

# **Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2024)**

Languages of Latin America

Chapter 1: Long Papers

Mexico City, Mexico  
16 – 21 June 2024

Volume 1 of 12

ISBN: 978-1-7138-9894-8

**Printed from e-media with permission by:**

Curran Associates, Inc.  
57 Morehouse Lane  
Red Hook, NY 12571



**Some format issues inherent in the e-media version may also appear in this print version.**

Copyright© (2024) by the Association for Computational Linguistics  
All rights reserved.

Printed with permission by Curran Associates, Inc. (2024)

For permission requests, please contact the Association for Computational Linguistics  
at the address below.

Association for Computational Linguistics  
209 N. Eighth Street  
Stroudsburg, Pennsylvania 18360

Phone: 1-570-476-8006

Fax: 1-570-476-0860

[acl@aclweb.org](mailto:acl@aclweb.org)

**Additional copies of this publication are available from:**

Curran Associates, Inc.  
57 Morehouse Lane  
Red Hook, NY 12571 USA  
Phone: 845-758-0400  
Fax: 845-758-2633  
Email: [curran@proceedings.com](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

## Table of Contents

<i>Named Entity Recognition Under Domain Shift via Metric Learning for Life Sciences</i> Hongyi Liu, Qingyun Wang, Payam Karisani and Heng Ji .....	1
<i>Text Diffusion Model with Encoder-Decoder Transformers for Sequence-to-Sequence Generation</i> Hongyi Yuan, Zheng Yuan, Chuanqi Tan, Fei Huang and Songfang Huang .....	22
<i>An Interactive Framework for Profiling News Media Sources</i> Nikhil Mehta and Dan Goldwasser .....	40
<i>Assessing Logical Puzzle Solving in Large Language Models: Insights from a Minesweeper Case Study</i> Yinghao Li, Haorui Wang and Chao Zhang .....	59
<i>TelME: Teacher-leading Multimodal Fusion Network for Emotion Recognition in Conversation</i> Taeyang Yun, Hyunkuk Lim, Jeonghwan Lee and Min Song .....	82
<i>Effective and Efficient Conversation Retrieval for Dialogue State Tracking with Implicit Text Summaries</i> Seanie Lee, Jianpeng Cheng, Joris Driesen, Alexandru Coca and Anders Johannsen .....	96
<i>Promptly Predicting Structures: The Return of Inference</i> Maitrey Mehta, Valentina Pyatkin and Vivek Srikumar .....	112
<i>On Linearizing Structured Data in Encoder-Decoder Language Models: Insights from Text-to-SQL</i> Yutong Shao and Ndapa Nakashole .....	131
<i>Extractive Summarization with Text Generator</i> Thang Le and Anh Tuan Luu .....	157
<i>Self-generated Replay Memories for Continual Neural Machine Translation</i> Michele Resta and Davide Bacciu .....	175
<i>Measuring and Improving Chain-of-Thought Reasoning in Vision-Language Models</i> Yangyi Chen, Karan Sikka, Michael Cogswell, Heng Ji and Ajay Divakaran .....	192
<i>Building Knowledge-Guided Lexica to Model Cultural Variation</i> Shreya Havaldar, Salvatore Giorgi, Sunny Rai, Thomas Talhelm, Sharath Chandra Guntuku and Lyle Ungar .....	211
<i>Adaptive Rank Selections for Low-Rank Approximation of Language Models</i> Shangqian Gao, Ting Hua, Yen-Chang Hsu, Yilin Shen and Hongxia Jin .....	227
<i>An Empirical Study of Consistency Regularization for End-to-End Speech-to-Text Translation</i> Pengzhi Gao, Ruiqing Zhang, Zhongjun He, Hua Wu and Haifeng Wang .....	242
<i>Unleashing the Emergent Cognitive Synergy in Large Language Models: A Task-Solving Agent through Multi-Persona Self-Collaboration</i> Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei and Heng Ji .....	257
<i>FPT: Feature Prompt Tuning for Few-shot Readability Assessment</i> Ziyang Wang, Sanwoo Lee, Hsiu-Yuan Huang and Yunfang Wu .....	280
<i>Self-Prompting Large Language Models for Zero-Shot Open-Domain QA</i> Junlong Li, Jinyuan Wang, Zhuosheng Zhang and hai zhao .....	296

<i>Head-to-Tail: How Knowledgeable are Large Language Models (LLMs)? A.K.A. Will LLMs Replace Knowledge Graphs?</i>	
Kai Sun, Yifan Ethan Xu, Hanwen Zha, Yue Liu and Xin Luna Dong .....	311
<i>kNN-ICL: Compositional Task-Oriented Parsing Generalization with Nearest Neighbor In-Context Learning</i>	
Wenting Zhao, Ye Liu, Yao Wan, Yibo Wang, Qingyang Wu, Zhongfen Deng, Jiangshu Du, Shuaiqi Liu, Yunlong Xu and Philip S. Yu .....	326
<i>ARES: An Automated Evaluation Framework for Retrieval-Augmented Generation Systems</i>	
Jon Saad-Falcon, Omar Khattab, Christopher Potts and Matei Zaharia .....	338
<i>DEMO: A Statistical Perspective for Efficient Image-Text Matching</i>	
Fan Zhang, Xian-Sheng Hua, Chong Chen and Xiao Luo .....	355
<i>SeaEval for Multilingual Foundation Models: From Cross-Lingual Alignment to Cultural Reasoning</i>	
Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, AiTi Aw and Nancy F. Chen	370
<i>Volcano: Mitigating Multimodal Hallucination through Self-Feedback Guided Revision</i>	
Seongyun Lee, Sue Hyun Park, Yongrae Jo and Minjoon Seo .....	391
<i>LLMs Are Few-Shot In-Context Low-Resource Language Learners</i>	
Samuel Cahyawijaya, Holy Lovenia and Pascale Fung .....	405
<i>Simple and effective data augmentation for compositional generalization</i>	
Yuekun Yao and Alexander Koller .....	434
<i>Rethinking Tabular Data Understanding with Large Language Models</i>	
Tianyang Liu, Fei Wang and Muhao Chen .....	450
<i>From Shortcuts to Triggers: Backdoor Defense with Denoised PoE</i>	
Qin Liu, Fei Wang, Chaowei Xiao and Muhao Chen .....	483
<i>BookSQL: A Large Scale Text-to-SQL Dataset for Accounting Domain</i>	
Rahul Kumar, Amar Raja Dibbu, Shrutendra Harsola, Vignesh Subrahmaniam and Ashutosh Modi	497
<i>FLAP: Flow-Adhering Planning with Constrained Decoding in LLMs</i>	
Shamik Roy, Sailik Sengupta, Daniele Bonadiman, Saab Mansour and Arshit Gupta .....	517
<i>DuRE: Dual Contrastive Self Training for Semi-Supervised Relation Extraction</i>	
Yuxi Feng and Laks V. S. Lakshmanan .....	540
<i>Query-Efficient Textual Adversarial Example Generation for Black-Box Attacks</i>	
Zhen Yu, Zhenhua Chen and Kun He .....	556
<i>Embrace Divergence for Richer Insights: A Multi-document Summarization Benchmark and a Case Study on Summarizing Diverse Information from News Articles</i>	
Kung-Hsiang Huang, Philippe Laban, Alexander Fabbri, Prafulla Kumar Choubey, Shafiq Joty, Caiming Xiong and Chien-Sheng Wu .....	570
<i>AMRFact: Enhancing Summarization Factuality Evaluation with AMR-Driven Negative Samples Generation</i>	
Haoyi Qiu, Kung-Hsiang Huang, Jingnong Qu and Nanyun Peng .....	594
<i>PILOT: Legal Case Outcome Prediction with Case Law</i>	
Lang Cao, Zifeng Wang, Cao Xiao and Jimeng Sun .....	609

<i>ALoRA: Allocating Low-Rank Adaptation for Fine-tuning Large Language Models</i> Zequan Liu, Jiawen Lyn, Wei Zhu and Xing Tian . . . . .	622
<i>R-Spin: Efficient Speaker and Noise-invariant Representation Learning with Acoustic Pieces</i> Heng-Jui Chang and James R. Glass . . . . .	642
<i>InsCL: A Data-efficient Continual Learning Paradigm for Fine-tuning Large Language Models with Instructions</i> Yifan Wang, Yafei Liu, Chufan Shi, Haoling Li, Chen Chen, Haonan Lu and Yujiu Yang . . . . .	663
<i>Language Agnostic Code Embeddings</i> Saiteja Utpala, Alex Gu and Pin-Yu Chen . . . . .	678
<i>An Examination of the Compositionality of Large Generative Vision-Language Models</i> Teli Ma, Rong Li and Junwei Liang . . . . .	692
<i>Two Heads are Better than One: Nested PoE for Robust Defense Against Multi-Backdoors</i> Victoria Graf, Qin Liu and Muhao Chen . . . . .	706
<i>VertAttack: Taking Advantage of Text Classifiers' Horizontal Vision</i> Jonathan Rusert . . . . .	719
<i>KDMCSE: Knowledge Distillation Multimodal Sentence Embeddings with Adaptive Angular margin Contrastive Learning</i> Cong-Duy T Nguyen, Thong Thanh Nguyen, Xiaobao Wu and Anh Tuan Luu . . . . .	733
<i>The taste of IPA: Towards open-vocabulary keyword spotting and forced alignment in any language</i> Jian Zhu, Changbing Yang, Farhan Samir and Jahurul Islam . . . . .	750
<i>Think Before You Act: A Two-Stage Framework for Mitigating Gender Bias Towards Vision-Language Tasks</i> Yunqi Zhang, Songda Li, Chunyuan Deng, Luyi Wang and Hui Zhao . . . . .	773
<i>BeLLM: Backward Dependency Enhanced Large Language Model for Sentence Embeddings</i> Xianming LI and Jing Li . . . . .	792
<i>Assessing Factual Reliability of Large Language Model Knowledge</i> Weixuan Wang, Barry Haddow, Alexandra Birch and Wei Peng . . . . .	805
<i>Dial-MAE: ConTextual Masked Auto-Encoder for Retrieval-based Dialogue Systems</i> Zhenpeng Su, Xing W, Wei Zhou, Guangyuan Ma and Songlin Hu . . . . .	820
<i>Toolink: Linking Toolkit Creation and Using through Chain-of-Solving on Open-Source Model</i> Cheng Qian, Chenyan Xiong, Zhenghao Liu and Zhiyuan Liu . . . . .	831
<i>Create! Don't Repeat: A Paradigm Shift in Multi-Label Augmentation through Label Creative Generation</i> Letian Wang, Xianggen Liu and Jiancheng Lv . . . . .	855
<i>Neurocache: Efficient Vector Retrieval for Long-range Language Modeling</i> Ali Safaya and Deniz Yuret . . . . .	870
<i>Unveiling the Generalization Power of Fine-Tuned Large Language Models</i> Haoran Yang, Yumeng Zhang, Jiaqi Xu, Hongyuan Lu, Pheng-Ann Heng and Wai Lam . . . . .	884
<i>A Closer Look at the Self-Verification Abilities of Large Language Models in Logical Reasoning</i> Ruixin Hong, Hongming Zhang, Xinyu Pang, Dong Yu and Changshui Zhang . . . . .	900

<i>Exploring Self-supervised Logic-enhanced Training for Large Language Models</i> Fangkai Jiao, Zhiyang Teng, Bosheng Ding, Zhengyuan Liu, Nancy F. Chen and Shafiq Joty .	926
<i>MATHSENSEI: A Tool-Augmented Large Language Model for Mathematical Reasoning</i> Debrup Das, Debopriyo Banerjee, Somak Aditya and Ashish Kulkarni . . . . .	942
<i>CoUDA: Coherence Evaluation via Unified Data Augmentation</i> Dawei Zhu, Wenhao Wu, Yifan Song, Fangwei Zhu, Ziqiang Cao and Sujian Li . . . . .	967
<i>mEdIT: Multilingual Text Editing via Instruction Tuning</i> Vipul Raheja, Dimitris Alikaniotis, Vivek Kulkarni, Bashar Alhafni and Dhruv Kumar . . . . .	979
<i>Navigation as Attackers Wish? Towards Building Robust Embodied Agents under Federated Learning</i> Yunchao Zhang, Zonglin Di, Kaiwen Zhou, Cihang Xie and Xin Eric Wang. . . . .	1002
<i>In-context Learning and Gradient Descent Revisited</i> Gilad Deutch, Nadav Magar, Tomer Bar Natan and Guy Dar . . . . .	1017
<i>Corpus Considerations for Annotator Modeling and Scaling</i> Sarumi Olufunke Oluyemi, Béla Neuendorf, Joan Plepi, Lucie Flek, Jörg Schlötterer and Charles Welch. . . . .	1029
<i>On Large Language Models' Hallucination with Regard to Known Facts</i> Che Jiang, Biqing Qi, Xiangyu Hong, Dayuan Fu, Yang Cheng, Fandong Meng, Mo Yu, Bowen Zhou and Jie Zhou . . . . .	1041
<i>One-Size-Fits-All"? Examining Expectations around What Constitute Fair or GoodNLG System Behaviors</i> Li Lucy, Su Lin Blodgett, Milad Shokouhi, Hanna Wallach and Alexandra Olteanu . . . . .	1054
<i>Language Models Hallucinate, but May Excel at Fact Verification</i> Jian Guan, Jesse Dodge, David Wadden, Minlie Huang and Hao Peng . . . . .	1090
<i>A Rationale-centric Counterfactual Data Augmentation Method for Cross-Document Event Coreference Resolution</i> Bowen Ding, Qingkai Min, Shengkun Ma, Yingjie Li, Linyi Yang and Yue Zhang . . . . .	1112
<i>TrojFSP: Trojan Insertion in Few-shot Prompt Tuning</i> Mengxin Zheng, Jiaqi Xue, Xun Chen, Yanshan Wang, Qian Lou and Lei Jiang . . . . .	1141
<i>Ensuring Safe and High-Quality Outputs: A Guideline Library Approach for Language Models</i> Yi Luo, Zhenghao Lin, YuHao Zhang, Jiashuo Sun, Chen Lin, Chengjin Xu, Xiangdong Su, yelong shen, Jian Guo and Yeyun Gong . . . . .	1152
<i>X-PARADE: Cross-Lingual Textual Entailment and Information Divergence across Paragraphs</i> Juan Diego Rodriguez, Katrin Erk and Greg Durrett . . . . .	1198
<i>Topics, Authors, and Institutions in Large Language Model Research: Trends from 17K arXiv Papers</i> Rajiv Movva, Sidhika Balachandar, Kenny Peng, Gabriel Agostini, Nikhil Garg and Emma Pierson . . . . .	1223
<i>E<sup>5</sup>: Zero-shot Hierarchical Table Analysis using Augmented LLMs via Explain, Extract, Execute, Exhibit and Extrapolate</i> Zhehao Zhang, Yan Gao and Jian-Guang Lou . . . . .	1244
<i>S3Eval: A Synthetic, Scalable, Systematic Evaluation Suite for Large Language Model</i> Fangyu Lei, Qian Liu, Yiming Huang, Shizhu He, Jun Zhao and Kang Liu . . . . .	1259

