shweta Yadav, Yadollah Yaghoobzadeh, Lingyong Yan, Rui Yan, Eugene Yang, Hao Yang, Jie Yang, Jingkang Yang, Liang Yang, Min Yang, Yezhou Yang, Yinfei Yang, Zhen Yang, Zhenglu Yang, Zhengyuan Yang, Bingsheng Yao, Huaxiu Yao, Jin-Ge Yao, Wenlin Yao, Ziyu Yao, Andrew Yates, Qinyuan Ye, Wei Ye, Xi Ye, An-Zi Yen, Da Yin, Yongjing Yin, Jaehong Yoon, Naoki Yoshinaga, Koichiro Yoshino, Dian Yu, Dong Yu, Jianfei Yu, Jianxing Yu, Juntao Yu, Lei Yu, Liang-Chih Yu, Mo Yu, Tao Yu, Wenhao Yu, Yue Yu, Marcos Zampieri, Fabio Massimo Zanzotto, Weixin Zeng, Xingshan Zeng, Chrysoula Zerva, Biao Zhang, Chen Zhang, Chen Zhang, Jiajun Zhang, Jiaxin Zhang, Jingyi Zhang, Ke Zhang, Li Zhang, Meishan Zhang, Ming Zhang, Ningyu Zhang, Rui Zhang, Ruiyi Zhang, Shikun Zhang, Shiyue Zhang, Shuai Zhang, Shuai Zhang, Wei Zhang, Weinan Zhang, Wenxuan Zhang, Xiangliang Zhang, Xingxing Zhang, Yang Zhang, Yang Zhang, Yi Zhang, Yue Zhang, Zhe Zhang, Zhenyu Zhang, Zhirui Zhang, Zhuosheng Zhang, Chenye Zhao, Dongyan Zhao, Hai Zhao, Jieyu Zhao, Jun Zhao, Shiwan Zhao, Tiejun Zhao, Xin Zhao, Xinran Zhao, Yang Zhao, Chujie Zheng, Hai-Tao Zheng, Zaixiang Zheng, Ruiqi Zhong, Ben Zhou, Guangyou Zhou, Hao Zhou, Jie Zhou, Qingyu Zhou, Wangchunshu Zhou, Wenxuan Zhou, Yi Zhou, Yingbo Zhou, Hao Zhu, Qiaoming Zhu, Su Zhu, Xiaodan Zhu, Zining Zhu, Yuchen Zhuang, Yf-tah Ziser, Yuexian Zou

## Reviewers

Micheal Abaho, Ahmed Abbasi, Ahmed Abdelali, Adnen Abdessaied, Muhammad Abdul-Mageed, Amir Abdullah, Giuseppe Abrami, Mitchell Abrams, Abdalghani Abujabal, Artem Abzaliev, Manoj Acharya, Samuel Ackerman, Griffin Thomas Adams, Tosin Adewumi, Suman Adhya, Muhammad Farid Adilazuarda, Somak Aditya, Stergos D. Afantenos, Amit Agarwal, Divyansh Agarwal, Ishika Agarwal, Milind Agarwal, Rajat Agarwal, Sumit Agarwal, Rodrigo Agerri, Arshiya Aggarwal, Jai Aggarwal, Karan Aggarwal, Aishwarya Agrawal, Ameeta Agrawal, Kumar Krishna Agrawal, Thomas Dybdahl Ahle, Ibrahim Said Ahmad, Iftakhar Ahmad, Sina Ahmadi, Arash Ahmadian, Murtadha Ahmed, Shafiuddin Rehan Ahmed, Zeeshan Ahmed, Hyemin Ahn, Jaewoo Ahn, Sumyeong Ahn, Wonhyuk Ahn, Yong-Yeol Ahn, Kian Ahrabian, Sanchit Ahuja, Xi Ai, Ankit Aich, Taichi Aida, Akiko Aizawa, Alham Fikri Aji, Yamen Ajour, Reina Akama, Kosuke Akimoto, Nil-Jana Akpınar, Taha Aksu, Syeda Sabrina Akter, Ekin Akyürek, Hend Al-Khalifa, Badr AlKhamissi, Sultan AlRashed, Ahmed Alajrami, Firoj Alam, Md Mahfuz Ibn Alam, Parand A. Alamdari, Belen Alastruey, Abeer Aldayel, Georgios Alexandridis, David Alfter, Asaad Alghamdi, Israa Alghanmi, Tariq Alhindi, Ahmed Ali, Ammar Ali, Sarwan Ali, Alexandre Allauzen, Emily Allaway, Sawsan Alqahtani, Duarte Miguel Alves, Zaid Alyafeai, Norah A. Alzahrani, Alfonso Amayuelas, Enrique Amigo, Saadullah Amin, Silvio Amir, Haozhe An, Jisun An, Kai-kai An, Shengnan An, Wenbin An, Ateret Anaby Tavor, Ashish Anand, Avishek Anand, Sophia Ananiadou, Aditya Anantharaman, Nicholas Andrews, Piotr Andruszkiewicz, Rico Angell, Tatiana Anikina, Raghuram Mandyam Annasamy, Alan Ansell, Luca Anselma, Diego Antognini, Jean-Yves Antoine, Abrar Anwar, Tatsuya Aoyama, Mario Ezra Aragon, Arash Ardakani, Arturo

Argueta, Sercan O Arik, Kristijan Armeni, Akhil Arora, Arnav Arora, Daman Arora, Jatin Arora, David Arps, Ekaterina Artemova, Viraat Aryabumi, Abi Aryan, Masaki Asada, Daiki Asami, Benjamin Ascoli, Elliott Ash, Nicholas Asher, Hadi Askari, Thushari Atapattu, Mohammad Atari, Ali Athar, Avinash Atreya, Yash Kumar Atri, Giuseppe Attanasio, Omar Attia, Tal August, AiTi Aw, Hammad Ayyubi, Hosein Azarbonyad, Abdul Hameed Azeemi, Matthias Aßenmacher, Elena Álvarez-Mellado

Nikolay Babakov, Petr Babkin, Stephen Bach, Seohyun Back, Jinheon Baek, Niyati Bafna, Ashutosh Baheti, Mehdi Bahrami, Ge Bai, Jiabin Bai, Jing Bai, Jun Bai, Long Bai, Yu Bai, Yushi Bai, Yuyang Bai, Divya Jyoti Bajpai, Nihal Balani, Ananth Balashankar, Nishant Balepur, Esma Balkir, Simone Balloccu, Lucas Bandarkar, Somnath Banerjee, Namoo Bang, Yejin Bang, Hritik Bansal, Naman Bansal, Srijan Bansal, Forrest Sheng Bao, Guangsheng Bao, Junwei Bao, Xiaoyi Bao, Yu Bao, Yuwei Bao, Francesco Barbieri, Eduard Barbu, Fazl Barez, Elham Barezi, Leslie Barrett, Joe Barrow, Max Bartolo, Marco Basaldella, Elias Bassani, Elisa Bassignana, Osbert Bastani, Samyadeep Basu, Somnath Basu Roy Chowdhury, Luke Bates, Riza Batista-Navarro, Soumya Batra, John Bauer, Tim Baumgärtner, Connor Baumler, Karin Becker, Melika Behjati, Hamid Beigy, Ahmad Beirami, Meriem Beloucif, Boualem Benatallah, Luciana Bencke, Michael Bendersky, Luca Benedetto, Christian Bentz, Gábor Berend, Ben Bergen, Maria Berger, Ian Berlot-Attwell, Guillaume Bernard, Timothée Bernard, Gabriel Bernier-Colborne, Dario Bertero, Amanda Bertsch, Marie Bexte, Anne Beyer, Sourav Sanjukta Bhabesh, Akshita Bhagia, Milan Bhan, Rishabh Bhardwaj, Suma Bhat, Rohan Bhatia, Sumit Bhatia, Shaily Bhatt, Abhik Bhattacharjee, Arnab Bhattacharya, Indrajit Bhattacharya, Sunit Bhattacharya, Prमित Bhattacharyya, Pushpak Bhattacharyya, Sree Bhattacharyya, Ayan Kumar Bhowmick, Mukul Bhutani, Guanqun Bi, Keping Bi, Zhen Bi, Laura Biester, Emil Biju, Yuri Bizzoni, Jakob Brandt Bjorner, Su Lin Blodgett, Jelke Bloem, Benedikt Boenninghoff, Ben Bogin, Ondřej Bojar, Daniele Bonadiman, Helena Bonaldi, Philipp Borchert, Shikha Bordia, Nadav Borenstein, Emanuela Boros, Antoine Bosselut, Kaj Bostrom, Gilles Boulianne, Zied Bouraoui, Andrey Bout, Kevin Bowden, Marco Braga, Ana Brassard, Daniel Braun, Martin Briesch, Thomas Brovelli, Dominique Brunato, Alessio Brutti, Jiajun Bu, Xingyuan Bu, Yuheng Bu, Weerayut Buaphet, Jan Buchmann, Nghi D. Q. Bui, Aljoscha Burchardt, Laurie Burchell, Evgeny Burnaev, Davide Buscaldi, Miriam Butt, Paula Buttery, Jan Buys, Happy Buzaaba, Lisa Bylinina, Bill Byrne, Necva Bölücü

Sky CH-Wang, Isabel Alyssa Cachola, Aoife Cahill, Samuel Cahyawijaya, Hengyi Cai, Li Cai, Longjun Cai, Yujun Cai, Zefan Cai, Agostina Calabrese, Cristian Buc Calderon, Chris Callison-Burch, Jose Camacho-Collados, Giovanni Campagna, Stefano Campese, Leonardo Campillos-Llanos, Bowen Cao, Boxi Cao, Haoyu Cao, Jiangxia Cao, Jie Cao, Juan Cao, Kai Cao, Lele Cao, Meng Cao, Minxuan Cao, Qi Cao, Qingxing Cao, Rui Cao, Yang Trista Cao, Yihan Cao, Yixuan Cao, Yu Cao, Yue Cao, Zhao Cao, Victor Carbune, Ronald Cardenas, Rémi Cardon, Boaz Carmeli, Danilo Carvalho, Silvia Casola, Federico Cassano, Pierluigi Cassotti, Thiago Castro Ferreira, Arie Cattan, Paulo Cavalin, James Caverlee, Francesco Cazzaro, Alba Cercas Curry, Christophe Cerisara, Tanise Ceron, Aman Chadha, Hyunjoo Chae, Yekun Chai, Sunandan Chakraborty, Bharathi Raja Chakravarthi, Chi-Min Chan, Chunkit Chan, David Chan, Hou Pong Chan, Sachin Chanchani, Allen Chang, Buru Chang, Haw-Shiuan Chang, Heng Chang, Kevin Chen-Chuan Chang, Maria Chang, Shuaichen Chang, Ting-Yun Chang, Tyler A. Chang, Yapei Chang, Yin-Wen Chang, Lucas Georges Gabriel Charpentier, Agneet Chatterjee, Soumya Chatterjee, Akshay Chaturvedi, Kushal Chawla, Tong Che, Emmanuel Chemla, Andong Chen, Beiduo Chen, Berlin Chen, Bo Chen, Boxing Chen, Canyu Chen, Catherine Chen, Cen Chen, Chen Chen, Chen Chen, Chien Hung Chen, Daoyuan Chen, Deli Chen, DeLong Chen, Derek Chen, Guanhua Chen, Guanyi Chen, Guanzheng Chen, Guiming Hardy Chen, Haihua Chen, Hailin Chen, Hang Chen, Hanjie Chen, Hao Chen, Haonan Chen, Hong-You Chen, Hui Chen, Jiangjie Chen, Jiaqi Chen, Jiawei Chen, Jiaxiang Chen, Jiayi Chen, Jiefeng Chen, Jindong Chen, Jingchang Chen, Jiuhai Chen, Jiyu Chen, Joya Chen,

Junfan Chen, Junjie Chen, Junying Chen, Junzhe Chen, Justin Chen, Kai Chen, Kejiang Chen, Liang Chen, Liang Chen, Liangyu Chen, Lichang Chen, Lihu Chen, Ling-Hao Chen, Maximillian Chen, Meiqi Chen, Meng Chen, Minshuo Chen, Nuo Chen, Nuo Chen, Nuo Chen, Pinzhen Chen, Po-Chun Chen, Qianglong Chen, Qiguang Chen, Ruijun Chen, Shan Chen, Shuang Chen, Shuguang Chen, Si Chen, Sihao Chen, Sishuo Chen, Tao Chen, Tao Chen, Tong Chen, Wei-Fan Chen, Weize Chen, Wenqing Chen, William Chen, Xiang Chen, Xiaohan Chen, Xiaoyang Chen, Xinyue Chen, Xuanjun Chen, Yanda Chen, Yang Chen, Yi-Pei Chen, Yi-Shin Chen, Yida Chen, Yifang Chen, Yiming Chen, Yingfa Chen, Yiru Chen, Yiwen Chen, Yubo Chen, Yulin Chen, Yun Chen, Yuyan Chen, Zekai Chen, Zeming Chen, Zhengyu Chen, Zhipeng Chen, Zhongwu Chen, Zhouxi Chen, Zhuang Chen, Zhuo Chen, Ziru Chen, Ziyang Chen, Daixuan Cheng, Fan Cheng, Fei Cheng, Gong Cheng, Hong Cheng, Jiale Cheng, Jiali Cheng, Jianpeng Cheng, Minhao Cheng, Pengxiang Cheng, Qinyuan Cheng, Sijie Cheng, Wei Cheng, Xiaoxia Cheng, Xiaoxue Cheng, Xize Cheng, Xuxin Cheng, Yi Cheng, Zhi-Qi Cheng, Zhoujun Cheng, Zifeng Cheng, Colin Cherry, Emmanuele Chersoni, Jackie CK Cheung, Yew Ken Chia, Cheng-Han Chiang, David Chiang, Ting-Rui Chiang, Jenny Chim, Patricia Chiril, Nadezhda Chirkova, Elena Chistova, Hyundong Justin Cho, Hyunsoo Cho, Sukmin Cho, Won Ik Cho, Yae Jee Cho, Young Min Cho, Rochelle Choenni, Jaesik Choi, Jason Ingyu Choi, Joon-Young Choi, Juhwan Choi, Jungwook Choi, Minje Choi, Minseok Choi, Sehyun Choi, Seungtaek Choi, Yong Suk Choi, YunSeok Choi, Shamil Chollampatt, Nurendra Choudhary, Sagnik Ray Choudhury, Fenia Christopoulou, Xu Chu, Yun-Wei Chu, Zheng Chu, Zhixuan Chu, Yu-Neng Chuang, Yung-Sung Chuang, Ekapol Chuangsuwanich, Tsz Ting Chung, Mark Cieliebak, Philipp Cimiano, Christian Clark, Elizabeth Clark, Charles L. A. Clarke, Christopher Clarke, Chloé Clavel, Éric Villemonte De La Clergerie, Stéphane Clinchant, João Coelho, Marco Cognetta, Nachshon Cohen, Shay B Cohen, Nigel Collier, Pierre Colombo, Pedro Colon-Hernandez, Andrei Catalin Coman, Xin Cong, Yan Cong, Hugo Contant, Anna Corazza, Luciano Del Corro, Mickael Coustaty, Fabio Cozman, Mathias Creutz, Maxwell Crouse, Emilio Villa Cueva, JJ Cui, Peng Cui, Shiyao Cui, Wanyun Cui, Wendi Cui, Washington Cunha, Erion Çano

Ido Dagan, Nico Daheim, Hanjun Dai, Hong-Jie Dai, Hongliang Dai, Jianbo Dai, Lirong Dai, Qin Dai, Shih-Chieh Dai, Wen Dai, Xiang Dai, Yinpei Dai, Yong Dai, Dhairya Dalal, David Dale, Majid Daliri, Mariana Damova, Soham Dan, Yuhao Dan, Sandipan Dandapat, John Dang, Rumen Dangovski, Debarati Das, Mayukh Das, Souvik Das, Mihai Dascalu, Sourish Dasgupta, Hal Daurmé Iii, Brian Davis, Ernest Davis, Heidar Davoudi, Hillary Dawkins, Gaël De Chalendar, Gerard De Melo, Alexandra DeLucia, Steve DeNeefe, Nicholas Deas, Alok Debnath, Mathieu Dehouck, Björn Deiseroth, Jean-Benoit Delbrouck, Pieter Delobelle, Gianluca Demartini, Daryna Dementieva, David Demeter, Çağatay Demiralp, Cheng Deng, Chenlong Deng, Chunyuan Deng, Mingkai Deng, Naihao Deng, Yue Deng, Zhi-Hong Deng, Pavel Denisov, Franck DERNONCOURT, Aniket Derooy, Ameet Deshpande, Vijeta Deshpande, Nina Dethlefs, Suvodip Dey, Jwala Dhamala, Prajit Dhar, Luigi Di Caro, Luca Di Liello, Mona T. Diab, Fernando Diaz, Harshita Diddee, Jana Diesner, Stefan Dietze, Dimitris Dimakopoulos, Elizabeth A Dinella, Bosheng Ding, Caiwen Ding, Hantian Ding, Peng Ding, Wenjian Ding, Wentao Ding, Wenxuan Ding, Yifeng Ding, Yihao Ding, Yuyang Ding, Song Dingjie, Saket Dingliwal, Tuan Dinh, Tanvi Dinkar, Anuj Diwan, Amirbek Djanibekov, Nemanja Djuric, Anna Dmitrieva, Heejin Do, Phong Nguyen-Thuan Do, Sumanth Doddapaneni, Charles Dognin, Pierre Dognin, Ljiljana Dolamic, Miguel Domingo, Shachar Don-Yehiya, Bowen Dong, Guanting Dong, Hang Dong, Haoyu Dong, Kuicai Dong, Manqing Dong, Ming Dong, Mingwen Dong, Peijie Dong, Wenchao Dong, Xiangjue Dong, Yuxiao Dong, Yuyang Dong, Giovanna Maria Dora Dore, Bonaventure F. P. Dossou, Jordan Dotzel, Guangyao Dou, Longxu Dou, Qingyun Dou, Shihan Dou, Yao Dou, Jad Doughman, Mark Dras, Mark Dredze, Felix Drinkall, Rotem Dror, Haowei Du, Mengfei Du, Mengnan Du, Tianyu Du, Weihong Du, Weiyu Du, Wenyu Du, Yufeng Du, Yupei Du, Zhengxiao Du, Hanyu Duan, Junwen Duan, Sufeng Duan, Kumar Avinava Dubey, Sai Meher Karthik Duddu, Richard Dufour, Philipp Dufter, Liam Dugan, Lavinia Dunagan, Marcel Dunaiski, Ewan Dunbar, Emmanuel Dupoux, Ritam Dutt, Subhabrata

Dutta, Jane Dwivedi-Yu, Nouha Dziri, Hervé Déjean, Esra Dönmez, Iria de-Dios-Flores

Abteen Ebrahimi, Ali Edalati, Lilach Eden, Tobias Eder, Aleksandra Edwards, Steffen Eger, Koji Eguchi, Yo Ehara, Max Eichler, Hafsteinn Einarsson, Roald Eiselen, Roxanne El Baff, Mohamed Elaraby, Maha Elbayad, Mohamed Elgaar, Desmond Elliott, AbdelRahim A. Elmadany, Micha El-sner, Ali Emami, Yahya Emara, Elena V. Epure, Aykut Erdem, Erkut Erdem, Justus-Jonas Erker, Beyza Ermis, Patrick Ernst, Carlos Escolano, Andrea Esuli, Ana Ezquerro

Alex Fabrikant, Marzieh Fadaee, Fahim Faisal, Neele Falk, Chuchu Fan, Ge Fan, Kai Fan, Rui Fan, Shengda Fan, Ting-Han Fan, Yao-Chung Fan, Yixing Fan, Biaoyan Fang, Haishuo Fang, Jinyuan Fang, Qingkai Fang, Tao Fang, Tianqing Fang, Yanbo Fang, Yin Fang, Yuejian Fang, Yuwei Fang, Youmna Farag, Marco Farina, António Farinhas, Nawshad Farruque, Farima Fatahi Bayat, Adam Faulkner, Pedro Faustini, Mohsen Fayyaz, Wilson Fearn, Joshua Feinglass, Sophie Fellenz, Dongji Feng, Jiangtao Feng, Jiazhan Feng, Jingrong Feng, Qianyu Feng, Shangbin Feng, Shi Feng, Steven Y. Feng, Wanyong Feng, Yanlin Feng, Yansong Feng, Yi Feng, Yi Feng, Yu Feng, Yujie Feng, Zeyu Feng, Zihao Feng, Alena Fenogenova, Elisa Ferracane, Javier Ferrando, Olivier Ferret, Besnik Fetahu, Constanza Fierro, James D. Finch, Matthew Finlayson, Hamed Firooz, Jack Fitzgerald, Margaret M. Fleck, Eve Fleisig, Riccardo Fogliato, Antske Fokkens, José A.r. Fonollosa, Brian Formento, Negar Foroutan, James R. Foulds, Meaghan Fowlie, Thomas François, Kathleen C. Fraser, Diego Frassinelli, Dayne Freitag, Andre Freitas, Harvey Yiyun Fu, Jinlan Fu, Luoyi Fu, Shuai Fu, Xingyu Fu, Xiyan Fu, Yi-Fu Fu, Yingxue Fu, Yu Fu, Dennis Fucci, Yoshinari Fujinuma, Fumiyo Fukumoto, Kotaro Funakoshi, Hiroaki Funayama

Kamel Gaanoun, Matteo Gabburo, Marco Gaido, Jay Gala, Andrew Gambardella, Leilei Gan, Yiming Gan, Yujian Gan, Zhe Gan, Sudeep Gandhe, Vineet Gandhi, Prakhar Ganesh, Revanth Gangi Reddy, Niloy Ganguly, William Gantt, Chaochen Gao, Hang Gao, Hongcheng Gao, Jinglong Gao, Jingsheng Gao, Jun Gao, Jun Gao, Mingqi Gao, Pengzhi Gao, Shen Gao, Songyang Gao, Tianyu Gao, Wei Gao, Yan Gao, Yanjun Gao, Yifan Gao, Yingbo Gao, Ze-Feng Gao, Utpal Garain, Marcos Garcia, Iker García-Ferrero, Claire Gardent, Shubham Garg, Kiril Gashteovski, Albert Gatt, Joseph Gatto, Manas Gaur, Vagrant Gautam, Jiabin Ge, Suyu Ge, Xiou Ge, Yubin Ge, Gregor Geigle, Matthieu Geist, Aryo Pradipta Gema, Josef Van Genabith, Jiahui Geng, Saibo Geng, Kallirroi Georgila, Ariel Gera, Stefan Gerdjikov, Sadaf Ghaffari, Omid Ghahroodi, Sarik Ghazarian, Mozhdeh Gheini, Iacopo Ghinassi, Kripabandhu Ghosh, Reshmi Ghosh, Satanu Ghosh, Sayan Ghosh, Sayan Ghosh, Sayontan Ghosh, Shinjini Ghosh, Soumitra Ghosh, Sourav Ghosh, Sreyan Ghosh, Sucheta Ghosh, Caroline Gihlstorf, Michael Ginn, Salvatore Giorgi, Leander Girrback, Alex Gittens, Michael Glass, Amir Globerson, Hyojun Go, Ameya Godbole, Anmol Goel, Vaibhava Goel, Preni Golazizian, Yoav Goldberg, Jonas Golde, Tomas Goldsack, Janis Goldzycher, Olga Golovneva, Helena Gomez Adorno, Chen Gong, Chengyue Gong, Jiaying Gong, Linyuan Gong, Shansan Gong, Zhuocheng Gong, Cesar Gonzalez-Gutierrez, Julio Gonzalo, Carlos-Emiliano González-Gallardo, Marcos André Gonçalves, Maharshi Gor, Philip John Gorinski, Dhiman Goswami, Koustava Goswami, Isao Goto, Yunhao Gou, Venkata S Govindarajan, Abhimanyu Goyal, Navita Goyal, Arianna Graciotti, Tommaso Green, Derek Greene, Matt Grenander, Thomas L. Griffiths, Cyril Grouin, Max Grusky, Adam Grycner, Marek Grzes, Chenxi Gu, Jiasheng Gu, Jindong Gu, Jing Gu, Naibin Gu, Nianlong Gu, Shuhao Gu, Tianlong Gu, Xiangming Gu, Yi Gu, Yuxian Gu, Yuxuan Gu, Zijin Gu, Renchu Guan, Ziyu Guan, Reto Gubelmann, Bhanu Prakash Reddy Guda, Imane Guellil, Marco Guerini, Nuno M Guerreiro, Camille Guinaudeau, Caglar Gulcehre, Varun Gumma, Aylin Ece Gunal, Kalpa Gunaratna, Beliz Gunel, Sharath Chandra Guntuku, Fang Guo, Han Guo, Hongcheng Guo, Jialiang Guo, Jiang Guo, Lingbing Guo, Meiqi Guo, Quan Guo, Ruohao Guo, Shoutao Guo, Wenya Guo, Xiaobo Guo, Xingang Guo, Xu Guo, Yanzhu Guo, Yi Guo, Yue Guo, Yue Guo, Yuhang Guo, Zhen Guo, Zhicheng Guo, Zhihui Guo, Ziyu Guo, Abhirut Gupta, Akshat Gupta, Ankita Gupta, Himanshu Gupta, Itika Gupta,



Prakhar Gupta, Priyanshu Gupta, Raghav Gupta, Umang Gupta, Sairam Gurajada, V Gurucharan, Sireesh Gururaja, Yoan Gutierrez, Jeremy Gwinnup

Matthias Hagen, Lovisa Hagström, Joonghyuk Hahn, Richard Hahnloser, Samar Haider, Huda Hakami, John T. Hale, Coleman Haley, Patrick Haller, Patrick Haller, Skyler Hallinan, Felix Hamborg, Hasan Abed Al Kader Hammoud, Cheng Han, Chengcheng Han, Chi Han, Guangzeng Han, Hojae Han, Jiale Han, Jiuzhou Han, Kelvin Han, Kyungsik Han, Namgi Han, Rujun Han, Shanshan Han, Shijie Han, Simeng Han, Wei Han, Wookje Han, Xiaochuang Han, Yo-Sub Han, Yucheng Han, Zhen Han, Allan Hanbury, Michael Hanna, Hongkun Hao, Jie Hao, Jie Hao, Shibo Hao, Mirazul Haque, Momchil Hardalov, Camille Harris, Ian Harris, Thomas Hartvigsen, Mohammad Hasan, Taku Hasegawa, Farnoosh Hashemi, Kazuma Hashimoto, Sabit Hassan, Oktie Hassanzadeh, Michael Hassid, Nabil Hathout, Hans Ole Hatzel, William N. Havard, Robert D. Hawkins, Yoshihiko Hayashi, Amir Hazem, Chaoqun He, Guoxiu He, Han He, Hangfeng He, Hongliang He, Jianfeng He, Jie He, Junwei He, Kun He, Pengfei He, Pengfei He, Qianyu He, Ruifang He, Shwai He, Tao He, Tianxing He, Tingting He, Xiangheng He, Xuanli He, Xuming He, Yujie He, Yuxin He, Zhaofeng He, Zhengfu He, Zhiwei He, Zihao He, Tobias Hecking, Michael A. Hedderich, Johannes Heinecke, Benjamin Heinzerling, Jindřich Helcl, James Henderson, Simon Hengchen, Amey Hengle, Aron Henriksson, Jonathan Herzig, Gerhard Heyer, Rem Hida, Tsutomu Hirao, Tatsuya Hiraoka, Toshio Hirasawa, Hlee Hlee, Namgyu Ho, Xanh Ho, William P Hogan, Carolin Holtermann, Nils Holzenberger, Lingzi Hong, Pengfei Hong, Ruixin Hong, Seokhee Hong, Xudong Hong, Or Honovich, Andrea Horbach, Thanapapas Horsuwan, Sho Hoshino, Mohammad Javad Hosseini, Pedram Hosseini, Guiyang Hou, Jinyu Hou, Jue Hou, Lei Hou, Yifan Hou, Yu Hou, Yuexian Hou, Zejiang Hou, Zhaoyi Joey Hou, Phillip Howard, David M Howcroft, Alexander Hoyle, Claudiu Daniel Hromei, Jennifer Hsia, Cheng-Yu Hsieh, Benjamin Hsu, Yen-Chang Hsu, Yu-Yin Hsu, Anwen Hu, Bozhen Hu, Chi Hu, Dou Hu, Guangneng Hu, Hailin Hu, Jinpeng Hu, Jinyi Hu, Linmei Hu, Minghao Hu, Shengding Hu, Songlin Hu, Tao Hu, Tiancheng Hu, Tianxi Hu, Wei Hu, Wenpeng Hu, Wenxiang Hu, Xiang Hu, Xiangkun Hu, Xinyu Hu, Yebowen Hu, Yibo Hu, Yiyan Hu, Yuchen Hu, Zhe Hu, Zhengyu Hu, Zhiqiang Hu, Zhiyuan Hu, Hang Hua, Tianze Hua, Ting Hua, Wenyue Hua, Yilun Hua, Yuncheng Hua, Baizhou Huang, Baorong Huang, Chao-Wei Huang, Chen Huang, Chengsong Huang, Chenyang Huang, Chieh-Yang Huang, Churen Huang, Dong Huang, Fei Huang, Guangyan Huang, Guanhua Huang, Haibin Huang, Haojing Huang, Hen-Hsen Huang, Heyan Huang, Hui Huang, Ian Huang, James Y. Huang, Jen-tse Huang, Jiaxin Huang, Jimin Huang, Jin-Xia Huang, Li Huang, Minghui Huang, Qingbao Huang, Qiushi Huang, Quzhe Huang, Rongjie Huang, Ruizhe Huang, Shaoyao Huang, Shen Huang, Shijue Huang, Shiyuan Huang, Tenghao Huang, Xiao Huang, Xiaolei Huang, Xijie Huang, Xinting Huang, Xiusheng Huang, Yichong Huang, Yin Jou Huang, Yizheng Huang, Yufei Huang, Yuzhen Huang, Zhenya Huang, Zhichao Huang, Zhiqi Huang, Zhongqiang Huang, Zhongzhan Huang, Zijie Huang, Patrick Huber, Pere-Lluís Huguet Cabot, Bo Hui, Tingfeng Hui, Siyu Huo, Fantine Huot, Lluís-F. Hurtado, Ben Hutchinson, Du Q. Huynh, Tin Van Huynh, Rebecca Hwa, Dae Yon Hwang, EunJeong Hwang, Jena D. Hwang, Seonjeong Hwang, Dongmin Hyun, Katharina Hämmerl

Robert L. Logan IV, Ignacio Iacobacci, Andreea Iana, Timour Igamberdiev, Oana Ignat, Taichi Iki, Evgenia Ilia, Filip Ilievski, Nikolai Ilinykh, Dmitry Ilvovsky, Vaiva Imbrasaitė, Mert Inan, Sathish Reddy Indurthi, Go Inoue, Kazuki Irie, Yoichi Ishibashi, Tatsuya Ishigaki, Etsuko Ishii, Raima Islam, Riashat Islam, Tunazzina Islam, Hayate Iso, Masaru Isonuma, Abraham Israeli, Itay Itzhak, Hamish Ivison, Alexandra Ivoylova, Yusuke Iwasawa, Bhavani Iyer

Alon Jacovi, Guillaume Jacquet, Aditya Ajay Jadhav, Labiba Jahan, Naman Jain, Nihal Jain, Parag Jain, Rishabh Jain, Vinija Jain, Ajay Kumar Jaiswal, Masoud Jalili Sabet, Abhik Jana, Eugene Jang, Hyewon Jang, Jin Yea Jang, Yoonna Jang, Ali Jannesari, Adam Jatowt, Sujay Kumar Jau-

har, Tommi Jauhiainen, Inigo Jauregi Unanue, Dávid Javorský, Ganesh Jawahar, Sébastien Jean, Fran Jelenić, Christopher William Jenkins, Eojin Jeon, Yongkweon Jeon, Minbyul Jeong, Soyeong Jeong, Sullam Jeoung, Prince Jha, Somesh Jha, Shaoxiong Ji, Tao Ji, Tianbo Ji, Wei Ji, Yandong Ji, Yixin Ji, Yuelu Ji, Ziwei Ji, Chen Jia, Furong Jia, Qi Jia, Qinjin Jia, Ruoxi Jia, Xu Jia, Yuxiang Jia, Zixia Jia, Ping Jian, Xiangru Jian, Yiren Jian, Chao Jiang, Chunyang Jiang, Dongfu Jiang, Fan Jiang, Feng Jiang, Gongyao Jiang, Hang Jiang, Huiqiang Jiang, Jing Jiang, Jiyue Jiang, Jyun-Yu Jiang, Lei Jiang, Lili Jiang, Meng Jiang, Ming Jiang, Nan Jiang, Nan-Jiang Jiang, Ning Jiang, Song Jiang, Tianyu Jiang, Xiaotong Jiang, Xin Jiang, Xuhui Jiang, Ye Jiang, Yuxin Jiang, Zhengping Jiang, Zhihua Jiang, Zhiying Jiang, Zhongtao Jiang, Zhuoren Jiang, Zifan Jiang, Ziyang Jiang, Cathy Jiao, Fangkai Jiao, Wenxiang Jiao, Yizhu Jiao, Cheng Jiayang, Bernal Jimenez Gutierrez, Bowen Jin, Feihu Jin, Hongye Jin, Li Jin, Lianwen Jin, Lisa Jin, Peng Jin, Peng Jin, Qiao Jin, Renren Jin, Xiaolong Jin, Xiaomeng Jin, Xisen Jin, Yiping Jin, Yiqiao Jin, Zhijing Jin, Zhuoran Jin, Zijian Jin, Ishan Jindal, Baoyu Jing, Liqiang Jing, Hwiyeol Jo, Jaehyeong Jo, Shailza Jolly, Abhinav Joshi, Aditya Joshi, Anant Joshi, Brihi Joshi, Harshit Joshi, Pratik Joshi, Yiming Ju, Josip Jukić, Dongwon Jung, Simeon Junker, Juraj Juraska, David Jurgens, Gerhard Jäger

Karthikeyan K, Mohsinul Kabir, Anubha Kabra, Lucie-Aimée Kaffee, Indika Kahanda, Hour Kaing, Kaisla Kajava, Ivana Kajic, Dariusz Kajtoch, Mihir Kale, Oren Kalinsky, Ryo Kamoi, Jaap Kamps, Andrey Kan, Min-Yen Kan, Christopher Kanan, Hiroshi Kanayama, Teja Kanchinadam, Kamil Kanclerz, Feiyang Kang, Junmo Kang, Minki Kang, Xiaomian Kang, Xiaoxi Kang, Zhao Kang, Karthik Reddy Kanjula, Nithish Kannen, Surya Kanoria, Ben Kao, Debanjana Kar, Pinar Karagoz, Antonia Karamolegkou, Amir Hossein Kargar, Priyanka Kargupta, Taelin Karidi, Constantinos Karouzos, Fnu Kartik, Omid Kashefi, Marc A. Kastner, Tushar Kataria, Denys Katerenchuk, Anisia Katinskaia, Uri Katz, Navdeep Kaur, Pride Kavumba, Daisuke Kawahara, Noriaki Kawamae, Yoshifumi Kawasaki, Hideto Kazawa, Ashkan Kazemi, Abe Kazemzadeh, Nazmul Kazi, Cai Ke, KediChen KediChen, Sedrick Keh, Amr Keleg, Mikaela Keller, Roman Kern, Natthawut Kertkeidkachorn, Phillip Keung, Kurt Keutzer, Tannon Kew, Lee Kezar, Baber Khalid, Muhammad Khalifa, Talaat Khalil, Abdul Rafae Khan, Haidar Khan, Mohammad Aflah Khan, Shima Khanehzar, Simran Khanuja, Omar Khattab, Subhendu Khatuya, Vivek Khetan, Ashiqur R. KhudaBukhsh, Urja Khurana, Varun Khurana, Dayeon Ki, Bugeun Kim, Byeongwook Kim, Byoungjip Kim, Dahyun Kim, Dong-Jin Kim, Donghyun Kim, Doyoung Kim, Gangwoo Kim, Geewook Kim, Gene Louis Kim, Gyuwan Kim, Hwichan Kim, Hyunjae Kim, Jaehyung Kim, Jaemin Kim, Jangwon Kim, Jeonghoon Kim, Jeonghwan Kim, Jiho Kim, Jihyuk Kim, Jin-Dong Kim, Jinsung Kim, Jiseon Kim, Jongho Kim, Jung-jae Kim, Junho Kim, Junu Kim, Junyeob Kim, Junyeong Kim, Juyong Kim, Kang-Min Kim, Michelle YoungJin Kim, Minbeom Kim, Minsoo Kim, Minsoo Kim, Misuk Kim, Seong Tae Kim, Seungbae Kim, Seungone Kim, Sungchul Kim, Sungdong Kim, Taehwan Kim, Taehyeon Kim, Taesup Kim, Takyoung Kim, Woojeong Kim, Ye- chan Kim, Yejin Kim, Yoon Kim, YoungBin Kim, Youngwoo Kim, Yumin Kim, Yasutomo Kimura, Tracy Holloway King, Alon Kipnis, Svetlana Kiritchenko, Christo Kirov, Hirokazu Kiyomaru, Shun Kiyono, Christopher Klamm, Jens Kleesiek, Ayal Klein, Jacques Klein, Tassilo Klein, Jan-Christoph Klie, Mateusz Klimaszewski, Julien Kloetzer, Marius Kloft, Rene Knaebel, Dohwan Ko, Jongwoo Ko, Miyoung Ko, Sang-Ki Ko, Youngjoong Ko, Ichiro Kobayashi, Naoki Kobayashi, Thomas H Kober, Elena Kochkina, Prashant Kodali, Jordan Kodner, Konstantinos Kogkalidis, Hyukhun Koh, Ryosuke Kohita, Satoshi Koide, Ryuto Koike, Mare Koit, Takeshi Kojima, Jachym Kolar, Alexander Koller, Kanako Komiya, Sai Koneru, Fang Kong, Zhenglun Kong, Myoung-Wan Koo, Seonmin Koo, Bevan Koopman, Michalis Korakakis, Yuta Koreeda, Mandy Barrett Korpusik, Katerina Korre, Leila Kosseim, Suhas Kotha, Fajri Koto, Ziyi Kou, George Kour, Punit Singh Koura, Adriana Kovashka, Pranoy Kovuri, Ivan Koychev, Elisa Kreiss, Matt Kretchmar, Kundan Krishna, Satyapriya Krishna, Adit Krishnan, Reno Kriz, Jason Krone, Haoyu Kuang, Kun Kuang, Andrei Kucharavy, Sneha Kudugunta, Alapan Kuila, Adithya Kulkarni, Mayank Kulkarni, Nitish Kulkarni, Devang Kulshreshtha, Rishabh Kumar, Sawan Kumar, Shanu Kumar, Shivani Kumar,

Sonal Kumar, Sujay S Kumar, Vaibhav Kumar, Varun Kumar, Vineet Kumar, Gitanjali Kumari, Lilly Kumari, Florian Kunneman, Yuri Kuratov, Shuhei Kurita, Kemal Kurniawan, Robin Kurtz, Wojciech Kusa, Guy Kushilevitz, Laida Kushnareva, Andrey Kutuzov, Saar Kuzi, Gleb Kuzmin, Henry Kvinge, Jin Myung Kwak, Wai-Chung Kwan, Deuksin Kwon, Jingun Kwon, Yeong-Dae Kwon, Abdullatif Köksal, Sandra Kübler

Peifeng LI, Xianming LI, Yucheng LI, Zexin LU, Emanuele La Malfa, Jack LaViolette, Sofie Labat, Vincent Labatut, Matthieu Labeau, Yanis Labrak, Tiziano Labruna, Avishek Lahiri, Salem Lahlou, Allison Lahnala, Bolin Lai, Huiyuan Lai, Viet Dac Lai, Zhixin Lai, Kushal Lakhota, Yash Kumar Lal, John P. Lator, Jessica Lam, Monica Lam, Hemank Lamba, Andrew Lan, Yunshi Lan, Zhenzhong Lan, Paul Landes, Ian Lane, Qiwei Lang, Mateusz Lango, Gabriella Lapesa, Daniil Larionov, Samuel Larkin, Stefan Larson, Md Tahmid Rahman Laskar, Sahinur Rahman Laskar, Mark Last, Jey Han Lau, Leo Laugier, Antonio Laverghetta Jr., Duc-Trong Le, Dung D. Le, Duy Le, Phong Le, Phuong-Hang Le, Thai Le, Kevin Leach, Matthew Lease, Bruce W. Lee, Changmin Lee, Chen-Yu Lee, Cheryl Lee, Chungman Lee, Daeun Lee, Dohyeon Lee, Dong Won Lee, Dongha Lee, Dongkyu Lee, Dongwon Lee, En-Shiun Annie Lee, Gyubok Lee, Hongrae Lee, Hung-yi Lee, Hyunju Lee, Jaeseong Lee, Janghwan Lee, Jay-Yoon Lee, Ji-Ung Lee, Jongwuk Lee, Jooyoung Lee, Jung Hyun Lee, Jungseob Lee, Kang-il Lee, Katherine Lee, Kyumin Lee, Mingyu Lee, Minhwa Lee, Mong-Li Lee, Sang-goo Lee, Sanghack Lee, Seanie Lee, Seolhwa Lee, Seonghyeon Lee, Taesung Lee, Wee Sun Lee, Wonkwang Lee, Yongjae Lee, Young-Jun Lee, Youngwon Lee, Yukyung Lee, Yunsung Lee, Tobias Leemann, Joël Legrand, Jens Lehmann, Bin Lei, Chuan Lei, Fangyu Lei, Jie Lei, Shuo Lei, Yibin Lei, Juho Leinonen, Xingjian Leng, Yichong Leng, Elisa Leonardelli, Paul Lerner, Pietro Lesci, Ulf Leser, James Lester, Johannes Leveling, Rivka Levitan, Mosh Levy, Sharon Levy, Martha Lewis, Bangzheng Li, Baohang Li, Bei Li, Belinda Z. Li, Bobo Li, Bowen Li, Bryan Li, Changchun Li, Changmao Li, Changye Li, Chaozhuo Li, Chengming Li, Chengxi Li, Chengzu Li, Chong Li, Chuanyi Li, Chun-Liang Li, Chunping Li, Dawei Li, Dingcheng Li, Diya Li, Dong Li, Dongsheng Li, Dongyang Li, Dongyuan Li, Dongyue Li, Fanrong Li, Feng-Lin Li, Guohui Li, Haau-Sing Li, Hao Li, Haochen Li, Haonan Li, Haoran Li, Haoyuan Li, Hebi Li, Hongsheng Li, Huao Li, Hui Li, Huihan Li, Huiying Li, Irene Li, Jiachen Li, Jiahuan Li, Jialu Li, Jianjun Li, Jiayi Li, Jiayuan Li, Jiazhao Li, Jiazheng Li, Jierui Li, Jing Li, Jinpeng Li, Jinyuan Li, Judith Yue Li, Jundong Li, Junlong Li, Junyi Li, Lei Li, Lei Li, Li Li, Li Erran Li, Liang Li, Lin Li, Linjing Li, Lishuang Li, Lujun Li, Lusi Li, Mengze Li, Miao Li, Ming Li, Mingchen Li, Mingzhe Li, Minzhi Li, Moxin Li, Mukai Li, Panfeng Li, Ping Li, Qi Li, Qian Li, Qian Li, Qicheng Li, Qintong Li, Qiuchi Li, Ruifan Li, Ruizhe Li, Ruosen Li, Shaobo Li, Sheng Li, Shengjie Li, Shicheng Li, Shuai Li, Shuaiyi Li, Shuo Li, Shuyang Li, Shuyue Stella Li, Sihang Li, Siyan Li, Tianjian Li, Tianrui Li, Tianyi Li, Tongliang Li, Wei Li, Wei Li, Wen-Ding Li, Wenhao Li, Wenyan Li, Xia Li, Xian Li, Xiang Li, Xiang Li, Xiangci Li, Xiao Li, Xiaonan Li, Ximing Li, Xin Li, Xingxuan Li, Xinlin Li, Xintong Li, Xinze Li, Yafu Li, Yang Li, Yanhong Li, Yanzeng Li, Ying Li, Ying Li, Yinghao Li, Yinghui Li, Yingting Li, Yingya Li, Yinqiao Li, Yitong Li, Yiwei Li, Yiyuan Li, Yongqi Li, Yuan Li, Yunshui Li, Zajing Li, Zekun Li, Zexuan Li, Zhen Li, Zhongli Li, Zhoujun Li, Zhuang Li, Zhuohang Li, Zihao Li, Ziheng Li, Zixuan Li, Ziyang Li, Zongxi Li, Zuchao Li, Jianxun Lian, Zheng Lian, Chao-Chun Liang, Davis Liang, Jiaqing Liang, KE Liang, Tian Liang, Xiaobo Liang, Xiaozhuan Liang, Yuedi Liang, Yunlong Liang, Zhenwen Liang, Baohao Liao, Hao Liao, I-Bin Liao, Minpeng Liao, Qing Liao, Ruotong Liao, Xinting Liao, Xixian Liao, Alexander Libov, Jindřich Libovický, Veronica Liesaputra, Gilbert Lim, Jia Peng Lim, Jungwoo Lim, KyungTae Lim, Sungbin Lim, Tomasz Limisiewicz, Peerat Limkonchotiwat, Binghuai Lin, Bingqian Lin, Chih-Jen Lin, Haowei Lin, Hongzhan Lin, Hsien-chin Lin, Huawei Lin, Jianghao Lin, Ke Lin, Li Lin, Minhua Lin, Peiqin Lin, Qika Lin, Qingwei Lin, Ting-En Lin, Weizhe Lin, Xudong Lin, Xueyuan Lin, Yen-Ting Lin, Ying-Jia Lin, Yong Lin, Yu-Hsiang Lin, Zeqi Lin, Zhe Lin, Zhen Lin, Zheng Lin, Zhenxi Lin, Zhouhan Lin, Zicheng Lin, Zijia Lin, Zongyu Lin, Matthias Lindemann, Zhen-Hua Ling, Gili Lior, Aldo Lipani,

Tom Lippincott, Enrico Liscio, Marina Litvak, Aiwei Liu, Anqi Liu, Ben Liu, Bing Liu, Bo Liu, Bo Liu, Boyang Liu, Chen Cecilia Liu, Chengyuan Liu, Daizong Liu, Danni Liu, Dugang Liu, Emmy Liu, Fenglin Liu, Fuxiao Liu, Gongshen Liu, Guangliang Liu, Guangyi Liu, Hanmeng Liu, Haoyu Liu, Hongfu Liu, Hongyi Liu, Hui Liu, Hui Liu, Jiacheng Liu, Jiaheng Liu, Jiangming Liu, Jiateng Liu, Jiawei Liu, Jie-jyun Liu, Jiongnan Liu, Juhua Liu, Jun Liu, Junteng Liu, Lingqiao Liu, Maofu Liu, Ming Liu, Mingrui Liu, Minqian Liu, Nelson F. Liu, Ning Liu, Ollie Liu, Peiyang Liu, Peiyu Liu, Peiyu Liu, Peng Liu, Qian Liu, Qianchu Liu, Qiang Liu, Qin Liu, Qingbin Liu, Ran Liu, Ruiyang Liu, Shilei Liu, Shiwei Liu, Sijia Liu, Sijia Liu, Siyi Liu, Siyi Liu, Tengxiao Liu, Tianqi Liu, Tianyang Liu, Tianyu Liu, Timothy Liu, Tong Liu, Wei Liu, Wei Liu, Wei Liu, Weifeng Liu, Weisi Liu, Xiang Liu, Xiangyang Liu, Xiao Liu, Xiaolong Liu, Xiaoming Liu, Xiaoze Liu, Xiaozhong Liu, Xin Liu, Xiping Liu, Xiuwen Liu, Yanchen Liu, Yang Janet Liu, Yanming Liu, Ye Liu, Ye Liu, Yihong Liu, Yinan Liu, Yinhong Liu, Yinxiao Liu, Yixin Liu, Yongbin Liu, Yonghao Liu, Yongmei Liu, Yuanchao Liu, Yuanxing Liu, Yuchen Liu, Yuhan Liu, Yujian Liu, Zeming Liu, Zhe Liu, Zhengyuan Liu, Zhexiong Liu, Zheyuan Liu, Zhi Liu, Zhiwei Liu, Zihan Liu, Ziyi Liu, Zoey Liu, Zuozhu Liu, Zuxin Liu, Adian Liusie, Tien-Hong Lo, Tyler Loakman, Mengsay Loem, Lajanugen Logeswaran, Do Xuan Long, Quanyu Long, Siyu Long, Wanqiu Long, Xinwei Long, Yinghan Long, Yunfei Long, Henrique Lopes Cardoso, Alejo Lopez-Avila, Michela Lorandi, Chao Lou, Qian Lou, Natalia V Loukachevitch, Holy Lovenia, Bo-Ru Lu, Chaochao Lu, Di Lu, Hengtong Lu, Hongyuan Lu, Jianqiao Lu, Jiarui Lu, Jiaying Lu, Jinghui Lu, Jinliang Lu, Junru Lu, Junyu Lu, Kaiji Lu, Keming Lu, Peng Lu, Qingyu Lu, Qiu hao Lu, Sidi Lu, Weiming Lu, Wenpeng Lu, Xiaoyu Lu, Xin Lu, Xingyu Lu, Xinyu Lu, Xinyuan Lu, Xuesong Lu, Yi Lu, Yujie Lu, Yuyin Lu, Evan Lucas, Li Lucy, Bernd Ludwig, Dan Luo, Dongsheng Luo, Feng Luo, Ge Luo, Guoqing Luo, Haoran Luo, Hongyin Luo, Jiaming Luo, Jixiang Luo, Junyu Luo, Lannan Luo, Ling Luo, Linhao Luo, Man Luo, Xiang Luo, Xiangyang Luo, Xiao Luo, Yanchen Luo, Yiran Lawrence Luo, Yong Luo, Zhekun Luo, Ziyang Luo, Ronny Luss, Pedro Henrique Luz De Araujo, Ang Lv, Shangwen Lv, Xingtai Lv, Zheqi Lv, Chunchuan Lyu, Weimin Lyu, Xinglin Lyu, Xinxi Lyu, Yiwei Lyu, Xing Han Lù, Yunxin li

Yaswanth M, Meryem M'hamdi, Da Ma, Fukun Ma, Guangyuan Ma, Jie Ma, Jiefeng Ma, Liang Ma, Lizhi Ma, MingHua Ma, Mingyu Derek Ma, Qianli Ma, Shiqing Ma, Shuming Ma, Xinbei Ma, Xinyin Ma, Xueguang Ma, Yao Ma, Youmi Ma, Yubo Ma, Yunpu Ma, Zhengrui Ma, Ziqiao Ma, Sean MacAvaney, Jakub Macina, Dominik Macko, Mounica Maddela, Rahul Madhavan, Tharindu Madusanka, Aru Maekawa, Joao Magalhaes, Marco Maggini, Ian Magnusson, Ahmed Magooda, Khyati Mahajan, Shweti Mahajan, Saad Mahamood, Kishan Maharaj, Debanjan Mahata, Rahmad Mahendra, Ayush Maheshwari, Aishwarya Maheswaran, Florian Mai, Fred Mairhot, Mieradilijiang Maimaiti, Aviya Maimon, Gallil Maimon, Brian Mak, Ilya Makarov, Aaron Maladry, Luca Malagutti, Bhavitvya Malik, Vijit Malik, Sri Raghu Malireddi, Advait Malladi, Ganesh Satish Mallya, Christopher Malon, Valentin Malykh, Jonathan Mamou, Hieu Man, Pottsawee Manakul, Pranav Maneriker, Pranav Mani, Matteo Manica, Dinesh Manocha, Kelong Mao, Qianren Mao, Rui Mao, Shaoguang Mao, Wenji Mao, Xin Mao, Yi Mao, Yu Mao, Yunlong Mao, Zhendong Mao, Zhiming Mao, Ana Marasovic, Marta Marchiori Manerba, Katerina Margatina, Alda Mari, Ilia Markov, Magdalena Markowska, Edison Marrese-Taylor, Brais Martinez, Abelardo Carlos Martinez Lorenzo, Bruno Martins, Pekka Marttinen, Marcos Martínez Galindo, Claudia Marzi, Mihai Masala, Mostafa Masoudi, Ahmed Masry, Sarah Masud, Ved Mathai, Brodie Mather, Sandeep Mathias, Puneet Mathur, Shivam Mathur, Justus Mattern, Kaushal Kumar Maurya, Nickil Maveli, Costas Mavromatis, Diana Maynard, Mantas Mazeika, Alessandro Mazzei, Kate McCurdy, Daniel McDuff, Nicholas Meade, Houman Mehrafarin, Kai Mei, John Mendonça, De Meng, Rui Meng, Shiao Meng, Tao Meng, Yuanliang Meng, Zaiqiao Meng, Zhao Meng, Stefano Menini, Aditya Krishna Menon, Sachit Menon, Wolfgang Menzel, Robert Mercer, Paolo Merialdo, Paola Merlo, William Merrill, Jack Merullo, Enza Messina, Eleni Metheniti, Guillaume Metzler, Stefano Mezza, Chenggang Mi, Md Messal Monem Miah, Yisong Miao, Alessio Miaschi, Henryk

