

Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)

Miami, Florida, USA
12-16 November 2024

Volume 1 of 29

ISBN: 979-8-3313-0837-7

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. (2025)

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

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
Web: www.proceedings.com

TABLE OF CONTENTS

VOLUME 1

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

Fine-Tuning Or Retrieval? Comparing Knowledge Injection in LLMs	237
<i>Oded Ovadia, Menachem Brief, Moshik Mishaeli, Oren Elisha</i>	
Systematic Biases in LLM Simulations of Debates.....	251
<i>Amir Taubenfeld, Yaniv Dover, Roi Reichart, Ariel Goldstein</i>	
Studying and Mitigating Biases in Sign Language Understanding Models	268
<i>Katherine Atwell, Danielle Bragg, Malihe Alikhani</i>	
Uncertainty in Language Models: Assessment Through Rank-Calibration.....	284
<i>Xinmeng Huang, Shuo Li, Mengxin Yu, Matteo Sesia, Hamed Hassani, Insup Lee, Osbert Bastani, Edgar Dobriban</i>	
RoTBench: A Multi-Level Benchmark for Evaluating the Robustness of Large Language Models in Tool Learning	313
<i>Junjie Ye, Yilong Wu, Songyang Gao, Caishuang Huang, Sixian Li, Guanyu Li, Xiaoran Fan, Qi Zhang, Tao Gui, Xuanjing Huang</i>	
Learning Planning-Based Reasoning by Trajectories Collection and Process Reward Synthesizing.....	334
<i>Fangkai Jiao, Chengwei Qin, Zhengyuan Liu, Nancy F. Chen, Shafiq Joty</i>	
Scaling Properties of Speech Language Models.....	351
<i>Santiago Cuervo, Ricard Marxer</i>	
“We Demand Justice!”: Towards Social Context Grounding of Political Texts	362
<i>Rajkumar Pujari, Chengfei Wu, Dan Goldwasser</i>	
An Experimental Analysis on Evaluating Patent Citations.....	373
<i>Rabindra Nath Nandi, Suman Maity, Brian Uzzi, Sourav Medya</i>	
Fine-Tuning Large Language Models to Translate: Will a Touch of Noisy Data in Misaligned Languages Suffice?	388
<i>Dawei Zhu, Pinzhen Chen, Miaoran Zhang, Barry Haddow, Xiaoyu Shen, Dietrich Klakow</i>	
Consolidating Ranking and Relevance Predictions of Large Language Models Through Post-Processing.....	410
<i>Le Yan, Zhen Qin, Honglei Zhuang, Rolf Jagerman, Xuanhui Wang, Michael Bendersky, Harrie Oosterhuis</i>	
Strength Lies in Differences! Improving Strategy Planning for Non-Collaborative Dialogues Via Diversified User Simulation	424
<i>Tong Zhang, Chen Huang, Yang Deng, Hongru Liang, Jia Liu, Zujie Wen, Wenqiang Lei, Tat-Seng Chua</i>	
Impeding LLM-Assisted Cheating in Introductory Programming Assignments Via Adversarial Perturbation	445
<i>Saiful Islam Salim, Rubin Yuchan Yang, Alexander Cooper, Suryashree Ray, Saumya Debray, Sazzadur Rahaman</i>	
Clustering and Ranking: Diversity-Preserved Instruction Selection Through Expert-Aligned Quality Estimation.....	464
<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, Jingbo Zhu</i>	
On the Influence of Gender and Race in Romantic Relationship Prediction from Large Language Models.....	479
<i>Abhilasha Sancheti, Haozhe An, Rachel Rudinger</i>	

EmphAssess : A Prosodic Benchmark on Assessing Emphasis Transfer in Speech-To-Speech Models.....	495
<i>Maureen De Seyssel, Antony D’Avirro, Adina Williams, Emmanuel Dupoux</i>	
On Fake News Detection with LLM Enhanced Semantics Mining.....	508
<i>Xiaoxiao Ma, Yuchen Zhang, Kaize Ding, Jian Yang, Jia Wu, Hao Fan</i>	
On Sensitivity of Learning with Limited Labelled Data to the Effects of Randomness: Impact of Interactions and Systematic Choices	522
<i>Branislav Pecher, Ivan Srba, Maria Bielikova</i>	
Evaluating the Instruction-Following Robustness of Large Language Models to Prompt Injection	557
<i>Zekun Li, Baolin Peng, Pengcheng He, Xifeng Yan</i>	

VOLUME 2

A Study of Nationality Bias in Names and Perplexity Using Off-The-Shelf Affect-Related Tweet Classifiers	569
<i>Valentin Barriere, Sebastian Cifuentes</i>	
Mitigating the Alignment Tax of RLHF	580
<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, Tong Zhang</i>	
Evaluating Readability and Faithfulness of Concept-Based Explanations	607
<i>Meng Li, Haoran Jin, Ruixuan Huang, Zhihao Xu, Defu Lian, Zijia Lin, Di Zhang, Xiting Wang</i>	
Personality-Aware Student Simulation for Conversational Intelligent Tutoring Systems	626
<i>Zhengyuan Liu, Stella Xin Yin, Geyu Lin, Nancy F. Chen</i>	
MSI-Agent: Incorporating Multi-Scale Insight into Embodied Agents for Superior Planning and Decision-Making	643
<i>Dayuan Fu, Biqing Qi, Yihuai Gao, Che Jiang, Guanting Dong, Bowen Zhou</i>	
CoCoLoFa: A Dataset of News Comments with Common Logical Fallacies Written by LLM-Assisted Crowds	660
<i>Min-Hsuan Yeh, Ruyuan Wan, Ting-Hao Kenneth Huang</i>	
Tokenization is More than Compression	678
<i>Craig W Schmidt, Varshini Reddy, Haoran Zhang, Alec Alameddine, Omri Uzan, Yuval Pinter, Chris Tanner</i>	
FLIRT: Feedback Loop In-Context Red Teaming	703
<i>Ninareh Mehrabi, Palash Goyal, Christophe Dupuy, Qian Hu, Shalini Ghosh, Richard Zemel, Kai-Wei Chang, Aram Galstyan, Rahul Gupta</i>	
Successfully Guiding Humans with Imperfect Instructions by Highlighting Potential Errors and Suggesting Corrections.....	719
<i>Lingjun Zhao, Khanh Xuan Nguyen, Hal Daumé III</i>	
Parameter-Efficient Sparsity Crafting from Dense to Mixture-Of-Experts for Instruction Tuning on General Tasks	737
<i>Haoyuan Wu, Haisheng Zheng, Zhuolun He, Bei Yu</i>	

GeoGPT4V: Towards Geometric Multi-Modal Large Language Models with Geometric Image Generation	750
<i>Shihao Cai, Keqin Bao, Hangyu Guo, Jizhi Zhang, Jun Song, Bo Zheng</i>	
DyVo: Dynamic Vocabularies for Learned Sparse Retrieval with Entities	767
<i>Thong Nguyen, Shubham Chatterjee, Sean Macavaney, Iain Mackie, Jeff Dalton, Andrew Yates</i>	
Let the Expert Stick to His Last: Expert-Specialized Fine-Tuning for Sparse Architectural Large Language Models	784
<i>Zihan Wang, Deli Chen, Damai Dai, Runxin Xu, Zhuoshu Li, Yu Wu</i>	
LongEmbed: Extending Embedding Models for Long Context Retrieval.....	802
<i>Dawei Zhu, Liang Wang, Nan Yang, Yifan Song, Wenhao Wu, Furu Wei, Sujian Li</i>	
Making Large Language Models Better Reasoners with Orchestrated Streaming Experiences.....	817
<i>Xiangyang Liu, Junliang He, Xipeng Qiu</i>	
Overcome Noise and Bias: Segmentation-Aided Multi-Granularity Denoising and Debiasing for Enhanced Quarduples Extraction in Dialogue.....	839
<i>Xianlong Luo, Meng Yang, Yihao Wang</i>	
Integrating Plutchik’s Theory with Mixture of Experts for Enhancing Emotion Classification.....	857
<i>Dongjun Lim, Yun-Gyung Cheong</i>	
In-Context Contrastive Learning for Event Causality Identification.....	868
<i>Liang Chao, Wei Xiang, Bang Wang</i>	
What’s Mine Becomes Yours: Defining, Annotating and Detecting Context-Dependent Paraphrases in News Interview Dialogs	882
<i>Anna Wegmann, Tijs A. Van Den Broek, Dong Nguyen</i>	
Language Models Learn Rare Phenomena from Less Rare Phenomena: The Case of the Missing AANNs.....	913
<i>Kanishka Misra, Kyle Mahowald</i>	
Large Language Models for Data Annotation and Synthesis: A Survey.....	930
<i>Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansooreh Karami, Jundong Li, Lu Cheng, Huan Liu</i>	
Chain-Of-Dictionary Prompting Elicits Translation in Large Language Models	958
<i>Hongyuan Lu, Haoran Yang, Haoyang Huang, Dongdong Zhang, Wai Lam, Furu Wei</i>	
AdaZeta: Adaptive Zeroth-Order Tensor-Train Adaption for Memory-Efficient Large Language Models Fine-Tuning	977
<i>Yifan Yang, Kai Zhen, Ershad Banijamali, Athanasios Mouchtaris, Zheng Zhang</i>	
RoseLoRA: Row and Column-Wise Sparse Low-Rank Adaptation of Pre-Trained Language Model for Knowledge Editing and Fine-Tuning.....	996
<i>Haoyu Wang, Tianci Liu, Ruirui Li, Monica Xiao Cheng, Tuo Zhao, Jing Gao</i>	
BlendFilter: Advancing Retrieval-Augmented Large Language Models Via Query Generation Blending and Knowledge Filtering.....	1009
<i>Haoyu Wang, Ruirui Li, Haoming Jiang, Jinjin Tian, Zhengyang Wang, Chen Luo, Xianfeng Tang, Monica Xiao Cheng, Tuo Zhao, Jing Gao</i>	

HEART-Felt Narratives: Tracing Empathy and Narrative Style in Personal Stories with LLMs	1026
<i>Jocelyn J Shen, Joel Mire, Hae Won Park, Cynthia Breazeal, Maarten Sap</i>	
Eliminating Biased Length Reliance of Direct Preference Optimization Via Down-Sampled KL Divergence.....	1047
<i>Junru Lu, Jiazheng Li, Siyu An, Meng Zhao, Yulan He, Di Yin, Xing Sun</i>	
Bridging Cultures in the Kitchen: A Framework and Benchmark for Cross-Cultural Recipe Retrieval	1068
<i>Tianyi Hu, Maria Maistro, Daniel Herscovich</i>	
RULE: Reliable Multimodal RAG for Factuality in Medical Vision Language Models.....	1081
<i>Peng Xia, Kangyu Zhu, Haoran Li, Hongtu Zhu, Yun Li, Gang Li, Linjun Zhang, Huaxiu Yao</i>	
CryptoTrade: A Reflective LLM-Based Agent to Guide Zero-Shot Cryptocurrency Trading.....	1094
<i>Yuan Li, Bingqiao Luo, Qian Wang, Nuo Chen, Xu Liu, Bingsheng He</i>	
A Survey on In-Context Learning	1107
<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, Zhifang Sui</i>	
DocHieNet: A Large and Diverse Dataset for Document Hierarchy Parsing	1129
<i>Hangdi Xing, Changxu Cheng, Feiyu Gao, Zirui Shao, Zhi Yu, Jiajun Bu, Qi Zheng, Cong Yao</i>	
AMR-Evol: Adaptive Modular Response Evolution Elicits Better Knowledge Distillation for Large Language Models in Code Generation	1143
<i>Ziyang Luo, Xin Li, Hongzhan Lin, Jing Ma, Lidong Bing</i>	
EFUF: Efficient Fine-Grained Unlearning Framework for Mitigating Hallucinations in Multimodal Large Language Models	1167
<i>Shangyu Xing, Fei Zhao, Zhen Wu, Tuo An, Weihao Chen, Chunhui Li, Jianbing Zhang, Xinyu Dai</i>	
Rethinking Pruning Large Language Models: Benefits and Pitfalls of Reconstruction Error Minimization	1182
<i>Sungbin Shin, Wonpyo Park, Jaeho Lee, Namhoon Lee</i>	
LLMs Are Zero-Shot Context-Aware Simultaneous Translators.....	1192
<i>Roman Koshkin, Katsuhito Sudoh, Satoshi Nakamura</i>	
AgentReview: Exploring Peer Review Dynamics with LLM Agents	1208
<i>Yiqiao Jin, Qinlin Zhao, Yiyang Wang, Hao Chen, Kaijie Zhu, Yijia Xiao, Jindong Wang</i>	
ChatRetriever: Adapting Large Language Models for Generalized and Robust Conversational Dense Retrieval	1227
<i>Kelong Mao, Chenlong Deng, Haonan Chen, Fengran Mo, Zheng Liu, Tetsuya Sakai, Zhicheng Dou</i>	
Fairer Preferences Elicit Improved Human-Aligned Large Language Model Judgments.....	1241
<i>Han Zhou, Xingchen Wan, Yinhong Liu, Nigel Collier, Ivan Vulic, Anna Korhonen</i>	
Learning Interpretable Legal Case Retrieval Via Knowledge-Guided Case Reformulation	1253
<i>Chenlong Deng, Kelong Mao, Zhicheng Dou</i>	
Effective Demonstration Annotation for In-Context Learning Via Language Model-Based Determinantal Point Process.....	1266
<i>Peng Wang, Xiaobin Wang, Chao Lou, Shengyu Mao, Pengjun Xie, Yong Jiang</i>	

Pre-Trained Language Models Do Not Help Auto-Regressive Text-To-Image Generation	1281
<i>Yuhui Zhang, Brandon McKinzie, Zhe Gan, Vaishaal Shankar, Alexander T Toshev</i>	
QUSELECT: Selective Decoding for Questions Under Discussion Parsing	1288
<i>Ashima Suvarna, Xiao Liu, Tanmay Parekh, Kai-Wei Chang, Nanyun Peng</i>	
Mitigating Language Bias of LMMs in Social Intelligence Understanding with Virtual Counterfactual Calibration	1300
<i>Peng Chen, Xiao-Yu Guo, Yuan-Fang Li, Xiaowang Zhang, Zhiyong Feng</i>	
Model Balancing Helps Low-Data Training and Fine-Tuning	1311
<i>Zihang Liu, Yuanzhe Hu, Tianyu Pang, Yefan Zhou, Pu Ren, Yaoqing Yang</i>	
Reuse Your Rewards: Reward Model Transfer for Zero-Shot Cross-Lingual Alignment	1332
<i>Zhaofeng Wu, Ananth Balashankar, Yoon Kim, Jacob Eisenstein, Ahmad Beirami</i>	

VOLUME 3

Large Language Models as Foundations for Next-Gen Dense Retrieval: A Comprehensive Empirical Assessment.....	1354
<i>Kun Luo, Minghao Qin, Zheng Liu, Shitao Xiao, Jun Zhao, Kang Liu</i>	
A New Pipeline for Knowledge Graph Reasoning Enhanced by Large Language Models Without Fine-Tuning	1366
<i>Zhongwu Chen, Long Bai, Zixuan Li, Zhen Huang, Xiaolong Jin, Yong Dou</i>	
Towards Tool Use Alignment of Large Language Models	1382
<i>Zhi-Yuan Chen, Shiqi Shen, Guangyao Shen, Gong Zhi, Xu Chen, Yankai Lin</i>	
DecorateLM: Data Engineering Through Corpus Rating, Tagging, and Editing with Language Models.....	1401
<i>Ranchi Zhao, Zhen Leng Thai, Yifan Zhang, Shengding Hu, Jie Zhou, Yunqi Ba, Jie Cai, Zhiyuan Liu, Maosong Sun</i>	
Lookback Lens: Detecting and Mitigating Contextual Hallucinations in Large Language Models Using Only Attention Maps.....	1419
<i>Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim, James R. Glass</i>	
Controllable Preference Optimization: Toward Controllable Multi-Objective Alignment.....	1437
<i>Yiju Guo, Ganqu Cui, Lifan Yuan, Ning Ding, Zexu Sun, Bowen Sun, Huimin Chen, Ruobing Xie, Jie Zhou, Yankai Lin, Zhiyuan Liu, Maosong Sun</i>	
Mitigating Matthew Effect: Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning for Conversational Recommendation.....	1455
<i>Yongsen Zheng, Ruilin Xu, Guohua Wang, Liang Lin, Kwok-Yan Lam</i>	
Advancing Event Causality Identification Via Heuristic Semantic Dependency Inquiry Network.....	1467
<i>Haoran Li, Qiang Gao, Hongmei Wu, Li Huang</i>	
Exploring Union and Intersection of Visual Regions for Generating Questions, Answers, and Distractors	1479
<i>Wenjian Ding, Yao Zhang, Jun Wang, Adam Jatowt, Zhenglu Yang</i>	
UniFashion: A Unified Vision-Language Model for Multimodal Fashion Retrieval and Generation	1490
<i>Xiangyu Zhao, Yuehan Zhang, Wenlong Zhang, Xiao-Ming Wu</i>	

Tracking the Perspectives of Interacting Language Models	1508
<i>Hayden Helm, Brandon Duderstadt, Youngser Park, Carey Priebe</i>	
MAR: Matching-Augmented Reasoning for Enhancing Visual-Based Entity Question Answering	1520
<i>Zhengxuan Zhang, Yin Wu, Yuyu Luo, Nan Tang</i>	
Can Large Language Models Always Solve Easy Problems if They Can Solve Harder Ones?	1531
<i>Zhe Yang, Yichang Zhang, Tianyu Liu, Jian Yang, Junyang Lin, Chang Zhou, Zhifang Sui</i>	
Watch Every Step! LLM Agent Learning Via Iterative Step-Level Process Refinement	1556
<i>Weimin Xiong, Yifan Song, Xiutian Zhao, Wenhao Wu, Xun Wang, Ke Wang, Cheng Li, Wei Peng, Sujian Li</i>	
Standardize: Aligning Language Models with Expert-Defined Standards for Content Generation	1573
<i>Joseph Marvin Imperial, Gail Forey, Harish Tayyar Madabushi</i>	
Cross-Domain NER with Generated Task-Oriented Knowledge: An Empirical Study from Information Density Perspective	1595
<i>Zhihao Zhang, Sophia Yat Mei Lee, Junshuang Wu, Dong Zhang, Shoushan Li, Erik Cambria, Guodong Zhou</i>	
Glue Pizza and Eat Rocks - Exploiting Vulnerabilities in Retrieval-Augmented Generative Models.....	1610
<i>Zhen Tan, Chengshuai Zhao, Raha Moraffah, Yifan Li, Song Wang, Jundong Li, Tianlong Chen, Huan Liu</i>	
Predicate Debiasing in Vision-Language Models Integration for Scene Graph Generation Enhancement	1627
<i>Yuxuan Wang, Xiaoyuan Liu</i>	
SHIELD: Evaluation and Defense Strategies for Copyright Compliance in LLM Text Generation	1640
<i>Xiaoze Liu, Ting Sun, Tianyang Xu, Feijie Wu, Cunxiang Wang, Xiaoqian Wang, Jing Gao</i>	
MatchTime: Towards Automatic Soccer Game Commentary Generation.....	1671
<i>Jiayuan Rao, Haoning Wu, Chang Liu, Yanfeng Wang, Weidi Xie</i>	
Rethinking Token Reduction for State Space Models	1686
<i>Zheng Zhan, Yushu Wu, Zhenglun Kong, Changdi Yang, Yifan Gong, Xuan Shen, Xue Lin, Pu Zhao, Yanzhi Wang</i>	
Triad: A Framework Leveraging a Multi-Role LLM-Based Agent to Solve Knowledge Base Question Answering	1698
<i>Chang Zong, Yuchen Yan, Weiming Lu, Jian Shao, Yongfeng Huang, Heng Chang, Yueting Zhuang</i>	
MetaGPT: Merging Large Language Models Using Model Exclusive Task Arithmetic	1711
<i>Yuyan Zhou, Liang Song, Bingning Wang, Weipeng Chen</i>	
Event Causality Identification with Synthetic Control	1725
<i>Haoyu Wang, Fengze Liu, Jiayao Zhang, Dan Roth, Kyle Richardson</i>	
Retrieved Sequence Augmentation for Protein Representation Learning.....	1738
<i>Chang Ma, Haiteng Zhao, Lin Zheng, Jiayi Xin, Qintong Li, Lijun Wu, Zhihong Deng, Yang Young Lu, Qi Liu, Sheng Wang, Lingpeng Kong</i>	
HELPD: Mitigating Hallucination of LVLMS by Hierarchical Feedback Learning with Vision-Enhanced Penalty Decoding	1768
<i>Fan Yuan, Chi Qin, Xiaogang Xu, Piji Li</i>	

TopViewRS: Vision-Language Models as Top-View Spatial Reasoners.....	1786
<i>Chengzu Li, Caiqi Zhang, Han Zhou, Nigel Collier, Anna Korhonen, Ivan Vulic</i>	
DA3: A Distribution-Aware Adversarial Attack Against Language Models.....	1808
<i>Yibo Wang, Xiangjue Dong, James Caverlee, Philip S. Yu</i>	
Evaluating Psychological Safety of Large Language Models	1826
<i>Xingxuan Li, Yutong Li, Lin Qiu, Shafiq Joty, Lidong Bing</i>	
An Effective Deployment of Diffusion LM for Data Augmentation in Low-Resource Sentiment Classification	1844
<i>Zhuowei Chen, Lianxi Wang, Yuben Wu, Xinfeng Liao, Yujia Tian, Junyang Zhong</i>	
Self-Bootstrapped Visual-Language Model for Knowledge Selection and Question Answering.....	1857
<i>Dongze Hao, Qunbo Wang, Longteng Guo, Jie Jiang, Jing Liu</i>	
PsFuture: A Pseudo-Future-Based Zero-Shot Adaptive Policy for Simultaneous Machine Translation.....	1869
<i>Libo Zhao, Jing Li, Ziqian Zeng</i>	
TinyChart: Efficient Chart Understanding with Program-Of-Thoughts Learning and Visual Token Merging	1882
<i>Liang Zhang, Anwen Hu, Haiyang Xu, Ming Yan, Yichen Xu, Qin Jin, Ji Zhang, Fei Huang</i>	
Do We Need Language-Specific Fact-Checking Models? the Case of Chinese.....	1899
<i>Caiqi Zhang, Zhijiang Guo, Andreas Vlachos</i>	
Enhancing Advanced Visual Reasoning Ability of Large Language Models	1915
<i>Zhiyuan Li, Dongnan Liu, Chaoyi Zhang, Heng Wang, Tengfei Xue, Weidong Cai</i>	
CMD: A Framework for Context-Aware Model self-Detoxification	1930
<i>Zecheng Tang, Keyan Zhou, Juntao Li, Yuyang Ding, Pinzheng Wang, Yan Bowen, Renjie Hua, Min Zhang</i>	
Embedding and Gradient Say Wrong: A White-Box Method for Hallucination Detection	1950
<i>Xiaomeng Hu, Yiming Zhang, Ru Peng, Haozhe Zhang, Chenwei Wu, Gang Chen, Junbo Zhao</i>	
TCSinger: Zero-Shot Singing Voice Synthesis with Style Transfer and Multi-Level Style Control	1960
<i>Yu Zhang, Ziyue Jiang, Ruiqi Li, Changhao Pan, Jinzheng He, Rongjie Huang, Chuxin Wang, Zhou Zhao</i>	
Be Helpful but Don't Talk Too Much - Enhancing Helpfulness in Conversations Through Relevance in Multi-Turn Emotional Support	1976
<i>Junlin Li, Bo Peng, Yu-Yin Hsu, Chu-Ren Huang</i>	
Aligning Language Models to Explicitly Handle Ambiguity	1989
<i>Hyuhng Joon Kim, Youna Kim, Cheonbok Park, Junyeob Kim, Choonghyun Park, Kang Min Yoo, Sang-Goo Lee, Taek Kim</i>	
Tag-Grounded Visual Instruction Tuning with Retrieval Augmentation	2008
<i>Daiqing Qi, Handong Zhao, Zijun Wei, Sheng Li</i>	
GLaPE: Gold Label-Agnostic Prompt Evaluation for Large Language Models	2027
<i>Xuanchang Zhang, Zhuosheng Zhang, Hai Zhao</i>	
Decoding the Echoes of Vision from fMRI: Memory Disentangling for Past Semantic Information	2040
<i>Runze Xia, Congchi Yin, Piji Li</i>	

Optimizing Code Retrieval: High-Quality and Scalable Dataset Annotation Through Large Language Models	2053
<i>Rui Li, Qi Liu, Liyang He, Zheng Zhang, Hao Zhang, Shengyu Ye, Junyu Lu, Zhenya Huang</i>	
Towards Difficulty-Agnostic Efficient Transfer Learning for Vision-Language Models	2066
<i>Yongjin Yang, Jongwoo Ko, Se-Young Yun</i>	
Advancing Process Verification for Large Language Models Via Tree-Based Preference Learning.....	2086
<i>Mingqian He, Yongliang Shen, Wenqi Zhang, Zeqi Tan, Weiming Lu</i>	
An Inversion Attack Against Obfuscated Embedding Matrix in Language Model Inference	2100
<i>Yu Lin, Qizhi Zhang, Quanwei Cai, Jue Hong, Wu Ye, Huiqi Liu, Bing Duan</i>	
VideoScore: Building Automatic Metrics to Simulate Fine-Grained Human Feedback for Video Generation	2105
<i>Xuan He, Dongfu Jiang, Ge Zhang, Max Ku, Achint Soni, Sherman Siu, Haonan Chen, Abhramil Chandra, Ziyang Jiang, Aaran Arulraj, Kai Wang, Quy Duc Do, Yuansheng Ni, Bohan Lyu, Yaswanth Narsupalli, Rongqi Fan, Zhiheng Lyu, Bill Yuchen Lin, Wenhui Chen</i>	
LogicAsker: Evaluating and Improving the Logical Reasoning Ability of Large Language Models.....	2124
<i>Yuxuan Wan, Wenxuan Wang, Yiliu Yang, Youliang Yuan, Jen-Tse Huang, Pinjia He, Wenxiang Jiao, Michael Lyu</i>	

VOLUME 4

Integrating Structural Semantic Knowledge for Enhanced Information Extraction Pre-Training.....	2156
<i>Xiaoyang Yi, Yuru Bao, Jian Zhang, Yifang Qin, Faxin Lin</i>	
FuseGen: PLM Fusion for Data-Generation Based Zero-Shot Learning	2172
<i>Tianyuan Zou, Yang Liu, Peng Li, Jianqing Zhang, Jingjing Liu, Ya-Qin Zhang</i>	
I Need Help! Evaluating LLM’s Ability to Ask for Users’ Support: A Case Study on Text-To-SQL Generation	2191
<i>Cheng-Kuang Wu, Zhi Rui Tam, Chao-Chung Wu, Chieh-Yen Lin, Hung-Yi Lee, Yun-Nung Chen</i>	
Oddballs and Misfits: Detecting Implicit Abuse in Which Identity Groups Are Depicted as Deviating from the Norm	2200
<i>Michael Wiegand, Josef Ruppenhofer</i>	
By My Eyes: Grounding Multimodal Large Language Models with Sensor Data Via Visual Prompting.....	2219
<i>Hyungjun Yoon, Biniyam Aschalew Tolera, Taesik Gong, Kimin Lee, Sung-Ju Lee</i>	
Prefixing Attention Sinks Can Mitigate Activation Outliers for Large Language Model Quantization	2242
<i>Seungwoo Son, Wonpyo Park, Woohyun Han, Kyuyeun Kim, Jaeho Lee</i>	
CHIQ: Contextual History Enhancement for Improving Query Rewriting in Conversational Search	2253
<i>Fengran Mo, Abbas Ghaddar, Kelong Mao, Mehdi Rezagholizadeh, Boxing Chen, Qun Liu, Jian-Yun Nie</i>	
Towards Low-Resource Harmful Meme Detection with LMM Agents	2269
<i>Jianzhao Huang, Hongzhan Lin, Liu Ziyang, Ziyang Luo, Guang Chen, Jing Ma</i>	

VIVA: A Benchmark for Vision-Grounded Decision-Making with Human Values.....	2294
<i>Zhe Hu, Yixiao Ren, Jing Li, Yu Yin</i>	
Direct Multi-Turn Preference Optimization for Language Agents	2312
<i>Wentao Shi, Mengqi Yuan, Junkang Wu, Qifan Wang, Fuli Feng</i>	
Self-Refine Instruction-Tuning for Aligning Reasoning in Language Models.....	2325
<i>Leonardo Ranaldi, Andre Freitas</i>	
In Search of the Long-Tail: Systematic Generation of Long-Tail Inferential Knowledge Via Logical Rule Guided Search.....	2348
<i>Huihan Li, Yuting Ning, Zeyi Liao, Siyuan Wang, Xiang Lorraine Li, Ximing Lu, Wenting Zhao, Faeze Brahma, Yejin Choi, Xiang Ren</i>	
AutoScraper: A Progressive Understanding Web Agent for Web Scraper Generation	2371
<i>Wenhao Huang, Zhouhong Gu, Chenghao Peng, Jiaqing Liang, Zhixu Li, Yanghua Xiao, Liqian Wen, Zulong Chen</i>	
Backward Lens: Projecting Language Model Gradients into the Vocabulary Space.....	2390
<i>Shahar Katz, Yonatan Belinkov, Mor Geva, Lior Wolf</i>	
Selective Vision is the Challenge for Visual Reasoning: A Benchmark for Visual Argument Understanding	2423
<i>Jiwan Chung, Sungjae Lee, Minseo Kim, Seungju Han, Ashkan Yousefpour, Jack Hessel, Youngjae Yu</i>	
Can Visual Language Models Resolve Textual Ambiguity with Visual Cues? Let Visual Puns Tell You!	2452
<i>Jiwan Chung, Seungwon Lim, Jaehyun Jeon, Seungbeen Lee, Youngjae Yu</i>	
Reusing Transferable Weight Increments for Low-Resource Style Generation	2470
<i>Chunzhen Jin, Eliot Huang, Heng Chang, Yaqi Wang, Peng Cao, Osmar Zaiane</i>	
Large Language Model as an Assignment Evaluator: Insights, Feedback, and Challenges in a 1000+ Student Course	2489
<i>Cheng-Han Chiang, Wei-Chih Chen, Chun-Yi Kuan, Chienchou Yang, Hung-Yi Lee</i>	
Seemingly Plausible Distractors in Multi-Hop Reasoning: Are Large Language Models Attentive Readers?	2514
<i>Neeladri Bhuiya, Viktor Schlegel, Stefan Winkler</i>	
Instruction Pre-Training: Language Models Are Supervised Multitask Learners	2529
<i>Daixuan Cheng, Yuxian Gu, Shaohan Huang, Junyu Bi, Minlie Huang, Furu Wei</i>	
LEMoE: Advanced Mixture of Experts Adaptor for Lifelong Model Editing of Large Language Models.....	2551
<i>Renzhi Wang, Piji Li</i>	
Collaborative Performance Prediction for Large Language Models	2576
<i>Qiyuan Zhang, Fuyuan Lyu, Xue Liu, Chen Ma</i>	
Surveying the Dead Minds: Historical-Psychological Text Analysis with Contextualized Construct Representation (CCR) for Classical Chinese.....	2597
<i>Yuqi Chen, Sixuan Li, Ying Li, Mohammad Atari</i>	
Knowledge Verification to Nip Hallucination in the Bud.....	2616
<i>Fanqi Wan, Xinting Huang, Leyang Cui, Xiaojun Quan, Wei Bi, Shuming Shi</i>	

