

Findings 2022

Findings of the Association for Computational Linguistics

Findings of ACL: EMNLP 2022

December 7-11, 2022

©2022 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 978-1-959429-43-2

Message from the Program Chairs

Welcome to EMNLP 2022, which is one of the most-attended conferences in the field of natural language processing, held in “hybrid” mode this year serving both virtual and in-person participants in Abu Dhabi!

Submission and Acceptance EMNLP 2022 received 4190 full paper submissions, the largest number to date. This number includes 275 ARR papers that were committed to EMNLP (see further discussion of ARR below). 225 papers were desk rejected for various reasons (missing limitation section, anonymity policy, multiple submission policy or formatting violations), leaving us with 3965 submissions that were fully reviewed. Despite the record-breaking number of submissions, based on the reviewers, area chairs and senior area chair comments, we kept the EMNLP 2022 acceptance rate similar to previous events, and accepted **829** papers to the main conference. Out of these, 175 are oral presentations and 654 poster presentations. Similarly to prior years, we also accepted 549 papers for “Findings of EMNLP”. The EMNLP 2022 program also features 39 papers from the Transactions of the Association for Computational Linguistics (TACL) and Computational Linguistics (CL) journals. More statistics on the accepted papers can be found below.

	Long	Short	Total
Submitted (Including ARR commits)	3242	948	4190
Accepted as Oral	163	12	175
Accepted as Poster	552	102	654
Acceptance Rate (main conference)	22%	12%	20%
Accepted to Findings	453	96	549
Acceptance Rate (Findings)	14%	10%	13%
Presented TACL papers	–	–	27
Presented CL papers	–	–	12

Limitations Section One innovation of EMNLP 2022 is the requirement that each submitted paper must include an explicitly named Limitations section, discussing the limitations of the work. This discussion does not count towards the page limit, and we asked reviewers to not use the mentioned limitations as reasons to reject the paper, unless there is a really good reason to. The effect was overall positive: most papers included a limitations section, and many of them were informative. We hope to see this requirement continue in the future conferences.

Tracks To ensure a smooth process, the submissions to EMNLP 2022 were divided into 26 tracks. The tracks mostly followed these of previous EMNLP conferences, reflecting the “standard” divisions in the field. We did however make the following changes: the “Machine Translation and Multilinguality” track was split into two separate tracks (“Machine Translation” and “Multilinguality”); the “Syntax: Tagging, Chunking and Parsing” track was renamed to “Syntax, Parsing and their Applications”; and we added 4 additional tracks, reflecting upcoming trends in the research landscape: Commonsense Reasoning; Language Models and Analysis of Language Models; Efficient Methods for NLP; and Semi-supervised and Weakly-supervised Methods. Finally, we also solicited papers for a “Theme Track”, discussing Open questions, major obstacles, and unresolved issues in NLP. Of the 26 tracks, the Resources and Evaluation, NLP Applications, Machine Learning for NLP, Dialogue and Interactive Systems, Language Modeling and Analysis of Language Models, Speech, Vision, Robotics, Multimodal Grounding and Information Extraction tracks were the most popular with over 200 submissions per track.

Program committee structure & reviewing Similar to prior NLP conferences, we adopted the hierarchical program committee structure, where for each area we invited between 1 to 5 Senior Area Chairs

(SACs), who worked with a team of Area Chairs (ACs), and an army of reviewers. We relied on statistics from prior years to estimate how many SACs, ACs and reviewers will be needed and ended up with 70 SACs and 297 ACs. For the reviewers, we used the reviewer lists from prior EMNLP conferences, solicited volunteer reviewers, and also invited all EMNLP 2022 authors to serve as reviewers. To this end, we attempted to recruit the most competent and matching reviewers by making a Google Form in reviewer recruitment, which was publicized through different channels and is linked to the author information page for all the submission authors to fill. We then provided the resulting information to the program committee for making reviewer assignments. This resulted in a reviewer pool of 4647 reviewers, of which 3828 reviewers were assigned at least one paper to review. For each submission, we assigned three reviewers and an AC. The initial paper assignment was made using an automatic algorithm that matches the abstracts with ACs/reviewers' past publication records, then the assignments were further refined by the SACs/PCs. We tried to avoid the Toronto matching score which had its limitations. In the end, most reviewers were assigned less than 6 papers, with a few reviewers working on 10 assignments and a large number of reviewers working on 1 assignment only. In the Google Form, we also asked whether the reviewer will volunteer for emergency review, which turned out to be very useful. We adapted the review forms from NAACL 2022, and ACL 2022 and EMNLP 2021. Besides the overall recommendation, reviewers were asked to evaluate how reproducible the results in the paper were, and whether there was any ethical concern. To ensure the review quality, we provided detailed guidelines about what reviewers should and shouldn't do in a review. We made final decisions according to the rankings and SAC recommendations. Our final decisions were made not just on the review scores, but also took into account the reviews, author responses, discussions among reviewers, meta-reviews and SAC/AC recommendations.

Ethics committee We also formed an Ethics Committee (EC) dedicated to ethical issues. After the technical reviews, but before the author-response and discussion phases, the ethics committee considered 150 papers that were flagged by the technical reviewing committee for ethical concerns. The two EC chairs went over the papers, to determine whether a full EC review would be required. As a result, 22 papers received two dedicated ethics reviews from a committee of 41 reviewers recruited by the EC chairs. An additional 22 papers raised one or more recurring issues (e.g., no information on annotator payment), which did not require a full ethics review, but were raised to the authors by the EC chairs. For any paper that was recommended to be accepted based on technical reviews and that had been referred to the EC, the EC chairs recommended one of the following to the PC chairs: (a) accept with comments (authors should take concerns raised in the ethical review into account in the camera-ready version; 29 papers); (b) conditionally accept (the ethical issues must be addressed in the camera-ready version, to be verified by the EC prior to final acceptance; 15 papers); and (c) reject due to ethical issues (0 papers). The authors of all conditionally accepted papers submitted the camera-ready version and a short response that explained how they had made the changes requested by the EC meta-reviews. The EC chairs double-checked these revised submissions and responses, and confirmed that the ethical concerns had been addressed. As a result, all conditionally accepted papers were accepted to the main conference or Findings. The EC chairs thank their committee for the excellent work.

ACL Rolling Review ACL Rolling Review (ARR) is an initiative of the Association for Computational Linguistics, where the reviewing and acceptance of papers to publication venues are done in a two-step process: (1) centralized rolling review and (2) the ability to commit the reviewed papers to be considered for publication by a publication venue. For EMNLP 2022, we decided to run a process which is separate from ARR, but allows for ARR submissions. Specifically, authors could either submit papers to EMNLP 2022 directly, or commit ARR reviewed papers by a certain date. We coordinated with the ARR team to extract the submission, review and meta-review from the OpenReview system, according to a submission link that the author provides when committing their ARR submission to EMNLP. The ARR commission deadline was set one month after the direct submission deadline since the ARR submissions already have their reviews and meta-recommendation. These ARR papers were then ranked by the SACs of the given tracks, together with the direct submissions in the track, and based on the reviews and meta-reviews

from ARR. Overall, EMNLP had 275 papers committed from ARR, of these 97 were accepted to the main conference and 73 were accepted to Findings of EMNLP.

Best paper selection Based on the nominations from SACs and ACs, we identified 11 candidates for the best papers and outstanding papers award. These papers are assessed by the Best Paper Award Committee. The award winners will be announced and present their works at the closing ceremony.

Presentation Mode We attempted the decision for oral vs poster presentations not to be made based on the quality/merit of the papers, but rather on the authors' interest in the presentation mode, and our understanding of what would be the best format for presentation of each individual paper.

Keynote talks Another highlight of our program is the three exciting keynote talks, presented by prof. Gary Marcus, NYU (Emeritus) on "Towards a Foundation for AGI"; prof. Neil Cohn, Tilburg University, Department of Communication and Cognition on "The multimodal language faculty and the visual languages of comics" & Dr. Mona Diab from Meta Responsible AI on "Towards a Responsible NLP: Walking the walk".

Gratitude We would like to thank the following people for their support & contributions:

- Our General Chair, Noah Smith, who led the whole organizing team, and helped with many of the decision processes;
- 70 SACs who helped us throughout the entire review process, from assigning papers, checking review quality, making final recommendation, suggesting presentation formats to recommending best paper candidates;
- 297 ACs who checked the initial submissions, led paper discussions, wrote meta reviews, ensured review quality and suggested best paper candidates;
- 3828 reviewers who reviewed the papers and actively participated in paper discussions; special thanks to those who stepped in at the last minute to serve as emergency reviewers;
- Reviewing Mentoring Chairs, David Mimno and Yanyan Lan and the team of mentors they recruited, for mentoring new reviewers.
- 36 Ethics Committee members, chaired by Lea Frermann and Margot Mieskes, for their hard work to provide ethical reviews and meta-reviews for all papers with serious ethical issues, and ensure that all the conditionally accepted papers have addressed the ethical issues appropriately;
- Best Paper Award Committee: Claire Cardie, Ellen Riloff, Hal Daume III, Rada Mihalcea, Raymond Mooney and Preslav Nakov for selecting the best papers;
- Jordan Boyd-Graber, the Virtual Posters chair, for organizing and managing the virtual poster sessions.
- Publication Chairs Ryan Cotterell, Steffen Eger and Sam Wiseman for completing the final proceedings within a short period;
- ACL Anthology Director Matt Post and his team, for his help in the production of the conference proceedings and maintenance of the ACL Anthology;
- TACL editor-in-chief Brian Roark and TACL Editorial Assistant Cindy Robinson, and CL Editor-in-Chief Hwee Tou Ng for coordinating TACL and CL presentations with us;

- The ARR team for their continued effort in running ARR, and for coordination with us. Particular thanks to Harold Rubio and Goran Glavas for multiple rounds of technical help in extracting data from the ARR OpenReview system.
- Rich Gerber at SoftConf, for setting up EMNLP 2022 conference site and quickly responding to our emails and resolving any issues we encountered with the START system;
- Website chairs Fajie Yuan and Zhaochun Ren and their team for continued effort in prompt updates to the website.
- Publicity chairs Eunsol Choi and Wajdi Zaghrouani for publicizing the conference and handling communications on social media.
- Sol Rosenberg, Damira Mršić and the whole Underline team, for helping to manage the logistics of both the virtual and online conference.
- Jenn Rachford and Priscilla Rasmussen for their professional and very valuable help in organizing the logistics of the conference.
- Nizar Habash and the rest of the Local Organizing Committee, for various discussions on organizing EMNLP, and making the local arrangements.
- 11854 authors for submitting their work to EMNLP 2022.

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

Yoav Goldberg, Bar Ilan University and Allen Institute for AI
Zornitsa Kozareva, SliceX AI
Yue Zhang, Westlake University
EMNLP 2022 Program Co-Chairs

Table of Contents

<i>LogicSolver: Towards Interpretable Math Word Problem Solving with Logical Prompt-enhanced Learning</i>	
Zhicheng Yang, Jinghui Qin, Jiaqi Chen, Liang Lin and Xiaodan Liang	1
<i>Commonsense Knowledge Saliency Evaluation with a Benchmark Dataset in E-commerce</i>	
Yincen Qu, Ningyu Zhang, Hui Chen, zelin Dai, Chengming Wang, Xiaoyu Wang, Qiang Chen and Huajun Chen	14
<i>Automatic Rule Induction for Efficient Semi-Supervised Learning</i>	
Reid Pryzant, Ziyi Yang, Yichong Xu, Chenguang Zhu and Michael Zeng	28
<i>Improving Semantic Matching through Dependency-Enhanced Pre-trained Model with Adaptive Fusion</i>	
Jian Song, Di Liang, Rumei Li, Yuntao Li, Sirui Wang, Minlong Peng, Wei Wu and Yongxin Yu	45
<i>Sparse Mixers: Combining MoE and Mixing to build a more efficient BERT</i>	
James Lee-Thorp and Joshua Ainslie	58
<i>KE-GCL: Knowledge Enhanced Graph Contrastive Learning for Commonsense Question Answering</i>	
Lihui Zhang and Ruifan Li	76
<i>Acceptability Judgements via Examining the Topology of Attention Maps</i>	
Daniil Cherniavskii, Eduard Tulchinskii, Vladislav Mikhailov, Irina Proskurina, Laida Kushnareva, Ekaterina Artemova, Serguei Barannikov, Irina Piontkovskaya, Dmitri Piontkovski and Evgeny Burnaev	88
<i>Clip-Tuning: Towards Derivative-free Prompt Learning with a Mixture of Rewards</i>	
Yekun Chai, Shuohuan Wang, Yu Sun, Hao Tian, Hua Wu and Haifeng Wang	108
<i>Soft-Labelled Contrastive Pre-Training for Function-Level Code Representation</i>	
Xiaonan Li, Daya Guo, Yeyun Gong, Yun Lin, Yelong Shen, Xipeng Qiu, Daxin Jiang, Weizhu Chen and Nan Duan	118
<i>Conditioned Masked Language and Image Modeling for Image-Text Dense Retrieval</i>	
Ziyang Luo, Yadong Xi, Rongsheng Zhang, GongZheng Li, Zeng Zhao and Jing Ma	130
<i>Does Simultaneous Speech Translation need Simultaneous Models?</i>	
Sara Papi, Marco Gaido, Matteo Negri and Marco Turchi	141
<i>Utilizing Language-Image Pretraining for Efficient and Robust Bilingual Word Alignment</i>	
Tuan Dinh, Jy-yong Sohn, Shashank Rajput, Timothy Ossowski, Yifei Ming, Junjie Hu, Dimitris Papailiopoulos and Kangwook Lee	154
<i>Grape: Knowledge Graph Enhanced Passage Reader for Open-domain Question Answering</i>	
Mingxuan Ju, Wenhao Yu, Tong Zhao, Chuxu Zhang and Yanfang Ye	169
<i>NarraSum: A Large-Scale Dataset for Abstractive Narrative Summarization</i>	
Chao Zhao, Faeze Brahman, Kaiqiang Song, Wenlin Yao, Dian Yu and Snigdha Chaturvedi .	182
<i>NMTScore: A Multilingual Analysis of Translation-based Text Similarity Measures</i>	
Jannis Vamvas and Rico Sennrich	198
<i>Language Models Understand Us, Poorly</i>	
Jared Moore	214

