

61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)

Long Papers

Toronto, Canada
9 – 14 July 2023

Volume 1 of 20

ISBN: 978-1-7138-8173-5

Printed from e-media with permission by:

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



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

Copyright© (2023) by the Association for Computational Linguistics
All rights reserved.

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

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

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

Phone: 1-570-476-8006

Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

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

Table of Contents

<i>One Cannot Stand for Everyone! Leveraging Multiple User Simulators to train Task-oriented Dialogue Systems</i> Yajiao LIU, Xin Jiang, Yichun Yin, Yasheng Wang, Fei Mi, Qun Liu, Xiang Wan and Benyou Wang	1
<i>SafeConv: Explaining and Correcting Conversational Unsafe Behavior</i> Mian Zhang, Lifeng Jin, Linfeng Song, Haitao Mi, Wenliang Chen and Dong Yu	22
<i>Detecting and Mitigating Hallucinations in Machine Translation: Model Internal Workings Alone Do Well, Sentence Similarity Even Better</i> David Dale, Elena Voita, Loic Barrault and Marta R. Costa-jussà	36
<i>Explainable Recommendation with Personalized Review Retrieval and Aspect Learning</i> Hao Cheng, Shuo Wang, Wensheng Lu, Wei Zhang, Mingyang Zhou, Kezhong Lu and Hao Liao	51
<i>Binary and Ternary Natural Language Generation</i> Zechun Liu, Barlas Oguz, Aasish Pappu, Yangyang Shi and Raghuraman Krishnamoorthi	65
<i>Span-Selective Linear Attention Transformers for Effective and Robust Schema-Guided Dialogue State Tracking</i> Björn Bebensee and Haejun Lee	78
<i>EM Pre-training for Multi-party Dialogue Response Generation</i> Yiyang Li and Hai Zhao	92
<i>ACLM: A Selective-Denoising based Generative Data Augmentation Approach for Low-Resource Complex NER</i> Sreyan Ghosh, Utkarsh Tyagi, Manan Suri, Sonal Kumar, Ramaneswaran S and Dinesh Manocha	104
<i>Natural Language to Code Generation in Interactive Data Science Notebooks</i> Pengcheng Yin, Wen-Ding Li, Kefan Xiao, Abhishek K Rao, Yeming Wen, Kensen Shi, Joshua Howland, Paige Bailey, Michele Catasta, Henryk Michalewski, Oleksandr Polozov and Charles Sutton	126
<i>Subset Retrieval Nearest Neighbor Machine Translation</i> Hiroyuki Deguchi, Taro Watanabe, Yusuke Matsui, Masao Utiyama, Hideki Tanaka and Eiichiro Sumita	174
<i>MIL-Decoding: Detoxifying Language Models at Token-Level via Multiple Instance Learning</i> Xu Zhang and Xiaojun Wan	190
<i>Dependency resolution at the syntax-semantics interface: psycholinguistic and computational insights on control dependencies</i> Iria de-Dios-Flores, Juan Pablo Garcia Amboage and Marcos Garcia	203
<i>Open-ended Long Text Generation via Masked Language Modeling</i> Xiaobo Liang, Zecheng Tang, Juntao Li and Min Zhang	223
<i>A Method for Studying Semantic Construal in Grammatical Constructions with Interpretable Contextual Embedding Spaces</i> Gabiella Chronis, Kyle Mahowald and Katrin Erk	242