QUITE: Quantifying Uncertainty in Natural Language Text in Bayesian Reasoning Scenarios.....	2634
<i>Timo Pierre Schrader, Lukas Lange, Simon Razniewski, Annemarie Friedrich</i>	
African Or European Swallow? Benchmarking Large Vision-Language Models for Fine-Grained Object Classification	2653
<i>Gregor Geigle, Radu Timofte, Goran Glavaš</i>	
Whispers that Shake Foundations: Analyzing and Mitigating False Premise Hallucinations in Large Language Models	2670
<i>Hongbang Yuan, Pengfei Cao, Zhuoran Jin, Yubo Chen, Daojian Zeng, Kang Liu, Jun Zhao</i>	
To Word Senses and Beyond: Inducing Concepts with Contextualized Language Models.....	2684
<i>Bastien Liétard, Pascal Denis, Mikaela Keller</i>	
ASETf: A Novel Method for Jailbreak Attack on LLMs Through Translate Suffix Embeddings	2697
<i>Hao Wang, Hao Li, Minlie Huang, Lei Sha</i>	
An Electoral Approach to Diversify LLM-Based Multi-Agent Collective Decision-Making.....	2712
<i>Xiutian Zhao, Ke Wang, Wei Peng</i>	
Does Object Grounding Really Reduce Hallucination of Large Vision-Language Models?.....	2728
<i>Gregor Geigle, Radu Timofte, Goran Glavaš</i>	
Take off the Training Wheels! Progressive In-Context Learning for Effective Alignment	2743
<i>Zhenyu Liu, Dongfang Li, Xinshuo Hu, Xinpeng Zhao, Yubin Chen, Baotian Hu, Min Zhang</i>	
MoDULA: Mixture of Domain-Specific and Universal LoRA for Multi-Task Learning.....	2758
<i>Yufei Ma, Zihan Liang, Huangyu Dai, Ben Chen, Dehong Gao, Zhuoran Ran, Wang Zihan, Linbo Jin, Wen Jiang, Guannan Zhang, Xiaoyan Cai, Libin Yang</i>	
Message Passing on Semantic-Anchor-Graphs for Fine-Grained Emotion Representation Learning and Classification	2771
<i>Pinyi Zhang, Jingyang Chen, Junchen Shen, Zijie Zhai, Ping Li, Jie Zhang, Kai Zhang</i>	
PhiloGPT: A Philology-Oriented Large Language Model for Ancient Chinese Manuscripts with Dunhuang as Case Study	2784
<i>Yuqing Zhang, Baoyi He, Yihan Chen, Hangqi Li, Han Yue, Shengyu Zhang, Huaiyong Dou, Junchi Yan, Zemin Liu, Yongquan Zhang, Fei Wu</i>	
Alignment-Enhanced Decoding: Defending Jailbreaks Via Token-Level Adaptive Refining of Probability Distributions	2802
<i>Quan Liu, Zhenhong Zhou, Longzhu He, Yi Liu, Wei Zhang, Sen Su</i>	
MiniConGTS: A Near Ultimate Minimalist Contrastive Grid Tagging Scheme for Aspect Sentiment Triplet Extraction.....	2817
<i>Qiao Sun, Liuja Yang, Minghao Ma, Nanyang Ye, Qinying Gu</i>	
Evaluating Large Language Models Via Linguistic Profiling	2835
<i>Alessio Miaschi, Felice Dell'Orletta, Giulia Venturi</i>	
With Ears to See and Eyes to Hear: Sound Symbolism Experiments with Multimodal Large Language Models	2849
<i>Tyler Loakman, Yucheng Li, Chenghua Lin</i>	
KB-Plugin: A Plug-And-Play Framework for Large Language Models to Induce Programs Over Low-Resourced Knowledge Bases.....	2868
<i>Jiajie Zhang, Shulin Cao, Linmei Hu, Ling Feng, Lei Hou, Juanzi Li</i>	

Understanding Higher-Order Correlations Among Semantic Components in Embeddings.....	2883
<i>Momose Oyama, Hiroaki Yamagiwa, Hidetoshi Shimodaira</i>	
DGLF: A Dual Graph-Based Learning Framework for Multi-Modal Sarcasm Detection.....	2900
<i>Zhihong Zhu, Kefan Shen, Zhaorun Chen, Yunyan Zhang, Yuyan Chen, Xiaoqi Jiao, Zhongwei Wan, Shaorong Xie, Wei Liu, Xian Wu, Yefeng Zheng</i>	
Evaluating D-MERIT of Partial-Annotation on Information Retrieval.....	2913
<i>Royi Rassin, Yaron Fairstein, Oren Kalinsky, Guy Kushilevitz, Nachshon Cohen, Alexander Libov, Yoav Goldberg</i>	
Verification and Refinement of Natural Language Explanations Through LLM-Symbolic Theorem Proving	2933
<i>Xin Quan, Marco Valentino, Louise A. Dennis, Andre Freitas</i>	

VOLUME 5

Calibrating the Confidence of Large Language Models by Eliciting Fidelity	2959
<i>Mozhi Zhang, Mianqiu Huang, Rundong Shi, Linsen Guo, Chong Peng, Peng Yan, Yaqian Zhou, Xipeng Qiu</i>	
The Accuracy Paradox in RLHF: When Better Reward Models Don't Yield Better Language Models.....	2980
<i>Yanjun Chen, Dawei Zhu, Yirong Sun, Xinghao Chen, Wei Zhang, Xiaoyu Shen</i>	
How Hard is This Test Set? NLI Characterization by Exploiting Training Dynamics	2990
<i>Adrian Cosma, Stefan Ruseti, Mihai Dascalu, Cornelia Caragea</i>	
Zero-Shot Cross-Lingual Transfer for Synthetic Data Generation in Grammatical Error Detection.....	3002
<i>Gaetan Lopez Latouche, Marc-André Carbonneau, Benjamin Swanson</i>	
CUTE: Measuring LLMs' Understanding of Their Tokens.....	3017
<i>Lukas Edman, Helmut Schmid, Alexander Fraser</i>	
SEER: Self-Aligned Evidence Extraction for Retrieval-Augmented Generation.....	3027
<i>Xinping Zhao, Dongfang Li, Yan Zhong, Boren Hu, Yibin Chen, Baotian Hu, Min Zhang</i>	
On the Role of Context in Reading Time Prediction	3042
<i>Andreas Opedal, Eleanor Chodroff, Ryan Cotterell, Ethan Wilcox</i>	
BC-Prover: Backward Chaining Prover for Formal Theorem Proving	3059
<i>Yuhang He, Jihai Zhang, Jianzhu Bao, Fangquan Lin, Cheng Yang, Bing Qin, Ruifeng Xu, Wotao Yin</i>	
From Insights to Actions: The Impact of Interpretability and Analysis Research on NLP	3078
<i>Marius Mosbach, Vagrant Gautam, Tomás Vergara Browne, Dietrich Klakow, Mor Geva</i>	
Autoregressive Pre-Training on Pixels and Texts.....	3106
<i>Yekun Chai, Qingyi Liu, Jingwu Xiao, Shuohuan Wang, Yu Sun, Hua Wu</i>	
On Training Data Influence of GPT Models	3126
<i>Yekun Chai, Qingyi Liu, Shuohuan Wang, Yu Sun, Qiwei Peng, Hua Wu</i>	
Understanding “Democratization” in NLP and ML Research.....	3151
<i>Arjun Subramonian, Vagrant Gautam, Dietrich Klakow, Zeerak Talat</i>	

DocKD: Knowledge Distillation from LLMs for Open-World Document Understanding Models	3167
<i>Sungnyun Kim, Haofu Liao, Srikar Appalaraju, Peng Tang, Zhuowen Tu, Ravi Kumar Satzoda, R. Manmatha, Vijay Mahadevan, Stefano Soatto</i>	
Cross-Lingual Transfer for Automatic Question Generation by Learning Interrogative Structures in Target Languages.....	3194
<i>Seonjeong Hwang, Yunsu Kim, Gary Lee</i>	
ScalingFilter: Assessing Data Quality Through Inverse Utilization of Scaling Laws	3209
<i>Ruihang Li, Yixuan Wei, Miaosen Zhang, Nenghai Yu, Han Hu, Houwen Peng</i>	
Word Alignment as Preference for Machine Translation.....	3223
<i>Qiyu Wu, Masaaki Nagata, Zhongtao Miao, Yoshimasa Tsuruoka</i>	
Improving Multi-Party Dialogue Generation Via Topic and Rhetorical Coherence.....	3240
<i>Yaxin Fan, Peifeng Li, Qiaoming Zhu</i>	
SEEKR: Selective Attention-Guided Knowledge Retention for Continual Learning of Large Language Models	3254
<i>Jinghan He, Haiyun Guo, Kuan Zhu, Zihan Zhao, Ming Tang, Jinqiao Wang</i>	
Neuron-Level Knowledge Attribution in Large Language Models	3267
<i>Zeping Yu, Sophia Ananiadou</i>	
How Do Large Language Models Learn In-Context? Query and Key Matrices of In-Context Heads Are Two Towers for Metric Learning.....	3281
<i>Zeping Yu, Sophia Ananiadou</i>	
Interpreting Arithmetic Mechanism in Large Language Models Through Comparative Neuron Analysis.....	3293
<i>Zeping Yu, Sophia Ananiadou</i>	
Pixology: Probing the Linguistic and Visual Capabilities of Pixel-Based Language Models	3307
<i>Kushal Tatariya, Vladimir Araujo, Thomas Bauwens, Miryam De Lhoneux</i>	
GoldCoin: Grounding Large Language Models in Privacy Laws Via Contextual Integrity Theory	3321
<i>Wei Fan, Haoran Li, Zheyang Deng, Weiqi Wang, Yangqiu Song</i>	
Noise, Novels, Numbers. a Framework for Detecting and Categorizing Noise in Danish and Norwegian Literature.....	3344
<i>Ali Al-Laith, Daniel Hershcovich, Jens Bjerring-Hansen, Jakob Ingemann Parby, Alexander Conroy, Timothy R Tangherlini</i>	
QUIK: Towards End-To-End 4-Bit Inference on Generative Large Language Models.....	3355
<i>Saleh Ashkboos, Iliia Markov, Elias Frantar, Tingxuan Zhong, Xincheng Wang, Jie Ren, Torsten Hoefler, Dan Alistarh</i>	
Fine-Grained Prediction of Reading Comprehension from Eye Movements	3372
<i>Omer Shubi, Yoav Meiri, Cfir Avraham Hadar, Yevgeni Berzak</i>	
EfficientRAG: Efficient Retriever for Multi-Hop Question Answering.....	3392
<i>Ziyuan Zhuang, Zhiyang Zhang, Sitao Cheng, Fangkai Yang, Jia Liu, Shujian Huang, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, Qi Zhang</i>	
Unsupervised Human Preference Learning	3412
<i>Sumuk Shashidhar, Abhinav Chinta, Vaibhav Sahai, Dilek Hakkani Tur</i>	

Is Safer Better? the Impact of Guardrails on the Argumentative Strength of LLMs in Hate Speech Countering.....	3446
<i>Helena Bonaldi, Greta Damo, Nicolás Benjamín Ocampo, Elena Cabrio, Serena Villata, Marco Guerini</i>	
Leading Whitespaces of Language Models' Subword Vocabulary Pose a Confound for Calculating Word Probabilities	3464
<i>Byung-Doh Oh, William Schuler</i>	
LLM4Decompile: Decompiling Binary Code with Large Language Models	3473
<i>Hanzhuo Tan, Qi Luo, Jing Li, Yuqun Zhang</i>	
From Bottom to Top: Extending the Potential of Parameter Efficient Fine-Tuning	3488
<i>Jihao Gu, Zelin Wang, Yibo Zhang, Ziji Zhang, Ping Gong</i>	
CoTKR: Chain-Of-Thought Enhanced Knowledge Rewriting for Complex Knowledge Graph Question Answering	3501
<i>Yike Wu, Yi Huang, Nan Hu, Yuncheng Hua, Guilin Qi, Jiaoyan Chen, Jeff Z. Pan</i>	
MTLS: Making Texts into Linguistic Symbols	3521
<i>Wenlong Fei, Xiaohua Wang, Min Hu, Qingyu Zhang, Hongbo Li</i>	
D2R: Dual-Branch Dynamic Routing Network for Multimodal Sentiment Detection.....	3536
<i>Yifan Chen, Kuntao Li, Weixing Mai, Qiaofeng Wu, Yun Xue, Fenghuan Li</i>	
A Generic Method for Fine-Grained Category Discovery in Natural Language Texts.....	3548
<i>Chang Tian, Matthew B. Blaschko, Wenpeng Yin, Mingzhe Xing, Yinliang Yue, Marie- Francine Moens</i>	
Toxicity Detection is NOT All You Need: Measuring the Gaps to Supporting Volunteer Content Moderators Through a User-Centric Method	3567
<i>Yang Trista Cao, Lovely-Frances Domingo, Sarah Gilbert, Michelle L. Mazurek, Katie Shilton, Hal Daumé III</i>	
A User-Centric Multi-Intent Benchmark for Evaluating Large Language Models.....	3588
<i>Jiayin Wang, Fengran Mo, Weizhi Ma, Peijie Sun, Min Zhang, Jian-Yun Nie</i>	
Decompose and Compare Consistency: Measuring VLMs' Answer Reliability Via Task- Decomposition Consistency Comparison	3613
<i>Qian Yang, Weixiang Yan, Aishwarya Agrawal</i>	
Learn to Refuse: Making Large Language Models More Controllable and Reliable Through Knowledge Scope Limitation and Refusal Mechanism.....	3628
<i>Lang Cao</i>	
VGBench: Evaluating Large Language Models on Vector Graphics Understanding and Generation.....	3647
<i>Bocheng Zou, Mu Cai, Jianrui Zhang, Yong Jae Lee</i>	
What Do Large Language Models Need for Machine Translation Evaluation?	3660
<i>Shenbin Qian, Archchana Sindhujan, Minnie Kabra, Diptesh Kanojia, Constantin Orasan, Tharindu Ranasinghe, Fred Blain</i>	
Performance-Guided LLM Knowledge Distillation for Efficient Text Classification at Scale	3675
<i>Flavio Di Palo, Prateek Singhi, Bilal H Fadhallah</i>	

External Knowledge-Driven Argument Mining: Leveraging Attention-Enhanced Multi-Network Models.....	3688
<i>Debela Gemechu, Chris Reed</i>	
C3PA: An Open Dataset of Expert-Annotated and Regulation-Aware Privacy Policies to Enable Scalable Regulatory Compliance Audits	3710
<i>Maaz Bin Musa, Steven M. Winston, Garrison Allen, Jacob Schiller, Kevin Moore, Sean Quick, Johnathan Melvin, Padmini Srinivasan, Mihailis E. Diamantis, Rishab Nithyanand</i>	
M2PT: Multimodal Prompt Tuning for Zero-Shot Instruction Learning.....	3723
<i>Taowen Wang, Yiyang Liu, James Chenhao Liang, Junhan Zhao, Yiming Cui, Yuning Mao, Shaoliang Nie, Jiahao Liu, Fuli Feng, Zenglin Xu, Cheng Han, Lifu Huang, Qifan Wang, Dongfang Liu</i>	

VOLUME 6

Text Grafting: Near-Distribution Weak Supervision for Minority Classes in Text Classification.....	3741
<i>Letian Peng, Yi Gu, Chengyu Dong, Zihan Wang, Jingbo Shang</i>	
Incubating Text Classifiers Following User Instruction with Nothing but LLM.....	3753
<i>Letian Peng, Zilong Wang, Jingbo Shang</i>	
PTD-SQL: Partitioning and Targeted Drilling with LLMs in Text-To-SQL	3767
<i>Ruilin Luo, Liyuan Wang, Binghui Lin, Zicheng Lin, Yujia Yang</i>	
Conditional and Modal Reasoning in Large Language Models	3800
<i>Wesley H. Holliday, Matthew Mandelkern, Cedegao E. Zhang</i>	
Advancing Large Language Model Attribution Through Self-Improving.....	3822
<i>Lei Huang, Xiaocheng Feng, Weitao Ma, Liang Zhao, Yuchun Fan, Weihong Zhong, Dongliang Xu, Qing Yang, Hongtao Liu, Bing Qin</i>	
AlignCap: Aligning Speech Emotion Captioning to Human Preferences	3837
<i>Ziqi Liang, Haoxiang Shi, Hanhui Chen</i>	
Interpretability-Based Tailored Knowledge Editing in Transformers.....	3847
<i>Yihuai Hong, Aldo Lipani</i>	
Prompt Optimization in Multi-Step Tasks (PROMST): Integrating Human Feedback and Heuristic-Based Sampling.....	3859
<i>Yongchao Chen, Jacob Arkin, Yilun Hao, Yang Zhang, Nicholas Roy, Chuchu Fan</i>	
Empowering Large Language Model for Continual Video Question Answering with Collaborative Prompting	3921
<i>Chen Cai, Zheng Wang, Jianjun Gao, Wenyang Liu, Ye Lu, Runzhong Zhang, Kim-Hui Yap</i>	
Dissecting Fine-Tuning Unlearning in Large Language Models.....	3933
<i>Yihuai Hong, Yuelin Zou, Lijie Hu, Ziqian Zeng, Di Wang, Haiqin Yang</i>	
Dancing in Chains: Reconciling Instruction Following and Faithfulness in Language Models.....	3942
<i>Zhengxuan Wu, Yuhao Zhang, Peng Qi, Yumo Xu, Rujun Han, Yian Zhang, Jifan Chen, Bonan Min, Zhiheng Huang</i>	
Where is the Signal in Tokenization Space?.....	3966
<i>Renato Geh, Honghua Zhang, Kareem Ahmed, Benjie Wang, Guy Van Den Broeck</i>	

Private Language Models Via Truncated Laplacian Mechanism.....	3980
<i>Tianhao Huang, Tao Yang, Ivan Habernal, Lijie Hu, Di Wang</i>	
Estimating Knowledge in Large Language Models Without Generating a Single Token	3994
<i>Daniela Gottesman, Mor Geva</i>	
Consistent Autoformalization for Constructing Mathematical Libraries.....	4020
<i>Lan Zhang, Xin Quan, Andre Freitas</i>	
When Context Leads but Parametric Memory Follows in Large Language Models.....	4034
<i>Yufei Tao, Adam Hiatt, Erik Haake, Antonie J. Jetter, Ameeta Agrawal</i>	
Semantic Training Signals Promote Hierarchical Syntactic Generalization in Transformers.....	4059
<i>Aditya Yedetore, Najoung Kim</i>	
When is Multilinguality a Curse? Language Modeling for 250 High- And Low-Resource Languages.....	4074
<i>Tyler A. Chang, Catherine Arnett, Zhuowen Tu, Ben Bergen</i>	
Teaching Embodied Reinforcement Learning Agents: Informativeness and Diversity of Language Use.....	4097
<i>Jiajun Xi, Yinong He, Jianing Yang, Yinpei Dai, Joyce Chai</i>	
MiTTeNS: A Dataset for Evaluating Gender Mistranslation.....	4115
<i>Kevin Robinson, Sneha Kudugunta, Romina Stella, Sunipa Dev, Jasmijn Bastings</i>	
Teaching LLMs to Abstain Across Languages Via Multilingual Feedback.....	4125
<i>Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Orevaoghene Ahia, Shuyue Stella Li, Vidhisha Balachandran, Sunayana Sitaram, Yulia Tsvetkov</i>	
Modular Pluralism: Pluralistic Alignment Via Multi-LLM Collaboration	4151
<i>Shangbin Feng, Taylor Sorensen, Yuhan Liu, Jillian Fisher, Chan Young Park, Yejin Choi, Yulia Tsvetkov</i>	
StyleRemix: Interpretable Authorship Obfuscation Via Distillation and Perturbation of Style Elements	4172
<i>Jillian Fisher, Skyler Hallinan, Ximing Lu, Mitchell L Gordon, Zaid Harchaoui, Yejin Choi</i>	
I Could've Asked That: Reformulating Unanswerable Questions.....	4207
<i>Wenting Zhao, Ge Gao, Claire Cardie, Alexander M Rush</i>	
STOP! Benchmarking Large Language Models with Sensitivity Testing on Offensive Progressions	4221
<i>Robert Morabito, Sangmitra Madhusudan, Tyler McDonald, Ali Emami</i>	
Hidden Persuaders: LLMs' Political Leaning and Their Influence on Voters.....	4244
<i>Yujin Potter, Shiyang Lai, Junsol Kim, James Evans, Dawn Song</i>	
SOUL: Unlocking the Power of Second-Order Optimization for LLM Unlearning	4276
<i>Jinghan Jia, Yihua Zhang, Yimeng Zhang, Jiancheng Liu, Bharat Runwal, James Diffenderfer, Bhavya Kailkhura, Sijia Liu</i>	
When Reasoning Meets Information Aggregation: A Case Study with Sports Narratives	4293
<i>Yebowen Hu, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Wenlin Yao, Hassan Foroosh, Dong Yu, Fei Liu</i>	
An Analysis of Multilingual FActScore	4309
<i>Vu Trong Kim, Michael Krumbick, Varshini Reddy, Franck Dernoncourt, Viet Dac Lai</i>	

Prometheus 2: An Open Source Language Model Specialized in Evaluating Other Language Models.....	4334
<i>Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee, Minjoon Seo</i>	
RAG-QA Arena: Evaluating Domain Robustness for Long-Form Retrieval Augmented Question Answering	4354
<i>Rujun Han, Yuhao Zhang, Peng Qi, Yumo Xu, Jenyuan Wang, Lan Liu, William Yang Wang, Bonan Min, Vittorio Castelli</i>	
PromptReps: Prompting Large Language Models to Generate Dense and Sparse Representations for Zero-Shot Document Retrieval	4375
<i>Shengyao Zhuang, Xueguang Ma, Bevan Koopman, Jimmy Lin, Guido Zuccon</i>	
Voices Unheard: NLP Resources and Models for Yorùbá Regional Dialects	4392
<i>Orevaoghene Ahia, Anuoluwapo Aremu, Diana Abagyan, Hila Gonen, David Ifeoluwa Adelani, Daud Abolade, Noah A. Smith, Yulia Tsvetkov</i>	
ARES: Alternating Reinforcement Learning and Supervised Fine-Tuning for Enhanced Multi-Modal Chain-Of-Thought Reasoning Through Diverse AI Feedback	4410
<i>Ju-Seung Byun, Jiyun Chun, Jihyung Kil, Andrew Perrault</i>	
Order of Magnitude Speedups for LLM Membership Inference	4431
<i>Rongting Zhang, Martin Andres Bertran, Aaron Roth</i>	
VIMI: Grounding Video Generation Through Multi-Modal Instruction	4444
<i>Yuwei Fang, Willi Menapace, Aliaksandr Siarohin, Tsai-Shien Chen, Kuan-Chieh Wang, Ivan Skorokhodov, Graham Neubig, Sergey Tulyakov</i>	
F2RL: Factuality and Faithfulness Reinforcement Learning Framework for Claim-Guided Evidence-Supported Counterspeech Generation	4457
<i>Haiyang Wang, Yuchen Pan, Xin Song, Xuechen Zhao, Minghao Hu, Bin Zhou</i>	
Deciphering Rumors: A Multi-Task Learning Approach with Intent-Aware Hierarchical Contrastive Learning	4471
<i>Chang Yang, Peng Zhang, Hui Gao, Jing Zhang</i>	
Visual Prompting in LLMs for Enhancing Emotion Recognition	4484
<i>Qixuan Zhang, Zhifeng Wang, Dylan Zhang, Wenjia Niu, Sabrina Caldwell, Tom Gedeon, Yang Liu, Zhenyue Qin</i>	
IDEAW: Robust Neural Audio Watermarking with Invertible Dual-Embedding.....	4500
<i>Pengcheng Li, Xulong Zhang, Jing Xiao, Jianzong Wang</i>	
Leveraging Conflicts in Social Media Posts: Unintended Offense Dataset.....	4512
<i>Che Wei Tsai, Yen-Hao Huang, Tsu-Keng Liao, Didier Fernando Salazar Estrada, Retmani Latifah, Yi-Shin Chen</i>	
Outcome-Constrained Large Language Models for Countering Hate Speech.....	4523
<i>Lingzi Hong, Pengcheng Luo, Eduardo Blanco, Xiaoying Song</i>	
VOLUME 7	
Multiple Sources Are Better than One: Incorporating External Knowledge in Low-Resource Glossing.....	4537
<i>Changbing Yang, Garrett Nicolai, Miikka Silfverberg</i>	

Adaptive Immune-Based Sound-Shape Code Substitution for Adversarial Chinese Text Attacks	4553
<i>Ao Wang, Xinghao Yang, Chen Li, Bao-Di Liu, Weifeng Liu</i>	
Bootstrapped Policy Learning for Task-Oriented Dialogue Through Goal Shaping	4566
<i>Yangyang Zhao, Ben Niu, Mehdi Dastani, Shihan Wang</i>	
PsyGUARD: An Automated System for Suicide Detection and Risk Assessment in Psychological Counseling.....	4581
<i>Huachuan Qiu, Lizhi Ma, Zhenzhong Lan</i>	
World to Code: Multi-Modal Data Generation Via Self-Instructed Compositional Captioning and Filtering	4608
<i>Jiacong Wang, Bohong Wu, Haiyong Jiang, Zhou Xun, Xin Xiao, Haoyuan Guo, Jun Xiao</i>	
DVD: Dynamic Contrastive Decoding for Knowledge Amplification in Multi-Document Question Answering	4624
<i>Jing Jin, Houfeng Wang, Hao Zhang, Xiaoguang Li, Zhijiang Guo</i>	
How Do Humans Write Code? Large Models Do it the Same Way Too	4638
<i>Long Li, Xuzheng He, Haozhe Wang, Linlin Wang, Liang He</i>	
Retrospec: Language Agent Meets Offline Reinforcement Learning Critic.....	4650
<i>Yufei Xiang, Yiqun Shen, Yeqin Zhang, Nguyen Cam-Tu</i>	
Forgetting Curve: A Reliable Method for Evaluating Memorization Capability for Long-Context Models.....	4667
<i>Xinyu Liu, Runsong Zhao, Pengcheng Huang, Chunyang Xiao, Bei Li, Jingang Wang, Tong Xiao, Jingbo Zhu</i>	
Retrieve-Plan-Generation: An Iterative Planning and Answering Framework for Knowledge-Intensive LLM Generation	4683
<i>Yuanjie Lyu, Zihan Niu, Zheyong Xie, Chao Zhang, Tong Xu, Yang Wang, Enhong Chen</i>	
CoEvol: Constructing Better Responses for Instruction Finetuning Through Multi-Agent Cooperation	4703
<i>Renhao Li, Minghuan Tan, Derek F. Wong, Min Yang</i>	
A Peek into Token Bias: Large Language Models Are Not Yet Genuine Reasoners	4722
<i>Bowen Jiang, Yangxinyu Xie, Zhuoqun Hao, Xiaomeng Wang, Tanwi Mallick, Weijie J Su, Camillo Jose Taylor, Dan Roth</i>	
Bayesian Calibration of Win Rate Estimation with LLM Evaluators.....	4757
<i>Yicheng Gao, Gonghan Xu, Zhe Wang, Arman Cohan</i>	
MuMath-Code: Combining Tool-Use Large Language Models with Multi-Perspective Data Augmentation for Mathematical Reasoning	4770
<i>Shuo Yin, Weihao You, Zhilong Ji, Guoqiang Zhong, Jinfeng Bai</i>	
Seeing the Forest Through the Trees: Data Leakage from Partial Transformer Gradients	4786
<i>Weijun Li, Qiongkai Xu, Mark Dras</i>	
RWKV-CLIP: A Robust Vision-Language Representation Learner	4799
<i>Tiancheng Gu, Kaicheng Yang, Xiang An, Ziyong Feng, Dongnan Liu, Weidong Cai, Jiankang Deng</i>	
KidLM: Advancing Language Models for Children – Early Insights and Future Directions.....	4813
<i>Mir Tafseer Nayeem, Davood Rafiei</i>	

Using Language Models to Disambiguate Lexical Choices in Translation	4837
<i>Josh Barua, Sanjay Subramanian, Kayo Yin, Alane Suhr</i>	
How Does the Disclosure of AI Assistance Affect the Perceptions of Writing?	4849
<i>Zhuoyan Li, Chen Liang, Jing Peng, Ming Yin</i>	
An Unsupervised Approach to Achieve Supervised-Level Explainability in Healthcare Records	4869
<i>Joakim Edin, Maria Maistro, Lars Maaløe, Lasse Borgholt, Jakob Drachmann Havtorn, Tuukka Ruotsalo</i>	
Crafting Personalized Agents Through Retrieval-Augmented Generation on Editable Memory Graphs	4891
<i>Zheng Wang, Zhongyang Li, Zeren Jiang, Dandan Tu, Wei Shi</i>	
EVEDIT: Event-Based Knowledge Editing for Deterministic Knowledge Propagation.....	4907
<i>Jiateng Liu, Pengfei Yu, Yuji Zhang, Sha Li, Zixuan Zhang, Ruhi Sarikaya, Kevin Small, Heng Ji</i>	
Modeling Nonnative Sentence Processing with L2 Language Models	4927
<i>Tatsuya Aoyama, Nathan Schneider</i>	
From the Least to the Most: Building a Plug-And-Play Visual Reasoner Via Data Synthesis	4941
<i>Chuanqi Cheng, Jian Guan, Wei Wu, Rui Yan</i>	
Quality Matters: Evaluating Synthetic Data for Tool-Using LLMs.....	4958
<i>Shadi Iskander, Sofia Tolmach, Ori Shapira, Nachshon Cohen, Zohar Karnin</i>	
Cross-Domain Audio Deepfake Detection: Dataset and Analysis	4977
<i>Yuang Li, Min Zhang, Mengxin Ren, Xiaosong Qiao, Miaomiao Ma, Daimeng Wei, Hao Yang</i>	
MaPPER: Multimodal Prior-Guided Parameter Efficient Tuning for Referring Expression Comprehension.....	4984
<i>Ting Liu, Zunnan Xu, Yue Hu, Liangtao Shi, Zhiqiang Wang, Quanjun Yin</i>	
Hierarchical Deconstruction of LLM Reasoning: A Graph-Based Framework for Analyzing Knowledge Utilization.....	4995
<i>Miyoung Ko, Sue Hyun Park, Joonsuk Park, Minjoon Seo</i>	
Aligning Translation-Specific Understanding to General Understanding in Large Language Models	5028
<i>Yichong Huang, Baohang Li, Xiaocheng Feng, Wenshuai Huo, Chengpeng Fu, Ting Liu, Bing Qin</i>	
FOOL ME if YOU CAN! an Adversarial Dataset to Investigate the Robustness of LMs in Word Sense Disambiguation	5042
<i>Mohamad Ballout, Anne Dedert, Nohayr Muhammad Abdelmoneim, Ulf Krumnack, Gunther Heidemann, Kai-Uwe Kühnberger</i>	
Concept-Skill Transferability-Based Data Selection for Large Vision-Language Models	5060
<i>Jaewoo Lee, Boyang Li, Sung Ju Hwang</i>	
LLMs Assist NLP Researchers: Critique Paper (Meta-)Reviewing	5081
<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, Wenpeng Yin</i>	

Academics Can Contribute to Domain-Specialized Language Models.....	5100
<i>Mark Dredze, Genta Indra Winata, Prabhanjan Kambadur, Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, David S Rosenberg, Sebastian Gehrmann</i>	
Beyond Reference: Evaluating High Quality Translations Better than Human References	5111
<i>Keonwoong Noh, Seokjin Oh, Woohwan Jung</i>	
Unveiling the Lexical Sensitivity of LLMs: Combinatorial Optimization for Prompt Enhancement	5128
<i>Pengwei Zhan, Zhen Xu, Qian Tan, Jie Song, Ru Xie</i>	
SEACrowd: A Multilingual Multimodal Data Hub and Benchmark Suite for Southeast Asian Languages.....	5155
<i>Holy Lovenia, Rahmad Mahendra, Salsabil Maulana Akbar, Lester James Validad Miranda, Jennifer Santoso, Elyanah Aco, Akhdan Fadhillah, Jonibek Mansurov, Joseph Marvin Imperial, Onno P. Kampman, Joel Ruben Antony Moniz, Muhammad Ravi Shulthan Habibi, Frederikus Hudi, Jann Railey 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, Samuel Cahyawijaya</i>	
Induct-Learn: Short Phrase Prompting with Instruction Induction.....	5204
<i>Po-Chun Chen, Sheng-Lun Wei, Hen-Hsen Huang, Hsin-Hsi Chen</i>	
Multi-Granularity History and Entity Similarity Learning for Temporal Knowledge Graph Reasoning	5232
<i>Shi Mingcong, Chunjiang Zhu, Detian Zhang, Shiting Wen, Li Qing</i>	
LUQ: Long-Text Uncertainty Quantification for LLMs.....	5244
<i>Caiqi Zhang, Fangyu Liu, Marco Basaldella, Nigel Collier</i>	
Pretraining Data Detection for Large Language Models: A Divergence-Based Calibration Method	5263
<i>Weichao Zhang, Ruqing Zhang, Jiafeng Guo, Maarten De Rijke, Yixing Fan, Xueqi Cheng</i>	
Scaling Synthetic Logical Reasoning Datasets with Context-Sensitive Declarative Grammars	5275
<i>Damien Sileo</i>	
Improving Spoken Language Modeling with Phoneme Classification: A Simple Fine-Tuning Approach	5284
<i>Maxime Poli, Emmanuel Chemla, Emmanuel Dupoux</i>	
Safely Learning with Private Data: A Federated Learning Framework for Large Language Model	5293
<i>Jia-Ying Zheng, Hainan Zhang, Lingxiang Wang, Wangjie Qiu, Hong-Wei Zheng, Zhi-Ming Zheng</i>	
Formality is Favored: Unraveling the Learning Preferences of Large Language Models on Data with Conflicting Knowledge	5307
<i>Jiahuan Li, Yiqing Cao, Shujian Huang, Jiajun Chen</i>	
How Does the Textual Information Affect the Retrieval of Multimodal In-Context Learning?	5321
<i>Yang Luo, Zangwei Zheng, Zirui Zhu, Yang You</i>	