<i>Dialogue Meaning Representation for Task-Oriented Dialogue Systems</i>	
Xiangkun Hu, Junqi Dai, Hang Yan, Yi Zhang, Qipeng Guo, Xipeng Qiu and Zheng Zhang ..	223
<i>Learning from the Dictionary: Heterogeneous Knowledge Guided Fine-tuning for Chinese Spell Checking</i>	
Yinghui Li, Shirong Ma, Qingyu Zhou, Zhongli Li, Li Yangning, Shulin Huang, Ruiyang Liu, Chao Li, Yunbo Cao and Haitao Zheng	238
<i>Salient Phrase Aware Dense Retrieval: Can a Dense Retriever Imitate a Sparse One?</i>	
Xilun Chen, Kushal Lakhotia, Barlas Oguz, Anchit Gupta, Patrick Lewis, Stan Peshterliev, Yashar Mehdad, Sonal Gupta and Wen-tau Yih	250
<i>SMARTAVE: Structured Multimodal Transformer for Product Attribute Value Extraction</i>	
Qifan Wang, Li Yang, Jingang Wang, Jitin Krishnan, Bo Dai, Sinong Wang, Zenglin Xu, Madian Khabsa and Hao Ma	263
<i>When Language Model Meets Private Library</i>	
Daoguang Zan, Bei Chen, Zeqi Lin, Bei Guan, Wang Yongji and Jian-Guang LOU	277
<i>Cross-Domain Sentiment Classification using Semantic Representation</i>	
Shichen Li, Zhongqing Wang, Xiaotong Jiang and Guodong Zhou	289
<i>Yes-Yes-Yes: Proactive Data Collection for ACL Rolling Review and Beyond</i>	
Nils Dycke, Ilia Kuznetsov and Iryna Gurevych	300
<i>AssistSR: Task-oriented Video Segment Retrieval for Personal AI Assistant</i>	
Weixian Lei, DIFEI GAO, Yuxuan Wang, Dongxing Mao, Zihan Liang, Lingmin Ran and Mike Zheng Shou	319
<i>Dim-Krum: Backdoor-Resistant Federated Learning for NLP with Dimension-wise Krum-Based Aggregation</i>	
Zhiyuan Zhang, Qi Su and Xu Sun	339
<i>Fine-mixing: Mitigating Backdoors in Fine-tuned Language Models</i>	
Zhiyuan Zhang, Lingjuan Lyu, Xingjun Ma, Chenguang Wang and Xu Sun	355
<i>Language Models that Seek for Knowledge: Modular Search & Generation for Dialogue and Prompt Completion</i>	
Kurt Shuster, Mojtaba Komeili, Leonard Adolphs, Stephen Roller, Arthur Szlam and Jason Weston	373
<i>Stretching Sentence-pair NLI Models to Reason over Long Documents and Clusters</i>	
Tal Schuster, Sihao Chen, Senaka Buthpitiya, Alex Fabrikant and Donald Metzler	394
<i>Towards Realistic Low-resource Relation Extraction: A Benchmark with Empirical Baseline Study</i>	
Xin Xu, Xiang Chen, Ningyu Zhang, Xin Xie, Xi Chen and Huajun Chen	413
<i>CLLE: A Benchmark for Continual Language Learning Evaluation in Multilingual Machine Translation</i>	
Han Zhang, Sheng Zhang, Yang Xiang, Bin Liang, Jinsong Su, Zhongjian Miao, Hui Wang and Ruifeng Xu	428
<i>Lexicon-Enhanced Self-Supervised Training for Multilingual Dense Retrieval</i>	
Houxing Ren, Linjun Shou, Jian Pei, Ning Wu, Ming Gong and Daxin Jiang	444
<i>Improve Interpretability of Neural Networks via Sparse Contrastive Coding</i>	
Junhong Liu, Yijie Lin, Liang Jiang, Jia Liu, Zujie Wen and Xi Peng	460

<i>LEMON: Language-Based Environment Manipulation via Execution-Guided Pre-training</i> Qi Shi, Qian Liu, Bei Chen, Yu Zhang, Ting Liu and Jian-Guang LOU	471
<i>CROP: Zero-shot Cross-lingual Named Entity Recognition with Multilingual Labeled Sequence Translation</i> Jian Yang, Shaohan Huang, Shuming Ma, Yuwei Yin, Li Dong, Dongdong Zhang, Hongcheng Guo, Zhoujun Li and Furu Wei	486
<i>Handling and Presenting Harmful Text in NLP Research</i> Hannah Kirk, Abeba Birhane, Bertie Vidgen and Leon Derczynski	497
<i>Multimodal Contrastive Learning via Uni-Modal Coding and Cross-Modal Prediction for Multimodal Sentiment Analysis</i> Ronghao Lin and Haifeng Hu	511
<i>Towards Unified Prompt Tuning for Few-shot Text Classification</i> Jianing Wang, Chengyu Wang, Fuli Luo, Chuanqi Tan, Minghui Qiu, Fei Yang, Qiuhui Shi, Songfang Huang and Ming Gao	524
<i>Can language models learn from explanations in context?</i> Andrew Lampinen, Ishita Dasgupta, Stephanie Chan, Kory Mathewson, MH Tessler, Antonia Creswell, James McClelland, Jane Wang and Felix Hill	537
<i>GNN-encoder: Learning a Dual-encoder Architecture via Graph Neural Networks for Dense Passage Retrieval</i> Jiduan Liu, Jiahao Liu, Yang Yang, Jingang Wang, Wei Wu, Dongyan Zhao and Rui Yan	564
<i>Linguistic Rules-Based Corpus Generation for Native Chinese Grammatical Error Correction</i> Shirong Ma, Yinghui Li, Rongyi Sun, Qingyu Zhou, Shulin Huang, Ding Zhang, Li Yangning, Ruiyang Liu, Zhongli Li, Yunbo Cao, Haitao Zheng and Ying Shen	576
<i>Rethinking the Video Sampling and Reasoning Strategies for Temporal Sentence Grounding</i> Jiahao Zhu, Daizong Liu, Pan Zhou, Xing Di, Yu Cheng, Song Yang, Wenzheng Xu, Zichuan Xu, Yao Wan, Lichao Sun and Zeyu Xiong	590
<i>System 1 + System 2 = Better World: Neural-Symbolic Chain of Logic Reasoning</i> Wenyue Hua and Yongfeng Zhang	601
<i>Efficient Federated Learning on Knowledge Graphs via Privacy-preserving Relation Embedding Aggregation</i> Kai Zhang, Yu Wang, Hongyi Wang, Lifu Huang, Carl Yang, Xun Chen and Lichao Sun	613
<i>TextHacker: Learning based Hybrid Local Search Algorithm for Text Hard-label Adversarial Attack</i> Zhen Yu, Xiaosen Wang, Wanxiang Che and Kun He	622
<i>Visualizing the Obvious: A Concreteness-based Ensemble Model for Noun Property Prediction</i> Yue Yang, Artemis Panagopoulou, Marianna Apidianaki, Mark Yatskar and Chris Callison-Burch	638
<i>It's Better to Teach Fishing than Giving a Fish: An Auto-Augmented Structure-aware Generative Model for Metaphor Detection</i> Huawen Feng and Qianli Ma	656
<i>Expose Backdoors on the Way: A Feature-Based Efficient Defense against Textual Backdoor Attacks</i> Sishuo Chen, Wenkai Yang, Zhiyuan Zhang, Xiaohan Bi and Xu Sun	668

<i>Diving Deep into Modes of Fact Hallucinations in Dialogue Systems</i> Souvik Das, Sougata Saha and Rohini Srihari	684
<i>Representation Learning for Resource-Constrained Keyphrase Generation</i> Di Wu, Wasi Ahmad, Sunipa Dev and Kai-Wei Chang	700
<i>Systematicity in GPT-3’s Interpretation of Novel English Noun Compounds</i> Siyun Li, Riley Carlson and Christopher Potts	717
<i>CARE: Causality Reasoning for Empathetic Responses by Conditional Graph Generation</i> Jiashuo WANG, Yi Cheng and Wenjie Li	729
<i>TransAdv: A Translation-based Adversarial Learning Framework for Zero-Resource Cross-Lingual Named Entity Recognition</i> Yichun Zhao, Jintao Du, Gongshen Liu and Huijia Zhu	742
<i>BARLE: Background-Aware Representation Learning for Background Shift Out-of-Distribution Detection</i> Hanyu Duan, Yi Yang, Ahmed Abbasi and Kar Yan Tam	750
<i>What Language Model to Train if You Have One Million GPU Hours?</i> Teven Le Scao, Thomas Wang, Daniel Hesslow, Stas Bekman, M Saiful Bari, Stella Biderman, Hady Elsahar, Niklas Muennighoff, Jason Phang, Ofir Press, Colin Raffel, Victor Sanh, Sheng Shen, Lintang Sutawika, Jaesung Tae, Zheng Xin Yong, Julien Launay and Iz Beltagy	765
<i>Enhancing Out-of-Distribution Detection in Natural Language Understanding via Implicit Layer Ensemble</i> Hyunsoo Cho, Choonghyun Park, Jaewook Kang, Kang Min Yoo, Taek Kim and Sang-goo Lee	783
<i>Contrastive Demonstration Tuning for Pre-trained Language Models</i> Xiaozhuan Liang, Ningyu Zhang, Siyuan Cheng, Zhenru Zhang, Chuanqi Tan and Huajun Chen	799
<i>Detect-Localize-Repair: A Unified Framework for Learning to Debug with CodeT5</i> Nghie Bui, Yue Wang and Steven C.H. Hoi	812
<i>Influence Functions for Sequence Tagging Models</i> Sarathak Jain, Varun Manjunatha, Byron Wallace and Ani Nenkova	824
<i>Impact of Pretraining Term Frequencies on Few-Shot Numerical Reasoning</i> Yasaman Razeghi, Robert L Logan IV, Matt Gardner and Sameer Singh	840
<i>Syntactic and Semantic Uniformity for Semantic Parsing and Task-Oriented Dialogue Systems</i> Bowen Chen and Yusuke Miyao	855
<i>Knowledge-Rich Self-Supervision for Biomedical Entity Linking</i> Sheng Zhang, Hao Cheng, Shikhar Vashishth, Cliff Wong, Jinfeng Xiao, Xiaodong Liu, Tristan Naumann, Jianfeng Gao and Hoifung Poon	868
<i>ARTIST: A Transformer-based Chinese Text-to-Image Synthesizer Digesting Linguistic and World Knowledge</i> Tingting Liu, Chengyu Wang, Xiangru Zhu, Lei Li, Minghui Qiu, Jun Huang, Ming Gao and Yanghua Xiao	881
<i>From Spelling to Grammar: A New Framework for Chinese Grammatical Error Correction</i> Xiuyu Wu and Yunfang Wu	889

<i>Language Models Are Poor Learners of Directional Inference</i>	
Tianyi Li, Mohammad Javad Hosseini, Sabine Weber and Mark Steedman	903
<i>Wish I Can Feel What You Feel: A Neural Approach for Empathetic Response Generation</i>	
Yangbin Chen and Chunfeng Liang	922
<i>Measuring and Improving Semantic Diversity of Dialogue Generation</i>	
Seungju Han, Beomsu Kim and Buru Chang	934
<i>Plug-and-Play VQA: Zero-shot VQA by Conjoining Large Pretrained Models with Zero Training</i>	
Anthony Meng Huat Tiong, Junnan Li, Boyang Li, Silvio Savarese and Steven C.H. Hoi	951
<i>TSGP: Two-Stage Generative Prompting for Unsupervised Commonsense Question Answering</i>	
Yueqing Sun, Yu Zhang, Le Qi and Qi Shi	968
<i>Subword-Delimited Downsampling for Better Character-Level Translation</i>	
Lukas Edman, Antonio Toral and Gertjan van Noord	981
<i>Autoregressive Structured Prediction with Language Models</i>	
Tianyu Liu, Yuchen Eleanor Jiang, Nicholas Monath, Ryan Cotterell and Mrinmaya Sachan	993
<i>XDoc: Unified Pre-training for Cross-Format Document Understanding</i>	
Jingye Chen, Tengchao Lv, Lei Cui, Cha Zhang and Furu Wei	1006
<i>A Few More Examples May Be Worth Billions of Parameters</i>	
Yuval Kirstain, Patrick Lewis, Sebastian Riedel and Omer Levy	1017
<i>MCP: Self-supervised Pre-training for Personalized Chatbots with Multi-level Contrastive Sampling</i>	
Zhaoheng Huang, Zhicheng Dou, Yutao Zhu and Zhengyi Ma	1030
<i>ExpertPLM: Pre-training Expert Representation for Expert Finding</i>	
Qiyao Peng and Hongtao Liu	1043
<i>You Truly Understand What I Need : Intellectual and Friendly Dialog Agents grounding Persona and Knowledge</i>	
Jungwoo Lim, Myugnhoon Kang, Yuna Hur, Seung Won Jeong, Jinsung Kim, Yoonna Jang, Dongyub Lee, Hyesung Ji, DongHoon Shin, Seungryong Kim and Heuseok Lim	1053
<i>Faithful to the Document or to the World? Mitigating Hallucinations via Entity-Linked Knowledge in Abstractive Summarization</i>	
Yue Dong, John Wieting and Pat Verga	1067
<i>RL with KL penalties is better viewed as Bayesian inference</i>	
Tomasz Korbak, Ethan Perez and Christopher Buckley	1083
<i>Evaluating Token-Level and Passage-Level Dense Retrieval Models for Math Information Retrieval</i>	
Wei Zhong, Jheng-Hong Yang, YUQING XIE and Jimmy Lin	1092
<i>Multi-View Reasoning: Consistent Contrastive Learning for Math Word Problem</i>	
Wenqi Zhang, Yongliang Shen, Yanna Ma, Xiaoxia Cheng, Zeqi Tan, Qingpeng Nong and Weiming Lu	1103
<i>Few-shot initializing of Active Learner via Meta-Learning</i>	
Zi Long Zhu, Vikrant Yadav, Zubair Afzal and George Tsatsaronis	1117
<i>Bootstrapping meaning through listening: Unsupervised learning of spoken sentence embeddings</i>	
Jian Zhu, Zuoyu Tian, Yadong Liu, Cong Zhang and Chia-Wen Lo	1134

<i>Progressive Sentiment Analysis for Code-Switched Text Data</i> Sudhanshu Ranjan, Dheeraj Mekala and Jingbo Shang	1155
<i>Knowledge Stimulated Contrastive Prompting for Low-Resource Stance Detection</i> Kai Zheng, Qingfeng Sun, Yaming Yang and Fei Xu	1168
<i>WSpeller: Robust Word Segmentation for Enhancing Chinese Spelling Check</i> Fangfang Li, Youran Shan, Junwen Duan, Xingliang Mao and Minlie Huang	1179
<i>Extracting Trigger-sharing Events via an Event Matrix</i> Jun Xu, Weidi Xu, Mengshu Sun, Taifeng Wang and Wei Chu	1189
<i>TranS: Transition-based Knowledge Graph Embedding with Synthetic Relation Representation</i> Xuanyu Zhang, Qing Yang and Dongliang Xu	1202
<i>Sequential Topic Selection Model with Latent Variable for Topic-Grounded Dialogue</i> Xiao-Fei Wen, Wei Wei and Xian-Ling Mao	1209
<i>Robust Task-Oriented Dialogue Generation with Contrastive Pre-training and Adversarial Filtering</i> Shiquan Yang, Xinting Huang, Jey Han Lau and Sarah Erfani	1220
<i>STAR: SQL Guided Pre-Training for Context-dependent Text-to-SQL Parsing</i> Zefeng Cai, Xiangyu Li, Binyuan Hui, Min Yang, Bowen Li, Binhua Li, Zheng Cao, weijie li, Fei Huang, Luo Si and Yongbin Li	1235
<i>Is MultiWOZ a Solved Task? An Interactive TOD Evaluation Framework with User Simulator</i> Qinyuan Cheng, Linyang Li, Guofeng Quan, Feng Gao, xiaofeng mou and Xipeng Qiu	1248
<i>Translating Hanja Historical Documents to Contemporary Korean and English</i> Juhee Son, Jiho Jin, Haneul Yoo, JinYeong Bak, Kyunghyun Cho and Alice Oh	1260
<i>Exploring Compositional Image Retrieval with Hybrid Compositional Learning and Heuristic Negative Mining</i> Chao Wang, Ehsan Nezhadarya, Tanmana Sadhu and Shengdong Zhang	1273
<i>Outlier Dimensions that Disrupt Transformers are Driven by Frequency</i> Giovanni Puccetti, Anna Rogers, Aleksandr Drozd and Felice Dell’Orletta	1286
<i>MiST: a Large-Scale Annotated Resource and Neural Models for Functions of Modal Verbs in English Scientific Text</i> Sophie Henning, Nicole Macher, Stefan Gr̄unewald and Annemarie Friedrich	1305
<i>Late Prompt Tuning: A Late Prompt Could Be Better Than Many Prompts</i> Xiangyang Liu, Tianxiang Sun, Xuanjing Huang and Xipeng Qiu	1325
<i>MICO: A Multi-alternative Contrastive Learning Framework for Commonsense Knowledge Representation</i> Ying Su, Zihao Wang, Tianqing Fang, Hongming Zhang, Yangqiu Song and Tong Zhang ...	1339
<i>Leveraging Only the Category Name for Aspect Detection through Prompt-based Constrained Clustering</i> Yazheng Li, Pengyun Wang, Yasheng Wang, Yong Dai, Yadao Wang, Lujia Pan and Zenglin Xu	1352
<i>Controllable Factuality in Document-Grounded Dialog Systems Using a Noisy Channel Model</i> Nico Daheim, David Thulke, Christian Dugast and Hermann Ney	1365