<i>MMC: Advancing Multimodal Chart Understanding with Large-scale Instruction Tuning</i>	
Fuxiao Liu, Xiaoyang Wang, Wenlin Yao, Jianshu Chen, Kaiqiang Song, Sangwoo Cho, Yaser Yacoob and Dong Yu .....	1287
<i>Visual Grounding Helps Learn Word Meanings in Low-Data Regimes</i>	
Chengxu Zhuang, Evelina Fedorenko and Jacob Andreas .....	1311
<i>Accurate Knowledge Distillation via n-best Reranking</i>	
Hendra Setiawan .....	1330
<i>AutoPRM: Automating Procedural Supervision for Multi-Step Reasoning via Controllable Question Decomposition</i>	
Zhaorun Chen, Zhuokai Zhao, Zhihong Zhu, Ruiqi Zhang, Xiang Li, Bhiksha Raj and Huaxiu Yao .....	1346
<i>SEMQA: Semi-Extractive Multi-Source Question Answering</i>	
Tal Schuster, Adam D Lelkes, Haitian Sun, Jai Gupta, Jonathan Berant, William W. Cohen and Donald Metzler .....	1363
<i>Fine-Tuning Language Models with Reward Learning on Policy</i>	
Hao Lang, Fei Huang and Yongbin Li .....	1382
<i>A Universal Dependencies Treebank for Highland Puebla Nahuatl</i>	
Robert Pugh and Francis M. Tyers .....	1393
<i>COPAL-ID: Indonesian Language Reasoning with Local Culture and Nuances</i>	
Haryo Akbarianto Wibowo, Erland Hilman Fuadi, Made Nindyatama Nityasya, Radityo Eko Prasajo and Alham Fikri Aji .....	1404
<i>IterAlign: Iterative Constitutional Alignment of Large Language Models</i>	
Xiusi Chen, Hongzhi Wen, Sreyashi Nag, Chen Luo, Qingyu Yin, Ruirui Li, Zheng Li and Wei Wang .....	1423
<i>OrchestraLLM: Efficient Orchestration of Language Models for Dialogue State Tracking</i>	
Chia-Hsuan Lee, Hao Cheng and Mari Ostendorf .....	1434
<i>Multi-Operational Mathematical Derivations in Latent Space</i>	
Marco Valentino, Jordan Meadows, Lan Zhang and Andre Freitas .....	1446
<i>Large Language Models Help Humans Verify Truthfulness – Except When They Are Convincingly Wrong</i>	
Chenglei Si, Navita Goyal, Tongshuang Wu, Chen Zhao, Shi Feng, Hal Daumé Iii and Jordan Lee Boyd-Graber .....	1459
<i>XferBench: a Data-Driven Benchmark for Emergent Language</i>	
Brendon Boldt and David R Mortensen .....	1475
<i>Evaluating Large Language Models as Generative User Simulators for Conversational Recommendation</i>	
Se-eun Yoon, Zhankui He, Jessica Maria Echterhoff and Julian McAuley .....	1490
<i>A Symbolic Framework for Evaluating Mathematical Reasoning and Generalisation with Transformers</i>	
Jordan Meadows, Marco Valentino, Damien Teney and Andre Freitas .....	1505
<i>Identifying Linear Relational Concepts in Large Language Models</i>	
David Chanin, Anthony Hunter and Oana-Maria Camburu .....	1524

<i>Benchmark Transparency: Measuring the Impact of Data on Evaluation</i> Venelin Kovatchev and Matthew Lease .....	1536
<i>JAMDEC: Unsupervised Authorship Obfuscation using Constrained Decoding over Small Language Models</i> Jillian Fisher, Ximing Lu, Jaehun Jung, Liwei Jiang, Zaid Harchaoui and Yejin Choi .....	1552
<i>REST: Retrieval-Based Speculative Decoding</i> Zhenyu He, Zexuan Zhong, Tianle Cai, Jason D. Lee and Di He .....	1582
<i>Sub-Sentence Encoder: Contrastive Learning of Propositional Semantic Representations</i> Sihao Chen, Hongming Zhang, Tong Chen, Ben Zhou, Wenhao Yu, Dian Yu, Baolin Peng, Hongwei Wang, Dan Roth and Dong Yu .....	1596
<i>MSciNLI: A Diverse Benchmark for Scientific Natural Language Inference</i> Mobashir Sadat and Cornelia Caragea .....	1610
<i>Causal Inference for Human-Language Model Collaboration</i> Bohan Zhang, Yixin Wang and Paramveer Dhillon .....	1630
<i>SELF-GUARD: Empower the LLM to Safeguard Itself</i> Zezhong Wang, Fangkai Yang, Lu Wang, Pu Zhao, Hongru Wang, Liang Chen, Qingwei Lin and Kam-Fai Wong .....	1648
<i>COSIGN: Contextual Facts Guided Generation for Knowledge Graph Completion</i> Jinpeng Li, Hang Yu, Xiangfeng Luo and Qian Liu .....	1669
<i>Toward Informal Language Processing: Knowledge of Slang in Large Language Models</i> Zhewei Sun, Qian Hu, Rahul Gupta, Richard Zemel and Yang Xu .....	1683
<i>Ghostbuster: Detecting Text Ghostwritten by Large Language Models</i> Vivek Verma, Eve Fleisig, Nicholas Tomlin and Dan Klein .....	1702
<i>End-to-End Beam Retrieval for Multi-Hop Question Answering</i> Jiahao Zhang, Haiyang Zhang, Dongmei Zhang, Liu Yong and Shen Huang .....	1718
<i>Leveraging Generative Large Language Models with Visual Instruction and Demonstration Retrieval for Multimodal Sarcasm Detection</i> Binghao Tang, Boda Lin, Haolong Yan and Si Li .....	1732
<i>Multi-Scale Prompt Memory-Augmented Model for Black-Box Scenarios</i> Xiaojun Kuang, C. L. Philip Chen, Shuzhen Li and Tong Zhang .....	1743
<i>Ungrammatical-syntax-based In-context Example Selection for Grammatical Error Correction</i> Chenming Tang, Fanyi Qu and Yunfang Wu .....	1758
<i>BUFFET: Benchmarking Large Language Models for Few-shot Cross-lingual Transfer</i> Akari Asai, Sneha Kudugunta, Xinyan Velocity Yu, Terra Blevins, Hila Gonen, Machel Reid, Yulia Tsvetkov, Sebastian Ruder and Hannaneh Hajishirzi .....	1771
<i>TISE: A Tripartite In-context Selection Method for Event Argument Extraction</i> Yanhe Fu, Yanan Cao, Qingyue Wang and Yi Liu .....	1801
<i>Reasoning or Reciting? Exploring the Capabilities and Limitations of Language Models Through Counterfactual Tasks</i> Zhaofeng Wu, Linlu Qiu, Alexis Ross, Ekin Akyürek, Boyuan Chen, Bailin Wang, Najoung Kim, Jacob Andreas and Yoon Kim .....	1819



<i>TRUE-UIE: Two Universal Relations Unify Information Extraction Tasks</i>	
Yucheng Wang, Bowen Yu, Yilin Liu and Shudong Lu . . . . .	1863
<i>zrLLM: Zero-Shot Relational Learning on Temporal Knowledge Graphs with Large Language Models</i>	
Zifeng Ding, Heling Cai, Jingpei Wu, Yunpu Ma, Ruotong Liao, Bo Xiong and Volker Tresp	1877
<i>Embodied Executable Policy Learning with Language-based Scene Summarization</i>	
Jielin Qiu, Mengdi Xu, William Han, Seungwhan Moon and Ding Zhao . . . . .	1896
<i>Metacognitive Prompting Improves Understanding in Large Language Models</i>	
Yuqing Wang and Yun Zhao . . . . .	1914
<i>MART: Improving LLM Safety with Multi-round Automatic Red-Teaming</i>	
Suyu Ge, Chunting Zhou, Rui Hou, Madian Khabza, Yi-Chia Wang, Qifan Wang, Jiawei Han and Yuning Mao . . . . .	1927
<i>DialogCC: An Automated Pipeline for Creating High-Quality Multi-Modal Dialogue Dataset</i>	
Young-Jun Lee, Byungsoo Ko, Han-Gyu Kim, Jonghwan Hyeon and Ho-Jin Choi . . . . .	1938
<i>Routing to the Expert: Efficient Reward-guided Ensemble of Large Language Models</i>	
Keming Lu, Hongyi Yuan, Runji Lin, Junyang Lin, Zheng Yuan, Chang Zhou and Jingren Zhou	1964
<i>Automatic Generation of Model and Data Cards: A Step Towards Responsible AI</i>	
Jiarui Liu, Wenkai Li, Zhijing Jin and Mona T. Diab . . . . .	1975
<i>FUN with Fisher: Improving Generalization of Adapter-Based Cross-lingual Transfer with Scheduled Unfreezing</i>	
Chen Cecilia Liu, Jonas Pfeiffer, Ivan Vulić and Iryna Gurevych . . . . .	1998
<i>Are Multilingual LLMs Culturally-Diverse Reasoners? An Investigation into Multicultural Proverbs and Sayings</i>	
Chen Cecilia Liu, Fajri Koto, Timothy Baldwin and Iryna Gurevych . . . . .	2016
<i>The Colorful Future of LLMs: Evaluating and Improving LLMs as Emotional Supporters for Queer Youth</i>	
Shir Lissak, Nitay Calderon, Geva Shenkman, Yaakov Ophir, Eyal Fruchter, Anat Brunstein Klokem and Roi Reichart . . . . .	2040
<i>IPED: An Implicit Perspective for Relational Triple Extraction based on Diffusion Model</i>	
Jianli Zhao, Changhao Xu and Bin. Jiang . . . . .	2080
<i>QualEval: Qualitative Evaluation for Model Improvement</i>	
Vishvak Murahari, Ameet Deshpande, Peter Clark, Tanmay Rajpurohit, Ashish Sabharwal, Karthik R Narasimhan and Ashwin Kalyan . . . . .	2093
<i>Quantum-inspired Language Model with Lindblad Master Equation and Interference Measurement for Sentiment Analysis</i>	
Kehuan Yan, Peichao Lai and yilei wang . . . . .	2112
<i>VisLingInstruct: Elevating Zero-Shot Learning in Multi-Modal Language Models with Autonomous Instruction Optimization</i>	
Dongsheng Zhu, Daniel Tang, Weidong Han, Jinghui Lu, Yukun Zhao, Guoliang Xing, Junfeng Wang and Dawei Yin . . . . .	2122
<i>A Wolf in Sheep's Clothing: Generalized Nested Jailbreak Prompts can Fool Large Language Models Easily</i>	
Peng Ding, Jun Kuang, Dan Ma, Xuezhi Cao, Yunsen Xian, Jiajun Chen and Shujian Huang	2136