VOLUME 8

How Far Can We Extract Diverse Perspectives from Large Language Models?.....	5336
<i>Shirley Anugrah Hayati, Minhwa Lee, Dheeraj Rajagopal, Dongyeop Kang</i>	
EXPLORA: Efficient Exemplar Subset Selection for Complex Reasoning.....	5367
<i>Kiran Purohit, Venkatesh V, Raghuram Devalla, Krishna Mohan Yerragorla, Sourangshu Bhattacharya, Avishek Anand</i>	
An LLM Feature-Based Framework for Dialogue Constructiveness Assessment.....	5389
<i>Lexin Zhou, Youmna Farag, Andreas Vlachos</i>	
Relevance is a Guiding Light: Relevance-Aware Adaptive Learning for End-To-End Task-Oriented Dialogue System.....	5410
<i>Zhanpeng Chen, Zhihong Zhu, Wanshi Xu, Xianwei Zhuang, Yuexian Zou</i>	
Dialog2Flow: Pre-Training Soft-Contrastive Action-Driven Sentence Embeddings for Automatic Dialog Flow Extraction	5421
<i>Sergio Burdisso, Srikanth Madikeri, Petr Motlicek</i>	
Words Worth a Thousand Pictures: Measuring and Understanding Perceptual Variability in Text-To-Image Generation	5441
<i>Raphael Tang, Crystina Zhang, Lixinyu Xu, Yao Lu, Wenyan Li, Pontus Stenetorp, Jimmy Lin, Ferhan Ture</i>	
Investigating LLMs as Voting Assistants Via Contextual Augmentation: A Case Study on the European Parliament Elections 2024.....	5455
<i>Ilias Chalkidis</i>	
Adaption-Of-Thought: Learning Question Difficulty Improves Large Language Models for Reasoning	5468
<i>Mayi Xu, Yongqi Li, Ke Sun, Tiejun Qian</i>	
LogicST: A Logical Self-Training Framework for Document-Level Relation Extraction with Incomplete Annotations.....	5496
<i>Shengda Fan, Yanting Wang, Shasha Mo, Jianwei Niu</i>	
Concept Space Alignment in Multilingual LLMs.....	5511
<i>Qiwei Peng, Anders Søgaard</i>	
Predicting Rewards Alongside Tokens: Non-Disruptive Parameter Insertion for Efficient Inference Intervention in Large Language Model	5527
<i>Chenhan Yuan, Fei Huang, Ru Peng, Keming Lu, Bowen Yu, Chang Zhou, Jingren Zhou</i>	
NLEBench+NorGLM: A Comprehensive Empirical Analysis and Benchmark Dataset for Generative Language Models in Norwegian	5543
<i>Peng Liu, Lemei Zhang, Terje Farup, Even W. Lauvrak, Jon Espen Ingvaldsen, Simen Eide, Jon Atle Gulla, Zhirong Yang</i>	
RSA-Control: A Pragmatics-Grounded Lightweight Controllable Text Generation Framework	5561
<i>Yifan Wang, Vera Demberg</i>	
Scaling Laws Across Model Architectures: A Comparative Analysis of Dense and MoE Models in Large Language Models	5583
<i>Siqi Wang, Zhengyu Chen, Bei Li, Keqing He, Min Zhang, Jingang Wang</i>	

Synergizing In-Context Learning with Hints for End-To-End Task-Oriented Dialog Systems.....	5596
<i>Vishal Vivek Saley, Rocktim Jyoti Das, Dinesh Raghu, Mausam .</i>	
REAR: A Relevance-Aware Retrieval-Augmented Framework for Open-Domain Question Answering	5613
<i>Yuhao Wang, Ruiyang Ren, Junyi Li, Xin Zhao, Jing Liu, Ji-Rong Wen</i>	
Leave No Document Behind: Benchmarking Long-Context LLMs with Extended Multi-Doc QA	5627
<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, Yongbin Li</i>	
On Mitigating Performance Disparities in Multilingual Speech Recognition	5647
<i>Monorama Swain, Anna Katrine Van Zee, Anders Søgaard</i>	
Thinking Outside of the Differential Privacy Box: A Case Study in Text Privatization with Language Model Prompting	5656
<i>Stephen Meisenbacher, Florian Matthes</i>	
To Preserve Or to Compress: An In-Depth Study of Connector Selection in Multimodal Large Language Models	5666
<i>Junyan Lin, Haoran Chen, Dawei Zhu, Xiaoyu Shen</i>	
What is "Typological Diversity" in NLP?	5681
<i>Esther Ploeger, Wessel Poelman, Miryam De Lhoneux, Johannes Bjerva</i>	
The Computational Anatomy of Humility: Modeling Intellectual Humility in Online Public Discourse.....	5701
<i>Xiaobo Guo, Neil Potnis, Melody Yu, Nabeel Gillani, Soroush Vosoughi</i>	
Consistent Bidirectional Language Modelling: Expressive Power and Representational Conciseness.....	5724
<i>Georgi Shopov, Stefan Gerdjikov</i>	
Benchmarking Vision Language Models for Cultural Understanding.....	5769
<i>Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd Van Steenkiste, Lisa Anne Hendricks, Karolina Stanczak, Aishwarya Agrawal</i>	
Methods of Automatic Matrix Language Determination for Code-Switched Speech	5791
<i>Olga Iakovenko, Thomas Hain</i>	
Analyzing Key Factors Influencing Emotion Prediction Performance of VLLMs in Conversational Contexts.....	5801
<i>Jaewook Lee, Yeajin Jang, Hongjin Kim, Woojin Lee, Harksoo Kim</i>	
Context-Aware Assistant Selection for Improved Inference Acceleration with Large Language Models.....	5817
<i>Jerry Huang, Prasanna Parthasarathi, Mehdi Rezagholizadeh, Sarath Chandar</i>	
Teaching Small Language Models Reasoning Through Counterfactual Distillation.....	5831
<i>Tao Feng, Yicheng Li, Li Chenglin, Hao Chen, Fei Yu, Yin Zhang</i>	
Pretraining Language Models Using Translationese	5843
<i>Meet Doshi, Raj Dabre, Pushpak Bhattacharyya</i>	
Quantifying the Gaps Between Translation and Native Perception in Training for Multimodal, Multilingual Retrieval	5863
<i>Kyle Buettner, Adriana Kovashka</i>	

MTA4DPR: Multi-Teaching-Assistants Based Iterative Knowledge Distillation for Dense Passage Retrieval	5871
<i>Qixi Lu, Endong Xun, Gongbo Tang</i>	
Fine-Grained Detection of Solidarity for Women and Migrants in 155 Years of German Parliamentary Debates.....	5884
<i>Aida Kostikova, Dominik Beese, Benjamin Paassen, Ole Pütz, Gregor Wiedemann, Steffen Eger</i>	
CItruS: Chunked Instruction-Aware State Eviction for Long Sequence Modeling	5908
<i>Yu Bai, Xiyuan Zou, Heyan Huang, Sanxing Chen, Marc-Antoine Rondeau, Yang Gao, Jackie Ck Cheung</i>	
Story Embeddings — Narrative-Focused Representations of Fictional Stories	5931
<i>Hans Ole Hatzel, Chris Biemann</i>	
C-LLM: Learn to Check Chinese Spelling Errors Character by Character	5944
<i>Kunting Li, Yong Hu, Liang He, Fandong Meng, Jie Zhou</i>	
PSC: Extending Context Window of Large Language Models Via Phase Shift Calibration	5958
<i>Wenqiao Zhu, Chao Xu, Lulu Wang, Jun Wu</i>	
Video-LLaVA: Learning United Visual Representation by Alignment Before Projection	5971
<i>Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, Li Yuan</i>	
SaySelf: Teaching LLMs to Express Confidence with Self-Reflective Rationales	5985
<i>Tianyang Xu, Shujin Wu, Shizhe Diao, Xiaoze Liu, Xingyao Wang, Yangyi Chen, Jing Gao</i>	
Mitigating Frequency Bias and Anisotropy in Language Model Pre-Training with Syntactic Smoothing	5999
<i>Richard Diehl Martinez, Zebulon Goriely, Andrew Caines, Paula Buttery, Lisa Beinborn</i>	
ToxiCloakCN: Evaluating Robustness of Offensive Language Detection in Chinese with Cloaking Perturbations.....	6012
<i>Yunze Xiao, Yujia Hu, Kenny Tsu Wei Choo, Roy Ka-Wei Lee</i>	
Boosting Scientific Concepts Understanding: Can Analogy from Teacher Models Empower Student Models?	6026
<i>Siyu Yuan, Cheng Jiayang, Lin Qiu, Deqing Yang</i>	
Model Internals-Based Answer Attribution for Trustworthy Retrieval-Augmented Generation	6037
<i>Jirui Qi, Gabriele Sarti, Raquel Fernández, Arianna Bisazza</i>	
Do Large Language Models Know How Much They Know?	6054
<i>Gabriele Prato, Jerry Huang, Prasanna Parthasarathi, Shagun Sodhani, Sarath Chandar</i>	
Investigating Mysteries of CoT-Augmented Distillation.....	6071
<i>Somin Wadhwa, Silvio Amir, Byron C Wallace</i>	
SciPrompt: Knowledge-Augmented Prompting for Fine-Grained Categorization of Scientific Topics	6087
<i>Zhiwen You, Kanyao Han, Haotian Zhu, Bertram Ludascher, Jana Diesner</i>	
Distilling Knowledge from Text-To-Image Generative Models Improves Visio-Linguistic Reasoning in CLIP	6105
<i>Samyadeep Basu, Shell Xu Hu, Maziar Sanjabi, Daniela Massiceti, Soheil Feizi</i>	

Learning from Natural Language Explanations for Generalizable Entity Matching6114
Somin Wadhwa, Adit Krishnan, Runhui Wang, Byron C Wallace, Luyang Kong

VOLUME 9

Do You Know What You Are Talking About? Characterizing Query-Knowledge Relevance for
Reliable Retrieval Augmented Generation 6130
*Zhuohang Li, Jiaxin Zhang, Chao Yan, Kamalika Das, Sricharan Kumar, Murat Kantarcioglu,
Bradley A. Malin*

On the Reliability of Psychological Scales on Large Language Models 6152
Jen-Tse Huang, Wenxiang Jiao, Man Ho Lam, Eric John Li, Wenxuan Wang, Michael Lyu

Contrastive Entity Coreference and Disambiguation for Historical Texts..... 6174
Abhishek Arora, Emily Silcock, Melissa Dell, Leander Heldring

Finer: Investigating and Enhancing Fine-Grained Visual Concept Recognition in Large Vision
Language Models 6187
Jeonghwan Kim, Heng Ji

Evaluating LLMs for Targeted Concept Simplification for Domain-Specific Texts 6208
Sumit Asthana, Hannah Rashkin, Elizabeth Clark, Fantine Huot, Mirella Lapata

VLFeedback: A Large-Scale AI Feedback Dataset for Large Vision-Language Models Alignment 6227
*Lei Li, Zhihui Xie, Mukai Li, Shunian Chen, Peiyi Wang, Liang Chen, Yazheng Yang, Benyou
Wang, Lingpeng Kong, Qi Liu*

Focused Large Language Models Are Stable Many-Shot Learners..... 6247
*Peiwen Yuan, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Yueqi Zhang, Chuyi Tan, Boyuan Pan,
Heda Wang, Yao Hu, Kan Li*

Reconsidering Sentence-Level Sign Language Translation 6262
Garrett Tanzer, Maximus Shengelia, Ken Harrenstien, David Uthus

GAMA: A Large Audio-Language Model with Advanced Audio Understanding and Complex
Reasoning Abilities..... 6288
*Sreyan Ghosh, Sonal Kumar, Ashish Seth, Chandra Kiran Reddy Evuru, Utkarsh Tyagi, S
Sakshi, Oriol Nieto, Ramani Duraiswami, Dinesh Manocha*

Verba Volant, Scripta Volant? Don't Worry! There Are Computational Solutions for Protoword
Reconstruction..... 6314
*Liviu P Dinu, Ana Sabina Uban, Alina Maria Cristea, Ioan-Bogdan Iordache, Teodor-George
Marchitan, Simona Georgescu, Laurentiu Zoicas*

ChatGPT Doesn't Trust Chargers Fans: Guardrail Sensitivity in Context..... 6327
Victoria R Li, Yida Chen, Naomi Saphra

Personas as a Way to Model Truthfulness in Language Models..... 6346
Nitish Joshi, Javier Rando, Abulhair Saparov, Najoung Kim, He He

Satyrn: A Platform for Analytics Augmented Generation..... 6360
*Marko Sterbentz, Cameron Barrie, Shubham Shahi, Abhratanu Dutta, Donna Hooshmand,
Harper Pack, Kristian J Hammond*

EH-MAM: Easy-To-Hard Masked Acoustic Modeling for Self-Supervised Speech Representation Learning	6386
<i>Ashish Seth, Ramaneswaran Selvakumar, S Sakshi, Sonal Kumar, Sreyan Ghosh, Dinesh Manocha</i>	
EPO: Hierarchical LLM Agents with Environment Preference Optimization.....	6401
<i>Qi Zhao, Haotian Fu, Chen Sun, George Konidaris</i>	
Detection and Measurement of Syntactic Templates in Generated Text.....	6416
<i>Chantal Shaib, Yanai Elazar, Junyi Jessy Li, Byron C Wallace</i>	
UOUO: Uncontextualized Uncommon Objects for Measuring Knowledge Horizons of Vision Language Models	6432
<i>Xinyu Pi, Mingyuan Wu, Jize Jiang, Haozhen Zheng, Beitong Tian, Chengxiang Zhai, Klara Nahrstedt, Zhiting Hu</i>	
Optimized Speculative Sampling for GPU Hardware Accelerators	6442
<i>Dominik Wagner, Seanie Lee, Ilya Baumann, Philipp Seeberger, Korbinian Riedhammer, Tobias Bocklet</i>	
Personalized Pieces: Efficient Personalized Large Language Models Through Collaborative Efforts	6459
<i>Zhaoxuan Tan, Zheyuan Liu, Meng Jiang</i>	
Democratizing Large Language Models Via Personalized Parameter-Efficient Fine-Tuning.....	6476
<i>Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin, Meng Jiang</i>	
Unifying Multimodal Retrieval Via Document Screenshot Embedding.....	6492
<i>Xueguang Ma, Sheng-Chieh Lin, Minghan Li, Wenhui Chen, Jimmy Lin</i>	
Neuron Specialization: Leveraging Intrinsic Task Modularity for Multilingual Machine Translation.....	6506
<i>Shaomu Tan, Di Wu, Christof Monz</i>	
An Audit on the Perspectives and Challenges of Hallucinations in NLP	6528
<i>Pranav Narayanan Venkit, Tatiana Chakravorti, Vipul Gupta, Heidi Biggs, Mukund Srinath, Koustava Goswami, Sarah Rajtmajer, Shomir Wilson</i>	
Discovering Knowledge-Critical Subnetworks in Pretrained Language Models	6549
<i>Deniz Bayazit, Negar Foroutan, Zeming Chen, Gail Weiss, Antoine Bosselut</i>	
Reconstruct Your Previous Conversations! Comprehensively Investigating Privacy Leakage Risks in Conversations with GPT Models.....	6584
<i>Junjie Chu, Zeyang Sha, Michael Backes, Yang Zhang</i>	
Right for Right Reasons: Large Language Models for Verifiable Commonsense Knowledge Graph Question Answering	6601
<i>Armin Toroghi, Willis Guo, Mohammad Mahdi Abdollah Pour, Scott Sanner</i>	
Verifiable, Debuggable, and Repairable Commonsense Logical Reasoning Via LLM-Based Theory Resolution.....	6634
<i>Armin Toroghi, Willis Guo, Ali Pesaranhader, Scott Sanner</i>	
Understanding and Mitigating Language Confusion in LLMs.....	6653
<i>Kelly Marchisio, Wei-Yin Ko, Alexandre Berard, Théo Dehaze, Sebastian Ruder</i>	
Can Large Language Models Learn Independent Causal Mechanisms?	6678
<i>Gael Gendron, Bao Trung Nguyen, Alex Yuxuan Peng, Michael Witbrock, Gillian Dobbie</i>	

MirrorStories: Reflecting Diversity Through Personalized Narrative Generation with Large Language Models	6702
<i>Sarfarozi Yunusov, Hamza Sidat, Ali Emami</i>	
InterIntent: Investigating Social Intelligence of LLMs Via Intention Understanding in an Interactive Game Context.....	6718
<i>Ziyi Liu, Abhishek Anand, Pei Zhou, Jen-Tse Huang, Jieyu Zhao</i>	
Locating Information Gaps and Narrative Inconsistencies Across Languages: A Case Study of LGBT People Portrayals on Wikipedia	6747
<i>Farhan Samir, Chan Young Park, Anjalie Field, Vered Shwartz, Yulia Tsvetkov</i>	
From Local Concepts to Universals: Evaluating the Multicultural Understanding of Vision-Language Models	6763
<i>Mehar Bhatia, Sahithya Ravi, Aditya Chinchure, Eunjeong Hwang, Vered Shwartz</i>	
Dynamic Multi-Reward Weighting for Multi-Style Controllable Generation	6783
<i>Karin De Langis, Ryan Koo, Dongyeop Kang</i>	
MMNeuron: Discovering Neuron-Level Domain-Specific Interpretation in Multimodal Large Language Model.....	6801
<i>Jiahao Huo, Yibo Yan, Boren Hu, Yutao Yue, Xuming Hu</i>	
Learning to Extract Structured Entities Using Language Models	6817
<i>Haolun Wu, Ye Yuan, Liana Mikaelyan, Alexander Meulemans, Xue Liu, James Hensman, Bhaskar Mitra</i>	
Efficient LLM Comparative Assessment: A Product of Experts Framework for Pairwise Comparisons.....	6835
<i>Adian Liusie, Vatsal Raina, Yassir Fathullah, Mark Gales</i>	
A Survey of AMR Applications.....	6856
<i>Shira Wein, Juri Opitz</i>	
Beyond Embeddings: The Promise of Visual Table in Visual Reasoning.....	6876
<i>Yiwu Zhong, Zi-Yuan Hu, Michael Lyu, Liwei Wang</i>	
CareCorpus+: Expanding and Augmenting Caregiver Strategy Data to Support Pediatric Rehabilitation	6912
<i>Shahla Farzana, Ivana Lucero, Vivian Villegas, Vera C Kaelin, Mary Khetani, Natalie Parde</i>	

VOLUME 10

Taylor Unswift: Secured Weight Release for Large Language Models Via Taylor Expansion.....	6928
<i>Guanchu Wang, Yu-Neng Chuang, Ruixiang Tang, Shaochen Zhong, Jiayi Yuan, Hongye Jin, Zirui Liu, Vipin Chaudhary, Shuai Xu, James Caverlee, Xia Hu</i>	
TimeR4 : Time-Aware Retrieval-Augmented Large Language Models for Temporal Knowledge Graph Question Answering	6942
<i>Xinying Qian, Ying Zhang, Yu Zhao, Baohang Zhou, Xuhui Sui, Li Zhang, Kehui Song</i>	
Knowledge-Centric Hallucination Detection	6953
<i>Xiangkun Hu, Dongyu Ru, Lin Qiu, Qipeng Guo, Tianhang Zhang, Yang Xu, Yun Luo, Pengfei Liu, Yue Zhang, Zheng Zhang</i>	

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

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

Rethinking the Reversal Curse of LLMs: A Prescription from Human Knowledge Reversal.....	7518
<i>Zhicong Lu, Li Jin, Peiguang Li, Yu Tian, Linhao Zhang, Sirui Wang, Guangluan Xu, Changyuan Tian, Xunliang Cai</i>	
More than Catastrophic Forgetting: Integrating General Capabilities for Domain-Specific LLMs	7531
<i>Chengyuan Liu, Yangyang Kang, Shihang Wang, Lizhi Qing, Fubang Zhao, Chao Wu, Changlong Sun, Kun Kuang, Fei Wu</i>	
Muting Whisper: A Universal Acoustic Adversarial Attack on Speech Foundation Models	7549
<i>Vyas Raina, Rao Ma, Charles McGhee, Kate Knill, Mark Gales</i>	
GENRA: Enhancing Zero-Shot Retrieval with Rank Aggregation	7566
<i>Georgios Katsimpras, Georgios Paliouras</i>	
XplainLLM: A Knowledge-Augmented Dataset for Reliable Grounded Explanations in LLMs.....	7578
<i>Zichen Chen, Jianda Chen, Ambuj Singh, Misha Sra</i>	
Divide and Conquer Radiology Report Generation Via Observation Level Fine-Grained Pretraining and Prompt Tuning	7597
<i>Yuanpin Zhou, Huogen Wang</i>	
SURf: Teaching Large Vision-Language Models to Selectively Utilize Retrieved Information	7611
<i>Jiashuo Sun, Jihai Zhang, Yucheng Zhou, Zhaochen Su, Xiaoye Qu, Yu Cheng</i>	
UNO Arena for Evaluating Sequential Decision-Making Capability of Large Language Models	7630
<i>Zhanyue Qin, Haochuan Wang, Deyuan Liu, Ziyang Song, Cunhang Fan, Zhao Lv, Jinlin Wu, Zhen Lei, Zhiying Tu, Dianhui Chu, Xiaoyan Yu, Dianbo Sui</i>	
Middleware for LLMs: Tools Are Instrumental for Language Agents in Complex Environments	7646
<i>Yu Gu, Yiheng Shu, Hao Yu, Xiao Liu, Yuxiao Dong, Jie Tang, Jayanth Srinivasa, Hugo Latapie, Yu Su</i>	
MORPHEUS: Modeling Role from Personalized Dialogue History by Exploring and Utilizing Latent Space	7664
<i>Yihong Tang, Bo Wang, Dongming Zhao, Jinxiaojia Jinxiaojia, Zhangjijun Zhangjijun, Ruifang He, Yuexian Hou</i>	
KnowledgeSG: Privacy-Preserving Synthetic Text Generation with Knowledge Distillation from Server	7677
<i>Wenhao Wang, Xiaoyu Liang, Rui Ye, Jingyi Chai, Siheng Chen, Yanfeng Wang</i>	
DAMRO: Dive into the Attention Mechanism of LVLm to Reduce Object Hallucination.....	7696
<i>Xuan Gong, Tianshi Ming, Xinpeng Wang, Zhihua Wei</i>	
Unlocking the Future: Exploring Look-Ahead Planning Mechanistic Interpretability in Large Language Models	7713
<i>Tianyi Men, Pengfei Cao, Zhuoran Jin, Yubo Chen, Kang Liu, Jun Zhao</i>	

VOLUME 11

Breaking Language Barriers: Cross-Lingual Continual Pre-Training at Scale.....	7725
<i>Wenzhen Zheng, Wenbo Pan, Xu Xu, Libo Qin, Li Yue, Ming Zhou</i>	
An Empirical Study of Multilingual Reasoning Distillation for Question Answering	7739
<i>Patomporn Payoungkhamdee, Peerat Limkonchotiwat, Jinheon Baek, Potsawee Manakul, Can Udomcharoenchaikit, Ekapol Chuangsuwanich, Sarana Nutanong</i>	

Can Large Language Models Faithfully Express Their Intrinsic Uncertainty in Words?	7752
<i>Gal Yona, Roei Aharoni, Mor Geva</i>	
Does Fine-Tuning LLMs on New Knowledge Encourage Hallucinations?	7765
<i>Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart, Jonathan Herzig</i>	
Bridging Modalities: Enhancing Cross-Modality Hate Speech Detection with Few-Shot In-Context Learning	7785
<i>Ming Shan Hee, Aditi Kumaresan, Roy Ka-Wei Lee</i>	
MIND: Multimodal Shopping Intention Distillation from Large Vision-Language Models for E-Commerce Purchase Understanding	7800
<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, Yangqiu Song</i>	
ECON: On the Detection and Resolution of Evidence Conflicts	7816
<i>Cheng Jiayang, Chunkit Chan, Qianqian Zhuang, Lin Qiu, Tianhang Zhang, Tengxiao Liu, Yangqiu Song, Yue Zhang, Pengfei Liu, Zheng Zhang</i>	
“Image, Tell Me Your Story!” Predicting the Original Meta-Context of Visual Misinformation	7845
<i>Jonathan Tonglet, Marie-Francine Moens, Iryna Gurevych</i>	
Improving Retrieval-Augmented Text-To-SQL with AST-Based Ranking and Schema Pruning	7865
<i>Zhili Shen, Pavlos Vougiouklis, Chenxin Diao, Kaustubh Vyas, Yuanyi Ji, Jeff Z. Pan</i>	
Mixture-Of-Subspaces in Low-Rank Adaptation	7880
<i>Taiqiang Wu, Jiahao Wang, Zhe Zhao, Ngai Wong</i>	
PARIKSHA: A Large-Scale Investigation of Human-LLM Evaluator Agreement on Multilingual and Multi-Cultural Data	7900
<i>Ishaan Watts, Varun Gumma, Aditya Yadavalli, Vivek Seshadri, Manohar Swaminathan, Sunayana Sitaram</i>	
LawBench: Benchmarking Legal Knowledge of Large Language Models	7933
<i>Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Alan Huang, Songyang Zhang, Kai Chen, Zhixin Yin, Zongwen Shen, Jidong Ge, Vincent Ng</i>	
Efficient Performance Tracking: Leveraging Large Language Models for Automated Construction of Scientific Leaderboards	7963
<i>Furkan Sahinuç, Thy Thy Tran, Yulia Grishina, Yufang Hou, Bei Chen, Iryna Gurevych</i>	
Efficient Vision-Language Pre-Training Via Domain-Specific Learning for Human Activities	7978
<i>Adrian Bulat, Yassine Ouali, Ricardo Guerrero, Brais Martinez, Georgios Tzimiropoulos</i>	
Empowering Backbone Models for Visual Text Generation with Input Granularity Control and Glyph-Aware Training	8001
<i>Wenbo Li, Guohao Li, Zhibin Lan, Xue Xu, Wanru Zhuang, Jiachen Liu, Xinyan Xiao, Jinsong Su</i>	
Evaluating Character Understanding of Large Language Models Via Character Profiling from Fictional Works	8015
<i>Xinfeng Yuan, Siyu Yuan, Yuhan Cui, Tianhe Lin, Xintao Wang, Rui Xu, Jiangjie Chen, Deqing Yang</i>	

Getting More from Less: Large Language Models Are Good Spontaneous Multilingual Learners	8037
<i>Shimao Zhang, Changjiang Gao, Wenhao Zhu, Jiajun Chen, Xin Huang, Xue Han, Junlan Feng, Chao Deng, Shujian Huang</i>	
AdaSwitch: Adaptive Switching Between Small and Large Agents for Effective Cloud-Local Collaborative Learning.....	8052
<i>Hao Sun, Jiayi Wu, Hengyi Cai, Xiaochi Wei, Yue Feng, Bo Wang, Shuaiqiang Wang, Yan Zhang, Dawei Yin</i>	
CoBa: Convergence Balancer for Multitask Finetuning of Large Language Models.....	8063
<i>Zi Gong, Hang Yu, Cong Liao, Bingchang Liu, Chaoyu Chen, Jianguo Li</i>	
MDPO: Conditional Preference Optimization for Multimodal Large Language Models.....	8078
<i>Fei Wang, Wenxuan Zhou, James Y. Huang, Nan Xu, Sheng Zhang, Hoifung Poon, Muhao Chen</i>	
Data Advisor: Dynamic Data Curation for Safety Alignment of Large Language Models	8089
<i>Fei Wang, Ninareh Mehrabi, Palash Goyal, Rahul Gupta, Kai-Wei Chang, Aram Galstyan</i>	
Language-To-Code Translation with a Single Labeled Example	8101
<i>Kaj Bostrom, Harsh Jhamtani, Hao Fang, Sam Thomson, Richard Shin, Patrick Xia, Benjamin Van Durme, Jason Eisner, Jacob Andreas</i>	
Attribute Or Abstain: Large Language Models as Long Document Assistants	8113
<i>Jan Buchmann, Xiao Liu, Iryna Gurevych</i>	
FEDKIM: Adaptive Federated Knowledge Injection into Medical Foundation Models	8141
<i>Xiaochen Wang, Jiaqi Wang, Houping Xiao, Jinghui Chen, Fenglong Ma</i>	
Retrieved In-Context Principles from Previous Mistakes	8155
<i>Hao Sun, Yong Jiang, Bo Wang, Yingyan Hou, Yan Zhang, Pengjun Xie, Fei Huang</i>	
EmoKnob: Enhance Voice Cloning with Fine-Grained Emotion Control	8170
<i>Haozhe Chen, Run Chen, Julia Hirschberg</i>	
VPTQ: Extreme Low-Bit Vector Post-Training Quantization for Large Language Models.....	8181
<i>Yifei Liu, Jicheng Wen, Yang Wang, Shengyu Ye, Li Lina Zhang, Ting Cao, Cheng Li, Mao Yang</i>	
An L* Algorithm for Deterministic Weighted Regular Languages	8197
<i>Clemente Pasti, Talu Karagöz, Franz Nowak, Anej Svete, Reda Boumasmoud, Ryan Cotterell</i>	
Towards Verifiable Text Generation with Evolving Memory and Self-Reflection	8211
<i>Hao Sun, Hengyi Cai, Bo Wang, Yingyan Hou, Xiaochi Wei, Shuaiqiang Wang, Yan Zhang, Dawei Yin</i>	
Pelican: Correcting Hallucination in Vision-LLMs Via Claim Decomposition and Program of Thought Verification.....	8228
<i>Pritish Sahu, Karan Sikka, Ajay Divakaran</i>	
Resampled Datasets Are Not Enough: Mitigating Societal Bias Beyond Single Attributes	8249
<i>Yusuke Hirota, Jerone Andrews, Dora Zhao, Orestis Papakyriakopoulos, Apostolos Modas, Yuta Nakashima, Alice Xiang</i>	
RealVul: Can We Detect Vulnerabilities in Web Applications with LLM?	8268
<i>Di Cao, Yong Liao, Xiuwei Shang</i>	