Michalewski, Erik Miebling, Rada Mihalcea, Vladislav Mikhailov, Elena Mikhalkova, Jeremiah Milbauer, Filip Miletić, Evangelos Milios, Simon Mille, Alice Millour, Do June Min, Qingkai Min, Pasquale Minervini, Bhavnick Singh Minhas, Benjamin Minixhofer, Einat Minkov, Hideya Mino, Seyed Abolghasem Mirroshandel, Anand Mishra, Kshitij Mishra, Pruthwik Mishra, Amita Misra, Teruko Mitamura, Cassie S. Mitchell, Paul Mithun, Ashish Mittal, Shubham Mittal, Yusuke Miyao, Taro Miyazaki, Masahiro Mizukami, Fengran Mo, Tong Mo, Wenjie Jacky Mo, Natwar Modani, Ali Modarressi, Sandip Modha, Abiodun Modupe, Lucas Moeller, Alireza Mohammadshahi, Wafaa Mohammed, Shrestha Mohanty, Hosein Mohebbi, Jisoo Mok, Diego Molla, Masoud Monajatipoor, Ishani Mondal, Sneha Mondal, Arturo Montejo-Ráez, Christof Monz, Han Cheol Moon, Hyeongdon Moon, Hyeonseok Moon, Jong Hak Moon, Sangwhan Moon, Ray Mooney, Sukanya Moorthy, Ibraheem Muhammad Moosa, Steven Moran, Junichiro Mori, Yusuke Mori, Makoto Morishita, Robert Moro, John Xavier Morris, Jacob Morrison, Daniil Moskovskiy, Aida Mostafazadeh Davani, Pablo Mosteiro, Yutao Mou, Seyed Mahed Mousavi, Basel Mousi, Rajiv Movva, Khalil Mrini, Frank Martin Mtumbuka, Feiteng Mu, Yida Mu, Yongyu Mu, Aaron Mueller, Animesh Mukherjee, Anjishnu Mukherjee, Sagnik Mukherjee, Sourabrata Mukherjee, Tanmoy Mukherjee, Medet Mukushev, Matthew Mulholland, Ankan Mullick, Alif Munim, Munkhtulga Munkhtulga, Juan Pablo Munoz, Vishvak Murahari, Koji Murakami, Soichiro Murakami, Lidiya Murakhovs'ka, Goran Muric, Kenton Murray, Pradeep K. Murukannaiah, John Murzaku, Zairah Mustahsan, Arianna Muti, Ricardo Muñoz Sánchez, Agnieszka Mykowiecka, Sheshera Mysore

Siddharth N, CheolWon Na, Clara Na, Seung-Hoon Na, Mohamed Nadif, Arijit Nag, Atharva Naik, Inderjeet Jayakumar Nair, Sathvik Nair, Saeed Najafi, Tetsuji Nakagawa, Satoshi Nakamura, Yuta Nakashima, Satya Sai Srinath Namburi Gnvv, Ananjan Nandi, Abhilash Nandy, Mukuntha Narayanan S, Pranav Narayanan Venkit, Ali Naseh, Sudip Kumar Naskar, Abhijnan Nath, Meith Hitesh Navlakha, Nihal V. Nayak, Shraavan Nayak, Siddharth Nayak, Mir Tafseer Nayeem, Hamada Nayel, Yaroslav Nechaev, Vasudevan Nedumpozhimana, Seyed Parsa Neshaei, Graham Neubig, Benjamin Newman, See-Kiong Ng, Youyang Ng, Chau Nguyen, Cong-Duy T Nguyen, Dong Nguyen, Duc-Vu Nguyen, Kiem-Hieu Nguyen, Kiet Van Nguyen, Long HB Nguyen, Minh-Tien Nguyen, Ngan Luu-Thuy Nguyen, Phuong Minh Nguyen, Thanh-Tung Nguyen, Thi-Nhung Nguyen, Vincent Nguyen, Hoang-Quoc Nguyen-Son, Ansong Ni, Jingwei Ni, Shiwen Ni, Xuanfan Ni, Zhaoheng Ni, Iftitahu Ni'mah, Garrett Nicolai, Vlad Niculae, Ercong Nie, Jian-Yun Nie, Lunyiu Nie, Pengyu Nie, Yuxiang Nie, Kristoffer Nielbo, Hellina Hailu Nigatu, Siegfried Nijssen, Malvina Nikandrou, Irina Nikishina, Christina Niklaus, Joel Niklaus, Sergey Nikolenko, Vassilina Nikoulina, Yue Ning, Kosuke Nishida, Noriki Nishida, Sergiu Nisioi, Di Niu, Guanglin Niu, Jingcheng Niu, Tianhao Niu, Bill Noble, Tadashi Nomoto, Enrique Noriega-Atala, Kai North, Armineh Nourbakhsh, Damien Nouvel, Michal Novák, Franz Nowak, Giorgio Di Nunzio, Sarana Nutanong, Aurélie Névél

Brendan O'Connor, Daisuke Oba, Sebastian Ochs, Bahadorreza Ofoghi, Perez Ogayo, Byung-Doh Oh, Hanseok Oh, Minsik Oh, Yui Oka, Eda Okur, Arlindo L. Oliveira, Hugo Gonçalo Oliveira, Byung-Won On, Ethel Chua Joy Ong, Yasumasa Onoe, Subba Reddy Oota, Andreas Opedal, Juri Opitz, Riccardo Orlando, Pedro Ortiz Suarez, Thomas Ortner, Yohei Oseki, Simon Ostermann, Yulia Otmakhova, Wolfgang Otto, Jiao Ou, Hiroki Ouchi, Siqi Ouyang, Siru Ouyang, Mete Ozay, Kadir Bulut Ozler

Inkit Padhi, Ankur Padia, Aishwarya Padmakumar, Vishakh Padmakumar, Sebastian Padó, Aline Paes, Patrizia Paggio, Vaishali Pal, Hamid Palangi, Chester Palen-Michel, Jeff Z. Pan, Jiayi Pan, Jinsheng Pan, Rong Pan, Rui Pan, Youcheng Pan, Yuchen Pan, Alexander Panchenko, Subhadarshi Panda, Rohan Pandey, Saurabh Kumar Pandey, Shivasankaran Vanaja Pandi, Shrey Pandit, Deepak Pandita, Chenxi Pang, Liang Pang, Aleksandr Panov, Georgios Pantazopoulos, Haris Papa-georgiou, Sara Papi, Paolo Papotti, Emerson Cabrera Paraiso, Sethupathy Parameswaran, Ashwin

Paranjape, Georgios Paraskevopoulos, Letitia Parcalabescu, Amit Parekh, Tanmay Parekh, Rahil Parikh, Soham Parikh, ChaeHun Park, Chan Young Park, Chanjun Park, Cheonbok Park, Choonghyun Park, Eunhwan Park, Hyunji Hayley Park, Jinyoung Park, Jungyeul Park, Seo Yeon Park, Shinwoo Park, Sumin Park, Mihir Parmar, Patrick Paroubek, Prasanna Parthasarathi, Md Rizwan Parvez, Nivranshu Pasricha, Ajay Patel, Nilay Patel, Manasi Patwardhan, Bibek Paudel, Felipe Paula, Amalie Brogaard Pauli, Adam Pauls, Jay Pavagadhi, Sachin Pawar, Ali Payani, Jiixin Pei, Qizhi Pei, Shichao Pei, Yulong Pei, Kellin Pelrine, Baolin Peng, Bo Peng, Bo Peng, Letian Peng, Min Peng, Qiwei Peng, Qiyao Peng, Shuyuan Peng, Siyao Peng, Wei Peng, Xiangyu Peng, Xueping Peng, Yifan Peng, Oren Pereg, E. Margaret Perkoff, Yotam Perlitz, Ali Pesaranghader, Denis Peskoff, Molly Petersen, Pavel Petrushkov, Suzanne Petryk, Samuel Pfrommer, Chau Minh Pham, Philip Pham, Van-Cuong Pham, Minh Hieu Phan, Van-Thuy Phi, Fred Philippy, Renjie Pi, Tanzir Pial, Francesco Piccinno, Stjepan Picek, Andrea Piergentili, Janet B. Pierrehumbert, Matúš Pikuliak, Sameer Pimparkhede, Mathis Pink, Yuval Pinter, Andrew Piper, Rajesh Piryani, Lidia Pivovarova, Barbara Plank, Flor Miriam Plaza-del-Arco, Moritz Plenz, Esther Ploeger, Christian Poelitz, Gabriel Poesia, Alex Polozov, Tim Polzehl, Nicholas Popovic, Ian Porada, Soujanya Poria, Matt Post, Clifton A Poth, Martin Potthast, Christopher Potts, Pascal Poupart, Mohammadreza Pourreza, Shrimai Prabhumoye, Pradip Pramanick, Archiki Prasad, Adithya Pratapa, Judita Preiss, Irina Proskurina, Dongqi Pu, Jiashu Pu, Shi Pu, Xiao Pu, Yewen Pu, Giovanni Puccetti, Giulia Pucci, Ratish Puduppully, Haritz Puerto, Rajkumar Pujari, Jitesh Punjabi, Sukannya Purkayastha, Matthew Purver, Ayu Purwarianti, Rifki Afina Putri, Adarsh Pyarelal

Ayesha Qamar, Ehsan Qasemi, Guilin Qi, Jiexing Qi, Yunjia Qi, Chen Qian, Cheng Qian, Hongjin Qian, Kun Qian, Lihua Qian, Shenbin Qian, Yusu Qian, Zhenxing Qian, Jipeng Qiang, Minjie Qiang, Yao Qiang, Lingfeng Qiao, Shuofei Qiao, Bowen Qin, Chengwei Qin, Jinghui Qin, Libo Qin, Yao Qin, Yulei Qin, Chen Qiu, Jieliu Qiu, Shuwen Qiu, Xin Ying Qiu, Zexuan Qiu, Zhaopeng Qiu, Muhammad Reza Qorib, Jin Qu, Tho Quan, Ariadna Quattoni

Rakesh R Menon, Ella Rabinovich, Edward Raff, Davood Rafiei, Dinesh Raghu, Tathagata Raha, Hossein A. Rahmani, Daking Rai, Sunny Rai, Md Nishat Raihan, Vatsal Raina, Vyas Raina, Lisa Raithel, Navid Rajabi, Hossein Rajaby Faghihi, Sara Rajaei, Sanguthevar Rajasekaran, Ori Ram, Anil Ramakrishna, Savitha Ramasamy, Sahana Ramnath, Rita Ramos, Shihao Ran, Tharindu Ranasinghe, Poorva Rane, Anku Rani, Jaspreet Ranjit, Peter A. Rankel, Dongning Rao, Jinneng Rao, Sudha Rao, Essam Rashed, Farzana Rashid, Abhinav Rastogi, Vipul Kumar Rathore, Shauli Ravfogel, Sahithya Ravi, Ambrish Rawat, Vipula Rawte, Baishakhi Ray, Soumya Ray, Agha Ali Raza, Anastasia Razdaibiedina, Yasaman Razeghi, Evgeniia Razumovskaia, Anton Razzhigaev, Chris Reed, Michaela Regneri, Ricardo Rei, Navid Rekabsaz, Jie Ren, Jie Ren, Mengjie Ren, Pengjie Ren, Qingyang Ren, Yubing Ren, Yujie Ren, Steven J Rennie, Rezvaneh Rezapour, Ryokan Ri, Stephen D. Richardson, Caitlin Laura Richter, Korbinian Riedhammer, Jonas Rieger, Matīss Rikters, Eric Ringger, Anthony Rios, Yong Man Ro, Gabriel Roccabruna, Alvaro Rodrigo, Christophe Rodrigues, Paul Rodrigues, Sergio José Rodríguez Méndez, Melissa Roemmele, Ana-Cristina Rogoz, Paul Roit, Muhammad Roman, Angelika Romanou, Srikanth Ronanki, Wenge Rong, Tanya Roosta, Rudolf Rosa, Domenic Rosati, Germán Rosati, Carolyn Rose, Alexis Ross, Candace Ross, Gaetano Rossiello, Mohammad Rostami, Michael Roth, Sascha Rothe, Guy Rotman, Mozhdah Rouhsedaghat, Dmitri Roussinov, Aniruddha Roy, Aurko Roy, Kalyani Roy, Kashob Kumar Roy, Sumegh Roychowdhury, Jos Rozen, William Rudman, Koustav Rudra, Frank Rudzicz, Federico Ruggeri, Ramon Ruiz-Dolz, Bharat Runwal, Josef Ruppenhofer, Jonathan Rusert, Stefan Ruseti, Giuseppe Russo, Phillip Rust, Pavel Rychlý, Susanna Rücker

Dan SU, Hongjin SU, Arkadiy Saakyan, Anusha Sabbineni, Ashish Sabharwal, Kaushik Ram Sadagopan, Mobashir Sadat, Zahra Sadeghi, Nafis Sadeq, Arka Sadhu, Niloofer Safi Samghabadi, Alsu Sagirova, Benoît Sagot, Punyajoy Saha, Sougata Saha, Sriparna Saha, Oscar Sainz, Frede-

ric Sala, Emily Saldanha, Chris Samarinas, Farhan Samir, Abraham Sanders, Jivnesh Sandhan, Piotr Sankowski, Scott Sanner, Leonid Sanochkin, Brenda Salenave Santana, Debarshi Kumar Sanyal, Soumya Sanyal, Abulhair Saparov, Hitesh Sapkota, Sunita Sarawagi, Rupak Sarkar, Gabriele Sarti, Sheikh Muhammad Sarwar, Shota Sasaki, Ashutosh Sathe, Prasanna Sattigeri, Sunil Saumya, Danielle Saunders, Andrey Savchenko, Germans Savcisens, Guergana K Savova, Apoorv Saxena, Michael Saxon, Salim Sazed, Harrisen Scells, Thomas Schaaf, Shigehiko Schamoni, Benjamin Schiller, David Schlangen, Jörg Schlötterer, Helmut Schmid, Fabian David Schmidt, Martin Schmitt, Felix Schneider, Florian Schneider, Nathan Schneider, Stephanie Schoch, Timo Pierre Schrader, Christopher Schröder, Hinrich Schuetze, Hendrik Schuff, Björn Schuller, Elliot Schumacher, Carolin M. Schuster, Tal Schuster, Idan Schwartz, Stefan Schweter, Pola Schwöbel, Alessandro Scirè, Anastasiia Sedova, Avi Segal, Encarna Segarra, Amit Seker, Toshiyuki Sekiya, Ramaneswaran Selvakumar, David Semedo, Sina Semnani, Indira Sen, Jaydeep Sen, Procheta Sen, Apurbalal Senapati, Meghdu Sengupta, Jaehyung Seo, Ovidiu Serban, Oleg Serikov, Sofia Serrano, Agrima Seth, Rita Sevastjanova, Shaden Shaar, Agam Shah, Raj Sanjay Shah, Rajiv Ratn Shah, Shalin Shah, Md Shihab Shahriar, Rifat Shahriyar, Sagi Shaier, Omar Shaikh, Valerie Shalin, Weiqiao Shan, Chao Shang, Hengchao Shang, Junyuan Shang, Lin Shang, Mingyue Shang, Abhilash Shankarampeta, Liangying Shao, Wei Shao, Yijia Shao, Yunfan Shao, Ori Shapira, Matthew Shardlow, Ehsan Shareghi, Aditya Sharma, Ashish Sharma, Drishti Sharma, Harsh Sharma, Kartik Sharma, Mayukh Sharma, Raksha Sharma, Shivam Sharma, Soumya Sharma, Serge Sharoff, Jingyuan Selena She, Ryan Shea, Zaid Sheikh, Ravi Shekhar, Artem Shelmanov, Bowen Shen, Guangyu Shen, Hua Shen, Jiabin Shen, Lei Shen, Lingfeng Shen, Mingwei Shen, Qinlan Shen, Tianhao Shen, Weizhou Shen, Xiangqing Shen, Yikang Shen, Ying Shen, Zejiang Shen, Jiawei Sheng, Quan Z. Sheng, Pranav Shetty, Chen Shi, Chufan Shi, Haizhou Shi, Haochen Shi, Kaize Shi, Kejian Shi, Ning Shi, Qi Shi, Wenqi Shi, Xiaoming Shi, Zhan Shi, Zhengliang Shi, Zhouxing Shi, Zijing Shi, Tomohide Shibata, Kyuhong Shim, Hidetoshi Shimodaira, Erica Kido Shimomoto, Andrew Shin, Haebin Shin, Joongbo Shin, Myeongcheol Shin, Kazutoshi Shinoda, Kiyoaki Shirai, Yow-Ting Shiue, Namrata Shivagunde, Esther Shizgal, Hossein Sholehrasa, Lidan Shou, Akshat Shrivastava, Ritvik Shrivastava, Kai Shu, Raphael Shu, Yiheng Shu, Omer Shubi, Zeren Shui, Shreya Shukla, KaShun Shum, Ilia Shumailov, Chenglei Si, Jiasheng Si, Shijing Si, Shuzheng Si, Suzanna Sia, Anthony Sicilia, A.b. Siddique, Sandipan Sikdar, Karan Sikka, Carina Silberer, Tejpalsingh Siledar, Damien Sileo, Miiikka Silfverberg, Chaklam Silpasuwanchai, João Silva, Purificação Silvano, Fabrizio Silvestri, Robert Sim, Patrick Simianer, Gabriel Simmons, Antoine Simoulin, Edwin Simpson, Apoorva Singh, Harman Singh, Mukul Singh, Nikhil Singh, Rituraj Singh, Aryan Singhal, Saksham Singhal, Iustin Sirbu, Jasivan Alex Sivakumar, Konstantinos Skianis, Steven Skiena, Aviv Slobodkin, David A. Smith, Răzvan-Alexandru Smădu, Daria Soboleva, Mohammad Golam Sohrab, Petr Sojka, Amir Soleimani, Junyoung Son, Chiyu Song, Dandan Song, Dawn Song, Demin Song, Guojie Song, Haiyue Song, Hwanjun Song, Jongyoon Song, Kaitao Song, Kyungwoo Song, Linqi Song, Mingyang Song, Ran Song, Ruihua Song, Shuangyong Song, Xiaoshuai Song, Xin Song, Yixiao Song, Yu Song, Yurun Song, Zhenqiao Song, Ziang Song, Nikita Soni, Sandeep Soni, Ekta Sood, Anna Sotnikova, Sajad Sotudeh, Laure Soulier, Gerasimos Spanakis, Alexander Spangher, Andreas Spitz, Makesh Narsimhan Sreedhar, Rohit Sridhar, Neha Srikanth, Mukund Srinath, Krishna Srinivasan, Padmini Srinivasan, Tejas Srinivasan, Aarohi Srivastava, Saurabh Srivastava, Joe Stacey, Jacopo Staiano, Ieva Staliunaite, Karolina Stanczak, Marija Stanojevic, Gabriel Stanovsky, Katherine Stasaski, Mark Steedman, Sebastian Steindl, Andreas Stephan, Simon Stepputtis, Hannah Sterz, Suzanne Stevenson, Brandon M. Stewart, Ian Stewart, Asa Cooper Stickland, Niklas Stoehr, Alessandro Stolfo, Shane Storks, Phillip Benjamin Ströbel, Rickard Stureborg, Sebastian Stüker, Fangfang Su, Hsuan Su, Hung-Ting Su, Ruolin Su, Weihang Su, Xiangdong Su, Xin Su, Ying Su, Yusheng Su, Zhaochen Su, Zhaolun Su, Melanie Subbiah, Anand Subramanian, Sanjay Subramanian, Shivashankar Subramanian, Katsuhito Sudoh, Hiroaki Sugiyama, Daewon Suh, Yoshi Suhara, Albert Sun, Anna Sun, Bin Sun, Changzhi Sun, Chen Sun, Chenkai Sun, Fei Sun, Haifeng Sun, Haitian Sun, Hao Sun,

Haotian Sun, Jiashuo Sun, Jingwei Sun, Jingyi Sun, Jingyuan Sun, Jun Sun, Kaiser Sun, Kexuan Sun, Lichao Sun, Mingming Sun, Qiushi Sun, Renliang Sun, Rui Sun, Si Sun, Siqi Sun, Wei Sun, Weiqi Sun, Xiaobing Sun, Yang Sun, Yuanyuan Sun, Yutao Sun, Zequn Sun, Zhaoyue Sun, Zhenjie Sun, Zhewei Sun, Anirudh Sundar, Jayaprakash Sundararaj, Barkavi Sundararajan, Dhanasekar Sundararaman, Megha Sundriyal, Mujeen Sung, Yoo Yeon Sung, Yun-Hsuan Sung, Marek Suppa, Mihai Surdeanu, Lintang Sutawika, Reem Suwaileh, Anej Svete, Monorama Swain, Maciej Szankin, Stan Szpakowicz, Piotr Szymański, Jorge Sánchez, Felipe Sánchez-Martínez, Jonne Sälevä, Jan Šnajder, Michal Štefánik

Shabnam Tafreshi, Andrea Tagarelli, Dima Taji, Makoto Takamoto, Sho Takase, Takehiro Takayanagi, Kunihiro Takeoka, Koichi Takeuchi, Yik-Cheung Tam, Fabio Tamburini, Pradyumna Tambwekar, Aleš Tamchyna, Chao-Hong Tan, Fuwen Tan, Hanzhuo Tan, Hao Tan, Haochen Tan, Hongye Tan, Liling Tan, Minghuan Tan, Mingkui Tan, Samson Tan, Shaomu Tan, Weiting Tan, Xiaoyu Tan, Xingwei Tan, Zeqi Tan, Zhaoxuan Tan, Zhen Tan, Zhixing Tan, Niket Tandon, Binghao Tang, Bo Tang, Buzhou Tang, Chen Tang, Jiliang Tang, Keke Tang, Pengfei Tang, Ruixiang Tang, Wei Tang, Xiangru Tang, Xuemei Tang, Yixuan Tang, Yun Tang, Yuqing Tang, Zecheng Tang, Zheng Tang, Ziwei Tang, Kumar Tanmay, Chaofan Tao, Chongyang Tao, Mingxu Tao, Wei Tao, Xijia Tao, Zhuo Tao, Allahsera Auguste Tapo, Panuthep Tasawong, Sandeep Tata, Christopher Tauchmann, Anuja Tayal, Seyed Nima Tayarani Bathaie, Sayed Mohammadreza Tayaranian Hosseini, Stephen Eugene Taylor, Andon Tchechmedjiev, Simone Tedeschi, Atula Tejaswi, Christopher Tensmeyer, Daniela Teodorescu, Katherine Thai, Megh Thakkar, Nandan Thakur, Luan Thanh Nguyen, Surendrabikram Thapa, Raghuv eer Thirukovalluru, Krishnaprasad Thirunarayan, Pierre Thodoroff, Alex Thomo, James Thorne, Tristan Thrush, David Thulke, Ran Tian, Yanzhi Tian, Yuan Tian, Yufei Tian, Anna Tigonova, Maria Tikhonova, Tiago Timponi Torrent, Devesh Tiwari, Erik Tjong Kim Sang, Lingbo Tong, MeiHan Tong, Terry Tong, Xiaoyu Tong, Yu Tong, Atnafu Lambebo Tonja, Cagri Toraman, Lucas Torroba Hennigen, Shubham Toshniwal, Benjamin Towle, Masashi Toyoda, Khanh Quoc Tran, Khiem Vinh Tran, Khoi-Nguyen Tran, Thy Thy Tran, Filip Trhlík, Rocco Tripodi, Raphael Troncy, Sergey Troshin, Sean Trott, Yun-Da Tsai, Dimitrios Tsarapatsanis, Liang-Hsuan Tseng, Tornike Tsereteli, Ioannis Tsiamas, Eleftheria Tspidi, Oren Tsur, Geng Tu, Jingxuan Tu, Lifu Tu, Rong-Cheng Tu, Shangqing Tu, Yi Tu, Zhucheng Tu, Ada Tur, Dilek Hakkani Tur, Elsbeth Turcan, Nicolas Turenne, Elena Tutubalina, Tinne Tuytelaars, Aman Tyagi, Utkarsh Tyagi, Gareth Tyson

Adaku Uchendu, Md Nayem Uddin, Can Udomcharoenchaikit, Stefan Ultes, Rishabh Upadhyay, Sagor Uprety, Ashok Urlana, Ricardo Usbeck, Yoshitaka Ushiku, David Uthus, Saiteja Utpala

Venktesh V, Adithya V Ganesan, Robert Vacareanu, Saujas Vaduguru, Teemu Vahtola, Nirvi K Vakharia, Nidhi Vakil, Thomas Vakili, Alex I. Valencia-Valencia, Jannis Vamvas, Michiel Van Der Meer, Rik Van Noord, David Vandyke, Daniel Varab, Francielle Vargas, Amey Varhade, Maya Varma, Lav R. Varshney, Yerram Varun, Oleg Vasilyev, Javier Vazquez-Corral, Eva Maria Vecchi, Hariram Veeramani, Aditya Srikanth Veerubhotla, Sarmishta Velury, Saranya Venkatraman, Giulia Venturi, Suzan Verberne, Pat Verga, Rakesh M Verma, Federica Vezzani, Lovekesh Vig, David Vilares, Dan Vilenchik, Krishnapriya Vishnubhotla, Vijay Viswanathan, Jean-Noël Vittaut, Gerson Vizcarra, Andreas Vlachos, Juraj Vladika, Rob Voigt, Sergey Volokhin, Pius Von Däniken, Piek Vossen, Thuy-Trang Vu, Yogarshi Vyas

Xing W, Yuiga Wada, Somin Wadhwa, Eitan Wagner, Abdul Waheed, Jan Philip Wahle, Hiromi Wakaki, Andreas Waldis, Changxuan Wan, David Wan, Herun Wan, Huaiyu Wan, Qizhi Wan, Ruyuan Wan, Xingchen Wan, Yingjia Wan, Yixin Wan, Zhen Wan, Zhongwei Wan, Ante Wang, Bang Wang, Baoxun Wang, Barry Wang, Benyou Wang, Bin Wang, Bing Wang, Bingning Wang, Bingqing Wang, Bo Wang, Bo Wang, Boshi Wang, Bowen Wang, Cheng Wang, Chenglong



Wang, Cunxiang Wang, Daling Wang, Danding Wang, Danqing Wang, Darren Yow-Bang Wang, Di Wang, Dingmin Wang, Fan Wang, Fanyu Wang, Gengyu Wang, Hai Wang, Han Wang, Hao Wang, Haobo Wang, Haodong Wang, Haoyu Wang, Haoyu Wang, Hongru Wang, Hongyu Wang, Huadong Wang, Huazheng Wang, Huazheng Wang, Huimin Wang, Ji Wang, Jiaan Wang, Jian Wang, Jianing Wang, Jianyong Wang, Jianzong Wang, Jiapeng Wang, Jin Wang, Jingang Wang, Jixuan Wang, Jun Wang, Jun Wang, Jun Wang, Junlin Wang, Junxiong Wang, Kexiang Wang, Kexin Wang, Kexin Wang, Keze Wang, Lean Wang, Lei Wang, Li Wang, Liang Wang, Lingzhi Wang, Liwei Wang, Longzheng Wang, Mengru Wang, Mengxiang Wang, Mingyang Wang, Pan Wang, Peiyi Wang, Peng Wang, Pengcheng Wang, Qifan Wang, Qingyue Wang, Qingyun Wang, Quan Wang, Renxi Wang, Rui Wang, Rui Wang, Rui Wang, Ruijie Wang, Ruixiang Wang, Ruixiao Wang, Runhui Wang, Ruoyao Wang, Ryan Yixiang Wang, Sen Wang, Shangfei Wang, Shihan Wang, Shuohuan Wang, Shuting Wang, Sijia Wang, Siqi Wang, Siyuan Wang, Song Wang, Song Wang, Songhe Wang, Tianchun Wang, Tianduo Wang, Wei Wang, Wei Wang, Weixuan Wang, Weiyue Wang, Wen Wang, Wenjie Wang, Wenjie Wang, Wenxuan Wang, Xi Wang, Xi Wang, Xiao Wang, Xiaoli Wang, Xiaoou Wang, Xiaosen Wang, Xinbing Wang, Xindi Wang, Xinpeng Wang, Xintong Wang, Xinyou Wang, Xinyu Wang, Xinyuan Wang, Xinze Wang, Xiting Wang, Xuan Wang, Yanfeng Wang, Yang Wang, Yanguang Wang, Yanhao Wang, Yanling Wang, Yaqiang Wang, Yau-Shian Wang, Ye Wang, Yejie Wang, Yibin Wang, Yibo Wang, Yichen Wang, Yijue Wang, Yile Wang, Yiming Wang, Yimu Wang, Yining Wang, Yiran Wang, Yiwei Wang, Yu Wang, Yuan-fang Wang, Yueqian Wang, Yufei Wang, Yuhang Wang, Yun-Cheng Wang, Yuqing Wang, Yuxia Wang, Yuxin Wang, Yuxuan Wang, Yuxuan Wang, Zekun Wang, Zezhong Wang, Zhaowei Wang, Zhaoyang Wang, Zhen Wang, Zhexuan Wang, ZhiZheng Wang, Zhilin Wang, Zhiruo Wang, Zhuoer Wang, Zihan Wang, Zihan Wang, Zihao Wang, Zihao Wang, Zihua Wang, Zijie Wang, Zilong Wang, Zimu Wang, Ziqi Wang, Dittaya Wanvarie, Azmine Touthik Wasi, Taro Watanabe, Julia Watson, Ishaan Watts, Tharindu Cyril Weerasooriya, Jerry Wei, Kaiwen Wei, Kangda Wei, Penghui Wei, Sheng-Lun Wei, Victor Junqiu Wei, Wei Wei, Xiang Wei, Xiao Wei, Maxwell Weinzierl, Leonie Weissweiler, Anuradha Welivita, Orion Weller, Bosi Wen, Liang Wen, Yuqiao Wen, Zeyi Wen, Zihao Wen, Zihua Wen, Chris Wendler, James Bradley Wendt, Yixuan Weng, Zhaotian Weng, Peter West, Robert West, Spencer Whitehead, Haryo Akbarianto Wibowo, Philipp Wicke, Gregor Wiedemann, Michael Wiegand, Matti Wiegmann, Adam Wiemerslage, John Frederick Wieting, Ethan Wilcox, Bryan Wilie, Miles Williams, Dane A Williamson, Steven R Wilson, Shuly Wintner, Michael Witbrock, Lior Wolf, Christian Wolff, Ka-Chun Wong, Tak-Lam Wong, Zach Wood-Doughty, Dustin Wright, Anna Wróblewska, Anne Wu, Bohong Wu, Bowen Wu, Chen Wu, Chen Henry Wu, Chenglin Wu, Chenming Wu, Di Wu, Di Wu, Fanyou Wu, Han Wu, Haoran Wu, Hongqiu Wu, Huijia Wu, Ji Wu, Jialong Wu, Jiaman Wu, Jiawei Wu, Jiaying Wu, Jincenzi Wu, Jinge Wu, Jun Wu, Junda Wu, Junjie Wu, Junran Wu, Junshuang Wu, KeChun Wu, Kebin Wu, Linjuan Wu, Linzhi Wu, Minghao Wu, Ning Wu, Qianhui Wu, Qingyang Wu, Qiyu Wu, Ruituo Wu, Shengguang Wu, Shengqiong Wu, Shujin Wu, Siwei Wu, Sixing Wu, Taiqiang Wu, Ting-Wei Wu, Tongtong Wu, Wei Wu, Wei Wu, Wenshan Wu, Xiao-Ming Wu, Xiaobao Wu, Xin Wu, Xiyang Wu, Xuansheng Wu, Xueqing Wu, Xuyang Wu, Yin Wu, Yiquan Wu, Yulong Wu, Yuping Wu, Yuxia Wu, Zhaofeng Wu, Zhen Wu, Zhengxuan Wu, Zhijing Wu, Zhuofeng Wu, Zijun Wu, Zirui Wu, Zongyu Wu, Amelie Wuehrl

Fuqi XU, Zhiheng Xi, Heming Xia, Menglin Xia, Tingyu Xia, Zhaoyang Xia, Jiannan Xiang, Lu Xiang, Tong Xiang, Wei Xiang, Bei Xiao, Cao Xiao, Chaojun Xiao, Chenghao Xiao, Chuan Xiao, Chunyang Xiao, Junbin Xiao, Likang Xiao, Qingfa Xiao, Shitao Xiao, Teng Xiao, Yao Xiao, Yisheng Xiao, Ziang Xiao, Zilin Xiao, Boyi Xie, Huiyuan Xie, Qianqian Xie, Roy Xie, Rui Xie, Sean Xie, Tianbao Xie, Tingyu Xie, Yaqi Xie, Yiqing Xie, Yong Xie, Yuexiang Xie, Yuqiang Xie, Yuqing Xie, Yuxi Xie, Zhihui Xie, Zhouhang Xie, Linzi Xing, Yue Xing, Jing Xiong, Jinjun Xiong, Weimin Xiong, Yizhe Xiong, Bo Xu, Boyan Xu, Chen Xu, Chengjin Xu, Chunpu Xu, Fanyuan Xu, Guangyue Xu, Guohai Xu, Haiming Xu, Hainiu Xu, Hanzi Xu, Haoran Xu, Haotian Xu,

Hongzhi Xu, Hu Xu, Jiacheng Xu, Jiahao Xu, Jialiang Xu, Jiashu Xu, Jing Xu, Jitao Xu, Jun Xu, Jun Xu, Kaidi Xu, Liyan Xu, Lvxiaowei Xu, Ming Xu, Nan Xu, Pengyu Xu, Ran Xu, Shanshan Xu, Tong Xu, Wang Xu, Weijia Xu, Weijie Xu, Wenda Xu, Wenduan Xu, Wenjie Xu, Xiao Xu, Yan Xu, Yao Xu, Yi Xu, Yifan Xu, Yifan Xu, Yige Xu, Zhangchen Xu, Zhenran Xu, Zhichao Xu, Zhikun Xu, Zhiyang Xu, Zihang Xu, Ziwei Xu, Junyu Xuan, Keyang Xuan, Boyang Xue, Huiyin Xue, Le Xue, Wei Xue

Tiezheng YU, Shuntaro Yada, Ashok Yadav, Prateek Yadav, Vikas Yadav, Aditya Yadavalli, Atsuki Yamaguchi, Ivan P. Yamshchikov, Da Yan, Hanqi Yan, Haolong Yan, Jiahuan Yan, Jing Nathan Yan, Jun Yan, Lingyong Yan, Ming Yan, Ming Yan, Shiyang Yan, Tianyi Lorena Yan, Weixiang Yan, Xifeng Yan, Xin Yan, An Yang, Changbing Yang, Chenyang Yang, Daniel Yang, Deqing Yang, Dingyi Yang, Guanqun Yang, Haojin Yang, Heng Yang, Hsiu-Yu Yang, Huichen Yang, Jian Yang, Jiaoyun Yang, Jinfan Yang, Jingfeng Yang, Jinrui Yang, June Yong Yang, Kaicheng Yang, Kailai Yang, Kejuan Yang, Kichang Yang, Li Yang, Linyi Yang, Longfei Yang, Meng Yang, Muyun Yang, Nakyeong Yang, Qikai Yang, Sen Yang, Shu-wen Yang, Sohee Yang, Songhua Yang, Tao Yang, Tianyu Yang, Tianyu Yang, Wei Yang, Wenkai Yang, Wenmian Yang, Xianjun Yang, Xiaocui Yang, Xincheng Yang, Yahan Yang, Yaming Yang, Yaodong Yang, Yating Yang, Yi Yang, Yifan Yang, Yifan Yang, Yizhe Yang, Yongjin Yang, Yuanhang Yang, Yue Yang, Yuhao Yang, Yujiu Yang, Zhenhua Yang, Zhicheng Yang, Zhuoyi Yang, Ziyi Yang, Zonglin Yang, Ken Yano, Binwei Yao, Fangzhou Yao, Feng Yao, Jiarui Yao, Junfeng Yao, Peiran Yao, Xin Yao, Yunzhi Yao, Zijun Yao, Zonghai Yao, Michal Yarom, Majid Yazdani, Bingyang Ye, Chenchen Ye, Fanghua Ye, Hai Ye, Haotian Ye, Jiacheng Ye, Jingheng Ye, Jinhui Ye, Junjie Ye, Meng Ye, Rong Ye, Tong Ye, Xuesong Ye, Yang Ye, Eric Yeh, Sung-Lin Yeh, Jinyoung Yeo, Akhila Yerukola, Dit-Yan Yeung, Jingwei Yi, Xiaoyuan Yi, Li S. Yifei, Congchi Yin, Fangcong Yin, Ming Yin, Xunjian Yin, Xuwang Yin, Yu Yin, Yuwei Yin, Zhangyue Yin, Michael Miller Yoder, Sho Yokoi, Gal Yona, Zheng Xin Yong, Chang D. Yoo, KiYoon Yoo, Eunseop Yoon, Hee Suk Yoon, Soyoun Yoon, Sungroh Yoon, Sunjae Yoon, Susik Yoon, Ori Yoran, Issei Yoshida, Ryo Yoshida, Hiyori Yoshikawa, Paul Youssef, Dian Yu, Dian Yu, Dianhai Yu, Dingyao Yu, Donghan Yu, Erxin Yu, Fangxu Yu, Guoxin Yu, Haiyang Yu, Haiyang Yu, Hang Yu, Haofei Yu, Hwanjo Yu, Jifan Yu, Junjie Yu, Keunwoo Peter Yu, Mengxia Yu, Min Yu, Peilin Yu, Pengfei Yu, Ping Yu, Qi Yu, Qinan Yu, Shi Yu, Shoubin Yu, Sicheng Yu, Tengfei Yu, Tianshu Yu, Tong Yu, Xiao Yu, Xiaodong Yu, Xiaojing Yu, Xincheng Yu, Xinyan Velocity Yu, Yahan Yu, Zhuohao Yu, Bo Yuan, Fan Yuan, Fei Yuan, Hongyi Yuan, Jiayi Yuan, Lifan Yuan, Moy Yuan, Ruifeng Yuan, Siyu Yuan, Weizhe Yuan, Ye Yuan, Ye Yuan, Yifei Yuan, Youliang Yuan, Zheng Yuan, Zheng Yuan, Li Yuan, Linan Yue, Murong Yue, Xiang Yue, Zhenrui Yue, Zihao Yue, Hyokun Yun, Se-Young Yun, Tian Yun, Aizihaijiang Yusufu

Ted Zadouri, Matei Zaharia, Hamada M Zahera, Mohd Zaki, Mahdi Zakizadeh, Marcel Zalmanovici, Kerem Zaman, Roberto Zamparelli, Changtong Zan, Daoguang Zan, Xiaoxue Zang, Klim Zaporozjets, Urchade Zaratiana, Guido Zarrella, Sina Zarrieß, Jiali Zeng, Jiaqi Zeng, Min Zeng, Qingcheng Zeng, Zhengran Zeng, Ziqian Zeng, Luke Zettlemoyer, Rodolfo Zevallos, Hanwen Zha, Haolan Zhan, Pengwei Zhan, Runzhe Zhan, Baohua Zhang, Bo-Wen Zhang, Bohan Zhang, Bowen Zhang, Caiqi Zhang, Ce Zhang, Cenyuan Zhang, Chao Zhang, Chaoli Zhang, Chaoran Zhang, Chen Zhang, Chen Zhang, Cheng Zhang, Chi Zhang, Chiyu Zhang, Chong Zhang, Chunxia Zhang, Crystina Zhang, Daoan Zhang, Dong Zhang, Duzhen Zhang, Feng Zhang, Hainan Zhang, Hanlin Zhang, Hao Zhang, Haopeng Zhang, Haotian Zhang, Hongbin Zhang, Hongming Zhang, Hongyang R. Zhang, Hongyu Zhang, Huajian Zhang, Jiajie Zhang, Jianyang Zhang, Jiabin Zhang, Jihai Zhang, Jinchuan Zhang, Jing Zhang, Jingqing Zhang, Jipeng Zhang, Jiwen Zhang, Junjie Zhang, Junlei Zhang, Junzhe Zhang, Kai Zhang, Kai Zhang, Kai Zhang, Kai Zhang, Kaiyan Zhang, Le Zhang, Lei Zhang, Li Lyna Zhang, Liang Zhang, Lichao Zhang, Licheng Zhang, Lin Zhang, Linfan Zhang, Linhai Zhang, Linhao Zhang, Liwen Zhang, Longhui Zhang, Miaoran Zhang, Mike Zhang,

Min Zhang, Mingyang Zhang, Pei Zhang, Peitian Zhang, Qi Zhang, Qiang Zhang, Qinglin Zhang, Richong Zhang, Rongzhi Zhang, Ruoyu Zhang, Shaolei Zhang, Sheng Zhang, Shengqiang Zhang, Shengyu Zhang, Shujian Zhang, Shunyu Zhang, Shurui Zhang, Songming Zhang, Taolin Zhang, Tianhang Zhang, Tianlin Zhang, Tianyi Zhang, Tianyi Zhang, Wei Emma Zhang, Wen Zhang, Wenqi Zhang, Wenqian Zhang, Wenqiao Zhang, Wentao Zhang, Wenxuan Zhang, Xi Zhang, Xiang Zhang, Xiaodan Zhang, Xiaotong Zhang, Xiaowang Zhang, Xin Zhang, Xinghua Zhang, Xinliang Frederick Zhang, Xuan Zhang, Xuanming Zhang, Xuanyu Zhang, Xuejie Zhang, Xuhong Zhang, Xulong Zhang, Yan Zhang, Yan Zhang, Yangjun Zhang, Yanzhe Zhang, Yao Zhang, Yaping Zhang, Yazhou Zhang, Yi Zhang, Yi Zhang, YiFan Zhang, Yian Zhang, Yice Zhang, Yichi Zhang, Ying Zhang, Ying Zhang, Yiqun Zhang, Yizhou Zhang, Yong Zhang, Yonggang Zhang, Yu Zhang, Yu Zhang, Yuan Zhang, Yuansen Zhang, Yuanzhe Zhang, Yubo Zhang, Yuhao Zhang, Yuhao Zhang, Yuhui Zhang, Yuji Zhang, Yunjian Zhang, Yunxiang Zhang, Yunyi Zhang, Yuqi Zhang, Yusen Zhang, Yuwei Zhang, Yuxiang Zhang, Zecheng Zhang, Zeliang Zhang, Zeyu Zhang, Zhehao Zhang, Zhengyan Zhang, Zhenru Zhang, Zhenyu Zhang, Zhexin Zhang, Zheyuan Zhang, Zhihan Zhang, Zhiyuan Zhang, Zhong Zhang, Zhongbao Zhang, Zhongping Zhang, Zhuo Zhang, Zihan Zhang, Ziheng Zhang, Zixuan Zhang, Zizheng Zhang, Qilong Zhangli, Bowen Zhao, Chao Zhao, Chen Zhao, Chenyang Zhao, Haiyan Zhao, Handong Zhao, Hao Zhao, Hongke Zhao, Huan Zhao, Jiahao Zhao, Jianyu Zhao, Jiaqi Zhao, Jiaxu Zhao, Jin Zhao, Junbo Zhao, Junchen Zhao, Justin Zhao, Kai Zhao, Liang Zhao, Liang Zhao, Libo Zhao, Mengjie Zhao, Minyi Zhao, Pu Zhao, Qinghua Zhao, Runcong Zhao, Shuai Zhao, Tiancheng Zhao, Tuo Zhao, Weilin Zhao, Weixiang Zhao, Wenbo Zhao, Wenlong Zhao, Wenting Zhao, Xiang Zhao, Xingyi Zhao, Xiutian Zhao, Xuandong Zhao, Xuechen Zhao, Xujiang Zhao, Yang Zhao, Yangyang Zhao, Yanyan Zhao, Yi Zhao, Yilun Zhao, Yiran Zhao, Yiren Zhao, Yizhou Zhao, Zheng Zhao, Zhenjie Zhao, Zhuokai Zhao, Ce Zheng, Chen Zheng, Jonathan Zheng, Junhao Zheng, Kai Zheng, Lecheng Zheng, Rui Zheng, Shen Zheng, Siqi Zheng, Xiaochen Zheng, Xiaoqing Zheng, Yandan Zheng, Yefeng Zheng, Yinhe Zheng, Yuanhang Zheng, Zaiyi Zheng, Zhisheng Zheng, Guoqiang Zhong, Ming Zhong, Qihuang Zhong, Ruiqi Zhong, Shanshan Zhong, Wanjun Zhong, Weihong Zhong, Yang Zhong, Yaoyao Zhong, Yiran Zhong, Baohang Zhou, Changzhi Zhou, Chuang Zhou, Deyu Zhou, Dong Zhou, Fan Zhou, Han Zhou, Hanzhang Zhou, Haoyi Zhou, Houquan Zhou, Ji-Zhe Zhou, Jiaming Zhou, Jianing Zhou, Jiawei Zhou, Jinfeng Zhou, Jingbo Zhou, Joey Tianyi Zhou, Kaitlyn Zhou, Li Zhou, Mengyu Zhou, Naitian Zhou, Shiji Zhou, Shilin Zhou, Sizhe Zhou, Wei Zhou, Wei Zhou, Xiabing Zhou, Xiaobing Zhou, Xiaofei Zhou, Xuhui Zhou, Yan Zhou, Yang Zhou, Yangqiaoyu Zhou, Yaqian Zhou, Yichao Zhou, Yichu Zhou, Yilun Zhou, Yiyang Zhou, Yongwei Zhou, Yucheng Zhou, Yuhang Zhou, Yunhua Zhou, Yuxiang Zhou, Yuxuan Zhou, Zhijie Zhou, Dawei Zhu, Derui Zhu, Dongsheng Zhu, Dongxiao Zhu, Fengbin Zhu, Henghui Zhu, Jennifer Zhu, Jian Zhu, Kenny Q. Zhu, Kunlun Zhu, Lixing Zhu, Luyao Zhu, Ming Zhu, Muhua Zhu, Qingfu Zhu, Qinglin Zhu, Qingqing Zhu, Shaolin Zhu, Shengqi Zhu, Suyang Zhu, Tinghui Zhu, Tong Zhu, Wang Zhu, Wenhao Zhu, Wenhong Zhu, Xiaochen Zhu, Xinyu Zhu, Xuekai Zhu, Yingjie Zhu, Yongxin Zhu, Yutao Zhu, Zhangchi Zhu, Zhihong Zhu, Zhiyuan Zhu, Zining Zhu, Ziwei Zhu, Chengxu Zhuang, Fuzhen Zhuang, Honglei Zhuang, Xinlin Zhuang, Yuan Zhuang, Yuchen Zhuang, Yueting Zhuang, Jingming Zhuo, Terry Yue Zhuo, Heike Zinsmeister, Ran Zmigrod, Qing Zong, Shi Zong, Henry Peng Zou, Weijin Zou, Yicheng Zou, Yuexian Zou, Amal Zouaq, Vilém Zouhar, Guido Zuccon, Andrea Zugarini, Xinyu Zuo, Naama Zwerdling, Maïke Züfle

## Outstanding Reviewers

Alafate Abulimiti, Muhammad Farid Adilazuarda, Amit Agarwal, Sumit Agarwal, Jai Aggarwal, Mohamed Aghzal, Aishwarya Agrawal, Kumar Krishna Agrawal, Sweta Agrawal, Carlos Alejandro Aguirre, Jaewoo Ahn, Kian Ahrabian, Xi Ai, Anirudh Ajith, Mahmud Elahi Akhter, Kosuke Akimoto, Ekin Akyürek, Abeer Aldayel, Duarte Miguel Alves, Silvio Amir, Wenbin An, Aparna

Ananthasubramaniam, Tatiana Anikina, Jean-Yves Antoine, Vladimir Araujo, Jordi Armengol-Estapé, David Arps, Nicholas Asher, Ali Athar, Tal August, Hammad Ayyubi

Nikolay Babakov, Selene Baez Santamaria, Ashutosh Baheti, Min Bai, Ananth Balashankar, Nishant Balepur, Naman Bansal, Edoardo Barba, Leslie Barrett, Elisa Bassignana, Barend Beekhuizen, Catarina G Belém, Christian Bentz, Timothée Bernard, Gabriel Bernier-Colborne, Max Berendorf, Amanda Bertsch, Marie Bexte, Anne Beyer, Indrajit Bhattacharya, Rajarshi Bhowmik, Russa Biswas, Johanna Björklund, Daniele Bonadiman, Philipp Borchert, Antoine Bosselut, Kaj Bostrom, Gilles Boulianne, Andrey Bout, Joseph Spartacus Boyle, Ana Brassard, Sergio Burdisso, Paula Buttery, Lisa Bylinina

Cristian Buc Calderon, Nitay Calderon, Iacer Calixto, Marc Canby, Rui Cao, Yu Cao, Boaz Carmeli, Julie Carson-Berndsen, Federico Cassano, Francesco Cazzaro, Christophe Cerisara, Alberto Cetoli, Taehun Cha, Aman Chadha, Hyungjoo Chae, Chunkit Chan, David Chan, Agneet Chatterjee, Boxing Chen, Canyu Chen, Catherine Chen, Cen Chen, Guanyi Chen, Guanzheng Chen, Guiming Hardy Chen, Haonan Chen, Jiangjie Chen, Jindong Chen, Maximillian Chen, Nuo Chen, Shiwei Chen, Sishuo Chen, Wei-Lin Chen, Yanda Chen, Yilong Chen, Yiming Chen, Yingfa Chen, Zhengyu Chen, Ziru Chen, Hong Cheng, Junyan Cheng, Xiaoxia Cheng, Xize Cheng, Jackie CK Cheung, Cheng-Han Chiang, Ting-Rui Chiang, Jungwook Choi, Yung-Sung Chuang, Ekapol Chuangsuwanich, Mark Cieliebak, Philipp Cimiano, Christopher Clark, Elizabeth Clark, Charles L. A. Clarke, Marco Cagnetta, Amir David Nissan Cohen, Nachshon Cohen, Gao Cong, Fabio Cozman, Mathias Creutz, Emilio Villa Cueva, Peng Cui, Wendi Cui

Ido Dagan, Lirong Dai, Yong Dai, Dipankar Das, Gaël De Chalendar, Hiroyuki Deguchi, Naihao Deng, Yue Deng, Luca Di Liello, Wentao Ding, Yifan Ding, Tuan Dinh, Anuj Diwan, Quyet V. Do, Miguel Domingo, Haoyu Dong, Ming Dong, Yuxiao Dong, Jordan Dotzel, Mark Dras, Mark Dredze, Rotem Dror, Yichao Du, Yufeng Du, Liam Dugan, Ewan Dunbar