<i>P<sup>3</sup>Sum: Preserving Author’s Perspective in News Summarization with Diffusion Language Models</i> Yuhan Liu, Shangbin Feng, Xiaochuang Han, Vidhisha Balachandran, Chan Young Park, Sachin Kumar and Yulia Tsvetkov .....	2154
<i>Bridging the Novice-Expert Gap via Models of Decision-Making: A Case Study on Remediating Math Mistakes</i> Rose E Wang, Qingyang Zhang, Carly D Robinson, Susanna Loeb and Dorottya Demszky .	2174
<i>RST-LoRA: A Discourse-Aware Low-Rank Adaptation for Long Document Abstractive Summarization</i> Dongqi Pu and Vera Demberg .....	2200
<i>Strings from the Library of Babel: Random Sampling as a Strong Baseline for Prompt Optimisation</i> Yao Lu, Jiayi Wang, Raphael Tang, Sebastian Riedel and Pontus Stenetorp .....	2221
<i>ReTA: Recursively Thinking Ahead to Improve the Strategic Reasoning of Large Language Models</i> Jinhao Duan, Shiqi Wang, James Diffenderfer, Lichao Sun, Tianlong Chen, Bhavya Kailkhura and Kaidi Xu .....	2232
<i>Fact Checking Beyond Training Set</i> Payam Karisani and Heng Ji .....	2247
<i>Program-Aided Reasoners (Better) Know What They Know</i> Anubha Kabra, Sanketh Rangreji, Yash Mathur, Aman Madaan, Emmy Liu and Graham Neubig	2262
<i>The Perspectivist Paradigm Shift: Assumptions and Challenges of Capturing Human Labels</i> Eve Fleisig, Su Lin Blodgett, Dan Klein and Zeerak Talat .....	2279
<i>Principles from Clinical Research for NLP Model Generalization</i> Aparna Elangovan, Jiayuan He, Yuan Li and Karin Verspoor .....	2293
<i>First Tragedy, then Parse: History Repeats Itself in the New Era of Large Language Models</i> Naomi Saphra, Eve Fleisig, Kyunghyun Cho and Adam Lopez .....	2310
<i>Found in the Middle: Permutation Self-Consistency Improves Listwise Ranking in Large Language Models</i> Raphael Tang, Crystina Zhang, Xueguang Ma, Jimmy Lin and Ferhan Ture .....	2327
<i>From Language Modeling to Instruction Following: Understanding the Behavior Shift in LLMs after Instruction Tuning</i> Xuansheng Wu, Wenlin Yao, Jianshu Chen, Xiaoman Pan, Xiaoyang Wang, Ninghao Liu and Dong Yu .....	2341
<i>POLYIE: A Dataset of Information Extraction from Polymer Material Scientific Literature</i> Jerry Junyang Cheung, Yuchen Zhuang, Yinghao Li, Pranav Shetty, Wantian Zhao, Sanjeev Gram-purohit, Rampi Ramprasad and Chao Zhang .....	2370
<i>LLM-based Medical Assistant Personalization with Short- and Long-Term Memory Coordination</i> Kai Zhang, Yangyang Kang, Fubang Zhao and Xiaozhong Liu .....	2386
<i>SumTra: A Differentiable Pipeline for Few-Shot Cross-Lingual Summarization</i> Jacob Parnell, Inigo Jauregi Unanue and Massimo Piccardi .....	2399
<i>KTRL+F: Knowledge-Augmented In-Document Search</i> Hanseok Oh, Haebin Shin, Miyoung Ko, Hyunji Lee and Minjoon Seo .....	2416

<i>How Well Do Large Language Models Truly Ground?</i>	
Hyunji Lee, Se June Joo, Chaeun Kim, Joel Jang, Doyoung Kim, Kyoung-Woon On and Minjoon Seo .....	2437
<i>ALBA: Adaptive Language-Based Assessments for Mental Health</i>	
Vasudha Varadarajan, Sverker Sikström, Oscar Kjell and H. Schwartz .....	2466
<i>FREB-TQA: A Fine-Grained Robustness Evaluation Benchmark for Table Question Answering</i>	
Wei Zhou, Mohsen Mesgar, Heike Adel and Annemarie Friedrich .....	2479
<i>MILL: Mutual Verification with Large Language Models for Zero-Shot Query Expansion</i>	
Pengyue Jia, Yiding Liu, Xiangyu Zhao, Xiaopeng LI, Changying Hao, Shuaiqiang Wang and Dawei Yin .....	2498
<i>Efficient Benchmarking (of Language Models)</i>	
Yotam Perlitz, Elron Bandel, Ariel Gera, Ofir Arviv, Liat Ein-Dor, Eyal Shnarch, Noam Slonim, Michal Shmueli-Scheuer and Leshem Choshen .....	2519
<i>ReFACT: Updating Text-to-Image Models by Editing the Text Encoder</i>	
Dana Arad, Hadas Orgad and Yonatan Belinkov .....	2537
<i>A Likelihood Ratio Test of Genetic Relationship among Languages</i>	
V.S.D.S.Mahesh Akavarapu and Arnab Bhattacharya .....	2559
<i>PaD: Program-aided Distillation Can Teach Small Models Reasoning Better than Chain-of-thought Fine-tuning</i>	
Xuekai Zhu, Biqing Qi, Kaiyan Zhang, Xinwei Long, Zhouhan Lin and Bowen Zhou .....	2571
<i>MEGAVERSE: Benchmarking Large Language Models Across Languages, Modalities, Models and Tasks</i>	
Sanchit Ahuja, Divyanshu Aggarwal, Varun Gumma, Ishaan Watts, Ashutosh Sathe, Millicent Ochieng, Rishav Hada, Prachi Jain, Mohamed Ahmed, Kalika Bali and Sunayana Sitaram .....	2598
<i>Unlocking Emergent Modularity in Large Language Models</i>	
Zihan Qiu, Zeyu Huang and Jie Fu .....	2638
<i>A School Student Essay Corpus for Analyzing Interactions of Argumentative Structure and Quality</i>	
Maja Stahl, Nadine Michel, Sebastian Kilsbach, Julian Schmidtke, Sara Rezat and Henning Wachsmuth .....	2661
<i>Adjusting Interpretable Dimensions in Embedding Space with Human Judgments</i>	
Katrin Erk and Marianna Apidianaki .....	2675
<i>PatentEval: Understanding Errors in Patent Generation</i>	
You Zuo, Kim Gerdes, Éric Villemonte de la Clergerie and Benoît Sagot .....	2687
<i>Contextual Refinement of Translations: Large Language Models for Sentence and Document-Level Post-Editing</i>	
Sai Koneru, Miriam Exel, Matthias Huck and Jan Niehues .....	2711
<i>Metaphor Detection with Context Enhancement and Curriculum Learning</i>	
Kaidi Jia and Rongsheng Li .....	2726
<i>What Causes the Failure of Explicit to Implicit Discourse Relation Recognition?</i>	
Wei Liu, Stephen Wan and Michael Strube .....	2738

<i>UniverSLU: Universal Spoken Language Understanding for Diverse Tasks with Natural Language Instructions</i>	
Siddhant Arora, Hayato Futami, Jee-weon Jung, Yifan Peng, Roshan Sharma, Yosuke Kashiwagi, Emiru Tsunoo, Karen Livescu and Shinji Watanabe . . . . .	2754
<i>How Trustworthy are Open-Source LLMs? An Assessment under Malicious Demonstrations Shows their Vulnerabilities</i>	
Lingbo Mo, Boshi Wang, Muhao Chen and Huan Sun . . . . .	2775
<i>Paraphrase and Solve: Exploring and Exploiting the Impact of Surface Form on Mathematical Reasoning in Large Language Models</i>	
Yue Zhou, Yada Zhu, Diego Antognini, Yoon Kim and Yang Zhang . . . . .	2793
<i>TriSum: Learning Summarization Ability from Large Language Models with Structured Rationale</i>	
Pengcheng Jiang, Cao Xiao, Zifeng Wang, Parminder Bhatia, Jimeng Sun and Jiawei Han . . . . .	2805
<i>GenRES: Rethinking Evaluation for Generative Relation Extraction in the Era of Large Language Models</i>	
Pengcheng Jiang, Jiacheng Lin, Zifeng Wang, Jimeng Sun and Jiawei Han . . . . .	2820
<i>Curated Datasets and Neural Models for Machine Translation of Informal Registers between Mayan and Spanish Vernaculars</i>	
Andrés Lou, Juan Antonio Pérez-Ortiz, Felipe Sánchez-Martínez and Víctor M. Sánchez-Cartagena . . . . .	2838
<i>The Effect of Data Partitioning Strategy on Model Generalizability: A Case Study of Morphological Segmentation</i>	
Zoey Liu and Bonnie J Dorr . . . . .	2851
<i>Measuring Entrainment in Spontaneous Code-switched Speech</i>	
Debasmita Bhattacharya, Siying Ding, Alayna Nguyen and Julia Hirschberg . . . . .	2865
<i>A Survey of Meaning Representations – From Theory to Practical Utility</i>	
Zacchary Sadeddine, Juri Opitz and Fabian M. Suchanek . . . . .	2877
<i>Mitigating Language-Level Performance Disparity in mPLMs via Teacher Language Selection and Cross-lingual Self-Distillation</i>	
Haozhe Zhao, Zefan Cai, Shuzheng Si, Liang Chen, Yufeng He, Kaikai An and Baobao Chang . . . . .	2893
<i>Evaluating In-Context Learning of Libraries for Code Generation</i>	
Arkil Patel, Siva Reddy, Dzmitry Bahdanau and Pradeep Dasigi . . . . .	2908
<i>Visually-Aware Context Modeling for News Image Captioning</i>	
Tingyu Qu, Tinne Tuytelaars and Marie-Francine Moens . . . . .	2927
<i>Regularized Conventions: Equilibrium Computation as a Model of Pragmatic Reasoning</i>	
Athul Paul Jacob, Gabriele Farina and Jacob Andreas . . . . .	2944
<i>TopicGPT: A Prompt-based Topic Modeling Framework</i>	
Chau Minh Pham, Alexander Hoyle, Simeng Sun, Philip Resnik and Mohit Iyyer . . . . .	2956
<i>ChatGPT as an Attack Tool: Stealthy Textual Backdoor Attack via Blackbox Generative Model Trigger</i>	
Jiazhao Li, Yijin Yang, Zhuofeng Wu, V.G.Vinod Vydiswaran and Chaowei Xiao . . . . .	2985
<i>Social Meme-ing: Measuring Linguistic Variation in Memes</i>	
Naitian Zhou, David Jurgens and David Bamman . . . . .	3005

<i>ExpertQA: Expert-Curated Questions and Attributed Answers</i> Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar and Dan Roth	3025
<i>What if you said that differently?: How Explanation Formats Affect Human Feedback Efficacy and User Perception</i> Chaitanya Malaviya, Subin Lee, Dan Roth and Mark Yatskar	3046
<i>When Life Gives You Lemons, Make Cherryade: Converting Feedback from Bad Responses into Good Labels</i> Weiyang Shi, Emily Dinan, Kurt Shuster, Jason E Weston and Jing Xu	3066
<i>Kreyòl-MT: Building MT for Latin American, Caribbean and Colonial African Creole Languages</i> Nathaniel Romney Robinson, Raj Dabre, Ammon C Shurtz, Rasul Dent, Onenamiyi Onesi, Claire Bizon Monroc, Loïc Grobol, Hasan Muhammad, Ashi Garg, Naome A Etori, Vijay Murari Tiyyala, Olanrewaju Samuel, Matthew Dean Stutzman, Bismarck Bamfo Odoom, Sanjeev Khudanpur, Stephen D. Richardson and Kenton Murray	3083
<i>Instructions as Backdoors: Backdoor Vulnerabilities of Instruction Tuning for Large Language Models</i> Jiashu Xu, Mingyu Derek Ma, Fei Wang, Chaowei Xiao and Muhao Chen	3111
<i>Modeling Empathetic Alignment in Conversation</i> Jiamin Yang and David Jurgens	3127
<i>Native Language Identification in Texts: A Survey</i> Dhiman Goswami, Sharanya Thilagan, Kai North, Shervin Malmasi and Marcos Zampieri	3149
<i>LoRETTA: Low-Rank Economic Tensor-Train Adaptation for Ultra-Low-Parameter Fine-Tuning of Large Language Models</i> Yifan Yang, Jiajun Zhou, Ngai Wong and Zheng Zhang	3161
<i>Which One? Leveraging Context Between Objects and Multiple Views for Language Grounding</i> Chancharik Mitra, Abrar Anwar, Rodolfo Corona, Dan Klein, Trevor Darrell and Jesse Thomason	3177
<i>Do Localization Methods Actually Localize Memorized Data in LLMs? A Tale of Two Benchmarks</i> Ting-Yun Chang, Jesse Thomason and Robin Jia	3190
<i>PromptFix: Few-shot Backdoor Removal via Adversarial Prompt Tuning</i> Tianrong Zhang, Zhaohan Xi, Ting Wang, Prasenjit Mitra and Jinghui Chen	3212
<i>Comparing Explanation Faithfulness between Multilingual and Monolingual Fine-tuned Language Models</i> Zhixue Zhao and Nikolaos Aletras	3226
<i>A Pretrainer’s Guide to Training Data: Measuring the Effects of Data Age, Domain Coverage, Quality, &amp; Toxicity</i> Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimno and Daphne Ippolito	3245
<i>Instructional Fingerprinting of Large Language Models</i> Jiashu Xu, Fei Wang, Mingyu Derek Ma, Pang Wei Koh, Chaowei Xiao and Muhao Chen	3277
<i>Reinforced Multiple Instance Selection for Speaker Attribute Prediction</i> Alireza Salkhordeh Ziabari, Ali Omrani, Parsa Hejabi, Preni Golazizian, Brendan Kennedy, Payam Piray and Morteza Dehghani	3307