<i>Holographic CCG Parsing</i>	
Ryosuke Yamaki, Tadahiro Taniguchi and Daichi Mochihashi	262
<i>Prompts Can Play Lottery Tickets Well: Achieving Lifelong Information Extraction via Lottery Prompt Tuning</i>	
Zujie Liang, feng wei, Yin Jie, YUXI QIAN, Zhenghong Hao and Bing Han	277
<i>Retrieve-and-Sample: Document-level Event Argument Extraction via Hybrid Retrieval Augmentation</i>	
Yubing Ren, Yanan Cao, Ping Guo, Fang Fang, Wei Ma and Zheng Lin	293
<i>WeCheck: Strong Factual Consistency Checker via Weakly Supervised Learning</i>	
Wenhao Wu, Wei Li, Xinyan Xiao, Jiachen Liu, Sujian Li and Yajuan Lyu	307
<i>AMR-based Network for Aspect-based Sentiment Analysis</i>	
Fukun Ma, Xuming Hu, Aiwei Liu, Yawen Yang, Shuang Li, Philip S. Yu and Lijie Wen	322
<i>Text Adversarial Purification as Defense against Adversarial Attacks</i>	
Linyang Li, Demin Song and Xipeng Qiu	338
<i>SPEECH: Structured Prediction with Energy-Based Event-Centric Hyperspheres</i>	
Shumin Deng, Shengyu Mao, Ningyu Zhang and Bryan Hooi	351
<i>Rule By Example: Harnessing Logical Rules for Explainable Hate Speech Detection</i>	
Christopher Clarke, Matthew Hall, Gaurav Mittal, Ye Yu, Sandra Sajeev, Jason Mars and Mei Chen	364
<i>What about em"? How Commercial Machine Translation Fails to Handle (Neo-)Pronouns</i>	
Anne Lauscher, Debora Nozza, Ehm Miltersen, Archie Crowley and Dirk Hovy	377
<i>What Is Overlap Knowledge in Event Argument Extraction? APE: A Cross-datasets Transfer Learning Model for EAE</i>	
Kaihong Zhang, Kai Shuang, Xinyue Yang, Xuyang Yao and Jinyu Guo	393
<i>Tailor: A Soft-Prompt-Based Approach to Attribute-Based Controlled Text Generation</i>	
Kexin Yang, Dayiheng Liu, Wenqiang Lei, Baosong Yang, Mingfeng Xue, Boxing Chen and Jun Xie	410
<i>Knowledge of cultural moral norms in large language models</i>	
Aida Ramezani and Yang Xu	428
<i>Songs Across Borders: Singable and Controllable Neural Lyric Translation</i>	
Longshen Ou, Xichu Ma, Min-Yen Kan and Ye Wang	447
<i>Fantastic Expressions and Where to Find Them: Chinese Simile Generation with Multiple Constraints</i>	
Kexin Yang, Dayiheng Liu, Wenqiang Lei, Baosong Yang, Xiangpeng Wei, Zhengyuan Liu and Jun Xie	468
<i>Revealing Single Frame Bias for Video-and-Language Learning</i>	
Jie Lei, Tamara Berg and Mohit Bansal	487
<i>Learning with Partial Annotations for Event Detection</i>	
Jian Liu, Dianbo Sui, Kang Liu, Haoyan Liu and Zhe Zhao	508
<i>World-to-Words: Grounded Open Vocabulary Acquisition through Fast Mapping in Vision-Language Models</i>	
Ziqiao Ma, Jiayi Pan and Joyce Chai	524

<i>A Causal Framework to Quantify the Robustness of Mathematical Reasoning with Language Models</i> Alessandro Stolfo, Zhijing Jin, Kumar Shridhar, Bernhard Schoelkopf and Mrinmaya Sachan	545
<i>Evaluating Open-Domain Dialogues in Latent Space with Next Sentence Prediction and Mutual Information</i> Kun Zhao, Bohao Yang, Chenghua Lin, Wenge Rong, Aline Villavicencio and Xiaohui Cui	562
<i>Increasing Diversity While Maintaining Accuracy: Text Data Generation with Large Language Models and Human Interventions</i> John Joon Young Chung, Ece Kamar and Saleema Amershi	575
<i>Pruning Pre-trained Language Models Without Fine-Tuning</i> Ting Jiang, deqing wang, Fuzhen Zhuang, Ruobing Xie and Feng Xia	594
<i>When Does Translation Require Context? A Data-driven, Multilingual Exploration</i> Patrick Fernandes, Kayo Yin, Emmy Liu, André Martins and Graham Neubig	606
<i>Causal Intervention and Counterfactual Reasoning for Multi-modal Fake News Detection</i> Ziwei Chen, Linmei Hu, Weixin Li, Yingxia Shao and Liqiang Nie	627
<i>LexSym: Compositionality as Lexical Symmetry</i> Ekin Akyurek and Jacob Andreas	639
<i>Layer-wise Fusion with Modality Independence Modeling for Multi-modal Emotion Recognition</i> Jun Sun, Shoukang Han, Yu-Ping Ruan, Xiaoning Zhang, Shu-Kai Zheng, Yulong Liu, Yuxin Huang and Taihao Li	658
<i>CASN: Class-Aware Score Network for Textual Adversarial Detection</i> Rong Bao, Rui Zheng, Liang Ding, Qi Zhang and Dacheng Tao	671
<i>Do Androids Laugh at Electric Sheep? Humor Understanding Benchmarks from The New Yorker Caption Contest</i> Jack Hessel, Ana Marasovic, Jena D. Hwang, Lillian Lee, Jeff Da, Rowan Zellers, Robert Mankoff and Yejin Choi	688
<i>Making More of Little Data: Improving Low-Resource Automatic Speech Recognition Using Data Augmentation</i> Martijn Bartelds, Nay San, Bradley McDonnell, Dan Jurafsky and Martijn Wieling	715
<i>CLCL: Non-compositional Expression Detection with Contrastive Learning and Curriculum Learning</i> Jianing Zhou, Ziheng Zeng and Suma Bhat	730
<i>Multi-VALUE: A Framework for Cross-Dialectal English NLP</i> Caleb Ziems, William Held, Jingfeng Yang, Jwala Dhamala, Rahul Gupta and Diyi Yang	744
<i>Self-Edit: Fault-Aware Code Editor for Code Generation</i> Kechi Zhang, Zhuo Li, Jia Li, Ge Li and Zhi Jin	769
<i>Cold Fusion: Collaborative Descent for Distributed Multitask Finetuning</i> Shachar Don-Yehiya, Elad Venezian, Colin Raffel, Noam Slonim and Leshem Choshen	788
<i>Test-time Adaptation for Machine Translation Evaluation by Uncertainty Minimization</i> Runzhe Zhan, Xuebo Liu, Derek F. Wong, Cuilian Zhang, Lidia S. Chao and Min Zhang	807
<i>Multi-CLS BERT: An Efficient Alternative to Traditional Ensembling</i> Haw-Shiuan Chang, Ruei-Yao Sun, Kathryn Ricci and Andrew McCallum	821