<i>Transformer Language Models without Positional Encodings Still Learn Positional Information</i> Adi Haviv, Ori Ram, Ofir Press, Peter Izsak and Omer Levy	1382
<i>Beyond Model Interpretability: On the Faithfulness and Adversarial Robustness of Contrastive Textual Explanations</i> Julia El Zini and Mariette Awad	1391
<i>How Much Does Attention Actually Attend? Questioning the Importance of Attention in Pretrained Transformers</i> Michael Hassid, Hao Peng, Daniel Rotem, Jungo Kasai, Ivan Montero, Noah A. Smith and Roy Schwartz	1403
<i>What Has Been Enhanced in my Knowledge-Enhanced Language Model?</i> Yifan Hou, Guoji Fu and Mrinmaya Sachan	1417
<i>Towards Generalized Open Information Extraction</i> Bowen Yu, Zhenyu Zhang, Jingyang Li, Haiyang Yu, Tingwen Liu, Jian Sun, Yongbin Li and Bin Wang	1439
<i>BioLORD: Learning Ontological Representations from Definitions for Biomedical Concepts and their Textual Descriptions</i> François Remy, Kris Demuynck and Thomas Demeester	1454
<i>Improving the Extraction of Supertags for Constituency Parsing with Linear Context-Free Rewriting Systems</i> Thomas Ruprecht	1466
<i>Mask More and Mask Later: Efficient Pre-training of Masked Language Models by Disentangling the [MASK] Token</i> Baohao Liao, David Thulke, Sanjika Hewavitharana, Hermann Ney and Christof Monz	1478
<i>SMSMix: Sense-Maintained Sentence Mixup for Word Sense Disambiguation</i> Hee Suk Yoon, Eunseop Yoon, John Harvill, Sunjae Yoon, Mark Hasegawa-Johnson and Chang Yoo	1493
<i>On the Effectiveness of Automated Metrics for Text Generation Systems</i> Pius von Däniken, Jan Deriu, Don Tuggener and Mark Cieliebak	1503
<i>Residual Learning of Neural Text Generation with n-gram Language Model</i> Huayang Li, Deng Cai, Jin Xu and Taro Watanabe	1523
<i>DiffG-RL: Leveraging Difference between Environment State and Common Sense</i> Tsunehiko Tanaka, Daiki Kimura and Michiaki Tatsubori	1534
<i>Unsupervised Syntactically Controlled Paraphrase Generation with Abstract Meaning Representations</i> Kuan-Hao Huang, Varun Iyer, Anoop Kumar, Sriram Venkatapathy, Kai-Wei Chang and Aram Galstyan	1547
<i>Can AMR Assist Legal and Logical Reasoning?</i> Nikolaus Schrack, Ruixiang Cui, Hugo López and Daniel Hershcovich	1555
<i>Data Selection Curriculum for Neural Machine Translation</i> Tasnim Mohiuddin, Philipp Koehn, Vishrav Chaudhary, James Cross, Shruti Bhosale and Shafiq Joty	1569
<i>Text Editing as Imitation Game</i> Ning Shi, Bin Tang, Bo Yuan, Longtao Huang, Yewen Pu, Jie Fu and Zhouhan Lin	1583

<i>Seeded Hierarchical Clustering for Expert-Crafted Taxonomies</i>	
Anish Saha, Amith Ananthram, Emily Allaway, Heng Ji and Kathleen McKeown	1595
<i>Knowledge Graph Generation From Text</i>	
Igor Melnyk, Pierre Dognin and Payel Das	1610
<i>DialogueGAT: A Graph Attention Network for Financial Risk Prediction by Modeling the Dialogues in Earnings Conference Calls</i>	
Yunxin Sang and Yang Bao	1623
<i>Investigating Ensemble Methods for Model Robustness Improvement of Text Classifiers</i>	
Jieyu Zhao, Xuezhi Wang, Yao Qin, Jilin Chen and Kai-Wei Chang	1634
<i>Adaptive Ranking-based Sample Selection for Weakly Supervised Class-imbalanced Text Classification</i>	
Linxin Song, Jieyu Zhang, Tianxiang Yang and Masayuki Goto	1641
<i>ComFact: A Benchmark for Linking Contextual Commonsense Knowledge</i>	
Silin Gao, Jena D. Hwang, Saya Kanno, Hiromi Wakaki, Yuki Mitsufuji and Antoine Bosselut	1656
<i>Learning to Perform Complex Tasks through Compositional Fine-Tuning of Language Models</i>	
Victor Bursztny, David Demeter, Doug Downey and Larry Birnbaum	1676
<i>Topic Taxonomy Expansion via Hierarchy-Aware Topic Phrase Generation</i>	
Dongha Lee, Jiaming Shen, Seonghyeon Lee, Susik Yoon, Hwanjo Yu and Jiawei Han	1687
<i>Language as a fingerprint: Self-supervised learning of user encodings using transformers</i>	
Roberta Rocca and Tal Yarkoni	1701
<i>Hyperdecoders: Instance-specific decoders for multi-task NLP</i>	
Hamish Ivison and Matthew Peters	1715
<i>Evaluating the Faithfulness of Importance Measures in NLP by Recursively Masking Allegedly Important Tokens and Retraining</i>	
Andreas Madsen, Nicholas Meade, Vaibhav Adlakha and Siva Reddy	1731
<i>Towards Explaining Subjective Ground of Individuals on Social Media</i>	
Younghun Lee and Dan Goldwasser	1752
<i>Knowledge Injected Prompt Based Fine-tuning for Multi-label Few-shot ICD Coding</i>	
Zhichao Yang, Shufan Wang, Bhanu Pratap Singh Rawat, Avijit Mitra and hong yu	1767
<i>Do Language Models Understand Measurements?</i>	
Sungjin Park, Seungwoo Ryu and Edward Choi	1782
<i>Reconciliation of Pre-trained Models and Prototypical Neural Networks in Few-shot Named Entity Recognition</i>	
Youcheng Huang, Wenqiang Lei, Jie Fu and Jiancheng Lv	1793
<i>HCL-TAT: A Hybrid Contrastive Learning Method for Few-shot Event Detection with Task-Adaptive Threshold</i>	
Ruihan Zhang, Wei Wei, Xian-Ling Mao, rui fang and Dangyang Chen	1808
<i>Doc2Bot: Accessing Heterogeneous Documents via Conversational Bots</i>	
Haomin Fu, Yeqin Zhang, Haiyang Yu, Jian Sun, Fei Huang, Luo Si, Yongbin Li and Cam Tu Nguyen	1820

<i>DualNER: A Dual-Teaching framework for Zero-shot Cross-lingual Named Entity Recognition</i>	
Jiali Zeng, Yufan Jiang, Yongjing Yin, Xu Wang, Binghuai Lin and Yunbo Cao	1837
<i>Knowledge-augmented Self-training of A Question Rewriter for Conversational Knowledge Base Question Answering</i>	
Xirui Ke, Jing Zhang, Xin Lv, Yiqi Xu, Shulin Cao, Cuiping Li, Hong Chen and Juanzi Li .	1844
<i>Extractive Summarization of Legal Decisions using Multi-task Learning and Maximal Marginal Relevance</i>	
Abhishek Agarwal, Shanshan Xu and Matthias Grabmair	1857
<i>MovieUN: A Dataset for Movie Understanding and Narrating</i>	
Qi Zhang, Zihao Yue, Anwen Hu, Ziheng Wang and Qin Jin.....	1873
<i>ASDOT: Any-Shot Data-to-Text Generation with Pretrained Language Models</i>	
Jiannan Xiang, Zhengzhong Liu, Yucheng Zhou, Eric Xing and Zhiting Hu	1886
<i>FCGEC: Fine-Grained Corpus for Chinese Grammatical Error Correction</i>	
Lvxiaowei Xu, Jianwang Wu, Jiawei Peng, Jiayu Fu and Ming Cai	1900
<i>Audience-Centric Natural Language Generation via Style Infusion</i>	
Samraj Moorjani, Adit Krishnan, Hari Sundaram, Ewa Maslowska and Aravind Sankar ...	1919
<i>DocFin: Multimodal Financial Prediction and Bias Mitigation using Semi-structured Documents</i>	
Puneet Mathur, Mihir Goyal, Ramit Sawhney, Ritik Mathur, Jochen Leidner, Franck Dernoncourt and Dinesh Manocha	1933
<i>Not Just Plain Text! Fuel Document-Level Relation Extraction with Explicit Syntax Refinement and Subsentence Modeling</i>	
Zhichao Duan, Xiuxing Li, Zhenyu Li, Zhuo Wang and Jianyong Wang	1941
<i>Self-supervised Rewiring of Pre-trained Speech Encoders: Towards Faster Fine-tuning with Less Labels in Speech Processing</i>	
Hao Yang, Jinming Zhao, Gholamreza Haffari and Ehsan Shareghi.....	1952
<i>RedApt: An Adaptor for wav2vec 2 Encoding</i>	
<i>Faster and Smaller Speech Translation without Quality Compromise</i>	
Jinming Zhao, Hao Yang, Gholamreza Haffari and Ehsan Shareghi.....	1960
<i>How sensitive are translation systems to extra contexts? Mitigating gender bias in Neural Machine Translation models through relevant contexts.</i>	
Shanya Sharma, Manan Dey and Koustuv Sinha	1968
<i>PM²F²N: Patient Multi-view Multi-modal Feature Fusion Networks for Clinical Outcome Prediction</i>	
Ying Zhang, Baohang Zhou, Kehui Song, Xuhui Sui, Guoqing Zhao, Ning Jiang and Xiaojie Yuan	1985
<i>Long Text and Multi-Table Summarization: Dataset and Method</i>	
Shuaiqi LIU, Jiannong Cao, Ruosong Yang and Zhiyuan Wen	1995
<i>MatRank: Text Re-ranking by Latent Preference Matrix</i>	
Jinwen Luo, Jiuding Yang, Weidong Guo, Chenglin Li, Di Niu and Yu Xu	2011
<i>Can Language Models Serve as Temporal Knowledge Bases?</i>	
Ruilin Zhao, Feng Zhao, Guandong Xu, Sixiao Zhang and Hai Jin	2024
<i>Are Large Pre-Trained Language Models Leaking Your Personal Information?</i>	
Jie Huang, Hanyin Shao and Kevin Chen-Chuan Chang.....	2038

<i>Self-Distillation with Meta Learning for Knowledge Graph Completion</i> Yunshui Li, Junhao Liu, Min Yang and Chengming Li	2048
<i>CQR-SQL: Conversational Question Reformulation Enhanced Context-Dependent Text-to-SQL Parsers</i> Dongling Xiao, LinZheng Chai, Qian-Wen Zhang, Zhao Yan, Zhoujun Li and Yunbo Cao ..	2055
<i>Assisting the Human Fact-Checkers: Detecting All Previously Fact-Checked Claims in a Document</i> Shaden Shaar, Nikola Georgiev, Firoj Alam, Giovanni Da San Martino, Aisha Mohamed and Preslav Nakov	2069
<i>No Word Embedding Model Is Perfect: Evaluating the Representation Accuracy for Social Bias in the Media</i> Maximilian Spliethöver, Maximilian Keiff and Henning Wachsmuth	2081
<i>Scientific and Creative Analogies in Pretrained Language Models</i> Tamara Czinczoll, Helen Yannakoudakis, Pushkar Mishra and Ekaterina Shutova	2094
<i>Bitext Mining Using Distilled Sentence Representations for Low-Resource Languages</i> Kevin Heffernan, Onur Çelebi and Holger Schwenk	2101
<i>Towards Generalizable and Robust Text-to-SQL Parsing</i> Chang Gao, Bowen Li, Wenxuan Zhang, Wai Lam, Binhua Li, Fei Huang, Luo Si and Yongbin Li 2113	
<i>EdiT5: Semi-Autoregressive Text Editing with T5 Warm-Start</i> Jonathan Mallinson, Jakub Adamek, Eric Malmi and Aliaksei Severyn	2126
<i>A Critical Reflection and Forward Perspective on Empathy and Natural Language Processing</i> Allison Lahkala, Charles Welch, David Jurgens and Lucie Flek	2139
<i>A Neural-Symbolic Approach to Natural Language Understanding</i> Zhixuan Liu, Zihao Wang, Yuan Lin and Hang Li	2159
<i>Social-aware Sparse Attention Network for Session-based Social Recommendation</i> Kai Ouyang, Xianghong Xu, Chen Tang, Wang Chen and Haitao Zheng	2173
<i>SparseAdapter: An Easy Approach for Improving the Parameter-Efficiency of Adapters</i> Shwai He, Liang Ding, Daize Dong, Jeremy Zhang and Dacheng Tao	2184
<i>Measurement Extraction with Natural Language Processing: A Review</i> Jan Göpfert, Patrick Kuckertz, Jann Weinand, Leander Kotzur and Detlef Stolten	2191
<i>Summarizing Procedural Text: Data and Approach</i> Shen Gao, Haotong Zhang, Xiuying Chen, Rui Yan and Dongyan Zhao	2216
<i>Snapshot-Guided Domain Adaptation for ELECTRA</i> Daixuan Cheng, Shaohan Huang, Jianfeng Liu, Yuefeng Zhan, Hao Sun, Furu Wei, Denvy Deng and Qi Zhang	2226
<i>Exploiting Labeled and Unlabeled Data via Transformer Fine-tuning for Peer-Review Score Prediction</i> Panitan Muangkammuen, Fumiyo Fukumoto, Jiyi Li and Yoshimi Suzuki	2233
<i>HARALD: Augmenting Hate Speech Data Sets with Real Data</i> Tal Ilan and Dan Vilenchik	2241
<i>Wait-info Policy: Balancing Source and Target at Information Level for Simultaneous Machine Tran- slation</i> Shaolei Zhang, Shoutao Guo and Yang Feng	2249