<i>DynaMo: Accelerating Language Model Inference with Dynamic Multi-Token Sampling</i> Shikhar Tuli, Chi-Heng Lin, Yen-Chang Hsu, Niraj Jha, Yilin Shen and Hongxia Jin . . . . .	3322
<i>Few-shot Knowledge Graph Relational Reasoning via Subgraph Adaptation</i> Haochen Liu, Song Wang, Chen Chen and Jundong Li . . . . .	3346
<i>Uncertainty Quantification for In-Context Learning of Large Language Models</i> Chen Ling, Xujiang Zhao, Xuchao Zhang, Wei Cheng, Yanchi Liu, Yiyun Sun, Mika Oishi, Takao Osaki, Katsushi Matsuda, Jie Ji, Guangji Bai, Liang Zhao and Haifeng Chen . . . . .	3357
<i>HelpSteer: Multi-attribute Helpfulness Dataset for SteerLM</i> Zhilin Wang, Yi Dong, Jiaqi Zeng, Virginia Adams, Makesh Narsimhan Sreedhar, Daniel Egert, Olivier Delalleau, Jane Polak Scowcroft, Neel Kant, Aidan M Swope and Oleksii Kuchaiev . . . . .	3371
<i>A Preference-driven Paradigm for Enhanced Translation with Large Language Models</i> Dawei Zhu, Sony Trenous, Xiaoyu Shen, Dietrich Klakow, Bill Byrne and Eva Hasler . . . . .	3385
<i>Fair Abstractive Summarization of Diverse Perspectives</i> Yusen Zhang, Nan Zhang, Yixin Liu, Alexander Fabbri, Junru Liu, Ryo Kamoi, Xiaoxin Lu, Caiming Xiong, Jieyu Zhao, Dragomir Radev, Kathleen McKeown and Rui Zhang . . . . .	3404
<i>What Are We Measuring When We Evaluate Large Vision-Language Models? An Analysis of Latent Factors and Biases</i> Anthony Tiong, Junqi Zhao, Boyang Li, Junnan Li, Steven Hoi and Caiming Xiong . . . . .	3427
<i>Show Your Work with Confidence: Confidence Bands for Tuning Curves</i> Nicholas Lourie, Kyunghyun Cho and He He . . . . .	3455
<i>GRASP: A Disagreement Analysis Framework to Assess Group Associations in Perspectives</i> Vinodkumar Prabhakaran, Christopher M Homan, Lora Aroyo, Aida Mostafazadeh Davani, Alicia Parrish, Alex Taylor, Mark Diaz, Ding Wang and Gregory Serapio-García . . . . .	3473
<i>Event Causality Is Key to Computational Story Understanding</i> Yidan Sun, Qin Chao and Boyang Li . . . . .	3493
<i>Subspace Representations for Soft Set Operations and Sentence Similarities</i> Yoichi Ishibashi, Sho Yokoi, Katsuhito Sudoh and Satoshi Nakamura . . . . .	3512
<i>My Heart Skipped a Beat! Recognizing Expressions of Embodied Emotion in Natural Language</i> Yuan Zhuang, Tianyu Jiang and Ellen Riloff . . . . .	3525
<i>Low-Cost Generation and Evaluation of Dictionary Example Sentences</i> Bill Cai, Ng Boon Liang Clarence, Daniel Tan Wee Liang and Shelvia Hotama . . . . .	3538
<i>Making Language Models Better Tool Learners with Execution Feedback</i> Shuofei Qiao, Honghao Gui, Chengfei Lv, Qianghuai Jia, Huajun Chen and Ningyu Zhang . . . . .	3550
<i>Complex Claim Verification with Evidence Retrieved in the Wild</i> Jifan Chen, Grace Kim, Aniruddh Sriram, Greg Durrett and Eunsol Choi . . . . .	3569
<i>Multimodal Multi-loss Fusion Network for Sentiment Analysis</i> Zehui Wu, Ziwei Gong, Jaywon Koo and Julia Hirschberg . . . . .	3588
<i>Confronting LLMs with Traditional ML: Rethinking the Fairness of Large Language Models in Tabular Classifications</i> Yanchen Liu, Srishti Gautam, Jiaqi Ma and Himabindu Lakkaraju . . . . .	3603

<i>Analyzing the Use of Metaphors in News Editorials for Political Framing</i> Meghdut Sengupta, Roxanne El Baff, Milad Alshomary and Henning Wachsmuth . . . . .	3621
<i>SharpSeq: Empowering Continual Event Detection through Sharpness-Aware Sequential-task Learning</i> Thanh-Thien Le, Viet Dao, Linh Van Nguyen, Thi-Nhung Nguyen, Linh Van Ngo and Thien Huu Nguyen . . . . .	3632
<i>Dissecting Paraphrases: The Impact of Prompt Syntax and supplementary Information on Knowledge Retrieval from Pretrained Language Models</i> Stephan Linzbach, Dimitar Dimitrov, Laura Kallmeyer, Kilian Evang, Hajira Jabeen and Stefan Dietze . . . . .	3645
<i>Know When To Stop: A Study of Semantic Drift in Text Generation</i> Ava Spataru . . . . .	3656
<i>Curriculum Masking in Vision-Language Pretraining to Maximize Cross Modal Interaction</i> Kraig Yuheng Tou and Zijun Sun . . . . .	3672
<i>Elote, Choclo and Mazorca: on the Varieties of Spanish</i> Cristina España-Bonet and Alberto Barrón-Cedeño . . . . .	3689
<i>Ada-LEval: Evaluating long-context LLMs with length-adaptable benchmarks</i> Chonghua Wang, Haodong Duan, Songyang Zhang, Dahua Lin and Kai Chen . . . . .	3712
<i>A Zero-Shot Monolingual Dual Stage Information Retrieval System for Spanish Biomedical Systematic Literature Reviews</i> Regina Ofori-Boateng, Magaly Aceves-Martins, Nirmalie Wiratunga and Carlos Francisco Moreno-Garcia . . . . .	3725
<i>LayoutPointer: A Spatial-Context Adaptive Pointer Network for Visual Information Extraction</i> Huang Siyuan, yongping xiong and Wu Guibin . . . . .	3737
<i>Long-form evaluation of model editing</i> Domenic Rosati, Robie Gonzales, Jinkun Chen, Xuemin Yu, Yahya Kayani, Frank Rudzicz and Hassan Sajjad . . . . .	3749
<i>Analyzing the Role of Semantic Representations in the Era of Large Language Models</i> Zhijing Jin, Yuen Chen, Fernando Gonzalez Adauto, Jiarui Liu, Jiayi Zhang, Julian Michael, Bernhard Schölkopf and Mona T. Diab . . . . .	3781
<i>TRAQ: Trustworthy Retrieval Augmented Question Answering via Conformal Prediction</i> Shuo Li, Sangdon Park, Insup Lee and Osbert Bastani . . . . .	3799
<i>MapGuide: A Simple yet Effective Method to Reconstruct Continuous Language from Brain Activities</i> Xinpei Zhao, Jingyuan Sun, Shaonan Wang, Jing Ye, Xhz Xhz and Chengqing Zong . . . . .	3822
<i>On-the-fly Definition Augmentation of LLMs for Biomedical NER</i> Monica Munnangi, Sergey Feldman, Byron C Wallace, Silvio Amir, Tom Hope and Aakanksha Naik . . . . .	3833
<i>This Land is Your, My Land: Evaluating Geopolitical Bias in Language Models through Territorial Disputes</i> Bryan Li, Samar Haider and Chris Callison-Burch . . . . .	3855
<i>Set-Aligning Framework for Auto-Regressive Event Temporal Graph Generation</i> Xingwei Tan, Yuxiang Zhou, Gabriele Pergola and Yulan He . . . . .	3872

<i>LanguageFlow: Advancing Diffusion Language Generation with Probabilistic Flows</i> Shujian Zhang, Lemeng Wu, Chengyue Gong and Xingchao Liu . . . . .	3893
<i>Towards Improved Multi-Source Attribution for Long-Form Answer Generation</i> Nilay Patel, Shivashankar Subramanian, Siddhant Garg, Pratyay Banerjee and Amita Misra	3906
<i>Synthetic Query Generation for Privacy-Preserving Deep Retrieval Systems using Differentially Private Language Models</i> Aldo Gael Carranza, Reza Farahani, Natalia Ponomareva, Alexey Kurakin, Matthew Jagielski and Milad Nasr . . . . .	3920
<i>Okay, Let's Do This! Modeling Event Coreference with Generated Rationales and Knowledge Distillation</i> Abhijnan Nath, Shadi Manafi Avari, Avyakta Chelle and Nikhil Krishnaswamy . . . . .	3931
<i>Can Knowledge Graphs Reduce Hallucinations in LLMs? : A Survey</i> Garima Agrawal, Tharindu Sandaruwan Kumaraage, Zeyad Alghamdi and huan liu . . . . .	3947
<i>Pedagogically Aligned Objectives Create Reliable Automatic Cloze Tests</i> Brian David Ondov, Kush Attal and Dina Demner-Fushman . . . . .	3961
<i>Take One Step at a Time to Know Incremental Utility of Demonstration: An Analysis on Reranking for Few-Shot In-Context Learning</i> Kazuma Hashimoto, Karthik Raman and Michael Bendersky . . . . .	3973
<i>LM-Infinite: Zero-Shot Extreme Length Generalization for Large Language Models</i> Chi Han, Qifan Wang, Hao Peng, Wenhan Xiong, Yu Chen, Heng Ji and Sinong Wang . . . . .	3991
<i>CONSCENDI: A Contrastive and Scenario-Guided Distillation Approach to Guardrail Models for Virtual Assistants</i> Albert Sun, Varun Nair, Elliot Schumacher and Anitha Kannan . . . . .	4009
<i>Advancing Beyond Identification: Multi-bit Watermark for Large Language Models</i> KiYoon Yoo, Wonhyuk Ahn and Nojun Kwak . . . . .	4031
<i>HTCCN: Temporal Causal Convolutional Networks with Hawkes Process for Extrapolation Reasoning in Temporal Knowledge Graphs</i> Tingxuan Chen, Jun Long, Liu Yang, Zidong Wang, Yongheng Wang and Xiongnan Jin . . . . .	4056
<i>SemStamp: A Semantic Watermark with Paraphrastic Robustness for Text Generation</i> Abe Bohan Hou, Jingyu Zhang, Tianxing He, Yichen Wang, Yung-Sung Chuang, Hongwei Wang, Lingfeng Shen, Benjamin Van Durme, Daniel Khashabi and Yulia Tsvetkov . . . . .	4067
<i>Media Bias Detection Across Families of Language Models</i> Iffat Maab, Edison Marrese-Taylor, Sebastian Padó and Yutaka Matsuo . . . . .	4083
<i>Better Zero-Shot Reasoning with Role-Play Prompting</i> Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Enzhi Wang and Xiaohang Dong . . . . .	4099
<i>Event-Content-Oriented Dialogue Generation in Short Video</i> Fenghua Cheng, Xue Li, Zi Huang, Jinxiang Wang and Sen Wang . . . . .	4114
<i>DoG-Instruct: Towards Premium Instruction-Tuning Data via Text-Grounded Instruction Wrapping</i> Yongrui Chen, Haiyun Jiang, Xinting Huang, Shuming Shi and Guilin Qi . . . . .	4125
<i>Beyond Borders: Investigating Cross-Jurisdiction Transfer in Legal Case Summarization</i> Santosh T.y.s.s, Vatsal Venkatkrishna, Saptarshi Ghosh and Matthias Grabmair . . . . .	4136



<i>EDC: Effective and Efficient Dialog Comprehension For Dialog State Tracking</i> Qifan Lu, Bhaskar Ramasubramanian and Radha Poovendran .....	4151
<i>Automatic Restoration of Diacritics for Speech Data Sets</i> Sara Shatnawi, Sawsan Alqahtani and Hanan Aldarmaki .....	4166
<i>XNLIeu: a dataset for cross-lingual NLI in Basque</i> Maite Heredia, Julen Etxaniz, Muitze Zulaika, Xabier Saralegi, Jeremy Barnes and Aitor Soroa 4177	
<i>MDR: Model-Specific Demonstration Retrieval at Inference Time for In-Context Learning</i> Huazheng Wang, Jinming Wu, Haifeng Sun, Zixuan Xia, Daixuan Cheng, Jingyu Wang, Qi Qi and Jianxin Liao .....	4189
<i>Exploring Cross-Cultural Differences in English Hate Speech Annotations: From Dataset Construction to Analysis</i> Nayeon Lee, Chani Jung, Junho Myung, Jiho Jin, Jose Camacho-Collados, Juho Kim and Alice Oh .....	4205
<i>Enhancing Contextual Understanding in Large Language Models through Contrastive Decoding</i> Zheng Zhao, Emilio Monti, Jens Lehmann and Haytham Assem .....	4225
<i>Generalizable Sarcasm Detection is Just Around the Corner, of Course!</i> Hyeon Jang and Diego Frassinelli .....	4238
<i>Encoding of lexical tone in self-supervised models of spoken language</i> Gaofei Shen, Michaela Watkins, Afra Alishahi, Arianna Bisazza and Grzegorz Chrupała ...	4250
<i>A Systematic Comparison of Contextualized Word Embeddings for Lexical Semantic Change</i> Francesco Periti and Nina Tahmasebi .....	4262
<i>iACOS: Advancing Implicit Sentiment Extraction with Informative and Adaptive Negative Examples</i> Xiancai XU, Jia-Dong Zhang, Lei Xiong and Zhishang Liu .....	4283
<i>Rectifying Demonstration Shortcut in In-Context Learning</i> Joonwon Jang, Sanghwan Jang, Wonbin Kweon, Minjin Jeon and Hwanjo Yu .....	4294
<i>Universal NER: A Gold-Standard Multilingual Named Entity Recognition Benchmark</i> Stephen Mayhew, Terra Blevins, Shuheng Liu, Marek Suppa, Hila Gonen, Joseph Marvin Imper- rial, Börje F. Karlsson, Peiqin Lin, Nikola Ljubešić, Lester James Validad Miranda, Barbara Plank, Arij Riabi and Yuval Pinter .....	4322
<i>ODD: A Benchmark Dataset for the Natural Language Processing Based Opioid Related Aberrant Behavior Detection</i> Sunjae Kwon, Xun Wang, Weisong Liu, Emily Druhl, Minhee L. Sung, Joel I. Reisman, Wenjun Li, Robert D Kerns, William C. Becker and hong yu .....	4338
<i>A Comprehensive Study of Gender Bias in Chemical Named Entity Recognition Models</i> Xingmeng Zhao, Ali Shahzad Niazi and Anthony Rios .....	4360
<i>The Promises and Pitfalls of Using Language Models to Measure Instruction Quality in Education</i> Paiheng Xu, Jing Liu, Nathan D Jones, Julie Cohen and Wei Ai .....	4375
<i>Differentially Private Next-Token Prediction of Large Language Models</i> James Flemings, Meisam Razaviyayn and Murali Annavam .....	4390