Unsupervised End-To-End Task-Oriented Dialogue with LLMs: The Power of the Noisy Channel	8283
<i>Brendan King, Jeffrey Flanigan</i>	
Humans Or LLMs as the Judge? a Study on Judgement Bias	8301
<i>Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, Benyou Wang</i>	
WPO: Enhancing RLHF with Weighted Preference Optimization	8328
<i>Wenxuan Zhou, Ravi Agrawal, Shujian Zhang, Sathish Reddy Indurthi, Sanqiang Zhao, Kaiqiang Song, Silei Xu, Chenguang Zhu</i>	
Walking in Others' Shoes: How Perspective-Taking Guides Large Language Models in Reducing Toxicity and Bias	8341
<i>Rongwu Xu, Zian Zhou, Tianwei Zhang, Zehan Qi, Su Yao, Ke Xu, Wei Xu, Han Qiu</i>	
MetaReflection: Learning Instructions for Language Agents Using Past Reflections.....	8369
<i>Priyanshu Gupta, Shashank Kirtania, Ananya Singha, Sumit Gulwani, Arjun Radhakrishna, Gustavo Soares, Sherry Shi</i>	
Stepwise Verification and Remediation of Student Reasoning Errors with Large Language Model Tutors.....	8386
<i>Nico Daheim, Jakub Macina, Manu Kapur, Iryna Gurevych, Mrinmaya Sachan</i>	
On Eliciting Syntax from Language Models Via Hashing	8412
<i>Yiran Wang, Masao Utiyama</i>	
CliMedBench: A Large-Scale Chinese Benchmark for Evaluating Medical Large Language Models in Clinical Scenarios.....	8428
<i>Zetian Ouyang, Yishuai Qiu, Linlin Wang, Gerard De Melo, Ya Zhang, Yanfeng Wang, Liang He</i>	
The Best Defense is Attack: Repairing Semantics in Textual Adversarial Examples.....	8439
<i>Heng Yang, Ke Li</i>	
CSSL: Contrastive Self-Supervised Learning for Dependency Parsing on Relatively Free Word Ordered and Morphologically Rich Low Resource Languages.....	8458
<i>Pretam Ray, Jivnesh Sandhan, Amrith Krishna, Pawan Goyal</i>	
Perceptions of Linguistic Uncertainty by Language Models and Humans.....	8467
<i>Catarina G Belém, Markelle Kelly, Mark Steyvers, Sameer Singh, Padhraic Smyth</i>	
Explaining and Improving Contrastive Decoding by Extrapolating the Probabilities of a Huge and Hypothetical LM	8503
<i>Haw-Shiuan Chang, Nanyun Peng, Mohit Bansal, Anil Ramakrishna, Tagyoung Chung</i>	

VOLUME 12

Zero-Shot Cross-Domain Dialogue State Tracking Via Context-Aware Auto-Prompting and Instruction-Following Contrastive Decoding	8527
<i>Xiaoyu Dong, Yujie Feng, Zexin Lu, Guangyuan Shi, Xiao-Ming Wu</i>	
Knowledge Conflicts for LLMs: A Survey.....	8541
<i>Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, Wei Xu</i>	
MisinfoEval: Generative AI in the Era of “Alternative Facts”	8566
<i>Saadia Gabriel, Liang Lyu, James Siderius, Marzyeh Ghassemi, Jacob Andreas, Asuman E. Ozdaglar</i>	

MEANT: Multimodal Encoder for Antecedent Information	8579
<i>Benjamin Irving, Annika Marie Schoene</i>	
A Thorough Examination of Decoding Methods in the Era of LLMs	8601
<i>Chufan Shi, Haoran Yang, Deng Cai, Zhisong Zhang, Yifan Wang, Yujiu Yang, Wai Lam</i>	
AGRaME: Any-Granularity Ranking with Multi-Vector Embeddings.....	8630
<i>Revanth Gangi Reddy, Omar Attia, Yunyao Li, Heng Ji, Saloni Potdar</i>	
FIRST: Faster Improved Listwise Reranking with Single Token Decoding.....	8642
<i>Revanth Gangi Reddy, Jaehyeok Doo, Yifei Xu, Md Arafat Sultan, Deevya Swain, Avirup Sil, Heng Ji</i>	
Exploring Nested Named Entity Recognition with Large Language Models: Methods, Challenges, and Insights.....	8653
<i>Hongjin Kim, Jai-Eun Kim, Harksoo Kim</i>	
ReCaLL: Membership Inference Via Relative Conditional Log-Likelihoods	8671
<i>Roy Xie, Junlin Wang, Ruomin Huang, Minxing Zhang, Rong Ge, Jian Pei, Neil Zhenqiang Gong, Bhuwan Dhingra</i>	
“Flex Tape Can’t Fix That”: Bias and Misinformation in Edited Language Models.....	8690
<i>Karina H Halevy, Anna Sotnikova, Badr Alkhamissi, Syrielle Montariol, Antoine Bosselut</i>	
Revisiting Who’s Harry Potter: Towards Targeted Unlearning from a Causal Intervention Perspective.....	8708
<i>Yujian Liu, Yang Zhang, Tommi Jaakkola, Shiyu Chang</i>	
LIONS: An Empirically Optimized Approach to Align Language Models.....	8732
<i>Xiao Yu, Qingyang Wu, Yu Li, Zhou Yu</i>	
Jellyfish: Instruction-Tuning Local Large Language Models for Data Preprocessing	8754
<i>Haochen Zhang, Yuyang Dong, Chuan Xiao, Masafumi Oyamada</i>	
A Comprehensive Survey of Scientific Large Language Models and Their Applications in Scientific Discovery	8783
<i>Yu Zhang, Xiusi Chen, Bowen Jin, Sheng Wang, Shuiwang Ji, Wei Wang, Jiawei Han</i>	
MiniCheck: Efficient Fact-Checking of LLMs on Grounding Documents	8818
<i>Liyan Tang, Philippe Laban, Greg Durrett</i>	
Beyond Label Attention: Transparency in Language Models for Automated Medical Coding Via Dictionary Learning.....	8848
<i>John Wu, David Wu, Jimeng Sun</i>	
MOSEL: Inference Serving Using Dynamic Modality Selection.....	8872
<i>Bodun Hu, Le Xu, Jeongyoon Moon, Neeraja J Yadwadkar, Aditya Akella</i>	
From RAG to Riches: Retrieval Interlaced with Sequence Generation.....	8887
<i>Palak Jain, Livio Baldini Soares, Tom Kwiatkowski</i>	
Task Arithmetic Can Mitigate Synthetic-To-Real Gap in Automatic Speech Recognition	8905
<i>Hsuan Su, Hua Farn, Fan-Yun Sun, Shang-Tse Chen, Hung-Yi Lee</i>	
Learning to Correct for QA Reasoning with Black-Box LLMs.....	8916
<i>Jaehyung Kim, Dongyoung Kim, Yiming Yang</i>	

AssistantBench: Can Web Agents Solve Realistic and Time-Consuming Tasks?	8938
<i>Ori Yoran, Samuel Joseph Amouyal, Chaitanya Malaviya, Ben Bogin, Ofir Press, Jonathan Berant</i>	
PostMark: A Robust Blackbox Watermark for Large Language Models.....	8969
<i>Yapei Chang, Kalpesh Krishna, Amir Houmansadr, John Frederick Wieting, Mohit Iyer</i>	
Assessing “Implicit” Retrieval Robustness of Large Language Models	8988
<i>Xiaoyu Shen, Rexhina Blloshmi, Dawei Zhu, Jiahuan Pei, Wei Zhang</i>	
On the Relationship Between Truth and Political Bias in Language Models.....	9004
<i>Suyash Fulay, William Brannon, Shrestha Mohanty, Cassandra Overney, Elinor Poole-Dayan, Deb Roy, Jad Kabbara</i>	
Can Active Label Correction Improve LLM-Based Modular AI Systems?.....	9019
<i>Karan Taneja, Ashok Goel</i>	
Statistical Uncertainty in Word Embeddings: GloVe-V	9032
<i>Andrea Vallebuono, Cassandra Handan-Nader, Christopher D Manning, Daniel E. Ho</i>	
Annotation Alignment: Comparing LLM and Human Annotations of Conversational Safety.....	9048
<i>Rajiv Movva, Pang Wei Koh, Emma Pierson</i>	
DiVERT: Distractor Generation with Variational Errors Represented as Text for Math Multiple-Choice Questions.....	9063
<i>Nigel Fernandez, Alexander Scarlatos, Wanyong Feng, Simon Woodhead, Andrew Lan</i>	
The Factuality Tax of Diversity-Intervened Text-To-Image Generation: Benchmark and Fact-Augmented Intervention.....	9082
<i>Yixin Wan, Di Wu, Haoran Wang, Kai-Wei Chang</i>	
CleanGen: Mitigating Backdoor Attacks for Generation Tasks in Large Language Models.....	9101
<i>Yuetai Li, Zhangchen Xu, Fengqing Jiang, Luyao Niu, Dinuka Sahabandu, Bhaskar Ramasubramanian, Radha Poovendran</i>	
Enhancing Reinforcement Learning with Dense Rewards from Language Model Critic	9119
<i>Meng Cao, Lei Shu, Lei Yu, Yun Zhu, Nevan Wichers, Yinxiao Liu, Lei Meng</i>	
Words Matter: Reducing Stigma in Online Conversations About Substance Use with Large Language Models	9139
<i>Layla Bouzoubaa, Elham Aghakhani, Rezvaneh Rezapour</i>	
Efficient Sequential Decision Making with Large Language Models	9157
<i>Dingyang Chen, Qi Zhang, Yinglun Zhu</i>	
SignCLIP: Connecting Text and Sign Language by Contrastive Learning	9171
<i>Zifan Jiang, Gerard Sant, Amit Moryossef, Mathias Müller, Rico Sennrich, Sarah Ebling</i>	
APPLS: Evaluating Evaluation Metrics for Plain Language Summarization.....	9194
<i>Yue Guo, Tal August, Gondy Leroy, Trevor Cohen, Lucy Lu Wang</i>	
Ontologically Faithful Generation of Non-Player Character Dialogues.....	9212
<i>Nathaniel Weir, Ryan Thomas, Randolph D’Amore, Kellie Hill, Benjamin Van Durme, Harsh Jhamtani</i>	
LLM See, LLM Do: Leveraging Active Inheritance to Target Non-Differentiable Objectives.....	9243
<i>Luísa Shimabucoro, Sebastian Ruder, Julia Kreutzer, Marzieh Fadaee, Sara Hooker</i>	

RuBLiMP: Russian Benchmark of Linguistic Minimal Pairs 9268
Ekaterina Taktasheva, Maxim Bazhukov, Kirill Koncha, Alena Fenogenova, Ekaterina Artemova, Vladislav Mikhailov

Text-Tuple-Table: Towards Information Integration in Text-To-Table Generation Via Global Tuple
Extraction 9300
*Zheyue Deng, Chunkit Chan, Weiqi Wang, Yuxi Sun, Wei Fan, Tianshi Zheng, Yauwai Yim,
Yangqiu Song*

VOLUME 13

Toward Compositional Behavior in Neural Models: A Survey of Current Views 9323
Kate McCurdy, Paul Soulos, Paul Smolensky, Roland Fernandez, Jianfeng Gao

Optimizing Instructions and Demonstrations for Multi-Stage Language Model Programs 9340
Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, Omar Khattab

Reverse-Engineering the Reader 9367
Samuel Kieglend, Ethan Wilcox, Afra Amini, David Robert Reich, Ryan Cotterell

Synchronous Faithfulness Monitoring for Trustworthy Retrieval-Augmented Generation..... 9390
Di Wu, Jia-Chen Gu, Fan Yin, Nanyun Peng, Kai-Wei Chang

Structure Guided Prompt: Instructing Large Language Model in Multi-Step Reasoning by
Exploring Graph Structure of the Text 9407
Kewei Cheng, Nesreen K. Ahmed, Theodore L. Willke, Yizhou Sun

Less is More: Parameter-Efficient Selection of Intermediate Tasks for Transfer Learning 9431
David Schulte, Felix Hamborg, Alan Akbik

The Effects of Distance on NPI Illusive Effects in BERT 9443
So Young Lee, Mai Ha Vu

Enhancing Systematic Decompositional Natural Language Inference Using Informal Logic 9458
Nathaniel Weir, Kate Sanders, Orion Weller, Shreya Sharma, Dongwei Jiang, Zhengping Jiang, Bhavana Dalvi Mishra, Oyvind Tafjord, Peter Jansen, Peter Clark, Benjamin Van Durme

Susu Box Or Piggy Bank: Assessing Cultural Commonsense Knowledge Between Ghana and the
US..... 9483
Christabel Acquaye, Haozhe An, Rachel Rudinger

Read Anywhere Pointed: Layout-Aware GUI Screen Reading with Tree-Of-Lens Grounding 9503
*Yue Fan, Lei Ding, Ching-Chen Kuo, Shan Jiang, Yang Zhao, Xinze Guan, Jie Yang, Yi Zhang,
Xin Eric Wang*

Ranking Manipulation for Conversational Search Engines..... 9523
Samuel Pfommer, Yatong Bai, Tanmay Gautam, Somayeh Sojoudi

Fast Forwarding Low-Rank Training 9553
Adir Rahamim, Naomi Saphra, Sara Kangaslahti, Yonatan Belinkov

Precise Model Benchmarking with Only a Few Observations 9563
Riccardo Fogliato, Pratik Patil, Nil-Jana Akpinar, Mathew Monfort

Attribute Diversity Determines the Systematicity Gap in VQA.....	9576
<i>Ian Berlot-Attwell, Kumar Krishna Agrawal, Annabelle Michael Carrell, Yash Sharma, Naomi Saphra</i>	
ArxivDIGESTables: Synthesizing Scientific Literature into Tables Using Language Models.....	9612
<i>Benjamin Newman, Yoonjoo Lee, Aakanksha Naik, Pao Siangliulue, Raymond Fok, Juho Kim, Daniel S Weld, Joseph Chee Chang, Kyle Lo</i>	
Development of Cognitive Intelligence in Pre-Trained Language Models	9632
<i>Raj Sanjay Shah, Khushi Bhardwaj, Sashank Varma</i>	
Modeling Layout Reading Order as Ordering Relations for Visually-Rich Document Understanding.....	9658
<i>Chong Zhang, Yi Tu, Yixi Zhao, Chenshu Yuan, Huan Chen, Yue Zhang, Mingxu Chai, Ya Guo, Huijia Zhu, Qi Zhang, Tao Gui</i>	
Birdie: Advancing State Space Language Modeling with Dynamic Mixtures of Training Objectives.....	9679
<i>Sam Blouir, Jimmy T. H. Smith, Antonios Anastasopoulos, Amarda Shehu</i>	
Is it Good Data for Multilingual Instruction Tuning Or Just Bad Multilingual Evaluation for Large Language Models?	9706
<i>Pinzhen Chen, Simon Yu, Zhicheng Guo, Barry Haddow</i>	
Token Erasure as a Footprint of Implicit Vocabulary Items in LLMs	9727
<i>Sheridan Feucht, David Atkinson, Byron C Wallace, David Bau</i>	
TravelER: A Modular Multi-LMM Agent Framework for Video Question-Answering.....	9740
<i>Chuyi Shang, Amos You, Sanjay Subramanian, Trevor Darrell, Roei Herzig</i>	
Evaluating the Effectiveness of Large Language Models in Establishing Conversational Grounding.....	9767
<i>Biswesh Mohapatra, Manav Nitin Kapadnis, Laurent Romary, Justine Cassell</i>	
Unlocking Memorization in Large Language Models with Dynamic Soft Prompting	9782
<i>Zhepeng Wang, Runxue Bao, Yawen Wu, Jackson Taylor, Cao Xiao, Feng Zheng, Weiwen Jiang, Shangqian Gao, Yanfu Zhang</i>	
If CLIP Could Talk: Understanding Vision-Language Model Representations Through Their Preferred Concept Descriptions.....	9797
<i>Reza Esfandiarpour, Cristina Menghini, Stephen Bach</i>	
Extract, Define, Canonicalize: An LLM-Based Framework for Knowledge Graph Construction	9820
<i>Bowen Zhang, Harold Soh</i>	
MQuinE: A Cure for “Z-Paradox” in Knowledge Graph Embedding	9837
<i>Yang Liu, Huang Fang, Yunfeng Cai, Mingming Sun</i>	
Can Transformers Learn N-Gram Language Models?	9851
<i>Anej Svete, Nadav Borenstein, Mike Zhou, Isabelle Augenstein, Ryan Cotterell</i>	
StablePrompt : Automatic Prompt Tuning Using Reinforcement Learning for Large Language Model	9868
<i>Minchan Kwon, Gaeun Kim, Jongsuk Kim, Haeil Lee, Junmo Kim</i>	
Summary of a Haystack: A Challenge to Long-Context LLMs and RAG Systems	9885
<i>Philippe Laban, Alexander Fabbri, Caiming Xiong, Chien-Sheng Wu</i>	
Multi-Pass Decoding for Grammatical Error Correction.....	9904
<i>Xiaoying Wang, Lingling Mu, Jingyi Zhang, Hongfei Xu</i>	

Into the Unknown Unknowns: Engaged Human Learning Through Participation in Language Model Agent Conversations	9917
<i>Yucheng Jiang, Yijia Shao, Dekun Ma, Sina Semnani, Monica Lam</i>	
SCOI: Syntax-Augmented Coverage-Based In-Context Example Selection for Machine Translation	9956
<i>Chenming Tang, Zhixiang Wang, Yunfang Wu</i>	
Efficient Temporal Extrapolation of Multimodal Large Language Models with Temporal Grounding Bridge.....	9972
<i>Yuxuan Wang, Yueqian Wang, Pengfei Wu, Jianxin Liang, Dongyan Zhao, Yang Liu, Zilong Zheng</i>	
STORYSUMM: Evaluating Faithfulness in Story Summarization	9988
<i>Melanie Subbiah, Faisal Ladhak, Akankshya Mishra, Griffin Thomas Adams, Lydia Chilton, Kathleen McKeown</i>	
MMoE: Enhancing Multimodal Models with Mixtures of Multimodal Interaction Experts	10006
<i>Haofei Yu, Zhengyang Qi, Lawrence Keunho Jang, Russ Salakhutdinov, Louis-Philippe Morency, Paul Pu Liang</i>	
OmAgent: A Multi-Modal Agent Framework for Complex Video Understanding with Task Divide-And-Conquer	10031
<i>Lu Zhang, Tiancheng Zhao, Heting Ying, Yibo Ma, Kyusong Lee</i>	
Enhancing Pre-Trained Generative Language Models with Question Attended Span Extraction on Machine Reading Comprehension.....	10046
<i>Lin Ai, Zheng Hui, Zizhou Liu, Julia Hirschberg</i>	
CommonIT: Commonality-Aware Instruction Tuning for Large Language Models Via Data Partitions.....	10064
<i>Jun Rao, Xuebo Liu, Lian Lian, Shengjun Cheng, Yunjie Liao, Min Zhang</i>	
ESC: Efficient Speech Coding with Cross-Scale Residual Vector Quantized Transformers.....	10084
<i>Yuzhe Gu, Enmao Diao</i>	
Breaking ReLU Barrier: Generalized MoEfication for Dense Pretrained Models	10097
<i>Jaeseong Lee, Seung-Won Hwang, Wonpyo Park, Mingi Ji</i>	

VOLUME 14

Detecting Subtle Differences Between Human and Model Languages Using Spectrum of Relative Likelihood	10108
<i>Yang Xu, Yu Wang, Hao An, Zhichen Liu, Yongyuan Li</i>	
Optimizing Language Models with Fair and Stable Reward Composition in Reinforcement Learning	10122
<i>Jiahui Li, Hanlin Zhang, Fengda Zhang, Tai-Wei Chang, Kun Kuang, Long Chen, Jun Zhou</i>	
Fine-Grained Pluggable Gradient Ascent for Knowledge Unlearning in Language Models.....	10141
<i>Xiaohua Feng, Chaochao Chen, Yuyuan Li, Zibin Lin</i>	
ARM: An Alignment-And-Replacement Module for Chinese Spelling Check Based on LLMs	10156
<i>Changchun Liu, Kai Zhang, Junzhe Jiang, Zirui Liu, Hanqing Tao, Min Gao, Enhong Chen</i>	
On the In-Context Generation of Language Models	10169
<i>Zhongtao Jiang, Yanzhe Zhang, Kun Luo, Xiaowei Yuan, Jun Zhao, Kang Liu</i>	

Atomic Inference for NLI with Generated Facts as Atoms	10188
<i>Joe Stacey, Pasquale Minervini, Haim Dubossarsky, Oana-Maria Camburu, Marek Rei</i>	
Towards Robust Speech Representation Learning for Thousands of Languages	10205
<i>William Chen, Wangyou Zhang, Yifan Peng, Xinjian Li, Jinchuan Tian, Jiatong Shi, Xuankai Chang, Soumi Maiti, Karen Livescu, Shinji Watanabe</i>	
I Learn Better if You Speak My Language: Understanding the Superior Performance of Fine-Tuning Large Language Models with LLM-Generated Responses	10225
<i>Xuan Ren, Biao Wu, Lingqiao Liu</i>	
PreAlign: Boosting Cross-Lingual Transfer by Early Establishment of Multilingual Alignment	10246
<i>Jiahuan Li, Shujian Huang, Aarron Ching, Xinyu Dai, Jiajun Chen</i>	
An Image Speaks a Thousand Words, but Can Everyone Listen? on Image Transcreation for Cultural Relevance	10258
<i>Simran Khanuja, Sathyanarayanan Ramamoorthy, Yueqi Song, Graham Neubig</i>	
When Parts Are Greater than Sums: Individual LLM Components Can Outperform Full Models	10280
<i>Ting-Yun Chang, Jesse Thomason, Robin Jia</i>	
Multimodal Clickbait Detection by De-Confounding Biases Using Causal Representation Inference	10300
<i>Jianxing Yu, Shiqi Wang, Han Yin, Zhenlong Sun, Ruobing Xie, Bo Zhang, Yanghui Rao</i>	
Matryoshka-Adaptor: Unsupervised and Supervised Tuning for Smaller Embedding Dimensions	10318
<i>Jinsung Yoon, Rajarishi Sinha, Sercan O Arik, Tomas Pfister</i>	
KNN-Instruct: Automatic Instruction Construction with K Nearest Neighbor Deduction	10337
<i>Jianshang Kou, Benfeng Xu, Chiwei Zhu, Zhendong Mao</i>	
Contextualized Sequence Likelihood: Enhanced Confidence Scores for Natural Language Generation	10351
<i>Zhen Lin, Shubhendu Trivedi, Jimeng Sun</i>	
MixGR: Enhancing Retriever Generalization for Scientific Domain Through Complementary Granularity	10369
<i>Fengyu Cai, Xinran Zhao, Tong Chen, Sihao Chen, Hongming Zhang, Iryna Gurevych, Heinz Koepl</i>	
CARER - ClinicAI Reasoning-Enhanced Representation for Temporal Health Risk Prediction	10392
<i>Tuan Dung Nguyen, Thanh Trung Huynh, Minh Hieu Phan, Quoc Viet Hung Nguyen, Phi Le Nguyen</i>	
“In-Dialogues We Learn”: Towards Personalized Dialogue Without Pre-Defined Profiles Through In-Dialogue Learning	10408
<i>Chuanqi Cheng, Quan Tu, Wei Wu, Shuo Shang, Cunli Mao, Zhengtao Yu, Rui Yan</i>	
Encourage Or Inhibit Monosemanticity? Revisit Monosemanticity from a Feature Decorrelation Perspective	10423
<i>Hanqi Yan, Yanzheng Xiang, Guangyi Chen, Yifei Wang, Lin Gui, Yulan He</i>	
Enhancing Language Model Factuality Via Activation-Based Confidence Calibration and Guided Decoding	10436
<i>Xin Liu, Farima Fatahi Bayat, Lu Wang</i>	

Reasoning Robustness of LLMs to Adversarial Typographical Errors.....	10449
<i>Esther Gan, Yiran Zhao, Liying Cheng, Mao Yancan, Anirudh Goyal, Kenji Kawaguchi, Min-Yen Kan, Michael Shieh</i>	
InferAligner: Inference-Time Alignment for Harmlessness Through Cross-Model Guidance.....	10460
<i>Pengyu Wang, Dong Zhang, Linyang Li, Chenkun Tan, Xinghao Wang, Mozhi Zhang, Ke Ren, Botian Jiang, Xipeng Qiu</i>	
Belief Revision: The Adaptability of Large Language Models Reasoning	10480
<i>Bryan Wilie, Samuel Cahyawijaya, Etsuko Ishii, Junxian He, Pascale Fung</i>	
Fisher Information-Based Efficient Curriculum Federated Learning with Large Language Models	10497
<i>Ji Liu, Jiaxiang Ren, Ruoming Jin, Zijie Zhang, Yang Zhou, Patrick Valduriez, Dejing Dou</i>	
Bio-RFX: Refining Biomedical Extraction Via Advanced Relation Classification and Structural Constraints.....	10524
<i>Minjia Wang, Fangzhou Liu, Xiuxing Li, Bowen Dong, Zhenyu Li, Tengyu Pan, Jianyong Wang</i>	
Decoding Matters: Addressing Amplification Bias and Homogeneity Issue in Recommendations for Large Language Models	10540
<i>Keqin Bao, Jizhi Zhang, Yang Zhang, Xinyue Huo, Chong Chen, Fuli Feng</i>	
LLMs Are Prone to Fallacies in Causal Inference.....	10553
<i>Nitish Joshi, Abulhair Saparov, Yixin Wang, He He</i>	
Roleplay-Doh: Enabling Domain-Experts to Create LLM-Simulated Patients Via Eliciting and Adhering to Principles.....	10570
<i>Ryan Louie, Ananjan Nandi, William Fang, Cheng Chang, Emma Brunskill, Diyi Yang</i>	
The Lou Dataset - Exploring the Impact of Gender-Fair Language in German Text Classification.....	10604
<i>Andreas Waldis, Joel Birrer, Anne Lauscher, Iryna Gurevych</i>	
When Generative Adversarial Networks Meet Sequence Labeling Challenges	10625
<i>Yu Tong, Ge Chen, Guokai Zheng, Rui Li, Jiang Dazhi</i>	
Evidence-Focused Fact Summarization for Knowledge-Augmented Zero-Shot Question Answering	10636
<i>Sungho Ko, Hyunjin Cho, Hyungjoo Chae, Jinyoung Yeo, Dongha Lee</i>	
Speechworthy Instruction-Tuned Language Models	10652
<i>Hyundong Justin Cho, Nicolaas Paul Jedema, Leonardo F. R. Ribeiro, Karishma Sharma, Pedro Szekely, Alessandro Moschitti, Ruben Janssen, Jonathan May</i>	
Data, Data Everywhere: A Guide for Pretraining Dataset Construction.....	10671
<i>Jupinder Parmar, Shrimai Prabhunoye, Joseph Jennings, Bo Liu, Aastha Jhunjhunwala, Zhilin Wang, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro</i>	
Fine-Tuning and Prompt Optimization: Two Great Steps that Work Better Together	10696
<i>Dilara Soylu, Christopher Potts, Omar Khattab</i>	
Demystifying Verbatim Memorization in Large Language Models	10711
<i>Jing Huang, Diyi Yang, Christopher Potts</i>	
AmbigNLG: Addressing Task Ambiguity in Instruction for NLG	10733
<i>Ayana Niwa, Hayate Iso</i>	

Distributional Properties of Subword Regularization.....	10753
<i>Marco Cognetta, Vilém Zouhar, Naoaki Okazaki</i>	
DataTales: A Benchmark for Real-World Intelligent Data Narration	10764
<i>Yajing Yang, Qian Liu, Min-Yen Kan</i>	
Towards Fast Multilingual LLM Inference: Speculative Decoding and Specialized Drafters	10789
<i>Euiin Yi, Taehyeon Kim, Hongseok Jeung, Du-Seong Chang, Se-Young Yun</i>	
GlobeSumm: A Challenging Benchmark Towards Unifying Multi-Lingual, Cross-Lingual and Multi-Document News Summarization.....	10803
<i>Yangfan Ye, Xiachong Feng, Xiaocheng Feng, Weitao Ma, Libo Qin, Dongliang Xu, Qing Yang, Hongtao Liu, Bing Qin</i>	
Breaking the Curse of Multilinguality with Cross-Lingual Expert Language Models	10822
<i>Terra Blevins, Tomasz Limisiewicz, Suchin Gururangan, Margaret Li, Hila Gonen, Noah A. Smith, Luke Zettlemoyer</i>	
More Insightful Feedback for Tutoring: Enhancing Generation Mechanisms and Automatic Evaluation.....	10838
<i>Wencke Liermann, Jin-Xia Huang, Yohan Lee, Kong Joo Lee</i>	
Stable Language Model Pre-Training by Reducing Embedding Variability.....	10852
<i>Woojin Chung, Jiwoo Hong, Na Min An, James Thorne, Se-Young Yun</i>	
What is Lost in Normalization? Exploring Pitfalls in Multilingual ASR Model Evaluations	10864
<i>Kavya Manohar, Leena G Pillai</i>	
Diversity Over Size: On the Effect of Sample and Topic Sizes for Topic-Dependent Argument Mining Datasets.....	10870
<i>Benjamin Schiller, Johannes Daxenberger, Andreas Waldis, Iryna Gurevych</i>	
Kiss Up, Kick Down: Exploring Behavioral Changes in Multi-Modal Large Language Models with Assigned Visual Personas.....	10888
<i>Seungjong Sun, Eungu Lee, Seo Yeon Baek, Seunghyun Hwang, Wonbyung Lee, Dongyan Nan, Bernard J Jansen, Jang Hyun Kim</i>	

VOLUME 15

ATM: Adversarial Tuning Multi-Agent System Makes a Robust Retrieval-Augmented Generator	10902
<i>Junda Zhu, Lingyong Yan, Haibo Shi, Dawei Yin, Lei Sha</i>	
Dynamic Multi-Granularity Attribution Network for Aspect-Based Sentiment Analysis	10920
<i>Yanjiang Chen, Kai Zhang, Feng Hu, Xianquan Wang, Ruikang Li, Qi Liu</i>	
Unlabeled Debiasing in Downstream Tasks Via Class-Wise Low Variance Regularization	10932
<i>Shahed Masoudian, Markus Frohmann, Navid Rekabsaz, Markus Schedl</i>	
Large Language Models Know What is Key Visual Entity: An LLM-Assisted Multimodal Retrieval for VQA.....	10939
<i>Pu Jian, Donglei Yu, Jiajun Zhang</i>	
Towards Probing Speech-Specific Risks in Large Multimodal Models: A Taxonomy, Benchmark, and Insights.....	10957
<i>Hao Yang, Lizhen Qu, Ehsan Shareghi, Reza Haf</i>	

Self-AMPLIFY: Improving Small Language Models with Self Post Hoc Explanations	10974
<i>Milan Bhan, Jean-Noël Vittaut, Nicolas Chesneau, Marie-Jeanne Lesot</i>	
What Are the Generator Preferences for End-To-End Task-Oriented Dialog System?	10992
<i>Wanshi Xu, Xianwei Zhuang, Zhanpeng Chen, Zhihong Zhu, Xuxin Cheng, Yuexian Zou</i>	
Paraphrase Types Elicit Prompt Engineering Capabilities	11004
<i>Jan Philip Wahle, Terry Ruas, Yang Xu, Bela Gipp</i>	
VLEU: A Method for Automatic Evaluation for Generalizability of Text-To-Image Models	11034
<i>Jingtao Cao, Zhang Zheng, Hongru Wang, Kam-Fai Wong</i>	
Towards Online Continuous Sign Language Recognition and Translation	11050
<i>Ronglai Zuo, Fangyun Wei, Brian Mak</i>	
Mitigate Extrinsic Social Bias in Pre-Trained Language Models Via Continuous Prompts Adjustment	11068
<i>Yiwei Dai, Hengrui Gu, Ying Wang, Xin Wang</i>	
Integrating Argumentation and Hate-Speech-Based Techniques for Countering Misinformation	11109
<i>Sougata Saha, Rohini Srihari</i>	
BPO: Staying Close to the Behavior LLM Creates Better Online LLM Alignment.....	11125
<i>Wenda Xu, Jiachen Li, William Yang Wang, Lei Li</i>	
One2Set + Large Language Model: Best Partners for Keyphrase Generation	11140
<i>Liangying Shao, Liang Zhang, Minlong Peng, Guoqi Ma, Hao Yue, Mingming Sun, Jinsong Su</i>	
Unlocking Markets: A Multilingual Benchmark to Cross-Market Question Answering	11154
<i>Yifei Yuan, Yang Deng, Anders Søgaard, Mohammad Aliannejadi</i>	
ORPO: Monolithic Preference Optimization Without Reference Model	11170
<i>Jiwoo Hong, Noah Lee, James Thorne</i>	
A Multi-Perspective Analysis of Memorization in Large Language Models.....	11190
<i>Bowen Chen, Namgi Han, Yusuke Miyao</i>	
Do LLMs Suffer from Multi-Party Hangover? a Diagnostic Approach to Addressee Recognition and Response Selection in Conversations	11210
<i>Nicolò Penzo, Maryam Sajedinia, Bruno Lepri, Sara Tonelli, Marco Guerini</i>	
Code Prompting Elicits Conditional Reasoning Abilities in Text+Code LLMs	11234
<i>Haritz Puerto, Martin Tutek, Somak Aditya, Xiaodan Zhu, Iryna Gurevych</i>	
Unveiling the Role of Pretraining in Direct Speech Translation	11259
<i>Belen Alastruey, Gerard I. Gállego, Marta R. Costa-Jussà</i>	
PCQPR: Proactive Conversational Question Planning with Reflection	11266
<i>Shasha Guo, Lizi Liao, Jing Zhang, Cuiping Li, Hong Chen</i>	
CodeAgent: Autonomous Communicative Agents for Code Review	11279
<i>Xunzhu Tang, Kisub Kim, Yewei Song, Cedric Lothritz, Bei Li, Saad Ezzini, Haoye Tian, Jacques Klein, Tegawendé F. Bissyandé</i>	
TroL: Traversal of Layers for Large Language and Vision Models.....	11314
<i>Byung-Kwan Lee, Sangyun Chung, Chae Won Kim, Beomchan Park, Yong Man Ro</i>	