<i>Turning Fixed to Adaptive: Integrating Post-Evaluation into Simultaneous Machine Translation</i> Shoutao Guo, Shaolei Zhang and Yang Feng	2264
<i>Alleviating Sparsity of Open Knowledge Graphs with Ternary Contrastive Learning</i> Qian Li, Shafiq Joty, Daling Wang, Shi Feng and Yifei Zhang	2279
<i>Using Developer Discussions to Guide Fixing Bugs in Software</i> Sheena Panthaplackel, Milos Gligoric, Junyi Jessy Li and Raymond Mooney	2292
<i>AutoCAD: Automatically Generate Counterfactuals for Mitigating Shortcut Learning</i> Jiaxin Wen, Yeshuang Zhu, Jinchao Zhang, Jie Zhou and Minlie Huang	2302
<i>A Multi-Modal Knowledge Graph for Classical Chinese Poetry</i> Yuqing Li, Yuxin Zhang, Bin Wu, Ji-Rong Wen, Ruihua Song and Ting Bai	2318
<i>Assessing Non-autoregressive Alignment in Neural Machine Translation via Word Reordering</i> Chun-Hin Tse, Ester Leung and William K. Cheung	2327
<i>Syntax-guided Localized Self-attention by Constituency Syntactic Distance</i> Shengyuan Hou, Jushi Kai, Haotian Xue, Bingyu Zhu, Bo Yuan, Longtao Huang, Xinbing Wang and Zhouhan Lin	2334
<i>CodeExp: Explanatory Code Document Generation</i> Haotian Cui, Chenglong Wang, Junjie Huang, Jeevana Priya Inala, Todd Mytkowicz, Bo Wang, Jianfeng Gao and Nan Duan	2342
<i>PAUQ: Text-to-SQL in Russian</i> Daria Bakshandaeva, Oleg Somov, Ekaterina Dmitrieva, Vera Davydova and Elena Tutubalina	2355
<i>Event-Centric Question Answering via Contrastive Learning and Invertible Event Transformation</i> Junru Lu, Xingwei Tan, Gabriele Pergola, Lin Gui and Yulan He	2377
<i>Label-Driven Denoising Framework for Multi-Label Few-Shot Aspect Category Detection</i> Fei Zhao, Yuchen Shen, Zhen Wu and Xinyu Dai	2390
<i>Visual Named Entity Linking: A New Dataset and A Baseline</i> Wen Sun, Yixing Fan, Jiafeng Guo, Ruqing Zhang and Xueqi Cheng	2403
<i>MAGMA – Multimodal Augmentation of Generative Models through Adapter-based Finetuning</i> Constantin Eichenberg, Sidney Black, Samuel Weinbach, Letitia Parcalabescu and Anette Frank	2416
<i>Towards Tracing Knowledge in Language Models Back to the Training Data</i> Ekin Akyurek, Tolga Bolukbasi, Frederick Liu, Binbin Xiong, Ian Tenney, Jacob Andreas and Kelvin Guu	2429
<i>ReaRev: Adaptive Reasoning for Question Answering over Knowledge Graphs</i> Costas Mavromatis and George Karypis	2447
<i>Understanding Social Media Cross-Modality Discourse in Linguistic Space</i> Chunpu Xu, Hanzhuo Tan, Jing Li and Piji Li	2459
<i>TAPE: Assessing Few-shot Russian Language Understanding</i> Ekaterina Taktasheva, Alena Fenogenova, Denis Shevelev, Nadezhda Katricheva, Maria Tikhono- va, Albina Akhmetgareeva, Oleg Zinkevich, Anastasiia Bashmakova, Svetlana Iordanskaia, Valentina Kurenschikova, Alena Spiridonova, Ekaterina Artemova, Tatiana Shavrina and Vladislav Mikhailov	2472

<i>A Hierarchical N-Gram Framework for Zero-Shot Link Prediction</i>	
Mingchen Li, Junfan Chen, Samuel Mensah, Nikolaos Aletras, Xiulong Yang and yang ye .	2498
<i>Quadapter: Adapter for GPT-2 Quantization</i>	
Minseop Park, Jaeseong You, Markus Nagel and Simyung Chang	2510
<i>BanglaRQA: A Benchmark Dataset for Under-resourced Bangla Language Reading Comprehension-based Question Answering with Diverse Question-Answer Types</i>	
Syed Mohammed Sartaj Ekram, Adham Arik Rahman, Md. Sajid Altaf, Mohammed Saidul Islam, Mehrab Mustafy Rahman, Md Mezbaur Rahman, Md Azam Hossain and Abu Raihan Mostofa Kamal	2518
<i>Chaining Simultaneous Thoughts for Numerical Reasoning</i>	
Zhihong Shao, Fei Huang and Minlie Huang	2533
<i>Inferring Implicit Relations in Complex Questions with Language Models</i>	
Uri Katz, Mor Geva and Jonathan Berant	2548
<i>Eliciting and Understanding Cross-task Skills with Task-level Mixture-of-Experts</i>	
Qinyuan Ye, Juan Zha and Xiang Ren	2567
<i>On the Curious Case of l2 norm of Sense Embeddings</i>	
Yi Zhou and Danushka Bollegala	2593
<i>Partially-Random Initialization: A Smoking Gun for Binarization Hypothesis of BERT</i>	
Arash Ardakani	2603
<i>Prompt Consistency for Zero-Shot Task Generalization</i>	
Chunting Zhou, Junxian He, Xuezhe Ma, Taylor Berg-Kirkpatrick and Graham Neubig	2613
<i>In-Context Learning for Few-Shot Dialogue State Tracking</i>	
Yushi Hu, Chia-Hsuan Lee, Tianbao Xie, Tao Yu, Noah A. Smith and Mari Ostendorf	2627
<i>On Advances in Text Generation from Images Beyond Captioning: A Case Study in Self-Rationalization</i>	
Shruti Palaskar, Akshita Bhagia, Yonatan Bisk, Florian Metze, Alan W Black and Ana Marasovic	2644
<i>The challenges of temporal alignment on Twitter during crises</i>	
Aniket Pramanick, Tilman Beck, Kevin Stowe and Iryna Gurevych	2658
<i>Experimental Standards for Deep Learning in Natural Language Processing Research</i>	
Dennis Ulmer, Elisa Bassignana, Max Müller-Eberstein, Daniel Varab, Mike Zhang, Rob van der Goot, Christian Hardmeier and Barbara Plank	2673
<i>Few-Shot Anaphora Resolution in Scientific Protocols via Mixtures of In-Context Experts</i>	
Nghia T. Le, Fan Bai and Alan Ritter	2693
<i>Exploring Predictive Uncertainty and Calibration in NLP: A Study on the Impact of Method & Data Scarcity</i>	
Dennis Ulmer, Jes Frellsen and Christian Hardmeier	2707
<i>Conditional Supervised Contrastive Learning for Fair Text Classification</i>	
Jianfeng Chi, William Shand, Yaodong Yu, Kai-Wei Chang, Han Zhao and Yuan Tian	2736
<i>SpaBERT: A Pretrained Language Model from Geographic Data for Geo-Entity Representation</i>	
Zekun Li, Jina Kim, Yao-Yi Chiang and Muhao Chen	2757

<i>Self-training with Two-phase Self-augmentation for Few-shot Dialogue Generation</i>	
Wanyu Du, Hanjie Chen and Yangfeng Ji	2770
<i>Is NLP Ready for Standardization?</i>	
Lauriane Aufrant	2785
<i>Probing for Incremental Parse States in Autoregressive Language Models</i>	
Tiwalayo Eisape, Vineet Gangireddy, Roger Levy and Yoon Kim	2801
<i>Re-Examining Calibration: The Case of Question Answering</i>	
Chenglei Si, Chen Zhao, Sewon Min and Jordan Boyd-Graber	2814
<i>Accelerating Learned Sparse Indexes Via Term Impact Decomposition</i>	
Joel Mackenzie, Antonio Mallia, Alistair Moffat and Matthias Petri	2830
<i>Do Text-to-Text Multi-Task Learners Suffer from Task Conflict?</i>	
David Mueller, Nicholas Andrews and Mark Dredze	2843
<i>MANTA: Efficient Gradient-Based Tokenization for End-to-End Robust Language Modeling</i>	
Nathan Godey, Roman Castagné, Éric de la Clergerie and Benoît Sagot	2859
<i>Towards Intelligent Clinically-Informed Language Analyses of People with Bipolar Disorder and Schizophrenia</i>	
Ankit Aich, Avery Quynh, Varsha Badal, Amy Pinkham, Philip Harvey, Colin Depp and Natalie Parde	2871
<i>Calibrating Trust of Multi-Hop Question Answering Systems with Decompositional Probes</i>	
Kaige Xie, Sarah Wiegrefe and Mark Riedl	2888
<i>CheckHARD: Checking Hard Labels for Adversarial Text Detection, Prediction Correction, and Perturbed Word Suggestion</i>	
Hoang-Quoc Nguyen-Son, Huy Quang Ung, Seira Hidano, Kazuhide Fukushima and Shinsaku Kiyomoto	2903
<i>Mitigating Covertly Unsafe Text within Natural Language Systems</i>	
Alex Mei, Anisha Kabir, Sharon Levy, Melanie Subbiah, Emily Allaway, John Judge, Desmond Patton, Bruce Bimber, Kathleen McKeown and William Yang Wang	2914
<i>“I Know Who You Are”: Character-Based Features for Conversational Humor Recognition in Chinese</i>	
Wenbo Shang, Jiangjiang Zhao, Zezhong WANG, Binyang Li, Fangchun Yang and Kam-Fai Wong	2927
<i>DebiasGAN: Eliminating Position Bias in News Recommendation with Adversarial Learning</i>	
Chuhan Wu, Fangzhao Wu, Xiangnan He and Yongfeng Huang	2933
<i>Generating Multiple-Length Summaries via Reinforcement Learning for Unsupervised Sentence Summarization</i>	
Dongmin Hyun, Xiting Wang, Chayoung Park, Xing Xie and Hwanjo Yu	2939
<i>Multilingual Sentence Transformer as A Multilingual Word Aligner</i>	
Weikang Wang, Guanhua Chen, Hanqing Wang, Yue Han and Yun Chen	2952
<i>CORE: A Retrieve-then-Edit Framework for Counterfactual Data Generation</i>	
Tanay Dixit, Bhargavi Paranjape, Hannaneh Hajishirzi and Luke Zettlemoyer	2964
<i>Conversation Disentanglement with Bi-Level Contrastive Learning</i>	
Chengyu Huang, Zheng Zhang, Hao Fei and Lizi Liao	2985

<i>You can't pick your neighbors, or can you? When and How to Rely on Retrieval in the kNN-LM</i> Andrew Drozdov, Shufan Wang, Razieh Rahimi, Andrew McCallum, Hamed Zamani and Mohit Iyyer.....	2997
<i>StuBot: Learning by Teaching a Conversational Agent Through Machine Reading Comprehension</i> Nayoung Jin and Hana Lee	3008
<i>Improved Universal Sentence Embeddings with Prompt-based Contrastive Learning and Energy-based Learning</i> Yuxin Jiang, Linhan Zhang and Wei Wang.....	3021
<i>RaP: Redundancy-aware Video-language Pre-training for Text-Video Retrieval</i> Xing Wu, Chaochen Gao, Zijia Lin, Zhongyuan Wang, Jizhong Han and Songlin Hu	3036
<i>FCGCL: Fine- and Coarse-Granularity Contrastive Learning for Speech Translation</i> Hao Zhang, Nianwen Si, Yaqi Chen, Zhen Li, Tong Niu, Xukui Yang and Dan Qu.....	3048
<i>InfoCSE: Information-aggregated Contrastive Learning of Sentence Embeddings</i> Xing Wu, Chaochen Gao, Zijia Lin, Jizhong Han, Zhongyuan Wang and Songlin Hu	3060
<i>Benchmarking Language Models for Code Syntax Understanding</i> Da Shen, Xinyun Chen, Chenguang Wang, Koushik Sen and Dawn Song	3071
<i>Learning When and What to Quote: A Quotation Recommender System with Mutual Promotion of Recommendation and Generation</i> Lingzhi Wang, Xingshan Zeng and Kam-Fai Wong	3094
<i>Think Beyond Words: Exploring Context-Relevant Visual Commonsense for Diverse Dialogue Generation</i> Yiting Liu, Liang Li, Beichen Zhang and Qingming Huang.....	3106
<i>Gender Bias in Meta-Embeddings</i> Masahiro Kaneko, Danushka Bollegala and Naoaki Okazaki	3118
<i>Third-Party Aligner for Neural Word Alignments</i> Jinpeng Zhang, Chuanqi Dong, Xiangyu Duan, Yuqi Zhang and Min Zhang	3134
<i>QaDialMoE: Question-answering Dialogue based Fact Verification with Mixture of Experts</i> Longzheng Wang, Peng Zhang, Xiaoyu Lu, Lei Zhang, Chaoyang Yan and Chuang Zhang ..	3146
<i>Multimodal Knowledge Learning for Named Entity Disambiguation</i> Zhang Dongjie and Longtao Huang.....	3160
<i>Generative Prompt Tuning for Relation Classification</i> Jiale Han, Shuai Zhao, Bo Cheng, shengkun ma and Wei Lu.....	3170
<i>Formulating Few-shot Fine-tuning Towards Language Model Pre-training: A Pilot Study on Named Entity Recognition</i> Zihan Wang, Kewen Zhao, Zilong Wang and Jingbo Shang.....	3186
<i>Masked Language Models Know Which are Popular: A Simple Ranking Strategy for Commonsense Question Answering</i> Xuan Luo, Chuang Fan, Yice Zhang, Wanguo Jiang, Bing Qin and Ruifeng Xu.....	3200
<i>DialogUSR: Complex Dialogue Utterance Splitting and Reformulation for Multiple Intent Detection</i> Haoran Meng, Zheng Xin, Tianyu Liu, Zizhen Wang, he feng, Binghuai Lin, Xuemin Zhao, Yunbo Cao and Zhifang Sui	3214

<i>Low-resource Interactive Active Labeling for Fine-tuning Language Models</i>	
Seiji Maekawa, Dan Zhang, Hannah Kim, Sajjadur Rahman and Estevam Hruschka	3230
<i>Getting the Most out of Simile Recognition</i>	
Xiaoyue Wang, Linfeng Song, Xin Liu, Chulun Zhou, Hualin Zeng and Jinsong Su	3243
<i>A Unified Framework for Pun Generation with Humor Principles</i>	
Yufei Tian, Divyanshu Sheth and Nanyun Peng	3253
<i>Improving English-Arabic Transliteration with Phonemic Memories</i>	
Yuanhe Tian, Renze Lou, Xiangyu Pang, Lianxi Wang, Shengyi JIANG and Yan Song	3262
<i>Mix-and-Match: Scalable Dialog Response Retrieval using Gaussian Mixture Embeddings</i>	
Gaurav Pandey, Danish Contractor and Sachindra Joshi	3273
<i>AlphaTuning: Quantization-Aware Parameter-Efficient Adaptation of Large-Scale Pre-Trained Language Models</i>	
Se Jung Kwon, Jeonghoon Kim, Jeongin Bae, Kang Min Yoo, Jin-Hwa Kim, Baeseong Park, Byeongwook Kim, Jung-Woo Ha, Nako Sung and Dongsoo Lee	3288
<i>Learning Invariant Representation Improves Robustness for MRC Models</i>	
Yu Hai, Liang Wen, Haoran Meng, Tianyu Liu and Houfeng Wang	3306
<i>ER-Test: Evaluating Explanation Regularization Methods for Language Models</i>	
Brihi Joshi, Aaron Chan, Ziyi Liu, Shaoliang Nie, Maziar Sanjabi, Hamed Firooz and Xiang Ren	3315
<i>Learning Cooperative Interactions for Multi-Overlap Aspect Sentiment Triplet Extraction</i>	
Shiman Zhao, Wei Chen and Tengjiao Wang	3337
<i>Different Tunes Played with Equal Skill: Exploring a Unified Optimization Subspace for Parameter-Efficient Tuning</i>	
Jing Yi, Weize Chen, Yujia Qin, Yankai Lin, Ning Ding, Xu Han, Zhiyuan Liu, Maosong Sun and Jie Zhou	3348
<i>Explainable Slot Type Attentions to Improve Joint Intent Detection and Slot Filling</i>	
Kalpa Gunaratna, Vijay Srinivasan, Akhila Yerukola and Hongxia Jin	3367
<i>PseudoReasoner: Leveraging Pseudo Labels for Commonsense Knowledge Base Population</i>	
Tianqing Fang, Quyet V. Do, Hongming Zhang, Yangqiu Song, Ginny Y. Wong and Simon See	3379
<i>History-Aware Hierarchical Transformer for Multi-session Open-domain Dialogue System</i>	
Tong Zhang, Yong Liu, Boyang Li, Zhiwei Zeng, Pengwei Wang, Yuan You, Chunyan Miao and Lizhen Cui	3395
<i>Guiding Abstractive Dialogue Summarization with Content Planning</i>	
Ye Wang, Xiaojun Wan and Zhiping Cai	3408
<i>Truncation Sampling as Language Model Desmoothing</i>	
John Hewitt, Christopher Manning and Percy Liang	3414
<i>Knowledge-grounded Dialog State Tracking</i>	
Dian Yu, Mingqiu Wang, Yuan Cao, Laurent El Shafey, Izhak Shafran and Hagen Soltau	3428
<i>Context-aware Information-theoretic Causal De-biasing for Interactive Sequence Labeling</i>	
Junda Wu, Rui Wang, Tong Yu, Ruiyi Zhang, Handong Zhao, Shuai Li, Ricardo Henao and Ani Nenkova	3436