<i>Improving Adversarial Data Collection by Supporting Annotators: Lessons from GAHD, a German Hate Speech Dataset</i>	
Janis Goldzycher, Paul Röttger and Gerold Schneider . . . . .	4405
<i>Memory Augmented Language Models through Mixture of Word Experts</i>	
Cicero Nogueira dos Santos, James Lee-Thorp, Isaac Noble, Chung-Ching Chang and David Uthus . . . . .	4425
<i>Impossible Distillation for Paraphrasing and Summarization: How to Make High-quality Lemonade out of Small, Low-quality Model</i>	
Jaehun Jung, Peter West, Liwei Jiang, Faeze Brahman, Ximing Lu, Jillian Fisher, Taylor Sorensen and Yejin Choi . . . . .	4439
<i>TofuEval: Evaluating Hallucinations of LLMs on Topic-Focused Dialogue Summarization</i>	
Liyan Tang, Igor Shalyminov, Amy Wing-mei Wong, Jon Burnsky, Jake W. Vincent, Yu’an Yang, Siffi Singh, Song Feng, Hwanjun Song, Hang Su, Lijia Sun, Yi Zhang, Saab Mansour and Kathleen McKeown . . . . .	4455
<i>MOKA: Moral Knowledge Augmentation for Moral Event Extraction</i>	
Xinliang Frederick Zhang, Winston Wu, Nicholas Beauchamp and Lu Wang . . . . .	4481
<i>Fixing Rogue Memorization in Many-to-One Multilingual Translators of Extremely-Low-Resource Languages by Rephrasing Training Samples</i>	
Paulo Cavalin, Pedro Henrique Leite da Silva Pires Domingues, Claudio Pinhanez and Julio Nogima . . . . .	4503
<i>Backdoor Attacks on Multilingual Machine Translation</i>	
Jun Wang, Qiongkai Xu, Xuanli He, Benjamin I. P. Rubinstein and Trevor Cohn . . . . .	4515
<i>Personalized Jargon Identification for Enhanced Interdisciplinary Communication</i>	
Yue Guo, Joseph Chee Chang, Maria Antoniak, Erin Bransom, Trevor Cohen, Lucy Lu Wang and Tal August . . . . .	4535
<i>Flames: Benchmarking Value Alignment of LLMs in Chinese</i>	
Kexin Huang, Xiangyang Liu, Qianyu Guo, Tianxiang Sun, Jiawei Sun, Yaru Wang, Zeyang Zhou, Yixu Wang, Yan Teng, Xipeng Qiu, Yingchun Wang and Dahua Lin . . . . .	4551
<i>Mitigating Bias for Question Answering Models by Tracking Bias Influence</i>	
Mingyu Derek Ma, Jiun-Yu Kao, Arpit Gupta, Yu-Hsiang Lin, Wenbo Zhao, Tagyoung Chung, Wei Wang, Kai-Wei Chang and Nanyun Peng . . . . .	4592
<i>Extending CLIP’s Image-Text Alignment to Referring Image Segmentation</i>	
Seoyeon Kim, Minguk Kang, Dongwon Kim, Jaesik Park and Suha Kwak . . . . .	4611
<i>Generating Attractive and Authentic Copywriting from Customer Reviews</i>	
Yu-Xiang Lin and Wei-Yun Ma . . . . .	4629
<i>Effective Long-Context Scaling of Foundation Models</i>	
Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajjwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oguz, Madian Khabza, Han Fang, Yashar Mehdad, Sharan Narang, Kshitiz Malik, Angela Fan, Shruti Bhosale, Sergey Edunov, Mike Lewis, Sinong Wang and Hao Ma . . . . .	4643
<i>Empowering Diffusion Models on the Embedding Space for Text Generation</i>	
Zhujin Gao, Junliang Guo, Xu Tan, Yongxin Zhu, Fang Zhang, Jiang Bian and Linli Xu . . . . .	4664

<i>Aligning as Debiasing: Causality-Aware Alignment via Reinforcement Learning with Interventional Feedback</i>	
Yu Xia, Tong Yu, Zhankui He, Handong Zhao, Julian McAuley and Shuai Li . . . . .	4684
<i>Fake Alignment: Are LLMs Really Aligned Well?</i>	
Yixu Wang, Yan Teng, Kexin Huang, Chengqi Lyu, Songyang Zhang, Wenwei Zhang, Xingjun Ma, Yu-Gang Jiang, Yu Qiao and Yingchun Wang . . . . .	4696
<i>Visually Guided Generative Text-Layout Pre-training for Document Intelligence</i>	
Zhiming Mao, Haoli Bai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu and Kam-Fai Wong . .	4713
<i>HILL: Hierarchy-aware Information Lossless Contrastive Learning for Hierarchical Text Classification</i>	
He Zhu, Junran Wu, Ruomei Liu, Yue Hou, Ze Yuan, Shangzhe Li, Yicheng Pan and Ke Xu	4731
<i>Investigating the Emergent Audio Classification Ability of ASR Foundation Models</i>	
Rao Ma, Adian Liusie, Mark Gales and Kate Knill . . . . .	4746
<i>In-context Learning Generalizes, But Not Always Robustly: The Case of Syntax</i>	
Aaron Mueller, Albert Webson, Jackson Petty and Tal Linzen . . . . .	4761
<i>Prompt-Singer: Controllable Singing-Voice-Synthesis with Natural Language Prompt</i>	
Yongqi Wang, Ruofan Hu, Rongjie Huang, Zhiqing Hong, Ruiqi Li, Wenrui Liu, Fuming You, Tao Jin and Zhou Zhao . . . . .	4780
<i>Lost in Transcription: Identifying and Quantifying the Accuracy Biases of Automatic Speech Recognition Systems Against Disfluent Speech</i>	
Dena Mujtaba, Nihar R. Mahapatra, Megan Arney, J Scott Yaruss, Hope Gerlach-Houck, Caryn Herring and Jia Bin . . . . .	4795
<i>MAFALDA: A Benchmark and Comprehensive Study of Fallacy Detection and Classification</i>	
Chadi Helwe, Tom Calamai, Pierre-Henri Paris, Chloé Clavel and Fabian M. Suchanek . . . .	4810
<i>Diffusion Glancing Transformer for Parallel Sequence-to-Sequence Learning</i>	
Lihua Qian, Mingxuan Wang, Yang Liu and Hao Zhou . . . . .	4846
<i>No Context Needed: Contextual Quandary In Idiomatic Reasoning With Pre-Trained Language Models</i>	
Kellen Tan Cheng and Suma Bhat . . . . .	4863
<i>Multi-stage Retrieve and Re-rank Model for Automatic Medical Coding Recommendation</i>	
Xindi Wang, Robert Mercer and Frank Rudzicz . . . . .	4881
<i>Anisotropy is Not Inherent to Transformers</i>	
Anemily Machina and Robert Mercer . . . . .	4892
<i>Finding Replicable Human Evaluations via Stable Ranking Probability</i>	
Parker Riley, Daniel Deutsch, George Foster, Viresh Ratnakar, Ali Dabirmoghaddam and Markus Freitag . . . . .	4908
<i>Stealthy and Persistent Misalignment on Large Language Models via Backdoor Injections</i>	
Yuanpu Cao, Bochuan Cao and Jinghui Chen . . . . .	4920
<i>Generalizable and Stable Finetuning of Pretrained Language Models on Low-Resource Texts</i>	
Sai Ashish Somayajula, Youwei Liang, Li Zhang, Abhishek Singh and Pengtao Xie . . . . .	4936
<i>Detecting Bipolar Disorder from Misdiagnosed Major Depressive Disorder with Mood-Aware Multi-Task Learning</i>	
Daemun Lee, Hyolim Jeon, Sejung Son, Chaewon Park, Ji hyun An, Seungbae Kim and Jinyoung Han . . . . .	4954

<i>Leveraging Code to Improve In-Context Learning for Semantic Parsing</i>	
Ben Bogin, Shivanshu Gupta, Peter Clark and Ashish Sabharwal .....	4971
<i>Improving Pre-trained Language Model Sensitivity via Mask Specific losses: A case study on Biomedical NER</i>	
Micheal Abaho, Danushka Bollegala, Gary Leeming, Dan W Joyce and Iain Edward Buchan	5013
<i>Language Models Implement Simple Word2Vec-style Vector Arithmetic</i>	
Jack Merullo, Carsten Eickhoff and Ellie Pavlick .....	5030
<i>AutoLoRA: Automatically Tuning Matrix Ranks in Low-Rank Adaptation Based on Meta Learning</i>	
Ruiyi Zhang, Rushi Qiang, Sai Ashish Somayajula and Pengtao Xie .....	5048
<i>SportQA: A Benchmark for Sports Understanding in Large Language Models</i>	
Haotian Xia, Zhengbang Yang, Yuqing Wang, Rhys Tracy, Yun Zhao, Dongdong Huang, Zezhi Chen, Yan Zhu, Yuan-fang Wang and Weining Shen .....	5061
<i>Revisiting subword tokenization: A case study on affixal negation in large language models</i>	
Thinh Hung Truong, Yulia Otmakhova, Karin Verspoor, Trevor Cohn and Timothy Baldwin	5082
<i>Generating Mental Health Transcripts with SAPE (Spanish Adaptive Prompt Engineering)</i>	
Daniel Cabrera Lozoya, Alejandro Casar Berazaluce, Juan Alberto Barajas Perches, Eloy Hernández Lúa, Mike Conway and Simon D’Alfonso .....	5096
<i>Where are you from? Geolocating Speech and Applications to Language Identification</i>	
Patrick Foley, Matthew Wiesner, Bismarck Bamfo Odoom, Leibny Paola Garcia Perera, Kenton Murray and Philipp Koehn .....	5114
<i>Teaching Language Models to Self-Improve through Interactive Demonstrations</i>	
Xiao Yu, Baolin Peng, Michel Galley, Jianfeng Gao and Zhou Yu .....	5127
<i>MAGID: An Automated Pipeline for Generating Synthetic Multi-modal Datasets</i>	
Hossein Aboutalebi, Hwanjun Song, Yusheng Xie, Arshit Gupta, Lijia Sun, Hang Su, Igor Shalymov, Nikolaos Pappas, Siffi Singh and Saab Mansour .....	5150
<i>Zero-shot Generative Linguistic Steganography</i>	
Ke Lin, Yiyang Luo, Zijian Zhang and Luo ping .....	5168
<i>Does GPT-4 pass the Turing test?</i>	
Cameron Robert Jones and Ben Bergen .....	5183
<i>Polarity Calibration for Opinion Summarization</i>	
Yuanyuan Lei, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Ruihong Huang and Dong Yu	5211
<i>Sentence-level Media Bias Analysis with Event Relation Graph</i>	
Yuanyuan Lei and Ruihong Huang .....	5225
<i>EMONA: Event-level Moral Opinions in News Articles</i>	
Yuanyuan Lei, Md Messal Monem Miah, Ayesha Qamar, Sai Ramana Reddy , Jonathan Tong, Haotian Xu and Ruihong Huang .....	5239
<i>DLM: A Decoupled Learning Model for Long-tailed Polyphone Disambiguation in Mandarin</i>	
Beibei Gao, Yangsen Zhang, Ga Xiang and Yushan Jiang .....	5252

<i>You don't need a personality test to know these models are unreliable: Assessing the Reliability of Large Language Models on Psychometric Instruments</i>	
Bangzhao Shu, Lechen Zhang, Minje Choi, Lavinia Dunagan, Lajanugen Logeswaran, Moontae Lee, Dallas Card and David Jurgens . . . . .	5263
<i>CASA: Causality-driven Argument Sufficiency Assessment</i>	
Xiao Liu, Yansong Feng and Kai-Wei Chang . . . . .	5282
<i>MacGyver: Are Large Language Models Creative Problem Solvers?</i>	
Yufei Tian, Abhilasha Ravichander, Lianhui Qin, Ronan Le Bras, Raja Marjeh, Nanyun Peng, Yejin Choi, Thomas L. Griffiths and Faeze Brahman . . . . .	5303
<i>To Translate or Not to Translate: A Systematic Investigation of Translation-Based Cross-Lingual Transfer to Low-Resource Languages</i>	
Benedikt Ebing and Goran Glavaš . . . . .	5325
<i>Enhancing Large Language Models Against Inductive Instructions with Dual-critique Prompting</i>	
Rui Wang, Hongru Wang, Fei Mi, Boyang Xue, Yi Chen, Kam-Fai Wong and Ruifeng Xu . . . . .	5345
<i>GLiNER: Generalist Model for Named Entity Recognition using Bidirectional Transformer</i>	
Urchade Zaratiana, Nadi Tomeh, Pierre Holat and Thierry Charnois . . . . .	5364
<i>XSTest: A Test Suite for Identifying Exaggerated Safety Behaviours in Large Language Models</i>	
Paul Röttger, Hannah Rose Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi and Dirk Hovy . . . . .	5377
<i>Carpe diem: On the Evaluation of World Knowledge in Lifelong Language Models</i>	
Yujin Kim, Jaehong Yoon, Seonghyeon Ye, Sangmin Bae, Namgyu Ho, Sung Ju Hwang and Se-Young Yun . . . . .	5401
<i>Fine-grained Gender Control in Machine Translation with Large Language Models</i>	
Minwoo Lee, Hyukhun Koh, Minsung Kim and Kyomin Jung . . . . .	5416
<i>DialogVCS: Robust Natural Language Understanding in Dialogue System Upgrade</i>	
Zefan Cai, Xin Zheng, Tianyu Liu, Haoran Meng, Jiaqi Han, Gang Yuan, Binghuai Lin, Baobao Chang and Yunbo Cao . . . . .	5431
<i>LLattribution: LLM-Verified Retrieval for Verifiable Generation</i>	
Xiaonan Li, Changtai Zhu, Linyang Li, Zhangyue Yin, Tianxiang Sun and Xipeng Qiu . . . . .	5453
<i>Mapping Long-term Causalities in Psychiatric Symptomatology and Life Events from Social Media</i>	
Siyuan Chen, Meilin Wang, Minghao Lv, Zhiling Zhang, Juqianqian Juqianqian, Dejyangla Dejyangla, Yujia Peng, Kenny Q. Zhu and Mengyue Wu . . . . .	5472
<i>Multimodal Chart Retrieval: A Comparison of Text, Table and Image Based Approaches</i>	
Averi Nowak, Francesco Piccinno and Yasemin Altun . . . . .	5488
<i>Retrieval Helps or Hurts? A Deeper Dive into the Efficacy of Retrieval Augmentation to Language Models</i>	
Seiji Maekawa, Hayate Iso, Sairam Gurajada and Nikita Bhutani . . . . .	5506
<i>AudioChatLlama: Towards General-Purpose Speech Abilities for LLMs</i>	
Yassir Fathullah, Chunyang Wu, Egor Lakomkin, Ke Li, Junteng Jia, Yuan Shangguan, Jay Mahadeokar, Ozlem Kalinli, Christian Fuegen and Mike Seltzer . . . . .	5522
<i>Whispers of Doubt Amidst Echoes of Triumph in NLP Robustness</i>	
Ashim Gupta, Rishanth Rajendhran, Nathan Stringham, Vivek Srikumar and Ana Marasovic . . . . .	5533