MMTE: Corpus and Metrics for Evaluating Machine Translation Quality of Metaphorical Language	11343
<i>Shun Wang, Ge Zhang, Han Wu, Tyler Loakman, Wenhao Huang, Chenghua Lin</i>	
Revisiting Supertagging for Faster HPSG Parsing	11359
<i>Olga Zamaraeva, Carlos Gómez-Rodríguez</i>	
Improve Dense Passage Retrieval with Entailment Tuning.....	11375
<i>Lu Dai, Hao Liu, Hui Xiong</i>	
ToolBeHonest: A Multi-Level Hallucination Diagnostic Benchmark for Tool-Augmented Large Language Models	11388
<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, Hayato Yamana</i>	
TEMA: Token Embeddings Mapping for Enriching Low-Resource Language Models	11423
<i>Rodolfo Zevallos, Níria Bel, Mireia Farrús</i>	
DECOR: Improving Coherence in L2 English Writing with a Novel Benchmark for Incoherence Detection, Reasoning, and Rewriting	11436
<i>Xuanming Zhang, Anthony Diaz, Zixun Chen, Qingyang Wu, Kun Qian, Erik Voss, Zhou Yu</i>	
Text2Chart31: Instruction Tuning for Chart Generation with Automatic Feedback.....	11459
<i>Fatemeh Pesaran Zadeh, Juyeon Kim, Jin-Hwa Kim, Gunhee Kim</i>	
PrExMe! Large Scale Prompt Exploration of Open Source LLMs for Machine Translation and Summarization Evaluation	11481
<i>Christoph Leiter, Steffen Eger</i>	
Universal Vulnerabilities in Large Language Models: Backdoor Attacks for In-Context Learning.....	11507
<i>Shuai Zhao, Meihuizi Jia, Anh Tuan Luu, Fengjun Pan, Jinming Wen</i>	
Repairs in a Block World: A New Benchmark for Handling User Corrections with Multi-Modal Language Models	11523
<i>Javier Chiyah-Garcia, Alessandro Suglia, Arash Eshghi</i>	
Beyond the Turn-Based Game: Enabling Real-Time Conversations with Duplex Models	11543
<i>Xinrong Zhang, Yingfa Chen, Shengding Hu, Xu Han, Zihang Xu, Yuanwei Xu, Weilin Zhao, Maosong Sun, Zhiyuan Liu</i>	
Strengthening Structural Inductive Biases by Pre-Training to Perform Syntactic Transformations.....	11558
<i>Matthias Lindemann, Alexander Koller, Ivan Titov</i>	
Puzzle Solving Using Reasoning of Large Language Models: A Survey	11574
<i>Panagiotis Giadikiaroglou, Maria Lymperaiou, Giorgos Filandrianos, Giorgos Stamou</i>	
SciEx: Benchmarking Large Language Models on Scientific Exams with Human Expert Grading and Automatic Grading.....	11592
<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 Stiefelwagen, Carsten Dachsbacher, Klemens Böhm, Jan Niehues</i>	
Red Teaming Language Models for Processing Contradictory Dialogues	11611
<i>Xiaofei Wen, Bangzheng Li, Tenghao Huang, Muhao Chen</i>	

Fishing for Magikarp: Automatically Detecting Under-Trained Tokens in Large Language Models	11631
<i>Sander Land, Max Bartolo</i>	
Reasoning Or a Semblance of It? a Diagnostic Study of Transitive Reasoning in LLMs	11647
<i>Houman Mehrfarin, Arash Eshghi, Ioannis Konstas</i>	
Pragmatic Norms Are All You Need – Why the Symbol Grounding Problem Does Not Apply to LLMs.....	11663
<i>Reto Gubelmann</i>	
Major Entity Identification: A Generalizable Alternative to Coreference Resolution	11679
<i>Kawshik Manikantan Sundar, Shubham Toshniwal, Makarand Tapaswi, Vineet Gandhi</i>	
Enhancing High-Order Interaction Awareness in LLM-Based Recommender Model	11696
<i>Xinfeng Wang, Jin Cui, Fumiyo Fukumoto, Yoshimi Suzuki</i>	
What Are the Odds? Language Models Are Capable of Probabilistic Reasoning	11712
<i>Akshay Paruchuri, Jake Garrison, Shun Liao, John B Hernandez, Jacob Sunshine, Tim Althoff, Xin Liu, Daniel McDuff</i>	

VOLUME 16

MARE: Multi-Aspect Rationale Extractor on Unsupervised Rationale Extraction.....	11734
<i>Han Jiang, Junwen Duan, Zhe Qu, Jianxin Wang</i>	
LoRA-Guard: Parameter-Efficient Guardrail Adaptation for Content Moderation of Large Language Models	11746
<i>Hayder Elesedy, Pedro M Esperanca, Silviu Vlad Oprea, Mete Ozay</i>	
“A Good Pun is Its Own Reword”: Can Large Language Models Understand Puns?	11766
<i>Zhijun Xu, Siyu Yuan, Lingjie Chen, Deqing Yang</i>	
QGEval: Benchmarking Multi-Dimensional Evaluation for Question Generation	11783
<i>Weiping Fu, Bifan Wei, Jianxiang Hu, Zhongmin Cai, Jun Liu</i>	
Dependency Graph Parsing as Sequence Labeling.....	11804
<i>Ana Ezquerro, David Vilares, Carlos Gómez-Rodríguez</i>	
NuNER: Entity Recognition Encoder Pre-Training Via LLM-Annotated Data	11829
<i>Sergei Bogdanov, Alexandre Constantin, Timothée Bernard, Benoit Crabbé, Etienne P Bernard</i>	
Towards a Greek Proverb Atlas: Computational Spatial Exploration and Attribution of Greek Proverbs.....	11842
<i>John Pavlopoulos, Panos Louridas, Panagiotis Filos</i>	
Unraveling Babel: Exploring Multilingual Activation Patterns of LLMs and Their Applications	11855
<i>Weize Liu, Yinlong Xu, Hongxia Xu, Jintai Chen, Xuming Hu, Jian Wu</i>	
Advancing Semantic Textual Similarity Modeling: A Regression Framework with Translated ReLU and Smooth K2 Loss	11882
<i>Bowen Zhang, Chunping Li</i>	
Rationalizing Transformer Predictions Via End-To-End Differentiable Self-Training.....	11894
<i>Marc Felix Brinner, Sina Zarriß</i>	

Segment Any Text: A Universal Approach for Robust, Efficient and Adaptable Sentence Segmentation	11908
<i>Markus Frohmann, Igor Sterner, Ivan Vulic, Benjamin Minixhofer, Markus Schedl</i>	
Applying Contrastive Learning to Code Vulnerability Type Classification	11942
<i>Chen Ji, Su Yang, Hongyu Sun, Yuqing Zhang</i>	
TheoremLlama: Transforming General-Purpose LLMs into Lean4 Experts	11953
<i>Ruida Wang, Jipeng Zhang, Yizhen Jia, Rui Pan, Shizhe Diao, Renjie Pi, Tong Zhang</i>	
Multi-Level Cross-Modal Alignment for Speech Relation Extraction	11975
<i>Liang Zhang, Zhen Yang, Biao Fu, Ziyao Lu, Liangying Shao, Shiyu Liu, Fandong Meng, Jie Zhou, Xiaoli Wang, Jinsong Su</i>	
Self-Training for Sample-Efficient Active Learning for Text Classification with Pre-Trained Language Models	11987
<i>Christopher Schröder, Gerhard Heyer</i>	
PANDA: Persona Attributes Navigation for Detecting and Alleviating Overuse Problem in Large Language Models	12005
<i>Jinsung Kim, Seonmin Koo, Heuseok Lim</i>	
The Multilingual Alignment Prism: Aligning Global and Local Preferences to Reduce Harm.....	12027
<i>Aakanksha, Arash Ahmadian, Beyza Ermis, Seraphina Goldfarb-Tarrant, Julia Kreutzer, Marzieh Fadaee, Sara Hooker</i>	
Subword Segmentation in LLMs: Looking at Inflection and Consistency	12050
<i>Marion Di Marco, Alexander Fraser</i>	
Explicit, Implicit, and Scattered: Revisiting Event Extraction to Capture Complex Arguments.....	12061
<i>Omar Sharif, Joseph Gatto, Madhusudan Basak, Sarah Masud Preum</i>	
Let Me Teach You: Pedagogical Foundations of Feedback for Language Models	12082
<i>Beatriz Borges, Niket Tandon, Tanja Käser, Antoine Bosselut</i>	
Unknown Claims: Generation of Fact-Checking Training Examples from Unstructured and Structured Data.....	12105
<i>Jean-Flavien Bussotti, Luca Ragazzi, Giacomo Frisoni, Gianluca Moro, Paolo Papotti</i>	
TL-CL: Task and Language Incremental Continual Learning	12123
<i>Shrey Satapara, P. K. Srijith</i>	
Medical Adaptation of Large Language and Vision-Language Models: Are We Making Progress?.....	12143
<i>Daniel P Jeong, Saurabh Garg, Zachary Chase Lipton, Michael Oberst</i>	
Empowering Multi-Step Reasoning Across Languages Via Program-Aided Language Models.....	12171
<i>Leonardo Ranaldi, Giulia Pucci, Barry Haddow, Alexandra Birch</i>	
Do LLMs Overcome Shortcut Learning? an Evaluation of Shortcut Challenges in Large Language Models.....	12188
<i>Yu Yuan, Lili Zhao, Kai Zhang, Guangting Zheng, Qi Liu</i>	
ControlMath: Controllable Data Generation Promotes Math Generalist Models.....	12201
<i>Nuo Chen, Ning Wu, Jianhui Chang, Linjun Shou, Jia Li</i>	
Where Am I From? Identifying Origin of LLM-Generated Content	12218
<i>Liyang Li, Yihan Bai, Minhao Cheng</i>	

ReadMe++: Benchmarking Multilingual Language Models for Multi-Domain Readability Assessment	12230
<i>Tarek Naous, Michael J Ryan, Anton Lavrouk, Mohit Chandra, Wei Xu</i>	
GlossLM: A Massively Multilingual Corpus and Pretrained Model for Interlinear Glossed Text.....	12267
<i>Michael Ginn, Lindia Tjuatja, Taiqi He, Enora Rice, Graham Neubig, Alexis Palmer, Lori Levin</i>	
GDTB: Genre Diverse Data for English Shallow Discourse Parsing Across Modalities, Text Types, and Domains.....	12287
<i>Yang Janet Liu, Tatsuya Aoyama, Wesley Scivetti, Yilun Zhu, Shabnam Behzad, Lauren Elizabeth Levine, Jessica Lin, Devika Tiwari, Amir Zeldes</i>	
RA2FD: Distilling Faithfulness into Efficient Dialogue Systems	12304
<i>Zhiyuan Zhu, Yusheng Liao, Chenxin Xu, Yunfeng Guan, Yanfeng Wang, Yu Wang</i>	
Subjective Topic Meets LLMs: Unleashing Comprehensive, Reflective and Creative Thinking Through the Negation of Negation	12318
<i>Fangrui Lv, Kaixiong Gong, Jian Liang, Xinyu Pang, Changshui Zhang</i>	
Experimental Contexts Can Facilitate Robust Semantic Property Inference in Language Models, but Inconsistently	12342
<i>Kanishka Misra, Allyson Ettinger, Kyle Mahowald</i>	
Leveraging Estimated Transferability Over Human Intuition for Model Selection in Text Ranking	12356
<i>Jun Bai, Zhuofan Chen, Zhenzi Li, Hanhua Hong, Jianfei Zhang, Chen Li, Chenghua Lin, Wenge Rong</i>	
Unveiling In-Context Learning: A Coordinate System to Understand Its Working Mechanism	12375
<i>Anhao Zhao, Fanghua Ye, Jinlan Fu, Xiaoyu Shen</i>	
Self-Powered LLM Modality Expansion for Large Speech-Text Models	12401
<i>Tengfei Yu, Xuebo Liu, Zhiyi Hou, Liang Ding, Dacheng Tao, Min Zhang</i>	
ABSEval: An Agent-Based Framework for Script Evaluation.....	12418
<i>Sirui Liang, Baoli Zhang, Jun Zhao, Kang Liu</i>	
Latent Concept-Based Explanation of NLP Models.....	12435
<i>Xuemin Yu, Fahim Dalvi, Nadir Durrani, Marzia Nouri, Hassan Sajjad</i>	
Decoding with Limited Teacher Supervision Requires Understanding When to Trust the Teacher	12460
<i>Hyunjong Ok, Jegwang Ryu, Jaeho Lee</i>	
Enhancing Data Quality Through Simple De-Duplication: Navigating Responsible Computational Social Science Research	12477
<i>Yida Mu, Mali Jin, Xingyi Song, Nikolaos Aletras</i>	
The Mystery of the Pathological Path-Star Task for Language Models	12493
<i>Arvid Frydenlund</i>	
Voices in a Crowd: Searching for Clusters of Unique Perspectives	12517
<i>Nikolas Vitsakis, Amit Parekh, Ioannis Konstas</i>	

VOLUME 17

Neeko: Leveraging Dynamic LoRA for Efficient Multi-Character Role-Playing Agent.....	12540
<i>Xiaoyan Yu, Tongxu Luo, Yifan Wei, Fangyu Lei, Yiming Huang, Hao Peng, Liehuang Zhu</i>	
SLANG: New Concept Comprehension of Large Language Models.....	12558
<i>Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Xueqi Cheng</i>	
Towards Interpretable Sequence Continuation: Analyzing Shared Circuits in Large Language Models.....	12576
<i>Michael Lan, Philip Torr, Fazl Barez</i>	
Why Does New Knowledge Create Messy Ripple Effects in LLMs?	12602
<i>Jiaxin Qin, Zixuan Zhang, Chi Han, Pengfei Yu, Manling Li, Heng Ji</i>	
Lifelong Event Detection Via Optimal Transport	12610
<i>Viet Dao, Van-Cuong Pham, Quyen Tran, Thanh-Thien Le, Linh Van Ngo, Thien Huu Nguyen</i>	
SUPER: Evaluating Agents on Setting Up and Executing Tasks from Research Repositories	12622
<i>Ben Bogin, Kejuan Yang, Shashank Gupta, Kyle Richardson, Erin Bransom, Peter Clark, Ashish Sabharwal, Tushar Khot</i>	
FIRST: Teach a Reliable Large Language Model Through Efficient Trustworthy Distillation	12646
<i>Kashun Shum, Minrui Xu, Jianshu Zhang, Zixin Chen, Shizhe Diao, Hanze Dong, Jipeng Zhang, Muhammad Omer Raza</i>	
Domain Adapted Machine Translation: What Does Catastrophic Forgetting Forget and Why?	12660
<i>Danielle Saunders, Steve Deneefe</i>	
Enhancing AI Assisted Writing with One-Shot Implicit Negative Feedback	12672
<i>Benjamin Towle, Ke Zhou</i>	
Atomic Self-Consistency for Better Long Form Generations	12681
<i>Raghuv eer Thirukovalluru, Yukun Huang, Bhuwan Dhingra</i>	
“Global is Good, Local is Bad?”: Understanding Brand Bias in LLMs	12695
<i>Mahammed Kamruzzaman, Hieu Minh Nguyen, Gene Louis Kim</i>	
Optimizing Rare Word Accuracy in Direct Speech Translation with a Retrieval-And-Demonstration Approach	12703
<i>Siqi Li, Danni Liu, Jan Niehues</i>	
ACE: A LLM-Based Negotiation Coaching System	12720
<i>Ryan Shea, Aymen Kallala, Xin Lucy Liu, Michael W. Morris, Zhou Yu</i>	
TransferTOD: A Generalizable Chinese Multi-Domain Task-Oriented Dialogue System with Transfer Capabilities.....	12750
<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, Xuanjing Huang</i>	
PATIENT- ψ : Using Large Language Models to Simulate Patients for Training Mental Health Professionals.....	12772
<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, Zhiyu Chen</i>	

DKEC: Domain Knowledge Enhanced Multi-Label Classification for Diagnosis Prediction	12798
<i>Xueren Ge, Abhishek Satpathy, Ronald Dean Williams, John Stankovic, Homa Alemzadeh</i>	
ModSCAN: Measuring Stereotypical Bias in Large Vision-Language Models from Vision and Language Modalities	12814
<i>Yukun Jiang, Zheng Li, Xinyue Shen, Yugeng Liu, Michael Backes, Yang Zhang</i>	
Large Language Models Can Self-Correct with Key Condition Verification	12846
<i>Zhenyu Wu, Qingkai Zeng, Zhihan Zhang, Zhaoxuan Tan, Chao Shen, Meng Jiang</i>	
Learning to Write Rationally: How Information is Distributed in Non-Native Speakers' Essays	12868
<i>Zixin Tang, Janet Van Hell</i>	
Defending Against Social Engineering Attacks in the Age of LLMs	12880
<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, Julia Hirschberg</i>	
Heterogeneous LoRA for Federated Fine-Tuning of On-Device Foundation Models	12903
<i>Yae Jee Cho, Luyang Liu, Zheng Xu, Aldi Fahrezi, Gauri Joshi</i>	
Make Some Noise: Unlocking Language Model Parallel Inference Capability Through Noisy Training	12914
<i>Yixuan Wang, Xianzhen Luo, Fuxuan Wei, Yijun Liu, Qingfu Zhu, Xuanyu Zhang, Qing Yang, Dongliang Xu, Wanxiang Che</i>	
Target-Aware Language Modeling Via Granular Data Sampling	12927
<i>Ernie Chang, Pin-Jie Lin, Yang Li, Changsheng Zhao, Daeil Kim, Rastislav Rabatin, Zechun Liu, Yangyang Shi, Vikas Chandra</i>	
SPEED++: A Multilingual Event Extraction Framework for Epidemic Prediction and Preparedness	12936
<i>Tanmay Parekh, Jeffrey Kwan, Jiarui Yu, Sparsh Johri, Hyosang Ahn, Sreya Muppalla, Kai-Wei Chang, Wei Wang, Nanyun Peng</i>	
CoGen: Learning from Feedback with Coupled Comprehension and Generation	12966
<i>Mustafa Omer Gul, Yoav Artzi</i>	
UNICORN: A Unified Causal Video-Oriented Language-Modeling Framework for Temporal Video-Language Tasks	12983
<i>Yuanhao Xiong, Yixin Nie, Haotian Liu, Boxin Wang, Jun Chen, Rong Jin, Cho-Jui Hsieh, Lorenzo Torresani, Jie Lei</i>	
Story Morals: Surfacing Value-Driven Narrative Schemas Using Large Language Models	12998
<i>David G Hobson, Haiqi Zhou, Derek Ruths, Andrew Piper</i>	
OATH-Frames: Characterizing Online Attitudes Towards Homelessness with LLM Assistants	13033
<i>Jaspreet Ranjit, Brihi Joshi, Rebecca Dorn, Laura Petry, Olga Koumoundouros, Jayne Bottarini, Peichen Liu, Eric Rice, Swabha Swayamdipta</i>	
AnaloBench: Benchmarking the Identification of Abstract and Long-Context Analogies	13060
<i>Xiao Ye, Andrew Wang, Jacob Choi, Yining Lu, Shreya Sharma, Lingfeng Shen, Vijay Murari Tiyyala, Nicholas Andrews, Daniel Khashabi</i>	
SciER: An Entity and Relation Extraction Dataset for Datasets, Methods, and Tasks in Scientific Documents	13083
<i>Qi Zhang, Zhijia Chen, Huitong Pan, Cornelia Caragea, Longin Jan Latecki, Eduard Dragut</i>	

Analysis of Plan-Based Retrieval for Grounded Text Generation	13101
<i>Ameya Godbole, Nicholas Monath, Seungyeon Kim, Ankit Singh Rawat, Andrew McCallum, Manzil Zaheer</i>	
Detecting Errors Through Ensembling Prompts (DEEP): An End-To-End LLM Framework for Detecting Factual Errors	13120
<i>Alex Chandler, Devesh Surve, Hui Su</i>	
RLHF Can Speak Many Languages: Unlocking Multilingual Preference Optimization for LLMs	13134
<i>John Dang, Arash Ahmadian, Kelly Marchisio, Julia Kreutzer, Ahmet Üstün, Sara Hooker</i>	
Boosting Logical Fallacy Reasoning in LLMs Via Logical Structure Tree	13157
<i>Yuanyuan Lei, Ruihong Huang</i>	
Chain and Causal Attention for Efficient Entity Tracking	13174
<i>Erwan Fagnou, Paul Caillon, Blaise Delattre, Alexandre Allauzen</i>	
BEEAR: Embedding-Based Adversarial Removal of Safety Backdoors in Instruction-Tuned Language Models	13189
<i>Yi Zeng, Weiyu Sun, Tran Huynh, Dawn Song, Bo Li, Ruoxi Jia</i>	
A Bayesian Approach to Harnessing the Power of LLMs in Authorship Attribution	13216
<i>Zhengmian Hu, Tong Zheng, Heng Huang</i>	
FAC2E: Better Understanding Large Language Model Capabilities by Dissociating Language and Cognition	13228
<i>Xiaoqiang Wang, Lingfei Wu, Tengfei Ma, Bang Liu</i>	
OpenSep: Leveraging Large Language Models with Textual Inversion for Open World Audio Separation	13244
<i>Tanvir Mahmud, Diana Marculescu</i>	
Language Concept Erasure for Language-Invariant Dense Retrieval	13261
<i>Zhiqi Huang, Puxuan Yu, Shauli Ravfogel, James Allan</i>	
Learning Personalized Alignment for Evaluating Open-Ended Text Generation	13274
<i>Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li, Yuandong Tian</i>	
Large Language Models Are Involuntary Truth-Tellers: Exploiting Fallacy Failure for Jailbreak Attacks	13293
<i>Yue Zhou, Henry Peng Zou, Barbara Di Eugenio, Yang Zhang</i>	
Turn Waste into Worth: Rectifying Top-K Router of MoE	13305
<i>Zhiyuan Zeng, Qipeng Guo, Zhaoye Fei, Zhangyue Yin, Yunhua Zhou, Linyang Li, Tianxiang Sun, Hang Yan, Dahua Lin, Xipeng Qiu</i>	

VOLUME 18

Null-Shot Prompting: Rethinking Prompting Large Language Models with Hallucination	13321
<i>Pittawat Taveekitworachai, Febri Abdullah, Ruck Thawonmas</i>	
CommVQA: Situating Visual Question Answering in Communicative Contexts	13362
<i>Nandita Shankar Naik, Christopher Potts, Elisa Kreiss</i>	

Ouroboros: Generating Longer Drafts Phrase by Phrase for Faster Speculative Decoding	13378
<i>Weilin Zhao, Yuxiang Huang, Xu Han, Wang Xu, Chaojun Xiao, Xinrong Zhang, Yewei Fang, Kaihuo Zhang, Zhiyuan Liu, Maosong Sun</i>	
1+1>2: Can Large Language Models Serve as Cross-Lingual Knowledge Aggregators?.....	13394
<i>Yue Huang, Chenrui Fan, Yuan Li, Siyuan Wu, Tianyi Zhou, Xiangliang Zhang, Lichao Sun</i>	
How to Leverage Demonstration Data in Alignment for Large Language Model? a Self-Imitation Learning Perspective	13413
<i>Teng Xiao, Mingxiao Li, Yige Yuan, Huaisheng Zhu, Chao Cui, Vasant G Honavar</i>	
Style-Specific Neurons for Steering LLMs in Text Style Transfer.....	13427
<i>Wen Lai, Viktor Hangya, Alexander Fraser</i>	
Adaptive Query Rewriting: Aligning Rewriters Through Marginal Probability of Conversational Answers.....	13444
<i>Tianhua Zhang, Kun Li, Hongyin Luo, Xixin Wu, James R. Glass, Helen M. Meng</i>	
Grasping the Essentials: Tailoring Large Language Models for Zero-Shot Relation Extraction.....	13462
<i>Sizhe Zhou, Yu Meng, Bowen Jin, Jiawei Han</i>	
DA-Code: Agent Data Science Code Generation Benchmark for Large Language Models	13487
<i>Yiming Huang, Jianwen Luo, Yan Yu, Yitong Zhang, Fangyu Lei, Yifan Wei, Shizhu He, Lifu Huang, Xiao Liu, Jun Zhao, Kang Liu</i>	
Leveraging Context-Aware Prompting for Commit Message Generation.....	13522
<i>Zhihua Jiang, Jianwei Chen, Dongning Rao, Guanghui Ye</i>	
Linguistic Bias in ChatGPT: Language Models Reinforce Dialect Discrimination	13541
<i>Eve Fleisig, Genevieve Smith, Madeline Bossi, Ishita Rustagi, Xavier Yin, Dan Klein</i>	
Lifelong Knowledge Editing for LLMs with Retrieval-Augmented Continuous Prompt Learning.....	13565
<i>Qizhou Chen, Taolin Zhang, Xiaofeng He, Dongyang Li, Chengyu Wang, Longtao Huang, Hui Xue'</i>	
A Learning Rate Path Switching Training Paradigm for Version Updates of Large Language Models.....	13581
<i>Zhihao Wang, Shiyu Liu, Jianheng Huang, Wang Zheng, Yixuan Liao, Xiaoxin Chen, Junfeng Yao, Jinsong Su</i>	
Zero-Shot Cross-Lingual NER Using Phonemic Representations for Low-Resource Languages	13595
<i>Jimin Sohn, Haeji Jung, Alex Cheng, Joeeon Kang, Yilin Du, David R Mortensen</i>	
An Analysis and Mitigation of the Reversal Curse.....	13603
<i>Ang Lv, Kaiyi Zhang, Shufang Xie, Quan Tu, Yuhan Chen, Ji-Rong Wen, Rui Yan</i>	
Exploring the Practicality of Generative Retrieval on Dynamic Corpora	13616
<i>Chaeun Kim, Soyoung Yoon, Hyunji Lee, Joel Jang, Sohee Yang, Minjoon Seo</i>	
OneNet: A Fine-Tuning Free Framework for Few-Shot Entity Linking Via Large Language Model Prompting	13634
<i>Xukai Liu, Ye Liu, Kai Zhang, Kehang Wang, Qi Liu, Enhong Chen</i>	
Don't Just Say "I Don't Know"! Self-Aligning Large Language Models for Responding to Unknown Questions with Explanations.....	13652
<i>Yang Deng, Yong Zhao, Moxin Li, See-Kiong Ng, Tat-Seng Chua</i>	

Fewer is More: Boosting Math Reasoning with Reinforced Context Pruning	13674
<i>Xijie Huang, Li Lyna Zhang, Kwang-Ting Cheng, Fan Yang, Mao Yang</i>	
Large Language Models Are Poor Clinical Decision-Makers: A Comprehensive Benchmark	13696
<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, David A. Clifton</i>	
Holistic Automated Red Teaming for Large Language Models Through Top-Down Test Case Generation and Multi-Turn Interaction.....	13711
<i>Jinchuan Zhang, Yan Zhou, Yaxin Liu, Ziming Li, Songlin Hu</i>	
Householder Pseudo-Rotation: A Novel Approach to Activation Editing in LLMs with Direction-Magnitude Perspective	13737
<i>Van-Cuong Pham, Thien Huu Nguyen</i>	
DynamicER: Resolving Emerging Mentions to Dynamic Entities for RAG.....	13752
<i>Jinyoung Kim, Dayoon Ko, Gunhee Kim</i>	
Preserving Generalization of Language Models in Few-Shot Continual Relation Extraction.....	13771
<i>Quyên Tran, Nguyen Xuan Thanh, Nguyen Hoang Anh, Nam Le Hai, Trung Le, Linh Van Ngo, Thien Huu Nguyen</i>	
A Systematic Survey and Critical Review on Evaluating Large Language Models: Challenges, Limitations, and Recommendations	13785
<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, Jimmy Huang</i>	
Consecutive Batch Model Editing with HooK Layers	13817
<i>Shuaiyi Li, Yang Deng, Deng Cai, Hongyuan Lu, Liang Chen, Wai Lam</i>	
Topic-Oriented Open Relation Extraction with a Priori Seed Generation.....	13834
<i>Linyi Ding, Jinfeng Xiao, Sizhe Zhou, Chaoqi Yang, Jiawei Han</i>	
Related Work and Citation Text Generation: A Survey	13846
<i>Xiangci Li, Jessica Ouyang</i>	
Curriculum Consistency Learning for Conditional Sentence Generation.....	13865
<i>Liangxin Liu, Xuebo Liu, Lian Lian, Shengjun Cheng, Jun Rao, Tengfei Yu, Hexuan Deng, Min Zhang</i>	
A Systematic Analysis of Large Language Models as Soft Reasoners: The Case of Syllogistic Inferences	13882
<i>Leonardo Bertolazzi, Albert Gatt, Raffaella Bernardi</i>	
Pre-Training Cross-Lingual Open Domain Question Answering with Large-Scale Synthetic Supervision.....	13906
<i>Fan Jiang, Tom Drummond, Trevor Cohn</i>	
MOSEL: 950,000 Hours of Speech Data for Open-Source Speech Foundation Model Training on EU Languages	13934
<i>Marco Gaido, Sara Papi, Luisa Bentivogli, Alessio Brutti, Mauro Cettolo, Roberto Gretter, Marco Matassoni, Mohamed Nabih, Matteo Negri</i>	
Improving Knowledge Graph Completion with Structure-Aware Supervised Contrastive Learning.....	13948
<i>Jiashi Lin, Lifang Wang, Xinyu Lu, Zhongtian Hu, Wei Zhang, Wenxuan Lu</i>	

Contribution of Linguistic Typology to Universal Dependency Parsing: An Empirical Investigation.....	13960
<i>Ali Basirat, Navid Baradaran Hemmati</i>	
TRoTR: A Framework for Evaluating the Re-Contextualization of Text Reuse	13972
<i>Francesco Periti, Pierluigi Cassotti, Stefano Montanelli, Nina Tahmasebi, Dominik Schlechtweg</i>	
Structured Optimal Brain Pruning for Large Language Models.....	13991
<i>Jiateng Wei, Quan Lu, Ning Jiang, Siqi Li, Jingyang Xiang, Jun Chen, Yong Liu</i>	
Automatically Generated Definitions and Their Utility for Modeling Word Meaning.....	14008
<i>Francesco Periti, David Alfier, Nina Tahmasebi</i>	
How Do Your Code LLMs Perform? Empowering Code Instruction Tuning with Really Good Data.....	14027
<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, Weiran Xu</i>	
MAIR: A Massive Benchmark for Evaluating Instructed Retrieval	14044
<i>Weiwei Sun, Zhengliang Shi, Wu Jiu Long, Lingyong Yan, Xinyu Ma, Yiding Liu, Min Cao, Dawei Yin, Zhaochun Ren</i>	
Rethinking the Evaluation of In-Context Learning for LLMs.....	14068
<i>Guoxin Yu, Lemao Liu, Mo Yu, Yue Yu, Xiang Ao</i>	
Cluster-Norm for Unsupervised Probing of Knowledge	14083
<i>Walter Laurito, Sharan Maiya, Grégoire Dhimoïla, Owen Ho Wan Yeung, Kaarel Hänni</i>	
Hopping Too Late: Exploring the Limitations of Large Language Models on Multi-Hop Queries	14113
<i>Eden Biran, Daniela Gottesman, Sohee Yang, Mor Geva, Amir Globerson</i>	