<i>Simple but Challenging: Natural Language Inference Models Fail on Simple Sentences</i> Cheng Luo, Wei Liu, Jieyu Lin, Jiajie Zou, Ming Xiang and Nai Ding	3449
<i>DORE: Document Ordered Relation Extraction based on Generative Framework</i> Qipeng Guo, Yuqing Yang, Hang Yan, Xipeng Qiu and Zheng Zhang	3463
<i>Explicit Role Interaction Network for Event Argument Extraction</i> Nan Ding, Chunming Hu, Kai Sun, Samuel Mensah and Richong Zhang	3475
<i>Few-Shot Out-of-Domain Transfer Learning of Natural Language Explanations in a Label-Abundant Setup</i> Yordan Yordanov, Vid Kocijan, Thomas Lukasiewicz and Oana-Maria Camburu	3486
<i>RoChBert: Towards Robust BERT Fine-tuning for Chinese</i> Zihan Zhang, Jinfeng Li, Ning Shi, Bo Yuan, Xiangyu Liu, Rong Zhang, hui xue, DONGHONG SUN and Chao Zhang	3502
<i>Lexical Entailment with Hierarchy Representations by Deep Metric Learning</i> Naomi Sato, Masaru Isonuma, Kimitaka Asatani, Shoya Ishizuka, Aori Shimizu and Ichiro Sakata	3517
<i>Improving the Sample Efficiency of Prompt Tuning with Domain Adaptation</i> Xu Guo, Boyang Li and Han Yu	3523
<i>McPhraSy: Multi-Context Phrase Similarity and Clustering</i> Amir Cohen, Hila Gonen, Ori Shapira, Ran Levy and Yoav Goldberg	3538
<i>CANarEx: Contextually Aware Narrative Extraction for Semantically Rich Text-as-data Applications</i> Nandini Anantharama, Simon Angus and Lachlan O’Neill	3551
<i>Narrate Dialogues for Better Summarization</i> Ruo Chen Xu, Chenguang Zhu and Michael Zeng	3565
<i>Towards Identifying Social Bias in Dialog Systems: Framework, Dataset, and Benchmark</i> Jingyan ZHOU, Jiawen Deng, Fei Mi, Yitong Li, Yasheng Wang, Minlie Huang, Xin Jiang, Qun Liu and Helen Meng	3576
<i>CrossRE: A Cross-Domain Dataset for Relation Extraction</i> Elisa Bassignana and Barbara Plank	3592
<i>Probing Structural Knowledge from Pre-trained Language Model for Argumentation Relation Classification</i> Yang Sun, Bin Liang, Jianzhu Bao, Min Yang and Ruifeng Xu	3605
<i>LogicNMR: Probing the Non-monotonic Reasoning Ability of Pre-trained Language Models</i> Yeliang Xiu, Zhanhao Xiao and Yongmei Liu	3616
<i>Cheater’s Bowl: Human vs. Computer Search Strategies for Open-Domain QA</i> Wanrong He, Andrew Mao and Jordan Boyd-Graber	3627
<i>FRSUM: Towards Faithful Abstractive Summarization via Enhancing Factual Robustness</i> Wenhao Wu, Wei Li, Jiachen Liu, Xinyan Xiao, Ziqiang Cao, Sujian Li and Hua Wu	3640
<i>PoeLM: A Meter- and Rhyme-Controllable Language Model for Unsupervised Poetry Generation</i> Aitor Ormazabal, Mikel Artetxe, Manex Agirrezabal, Aitor Soroa and Eneko Agirre	3655
<i>ProGen: Progressive Zero-shot Dataset Generation via In-context Feedback</i> Jiacheng Ye, Jiahui Gao, Zhiyong Wu, Jiangtao Feng, Tao Yu and Lingpeng Kong	3671

<i>Constructing Highly Inductive Contexts for Dialogue Safety through Controllable Reverse Generation</i> Zhexin Zhang, Jiale Cheng, Hao Sun, Jiawen Deng, Fei Mi, Yasheng Wang, Lifeng Shang and Minlie Huang	3684
<i>Language Prior Is Not the Only Shortcut: A Benchmark for Shortcut Learning in VQA</i> Qingyi Si, Fandong Meng, Mingyu Zheng, Zheng Lin, Yuanxin LIU, Peng Fu, Yanan Cao, Weiping Wang and Jie Zhou	3698
<i>Bridging the Training-Inference Gap for Dense Phrase Retrieval</i> Gyuwan Kim, Jinhyuk Lee, Barlas Oguz, Wenhan Xiong, Yizhe Zhang, Yashar Mehdad and William Yang Wang	3713
<i>Beyond Counting Datasets: A Survey of Multilingual Dataset Construction and Necessary Resources</i> Xinyan Yu, Trina Chatterjee, Akari Asai, Junjie Hu and Eunsol Choi	3725
<i>ERNIE-Layout: Layout Knowledge Enhanced Pre-training for Visually-rich Document Understanding</i> Qiming Peng, Yinxu Pan, Wenjin Wang, Bin Luo, Zhenyu Zhang, Zhengjie Huang, Yuhui Cao, Weichong Yin, Yongfeng Chen, Yin Zhang, Shikun Feng, Yu Sun, Hao Tian, Hua Wu and Haifeng Wang	3744
<i>Do Charge Prediction Models Learn Legal Theory?</i> Zhenwei An, Quzhe Huang, Cong Jiang, Yansong Feng and Dongyan Zhao	3757
<i>Keep Me Updated! Memory Management in Long-term Conversations</i> Sanghwan Bae, Donghyun Kwak, Soyoung Kang, Min Young Lee, Sungdong Kim, Yui Jeong, Hyeri Kim, Sang-Woo Lee, Woomyoung Park and Nako Sung	3769
<i>A Unified Dialogue User Simulator for Few-shot Data Augmentation</i> Dazhen Wan, Zheng Zhang, Qi Zhu, Lizi Liao and Minlie Huang	3788
<i>An Error-Guided Correction Model for Chinese Spelling Error Correction</i> Rui Sun, Xiuyu Wu and Yunfang Wu	3800
<i>Describing Sets of Images with Textual-PCA</i> Oded Hupert, Idan Schwartz and Lior Wolf	3811
<i>Learning to Model Editing Processes</i> Machel Reid and Graham Neubig	3822
<i>PALT: Parameter-Lite Transfer of Language Models for Knowledge Graph Completion</i> Jianhao Shen, Chenguang Wang, Ye Yuan, Jiawei Han, Heng Ji, Koushik Sen, Ming Zhang and Dawn Song	3833
<i>Prompt-based Connective Prediction Method for Fine-grained Implicit Discourse Relation Recognition</i> Hao Zhou, Man Lan, Yuanbin Wu, Yuefeng Chen and Meirong Ma	3848
<i>On Utilizing Constituent Language Resources to Improve Downstream Tasks in Hinglish</i> Vishwajeet Kumar, Rudra Murthy and Tejas Dhamecha	3859
<i>SYGMA: A System for Generalizable and Modular Question Answering Over Knowledge Bases</i> Sumit Neelam, Udit Sharma, Hima Karanam, Shajith Iqbal, Pavan Kapanipathi, Ibrahim Abdelaziz, Nandana Mihindukulasooriya, Young-Suk Lee, Santosh Srivastava, Cezar Pendus, Saswati Dana, Dinesh Garg, Achille Fokoue, G P Shrivatsa Bhargav, Dinesh Khandelwal, Srinivas Ravishankar, Sairam Gurajada, Maria Chang, Rosario Uceda-Sosa, Salim Roukos, Alexander Gray, Guilherme Lima, Ryan Riegel, Francois Luus and L V Subramaniam	3866

<i>Instance-Guided Prompt Learning for Few-Shot Text Matching</i>	
Jia Du, Xuanyu Zhang, siyi wang, Kai Wang, Yanquan Zhou, Lei Li, Qing Yang and Dongliang Xu.....	3880
<i>M3: Multi-level dataset for Multi-document summarisation of Medical studies</i>	
Yulia Otmakhova, Karin Verspoor, Timothy Baldwin, Antonio Jimeno Yepes and Jey Han Lau	3887
<i>Adapters for Enhanced Modeling of Multilingual Knowledge and Text</i>	
Yifan Hou, Wenxiang Jiao, Meizhen Liu, Carl Allen, Zhaopeng Tu and Mrinmaya Sachan .	3902
<i>SepLL: Separating Latent Class Labels from Weak Supervision Noise</i>	
Andreas Stephan, Vasiliki Kougia and Benjamin Roth	3918
<i>Probing Relational Knowledge in Language Models via Word Analogies</i>	
Kiamehr Rezaee and Jose Camacho-Collados	3930
<i>Semi-Supervised Lifelong Language Learning</i>	
Yingxiu Zhao, Yinhe Zheng, Bowen Yu, Zhiliang Tian, Dongkyu Lee, Jian Sun, Yongbin Li and Nevin L. Zhang	3937
<i>Parameter-free Automatically Prompting: A Latent Pseudo Label Mapping Model for Prompt-based Learning</i>	
Jirui Qi, Richong Zhang, Junfan Chen, Jaein Kim and Yongyi Mao	3952
<i>Exploring Logographic Image for Chinese Aspect-based Sentiment Classification</i>	
Xiabing Zhou, Renjie Feng, Xiaotong Jiang and Zhongqing Wang	3963
<i>On the Role of Bidirectionality in Language Model Pre-Training</i>	
Mikel Artetxe, Jingfei Du, Naman Goyal, Luke Zettlemoyer and Veselin Stoyanov	3973
<i>You Are What You Talk About: Inducing Evaluative Topics for Personality Analysis</i>	
Josip Jukić, Iva Vukojević and Jan Snajder	3986
<i>CAT-probing: A Metric-based Approach to Interpret How Pre-trained Models for Programming Language Attend Code Structure</i>	
Nuo Chen, Qiushi Sun, Renyu Zhu, Xiang Li, Xuesong Lu and Ming Gao	4000
<i>Learning to Revise References for Faithful Summarization</i>	
Griffin Adams, Han-Chin Shing, Qing Sun, Christopher Winestock, Kathleen McKeown and Noémie Elhadad	4009
<i>Towards Intention Understanding in Suicidal Risk Assessment with Natural Language Processing</i>	
Shaoxiong Ji	4028
<i>On the Impact of Temporal Concept Drift on Model Explanations</i>	
Zhixue Zhao, George Chrysostomou, Kalina Bontcheva and Nikolaos Aletras	4039
<i>Text-Only Training for Image Captioning using Noise-Injected CLIP</i>	
David Nukrai, Ron Mokady and Amir Globerson	4055
<i>Improving Sharpness-Aware Minimization with Fisher Mask for Better Generalization on Language Models</i>	
Qihuang Zhong, Liang Ding, Li Shen, Peng Mi, Juhua Liu, Bo Du and Dacheng Tao	4064
<i>TINA: Textual Inference with Negation Augmentation</i>	
Chadi Helwe, Simon Coumes, Chloé Clavel and Fabian Suchanek	4086

<i>Improving Bilingual Lexicon Induction with Cross-Encoder Reranking</i>	
Yaoyiran Li, Fangyu Liu, Ivan Vulić and Anna Korhonen	4100
<i>Mixed-modality Representation Learning and Pre-training for Joint Table-and-Text Retrieval in OpenQA</i>	
Junjie Huang, Wanjun Zhong, Qian Liu, Ming Gong, Daxin Jiang and Nan Duan	4117
<i>The Effects of Corpus Choice and Morphosyntax on Multilingual Space Induction</i>	
Vinit Ravishankar and Joakim Nivre	4130
<i>Modeling Complex Dialogue Mappings via Sentence Semantic Segmentation Guided Conditional Variational Auto-Encoder</i>	
Bin Sun, Shaoxiong Feng, Yiwei Li, Weichao Wang, Fei Mi, Yitong Li and Kan Li	4140
<i>Graph Embeddings for Argumentation Quality Assessment</i>	
Santiago Marro, Elena Cabrio and Serena Villata	4154
<i>SMiLE: Schema-augmented Multi-level Contrastive Learning for Knowledge Graph Link Prediction</i>	
Miao Peng, Ben Liu, Qianqian Xie, Wenjie Xu, Hua Wang and Min Peng	4165
<i>Multilingual Multimodal Learning with Machine Translated Text</i>	
Chen Qiu, Dan Oneață, Emanuele Bugliarello, Stella Frank and Desmond Elliott	4178
<i>Learning From the Source Document: Unsupervised Abstractive Summarization</i>	
Haojie Zhuang, Wei Emma Zhang, Jian Yang, Congbo Ma, Yutong Qu and Quan Z. Sheng ..	4194
<i>How to Do Things without Words: Modeling Semantic Drift of Emoji</i>	
Eyal Arviv and Oren Tsur	4206
<i>Mind Your Bias: A Critical Review of Bias Detection Methods for Contextual Language Models</i>	
Silke Husse and Andreas Spitz	4212
<i>ZeroPrompt: Scaling Prompt-Based Pretraining to 1,000 Tasks Improves Zero-Shot Generalization</i>	
Hanwei Xu, Yujun Chen, Yulun Du, Nan Shao, wang yanggang, Haiyu Li and Zhilin Yang ..	4235
<i>Semantic Role Labeling Meets Definition Modeling: Using Natural Language to Describe Predicate-Argument Structures</i>	
Simone Conia, Edoardo Barba, Alessandro Scirè and Roberto Navigli	4253
<i>Is anisotropy really the cause of BERT embeddings not being semantic?</i>	
Alejandro Fuster Baggetto and Victor Fresno	4271
<i>m⁴ Adapter: Multilingual Multi-Domain Adaptation for Machine Translation with a Meta-Adapter</i>	
Wen Lai, Alexandra Chronopoulou and Alexander Fraser	4282
<i>Textual Enhanced Contrastive Learning for Solving Math Word Problems</i>	
Yibin Shen, Qianying Liu, Zhuoyuan Mao, Fei Cheng and Sadao Kurohashi	4297
<i>What Do Compressed Multilingual Machine Translation Models Forget?</i>	
Alireza Mohammadshahi, Vassilina Nikoulina, Alexandre Berard, Caroline Brun, James Henderson and Laurent Besacier	4308
<i>Controllable Dialogue Simulation with In-context Learning</i>	
Zekun Li, Wenhu Chen, Shiyang Li, Hong Wang, Jing Qian and Xifeng Yan	4330
<i>Improving the Factual Correctness of Radiology Report Generation with Semantic Rewards</i>	
Jean-Benoit Delbrouck, Pierre Chambon, Christian Bluethgen, Emily Tsai, Omar Almusa and Curtis Langlotz	4348