Lukas Edman, Hafsteinn Einarsson, Roald Eiselen, Roxanne El Baff, Mohamed Elgaar, Aykut Erdem, Erkut Erdem

Fahim Faisal, Neele Falk, Yao-Chung Fan, Tao Fang, Dongji Feng, Shangbin Feng, Yujie Feng, Javier Ferrando, James D. Finch, Antske Fokkens, Negar Foroutan, James R. Foulds, Meaghan Fowlie, Dayne Freitag, Giacomo Frisoni, Dennis Fucci, Atsushi Fujita

Kamel Gaanoun, Jay Gala, Vineet Gandhi, Prakhar Ganesh, William Gantt, Difei Gao, Hongcheng Gao, Jinglong Gao, Gregor Geigle, Matthieu Geist, Aryo Pradipta Gema, Ariel Gera, Harritxu Gete, Iacopo Ghinassi, Sourav Ghosh, Salvatore Giorgi, Max Glockner, Ann-Sophie Gnehm, Ameya Godbole, Charles Godfrey, Jonas Golde, Sujatha Das Gollapalli, Chen Gong, Shansan Gong, Zhuocheng Gong, Julio Gonzalo, Maharshi Gor, Venkata S Govindarajan, Abhimanyu Goyal, Navita Goyal, Derek Greene, Matt Grenander, Milan Gritta, Naibin Gu, Nianlong Gu, Yu Gu, Zijin Gu, Sagar Gubbi Venkatesh, Reto Gubelmann, Marco Guerini, Nuno M Guerreiro, Michael Guerzhoy, Varun Gumma, Chunxi Guo, Xu Guo, Ashim Gupta, Prakhar Gupta

Lovisa Hagström, Joonghyuk Hahn, Rajarshi Haldar, Coleman Haley, Chi Han, Michael Hanna, Fabrice Y Harel-Canada, Keith Harrigian, Maram Hasanain, Kazuma Hashimoto, Oktie Hassanzadeh, Katsuhiko Hayashi, Han He, Junqing He, Xiangheng He, Xiaodong He, Xuming He, Jack Hessel, Musashi Hinck, Namgyu Ho, Pengfei Hong, Ruixin Hong, Seokhee Hong, Andrea Horbach, Sho Hoshino, Mohammad Javad Hosseini, Guiyang Hou, Yuexian Hou, Zhaoyi Joey Hou, Jinpeng Hu, Linmei Hu, Minghao Hu, Tiancheng Hu, Wenxiang Hu, Yibo Hu, Zhiyuan Hu, Dong Huang, Guanhua Huang, Jen-tse Huang, Xinting Huang, Pere-Lluís Hugué Cabot, Lluís-F. Hur

tado, Ben Hutchinson, Katharina Hämmerl

Robert L. Logan IV, Nikolai Ilinykh, Mert Inan, Go Inoue, Ioan-Bogdan Iordache, Tatsuya Ishigaki, Hamish Ivison, Bhavani Iyer

Naman Jain, Fran Jelenić, Ananya Harsh Jha, Qi Jia, Haiyun Jiang, Xin Jiang, Zhuoren Jiang, Wenxiang Jiao, Peng Jin, Xiaolong Jin, Wang Jing, Abhinav Joshi, Nitish Joshi, Rishabh Joshi, David Jurgens

Marek Kadlčík, Dariusz Kajtoch, Nishant Kambhatla, Ryo Kamoi, Nikhil Kandpal, Feiyang Kang, Jaewoo Kang, Zhao Kang, Raghav Kapoor, Debanjana Kar, Antonia Karamolegkou, Priyanka Kar Gupta, Tushar Kataria, Anisia Katinskaia, Hideto Kazawa, Lee Kezar, Abdul Rafae Khan, Omar Khattab, Halil Kilicoglu, Gangwo Kim, Geewook Kim, Gene Louis Kim, Jongho Kim, Juyong Kim, Minsoo Kim, Misuk Kim, Youna Kim, Yasutomo Kimura, Brendan King, Hannah Rose Kirk, Jens Kleesiek, Mateusz Klimaszewski, Marius Kloft, Filip Klubicka, Rebecca Knowles, Ichiro Kobayashi, Ari Kobren, Prashant Kodali, Konstantinos Kogkalidis, Ryosuke Kohita, Takeshi Kojima, Jun Kong, Bevan Koopman, Michalis Korakakis, Fajri Koto, Neema Kotonya, Elisa Kreiss, Satyapriya Krishna, Mateusz Krubiński, Nitish Kulkarni, Kemal Kurniawan, Iliia Kuznetsov, Jingu Kwon

Peifeng LI, Matthieu Labeau, Salem Lahlou, Allison Lahnala, Tian Lan, Daniil Larionov, Stefan Larson, Hung Le, Bruce W. Lee, Chia-Hsuan Lee, En-Shiun Annie Lee, Jaeho Lee, Jaeseong Lee, Jay-Yoon Lee, Kyumin Lee, Lillian Lee, Yongjae Lee, Young-Jun Lee, Chuan Lei, Pietro Lesci, Bei Li, Bowen Li, Bryan Li, Changye Li, Chaozhuo Li, Chengkai Li, Chengxi Li, Diya Li, Dong Li, Fanrong Li, Haochen Li, Haoxin Li, Haoyuan Li, Hui Li, Jianjun Li, Jiakuan Li, Jiazheng Li, Jieyu Li, Judith Yue Li, Junlong Li, Kun Li, Li Li, Liang Li, Lujun Li, Ming Li, Qi Li, Ruizhe Li, Shuyue Stella Li, Xiangci Li, Xinlin Li, Xinze Li, Yinghui Li, Yingya Li, Davis Liang, Baohao Liao, Hsien-chin Lin, Weizhe Lin, Zhen Lin, Matthias Lindemann, Gili Lior, Marina Litvak, Andy Liu, Bingshuai Liu, Emmy Liu, Hui Liu, Lema Liu, Peilin Liu, Ran Liu, Shudong Liu, Sijia Liu, Siyang Liu, Xiang Liu, Yang Janet Liu, Yihong Liu, Yinhong Liu, YongKang Liu, Yongbin Liu, Zoey Liu, Xinwei Long, Yunfei Long, Henrique Lopes Cardoso, Cedric Lothritz, Chao Lou, Bo-Ru Lu, Di Lu, Weiming Lu, Bernd Ludwig, Ling Luo, Chunchuan Lyu, Xing Han Lù

Meryem M'hamdi, Da Ma, Jiefeng Ma, Xinbei Ma, Xueguang Ma, Yun Ma, Zhengrui Ma, Jakub Macina, Rahul Madhavan, Aru Maekawa, Ian Magnusson, Adyasha Maharana, Sainik Kumar Mahata, Aaron Maladry, Sri Raghu Malireddi, Christopher Malon, Pranav Mani, Jiabin Mao, Kelong Mao, Qianren Mao, Yunlong Mao, Zhiming Mao, Marion Di Marco, Magdalena Markowska, Edison Marrese-Taylor, Ved Mathai, Takuya Matsuzaki, Florian Matthes, Kaushal Kumar Maurya, R. Thomas McCoy, Daniel McDuff, Nick McKenna, Nicholas Meade, Nikhil Mehta, John Mendonça, Yuanliang Meng, William Merrill, Vladislav Mikhailov, Elena Mikhalkova, Jeremiah Milbauer, Koji Mineshima, Benjamin Minixhofer, Yusuke Miyao, Taro Miyazaki, Ali Modarressi, Lucas Moeller, Hyeongdon Moon, Hyeonseok Moon, Robert Moro, David R Mortensen, Xinyi Mou, Basel Mousi, Yongyu Mu, Anjishnu Mukherjee, Ricardo Muñoz Sánchez

Siddharth N, Clara Na, Atharva Naik, Inderjeet Jayakumar Nair, Ananjan Nandi, Mukuntha Narayanan S, Pranav Narayanan Venkit, Shravan Nayak, Seyed Parsa Neshaei, See-Kiong Ng, Youyang Ng, Kiet Van Nguyen, Vlad Niculae, Pengyu Nie, Irina Nikishina, Jinzhong Ning, Bill Noble, Enrique Noriega-Atala

Shu Okabe, Eda Okur, Hugo Gonçalo Oliveira, Andreas Opedal, Silviu Vlad Oprea, Simon Ostermann, Yulia Otmakhova

Vishakh Padmakumar, Kuntal Kumar Pal, Rong Pan, Rui Pan, Haris Papageorgiou, Duccio Papadopulo, Letitia Parcalabescu, ChaeHun Park, Choonghyun Park, Kunwoo Park, Mihir Parmar, Eliana Pastor, Panupong Pasupat, Braja Patra, Manasi Patwardhan, Amalie Brogaard Pauli, Kellin Pelrine, Bo Peng, Letian Peng, Jan-Thorsten Peter, Suzanne Petryk, Shubham Milind Phal, Quang Pham, Tanzir Pial, Barbara Plank, Adam Poliak, Girishkumar Ponkiya, Amrit Poudel, Aniket Pramanick, Massimiliano Pronesti, Rifki Afina Putri, Adarsh Pyarelal

Dong Qian, Linlu Qiu, Muhammad Reza Qorib

Rafael Rafailov, Edward Raff, Dinesh Raghu, Vyas Raina, Lisa Raitel, Owen Rambow, Rita Ramos, Hannah Rashkin, Abhinav Rastogi, Jishnu Ray Chowdhury, Yuval Reif, Jie Ren, Pengjie Ren, Stephen D. Richardson, Caitlin Laura Richter, Korbinian Riedhammer, Angus Roberts, Melissa Roemmele, Domenic Rosati, Candace Ross, Kashob Kumar Roy, Nikolai Rozanov, Benjamin I. P. Rubinstein, Daniel Ruffinelli, Federico Ruggeri, Stefan Ruseti

Arkadiy Saakyan, Ashish Sabharwal, Kaushik Ram Sadagopan, Yusuke Sakai, Alejandro R. Salamanca, Abhilasha Sancheti, Vicente Ivan Sanchez Carmona, Abraham Sanders, Kate Sanders, Irina Saparina, Abulhair Saparov, Xabier Saralegi, Gabriele Sarti, Guergana K Savova, Michael Saxon, Yves Scherrer, Florian Schneider, Nathan Schneider, Hinrich Schuetze, Anastasiia Sedova, Amit Seker, Oleg Serikov, Sofia Serrano, Rita Sevastjanova, Shaden Shaar, Md Shihab Shahriar, Valerie Shalin, Natalie Shapira, Ori Shapira, Ashish Sharma, Drishti Sharma, Harsh Sharma, Kartik Sharma, Shivam Sharma, Serge Sharoff, Yikang Shen, Qiang Sheng, Quan Z. Sheng, Pranav Shetty, Ning Shi, Wenqi Shi, Hidetoshi Shimodaira, Kazutoshi Shinoda, Jiasheng Si, Karan Sikka, Carina Silberer, Fabrizio Silvestri, Gabriel Simmons, Nikhil Singh, Christopher T. Small, Dawn Song, Taylor Sorensen, Laure Soulier, Gerasimos Spanakis, Alexander Spangher, Lucia Specia, Andreas Spitz, Neha Srikanth, Padmini Srinivasan, Saurabh Srivastava, Joe Stacey, Felix Stahlberg, Katherine Stasaski, Julius Steen, Hannah Sterz, Brandon M. Stewart, Rickard Stureborg, Sebastian Stüker, Melanie Subbiah, Katsuhito Sudoh, Jimin Sun, Shichao Sun, Weiqi Sun, Yutao Sun, Jayaprakash Sundararaj, Masahiro Suzuki, Anej Svete, Maciej Szankin, Eduardo Sánchez, Felipe Sánchez-Martínez, Michal Štefánik

Sho Takase, Kunihiro Takeoka, Pradyumna Tambwekar, Fiona Anting Tan, Fuwen Tan, Liling Tan, Yixuan Tang, Parth Thakkar, Sarubi Thillainathan, Tristan Thrush, Ran Tian, Yufei Tian, MeiHan Tong, Khanh Quoc Tran, David Traum, Sean Trott, Tornike Tsereteli

Adaku Uchendu, Can Udomcharoenchaikit, Ashok Urlana, Ricardo Usbeck

Venktesh V, Saujas Vaduguru, Josef Valvoda, Jannis Vamvas, Lindsey Morgan Vanderlyn, Daniel Varab, Praveen Srinivasa Varadhan, Justin Vasselli, Giulia Venturi, Sahil Verma, Krishnapriya Vishnubhotla, Vijay Viswanathan, Nikolas Vitsakis, Juraj Vladika, Huy-Hien Vu

Eitan Wagner, Jan Philip Wahle, Hiromi Wakaki, Andreas Waldis, Thiemo Wambsganss, Ruyuan Wan, Xingchen Wan, Yingjia Wan, Cheng Wang, Chenglong Wang, Danqing Wang, Guanchu Wang, Kexiang Wang, RuiBo Wang, Ruoyao Wang, Shuohuan Wang, Songhe Wang, Wei Wang, Xiao Wang, Yang Wang, Yichen Wang, Yiming Wang, Yiran Wang, Yu Wang, Zehao Wang, Zhaowei Wang, Zhilin Wang, Zihan Wang, Zijie Wang, Taro Watanabe, Tharindu Cyril Weerasooriya, Sheng-Lun Wei, Jiabin Wen, Zhihua Wen, Yixuan Weng, Robert West, Alexander Wettig, Spencer Whitehead, Michael Wiegand, Matti Wiegmann, Adam Wiemerslage, Bram Willemsen, Shomir Wilson, Tomer Wolfson, Zach Wood-Doughty, Anne Wu, Di Wu, Feijie Wu, Ji Wu, Jinge Wu, Linjuan Wu, Xinyi Wu, Zhijing Wu, Zirui Wu

Heming Xia, Wei Xiang, Likang Xiao, Kaige Xie, Yuexiang Xie, Yuqing Xie, Yuxi Xie, Chunlei Xin, Weimin Xiong, Hongzhi Xu, Wenda Xu, Xin Xu, Yi Xu, Yige Xu, Zhenran Xu, Huiyin Xue

Jing Nathan Yan, Xin Yan, Chenghao Yang, Hsiu-Yu Yang, Jian Yang, Mingming Yang, Songhua Yang, Wenkai Yang, Yaming Yang, Yazheng Yang, Yuqing Yang, Zhicheng Yang, Zhihao Yang, Peiran Yao, Yixiang Yao, Zijun Yao, Michal Yarom, Yang Ye, Akhila Yerukola, Congchi Yin, Ming Yin, Michael Miller Yoder, Gal Yona, Zheng Xin Yong, Soyoun Yoon, Issei Yoshida, Hiyori Yoshikawa, Haofei Yu, Pengfei Yu, Sangwon Yu, Tianshu Yu, Xiao Yu, Xinyan Velocity Yu, Jiahong Yuan, Moy Yuan, Wei Yuan, Yifei Yuan, Zheng Yuan

Runzhe Zhan, Caiqi Zhang, Chen Zhang, Chong Zhang, Feng Zhang, Haoran Ranran Zhang, Jinchuan Zhang, Jingqing Zhang, Jiwen Zhang, Junzhe Zhang, Kai Zhang, Lechen Zhang, Michael JQ Zhang, Peiyuan Zhang, Ruochen Zhang, Xinliang Frederick Zhang, Yanzhe Zhang, YiFan Zhang, Yice Zhang, Yiming Zhang, Ying Zhang, Yu Zhang, Yuhui Zhang, Yunyi Zhang, Yusen Zhang, Yuwei Zhang, Zhengyan Zhang, Zhenyu Zhang, Zhihan Zhang, Bowen Zhao, Jiaxu Zhao, Junbo Zhao, Pu Zhao, Tuo Zhao, Xingmeng Zhao, Yi Zhao, Ce Zheng, Chen Zheng, Kai Zheng, Lecheng Zheng, Gengze Zhou, Hanzhang Zhou, Jiawei Zhou, Jing Zhou, Junkai Zhou, Naitian Zhou, Yang Zhou, Yaqian Zhou, Yichao Zhou, Dawei Zhu, Fangqi Zhu, Lixing Zhu, Shengqi Zhu, Tongyao Zhu, Xiaochen Zhu, Chengxu Zhuang, Xinlin Zhuang, Roger Zimmermann, Ran Zmigrod, Yueshan Zou, Pierre Zweigenbaum, Maiké Züfle

# Keynote Talk

## Open-Source and Science in the Era of Foundation Models

**Percy Liang**  
Stanford University



**Tue, November 12th, 2024 – Time: 09:30 – 10:30 – Room: James Knight Center**

**Abstract:** As capabilities of foundation models skyrocket, openness plummets. In this talk, I argue that open-source models are essential for the long-term goal of building a rigorous foundation for AI. Greater access—from API to open-weight to open-source—enables deeper forms of research. API access allows us to push the frontier of agents, and I will present our recent work on simulation and problem-solving agents. Open weights enables reproducible research on safety, interpretability, and more generally, “model forensics”. Open-source unlocks fundamental innovations in architectures, training procedures, and data curation methods. Of course, the key obstacle for building open-source models is the resources required (data, compute, and research/engineering). I will conclude with some promising directions that leverage the community that bring us closer to the vision of open-source foundation models.

**Bio:** Percy Liang is an Associate Professor of Computer Science at Stanford University (B.S. from MIT, 2004; Ph.D. from UC Berkeley, 2011) and the director of the Center for Research on Foundation Models (CRFM). He is currently focused on making foundation models (in particular, language models) more accessible through open-source and understandable through rigorous benchmarking. In the past, he has worked on many topics centered on machine learning and natural language processing, including robustness, interpretability, human interaction, learning theory, grounding, semantics, and reasoning. He is also a strong proponent of reproducibility through the creation of CodaLab Worksheets. His awards include the Presidential Early Career Award for Scientists and Engineers (2019), IJCAI Computers and Thought Award (2016), an NSF CAREER Award (2016), a Sloan Research Fellowship (2015), a Microsoft Research Faculty Fellowship (2014), and paper awards at ACL, EMNLP, ICML, COLT, ISMIR, CHI, UIST, and RSS.



# Keynote Talk

## My Journey in AI Safety and Alignment

Anca Dragan

University of California Berkeley and Google Deepmind



**Wed, November 13th, 2024 – Time: 09:00 – 10:00 – Room: James Knight Center**

**Abstract:** For nearly a decade now, the problem that has been top of mind for me is how we might enable AI systems to robustly optimize for what people want, and to avoid causing harm – from robots and self-driving cars, to assistive devices and deep brain stimulation, to theory and toy models, to large language models and now Gemini. In this talk, I’ll take the opportunity to share a bit about my journey in this space, what lessons I’ve learned, and how we’re approaching the safety and alignment of frontier models at Google DeepMind.

**Bio:** Anca Dragan is an Associate Professor in the EECS Department at UC Berkeley, currently on leave to head AI Safety and Alignment at Google DeepMind. The goal of her research at UC Berkeley has been to enable AI agents (from robots to cars to LLMs to recommender systems) to work with, around, and in support of people. Anca runs the InterACT Lab, where they focus on algorithms for human-AI and human-robot interaction. One of the core problems the Lab has worked on since its inception is AI alignment: getting AI agents to do what people actually want – this has meant learning reward functions interactively, from diverse human feedback forms, across different modalities, while maintaining uncertainty. They have also contributed to algorithms for human-AI collaboration and coordination, like agents fluently working together with human-driven avatars in games, assistance and adaption in brain-machine interfaces, and autonomous cars sharing the road with human drivers. At Google DeepMind, Anca currently leads a collection of teams responsible for both the safety of the current Gemini models, and preparing for Gemini capabilities to keep advancing and ensuring that safety advances hand-in-hand. This means ensuring Gemini models are and will be aligned with human goals and values, including avoiding present-day harms and catastrophic risks, enabling models to better and more robustly understand human preferences, enabling informed oversight, increasing robustness to adversarial attacks, and accounting for the plurality of human values and viewpoints. Previously, she helped found and serve on the steering committee for the Berkeley AI Research (BAIR) Lab. Anca has been (and still is) a co-PI of the Center for Human-Compatible AI. She has consulted for Waymo for the past 6 years, helping with the roadmap for how to deploy an increasingly learning-based safety-critical system. She’s been honored by the Sloan Fellowship, MIT TR35, the Okawa award, an NSF CAREER award, and the PECASE award. Anca takes the most pride in my former students, who have gone on to faculty positions at MIT, Stanford, CMU, and Princeton, and to industry positions at Google DeepMind, Waymo, and Meta.

# Keynote Talk

## Bayes in the age of intelligent machines

Tom Griffiths  
Princeton University



**Thu, November 14th, 2024 – Time: 09:00 – 10:00 – Room: James Knight Center**

**Abstract:** Recent rapid progress in the creation of artificial intelligence (AI) systems has been driven in large part by innovations in architectures and algorithms for developing large scale artificial neural networks. As a consequence, it's natural to ask what role abstract principles of intelligence — such as Bayes' rule — might play in developing intelligent machines. In this talk, I will argue that there is a new way in which Bayes can be used in the context of AI, more akin to how it is used in cognitive science: providing an abstract description of how agents should solve certain problems and hence a tool for understanding their behavior. This new role is motivated in large part by the fact that we have succeeded in creating intelligent systems that we do not fully understand, making the problem for the machine learning researcher more closely parallel that of the cognitive scientist. I will talk about how this perspective can help us think about making machines with better informed priors about the world and give us insight into their behavior by directly creating cognitive models of neural networks.

**Bio:** Tom Griffiths is the Henry R. Luce Professor of Information Technology, Consciousness and Culture in the Departments of Psychology and Computer Science at Princeton University, where he is also the Director of the Princeton Laboratory for Artificial Intelligence. His research explores connections between human and machine learning, using ideas from statistics and artificial intelligence to understand how people solve the challenging computational problems they encounter in everyday life. Tom completed his PhD in Psychology at Stanford University in 2005, and taught at Brown University and the University of California, Berkeley before moving to Princeton. He has received awards for his research from organizations ranging from the American Psychological Association to the National Academy of Sciences and is a co-author of the book *Algorithms to Live By*, introducing ideas from computer science and cognitive science to a general audience.

## Table of Contents

<i>UniGen: Universal Domain Generalization for Sentiment Classification via Zero-shot Dataset Generation</i>	
Juhwan Choi, Yeonghwa Kim, Seunguk Yu, JungMin Yun and YoungBin Kim . . . . .	1
<i>Multi-News+: Cost-efficient Dataset Cleansing via LLM-based Data Annotation</i>	
Juhwan Choi, JungMin Yun, Kyohoon Jin and YoungBin Kim . . . . .	15
<i>FIZZ: Factual Inconsistency Detection by Zoom-in Summary and Zoom-out Document</i>	
Joonho Yang, Seunghyun Yoon, ByeongJeong Kim and Hwanhee Lee . . . . .	30
<i>Prompts have evil twins</i>	
Rimon Melamed, Lucas Hurley McCabe, Tanay Wakhare, Yejin Kim, H. Howie Huang and Enric Boix-Adserà . . . . .	46
<i>Table Question Answering for Low-resourced Indic Languages</i>	
Vaishali Pal, Evangelos Kanoulas, Andrew Yates and Maarten de Rijke . . . . .	75
<i>ImageInWords: Unlocking Hyper-Detailed Image Descriptions</i>	
Roopal Garg, Andrea Burns, Burcu Karagol Ayan, Yonatan Bitton, Ceslee Montgomery, Yasumasa Onoe, Andrew Bunner, Ranjay Krishna, Jason Michael Baldrige and Radu Soricut . . . . .	93
<i>LLM-Based Agent Society Investigation: Collaboration and Confrontation in Avalon Gameplay</i>	
Yihuai Lan, Zhiqiang Hu, Lei Wang, Yang Wang, Deheng Ye, Peilin Zhao, Ee-Peng Lim, Hui Xiong and Hao Wang . . . . .	128
<i>When LLMs Meets Acoustic Landmarks: An Efficient Approach to Integrate Speech into Large Language Models for Depression Detection</i>	
Xiangyu Zhang, Hexin Liu, Kaishuai Xu, Qiquan Zhang, Daijiao Liu, Beena Ahmed and Julien Epps . . . . .	146
<i>Speaking in Wavelet Domain: A Simple and Efficient Approach to Speed up Speech Diffusion Model</i>	
Xiangyu Zhang, Daijiao Liu, Hexin Liu, Qiquan Zhang, Hanyu Meng, Leibny Paola Garcia Perera, EngSiong Chng and Lina Yao . . . . .	159
<i>Hateful Word in Context Classification</i>	
Sanne Hoeken, Sina Zarrieß and Özge Alacam . . . . .	172
<i>Eyes Don't Lie: Subjective Hate Annotation and Detection with Gaze</i>	
Özge Alacam, Sanne Hoeken and Sina Zarrieß . . . . .	187
<i>NúmeroLogic: Number Encoding for Enhanced LLMs' Numerical Reasoning</i>	
Eli Schwartz, Leshem Choshen, Joseph Shtok, Sivan Doveh, Leonid Karlinsky and Assaf Arbelle . . . . .	206
<i>ThinkingFair and Slow: On the Efficacy of Structured Prompts for Debiasing Language Models</i>	
Shaz Furniturewala, Surgan Jandial, Abhinav Java, Pragyana Banerjee, Simra Shahid, Sumit Bhatia and Kokil Jaidka . . . . .	213
<i>A Usage-centric Take on Intent Understanding in E-Commerce</i>	
Wendi Zhou, Tianyi Li, Pavlos Vougiouklis, Mark Steedman and Jeff Z. Pan . . . . .	228
<i>Fine-Tuning or Retrieval? Comparing Knowledge Injection in LLMs</i>	
Oded Ovadia, Menachem Brief, Moshik Mishaeli and Oren Elisha . . . . .	237

<i>Systematic Biases in LLM Simulations of Debates</i> Amir Taubenfeld, Yaniv Dover, Roi Reichart and Ariel Goldstein .....	251
<i>Studying and Mitigating Biases in Sign Language Understanding Models</i> Katherine Atwell, Danielle Bragg and Malihe Alikhani .....	268
<i>Uncertainty in Language Models: Assessment through Rank-Calibration</i> Xinmeng Huang, Shuo Li, Mengxin Yu, Matteo Sesia, Hamed Hassani, Insup Lee, Osbert Bastani and Edgar Dobriban .....	284
<i>RoTBench: A Multi-Level Benchmark for Evaluating the Robustness of Large Language Models in Tool Learning</i> Junjie Ye, Yilong Wu, Songyang Gao, Caishuang Huang, Sixian Li, Guanyu Li, Xiaoran Fan, Qi Zhang, Tao Gui and Xuanjing Huang .....	313
<i>Learning Planning-based Reasoning by Trajectories Collection and Process Reward Synthesizing</i> Fangkai Jiao, Chengwei Qin, Zhengyuan Liu, Nancy F. Chen and Shafiq Joty .....	334
<i>Scaling Properties of Speech Language Models</i> Santiago Cuervo and Ricard Marxer .....	351
<i>We Demand Justice!": Towards Social Context Grounding of Political Texts</i> Rajkumar Pujari, Chengfei Wu and Dan Goldwasser .....	362
<i>An Experimental Analysis on Evaluating Patent Citations</i> Rabindra Nath Nandi, Suman Maity, Brian Uzzi and Sourav Medya .....	373
<i>Fine-Tuning Large Language Models to Translate: Will a Touch of Noisy Data in Misaligned Languages Suffice?</i> Dawei Zhu, Pinzhen Chen, Miaoran Zhang, Barry Haddow, Xiaoyu Shen and Dietrich Klakow	388
<i>Consolidating Ranking and Relevance Predictions of Large Language Models through Post-Processing</i> Le Yan, Zhen Qin, Honglei Zhuang, Rolf Jagerman, Xuanhui Wang, Michael Bendersky and Harrie Oosterhuis .....	410
<i>Strength Lies in Differences! Improving Strategy Planning for Non-collaborative Dialogues via Diversified User Simulation</i> Tong Zhang, Chen Huang, Yang Deng, Hongru Liang, Jia Liu, Zujie Wen, Wenqiang Lei and Tat-Seng Chua .....	424
<i>Impeding LLM-assisted Cheating in Introductory Programming Assignments via Adversarial Perturbation</i> Saiful Islam Salim, Rubin Yuchan Yang, Alexander Cooper, Suryashree Ray, Saumya Debray and Sazzadur Rahaman .....	445
<i>Clustering and Ranking: Diversity-preserved Instruction Selection through Expert-aligned Quality Estimation</i> Yuan Ge, Yilun Liu, Chi Hu, Weibin Meng, Shimin Tao, Xiaofeng Zhao, Mahong Xia, Zhang Li, Boxing Chen, Hao Yang, Bei Li, Tong Xiao and JingBo Zhu .....	464
<i>On the Influence of Gender and Race in Romantic Relationship Prediction from Large Language Models</i> Abhilasha Sancheti, Haozhe An and Rachel Rudinger .....	479
<i>EmphAssess : a Prosodic Benchmark on Assessing Emphasis Transfer in Speech-to-Speech Models</i> Maureen de Seyssel, Antony D'Avirro, Adina Williams and Emmanuel Dupoux .....	495

<i>On Fake News Detection with LLM Enhanced Semantics Mining</i>	
Xiaoxiao Ma, Yuchen Zhang, Kaize Ding, Jian Yang, Jia Wu and Hao Fan . . . . .	508
<i>On Sensitivity of Learning with Limited Labelled Data to the Effects of Randomness: Impact of Interactions and Systematic Choices</i>	
Branislav Pecher, Ivan Srba and Maria Bielikova . . . . .	522
<i>Evaluating the Instruction-Following Robustness of Large Language Models to Prompt Injection</i>	
Zekun Li, Baolin Peng, Pengcheng He and Xifeng Yan . . . . .	557
<i>A Study of Nationality Bias in Names and Perplexity using Off-the-Shelf Affect-related Tweet Classifiers</i>	
Valentin Barriere and Sebastian Cifuentes . . . . .	569
<i>Mitigating the Alignment Tax of RLHF</i>	
Yong Lin, Hangyu Lin, Wei Xiong, Shizhe Diao, Jianmeng Liu, Jipeng Zhang, Rui Pan, Haoxiang Wang, Wenbin Hu, Hanning Zhang, Hanze Dong, Renjie Pi, Han Zhao, Nan Jiang, Heng Ji, Yuan Yao and Tong Zhang . . . . .	580
<i>Evaluating Readability and Faithfulness of Concept-based Explanations</i>	
Meng Li, Haoran Jin, Ruixuan Huang, Zhihao Xu, Defu Lian, Zijia Lin, Di Zhang and Xiting Wang . . . . .	607
<i>Personality-aware Student Simulation for Conversational Intelligent Tutoring Systems</i>	
Zhengyuan Liu, Stella Xin Yin, Geyu Lin and Nancy F. Chen . . . . .	626
<i>MSI-Agent: Incorporating Multi-Scale Insight into Embodied Agents for Superior Planning and Decision-Making</i>	
Dayuan Fu, Biqing Qi, Yihuai Gao, Che Jiang, Guanting Dong and Bowen Zhou . . . . .	643
<i>CoCoLoFa: A Dataset of News Comments with Common Logical Fallacies Written by LLM-Assisted Crowds</i>	
Min-Hsuan Yeh, Ruyuan Wan and Ting-Hao Kenneth Huang . . . . .	660
<i>Tokenization Is More Than Compression</i>	
Craig W Schmidt, Varshini Reddy, Haoran Zhang, Alec Alameddine, Omri Uzan, Yuval Pinter and Chris Tanner . . . . .	678
<i>FLIRT: Feedback Loop In-context Red Teaming</i>	
Ninareh Mehrabi, Palash Goyal, Christophe Dupuy, Qian Hu, Shalini Ghosh, Richard Zemel, Kai-Wei Chang, Aram Galstyan and Rahul Gupta . . . . .	703
<i>Successfully Guiding Humans with Imperfect Instructions by Highlighting Potential Errors and Suggesting Corrections</i>	
Lingjun Zhao, Khanh Xuan Nguyen and Hal Daumé Iii . . . . .	719
<i>Parameter-Efficient Sparsity Crafting from Dense to Mixture-of-Experts for Instruction Tuning on General Tasks</i>	
Haoyuan Wu, Haisheng Zheng, Zhuolun He and Bei Yu . . . . .	737
<i>GeoGPT4V: Towards Geometric Multi-modal Large Language Models with Geometric Image Generation</i>	
Shihao Cai, Keqin Bao, Hangyu Guo, Jizhi Zhang, Jun Song and Bo Zheng . . . . .	750
<i>DyVo: Dynamic Vocabularies for Learned Sparse Retrieval with Entities</i>	
Thong Nguyen, Shubham Chatterjee, Sean MacAvaney, Iain Mackie, Jeff Dalton and Andrew Yates . . . . .	767

<i>Let the Expert Stick to His Last: Expert-Specialized Fine-Tuning for Sparse Architectural Large Language Models</i>	
Zihan Wang, Deli Chen, Damai Dai, Runxin Xu, Zhuoshu Li and Yu Wu . . . . .	784
<i>LongEmbed: Extending Embedding Models for Long Context Retrieval</i>	
Dawei Zhu, Liang Wang, Nan Yang, Yifan Song, Wenhao Wu, Furu Wei and Sujian Li . . . . .	802
<i>Making Large Language Models Better Reasoners with Orchestrated Streaming Experiences</i>	
Xiangyang Liu, Junliang He and Xipeng Qiu . . . . .	817
<i>Overcome Noise and Bias: Segmentation-Aided Multi-Granularity Denoising and Debiasing for Enhanced Quarduples Extraction in Dialogue</i>	
Xianlong Luo, Meng Yang and Yihao Wang . . . . .	839
<i>Integrating Plutchik’s Theory with Mixture of Experts for Enhancing Emotion Classification</i>	
Dongjun Lim and Yun-Gyung Cheong . . . . .	857
<i>In-context Contrastive Learning for Event Causality Identification</i>	
Liang Chao, Wei Xiang and Bang Wang . . . . .	868
<i>What’s Mine becomes Yours: Defining, Annotating and Detecting Context-Dependent Paraphrases in News Interview Dialogs</i>	
Anna Wegmann, Tijs A. Van Den Broek and Dong Nguyen . . . . .	882
<i>Language Models Learn Rare Phenomena from Less Rare Phenomena: The Case of the Missing AANNs</i>	
Kanishka Misra and Kyle Mahowald . . . . .	913
<i>Large Language Models for Data Annotation and Synthesis: A Survey</i>	
Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansoor Karami, Jundong Li, Lu Cheng and Huan Liu . . . . .	930
<i>Chain-of-Dictionary Prompting Elicits Translation in Large Language Models</i>	
Hongyuan Lu, Haoran Yang, Haoyang Huang, Dongdong Zhang, Wai Lam and Furu Wei . . . . .	958
<i>AdaZeta: Adaptive Zeroth-Order Tensor-Train Adaption for Memory-Efficient Large Language Models Fine-Tuning</i>	
Yifan Yang, Kai Zhen, Ershad Banijamali, Athanasios Mouchtaris and Zheng Zhang . . . . .	977
<i>RoseLoRA: Row and Column-wise Sparse Low-rank Adaptation of Pre-trained Language Model for Knowledge Editing and Fine-tuning</i>	
Haoyu Wang, Tianci Liu, Ruirui Li, Monica Xiao Cheng, Tuo Zhao and Jing Gao . . . . .	996
<i>BlendFilter: Advancing Retrieval-Augmented Large Language Models via Query Generation Blending and Knowledge Filtering</i>	
Haoyu Wang, Ruirui Li, Haoming Jiang, Jinjin Tian, Zhengyang Wang, Chen Luo, Xianfeng Tang, Monica Xiao Cheng, Tuo Zhao and Jing Gao . . . . .	1009
<i>HEART-felt Narratives: Tracing Empathy and Narrative Style in Personal Stories with LLMs</i>	
Jocelyn J Shen, Joel Mire, Hae Won Park, Cynthia Breazeal and Maarten Sap . . . . .	1026
<i>Eliminating Biased Length Reliance of Direct Preference Optimization via Down-Sampled KL Divergence</i>	
Junru Lu, Jiazheng Li, Siyu An, Meng Zhao, Yulan He, di Yin and Xing Sun . . . . .	1047
<i>Bridging Cultures in the Kitchen: A Framework and Benchmark for Cross-Cultural Recipe Retrieval</i>	
Tianyi Hu, Maria Maistro and Daniel Hershcovich . . . . .	1068

<i>RULE: Reliable Multimodal RAG for Factuality in Medical Vision Language Models</i>	
Peng Xia, Kangyu Zhu, Haoran Li, Hongtu Zhu, Yun Li, Gang Li, Linjun Zhang and Huaxiu Yao	
1081	
<i>CryptoTrade: A Reflective LLM-based Agent to Guide Zero-shot Cryptocurrency Trading</i>	
Yuan Li, Bingqiao Luo, Qian Wang, Nuo Chen, Xu Liu and Bingsheng He	1094
<i>A Survey on In-context Learning</i>	
Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li and Zhifang Sui	1107
<i>DocHieNet: A Large and Diverse Dataset for Document Hierarchy Parsing</i>	
Hangdi Xing, Changxu Cheng, Feiyu Gao, Zirui Shao, Zhi Yu, Jiajun Bu, Qi Zheng and Cong Yao	1129
<i>AMR-Evol: Adaptive Modular Response Evolution Elicits Better Knowledge Distillation for Large Language Models in Code Generation</i>	
Ziyang Luo, Xin Li, Hongzhan Lin, Jing Ma and Lidong Bing	1143
<i>EFUF: Efficient Fine-Grained Unlearning Framework for Mitigating Hallucinations in Multimodal Large Language Models</i>	
Shangyu Xing, Fei Zhao, Zhen Wu, Tuo An, Weihao Chen, Chunhui Li, Jianbing Zhang and Xinyu Dai	1167
<i>Rethinking Pruning Large Language Models: Benefits and Pitfalls of Reconstruction Error Minimization</i>	
Sungbin Shin, Wonpyo Park, Jaeho Lee and Namhoon Lee	1182
<i>LLMs Are Zero-Shot Context-Aware Simultaneous Translators</i>	
Roman Koshkin, Katsuhito Sudoh and Satoshi Nakamura	1192
<i>AgentReview: Exploring Peer Review Dynamics with LLM Agents</i>	
Yiqiao Jin, Qinlin Zhao, Yiyang Wang, Hao Chen, Kaijie Zhu, Yijia Xiao and Jindong Wang	1208
<i>ChatRetriever: Adapting Large Language Models for Generalized and Robust Conversational Dense Retrieval</i>	
Kelong Mao, Chenlong Deng, Haonan Chen, Fengran Mo, Zheng Liu, Tetsuya Sakai and Zhicheng Dou	1227
<i>Fairer Preferences Elicit Improved Human-Aligned Large Language Model Judgments</i>	
Han Zhou, Xingchen Wan, Yinhong Liu, Nigel Collier, Ivan Vulić and Anna Korhonen	1241
<i>Learning Interpretable Legal Case Retrieval via Knowledge-Guided Case Reformulation</i>	
Chenlong Deng, Kelong Mao and Zhicheng Dou	1253
<i>Effective Demonstration Annotation for In-Context Learning via Language Model-Based Determinantal Point Process</i>	
Peng Wang, Xiaobin Wang, Chao Lou, Shengyu Mao, Pengjun Xie and Yong Jiang	1266
<i>Pre-trained Language Models Do Not Help Auto-regressive Text-to-Image Generation</i>	
Yuhui Zhang, Brandon McKinzie, Zhe Gan, Vaishaal Shankar and Alexander T Toshev	1281
<i>QUSELECT: Selective Decoding for Questions Under Discussion Parsing</i>	
Ashima Suvarna, Xiao Liu, Tanmay Parekh, Kai-Wei Chang and Nanyun Peng	1288
<i>Mitigating Language Bias of LMMs in Social Intelligence Understanding with Virtual Counterfactual Calibration</i>	
Peng Chen, Xiao-Yu Guo, Yuan-Fang Li, Xiaowang Zhang and Zhiyong Feng	1300

<i>Model Balancing Helps Low-data Training and Fine-tuning</i>	
Zihang Liu, Yuanzhe Hu, Tianyu Pang, Yefan Zhou, Pu Ren and Yaoqing Yang . . . . .	1311
<i>Reuse Your Rewards: Reward Model Transfer for Zero-Shot Cross-Lingual Alignment</i>	
Zhaofeng Wu, Ananth Balashankar, Yoon Kim, Jacob Eisenstein and Ahmad Beirami . . . . .	1332
<i>Large Language Models as Foundations for Next-Gen Dense Retrieval: A Comprehensive Empirical Assessment</i>	
Kun Luo, Minghao Qin, Zheng Liu, Shitao Xiao, Jun Zhao and Kang Liu . . . . .	1354
<i>A New Pipeline for Knowledge Graph Reasoning Enhanced by Large Language Models Without Fine-Tuning</i>	
Zhongwu Chen, Long Bai, Zixuan Li, Zhen Huang, Xiaolong Jin and Yong Dou . . . . .	1366
<i>Towards Tool Use Alignment of Large Language Models</i>	
Zhi-Yuan Chen, Shiqi Shen, Guangyao Shen, Gong Zhi, Xu Chen and Yankai Lin . . . . .	1382
<i>DecorateLM: Data Engineering through Corpus Rating, Tagging, and Editing with Language Models</i>	
Ranchi Zhao, Zhen Leng Thai, Yifan Zhang, Shengding Hu, Jie Zhou, Yunqi Ba, Jie Cai, Zhiyuan Liu and Maosong Sun . . . . .	1401
<i>Lookback Lens: Detecting and Mitigating Contextual Hallucinations in Large Language Models Using Only Attention Maps</i>	
Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim and James R. Glass	1419
<i>Controllable Preference Optimization: Toward Controllable Multi-Objective Alignment</i>	
Yiju Guo, Ganqu Cui, Lifan Yuan, Ning Ding, Zexu Sun, Bowen Sun, Huimin Chen, Ruobing Xie, Jie Zhou, Yankai Lin, Zhiyuan Liu and Maosong Sun . . . . .	1437
<i>Mitigating Matthew Effect: Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning for Conversational Recommendation</i>	
Yongsen Zheng, Ruilin Xu, Guohua Wang, Liang Lin and Kwok-Yan Lam . . . . .	1455
<i>Advancing Event Causality Identification via Heuristic Semantic Dependency Inquiry Network</i>	
Haoran Li, Qiang Gao, Hongmei Wu and Li Huang . . . . .	1467
<i>Exploring Union and Intersection of Visual Regions for Generating Questions, Answers, and Distractors</i>	
Wenjian Ding, Yao Zhang, Jun Wang, Adam Jatowt and Zhenglu Yang . . . . .	1479
<i>UniFashion: A Unified Vision-Language Model for Multimodal Fashion Retrieval and Generation</i>	
Xiangyu Zhao, Yuehan Zhang, Wenlong Zhang and Xiao-Ming Wu . . . . .	1490
<i>Tracking the perspectives of interacting language models</i>	
Hayden Helm, Brandon Duderstadt, Youngser Park and Carey Priebe . . . . .	1508
<i>MAR: Matching-Augmented Reasoning for Enhancing Visual-based Entity Question Answering</i>	
Zhengxuan Zhang, Yin WU, Yuyu Luo and Nan Tang . . . . .	1520
<i>Can Large Language Models Always Solve Easy Problems if They Can Solve Harder Ones?</i>	
Zhe Yang, Yichang Zhang, Tianyu Liu, Jian Yang, Junyang Lin, Chang Zhou and Zhifang Sui	1531
<i>Watch Every Step! LLM Agent Learning via Iterative Step-level Process Refinement</i>	
Weimin Xiong, Yifan Song, Xiutian Zhao, Wenhao Wu, Xun Wang, Ke Wang, Cheng LI, Wei Peng and Sujian Li . . . . .	1556



<i>Standardize: Aligning Language Models with Expert-Defined Standards for Content Generation</i> Joseph Marvin Imperial, Gail Forey and Harish Tayyar Madabushi .....	1573
<i>Cross-domain NER with Generated Task-Oriented Knowledge: An Empirical Study from Information Density Perspective</i> Zhihao Zhang, Sophia Yat Mei Lee, Junshuang Wu, Dong Zhang, Shoushan Li, Erik Cambria and Guodong Zhou .....	1595
<i>Glue pizza and eat rocks - Exploiting Vulnerabilities in Retrieval-Augmented Generative Models</i> Zhen Tan, Chengshuai Zhao, Raha Moraffah, Yifan Li, Song Wang, Jundong Li, Tianlong Chen and Huan Liu .....	1610
<i>Predicate Debiasing in Vision-Language Models Integration for Scene Graph Generation Enhancement</i> Yuxuan Wang and Xiaoyuan Liu .....	1627
<i>SHIELD: Evaluation and Defense Strategies for Copyright Compliance in LLM Text Generation</i> Xiaoze Liu, Ting Sun, Tianyang Xu, Feijie Wu, Cunxiang Wang, Xiaoqian Wang and Jing Gao	1640
<i>MatchTime: Towards Automatic Soccer Game Commentary Generation</i> Jiayuan Rao, Haoning Wu, Chang Liu, Yanfeng Wang and Weidi Xie .....	1671
<i>Rethinking Token Reduction for State Space Models</i> Zheng Zhan, Yushu Wu, Zhenglun Kong, Changdi Yang, Yifan Gong, Xuan Shen, Xue Lin, Pu Zhao and Yanzhi Wang .....	1686
<i>Triad: A Framework Leveraging a Multi-Role LLM-based Agent to Solve Knowledge Base Question Answering</i> Chang Zong, Yuchen Yan, Weiming Lu, Jian Shao, Yongfeng Huang, Heng Chang and Yueting Zhuang .....	1698
<i>MetaGPT: Merging Large Language Models Using Model Exclusive Task Arithmetic</i> Yuyan Zhou, Liang Song, Bingning Wang and Weipeng Chen .....	1711
<i>Event Causality Identification with Synthetic Control</i> Haoyu Wang, Fengze Liu, Jiayao Zhang, Dan Roth and Kyle Richardson .....	1725
<i>Retrieved Sequence Augmentation for Protein Representation Learning</i> Chang Ma, Haiteng Zhao, Lin Zheng, Jiayi Xin, Qintong Li, Lijun Wu, Zhihong Deng, Yang Young Lu, Qi Liu, Sheng Wang and Lingpeng Kong .....	1738
<i>HELPD: Mitigating Hallucination of LLMs by Hierarchical Feedback Learning with Vision-enhanced Penalty Decoding</i> Fan Yuan, Chi Qin, Xiaogang Xu and Piji Li .....	1768
<i>TopViewRS: Vision-Language Models as Top-View Spatial Reasoners</i> Chengzu Li, Caiqi Zhang, Han Zhou, Nigel Collier, Anna Korhonen and Ivan Vulić .....	1786
<i>DA<sup>3</sup>: A Distribution-Aware Adversarial Attack against Language Models</i> Yibo Wang, Xiangjue Dong, James Caverlee and Philip S. Yu .....	1808
<i>Evaluating Psychological Safety of Large Language Models</i> Xingxuan Li, Yutong Li, Lin Qiu, Shafiq Joty and Lidong Bing .....	1826
<i>An Effective Deployment of Diffusion LM for Data Augmentation in Low-Resource Sentiment Classification</i> Zhuowei Chen, Lianxi Wang, Yuben Wu, Xinfeng Liao, Yujia Tian and Junyang Zhong ....	1844

<i>Self-Bootstrapped Visual-Language Model for Knowledge Selection and Question Answering</i> Dongze Hao, Qunbo Wang, Longteng Guo, Jie Jiang and Jing Liu .....	1857
<i>PsFuture: A Pseudo-Future-based Zero-Shot Adaptive Policy for Simultaneous Machine Translation</i> Libo Zhao, Jing Li and Ziqian Zeng .....	1869
<i>TinyChart: Efficient Chart Understanding with Program-of-Thoughts Learning and Visual Token Merging</i> Liang Zhang, Anwen Hu, Haiyang Xu, Ming Yan, Yichen Xu, Qin Jin, Ji Zhang and Fei Huang	1882
<i>Do We Need Language-Specific Fact-Checking Models? The Case of Chinese</i> Caiqi Zhang, Zhijiang Guo and Andreas Vlachos .....	1899
<i>Enhancing Advanced Visual Reasoning Ability of Large Language Models</i> Zhiyuan Li, Dongnan Liu, Chaoyi Zhang, Heng Wang, Tengfei Xue and Weidong Cai .....	1915
<i>CMD: a framework for Context-aware Model self-Detoxification</i> Zecheng Tang, Keyan Zhou, Juntao Li, Yuyang Ding, Pinzheng Wang, Yan Bowen, Renjie Hua and Min Zhang .....	1930
<i>Embedding and Gradient Say Wrong: A White-Box Method for Hallucination Detection</i> Xiaomeng Hu, Yiming Zhang, Ru Peng, Haozhe Zhang, Chenwei Wu, Gang Chen and Junbo Zhao .....	1950
<i>TCSinger: Zero-Shot Singing Voice Synthesis with Style Transfer and Multi-Level Style Control</i> Yu Zhang, Ziyue Jiang, Ruiqi Li, Changhao Pan, Jinzheng He, Rongjie Huang, Chuxin Wang and Zhou Zhao .....	1960
<i>Be Helpful but Don't Talk too Much - Enhancing Helpfulness in Conversations through Relevance in Multi-Turn Emotional Support</i> Junlin Li, Bo Peng, Yu-Yin Hsu and Chu-Ren Huang .....	1976
<i>Aligning Language Models to Explicitly Handle Ambiguity</i> Hyuhng Joon Kim, Youna Kim, Cheonbok Park, Junyeob Kim, Choonghyun Park, Kang Min Yoo, Sang-goo Lee and Taeuk Kim .....	1989
<i>Tag-grounded Visual Instruction Tuning with Retrieval Augmentation</i> Daiqing Qi, Handong Zhao, Zijun Wei and Sheng Li .....	2008
<i>GLaPE: Gold Label-agnostic Prompt Evaluation for Large Language Models</i> Xuanchang Zhang, Zhuosheng Zhang and Hai Zhao .....	2027
<i>Decoding the Echoes of Vision from fMRI: Memory Disentangling for Past Semantic Information</i> Runze Xia, Congchi Yin and Piji Li .....	2040
<i>Optimizing Code Retrieval: High-Quality and Scalable Dataset Annotation through Large Language Models</i> Rui Li, Qi Liu, Liyang He, Zheng Zhang, Hao Zhang, Shengyu Ye, Junyu Lu and Zhenya Huang	2053
<i>Towards Difficulty-Agnostic Efficient Transfer Learning for Vision-Language Models</i> Yongjin Yang, Jongwoo Ko and Se-Young Yun .....	2066
<i>Advancing Process Verification for Large Language Models via Tree-Based Preference Learning</i> Mingqian He, Yongliang Shen, Wenqi Zhang, Zeqi Tan and Weiming Lu .....	2086