VOLUME 19

Enhancing Training Data Attribution for Large Language Models with Fitting Error Consideration.....	14131
<i>Kangxi Wu, Liang Pang, Huawei Shen, Xueqi Cheng</i>	
Where Am I? Large Language Models Wandering Between Semantics and Structures in Long Contexts.....	14144
<i>Seonmin Koo, Jinsung Kim, Youngjoon Jang, Chanjun Park, Heuseok Lim</i>	
KARL: Knowledge-Aware Retrieval and Representations Aid Retention and Learning in Students.....	14161
<i>Matthew Shu, Nishant Balepur, Shi Feng, Jordan Lee Boyd-Graber</i>	
Large Language Models Can Be Contextual Privacy Protection Learners.....	14179
<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, Wei Cheng</i>	
A SMART Mnemonic Sounds Like “Glue Tonic”: Mixing LLMs with Student Feedback to Make Mnemonic Learning Stick	14202
<i>Nishant Balepur, Matthew Shu, Alexander Hoyle, Alison Robey, Shi Feng, Seraphina Goldfarb-Tarrant, Jordan Lee Boyd-Graber</i>	
Mixture-Of-Skills: Learning to Optimize Data Usage for Fine-Tuning Large Language Models	14226
<i>Minghao Wu, Thuy-Trang Vu, Lizhen Qu, Reza Haf</i>	

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

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

BiasAlert: A Plug-And-Play Tool for Social Bias Detection in LLMs	14778
<i>Zhiting Fan, Ruizhe Chen, Ruiling Xu, Zuozhu Liu</i>	
VHASR: A Multimodal Speech Recognition System with Vision Hotwords.....	14791
<i>Jiliang Hu, Zuchao Li, Ping Wang, Haojun Ai, Lefei Zhang, Hai Zhao</i>	
A Probability–Quality Trade-Off in Aligned Language Models and Its Relation to Sampling Adaptors	14805
<i>Naaman Tan, Josef Valvoda, Tianyu Liu, Anej Svete, Yanxia Qin, Min-Yen Kan, Ryan Cotterell</i>	
Bridging Local Details and Global Context in Text-Attributed Graphs	14830
<i>Yaoke Wang, Yun Zhu, Wenqiao Zhang, Yueting Zhuang, Liyunfei Liyunfei, Siliang Tang</i>	
Building Resources for Emakhuwa: Machine Translation and News Classification Benchmarks	14842
<i>Felermino D. M. A. Ali, Henrique Lopes Cardoso, Rui Sousa-Silva</i>	
RepMatch: Quantifying Cross-Instance Similarities in Representation Space.....	14858
<i>Mohammad Reza Modarres, Sina Abbasi, Mohammad Taher Pilehvar</i>	
Commonsense Knowledge Editing Based on Free-Text in LLMs	14870
<i>Xiusheng Huang, Yequan Wang, Jun Zhao, Kang Liu</i>	
A Closer Look at Multidimensional Online Political Incivility.....	14881
<i>Sagi Pendzel, Nir Lotan, Alon Zoizner, Einat Minkov</i>	
Leveraging BERT and TFIDF Features for Short Text Clustering Via Alignment-Promoting Co- Training	14897
<i>Zetong Li, Qinliang Su, Shijing Si, Jianxing Yu</i>	
Applying Intrinsic Debiasing on Downstream Tasks: Challenges and Considerations for Machine Translation.....	14914
<i>Bar Iluz, Yanai Elazar, Asaf Yehudai, Gabriel Stanovsky</i>	
Unsupervised Named Entity Disambiguation for Low Resource Domains.....	14922
<i>Debarghya Datta, Soumajit Pramanik</i>	

VOLUME 20

SparseGrad: A Selective Method for Efficient Fine-Tuning of MLP Layers.....	14929
<i>Viktoria A. Chekalina, Anna Rudenko, Gleb Mezentsev, Aleksandr Mikhalev, Alexander Panchenko, Ivan Oseledets</i>	
MoCoKGC: Momentum Contrast Entity Encoding for Knowledge Graph Completion.....	14940
<i>Qingyang Li, Yanru Zhong, Yuchu Qin</i>	
ActPlan-1K: Benchmarking the Procedural Planning Ability of Visual Language Models in Household Activities	14953
<i>Ying Su, Zhan Ling, Haochen Shi, Cheng Jiayang, Yauwai Yim, Yangqiu Song</i>	
Shortcuts Arising from Contrast: Towards Effective and Lightweight Clean-Label Attacks in Prompt-Based Learning.....	14966
<i>Xiaopeng Xie, Ming Yan, Xiwen Zhou, Chenlong Zhao, Suli Wang, Yong Zhang, Joey Tianyi Zhou</i>	
GRASS: Compute Efficient Low-Memory LLM Training with Structured Sparse Gradients.....	14978
<i>Aashiq Muhamed, Oscar Li, David Woodruff, Mona T. Diab, Virginia Smith</i>	

RaTEScore: A Metric for Radiology Report Generation.....	15004
<i>Weike Zhao, Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, Weidi Xie</i>	
HalluMeasure: Fine-Grained Hallucination Measurement Using Chain-Of-Thought Reasoning.....	15020
<i>Shayan Ali Akbar, Md Mosharaf Hossain, Tess Wood, Si-Chi Chin, Erica M Salinas, Victor Alvarez, Erwin Cornejo</i>	
Learning to Rank Salient Content for Query-Focused Summarization.....	15038
<i>Sajad Sotudeh, Nazli Goharian</i>	
Are Large Language Models Good Classifiers? a Study on Edit Intent Classification in Scientific Document Revisions.....	15049
<i>Qian Ruan, Ilia Kuznetsov, Iryna Gurevych</i>	
LitSearch: A Retrieval Benchmark for Scientific Literature Search.....	15068
<i>Anirudh Ajith, Mengzhou Xia, Alexis Chevalier, Tanya Goyal, Danqi Chen, Tianyu Gao</i>	
Open-World Multi-Label Text Classification with Extremely Weak Supervision.....	15084
<i>Xintong Li, Jinya Jiang, Ria Dharmani, Jayanth Srinivasa, Gaowen Liu, Jingbo Shang</i>	
LLMs Learn Governing Principles of Dynamical Systems, Revealing an In-Context Neural Scaling Law.....	15097
<i>Toni J. B. Liu, Nicolas Boulle, Raphaël Sarfati, Christopher Earls</i>	
AKEW: Assessing Knowledge Editing in the Wild.....	15118
<i>Xiaobao Wu, Liangming Pan, William Yang Wang, Anh Tuan Luu</i>	
CopyBench: Measuring Literal and Non-Literal Reproduction of Copyright-Protected Text in Language Model Generation.....	15134
<i>Tong Chen, Akari Asai, Niloofar Miresghallah, Sewon Min, James Grimmermann, Yejin Choi, Hannaneh Hajishirzi, Luke Zettlemoyer, Pang Wei Koh</i>	
Dense X Retrieval: What Retrieval Granularity Should We Use?.....	15159
<i>Tong Chen, Hongwei Wang, Sihao Chen, Wenhao Yu, Kaixin Ma, Xinran Zhao, Hongming Zhang, Dong Yu</i>	
Decoding Susceptibility: Modeling Misbelief to Misinformation Through a Computational Approach.....	15178
<i>Yanchen Liu, Mingyu Derek Ma, Wenna Qin, Azure Zhou, Jiaao Chen, Weiyan Shi, Wei Wang, Diyi Yang</i>	
Layer by Layer: Uncovering Where Multi-Task Learning Happens in Instruction-Tuned Large Language Models.....	15195
<i>Zheng Zhao, Yftah Ziser, Shay B Cohen</i>	
XDetox: Text Detoxification with Token-Level Toxicity Explanations.....	15215
<i>Beomseok Lee, Hyunwoo Kim, Keon Kim, Yong Suk Choi</i>	
Optimizing Chinese Lexical Simplification Across Word Types: A Hybrid Approach.....	15227
<i>Zihao Xiao, Jiefu Gong, Shijin Wang, Wei Song</i>	
Control Large Language Models Via Divide and Conquer.....	15240
<i>Bingxuan Li, Yiwei Wang, Tao Meng, Kai-Wei Chang, Nanyun Peng</i>	
Joint Pre-Encoding Representation and Structure Embedding for Efficient and Low-Resource Knowledge Graph Completion.....	15257
<i>Chenyu Qiu, Pengjiang Qian, Chuang Wang, Jian Yao, Li Liu, Fang Wei, Eddie Y. K. Eddie</i>	

Improving Discriminative Capability of Reward Models in RLHF Using Contrastive Learning	15270
<i>Lu Chen, Rui Zheng, Binghai Wang, Senjie Jin, Caishuang Huang, Junjie Ye, Zhihao Zhang, Yuhao Zhou, Zhiheng Xi, Tao Gui, Qi Zhang, Xuanjing Huang</i>	
RoCEL: Advancing Table Entity Linking Through Distinctive Row and Column Contexts	15284
<i>Yuanzheng Wang, Yixing Fan, Jiafeng Guo, Ruqing Zhang, Xueqi Cheng</i>	
Exploring the Role of Reasoning Structures for Constructing Proofs in Multi-Step Natural Language Reasoning with Large Language Models.....	15299
<i>Zi'Ou Zheng, Christopher Malon, Martin Renqiang Min, Xiaodan Zhu</i>	
Efficient Overshadowed Entity Disambiguation by Mitigating Shortcut Learning.....	15313
<i>Panuthep Tasawong, Peerat Limkonchotiwat, Potsawee Manakul, Can Udomcharoenchaikit, Ekapol Chuangsuwanich, Sarana Nutanong</i>	
AppBench: Planning of Multiple APIs from Various APPs for Complex User Instruction.....	15322
<i>Hongru Wang, Rui Wang, Boyang Xue, Heming Xia, Jingtao Cao, Zeming Liu, Jeff Z. Pan, Kam-Fai Wong</i>	
Not Everything is All You Need: Toward Low-Redundant Optimization for Large Language Model Alignment.....	15337
<i>Zhipeng Chen, Kun Zhou, Xin Zhao, Jingyuan Wang, Ji-Rong Wen</i>	
AudioVSR: Enhancing Video Speech Recognition with Audio Data.....	15352
<i>Xiaoda Yang, Xize Cheng, Jiaqi Duan, Hongshun Qiu, Minjie Hong, Minghui Fang, Shengpeng Ji, Jialong Zuo, Zhiqing Hong, Zhimeng Zhang, Tao Jin</i>	
ECCO: Can We Improve Model-Generated Code Efficiency Without Sacrificing Functional Correctness?	15362
<i>Siddhant Waghjale, Vishruth Veerendranath, Zhiruo Wang, Daniel Fried</i>	
Ladder: A Model-Agnostic Framework Boosting LLM-Based Machine Translation to the Next Level.....	15377
<i>Zhaopeng Feng, Ruizhe Chen, Yan Zhang, Zijie Meng, Zuo Zhu Liu</i>	
Re-ReST: Reflection-Reinforced Self-Training for Language Agents	15394
<i>Zi-Yi Dou, Cheng-Fu Yang, Xueqing Wu, Kai-Wei Chang, Nanyun Peng</i>	
Effective Synthetic Data and Test-Time Adaptation for OCR Correction	15412
<i>Shuhao Guan, Cheng Xu, Moule Lin, Derek Greene</i>	
SRF: Enhancing Document-Level Relation Extraction with a Novel Secondary Reasoning Framework.....	15426
<i>Fu Zhang, Qi Miao, Jingwei Cheng, Hongsen Yu, Yi Yan, Xin Li, Yongxue Wu</i>	
FineCops-Ref: A New Dataset and Task for Fine-Grained Compositional Referring Expression Comprehension.....	15440
<i>Junzhuo Liu, Xuzheng Yang, Weiwei Li, Peng Wang</i>	
Exploring the Learning Capabilities of Language Models Using LEVERWORLDS	15458
<i>Eitan Wagner, Amir Feder, Omri Abend</i>	
CONTESTS: A Framework for Consistency Testing of Span Probabilities in Language Models.....	15469
<i>Eitan Wagner, Yuli Slavutsky, Omri Abend</i>	

DocEdit-V2: Document Structure Editing Via Multimodal LLM Grounding	15485
<i>Manan Suri, Puneet Mathur, Franck Deroncourt, Rajiv Jain, Vlad I Morariu, Ramit Sawhney, Preslav Nakov, Dinesh Manocha</i>	
DogeRM: Equipping Reward Models with Domain Knowledge Through Model Merging.....	15506
<i>Tzu-Han Lin, Chen-An Li, Hung-Yi Lee, Yun-Nung Chen</i>	
Understanding Slang with LLMs: Modelling Cross-Cultural Nuances Through Paraphrasing.....	15525
<i>Ifeoluwa Wuraola, Nina Dethlefs, Daniel Marciniak</i>	
Unlocking Anticipatory Text Generation: A Constrained Approach for Large Language Models Decoding	15532
<i>Lifu Tu, Semih Yavuz, Jin Qu, Jiacheng Xu, Rui Meng, Caiming Xiong, Yingbo Zhou</i>	
Re-Reading Improves Reasoning in Large Language Models	15549
<i>Xiaohan Xu, Chongyang Tao, Tao Shen, Can Xu, Hongbo Xu, Guodong Long, Jian-Guang Lou, Shuai Ma</i>	
Adaptive Axes: A Pipeline for In-Domain Social Stereotype Analysis	15576
<i>Qingcheng Zeng, Mingyu Jin, Rob Voigt</i>	
ERVQA: A Dataset to Benchmark the Readiness of Large Vision Language Models in Hospital Environments.....	15594
<i>Sourjyadip Ray, Kushal Gupta, Soumi Kundu, Dr Payal Arvind Kasat, Somak Aditya, Pawan Goyal</i>	
Human-LLM Hybrid Text Answer Aggregation for Crowd Annotations	15609
<i>Jiyi Li</i>	
Improve Student’s Reasoning Generalizability Through Cascading Decomposed CoTs Distillation	15623
<i>Chengwei Dai, Kun Li, Wei Zhou, Songlin Hu</i>	
Revisiting Supervised Contrastive Learning for Microblog Classification	15644
<i>Junbo Huang, Ricardo Usbeck</i>	
BaitAttack: Alleviating Intention Shift in Jailbreak Attacks Via Adaptive Bait Crafting.....	15654
<i>Rui Pu, Chaozhuo Li, Rui Ha, Litian Zhang, Lirong Qiu, Xi Zhang</i>	
Images Speak Louder than Words: Understanding and Mitigating Bias in Vision-Language Model from a Causal Mediation Perspective	15669
<i>Zhaotian Weng, Zijun Gao, Jerone Andrews, Jieyu Zhao</i>	
Mitigating the Language Mismatch and Repetition Issues in LLM-Based Machine Translation Via Model Editing.....	15681
<i>Weichuan Wang, Zhaoyi Li, Defu Lian, Chen Ma, Linqi Song, Ying Wei</i>	

VOLUME 21

SciAgent: Tool-Augmented Language Models for Scientific Reasoning.....	15701
<i>Yubo Ma, Zhibin Gou, Junheng Hao, Ruochen Xu, Shuohang Wang, Liangming Pan, Yujiu Yang, Yixin Cao, Aixin Sun</i>	
Global Reward to Local Rewards: Multimodal-Guided Decomposition for Improving Dialogue Agents.....	15737
<i>Dong Won Lee, Hae Won Park, Yoon Kim, Cynthia Breazeal, Louis-Philippe Morency</i>	

Towards Measuring and Modeling “Culture” in LLMs: A Survey	15763
<i>Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Shivdutt Singh, Alham Fikri Aji, Jacki O’Neill, Ashutosh Modi, Monojit Choudhury</i>	
ESC-Eval: Evaluating Emotion Support Conversations in Large Language Models	15785
<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, Yingchun Wang</i>	
Cultural Conditioning Or Placebo? on the Effectiveness of Socio-Demographic Prompting	15811
<i>Sagnik Mukherjee, Muhammad Farid Adilazuarda, Sunayana Sitaram, Kalika Bali, Alham Fikri Aji, Monojit Choudhury</i>	
Text Fluoroscopy: Detecting LLM-Generated Text Through Intrinsic Features	15838
<i>Xiao Yu, Kejiang Chen, Qi Yang, Weiming Zhang, Nenghai Yu</i>	
Hate Personified: Investigating the Role of LLMs in Content Moderation.....	15847
<i>Sarah Masud, Sahajpreet Singh, Viktor Hangya, Alexander Fraser, Tanmoy Chakraborty</i>	
Temporally Consistent Factuality Probing for Large Language Models	15864
<i>Ashutosh Bajpai, Aaryan Goyal, Atif Anwer, Tanmoy Chakraborty</i>	
A Comparison of Language Modeling and Translation as Multilingual Pretraining Objectives	15882
<i>Zihao Li, Shaoxiong Ji, Timothee Mickus, Vincent Segonne, Jörg Tiedemann</i>	
Can LLMs Replace Neil deGrasse Tyson? Evaluating the Reliability of LLMs as Science Communicators	15895
<i>Prasoon Bajpai, Niladri Chatterjee, Subhabrata Dutta, Tanmoy Chakraborty</i>	
LLaMA-MoE: Building Mixture-Of-Experts from LLaMA with Continual Pre-Training.....	15913
<i>Tong Zhu, Xiaoye Qu, Daize Dong, Jiacheng Ruan, Jingqi Tong, Conghui He, Yu Cheng</i>	
Themis: A Reference-Free NLG Evaluation Language Model with Flexibility and Interpretability.....	15924
<i>Xinyu Hu, Li Lin, Mingqi Gao, Xunjian Yin, Xiaojun Wan</i>	
Mitigating Training Imbalance in LLM Fine-Tuning Via Selective Parameter Merging	15952
<i>Yiming Ju, Ziyi Ni, Xingrun Xing, Zhixiong Zeng, Hanyu Zhao, Siqi Fan, Zheng Zhang</i>	
Generating Demonstrations for In-Context Compositional Generalization in Grounded Language Learning	15960
<i>Sam Spilsbury, Pekka Marttinen, Alexander Ilin</i>	
FAME: Towards Factual Multi-Task Model Editing	15992
<i>Li Zeng, Yingyu Shan, Zeming Liu, Jiashu Yao, Yuhang Guo</i>	
MLLM-Protector: Ensuring MLLM’s Safety Without Hurting Performance.....	16012
<i>Renjie Pi, Tianyang Han, Jianshu Zhang, Yueqi Xie, Rui Pan, Qing Lian, Hanze Dong, Jipeng Zhang, Tong Zhang</i>	
Leveraging Large Language Models for NLG Evaluation: Advances and Challenges	16028
<i>Zhen Li, Xiaohan Xu, Tao Shen, Can Xu, Jia-Chen Gu, Yuxuan Lai, Chongyang Tao, Shuai Ma</i>	
InfiniPot: Infinite Context Processing on Memory-Constrained LLMs	16046
<i>Minsoo Kim, Kyuhong Shim, Jungwook Choi, Simyung Chang</i>	
VideoCLIP-XL: Advancing Long Description Understanding for Video CLIP Models	16061
<i>Jiapeng Wang, Chengyu Wang, Kunzhe Huang, Jun Huang, Lianwen Jin</i>	

CorrSynth - A Correlated Sampling Method for Diverse Dataset Generation from LLMs	16076
<i>Suhas S Kowshik, Abhishek Divekar, Vijit Malik</i>	
Defining Knowledge: Bridging Epistemology and Large Language Models.....	16096
<i>Constanza Fierro, Ruchira Dhar, Filippos Stamatiou, Nicolas Garneau, Anders Søgaard</i>	
TKGT: Redefinition and a New Way of Text-To-Table Tasks Based on Real World Demands and Knowledge Graphs Augmented LLMs.....	16112
<i>Peiwen Jiang, Xinbo Lin, Zibo Zhao, Ruhui Ma, Yvonne Jie Chen, Jinhua Cheng</i>	
Free Your Mouse! Command Large Language Models to Generate Code to Format Word Documents.....	16127
<i>Shihao Rao, Liang Li, Jiapeng Liu, Guan Weixin, Xiyao Gao, Bing Lim, Can Ma</i>	
CMR Scaling Law: Predicting Critical Mixture Ratios for Continual Pre-Training of Language Models.....	16143
<i>Jiawei Gu, Zacc Yang, Chuanghao Ding, Rui Zhao, Fei Tan</i>	
The Instinctive Bias: Spurious Images Lead to Illusion in MLLMs.....	16163
<i>Tianyang Han, Qing Lian, Rui Pan, Renjie Pi, Jipeng Zhang, Shizhe Diao, Yong Lin, Tong Zhang</i>	
Rationale-Aware Answer Verification by Pairwise Self-Evaluation.....	16178
<i>Akira Kawabata, Saku Sugawara</i>	
On the Robustness of Editing Large Language Models	16197
<i>Xinbei Ma, Tianjie Ju, Jiyang Qiu, Zhuosheng Zhang, Hai Zhao, Lifeng Liu, Yulong Wang</i>	
IM-BERT: Enhancing Robustness of BERT Through the Implicit Euler Method.....	16217
<i>Mihyeon Kim, Juhyoung Park, Youngbin Kim</i>	
Distract Large Language Models for Automatic Jailbreak Attack.....	16230
<i>Zeguan Xiao, Yan Yang, Guanhua Chen, Yun Chen</i>	
Exploring Space Efficiency in a Tree-Based Linear Model for Extreme Multi-Label Classification	16245
<i>He-Zhe Lin, Cheng-Hung Liu, Chih-Jen Lin</i>	
WorryWords: Norms of Anxiety Association for Over 44k English Words	16261
<i>Saif M. Mohammad</i>	
Finding Blind Spots in Evaluator LLMs with Interpretable Checklists.....	16279
<i>Sumanth Doddapaneni, Mohammed Safi Ur Rahman Khan, Sshubam Verma, Mitesh M Khapra</i>	
LONGAGENT: Achieving Question Answering for 128k-Token-Long Documents Through Multi-Agent Collaboration	16310
<i>Jun Zhao, Can Zu, Xu Hao, Yi Lu, Wei He, Yiwen Ding, Tao Gui, Qi Zhang, Xuanjing Huang</i>	
AutoPersuade: A Framework for Evaluating and Explaining Persuasive Arguments.....	16325
<i>Till Raphael Saenger, Musashi Hinck, Justin Grimmer, Brandon M. Stewart</i>	
Towards Cross-Cultural Machine Translation with Retrieval-Augmented Generation from Multilingual Knowledge Graphs	16343
<i>Simone Conia, Daniel Lee, Min Li, Umar Farooq Minhas, Saloni Potdar, Yunyao Li</i>	

Exploring the Compositional Deficiency of Large Language Models in Mathematical Reasoning Through Trap Problems.....	16361
<i>Jun Zhao, Jingqi Tong, Yurong Mou, Ming Zhang, Qi Zhang, Xuanjing Huang</i>	
Scaling Laws for Linear Complexity Language Models.....	16377
<i>Xuyang Shen, Dong Li, Ruitao Leng, Zhen Qin, Weigao Sun, Yiran Zhong</i>	
Autoregressive Multi-Trait Essay Scoring Via Reinforcement Learning with Scoring-Aware Multiple Rewards.....	16427
<i>Heejin Do, Sangwon Ryu, Gary Lee</i>	
Intrinsic Self-Correction for Enhanced Morality: An Analysis of Internal Mechanisms and the Superficial Hypothesis.....	16439
<i>Guangliang Liu, Haitao Mao, Jiliang Tang, Kristen Johnson</i>	
ATAP: Automatic Template-Augmented Commonsense Knowledge Graph Completion Via Pre-Trained Language Models.....	16456
<i>Fu Zhang, Yifan Ding, Jingwei Cheng</i>	
LM2: A Simple Society of Language Models Solves Complex Reasoning.....	16473
<i>Gurusha Juneja, Subhabrata Dutta, Tanmoy Chakraborty</i>	
Towards a Similarity-Adjusted Surprisal Theory.....	16485
<i>Clara Meister, Mario Giulianelli, Tiago Pimentel</i>	
Multi-Level Information Retrieval Augmented Generation for Knowledge-Based Visual Question Answering.....	16499
<i>Omar Adjali, Olivier Ferret, Sahar Ghannay, Hervé Le Borgne</i>	

VOLUME 22

Can We Trust the Performance Evaluation of Uncertainty Estimation Methods in Text Summarization?.....	16514
<i>Jianfeng He, Runing Yang, Linlin Yu, Changbin Li, Ruoxi Jia, Feng Chen, Ming Jin, Chang-Tien Lu</i>	
Is it Really Long Context if All You Need is Retrieval? Towards Genuinely Difficult Long Context NLP.....	16576
<i>Omer Goldman, Alon Jacovi, Aviv Slobodkin, Aviya Maimon, Ido Dagan, Reut Tsarfaty</i>	
BPE Gets Picky: Efficient Vocabulary Refinement During Tokenizer Training.....	16587
<i>Pavel Chizhov, Catherine Arnett, Elizaveta Korotkova, Ivan P. Yamshchikov</i>	
SEGMENT+: Long Text Processing with Short-Context Language Models.....	16605
<i>Wei Shi, Shuang Li, Kerun Yu, Jinglei Chen, Zujie Liang, Xinhui Wu, Yuxi Qian, Feng Wei, Bo Zheng, Jiaqing Liang, Jiangjie Chen, Yanghua Xiao</i>	
Explicit Memory Learning with Expectation Maximization.....	16618
<i>Zhangyue Yin, Qiushi Sun, Qipeng Guo, Zhiyuan Zeng, Qinyuan Cheng, Xipeng Qiu, Xuanjing Huang</i>	
Closing the Loop: Learning to Generate Writing Feedback Via Language Model Simulated Student Revisions.....	16636
<i>Inderjeet Jayakumar Nair, Jiaye Tan, Xiaotian Su, Anne Gere, Xu Wang, Lu Wang</i>	

Small LLMs Are Weak Tool Learners: A Multi-LLM Agent.....	16658
<i>Weizhou Shen, Chenliang Li, Hongzhan Chen, Ming Yan, Xiaojun Quan, Hehong Chen, Ji Zhang, Fei Huang</i>	
Interpreting Context Look-Ups in Transformers: Investigating Attention-MLP Interactions.....	16681
<i>Clement Neo, Shay B Cohen, Fazl Barez</i>	
Still Not Quite There! Evaluating Large Language Models for Comorbid Mental Health Diagnosis.....	16698
<i>Amey Hengle, Atharva Kulkarni, Shantanu Deepak Patankar, Madhumitha Chandrasekaran, Sneha D'Silva, Jemima S. Jacob, Rashmi Gupta</i>	
The Odyssey of Commonsense Causality: From Foundational Benchmarks to Cutting-Edge Reasoning	16722
<i>Shaobo Cui, Zhijing Jin, Bernhard Schölkopf, Boi Faltings</i>	
Investigating Large Language Models for Complex Word Identification in Multilingual and Multidomain Setups	16764
<i>Razvan-Alexandru Smadu, David-Gabriel Ion, Dumitru-Clementin Cercel, Florin Pop, Mihaela-Claudia Cercel</i>	
Model Editing Harms General Abilities of Large Language Models: Regularization to the Rescue	16801
<i>Jia-Chen Gu, Hao-Xiang Xu, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang, Nanyun Peng</i>	
Are Large Language Models In-Context Personalized Summarizers? Get an iCOPERNICUS Test Done!.....	16820
<i>Divya Patel, Pathik Patel, Ankush Chander, Sourish Dasgupta, Tanmoy Chakraborty</i>	
MediTOD: An English Dialogue Dataset for Medical History Taking with Comprehensive Annotations	16843
<i>Vishal Vivek Saley, Goonjan Saha, Rocktim Jyoti Das, Dinesh Raghu, Mausam .</i>	
YesBut: A High-Quality Annotated Multimodal Dataset for Evaluating Satire Comprehension Capability of Vision-Language Models.....	16878
<i>Abhilash Nandy, Yash Agarwal, Ashish Patwa, Millon Madhur Das, Aman Bansal, Ankit Raj, Pawan Goyal, Niloy Ganguly</i>	
Working Memory Identifies Reasoning Limits in Language Models.....	16896
<i>Chunhui Zhang, Yiren Jian, Zhongyu Ouyang, Soroush Vosoughi</i>	
RAFT: Realistic Attacks to Fool Text Detectors.....	16923
<i>James Liyuan Wang, Ran Li, Junfeng Yang, Chengzhi Mao</i>	
LLM-Evolve: Evaluation for LLM's Evolving Capability on Benchmarks	16937
<i>Jiaxuan You, Mingjie Liu, Shrimai Prabhumoye, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro</i>	
FFN-SkipLLM: A Hidden Gem for Autoregressive Decoding with Adaptive Feed Forward Skipping	16943
<i>Ajay Kumar Jaiswal, Bodun Hu, Lu Yin, Yeonju Ro, Tianlong Chen, Shiwei Liu, Aditya Akella</i>	
LLM-Based Code-Switched Text Generation for Grammatical Error Correction.....	16957
<i>Tom Potter, Zheng Yuan</i>	
Deciphering the Interplay of Parametric and Non-Parametric Memory in Retrieval-Augmented Language Models	16966
<i>Mehrdad Farahani, Richard Johansson</i>	

On Efficient Language and Vision Assistants for Visually-Situated Natural Language Understanding: What Matters in Reading and Reasoning	16978
<i>Geewook Kim, Minjoon Seo</i>	
Community-Cross-Instruct: Unsupervised Instruction Generation for Aligning Large Language Models to Online Communities.....	17001
<i>Zihao He, Minh Duc Chu, Rebecca Dorn, Siyi Guo, Kristina Lerman</i>	
Mathador-LM: A Dynamic Benchmark for Mathematical Reasoning on Large Language Models	17020
<i>Eldar Kurtic, Amir Moeini, Dan Alistarh</i>	
Reasoning Paths with Reference Objects Elicit Quantitative Spatial Reasoning in Large Vision-Language Models	17028
<i>Yuan-Hong Liao, Rafid Mahmood, Sanja Fidler, David Acuna</i>	
One Thousand and One Pairs: A “novel” Challenge for Long-Context Language Models	17048
<i>Marzena Karpinska, Katherine Thai, Kyle Lo, Tanya Goyal, Mohit Iyer</i>	
Foundational Autoraters: Taming Large Language Models for Better Automatic Evaluation.....	17086
<i>Tu Vu, Kalpesh Krishna, Salaheddin Alzubi, Chris Tar, Manaal Faruqi, Yun-Hsuan Sung</i>	
Do LLMs Learn a True Syntactic Universal?.....	17106
<i>John T. Hale, Miloš Stanojevic</i>	
GDPO: Learning to Directly Align Language Models with Diversity Using GFlowNets.....	17120
<i>Oh Joon Kwon, Daiki E. Matsunaga, Kee-Eung Kim</i>	
How Susceptible Are Large Language Models to Ideological Manipulation?	17140
<i>Kai Chen, Zihao He, Jun Yan, Taiwei Shi, Kristina Lerman</i>	
Measuring Psychological Depth in Language Models	17162
<i>Fabrice Y Harel-Canada, Hanyu Zhou, Sreya Muppalla, Zeynep Senahan Yildiz, Miryung Kim, Amit Sahai, Nanyun Peng</i>	
Media Attitude Detection Via Framing Analysis with Events and Their Relations.....	17197
<i>Jin Zhao, Jingxuan Tu, Han Du, Nianwen Xue</i>	
Fill in the Gaps: Model Calibration and Generalization with Synthetic Data	17211
<i>Yang Ba, Michelle V Mancenido, Rong Pan</i>	
Adaptive Question Answering: Enhancing Language Model Proficiency for Addressing Knowledge Conflicts with Source Citations.....	17226
<i>Sagi Shaier, Ari Kobren, Philip V. Ogren</i>	
Granular Privacy Control for Geolocation with Vision Language Models.....	17240
<i>Ethan Mendes, Yang Chen, James Hays, Sauvik Das, Wei Xu, Alan Ritter</i>	
MedReadMe: A Systematic Study for Fine-Grained Sentence Readability in Medical Domain.....	17293
<i>Chao Jiang, Wei Xu</i>	

VOLUME 23

MemeCLIP: Leveraging CLIP Representations for Multimodal Meme Classification	17320
<i>Siddhant Bikram Shah, Shuvam Shiwakoti, Maheep Chaudhary, Haohan Wang</i>	