<i>Recursive Neural Networks with Bottlenecks Diagnose (Non-)Compositionality</i> Verna Dankers and Ivan Titov	4361
<i>HumSet: Dataset of Multilingual Information Extraction and Classification for Humanitarian Crises Response</i> Selim Fekih, Nicolo' Tamagnone, Benjamin Minixhofer, Ranjan Shrestha, Ximena Contla, Ewan Oglethorpe and Navid Rekabsaz	4379
<i>Viterbi Decoding of Directed Acyclic Transformer for Non-Autoregressive Machine Translation</i> Chenze Shao, Zhengrui Ma and Yang Feng	4390
<i>Lexical Generalization Improves with Larger Models and Longer Training</i> Elron Bandel, Yoav Goldberg and Yanai Elazar	4398
<i>Realistic Data Augmentation Framework for Enhancing Tabular Reasoning</i> Dibyakanti Kumar, Vivek Gupta, Soumya Sharma and Shuo Zhang	4411
<i>Inducing Generalizable and Interpretable Lexica</i> Yilin Geng, Zetian Wu, Roshan Santhosh, Tejas Srivastava, Lyle Ungar and João Sedoc	4430
<i>The Curious Case of Absolute Position Embeddings</i> Koustuv Sinha, Amirhossein Kazemnejad, Siva Reddy, Joelle Pineau, Dieuwke Hupkes and Adina Williams	4449
<i>Goal-oriented Vision-and-Dialog Navigation via Reinforcement Learning</i> Yan Cao, Keting Lu, David DeFazio and Shiqi Zhang	4473
<i>Leveraging Data Recasting to Enhance Tabular Reasoning</i> Aashna Jena, Vivek Gupta, Manish Shrivastava and Julian Eisenschlos	4483
<i>Thinking about GPT-3 In-Context Learning for Biomedical IE? Think Again</i> Bernal Jimenez Gutierrez, Nikolas McNeal, Clayton Washington, You Chen, Lang Li, Huan Sun and Yu Su	4497
<i>Attention weights accurately predict language representations in the brain</i> Mathis Lamarre, Catherine Chen and Fatma Deniz	4513
<i>Improving HowNet-Based Chinese Word Sense Disambiguation with Translations</i> Xiang Zhang, Bradley Hauer and Grzegorz Kondrak	4530
<i>Mask-then-Fill: A Flexible and Effective Data Augmentation Framework for Event Extraction</i> Jun Gao, Changlong Yu, Wei Wang, Huan Zhao and Ruifeng Xu	4537
<i>MOBA-E2C: Generating MOBA Game Commentaries via Capturing Highlight Events from the Meta-Data</i> Dawei Zhang, Sixing Wu, Yao Guo and Xiangqun Chen	4545
<i>Enhancing Automatic Readability Assessment with Pre-training and Soft Labels for Ordinal Regression</i> Jinshan Zeng, Yudong Xie, Xianglong Yu, John Lee and Ding-Xuan Zhou	4557
<i>Opening up Minds with Argumentative Dialogues</i> Youmna Farag, Charlotte Brand, Jacopo Amidei, Paul Piwek, Tom Stafford, Svetlana Stoyanchev and Andreas Vlachos	4569
<i>You Are My Type! Type Embeddings for Pre-trained Language Models</i> Mohammed Saeed and Paolo Papotti	4583

<i>Generating Textual Adversaries with Minimal Perturbation</i>	
Xingyi Zhao, Lu Zhang, Depeng Xu and Shuhan Yuan	4599
<i>SensePOLAR: Word sense aware interpretability for pre-trained contextual word embeddings</i>	
Jan Engler, Sandipan Sikdar, Marlene Lutz and Markus Strohmaier	4607
<i>Contextualizing Language Models for Norms Diverging from Social Majority</i>	
Niklas Kiehne, Hermann Kroll and Wolf-Tilo Balke	4620
<i>Empathetic Dialogue Generation via Sensitive Emotion Recognition and Sensible Knowledge Selection</i>	
Lanrui Wang, Jiangnan Li, Zheng Lin, Fandong Meng, Chenxu Yang, Weiping Wang and Jie Zhou	4634
<i>Joint Multilingual Knowledge Graph Completion and Alignment</i>	
Vinh Tong, Dat Quoc Nguyen, Trung Thanh Huynh, Tam Thanh Nguyen, Quoc Viet Hung Nguyen and Mathias Niepert	4646
<i>A Framework for Automatic Generation of Spoken Question-Answering Data</i>	
Merve Ünlü Menevşe, Yusufcan Manav, Ebru Arisoy and Arzucan Özgür	4659
<i>Readability Controllable Biomedical Document Summarization</i>	
Zheheng Luo, Qianqian Xie and Sophia Ananiadou	4667
<i>Beyond Additive Fusion: Learning Non-Additive Multimodal Interactions</i>	
Torsten Wörtwein, Lisa Sheeber, Nicholas Allen, Jeffrey Cohn and Louis-Philippe Morency	4681
<i>Generalization Differences between End-to-End and Neuro-Symbolic Vision-Language Reasoning Systems</i>	
Wang Zhu, Jesse Thomason and Robin Jia	4697
<i>Learning to Model Multimodal Semantic Alignment for Story Visualization</i>	
Bowen Li and Thomas Lukasiewicz	4712
<i>SciFact-Open: Towards open-domain scientific claim verification</i>	
David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang and Hannaneh Hajishirzi	4719
<i>COMET-QE and Active Learning for Low-Resource Machine Translation</i>	
Everlyn Chimoto and Bruce Bassett	4735
<i>MedicalSum: A Guided Clinical Abstractive Summarization Model for Generating Medical Reports from Patient-Doctor Conversations</i>	
George Michalopoulos, Kyle Williams, Gagandeep Singh and Thomas Lin	4741
<i>Leveraging Training Dynamics and Self-Training for Text Classification</i>	
Tiberiu Sosea and Cornelia Caragea	4750
<i>Learning to Infer from Unlabeled Data: A Semi-supervised Learning Approach for Robust Natural Language Inference</i>	
Mobashir Sadat and Cornelia Caragea	4763
<i>Unsupervised Text Deidentification</i>	
John Morris, Justin Chiu, Ramin Zabih and Alexander Rush	4777
<i>Federated Continual Learning for Text Classification via Selective Inter-client Transfer</i>	
Yatin Chaudhary, Pranav Rai, Matthias Schubert, Hinrich Schütze and Pankaj Gupta	4789

<i>DOROTHIE: Spoken Dialogue for Handling Unexpected Situations in Interactive Autonomous Driving Agents</i>	
Ziqiao Ma, Benjamin VanDerPloeg, Cristian-Paul Bara, Yidong Huang, Eui-In Kim, Felix Gervits, Matthew Marge and Joyce Chai	4800
<i>He Said, She Said: Style Transfer for Shifting the Perspective of Dialogues</i>	
Amanda Bertsch, Graham Neubig and Matthew R. Gormley	4823
<i>Dynamic Augmentation Data Selection for Few-shot Text Classification</i>	
Guangliang Liu, Lifeng Jin, Owen Yuan and Jiayu Zhou	4841
<i>KPDROP: Improving Absent Keyphrase Generation</i>	
Jishnu Ray Chowdhury, Seo Yeon Park, Tuhin Kundu and Cornelia Caragea	4853
<i>Natural Language Deduction through Search over Statement Compositions</i>	
Kaj Bostrom, Zayne Sprague, Swarat Chaudhuri and Greg Durrett	4871
<i>EnDex: Evaluation of Dialogue Engagingness at Scale</i>	
Guangxuan Xu, Ruibo Liu, Fabrice Harel-Canada, Nischal Reddy Chandra and Nanyun Peng	4884
<i>LOPS: Learning Order Inspired Pseudo-Label Selection for Weakly Supervised Text Classification</i>	
Dheeraj Mekala, Chengyu Dong and Jingbo Shang	4894
<i>Train Flat, Then Compress: Sharpness-Aware Minimization Learns More Compressible Models</i>	
Clara Na, Sanket Vaibhav Mehta and Emma Strubell	4909
<i>Structural Contrastive Representation Learning for Zero-shot Multi-label Text Classification</i>	
Tianyi Zhang, Zhaozhuo Xu, Tharun Medini and Anshumali Shrivastava	4937
<i>Improving Generalization of Pre-trained Language Models via Stochastic Weight Averaging</i>	
Peng Lu, Ivan Kobyzev, Mehdi Rezagholizadeh, Ahmad Rashid, Ali Ghodsi and Phillippe Langlais	4948
<i>Learn What Is Possible, Then Choose What Is Best: Disentangling One-To-Many Relations in Language Through Text-based Games</i>	
Benjamin Towle and Ke Zhou	4955
<i>Structurally Diverse Sampling for Sample-Efficient Training and Comprehensive Evaluation</i>	
Shivanshu Gupta, Sameer Singh and Matt Gardner	4966
<i>Unsupervised Multi-Granularity Summarization</i>	
Ming Zhong, Yang Liu, Suyu Ge, Yuning Mao, Yizhu Jiao, Xingxing Zhang, Yichong Xu, Chenguang Zhu, Michael Zeng and Jiawei Han	4980
<i>HeLo: Learning-Free Lookahead Decoding for Conversation Infilling</i>	
Ivan Lee and Taylor Berg-Kirkpatrick	4996
<i>Invernet: An Inversion Attack Framework to Infer Fine-Tuning Datasets through Word Embeddings</i>	
Ishrak Hayet, Zijun Yao and Bo Luo	5009
<i>LawngNLI: A Long-Premise Benchmark for In-Domain Generalization from Short to Long Contexts and for Implication-Based Retrieval</i>	
William Bruno and Dan Roth	5019
<i>Distillation-Resistant Watermarking for Model Protection in NLP</i>	
Xuandong Zhao, Lei Li and Yu-Xiang Wang	5044

<i>NeuroCounterfactuals: Beyond Minimal-Edit Counterfactuals for Richer Data Augmentation</i> Phillip Howard, Gadi Singer, Vasudev Lal, Yejin Choi and Swabha Swayamdipta	5056
<i>Don't Just Clean It, Proxy Clean It: Mitigating Bias by Proxy in Pre-Trained Models</i> Swetasudha Panda, Ari Kobren, Michael Wick and Qinlan Shen	5073
<i>The Undesirable Dependence on Frequency of Gender Bias Metrics Based on Word Embeddings</i> Francisco Valentini, Germán Rosati, Diego Fernandez Slezak and Edgar Altszyler	5086
<i>BioNLI: Generating a Biomedical NLI Dataset Using Lexico-semantic Constraints for Adversarial Examples</i> Mohaddeseh Bastan, Mihai Surdeanu and Niranjan Balasubramanian	5093
<i>Self-supervised Cross-modal Pretraining for Speech Emotion Recognition and Sentiment Analysis</i> Iek-Heng Chu, Ziyi Chen, Xinlu Yu, Mei Han, Jing Xiao and Peng Chang	5105
<i>Multimodal Conversation Modelling for Topic Derailment Detection</i> Zhenhao Li, Marek Rei and Lucia Specia	5115
<i>Active Learning for Abstractive Text Summarization</i> Akim Tsvigun, Ivan Lysenko, Danila Sedashov, Ivan Lazichny, Eldar Damirov, Vladimir Karlov, Artemy Belousov, Leonid Sanochkin, Maxim Panov, Alexander Panchenko, Mikhail Burtsev and Artem Shelmanov	5128
<i>Finding Memo: Extractive Memorization in Constrained Sequence Generation Tasks</i> Vikas Raunak and Arul Menezes	5153
<i>SALTED: A Framework for SAlient Long-tail Translation Error Detection</i> Vikas Raunak, Matt Post and Arul Menezes	5163
<i>Discord Questions: A Computational Approach To Diversity Analysis in News Coverage</i> Philippe Laban, Chien-Sheng Wu, Lidiya Murakhov's'ka, Xiang Chen and Caiming Xiong .	5180
<i>FocusQA: Open-Domain Question Answering with a Context in Focus</i> Gianni Barlacchi, Ivano Lauriola, Alessandro Moschitti, Marco Del Tredici, Xiaoyu Shen, Thuy Vu, Bill Byrne and Adrià de Gispert	5195
<i>Challenges and Opportunities in Information Manipulation Detection: An Examination of Wartime Russian Media</i> Chan Young Park, Julia Mendelsohn, Anjalie Field and Yulia Tsvetkov	5209
<i>Disentangling Task Relations for Few-shot Text Classification via Self-Supervised Hierarchical Task Clustering</i> Juan Zha, Zheng Li, Ying Wei and Yu Zhang	5236
<i>XRICL: Cross-lingual Retrieval-Augmented In-Context Learning for Cross-lingual Text-to-SQL Semantic Parsing</i> Peng Shi, Rui Zhang, He Bai and Jimmy Lin	5248
<i>Continuation KD: Improved Knowledge Distillation through the Lens of Continuation Optimization</i> Aref Jafari, Ivan Kobyzev, Mehdi Rezagholizadeh, Pascal Poupart and Ali Ghodsi	5260
<i>Detecting Dementia from Long Neuropsychological Interviews</i> Nauman Dawalatabad, Yuan Gong, Sameer Khurana, Rhoda Au and James Glass	5270
<i>Sarcasm Detection is Way Too Easy! An Empirical Comparison of Human and Machine Sarcasm Detection</i> Ibrahim Abu Farha, Steven Wilson, Silviu Oprea and Walid Magdy	5284