<i>An Inversion Attack Against Obfuscated Embedding Matrix in Language Model Inference</i>	
Yu Lin, Qizhi Zhang, Quanwei Cai, Jue Hong, Wu Ye, Huiqi Liu and Bing Duan . . . . .	2100
<i>VideoScore: Building Automatic Metrics to Simulate Fine-grained Human Feedback for Video Generation</i>	
Xuan He, Dongfu Jiang, Ge Zhang, Max Ku, Achint Soni, Sherman Siu, Haonan Chen, Abhranil Chandra, Ziyang Jiang, Aaran Arulraj, Kai Wang, Quy Duc Do, Yuansheng Ni, Bohan Lyu, Yaswanth Narsupalli, Rongqi Fan, Zhiheng Lyu, Bill Yuchen Lin and Wenhui Chen . . . . .	2105
<i>LogicAsker: Evaluating and Improving the Logical Reasoning Ability of Large Language Models</i>	
Yuxuan Wan, Wenxuan Wang, Yiliu Yang, Youliang Yuan, Jen-tse Huang, Pinjia He, Wenxiang Jiao and Michael Lyu . . . . .	2124
<i>Integrating Structural Semantic Knowledge for Enhanced Information Extraction Pre-training</i>	
Xiaoyang Yi, Yuru Bao, Jian Zhang, Yifang Qin and Faxin Lin . . . . .	2156
<i>FuseGen: PLM Fusion for Data-generation based Zero-shot Learning</i>	
Tianyuan Zou, Yang Liu, Peng Li, Jianqing Zhang, Jingjing Liu and Ya-Qin Zhang . . . . .	2172
<i>I Need Help! Evaluating LLM’s Ability to Ask for Users’ Support: A Case Study on Text-to-SQL Generation</i>	
Cheng-Kuang Wu, Zhi Rui Tam, Chao-Chung Wu, Chieh-Yen Lin, Hung-yi Lee and Yun-Nung Chen . . . . .	2191
<i>Oddballs and Misfits: Detecting Implicit Abuse in Which Identity Groups are Depicted as Deviating from the Norm</i>	
Michael Wiegand and Josef Ruppenhofer . . . . .	2200
<i>By My Eyes: Grounding Multimodal Large Language Models with Sensor Data via Visual Prompting</i>	
Hyungjun Yoon, Biniyam Aschalew Tolera, Taesik Gong, Kimin Lee and Sung-Ju Lee . . . . .	2219
<i>Prefixing Attention Sinks can Mitigate Activation Outliers for Large Language Model Quantization</i>	
Seungwoo Son, Wonpyo Park, Woohyun Han, Kyuyeun Kim and Jaeho Lee . . . . .	2242
<i>CHIQ: Contextual History Enhancement for Improving Query Rewriting in Conversational Search</i>	
Fengran Mo, Abbas Ghaddar, Kelong Mao, Mehdi Rezagholizadeh, Boxing Chen, Qun Liu and Jian-Yun Nie . . . . .	2253
<i>Towards Low-Resource Harmful Meme Detection with LMM Agents</i>	
Jianzhao Huang, Hongzhan Lin, Liu Ziyang, Ziyang Luo, Guang Chen and Jing Ma . . . . .	2269
<i>VIVA: A Benchmark for Vision-Grounded Decision-Making with Human Values</i>	
Zhe Hu, Yixiao Ren, Jing Li and Yu Yin . . . . .	2294
<i>Direct Multi-Turn Preference Optimization for Language Agents</i>	
Wentao Shi, Mengqi Yuan, Junkang Wu, Qifan Wang and Fuli Feng . . . . .	2312
<i>Self-Refine Instruction-Tuning for Aligning Reasoning in Language Models</i>	
Leonardo Ranaldi and Andre Freitas . . . . .	2325
<i>In Search of the Long-Tail: Systematic Generation of Long-Tail Inferential Knowledge via Logical Rule Guided Search</i>	
Huihan Li, Yuting Ning, Zeyi Liao, Siyuan Wang, Xiang Lorraine Li, Ximing Lu, Wenting Zhao, Faeze Brahma, Yejin Choi and Xiang Ren . . . . .	2348
<i>AutoScraper: A Progressive Understanding Web Agent for Web Scraper Generation</i>	
Wenhao Huang, Zhouhong Gu, Chenghao Peng, Jiaqing Liang, Zhixu Li, Yanghua Xiao, Liqian Wen and Zulong Chen . . . . .	2371

<i>Backward Lens: Projecting Language Model Gradients into the Vocabulary Space</i> Shahar Katz, Yonatan Belinkov, Mor Geva and Lior Wolf . . . . .	2390
<i>Selective Vision is the Challenge for Visual Reasoning: A Benchmark for Visual Argument Understanding</i> Jiwan Chung, Sungjae Lee, Minseo Kim, Seungju Han, Ashkan Yousefpour, Jack Hessel and Youngjae Yu . . . . .	2423
<i>Can visual language models resolve textual ambiguity with visual cues? Let visual puns tell you!</i> Jiwan Chung, Seungwon Lim, Jaehyun Jeon, Seungbeen Lee and Youngjae Yu . . . . .	2452
<i>Reusing Transferable Weight Increments for Low-resource Style Generation</i> Chunzhen Jin, Eliot Huang, Heng Chang, Yaqi Wang, Peng Cao and Osmar Zaiane . . . . .	2470
<i>Large Language Model as an Assignment Evaluator: Insights, Feedback, and Challenges in a 1000+ Student Course</i> Cheng-Han Chiang, Wei-Chih Chen, Chun-Yi Kuan, Chienchou Yang and Hung-yi Lee . . . . .	2489
<i>Seemingly Plausible Distractors in Multi-Hop Reasoning: Are Large Language Models Attentive Readers?</i> Neeladri Bhuiya, Viktor Schlegel and Stefan Winkler . . . . .	2514
<i>Instruction Pre-Training: Language Models are Supervised Multitask Learners</i> Daixuan Cheng, Yuxian Gu, Shaohan Huang, Junyu Bi, Minlie Huang and Furu Wei . . . . .	2529
<i>LEMoe: Advanced Mixture of Experts Adaptor for Lifelong Model Editing of Large Language Models</i> Renzhi Wang and Piji Li . . . . .	2551
<i>Collaborative Performance Prediction for Large Language Models</i> Qiyuan Zhang, Fuyuan Lyu, Xue Liu and Chen Ma . . . . .	2576
<i>Surveying the Dead Minds: Historical-Psychological Text Analysis with Contextualized Construct Representation (CCR) for Classical Chinese</i> Yuqi Chen, Sixuan Li, Ying Li and Mohammad Atari . . . . .	2597
<i>Knowledge Verification to Nip Hallucination in the Bud</i> Fanqi Wan, Xinting Huang, Leyang Cui, Xiaojun Quan, Wei Bi and Shuming Shi . . . . .	2616
<i>QUITE: Quantifying Uncertainty in Natural Language Text in Bayesian Reasoning Scenarios</i> Timo Pierre Schrader, Lukas Lange, Simon Razniewski and Annemarie Friedrich . . . . .	2634
<i>African or European Swallow? Benchmarking Large Vision-Language Models for Fine-Grained Object Classification</i> Gregor Geigle, Radu Timofte and Goran Glavaš . . . . .	2653
<i>Whispers that Shake Foundations: Analyzing and Mitigating False Premise Hallucinations in Large Language Models</i> Hongbang Yuan, Pengfei Cao, Zhuoran Jin, Yubo Chen, Daojian Zeng, Kang Liu and Jun Zhao . . . . .	2670
<i>To Word Senses and Beyond: Inducing Concepts with Contextualized Language Models</i> Bastien Liétard, Pascal Denis and Mikaela Keller . . . . .	2684
<i>ASETf: A Novel Method for Jailbreak Attack on LLMs through Translate Suffix Embeddings</i> Hao Wang, Hao Li, Minlie Huang and Lei Sha . . . . .	2697
<i>An Electoral Approach to Diversify LLM-based Multi-Agent Collective Decision-Making</i> Xiutian Zhao, Ke Wang and Wei Peng . . . . .	2712

<i>Does Object Grounding Really Reduce Hallucination of Large Vision-Language Models?</i> Gregor Geigle, Radu Timofte and Goran Glavaš .....	2728
<i>Take Off the Training Wheels! Progressive In-Context Learning for Effective Alignment</i> Zhenyu Liu, Dongfang Li, Xinshuo Hu, Xinping Zhao, Yibin Chen, Baotian Hu and Min Zhang	2743
<i>MoDULA: Mixture of Domain-Specific and Universal LoRA for Multi-Task Learning</i> Yufei Ma, Zihan Liang, Huangyu Dai, Ben Chen, Dehong Gao, Zhuoran Ran, Wang Zihan, Linbo Jin, Wen Jiang, Guannan Zhang, Xiaoyan Cai and Libin Yang .....	2758
<i>Message Passing on Semantic-Anchor-Graphs for Fine-grained Emotion Representation Learning and Classification</i> Pinyi Zhang, Jingyang Chen, Junchen Shen, Zijie Zhai, Ping Li, Jie Zhang and Kai Zhang ..	2771
<i>PhiloGPT: A Philology-Oriented Large Language Model for Ancient Chinese Manuscripts with Dunhuang as Case Study</i> Yuqing Zhang, Baoyi He, Yihan Chen, Hangqi Li, Han Yue, Shengyu Zhang, Huaiyong Dou, Junchi Yan, Zemin Liu, Yongquan Zhang and Fei Wu .....	2784
<i>Alignment-Enhanced Decoding: Defending Jailbreaks via Token-Level Adaptive Refining of Probability Distributions</i> Quan Liu, Zhenhong Zhou, Longzhu He, Yi Liu, Wei Zhang and Sen Su .....	2802
<i>MiniConGTS: A Near Ultimate Minimalist Contrastive Grid Tagging Scheme for Aspect Sentiment Triplet Extraction</i> Qiao Sun, Liujia Yang, Minghao Ma, Nanyang Ye and Qinying Gu .....	2817
<i>Evaluating Large Language Models via Linguistic Profiling</i> Alessio Miaschi, Felice Dell’Orletta and Giulia Venturi .....	2835
<i>With Ears to See and Eyes to Hear: Sound Symbolism Experiments with Multimodal Large Language Models</i> Tyler Loakman, Yucheng LI and Chenghua Lin .....	2849
<i>KB-Plugin: A Plug-and-play Framework for Large Language Models to Induce Programs over Low-resourced Knowledge Bases</i> Jiajie Zhang, Shulin Cao, Linmei Hu, Ling Feng, Lei Hou and Juanzi Li .....	2868
<i>Understanding Higher-Order Correlations Among Semantic Components in Embeddings</i> Momose Oyama, Hiroaki Yamagiwa and Hidetoshi Shimodaira .....	2883
<i>DGLF: A Dual Graph-based Learning Framework for Multi-modal Sarcasm Detection</i> Zhihong Zhu, Kefan Shen, Zhaorun Chen, Yunyan Zhang, Yuyan Chen, Xiaoqi Jiao, Zhongwei Wan, Shaorong Xie, Wei Liu, Xian Wu and Yefeng Zheng .....	2900
<i>Evaluating D-MERIT of Partial-annotation on Information Retrieval</i> Royi Rassin, Yaron Fairstein, Oren Kalinsky, Guy Kushilevitz, Nachshon Cohen, Alexander Libov and Yoav Goldberg .....	2913
<i>Verification and Refinement of Natural Language Explanations through LLM-Symbolic Theorem Proving</i> Xin Quan, Marco Valentino, Louise A. Dennis and Andre Freitas .....	2933
<i>Calibrating the Confidence of Large Language Models by Eliciting Fidelity</i> Mozhi Zhang, Mianqiu Huang, Rundong Shi, Linsen Guo, Chong Peng, Peng Yan, Yaqian Zhou and Xipeng Qiu .....	2959

<i>The Accuracy Paradox in RLHF: When Better Reward Models Don't Yield Better Language Models</i> Yanjun Chen, Dawei Zhu, Yirong Sun, Xinghao Chen, Wei Zhang and Xiaoyu Shen . . . . .	2980
<i>How Hard is this Test Set? NLI Characterization by Exploiting Training Dynamics</i> Adrian Cosma, Stefan Ruseti, Mihai Dascalu and Cornelia Caragea . . . . .	2990
<i>Zero-shot Cross-Lingual Transfer for Synthetic Data Generation in Grammatical Error Detection</i> Gaetan Lopez Latouche, Marc-André Carbonneau and Benjamin Swanson . . . . .	3002
<i>CUTE: Measuring LLMs' Understanding of Their Tokens</i> Lukas Edman, Helmut Schmid and Alexander Fraser . . . . .	3017
<i>SEER: Self-Aligned Evidence Extraction for Retrieval-Augmented Generation</i> Xinping Zhao, Dongfang Li, Yan Zhong, Boren Hu, Yibin Chen, Baotian Hu and Min Zhang	3027
<i>On the Role of Context in Reading Time Prediction</i> Andreas Opedal, Eleanor Chodroff, Ryan Cotterell and Ethan Wilcox . . . . .	3042
<i>BC-Prover: Backward Chaining Prover for Formal Theorem Proving</i> Yuhang He, Jihai Zhang, Jianzhu Bao, Fangquan Lin, Cheng Yang, Bing Qin, Ruifeng Xu and Wotao Yin . . . . .	3059
<i>From Insights to Actions: The Impact of Interpretability and Analysis Research on NLP</i> Marius Mosbach, Vagrant Gautam, Tomás Vergara Browne, Dietrich Klakow and Mor Geva	3078
<i>Autoregressive Pre-Training on Pixels and Texts</i> Yekun Chai, Qingyi Liu, Jingwu Xiao, Shuohuan Wang, Yu Sun and Hua Wu . . . . .	3106
<i>On Training Data Influence of GPT Models</i> Yekun Chai, Qingyi Liu, Shuohuan Wang, Yu Sun, Qiwei Peng and Hua Wu . . . . .	3126
<i>Understanding Democratization in NLP and ML Research</i> Arjun Subramonian, Vagrant Gautam, Dietrich Klakow and Zeerak Talat . . . . .	3151
<i>DocKD: Knowledge Distillation from LLMs for Open-World Document Understanding Models</i> Sungnyun Kim, Haofu Liao, Srikar Appalaraju, Peng Tang, Zhuowen Tu, Ravi Kumar Satzoda, R. Manmatha, Vijay Mahadevan and Stefano Soatto . . . . .	3167
<i>Cross-lingual Transfer for Automatic Question Generation by Learning Interrogative Structures in Target Languages</i> Seonjeong Hwang, Yunsu Kim and Gary Lee . . . . .	3194
<i>ScalingFilter: Assessing Data Quality through Inverse Utilization of Scaling Laws</i> Ruihang Li, Yixuan Wei, Miaosen Zhang, Nenghai Yu, Han Hu and Houwen Peng . . . . .	3209
<i>Word Alignment as Preference for Machine Translation</i> Qiyu Wu, Masaaki Nagata, Zhongtao Miao and Yoshimasa Tsuruoka . . . . .	3223
<i>Improving Multi-party Dialogue Generation via Topic and Rhetorical Coherence</i> Yaxin Fan, Peifeng LI and Qiaoming Zhu . . . . .	3240
<i>SEEKR: Selective Attention-Guided Knowledge Retention for Continual Learning of Large Language Models</i> Jinghan He, Haiyun Guo, Kuan Zhu, Zihan Zhao, Ming Tang and Jinqiao Wang . . . . .	3254
<i>Neuron-Level Knowledge Attribution in Large Language Models</i> Zeping Yu and Sophia Ananiadou . . . . .	3267

<i>How do Large Language Models Learn In-Context? Query and Key Matrices of In-Context Heads are Two Towers for Metric Learning</i>	
Zeping Yu and Sophia Ananiadou .....	3281
<i>Interpreting Arithmetic Mechanism in Large Language Models through Comparative Neuron Analysis</i>	
Zeping Yu and Sophia Ananiadou .....	3293
<i>Pixology: Probing the Linguistic and Visual Capabilities of Pixel-based Language Models</i>	
Kushal Tatariya, Vladimir Araujo, Thomas Bauwens and Miryam De Lhoneux .....	3307
<i>GoldCoin: Grounding Large Language Models in Privacy Laws via Contextual Integrity Theory</i>	
Wei Fan, Haoran Li, Zheyang Deng, Weiqi Wang and Yangqiu Song .....	3321
<i>Noise, Novels, Numbers. A Framework for Detecting and Categorizing Noise in Danish and Norwegian Literature</i>	
Ali Al-Laith, Daniel Hershcovich, Jens Bjerring-Hansen, Jakob Ingemann Parby, Alexander Conroy and Timothy R Tangherlini .....	3344
<i>QUIK: Towards End-to-end 4-Bit Inference on Generative Large Language Models</i>	
Saleh Ashkboos, Iliia Markov, Elias Frantar, Tingxuan Zhong, Xincheng Wang, Jie Ren, Torsten Hoefler and Dan Alistarh .....	3355
<i>Fine-Grained Prediction of Reading Comprehension from Eye Movements</i>	
Omer Shubi, Yoav Meiri, Cfir Avraham Hadar and Yevgeni Berzak .....	3372
<i>EfficientRAG: Efficient Retriever for Multi-Hop Question Answering</i>	
Ziyuan Zhuang, Zhiyang Zhang, Sitao Cheng, Fangkai Yang, Jia Liu, Shujian Huang, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang and Qi Zhang .....	3392
<i>Unsupervised Human Preference Learning</i>	
Sumuk Shashidhar, Abhinav Chinta, Vaibhav Sahai and Dilek Hakkani Tur .....	3412
<i>Is Safer Better? The Impact of Guardrails on the Argumentative Strength of LLMs in Hate Speech Countering</i>	
Helena Bonaldi, Greta Damo, Nicolás Benjamín Ocampo, Elena Cabrio, Serena Villata and Marco Guerini .....	3446
<i>Leading Whitespaces of Language Models' Subword Vocabulary Pose a Confound for Calculating Word Probabilities</i>	
Byung-Doh Oh and William Schuler .....	3464
<i>LLM4Decompile: Decompiling Binary Code with Large Language Models</i>	
Hanzhuo Tan, Qi Luo, Jing Li and Yuqun Zhang .....	3473
<i>From Bottom to Top: Extending the Potential of Parameter Efficient Fine-Tuning</i>	
Jihao Gu, Zelin Wang, Yibo Zhang, Ziji Zhang and Ping Gong .....	3488
<i>CoTKR: Chain-of-Thought Enhanced Knowledge Rewriting for Complex Knowledge Graph Question Answering</i>	
Yike Wu, Yi Huang, Nan Hu, Yuncheng Hua, Guilin Qi, Jiaoyan Chen and Jeff Z. Pan .....	3501
<i>MTLS: Making Texts into Linguistic Symbols</i>	
Wenlong Fei, Xiaohua Wang, Min Hu, Qingyu Zhang and Hongbo Li .....	3521
<i>D2R: Dual-Branch Dynamic Routing Network for Multimodal Sentiment Detection</i>	
Yifan Chen, Kuntao Li, Weixing Mai, Qiaofeng Wu, Yun Xue and Fenghuan Li .....	3536

<i>A Generic Method for Fine-grained Category Discovery in Natural Language Texts</i>	
Chang Tian, Matthew B. Blaschko, Wenpeng Yin, Mingzhe Xing, Yinliang Yue and Marie-Francine Moens . . . . .	3548
<i>Toxicity Detection is NOT all you Need: Measuring the Gaps to Supporting Volunteer Content Moderators through a User-Centric Method</i>	
Yang Trista Cao, Lovely-Frances Domingo, Sarah Gilbert, Michelle L. Mazurek, Katie Shilton and Hal Daumé Iii . . . . .	3567
<i>A User-Centric Multi-Intent Benchmark for Evaluating Large Language Models</i>	
Jiayin Wang, Fengran Mo, Weizhi Ma, Peijie Sun, Min Zhang and Jian-Yun Nie . . . . .	3588
<i>Decompose and Compare Consistency: Measuring VLMs’ Answer Reliability via Task-Decomposition Consistency Comparison</i>	
Qian Yang, Weixiang Yan and Aishwarya Agrawal . . . . .	3613
<i>Learn to Refuse: Making Large Language Models More Controllable and Reliable through Knowledge Scope Limitation and Refusal Mechanism</i>	
Lang Cao . . . . .	3628
<i>VGBench: A Comprehensive Benchmark of Vector Graphics Understanding and Generation for Large Language Models</i>	
Bocheng Zou, Mu Cai, Jianrui Zhang and Yong Jae Lee . . . . .	3647
<i>What do Large Language Models Need for Machine Translation Evaluation?</i>	
Shenbin Qian, Archchana Sindhujan, Minnie Kabra, Diptesh Kanojia, Constantin Orasan, Tharindu Ranasinghe and Fred Blain . . . . .	3660
<i>Performance-Guided LLM Knowledge Distillation for Efficient Text Classification at Scale</i>	
Flavio Di Palo, Prateek Singhi and Bilal H Fadlallah . . . . .	3675
<i>External Knowledge-Driven Argument Mining: Leveraging Attention-Enhanced Multi-Network Models</i>	
Debelá Gemechu and Chris Reed . . . . .	3688
<i>C3PA: An Open Dataset of Expert-Annotated and Regulation-Aware Privacy Policies to Enable Scalable Regulatory Compliance Audits</i>	
Maaz Bin Musa, Steven M. Winston, Garrison Allen, Jacob Schiller, Kevin Moore, Sean Quick, Johnathan Melvin, Padmini Srinivasan, Mihailis E. Diamantis and Rishab Nithyanand . . . . .	3710
<i>M<sup>2</sup>PT: Multimodal Prompt Tuning for Zero-shot Instruction Learning</i>	
Taowen Wang, Yiyang Liu, James Chenhao Liang, Junhan Zhao, Yiming Cui, Yuning Mao, Shao-liang Nie, Jiahao Liu, Fuli Feng, Zenglin Xu, Cheng Han, Lifu Huang, Qifan Wang and Dongfang Liu . . . . .	3723
<i>Text Grafting: Near-Distribution Weak Supervision for Minority Classes in Text Classification</i>	
Letian Peng, Yi Gu, Chengyu Dong, Zihan Wang and Jingbo Shang . . . . .	3741
<i>Incubating Text Classifiers Following User Instruction with Nothing but LLM</i>	
Letian Peng, Zilong Wang and Jingbo Shang . . . . .	3753
<i>PTD-SQL: Partitioning and Targeted Drilling with LLMs in Text-to-SQL</i>	
Ruilin Luo, Liyuan Wang, Binghuai Lin, Zicheng Lin and Yujiu Yang . . . . .	3767
<i>Conditional and Modal Reasoning in Large Language Models</i>	
Wesley H. Holliday, Matthew Mandelkern and Cedegao E. Zhang . . . . .	3800



<i>Advancing Large Language Model Attribution through Self-Improving</i>	
Lei Huang, Xiaocheng Feng, Weitao Ma, Liang Zhao, Yuchun Fan, Weihong Zhong, Dongliang Xu, Qing Yang, Hongtao Liu and Bing Qin . . . . .	3822
<i>AlignCap: Aligning Speech Emotion Captioning to Human Preferences</i>	
Ziqi Liang, Haoxiang Shi and Hanhui Chen . . . . .	3837
<i>Interpretability-based Tailored Knowledge Editing in Transformers</i>	
Yihuai Hong and Aldo Lipani . . . . .	3847
<i>Prompt Optimization in Multi-Step Tasks (PROMST): Integrating Human Feedback and Heuristic-based Sampling</i>	
Yongchao Chen, Jacob Arkin, Yilun Hao, Yang Zhang, Nicholas Roy and Chuchu Fan . . . . .	3859
<i>Empowering Large Language Model for Continual Video Question Answering with Collaborative Prompting</i>	
Chen Cai, Zheng Wang, Jianjun Gao, Wenyang Liu, Ye Lu, Runzhong Zhang and Kim-Hui Yap	3921
<i>Dissecting Fine-Tuning Unlearning in Large Language Models</i>	
Yihuai Hong, Yuelin Zou, Lijie Hu, Ziqian Zeng, Di Wang and Haiqin Yang . . . . .	3933
<i>Dancing in Chains: Reconciling Instruction Following and Faithfulness in Language Models</i>	
Zhengxuan Wu, Yuhao Zhang, Peng Qi, Yumo Xu, Rujun Han, Yian Zhang, Jifan Chen, Bonan Min and Zhiheng Huang . . . . .	3942
<i>Where is the signal in tokenization space?</i>	
Renato Geh, Honghua Zhang, Kareem Ahmed, Benjie Wang and Guy Van Den Broeck . . . . .	3966
<i>Private Language Models via Truncated Laplacian Mechanism</i>	
Tianhao Huang, Tao Yang, Ivan Habernal, Lijie Hu and Di Wang . . . . .	3980
<i>Estimating Knowledge in Large Language Models Without Generating a Single Token</i>	
Daniela Gottesman and Mor Geva . . . . .	3994
<i>Consistent Autoformalization for Constructing Mathematical Libraries</i>	
Lan Zhang, Xin Quan and Andre Freitas . . . . .	4020
<i>When Context Leads but Parametric Memory Follows in Large Language Models</i>	
Yufei Tao, Adam Hiatt, Erik Haake, Antonie J. Jetter and Ameeta Agrawal . . . . .	4034
<i>Semantic Training Signals Promote Hierarchical Syntactic Generalization in Transformers</i>	
Aditya Yedetore and Najoung Kim . . . . .	4059
<i>When Is Multilinguality a Curse? Language Modeling for 250 High- and Low-Resource Languages</i>	
Tyler A. Chang, Catherine Arnett, Zhuowen Tu and Ben Bergen . . . . .	4074
<i>Teaching Embodied Reinforcement Learning Agents: Informativeness and Diversity of Language Use</i>	
Jiajun Xi, Yinong He, Jianing Yang, Yinpei Dai and Joyce Chai . . . . .	4097
<i>MiTTeNS: A Dataset for Evaluating Gender Mistranslation</i>	
Kevin Robinson, Sneha Kudugunta, Romina Stella, Sunipa Dev and Jasmijn Bastings . . . . .	4115
<i>Teaching LLMs to Abstain across Languages via Multilingual Feedback</i>	
Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Orevaoghene Ahia, Shuyue Stella Li, Vidhisha Balachandran, Sunayana Sitaram and Yulia Tsvetkov . . . . .	4125

<i>Modular Pluralism: Pluralistic Alignment via Multi-LLM Collaboration</i>	
Shangbin Feng, Taylor Sorensen, Yuhan Liu, Jillian Fisher, Chan Young Park, Yejin Choi and Yulia Tsvetkov . . . . .	4151
<i>StyleRemix: Interpretable Authorship Obfuscation via Distillation and Perturbation of Style Elements</i>	
Jillian Fisher, Skyler Hallinan, Ximing Lu, Mitchell L Gordon, Zaid Harchaoui and Yejin Choi	4172
<i>I Could've Asked That: Reformulating Unanswerable Questions</i>	
Wenting Zhao, Ge Gao, Claire Cardie and Alexander M Rush . . . . .	4207
<i>STOP! Benchmarking Large Language Models with Sensitivity Testing on Offensive Progressions</i>	
Robert Morabito, Sangmitra Madhusudan, Tyler McDonald and Ali Emami . . . . .	4221
<i>Hidden Persuaders: LLMs' Political Leaning and Their Influence on Voters</i>	
Yujin Potter, Shiyang Lai, Junsol Kim, James Evans and Dawn Song . . . . .	4244
<i>SOUL: Unlocking the Power of Second-Order Optimization for LLM Unlearning</i>	
Jinghan Jia, Yihua Zhang, Yimeng Zhang, Jiancheng Liu, Bharat Runwal, James Diffenderfer, Bhavya Kailkhura and Sijia Liu . . . . .	4276
<i>When Reasoning Meets Information Aggregation: A Case Study with Sports Narratives</i>	
Yebowen Hu, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Wenlin Yao, Hassan Foroosh, Dong Yu and Fei Liu . . . . .	4293
<i>An Analysis of Multilingual FACTScore</i>	
Vu Trong Kim, Michael Krumdick, Varshini Reddy, Franck Dernoncourt and Viet Dac Lai .	4309
<i>Prometheus 2: An Open Source Language Model Specialized in Evaluating Other Language Models</i>	
Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee and Minjoon Seo . . . . .	4334
<i>RAG-QA Arena: Evaluating Domain Robustness for Long-form Retrieval Augmented Question Answering</i>	
Rujun Han, Yuhao Zhang, Peng Qi, Yumo Xu, Jenyuan Wang, Lan Liu, William Yang Wang, Bonan Min and Vittorio Castelli . . . . .	4354
<i>PromptReps: Prompting Large Language Models to Generate Dense and Sparse Representations for Zero-Shot Document Retrieval</i>	
Shengyao Zhuang, Xueguang Ma, Bevan Koopman, Jimmy Lin and Guido Zuccon . . . . .	4375
<i>Voices Unheard: NLP Resources and Models for Yorùbá Regional Dialects</i>	
Orevaoghene Ahia, Anuoluwapo Aremu, Diana Abagyan, Hila Gonen, David Ifeoluwa Adelani, Daud Abolade, Noah A. Smith and Yulia Tsvetkov . . . . .	4392
<i>ARES: Alternating Reinforcement Learning and Supervised Fine-Tuning for Enhanced Multi-Modal Chain-of-Thought Reasoning Through Diverse AI Feedback</i>	
Ju-Seung Byun, Jiyun Chun, Jihyung Kil and Andrew Perrault . . . . .	4410
<i>Order of Magnitude Speedups for LLM Membership Inference</i>	
Rongting Zhang, Martin Andres Bertran and Aaron Roth . . . . .	4431
<i>VIMI: Grounding Video Generation through Multi-modal Instruction</i>	
Yuwei Fang, Willi Menapace, Aliaksandr Siarohin, Tsai-Shien Chen, Kuan-Chieh Wang, Ivan Skorokhodov, Graham Neubig and Sergey Tulyakov . . . . .	4444

<i>F<sup>2</sup>RL: Factuality and Faithfulness Reinforcement Learning Framework for Claim-Guided Evidence-Supported Counterspeech Generation</i>	
Haiyang Wang, Yuchen Pan, Xin Song, Xuechen Zhao, Minghao Hu and Bin Zhou . . . . .	4457
<i>Deciphering Rumors: A Multi-Task Learning Approach with Intent-aware Hierarchical Contrastive Learning</i>	
Chang Yang, Peng Zhang, Hui Gao and Jing Zhang . . . . .	4471
<i>Visual Prompting in LLMs for Enhancing Emotion Recognition</i>	
Qixuan Zhang, Zhifeng Wang, Dylan Zhang, Wenjia Niu, Sabrina Caldwell, Tom Gedeon, Yang Liu and Zhenyue Qin . . . . .	4484
<i>IDEAW: Robust Neural Audio Watermarking with Invertible Dual-Embedding</i>	
Pengcheng Li, Xulong Zhang, Jing Xiao and Jianzong Wang . . . . .	4500
<i>Leveraging Conflicts in Social Media Posts: Unintended Offense Dataset</i>	
Che Wei Tsai, Yen-Hao Huang, Tsu-Keng Liao, Didier Fernando Salazar Estrada, Retnani Latifah and Yi-Shin Chen . . . . .	4512
<i>Outcome-Constrained Large Language Models for Countering Hate Speech</i>	
Lingzi Hong, Pengcheng Luo, Eduardo Blanco and Xiaoying Song . . . . .	4523
<i>Multiple Sources are Better Than One: Incorporating External Knowledge in Low-Resource Glossing</i>	
Changbing Yang, Garrett Nicolai and Miikka Silfverberg . . . . .	4537
<i>Adaptive Immune-based Sound-Shape Code Substitution for Adversarial Chinese Text Attacks</i>	
Ao Wang, Xinghao Yang, Chen Li, Bao-di Liu and Weifeng Liu . . . . .	4553
<i>Bootstrapped Policy Learning for Task-oriented Dialogue through Goal Shaping</i>	
Yangyang Zhao, Ben Niu, Mehdi Dastani and Shihan Wang . . . . .	4566
<i>PsyGUARD: An Automated System for Suicide Detection and Risk Assessment in Psychological Counseling</i>	
Huachuan Qiu, Lizhi Ma and Zhenzhong Lan . . . . .	4581
<i>World to Code: Multi-modal Data Generation via Self-Instructed Compositional Captioning and Filtering</i>	
Jiacong Wang, Bohong Wu, Haiyong Jiang, Zhou Xun, Xin Xiao, Haoyuan Guo and Jun Xiao	4608
<i>DVD: Dynamic Contrastive Decoding for Knowledge Amplification in Multi-Document Question Answering</i>	
Jing Jin, Houfeng Wang, Hao Zhang, Xiaoguang Li and Zhijiang Guo . . . . .	4624
<i>How Do Humans Write Code? Large Models Do It the Same Way Too</i>	
Long Li, Xuzheng He, Haozhe Wang, Linlin Wang and Liang He . . . . .	4638
<i>Retrosplex: Language Agent Meets Offline Reinforcement Learning Critic</i>	
Yufei Xiang, Yiqun Shen, Yeqin Zhang and Nguyen Cam-Tu . . . . .	4650
<i>Forgetting Curve: A Reliable Method for Evaluating Memorization Capability for Long-Context Models</i>	
Xinyu Liu, Runsong Zhao, Pengcheng Huang, Chunyang Xiao, Bei Li, Jingang Wang, Tong Xiao and JingBo Zhu . . . . .	4667
<i>Retrieve-Plan-Generation: An Iterative Planning and Answering Framework for Knowledge-Intensive LLM Generation</i>	
Yuanjie Lyu, Zihan Niu, Zheyong Xie, Chao Zhang, Tong Xu, Yang Wang and Enhong Chen	4683

<i>CoEvol: Constructing Better Responses for Instruction Finetuning through Multi-Agent Cooperation</i> Renhao Li, Minghuan Tan, Derek F. Wong and Min Yang .....	4703
<i>A Peek into Token Bias: Large Language Models Are Not Yet Genuine Reasoners</i> Bowen Jiang, Yangxinyu Xie, Zhuoqun Hao, Xiaomeng Wang, Tanwi Mallick, Weijie J Su, Camillo Jose Taylor and Dan Roth .....	4722
<i>Bayesian Calibration of Win Rate Estimation with LLM Evaluators</i> Yicheng Gao, Gonghan Xu, Zhe Wang and Arman Cohan .....	4757
<i>MuMath-Code: Combining Tool-Use Large Language Models with Multi-perspective Data Augmentation for Mathematical Reasoning</i> Shuo Yin, Weihao You, Zhilong Ji, Guoqiang Zhong and Jinfeng Bai .....	4770
<i>Seeing the Forest through the Trees: Data Leakage from Partial Transformer Gradients</i> Weijun Li, Qionghai Xu and Mark Dras .....	4786
<i>RWKV-CLIP: A Robust Vision-Language Representation Learner</i> Tiancheng Gu, Kaicheng Yang, Xiang An, Ziyong Feng, Dongnan Liu, Weidong Cai and Jiankang Deng .....	4799
<i>KidLM: Advancing Language Models for Children – Early Insights and Future Directions</i> Mir Tafseer Nayeem and Davood Rafiei .....	4813
<i>Using Language Models to Disambiguate Lexical Choices in Translation</i> Josh Barua, Sanjay Subramanian, Kayo Yin and Alane Suhr .....	4837
<i>How Does the Disclosure of AI Assistance Affect the Perceptions of Writing?</i> Zhuoyan Li, Chen Liang, Jing Peng and Ming Yin .....	4849
<i>An Unsupervised Approach to Achieve Supervised-Level Explainability in Healthcare Records</i> Joakim Edin, Maria Maistro, Lars Maaløe, Lasse Borgholt, Jakob Drachmann Havtorn and Tuukka Ruotsalo .....	4869
<i>Crafting Personalized Agents through Retrieval-Augmented Generation on Editable Memory Graphs</i> Zheng Wang, Zhongyang Li, Zeren Jiang, Dandan Tu and Wei Shi .....	4891
<i>EVEDIT: Event-based Knowledge Editing for Deterministic Knowledge Propagation</i> Jiateng Liu, Pengfei Yu, Yuji Zhang, Sha Li, Zixuan Zhang, Ruhi Sarikaya, Kevin Small and Heng Ji .....	4907
<i>Modeling Nonnative Sentence Processing with L2 Language Models</i> Tatsuya Aoyama and Nathan Schneider .....	4927
<i>From the Least to the Most: Building a Plug-and-Play Visual Reasoner via Data Synthesis</i> Chuanqi Cheng, Jian Guan, Wei Wu and Rui Yan .....	4941
<i>Quality Matters: Evaluating Synthetic Data for Tool-Using LLMs</i> Shadi Iskander, Sofia Tolmach, Ori Shapira, Nachshon Cohen and Zohar Karnin .....	4958
<i>Cross-Domain Audio Deepfake Detection: Dataset and Analysis</i> Yuang Li, Min Zhang, Mengxin Ren, Xiaosong Qiao, Miaomiao Ma, Daimeng Wei and Hao Yang	4977
<i>MaPPER: Multimodal Prior-guided Parameter Efficient Tuning for Referring Expression Comprehension</i> Ting Liu, Zunnan Xu, Yue Hu, Liangtao Shi, Zhiqiang Wang and Qunjun Yin .....	4984

<i>Hierarchical Deconstruction of LLM Reasoning: A Graph-Based Framework for Analyzing Knowledge Utilization</i>	
Miyoung Ko, Sue Hyun Park, Joonsuk Park and Minjoon Seo .....	4995
<i>Aligning Translation-Specific Understanding to General Understanding in Large Language Models</i>	
Yichong Huang, Baohang Li, Xiaocheng Feng, Wenshuai Huo, Chengpeng Fu, Ting Liu and Bing Qin .....	5028
<i>FOOL ME IF YOU CAN! An Adversarial Dataset to Investigate the Robustness of LMs in Word Sense Disambiguation</i>	
Mohamad Ballout, Anne Dedert, Nohayr Muhammad Abdelmoneim, Ulf Krumnack, Gunther Heidemann and Kai-Uwe Kühnberger .....	5042
<i>Concept-skill Transferability-based Data Selection for Large Vision-Language Models</i>	
Jaewoo Lee, Boyang Li and Sung Ju Hwang .....	5060
<i>LLMs Assist NLP Researchers: Critique Paper (Meta-)Reviewing</i>	
Jiangshu Du, Yibo Wang, Wenting Zhao, Zhongfen Deng, Shuaiqi Liu, Renze Lou, Henry Peng Zou, Pranav Narayanan Venkit, Nan Zhang, Mukund Srinath, Haoran Ranran Zhang, Vipul Gupta, Yinghui Li, Tao Li, Fei Wang, Qin Liu, Tianlin Liu, Pengzhi Gao, Congying Xia, Chen Xing, Cheng Jiayang, Zhaowei Wang, Ying Su, Raj Sanjay Shah, Ruohao Guo, Jing Gu, Haoran Li, Kangda Wei, Zihao Wang, Lu Cheng, Surangika Ranathunga, Meng Fang, Jie Fu, Fei Liu, Ruihong Huang, Eduardo Blanco, Yixin Cao, Rui Zhang, Philip S. Yu and Wenpeng Yin .....	5081
<i>Academics Can Contribute to Domain-Specialized Language Models</i>	
Mark Dredze, Genta Indra Winata, Prabhanjan Kambadur, Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, David S Rosenberg and Sebastian Gehrmann .....	5100
<i>Beyond Reference: Evaluating High Quality Translations Better than Human References</i>	
Keonwoong Noh, Seokjin Oh and Woohwan Jung .....	5111
<i>Unveiling the Lexical Sensitivity of LLMs: Combinatorial Optimization for Prompt Enhancement</i>	
Pengwei Zhan, Zhen Xu, Qian Tan, Jie Song and Ru Xie .....	5128
<i>SEACrowd: A Multilingual Multimodal Data Hub and Benchmark Suite for Southeast Asian Languages</i>	
Holy Lovenia, Rahmad Mahendra, Salsabil Maulana Akbar, Lester James Validad Miranda, Jennifer Santoso, Elyanah Aco, Akhdan Fadhilah, Jonibek Mansurov, Joseph Marvin Imperial, Onno P. Kampman, Joel Ruben Antony Moniz, Muhammad Ravi Shulthan Habibi, Frederikus Hudi, Jann Ralley Montalan, Ryan Ignatius Hadiwijaya, Joanito Agili Lopo, William Nixon, Börje F. Karlsson, James Jaya, Ryandito Diandaru, Yuze Gao, Patrick Amadeus Irawan, Bin Wang, Jan Christian Blaise Cruz, Chenxi Whitehouse, Ivan Halim Parmonangan, Maria Khelli, Wenyu Zhang, Lucky Susanto, Reynard Adha Ryanda, Sonny Lazuardi Hermawan, Dan John Velasco, Muhammad Dehan Al Kautsar, Willy Fitra Hendria, Yasmin Moslem, Noah Flynn, Muhammad Farid Adilazuarda, Haochen Li, Johannes Lee, R. Damanhuri, Shuo Sun, Muhammad Reza Qorib, Amirbek Djanibekov, Wei Qi Leong, Quyet V. Do, Niklas Muennighoff, Tanrada Pansuwan, Ilham Firdausi Putra, Yan Xu, Tai Ngee Chia, Ayu Purwarianti, Sebastian Ruder, William Chandra Tjhi, Peerat Limkonchotiwat, Alham Fikri Aji, Sedrick Keh, Genta Indra Winata, Ruochen Zhang, Fajri Koto, Zheng Xin Yong and Samuel Cahyawijaya .....	5155
<i>Induct-Learn: Short Phrase Prompting with Instruction Induction</i>	
Po-Chun Chen, Sheng-Lun Wei, Hen-Hsen Huang and Hsin-Hsi Chen .....	5204
<i>Multi-Granularity History and Entity Similarity Learning for Temporal Knowledge Graph Reasoning</i>	
Shi Mingcong, Chunjiang Zhu, Detian Zhang, Shiting Wen and Li Qing .....	5232
<i>LUQ: Long-text Uncertainty Quantification for LLMs</i>	
Caiqi Zhang, Fangyu Liu, Marco Basaldella and Nigel Collier .....	5244

<i>Pretraining Data Detection for Large Language Models: A Divergence-based Calibration Method</i> Weichao Zhang, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan and Xueqi Cheng	5263
<i>Scaling Synthetic Logical Reasoning Datasets with Context-Sensitive Declarative Grammars</i> Damien Sileo	5275
<i>Improving Spoken Language Modeling with Phoneme Classification: A Simple Fine-tuning Approach</i> Maxime Poli, Emmanuel Chemla and Emmanuel Dupoux	5284
<i>Safely Learning with Private Data: A Federated Learning Framework for Large Language Model</i> Jia-Ying Zheng, Hainan Zhang, Lingxiang Wang, Wangjie Qiu, Hong-Wei Zheng and Zhi-Ming Zheng	5293
<i>Formality is Favored: Unraveling the Learning Preferences of Large Language Models on Data with Conflicting Knowledge</i> Jiahuan Li, Yiqing Cao, Shujian Huang and Jiajun Chen	5307
<i>How Does the Textual Information Affect the Retrieval of Multimodal In-Context Learning?</i> Yang Luo, Zangwei Zheng, Zirui Zhu and Yang You	5321
<i>How Far Can We Extract Diverse Perspectives from Large Language Models?</i> Shirley Anugrah Hayati, Minhwa Lee, Dheeraj Rajagopal and Dongyeop Kang	5336
<i>EXPLORA: Efficient Exemplar Subset Selection for Complex Reasoning</i> Kiran Purohit, Venkatesh V, Raghuram Devalla, Krishna Mohan Yerragorla, Sourangshu Bhattacharya and Avishek Anand	5367
<i>An LLM Feature-based Framework for Dialogue Constructiveness Assessment</i> Lexin Zhou, Youmna Farag and Andreas Vlachos	5389
<i>Relevance Is a Guiding Light: Relevance-aware Adaptive Learning for End-to-end Task-oriented Dialogue System</i> Zhanpeng Chen, Zhihong Zhu, Wanshi Xu, Xianwei Zhuang and Yuexian Zou	5410
<i>Dialog2Flow: Pre-training Soft-Contrastive Action-Driven Sentence Embeddings for Automatic Dialog Flow Extraction</i> Sergio Burdisso, Srikanth Madikeri and Petr Motlicek	5421
<i>Words Worth a Thousand Pictures: Measuring and Understanding Perceptual Variability in Text-to-Image Generation</i> Raphael Tang, Crystina Zhang, Lixinyu Xu, Yao Lu, Wenyan Li, Pontus Stenetorp, Jimmy Lin and Ferhan Ture	5441
<i>Investigating LLMs as Voting Assistants via Contextual Augmentation: A Case Study on the European Parliament Elections 2024</i> Ilias Chalkidis	5455
<i>Adaption-of-Thought: Learning Question Difficulty Improves Large Language Models for Reasoning</i> Mayi Xu, Yongqi Li, Ke Sun and Tiejun Qian	5468
<i>LogicST: A Logical Self-Training Framework for Document-Level Relation Extraction with Incomplete Annotations</i> Shengda Fan, Yanting Wang, Shasha Mo and Jianwei Niu	5496
<i>Concept Space Alignment in Multilingual LLMs</i> Qiwei Peng and Anders Søgaard	5511

<i>Predicting Rewards Alongside Tokens: Non-disruptive Parameter Insertion for Efficient Inference Intervention in Large Language Model</i>	
Chenhan Yuan, Fei Huang, Ru Peng, Keming Lu, Bowen Yu, Chang Zhou and Jingren Zhou	5527
<i>NLEBench+NorGLM: A Comprehensive Empirical Analysis and Benchmark Dataset for Generative Language Models in Norwegian</i>	
Peng Liu, Lemei Zhang, Terje Farup, Even W. Lauvrak, Jon Espen Ingvaldsen, Simen Eide, Jon Atle Gulla and Zhirong Yang	5543
<i>RSA-Control: A Pragmatics-Grounded Lightweight Controllable Text Generation Framework</i>	
Yifan Wang and Vera Demberg	5561
<i>Scaling Laws Across Model Architectures: A Comparative Analysis of Dense and MoE Models in Large Language Models</i>	
Siqi Wang, Zhengyu Chen, Bei Li, Keqing He, Min Zhang and Jingang Wang	5583
<i>Synergizing In-context Learning with Hints for End-to-end Task-oriented Dialog Systems</i>	
Vishal Vivek Saley, Rocktim Jyoti Das, Dinesh Raghu and Mausam	5596
<i>REAR: A Relevance-Aware Retrieval-Augmented Framework for Open-Domain Question Answering</i>	
Yuhao Wang, Ruiyang Ren, Junyi Li, Xin Zhao, Jing Liu and Ji-Rong Wen	5613
<i>Leave No Document Behind: Benchmarking Long-Context LLMs with Extended Multi-Doc QA</i>	
Minzheng Wang, Longze Chen, Fu Cheng, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu, Nan Xu, Lei Zhang, Run Luo, Yunshui Li, Min Yang, Fei Huang and Yongbin Li	5627
<i>On Mitigating Performance Disparities in Multilingual Speech Recognition</i>	
Monorama Swain, Anna Katrine Van Zee and Anders Søgaard	5647
<i>Thinking Outside of the Differential Privacy Box: A Case Study in Text Privatization with Language Model Prompting</i>	
Stephen Meisenbacher and Florian Matthes	5656
<i>To Preserve or To Compress: An In-Depth Study of Connector Selection in Multimodal Large Language Models</i>	
Junyan Lin, Haoran Chen, Dawei Zhu and Xiaoyu Shen	5666
<i>What is "Typological Diversity" in NLP?</i>	
Esther Ploeger, Wessel Poelman, Miryam De Lhoneux and Johannes Bjerva	5681
<i>The Computational Anatomy of Humility: Modeling Intellectual Humility in Online Public Discourse</i>	
Xiaobo Guo, Neil Potnis, Melody Yu, Nabeel Gillani and Soroush Vosoughi	5701
<i>Consistent Bidirectional Language Modelling: Expressive Power and Representational Conciseness</i>	
Georgi Shopov and Stefan Gerdjikov	5724
<i>Benchmarking Vision Language Models for Cultural Understanding</i>	
Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd Van Steenkiste, Lisa Anne Hendricks, Karolina Stanczak and Aishwarya Agrawal	5769
<i>Methods of Automatic Matrix Language Determination for Code-Switched Speech</i>	
Olga Iakovenko and Thomas Hain	5791
<i>Analyzing Key Factors Influencing Emotion Prediction Performance of VLLMs in Conversational Contexts</i>	
Jaewook Lee, Yeajin Jang, Hongjin Kim, Woojin Lee and Harksoo Kim	5801

<i>Context-Aware Assistant Selection for Improved Inference Acceleration with Large Language Models</i> Jerry Huang, Prasanna Parthasarathi, Mehdi Rezagholizadeh and Sarath Chandar . . . . .	5817
<i>Teaching Small Language Models Reasoning through Counterfactual Distillation</i> Tao Feng, Yicheng Li, Li Chenglin, Hao Chen, Fei Yu and Yin Zhang . . . . .	5831
<i>Pretraining Language Models Using Translationese</i> Meet Doshi, Raj Dabre and Pushpak Bhattacharyya . . . . .	5843
<i>Quantifying the Gaps Between Translation and Native Perception in Training for Multimodal, Multilingual Retrieval</i> Kyle Buettner and Adriana Kovashka . . . . .	5863
<i>MTA4DPR: Multi-Teaching-Assistants Based Iterative Knowledge Distillation for Dense Passage Retrieval</i> Qixi Lu, Endong Xun and Gongbo Tang . . . . .	5871
<i>Fine-Grained Detection of Solidarity for Women and Migrants in 155 Years of German Parliamentary Debates</i> Aida Kostikova, Dominik Beese, Benjamin Paassen, Ole Pütz, Gregor Wiedemann and Steffen Eger . . . . .	5884
<i>CItruS: Chunked Instruction-aware State Eviction for Long Sequence Modeling</i> Yu Bai, Xiyuan Zou, Heyan Huang, Sanxing Chen, Marc-Antoine Rondeau, Yang Gao and Jackie CK Cheung . . . . .	5908
<i>Story Embeddings — Narrative-Focused Representations of Fictional Stories</i> Hans Ole Hatzel and Chris Biemann . . . . .	5931
<i>C-LLM: Learn to Check Chinese Spelling Errors Character by Character</i> Kunting Li, Yong Hu, Liang He, Fandong Meng and Jie Zhou . . . . .	5944
<i>PSC: Extending Context Window of Large Language Models via Phase Shift Calibration</i> Wenqiao Zhu, Chao Xu, Lulu Wang and Jun Wu . . . . .	5958
<i>Video-LLaVA: Learning United Visual Representation by Alignment Before Projection</i> Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin and Li Yuan . . . . .	5971
<i>SaySelf: Teaching LLMs to Express Confidence with Self-Reflective Rationales</i> Tianyang Xu, Shujin Wu, Shizhe Diao, Xiaoze Liu, Xingyao Wang, Yangyi Chen and Jing Gao . . . . .	5985
<i>Mitigating Frequency Bias and Anisotropy in Language Model Pre-Training with Syntactic Smoothing</i> Richard Diehl Martinez, Zebulun Goriely, Andrew Caines, Paula Buttery and Lisa Beinborn . . . . .	5999
<i>ToxiCloakCN: Evaluating Robustness of Offensive Language Detection in Chinese with Cloaking Perturbations</i> Yunze Xiao, Yujia Hu, Kenny Tsu Wei Choo and Roy Ka-Wei Lee . . . . .	6012
<i>Boosting Scientific Concepts Understanding: Can Analogy from Teacher Models Empower Student Models?</i> Siyu Yuan, Cheng Jiayang, Lin Qiu and Deqing Yang . . . . .	6026
<i>Model Internals-based Answer Attribution for Trustworthy Retrieval-Augmented Generation</i> Jirui Qi, Gabriele Sarti, Raquel Fernández and Arianna Bisazza . . . . .	6037
<i>Do Large Language Models Know How Much They Know?</i> Gabriele Prato, Jerry Huang, Prasanna Parthasarathi, Shagun Sodhani and Sarath Chandar . . . . .	6054