<i>On-the-fly Cross-lingual Masking for Multilingual Pre-training</i> Xi Ai and Bin Fang	855
<i>How About Kind of Generating Hedges using End-to-End Neural Models?</i> Alafate Abulimiti, Chloé Clavel and Justine Cassell	877
<i>DiffusionDB: A Large-scale Prompt Gallery Dataset for Text-to-Image Generative Models</i> Zijie J. Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover and Duen Horng Chau	893
<i>From Key Points to Key Point Hierarchy: Structured and Expressive Opinion Summarization</i> Arie Cattan, Lilach Eden, Yoav Kantor and Roy Bar-Haim	912
<i>When to Use What: An In-Depth Comparative Empirical Analysis of OpenIE Systems for Downstream Applications</i> Kevin Song Pei, Ishan Jindal, Kevin Chen-Chuan Chang, ChengXiang Zhai and Yunyao Li ..	929
<i>Subjective Crowd Disagreements for Subjective Data: Uncovering Meaningful CrowdOpinion with Population-level Learning</i> Tharindu Cyril Weerasooriya, Sarah Luger, Saloni Poddar, Ashiqur KhudaBukhsh and Christopher Homan	950
<i>Post-Abstention: Towards Reliably Re-Attempting the Abstained Instances in QA</i> Neeraj Varshney and Chitta Baral	967
<i>UniLG: A Unified Structure-aware Framework for Lyrics Generation</i> Tao Qian, Fan Lou, Jiatong Shi, Yuning Wu, Shuai Guo, Xiang Yin and Qin Jin	983
<i>FC-KBQA: A Fine-to-Coarse Composition Framework for Knowledge Base Question Answering</i> Lingxi Zhang, Jing Zhang, Yanling Wang, Shulin Cao, Xinmei Huang, Cuiping Li, Hong Chen and Juanzi Li	1002
<i>Does GPT-3 Grasp Metaphors? Identifying Metaphor Mappings with Generative Language Models</i> Lennart Wachowiak and Dagmar Gromann	1018
<i>Being Right for Whose Right Reasons?</i> Terne Sasha Thorn Jakobsen, Laura Cabello and Anders Søgaard	1033
<i>ALERT: Adapt Language Models to Reasoning Tasks</i> Ping Yu, Tianlu Wang, Olga Golovneva, Badr AlKhamissi, Siddharth Verma, Zhijing Jin, Gargi Ghosh, Mona Diab and Asli Celikyilmaz	1055
<i>Glott500: Scaling Multilingual Corpora and Language Models to 500 Languages</i> Ayyoob ImaniGooghari, Peiqin Lin, Amir Hossein Kargaran, Silvia Severini, Masoud Jalili Sabet, Nora Kassner, Chunlan Ma, Helmut Schmid, André Martins, François Yvon and Hinrich Schütze	1082
<i>Joint Constrained Learning with Boundary-adjusting for Emotion-Cause Pair Extraction</i> Huawen Feng, Junlong Liu, Junhao Zheng, Haibin Chen, Xichen Shang and Qianli Ma	1118
<i>Pretrained Bidirectional Distillation for Machine Translation</i> Yimeng Zhuang and Mei Tu	1132
<i>Pivotal Role of Language Modeling in Recommender Systems: Enriching Task-specific and Task-agnostic Representation Learning</i> Kyu Yong Shin, Hanock Kwak, Wonjae Kim, Jisu Jeong, Seungjae Jung, Kyungmin Kim, Jung-Woo Ha and Sang-Woo Lee	1146