FlipGuard: Defending Preference Alignment Against Update Regression with Constrained Optimization.....	17333
<i>Mingye Zhu, Yi Liu, Quan Wang, Junbo Guo, Zhendong Mao</i>	
StorySparkQA: Expert-Annotated QA Pairs with Real-World Knowledge for Children’s Story-Based Learning.....	17351
<i>Jiaju Chen, Yuxuan Lu, Shao Zhang, Bingsheng Yao, Yuanzhe Dong, Ying Xu, Yunyao Li, Qianwen Wang, Dakuo Wang, Yuling Sun</i>	
MedCoT: Medical Chain of Thought Via Hierarchical Expert.....	17371
<i>Jiaxiang Liu, Yuan Wang, Jiawei Du, Joey Tianyi Zhou, Zuozhu Liu</i>	
Varying Sentence Representations Via Condition-Specified Routers.....	17390
<i>Ziyong Lin, Quansen Wang, Zixia Jia, Zilong Zheng</i>	
Inductive-Deductive Strategy Reuse for Multi-Turn Instructional Dialogues.....	17402
<i>Jiao Ou, Jiayu Wu, Che Liu, Fuzheng Zhang, Di Zhang, Kun Gai</i>	
Information Flow Routes: Automatically Interpreting Language Models at Scale	17432
<i>Javier Ferrando, Elena Voita</i>	
A Simple Yet Effective Training-Free Prompt-Free Approach to Chinese Spelling Correction Based on Large Language Models	17446
<i>Houquan Zhou, Zhenghua Li, Bo Zhang, Chen Li, Shaopeng Lai, Ji Zhang, Fei Huang, Min Zhang</i>	
Representational Analysis of Binding in Language Models.....	17468
<i>Qin Dai, Benjamin Heinzerling, Kentaro Inui</i>	
CoSafe: Evaluating Large Language Model Safety in Multi-Turn Dialogue Coreference.....	17494
<i>Erxin Yu, Jing Li, Ming Liao, Siqi Wang, Gao Zuchen, Fei Mi, Lanqing Hong</i>	
ClimRetrieve: A Benchmarking Dataset for Information Retrieval from Corporate Climate Disclosures	17509
<i>Tobias Schimanski, Jingwei Ni, Roberto Spacey Martín, Nicola Ranger, Markus Leippold</i>	
Context-Aware Adapter Tuning for Few-Shot Relation Learning in Knowledge Graphs	17525
<i>Liu Ran, Zhongzhou Liu, Xiaoli Li, Yuan Fang</i>	
Zero-Shot Detection of LLM-Generated Text Using Token Cohesiveness	17538
<i>Shixuan Ma, Quan Wang</i>	
Dual-Oriented Disentangled Network with Counterfactual Intervention for Multimodal Intent Detection	17554
<i>Zhanpeng Chen, Zhihong Zhu, Xianwei Zhuang, Zhiqi Huang, Yuexian Zou</i>	
From LLMs to MLLMs: Exploring the Landscape of Multimodal Jailbreaking.....	17568
<i>Siyuan Wang, Zhuohan Long, Zhihao Fan, Zhongyu Wei</i>	
Symbolic Working Memory Enhances Language Models for Complex Rule Application	17583
<i>Siyuan Wang, Zhongyu Wei, Yejin Choi, Xiang Ren</i>	
LLoCO: Learning Long Contexts Offline	17605
<i>Sijun Tan, Xiuyu Li, Shishir G Patil, Ziyang Wu, Tianjun Zhang, Kurt Keutzer, Joseph E. Gonzalez, Raluca Ada Popa</i>	

Don't Forget Your Reward Values: Language Model Alignment Via Value-Based Calibration.....	17622
<i>Xin Mao, Feng-Lin Li, Huimin Xu, Wei Zhang, Wang Chen, Anh Tuan Luu</i>	
Mentor-KD: Making Small Language Models Better Multi-Step Reasoners.....	17643
<i>Hojae Lee, Junho Kim, Sangkeun Lee</i>	
Are Large Language Models Capable of Generating Human-Level Narratives?	17659
<i>Yufei Tian, Tenghao Huang, Miri Liu, Derek Jiang, Alexander Spangher, Muhao Chen, Jonathan May, Nanyun Peng</i>	
MP2D: An Automated Topic Shift Dialogue Generation Framework Leveraging Knowledge Graphs	17682
<i>Yerin Hwang, Yongil Kim, Yunah Jang, Jeessoo Bang, Hyunkyung Bae, Kyomin Jung</i>	
Can Large Language Models Enhance Predictions of Disease Progression? Investigating Through Disease Network Link Prediction.....	17703
<i>Haohui Lu, Usman Naseem</i>	
Searching for Best Practices in Retrieval-Augmented Generation.....	17716
<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, Xuanjing Huang</i>	
Moral Foundations of Large Language Models	17737
<i>Marwa Abdulhai, Gregory Serapio-García, Clement Crepy, Daria Valter, John Canny, Natasha Jaques</i>	
The Zeno's Paradox of 'Low-Resource' Languages.....	17753
<i>Hellina Hailu Nigatu, Atnafu Lambebo Tonja, Benjamin Rosman, Thamar Solorio, Monojit Choudhury</i>	
Knowledge Planning in Large Language Models for Domain-Aligned Counseling Summarization.....	17775
<i>Aseem Srivastava, Smriti Joshi, Tanmoy Chakraborty, Md Shad Akhtar</i>	
Enhancing Post-Hoc Attributions in Long Document Comprehension Via Coarse Grained Answer Decomposition.....	17790
<i>Pritika Ramu, Koustava Goswami, Apoorv Saxena, Balaji Vasanth Srinivasan</i>	
From Descriptive Richness to Bias: Unveiling the Dark Side of Generative Image Caption Enrichment	17807
<i>Yusuke Hirota, Ryo Hachiuma, Chao-Han Huck Yang, Yuta Nakashima</i>	
Pruning Via Merging: Compressing LLMs Via Manifold Alignment Based Layer Merging.....	17817
<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, Dianbo Sui</i>	
Embedded Named Entity Recognition Using Probing Classifiers.....	17830
<i>Nicholas Popovic, Michael Färber</i>	
Unleashing the Power of Emojis in Texts Via Self-Supervised Graph Pre-Training.....	17851
<i>Zhou Zhang, Dongzeng Tan, Jiaan Wang, Yilong Chen, Jiarong Xu</i>	
Data Contamination Can Cross Language Barriers.....	17864
<i>Feng Yao, Yufan Zhuang, Zihao Sun, Sunan Xu, Animesh Kumar, Jingbo Shang</i>	
Automated Essay Scoring: A Reflection on the State of the Art.....	17876
<i>Shengjie Li, Vincent Ng</i>	

Encouraging Divergent Thinking in Large Language Models Through Multi-Agent Debate.....	17889
<i>Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, Zhaopeng Tu</i>	
Unveiling and Consulting Core Experts in Retrieval-Augmented MoE-Based LLMs.....	17905
<i>Xin Zhou, Ping Nie, Yiwen Guo, Haojie Wei, Zhanqiu Zhang, Pasquale Minervini, Ruotian Ma, Tao Gui, Qi Zhang, Xuanjing Huang</i>	
CURE: Context- And Uncertainty-Aware Mental Disorder Detection.....	17924
<i>Migyeong Kang, Goun Choi, Hyolim Jeon, Ji Hyun An, Daejin Choi, Jinyoung Han</i>	
PepRec: Progressive Enhancement of Prompting for Recommendation.....	17941
<i>Yakun Yu, Shi-Ang Qi, Baochun Li, Di Niu</i>	
In-Context Compositional Generalization for Large Vision-Language Models.....	17954
<i>Chuanhao Li, Chenchen Jing, Zhen Li, Mingliang Zhai, Yuwei Wu, Yunde Jia</i>	
Improving Zero-Shot LLM Re-Ranker with Risk Minimization.....	17967
<i>Xiaowei Yuan, Zhao Yang, Yequan Wang, Jun Zhao, Kang Liu</i>	
Game on Tree: Visual Hallucination Mitigation Via Coarse-To-Fine View Tree and Game Theory.....	17984
<i>Xianwei Zhuang, Zhihong Zhu, Zhanpeng Chen, Yuxin Xie, Liming Liang, Yuexian Zou</i>	
Label Confidence Weighted Learning for Target-Level Sentence Simplification	18004
<i>Xin Ying Qiu, Jingshen Zhang</i>	
Quantum Recurrent Architectures for Text Classification.....	18020
<i>Wenduan Xu, Stephen Clark, Douglas Brown, Gabriel Matos, Konstantinos Meichanetzidis</i>	
Tree of Problems: Improving Structured Problem Solving with Compositionality.....	18028
<i>Armel Randy Zebaze, Benoît Sagot, Rachel Bawden</i>	
What the Harm? Quantifying the Tangible Impact of Gender Bias in Machine Translation with a Human-Centered Study	18048
<i>Beatrice Savoldi, Sara Papi, Matteo Negri, Ana Guerberof-Arenas, Luisa Bentivogli</i>	
Seg2Act: Global Context-Aware Action Generation for Document Logical Structuring.....	18077
<i>Zichao Li, Shaojie He, Meng Liao, Xuanang Chen, Yaojie Lu, Hongyu Lin, Yanxiong Lu, Xianpei Han, Le Sun</i>	
Is C4 Dataset Optimal for Pruning? an Investigation of Calibration Data for LLM Pruning.....	18089
<i>Abhinav Bandari, Lu Yin, Cheng-Yu Hsieh, Ajay Kumar Jaiswal, Tianlong Chen, Li Shen, Ranjay Krishna, Shiwei Liu</i>	
Revisiting the Robustness of Watermarking to Paraphrasing Attacks	18100
<i>Saksham Rastogi, Danish Pruthi</i>	

VOLUME 24

A Survey of Ontology Expansion for Conversational Understanding.....	18111
<i>Jinggui Liang, Yuxia Wu, Yuan Fang, Hao Fei, Lizi Liao</i>	
Calibrating Language Models with Adaptive Temperature Scaling	18128
<i>Johnathan Xie, Annie S Chen, Yoonho Lee, Eric Mitchell, Chelsea Finn</i>	

Which Programming Language and What Features at Pre-Training Stage Affect Downstream Logical Inference Performance?.....	18139
<i>Fumiya Uchiyama, Takeshi Kojima, Andrew Gambardella, Qi Cao, Yusuke Iwasawa, Yutaka Matsuo</i>	
Why Do Objects Have Many Names? a Study on Word Informativeness in Language Use and Lexical Systems.....	18150
<i>Eleonora Gualdoni, Gemma Boleda</i>	
Dual-Space Knowledge Distillation for Large Language Models.....	18164
<i>Songming Zhang, Xue Zhang, Zengkui Sun, Yufeng Chen, Jinan Xu</i>	
NoiseBench: Benchmarking the Impact of Real Label Noise on Named Entity Recognition.....	18182
<i>Elena Merdjanovska, Ansar Aynedinov, Alan Akbik</i>	
On the Universal Truthfulness Hyperplane Inside LLMs.....	18199
<i>Junteng Liu, Shiqi Chen, Yu Cheng, Junxian He</i>	
PairDistill: Pairwise Relevance Distillation for Dense Retrieval	18225
<i>Chao-Wei Huang, Yun-Nung Chen</i>	
User Inference Attacks on Large Language Models.....	18238
<i>Nikhil Kandpal, Krishna Pillutla, Alina Oprea, Peter Kairouz, Christopher A. Choquette-Choo, Zheng Xu</i>	
HiFT: A Hierarchical Full Parameter Fine-Tuning Strategy.....	18266
<i>Yongkang Liu, Yiqun Zhang, Qian Li, Tong Liu, Shi Feng, Daling Wang, Yifei Zhang, Hinrich Schuetze</i>	
Investigating and Mitigating Object Hallucinations in Pretrained Vision-Language (CLIP) Models	18288
<i>Yufang Liu, Tao Ji, Changzhi Sun, Yuanbin Wu, Aimin Zhou</i>	
Simultaneous Masking, Not Prompting Optimization: A Paradigm Shift in Fine-Tuning LLMs for Simultaneous Translation	18302
<i>Matthew Raffel, Victor Agostinelli, Lizhong Chen</i>	
ToolPlanner: A Tool Augmented LLM for Multi Granularity Instructions with Path Planning and Feedback.....	18315
<i>Qinzhao Wu, Wei Liu, Jian Luan, Bin Wang</i>	
Please Note that I'm Just an AI: Analysis of Behavior Patterns of LLMs in (Non-)offensive Speech Identification	18340
<i>Esra Dönmez, Thang Vu, Agnieszka Falenska</i>	
How to Compute the Probability of a Word	18358
<i>Tiago Pimentel, Clara Meister</i>	
A Linguistically-Motivated Evaluation Methodology for Unraveling Model's Abilities in Reading Comprehension Tasks.....	18376
<i>Elie Antoine, Frederic Bechet, Géraldine Damnati, Philippe Langlais</i>	
GuardBench: A Large-Scale Benchmark for Guardrail Models	18393
<i>Elias Bassani, Ignacio Sanchez</i>	

Generate-On-Graph: Treat LLM as Both Agent and KG for Incomplete Knowledge Graph Question Answering	18410
<i>Yao Xu, Shizhu He, Jiabei Chen, Zihao Wang, Yangqiu Song, Hanghang Tong, Guang Liu, Jun Zhao, Kang Liu</i>	
Language Models and Brains Align Due to More than Next-Word Prediction and Word-Level Information	18431
<i>Gabriele Merlin, Mariya Toneva</i>	
LLMEdgeRefine: Enhancing Text Clustering with LLM-Based Boundary Point Refinement	18455
<i>Zijin Feng, Luyang Lin, Lingzhi Wang, Hong Cheng, Kam-Fai Wong</i>	
CasiMedicos-Arg: A Medical Question Answering Dataset Annotated with Explanatory Argumentative Structures	18463
<i>Ekaterina Sviridova, Anar Yeginbergen, Ainara Estarrona, Elena Cabrio, Serena Villata, Rodrigo Agerri</i>	
A Simple and Effective L ₂ Norm-Based Strategy for KV Cache Compression.....	18476
<i>Alessio Devoto, Yu Zhao, Simone Scardapane, Pasquale Minervini</i>	
GOME: Grounding-Based Metaphor Binding with Conceptual Elaboration for Figurative Language Illustration	18500
<i>Linhao Zhang, Jintao Liu, Li Jin, Hao Wang, Kaiwen Wei, Guangluan Xu</i>	
D3CODE: Disentangling Disagreements in Data Across Cultures on Offensiveness Detection and Evaluation.....	18511
<i>Aida Mostafazadeh Davani, Mark Diaz, Dylan K Baker, Vinodkumar Prabhakaran</i>	
PALM: Few-Shot Prompt Learning for Audio Language Models.....	18527
<i>Asif Hanif, Maha Tufail Agro, Mohammad Areeb Qazi, Hanan Aldarmaki</i>	
Annotator-Centric Active Learning for Subjective NLP Tasks.....	18537
<i>Michiel Van Der Meer, Neele Falk, Pradeep K. Murukannaiah, Enrico Liscio</i>	
On the Proper Treatment of Tokenization in Psycholinguistics.....	18556
<i>Mario Giulianelli, Luca Malagutti, Juan Luis Gastaldi, Brian Dusell, Tim Vieira, Ryan Cotterell</i>	
Enhanced Hallucination Detection in Neural Machine Translation Through Simple Detector Aggregation	18573
<i>Anas Himmi, Guillaume Staerman, Marine Picot, Pierre Colombo, Nuno M Guerreiro</i>	
Jailbreaking LLMs with Arabic Transliteration and Arabizi	18584
<i>Mansour Al Ghanim, Saleh Almohaimeed, Mengxin Zheng, Yan Solihin, Qian Lou</i>	
Who is Better at Math, Jenny Or Jingzhen? Uncovering Stereotypes in Large Language Models	18601
<i>Zara Siddique, Liam Turner, Luis Espinosa-Anke</i>	
Instruction Matters: A Simple Yet Effective Task Selection for Optimized Instruction Tuning of Specific Tasks.....	18620
<i>Changho Lee, Janghoon Han, Seonghyeon Ye, Stanley Jungkyu Choi, Honglak Lee, Kyunghoon Bae</i>	
Recurrent Alignment with Hard Attention for Hierarchical Text Rating.....	18643
<i>Chenxi Lin, Ren Jiayu, Guoxiu He, Zhuoren Jiang, Haiyan Yu, Xiaomin Zhu</i>	

CHES: Optimizing LLM Inference Via Channel-Wise Thresholding and Selective Sparsification.....	18658
<i>Junhui He, Shangyu Wu, Weidong Wen, Chun Jason Xue, Qingan Li</i>	
Semformer: Transformer Language Models with Semantic Planning.....	18669
<i>Yongjing Yin, Junran Ding, Kai Song, Yue Zhang</i>	
DocCGen: Document-Based Controlled Code Generation	18681
<i>Sameer Pimparkhede, Mehant Kammakomati, Srikanth G. Tamilselvam, Prince Kumar, Ashok Pon Kumar, Pushpak Bhattacharyya</i>	
Semantics and Sentiment: Cross-Lingual Variations in Emoji Use.....	18698
<i>Giulio Zhou, Sydelle De Souza, Ella Markham, Oghenetekewe Kwakpovwe, Sumin Zhao</i>	
The Emergence of Compositional Languages in Multi-Entity Referential Games: From Image to Graph Representations.....	18713
<i>Daniel Akkerman, Phong Le, Raquel G. Alhama</i>	
Transformers Are Multi-State RNNs.....	18724
<i>Matanel Oren, Michael Hassid, Nir Yarden, Yossi Adi, Roy Schwartz</i>	
Evaluating Large Language Models Along Dimensions of Language Variation: A Systematic Investigation of Cross-Lingual Generalization	18742
<i>Niyati Bafna, Kenton Murray, David Yarowsky</i>	
Fuse to Forget: Bias Reduction and Selective Memorization Through Model Fusion	18763
<i>Kerem Zaman, Leshem Choshen, Shashank Srivastava</i>	
Collective Critics for Creative Story Generation.....	18784
<i>Minwook Bae, Hyounghun Kim</i>	
Surprise! Uniform Information Density Isn't the Whole Story: Predicting Surprisal Contours in Long-Form Discourse.....	18820
<i>Eleftheria Tsipidi, Franz Nowak, Ryan Cotterell, Ethan Wilcox, Mario Giulianelli, Alex Warstadt</i>	
Model-Based Preference Optimization in Abstractive Summarization Without Human Feedback	18837
<i>Jaepill Choi, Kyubyung Chae, Jiwoo Song, Yohan Jo, Taesup Kim</i>	
Are Data Augmentation Methods in Named Entity Recognition Applicable for Uncertainty Estimation?	18852
<i>Wataru Hashimoto, Hidetaka Kamigaito, Taro Watanabe</i>	
NeuroTrialNER: An Annotated Corpus for Neurological Diseases and Therapies in Clinical Trial Registries.....	18868
<i>Simona Emilova Doneva, Tilia Ellendorff, Beate Sick, Jean-Philippe Goldman, Amelia Elaine Cannon, Gerold Schneider, Benjamin Victor Ineichen</i>	
Fool Me Once? Contrasting Textual and Visual Explanations in a Clinical Decision-Support Setting.....	18891
<i>Maxime Guillaume Kayser, Bayar Menzat, Cornelius Emde, Bogdan Alexandru Bercean, Alex Novak, Abdalá Trinidad Espinosa Morgado, Bartłomiej Papież, Susanne Gaube, Thomas Lukasiewicz, Oana-Maria Camburu</i>	

VOLUME 25

Towards Faithful Knowledge Graph Explanation Through Deep Alignment in Commonsense Question Answering	18920
<i>Weihe Zhai, Arkaitz Zubiaga, Bingquan Liu, Chengjie Sun, Yalong Zhao</i>	
Generation with Dynamic Vocabulary.....	18931
<i>Yanting Liu, Tao Ji, Changzhi Sun, Yuanbin Wu, Xiaoling Wang</i>	
Argument Relation Classification Through Discourse Markers and Adversarial Training	18949
<i>Michele Luca Contalbo, Francesco Guerra, Matteo Paganelli</i>	
Getting the Most Out of Your Training Data: Exploring Unsupervised Tasks for Morphological Inflection	18955
<i>Abhishek Purushothama, Adam Wiemerslage, Katharina Von Der Wense</i>	
Link, Synthesize, Retrieve: Universal Document Linking for Zero-Shot Information Retrieval.....	18971
<i>Dae Yon Hwang, Bilal Taha, Harshit Pande, Yaroslav Nechaev</i>	
Efficient Unseen Language Adaptation for Multilingual Pre-Trained Language Models	18983
<i>Po-Heng Chen, Yun-Nung Chen</i>	
Prove Your Point!: Bringing Proof-Enhancement Principles to Argumentative Essay Generation	18995
<i>Ruiyu Xiao, Lei Wu, Yuhang Gou, Weinan Zhang, Ting Liu</i>	
TV-TREES: Multimodal Entailment Trees for Neuro-Symbolic Video Reasoning	19009
<i>Kate Sanders, Nathaniel Weir, Benjamin Van Durme</i>	
Unsupervised Extraction of Dialogue Policies from Conversations.....	19029
<i>Makesh Narsimhan Sreedhar, Traian Rebedea, Christopher Parisien</i>	
GRIZAL: Generative Prior-Guided Zero-Shot Temporal Action Localization	19046
<i>Onkar Kishor Susladkar, Gayatri Sudhir Deshmukh, Vandan Gorade, Sparsh Mittal</i>	
Preserving Multi-Modal Capabilities of Pre-Trained VLMs for Improving Vision-Linguistic Compositionality	19060
<i>Youngtaek Oh, Jae Won Cho, Dong-Jin Kim, In So Kweon, Junmo Kim</i>	
FoodieQA: A Multimodal Dataset for Fine-Grained Understanding of Chinese Food Culture.....	19077
<i>Wenyan Li, Crystina Zhang, Jiaang Li, Qiwei Peng, Raphael Tang, Li Zhou, Weijia Zhang, Guimin Hu, Yifei Yuan, Anders Søgaard, Daniel Hershcovich, Desmond Elliott</i>	
A Two-Step Approach for Data-Efficient French Pronunciation Learning.....	19096
<i>Hoyeon Lee, Hyeun Jang, Jonghwan Kim, Jaemin Kim</i>	
Exploring Intra and Inter-Language Consistency in Embeddings with ICA	19104
<i>Rongzhi Li, Takeru Matsuda, Hitomi Yanaka</i>	
DetoxLLM: A Framework for Detoxification with Explanations.....	19112
<i>Md Tawkat Islam Khondaker, Muhammad Abdul-Mageed, Laks V. S. Lakshmanan</i>	
Comparing a BERT Classifier and a GPT Classifier for Detecting Connective Language Across Multiple Social Media	19140
<i>Josephine Lukito, Bin Chen, Gina M. Masullo, Natalie Jomini Stroud</i>	

ShadowLLM: Predictor-Based Contextual Sparsity for Large Language Models	19154
<i>Yash Akhauri, Ahmed F Abouelhamayed, Jordan Dotzel, Zhiru Zhang, Alexander M Rush, Safeen Huda, Mohamed S. Abdelfattah</i>	
Emotion Granularity from Text: An Aggregate-Level Indicator of Mental Health	19168
<i>Krishnapriya Vishnubhotla, Daniela Teodorescu, Mallory J Feldman, Kristen Lindquist, Saif M. Mohammad</i>	
BLSP-Emo: Towards Empathetic Large Speech-Language Models	19186
<i>Chen Wang, Minpeng Liao, Zhongqiang Huang, Junhong Wu, Chengqing Zong, Jiajun Zhang</i>	
SynthesizRR: Generating Diverse Datasets with Retrieval Augmentation.....	19200
<i>Abhishek Divekar, Greg Durrett</i>	
Multimodal Self-Instruct: Synthetic Abstract Image and Visual Reasoning Instruction Using Language Model.....	19228
<i>Wenqi Zhang, Zhenglin Cheng, Yuanyu He, Mengna Wang, Yongliang Shen, Zeqi Tan, Guiyang Hou, Mingqian He, Yanna Ma, Weiming Lu, Yueting Zhuang</i>	
DataNarrative: Automated Data-Driven Storytelling with Visualizations and Texts.....	19253
<i>Mohammed Saidul Islam, Md Tahmid Rahman Laskar, Md Rizwan Parvez, Enamul Hoque, Shafiq Joty</i>	
DEM: Distribution Edited Model for Training with Mixed Data Distributions	19287
<i>Dhananjay Ram, Aditya Rawal, Momchil Hardalov, Nikolaos Pappas, Sheng Zha</i>	
Altogether: Image Captioning Via Re-Aligning Alt-Text.....	19302
<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, Christoph Feichtenhofer</i>	
VerifyMatch: A Semi-Supervised Learning Paradigm for Natural Language Inference with Confidence-Aware MixUp.....	19319
<i>Seo Yeon Park, Cornelia Caragea</i>	
CaT-Bench: Benchmarking Language Model Understanding of Causal and Temporal Dependencies in Plans	19336
<i>Yash Kumar Lal, Vanya Cohen, Nathanael Chambers, Niranjana Balasubramanian, Ray Mooney</i>	
Mitigating the Impact of Reference Quality on Evaluation of Summarization Systems with Reference-Free Metrics	19355
<i>Théo Gigant, Camille Guinaudeau, Marc Decombas, Frederic Dufaux</i>	
An Empirical Analysis of the Writing Styles of Persona-Assigned LLMs	19369
<i>Manuj Malik, Jing Jiang, Kian Ming A. Chai</i>	
Investigating the Role of Instruction Variety and Task Difficulty in Robotic Manipulation Tasks	19389
<i>Amit Parekh, Nikolas Vitsakis, Alessandro Suglia, Ioannis Konstas</i>	
GPT Vs RETRO: Exploring the Intersection of Retrieval and Parameter-Efficient Fine-Tuning	19425
<i>Aleksander Ficek, Jiaqi Zeng, Oleksii Kuchaiev</i>	
CoCoST: Automatic Complex Code Generation with Online Searching and Correctness Testing	19433
<i>Xinyi He, Jiaru Zou, Yun Lin, Mengyu Zhou, Shi Han, Zejian Yuan, Dongmei Zhang</i>	

Sequential API Function Calling Using GraphQL Schema.....	19452
<i>Avirup Saha, Lakshmi Mandal, Balaji Ganesan, Sambit Ghosh, Renuka Sindhgatta, Carlos Eberhardt, Dan Debrunner, Sameep Mehta</i>	
The Illusion of Competence: Evaluating the Effect of Explanations on Users' Mental Models of Visual Question Answering Systems.....	19459
<i>Judith Sieker, Simeon Junker, Ronja Utescher, Nazia Attari, Heiko Wersing, Hendrik Buschmeier, Sina Zarriß</i>	
Re-Evaluating Evaluation for Multilingual Summarization.....	19476
<i>Jessica Zosa Forde, Ruochen Zhang, Lintang Sutawika, Alham Fikri Aji, Samuel Cahyawijaya, Genta Indra Winata, Minghao Wu, Carsten Eickhoff, Stella Biderman, Ellie Pavlick</i>	
Video-Text Prompting for Weakly Supervised Spatio-Temporal Video Grounding.....	19494
<i>Heng Zhao, Zhao Yinjie, Bihan Wen, Yew-Soon Ong, Joey Tianyi Zhou</i>	
A Fast and Sound Tagging Method for Discontinuous Named-Entity Recognition.....	19506
<i>Caio Filippo Corro</i>	
Factuality of Large Language Models: A Survey.....	19519
<i>Yuxia Wang, Minghan Wang, Muhammad Arslan Manzoor, Fei Liu, Georgi Nenkov Georgiev, Rocktim Jyoti Das, Preslav Nakov</i>	
Discovering Biases in Information Retrieval Models Using Relevance Thesaurus as Global Explanation.....	19530
<i>Youngwoo Kim, Razieh Rahimi, James Allan</i>	
Adaptable Moral Stances of Large Language Models on Sexist Content: Implications for Society and Gender Discourse.....	19548
<i>Rongchen Guo, Isar Nejadgholi, Hillary Dawkins, Kathleen C. Fraser, Svetlana Kiritchenko</i>	
DISCERN: Decoding Systematic Errors in Natural Language for Text Classifiers.....	19565
<i>Rakesh R Menon, Shashank Srivastava</i>	
IntCoOp: Interpretability-Aware Vision-Language Prompt Tuning.....	19584
<i>Soumya Suvra Ghosal, Samyadeep Basu, Soheil Feizi, Dinesh Manocha</i>	
Scope-Enhanced Compositional Semantic Parsing for DRT.....	19602
<i>Xiulin Yang, Jonas Groschwitz, Alexander Koller, Johan Bos</i>	
The Generation Gap: Exploring Age Bias in the Value Systems of Large Language Models.....	19617
<i>Siyang Liu, Trisha Maturi, Bowen Yi, Siqu Shen, Rada Mihalcea</i>	
TempoFormer: A Transformer for Temporally-Aware Representations in Change Detection.....	19635
<i>Talia Tseriotou, Adam Tsakalidis, Maria Liakata</i>	
Pron Vs Prompt: Can Large Language Models Already Challenge a World-Class Fiction Author at Creative Text Writing?.....	19654
<i>Guillermo Marco, Julio Gonzalo, M. Teresa Mateo-Girona, Ramón Del Castillo Santos</i>	
Evaluating Diversity in Automatic Poetry Generation.....	19671
<i>Yanran Chen, Hannes Gröner, Sina Zarriß, Steffen Eger</i>	
Evaluating Short-Term Temporal Fluctuations of Social Biases in Social Media Data and Masked Language Models.....	19693
<i>Yi Zhou, Danushka Bollegala, Jose Camacho-Collados</i>	

VOLUME 26

Delving into Qualitative Implications of Synthetic Data for Hate Speech Detection.....	19709
<i>Camilla Casula, Sebastiano Vecellio Salto, Alan Ramponi, Sara Tonelli</i>	
Grounding Language in Multi-Perspective Referential Communication	19727
<i>Zineng Tang, Lingjun Mao, Alane Suhr</i>	
Threshold-Driven Pruning with Segmented Maximum Term Weights for Approximate Cluster- Based Sparse Retrieval.....	19742
<i>Yifan Qiao, Parker Carlson, Shanxiu He, Yingrui Yang, Tao Yang</i>	
Error Analysis of Multilingual Language Models in Machine Translation: A Case Study of English- Amharic Translation	19758
<i>Hizkiel Mitiku Alemayehu, Hamada M Zahera, Axel-Cyrille Ngonga Ngomo</i>	
MIPD: Exploring Manipulation and Intention in a Novel Corpus of Polish Disinformation.....	19769
<i>Arkadiusz Modzelewski, Giovanni Da San Martino, Pavel Savov, Magdalena Anna Wilczynska, Adam Wierzbicki</i>	
Unsupervised Discrete Representations of American Sign Language.....	19786
<i>Artem Abzaliev, Rada Mihalcea</i>	
Perceptions to Beliefs: Exploring Precursory Inferences for Theory of Mind in Large Language Models.....	19794
<i>Chani Jung, Dongkwan Kim, Jiho Jin, Jiseon Kim, Yeon Seonwoo, Yejin Choi, Alice Oh, Hyunwoo Kim</i>	
Towards Enhancing Coherence in Extractive Summarization: Dataset and Experiments with LLMs.....	19810
<i>Mihir Parmar, Hanieh Deilamsalehy, Franck Deroncourt, Seunghyun Yoon, Ryan A. Rossi, Trung Bui</i>	
Jump Starting Bandits with LLM-Generated Prior Knowledge	19821
<i>Parand A. Alamdari, Yanshuai Cao, Kevin H. Wilson</i>	
Adaptation Odyssey in LLMs: Why Does Additional Pretraining Sometimes Fail to Improve?	19834
<i>Firat Öncel, Matthias Bethge, Beyza Ermis, Mirco Ravanelli, Cem Subakan, Çagatay Yildiz</i>	
Not All Contexts Are Equal: Teaching LLMs Credibility-Aware Generation	19844
<i>Ruotong Pan, Boxi Cao, Hongyu Lin, Xianpei Han, Jia Zheng, Sirui Wang, Xunliang Cai, Le Sun</i>	
Virtual Personas for Language Models Via an Anthology of Backstories.....	19864
<i>Suhong Moon, Marwa Abdulhai, Minwoo Kang, Joseph Suh, Widyadewi Soedarmadji, Eran Kohen Behar, David Chan</i>	
Step-By-Step Reasoning to Solve Grid Puzzles: Where Do LLMs Falter?.....	19898
<i>Nemika Tyagi, Mihir Parmar, Mohith Kulkarni, Aswin Rrv, Nisarg Patel, Mutsumi Nakamura, Arindam Mitra, Chitta Baral</i>	
Reasoning in Token Economies: Budget-Aware Evaluation of LLM Reasoning Strategies	19916
<i>Junlin Wang, Siddhartha Jain, Dejiao Zhang, Baishakhi Ray, Varun Kumar, Ben Athiwaratkun</i>	
The Empirical Variability of Narrative Perceptions of Social Media Texts.....	19940
<i>Joel Mire, Maria Antoniak, Elliott Ash, Andrew Piper, Maarten Sap</i>	