<i>Sequential Compositional Generalization in Multimodal Models</i>	
Semih Yagcioglu, Osman Batur İnce, Aykut Erdem, Erkut Erdem, Desmond Elliott and Deniz Yuret .....	5591
<i>Generating Uncontextualized and Contextualized Questions for Document-Level Event Argument Extraction</i>	
Md Nayem Uddin, Enfa Rose George, Eduardo Blanco and Steven Corman.....	5612
<i>Evidence-Driven Retrieval Augmented Response Generation for Online Misinformation</i>	
Zhenrui Yue, Huimin Zeng, Yimeng Lu, Lanyu Shang, Yang Zhang and Dong Wang.....	5628
<i>Open-Vocabulary Federated Learning with Multimodal Prototyping</i>	
Huimin Zeng, Zhenrui Yue and Dong Wang .....	5644
<i>Exploring Key Point Analysis with Pairwise Generation and Graph Partitioning</i>	
Xiao Li, Yong Jiang, Shen Huang, Pengjun Xie, Gong Cheng and Fei Huang.....	5657
<i>Understanding the Capabilities and Limitations of Large Language Models for Cultural Commonsense</i>	
Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria and Rada Mihalcea .....	5668
<i>Code Models are Zero-shot Precondition Reasoners</i>	
Lajanugen Logeswaran, Sungryull Sohn, Yiwei Lyu, Anthony Zhe Liu, Dong-Ki Kim, Dongsub Shim, Moontae Lee and Honglak Lee .....	5681
<i>Contrastive and Consistency Learning for Neural Noisy-Channel Model in Spoken Language Understanding</i>	
Suyoung Kim, Jiyeon Hwang and Ho-Young Jung .....	5698
<i>Do Large Language Models Rank Fairly? An Empirical Study on the Fairness of LLMs as Rankers</i>	
Yuan Wang, Xuyang Wu, Hsin-Tai Wu, Zhiqiang Tao and Yi Fang .....	5712
<i>TabSQLify: Enhancing Reasoning Capabilities of LLMs Through Table Decomposition</i>	
Md Mahadi Hasan Nahid and Davood Rafiei.....	5725
<i>Contextual Label Projection for Cross-Lingual Structured Prediction</i>	
Tanmay Parekh, I-Hung Hsu, Kuan-Hao Huang, Kai-Wei Chang and Nanyun Peng .....	5738
<i>Event Detection from Social Media for Epidemic Prediction</i>	
Tanmay Parekh, Anh Mac, Jiarui Yu, Yuxuan Dong, Syed Shahriar, Bonnie Liu, Eric J Yang, Kuan-Hao Huang, Wei Wang, Nanyun Peng and Kai-Wei Chang .....	5758
<i>RESPROMPT: Residual Connection Prompting Advances Multi-Step Reasoning in Large Language Models</i>	
Song Jiang, Zahra Shakeri, Aaron Chan, Maziar Sanjabi, Hamed Firooz, Yinglong Xia, Bugra Akyildiz, Yizhou Sun, Jinchao Li, Qifan Wang and Asli Celikyilmaz .....	5784
<i>BPE-knockout: Pruning Pre-existing BPE Tokenisers with Backwards-compatible Morphological Semi-supervision</i>	
Thomas Bauwens and Pieter Delobelle.....	5810
<i>How are Prompts Different in Terms of Sensitivity?</i>	
Sheng Lu, Hendrik Schuff and Iryna Gurevych.....	5833
<i>LSTDial: Enhancing Dialogue Generation via Long- and Short-Term Measurement Feedback</i>	
Guanghui Ye, Huan Zhao, Zixing Zhang, Xupeng Zha and Zhihua Jiang .....	5857

<i>The ART of LLM Refinement: Ask, Refine, and Trust</i>	
Kumar Shridhar, Koustuv Sinha, Andrew Cohen, Tianlu Wang, Ping Yu, Ramakanth Pasunuru, Mrinmaya Sachan, Jason E Weston and Asli Celikyilmaz . . . . .	5872
<i>Modularized Multilingual NMT with Fine-grained Interlingua</i>	
Sungjun Lim, Yoonjung Choi and Sangha Kim . . . . .	5884
<i>ParallelPARC: A Scalable Pipeline for Generating Natural-Language Analogies</i>	
Oren Sultan, Yonatan Bitton, Ron Yosef and Dafna Shahaf . . . . .	5900
<i>AWESOME: GPU Memory-constrained Long Document Summarization using Memory Mechanism and Global Salient Content</i>	
Shuyang Cao and Lu Wang . . . . .	5925
<i>NLP Systems That Can't Tell Use from Mention Censor Counterspeech, but Teaching the Distinction Helps</i>	
Kristina Gligoric, Myra Cheng, Lucia Zheng, Esin Durmus and Dan Jurafsky . . . . .	5942
<i>Debiasing with Sufficient Projection: A General Theoretical Framework for Vector Representations</i>	
Enze Shi, Lei Ding, Linglong Kong and Bei Jiang . . . . .	5960
<i>Semi-Supervised Dialogue Abstractive Summarization via High-Quality Pseudolabel Selection</i>	
Jianfeng He, Hang Su, Jason Cai, Igor Shalyminov, Hwanjun Song and Saab Mansour . . . . .	5976
<i>AfriMTE and AfriCOMET: Enhancing COMET to Embrace Under-resourced African Languages</i>	
Jiayi Wang, David Ifeoluwa Adelani, Sweta Agrawal, Marek Masiak, Ricardo Rei, Eleftheria Briakou, Marine Carpuat, Xuanli He, Sofia Bourhim, Andiswa Bukula, Muhidin A. Mohamed, Temitayo Olatoye, Tosin Adewumi, Hamam Mokayed, Christine Mwase, Wangui Kimotho, Foutse Yuehghoh, Anuoluwapo Aremu, Jessica Ojo, Shamsuddeen Hassan Muhammad, Salomey Osei, Abdul-Hakeem Omotayo, Chiamaka Ijeoma Chukwuneke, Perez Ogayo, Oumaima Hourrane, Salma El Anigri, Lolwethu Ndolela, Thabiso Mangwana, Shafie Abdi Mohamed, Hassan Ayinde, Oluwabusayo Olufunke Awoyomi, Lama Alkhaled, sana Sabah al-azzawi, Naome A Etori, Millicent Ochieng, Clemencia Siro, Njoroge Kiragu, Eric Muchiri, Wangari Kimotho, Toadoum Sari Sakayo, Lyse Naomi Momo Wamba, Daud Abolade, Simbiat Ajao, Iyanuoluwa Shode, Ricky Sambo Macharm, Ruqayya Nasir Iro, Saheed Salahudeen Abdullahi, Stephen Edward Moore, Bernard Opoku, Zainab Akinjobi, Abeeb Afolabi, Nnaemeka Casmir Obiefuna, Onyekachi Ogbu, Sam Brian Ochieng', Verrah Akinyi Otiende, Chinedu Emmanuel Mbonu, Yao Lu and Pontus Stenetorp . . . . .	5997
<i>TableLlama: Towards Open Large Generalist Models for Tables</i>	
Tianshu Zhang, Xiang Yue, Yifei Li and Huan Sun . . . . .	6024
<i>PEMA: An Offsite-Tunable Plug-in External Memory Adaptation for Language Models</i>	
HyunJin Kim, Young Jin Kim and JinYeong Bak . . . . .	6045
<i>Backdoor-ing Instruction-Tuned Large Language Models with Virtual Prompt Injection</i>	
Jun Yan, Vikas Yadav, Shiyang Li, Lichang Chen, Zheng Tang, Hai Wang, Vijay Srinivasan, Xiang Ren and Hongxia Jin . . . . .	6065
<i>Exploring the Factual Consistency in Dialogue Comprehension of Large Language Models</i>	
Shuaijie She, Shujian Huang, Xingyun Wang, Yanke Zhou and Jiajun Chen . . . . .	6087
<i>Multilingual Pretraining and Instruction Tuning Improve Cross-Lingual Knowledge Alignment, But Only Shallowly</i>	
Changjiang Gao, Hongda Hu, Peng Hu, Jiajun Chen, Jixing Li and Shujian Huang . . . . .	6101

<i>A Study on the Calibration of In-context Learning</i>	
Hanlin Zhang, YiFan Zhang, Yaodong Yu, Dhruv Madeka, Dean Foster, Eric P. Xing, Himabindu Lakkaraju and Sham M. Kakade . . . . .	6118
<i>DialogBench: Evaluating LLMs as Human-like Dialogue Systems</i>	
Jiao Ou, Junda Lu, Che Liu, Yihong Tang, Fuzheng Zhang, Di Zhang and Kun Gai . . . . .	6137
<i>GINopic: Topic Modeling with Graph Isomorphism Network</i>	
Suman Adhya and Debarshi Kumar Sanyal . . . . .	6171
<i>CMB: A Comprehensive Medical Benchmark in Chinese</i>	
Xidong Wang, Guiming Hardy Chen, Song Dingjie, Zhang Zhiyi, Zhihong Chen, Qingying Xiao, Junying Chen, Feng Jiang, Jianquan Li, Xiang Wan, Benyou Wang and Haizhou Li . . . . .	6184
<i>Massive End-to-end Speech Recognition Models with Time Reduction</i>	
Weiran Wang, Rohit Prabhavalkar, Haozhe Shan, Zhong Meng, Dongseong Hwang, Qiuqia Li, Khe Chai Sim, Bo Li, James Qin, Xingyu Cai, Adam Stooke, Chengjian Zheng, Yanzhang He, Tara N Sainath and Pedro J Moreno Mengibar . . . . .	6206
<i>SlimFit: Memory-Efficient Fine-Tuning of Transformer-based Models Using Training Dynamics</i>	
Arash Ardakani, Altan Haan, Shangyin Tan, Doru Thom Popovici, Alvin Cheung, Costin Iancu and Koushik Sen . . . . .	6218
<i>Effective Large Language Model Adaptation for Improved Grounding and Citation Generation</i>	
Xi Ye, Ruoxi Sun, Sercan O Arik and Tomas Pfister . . . . .	6237
<i>Assisting in Writing Wikipedia-like Articles From Scratch with Large Language Models</i>	
Yijia Shao, Yucheng Jiang, Theodore A Kanell, Peter Xu, Omar Khattab and Monica Lam . . . . .	6252
<i>Grounding Gaps in Language Model Generations</i>	
Omar Shaikh, Kristina Gligoric, Ashna Khetan, Matthias Gerstgrasser, Diyi Yang and Dan Jurafsky . . . . .	6279
<i>When Does Monolingual Data Help Multilingual Translation: The Role of Domain and Model Scale</i>	
Christos Baziotis, Biao Zhang, Alexandra Birch and Barry Haddow . . . . .	6297
<i>ContraSim – Analyzing Neural Representations Based on Contrastive Learning</i>	
Adir Rahamim and Yonatan Belinkov . . . . .	6325
<i>Universal Prompt Optimizer for Safe Text-to-Image Generation</i>	
Zongyu Wu, Hongcheng Gao, Yueze Wang, Xiang Zhang and Suhang Wang . . . . .	6340
<i>Language Model Based Unsupervised Dependency Parsing with Conditional Mutual Information and Grammatical Constraints</i>	
Junjie Chen, Xiangheng He and Yusuke Miyao . . . . .	6355
<i>The Bias Amplification Paradox in Text-to-Image Generation</i>	
Preethi Seshadri, Sameer Singh and Yanai Elazar . . . . .	6367
<i>Grammar-based Data Augmentation for Low-Resource Languages: The Case of Guarani-Spanish Neural Machine Translation</i>	
Agustín Lucas, Alexis Baladón, Victoria Pardiñas, Marvin M. Agüero-Torales, Santiago Góngora and Luis Chiruzzo . . . . .	6385
<i>Global Gallery: The Fine Art of Painting Culture Portraits through Multilingual Instruction Tuning</i>	
Anjishnu Mukherjee, Aylin Caliskan, Ziwei Zhu and Antonios Anastasopoulos . . . . .	6398