<i>Improving Continual Relation Extraction by Distinguishing Analogous Semantics</i> Wenzheng Zhao, Yuaning Cui and Wei Hu	1162
<i>Improving Pretraining Techniques for Code-Switched NLP</i> Richeek Das, Sahasra Ranjan, Shreya Pathak and Preethi Jyothi	1176
<i>A Theory of Unsupervised Speech Recognition</i> Liming Wang, Mark Hasegawa-Johnson and Chang D. Yoo	1192
<i>ThinkSum: Probabilistic reasoning over sets using large language models</i> Batu M Ozturkler, Nikolay Malkin, Zhen Wang and Nebojsa Jojic	1216
<i>NLG Evaluation Metrics Beyond Correlation Analysis: An Empirical Metric Preference Checklist</i> Iftitahu Nimah, Meng Fang, Vlado Menkovski and Mykola Pechenizkiy	1240
<i>DialoGPS: Dialogue Path Sampling in Continuous Semantic Space for Data Augmentation in Multi-Turn Conversations</i> Ang Lv, Jinpeng Li, yuhan chen, GAO XING, Ji Zhang and Rui Yan	1267
<i>TECHS: Temporal Logical Graph Networks for Explainable Extrapolation Reasoning</i> Qika Lin, Jun Liu, Rui Mao, Fangzhi Xu and Erik Cambria	1281
<i>Consistency Regularization Training for Compositional Generalization</i> Yongjing Yin, Jiali Zeng, Yafu Li, Fandong Meng, Jie Zhou and Yue Zhang	1294
<i>NUWA-XL: Diffusion over Diffusion for eXtremely Long Video Generation</i> Shengming Yin, Chenfei Wu, Huan Yang, Jianfeng Wang, Xiaodong Wang, Minheng Ni, Zhenyuan Yang, Linjie Li, Shuguang Liu, Fan Yang, Jianlong Fu, Ming Gong, Lijuan Wang, Zicheng Liu, Houqiang Li and Nan Duan	1309
<i>Synthetic Text Generation with Differential Privacy: A Simple and Practical Recipe</i> Xiang Yue, Huseyin A Inan, Xuechen Li, Girish Kumar, Julia McAnallen, Hoda Shajari, Huan Sun, David Levitan and Robert Sim	1321
<i>A Close Look into the Calibration of Pre-trained Language Models</i> Yangyi Chen, Lifan Yuan, Ganqu CUI, Zhiyuan Liu and Heng Ji	1343
<i>DIONYSUS: A Pre-trained Model for Low-Resource Dialogue Summarization</i> Yu Li, Baolin Peng, Pengcheng He, Michel Galley, Zhou Yu and Jianfeng Gao	1368
<i>MS-DETR: Natural Language Video Localization with Sampling Moment-Moment Interaction</i> Wang Jing, Aixin Sun, Hao Zhang and Xiaoli Li	1387
<i>Diverse Demonstrations Improve In-context Compositional Generalization</i> Itay Levy, Ben Bogin and Jonathan Berant	1401
<i>Self-Adaptive In-Context Learning: An Information Compression Perspective for In-Context Example Selection and Ordering</i> Zhiyong Wu, Yaoxiang Wang, Jiacheng Ye and Lingpeng Kong	1423
<i>On the Efficacy of Sampling Adapters</i> Clara Meister, Tiago Pimentel, Luca Malagutti, Ethan Wilcox and Ryan Cotterell	1437
<i>Cross-Domain Data Augmentation with Domain-Adaptive Language Modeling for Aspect-Based Sentiment Analysis</i> Jianfei Yu, Qiankun Zhao and Rui Xia	1456

<i>Compositional Data Augmentation for Abstractive Conversation Summarization</i> Siru Ouyang, Jiaao Chen, Jiawei Han and Diyi Yang	1471
<i>PMAES: Prompt-mapping Contrastive Learning for Cross-prompt Automated Essay Scoring</i> Yuan Chen and Xia Li	1489
<i>Marked Personas: Using Natural Language Prompts to Measure Stereotypes in Language Models</i> Myra Cheng, Esin Durmus and Dan Jurafsky	1504
<i>On Prefix-tuning for Lightweight Out-of-distribution Detection</i> Yawen Ouyang, Yongchang Cao, Yuan Gao, Zhen Wu, Jianbing Zhang and Xinyu Dai	1533
<i>GEC-DePenD: Non-Autoregressive Grammatical Error Correction with Decoupled Permutation and Decoding</i> Konstantin Yakovlev, Alexander Podolskiy, Andrey Bout, Sergey I Nikolenko and Irina Piontkovskaya	1546
<i>Measuring Progress in Fine-grained Vision-and-Language Understanding</i> Emanuele Bugliarello, Laurent Sartran, Aishwarya Agrawal, Lisa Anne Hendricks and Aida Nematzadeh	1559
<i>Vision Meets Definitions: Unsupervised Visual Word Sense Disambiguation Incorporating Gloss Information</i> Sunjae Kwon, Rishabh Garodia, Minhwa Lee, Zhichao Yang and hong yu	1583
<i>Chain-of-Skills: A Configurable Model for Open-Domain Question Answering</i> Kaixin Ma, Hao Cheng, Yu Zhang, Xiaodong Liu, Eric Nyberg and Jianfeng Gao	1599
<i>Elaboration-Generating Commonsense Question Answering at Scale</i> Wenya Wang, Vivek Srikumar, Hannaneh Hajishirzi and Noah A. Smith	1619
<i>Neural Unsupervised Reconstruction of Protolanguage Word Forms</i> Andre W He, Nicholas Tomlin and Dan Klein	1636
<i>DaMSTF: Domain Adversarial Learning Enhanced Meta Self-Training for Domain Adaptation</i> Menglong Lu, Zhen Huang, Yunxiang Zhao, Zhiliang Tian, Yang Liu and Dongsheng Li	1650
<i>On Evaluating Multilingual Compositional Generalization with Translated Datasets</i> Zi Wang and Daniel Hershcovich	1669
<i>FAA: Fine-grained Attention Alignment for Cascade Document Ranking</i> Zhen Li, Chongyang Tao, Jiazhan Feng, Tao Shen, Dongyan Zhao, Xiubo Geng and Daxin Jiang	1688
<i>Fine-tuning Happens in Tiny Subspaces: Exploring Intrinsic Task-specific Subspaces of Pre-trained Language Models</i> Zhong Zhang, Bang Liu and Junming Shao	1701
<i>Facilitating Multi-turn Emotional Support Conversation with Positive Emotion Elicitation: A Reinforcement Learning Approach</i> Jinfeng Zhou, Zhuang Chen, Bo Wang and Minlie Huang	1714
<i>Query Enhanced Knowledge-Intensive Conversation via Unsupervised Joint Modeling</i> Mingzhu Cai, Siqi Bao, Xin Tian, Huang He, Fan Wang and Hua Wu	1730