<i>Investigating Mysteries of CoT-Augmented Distillation</i> Somin Wadhwa, Silvio Amir and Byron C Wallace .....	6071
<i>SciPrompt: Knowledge-augmented Prompting for Fine-grained Categorization of Scientific Topics</i> Zhiwen You, Kanyao Han, Haotian Zhu, Bertram Ludascher and Jana Diesner .....	6087
<i>Distilling Knowledge from Text-to-Image Generative Models Improves Visio-Linguistic Reasoning in CLIP</i> Samyadeep Basu, Shell Xu Hu, Maziar Sanjabi, Daniela Massiceti and Soheil Feizi.....	6105
<i>Learning from Natural Language Explanations for Generalizable Entity Matching</i> Somin Wadhwa, Adit Krishnan, Runhui Wang, Byron C Wallace and Luyang Kong.....	6114
<i>Do You Know What You Are Talking About? Characterizing Query-Knowledge Relevance For Reliable Retrieval Augmented Generation</i> Zhuohang Li, Jiaxin Zhang, Chao Yan, Kamalika Das, Sricharan Kumar, Murat Kantarcioglu and Bradley A. Malin .....	6130
<i>On the Reliability of Psychological Scales on Large Language Models</i> Jen-tse Huang, Wenxiang Jiao, Man Ho Lam, Eric John Li, Wenxuan Wang and Michael Lyu.....	6152
<i>Contrastive Entity Coreference and Disambiguation for Historical Texts</i> Abhishek Arora, Emily Silcock, Melissa Dell and Leander Heldring .....	6174
<i>Finer: Investigating and Enhancing Fine-Grained Visual Concept Recognition in Large Vision Language Models</i> Jeonghwan Kim and Heng Ji .....	6187
<i>Evaluating LLMs for Targeted Concept Simplification for Domain-Specific Texts</i> Sumit Asthana, Hannah Rashkin, Elizabeth Clark, Fantine Huot and Mirella Lapata .....	6208
<i>VLFeedback: A Large-Scale AI Feedback Dataset for Large Vision-Language Models Alignment</i> Lei Li, Zhihui Xie, Mukai Li, Shunian Chen, Peiyi Wang, Liang Chen, Yazheng Yang, Benyou Wang, Lingpeng Kong and Qi Liu .....	6227
<i>Focused Large Language Models are Stable Many-Shot Learners</i> Peiwen Yuan, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Yueqi Zhang, Chuyi Tan, Boyuan Pan, Heda Wang, Yao Hu and Kan Li .....	6247
<i>Reconsidering Sentence-Level Sign Language Translation</i> Garrett Tanzer, Maximus Shengelia, Ken Harrenstien and David Uthus .....	6262
<i>GAMA: A Large Audio-Language Model with Advanced Audio Understanding and Complex Reasoning Abilities</i> Sreyan Ghosh, Sonal Kumar, Ashish Seth, Chandra Kiran Reddy Evuru, Utkarsh Tyagi, S Sakshi, Oriol Nieto, Ramani Duraiswami and Dinesh Manocha .....	6288
<i>Verba volant, scripta volant? Don't worry! There are computational solutions for protoword reconstruction</i> Liviu P Dinu, Ana Sabina Uban, Alina Maria Cristea, Ioan-Bogdan Iordache, Teodor-George Marchitan, Simona Georgescu and Laurentiu Zoicas .....	6314
<i>ChatGPT Doesn't Trust Chargers Fans: Guardrail Sensitivity in Context</i> Victoria R Li, Yida Chen and Naomi Saphra .....	6327
<i>Personas as a Way to Model Truthfulness in Language Models</i> Nitish Joshi, Javier Rando, Abulhair Saparov, Najoung Kim and He He.....	6346

<i>Satyrn: A Platform for Analytics Augmented Generation</i>	
Marko Sterbentz, Cameron Barrie, Shubham Shahi, Abhratanu Dutta, Donna Hooshmand, Harper Pack and Kristian J Hammond . . . . .	6360
<i>EH-MAM: Easy-to-Hard Masked Acoustic Modeling for Self-Supervised Speech Representation Learning</i>	
Ashish Seth, Ramaneswaran Selvakumar, S Sakshi, Sonal Kumar, Sreyan Ghosh and Dinesh Manocha . . . . .	6386
<i>EPO: Hierarchical LLM Agents with Environment Preference Optimization</i>	
Qi Zhao, Haotian Fu, Chen Sun and George Konidaris . . . . .	6401
<i>Detection and Measurement of Syntactic Templates in Generated Text</i>	
Chantal Shaib, Yanai Elazar, Junyi Jessy Li and Byron C Wallace . . . . .	6416
<i>UOUO: Uncontextualized Uncommon Objects for Measuring Knowledge Horizons of Vision Language Models</i>	
Xinyu Pi, Mingyuan Wu, Jize Jiang, Haozhen Zheng, Beitong Tian, ChengXiang Zhai, Klara Nahrstedt and Zhiting Hu . . . . .	6432
<i>Optimized Speculative Sampling for GPU Hardware Accelerators</i>	
Dominik Wagner, Seanie Lee, Ilja Baumann, Philipp Seeberger, Korbinian Riedhammer and Tobias Bocklet . . . . .	6442
<i>Personalized Pieces: Efficient Personalized Large Language Models through Collaborative Efforts</i>	
Zhaoxuan Tan, Zheyuan Liu and Meng Jiang . . . . .	6459
<i>Democratizing Large Language Models via Personalized Parameter-Efficient Fine-tuning</i>	
Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin and Meng Jiang . . . . .	6476
<i>Unifying Multimodal Retrieval via Document Screenshot Embedding</i>	
Xueguang Ma, Sheng-Chieh Lin, Minghan Li, Wenhui Chen and Jimmy Lin . . . . .	6492
<i>Neuron Specialization: Leveraging Intrinsic Task Modularity for Multilingual Machine Translation</i>	
Shaomu Tan, Di Wu and Christof Monz . . . . .	6506
<i>An Audit on the Perspectives and Challenges of Hallucinations in NLP</i>	
Pranav Narayanan Venkit, Tatiana Chakravorti, Vipul Gupta, Heidi Biggs, Mukund Srinath, Koustava Goswami, Sarah Rajtmajer and Shomir Wilson . . . . .	6528
<i>Discovering Knowledge-Critical Subnetworks in Pretrained Language Models</i>	
Deniz Bayazit, Negar Foroutan, Zeming Chen, Gail Weiss and Antoine Bosselut . . . . .	6549
<i>Reconstruct Your Previous Conversations! Comprehensively Investigating Privacy Leakage Risks in Conversations with GPT Models</i>	
Junjie Chu, Zeyang Sha, Michael Backes and Yang Zhang . . . . .	6584
<i>Right for Right Reasons: Large Language Models for Verifiable Commonsense Knowledge Graph Question Answering</i>	
Armin Toroghi, Willis Guo, Mohammad Mahdi Abdollah Pour and Scott Sanner . . . . .	6601
<i>Verifiable, Debuggable, and Repairable Commonsense Logical Reasoning via LLM-based Theory Resolution</i>	
Armin Toroghi, Willis Guo, Ali Pesaraghader and Scott Sanner . . . . .	6634
<i>Understanding and Mitigating Language Confusion in LLMs</i>	
Kelly Marchisio, Wei-Yin Ko, Alexandre Berard, Théo Dehaze and Sebastian Ruder . . . . .	6653

<i>Can Large Language Models Learn Independent Causal Mechanisms?</i>	
Gael Gendron, Bao Trung Nguyen, Alex Yuxuan Peng, Michael Witbrock and Gillian Dobbie	
6678	
<i>MirrorStories: Reflecting Diversity through Personalized Narrative Generation with Large Language Models</i>	
Sarfarozi Yunusov, Hamza Sidat and Ali Emami	6702
<i>InterIntent: Investigating Social Intelligence of LLMs via Intention Understanding in an Interactive Game Context</i>	
Ziyi Liu, Abhishek Anand, Pei Zhou, Jen-tse Huang and Jieyu Zhao	6718
<i>Locating Information Gaps and Narrative Inconsistencies Across Languages: A Case Study of LGBT People Portrayals on Wikipedia</i>	
Farhan Samir, Chan Young Park, Anjalie Field, Vered Shwartz and Yulia Tsvetkov	6747
<i>From Local Concepts to Universals: Evaluating the Multicultural Understanding of Vision-Language Models</i>	
Mehar Bhatia, Sahithya Ravi, Aditya Chinchure, EunJeong Hwang and Vered Shwartz	6763
<i>Dynamic Multi-Reward Weighting for Multi-Style Controllable Generation</i>	
Karin De Langis, Ryan Koo and Dongyeop Kang	6783
<i>MMNeuron: Discovering Neuron-Level Domain-Specific Interpretation in Multimodal Large Language Model</i>	
Jiahao Huo, Yibo Yan, Boren Hu, Yutao Yue and Xuming Hu	6801
<i>Learning to Extract Structured Entities Using Language Models</i>	
Haolun Wu, Ye Yuan, Liana Mikaelyan, Alexander Meulemans, Xue Liu, James Hensman and Bhaskar Mitra	6817
<i>Efficient LLM Comparative Assessment: A Product of Experts Framework for Pairwise Comparisons</i>	
Adian Liusie, Vatsal Raina, Yassir Fathullah and Mark Gales	6835
<i>A Survey of AMR Applications</i>	
Shira Wein and Juri Opitz	6856
<i>Beyond Embeddings: The Promise of Visual Table in Visual Reasoning</i>	
Yiwu Zhong, Zi-Yuan Hu, Michael Lyu and Liwei Wang	6876
<i>CareCorpus+: Expanding and Augmenting Caregiver Strategy Data to Support Pediatric Rehabilitation</i>	
Shahla Farzana, Ivana Lucero, Vivian Villegas, Vera C Kaelin, Mary Khetani and Natalie Parde	6912
<i>Taylor Unswift: Secured Weight Release for Large Language Models via Taylor Expansion</i>	
Guanchu Wang, Yu-Neng Chuang, Ruixiang Tang, Shaochen Zhong, Jiayi Yuan, Hongye Jin, Zirui Liu, Vipin Chaudhary, Shuai Xu, James Caverlee and Xia Hu	6928
<i>TimeR<sup>4</sup> : Time-aware Retrieval-Augmented Large Language Models for Temporal Knowledge Graph Question Answering</i>	
Xinying Qian, Ying Zhang, Yu Zhao, Baohang Zhou, Xuhui Sui, Li Zhang and Kehui Song	6942
<i>Knowledge-Centric Hallucination Detection</i>	
Xiangkun Hu, Dongyu Ru, Lin Qiu, Qipeng Guo, Tianhang Zhang, Yang Xu, Yun Luo, Pengfei Liu, Yue Zhang and Zheng Zhang	6953

<i>Revealing the Parallel Multilingual Learning within Large Language Models</i>	
Yongyu Mu, Peinan Feng, Zhiquan Cao, Yuzhang Wu, Bei Li, Chenglong Wang, Tong Xiao, Kai Song, Tongran Liu, Chunliang Zhang and JingBo Zhu . . . . .	6976
<i>Automatic Instruction Evolving for Large Language Models</i>	
Weihao Zeng, Can Xu, Yingxiu Zhao, Jian-Guang Lou and Weizhu Chen . . . . .	6998
<i>RepEval: Effective Text Evaluation with LLM Representation</i>	
Shuqian Sheng, Yi Xu, Tianhang Zhang, Zanwei Shen, Luoyi Fu, Jiabin Ding, Lei Zhou, Xiaoying Gan, Xinbing Wang and Chenghu Zhou . . . . .	7019
<i>Generative Models for Automatic Medical Decision Rule Extraction from Text</i>	
Yuxin He, Buzhou Tang and Xiaoling Wang . . . . .	7034
<i>Encoding and Controlling Global Semantics for Long-form Video Question Answering</i>	
Thong Thanh Nguyen, Zhiyuan Hu, Xiaobao Wu, Cong-Duy T Nguyen, See-Kiong Ng and Anh Tuan Luu . . . . .	7049
<i>Towards Understanding Jailbreak Attacks in LLMs: A Representation Space Analysis</i>	
Yuping Lin, Pengfei He, Han Xu, Yue Xing, Makoto Yamada, Hui Liu and Jiliang Tang . . . . .	7067
<i>Enhancing Legal Case Retrieval via Scaling High-quality Synthetic Query-Candidate Pairs</i>	
Cheng Gao, Chaojun Xiao, Zhenghao Liu, Huimin Chen, Zhiyuan Liu and Maosong Sun . . . . .	7086
<i>Does Large Language Model Contain Task-Specific Neurons?</i>	
Ran Song, Shizhu He, Shuting Jiang, Yantuan Xian, Shengxiang Gao, Kang Liu and Zhengtao Yu . . . . .	7101
<i>Liar, Liar, Logical Mire: A Benchmark for Suppositional Reasoning in Large Language Models</i>	
Philipp Mondorf and Barbara Plank . . . . .	7114
<i>Advancing Test-Time Adaptation in Wild Acoustic Test Settings</i>	
Hongfu Liu, Hengguan Huang and Ye Wang . . . . .	7138
<i>Learning to Retrieve Iteratively for In-Context Learning</i>	
Yunmo Chen, Tongfei Chen, Harsh Jhamtani, Patrick Xia, Richard Shin, Jason Eisner and Benjamin Van Durme . . . . .	7156
<i>Taxonomy-guided Semantic Indexing for Academic Paper Search</i>	
SeongKu Kang, Yunyi Zhang, Pengcheng Jiang, Dongha Lee, Jiawei Han and Hwanjo Yu . . . . .	7169
<i>Python is Not Always the Best Choice: Embracing Multilingual Program of Thoughts</i>	
Xianzhen Luo, Qingfu Zhu, Zhiming Zhang, Libo Qin, Xuanyu Zhang, Qing Yang, Dongliang Xu and Wanxiang Che . . . . .	7185
<i>Advancing Adversarial Suffix Transfer Learning on Aligned Large Language Models</i>	
Hongfu Liu, Yuxi Xie, Ye Wang and Michael Shieh . . . . .	7213
<i>Incomplete Utterance Rewriting with Editing Operation Guidance and Utterance Augmentation</i>	
Zhiyu Cao, Peifeng LI, Yaxin Fan and Qiaoming Zhu . . . . .	7225
<i>FROG: Evaluating Fuzzy Reasoning of Generalized Quantifiers in LLMs</i>	
Yiyuan Li, Shichao Sun and Pengfei Liu . . . . .	7239
<i>Aligning Large Language Models with Diverse Political Viewpoints</i>	
Dominik Stambach, Philine Widmer, Eunjung Cho, Caglar Gulcehre and Elliott Ash . . . . .	7257

<i>You Gotta be a Doctor, Lin": An Investigation of Name-Based Bias of Large Language Models in Employment Recommendations</i>	
Huy Nghiem, John Prindle, Jieyu Zhao and Hal Daumé Iii . . . . .	7268
<i>Extending Context Window of Large Language Models from a Distributional Perspective</i>	
Yingsheng Wu, Yuxuan Gu, Xiaocheng Feng, Weihong Zhong, Dongliang Xu, Qing Yang, Hongtao Liu and Bing Qin . . . . .	7288
<i>Leveraging pre-trained language models for linguistic analysis: A case of argument structure constructions</i>	
Hakyung Sung and Kristopher Kyle . . . . .	7302
<i>MAGIC: Investigation of Large Language Model Powered Multi-Agent in Cognition, Adaptability, Rationality and Collaboration</i>	
Lin Xu, Zhiyuan Hu, Daquan Zhou, Hongyu Ren, Zhen Dong, Kurt Keutzer, See-Kiong Ng and Jiashi Feng . . . . .	7315
<i>Position Engineering: Boosting Large Language Models through Positional Information Manipulation</i>	
Zhiyuan He, Huiqiang Jiang, Zilong Wang, Yuqing Yang, Luna K. Qiu and Lili Qiu . . . . .	7333
<i>Towards Injecting Medical Visual Knowledge into Multimodal LLMs at Scale</i>	
Junying Chen, Chi Gui, Ruyi Ouyang, Anningzhe Gao, Shunian Chen, Guiming Hardy Chen, Xidong Wang, Zhenyang Cai, Ke Ji, Xiang Wan and Benyou Wang . . . . .	7346
<i>ADELIE: Aligning Large Language Models on Information Extraction</i>	
Yunjia Qi, Hao Peng, Xiaozhi Wang, Bin Xu, Lei Hou and Juanzi Li . . . . .	7371
<i>Unveiling Factual Recall Behaviors of Large Language Models through Knowledge Neurons</i>	
Yifei Wang, Yuheng Chen, Wanting Wen, Yu Sheng, Linjing Li and Daniel Dajun Zeng . . . . .	7388
<i>Lexically Grounded Subword Segmentation</i>	
Jindřich Libovický and Jindřich Helcl . . . . .	7403
<i>EAGLE-2: Faster Inference of Language Models with Dynamic Draft Trees</i>	
Yuhui Li, Fangyun Wei, Chao Zhang and Hongyang Zhang . . . . .	7421
<i>Do Text-to-Vis Benchmarks Test Real Use of Visualisations?</i>	
Hy Nguyen, Xuefei He, Andrew Reeson, Cecile Paris, Josiah Poon and Jonathan K. Kummerfeld	7433
<i>Gold Panning in Vocabulary: An Adaptive Method for Vocabulary Expansion of Domain-Specific LLMs</i>	
Chengyuan Liu, Shihang Wang, Lizhi Qing, Kun Kuang, Yangyang Kang, Changlong Sun and Fei Wu . . . . .	7442
<i>Strategic Demonstration Selection for Improved Fairness in LLM In-Context Learning</i>	
Jingyu Hu, Weiru Liu and Mengnan Du . . . . .	7460
<i>Multi-Dialect Vietnamese: Task, Dataset, Baseline Models and Challenges</i>	
Nguyen Van Dinh, Thanh Chi Dang, Luan Thanh Nguyen and Kiet Van Nguyen . . . . .	7476
<i>Is LLM-as-a-Judge Robust? Investigating Universal Adversarial Attacks on Zero-shot LLM Assessment</i>	
Vyas Raina, Adian Liusie and Mark Gales . . . . .	7499
<i>Rethinking the Reversal Curse of LLMs: a Prescription from Human Knowledge Reversal</i>	
Zhicong Lu, Li Jin, Peiguang Li, Yu Tian, Linhao Zhang, Sirui Wang, Guangluan Xu, Changyuan Tian and Xunliang Cai . . . . .	7518

<i>More Than Catastrophic Forgetting: Integrating General Capabilities For Domain-Specific LLMs</i> Chengyuan Liu, Yangyang Kang, Shihang Wang, Lizhi Qing, Fubang Zhao, Chao Wu, Changlong Sun, Kun Kuang and Fei Wu .....	7531
<i>Muting Whisper: A Universal Acoustic Adversarial Attack on Speech Foundation Models</i> Vyas Raina, Rao Ma, Charles McGhee, Kate Knill and Mark Gales .....	7549
<i>GENRA: Enhancing Zero-shot Retrieval with Rank Aggregation</i> Georgios Katsimpras and Georgios Paliouras .....	7566
<i>XplainLLM: A Knowledge-Augmented Dataset for Reliable Grounded Explanations in LLMs</i> Zichen Chen, Jianda Chen, Ambuj Singh and Misha Sra .....	7578
<i>Divide and Conquer Radiology Report Generation via Observation Level Fine-grained Pretraining and Prompt Tuning</i> Yuanpin Zhou and Huogen Wang .....	7597
<i>SURf: Teaching Large Vision-Language Models to Selectively Utilize Retrieved Information</i> Jiashuo Sun, Jihai Zhang, Yucheng Zhou, Zhaochen Su, Xiaoye Qu and Yu Cheng .....	7611
<i>UNO Arena for Evaluating Sequential Decision-Making Capability of Large Language Models</i> Zhanyue Qin, Haochuan Wang, Deyuan Liu, Ziyang Song, Cunhang Fan, Zhao Lv, Jinlin Wu, Zhen Lei, Zhiying Tu, Dianhui Chu, Xiaoyan Yu and Dianbo Sui .....	7630
<i>Middleware for LLMs: Tools Are Instrumental for Language Agents in Complex Environments</i> Yu Gu, Yiheng Shu, Hao Yu, Xiao Liu, Yuxiao Dong, Jie Tang, Jayanth Srinivasa, Hugo Latapie and Yu Su .....	7646
<i>MORPHEUS: Modeling Role from Personalized Dialogue History by Exploring and Utilizing Latent Space</i> Yihong Tang, Bo Wang, Dongming Zhao, Jinxiaojia Jinxiaojia, Zhangjijun Zhangjijun, Ruifang He and Yuexian Hou .....	7664
<i>KnowledgeSG: Privacy-Preserving Synthetic Text Generation with Knowledge Distillation from Server</i> WenHao Wang, Xiaoyu Liang, Rui Ye, Jingyi Chai, Siheng Chen and Yanfeng Wang .....	7677
<i>DAMRO: Dive into the Attention Mechanism of LLM to Reduce Object Hallucination</i> Xuan Gong, Tianshi Ming, Xinpeng Wang and Zhihua Wei .....	7696
<i>Unlocking the Future: Exploring Look-Ahead Planning Mechanistic Interpretability in Large Language Models</i> Tianyi Men, Pengfei Cao, Zhuoran Jin, Yubo Chen, Kang Liu and Jun Zhao .....	7713
<i>Breaking Language Barriers: Cross-Lingual Continual Pre-Training at Scale</i> Wenzhen Zheng, Wenbo Pan, Xu Xu, Libo Qin, Li Yue and Ming Zhou .....	7725
<i>An Empirical Study of Multilingual Reasoning Distillation for Question Answering</i> Patomporn Payoungkhamdee, Peerat Limkonchotiwat, Jinheon Baek, Potsawee Manakul, Can Udomcharoenchaikit, Ekapol Chuangsuwanich and Sarana Nutanong .....	7739
<i>Can Large Language Models Faithfully Express Their Intrinsic Uncertainty in Words?</i> Gal Yona, Roei Aharoni and Mor Geva .....	7752
<i>Does Fine-Tuning LLMs on New Knowledge Encourage Hallucinations?</i> Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart and Jonathan Herzig .....	7765

<i>Bridging Modalities: Enhancing Cross-Modality Hate Speech Detection with Few-Shot In-Context Learning</i>	
Ming Shan Hee, Aditi Kumaresan and Roy Ka-Wei Lee . . . . .	7785
<i>MIND: Multimodal Shopping Intention Distillation from Large Vision-language Models for E-commerce Purchase Understanding</i>	
Baixuan Xu, Weiqi Wang, Haochen Shi, Wenxuan Ding, Huihao Jing, Tianqing Fang, Jiaxin Bai, Xin Liu, Changlong Yu, Zheng Li, Chen Luo, Qingyu Yin, Bing Yin, Long Chen and Yangqiu Song	7800
<i>ECON: On the Detection and Resolution of Evidence Conflicts</i>	
Cheng Jiayang, Chunkit Chan, Qianqian Zhuang, Lin Qiu, Tianhang Zhang, Tengxiao Liu, Yangqiu Song, Yue Zhang, Pengfei Liu and Zheng Zhang . . . . .	7816
<i>“Image, Tell me your story!” Predicting the original meta-context of visual misinformation</i>	
Jonathan Tonglet, Marie-Francine Moens and Iryna Gurevych . . . . .	7845
<i>Improving Retrieval-augmented Text-to-SQL with AST-based Ranking and Schema Pruning</i>	
Zhili Shen, Pavlos Vougiouklis, Chenxin Diao, Kaustubh Vyas, Yuanyi Ji and Jeff Z. Pan . .	7865
<i>Mixture-of-Subspaces in Low-Rank Adaptation</i>	
Taiqiang Wu, Jiahao Wang, Zhe Zhao and Ngai Wong . . . . .	7880
<i>PARIKSHA: A Large-Scale Investigation of Human-LLM Evaluator Agreement on Multilingual and Multi-Cultural Data</i>	
Ishaan Watts, Varun Gumma, Aditya Yadavalli, Vivek Seshadri, Manohar Swaminathan and Sunayana Sitaram . . . . .	7900
<i>LawBench: Benchmarking Legal Knowledge of Large Language Models</i>	
Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Alan Huang, Songyang Zhang, Kai Chen, Zhixin Yin, Zongwen Shen, Jidong Ge and Vincent Ng . . . . .	7933
<i>Efficient Performance Tracking: Leveraging Large Language Models for Automated Construction of Scientific Leaderboards</i>	
Furkan Şahinuç, Thy Thy Tran, Yulia Grishina, Yufang Hou, Bei Chen and Iryna Gurevych	7963
<i>Efficient Vision-Language pre-training via domain-specific learning for human activities</i>	
Adrian Bulat, Yassine Ouali, Ricardo Guerrero, Brais Martinez and Georgios Tzimiropoulos	7978
<i>Empowering Backbone Models for Visual Text Generation with Input Granularity Control and Glyph-Aware Training</i>	
Wenbo Li, Guohao Li, Zhibin Lan, Xue Xu, Wanru Zhuang, Jiachen Liu, Xinyan Xiao and Jinsong Su . . . . .	8001
<i>Evaluating Character Understanding of Large Language Models via Character Profiling from Fictional Works</i>	
Xinfeng Yuan, Siyu Yuan, Yuhan Cui, Tianhe Lin, Xintao Wang, Rui Xu, Jiangjie Chen and Deqing Yang . . . . .	8015
<i>Getting More from Less: Large Language Models are Good Spontaneous Multilingual Learners</i>	
Shimao Zhang, Changjiang Gao, Wenhao Zhu, Jiajun Chen, Xin Huang, Xue Han, Junlan Feng, Chao Deng and Shujian Huang . . . . .	8037
<i>AdaSwitch: Adaptive Switching between Small and Large Agents for Effective Cloud-Local Collaborative Learning</i>	
Hao Sun, Jiayi Wu, Hengyi Cai, Xiaochi Wei, Yue Feng, Bo Wang, Shuaiqiang Wang, Yan Zhang and Dawei Yin . . . . .	8052

<i>CoBa: Convergence Balancer for Multitask Finetuning of Large Language Models</i> Zi Gong, Hang Yu, Cong Liao, Bingchang Liu, Chaoyu Chen and Jianguo Li .....	8063
<i>mDPO: Conditional Preference Optimization for Multimodal Large Language Models</i> Fei Wang, Wenxuan Zhou, James Y. Huang, Nan Xu, Sheng Zhang, Hoifung Poon and Muhao Chen.....	8078
<i>Data Advisor: Dynamic Data Curation for Safety Alignment of Large Language Models</i> Fei Wang, Ninareh Mehrabi, Palash Goyal, Rahul Gupta, Kai-Wei Chang and Aram Galstyan	8089
<i>Language-to-Code Translation with a Single Labeled Example</i> Kaj Bostrom, Harsh Jhamtani, Hao Fang, Sam Thomson, Richard Shin, Patrick Xia, Benjamin Van Durme, Jason Eisner and Jacob Andreas .....	8101
<i>Attribute or Abstain: Large Language Models as Long Document Assistants</i> Jan Buchmann, Xiao Liu and Iryna Gurevych .....	8113
<i>FEDKIM: Adaptive Federated Knowledge Injection into Medical Foundation Models</i> Xiaochen Wang, Jiaqi Wang, Houping Xiao, Jinghui Chen and Fenglong Ma.....	8141
<i>Retrieved In-Context Principles from Previous Mistakes</i> Hao Sun, Yong Jiang, Bo Wang, Yingyan Hou, Yan Zhang, Pengjun Xie and Fei Huang....	8155
<i>EmoKnob: Enhance Voice Cloning with Fine-Grained Emotion Control</i> Haozhe Chen, Run Chen and Julia Hirschberg .....	8170
<i>VPTQ: Extreme Low-bit Vector Post-Training Quantization for Large Language Models</i> Yifei Liu, Jicheng Wen, Yang Wang, Shengyu Ye, Li Lina Zhang, Ting Cao, Cheng Li and Mao Yang.....	8181
<i>An L* Algorithm for Deterministic Weighted Regular Languages</i> Clemente Pasti, Talu Karagöz, Franz Nowak, Anej Svete, Reda Boumasmoud and Ryan Cotterell	8197
<i>Towards Verifiable Text Generation with Evolving Memory and Self-Reflection</i> Hao Sun, Hengyi Cai, Bo Wang, Yingyan Hou, Xiaochi Wei, Shuaiqiang Wang, Yan Zhang and Dawei Yin.....	8211
<i>Pelican: Correcting Hallucination in Vision-LLMs via Claim Decomposition and Program of Thought Verification</i> Pritish Sahu, Karan Sikka and Ajay Divakaran .....	8228
<i>Resampled Datasets Are Not Enough: Mitigating Societal Bias Beyond Single Attributes</i> Yusuke Hirota, Jerone Andrews, Dora Zhao, Orestis Papakyriakopoulos, Apostolos Modas, Yuta Nakashima and Alice Xiang.....	8249
<i>RealVul: Can We Detect Vulnerabilities in Web Applications with LLM?</i> Di Cao, Yong Liao and Xiuwei Shang.....	8268
<i>Unsupervised End-to-End Task-Oriented Dialogue with LLMs: The Power of the Noisy Channel</i> Brendan King and Jeffrey Flanigan .....	8283
<i>Humans or LLMs as the Judge? A Study on Judgement Bias</i> Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang and Benyou Wang .....	8301
<i>WPO: Enhancing RLHF with Weighted Preference Optimization</i> Wenxuan Zhou, Ravi Agrawal, Shujian Zhang, Sathish Reddy Indurthi, Sanqiang Zhao, Kaiqiang Song, Silei Xu and Chenguang Zhu .....	8328



<i>Walking in Others' Shoes: How Perspective-Taking Guides Large Language Models in Reducing Toxicity and Bias</i>	
Rongwu Xu, Zian Zhou, Tianwei Zhang, Zehan Qi, SU Yao, Ke Xu, Wei Xu and Han Qiu . . .	8341
<i>MetaReflection: Learning Instructions for Language Agents using Past Reflections</i>	
Priyanshu Gupta, Shashank Kirtania, Ananya Singha, Sumit Gulwani, Arjun Radhakrishna, Gustavo Soares and Sherry Shi . . . . .	8369
<i>Stepwise Verification and Remediation of Student Reasoning Errors with Large Language Model Tutors</i>	
Nico Daheim, Jakub Macina, Manu Kapur, Iryna Gurevych and Mrinmaya Sachan . . . . .	8386
<i>On Eliciting Syntax from Language Models via Hashing</i>	
Yiran Wang and Masao Utiyama . . . . .	8412
<i>CliMedBench: A Large-Scale Chinese Benchmark for Evaluating Medical Large Language Models in Clinical Scenarios</i>	
Zetian Ouyang, Yishuai Qiu, Linlin Wang, Gerard De Melo, Ya Zhang, Yanfeng Wang and Liang He . . . . .	8428
<i>The Best Defense is Attack: Repairing Semantics in Textual Adversarial Examples</i>	
Heng Yang and Ke Li . . . . .	8439
<i>CSSL: Contrastive Self-Supervised Learning for Dependency Parsing on Relatively Free Word Ordered and Morphologically Rich Low Resource Languages</i>	
Pretam Ray, Jivnesh Sandhan, Amrith Krishna and Pawan Goyal . . . . .	8458
<i>Perceptions of Linguistic Uncertainty by Language Models and Humans</i>	
Catarina G Belém, Markelle Kelly, Mark Steyvers, Sameer Singh and Padhraic Smyth . . . . .	8467
<i>Explaining and Improving Contrastive Decoding by Extrapolating the Probabilities of a Huge and Hypothetical LM</i>	
Haw-Shiuan Chang, Nanyun Peng, Mohit Bansal, Anil Ramakrishna and Tagyoung Chung . . . . .	8503
<i>Zero-shot Cross-domain Dialogue State Tracking via Context-aware Auto-prompting and Instruction-following Contrastive Decoding</i>	
Xiaoyu Dong, Yujie Feng, Zexin LU, Guangyuan Shi and Xiao-Ming Wu . . . . .	8527
<i>Knowledge Conflicts for LLMs: A Survey</i>	
Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang and Wei Xu . . . . .	8541
<i>MisinfoEval: Generative AI in the Era of Alternative Facts"</i>	
Saadia Gabriel, Liang Lyu, James Siderius, Marzyeh Ghassemi, Jacob Andreas and Asuman E. Ozdaglar . . . . .	8566
<i>MEANT: Multimodal Encoder for Antecedent Information</i>	
Benjamin Irving and Annika Marie Schoene . . . . .	8579
<i>A Thorough Examination of Decoding Methods in the Era of LLMs</i>	
Chufan Shi, Haoran Yang, Deng Cai, Zhisong Zhang, Yifan Wang, Yujiu Yang and Wai Lam . . . . .	8601
<i>AGRaME: Any-Granularity Ranking with Multi-Vector Embeddings</i>	
Revanth Gangi Reddy, Omar Attia, Yunyao Li, Heng Ji and Saloni Potdar . . . . .	8630
<i>FIRST: Faster Improved Listwise Reranking with Single Token Decoding</i>	
Revanth Gangi Reddy, JaeHyeok Doo, Yifei Xu, Md Arafat Sultan, Deevya Swain, Avirup Sil and Heng Ji . . . . .	8642

<i>Exploring Nested Named Entity Recognition with Large Language Models: Methods, Challenges, and Insights</i>	
Hongjin Kim, Jai-Eun Kim and Harksoo Kim . . . . .	8653
<i>ReCaLL: Membership Inference via Relative Conditional Log-Likelihoods</i>	
Roy Xie, Junlin Wang, Ruomin Huang, Minxing Zhang, Rong Ge, Jian Pei, Neil Zhenqiang Gong and Bhuwan Dhingra . . . . .	8671
<i>“Flex Tape Can’t Fix That”: Bias and Misinformation in Edited Language Models</i>	
Karina H Halevy, Anna Sotnikova, Badr AlKhamissi, Syrielle Montariol and Antoine Bosselut	8690
<i>Revisiting Who’s Harry Potter: Towards Targeted Unlearning from a Causal Intervention Perspective</i>	
Yujian Liu, Yang Zhang, Tommi Jaakkola and Shiyu Chang . . . . .	8708
<i>LIONs: An Empirically Optimized Approach to Align Language Models</i>	
Xiao Yu, Qingyang Wu, Yu Li and Zhou Yu . . . . .	8732
<i>Jellyfish: Instruction-Tuning Local Large Language Models for Data Preprocessing</i>	
Haochen Zhang, Yuyang Dong, Chuan Xiao and Masafumi Oyamada . . . . .	8754
<i>A Comprehensive Survey of Scientific Large Language Models and Their Applications in Scientific Discovery</i>	
Yu Zhang, Xiusi Chen, Bowen Jin, Sheng Wang, Shuiwang Ji, Wei Wang and Jiawei Han . .	8783
<i>MiniCheck: Efficient Fact-Checking of LLMs on Grounding Documents</i>	
Liyan Tang, Philippe Laban and Greg Durrett . . . . .	8818
<i>Beyond Label Attention: Transparency in Language Models for Automated Medical Coding via Dictionary Learning</i>	
John Wu, David Wu and Jimeng Sun . . . . .	8848
<i>MOSEL: Inference Serving Using Dynamic Modality Selection</i>	
Bodun Hu, Le Xu, Jeongyoon Moon, Neeraja J Yadwadkar and Aditya Akella . . . . .	8872
<i>From RAG to Riches: Retrieval Interlaced with Sequence Generation</i>	
Palak Jain, Livio Baldini Soares and Tom Kwiatkowski . . . . .	8887
<i>Task Arithmetic can Mitigate Synthetic-to-Real Gap in Automatic Speech Recognition</i>	
Hsuan Su, Hua Farn, Fan-Yun Sun, Shang-Tse Chen and Hung-yi Lee . . . . .	8905
<i>Learning to Correct for QA Reasoning with Black-box LLMs</i>	
Jaehyung Kim, Dongyoung Kim and Yiming Yang . . . . .	8916
<i>AssistantBench: Can Web Agents Solve Realistic and Time-Consuming Tasks?</i>	
Ori Yoran, Samuel Joseph Amouyal, Chaitanya Malaviya, Ben Bogin, Ofir Press and Jonathan Berant . . . . .	8938
<i>PostMark: A Robust Blackbox Watermark for Large Language Models</i>	
Yapei Chang, Kalpesh Krishna, Amir Houmansadr, John Frederick Wieting and Mohit Iyyer	8969
<i>Assessing “Implicit” Retrieval Robustness of Large Language Models</i>	
Xiaoyu Shen, Rexhina Blloshmi, Dawei Zhu, Jiahuan Pei and Wei Zhang . . . . .	8988
<i>On the Relationship between Truth and Political Bias in Language Models</i>	
Suyash Fulay, William Brannon, Shrestha Mohanty, Cassandra Overney, Elinor Poole-Dayana, Deb Roy and Jad Kabbara . . . . .	9004

<i>Can Active Label Correction Improve LLM-based Modular AI Systems?</i> Karan Taneja and Ashok Goel .....	9019
<i>Statistical Uncertainty in Word Embeddings: GloVe-V</i> Andrea Vallebuono, Cassandra Handan-Nader, Christopher D Manning and Daniel E. Ho ..	9032
<i>Annotation alignment: Comparing LLM and human annotations of conversational safety</i> Rajiv Movva, Pang Wei Koh and Emma Pierson .....	9048
<i>DiVERT: Distractor Generation with Variational Errors Represented as Text for Math Multiple-choice Questions</i> Nigel Fernandez, Alexander Scarlatos, Wanyong Feng, Simon Woodhead and Andrew Lan .	9063
<i>The Factuality Tax of Diversity-Intervened Text-to-Image Generation: Benchmark and Fact-Augmented Intervention</i> Yixin Wan, Di Wu, Haoran Wang and Kai-Wei Chang .....	9082
<i>CleanGen: Mitigating Backdoor Attacks for Generation Tasks in Large Language Models</i> Yuetai Li, Zhangchen Xu, Fengqing Jiang, Luyao Niu, Dinuka Sahabandu, Bhaskar Ramasubramanian and Radha Poovendran .....	9101
<i>Enhancing Reinforcement Learning with Dense Rewards from Language Model Critic</i> Meng Cao, Lei Shu, Lei Yu, Yun Zhu, Nevan Wichers, Yinxiao Liu and Lei Meng .....	9119
<i>Words Matter: Reducing Stigma in Online Conversations about Substance Use with Large Language Models</i> Layla Bouzoubaa, Elham Aghakhani and Rezvaneh Rezapour .....	9139
<i>Efficient Sequential Decision Making with Large Language Models</i> Dingyang Chen, Qi Zhang and Yinglun Zhu .....	9157
<i>SignCLIP: Connecting Text and Sign Language by Contrastive Learning</i> Zifan Jiang, Gerard Sant, Amit Moryossef, Mathias Müller, Rico Sennrich and Sarah Ebling	9171
<i>APPLS: Evaluating Evaluation Metrics for Plain Language Summarization</i> Yue Guo, Tal August, Gondy Leroy, Trevor Cohen and Lucy Lu Wang .....	9194
<i>Ontologically Faithful Generation of Non-Player Character Dialogues</i> Nathaniel Weir, Ryan Thomas, Randolph d’Amore, Kellie Hill, Benjamin Van Durme and Harsh Jhamtani .....	9212
<i>LLM See, LLM Do: Leveraging Active Inheritance to Target Non-Differentiable Objectives</i> Luísa Shimabucoro, Sebastian Ruder, Julia Kreutzer, Marzieh Fadaee and Sara Hooker ....	9243
<i>RuBLiMP: Russian Benchmark of Linguistic Minimal Pairs</i> Ekaterina Taktasheva, Maxim Bazhukov, Kirill Koncha, Alena Fenogenova, Ekaterina Artemova and Vladislav Mikhailov .....	9268
<i>Text-Tuple-Table: Towards Information Integration in Text-to-Table Generation via Global Tuple Extraction</i> Zheye Deng, Chunkit Chan, Weiqi Wang, Yuxi Sun, Wei Fan, Tianshi Zheng, Yauwai Yim and Yangqiu Song .....	9300
<i>Toward Compositional Behavior in Neural Models: A Survey of Current Views</i> Kate McCurdy, Paul Soulos, Paul Smolensky, Roland Fernandez and Jianfeng Gao .....	9323

<i>Optimizing Instructions and Demonstrations for Multi-Stage Language Model Programs</i> Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia and Omar Khattab . . . . .	9340
<i>Reverse-Engineering the Reader</i> Samuel Kiegeand, Ethan Wilcox, Afra Amini, David Robert Reich and Ryan Cotterell . . . . .	9367
<i>Synchronous Faithfulness Monitoring for Trustworthy Retrieval-Augmented Generation</i> Di Wu, Jia-Chen Gu, Fan Yin, Nanyun Peng and Kai-Wei Chang . . . . .	9390
<i>Structure Guided Prompt: Instructing Large Language Model in Multi-Step Reasoning by Exploring Graph Structure of the Text</i> Kewei Cheng, Nesreen K. Ahmed, Theodore L. Willke and Yizhou Sun . . . . .	9407
<i>Less is More: Parameter-Efficient Selection of Intermediate Tasks for Transfer Learning</i> David Schulte, Felix Hamborg and Alan Akbik . . . . .	9431
<i>The effects of distance on NPI illusive effects in BERT</i> So Young Lee and Mai Ha Vu . . . . .	9443
<i>Enhancing Systematic Decompositional Natural Language Inference Using Informal Logic</i> Nathaniel Weir, Kate Sanders, Orion Weller, Shreya Sharma, Dongwei Jiang, Zhengping Jiang, Bhavana Dalvi Mishra, Oyvind Tafjord, Peter Jansen, Peter Clark and Benjamin Van Durme . . . . .	9458
<i>Susu Box or Piggy Bank: Assessing Cultural Commonsense Knowledge between Ghana and the US</i> Christabel Acquaye, Haozhe An and Rachel Rudinger . . . . .	9483
<i>Read Anywhere Pointed: Layout-aware GUI Screen Reading with Tree-of-Lens Grounding</i> Yue Fan, Lei Ding, Ching-Chen Kuo, Shan Jiang, Yang Zhao, Xinze Guan, Jie Yang, Yi Zhang and Xin Eric Wang . . . . .	9503
<i>Ranking Manipulation for Conversational Search Engines</i> Samuel Pfrommer, Yatong Bai, Tanmay Gautam and Somayeh Sojoudi . . . . .	9523
<i>Fast Forwarding Low-Rank Training</i> Adir Rahamim, Naomi Saphra, Sara Kangaslahti and Yonatan Belinkov . . . . .	9553
<i>Precise Model Benchmarking with Only a Few Observations</i> Riccardo Fogliato, Pratik Patil, Nil-Jana Akpınar and Mathew Monfort . . . . .	9563
<i>Attribute Diversity Determines the Systematicity Gap in VQA</i> Ian Berlot-Attwell, Kumar Krishna Agrawal, Annabelle Michael Carrell, Yash Sharma and Naomi Saphra . . . . .	9576
<i>ArxivDIGESTables: Synthesizing Scientific Literature into Tables using Language Models</i> Benjamin Newman, Yoonjoo Lee, Aakanksha Naik, Pao Siangliulue, Raymond Fok, Juho Kim, Daniel S Weld, Joseph Chee Chang and Kyle Lo . . . . .	9612
<i>Development of Cognitive Intelligence in Pre-trained Language Models</i> Raj Sanjay Shah, Khushi Bhardwaj and Sashank Varma . . . . .	9632
<i>Modeling Layout Reading Order as Ordering Relations for Visually-rich Document Understanding</i> Chong Zhang, Yi Tu, Yixi Zhao, Chenshu Yuan, Huan Chen, Yue Zhang, Mingxu Chai, Ya Guo, Huijia Zhu, Qi Zhang and Tao Gui . . . . .	9658
<i>Birdie: Advancing State Space Language Modeling with Dynamic Mixtures of Training Objectives</i> Sam Blouir, Jimmy T.h. Smith, Antonios Anastasopoulos and Amarda Shehu . . . . .	9679

<i>Is It Good Data for Multilingual Instruction Tuning or Just Bad Multilingual Evaluation for Large Language Models?</i>	
Pinzhen Chen, Simon Yu, Zhicheng Guo and Barry Haddow .....	9706
<i>Token Erasure as a Footprint of Implicit Vocabulary Items in LLMs</i>	
Sheridan Feucht, David Atkinson, Byron C Wallace and David Bau .....	9727
<i>TraveLER: A Modular Multi-LMM Agent Framework for Video Question-Answering</i>	
Chuyi Shang, Amos You, Sanjay Subramanian, Trevor Darrell and Roei Herzig .....	9740
<i>Evaluating the Effectiveness of Large Language Models in Establishing Conversational Grounding</i>	
Biswesh Mohapatra, Manav Nitin Kapadnis, Laurent Romary and Justine Cassell .....	9767
<i>Unlocking Memorization in Large Language Models with Dynamic Soft Prompting</i>	
Zhepeng Wang, Runxue Bao, Yawen Wu, Jackson Taylor, Cao Xiao, Feng Zheng, Weiwen Jiang, Shangqian Gao and Yanfu Zhang .....	9782
<i>If CLIP Could Talk: Understanding Vision-Language Model Representations Through Their Preferred Concept Descriptions</i>	
Reza Esfandiarpour, Cristina Menghini and Stephen Bach .....	9797
<i>Extract, Define, Canonicalize: An LLM-based Framework for Knowledge Graph Construction</i>	
Bowen Zhang and Harold Soh .....	9820
<i>MQuinE: a Cure for “Z-paradox” in Knowledge Graph Embedding</i>	
Yang Liu, Huang Fang, Yunfeng Cai and Mingming Sun .....	9837
<i>Can Transformers Learn n-gram Language Models?</i>	
Anej Svete, Nadav Borenstein, Mike Zhou, Isabelle Augenstein and Ryan Cotterell .....	9851
<i>StablePrompt : Automatic Prompt Tuning using Reinforcement Learning for Large Language Model</i>	
Minchan Kwon, Gaeun Kim, Jongsuk Kim, Haeil Lee and Junmo Kim .....	9868
<i>Summary of a Haystack: A Challenge to Long-Context LLMs and RAG Systems</i>	
Philippe Laban, Alexander Fabbri, Caiming Xiong and Chien-Sheng Wu .....	9885
<i>Multi-pass Decoding for Grammatical Error Correction</i>	
Xiaoying Wang, Lingling Mu, Jingyi Zhang and Hongfei Xu .....	9904
<i>Into the Unknown Unknowns: Engaged Human Learning through Participation in Language Model Agent Conversations</i>	
Yucheng Jiang, Yijia Shao, Dekun Ma, Sina Semnani and Monica Lam .....	9917
<i>SCOI: Syntax-augmented Coverage-based In-context Example Selection for Machine Translation</i>	
Chenming Tang, Zhixiang Wang and Yunfang Wu .....	9956
<i>Efficient Temporal Extrapolation of Multimodal Large Language Models with Temporal Grounding Bridge</i>	
Yuxuan Wang, Yueqian Wang, Pengfei Wu, Jianxin Liang, Dongyan Zhao, Yang Liu and Zilong Zheng .....	9972
<i>STORYSUMM: Evaluating Faithfulness in Story Summarization</i>	
Melanie Subbiah, Faisal Ladhak, Akankshya Mishra, Griffin Thomas Adams, Lydia Chilton and Kathleen McKeown .....	9988
<i>MMoE: Enhancing Multimodal Models with Mixtures of Multimodal Interaction Experts</i>	
Haofei Yu, Zhengyang Qi, Lawrence Keunho Jang, Russ Salakhutdinov, Louis-Philippe Morency and Paul Pu Liang .....	10006

<i>OmAgent: A Multi-modal Agent Framework for Complex Video Understanding with Task Divide-and-Conquer</i>	
Lu Zhang, Tiancheng Zhao, Heting Ying, Yibo Ma and Kyusong Lee . . . . .	10031
<i>Enhancing Pre-Trained Generative Language Models with Question Attended Span Extraction on Machine Reading Comprehension</i>	
Lin Ai, Zheng Hui, Zizhou Liu and Julia Hirschberg . . . . .	10046
<i>CommonIT: Commonality-Aware Instruction Tuning for Large Language Models via Data Partitions</i>	
Jun Rao, Xuebo Liu, Lian Lian, Shengjun Cheng, Yunjie Liao and Min Zhang . . . . .	10064
<i>ESC: Efficient Speech Coding with Cross-Scale Residual Vector Quantized Transformers</i>	
Yuzhe Gu and Enmao Diao . . . . .	10084
<i>Breaking ReLU Barrier: Generalized MoEfication for Dense Pretrained Models</i>	
Jaeseong Lee, Seung-won Hwang, Wonpyo Park and Mingi Ji . . . . .	10097
<i>Detecting Subtle Differences between Human and Model Languages Using Spectrum of Relative Likelihood</i>	
Yang Xu, Yu Wang, Hao An, Zhichen Liu and Yongyuan Li . . . . .	10108
<i>Optimizing Language Models with Fair and Stable Reward Composition in Reinforcement Learning</i>	
Jiahui Li, Hanlin Zhang, Fengda Zhang, Tai-Wei Chang, Kun Kuang, Long Chen and Jun Zhou	10122
<i>Fine-grained Pluggable Gradient Ascent for Knowledge Unlearning in Language Models</i>	
XiaoHua Feng, Chaochao Chen, Yuyuan Li and Zibin Lin . . . . .	10141
<i>ARM: An Alignment-and-Replacement Module for Chinese Spelling Check Based on LLMs</i>	
Changchun Liu, Kai Zhang, Junzhe Jiang, Zirui Liu, Hanqing Tao, Min Gao and Enhong Chen	10156
<i>On the In-context Generation of Language Models</i>	
Zhongtao Jiang, Yuanzhe Zhang, Kun Luo, Xiaowei Yuan, Jun Zhao and Kang Liu . . . . .	10169
<i>Atomic Inference for NLI with Generated Facts as Atoms</i>	
Joe Stacey, Pasquale Minervini, Haim Dubossarsky, Oana-Maria Camburu and Marek Rei	10188
<i>Towards Robust Speech Representation Learning for Thousands of Languages</i>	
William Chen, Wangyou Zhang, Yifan Peng, Xinjian Li, Jinchuan Tian, Jiatong Shi, Xuankai Chang, Soumi Maiti, Karen Livescu and Shinji Watanabe . . . . .	10205
<i>I Learn Better If You Speak My Language: Understanding the Superior Performance of Fine-Tuning Large Language Models with LLM-Generated Responses</i>	
Xuan Ren, Biao Wu and Lingqiao Liu . . . . .	10225
<i>PreAlign: Boosting Cross-Lingual Transfer by Early Establishment of Multilingual Alignment</i>	
Jiahuan Li, Shujian Huang, Aarron Ching, Xinyu Dai and Jiajun Chen . . . . .	10246
<i>An image speaks a thousand words, but can everyone listen? On image transcreation for cultural relevance</i>	
Simran Khanuja, Sathyanarayanan Ramamoorthy, Yueqi Song and Graham Neubig . . . . .	10258
<i>When Parts Are Greater Than Sums: Individual LLM Components Can Outperform Full Models</i>	
Ting-Yun Chang, Jesse Thomason and Robin Jia . . . . .	10280