<i>Cross-lingual Text-to-SQL Semantic Parsing with Representation Mixup</i>	
Peng Shi, Linfeng Song, Lifeng Jin, Haitao Mi, He Bai, Jimmy Lin and Dong Yu	5296
<i>JamPatoisNLI: A Jamaican Patois Natural Language Inference Dataset</i>	
Ruth-Ann Armstrong, John Hewitt and Christopher Manning	5307
<i>Are Neural Topic Models Broken?</i>	
Alexander Miserlis Hoyle, Rupak Sarkar, Pranav Goel and Philip Resnik	5321
<i>Know Thy Strengths: Comprehensive Dialogue State Tracking Diagnostics</i>	
Hyundong Cho, Chinnadhurai Sankar, Christopher Lin, Kaushik Sadagopan, Shahin Shayandeh, Asli Celikyilmaz, Jonathan May and Ahmad Beirami	5345
<i>Open-domain Question Answering via Chain of Reasoning over Heterogeneous Knowledge</i>	
Kaixin Ma, Hao Cheng, Xiaodong Liu, Eric Nyberg and Jianfeng Gao	5360
<i>Detecting Languages Unintelligible to Multilingual Models through Local Structure Probes</i>	
Louis Clouatre, Prasanna Parthasarathi, Amal Zouaq and Sarath Chandar	5375
<i>Cards Against AI: Predicting Humor in a Fill-in-the-blank Party Game</i>	
Dan Ofer and Dafna Shahaf	5397
<i>Open-Vocabulary Argument Role Prediction For Event Extraction</i>	
Yizhu Jiao, Sha Li, Yiqing Xie, Ming Zhong, Heng Ji and Jiawei Han	5404
<i>Token-level Sequence Labeling for Spoken Language Understanding using Compositional End-to-End Models</i>	
Siddhant Arora, Siddharth Dalmia, Brian Yan, Florian Metze, Alan W Black and Shinji Watanabe	5419
<i>Baked-in State Probing</i>	
Shubham Toshniwal, Sam Wiseman, Karen Livescu and Kevin Gimpel	5430
<i>ClinicalT5: A Generative Language Model for Clinical Text</i>	
Qiuhaio Lu, Dejing Dou and Thien Nguyen	5436
<i>Find Someone Who: Visual Commonsense Understanding in Human-Centric Grounding</i>	
Haoxuan You, Rui Sun, Zhecan Wang, Kai-Wei Chang and Shih-Fu Chang	5444
<i>CrisisLTLSum: A Benchmark for Local Crisis Event Timeline Extraction and Summarization</i>	
Hossein Rajaby Faghihi, Bashar Alhafni, Ke Zhang, Shihao Ran, Joel Tetreault and Alejandro Jaimes	5455
<i>Prompt-Tuning Can Be Much Better Than Fine-Tuning on Cross-lingual Understanding With Multilingual Language Models</i>	
Lifu Tu, Caiming Xiong and Yingbo Zhou	5478
<i>BERT Meets CTC: New Formulation of End-to-End Speech Recognition with Pre-trained Masked Language Model</i>	
Yosuke Higuchi, Brian Yan, Siddhant Arora, Tetsuji Ogawa, Tetsunori Kobayashi and Shinji Watanabe	5486
<i>EtriCA: Event-Triggered Context-Aware Story Generation Augmented by Cross Attention</i>	
Chen Tang, Chenghua Lin, Henglin Huang, Frank Guerin and Zhihao Zhang	5504
<i>LADIS: Language Disentanglement for 3D Shape Editing</i>	
Ian Huang, Panos Achlioptas, Tianyi Zhang, Sergei Tulyakov, Minhyuk Sung and Leonidas Guibas	5519

<i>Effective Pretraining Objectives for Transformer-based Autoencoders</i>	
Luca Di Liello, Matteo Gabburo and Alessandro Moschitti	5533
<i>Language Model Detoxification in Dialogue with Contextualized Stance Control</i>	
Jing Qian and Xifeng Yan	5548
<i>Multilingual SubEvent Relation Extraction: A Novel Dataset and Structure Induction Method</i>	
Viet Lai, Hieu Man, Linh Ngo, Franck Dernoncourt and Thien Nguyen	5559
<i>A Two-Stage Approach towards Generalization in Knowledge Base Question Answering</i>	
Srinivas Ravishankar, Dung Thai, Ibrahim Abdelaziz, Nandana Mihindukulasooriya, Tahira Naseem, Pavan Kapanipathi, Gaetano Rossiello and Achille Fokoue	5571
<i>Few-Shot (Dis)Agreement Identification in Online Discussions with Regularized and Augmented Meta-Learning</i>	
Yuanyuan Lei and Ruihong Huang	5581
<i>Data Cartography for Low-Resource Neural Machine Translation</i>	
Aquia Richburg and Marine Carpuat	5594
<i>Augmenting Multi-Turn Text-to-SQL Datasets with Self-Play</i>	
Qi Liu, Zihuiwen Ye, Tao Yu, Linfeng Song and Phil Blunsom	5608
<i>Prompt Compression and Contrastive Conditioning for Controllability and Toxicity Reduction in Language Models</i>	
David Wingate, Mohammad Shoeybi and Taylor Sorensen	5621
<i>NaturalAdversaries: Can Naturalistic Adversaries Be as Effective as Artificial Adversaries?</i>	
Saadia Gabriel, Hamid Palangi and Yejin Choi	5635
<i>Multi-Path Transformer is Better: A Case Study on Neural Machine Translation</i>	
Ye Lin, Shuhan Zhou, Yanyang Li, Anxiang Ma, Tong Xiao and Jingbo Zhu	5646
<i>Unsupervised Learning of Hierarchical Conversation Structure</i>	
Bo-Ru Lu, Yushi Hu, Hao Cheng, Noah A. Smith and Mari Ostendorf	5657
<i>Task Compass: Scaling Multi-task Pre-training with Task Prefix</i>	
Zhuosheng Zhang, Shuohang Wang, Yichong Xu, Yuwei Fang, Wenhao Yu, Yang Liu, Hai Zhao, Chenguang Zhu and Michael Zeng	5671
<i>Sharpness-Aware Minimization with Dynamic Reweighting</i>	
Wenxuan Zhou, Fangyu Liu, Huan Zhang and Muhao Chen	5686
<i>Predicting Long-Term Citations from Short-Term Linguistic Influence</i>	
Sandeep Soni, David Bamman and Jacob Eisenstein	5700
<i>Joint Audio/Text Training for Transformer Rescorer of Streaming Speech Recognition</i>	
Suyoun Kim, Ke Li, Lucas Kabela, Ron Huang, Jiedan Zhu, Ozlem Kalinli and Duc Le	5717
<i>TyDiP: A Dataset for Politeness Classification in Nine Typologically Diverse Languages</i>	
Anirudh Srinivasan and Eunsol Choi	5723
<i>Probing Cross-modal Semantics Alignment Capability from the Textual Perspective</i>	
Zheng Ma, Shi Zong, Mianzhi Pan, Jianbing Zhang, Shujian Huang, Xinyu Dai and Jiajun CHEN	5739

<i>Hidden State Variability of Pretrained Language Models Can Guide Computation Reduction for Transfer Learning</i>	
Shuo Xie, Jiahao Qiu, Ankita Pasad, Li Du, Qing Qu and Hongyuan Mei	5750
<i>Language Models as Agent Models</i>	
Jacob Andreas	5769
<i>Combinatory Grammar Tells Underlying Relevance among Entities</i>	
Yuanhe Tian and Yan Song	5780
<i>Leveraging Open Data and Task Augmentation to Automated Behavioral Coding of Psychotherapy Conversations in Low-Resource Scenarios</i>	
Zhuohao Chen, Nikolaos Flemotomos, Zac Imel, David Atkins and Shrikanth Narayanan ..	5787
<i>Learning to Detect Noisy Labels Using Model-Based Features</i>	
Zhihao Wang, Zongyu Lin, Junjie Wen, Xianxin Chen, Peiqi Liu, Guidong Zheng, Yujun Chen and Zhilin Yang	5796
<i>Keyphrase Generation Beyond the Boundaries of Title and Abstract</i>	
Krishna Garg, Jishnu Ray Chowdhury and Cornelia Caragea	5809
<i>Composition, Attention, or Both?</i>	
Ryo Yoshida and Yohei Oseki	5822
<i>CDGP: Automatic Cloze Distractor Generation based on Pre-trained Language Model</i>	
Shang-Hsuan Chiang, Ssu-Cheng Wang and Yao-Chung Fan	5835
<i>G3: Geolocation via Guidebook Grounding</i>	
Grace Luo, Giscard Biamby, Trevor Darrell, Daniel Fried and Anna Rohrbach	5841
<i>Controlling Bias Exposure for Fair Interpretable Predictions</i>	
Zexue He, Yu Wang, Julian McAuley and Bodhisattwa Prasad Majumder	5854
<i>Investigating the Benefits of Free-Form Rationales</i>	
Jiao Sun, Swabha Swayamdipta, Jonathan May and Xuezhe Ma	5867
<i>Data-Efficient Concept Extraction from Pre-trained Language Models for Commonsense Explanation Generation</i>	
Yanbo Fang and Yongfeng Zhang	5883
<i>Unsupervised Domain Adaptation for Joint Information Extraction</i>	
Nghia Ngo, Bonan Min and Thien Nguyen	5894
<i>Foiling Training-Time Attacks on Neural Machine Translation Systems</i>	
Jun Wang, Xuanli He, Benjamin Rubinstein and Trevor Cohn	5906
<i>Learning Action-Effect Dynamics for Hypothetical Vision-Language Reasoning Task</i>	
Shailaja Keyur Sampat, Pratyay Banerjee, Yezhou Yang and Chitta Baral	5914
<i>Named Entity and Relation Extraction with Multi-Modal Retrieval</i>	
Xinyu Wang, Jiong Cai, Yong Jiang, Pengjun Xie, Kewei Tu and Wei Lu	5925
<i>Calibrating Factual Knowledge in Pretrained Language Models</i>	
Qingxiu Dong, Damai Dai, Yifan Song, Jingjing Xu, Zhifang Sui and Lei Li	5937
<i>MCPG: A Flexible Multi-Level Controllable Framework for Unsupervised Paraphrase Generation</i>	
Yi Chen, Haiyun Jiang, Lemao Liu, Rui Wang, Shuming Shi and Ruifeng Xu	5948

<i>WordTies: Measuring Word Associations in Language Models via Constrained Sampling</i> Peiran Yao, Tobias Renwick and Denilson Barbosa	5959
<i>Exploring The Landscape of Distributional Robustness for Question Answering Models</i> Anas Awadalla, Mitchell Wortsman, Gabriel Ilharco, Sewon Min, Ian Magnusson, Hannaneh Hajishirzi and Ludwig Schmidt	5971
<i>Collaborative Reasoning on Multi-Modal Semantic Graphs for Video-Grounded Dialogue Generation</i> Xueliang Zhao, Yuxuan Wang, Chongyang Tao, Chenshuo Wang and Dongyan Zhao	5988
<i>Partitioned Gradient Matching-based Data Subset Selection for Compute-Efficient Robust ASR Training</i> Ashish Mittal, Durga Sivasubramanian, Rishabh Iyer, Preethi Jyothi and Ganesh Ramakrishnan	5999
<i>Adaptive Graph Convolutional Network for Knowledge Graph Entity Alignment</i> Renbo Zhu, Xukun Luo, Meng Ma and Ping Wang	6011
<i>Towards Robust NLG Bias Evaluation with Syntactically-diverse Prompts</i> Arshiya Aggarwal, Jiao Sun and Nanyun Peng	6022
<i>PcMSP: A Dataset for Scientific Action Graphs Extraction from Polycrystalline Materials Synthesis Procedure Text</i> Xianjun Yang, Ya Zhuo, Julia Zuo, Xinlu Zhang, Stephen Wilson and Linda Petzold	6033
<i>Validity Assessment of Legal Will Statements as Natural Language Inference</i> Alice Kwak, Jacob Israelsen, Clayton Morrison, Derek Bambauer and Mihai Surdeanu	6047
<i>AdaPrompt: Adaptive Model Training for Prompt-based NLP</i> Yulong Chen, Yang Liu, Li Dong, Shuhang Wang, Chenguang Zhu, Michael Zeng and Yue Zhang	6057
<i>Code Generation From Flowcharts with Texts: A Benchmark Dataset and An Approach</i> Zejie Liu, Xiaoyu Hu, Deyu Zhou, Lin Li, Xu Zhang and Yanzheng Xiang	6069
<i>Focus! Relevant and Sufficient Context Selection for News Image Captioning</i> Mingyang Zhou, Grace Luo, Anna Rohrbach and Zhou Yu	6078
<i>Generative Aspect-Based Sentiment Analysis with Contrastive Learning and Expressive Structure</i> Joseph Peper and Lu Wang	6089
<i>Semantic Dependency Parsing with Edge GNNs</i> Songlin Yang and Kewei Tu	6096
<i>Explore Unsupervised Structures in Pretrained Models for Relation Extraction</i> Xi Yang, Tao Ji and Yuanbin Wu	6103
<i>Identifying Human Strategies for Generating Word-Level Adversarial Examples</i> Maximilian Mozes, Bennett Kleinberg and Lewis Griffin	6118
<i>Refinement Matters: Textual Description Needs to be Refined for Zero-shot Learning</i> Chandan Gautam, Sethupathy Parameswaran, Vinay Verma, Suresh Sundaram and Savitha Ramasamy	6127
<i>SAT: Improving Semi-Supervised Text Classification with Simple Instance-Adaptive Self-Training</i> Hui Chen, Wei Han and Soujanya Poria	6141
<i>Answer Quality Aware Aggregation for Extractive QA Crowdsourcing</i> Peide Zhu, Zhen Wang, Claudia Hauff, Jie Yang and Avishek Anand	6147

<i>Search to Pass Messages for Temporal Knowledge Graph Completion</i> Zhen Wang, Haotong Du, Quanming Yao and Xuelong Li	6160
<i>Code Vulnerability Detection via Nearest Neighbor Mechanism</i> Qianjin Du, Xiaohui Kuang and Gang Zhao	6173
<i>Robust Question Answering against Distribution Shifts with Test-Time Adaption: An Empirical Study</i> Hai Ye, Yuyang Ding, Juntao Li and Hwee Tou Ng	6179
<i>ParaMac: A General Unsupervised Paraphrase Generation Framework Leveraging Semantic Constraints and Diversifying Mechanisms</i> Jinxin Liu, Jiaxin Shi, Ji Qi, Lei Hou, Juanzi Li and Qi Tian	6193
<i>Semi-supervised New Slot Discovery with Incremental Clustering</i> Yuxia Wu, Lizi Liao, Xueming Qian and Tat-Seng Chua	6207
<i>Con-NAT: Contrastive Non-autoregressive Neural Machine Translation</i> Hao Cheng and Zhihua Zhang	6219
<i>Improved Knowledge Distillation for Pre-trained Language Models via Knowledge Selection</i> Chenglong Wang, Yi Lu, Yongyu Mu, Yimin Hu, Tong Xiao and Jingbo Zhu	6232
<i>Syntactically Robust Training on Partially-Observed Data for Open Information Extraction</i> Ji Qi, Yuxiang Chen, Lei Hou, Juanzi Li and Bin Xu	6245
<i>A Benchmark and Dataset for Post-OCR text correction in Sanskrit</i> Ayush Maheshwari, Nikhil Singh, Amrith Krishna and Ganesh Ramakrishnan	6258
<i>Knowledge-Enhanced Self-Supervised Prototypical Network for Few-Shot Event Detection</i> Kailin Zhao, Xiaolong Jin, Long Bai, Jiafeng Guo and Xueqi Cheng	6266
<i>VarMAE: Pre-training of Variational Masked Autoencoder for Domain-adaptive Language Understanding</i> Dou Hu, Xiaolong Hou, Xiyang Du, Mengyuan Zhou, Lianxin Jiang, Yang Mo and Xiaofeng Shi	6276
<i>Exploring Methods for Building Dialects-Mandarin Code-Mixing Corpora: A Case Study in Taiwanese Hokkien</i> Sin-En Lu, Bo-Han Lu, Chao-Yi Lu and Richard Tzong-Han Tsai	6287
<i>Recurrence Boosts Diversity! Revisiting Recurrent Latent Variable in Transformer-Based Variational AutoEncoder for Diverse Text Generation</i> Jinyi Hu, Xiaoyuan Yi, Wenhao Li, Maosong Sun and Xing Xie	6306
<i>Tweet Based Reach Aware Temporal Attention Network for NFT Valuation</i> Ramit Sawhney, Megh Thakkar, Ritesh Soun, Atula Neerkaje, Vasu Sharma, Dipanwita Guhathakurta and Sudheer Chava	6321
<i>Entity Embedding Completion for Wide-Coverage Entity Disambiguation</i> Daisuke Oba, Ikuya Yamada, Naoki Yoshinaga and Masashi Toyoda	6333
<i>Entity-level Interaction via Heterogeneous Graph for Multimodal Named Entity Recognition</i> Gang Zhao, Guanting Dong, Yidong Shi, Haolong Yan, Weiran Xu and Si Li	6345
<i>Status Biases in Deliberation Online: Evidence from a Randomized Experiment on ChangeMyView</i> Emaad Manzoor, Yohan Jo and Alan Montgomery	6351