<i>Why Aren't We NER Yet? Artifacts of ASR Errors in Named Entity Recognition in Spontaneous Speech Transcripts</i>	
Piotr Szymański, Lukasz Augustyniak, Mikolaj Morzy, Adrian Szymczak, Krzysztof Surdyk and Piotr Żelasko	1746
<i>Precise Zero-Shot Dense Retrieval without Relevance Labels</i>	
Luyu Gao, Xueguang Ma, Jimmy Lin and Jamie Callan	1762
<i>White-Box Multi-Objective Adversarial Attack on Dialogue Generation</i>	
Yufei Li, Zexin Li, Yingfan Gao and Cong Liu	1778
<i>A Cautious Generalization Goes a Long Way: Learning Morphophonological Rules</i>	
Salam Khalifa, Sarah Payne, Jordan Kodner, Ellen Broselow and Owen Rambow	1793
<i>Few-shot Adaptation Works with Unpredictable Data</i>	
Jun Shern Chan, Michael Pieler, Jonathan Jao, Jérémy Scheurer and Ethan Perez	1806
<i>Cross-lingual Science Journalism: Select, Simplify and Rewrite Summaries for Non-expert Readers</i>	
Mehwish Fatima and Michael Strube	1843
<i>HuCurl: Human-induced Curriculum Discovery</i>	
Mohamed Elgaar and Hadi Amiri	1862
<i>kNN-TL: k-Nearest-Neighbor Transfer Learning for Low-Resource Neural Machine Translation</i>	
Shudong Liu, Xuebo Liu, Derek F. Wong, Zhaocong Li, Wenxiang Jiao, Lidia S. Chao and Min Zhang	1878
<i>Do language models have coherent mental models of everyday things?</i>	
Yuling Gu, Bhavana Dalvi Mishra and Peter Clark	1892
<i>Rogue Scores</i>	
Max Grusky	1914
<i>Instruction Induction: From Few Examples to Natural Language Task Descriptions</i>	
Or Honovich, Uri Shaham, Samuel R. Bowman and Omer Levy	1935
<i>In-Context Analogical Reasoning with Pre-Trained Language Models</i>	
Xiaoyang Hu, Shane Storcks, Richard L. Lewis and Joyce Chai	1953
<i>Peek Across: Improving Multi-Document Modeling via Cross-Document Question-Answering</i>	
Avi Caciularu, Matthew Peters, Jacob Goldberger, Ido Dagan and Arman Cohan	1970
<i>Tailoring Instructions to Student's Learning Levels Boosts Knowledge Distillation</i>	
Yuxin Ren, Zihan Zhong, Xingjian Shi, Yi Zhu, Chun Yuan and Mu Li	1990
<i>REV: Information-Theoretic Evaluation of Free-Text Rationales</i>	
Hanjie Chen, Faeze Brahman, Xiang Ren, Yangfeng Ji, Yejin Choi and Swabha Swayamdipta	2007
<i>ELQA: A Corpus of Metalinguistic Questions and Answers about English</i>	
Shabnam Behzad, Keisuke Sakaguchi, Nathan Schneider and Amir Zeldes	2031
<i>Divide, Conquer, and Combine: Mixture of Semantic-Independent Experts for Zero-Shot Dialogue State Tracking</i>	
Qingyue Wang, Liang Ding, Yanan Cao, Yibing Zhan, Zheng Lin, Shi Wang, Dacheng Tao and Li Guo	2048
<i>BIG-C: a Multimodal Multi-Purpose Dataset for Bemba</i>	
Claytone Sikasote, Eunice Mukonde, Md Mahfuz Ibn Alam and Antonios Anastasopoulos .	2062