<i>Multimodal Clickbait Detection by De-confounding Biases Using Causal Representation Inference</i> Jianxing Yu, Shiqi Wang, Han Yin, Zhenlong Sun, Ruobing Xie, Bo Zhang and Yanghui Rao	10300
<i>Matryoshka-Adaptor: Unsupervised and Supervised Tuning for Smaller Embedding Dimensions</i> Jinsung Yoon, Rajarishi Sinha, Sercan O Arik and Tomas Pfister	10318
<i>KNN-Instruct: Automatic Instruction Construction with K Nearest Neighbor Deduction</i> Jianshang Kou, Benfeng Xu, Chiwei Zhu and Zhendong Mao	10337
<i>Contextualized Sequence Likelihood: Enhanced Confidence Scores for Natural Language Generation</i> Zhen Lin, Shubhendu Trivedi and Jimeng Sun	10351
<i>MixGR: Enhancing Retriever Generalization for Scientific Domain through Complementary Granularity</i> Fengyu Cai, Xinran Zhao, Tong Chen, Sihao Chen, Hongming Zhang, Iryna Gurevych and Heinz Koepl	10369
<i>CARER - Clinical Reasoning-Enhanced Representation for Temporal Health Risk Prediction</i> Tuan Dung Nguyen, Thanh Trung Huynh, Minh Hieu Phan, Quoc Viet Hung Nguyen and Phi Le Nguyen	10392
<i>In-Dialogues We Learn": Towards Personalized Dialogue Without Pre-defined Profiles through In-Dialogue Learning</i> Chuanqi Cheng, Quan Tu, Wei Wu, Shuo Shang, Cunli Mao, Zhengtao Yu and Rui Yan	10408
<i>Encourage or Inhibit Monosemanticity? Revisit Monosemanticity from a Feature Decorrelation Perspective</i> Hanqi Yan, Yanzheng Xiang, Guangyi Chen, Yifei Wang, Lin Gui and Yulan He	10423
<i>Enhancing Language Model Factuality via Activation-Based Confidence Calibration and Guided Decoding</i> Xin Liu, Farima Fatahi Bayat and Lu Wang	10436
<i>Reasoning Robustness of LLMs to Adversarial Typographical Errors</i> Esther Gan, Yiran Zhao, Liying Cheng, Mao Yancan, Anirudh Goyal, Kenji Kawaguchi, Min-Yen Kan and Michael Shieh	10449
<i>InferAligner: Inference-Time Alignment for Harmlessness through Cross-Model Guidance</i> Pengyu Wang, Dong Zhang, Linyang Li, Chenkun Tan, Xinghao Wang, Mozhi Zhang, Ke Ren, Botian Jiang and Xipeng Qiu	10460
<i>Belief Revision: The Adaptability of Large Language Models Reasoning</i> Bryan Wilie, Samuel Cahyawijaya, Etsuko Ishii, Junxian He and Pascale Fung	10480
<i>Fisher Information-based Efficient Curriculum Federated Learning with Large Language Models</i> Ji Liu, Jiaxiang Ren, Ruoming Jin, Zijie Zhang, Yang Zhou, Patrick Valduriez and Dejing Dou	10497
<i>Bio-RFX: Refining Biomedical Extraction via Advanced Relation Classification and Structural Constraints</i> Minjia Wang, Fangzhou Liu, Xiuxing Li, Bowen Dong, Zhenyu Li, Tengyu Pan and Jianyong Wang	10524
<i>Decoding Matters: Addressing Amplification Bias and Homogeneity Issue in Recommendations for Large Language Models</i> Keqin Bao, Jizhi Zhang, Yang Zhang, Xinyue Huo, Chong Chen and Fuli Feng	10540

<i>LLMs Are Prone to Fallacies in Causal Inference</i>	
Nitish Joshi, Abulhair Saparov, Yixin Wang and He He .....	10553
<i>Roleplay-doh: Enabling Domain-Experts to Create LLM-simulated Patients via Eliciting and Adhering to Principles</i>	
Ryan Louie, Ananjan Nandi, William Fang, Cheng Chang, Emma Brunskill and Diyi Yang	10570
<i>The Lou Dataset - Exploring the Impact of Gender-Fair Language in German Text Classification</i>	
Andreas Waldis, Joel Birrer, Anne Lauscher and Iryna Gurevych .....	10604
<i>When Generative Adversarial Networks Meet Sequence Labeling Challenges</i>	
Yu Tong, Ge Chen, Guokai Zheng, Rui Li and Jiang Dazhi .....	10625
<i>Evidence-Focused Fact Summarization for Knowledge-Augmented Zero-Shot Question Answering</i>	
Sungho Ko, Hyunjin Cho, Hyungjoo Chae, Jinyoung Yeo and Dongha Lee .....	10636
<i>Speechworthy Instruction-tuned Language Models</i>	
Hyundong Justin Cho, Nicolaas Paul Jedema, Leonardo F. R. Ribeiro, Karishma Sharma, Pedro Szekely, Alessandro Moschitti, Ruben Janssen and Jonathan May .....	10652
<i>Data, Data Everywhere: A Guide for Pretraining Dataset Construction</i>	
Jupinder Parmar, Shrimai Prabhumoye, Joseph Jennings, Bo Liu, Aastha Jhunjhunwala, Zhilin Wang, Mostofa Patwary, Mohammad Shoeybi and Bryan Catanzaro .....	10671
<i>Fine-Tuning and Prompt Optimization: Two Great Steps that Work Better Together</i>	
Dilara Soylu, Christopher Potts and Omar Khattab .....	10696
<i>Demystifying Verbatim Memorization in Large Language Models</i>	
Jing Huang, Diyi Yang and Christopher Potts .....	10711
<i>AmbigNLG: Addressing Task Ambiguity in Instruction for NLG</i>	
Ayana Niwa and Hayate Iso .....	10733
<i>Distributional Properties of Subword Regularization</i>	
Marco Cognition, Vilém Zouhar and Naoaki Okazaki .....	10753
<i>DataTales: A Benchmark for Real-World Intelligent Data Narration</i>	
Yajing Yang, Qian Liu and Min-Yen Kan .....	10764
<i>Towards Fast Multilingual LLM Inference: Speculative Decoding and Specialized Drafters</i>	
Euiin Yi, Taehyeon Kim, Hongseok Jeung, Du-Seong Chang and Se-Young Yun .....	10789
<i>GlobeSumm: A Challenging Benchmark Towards Unifying Multi-lingual, Cross-lingual and Multi-document News Summarization</i>	
Yangfan Ye, Xiachong Feng, Xiaocheng Feng, Weitao Ma, Libo Qin, Dongliang Xu, Qing Yang, Hongtao Liu and Bing Qin .....	10803
<i>Breaking the Curse of Multilinguality with Cross-lingual Expert Language Models</i>	
Terra Blevins, Tomasz Limisiewicz, Suchin Gururangan, Margaret Li, Hila Gonen, Noah A. Smith and Luke Zettlemoyer .....	10822
<i>More Insightful Feedback for Tutoring: Enhancing Generation Mechanisms and Automatic Evaluation</i>	
Wencke Liermann, Jin-Xia Huang, Yohan Lee and Kong Joo Lee .....	10838
<i>Stable Language Model Pre-training by Reducing Embedding Variability</i>	
Woojin Chung, Jiwoo Hong, Na Min An, James Thorne and Se-Young Yun .....	10852



<i>What is lost in Normalization? Exploring Pitfalls in Multilingual ASR Model Evaluations</i> Kavya Manohar and Leena G Pillai .....	10864
<i>Diversity Over Size: On the Effect of Sample and Topic Sizes for Topic-Dependent Argument Mining Datasets</i> Benjamin Schiller, Johannes Daxenberger, Andreas Waldis and Iryna Gurevych .....	10870
<i>Kiss up, Kick down: Exploring Behavioral Changes in Multi-modal Large Language Models with Assigned Visual Personas</i> Seungjong Sun, Eungu Lee, Seo Yeon Baek, Seunghyun Hwang, Wonbyung Lee, Dongyan Nan, Bernard J Jansen and Jang Hyun Kim .....	10888
<i>ATM: Adversarial Tuning Multi-agent System Makes a Robust Retrieval-Augmented Generator</i> Junda Zhu, Lingyong Yan, Haibo Shi, Dawei Yin and Lei Sha .....	10902
<i>Dynamic Multi-granularity Attribution Network for Aspect-based Sentiment Analysis</i> Yanjiang Chen, Kai Zhang, Feng Hu, Xianquan Wang, Ruikang li and Qi Liu .....	10920
<i>Unlabeled Debiasing in Downstream Tasks via Class-wise Low Variance Regularization</i> Shahed Masoudian, Markus Frohmann, Navid Rekabsaz and Markus Schedl .....	10932
<i>Large Language Models Know What is Key Visual Entity: An LLM-assisted Multimodal Retrieval for VQA</i> Pu Jian, Donglei Yu and Jiajun Zhang .....	10939
<i>Towards Probing Speech-Specific Risks in Large Multimodal Models: A Taxonomy, Benchmark, and Insights</i> Hao Yang, Lizhen Qu, Ehsan Shareghi and Reza Haf .....	10957
<i>Self-AMPLIFY: Improving Small Language Models with Self Post Hoc Explanations</i> Milan Bhan, Jean-Noël Vittaut, Nicolas Chesneau and Marie-Jeanne Lesot .....	10974
<i>What are the Generator Preferences for End-to-end Task-Oriented Dialog System?</i> Wanshi Xu, Xianwei Zhuang, Zhanpeng Chen, Zhihong Zhu, Xuxin Cheng and Yuexian Zou	10992
<i>Paraphrase Types Elicit Prompt Engineering Capabilities</i> Jan Philip Wahle, Terry Ruas, Yang Xu and Bela Gipp .....	11004
<i>VLEU: a Method for Automatic Evaluation for Generalizability of Text-to-Image Models</i> Jingtao Cao, Zhang Zheng, Hongru Wang and Kam-Fai Wong .....	11034
<i>Towards Online Continuous Sign Language Recognition and Translation</i> Ronglai Zuo, Fangyun Wei and Brian Mak .....	11050
<i>Mitigate Extrinsic Social Bias in Pre-trained Language Models via Continuous Prompts Adjustment</i> Yiwei Dai, Hengrui Gu, Ying Wang and Xin Wang .....	11068
<i>Split and Merge: Aligning Position Biases in LLM-based Evaluators</i> Zongjie Li, Chaozheng Wang, Pingchuan Ma, Daoyuan Wu, Shuai Wang, Cuiyun Gao and Yang Liu .....	11084
<i>Integrating Argumentation and Hate-Speech-based Techniques for Countering Misinformation</i> Sougata Saha and Rohini Srihari .....	11109
<i>BPO: Staying Close to the Behavior LLM Creates Better Online LLM Alignment</i> Wenda Xu, Jiachen Li, William Yang Wang and Lei Li .....	11125

<i>One2Set + Large Language Model: Best Partners for Keyphrase Generation</i>	
Liangying Shao, Liang Zhang, Minlong Peng, Guoqi Ma, Hao Yue, Mingming Sun and Jinsong Su .....	11140
<i>Unlocking Markets: A Multilingual Benchmark to Cross-Market Question Answering</i>	
Yifei Yuan, Yang Deng, Anders Søgaard and Mohammad Aliannejadi .....	11154
<i>ORPO: Monolithic Preference Optimization without Reference Model</i>	
Jiwoo Hong, Noah Lee and James Thorne .....	11170
<i>A Multi-Perspective Analysis of Memorization in Large Language Models</i>	
Bowen Chen, Namgi Han and Yusuke Miyao .....	11190
<i>Do LLMs suffer from Multi-Party Hangover? A Diagnostic Approach to Addressee Recognition and Response Selection in Conversations</i>	
Nicolò Penzo, Maryam Sajedinia, Bruno Lepri, Sara Tonelli and Marco Guerini .....	11210
<i>Code Prompting Elicits Conditional Reasoning Abilities in Text+Code LLMs</i>	
Haritz Puerto, Martin Tutek, Somak Aditya, Xiaodan Zhu and Iryna Gurevych .....	11234
<i>Unveiling the Role of Pretraining in Direct Speech Translation</i>	
Belen Alastruey, Gerard I. Gállego and Marta R. Costa-jussà .....	11259
<i>PCQPR: Proactive Conversational Question Planning with Reflection</i>	
Shasha Guo, Lizi Liao, Jing Zhang, Cuiping Li and Hong Chen .....	11266
<i>CodeAgent: Autonomous Communicative Agents for Code Review</i>	
Xunzhu Tang, Kisub Kim, Yewei Song, Cedric Lothritz, Bei Li, Saad Ezzini, Haoye Tian, Jacques Klein and Tegawendé F. Bissyandé .....	11279
<i>TroL: Traversal of Layers for Large Language and Vision Models</i>	
Byung-Kwan Lee, Sangyun Chung, Chae Won Kim, Beomchan Park and Yong Man Ro ..	11314
<i>MMTE: Corpus and Metrics for Evaluating Machine Translation Quality of Metaphorical Language</i>	
Shun Wang, Ge Zhang, Han Wu, Tyler Loakman, Wenhao Huang and Chenghua Lin .....	11343
<i>Revisiting Supertagging for faster HPSG parsing</i>	
Olga Zamaraeva and Carlos Gómez-Rodríguez .....	11359
<i>Improve Dense Passage Retrieval with Entailment Tuning</i>	
Lu Dai, Hao Liu and Hui Xiong .....	11375
<i>ToolBeHonest: A Multi-level Hallucination Diagnostic Benchmark for Tool-Augmented Large Language Models</i>	
Yuxiang Zhang, Jing Chen, Junjie Wang, Yaxin Liu, Cheng Yang, Chufan Shi, Xinyu Zhu, Zihao Lin, Hanwen Wan, Yujiu Yang, Tetsuya Sakai, Tian Feng and Hayato Yamana .....	11388
<i>TEMA: Token Embeddings Mapping for Enriching Low-Resource Language Models</i>	
Rodolfo Zevallos, Núria Bel and Mireia Farrús .....	11423
<i>DECOR: Improving Coherence in L2 English Writing with a Novel Benchmark for Incoherence Detection, Reasoning, and Rewriting</i>	
Xuanming Zhang, Anthony Diaz, Zixun Chen, Qingyang Wu, Kun Qian, Erik Voss and Zhou Yu	11436
<i>Text2Chart31: Instruction Tuning for Chart Generation with Automatic Feedback</i>	
Fatemeh Pesaran Zadeh, Juyeon Kim, Jin-Hwa Kim and Gunhee Kim .....	11459

<i>PrExMe! Large Scale Prompt Exploration of Open Source LLMs for Machine Translation and Summarization Evaluation</i>	
Christoph Leiter and Steffen Eger .....	11481
<i>Universal Vulnerabilities in Large Language Models: Backdoor Attacks for In-context Learning</i>	
Shuai Zhao, Meihuizi Jia, Anh Tuan Luu, Fengjun Pan and Jinming Wen .....	11507
<i>Repairs in a Block World: A New Benchmark for Handling User Corrections with Multi-Modal Language Models</i>	
Javier Chiyah-Garcia, Alessandro Suglia and Arash Eshghi .....	11523
<i>Beyond the Turn-Based Game: Enabling Real-Time Conversations with Duplex Models</i>	
Xinrong Zhang, Yingfa Chen, Shengding Hu, Xu Han, Zihang Xu, Yuanwei Xu, Weilin Zhao, Maosong Sun and Zhiyuan Liu .....	11543
<i>Strengthening Structural Inductive Biases by Pre-training to Perform Syntactic Transformations</i>	
Matthias Lindemann, Alexander Koller and Ivan Titov .....	11558
<i>Puzzle Solving using Reasoning of Large Language Models: A Survey</i>	
Panagiotis Giadikiaroglou, Maria Lymperaioi, Giorgos Filandrianos and Giorgos Stamou ..	11574
<i>SciEx: Benchmarking Large Language Models on Scientific Exams with Human Expert Grading and Automatic Grading</i>	
Tu Anh Dinh, Carlos Mullov, Leonard Bärmann, Zhaolin Li, Danni Liu, Simon Reiß, Jueun Lee, Nathan Lerzer, Jianfeng Gao, Fabian Peller-Konrad, Tobias Röddiger, Alexander Waibel, Tamim Asfour, Michael Beigl, Rainer Stiefelhagen, Carsten Dachsbacher, Klemens Böhm and Jan Niehues	11592
<i>Red Teaming Language Models for Processing Contradictory Dialogues</i>	
Xiaofei Wen, Bangzheng Li, Tenghao Huang and Muhao Chen .....	11611
<i>Fishing for Magikarp: Automatically Detecting Under-trained Tokens in Large Language Models</i>	
Sander Land and Max Bartolo .....	11631
<i>Reasoning or a Semblance of it? A Diagnostic Study of Transitive Reasoning in LLMs</i>	
Houman Mehrfarin, Arash Eshghi and Ioannis Konstas .....	11647
<i>Pragmatic Norms Are All You Need – Why The Symbol Grounding Problem Does Not Apply to LLMs</i>	
Reto Gubelmann .....	11663
<i>Major Entity Identification: A Generalizable Alternative to Coreference Resolution</i>	
Kawshik Manikantan Sundar, Shubham Toshniwal, Makarand Tapaswi and Vineet Gandhi	11679
<i>Enhancing High-order Interaction Awareness in LLM-based Recommender Model</i>	
Xinfeng Wang, Jin Cui, Fumiyo Fukumoto and Yoshimi Suzuki .....	11696
<i>What Are the Odds? Language Models Are Capable of Probabilistic Reasoning</i>	
Akshay Paruchuri, Jake Garrison, Shun Liao, John B Hernandez, Jacob Sunshine, Tim Althoff, Xin Liu and Daniel McDuff .....	11712
<i>MARE: Multi-Aspect Rationale Extractor on Unsupervised Rationale Extraction</i>	
Han Jiang, Junwen Duan, Zhe Qu and Jianxin Wang .....	11734
<i>LoRA-Guard: Parameter-Efficient Guardrail Adaptation for Content Moderation of Large Language Models</i>	
Hayder Elesedy, Pedro M Esperanca, Silviu Vlad Oprea and Mete Ozay .....	11746

<i>A good pun is its own reword": Can Large Language Models Understand Puns?</i> Zhijun Xu, Siyu Yuan, Lingjie Chen and Deqing Yang .....	11766
<i>QGEval: Benchmarking Multi-dimensional Evaluation for Question Generation</i> Weiping Fu, Bifan Wei, Jianxiang Hu, Zhongmin Cai and Jun Liu .....	11783
<i>Dependency Graph Parsing as Sequence Labeling</i> Ana Ezquerro, David Vilares and Carlos Gómez-Rodríguez .....	11804
<i>NuNER: Entity Recognition Encoder Pre-training via LLM-Annotated Data</i> Sergei Bogdanov, Alexandre Constantin, Timothée Bernard, Benoit Crabbé and Etienne P Bernard	11829
<i>Towards a Greek Proverb Atlas: Computational Spatial Exploration and Attribution of Greek Proverbs</i> John Pavlopoulos, Panos Louridas and Panagiotis Filos .....	11842
<i>Unraveling Babel: Exploring Multilingual Activation Patterns of LLMs and Their Applications</i> Weize Liu, Yinlong Xu, Hongxia Xu, Jintai Chen, Xuming Hu and Jian Wu .....	11855
<i>Advancing Semantic Textual Similarity Modeling: A Regression Framework with Translated ReLU and Smooth K2 Loss</i> Bowen Zhang and Chunping Li .....	11882
<i>Rationalizing Transformer Predictions via End-To-End Differentiable Self-Training</i> Marc Felix Brinner and Sina Zarrieß .....	11894
<i>Segment Any Text: A Universal Approach for Robust, Efficient and Adaptable Sentence Segmentation</i> Markus Frohmann, Igor Sterner, Ivan Vulić, Benjamin Minixhofer and Markus Schedl .....	11908
<i>Applying Contrastive Learning to Code Vulnerability Type Classification</i> Chen Ji, Su Yang, Hongyu Sun and Yuqing Zhang .....	11942
<i>TheoremLlama: Transforming General-Purpose LLMs into Lean4 Experts</i> Ruida Wang, Jipeng Zhang, Yizhen Jia, Rui Pan, Shizhe Diao, Renjie Pi and Tong Zhang .	11953
<i>Multi-Level Cross-Modal Alignment for Speech Relation Extraction</i> Liang Zhang, Zhen Yang, Biao Fu, Ziyao Lu, Liangying Shao, Shiyu Liu, Fandong Meng, Jie Zhou, Xiaoli Wang and Jinsong Su .....	11975
<i>Self-Training for Sample-Efficient Active Learning for Text Classification with Pre-Trained Language Models</i> Christopher Schröder and Gerhard Heyer .....	11987
<i>PANDA: Persona Attributes Navigation for Detecting and Alleviating Overuse Problem in Large Language Models</i> Jinsung Kim, Seonmin Koo and Heuseok Lim .....	12005
<i>The Multilingual Alignment Prism: Aligning Global and Local Preferences to Reduce Harm</i> Aakanksha , Arash Ahmadian, Beyza Ermis, Seraphina Goldfarb-Tarrant, Julia Kreutzer, Marzieh Fadaee and Sara Hooker .....	12027
<i>Subword Segmentation in LLMs: Looking at Inflection and Consistency</i> Marion Di Marco and Alexander Fraser .....	12050
<i>Explicit, Implicit, and Scattered: Revisiting Event Extraction to Capture Complex Arguments</i> Omar Sharif, Joseph Gatto, Madhusudan Basak and Sarah Masud Preum .....	12061

<i>Let Me Teach You: Pedagogical Foundations of Feedback for Language Models</i> Beatriz Borges, Niket Tandon, Tanja Käser and Antoine Bosselut . . . . .	12082
<i>Unknown Claims: Generation of Fact-Checking Training Examples from Unstructured and Structured Data</i> Jean-Flavien Bussotti, Luca Ragazzi, Giacomo Frisoni, Gianluca Moro and Paolo Papotti .	12105
<i>TL-CL: Task And Language Incremental Continual Learning</i> Shrey Satapara and P. K. Srijith . . . . .	12123
<i>Medical Adaptation of Large Language and Vision-Language Models: Are We Making Progress?</i> Daniel P Jeong, Saurabh Garg, Zachary Chase Lipton and Michael Oberst . . . . .	12143
<i>Empowering Multi-step Reasoning across Languages via Program-Aided Language Models</i> Leonardo Ranaldi, Giulia Pucci, Barry Haddow and Alexandra Birch . . . . .	12171
<i>Do LLMs Overcome Shortcut Learning? An Evaluation of Shortcut Challenges in Large Language Models</i> Yu Yuan, Lili Zhao, Kai Zhang, Guangting Zheng and Qi Liu . . . . .	12188
<i>ControlMath: Controllable Data Generation Promotes Math Generalist Models</i> Nuo Chen, Ning Wu, Jianhui Chang, Linjun Shou and Jia Li . . . . .	12201
<i>Where Am I From? Identifying Origin of LLM-generated Content</i> Liyang LI, Yihan Bai and Minhao Cheng . . . . .	12218
<i>ReadMe++: Benchmarking Multilingual Language Models for Multi-Domain Readability Assessment</i> Tarek Naous, Michael J Ryan, Anton Lavrouk, Mohit Chandra and Wei Xu . . . . .	12230
<i>GlossLM: A Massively Multilingual Corpus and Pretrained Model for Interlinear Glossed Text</i> Michael Ginn, Lindia Tjuatja, Taiqi He, Enora Rice, Graham Neubig, Alexis Palmer and Lori Levin . . . . .	12267
<i>GDTB: Genre Diverse Data for English Shallow Discourse Parsing across Modalities, Text Types, and Domains</i> Yang Janet Liu, Tatsuya Aoyama, Wesley Scivetti, Yilun Zhu, Shabnam Behzad, Lauren Elizabeth Levine, Jessica Lin, Devika Tiwari and Amir Zeldes . . . . .	12287
<i>RA2FD: Distilling Faithfulness into Efficient Dialogue Systems</i> Zhiyuan Zhu, Yusheng Liao, Chenxin Xu, Yunfeng Guan, Yanfeng Wang and Yu Wang . . . . .	12304
<i>Subjective Topic meets LLMs: Unleashing Comprehensive, Reflective and Creative Thinking through the Negation of Negation</i> Fangrui Lv, Kaixiong Gong, Jian Liang, Xinyu Pang and Changshui Zhang . . . . .	12318
<i>Experimental Contexts Can Facilitate Robust Semantic Property Inference in Language Models, but Inconsistently</i> Kanishka Misra, Allyson Ettinger and Kyle Mahowald . . . . .	12342
<i>Leveraging Estimated Transferability Over Human Intuition for Model Selection in Text Ranking</i> Jun Bai, Zhuofan Chen, Zhenzi Li, Hanhua Hong, Jianfei Zhang, Chen Li, Chenghua Lin and Wenge Rong . . . . .	12356
<i>Unveiling In-Context Learning: A Coordinate System to Understand Its Working Mechanism</i> Anhao Zhao, Fanghua Ye, Jinlan Fu and Xiaoyu Shen . . . . .	12375
<i>Self-Powered LLM Modality Expansion for Large Speech-Text Models</i> Tengfei Yu, Xuebo Liu, Zhiyi Hou, Liang Ding, Dacheng Tao and Min Zhang . . . . .	12401

<i>ABSEval: An Agent-based Framework for Script Evaluation</i>	
Sirui Liang, Baoli Zhang, Jun Zhao and Kang Liu .....	12418
<i>Latent Concept-based Explanation of NLP Models</i>	
Xuemin Yu, Fahim Dalvi, Nadir Durrani, Marzia Nouri and Hassan Sajjad .....	12435
<i>Decoding with Limited Teacher Supervision Requires Understanding When to Trust the Teacher</i>	
Hyunjong Ok, Jegwang Ryu and Jaeho Lee .....	12460
<i>Enhancing Data Quality through Simple De-duplication: Navigating Responsible Computational Social Science Research</i>	
Yida Mu, Mali Jin, Xingyi Song and Nikolaos Aletras .....	12477
<i>The Mystery of the Pathological Path-star Task for Language Models</i>	
Arvid Frydenlund .....	12493
<i>Voices in a Crowd: Searching for clusters of unique perspectives</i>	
Nikolas Vitsakis, Amit Parekh and Ioannis Konstas .....	12517
<i>Neeko: Leveraging Dynamic LoRA for Efficient Multi-Character Role-Playing Agent</i>	
Xiaoyan Yu, Tongxu Luo, Yifan Wei, Fangyu Lei, Yiming Huang, Hao Peng and Liehuang Zhu	12540
<i>SLANG: New Concept Comprehension of Large Language Models</i>	
Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi and Xueqi Cheng .....	12558
<i>Towards Interpretable Sequence Continuation: Analyzing Shared Circuits in Large Language Models</i>	
Michael Lan, Philip Torr and Fazl Barez .....	12576
<i>Why Does New Knowledge Create Messy Ripple Effects in LLMs?</i>	
Jiaxin Qin, Zixuan Zhang, Chi Han, Pengfei Yu, Manling Li and Heng Ji .....	12602
<i>Lifelong Event Detection via Optimal Transport</i>	
Viet Dao, Van-Cuong Pham, Quyen Tran, Thanh-Thien Le, Linh Van Ngo and Thien Huu Nguyen	12610
<i>SUPER: Evaluating Agents on Setting Up and Executing Tasks from Research Repositories</i>	
Ben Bogin, Kejuan Yang, Shashank Gupta, Kyle Richardson, Erin Bransom, Peter Clark, Ashish Sabharwal and Tushar Khot .....	12622
<i>FIRST: Teach A Reliable Large Language Model Through Efficient Trustworthy Distillation</i>	
KaShun Shum, Minrui Xu, Jianshu Zhang, Zixin Chen, Shizhe Diao, Hanze Dong, Jipeng Zhang and Muhammad Omer Raza .....	12646
<i>Domain adapted machine translation: What does catastrophic forgetting forget and why?</i>	
Danielle Saunders and Steve DeNeefe .....	12660
<i>Enhancing AI Assisted Writing with One-Shot Implicit Negative Feedback</i>	
Benjamin Towle and Ke Zhou .....	12672
<i>Atomic Self-Consistency for Better Long Form Generations</i>	
Raghuveer Thirukovalluru, Yukun Huang and Bhuwan Dhingra .....	12681
<i>“Global is Good, Local is Bad?”: Understanding Brand Bias in LLMs</i>	
Mahammed Kamruzzaman, Hieu Minh Nguyen and Gene Louis Kim .....	12695

<i>Optimizing Rare Word Accuracy in Direct Speech Translation with a Retrieval-and-Demonstration Approach</i>	
Siqi Li, Danni Liu and Jan Niehues	12703
<i>ACE: A LLM-based Negotiation Coaching System</i>	
Ryan Shea, Aymen Kallala, Xin Lucy Liu, Michael W. Morris and Zhou Yu	12720
<i>TransferTOD: A Generalizable Chinese Multi-Domain Task-Oriented Dialogue System with Transfer Capabilities</i>	
Ming Zhang, Caishuang Huang, Yilong Wu, Shichun Liu, Huiyuan Zheng, Yurui Dong, Yujiong Shen, Shihan Dou, Jun Zhao, Junjie Ye, Qi Zhang, Tao Gui and Xuanjing Huang	12750
<i>PATIENT-<math>\psi</math>: Using Large Language Models to Simulate Patients for Training Mental Health Professionals</i>	
Ruiyi Wang, Stephanie Milani, Jamie C. Chiu, Jiayin Zhi, Shaun M. Eack, Travis Labrum, Samuel M Murphy, Nev Jones, Kate V Hardy, Hong Shen, Fei Fang and Zhiyu Chen	12772
<i>DKEC: Domain Knowledge Enhanced Multi-Label Classification for Diagnosis Prediction</i>	
Xueren Ge, Abhishek Satpathy, Ronald Dean Williams, John Stankovic and Homa Alemzadeh	12798
<i>ModSCAN: Measuring Stereotypical Bias in Large Vision-Language Models from Vision and Language Modalities</i>	
Yukun Jiang, Zheng Li, Xinyue Shen, Yugeng Liu, Michael Backes and Yang Zhang	12814
<i>Large Language Models Can Self-Correct with Key Condition Verification</i>	
Zhenyu Wu, Qingkai Zeng, Zhihan Zhang, Zhaoxuan Tan, Chao Shen and Meng Jiang	12846
<i>Learning to Write Rationally: How Information Is Distributed in Non-native Speakers' Essays</i>	
Zixin Tang and Janet Van Hell	12868
<i>Defending Against Social Engineering Attacks in the Age of LLMs</i>	
Lin Ai, Tharindu Sandaruwan Kumarage, Amrita Bhattacharjee, Zizhou Liu, Zheng Hui, Michael S. Davinroy, James Cook, Laura Cassani, Kirill Trapeznikov, Matthias Kirchner, Arslan Basharat, Anthony Hoogs, Joshua Garland, Huan Liu and Julia Hirschberg	12880
<i>Heterogeneous LoRA for Federated Fine-tuning of On-Device Foundation Models</i>	
Yae Jee Cho, Luyang Liu, Zheng Xu, Aldi Fahrezi and Gauri Joshi	12903
<i>Make Some Noise: Unlocking Language Model Parallel Inference Capability through Noisy Training</i>	
Yixuan Wang, Xianzhen Luo, Fuxuan Wei, Yijun Liu, Qingfu Zhu, Xuanyu Zhang, Qing Yang, Dongliang Xu and Wanxiang Che	12914
<i>Target-Aware Language Modeling via Granular Data Sampling</i>	
Ernie Chang, Pin-Jie Lin, Yang Li, Changsheng Zhao, Daeil Kim, Rastislav Rabatin, Zechun Liu, Yangyang Shi and Vikas Chandra	12927
<i>SPEED++: A Multilingual Event Extraction Framework for Epidemic Prediction and Preparedness</i>	
Tanmay Parekh, Jeffrey Kwan, Jiarui Yu, Sparsh Johri, Hyosang Ahn, Sreya Muppalla, Kai-Wei Chang, Wei Wang and Nanyun Peng	12936
<i>CoGen: Learning from Feedback with Coupled Comprehension and Generation</i>	
Mustafa Omer Gul and Yoav Artzi	12966
<i>UNICORN: A Unified Causal Video-Oriented Language-Modeling Framework for Temporal Video-Language Tasks</i>	
Yuanhao Xiong, Yixin Nie, Haotian Liu, Boxin Wang, Jun Chen, Rong Jin, Cho-Jui Hsieh, Lorenzo Torresani and Jie Lei	12983

<i>Story Morals: Surfacing value-driven narrative schemas using large language models</i> David G Hobson, Haiqi Zhou, Derek Ruths and Andrew Piper .....	12998
<i>OATH-Frames: Characterizing Online Attitudes Towards Homelessness with LLM Assistants</i> Jaspreet Ranjit, Brihi Joshi, Rebecca Dorn, Laura Petry, Olga Koumoundouros, Jayne Bottarini, Peichen Liu, Eric Rice and Swabha Swayamdipta .....	13033
<i>AnaloBench: Benchmarking the Identification of Abstract and Long-context Analogies</i> Xiao Ye, Andrew Wang, Jacob Choi, Yining Lu, Shreya Sharma, Lingfeng Shen, Vijay Murari Tiyiyala, Nicholas Andrews and Daniel Khashabi .....	13060
<i>SciER: An Entity and Relation Extraction Dataset for Datasets, Methods, and Tasks in Scientific Documents</i> Qi Zhang, Zhijia Chen, Huitong Pan, Cornelia Caragea, Longin Jan Latecki and Eduard Dragut	13083
<i>Analysis of Plan-based Retrieval for Grounded Text Generation</i> Ameya Godbole, Nicholas Monath, Seungyeon Kim, Ankit Singh Rawat, Andrew McCallum and Manzil Zaheer .....	13101
<i>Detecting Errors through Ensembling Prompts (DEEP): An End-to-End LLM Framework for Detecting Factual Errors</i> Alex Chandler, Devesh Surve and Hui Su .....	13120
<i>RLHF Can Speak Many Languages: Unlocking Multilingual Preference Optimization for LLMs</i> John Dang, Arash Ahmadian, Kelly Marchisio, Julia Kreutzer, Ahmet Üstün and Sara Hooker	13134
<i>Boosting Logical Fallacy Reasoning in LLMs via Logical Structure Tree</i> Yuanyuan Lei and Ruihong Huang .....	13157
<i>Chain and Causal Attention for Efficient Entity Tracking</i> Erwan Fagnou, Paul Caillon, Blaise Delattre and Alexandre Allauzen .....	13174
<i>BEEAR: Embedding-based Adversarial Removal of Safety Backdoors in Instruction-tuned Language Models</i> Yi Zeng, Weiyu Sun, Tran Huynh, Dawn Song, Bo Li and Ruoxi Jia .....	13189
<i>A Bayesian Approach to Harnessing the Power of LLMs in Authorship Attribution</i> Zhengmian Hu, Tong Zheng and Heng Huang .....	13216
<i>FAC<sup>2</sup>E: Better Understanding Large Language Model Capabilities by Dissociating Language and Cognition</i> Xiaoqiang Wang, Lingfei Wu, Tengfei Ma and Bang Liu .....	13228
<i>OpenSep: Leveraging Large Language Models with Textual Inversion for Open World Audio Separation</i> Tanvir Mahmud and Diana Marculescu .....	13244
<i>Language Concept Erasure for Language-invariant Dense Retrieval</i> Zhiqi Huang, Puxuan Yu, Shauli Ravfogel and James Allan .....	13261
<i>Learning Personalized Alignment for Evaluating Open-ended Text Generation</i> Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li and Yuandong Tian .....	13274
<i>Large Language Models Are Involuntary Truth-Tellers: Exploiting Fallacy Failure for Jailbreak Attacks</i> Yue Zhou, Henry Peng Zou, Barbara Di Eugenio and Yang Zhang .....	13293



<i>Turn Waste into Worth: Rectifying Top-k Router of MoE</i>	
Zhiyuan Zeng, Qipeng Guo, Zhaoye Fei, Zhangyue Yin, Yunhua Zhou, Linyang Li, Tianxiang Sun, Hang Yan, Dahua Lin and Xipeng Qiu . . . . .	13305
<i>Null-Shot Prompting: Rethinking Prompting Large Language Models With Hallucination</i>	
Pittawat Taveekitworachai, Febri Abdullah and Ruck Thawonmas . . . . .	13321
<i>CommVQA: Situating Visual Question Answering in Communicative Contexts</i>	
Nandita Shankar Naik, Christopher Potts and Elisa Kreiss . . . . .	13362
<i>Ouroboros: Generating Longer Drafts Phrase by Phrase for Faster Speculative Decoding</i>	
Weilin Zhao, Yuxiang Huang, Xu Han, Wang Xu, Chaojun Xiao, Xinrong Zhang, Yewei Fang, Kaihuo Zhang, Zhiyuan Liu and Maosong Sun . . . . .	13378
<i>1+1&gt;2: Can Large Language Models Serve as Cross-Lingual Knowledge Aggregators?</i>	
Yue Huang, Chenrui Fan, Yuan Li, Siyuan Wu, Tianyi Zhou, Xiangliang Zhang and Lichao Sun	13394
<i>How to Leverage Demonstration Data in Alignment for Large Language Model? A Self-Imitation Learning Perspective</i>	
Teng Xiao, Mingxiao Li, Yige Yuan, Huaisheng Zhu, Chao Cui and Vasant G Honavar . . .	13413
<i>Style-Specific Neurons for Steering LLMs in Text Style Transfer</i>	
Wen Lai, Viktor Hangya and Alexander Fraser . . . . .	13427
<i>Adaptive Query Rewriting: Aligning Rewriters through Marginal Probability of Conversational Answers</i>	
Tianhua Zhang, Kun LI, Hongyin Luo, Xixin Wu, James R. Glass and Helen M. Meng . . .	13444
<i>Grasping the Essentials: Tailoring Large Language Models for Zero-Shot Relation Extraction</i>	
Sizhe Zhou, Yu Meng, Bowen Jin and Jiawei Han . . . . .	13462
<i>DA-Code: Agent Data Science Code Generation Benchmark for Large Language Models</i>	
Yiming Huang, Jianwen Luo, Yan Yu, Yitong Zhang, Fangyu Lei, Yifan Wei, Shizhu He, Lifu Huang, Xiao Liu, Jun Zhao and Kang Liu . . . . .	13487
<i>Leveraging Context-Aware Prompting for Commit Message Generation</i>	
Zhihua Jiang, Jianwei Chen, Dongning Rao and Guanghui Ye . . . . .	13522
<i>Linguistic Bias in ChatGPT: Language Models Reinforce Dialect Discrimination</i>	
Eve Fleisig, Genevieve Smith, Madeline Bossi, Ishita Rustagi, Xavier Yin and Dan Klein .	13541
<i>Lifelong Knowledge Editing for LLMs with Retrieval-Augmented Continuous Prompt Learning</i>	
Qizhou Chen, Taolin Zhang, Xiaofeng He, Dongyang Li, Chengyu Wang, Longtao Huang and Hui Xue' . . . . .	13565
<i>A Learning Rate Path Switching Training Paradigm for Version Updates of Large Language Models</i>	
Zhihao Wang, Shiyu Liu, Jianheng Huang, Wang Zheng, YiXuan Liao, Xiaoxin Chen, Junfeng Yao and Jinsong Su . . . . .	13581
<i>Zero-Shot Cross-Lingual NER Using Phonemic Representations for Low-Resource Languages</i>	
Jimin Sohn, Haeji Jung, Alex Cheng, Joeon Kang, Yilin Du and David R Mortensen . . .	13595
<i>An Analysis and Mitigation of the Reversal Curse</i>	
Ang Lv, Kaiyi Zhang, Shufang Xie, Quan Tu, Yuhan Chen, Ji-Rong Wen and Rui Yan . . .	13603
<i>Exploring the Practicality of Generative Retrieval on Dynamic Corpora</i>	
Chaeun Kim, Soyoung Yoon, Hyunji Lee, Joel Jang, Sohee Yang and Minjoon Seo . . . . .	13616

<i>OneNet: A Fine-Tuning Free Framework for Few-Shot Entity Linking via Large Language Model Prompting</i>	Xukai Liu, Ye Liu, Kai Zhang, Kehang Wang, Qi Liu and Enhong Chen . . . . .	13634
<i>Don't Just Say I don't know"! Self-aligning Large Language Models for Responding to Unknown Questions with Explanations</i>	Yang Deng, Yong Zhao, Moxin Li, See-Kiong Ng and Tat-Seng Chua . . . . .	13652
<i>Fewer is More: Boosting Math Reasoning with Reinforced Context Pruning</i>	Xijie Huang, Li Lyna Zhang, Kwang-Ting Cheng, Fan Yang and Mao Yang . . . . .	13674
<i>Large Language Models Are Poor Clinical Decision-Makers: A Comprehensive Benchmark</i>	Fenglin Liu, Zheng Li, Hongjian Zhou, Qingyu Yin, Jingfeng Yang, Xianfeng Tang, Chen Luo, Ming Zeng, Haoming Jiang, Yifan Gao, Priyanka Nigam, Sreyashi Nag, Bing Yin, Yining Hua, Xuan Zhou, Omid Rohanian, Anshul Thakur, Lei Clifton and David A. Clifton . . . . .	13696
<i>Holistic Automated Red Teaming for Large Language Models through Top-Down Test Case Generation and Multi-turn Interaction</i>	Jinchuan Zhang, Yan Zhou, Yaxin Liu, Ziming Li and Songlin Hu . . . . .	13711
<i>Householder Pseudo-Rotation: A Novel Approach to Activation Editing in LLMs with Direction-Magnitude Perspective</i>	Van-Cuong Pham and Thien Huu Nguyen . . . . .	13737
<i>DynamicER: Resolving Emerging Mentions to Dynamic Entities for RAG</i>	Jinyoung Kim, Dayoon Ko and Gunhee Kim . . . . .	13752
<i>Preserving Generalization of Language models in Few-shot Continual Relation Extraction</i>	Quyen Tran, Nguyen Xuan Thanh, Nguyen Hoang Anh, Nam Le Hai, Trung Le, Linh Van Ngo and Thien Huu Nguyen . . . . .	13771
<i>A Systematic Survey and Critical Review on Evaluating Large Language Models: Challenges, Limitations, and Recommendations</i>	Md Tahmid Rahman Laskar, Sawsan Alqahtani, M Saiful Bari, Mizanur Rahman, Mohammad Abdullah Matin Khan, Haidar Khan, Israt Jahan, Amran Bhuiyan, Chee Wei Tan, Md Rizwan Parvez, Enamul Hoque, Shafiq Joty and Jimmy Huang . . . . .	13785
<i>Consecutive Batch Model Editing with HooK Layers</i>	Shuaiyi Li, Yang Deng, Deng Cai, Hongyuan Lu, Liang Chen and Wai Lam . . . . .	13817
<i>Topic-Oriented Open Relation Extraction with A Priori Seed Generation</i>	Linyi Ding, Jinfeng Xiao, Sizhe Zhou, Chaoqi Yang and Jiawei Han . . . . .	13834
<i>Related Work and Citation Text Generation: A Survey</i>	Xiangci Li and Jessica Ouyang . . . . .	13846
<i>Curriculum Consistency Learning for Conditional Sentence Generation</i>	Liangxin Liu, Xuebo Liu, Lian Lian, Shengjun Cheng, Jun Rao, Tengfei Yu, Hexuan Deng and Min Zhang . . . . .	13865
<i>A Systematic Analysis of Large Language Models as Soft Reasoners: The Case of Syllogistic Inferences</i>	Leonardo Bertolazzi, Albert Gatt and Raffaella Bernardi . . . . .	13882
<i>Pre-training Cross-lingual Open Domain Question Answering with Large-scale Synthetic Supervision</i>	Fan Jiang, Tom Drummond and Trevor Cohn . . . . .	13906

<i>MOSEL: 950,000 Hours of Speech Data for Open-Source Speech Foundation Model Training on EU Languages</i>	
Marco Gaido, Sara Papi, Luisa Bentivogli, Alessio Brutti, Mauro Cettolo, Roberto Gretter, Marco Matassoni, Mohamed Nabih and Matteo Negri . . . . .	13934
<i>Improving Knowledge Graph Completion with Structure-Aware Supervised Contrastive Learning</i>	
Jiashi Lin, Lifang Wang, Xinyu Lu, Zhongtian Hu, Wei Zhang and Wenxuan Lu . . . . .	13948
<i>Contribution of Linguistic Typology to Universal Dependency Parsing: An Empirical Investigation</i>	
Ali Basirat and Navid Baradaran Hemmati . . . . .	13960
<i>TRoTR: A Framework for Evaluating the Re-contextualization of Text Reuse</i>	
Francesco Periti, Pierluigi Cassotti, Stefano Montanelli, Nina Tahmasebi and Dominik Schlechtweg . . . . .	13972
<i>Structured Optimal Brain Pruning for Large Language Models</i>	
Jiateng Wei, Quan Lu, Ning Jiang, Siqi Li, Jingyang Xiang, Jun Chen and Yong Liu . . . . .	13991
<i>Automatically Generated Definitions and their utility for Modeling Word Meaning</i>	
Francesco Periti, David Alfter and Nina Tahmasebi . . . . .	14008
<i>How Do Your Code LLMs perform? Empowering Code Instruction Tuning with Really Good Data</i>	
Yejie Wang, Keqing He, Dayuan Fu, Zhuoma GongQue, Heyang Xu, Yanxu Chen, Zhexu Wang, Yujia Fu, Guanting Dong, Muxi Diao, Jingang Wang, Mengdi Zhang, Xunliang Cai and Weiran Xu	14027
<i>MAIR: A Massive Benchmark for Evaluating Instructed Retrieval</i>	
Weiwei Sun, Zhengliang Shi, Wu Jiu Long, Lingyong Yan, Xinyu Ma, Yiding Liu, Min Cao, Dawei Yin and Zhaochun Ren . . . . .	14044
<i>Rethinking the Evaluation of In-Context Learning for LLMs</i>	
Guoxin Yu, Lema Liu, Mo Yu, Yue Yu and Xiang Ao . . . . .	14068
<i>Cluster-Norm for Unsupervised Probing of Knowledge</i>	
Walter Laurito, Sharan Maiya, Grégoire Dhimoïla, Owen Ho Wan Yeung and Kaarel Hänni	14083
<i>Hopping Too Late: Exploring the Limitations of Large Language Models on Multi-Hop Queries</i>	
Eden Biran, Daniela Gottesman, Sohee Yang, Mor Geva and Amir Globerson . . . . .	14113
<i>Enhancing Training Data Attribution for Large Language Models with Fitting Error Consideration</i>	
Kangxi Wu, Liang Pang, Huawei Shen and Xueqi Cheng . . . . .	14131
<i>Where am I? Large Language Models Wandering between Semantics and Structures in Long Contexts</i>	
Seonmin Koo, Jinsung Kim, YoungJoon Jang, Chanjun Park and Heuseok Lim . . . . .	14144
<i>KARL: Knowledge-Aware Retrieval and Representations aid Retention and Learning in Students</i>	
Matthew Shu, Nishant Balepur, Shi Feng and Jordan Lee Boyd-Graber . . . . .	14161
<i>Large Language Models Can Be Contextual Privacy Protection Learners</i>	
Yijia Xiao, Yiqiao Jin, Yushi Bai, Yue Wu, Xianjun Yang, Xiao Luo, Wenchao Yu, Xujiang Zhao, Yanchi Liu, Quanquan Gu, Haifeng Chen, Wei Wang and Wei Cheng . . . . .	14179
<i>A SMART Mnemonic Sounds like Glue Tonic": Mixing LLMs with Student Feedback to Make Mnemonic Learning Stick</i>	
Nishant Balepur, Matthew Shu, Alexander Hoyle, Alison Robey, Shi Feng, Seraphina Goldfarb-Tarrant and Jordan Lee Boyd-Graber . . . . .	14202

<i>Mixture-of-Skills: Learning to Optimize Data Usage for Fine-Tuning Large Language Models</i> Minghao Wu, Thuy-Trang Vu, Lizhen Qu and Reza Haf .....	14226
<i>MolTRES: Improving Chemical Language Representation Learning for Molecular Property Prediction</i> Jun-Hyung Park, Yeachan Kim, Mingyu Lee, Hyuntae Park and SangKeun Lee .....	14241
<i>First Heuristic Then Rational: Dynamic Use of Heuristics in Language Model Reasoning</i> Yoichi Aoki, Keito Kudo, Tatsuki Kuribayashi, Shusaku Sone, Masaya Taniguchi, Keisuke Sakaguchi and Kentaro Inui .....	14255
<i>Tools Fail: Detecting Silent Errors in Faulty Tools</i> Jimin Sun, So Yeon Min, Yingshan Chang and Yonatan Bisk .....	14272
<i>Pcc-tuning: Breaking the Contrastive Learning Ceiling in Semantic Textual Similarity</i> Bowen Zhang and Chunping Li .....	14290
<i>Cross-lingual Back-Parsing: Utterance Synthesis from Meaning Representation for Zero-Resource Semantic Parsing</i> Deokhyung Kang, Seonjeong Hwang, Yunsu Kim and Gary Lee .....	14303
<i>Shaking Up VLMs: Comparing Transformers and Structured State Space Models for Vision &amp; Language Modeling</i> Georgios Pantazopoulos, Malvina Nikandrou, Alessandro Suglia, Oliver Lemon and Arash Eshghi	14318
<i>Are LLMs Good Zero-Shot Fallacy Classifiers?</i> Fengjun Pan, Xiaobao Wu, Zongrui Li and Anh Tuan Luu .....	14338
<i>The Mystery of In-Context Learning: A Comprehensive Survey on Interpretation and Analysis</i> Yuxiang Zhou, Jiazheng Li, Yanzheng Xiang, Hanqi Yan, Lin Gui and Yulan He .....	14365
<i>More DWUGs: Extending and Evaluating Word Usage Graph Datasets in Multiple Languages</i> Dominik Schlechtweg, Pierluigi Cassotti, Bill Noble, David Alfter, Sabine Schulte Im Walde and Nina Tahmasebi .....	14379
<i>Vision-Language Model Fine-Tuning via Simple Parameter-Efficient Modification</i> Ming Li, Jike Zhong, Chenxin Li, Liuzhuozheng Li, Nie Lin and Masashi Sugiyama .....	14394
<i>ECIS-VQG: Generation of Entity-centric Information-seeking Questions from Videos</i> Arpan Phukan, Manish Gupta and Asif Ekbal .....	14411
<i>Distractor Generation in Multiple-Choice Tasks: A Survey of Methods, Datasets, and Evaluation</i> Elaf Alhazmi, Quan Z. Sheng, Wei Emma Zhang, Munazza Zaib and Ahoud Alhazmi .....	14437
<i>Evaluating n-Gram Novelty of Language Models Using Rusty-DAWG</i> William Merrill, Noah A. Smith and Yanai Elazar .....	14459
<i>ASL STEM Wiki: Dataset and Benchmark for Interpreting STEM Articles</i> Kayo Yin, Chinmay Singh, Fyodor O Minakov, Vanessa Milan, Hal Daumé Iii, Cyril Zhang, Alex Xijie Lu and Danielle Bragg .....	14474
<i>Can Automatic Metrics Assess High-Quality Translations?</i> Sweta Agrawal, António Farinhas, Ricardo Rei and Andre Martins .....	14491
<i>Modeling User Preferences with Automatic Metrics: Creating a High-Quality Preference Dataset for Machine Translation</i> Sweta Agrawal, José G. C. De Souza, Ricardo Rei, António Farinhas, Gonçalo Faria, Patrick Fernandes, Nuno M Guerreiro and Andre Martins .....	14503