<i>Empathetic and Emotionally Positive Conversation Systems with an Emotion-specific Query-Response Memory</i>	
Zhiliang Tian, Yinliang Wang, YIPING SONG, Chi ZHANG, Dongkyu Lee, Yingxiu Zhao, Dongsheng Li and Nevin L. Zhang	6364
<i>Trial2Vec: Zero-Shot Clinical Trial Document Similarity Search using Self-Supervision</i>	
Zifeng Wang and Jimeng Sun	6377
<i>From Mimicking to Integrating: Knowledge Integration for Pre-Trained Language Models</i>	
Lei Li, Yankai Lin, Xuancheng Ren, Guangxiang Zhao, Peng Li, Jie Zhou and Xu Sun	6391
<i>Model and Data Transfer for Cross-Lingual Sequence Labelling in Zero-Resource Settings</i>	
Iker García-Ferrero, Rodrigo Aggerri and German Rigau	6403
<i>Early Guessing for Dialect Identification</i>	
Vani Kanjirang, Tanja Samardzic, Fabio Rinaldi and Ljiljana Dolamic	6417
<i>R-AT: Regularized Adversarial Training for Natural Language Understanding</i>	
Shiwen Ni, Jiawen Li and Hung-Yu Kao	6427
<i>Multi-View Active Learning for Short Text Classification in User-Generated Data</i>	
Payam Karisani, Negin Karisani and Li Xiong	6441
<i>Forging Multiple Training Objectives for Pre-trained Language Models via Meta-Learning</i>	
Hongqiu Wu, Ruixue Ding, Hai Zhao, Boli Chen, Pengjun Xie, Fei Huang and Min Zhang	6454
<i>ConGen: Unsupervised Control and Generalization Distillation For Sentence Representation</i>	
Peerat Limkonchotiwat, Wuttikorn Ponwitayarat, Lalita Lowphansirikul, Ekapol Chuangsuwanich and Sarana Nutanong	6467
<i>Large-Scale Differentially Private BERT</i>	
Rohan Anil, Badih Ghazi, Vineet Gupta, Ravi Kumar and Pasin Manurangsi	6481
<i>Improving Zero-Shot Multilingual Translation with Universal Representations and Cross-Mapping</i>	
Shuhao Gu and Yang Feng	6492
<i>Controllable Fake Document Infilling for Cyber Deception</i>	
Yibo Hu, Yu Lin, Erick Skorupa Parolin, Latifur Khan and Kevin Hamlen	6505
<i>Weakly Supervised Headline Dependency Parsing</i>	
Adrian Benton, Tianze Shi, Ozan Irsoy and Igor Malioutov	6520
<i>BOOKSUM: A Collection of Datasets for Long-form Narrative Summarization</i>	
Wojciech Kryscinski, Nazneen Rajani, Divyansh Agarwal, Caiming Xiong and Dragomir Radev	6536
<i>Not All Errors are Equal: Learning Text Generation Metrics using Stratified Error Synthesis</i>	
Wenda Xu, Yi-Lin Tuan, Yujie Lu, Michael Saxon, Lei Li and William Yang Wang	6559
<i>Summarization as Indirect Supervision for Relation Extraction</i>	
Keming Lu, I-Hung Hsu, Wenxuan Zhou, Mingyu Derek Ma and Muhao Chen	6575
<i>DIGAT: Modeling News Recommendation with Dual-Graph Interaction</i>	
Zhiming Mao, Jian Li, Hongru Wang, Xingshan Zeng and Kam-Fai Wong	6595
<i>SMASH: Improving SMALL Language Models' Few-SHOT Ability with Prompt-Based Distillation</i>	
Yueqian Wang, Chang Liu, Kai Chen, Xi Wang and Dongyan Zhao	6608

<i>Consecutive Question Generation via Dynamic Multitask Learning</i> Yunji Li, Sujian Li and Xing Shi	6620
<i>Subword Segmental Language Modelling for Nguni Languages</i> Francois Meyer and Jan Buys	6636
<i>Towards Robust Visual Question Answering: Making the Most of Biased Samples via Contrastive Learning</i> Qingyi Si, Yuanxin LIU, Fandong Meng, Zheng Lin, Peng Fu, Yanan Cao, Weiping Wang and Jie Zhou	6650
<i>P3LM: Probabilistically Permuted Prophet Language Modeling for Generative Pre-Training</i> Junwei Bao, Yifan Wang, ying jiangyong, Yeyun Gong, Jing Zhao, Youzheng Wu and Xiaodong He	6663
<i>Holistic Sentence Embeddings for Better Out-of-Distribution Detection</i> Sishuo Chen, Xiaohan Bi, Rundong Gao and Xu Sun	6676
<i>MuGER2: Multi-Granularity Evidence Retrieval and Reasoning for Hybrid Question Answering</i> Yingyao Wang, Junwei Bao, Chaoqun Duan, Youzheng Wu, Xiaodong He and Tiejun Zhao	6687
<i>EntityCS: Improving Zero-Shot Cross-lingual Transfer with Entity-Centric Code Switching</i> Chenxi Whitehouse, Fenia Christopoulou and Ignacio Iacobacci	6698
<i>MBTI Personality Prediction for Fictional Characters Using Movie Scripts</i> Yisi Sang, Xiangyang Mou, Mo Yu, Dakuo Wang, Jing Li and Jeffrey Stanton	6715
<i>A Simple and Strong Baseline for End-to-End Neural RST-style Discourse Parsing</i> Naoki Kobayashi, Tsutomu Hirao, Hidetaka Kamigaito, Manabu Okumura and Masaaki Nagata	6725
<i>Probing for Constituency Structure in Neural Language Models</i> David Arps, Younes Samih, Laura Kallmeyer and Hassan Sajjad	6738
<i>Table-To-Text generation and pre-training with TabT5</i> Ewa Andrejczuk, Julian Eisenschlos, Francesco Piccinno, Syrine Krichene and Yasemin Altun	6758
<i>A POMDP Dialogue Policy with 3-way Grounding and Adaptive Sensing for Learning through Communication</i> Maryam Zare, Alan Wagner and Rebecca Passonneau	6767
<i>PaCo: Preconditions Attributed to Commonsense Knowledge</i> Ehsan Qasemi, Filip Ilievski, Muhao Chen and Pedro Szekely	6781
<i>Improving Few-Shot Domain Transfer for Named Entity Disambiguation with Pattern Exploitation</i> Philip Blair and Kfir Bar	6797
<i>Capturing Topic Framing via Masked Language Modeling</i> Xiaobo Guo, Weicheng Ma and Soroush Vosoughi	6811
<i>WANLI: Worker and AI Collaboration for Natural Language Inference Dataset Creation</i> Alisa Liu, Swabha Swayamdipta, Noah A. Smith and Yejin Choi	6826
<i>Sequentially Controlled Text Generation</i> Alexander Spangher, Yao Ming, Xinyu Hua and Nanyun Peng	6848

<i>Revisiting the Roles of Text” in Text Games</i>	Yi Gu, Shunyu Yao, Chuang Gan, Josh Tenenbaum and Mo Yu	6867
<i>FPT: Improving Prompt Tuning Efficiency via Progressive Training</i>	Yufei Huang, Yujia Qin, Huadong Wang, Yichun Yin, Maosong Sun, Zhiyuan Liu and Qun Liu	6877
<i>Prompt-learning for Fine-grained Entity Typing</i>	Ning Ding, Yulin Chen, Xu Han, Guangwei Xu, Xiaobin Wang, Pengjun Xie, Haitao Zheng, Zhiyuan Liu, Juanzi Li and Hong-Gee Kim	6888
<i>TransLIST: A Transformer-Based Linguistically Informed Sanskrit Tokenizer</i>	Jivnesh Sandhan, Rathin Singha, Narein Rao, SUVENDU SAMANTA, Laxmidhar Behera and Pawan Goyal	6902
<i>Fair NLP Models with Differentially Private Text Encoders</i>	Gaurav Maheshwari, Pascal Denis, Mikaela Keller and Aurélien Bellet	6913
<i>Modeling Context With Linear Attention for Scalable Document-Level Translation</i>	Zhaofeng Wu, Hao Peng, Nikolaos Pappas and Noah A. Smith	6931
<i>What do Large Language Models Learn beyond Language?</i>	Avinash Madasu and Shashank Srivastava	6940
<i>CONSISTENT: Open-Ended Question Generation From News Articles</i>	Tuhin Chakrabarty, Justin Lewis and Smaranda Muresan	6954
<i>Efficient (Soft) Q-Learning for Text Generation with Limited Good Data</i>	Han Guo, Bowen Tan, Zhengzhong Liu, Eric Xing and Zhiting Hu	6969
<i>Lexi: Self-Supervised Learning of the UI Language</i>	Pratyay Banerjee, Shweti Mahajan, Kushal Arora, Chitta Baral and Oriana Riva	6992
<i>Inferring the Reader: Guiding Automated Story Generation with Commonsense Reasoning</i>	Xiangyu Peng, Siyan Li, Sarah Wiegrefe and Mark Riedl	7008
<i>How to Stop an Avalanche? JoDeM: Joint Decision Making through Compare and Contrast for Dialog State Tracking</i>	Haoming Wang and Wang Xin	7030
<i>Contrastive Learning with Prompt-derived Virtual Semantic Prototypes for Unsupervised Sentence Embedding</i>	Jiali Zeng, Yongjing Yin, Yufan Jiang, Shuangzhi Wu and Yunbo Cao	7042
<i>Weight Perturbation as Defense against Adversarial Word Substitutions</i>	Jianhan Xu, Linyang Li, Jiping Zhang, Xiaoqing Zheng, Kai-Wei Chang, Cho-Jui Hsieh and Xuanjing Huang	7054
<i>CORT: A New Baseline for Comparative Opinion Classification by Dual Prompts</i>	Yequan Wang, Hengran Zhang, Aixin Sun and Xuying Meng	7064
<i>APEACH: Attacking Pejorative Expressions with Analysis on Crowd-Generated Hate Speech Evaluation Datasets</i>	Kichang Yang, Wonjun Jang and Won Ik Cho	7076
<i>Guiding Neural Story Generation with Reader Models</i>	Xiangyu Peng, Kaige Xie, Amal Alabdulkarim, Harshith Kayam, Samihan Dani and Mark Riedl	7087

<i>Reason first, then respond: Modular Generation for Knowledge-infused Dialogue</i> Leonard Adolphs, Kurt Shuster, Jack Urbanek, Arthur Szlam and Jason Weston	7112
<i>Adapting Multilingual Models for Code-Mixed Translation</i> Aditya Vavre, Abhirut Gupta and Sunita Sarawagi	7133
<i>LPC: A Logits and Parameter Calibration Framework for Continual Learning</i> Xiaodi Li, Zhuoyi Wang, Dingcheng Li, Latifur Khan and Bhavani Thuraisingham	7142
<i>SlovakBERT: Slovak Masked Language Model</i> Matúš Pikuliak, Štefan Grivalský, Martin Konôpka, Miroslav Blšták, Martin Tamajka, Viktor Bachratý, Marian Simko, Pavol Balážik, Michal Trnka and Filip Uhlárik	7156
<i>Efficient Zero-shot Event Extraction with Context-Definition Alignment</i> Hongming Zhang, Wenlin Yao and Dong Yu	7169
<i>Logical Fallacy Detection</i> Zhijing Jin, Abhinav Lalwani, Tejas Vaidhya, Xiaoyu Shen, Yiwen Ding, Zhiheng Lyu, Mrinmaya Sachan, Rada Mihalcea and Bernhard Schoelkopf	7180
<i>Topic-Aware Response Generation in Task-Oriented Dialogue with Unstructured Knowledge Access</i> Yue Feng, Gerasimos Lampouras and Ignacio Iacobacci	7199
<i>Revisiting Transformer-based Models for Long Document Classification</i> Xiang Dai, Ilias Chalkidis, Sune Darkner and Desmond Elliott	7212
<i>Time-aware Prompting for Text Generation</i> Shuyang Cao and Lu Wang	7231
<i>Improving Scheduled Sampling with Elastic Weight Consolidation for Neural Machine Translation</i> Michalis Korakakis and Andreas Vlachos	7247
<i>Ensemble Transformer for Efficient and Accurate Ranking Tasks: an Application to Question Answering Systems</i> Yoshitomo Matsubara, Luca Soldaini, Eric Lind and Alessandro Moschitti	7259
<i>Uncertainty Quantification with Pre-trained Language Models: A Large-Scale Empirical Analysis</i> Yuxin Xiao, Paul Pu Liang, Umang Bhatt, Willie Neiswanger, Ruslan Salakhutdinov and Louis-Philippe Morency	7273
<i>How to Represent Context Better? An Empirical Study on Context Modeling for Multi-turn Response Selection</i> Jiazhan Feng, Chongyang Tao, Chang Liu, Rui Yan and Dongyan Zhao	7285
<i>CHIA: CHoosing Instances to Annotate for Machine Translation</i> Rajat Bhatnagar, Ananya Ganesh and Katharina Kann	7299
<i>Guiding Neural Machine Translation with Semantic Kernels</i> Ping Guo, Yue Hu, Xiangpeng Wei, Yubing Ren, Yunpeng Li, Luxi Xing and Yuqiang Xie	7316
<i>HiSMATCH: Historical Structure Matching based Temporal Knowledge Graph Reasoning</i> Zixuan Li, Zhongni Hou, Saiping Guan, Xiaolong Jin, Weihua Peng, Long Bai, Yajuan Lyu, Wei Li, Jiafeng Guo and Xueqi Cheng	7328
<i>Dependency Parsing via Sequence Generation</i> Boda Lin, Zijun Yao, Jiabin Shi, Shulin Cao, Binghao Tang, Si Li, Yong Luo, Juanzi Li and Lei Hou	7339

<i>Scaling Laws Under the Microscope: Predicting Transformer Performance from Small Scale Experiments</i>	
Maor Ivgi, Yair Carmon and Jonathan Berant	7354
<i>Analyzing the Limits of Self-Supervision in Handling Bias in Language</i>	
Lisa Bauer, Karthik Gopalakrishnan, Spandana Gella, Yang Liu, Mohit Bansal and Dilek Hakkani-Tur	7372
<i>Multiple Instance Learning for Offensive Language Detection</i>	
Jiexi Liu, Dehan Kong, Longtao Huang, Dinghui Mao and hui xue	7387
<i>Grounded Keys-to-Text Generation: Towards Factual Open-Ended Generation</i>	
Faeze Brahman, Baolin Peng, Michel Galley, Sudha Rao, Bill Dolan, Snigdha Chaturvedi and Jianfeng Gao	7397
<i>Human-in-the-Loop Hate Speech Classification in a Multilingual Context</i>	
Ana Kotarcic, Dominik Hangartner, Fabrizio Gilardi, Selina Kurer and Karsten Donnay	7414