<i>Schema-Guided User Satisfaction Modeling for Task-Oriented Dialogues</i>	
Yue Feng, Yunlong Jiao, Animesh Prasad, Nikolaos Aletras, Emine Yilmaz and Gabriella Kazai	
2079	
<i>Robust Multi-bit Natural Language Watermarking through Invariant Features</i>	
KiYoon Yoo, Wonhyuk Ahn, Jiho Jang and Nojun Kwak	2092
<i>KALM: Knowledge-Aware Integration of Local, Document, and Global Contexts for Long Document Understanding</i>	
Shangbin Feng, Zhaoxuan Tan, Wenqian Zhang, Zhenyu Lei and Yulia Tsvetkov	2116
<i>AtTGen: Attribute Tree Generation for Real-World Attribute Joint Extraction</i>	
Yanzeng Li, Bingcong Xue, Ruoyu Zhang and Lei Zou	2139
<i>Extractive is not Faithful: An Investigation of Broad Unfaithfulness Problems in Extractive Summarization</i>	
Shiyue Zhang, David Wan and Mohit Bansal	2153
<i>Improving Translation Quality Estimation with Bias Mitigation</i>	
Hui Huang, Shuangzhi Wu, Kehai Chen, Hui Di, Muyun Yang and Tiejun Zhao	2175
<i>Breeding Machine Translations: Evolutionary approach to survive and thrive in the world of automated evaluation</i>	
Josef Jon and Ondřej Bojar	2191
<i>MoralDial: A Framework to Train and Evaluate Moral Dialogue Systems via Moral Discussions</i>	
Hao Sun, Zhixin Zhang, Fei Mi, Yasheng Wang, Wei Liu, Jianwei Cui, Bin Wang, Qun Liu and Minlie Huang	2213
<i>Denoising Bottleneck with Mutual Information Maximization for Video Multimodal Fusion</i>	
Shaoxiang Wu, Damai Dai, Ziwei Qin, Tianyu Liu, Binghuai Lin, Yunbo Cao and Zhifang Sui	
2231	
<i>SimLM: Pre-training with Representation Bottleneck for Dense Passage Retrieval</i>	
Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder and Furu Wei	2244
<i>From Ultra-Fine to Fine: Fine-tuning Ultra-Fine Entity Typing Models to Fine-grained</i>	
Hongliang Dai and Ziqian Zeng	2259
<i>Controlling Learned Effects to Reduce Spurious Correlations in Text Classifiers</i>	
Parikshit Bansal and Amit Sharma	2271
<i>What Makes Pre-trained Language Models Better Zero-shot Learners?</i>	
Jinghui Lu, dongsheng zhu, weidong Han, Rui Zhao, Brian Mac Namee and Fei Tan	2288
<i>Z-ICL: Zero-Shot In-Context Learning with Pseudo-Demonstrations</i>	
Xinxi Lyu, Sewon Min, Iz Beltagy, Luke Zettlemoyer and Hannaneh Hajishirzi	2304
<i>Learning Optimal Policy for Simultaneous Machine Translation via Binary Search</i>	
Shoutao Guo, Shaolei Zhang and Yang Feng	2318
<i>Better Simultaneous Translation with Monotonic Knowledge Distillation</i>	
Shushu Wang, Jing Wu, Kai Fan, Wei Luo, Jun Xiao and Zhongqiang Huang	2334
<i>StoryARG: a corpus of narratives and personal experiences in argumentative texts</i>	
Neele Falk and Gabriella Lapesa	2350