<i>DC-Instruct: An Effective Framework for Generative Multi-intent Spoken Language Understanding</i> Bowen Xing, Lizi Liao, Minlie Huang and Ivor Tsang .....	14520
<i>KnowTuning: Knowledge-aware Fine-tuning for Large Language Models</i> Yougang Lyu, Lingyong Yan, Shuaiqiang Wang, Haibo Shi, Dawei Yin, Pengjie Ren, Zhumin Chen, Maarten de Rijke and Zhaochun Ren .....	14535
<i>SecCoder: Towards Generalizable and Robust Secure Code Generation</i> Boyu Zhang, Tianyu Du, Junkai Tong, Xuhong Zhang, Kingsum Chow, Sheng Cheng, Xun Wang and Jianwei Yin .....	14557
<i>Nash CoT: Multi-Path Inference with Preference Equilibrium</i> Ziqi Zhang, Cunxiang Wang, Xiao Xiong, Yue Zhang and Donglin Wang .....	14572
<i>Scalable Efficient Training of Large Language Models with Low-dimensional Projected Attention</i> Xingtai Lv, Ning Ding, Kaiyan Zhang, Ermo Hua, Ganqu Cui and Bowen Zhou .....	14588
<i>Small Agent Can Also Rock! Empowering Small Language Models as Hallucination Detector</i> Xiaoxue Cheng, Junyi Li, Xin Zhao, Hongzhi Zhang, Fuzheng Zhang, Di Zhang, Kun Gai and Ji-Rong Wen .....	14600
<i>Interpretable Composition Attribution Enhancement for Visio-linguistic Compositional Understanding</i> Wei Li, Zhen Huang, Xinmei Tian, Le Lu, Houqiang Li, Xu Shen and Jieping Ye .....	14616
<i>LLM Task Interference: An Initial Study on the Impact of Task-Switch in Conversational History</i> Akash Gupta, Ivaxi Sheth, Vyas Raina, Mark Gales and Mario Fritz .....	14633
<i>Social Bias Probing: Fairness Benchmarking for Language Models</i> Marta Marchiori Manerba, Karolina Stanczak, Riccardo Guidotti and Isabelle Augenstein ..	14653
<i>Chain-of-Note: Enhancing Robustness in Retrieval-Augmented Language Models</i> Wenhao Yu, Hongming Zhang, Xiaoman Pan, Peixin Cao, Kaixin Ma, Jian Li, Hongwei Wang and Dong Yu .....	14672
<i>DynaThink: Fast or Slow? A Dynamic Decision-Making Framework for Large Language Models</i> Jiabao Pan, Yan Zhang, Chen Zhang, Zuozhu Liu, Hongwei Wang and Haizhou Li .....	14686
<i>Revisiting Automated Evaluation for Long-form Table Question Answering</i> Yuqi Wang, Lyuhao Chen, Songcheng Cai, Zhijian Xu and Yilun Zhao .....	14696
<i>Weak Reward Model Transforms Generative Models into Robust Causal Event Extraction Systems</i> Italo Luis Da Silva, Hanqi Yan, Lin Gui and Yulan He .....	14707
<i>Learn Beyond The Answer: Training Language Models with Reflection for Mathematical Reasoning</i> Zhihan Zhang, Tao Ge, Zhenwen Liang, Wenhao Yu, Dian Yu, Mengzhao Jia, Dong Yu and Meng Jiang .....	14720
<i>FinDVer: Explainable Claim Verification over Long and Hybrid-content Financial Documents</i> Yilun Zhao, Yitao Long, Tintin Jiang, Chengye Wang, Weiyuan Chen, Hongjun Liu, Xiangru Tang, Yiming Zhang, Chen Zhao and Arman Cohan .....	14739
<i>Extracting Prompts by Inverting LLM Outputs</i> Collin Zhang, John Xavier Morris and Vitaly Shmatikov .....	14753
<i>BiasAlert: A Plug-and-play Tool for Social Bias Detection in LLMs</i> Zhiting Fan, Ruizhe Chen, Ruiling Xu and Zuozhu Liu .....	14778

<i>VHASR: A Multimodal Speech Recognition System With Vision Hotwords</i>	
Jiliang Hu, Zuchao Li, Ping Wang, Haojun Ai, Lefei Zhang and Hai Zhao .....	14791
<i>A Probability–Quality Trade-off in Aligned Language Models and its Relation to Sampling Adaptors</i>	
Naaman Tan, Josef Valvoda, Tianyu Liu, Anej Svete, Yanxia Qin, Min-Yen Kan and Ryan Cotterell	14805
<i>Bridging Local Details and Global Context in Text-Attributed Graphs</i>	
Yaoke Wang, Yun Zhu, Wenqiao Zhang, Yueting Zhuang, Liyunfei Liyunfei and Siliang Tang	14830
<i>Building Resources for Emakhuwa: Machine Translation and News Classification Benchmarks</i>	
Felermimo D. M. A. Ali, Henrique Lopes Cardoso and Rui Sousa-Silva.....	14842
<i>RepMatch: Quantifying Cross-Instance Similarities in Representation Space</i>	
Mohammad Reza Modarres, Sina Abbasi and Mohammad Taher Pilehvar .....	14858
<i>Commonsense Knowledge Editing Based on Free-Text in LLMs</i>	
Xiusheng Huang, Yequan Wang, Jun Zhao and Kang Liu .....	14870
<i>A Closer Look at Multidimensional Online Political Incivility</i>	
Sagi Pendzel, Nir Lotan, Alon Zoizner and Einat Minkov .....	14881
<i>Leveraging BERT and TFIDF Features for Short Text Clustering via Alignment-Promoting Co-Training</i>	
Zetong Li, Qinliang Su, Shijing Si and Jianxing Yu.....	14897
<i>Applying Intrinsic Debiasing on Downstream Tasks: Challenges and Considerations for Machine Translation</i>	
Bar Iluz, Yanai Elazar, Asaf Yehudai and Gabriel Stanovsky .....	14914
<i>Unsupervised Named Entity Disambiguation for Low Resource Domains</i>	
Debarghya Datta and Soumajit Pramanik.....	14922
<i>SparseGrad: A Selective Method for Efficient Fine-tuning of MLP Layers</i>	
Viktoriia A. Chekalina, Anna Rudenko, Gleb Mezentsev, Aleksandr Mikhalev, Alexander Panchenko and Ivan Oseledets .....	14929
<i>MoCoKGC: Momentum Contrast Entity Encoding for Knowledge Graph Completion</i>	
Qingyang Li, Yanru Zhong and Yuchu Qin .....	14940
<i>ActPlan-1K: Benchmarking the Procedural Planning Ability of Visual Language Models in Household Activities</i>	
Ying Su, Zhan Ling, Haochen Shi, Cheng Jiayang, Yauwai Yim and Yangqiu Song .....	14953
<i>Shortcuts Arising from Contrast: Towards Effective and Lightweight Clean-Label Attacks in Prompt-Based Learning</i>	
Xiaopeng Xie, Ming Yan, Xiwen Zhou, Chenlong Zhao, Suli Wang, Yong Zhang and Joey Tianyi Zhou .....	14966
<i>GRASS: Compute Efficient Low-Memory LLM Training with Structured Sparse Gradients</i>	
Aashiq Muhamed, Oscar Li, David Woodruff, Mona T. Diab and Virginia Smith .....	14978
<i>RaTEScore: A Metric for Radiology Report Generation</i>	
Weike Zhao, Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang and Weidi Xie .....	15004
<i>HalluMeasure: Fine-grained Hallucination Measurement Using Chain-of-Thought Reasoning</i>	
Shayan Ali Akbar, Md Mosharaf Hossain, Tess Wood, Si-Chi Chin, Erica M Salinas, Victor Alvarez and Erwin Cornejo .....	15020

<i>Learning to Rank Salient Content for Query-focused Summarization</i> Sajad Sotudeh and Nazli Goharian .....	15038
<i>Are Large Language Models Good Classifiers? A Study on Edit Intent Classification in Scientific Document Revisions</i> Qian Ruan, Ilia Kuznetsov and Iryna Gurevych .....	15049
<i>LitSearch: A Retrieval Benchmark for Scientific Literature Search</i> Anirudh Ajith, Mengzhou Xia, Alexis Chevalier, Tanya Goyal, Danqi Chen and Tianyu Gao	15068
<i>Open-world Multi-label Text Classification with Extremely Weak Supervision</i> Xintong Li, Jinya Jiang, Ria Dharmani, Jayanth Srinivasa, Gaowen Liu and Jingbo Shang .	15084
<i>LLMs learn governing principles of dynamical systems, revealing an in-context neural scaling law</i> Toni J.b. Liu, Nicolas Boulle, Raphaël Sarfati and Christopher Earls .....	15097
<i>AKEW: Assessing Knowledge Editing in the Wild</i> Xiaobao Wu, Liangming Pan, William Yang Wang and Anh Tuan Luu .....	15118
<i>CopyBench: Measuring Literal and Non-Literal Reproduction of Copyright-Protected Text in Language Model Generation</i> Tong Chen, Akari Asai, Niloofer Miresghallah, Sewon Min, James Grimmermann, Yejin Choi, Hannaneh Hajishirzi, Luke Zettlemoyer and Pang Wei Koh.....	15134
<i>Dense X Retrieval: What Retrieval Granularity Should We Use?</i> Tong Chen, Hongwei Wang, Sihao Chen, Wenhao Yu, Kaixin Ma, Xinran Zhao, Hongming Zhang and Dong Yu .....	15159
<i>Decoding Susceptibility: Modeling Misbelief to Misinformation Through a Computational Approach</i> Yanchen Liu, Mingyu Derek Ma, Wenna Qin, Azure Zhou, Jiaao Chen, Weiyan Shi, Wei Wang and Diyi Yang .....	15178
<i>Layer by Layer: Uncovering Where Multi-Task Learning Happens in Instruction-Tuned Large Language Models</i> Zheng Zhao, Yftah Ziser and Shay B Cohen .....	15195
<i>XDetox: Text Detoxification with Token-Level Toxicity Explanations</i> Beomseok Lee, Hyunwoo Kim, Keon Kim and Yong Suk Choi .....	15215
<i>Optimizing Chinese Lexical Simplification Across Word Types: A Hybrid Approach</i> ZiHao Xiao, Jiefu Gong, Shijin Wang and Wei Song.....	15227
<i>Control Large Language Models via Divide and Conquer</i> Bingxuan Li, Yiwei Wang, Tao Meng, Kai-Wei Chang and Nanyun Peng .....	15240
<i>Joint Pre-Encoding Representation and Structure Embedding for Efficient and Low-Resource Knowledge Graph Completion</i> Chenyu Qiu, Pengjiang Qian, Chuang Wang, Jian Yao, Li Liu, Fang Wei and Eddie Y.k. Eddie	15257
<i>Improving Discriminative Capability of Reward Models in RLHF Using Contrastive Learning</i> Lu Chen, Rui Zheng, Binghai Wang, Senjie Jin, Caishuang Huang, Junjie Ye, Zhihao Zhang, Yuhao Zhou, Zhiheng Xi, Tao Gui, Qi Zhang and Xuanjing Huang.....	15270
<i>RoCEL: Advancing Table Entity Linking through Distinctive Row and Column Contexts</i> Yuanzheng Wang, Yixing Fan, Jiafeng Guo, Ruqing Zhang and Xueqi Cheng .....	15284

<i>Exploring the Role of Reasoning Structures for Constructing Proofs in Multi-Step Natural Language Reasoning with Large Language Models</i>	
Zi'ou Zheng, Christopher Malon, Martin Renqiang Min and Xiaodan Zhu . . . . .	15299
<i>Efficient Overshadowed Entity Disambiguation by Mitigating Shortcut Learning</i>	
Panuthep Tasawong, Peerat Limkonchotiwat, Potsawee Manakul, Can Udomcharoenchaikit, Eka-pol Chuangsuwanich and Sarana Nutanong . . . . .	15313
<i>AppBench: Planning of Multiple APIs from Various APPs for Complex User Instruction</i>	
Hongru Wang, Rui Wang, Boyang Xue, Heming Xia, Jingtao Cao, Zeming Liu, Jeff Z. Pan and Kam-Fai Wong . . . . .	15322
<i>Not Everything is All You Need: Toward Low-Redundant Optimization for Large Language Model Alignment</i>	
Zhipeng Chen, Kun Zhou, Xin Zhao, Jingyuan Wang and Ji-Rong Wen . . . . .	15337
<i>AudioVSR: Enhancing Video Speech Recognition with Audio Data</i>	
Xiaoda Yang, Xize Cheng, Jiaqi Duan, Hongshun Qiu, Minjie Hong, Minghui Fang, Shengpeng Ji, Jialong Zuo, Zhiqing Hong, Zhimeng Zhang and Tao Jin . . . . .	15352
<i>ECCO: Can We Improve Model-Generated Code Efficiency Without Sacrificing Functional Correctness?</i>	
Siddhant Waghjale, Vishruth Veerendranath, Zhiruo Wang and Daniel Fried . . . . .	15362
<i>Ladder: A Model-Agnostic Framework Boosting LLM-based Machine Translation to the Next Level</i>	
Zhaopeng Feng, Ruizhe Chen, Yan Zhang, Zijie Meng and Zuozhu Liu . . . . .	15377
<i>Re-ReST: Reflection-Reinforced Self-Training for Language Agents</i>	
Zi-Yi Dou, Cheng-Fu Yang, Xueqing Wu, Kai-Wei Chang and Nanyun Peng . . . . .	15394
<i>Effective Synthetic Data and Test-Time Adaptation for OCR Correction</i>	
Shuhao Guan, Cheng Xu, Moule Lin and Derek Greene . . . . .	15412
<i>SRF: Enhancing Document-Level Relation Extraction with a Novel Secondary Reasoning Framework</i>	
Fu Zhang, Qi Miao, Jingwei Cheng, Hongsen Yu, Yi Yan, Xin Li and Yongxue Wu . . . . .	15426
<i>FineCops-Ref: A new Dataset and Task for Fine-Grained Compositional Referring Expression Comprehension</i>	
Junzhuo Liu, Xuzheng Yang, Weiwei LI and Peng Wang . . . . .	15440
<i>Exploring the Learning Capabilities of Language Models using LEVERWORLDS</i>	
Eitan Wagner, Amir Feder and Omri Abend . . . . .	15458
<i>CONTESTS: a Framework for Consistency Testing of Span Probabilities in Language Models</i>	
Eitan Wagner, Yuli Slavutsky and Omri Abend . . . . .	15469
<i>DocEdit-v2: Document Structure Editing Via Multimodal LLM Grounding</i>	
Manan Suri, Puneet Mathur, Franck Dernoncourt, Rajiv Jain, Vlad I Morariu, Ramit Sawhney, Preslav Nakov and Dinesh Manocha . . . . .	15485
<i>DogeRM: Equipping Reward Models with Domain Knowledge through Model Merging</i>	
Tzu-Han Lin, Chen-An Li, Hung-yi Lee and Yun-Nung Chen . . . . .	15506
<i>Understanding Slang with LLMs: Modelling Cross-Cultural Nuances through Paraphrasing</i>	
Ifeoluwa Wuraola, Nina Dethlefs and Daniel Marciniak . . . . .	15525



<i>Unlocking Anticipatory Text Generation: A Constrained Approach for Large Language Models Decoding</i>	Lifu Tu, Semih Yavuz, Jin Qu, Jiacheng Xu, Rui Meng, Caiming Xiong and Yingbo Zhou .	15532
<i>Re-Reading Improves Reasoning in Large Language Models</i>	Xiaohan Xu, Chongyang Tao, Tao Shen, Can Xu, Hongbo Xu, Guodong Long, Jian-Guang Lou and Shuai Ma.....	15549
<i>Adaptive Axes: A Pipeline for In-domain Social Stereotype Analysis</i>	Qingcheng Zeng, Mingyu Jin and Rob Voigt.....	15576
<i>ERVQA: A Dataset to Benchmark the Readiness of Large Vision Language Models in Hospital Environments</i>	Sourjyadip Ray, Kushal Gupta, Soumi Kundu, Dr Payal Arvind Kasat, Somak Aditya and Pawan Goyal.....	15594
<i>Human-LLM Hybrid Text Answer Aggregation for Crowd Annotations</i>	Jiyi Li.....	15609
<i>Improve Student’s Reasoning Generalizability through Cascading Decomposed CoTs Distillation</i>	Chengwei Dai, Kun Li, Wei Zhou and Songlin Hu.....	15623
<i>Revisiting Supervised Contrastive Learning for Microblog Classification</i>	Junbo Huang and Ricardo Usbeck.....	15644
<i>BaitAttack: Alleviating Intention Shift in Jailbreak Attacks via Adaptive Bait Crafting</i>	Rui Pu, Chaozhuo Li, Rui Ha, Litian Zhang, Lirong Qiu and Xi Zhang.....	15654
<i>Images Speak Louder than Words: Understanding and Mitigating Bias in Vision-Language Model from a Causal Mediation Perspective</i>	Zhaotian Weng, Zijun Gao, Jerone Andrews and Jieyu Zhao.....	15669
<i>Mitigating the Language Mismatch and Repetition Issues in LLM-based Machine Translation via Model Editing</i>	Weichuan Wang, Zhaoyi Li, Defu Lian, Chen Ma, Linqi Song and Ying Wei.....	15681
<i>SciAgent: Tool-augmented Language Models for Scientific Reasoning</i>	Yubo Ma, Zhibin Gou, Junheng Hao, Ruochen Xu, Shuohang Wang, Liangming Pan, Yujiu Yang, Yixin Cao and Aixin Sun.....	15701
<i>Global Reward to Local Rewards: Multimodal-Guided Decomposition for Improving Dialogue Agents</i>	Dong Won Lee, Hae Won Park, Yoon Kim, Cynthia Breazeal and Louis-Philippe Morency	15737
<i>Towards Measuring and Modeling “Culturein LLMs: A Survey</i>	Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Shivdutt Singh, Alham Fikri Aji, Jacki O’Neill, Ashutosh Modi and Monojit Choudhury.....	15763
<i>ESC-Eval: Evaluating Emotion Support Conversations in Large Language Models</i>	Haiquan Zhao, Lingyu Li, Shisong Chen, Shuqi Kong, Jiaan Wang, Kexin Huang, Tianle Gu, Yixu Wang, Jian Wang, Liang Dandan, Zhixu Li, Yan Teng, Yanghua Xiao and Yingchun Wang	15785
<i>Cultural Conditioning or Placebo? On the Effectiveness of Socio-Demographic Prompting</i>	Sagnik Mukherjee, Muhammad Farid Adilazuarda, Sunayana Sitaram, Kalika Bali, Alham Fikri Aji and Monojit Choudhury.....	15811
<i>Text Fluoroscopy: Detecting LLM-Generated Text through Intrinsic Features</i>	Xiao Yu, Kejiang Chen, Qi Yang, Weiming Zhang and Nenghai Yu.....	15838

<i>Hate Personified: Investigating the role of LLMs in content moderation</i>	
Sarah Masud, Sahajpreet Singh, Viktor Hangya, Alexander Fraser and Tanmoy Chakraborty	15847
<i>Temporally Consistent Factuality Probing for Large Language Models</i>	
Ashutosh Bajpai, Aaryan Goyal, Atif Anwer and Tanmoy Chakraborty	15864
<i>A Comparison of Language Modeling and Translation as Multilingual Pretraining Objectives</i>	
Zihao Li, Shaoxiong Ji, Timothee Mickus, Vincent Segonne and Jörg Tiedemann	15882
<i>Can LLMs replace Neil deGrasse Tyson? Evaluating the Reliability of LLMs as Science Communicators</i>	
Prasoon Bajpai, Niladri Chatterjee, Subhabrata Dutta and Tanmoy Chakraborty	15895
<i>LLaMA-MoE: Building Mixture-of-Experts from LLaMA with Continual Pre-Training</i>	
Tong Zhu, Xiaoye Qu, Daize Dong, Jiacheng Ruan, Jingqi Tong, Conghui He and Yu Cheng	15913
<i>Themis: A Reference-free NLG Evaluation Language Model with Flexibility and Interpretability</i>	
Xinyu Hu, Li Lin, Mingqi Gao, Xunjian Yin and Xiaojun Wan	15924
<i>Mitigating Training Imbalance in LLM Fine-Tuning via Selective Parameter Merging</i>	
Yiming Ju, Ziyi Ni, Xingrun Xing, Zhixiong Zeng, Hanyu Zhao, Siqi Fan and Zheng Zhang	15952
<i>Generating Demonstrations for In-Context Compositional Generalization in Grounded Language Learning</i>	
Sam Spilsbury, Pekka Marttinen and Alexander Ilin	15960
<i>FAME: Towards Factual Multi-Task Model Editing</i>	
Li Zeng, Yingyu Shan, Zeming Liu, Jiashu Yao and Yuhang Guo	15992
<i>MLLM-Protector: Ensuring MLLM’s Safety without Hurting Performance</i>	
Renjie Pi, Tianyang Han, Jianshu Zhang, Yueqi Xie, Rui Pan, Qing Lian, Hanze Dong, Jipeng Zhang and Tong Zhang	16012
<i>Leveraging Large Language Models for NLG Evaluation: Advances and Challenges</i>	
Zhen Li, Xiaohan Xu, Tao Shen, Can Xu, Jia-Chen Gu, Yuxuan Lai, Chongyang Tao and Shuai Ma	16028
<i>InfiniPot: Infinite Context Processing on Memory-Constrained LLMs</i>	
Minsoo Kim, Kyuhong Shim, Jungwook Choi and Simyung Chang	16046
<i>VideoCLIP-XL: Advancing Long Description Understanding for Video CLIP Models</i>	
Jiapeng Wang, Chengyu Wang, Kunzhe Huang, Jun Huang and Lianwen Jin	16061
<i>CorrSynth - A Correlated Sampling Method for Diverse Dataset Generation from LLMs</i>	
Suhas S Kowshik, Abhishek Divekar and Vijit Malik	16076
<i>Defining Knowledge: Bridging Epistemology and Large Language Models</i>	
Constanza Fierro, Ruchira Dhar, Filippos Stamatiou, Nicolas Garneau and Anders Søgaard	16096
<i>TKGT: Redefinition and A New Way of Text-to-Table Tasks Based on Real World Demands and Knowledge Graphs Augmented LLMs</i>	
Peiwen Jiang, Xinbo Lin, Zibo Zhao, Ruhui Ma, Yvonne Jie Chen and Jinhua Cheng	16112
<i>Free your mouse! Command Large Language Models to Generate Code to Format Word Documents</i>	
Shihao Rao, Liang Li, Jiapeng Liu, Guan Weixin, Xiyan Gao, Bing Lim and Can Ma	16127
<i>CMR Scaling Law: Predicting Critical Mixture Ratios for Continual Pre-training of Language Models</i>	
Jiawei Gu, Zacc Yang, Chuanghao Ding, Rui Zhao and Fei Tan	16143

<i>The Instinctive Bias: Spurious Images lead to Illusion in MLLMs</i> Tianyang Han, Qing Lian, Rui Pan, Renjie Pi, Jipeng Zhang, Shizhe Diao, Yong Lin and Tong Zhang .....	16163
<i>Rationale-Aware Answer Verification by Pairwise Self-Evaluation</i> Akira Kawabata and Saku Sugawara .....	16178
<i>On the Robustness of Editing Large Language Models</i> Xinbei Ma, Tianjie Ju, Jiyang Qiu, Zhuosheng Zhang, Hai Zhao, Lifeng Liu and Yulong Wang	16197
<i>IM-BERT: Enhancing Robustness of BERT through the Implicit Euler Method</i> MiHyeon Kim, Juhyoung Park and YoungBin Kim .....	16217
<i>Distract Large Language Models for Automatic Jailbreak Attack</i> Zeguan Xiao, Yan Yang, Guanhua Chen and Yun Chen .....	16230
<i>Exploring Space Efficiency in a Tree-based Linear Model for Extreme Multi-label Classification</i> He-Zhe Lin, Cheng-Hung Liu and Chih-Jen Lin .....	16245
<i>WorryWords: Norms of Anxiety Association for over 44k English Words</i> Saif M. Mohammad .....	16261
<i>Finding Blind Spots in Evaluator LLMs with Interpretable Checklists</i> Sumanth Doddapaneni, Mohammed Safi Ur Rahman Khan, Sshubam Verma and Mitesh M Khapra .....	16279
<i>LONGAGENT: Achieving Question Answering for 128k-Token-Long Documents through Multi-Agent Collaboration</i> Jun Zhao, Can Zu, Xu Hao, Yi Lu, Wei He, Yiwen Ding, Tao Gui, Qi Zhang and Xuanjing Huang	16310
<i>AutoPersuade: A Framework for Evaluating and Explaining Persuasive Arguments</i> Till Raphael Saenger, Musashi Hinck, Justin Grimmer and Brandon M. Stewart .....	16325
<i>Towards Cross-Cultural Machine Translation with Retrieval-Augmented Generation from Multilingual Knowledge Graphs</i> Simone Conia, Daniel Lee, Min Li, Umar Farooq Minhas, Saloni Potdar and Yunyao Li ...	16343
<i>Exploring the Compositional Deficiency of Large Language Models in Mathematical Reasoning Through Trap Problems</i> Jun Zhao, Jingqi Tong, Yurong Mou, Ming Zhang, Qi Zhang and Xuanjing Huang .....	16361
<i>Scaling Laws for Linear Complexity Language Models</i> Xuyang Shen, Dong Li, Ruitao Leng, Zhen Qin, Weigao Sun and Yiran Zhong .....	16377
<i>Autoregressive Multi-trait Essay Scoring via Reinforcement Learning with Scoring-aware Multiple Rewards</i> Heejin Do, Sangwon Ryu and Gary Lee .....	16427
<i>Intrinsic Self-correction for Enhanced Morality: An Analysis of Internal Mechanisms and the Superficial Hypothesis</i> Guangliang Liu, Haitao Mao, Jiliang Tang and Kristen Johnson .....	16439
<i>ATAP: Automatic Template-Augmented Commonsense Knowledge Graph Completion via Pre-Trained Language Models</i> Fu Zhang, Yifan Ding and Jingwei Cheng .....	16456

<i>LM2: A Simple Society of Language Models Solves Complex Reasoning</i> Gurusha Juneja, Subhabrata Dutta and Tanmoy Chakraborty .....	16473
<i>Towards a Similarity-adjusted Surprisal Theory</i> Clara Meister, Mario Giulianelli and Tiago Pimentel .....	16485
<i>Multi-Level Information Retrieval Augmented Generation for Knowledge-based Visual Question Answering</i> Adjali Omar, Olivier Ferret, Sahar Ghannay and Hervé Le Borgne .....	16499
<i>Can We Trust the Performance Evaluation of Uncertainty Estimation Methods in Text Summarization?</i> Jianfeng He, Runing Yang, Linlin Yu, Changbin Li, Ruoxi Jia, Feng Chen, Ming Jin and Chang-Tien Lu .....	16514
<i>Is It Really Long Context if All You Need Is Retrieval? Towards Genuinely Difficult Long Context NLP</i> Omer Goldman, Alon Jacovi, Aviv Slobodkin, Aviya Maimon, Ido Dagan and Reut Tsarfaty	16576
<i>BPE Gets Picky: Efficient Vocabulary Refinement During Tokenizer Training</i> Pavel Chizhov, Catherine Arnett, Elizaveta Korotkova and Ivan P. Yamshchikov .....	16587
<i>SEGMENT+: Long Text Processing with Short-Context Language Models</i> Wei Shi, Shuang Li, Kerun Yu, Jinglei Chen, Zujie Liang, Xinhui Wu, Yuxi Qian, Feng Wei, Bo Zheng, Jiaqing Liang, Jiangjie Chen and Yanghua Xiao .....	16605
<i>Explicit Memory Learning with Expectation Maximization</i> Zhangyue Yin, Qiushi Sun, Qipeng Guo, Zhiyuan Zeng, Qinyuan Cheng, Xipeng Qiu and Xuanjing Huang .....	16618
<i>Closing the Loop: Learning to Generate Writing Feedback via Language Model Simulated Student Revisions</i> Inderjeet Jayakumar Nair, Jiaye Tan, Xiaotian Su, Anne Gere, Xu Wang and Lu Wang ....	16636
<i>Small LLMs Are Weak Tool Learners: A Multi-LLM Agent</i> Weizhou Shen, Chenliang Li, Hongzhan Chen, Ming Yan, Xiaojun Quan, Hehong Chen, Ji Zhang and Fei Huang .....	16658
<i>Interpreting Context Look-ups in Transformers: Investigating Attention-MLP Interactions</i> Clement Neo, Shay B Cohen and Fazl Barez .....	16681
<i>Still Not Quite There! Evaluating Large Language Models for Comorbid Mental Health Diagnosis</i> Amey Hengle, Atharva Kulkarni, Shantanu Deepak Patankar, Madhumitha Chandrasekaran, Sneha D'silva, Jemima S. Jacob and Rashmi Gupta .....	16698
<i>The Odyssey of Commonsense Causality: From Foundational Benchmarks to Cutting-Edge Reasoning</i> Shaobo Cui, Zhijing Jin, Bernhard Schölkopf and Boi Faltings .....	16722
<i>Investigating Large Language Models for Complex Word Identification in Multilingual and Multidomain Setups</i> Răzvan-Alexandru Smădu, David-Gabriel Ion, Dumitru-Clementin Cercel, Florin Pop and Mihaela-Claudia Cercel .....	16764
<i>Model Editing Harms General Abilities of Large Language Models: Regularization to the Rescue</i> Jia-Chen Gu, Hao-Xiang Xu, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang and Nanyun Peng .....	16801
<i>Are Large Language Models In-Context Personalized Summarizers? Get an iCOPERNICUS Test Done!</i> Divya Patel, Pathik Patel, Ankush Chander, Sourish Dasgupta and Tanmoy Chakraborty ..	16820

<i>MediTOD: An English Dialogue Dataset for Medical History Taking with Comprehensive Annotations</i> Vishal Vivek Saley, Goonjan Saha, Rocktim Jyoti Das, Dinesh Raghu and Mausam . . . . .	16843
<i>***YesBut***: A High-Quality Annotated Multimodal Dataset for evaluating Satire Comprehension capability of Vision-Language Models</i> Abhilash Nandy, Yash Agarwal, Ashish Patwa, Millon Madhur Das, Aman Bansal, Ankit Raj, Pawan Goyal and Niloy Ganguly . . . . .	16878
<i>Working Memory Identifies Reasoning Limits in Language Models</i> Chunhui Zhang, Yiren Jian, Zhongyu Ouyang and Soroush Vosoughi . . . . .	16896
<i>RAFT: Realistic Attacks to Fool Text Detectors</i> James Liyuan Wang, Ran Li, Junfeng Yang and Chengzhi Mao . . . . .	16923
<i>LLM-Evolve: Evaluation for LLM’s Evolving Capability on Benchmarks</i> Jiaxuan You, Mingjie Liu, Shrimai Prabhumoye, Mostofa Patwary, Mohammad Shoeybi and Bryan Catanzaro . . . . .	16937
<i>FFN-SkipLLM: A Hidden Gem for Autoregressive Decoding with Adaptive Feed Forward Skipping</i> Ajay Kumar Jaiswal, Bodun Hu, Lu Yin, Yeonju Ro, Tianlong Chen, Shiwei Liu and Aditya Akella . . . . .	16943
<i>LLM-based Code-Switched Text Generation for Grammatical Error Correction</i> Tom Potter and Zheng Yuan . . . . .	16957
<i>Deciphering the Interplay of Parametric and Non-parametric Memory in Retrieval-augmented Language Models</i> Mehrdad Farahani and Richard Johansson . . . . .	16966
<i>On Efficient Language and Vision Assistants for Visually-Situated Natural Language Understanding: What Matters in Reading and Reasoning</i> Geewook Kim and Minjoon Seo . . . . .	16978
<i>Community-Cross-Instruct: Unsupervised Instruction Generation for Aligning Large Language Models to Online Communities</i> Zihao He, Minh Duc Chu, Rebecca Dorn, Siyi Guo and Kristina Lerman . . . . .	17001
<i>Mathador-LM: A Dynamic Benchmark for Mathematical Reasoning on Large Language Models</i> Eldar Kurtic, Amir Moeini and Dan Alistarh . . . . .	17020
<i>Reasoning Paths with Reference Objects Elicit Quantitative Spatial Reasoning in Large Vision-Language Models</i> Yuan-Hong Liao, Rafid Mahmood, Sanja Fidler and David Acuna . . . . .	17028
<i>One Thousand and One Pairs: A novel challenge for long-context language models</i> Marzena Karpinska, Katherine Thai, Kyle Lo, Tanya Goyal and Mohit Iyyer . . . . .	17048
<i>Foundational Autoraters: Taming Large Language Models for Better Automatic Evaluation</i> Tu Vu, Kalpesh Krishna, Salaheddin Alzubi, Chris Tar, Manaal Faruqui and Yun-Hsuan Sung	17086
<i>Do LLMs learn a true syntactic universal?</i> John T. Hale and Miloš Stanojević . . . . .	17106
<i>GDPO: Learning to Directly Align Language Models with Diversity Using GFlowNets</i> Oh Joon Kwon, Daiki E. Matsunaga and Kee-Eung Kim . . . . .	17120

<i>How Susceptible are Large Language Models to Ideological Manipulation?</i> Kai Chen, Zihao He, Jun Yan, Taiwei Shi and Kristina Lerman .....	17140
<i>Measuring Psychological Depth in Language Models</i> Fabrice Y Harel-Canada, Hanyu Zhou, Sreya Muppalla, Zeynep Senahan Yildiz, Miryung Kim, Amit Sahai and Nanyun Peng .....	17162
<i>Media Attitude Detection via Framing Analysis with Events and their Relations</i> Jin Zhao, Jingxuan Tu, Han Du and Nianwen Xue .....	17197
<i>Fill In The Gaps: Model Calibration and Generalization with Synthetic Data</i> Yang Ba, Michelle V Mancenido and Rong Pan .....	17211
<i>Adaptive Question Answering: Enhancing Language Model Proficiency for Addressing Knowledge Conflicts with Source Citations</i> Sagi Shaier, Ari Kobren and Philip V. Ogren .....	17226
<i>Granular Privacy Control for Geolocation with Vision Language Models</i> Ethan Mendes, Yang Chen, James Hays, Sauvik Das, Wei Xu and Alan Ritter .....	17240
<i>MedReadMe: A Systematic Study for Fine-grained Sentence Readability in Medical Domain</i> Chao Jiang and Wei Xu .....	17293
<i>MemeCLIP: Leveraging CLIP Representations for Multimodal Meme Classification</i> Siddhant Bikram Shah, Shuvam Shiwakoti, Maheep Chaudhary and Haohan Wang .....	17320
<i>FlipGuard: Defending Preference Alignment against Update Regression with Constrained Optimization</i> Mingye Zhu, Yi Liu, Quan Wang, Junbo Guo and Zhendong Mao .....	17333
<i>StorySparkQA: Expert-Annotated QA Pairs with Real-World Knowledge for Children’s Story-Based Learning</i> Jiaju Chen, Yuxuan Lu, Shao Zhang, Bingsheng Yao, Yuanzhe Dong, Ying Xu, Yunyao Li, Qia- nwen Wang, Dakuo Wang and Yuling Sun .....	17351
<i>MedCoT: Medical Chain of Thought via Hierarchical Expert</i> Jiaxiang Liu, Yuan Wang, Jiawei Du, Joey Tianyi Zhou and Zuozhu Liu .....	17371
<i>Varying Sentence Representations via Condition-Specified Routers</i> Ziyong Lin, Quansen Wang, Zixia Jia and Zilong Zheng .....	17390
<i>Inductive-Deductive Strategy Reuse for Multi-Turn Instructional Dialogues</i> Jiao Ou, Jiayu wu, Che Liu, Fuzheng Zhang, Di Zhang and Kun Gai .....	17402
<i>Information Flow Routes: Automatically Interpreting Language Models at Scale</i> Javier Ferrando and Elena Voita .....	17432
<i>A Simple yet Effective Training-free Prompt-free Approach to Chinese Spelling Correction Based on Large Language Models</i> Houquan Zhou, Zhenghua Li, Bo Zhang, Chen Li, Shaopeng Lai, Ji Zhang, Fei Huang and Min Zhang .....	17446
<i>Representational Analysis of Binding in Language Models</i> Qin Dai, Benjamin Heinzerling and Kentaro Inui .....	17468
<i>CoSafe: Evaluating Large Language Model Safety in Multi-Turn Dialogue Coreference</i> Erxin Yu, Jing Li, Ming Liao, Siqi Wang, Gao Zuchen, Fei Mi and Lanqing Hong .....	17494

<i>ClimRetrieve: A Benchmarking Dataset for Information Retrieval from Corporate Climate Disclosures</i> Tobias Schimanski, Jingwei Ni, Roberto Spacey Martín, Nicola Ranger and Markus Leippold	17509
<i>Context-Aware Adapter Tuning for Few-Shot Relation Learning in Knowledge Graphs</i> Liu Ran, Zhongzhou Liu, Xiaoli Li and Yuan Fang	17525
<i>Zero-Shot Detection of LLM-Generated Text using Token Cohesiveness</i> Shixuan Ma and Quan Wang	17538
<i>Dual-oriented Disentangled Network with Counterfactual Intervention for Multimodal Intent Detection</i> Zhanpeng Chen, Zhihong Zhu, Xianwei Zhuang, Zhiqi Huang and Yuexian Zou	17554
<i>From LLMs to MLLMs: Exploring the Landscape of Multimodal Jailbreaking</i> Siyuan Wang, Zhuohan Long, Zhihao Fan and Zhongyu Wei	17568
<i>Symbolic Working Memory Enhances Language Models for Complex Rule Application</i> Siyuan Wang, Zhongyu Wei, Yejin Choi and Xiang Ren	17583
<i>LLoCO: Learning Long Contexts Offline</i> Sijun Tan, Xiuyu Li, Shishir G Patil, Ziyang Wu, Tianjun Zhang, Kurt Keutzer, Joseph E. Gonzalez and Raluca Ada Popa	17605
<i>Don't Forget Your Reward Values: Language Model Alignment via Value-based Calibration</i> Xin Mao, Feng-Lin Li, Huimin Xu, Wei Zhang, Wang Chen and Anh Tuan Luu	17622
<i>Mentor-KD: Making Small Language Models Better Multi-step Reasoners</i> Hojae Lee, Junho Kim and SangKeun Lee	17643
<i>Are Large Language Models Capable of Generating Human-Level Narratives?</i> Yufei Tian, Tenghao Huang, Miri Liu, Derek Jiang, Alexander Spangher, Muhao Chen, Jonathan May and Nanyun Peng	17659
<i>MP2D: An Automated Topic Shift Dialogue Generation Framework Leveraging Knowledge Graphs</i> Yerin Hwang, Yongil Kim, Yunah Jang, Jeesoo Bang, Hyunkyung Bae and Kyomin Jung	17682
<i>Can Large Language Models Enhance Predictions of Disease Progression? Investigating Through Disease Network Link Prediction</i> Haohui Lu and Usman Naseem	17703
<i>Searching for Best Practices in Retrieval-Augmented Generation</i> Xiaohua Wang, Zhenghua Wang, Xuan Gao, Feiran Zhang, Yixin Wu, Zhibo Xu, Tianyuan Shi, Zhengyuan Wang, Shizheng Li, Qi Qian, Ruicheng Yin, Changze Lv, Xiaoqing Zheng and Xuanjing Huang	17716
<i>Moral Foundations of Large Language Models</i> Marwa Abdulhai, Gregory Serapio-García, Clement Crepy, Daria Valter, John Canny and Natasha Jaques	17737
<i>The Zeno's Paradox of 'Low-Resource' Languages</i> Hellina Hailu Nigatu, Atnafu Lambebo Tonja, Benjamin Rosman, Tamar Solorio and Monojit Choudhury	17753
<i>Knowledge Planning in Large Language Models for Domain-Aligned Counseling Summarization</i> Aseem Srivastava, Smriti Joshi, Tanmoy Chakraborty and Md Shad Akhtar	17775

<i>Enhancing Post-Hoc Attributions in Long Document Comprehension via Coarse Grained Answer Decomposition</i>	Pritika Ramu, Koustava Goswami, Apoorv Saxena and Balaji Vasani Srinivasan . . . . .	17790
<i>From Descriptive Richness to Bias: Unveiling the Dark Side of Generative Image Caption Enrichment</i>	Yusuke Hirota, Ryo Hachiuma, Chao-Han Huck Yang and Yuta Nakashima . . . . .	17807
<i>Pruning via Merging: Compressing LLMs via Manifold Alignment Based Layer Merging</i>	Deyuan Liu, Zhanyue Qin, Hairu Wang, Zhao Yang, Zecheng Wang, Fangying Rong, Qingbin Liu, Yanchao Hao, Bo Li, Xi Chen, Cunhang Fan, Zhao Lv, Dianhui Chu, Zhiying Tu and Dianbo Sui	17817
<i>Embedded Named Entity Recognition using Probing Classifiers</i>	Nicholas Popovic and Michael Färber . . . . .	17830
<i>Unleashing the Power of Emojis in Texts via Self-supervised Graph Pre-Training</i>	Zhou Zhang, Dongzeng Tan, Jiaan Wang, Yilong Chen and Jiarong Xu . . . . .	17851
<i>Data Contamination Can Cross Language Barriers</i>	Feng Yao, Yufan Zhuang, Zihao Sun, Sunan Xu, Animesh Kumar and Jingbo Shang . . . . .	17864
<i>Automated Essay Scoring: A Reflection on the State of the Art</i>	Shengjie Li and Vincent Ng . . . . .	17876
<i>Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate</i>	Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi and Zhaopeng Tu . . . . .	17889
<i>Unveiling and Consulting Core Experts in Retrieval-Augmented MoE-based LLMs</i>	Xin Zhou, Ping Nie, Yiwen Guo, Haojie Wei, Zhanqiu Zhang, Pasquale Minervini, Ruotian Ma, Tao Gui, Qi Zhang and Xuanjing Huang . . . . .	17905
<i>CURE: Context- and Uncertainty-Aware Mental Disorder Detection</i>	Migyeong Kang, Goun Choi, Hyolim Jeon, Ji Hyun An, Daejin Choi and Jinyoung Han . . . . .	17924
<i>PepRec: Progressive Enhancement of Prompting for Recommendation</i>	Yakun Yu, Shi-ang Qi, Baochun Li and Di Niu . . . . .	17941
<i>In-Context Compositional Generalization for Large Vision-Language Models</i>	Chuanhao Li, Chenchen Jing, Zhen Li, Mingliang Zhai, Yuwei Wu and Yunde Jia . . . . .	17954
<i>Improving Zero-shot LLM Re-Ranker with Risk Minimization</i>	Xiaowei Yuan, Zhao Yang, Yequan Wang, Jun Zhao and Kang Liu . . . . .	17967
<i>Game on Tree: Visual Hallucination Mitigation via Coarse-to-Fine View Tree and Game Theory</i>	Xianwei Zhuang, Zhihong Zhu, Zhanpeng Chen, Yuxin Xie, Liming Liang and Yuexian Zou	17984
<i>Label Confidence Weighted Learning for Target-level Sentence Simplification</i>	Xin Ying Qiu and Jingshen Zhang . . . . .	18004
<i>Quantum Recurrent Architectures for Text Classification</i>	Wenduan Xu, Stephen Clark, Douglas Brown, Gabriel Matos and Konstantinos Meichanetzidis	18020
<i>Tree of Problems: Improving structured problem solving with compositionality</i>	Armel Randy Zebaze, Benoît Sagot and Rachel Bawden . . . . .	18028



<i>What the Harm? Quantifying the Tangible Impact of Gender Bias in Machine Translation with a Human-centered Study</i>	Beatrice Savoldi, Sara Papi, Matteo Negri, Ana Guerberof-Arenas and Luisa Bentivogli . . .	18048
<i>Seg2Act: Global Context-aware Action Generation for Document Logical Structuring</i>	Zichao Li, Shaojie He, Meng Liao, Xuanang Chen, Yaojie Lu, Hongyu Lin, Yanxiong Lu, Xianpei Han and Le Sun . . . . .	18077
<i>Is C4 Dataset Optimal for Pruning? An Investigation of Calibration Data for LLM Pruning</i>	Abhinav Bandari, Lu Yin, Cheng-Yu Hsieh, Ajay Kumar Jaiswal, Tianlong Chen, Li Shen, Ranjay Krishna and Shiwei Liu . . . . .	18089
<i>Revisiting the Robustness of Watermarking to Paraphrasing Attacks</i>	Saksham Rastogi and Danish Pruthi . . . . .	18100
<i>A Survey of Ontology Expansion for Conversational Understanding</i>	Jinggui Liang, Yuxia Wu, Yuan Fang, Hao Fei and Lizi Liao . . . . .	18111
<i>Calibrating Language Models with Adaptive Temperature Scaling</i>	Johnathan Xie, Annie S Chen, Yoonho Lee, Eric Mitchell and Chelsea Finn . . . . .	18128
<i>Which Programming Language and What Features at Pre-training Stage Affect Downstream Logical Inference Performance?</i>	Fumiya Uchiyama, Takeshi Kojima, Andrew Gambardella, Qi Cao, Yusuke Iwasawa and Yutaka Matsuo . . . . .	18139
<i>Why do objects have many names? A study on word informativeness in language use and lexical systems</i>	Eleonora Gualdoni and Gemma Boleda . . . . .	18150
<i>Dual-Space Knowledge Distillation for Large Language Models</i>	Songming Zhang, Xue Zhang, Zengkui Sun, Yufeng Chen and Jinan Xu . . . . .	18164
<i>NoiseBench: Benchmarking the Impact of Real Label Noise on Named Entity Recognition</i>	Elena Merdjanovska, Ansar Aynetdinov and Alan Akbik . . . . .	18182
<i>On the Universal Truthfulness Hyperplane Inside LLMs</i>	Junteng Liu, Shiqi Chen, Yu Cheng and Junxian He . . . . .	18199
<i>PairDistill: Pairwise Relevance Distillation for Dense Retrieval</i>	Chao-Wei Huang and Yun-Nung Chen . . . . .	18225
<i>User Inference Attacks on Large Language Models</i>	Nikhil Kandpal, Krishna Pillutla, Alina Oprea, Peter Kairouz, Christopher A. Choquette-Choo and Zheng Xu . . . . .	18238
<i>HiFT: A Hierarchical Full Parameter Fine-Tuning Strategy</i>	YongKang Liu, Yiqun Zhang, Qian Li, Tong Liu, Shi Feng, Daling Wang, Yifei Zhang and Hinrich Schuetze . . . . .	18266
<i>Investigating and Mitigating Object Hallucinations in Pretrained Vision-Language (CLIP) Models</i>	Yufang Liu, Tao Ji, Changzhi Sun, Yuanbin Wu and Aimin Zhou . . . . .	18288
<i>Simultaneous Masking, Not Prompting Optimization: A Paradigm Shift in Fine-tuning LLMs for Simultaneous Translation</i>	Matthew Raffel, Victor Agostinelli and Lizhong Chen . . . . .	18302

<i>ToolPlanner: A Tool Augmented LLM for Multi Granularity Instructions with Path Planning and Feedback</i>	
Qinzhuo Wu, Wei Liu, Jian Luan and Bin Wang .....	18315
<i>Please note that I'm just an AI: Analysis of Behavior Patterns of LLMs in (Non-)offensive Speech Identification</i>	
Esra Dönmez, Thang Vu and Agnieszka Falenska .....	18340
<i>How to Compute the Probability of a Word</i>	
Tiago Pimentel and Clara Meister .....	18358
<i>A linguistically-motivated evaluation methodology for unraveling model's abilities in reading comprehension tasks</i>	
Elie Antoine, Frederic Bechet, Géraldine Damnati and Philippe Langlais .....	18376
<i>GuardBench: A Large-Scale Benchmark for Guardrail Models</i>	
Elias Bassani and Ignacio Sanchez .....	18393
<i>Generate-on-Graph: Treat LLM as both Agent and KG for Incomplete Knowledge Graph Question Answering</i>	
Yao Xu, Shizhu He, Jiabei Chen, Zihao Wang, Yangqiu Song, Hanghang Tong, Guang Liu, Jun Zhao and Kang Liu .....	18410
<i>Language models and brains align due to more than next-word prediction and word-level information</i>	
Gabriele Merlin and Mariya Toneva .....	18431
<i>LLMEdgeRefine: Enhancing Text Clustering with LLM-Based Boundary Point Refinement</i>	
Zijin Feng, Luyang Lin, Lingzhi Wang, Hong Cheng and Kam-Fai Wong .....	18455
<i>CasiMedicos-Arg: A Medical Question Answering Dataset Annotated with Explanatory Argumentative Structures</i>	
Ekaterina Sviridova, Anar Yeginbergen, Ainara Estarrona, Elena Cabrio, Serena Villata and Rodrigo Agerri .....	18463
<i>A Simple and Effective L<sub>2</sub> Norm-Based Strategy for KV Cache Compression</i>	
Alessio Devoto, Yu Zhao, Simone Scardapane and Pasquale Minervini .....	18476
<i>GOME: Grounding-based Metaphor Binding With Conceptual Elaboration For Figurative Language Illustration</i>	
Linhao Zhang, Jintao Liu, Li Jin, Hao Wang, Kaiwen Wei and Guangluan Xu .....	18500
<i>D3CODE: Disentangling Disagreements in Data across Cultures on Offensiveness Detection and Evaluation</i>	
Aida Mostafazadeh Davani, Mark Diaz, Dylan K Baker and Vinodkumar Prabhakaran .....	18511
<i>PALM: Few-Shot Prompt Learning for Audio Language Models</i>	
Asif Hanif, Maha Tufail Agro, Mohammad Areeb Qazi and Hanan Aldarmaki .....	18527
<i>Annotator-Centric Active Learning for Subjective NLP Tasks</i>	
Michiel Van Der Meer, Neele Falk, Pradeep K. Murukannaiah and Enrico Liscio .....	18537
<i>On the Proper Treatment of Tokenization in Psycholinguistics</i>	
Mario Giulianelli, Luca Malagutti, Juan Luis Gastaldi, Brian DuSell, Tim Vieira and Ryan Cottrell .....	18556
<i>Enhanced Hallucination Detection in Neural Machine Translation through Simple Detector Aggregation</i>	
Anas Himmi, Guillaume Staerman, Marine Picot, Pierre Colombo and Nuno M Guerreiro .....	18573

<i>Jailbreaking LLMs with Arabic Transliteration and Arabizi</i>	
Mansour Al Ghanim, Saleh Almohaimeed, Mengxin Zheng, Yan Solihin and Qian Lou . . .	18584
<i>Who is better at math, Jenny or Jingzhen? Uncovering Stereotypes in Large Language Models</i>	
Zara Siddique, Liam Turner and Luis Espinosa-Anke . . . . .	18601
<i>Instruction Matters: A Simple yet Effective Task Selection for Optimized Instruction Tuning of Specific Tasks</i>	
Changho Lee, Janghoon Han, Seonghyeon Ye, Stanley Jungkyu Choi, Honglak Lee and Kyung-hoon Bae . . . . .	18620
<i>Recurrent Alignment with Hard Attention for Hierarchical Text Rating</i>	
Chenxi Lin, Ren Jiayu, Guoxiu He, Zhuoren Jiang, Haiyan yu and Xiaomin Zhu . . . . .	18643
<i>CHESS: Optimizing LLM Inference via Channel-Wise Thresholding and Selective Sparsification</i>	
Junhui He, Shangyu Wu, Weidong Wen, Chun Jason Xue and Qingan Li . . . . .	18658
<i>Semformer: Transformer Language Models with Semantic Planning</i>	
Yongjing Yin, Junran Ding, Kai Song and Yue Zhang . . . . .	18669
<i>DocCGen: Document-based Controlled Code Generation</i>	
Sameer Pimparkhede, Mehant Kammakomati, Srikanth G. Tamilselvam, Prince Kumar, Ashok Pon Kumar and Pushpak Bhattacharyya . . . . .	18681
<i>Semantics and Sentiment: Cross-lingual Variations in Emoji Use</i>	
Giulio Zhou, Sydelle De Souza, Ella Markham, Oghenetekewe Kwakpovwe and Sumin Zhao	18698
<i>The Emergence of Compositional Languages in Multi-entity Referential Games: from Image to Graph Representations</i>	
Daniel Akkerman, Phong Le and Raquel G. Alhama . . . . .	18713
<i>Transformers are Multi-State RNNs</i>	
Matanel Oren, Michael Hassid, Nir Yarden, Yossi Adi and Roy Schwartz . . . . .	18724
<i>Evaluating Large Language Models along Dimensions of Language Variation: A Systematic Investigation of Cross-lingual Generalization</i>	
Niyati Bafna, Kenton Murray and David Yarowsky . . . . .	18742
<i>Fuse to Forget: Bias Reduction and Selective Memorization through Model Fusion</i>	
Kerem Zaman, Leshem Choshen and Shashank Srivastava . . . . .	18763
<i>Collective Critics for Creative Story Generation</i>	
Minwook Bae and Hyoungun Kim . . . . .	18784
<i>Surprise! Uniform Information Density Isn't the Whole Story: Predicting Surprisal Contours in Long-form Discourse</i>	
Eleftheria Tshipidi, Franz Nowak, Ryan Cotterell, Ethan Wilcox, Mario Giulianelli and Alex Warstadt . . . . .	18820
<i>Model-based Preference Optimization in Abstractive Summarization without Human Feedback</i>	
Jaepill Choi, Kyubyung Chae, Jiwoo Song, Yohan Jo and Taesup Kim . . . . .	18837
<i>Are Data Augmentation Methods in Named Entity Recognition Applicable for Uncertainty Estimation?</i>	
Wataru Hashimoto, Hidetaka Kamigaito and Taro Watanabe . . . . .	18852