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

AMPO: Automatic Multi-Branched Prompt Optimization.....	20267
<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, Linjun Yang</i>	
DeMPT: Decoding-Enhanced Multi-Phase Prompt Tuning for Making LLMs Be Better Context-Aware Translators.....	20280
<i>Xinglin Lyu, Junhui Li, Yanqing Zhao, Min Zhang, Daimeng Wei, Shimin Tao, Hao Yang, Min Zhang</i>	
DEFT-UCS: Data Efficient Fine-Tuning for Pre-Trained Language Models Via Unsupervised Core-Set Selection for Text-Editing.....	20296
<i>Devleena Das, Vivek Khetan</i>	
Unveiling Multi-Level and Multi-Modal Semantic Representations in the Human Brain Using Large Language Models.....	20313
<i>Yuko Nakagi, Takuya Matsuyama, Naoko Koide-Majima, Hiroto Q. Yamaguchi, Rieko Kubo, Shinji Nishimoto, Yu Takagi</i>	
“They Are Uncultured”: Unveiling Covert Harms and Social Threats in LLM Generated Conversations.....	20339
<i>Preetam Prabhu Srikar Dammu, Hayoung Jung, Anjali Singh, Monojit Choudhury, Tanu Mitra</i>	
Multi-Expert Prompting Improves Reliability, Safety and Usefulness of Large Language Models.....	20370
<i>Do Xuan Long, Duong Ngoc Yen, Anh Tuan Luu, Kenji Kawaguchi, Min-Yen Kan, Nancy F. Chen</i>	
Will LLMs Replace the Encoder-Only Models in Temporal Relation Classification?.....	20402
<i>Gabriel Roccabruna, Massimo Rizzoli, Giuseppe Riccardi</i>	
Eliciting In-Context Learning in Vision-Language Models for Videos Through Curated Data Distributional Properties.....	20416
<i>Keunwoo Peter Yu, Zheyuan Zhang, Fengyuan Hu, Shane Storks, Joyce Chai</i>	
Waterfall: Scalable Framework for Robust Text Watermarking and Provenance for LLMs.....	20432
<i>Gregory Kang Ruey Lau, Xinyuan Niu, Hieu Dao, Jiangwei Chen, Chuan-Sheng Foo, Bryan Kian Hsiang Low</i>	
MASIVE: Open-Ended Affective State Identification in English and Spanish.....	20467
<i>Nicholas Deas, Elsbeth Turcan, Ivan Ernesto Perez Mejia, Kathleen McKeown</i>	
You Make Me Feel Like a Natural Question: Training QA Systems on Transformed Trivia Questions.....	20486
<i>Tasnim Kabir, Yoo Yeon Sung, Saptarashmi Bandyopadhyay, Hao Zou, Abhrranil Chandra, Jordan Lee Boyd-Graber</i>	

VOLUME 27

AlphaLoRA: Assigning LoRA Experts Based on Layer Training Quality.....	20511
<i>Peijun Qing, Chongyang Gao, Yefan Zhou, Xingjian Diao, Yaoqing Yang, Soroush Vosoughi</i>	
Flee the Flaw: Annotating the Underlying Logic of Fallacious Arguments Through Templates and Slot-Filling.....	20524
<i>Irfan Robbani, Paul Reisert, Surawat Pothong, Naoya Inoue, Camélia Guerraoui, Wenzhi Wang, Shoichi Naito, Jungmin Choi, Kentaro Inui</i>	

Advancing Social Intelligence in AI Agents: Technical Challenges and Open Questions.....	20541
<i>Leena Mathur, Paul Pu Liang, Louis-Philippe Morency</i>	
RAT: Injecting Implicit Bias for Text-To-Image Prompt Refinement Models	20561
<i>Ziyi Kou, Shichao Pei, Meng Jiang, Xiangliang Zhang</i>	
Can LLM Generate Culturally Relevant Commonsense QA Data? Case Study in Indonesian and Sundanese.....	20571
<i>Rifki Afina Putri, Faiz Ghifari Haznitrama, Dea Adhista, Alice Oh</i>	
Can Language Models Induce Grammatical Knowledge from Indirect Evidence?.....	20591
<i>Miyu Oba, Yohei Oseki, Akiyo Fukatsu, Akari Haga, Hiroki Ouchi, Taro Watanabe, Saku Sugawara</i>	
Do LLMs Know to Respect Copyright Notice?	20604
<i>Jialiang Xu, Shenglan Li, Zhaozhuo Xu, Denghui Zhang</i>	
SpecHub: Provable Acceleration to Multi-Draft Speculative Decoding	20620
<i>Ryan Sun, Tianyi Zhou, Xun Chen, Lichao Sun</i>	
Interventional Speech Noise Injection for ASR Generalizable Spoken Language Understanding.....	20642
<i>Yeonjoon Jung, Jaeseong Lee, Seungtaek Choi, Dohyeon Lee, Minsoo Kim, Seung-Won Hwang</i>	
Rethinking the Role of Proxy Rewards in Language Model Alignment	20656
<i>Sungdong Kim, Minjoon Seo</i>	
Visual Text Matters: Improving Text-KVQA with Visual Text Entity Knowledge-Aware Large Multimodal Assistant.....	20675
<i>Abhirama Subramanyam Penamakuri, Anand Mishra</i>	
Beyond Correlation: Interpretable Evaluation of Machine Translation Metrics.....	20689
<i>Stefano Perrella, Lorenzo Proietti, Pere-Lluís Huguet Cabot, Edoardo Barba, Roberto Navigli</i>	
IFCap: Image-Like Retrieval and Frequency-Based Entity Filtering for Zero-Shot Captioning.....	20715
<i>Soeun Lee, Si-Woo Kim, Taewhan Kim, Dong-Jin Kim</i>	
Encoding Spreadsheets for Large Language Models.....	20728
<i>Haoyu Dong, Jianbo Zhao, Yuzhang Tian, Junyu Xiong, Mengyu Zhou, Yun Lin, José Cambroner, Yeye He, Shi Han, Dongmei Zhang</i>	
Let's Discuss! Quality Dimensions and Annotated Datasets for Computational Argument Quality Assessment.....	20749
<i>Rositsa V Ivanova, Thomas Huber, Christina Niklaus</i>	
Automatic Sentence Segmentation of Clinical Record Narratives in Real-World Data	20780
<i>Dongfang Xu, Davy Weissenbacher, Karen O'Connor, Siddharth Rawal, Graciela Gonzalez Hernandez</i>	
One-To-Many Communication and Compositionality in Emergent Communication	20794
<i>Heeyoung Lee</i>	
Bayesian Example Selection Improves In-Context Learning for Speech, Text and Visual Modalities	20812
<i>Siyin Wang, Chao-Han Huck Yang, Ji Wu, Chao Zhang</i>	

Investigating Multilingual Instruction-Tuning: Do Polyglot Models Demand for Multilingual Instructions?	20829
<i>Alexander Arno Weber, Klaudia Thellmann, Jan Ebert, Nicolas Flores-Herr, Jens Lehmann, Michael Fromm, Mehdi Ali</i>	
Multi-LogiEval: Towards Evaluating Multi-Step Logical Reasoning Ability of Large Language Models	20856
<i>Nisarg Patel, Mohith Kulkarni, Mihir Parmar, Aashna Budhiraja, Mutsumi Nakamura, Neeraj Varshney, Chitta Baral</i>	
Linear Layer Extrapolation for Fine-Grained Emotion Classification	20880
<i>Mayukh Sharma, Sean O'Brien, Julian McAuley</i>	
Task Oriented In-Domain Data Augmentation	20889
<i>Xiao Liang, Xinyu Hu, Simiao Zuo, Yeyun Gong, Qiang Lou, Yi Liu, Shao-Lun Huang, Jian Jiao</i>	
SciDQA: A Deep Reading Comprehension Dataset Over Scientific Papers	20908
<i>Shruti Singh, Nandan Sarkar, Arman Cohan</i>	
Mixture-Of-Modules: Reinventing Transformers as Dynamic Assemblies of Modules.....	20924
<i>Zhuocheng Gong, Ang Lv, Jian Guan, Wei Wu, Huishuai Zhang, Minlie Huang, Dongyan Zhao, Rui Yan</i>	
No Culture Left Behind: ArtELingo-28, a Benchmark of WikiArt with Captions in 28 Languages.....	20939
<i>Youssef Mohamed, Runjia Li, Ibrahim Said Ahmad, Kilichbek Haydarov, Philip Torr, Kenneth Church, Mohamed Elhoseiny</i>	
PREDICT: Multi-Agent-Based Debate Simulation for Generalized Hate Speech Detection.....	20963
<i>Someen Park, Jaehoon Kim, Seungwan Jin, Sohyun Park, Kyungsik Han</i>	
TokenVerse: Towards Unifying Speech and NLP Tasks Via Transducer-Based ASR	20988
<i>Shashi Kumar, Srikanth Madikeri, Juan Pablo Zuluaga Gomez, Iuliia Thorbecke, Esau Villatoro-Tello, Sergio Burdisso, Petr Motlicek, Karthik Pandia D S, Aravind Ganapathiraju</i>	
ApiQ: Finetuning of 2-Bit Quantized Large Language Model.....	20996
<i>Baohao Liao, Christian Herold, Shahram Khadivi, Christof Monz</i>	
Memorize Step by Step: Efficient Long-Context Prefilling with Incremental Memory and Decremental Chunk	21021
<i>Zhiyuan Zeng, Qipeng Guo, Xiaoran Liu, Zhangyue Yin, Wentao Shu, Mianqiu Huang, Bo Wang, Yunhua Zhou, Linlin Li, Qun Liu, Xipeng Qiu</i>	
A Morphology-Based Investigation of Positional Encodings.....	21035
<i>Poulami Ghosh, Shikhar Vashishth, Raj Dabre, Pushpak Bhattacharyya</i>	
I Love Pineapple on Pizza != I Hate Pineapple on Pizza: Stance-Aware Sentence Transformers for Opinion Mining	21046
<i>Vahid Ghafouri, Jose Such, Guillermo Suarez-Tangil</i>	
BiasWipe: Mitigating Unintended Bias in Text Classifiers Through Model Interpretability	21059
<i>Mamta Mamta, Rishikant Chigrupaatii, Asif Ekbal</i>	
ArMeme: Propagandistic Content in Arabic Memes	21071
<i>Firoj Alam, Abul Hasnat, Fatema Ahmad, Md. Arid Hasan, Maram Hasanain</i>	

Language is Scary When Over-Analyzed: Unpacking Implied Misogynistic Reasoning with Argumentation Theory-Driven Prompts	21091
<i>Arianna Muti, Federico Ruggeri, Khalid Al Khatib, Alberto Barrón-Cedeño, Tommaso Caselli</i>	
Thoughts to Target: Enhance Planning for Target-Driven Conversation	21108
<i>Zhonghua Zheng, Lizi Liao, Yang Deng, Ee-Peng Lim, Minlie Huang, Liqiang Nie</i>	
Scalable Data Ablation Approximations for Language Models Through Modular Training and Merging	21125
<i>Clara Na, Ian Magnusson, Ananya Harsh Jha, Tom Sherborne, Emma Strubell, Jesse Dodge, Pradeep Dasigi</i>	
Exploring Intrinsic Language-Specific Subspaces in Fine-Tuning Multilingual Neural Machine Translation	21142
<i>Zhe Cao, Zhi Qu, Hidetaka Kamigaito, Taro Watanabe</i>	
Attention Score is Not All You Need for Token Importance Indicator in KV Cache Reduction: Value Also Matters	21158
<i>Zhiyu Guo, Hidetaka Kamigaito, Taro Watanabe</i>	
Generative Subgraph Retrieval for Knowledge Graph–Grounded Dialog Generation	21167
<i>Jinyoung Park, Minseok Joo, Joo-Kyung Kim, Hyunwoo J. Kim</i>	
Adapters Mixup: Mixing Parameter-Efficient Adapters to Enhance the Adversarial Robustness of Fine-Tuned Pre-Trained Text Classifiers	21183
<i>Tuc Van Nguyen, Thai Le</i>	
Generalizing Clinical De-Identification Models by Privacy-Safe Data Augmentation Using GPT-4	21204
<i>Woojin Kim, Sungeun Hahm, Jaejin Lee</i>	
Connecting the Dots: Evaluating Abstract Reasoning Capabilities of LLMs Using the New York Times Connections Word Game	21219
<i>Prisha Samdarshi, Mariam Mustafa, Anushka Kulkarni, Raven Rothkopf, Tuhin Chakrabarty, Smaranda Muresan</i>	
GottBERT: A Pure German Language Model	21237
<i>Raphael Scheible, Johann Frei, Fabian Thomczyk, Henry He, Patric Tippmann, Jochen Knaus, Victor Jaravine, Frank Kramer, Martin Boeker</i>	
Computational Meme Understanding: A Survey	21251
<i>Khoi P. N. Nguyen, Vincent Ng</i>	
CoverICL: Selective Annotation for In-Context Learning Via Active Graph Coverage	21268
<i>Costas Mavromatis, Balasubramaniam Srinivasan, Zhengyuan Shen, Jiani Zhang, Huzefa Rangwala, Christos Faloutsos, George Karypis</i>	

VOLUME 28

Retrieval-Enriched Zero-Shot Image Classification in Low-Resource Domains	21287
<i>Nicola Dall’Asen, Yiming Wang, Enrico Fini, Elisa Ricci</i>	
I-AM-G: Interest Augmented Multimodal Generator for Item Personalization	21303
<i>Xianquan Wang, Likang Wu, Shukang Yin, Zhi Li, Yanjiang Chen, Hufeng Hufeng, Yu Su, Qi Liu</i>	

Twists, Humps, and Pebbles: Multilingual Speech Recognition Models Exhibit Gender Performance Gaps	21318
<i>Giuseppe Attanasio, Beatrice Savoldi, Dennis Fucci, Dirk Hovy</i>	
Enhancing Language Model Alignment: A Confidence-Based Approach to Label Smoothing	21341
<i>Baihe Huang, Hiteshi Sharma, Yi Mao</i>	
Contrastive Policy Gradient: Aligning LLMs on Sequence-Level Scores in a Supervised-Friendly Fashion	21353
<i>Yannis Flet-Berliac, Nathan Grinsztajn, Florian Strub, Eugene Choi, Bill Wu, Chris Cremer, Arash Ahmadian, Yash Chandak, Mohammad Gheshlaghi Azar, Olivier Pietquin, Matthieu Geist</i>	
Show and Guide: Instructional-Plan Grounded Vision and Language Model	21371
<i>Diogo Glória-Silva, David Semedo, Joao Magalhaes</i>	
Beyond Turn-Based Interfaces: Synchronous LLMs as Full-Duplex Dialogue Agents	21390
<i>Bandhav Veluri, Benjamin N Peloquin, Bokai Yu, Hongyu Gong, Shyamnath Gollakota</i>	
QuBE: Question-Based Belief Enhancement for Agentic LLM Reasoning	21403
<i>Minsoo Kim, Jongyoon Kim, Jihyuk Kim, Seung-Won Hwang</i>	
CompAct: Compressing Retrieved Documents Actively for Question Answering	21424
<i>Chanwoong Yoon, Taewhoo Lee, Hyeon Hwang, Minbyul Jeong, Jaewoo Kang</i>	
An Empirical Analysis on Spatial Reasoning Capabilities of Large Multimodal Models	21440
<i>Fatemeh Shiri, Xiao-Yu Guo, Mona Golestan Far, Xin Yu, Reza Haf, Yuan-Fang Li</i>	
Synthetic Knowledge Ingestion: Towards Knowledge Refinement and Injection for Enhancing Large Language Models	21456
<i>Jiaxin Zhang, Wendi Cui, Yiran Huang, Kamalika Das, Sricharan Kumar</i>	
Local Contrastive Editing of Gender Stereotypes	21474
<i>Marlene Lutz, Rochelle Choenni, Markus Strohmaier, Anne Lauscher</i>	
De-Identification of Sensitive Personal Data in Datasets Derived from IIT-CDIP	21494
<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, Kevin Leach</i>	
RAR: Retrieval-Augmented Retrieval for Code Generation in Low Resource Languages	21506
<i>Avik Dutta, Mukul Singh, Gust Verbruggen, Sumit Gulwani, Vu Le</i>	
STAR: SocioTechnical Approach to Red Teaming Language Models	21516
<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, William Isaac</i>	
Do Great Minds Think Alike? Investigating Human-AI Complementarity in Question Answering with CAIMIRA	21533
<i>Maharshi Gor, Hal Daumé III, Tianyi Zhou, Jordan Lee Boyd-Graber</i>	
Memory-Efficient Fine-Tuning of Transformers Via Token Selection	21565
<i>Antoine Simoulin, Namyong Park, Xiaoyi Liu, Grey Yang</i>	

Unveiling the Mystery of Visual Attributes of Concrete and Abstract Concepts: Variability, Nearest Neighbors, and Challenging Categories	21581
<i>Tarun Tater, Sabine Schulte Im Walde, Diego Frassinelli</i>	
Evaluating Large Language Models on Time Series Feature Understanding: A Comprehensive Taxonomy and Benchmark.....	21598
<i>Elizabeth Fons, Rachneet Kaur, Soham Palande, Zhen Zeng, Tucker Balch, Manuela Veloso, Svitlana Vyetrenko</i>	
Can LLMs Learn Uncertainty on Their Own? Expressing Uncertainty Effectively in a Self-Training Manner	21635
<i>Shudong Liu, Zhaocong Li, Xuebo Liu, Runzhe Zhan, Derek F. Wong, Lidia S. Chao, Min Zhang</i>	
Preference-Guided Reflective Sampling for Aligning Language Models	21646
<i>Hai Ye, Hwee Tou Ng</i>	
Metrics for What, Metrics for Whom: Assessing Actionability of Bias Evaluation Metrics in NLP	21669
<i>Pieter Delobelle, Giuseppe Attanasio, Debora Nozza, Su Lin Blodgett, Zeerak Talat</i>	
Is This the Real Life? is This Just Fantasy? the Misleading Success of Simulating Social Interactions with LLMs	21692
<i>Xuhui Zhou, Zhe Su, Tiwalayo Eisape, Hyunwoo Kim, Maarten Sap</i>	
A Simple LLM Framework for Long-Range Video Question-Answering	21715
<i>Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, Gedas Bertasius</i>	
Rebuilding ROME : Resolving Model Collapse During Sequential Model Editing	21738
<i>Akshat Gupta, Sidharth Baskaran, Gopala Anumanchipalli</i>	
Casablanca: Data and Models for Multidialectal Arabic Speech Recognition	21745
<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, Muhammad Abdul-Mageed</i>	
Safety Arithmetic: A Framework for Test-Time Safety Alignment of Language Models by Steering Parameters and Activations	21759
<i>Rima Hazra, Sayan Layek, Somnath Banerjee, Soujanya Poria</i>	
Communicating with Speakers and Listeners of Different Pragmatic Levels	21777
<i>Kata Naszadi, Frans A Oliehoek, Christof Monz</i>	
RECANTFormer: Referring Expression Comprehension with Varying Numbers of Targets	21784
<i>Bhathiya Hemanthage, Hakan Bilen, Phil Bartie, Christian Dondrup, Oliver Lemon</i>	
Sprout: Green Generative AI with Carbon-Efficient LLM Inference.....	21799
<i>Baolin Li, Yankai Jiang, Vijay Gadepally, Devesh Tiwari</i>	
Do LLMs Plan Like Human Writers? Comparing Journalist Coverage of Press Releases with LLMs.....	21814
<i>Alexander Spangher, Nanyun Peng, Sebastian Gehrmann, Mark Dredze</i>	

T-FREE: Subword Tokenizer-Free Generative LLMs Via Sparse Representations for Memory-Efficient Embeddings	21829
<i>Björn Deiseroth, Manuel Brack, Patrick Schramowski, Kristian Kersting, Samuel Weinbach</i>	
SpeechQE: Estimating the Quality of Direct Speech Translation.....	21852
<i>Hyojung Han, Kevin Duh, Marine Carpuat</i>	
Assessing and Verifying Task Utility in LLM-Powered Applications	21868
<i>Negar Arabzadeh, Siqing Huo, Nikhil Mehta, Qingyun Wu, Chi Wang, Ahmed Hassan Awadallah, Charles L. A. Clarke, Julia Kiseleva</i>	
Dynamic Rewarding with Prompt Optimization Enables Tuning-Free Self-Alignment of Language Models.....	21889
<i>Somanshu Singla, Zhen Wang, Tianyang Liu, Abdullah Ashfaq, Zhiting Hu, Eric P. Xing</i>	
Accurate and Data-Efficient Toxicity Prediction When Annotators Disagree	21910
<i>Harbani Jaggi, Kashyap Coimbatore Murali, Eve Fleisig, Erdem Biyik</i>	
Adversarial Text Generation Using Large Language Models for Dementia Detection	21918
<i>Youxiang Zhu, Nana Lin, Kiran Sandilya Balivada, Daniel Haehn, Xiaohui Liang</i>	
XCOMET-Lite: Bridging the Gap Between Efficiency and Quality in Learned MT Evaluation Metrics.....	21934
<i>Daniil Larionov, Mikhail Seleznyov, Vasilij Viskov, Alexander Panchenko, Steffen Eger</i>	
The Greatest Good Benchmark: Measuring LLMs' Alignment with Utilitarian Moral Dilemmas	21950
<i>Giovanni Franco Gabriel Marraffini, Andrés Cotton, Noe Fabian Hsueh, Axel Fridman, Juan Wisznia, Luciano Del Corro</i>	
FairFlow: Mitigating Dataset Biases Through Undecided Learning for Natural Language Understanding	21960
<i>Jiali Cheng, Hadi Amiri</i>	
Style-Shifting Behaviour of the Manosphere on Reddit.....	21976
<i>Jai Aggarwal, Suzanne Stevenson</i>	
The Death and Life of Great Prompts: Analyzing the Evolution of LLM Prompts from the Structural Perspective.....	21990
<i>Yihan Ma, Xinyue Shen, Yixin Wu, Boyang Zhang, Michael Backes, Yang Zhang</i>	
Holistic Evaluation for Interleaved Text-And-Image Generation.....	22002
<i>Minqian Liu, Zhiyang Xu, Zihao Lin, Trevor Ashby, Joy Rimchala, Jiaxin Zhang, Lifu Huang</i>	
FOLIO: Natural Language Reasoning with First-Order Logic.....	22017
<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, Dragomir Radev</i>	
The LLM Effect: Are Humans Truly Using LLMs, Or Are They Being Influenced by Them Instead?.....	22032
<i>Alexander Choi, Syeda Sabrina Akter, J. P. Singh, Antonios Anastasopoulos</i>	

VOLUME 29

Is Child-Directed Speech Effective Training Data for Language Models?	22055
<i>Steven Y. Feng, Noah Goodman, Michael Frank</i>	
RevMUX: Data Multiplexing with Reversible Adapters for Efficient LLM Batch Inference.....	22072
<i>Yige Xu, Xu Guo, Zhiwei Zeng, Chunyan Miao</i>	
Inference Helps PLMs' Conceptual Understanding: Improving the Abstract Inference Ability with Hierarchical Conceptual Entailment Graphs	22088
<i>Juncai Li, Ru Li, Xiaoli Li, Qinghua Chai, Jeff Z. Pan</i>	
M3Hop-CoT: Misogynous Meme Identification with Multimodal Multi-Hop Chain-Of-Thought.....	22105
<i>Gitanjali Kumari, Kirtan Jain, Asif Ekbal</i>	
GPT-4 Jailbreaks Itself with Near-Perfect Success Using Self-Explanation	22139
<i>Govind Ramesh, Yao Dou, Wei Xu</i>	
RE-RAG: Improving Open-Domain QA Performance and Interpretability with Relevance Estimator in Retrieval-Augmented Generation	22149
<i>Kiseung Kim, Jay-Yoon Lee</i>	
Evaluating Concurrent Robustness of Language Models Across Diverse Challenge Sets	22162
<i>Vatsal Gupta, Pranshu Pandya, Tushar Kataria, Vivek Gupta, Dan Roth</i>	
Simul-MuST-C: Simultaneous Multilingual Speech Translation Corpus Using Large Language Model	22185
<i>Mana Makinae, Yusuke Sakai, Hidetaka Kamigaito, Taro Watanabe</i>	
Is This a Bad Table? a Closer Look at the Evaluation of Table Generation from Text.....	22206
<i>Pritika Ramu, Aparna Garimella, Sambaran Bandyopadhyay</i>	
On the Fragility of Active Learners for Text Classification.....	22217
<i>Abhishek Ghose, Emma Thuong Nguyen</i>	
BMRetriever: Tuning Large Language Models as Better Biomedical Text Retrievers	22234
<i>Ran Xu, Wenqi Shi, Yue Yu, Yuchen Zhuang, Yanqiao Zhu, May Dongmei Wang, Joyce C. Ho, Chao Zhang, Carl Yang</i>	
Comparing Neighbors Together Makes it Easy: Jointly Comparing Multiple Candidates for Efficient and Effective Retrieval	22255
<i>Jonghyun Song, Cheyon Jin, Wenlong Zhao, Andrew McCallum, Jay-Yoon Lee</i>	
M3D: MultiModal MultiDocument Fine-Grained Inconsistency Detection	22270
<i>Chia-Wei Tang, Ting-Chih Chen, Kiet A. Nguyen, Kazi Sajeed Mehrab, Alvi Md Ishmam, Chris Thomas</i>	
MedAdapter: Efficient Test-Time Adaptation of Large Language Models Towards Medical Reasoning	22294
<i>Wenqi Shi, Ran Xu, Yuchen Zhuang, Yue Yu, Haotian Sun, Hang Wu, Carl Yang, May Dongmei Wang</i>	
EHRAgent: Code Empowers Large Language Models for Few-Shot Complex Tabular Reasoning on Electronic Health Records.....	22315
<i>Wenqi Shi, Ran Xu, Yuchen Zhuang, Yue Yu, Jieyu Zhang, Hang Wu, Yuanda Zhu, Joyce C. Ho, Carl Yang, May Dongmei Wang</i>	

SimLLM: Detecting Sentences Generated by Large Language Models Using Similarity Between the Generation and Its Re-Generation	22340
<i>Hoang-Quoc Nguyen-Son, Minh-Son Dao, Koji Zettsu</i>	
CELLO: Causal Evaluation of Large Vision-Language Models	22353
<i>Meiqi Chen, Bo Peng, Yan Zhang, Chaochao Lu</i>	
Simultaneous Interpretation Corpus Construction by Large Language Models in Distant Language Pair	22375
<i>Yusuke Sakai, Mana Makinae, Hidetaka Kamigaito, Taro Watanabe</i>	
Training-Free Deep Concept Injection Enables Language Models for Video Question Answering.....	22399
<i>Xudong Lin, Manling Li, Richard Zemel, Heng Ji, Shih-Fu Chang</i>	
MIBench: Evaluating Multimodal Large Language Models Over Multiple Images	22417
<i>Haowei Liu, Xi Zhang, Haiyang Xu, Yaya Shi, Chaoya Jiang, Ming Yan, Ji Zhang, Fei Huang, Chunfeng Yuan, Bing Li, Weiming Hu</i>	
ZEBRA: Zero-Shot Example-Based Retrieval Augmentation for Commonsense Question Answering	22429
<i>Francesco Maria Molfese, Simone Conia, Riccardo Orlando, Roberto Navigli</i>	
ABLE: Personalized Disability Support with Politeness and Empathy Integration	22445
<i>Kshitij Mishra, Manisha Burja, Asif Ekbal</i>	
Language Models as Compilers: Simulating Pseudocode Execution Improves Algorithmic Reasoning in Language Models	22471
<i>Hyungjoo Chae, Yeonghyeon Kim, Seungone Kim, Kai Tzu-Iunn Ong, Beong-Woo Kwak, Moohyeon Kim, Sunghwan Kim, Taeyoon Kwon, Jiwan Chung, Youngjae Yu, Jinyoung Yeo</i>	
Coffee-Gym: An Environment for Evaluating and Improving Natural Language Feedback on Erroneous Code	22503
<i>Hyungjoo Chae, Taeyoon Kwon, Seungjun Moon, Yongho Song, Dongjin Kang, Kai Tzu-Iunn Ong, Beong-Woo Kwak, Seonghyeon Bae, Seung-Won Hwang, Jinyoung Yeo</i>	
Improving Minimum Bayes Risk Decoding with Multi-Prompt.....	22525
<i>David Heineman, Yao Dou, Wei Xu</i>	
Deciphering Cognitive Distortions in Patient-Doctor Mental Health Conversations: A Multimodal LLM-Based Detection and Reasoning Framework	22546
<i>Gopendra Vikram Singh, Sai Vardhan Vemulapalli, Mauajama Firdaus, Asif Ekbal</i>	
Nearest Neighbor Normalization Improves Multimodal Retrieval.....	22571
<i>Neil Chowdhury, Franklin Wang, Sumedh Shenoy, Douwe Kiela, Sarah Schwettmann, Tristan Thrush</i>	
Rethinking Pragmatics in Large Language Models: Towards Open-Ended Evaluation and Preference Tuning.....	22583
<i>Shengguang Wu, Shusheng Yang, Zhenglun Chen, Qi Su</i>	
LongRAG: A Dual-Perspective Retrieval-Augmented Generation Paradigm for Long-Context Question Answering	22600
<i>Qingfei Zhao, Ruobing Wang, Yukuo Cen, Daren Zha, Shicheng Tan, Yuxiao Dong, Jie Tang</i>	
Context-Aware Watermark with Semantic Balanced Green-Red Lists for Large Language Models	22633
<i>Yuxuan Guo, Zhiliang Tian, Yiping Song, Tianlun Liu, Liang Ding, Dongsheng Li</i>	

Knowledge Graph Enhanced Large Language Model Editing	22647
<i>Mengqi Zhang, Xiaotian Ye, Qiang Liu, Pengjie Ren, Shu Wu, Zhumin Chen</i>	
‘Quis Custodiet Ipsos Custodes?’ Who Will Watch the Watchmen? on Detecting AI-Generated Peer-Reviews.....	22663
<i>Sandeep Kumar, Mohit Sahu, Vardhan Gacche, Tirthankar Ghosal, Asif Ekbal</i>	
Mitigating Open-Vocabulary Caption Hallucinations	22680
<i>Assaf Ben-Kish, Moran Yanuka, Morris Alper, Raja Giryes, Hadar Averbuch-Elor</i>	
Initialization of Large Language Models Via Reparameterization to Mitigate Loss Spikes	22699
<i>Kosuke Nishida, Kyosuke Nishida, Kuniko Saito</i>	
ALVIN: Active Learning Via INterpolation	22715
<i>Michalis Korakakis, Andreas Vlachos, Adrian Weller</i>	
Filtered Direct Preference Optimization.....	22729
<i>Tetsuro Morimura, Mitsuki Sakamoto, Yuu Jinnai, Kenshi Abe, Kaito Ariu</i>	
Instruction Fine-Tuning: Does Prompt Loss Matter?	22771
<i>Mathew Huerta-Enochian, Seung Yong Ko</i>	
Entity Insertion in Multilingual Linked Corpora: The Case of Wikipedia	22796
<i>Tomás Feith, Akhil Arora, Martin Gerlach, Debjit Paul, Robert West</i>	

Author Index