<i>Injecting knowledge into language generation: a case study in auto-charting after-visit care instructions from medical dialogue</i>	
Maksim Ereemeev, Ilya Valmianski, Xavier Amatriain and Anitha Kannan	2373
<i>Sequence Parallelism: Long Sequence Training from System Perspective</i>	
Shenggui Li, Fuzhao Xue, Chaitanya Baranwal, Yongbin Li and Yang You	2391
<i>MUSTIE: Multimodal Structural Transformer for Web Information Extraction</i>	
Qifan Wang, Jingang Wang, Xiaojun Quan, Fuli Feng, Zenglin Xu, Shaoliang Nie, Sinong Wang, Madian Khabsa, Hamed Firooz and Dongfang Liu	2405
<i>Augmentation-Adapted Retriever Improves Generalization of Language Models as Generic Plug-In</i>	
Zichun Yu, Chenyan Xiong, Shi Yu and Zhiyuan Liu	2421
<i>TableVLM: Multi-modal Pre-training for Table Structure Recognition</i>	
Leiyan Chen, Chengsong Huang, Xiaoqing Zheng, Jinshu Lin and Xuanjing Huang	2437
<i>Can NLI Provide Proper Indirect Supervision for Low-resource Biomedical Relation Extraction?</i>	
Jiashu Xu, Mingyu Derek Ma and Muhao Chen	2450
<i>Dynamic Routing Transformer Network for Multimodal Sarcasm Detection</i>	
Yuan Tian, Nan Xu, Ruike Zhang and Wenji Mao	2468
<i>What Are You Token About? Dense Retrieval as Distributions Over the Vocabulary</i>	
Ori Ram, Liat Bezael, Adi Zicher, Yonatan Belinkov, Jonathan Berant and Amir Globerson	2481
<i>Cold-Start Data Selection for Better Few-shot Language Model Fine-tuning: A Prompt-based Uncertainty Propagation Approach</i>	
Yue Yu, Rongzhi Zhang, Ran Xu, Jieyu Zhang, Jiaming Shen and Chao Zhang	2499
<i>Training-free Neural Architecture Search for RNNs and Transformers</i>	
Aaron Serianni and Jugal Kalita	2522
<i>CrossSum: Beyond English-Centric Cross-Lingual Summarization for 1,500+ Language Pairs</i>	
Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, Yuan-Fang Li, Yong-Bin Kang and Rifat Shahriyar	2541
<i>Improving Gradient Trade-offs between Tasks in Multi-task Text Classification</i>	
Heyan Chai, Jinhao Cui, Ye Wang, Min Zhang, Binxing Fang and Qing Liao	2565
<i>Bi-Phone: Modeling Inter Language Phonetic Influences in Text</i>	
Abhirut Gupta, Ananya B. Sai, Richard Sproat, Yuri Vasilevski, James S Ren, Ambarish Jash, Sukhdeep S Sodhi and Aravindan Raghuv eer	2580
<i>Cross2StrA: Unpaired Cross-lingual Image Captioning with Cross-lingual Cross-modal Structure-pivoted Alignment</i>	
Shengqiong Wu, Hao Fei, Wei Ji and Tat-Seng Chua	2593
<i>Plan-and-Solve Prompting: Improving Zero-Shot Chain-of-Thought Reasoning by Large Language Models</i>	
Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee and Ee-Peng Lim	2609
<i>RetroMAE-2: Duplex Masked Auto-Encoder For Pre-Training Retrieval-Oriented Language Models</i>	
Zheng Liu, Shitao Xiao, Yingxia Shao and Zhao Cao	2635

<i>DecompX: Explaining Transformers Decisions by Propagating Token Decomposition</i> Ali Modarressi, Mohsen Fayyaz, Ehsan Aghazadeh, Yadollah Yaghoobzadeh and Mohammad Taher Pilehvar	2649
<i>Symbolic Chain-of-Thought Distillation: Small Models Can Also Think Step-by-Step</i> Liunian Harold Li, Jack Hessel, Youngjae Yu, Xiang Ren, Kai-Wei Chang and Yejin Choi ..	2665
<i>Generating EDU Extracts for Plan-Guided Summary Re-Ranking</i> Griffin Adams, Alex Fabbri, Faisal Ladhak, Noémie Elhadad and Kathleen McKeown	2680
<i>A Survey on Asking Clarification Questions Datasets in Conversational Systems</i> Hossein A. Rahmani, Xi Wang, Yue Feng, Qiang Zhang, Emine Yilmaz and Aldo Lipani ..	2698
<i>Towards Understanding Chain-of-Thought Prompting: An Empirical Study of What Matters</i> Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer and Huan Sun	2717
<i>Small Data, Big Impact: Leveraging Minimal Data for Effective Machine Translation</i> Jean Maillard, Cynthia Gao, Elahe Kalbassi, Kaushik Ram Sadagopan, Vedanuj Goswami, Philipp Koehn, Angela Fan and Francisco Guzman	2740
<i>RMLM: A Flexible Defense Framework for Proactively Mitigating Word-level Adversarial Attacks</i> Zhaoyang Wang, Zhiyue Liu, Xiaopeng Zheng, Qinliang Su and Jiahai Wang	2757
<i>Gradient-based Intra-attention Pruning on Pre-trained Language Models</i> Ziqing Yang, Yiming Cui, Xin Yao and Shijin Wang	2775
<i>Learning to Substitute Spans towards Improving Compositional Generalization</i> Zhaoyi Li, Ying Wei and Defu Lian	2791
<i>DiffusEmp: A Diffusion Model-Based Framework with Multi-Grained Control for Empathetic Response Generation</i> Guanqun Bi, Lei Shen, Yanan Cao, Meng Chen, Yuqiang Xie, Zheng Lin and Xiaodong He	2812
<i>BREAK: Breaking the Dialogue State Tracking Barrier with Beam Search and Re-ranking</i> Seungpil Won, Heeyoung Kwak, Joongbo Shin, Janghoon Han and Kyomin Jung	2832
<i>Faithful Low-Resource Data-to-Text Generation through Cycle Training</i> Zhuoer Wang, Marcus Collins, Nikhita Vedula, Simone Filice, Shervin Malmasi and Oleg Rokhlenko	2847
<i>Towards Stable Natural Language Understanding via Information Entropy Guided Debiasing</i> Li Du, Xiao Ding, Zhouhao Sun, Ting Liu, Bing Qin and Jingshuo Liu	2868
<i>Dynamic and Efficient Inference for Text Generation via BERT Family</i> Xiaobo Liang, Juntao Li, Lijun Wu, Ziqiang Cao and Min Zhang	2883
<i>Learning to Generate Equitable Text in Dialogue from Biased Training Data</i> Anthony B Sicilia and Malihe Alikhani	2898
<i>Hierarchical Verbalizer for Few-Shot Hierarchical Text Classification</i> Ke Ji, Yixin Lian, Jingsheng Gao and Baoyuan Wang	2918
<i>Summary-Oriented Vision Modeling for Multimodal Abstractive Summarization</i> Yunlong Liang, Fandong Meng, Jinan Xu, Jiaan Wang, Yufeng Chen and Jie Zhou	2934