<i>Toward Interactive Regional Understanding in Vision-Large Language Models</i> Jungbeom Lee, Sanghyuk Chun and Sangdoo Yun .....	6416
<i>ScriptMix: Mixing Scripts for Low-resource Language Parsing</i> Jaeseong Lee, Dohyeon Lee and seung-won hwang .....	6430
<i>MT-PATCHER: Selective and Extendable Knowledge Distillation from Large Language Models for Machine Translation</i> Jiahuan Li, Shanbo Cheng, Shujian Huang and Jiajun Chen .....	6445
<i>ToXCL: A Unified Framework for Toxic Speech Detection and Explanation</i> Nhat M. Hoang, Do Xuan Long, Duc Anh Do, Duc Anh Vu and Anh Tuan Luu .....	6460
<i>LinkPrompt: Natural and Universal Adversarial Attacks on Prompt-based Language Models</i> Yue Xu and Wenjie Wang .....	6473
<i>CoE-SQL: In-Context Learning for Multi-Turn Text-to-SQL with Chain-of-Edits</i> Hanchong Zhang, Ruisheng Cao, Hongshen Xu, Lu Chen and Kai Yu .....	6487
<i>ContraDoc: Understanding Self-Contradictions in Documents with Large Language Models</i> Jierui Li, Vipul Raheja and Dhruv Kumar .....	6509
<i>Entity Disambiguation via Fusion Entity Decoding</i> Junxiong Wang, Ali Mousavi, Omar Attia, Ronak Pradeep, Saloni Potdar, Alexander M Rush, Umar Farooq Minhas and Yunyao Li .....	6524
<i>PlanRAG: A Plan-then-Retrieval Augmented Generation for Generative Large Language Models as Decision Makers</i> Myeonghwa Lee, Seonho An and Min-Soo Kim .....	6537
<i>GPTScore: Evaluate as You Desire</i> Jinlan Fu, See-Kiong Ng, Zhengbao Jiang and Pengfei Liu .....	6556
<i>A Survey of Confidence Estimation and Calibration in Large Language Models</i> Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov and Iryna Gurevych ..	6577
<i>Not All Metrics Are Guilty: Improving NLG Evaluation by Diversifying References</i> Tianyi Tang, Hongyuan Lu, Yuchen Eleanor Jiang, Haoyang Huang, Dongdong Zhang, Xin Zhao, Tom Kocmi and Furu Wei .....	6596
<i>Separation and Fusion: A Novel Multiple Token Linking Model for Event Argument Extraction</i> Jing Xu, Dandan Song, Siu Cheung Hui, Zhijing Wu, Meihuizi Jia, Hao Wang, Yanru Zhou, Changzhi Zhou and Ziyi Yang .....	6611
<i>The Integration of Semantic and Structural Knowledge in Knowledge Graph Entity Typing</i> Muzhi Li, Minda Hu, Irwin King and Ho-fung Leung .....	6625
<i>ComCLIP: Training-Free Compositional Image and Text Matching</i> Kenan Jiang, Xuehai He, Ruize Xu and Xin Eric Wang .....	6639
<i>ACLSum: A New Dataset for Aspect-based Summarization of Scientific Publications</i> Sotaro Takeshita, Tommaso Green, Ines Reinig, Kai Eckert and Simone Paolo Ponzetto . . .	6660
<i>XAL: EXplainable Active Learning Makes Classifiers Better Low-resource Learners</i> Yun Luo, Zhen Yang, Fandong Meng, Yingjie Li, Fang Guo, Qinglin Qi, Jie Zhou and Yue Zhang	6676

<i>LaDiC: Are Diffusion Models Really Inferior to Autoregressive Counterparts for Image-to-Text Generation?</i>	
Yuchi Wang, Shuhuai Ren, Rundong Gao, Linli Yao, Qingyan Guo, Kaikai An, Jianhong Bai and Xu Sun .....	6699
<i>Intent-conditioned and Non-toxic Counterspeech Generation using Multi-Task Instruction Tuning with RLAIF</i>	
Amey Hengle, Aswini Kumar Padhi, Sahajpreet Singh, Anil Bandhakavi, Md Shad Akhtar and Tanmoy Chakraborty .....	6716
<i>Attacks, Defenses and Evaluations for LLM Conversation Safety: A Survey</i>	
Zhichen Dong, Zhanhui Zhou, Chao Yang, Jing Shao and Yu Qiao .....	6734
<i>Mind’s Mirror: Distilling Self-Evaluation Capability and Comprehensive Thinking from Large Language Models</i>	
Weize Liu, Guocong Li, Kai Zhang, Bang Du, Qiyuan Chen, Xuming Hu, Hongxia Xu, Jintai Chen and Jian Wu .....	6748
<i>Divergent Token Metrics: Measuring degradation to prune away LLM components – and optimize quantization</i>	
Björn Deiseroth, Max Meuer, Nikolas Gritsch, Constantin Eichenberg, Patrick Schramowski, Matthias Aßenmacher and Kristian Kersting .....	6764
<i>Beyond Performance: Quantifying and Mitigating Label Bias in LLMs</i>	
Yuval Reif and Roy Schwartz .....	6784
<i>Instructing Large Language Models to Identify and Ignore Irrelevant Conditions</i>	
Zhenyu Wu, Chao Shen and Meng Jiang .....	6799
<i>Lower Bounds on the Expressivity of Recurrent Neural Language Models</i>	
Anej Svete, Franz Nowak, Anisha Mohamed Sahabdeen and Ryan Cotterell .....	6820
<i>Transformers Can Represent n-gram Language Models</i>	
Anej Svete and Ryan Cotterell .....	6841
<i>The Role of n-gram Smoothing in the Age of Neural Networks</i>	
Luca Malagutti, Andrius Buinovskij, Anej Svete, Clara Meister, Afra Amini and Ryan Cotterell	6875
<i>Reliability Estimation of News Media Sources: Birds of a Feather Flock Together</i>	
Sergio Burdisso, Dairazalia Sanchez-cortes, Esaú Villatoro-tello and Petr Motlicek .....	6893
<i>On the Multilingual Ability of Decoder-based Pre-trained Language Models: Finding and Controlling Language-Specific Neurons</i>	
Takeshi Kojima, Itsuki Okimura, Yusuke Iwasawa, Hitomi Yanaka and Yutaka Matsuo .....	6912
<i>NLP Progress in Indigenous Latin American Languages</i>	
Atnafu Lambebo Tonja, Fazlourrahman Balouchzahi, Sabur Butt, Olga Kolesnikova, Hector Ceballos, Alexander Gelbukh and Tamar Solorio .....	6965
<i>On the Effectiveness of Adversarial Robustness for Abuse Mitigation with Counterspeech</i>	
Yi-Ling Chung and Jonathan Bright .....	6981
<i>Leveraging the Structure of Pre-trained Embeddings to Minimize Annotation Effort</i>	
Cesar Gonzalez-Gutierrez and Ariadna Quattoni .....	6996

<i>UniArk: Improving Generalisation and Consistency for Factual Knowledge Extraction through Debiasing</i>	
Yijun Yang, Jie He, Pinzhen Chen, Victor Gutierrez Basulto and Jeff Z. Pan	7011
<i>Adaptive-RAG: Learning to Adapt Retrieval-Augmented Large Language Models through Question Complexity</i>	
Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang and Jong C. Park	7029
<i>Knowing What LLMs DO NOT Know: A Simple Yet Effective Self-Detection Method</i>	
Yukun Zhao, Lingyong Yan, Weiwei Sun, Guoliang Xing, Chong Meng, Shuaiqiang Wang, Zhi-cong Cheng, Zhaochun Ren and Dawei Yin	7044
<i>Are Large Language Model Temporally Grounded?</i>	
Yifu Qiu, Zheng Zhao, Yftah Ziser, Anna Korhonen, Edoardo Ponti and Shay B Cohen	7057
<i>Document Image Machine Translation with Dynamic Multi-pre-trained Models Assembling</i>	
Yupu Liang, Yaping Zhang, Cong MA, Zhiyang Zhang, Yang Zhao, Lu Xiang, Chengqing Zong and Yu Zhou	7077
<i>Elastic Weight Removal for Faithful and Abstractive Dialogue Generation</i>	
Nico Daheim, Nouha Dziri, Mrinmaya Sachan, Iryna Gurevych and Edoardo Ponti	7089
<i>R-Tuning: Instructing Large Language Models to Say ‘I Don’t Know’</i>	
Hanning Zhang, Shizhe Diao, Yong Lin, Yi Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji and Tong Zhang	7106
<i>Bridging the Gap between Different Vocabularies for LLM Ensemble</i>	
Yangyifan Xu, Jinliang Lu and Jiajun Zhang	7133
<i>KnowLA: Enhancing Parameter-efficient Finetuning with Knowledgeable Adaptation</i>	
Xindi Luo, Zequn Sun, Jing Zhao, Zhe Zhao and Wei Hu	7146
<i>Extremely Weakly-supervised Text Classification with Wordsets Mining and Sync-Denoising</i>	
Lysa Xiao	7160
<i>F-MALLOC: Feed-forward Memory Allocation for Continual Learning in Neural Machine Translation</i>	
Junhong Wu, Yuchen Liu and Chengqing Zong	7173
<i>Towards Reducing Diagnostic Errors with Interpretable Risk Prediction</i>	
Denis Jered McInerney, William Dickinson, Lucy C. Flynn, Andrea C Young, Geoffrey Young, Jan-Willem van de Meent and Byron C Wallace	7186
<i>Generalizable Multilingual Hate Speech Detection on Low Resource Indian Languages using Fair Selection in Federated Learning</i>	
Akshay Singh and Rahul Thakur	7204
<i>Key ingredients for effective zero-shot cross-lingual knowledge transfer in generative tasks</i>	
Nadezhda Chirkova and Vassilina Nikoulina	7215
<i>The Impact of Depth on Compositional Generalization in Transformer Language Models</i>	
Jackson Petty, Sjoerd van Steenkiste, Ishita Dasgupta, Fei Sha, Dan Garrette and Tal Linzen	7232
<i>Pregnant Questions: The Importance of Pragmatic Awareness in Maternal Health Question Answering</i>	
Neha Srikanth, Rupak Sarkar, Heran Mane, Elizabeth Aparicio, Quynh Nguyen, Rachel Rudinger and Jordan Lee Boyd-Graber	7246
<i>Towards Explainability in Legal Outcome Prediction Models</i>	
Josef Valvoda and Ryan Cotterell	7262

<i>The steerability of large language models toward data-driven personas</i> Junyi Li, Charith Peris, Ninareh Mehrabi, Palash Goyal, Kai-Wei Chang, Aram Galstyan, Richard Zemel and Rahul Gupta . . . . .	7283
<i>CCSum: A Large-Scale and High-Quality Dataset for Abstractive News Summarization</i> Xiang Jiang and Markus Dreyer . . . . .	7299
<i>Capturing Perspectives of Crowdsourced Annotators in Subjective Learning Tasks</i> Negar Mokhberian, Myrl G Marmarelis, Frederic Rene Hopp, Valerio Basile, Fred Morstatter and Kristina Lerman . . . . .	7330
<i>Improving Factual Accuracy of Neural Table-to-Text Output by Addressing Input Problems in ToTTo</i> Barkavi Sundararajan, Yaji Sripada and Ehud Reiter . . . . .	7343
<i>CERET: Cost-Effective Extrinsic Refinement for Text Generation</i> Jason Cai, Hang Su, Monica Sunkara, Igor Shalyminov and Saab Mansour . . . . .	7370
<i>Parameter-Efficient Instruction Tuning of Large Language Models For Extreme Financial Numeral Labelling</i> Subhendu Khatuya, Rajdeep Mukherjee, Akash Ghosh, Manjunath Hegde, Koustuv Dasgupta, Niloy Ganguly, Saptarshi Ghosh and Pawan Goyal . . . . .	7384
<i>Analysis of State-Level Legislative Process in Enhanced Linguistic and Nationwide Network Contexts</i> Maryam Davoodi and Dan Goldwasser . . . . .	7397
<i>DeMuX: Data-efficient Multilingual Learning</i> Simran Khanuja, Srinivas Gowriraj, Lucio M. Dery and Graham Neubig . . . . .	7416
<i>DUQGen: Effective Unsupervised Domain Adaptation of Neural Rankers by Diversifying Synthetic Query Generation</i> Ramraj Chandradevan, Kaustubh Dhole and Eugene Agichtein . . . . .	7430
<i>How did we get here? Summarizing conversation dynamics</i> Yilun Hua, Nicholas Chernogor, Yuzhe Gu, Seoyeon Julie Jeong, Miranda Luo and Cristian Danescu-Niculescu-Mizil . . . . .	7445
<i>Can Language Model Moderators Improve the Health of Online Discourse?</i> Hyundong Justin Cho, Shuai Liu, Taiwei Shi, Darpan Jain, Basem Rizk, Yuyang Huang, Zixun Lu, Nuan Wen, Jonathan Gratch, Emilio Ferrara and Jonathan May . . . . .	7471
<i>LeanReasoner: Boosting Complex Logical Reasoning with Lean</i> Dongwei Jiang, Marcio Fonseca and Shay B Cohen . . . . .	7490
<i>UICoder: Finetuning Large Language Models to Generate User Interface Code through Automated Feedback</i> Jason Wu . . . . .	7504
<i>Measuring Cross-lingual Transfer in Bytes</i> Leandro Rodrigues De Souza, Thales Sales Almeida, Roberto Lotufo and Rodrigo Frassetto Nogueira . . . . .	7519
<i>MisgenderMender: A Community-Informed Approach to Interventions for Misgendering</i> Tamanna Hossain, Sunipa Dev and Sameer Singh . . . . .	7531
<i>Interplay of Machine Translation, Diacritics, and Diacritization</i> Wei-Rui Chen, Ife Adebara and Muhammad Abdul-Mageed . . . . .	7552