<i>NeuroTrialNER: An Annotated Corpus for Neurological Diseases and Therapies in Clinical Trial Registries</i>	
Simona Emilova Doneva, Tilia Ellendorff, Beate Sick, Jean-Philippe Goldman, Amelia Elaine Cannon, Gerold Schneider and Benjamin Victor Ineichen . . . . .	18868
<i>Fool Me Once? Contrasting Textual and Visual Explanations in a Clinical Decision-Support Setting</i>	
Maxime Guillaume Kayser, Bayar Menzat, Cornelius Emde, Bogdan Alexandru Bercean, Alex Novak, Abdalá Trinidad Espinosa Morgado, Bartłomiej Papież, Susanne Gaube, Thomas Lukaszewicz and Oana-Maria Camburu . . . . .	18891
<i>Towards Faithful Knowledge Graph Explanation Through Deep Alignment in Commonsense Question Answering</i>	
Weihe Zhai, Arkaitz Zubiaga, Bingquan Liu, Chengjie Sun and Yalong Zhao . . . . .	18920
<i>Generation with Dynamic Vocabulary</i>	
Yanting Liu, Tao Ji, Changzhi Sun, Yuanbin Wu and Xiaoling Wang . . . . .	18931
<i>Argument Relation Classification through Discourse Markers and Adversarial Training</i>	
Michele Luca Contalbo, Francesco Guerra and Matteo Paganelli . . . . .	18949
<i>Getting The Most Out of Your Training Data: Exploring Unsupervised Tasks for Morphological Inflection</i>	
Abhishek Purushothama, Adam Wiemerslage and Katharina Von Der Wense . . . . .	18955
<i>Link, Synthesize, Retrieve: Universal Document Linking for Zero-Shot Information Retrieval</i>	
Dae Yon Hwang, Bilal Taha, Harshit Pande and Yaroslav Nechaev . . . . .	18971
<i>Efficient Unseen Language Adaptation for Multilingual Pre-Trained Language Models</i>	
Po-Heng Chen and Yun-Nung Chen . . . . .	18983
<i>Prove Your Point!: Bringing Proof-Enhancement Principles to Argumentative Essay Generation</i>	
Ruiyu Xiao, Lei Wu, Yuhang Gou, Weinan Zhang and Ting Liu . . . . .	18995
<i>TV-TREES: Multimodal Entailment Trees for Neuro-Symbolic Video Reasoning</i>	
Kate Sanders, Nathaniel Weir and Benjamin Van Durme . . . . .	19009
<i>Unsupervised Extraction of Dialogue Policies from Conversations</i>	
Makesh Narsimhan Sreedhar, Traian Rebedea and Christopher Parisien . . . . .	19029
<i>GRIZAL: Generative Prior-guided Zero-Shot Temporal Action Localization</i>	
Onkar Kishor Susladkar, Gayatri Sudhir Deshmukh, Vandan Gorade and Sparsh Mittal . . . . .	19046
<i>Preserving Multi-Modal Capabilities of Pre-trained VLMs for Improving Vision-Linguistic Compositionality</i>	
Youngtaek Oh, Jae Won Cho, Dong-Jin Kim, In So Kweon and Junmo Kim . . . . .	19060
<i>FoodieQA: A Multimodal Dataset for Fine-Grained Understanding of Chinese Food Culture</i>	
Wenyan Li, Crystina Zhang, Jiaang Li, Qiwei Peng, Raphael Tang, Li Zhou, Weijia Zhang, Guimin Hu, Yifei Yuan, Anders Søgaard, Daniel Hershcovich and Desmond Elliott . . . . .	19077
<i>A Two-Step Approach for Data-Efficient French Pronunciation Learning</i>	
Hoyeon Lee, Hyeon Jang, Jonghwan Kim and Jaemin Kim . . . . .	19096
<i>Exploring Intra and Inter-language Consistency in Embeddings with ICA</i>	
Rongzhi Li, Takeru Matsuda and Hitomi Yanaka . . . . .	19104
<i>DetoxLLM: A Framework for Detoxification with Explanations</i>	
Md Tawkat Islam Khondaker, Muhammad Abdul-Mageed and Laks V. S. Lakshmanan . . . . .	19112

<i>Comparing a BERT Classifier and a GPT classifier for Detecting Connective Language Across Multiple Social Media</i>	Josephine Lukito, Bin Chen, Gina M. Masullo and Natalie Jomini Stroud . . . . .	19140
<i>ShadowLLM: Predictor-based Contextual Sparsity for Large Language Models</i>	Yash Akhauri, Ahmed F AbouElhamayed, Jordan Dotzel, Zhiru Zhang, Alexander M Rush, Safeen Huda and Mohamed S. Abdelfattah . . . . .	19154
<i>Emotion Granularity from Text: An Aggregate-Level Indicator of Mental Health</i>	Krishnapriya Vishnubhotla, Daniela Teodorescu, Mallory J Feldman, Kristen Lindquist and Saif M. Mohammad . . . . .	19168
<i>BLSP-Emo: Towards Empathetic Large Speech-Language Models</i>	Chen Wang, Minpeng Liao, Zhongqiang Huang, Junhong Wu, Chengqing Zong and Jiajun Zhang	19186
<i>SynthesizRR: Generating Diverse Datasets with Retrieval Augmentation</i>	Abhishek Divekar and Greg Durrett . . . . .	19200
<i>Multimodal Self-Instruct: Synthetic Abstract Image and Visual Reasoning Instruction Using Language Model</i>	Wenqi Zhang, Zhenglin Cheng, Yuanyu He, Mengna Wang, Yongliang Shen, Zeqi Tan, Guiyang Hou, Mingqian He, Yanna Ma, Weiming Lu and Yueting Zhuang . . . . .	19228
<i>DataNarrative: Automated Data-Driven Storytelling with Visualizations and Texts</i>	Mohammed Saidul Islam, Md Tahmid Rahman Laskar, Md Rizwan Parvez, Enamul Hoque and Shafiq Joty . . . . .	19253
<i>DEM: Distribution Edited Model for Training with Mixed Data Distributions</i>	Dhananjay Ram, Aditya Rawal, Momchil Hardalov, Nikolaos Pappas and Sheng Zha . . . . .	19287
<i>Altogether: Image Captioning via Re-aligning Alt-text</i>	Hu Xu, Po-Yao Huang, Xiaoqing Tan, Ching-Feng Yeh, Jacob Kahn, Christine Jou, Gargi Ghosh, Omer Levy, Luke Zettlemoyer, Wen-tau Yih, Shang-Wen Li, Saining Xie and Christoph Feichtenhofer	19302
<i>VerifyMatch: A Semi-Supervised Learning Paradigm for Natural Language Inference with Confidence-Aware MixUp</i>	Seo Yeon Park and Cornelia Caragea . . . . .	19319
<i>CaT-Bench: Benchmarking Language Model Understanding of Causal and Temporal Dependencies in Plans</i>	Yash Kumar Lal, Vanya Cohen, Nathanael Chambers, Niranjana Balasubramanian and Ray Mooney	19336
<i>Mitigating the Impact of Reference Quality on Evaluation of Summarization Systems with Reference-Free Metrics</i>	Théo Gigant, Camille Guinaudeau, Marc Decombas and Frederic Dufaux . . . . .	19355
<i>An Empirical Analysis of the Writing Styles of Persona-Assigned LLMs</i>	Manuj Malik, Jing Jiang and Kian Ming A. Chai . . . . .	19369
<i>Investigating the Role of Instruction Variety and Task Difficulty in Robotic Manipulation Tasks</i>	Amit Parekh, Nikolas Vitsakis, Alessandro Suglia and Ioannis Konostas . . . . .	19389
<i>GPT vs RETRO: Exploring the Intersection of Retrieval and Parameter-Efficient Fine-Tuning</i>	Aleksander Ficek, Jiaqi Zeng and Oleksii Kuchaiev . . . . .	19425

<i>CoCoST: Automatic Complex Code Generation with Online Searching and Correctness Testing</i> Xinyi He, Jiaru Zou, Yun Lin, Mengyu Zhou, Shi Han, Zejian Yuan and Dongmei Zhang .	19433
<i>Sequential API Function Calling Using GraphQL Schema</i> Avirup Saha, Lakshmi Mandal, Balaji Ganesan, Sambit Ghosh, Renuka Sindhgatta, Carlos Eberhardt, Dan Debrunner and Sameep Mehta . . . . .	19452
<i>The Illusion of Competence: Evaluating the Effect of Explanations on Users' Mental Models of Visual Question Answering Systems</i> Judith Sieker, Simeon Junker, Ronja Utescher, Nazia Attari, Heiko Wersing, Hendrik Buschmeier and Sina Zarriß . . . . .	19459
<i>Re-Evaluating Evaluation for Multilingual Summarization</i> Jessica Zosa Forde, Ruochen Zhang, Lintang Sutawika, Alham Fikri Aji, Samuel Cahyawijaya, Genta Indra Winata, Minghao Wu, Carsten Eickhoff, Stella Biderman and Ellie Pavlick . . . . .	19476
<i>Video-Text Prompting for Weakly Supervised Spatio-Temporal Video Grounding</i> Heng Zhao, Zhao Yinjie, Bihan Wen, Yew-Soon Ong and Joey Tianyi Zhou . . . . .	19494
<i>A Fast and Sound Tagging Method for Discontinuous Named-Entity Recognition</i> Caio Filippo Corro . . . . .	19506
<i>Factuality of Large Language Models: A Survey</i> Yuxia Wang, Minghan Wang, Muhammad Arslan Manzoor, Fei Liu, Georgi Nenkov Georgiev, Rocktim Jyoti Das and Preslav Nakov . . . . .	19519
<i>Discovering Biases in Information Retrieval Models Using Relevance Thesaurus as Global Explanation</i> Youngwoo Kim, Razieh Rahimi and James Allan . . . . .	19530
<i>Adaptable Moral Stances of Large Language Models on Sexist Content: Implications for Society and Gender Discourse</i> Rongchen Guo, Isar Nejadgholi, Hillary Dawkins, Kathleen C. Fraser and Svetlana Kiritchenko	19548
<i>DISCERN: Decoding Systematic Errors in Natural Language for Text Classifiers</i> Rakesh R Menon and Shashank Srivastava . . . . .	19565
<i>IntCoOp: Interpretability-Aware Vision-Language Prompt Tuning</i> Soumya Suvra Ghosal, Samyadeep Basu, Soheil Feizi and Dinesh Manocha . . . . .	19584
<i>Scope-enhanced Compositional Semantic Parsing for DRT</i> Xiulin Yang, Jonas Groschwitz, Alexander Koller and Johan Bos . . . . .	19602
<i>The Generation Gap: Exploring Age Bias in the Value Systems of Large Language Models</i> Siyang Liu, Trisha Maturi, Bowen Yi, Siqi Shen and Rada Mihalcea . . . . .	19617
<i>TempoFormer: A Transformer for Temporally-aware Representations in Change Detection</i> Talia Tseriotou, Adam Tsakalidis and Maria Liakata . . . . .	19635
<i>Pron vs Prompt: Can Large Language Models already Challenge a World-Class Fiction Author at Creative Text Writing?</i> Guillermo Marco, Julio Gonzalo, M.Teresa Mateo-Girona and Ramón Del Castillo Santos	19654
<i>Evaluating Diversity in Automatic Poetry Generation</i> Yanran Chen, Hannes Gröner, Sina Zarriß and Steffen Eger . . . . .	19671

<i>Evaluating Short-Term Temporal Fluctuations of Social Biases in Social Media Data and Masked Language Models</i>	
Yi Zhou, Danushka Bollegala and Jose Camacho-Collados . . . . .	19693
<i>Delving into Qualitative Implications of Synthetic Data for Hate Speech Detection</i>	
Camilla Casula, Sebastiano Vecellio Salto, Alan Ramponi and Sara Tonelli . . . . .	19709
<i>Grounding Language in Multi-Perspective Referential Communication</i>	
Zineng Tang, Lingjun Mao and Alane Suhr . . . . .	19727
<i>Threshold-driven Pruning with Segmented Maximum Term Weights for Approximate Cluster-based Sparse Retrieval</i>	
Yifan Qiao, Parker Carlson, Shanxiu He, Yingrui Yang and Tao Yang . . . . .	19742
<i>Error Analysis of Multilingual Language Models in Machine Translation: A Case Study of English-Amharic Translation</i>	
Hizkiel Mitiku Alemayehu, Hamada M Zahera and Axel-Cyrille Ngonga Ngomo . . . . .	19758
<i>MIPD: Exploring Manipulation and Intention In a Novel Corpus of Polish Disinformation</i>	
Arkadiusz Modzelewski, Giovanni Da San Martino, Pavel Savov, Magdalena Anna Wilczyńska and Adam Wierzbicki . . . . .	19769
<i>Unsupervised Discrete Representations of American Sign Language</i>	
Artem Abzaliev and Rada Mihalcea . . . . .	19786
<i>Perceptions to Beliefs: Exploring Precursory Inferences for Theory of Mind in Large Language Models</i>	
Chani Jung, Dongkwan Kim, Jiho Jin, Jiseon Kim, Yeon Seonwoo, Yejin Choi, Alice Oh and Hyunwoo Kim . . . . .	19794
<i>Towards Enhancing Coherence in Extractive Summarization: Dataset and Experiments with LLMs</i>	
Mihir Parmar, Hanieh Deilamsalehy, Franck Dernoncourt, Seunghyun Yoon, Ryan A. Rossi and Trung Bui . . . . .	19810
<i>Jump Starting Bandits with LLM-Generated Prior Knowledge</i>	
Parand A. Alamdari, Yanshuai Cao and Kevin H. Wilson . . . . .	19821
<i>Adaptation Odyssey in LLMs: Why Does Additional Pretraining Sometimes Fail to Improve?</i>	
Firat Öncel, Matthias Bethge, Beyza Ermis, Mirco Ravanelli, Cem Subakan and Çağatay Yıldız	19834
<i>Not All Contexts Are Equal: Teaching LLMs Credibility-aware Generation</i>	
Ruotong Pan, Boxi Cao, Hongyu Lin, Xianpei Han, Jia Zheng, Sirui Wang, Xunliang Cai and Le Sun . . . . .	19844
<i>Virtual Personas for Language Models via an Anthology of Backstories</i>	
Suhong Moon, Marwa Abdulhai, Minwoo Kang, Joseph Suh, Widyadewi Soedarmadji, Eran Kohen Behar and David Chan . . . . .	19864
<i>Step-by-Step Reasoning to Solve Grid Puzzles: Where do LLMs Falter?</i>	
Nemika Tyagi, Mihir Parmar, Mohith Kulkarni, Aswin Rrv, Nisarg Patel, Mutsumi Nakamura, Arindam Mitra and Chitta Baral . . . . .	19898
<i>Reasoning in Token Economies: Budget-Aware Evaluation of LLM Reasoning Strategies</i>	
Junlin Wang, Siddhartha Jain, Dejiao Zhang, Baishakhi Ray, Varun Kumar and Ben Athiwaratkun	19916

<i>The Empirical Variability of Narrative Perceptions of Social Media Texts</i>	
Joel Mire, Maria Antoniak, Elliott Ash, Andrew Piper and Maarten Sap .....	19940
<i>Which questions should I answer? Saliency Prediction of Inquisitive Questions</i>	
Yating Wu, Ritika Rajesh Mangla, Alex Dimakis, Greg Durrett and Junyi Jessy Li .....	19969
<i>Revealing Personality Traits: A New Benchmark Dataset for Explainable Personality Recognition on Dialogues</i>	
Lei Sun, Jinming Zhao and Qin Jin .....	19988
<i>Continual Test-time Adaptation for End-to-end Speech Recognition on Noisy Speech</i>	
Guan-Ting Lin, Wei Ping Huang and Hung-yi Lee .....	20003
<i>Whiteboard-of-Thought: Thinking Step-by-Step Across Modalities</i>	
Sachit Menon, Richard Zemel and Carl Vondrick .....	20016
<i>CodeJudge: Evaluating Code Generation with Large Language Models</i>	
Weixi Tong and Tianyi Zhang .....	20032
<i>Self-Training Large Language and Vision Assistant for Medical Question Answering</i>	
Guohao Sun, Can Qin, Huazhu Fu, Linwei Wang and Zhiqiang Tao .....	20052
<i>SYNFAC-EDIT: Synthetic Imitation Edit Feedback for Factual Alignment in Clinical Summarization</i>	
Prakamya Mishra, Zonghai Yao, Parth Vashisht, Feiyun Ouyang , Beining Wang, Vidhi Dhaval Mody and Hong yu .....	20061
<i>Defending Jailbreak Prompts via In-Context Adversarial Game</i>	
Yujun Zhou, Yufei Han, Haomin Zhuang, Kehan Guo, Zhenwen Liang, Hongyan Bao and Xian-gliang Zhang .....	20084
<i>Detecting Online Community Practices with Large Language Models: A Case Study of Pro-Ukrainian Publics on Twitter</i>	
Kateryna Kasianenko, Shima Khanehzar, Stephen Wan, Ehsan Dehghan and Axel Bruns ..	20106
<i>Multilingual Topic Classification in X: Dataset and Analysis</i>	
Dimosthenis Antypas, Asahi Ushio, Francesco Barbieri and Jose Camacho-Collados .....	20136
<i>MT-Eval: A Multi-Turn Capabilities Evaluation Benchmark for Large Language Models</i>	
Wai-Chung Kwan, Xingshan Zeng, Yuxin Jiang, Yufei Wang, Liangyou Li, Lifeng Shang, Xin Jiang, Qun Liu and Kam-Fai Wong .....	20153
<i>Updating CLIP to Prefer Descriptions Over Captions</i>	
Amir Zur, Elisa Kreiss, Karel D'Oosterlinck, Christopher Potts and Atticus Geiger .....	20178
<i>CmdCaliper: A Semantic-Aware Command-Line Embedding Model and Dataset for Security Research</i>	
Sian-Yao Huang, Cheng-Lin Yang, Che-Yu Lin and Chun-Ying Huang .....	20188
<i>Back to School: Translation Using Grammar Books</i>	
Jonathan Hus and Antonios Anastasopoulos .....	20207
<i>VIEWS: Entity-Aware News Video Captioning</i>	
Hammad Ayyubi, Tianqi Liu, Arsha Nagrani, Xudong Lin, Mingda Zhang, Anurag Arnab, Feng Han, Yukun Zhu, Xuande Feng, Kevin Zhang, Jialu Liu and Shih-Fu Chang .....	20220
<i>Towards Aligning Language Models with Textual Feedback</i>	
Sauic Abadal Lloret, Shehzaad Dhuliawala, Keerthiram Murugesan and Mrinmaya Sachan.	20240



<i>AMPO: Automatic Multi-Branched Prompt Optimization</i>	
Sheng Yang, Yurong Wu, Yan Gao, Zineng Zhou, Bin Benjamin Zhu, Xiaodi Sun, Jian-Guang Lou, Zhiming Ding, Anbang Hu, Yuan Fang, Yunsong Li, Junyan Chen and Linjun Yang . . . . .	20267
<i>DeMPT: Decoding-enhanced Multi-phase Prompt Tuning for Making LLMs Be Better Context-aware Translators</i>	
Xinglin Lyu, Junhui Li, Yanqing Zhao, Min Zhang, Daimeng Wei, Shimin Tao, Hao Yang and Min Zhang . . . . .	20280
<i>DEFT-UCS: Data Efficient Fine-Tuning for Pre-Trained Language Models via Unsupervised Core-Set Selection for Text-Editing</i>	
Devleena Das and Vivek Khetan . . . . .	20296
<i>Unveiling Multi-level and Multi-modal Semantic Representations in the Human Brain using Large Language Models</i>	
Yuko Nakagi, Takuya Matsuyama, Naoko Koide-Majima, Hiroto Q. Yamaguchi, Rieko Kubo, Shinji Nishimoto and Yu Takagi . . . . .	20313
<i>They are uncultured": Unveiling Covert Harms and Social Threats in LLM Generated Conversations</i>	
Preetam Prabhu Srikar Dammu, Hayoung Jung, Anjali Singh, Monojit Choudhury and Tanu Mitra	20339
<i>Multi-expert Prompting Improves Reliability, Safety and Usefulness of Large Language Models</i>	
Do Xuan Long, Duong Ngoc Yen, Anh Tuan Luu, Kenji Kawaguchi, Min-Yen Kan and Nancy F. Chen . . . . .	20370
<i>Will LLMs Replace the Encoder-Only Models in Temporal Relation Classification?</i>	
Gabriel Roccabruna, Massimo Rizzoli and Giuseppe Riccardi . . . . .	20402
<i>Eliciting In-Context Learning in Vision-Language Models for Videos Through Curated Data Distributional Properties</i>	
Keunwoo Peter Yu, Zheyuan Zhang, Fengyuan Hu, Shane Storks and Joyce Chai . . . . .	20416
<i>Waterfall: Scalable Framework for Robust Text Watermarking and Provenance for LLMs</i>	
Gregory Kang Ruey Lau, Xinyuan Niu, Hieu Dao, Jiangwei Chen, Chuan-Sheng Foo and Bryan Kian Hsiang Low . . . . .	20432
<i>MASIVE: Open-Ended Affective State Identification in English and Spanish</i>	
Nicholas Deas, Elsbeth Turcan, Ivan Ernesto Perez Mejia and Kathleen McKeown . . . . .	20467
<i>You Make me Feel like a Natural Question: Training QA Systems on Transformed Trivia Questions</i>	
Tasnim Kabir, Yoo Yeon Sung, Saptarashmi Bandyopadhyay, Hao Zou, Abhranil Chandra and Jordan Lee Boyd-Graber . . . . .	20486
<i>AlphaLoRA: Assigning LoRA Experts Based on Layer Training Quality</i>	
Peijun Qing, Chongyang Gao, Yefan Zhou, Xingjian Diao, Yaoqing Yang and Soroush Vosoughi	20511
<i>Flee the Flaw: Annotating the Underlying Logic of Fallacious Arguments Through Templates and Slot-filling</i>	
Irfan Robbani, Paul Reisert, Surawat Pothong, Naoya Inoue, Camélia Guerraoui, Wenzhi Wang, Shoichi Naito, Jungmin Choi and Kentaro Inui . . . . .	20524
<i>Advancing Social Intelligence in AI Agents: Technical Challenges and Open Questions</i>	
Leena Mathur, Paul Pu Liang and Louis-Philippe Morency . . . . .	20541

<i>RAI: Injecting Implicit Bias for Text-To-Image Prompt Refinement Models</i>	
Ziyi Kou, Shichao Pei, Meng Jiang and Xiangliang Zhang .....	20561
<i>Can LLM Generate Culturally Relevant Commonsense QA Data? Case Study in Indonesian and Sundanese</i>	
Rifki Afina Putri, Faiz Ghifari Haznitrama, Dea Adhista and Alice Oh .....	20571
<i>Can Language Models Induce Grammatical Knowledge from Indirect Evidence?</i>	
Miyu Oba, Yohei Oseki, Akiyo Fukatsu, Akari Haga, Hiroki Ouchi, Taro Watanabe and Saku Sugawara .....	20591
<i>Do LLMs Know to Respect Copyright Notice?</i>	
Jialiang Xu, Shenglan LI, Zhaozhuo Xu and Denghui Zhang .....	20604
<i>SpecHub: Provable Acceleration to Multi-Draft Speculative Decoding</i>	
Ryan Sun, Tianyi Zhou, Xun Chen and Lichao Sun .....	20620
<i>Interventional Speech Noise Injection for ASR Generalizable Spoken Language Understanding</i>	
YeonJoon Jung, Jaeseong Lee, Seungtaek Choi, Dohyeon Lee, Minsoo Kim and Seung-won Hwang .....	20642
<i>Rethinking the Role of Proxy Rewards in Language Model Alignment</i>	
Sungdong Kim and Minjoon Seo .....	20656
<i>Visual Text Matters: Improving Text-KVQA with Visual Text Entity Knowledge-aware Large Multimodal Assistant</i>	
Abhirama Subramanyam Penamakuri and Anand Mishra .....	20675
<i>Beyond Correlation: Interpretable Evaluation of Machine Translation Metrics</i>	
Stefano Perrella, Lorenzo Proietti, Pere-Lluís Huguet Cabot, Edoardo Barba and Roberto Navigli	20689
<i>IFCap: Image-like Retrieval and Frequency-based Entity Filtering for Zero-shot Captioning</i>	
Soeun Lee, Si-Woo Kim, Taewhan Kim and Dong-Jin Kim .....	20715
<i>Encoding Spreadsheets for Large Language Models</i>	
Haoyu Dong, Jianbo Zhao, Yuzhang Tian, Junyu Xiong, Mengyu Zhou, Yun Lin, José Cambro-nero, Yeye He, Shi Han and Dongmei Zhang .....	20728
<i>Let's discuss! Quality Dimensions and Annotated Datasets for Computational Argument Quality Assessment</i>	
Rositsa V Ivanova, Thomas Huber and Christina Niklaus .....	20749
<i>Automatic sentence segmentation of clinical record narratives in real-world data</i>	
Dongfang Xu, Davy Weissenbacher, Karen O'Connor, Siddharth Rawal and Graciela Gonzalez Hernandez .....	20780
<i>One-to-Many Communication and Compositionality in Emergent Communication</i>	
Heeyoung Lee .....	20794
<i>Bayesian Example Selection Improves In-Context Learning for Speech, Text and Visual Modalities</i>	
Siyin Wang, Chao-Han Huck Yang, Ji Wu and Chao Zhang .....	20812
<i>Investigating Multilingual Instruction-Tuning: Do Polyglot Models Demand for Multilingual Instructions?</i>	
Alexander Arno Weber, Klaudia Thellmann, Jan Ebert, Nicolas Flores-Herr, Jens Lehmann, Michael Fromm and Mehdi Ali .....	20829

<i>Multi-LogiEval: Towards Evaluating Multi-Step Logical Reasoning Ability of Large Language Models</i> Nisarg Patel, Mohith Kulkarni, Mihir Parmar, Aashna Budhiraja, Mutsumi Nakamura, Neeraj Varshney and Chitta Baral . . . . .	20856
<i>Linear Layer Extrapolation for Fine-Grained Emotion Classification</i> Mayukh Sharma, Sean O’Brien and Julian McAuley . . . . .	20880
<i>Task Oriented In-Domain Data Augmentation</i> Xiao Liang, Xinyu Hu, Simiao Zuo, Yeyun Gong, Qiang Lou, Yi Liu, Shao-Lun Huang and Jian Jiao . . . . .	20889
<i>SciDQA: A Deep Reading Comprehension Dataset over Scientific Papers</i> Shruti Singh, Nandan Sarkar and Arman Cohan . . . . .	20908
<i>Mixture-of-Modules: Reinventing Transformers as Dynamic Assemblies of Modules</i> Zhuocheng Gong, Ang Lv, Jian Guan, Wei Wu, Huishuai Zhang, Minlie Huang, Dongyan Zhao and Rui Yan . . . . .	20924
<i>No Culture Left Behind: ArtELingo-28, a Benchmark of WikiArt with Captions in 28 Languages</i> Youssef Mohamed, Runjia Li, Ibrahim Said Ahmad, Kilichbek Haydarov, Philip Torr, Kenneth Church and Mohamed Elhoseiny . . . . .	20939
<i>PREDICT: Multi-Agent-based Debate Simulation for Generalized Hate Speech Detection</i> Someen Park, Jaehoon Kim, Seungwan Jin, Sohyun Park and Kyungsik Han . . . . .	20963
<i>TokenVerse: Towards Unifying Speech and NLP Tasks via Transducer-based ASR</i> Shashi Kumar, Srikanth Madikeri, Juan Pablo Zuluaga Gomez, Iuliia Thorbecke, Esaú Villatoro-tello, Sergio Burdisso, Petr Motlicek, Karthik Pandia D S and Aravind Ganapathiraju . . . . .	20988
<i>ApiQ: Finetuning of 2-Bit Quantized Large Language Model</i> Baohao Liao, Christian Herold, Shahram Khadivi and Christof Monz . . . . .	20996
<i>Memorize Step by Step: Efficient Long-Context Prefilling with Incremental Memory and Decremental Chunk</i> Zhiyuan Zeng, Qipeng Guo, Xiaoran Liu, Zhangyue Yin, Wentao Shu, Mianqiu Huang, Bo Wang, Yunhua Zhou, Linlin Li, Qun Liu and Xipeng Qiu . . . . .	21021
<i>A Morphology-Based Investigation of Positional Encodings</i> Poulami Ghosh, Shikhar Vashishth, Raj Dabre and Pushpak Bhattacharyya . . . . .	21035
<i>I love pineapple on pizza != I hate pineapple on pizza: Stance-Aware Sentence Transformers for Opinion Mining</i> Vahid Ghafouri, Jose Such and Guillermo Suarez-Tangil . . . . .	21046
<i>BiasWipe: Mitigating Unintended Bias in Text Classifiers through Model Interpretability</i> Mamta Mamta, Rishikant Chigrupaatii and Asif Ekbal . . . . .	21059
<i>ArMeme: Propagandistic Content in Arabic Memes</i> Firoj Alam, Abul Hasnat, Fatema Ahmad, Md. Arid Hasan and Maram Hasanain . . . . .	21071
<i>Language is Scary when Over-Analyzed: Unpacking Implied Misogynistic Reasoning with Argumentation Theory-Driven Prompts</i> Arianna Muti, Federico Ruggeri, Khalid Al Khatib, Alberto Barrón-Cedeño and Tommaso Caselli . . . . .	21091
<i>Thoughts to Target: Enhance Planning for Target-driven Conversation</i> Zhonghua Zheng, Lizi Liao, Yang Deng, Ee-Peng Lim, Minlie Huang and Liqiang Nie . . . . .	21108

<i>Scalable Data Ablation Approximations for Language Models through Modular Training and Merging</i> Clara Na, Ian Magnusson, Ananya Harsh Jha, Tom Sherborne, Emma Strubell, Jesse Dodge and Pradeep Dasigi .....	21125
<i>Exploring Intrinsic Language-specific Subspaces in Fine-tuning Multilingual Neural Machine Translation</i> Zhe Cao, Zhi Qu, Hidetaka Kamigaito and Taro Watanabe .....	21142
<i>Attention Score is not All You Need for Token Importance Indicator in KV Cache Reduction: Value Also Matters</i> Zhiyu Guo, Hidetaka Kamigaito and Taro Watanabe .....	21158
<i>Generative Subgraph Retrieval for Knowledge Graph–Grounded Dialog Generation</i> Jinyoung Park, Minseok Joo, Joo-Kyung Kim and Hyunwoo J. Kim .....	21167
<i>Adapters Mixup: Mixing Parameter-Efficient Adapters to Enhance the Adversarial Robustness of Fine-tuned Pre-trained Text Classifiers</i> Tuc Van Nguyen and Thai Le .....	21183
<i>Generalizing Clinical De-identification Models by Privacy-safe Data Augmentation using GPT-4</i> Woojin Kim, Sungeun Hahm and Jaejin Lee .....	21204
<i>Connecting the Dots: Evaluating Abstract Reasoning Capabilities of LLMs Using the New York Times Connections Word Game</i> Prisha Samdarshi, Mariam Mustafa, Anushka Kulkarni, Raven Rothkopf, Tuhin Chakrabarty and Smaranda Muresan .....	21219
<i>GottBERT: a pure German Language Model</i> Raphael Scheible, Johann Frei, Fabian Thomczyk, Henry He, Patric Tippmann, Jochen Knaus, Victor Jaravine, Frank Kramer and Martin Boeker .....	21237
<i>Computational Meme Understanding: A Survey</i> Khoi P. N. Nguyen and Vincent Ng .....	21251
<i>CoverICL: Selective Annotation for In-Context Learning via Active Graph Coverage</i> Costas Mavromatis, Balasubramaniam Srinivasan, Zhengyuan Shen, Jiani Zhang, Huzefa Rangwala, Christos Faloutsos and George Karypis .....	21268
<i>Retrieval-enriched zero-shot image classification in low-resource domains</i> Nicola Dall’Asen, Yiming Wang, Enrico Fini and Elisa Ricci .....	21287
<i>I-AM-G: Interest Augmented Multimodal Generator for Item Personalization</i> Xianquan Wang, Likang Wu, Shukang Yin, Zhi Li, Yanjiang Chen, Hufeng Hufeng, Yu Su and Qi Liu .....	21303
<i>Twists, Humps, and Pebbles: Multilingual Speech Recognition Models Exhibit Gender Performance Gaps</i> Giuseppe Attanasio, Beatrice Savoldi, Dennis Fucci and Dirk Hovy .....	21318
<i>Enhancing Language Model Alignment: A Confidence-Based Approach to Label Smoothing</i> Baihe Huang, Hiteshi Sharma and Yi Mao .....	21341
<i>Contrastive Policy Gradient: Aligning LLMs on sequence-level scores in a supervised-friendly fashion</i> Yannis Flet-Berliac, Nathan Grinsztajn, Florian Strub, Eugene Choi, Bill Wu, Chris Cremer, Arash Ahmadian, Yash Chandak, Mohammad Gheshlaghi Azar, Olivier Pietquin and Matthieu Geist ..	21353

<i>Show and Guide: Instructional-Plan Grounded Vision and Language Model</i> Diogo Glória-Silva, David Semedo and Joao Magalhaes .....	21371
<i>Beyond Turn-Based Interfaces: Synchronous LLMs as Full-Duplex Dialogue Agents</i> Bandhav Veluri, Benjamin N Peloquin, Bokai YU, Hongyu Gong and Shyamnath Gollakota	21390
<i>QuBE: Question-based Belief Enhancement for Agentic LLM Reasoning</i> Minsoo Kim, Jongyoon Kim, Jihyuk Kim and Seung-won Hwang.....	21403
<i>CompAct: Compressing Retrieved Documents Actively for Question Answering</i> Chanwoong Yoon, Taewhoo Lee, Hyeon Hwang, Minbyul Jeong and Jaewoo Kang .....	21424
<i>An Empirical Analysis on Spatial Reasoning Capabilities of Large Multimodal Models</i> Fatemeh Shiri, Xiao-Yu Guo, Mona Golestan Far, Xin Yu, Reza Haf and Yuan-Fang Li ...	21440
<i>Synthetic Knowledge Ingestion: Towards Knowledge Refinement and Injection for Enhancing Large Language Models</i> Jiaxin Zhang, Wendi Cui, Yiran Huang, Kamalika Das and Sricharan Kumar .....	21456
<i>Local Contrastive Editing of Gender Stereotypes</i> Marlene Lutz, Rochelle Choenni, Markus Strohmaier and Anne Lauscher .....	21474
<i>De-Identification of Sensitive Personal Data in Datasets Derived from IIT-CDIP</i> Stefan Larson, Nicole Cornehl Lima, Santiago Pedroza Diaz, Amogh Manoj Joshi, Siddharth Betala, Jamiu Tunde Suleiman, Yash Mathur, Kaushal Kumar Prajapati, Ramla Alakraa, Junjie Shen, Temi Okotore and Kevin Leach.....	21494
<i>RAR: Retrieval-augmented retrieval for code generation in low resource languages</i> Avik Dutta, Mukul Singh, Gust Verbruggen, Sumit Gulwani and Vu Le .....	21506
<i>STAR: SocioTechnical Approach to Red Teaming Language Models</i> Laura Weidinger, John F J Mellor, Bernat Guillén Pegueroles, Nahema Marchal, Ravin Kumar, Kristian Lum, Canfer Akbulut, Mark Diaz, A. Stevie Bergman, Mikel D. Rodriguez, Verena Rieser and William Isaac.....	21516
<i>Do great minds think alike? Investigating Human-AI Complementarity in Question Answering with CAIMIRA</i> Maharshi Gor, Hal Daumé Iii, Tianyi Zhou and Jordan Lee Boyd-Graber .....	21533
<i>Memory-Efficient Fine-Tuning of Transformers via Token Selection</i> Antoine Simoulin, Namyong Park, Xiaoyi Liu and Grey Yang .....	21565
<i>Unveiling the mystery of visual attributes of concrete and abstract concepts: Variability, nearest neighbors, and challenging categories</i> Tarun Tater, Sabine Schulte Im Walde and Diego Frassinelli.....	21581
<i>Evaluating Large Language Models on Time Series Feature Understanding: A Comprehensive Taxonomy and Benchmark</i> Elizabeth Fons, Rachneet Kaur, Soham Palande, Zhen Zeng, Tucker Balch, Manuela Veloso and Svitlana Vyetenko.....	21598
<i>Can LLMs Learn Uncertainty on Their Own? Expressing Uncertainty Effectively in A Self-Training Manner</i> Shudong Liu, Zhaocong Li, Xuebo Liu, Runzhe Zhan, Derek F. Wong, Lidia S. Chao and Min Zhang .....	21635

<i>Preference-Guided Reflective Sampling for Aligning Language Models</i>	
Hai Ye and Hwee Tou Ng .....	21646
<i>Metrics for What, Metrics for Whom: Assessing Actionability of Bias Evaluation Metrics in NLP</i>	
Pieter Delobelle, Giuseppe Attanasio, Debora Nozza, Su Lin Blodgett and Zeerak Talat ...	21669
<i>Is this the real life? Is this just fantasy? The Misleading Success of Simulating Social Interactions With LLMs</i>	
Xuhui Zhou, Zhe Su, Tiwalayo Eisape, Hyunwoo Kim and Maarten Sap .....	21692
<i>A Simple LLM Framework for Long-Range Video Question-Answering</i>	
Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal and Gedas Bertasius .....	21715
<i>Rebuilding ROME : Resolving Model Collapse during Sequential Model Editing</i>	
Akshat Gupta, Sidharth Baskaran and Gopala Anumanchipalli .....	21738
<i>Casablanca: Data and Models for Multidialectal Arabic Speech Recognition</i>	
Bashar Talafha, Karima Kadaoui, Samar Mohamed Magdy, Mariem Habiboullah, Chafei Mohamed Chafei, Ahmed Oumar El-Shangiti, Hiba Zayed, Mohamedou Cheikh Tourad, Rahaf Alhamouri, Rwa Assi, Aisha Alraeesi, Hour Mohamed, Fakhraddin Alwajih, Abdelrahman Mohamed, Abdellah El Mekki, El Moatez Billah Nagoudi, Benelhadj Djelloul Mama Saadia, Hamzah A. Alsayadi, Walid Al-Dhabyani, Sara Shatnawi, Yasir Ech-chammakhy, Amal Makouar, Yousra Berrachedi, Mustafa Jarrar, Shady Shehata, Ismail Berrada and Muhammad Abdul-Mageed .....	21745
<i>Safety Arithmetic: A Framework for Test-time Safety Alignment of Language Models by Steering Parameters and Activations</i>	
Rima Hazra, Sayan Layek, Somnath Banerjee and Soujanya Poria .....	21759
<i>Communicating with Speakers and Listeners of Different Pragmatic Levels</i>	
Kata Naszadi, Frans A Oliehoek and Christof Monz .....	21777
<i>RECANTFormer: Referring Expression Comprehension with Varying Numbers of Targets</i>	
Bhathiya Hemanthage, Hakan Bilen, Phil Bartie, Christian Dondrup and Oliver Lemon ...	21784
<i>Sprout: Green Generative AI with Carbon-Efficient LLM Inference</i>	
Baolin Li, Yankai Jiang, Vijay Gadepally and Devesh Tiwari .....	21799
<i>Do LLMs Plan Like Human Writers? Comparing Journalist Coverage of Press Releases with LLMs</i>	
Alexander Spangher, Nanyun Peng, Sebastian Gehrmann and Mark Dredze .....	21814
<i>T-FREE: Subword Tokenizer-Free Generative LLMs via Sparse Representations for Memory-Efficient Embeddings</i>	
Björn Deiseroth, Manuel Brack, Patrick Schramowski, Kristian Kersting and Samuel Weinbach	21829
<i>SpeechQE: Estimating the Quality of Direct Speech Translation</i>	
HyoJung Han, Kevin Duh and Marine Carpuat .....	21852
<i>Assessing and Verifying Task Utility in LLM-Powered Applications</i>	
Negar Arabzadeh, Siqing Huo, Nikhil Mehta, Qingyun Wu, Chi Wang, Ahmed Hassan Awadallah, Charles L. A. Clarke and Julia Kiseleva .....	21868
<i>Dynamic Rewarding with Prompt Optimization Enables Tuning-free Self-Alignment of Language Models</i>	
Somanshu Singla, Zhen Wang, Tianyang Liu, Abdullah Ashfaq, Zhiting Hu and Eric P. Xing	21889

<i>Accurate and Data-Efficient Toxicity Prediction when Annotators Disagree</i> Harbani Jaggi, Kashyap Coimbatore Murali, Eve Fleisig and Erdem Biyik . . . . .	21910
<i>Adversarial Text Generation using Large Language Models for Dementia Detection</i> Youxiang Zhu, Nana Lin, Kiran Sandilya Balivada, Daniel Haehn and Xiaohui Liang . . . . .	21918
<i>xCOMET-lite: Bridging the Gap Between Efficiency and Quality in Learned MT Evaluation Metrics</i> Daniil Larionov, Mikhail Seleznyov, Vasiliy Viskov, Alexander Panchenko and Steffen Eger	21934
<i>The Greatest Good Benchmark: Measuring LLMs' Alignment with Utilitarian Moral Dilemmas</i> Giovanni Franco Gabriel Marraffini, Andrés Cotton, Noe Fabian Hsueh, Axel Fridman, Juan Wisznia and Luciano Del Corro . . . . .	21950
<i>FairFlow: Mitigating Dataset Biases through Undecided Learning for Natural Language Understanding</i> Jiali Cheng and Hadi Amiri . . . . .	21960
<i>Style-Shifting Behaviour of the Manosphere on Reddit</i> Jai Aggarwal and Suzanne Stevenson . . . . .	21976
<i>The Death and Life of Great Prompts: Analyzing the Evolution of LLM Prompts from the Structural Perspective</i> Yihan Ma, Xinyue Shen, Yixin Wu, Boyang Zhang, Michael Backes and Yang Zhang . . . . .	21990
<i>Holistic Evaluation for Interleaved Text-and-Image Generation</i> Minqian Liu, Zhiyang Xu, Zihao Lin, Trevor Ashby, Joy Rimchala, Jiaxin Zhang and Lifu Huang	22002
<i>FOLIO: Natural Language Reasoning with First-Order Logic</i> Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, Lucy Sun, Alexander Wardle-Solano, Hannah Szabó, Ekaterina Zubova, Matthew Burtell, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, Ansong Ni, Linyong Nan, Jungo Kasai, Tao Yu, Rui Zhang, Alexander Fabbri, Wojciech Maciej Kryscinski, Semih Yavuz, Ye Liu, Xi Victoria Lin, Shafiq Joty, Yingbo Zhou, Caiming Xiong, Rex Ying, Arman Cohan and Dragomir Radev . . . . .	22017
<i>The LLM Effect: Are Humans Truly Using LLMs, or Are They Being Influenced By Them Instead?</i> Alexander Choi, Syeda Sabrina Akter, J.p. Singh and Antonios Anastasopoulos . . . . .	22032
<i>Is Child-Directed Speech Effective Training Data for Language Models?</i> Steven Y. Feng, Noah Goodman and Michael Frank . . . . .	22055
<i>RevMUX: Data Multiplexing with Reversible Adapters for Efficient LLM Batch Inference</i> Yige Xu, Xu Guo, Zhiwei Zeng and Chunyan Miao . . . . .	22072
<i>Inference Helps PLMs' Conceptual Understanding: Improving the Abstract Inference Ability with Hierarchical Conceptual Entailment Graphs</i> Juncai Li, Ru Li, Xiaoli Li, Qinghua Chai and Jeff Z. Pan . . . . .	22088
<i>M3Hop-CoT: Misogynous Meme Identification with Multimodal Multi-hop Chain-of-Thought</i> Gitanjali Kumari, Kirtan Jain and Asif Ekbal . . . . .	22105
<i>GPT-4 Jailbreaks Itself with Near-Perfect Success Using Self-Explanation</i> Govind Ramesh, Yao Dou and Wei Xu . . . . .	22139
<i>RE-RAG: Improving Open-Domain QA Performance and Interpretability with Relevance Estimator in Retrieval-Augmented Generation</i> Kiseung Kim and Jay-Yoon Lee . . . . .	22149

<i>Evaluating Concurrent Robustness of Language Models Across Diverse Challenge Sets</i> Vatsal Gupta, Pranshu Pandya, Tushar Kataria, Vivek Gupta and Dan Roth .....	22162
<i>Simul-MuST-C: Simultaneous Multilingual Speech Translation Corpus Using Large Language Model</i> Mana Makinae, Yusuke Sakai, Hidetaka Kamigaito and Taro Watanabe .....	22185
<i>Is This a Bad Table? A Closer Look at the Evaluation of Table Generation from Text</i> Pritika Ramu, Aparna Garimella and Sambaran Bandyopadhyay .....	22206
<i>On the Fragility of Active Learners for Text Classification</i> Abhishek Ghose and Emma Thuong Nguyen .....	22217
<i>BMRetriever: Tuning Large Language Models as Better Biomedical Text Retrievers</i> Ran Xu, Wenqi Shi, Yue Yu, Yuchen Zhuang, Yanqiao Zhu, May Dongmei Wang, Joyce C. Ho, Chao Zhang and Carl Yang .....	22234
<i>Comparing Neighbors Together Makes it Easy: Jointly Comparing Multiple Candidates for Efficient and Effective Retrieval</i> Jonghyun Song, Cheyon Jin, Wenlong Zhao, Andrew McCallum and Jay-Yoon Lee .....	22255
<i>M3D: MultiModal MultiDocument Fine-Grained Inconsistency Detection</i> Chia-Wei Tang, Ting-Chih Chen, Kiet A. Nguyen, Kazi Sajeed Mehrab, Alvi Md Ishmam and Chris Thomas .....	22270
<i>MedAdapter: Efficient Test-Time Adaptation of Large Language Models Towards Medical Reasoning</i> Wenqi Shi, Ran Xu, Yuchen Zhuang, Yue Yu, Haotian Sun, Hang Wu, Carl Yang and May Dong- mei Wang .....	22294
<i>EHRAgent: Code Empowers Large Language Models for Few-shot Complex Tabular Reasoning on Electronic Health Records</i> Wenqi Shi, Ran Xu, Yuchen Zhuang, Yue Yu, Jieyu Zhang, Hang Wu, Yuanda Zhu, Joyce C. Ho, Carl Yang and May Dongmei Wang .....	22315
<i>SimLLM: Detecting Sentences Generated by Large Language Models Using Similarity between the Generation and its Re-generation</i> Hoang-Quoc Nguyen-Son, Minh-Son Dao and Koji Zettsu .....	22340
<i>CELLO: Causal Evaluation of Large Vision-Language Models</i> Meiqi Chen, Bo Peng, Yan Zhang and Chaochao Lu .....	22353
<i>Simultaneous Interpretation Corpus Construction by Large Language Models in Distant Language Pair</i> Yusuke Sakai, Mana Makinae, Hidetaka Kamigaito and Taro Watanabe .....	22375
<i>Training-free Deep Concept Injection Enables Language Models for Video Question Answering</i> Xudong Lin, Manling Li, Richard Zemel, Heng Ji and Shih-Fu Chang .....	22399
<i>MI Bench: Evaluating Multimodal Large Language Models over Multiple Images</i> Haowei Liu, Xi Zhang, Haiyang Xu, Yaya Shi, Chaoya Jiang, Ming Yan, Ji Zhang, Fei Huang, Chunfeng Yuan, Bing Li and Weiming Hu .....	22417
<i>ZEBRA: Zero-Shot Example-Based Retrieval Augmentation for Commonsense Question Answering</i> Francesco Maria Molfese, Simone Conia, Riccardo Orlando and Roberto Navigli .....	22429
<i>ABLE: Personalized Disability Support with Politeness and Empathy Integration</i> Kshitij Mishra, Manisha Burja and Asif Ekbal .....	22445



<i>Language Models as Compilers: Simulating Pseudocode Execution Improves Algorithmic Reasoning in Language Models</i>	
Hyunjoo Chae, Yeonghyeon Kim, Seungone Kim, Kai Tzu-iunn Ong, Beong-woo Kwak, Moo-hyeon Kim, Sunghwan Kim, Taeyoon Kwon, Jiwan Chung, Youngjae Yu and Jinyoung Yeo . . . .	22471
<i>Coffee-Gym: An Environment for Evaluating and Improving Natural Language Feedback on Erroneous Code</i>	
Hyunjoo Chae, Taeyoon Kwon, Seungjun Moon, Yongho Song, Dongjin Kang, Kai Tzu-iunn Ong, Beong-woo Kwak, Seonghyeon Bae, Seung-won Hwang and Jinyoung Yeo . . . . .	22503
<i>Improving Minimum Bayes Risk Decoding with Multi-Prompt</i>	
David Heineman, Yao Dou and Wei Xu . . . . .	22525
<i>Deciphering Cognitive Distortions in Patient-Doctor Mental Health Conversations: A Multimodal LLM-Based Detection and Reasoning Framework</i>	
Gopendra Vikram Singh, Sai Vardhan Vemulapalli, Mauajama Firdaus and Asif Ekbal . . . .	22546
<i>Nearest Neighbor Normalization Improves Multimodal Retrieval</i>	
Neil Chowdhury, Franklin Wang, Sumedh Shenoy, Douwe Kiela, Sarah Schwettmann and Tristan Thrush . . . . .	22571
<i>Rethinking Pragmatics in Large Language Models: Towards Open-Ended Evaluation and Preference Tuning</i>	
Shengguang Wu, Shusheng Yang, Zhenglun Chen and Qi Su . . . . .	22583
<i>LongRAG: A Dual-Perspective Retrieval-Augmented Generation Paradigm for Long-Context Question Answering</i>	
Qingfei Zhao, Ruobing Wang, Yukuo Cen, Daren Zha, Shicheng Tan, Yuxiao Dong and Jie Tang	22600
<i>Context-aware Watermark with Semantic Balanced Green-red Lists for Large Language Models</i>	
Yuxuan Guo, Zhiliang Tian, Yiping Song, Tianlun Liu, Liang Ding and Dongsheng Li . . . .	22633
<i>Knowledge Graph Enhanced Large Language Model Editing</i>	
Mengqi Zhang, Xiaotian Ye, Qiang Liu, Pengjie Ren, Shu Wu and Zhumin Chen . . . . .	22647
<i>'Quis custodiet ipsos custodes?' Who will watch the watchmen? On Detecting AI-generated peer-reviews</i>	
Sandeep Kumar, Mohit Sahu, Vardhan Gacche, Tirthankar Ghosal and Asif Ekbal . . . . .	22663
<i>Mitigating Open-Vocabulary Caption Hallucinations</i>	
Assaf Ben-Kish, Moran Yanuka, Morris Alper, Raja Giryes and Hadar Averbuch-Elor . . . .	22680
<i>Initialization of Large Language Models via Reparameterization to Mitigate Loss Spikes</i>	
Kosuke Nishida, Kyosuke Nishida and Kuniko Saito . . . . .	22699
<i>ALVIN: Active Learning Via INterpolation</i>	
Michalis Korakakis, Andreas Vlachos and Adrian Weller . . . . .	22715
<i>Filtered Direct Preference Optimization</i>	
Tetsuro Morimura, Mitsuki Sakamoto, Yuu Jinnai, Kenshi Abe and Kaito Ariu . . . . .	22729
<i>Instruction Fine-Tuning: Does Prompt Loss Matter?</i>	
Mathew Huerta-Enochian and Seung Yong Ko . . . . .	22771
<i>Entity Insertion in Multilingual Linked Corpora: The Case of Wikipedia</i>	
Tomás Feith, Akhil Arora, Martin Gerlach, Debjit Paul and Robert West . . . . .	22796