<i>From Quantity to Quality: Boosting LLM Performance with Self-Guided Data Selection for Instruction Tuning</i>	
Ming Li, Yong Zhang, Zhitao Li, Jiuhai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou and Jing Xiao .....	7595
<i>Safer-Instruct: Aligning Language Models with Automated Preference Data</i>	
Taiwei Shi, Kai Chen and Jieyu Zhao .....	7629
<i>PELMS: Pre-training for Effective Low-Shot Multi-Document Summarization</i>	
Joseph J Peper, Wenzhao Qiu and Lu Wang .....	7645
<i>Deceptive Semantic Shortcuts on Reasoning Chains: How Far Can Models Go without Hallucination?</i>	
Bangzheng Li, Ben Zhou, Fei Wang, Xingyu Fu, Dan Roth and Muhao Chen .....	7668
<i>IndiSentiment140: Sentiment Analysis Dataset for Indian Languages with Emphasis on Low-Resource Languages using Machine Translation</i>	
Saurabh Kumar, Ranbir Singh Sanasam and Sukumar Nandi .....	7682
<i>Leveraging LLMs for Synthesizing Training Data Across Many Languages in Multilingual Dense Retrieval</i>	
Nandan Thakur, Jianmo Ni, Gustavo Hernandez Abrego, John Frederick Wieting, Jimmy Lin and Daniel Cer .....	7692
<i>SCANNER: Knowledge-Enhanced Approach for Robust Multi-modal Named Entity Recognition of Unseen Entities</i>	
Hyunjong Ok, Taeho Kil, Sukmin Seo and Jaeho Lee .....	7718
<i>A Theory Guided Scaffolding Instruction Framework for LLM-Enabled Metaphor Reasoning</i>	
Yuan Tian, Nan Xu and Wenji Mao .....	7731
<i>Learning to Compress Prompt in Natural Language Formats</i>	
Yu-Neng Chuang, Tianwei Xing, Chia-Yuan Chang, Zirui Liu, Xun Chen and Xia Hu .....	7749
<i>Automatic, Meta and Human Evaluation for Multimodal Summarization with Multimodal Output</i>	
Haojie Zhuang, Wei Emma Zhang, Leon Xie, Weitong Chen, Jian Yang and Quan Z. Sheng .....	7761
<i>Naive Bayes-based Context Extension for Large Language Models</i>	
Jianlin Su, Murtadha Ahmed, Bo Wen, Luo Ao, Mingren Zhu and Yunfeng Liu .....	7784
<i>Leitner-Guided Memory Replay for Cross-lingual Continual Learning</i>	
Meryem M’hamdi and Jonathan May .....	7801
<i>Multilingual Nonce Dependency Treebanks: Understanding how Language Models Represent and Process Syntactic Structure</i>	
David Arps, Laura Kallmeyer, Younes Samih and Hassan Sajjad .....	7815
<i>Actively Learn from LLMs with Uncertainty Propagation for Generalized Category Discovery</i>	
Jinggui Liang, Lizi Liao, Hao Fei, Bobo Li and Jing Jiang .....	7838
<i>Explaining Text Similarity in Transformer Models</i>	
Alexandros Vasileiou and Oliver Eberle .....	7852
<i>Large Language Models can Contrastively Refine their Generation for Better Sentence Representation Learning</i>	
Huiming Wang, Zhaodonghui Li, Liying Cheng, De Wen Soh and Lidong Bing .....	7867
<i>HIL: Hybrid Isotropy Learning for Zero-shot Performance in Dense retrieval</i>	
Jaeyoung Kim, Dohyeon Lee and seung-won hwang .....	7885

<i>SuperGLEBer: German Language Understanding Evaluation Benchmark</i> Jan Pfister and Andreas Hotho .....	7897
<i>"You are an expert annotator": Automatic Best–Worst-Scaling Annotations for Emotion Intensity Modeling</i> Christopher Bagdon, Prathamesh Karmalkar, Harsha Gurulingappa and Roman Klinger . . .	7917
<i>What Matters in Training a GPT4-Style Language Model with Multimodal Inputs?</i> Yan Zeng, Hanbo Zhang, Jiani zheng, Jiangnan Xia, Guoqiang Wei, Yang Wei, Yuchen Zhang, Tao Kong and Ruihua Song .....	7930
<i>Defining and Detecting Vulnerability in Human Evaluation Guidelines: A Preliminary Study Towards Reliable NLG Evaluation</i> Jie Ruan, WangWenqing WangWenqing and Xiaojun Wan .....	7958
<i>MOSAICo: a Multilingual Open-text Semantically Annotated Interlinked Corpus</i> Simone Conia, Edoardo Barba, Abelardo Carlos Martinez Lorenzo, Pere-Lluís Huguet Cabot, Riccardo Orlando, Luigi Procopio and Roberto Navigli .....	7983
<i>SemRoDe: Macro Adversarial Training to Learn Representations that are Robust to Word-Level Attacks</i> Brian Formento, Wenjie Feng, Chuan-Sheng Foo, Anh Tuan Luu and See-Kiong Ng .....	7998
<i>BUST: Benchmark for the evaluation of detectors of LLM-Generated Text</i> Joseph Cornelius, Oscar William Lithgow-Serrano, Sandra Mitrovic, Ljiljana Dolamic and Fabio Rinaldi .....	8022
<i>Improving In-context Learning of Multilingual Generative Language Models with Cross-lingual Alignment</i> Chong Li, Shaonan Wang, Jiajun Zhang and Chengqing Zong .....	8051
<i>MaCSC: Towards Multimodal-augmented Pre-trained Language Models via Conceptual Prototypes and Self-balancing Calibration</i> Xianwei Zhuang, Zhichang Wang, Xuxin Cheng, Yuxin Xie, Liming Liang and Yuexian Zou	8070
<i>Does Pre-trained Language Model Actually Infer Unseen Links in Knowledge Graph Completion?</i> Yusuke Sakai, Hidetaka Kamigaito, Katsuhiko Hayashi and Taro Watanabe .....	8084
<i>Discovering Lobby-Parliamentarian Alignments through NLP</i> Aswin Suresh, Lazar Radojević, Francesco Salvi, Antoine Magron, Victor Kristof and Matthias Grossglauser .....	8100
<i>IterCQR: Iterative Conversational Query Reformulation with Retrieval Guidance</i> Yunah Jang, Kang-il Lee, Hyunkyung Bae, Hwanhee Lee and Kyomin Jung .....	8114
<i>AceGPT, Localizing Large Language Models in Arabic</i> Huang Huang, Fei Yu, Jianqing Zhu, Xuening Sun, Hao Cheng, Song Dingjie, Zhihong Chen, Mosen Alharthi, Bang An, Juncai He, Ziche Liu, Junying Chen, Jianquan Li, Benyou Wang, Lian Zhang, Ruoyu Sun, Xiang Wan, Haizhou Li and Jinchao Xu .....	8132
<i>Improving Machine Translation with Human Feedback: An Exploration of Quality Estimation as a Reward Model</i> Zhiwei He, Xing Wang, Wenxiang Jiao, Zhuosheng Zhang, Rui Wang, Shuming Shi and Zhaopeng Tu .....	8157
<i>Depression Detection in Clinical Interviews with LLM-Empowered Structural Element Graph</i> Zhuang Chen, Jiawen Deng, Jinfeng Zhou, Jincenzi Wu, Tiejun Qian and Minlie Huang . . .	8174

<i>SQATIN: Supervised Instruction Tuning Meets Question Answering for Improved Dialogue NLU</i> Evgeniia Razumovskaia, Goran Glavaš, Anna Korhonen and Ivan Vulić .....	8188
<i>Enhancing Argument Summarization: Prioritizing Exhaustiveness in Key Point Generation and Introducing an Automatic Coverage Evaluation Metric</i> Mohammad Khosravani, Chenyang Huang and Amine Trabelsi .....	8205
<i>ARM: Alignment with Residual Energy-Based Model</i> Bo Pang, Caiming Xiong and Yingbo Zhou .....	8218
<i>HumanRankEval: Automatic Evaluation of LMs as Conversational Assistants</i> Milan Gritta, Gerasimos Lampouras and Ignacio Iacobacci .....	8230
<i>FAMuS: Frames Across Multiple Sources</i> Siddharth Vashishtha, Alexander Martin, William Gantt, Benjamin Van Durme and Aaron Steven White .....	8243
<i>Rationale-based Opinion Summarization</i> Haoyuan Li and Snigdha Chaturvedi .....	8267
<i>Mustango: Toward Controllable Text-to-Music Generation</i> Jan Melechovsky, Zixun Guo, Deepanway Ghosal, Navonil Majumder, Dorien Herremans and Soujanya Poria .....	8286
<i>Adaptive Cross-lingual Text Classification through In-Context One-Shot Demonstrations</i> Emilio Villa Cueva, Adrian Pastor Lopez Monroy, Fernando Sánchez-Vega and Thamar Solorio	8310
<i>CNER: Concept and Named Entity Recognition</i> Giuliano Martinelli, Francesco Maria Molfese, Simone Tedeschi, Alberte Fernández-Castro and Roberto Navigli .....	8329
<i>Branch-Solve-Merge Improves Large Language Model Evaluation and Generation</i> Swarnadeep Saha, Omer Levy, Asli Celikyilmaz, Mohit Bansal, Jason E Weston and Xian Li	8345
<i>REPLUG: Retrieval-Augmented Black-Box Language Models</i> Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer and Wen-tau Yih .....	8364
<i>David helps Goliath: Inference-Time Collaboration Between Small Specialized and Large General Diffusion LMs</i> Xiaochuang Han, Sachin Kumar, Yulia Tsvetkov and Marjan Ghazvininejad .....	8378
<i>Efficient End-to-End Visual Document Understanding with Rationale Distillation</i> Wang Zhu, Alekh Agarwal, Mandar Joshi, Robin Jia, Jesse Thomason and Kristina Toutanova	8394
<i>A Systematic Comparison of Syllogistic Reasoning in Humans and Language Models</i> Tiwalayo Eisape, Michael Henry Tessler, Ishita Dasgupta, Fei Sha, Sjoerd van Steenkiste and Tal Linzen .....	8418
<i>AnchorAL: Computationally Efficient Active Learning for Large and Imbalanced Datasets</i> Pietro Lesci and Andreas Vlachos .....	8438
<i>ICLE++: Modeling Fine-Grained Traits for Holistic Essay Scoring</i> Shengjie Li and Vincent Ng .....	8458

<i>UNcommonsense Reasoning: Abductive Reasoning about Uncommon Situations</i> Wenting Zhao, Justin T Chiu, Jena D. Hwang, Faeze Brahman, Jack Hessel, Sanjiban Choudhury, Yejin Choi, Xiang Lorraine Li and Alane Suhr .....	8479
<i>To Tell The Truth: Language of Deception and Language Models</i> Sanchaita Hazra and Bodhisattwa Prasad Majumder .....	8498
<i>Multilingual Models for ASR in Chibchan Languages</i> Rolando Coto-Solano, Tai Wan Kim, Alexander Jones and Sharid Loáiciga .....	8513
<i>LegalDiscourse: Interpreting When Laws Apply and To Whom</i> Alexander Spangher, Zihan Xue, Te-Lin Wu, Mark Hansen and Jonathan May .....	8528
<i>X-Eval: Generalizable Multi-aspect Text Evaluation via Augmented Instruction Tuning with Auxiliary Evaluation Aspects</i> Minqian Liu, Ying Shen, Zhiyang Xu, Yixin Cao, Eunah Cho, Vaibhav Kumar, Reza Ghanadan and Lifu Huang .....	8552
<i>Is Reference Necessary in the Evaluation of NLG Systems? When and Where?</i> Shuqian Sheng, Yi Xu, Luoyi Fu, Jiaxin Ding, Lei Zhou, Xinbing Wang and Chenghu Zhou	8572
<i>Semi-Structured Chain-of-Thought: Integrating Multiple Sources of Knowledge for Improved Language Model Reasoning</i> Xin Su, Tiej Le, Steven Bethard and Phillip Howard .....	8589
<i>Evaluating the Deductive Competence of Large Language Models</i> S M Seals and Valerie Shalin .....	8606
<i>Large Human Language Models: A Need and the Challenges</i> Nikita Soni, H. Schwartz, João Sedoc and Niranjana Balasubramanian .....	8623
<i>On Learning to Summarize with Large Language Models as References</i> Yixin Liu, Kejian Shi, Katherine S He, Longtian Ye, Alexander Fabbri, Pengfei Liu, Dragomir Radev and Arman Cohan .....	8639
<i>Hallucination Diversity-Aware Active Learning for Text Summarization</i> Yu Xia, Xu Liu, Tong Yu, Sungchul Kim, Ryan A. Rossi, Anup Rao, Tung Mai and Shuai Li	8657
<i>Keep it Private: Unsupervised Privatization of Online Text</i> Calvin Bao and Marine Carpuat .....	8670
<i>Tied-LoRA: Enhancing parameter efficiency of LoRA with Weight Tying</i> Adithya Renduchintala, Tugrul Konuk and Oleksii Kuchaiev .....	8686
<i>Investigating Data Contamination in Modern Benchmarks for Large Language Models</i> Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gerstein and Arman Cohan .....	8698
<i>Pre-trained Language Models for Entity Blocking: A Reproducibility Study</i> Runhui Wang and Yongfeng Zhang .....	8712
<i>RE<sup>2</sup>: Region-Aware Relation Extraction from Visually Rich Documents</i> Pritika Ramu, Sijia Wang, Lalla Mouatadid, Joy Rimchala and Lifu Huang .....	8723
<i>Mix-Initiative Response Generation with Dynamic Prefix Tuning</i> Yuxiang Nie, Heyan Huang, Xian-Ling Mao and Lizi Liao .....	8740



<i>Value FULCRA: Mapping Large Language Models to the Multidimensional Spectrum of Basic Human Value</i>	
Jing Yao, Xiaoyuan Yi, Yifan Gong, Xiting Wang and Xing Xie .....	8754
<i>IndiBias: A Benchmark Dataset to Measure Social Biases in Language Models for Indian Context</i>	
Nihar Ranjan Sahoo, Pranamya Prashant Kulkarni, Arif Ahmad, Tanu Goyal, Natjis Asad, Aparna Garimella and Pushpak Bhattacharyya .....	8778