<i>Helping a Friend or Supporting a Cause? Disentangling Active and Passive Cosponsorship in the U.S. Congress</i>	
Giuseppe Russo, Christoph Gote, Laurence Brandenberger, Sophia Johanna Schlosser and Frank Schweitzer	2952
<i>TREA: Tree-Structure Reasoning Schema for Conversational Recommendation</i>	
Wendi Li, Wei Wei, Xiaoye Qu, Xian-Ling Mao, Ye Yuan, Wenfeng Xie and Dangyang Chen	2970
<i>CATS: A Pragmatic Chinese Answer-to-Sequence Dataset with Large Scale and High Quality</i>	
Liang Li, Ruiying Geng, Chengyang Fang, Bing Li, Can Ma, Rongyu Cao, Binhua Li, Fei Huang and Yongbin Li	2983
<i>Multilingual Multifaceted Understanding of Online News in Terms of Genre, Framing, and Persuasion Techniques</i>	
Jakub Piskorski, Nicolas Stefanovitch, Nikolaos Nikolaidis, Giovanni Da San Martino and Preslav Nakov	3001
<i>Learning Action Conditions from Instructional Manuals for Instruction Understanding</i>	
Te-Lin Wu, Caiqi ZHANG, Qingyuan Hu, Alexander Spangher and Nanyun Peng	3023
<i>StoryWars: A Dataset and Instruction Tuning Baselines for Collaborative Story Understanding and Generation</i>	
Yulun Du and Lydia Chilton	3044
<i>Did You Read the Instructions? Rethinking the Effectiveness of Task Definitions in Instruction Learning</i>	
Fan Yin, Jesse Vig, Philippe Laban, Shafiq Joty, Caiming Xiong and Chien-Sheng Jason Wu	3063
<i>Do PLMs Know and Understand Ontological Knowledge?</i>	
Weiqi Wu, Chengyue Jiang, Yong Jiang, Pengjun Xie and Kewei Tu	3080
<i>CORE: Cooperative Training of Retriever-Reranker for Effective Dialogue Response Selection</i>	
Chongyang Tao, Jiazhan Feng, Tao Shen, Chang Liu, Juntao Li, Xiubo Geng and Daxin Jiang	3102
<i>Exploring How Generative Adversarial Networks Learn Phonological Representations</i>	
Jingyi Chen and Micha Elsner	3115
<i>Interpretable Word Sense Representations via Definition Generation: The Case of Semantic Change Analysis</i>	
Mario Giulianelli, Iris Luden, Raquel Fernandez and Andrey Kutuzov	3130
<i>Learning to Simulate Natural Language Feedback for Interactive Semantic Parsing</i>	
Hao Yan, Saurabh Srivastava, Yintao Tai, Sida I. Wang, Wen-tau Yih and Ziyu Yao	3149
<i>InfoMetIC: An Informative Metric for Reference-free Image Caption Evaluation</i>	
Anwen Hu, Shizhe Chen, Liang Zhang and Qin Jin	3171
<i>An Invariant Learning Characterization of Controlled Text Generation</i>	
Carolina Zheng, Claudia Shi, Keyon Vafa, Amir Feder and David Blei	3186
<i>HistRED: A Historical Document-Level Relation Extraction Dataset</i>	
Soyoung Yang, Minseok Choi, Youngwoo Cho and Jaegul Choo	3207
<i>A Critical Evaluation of Evaluations for Long-form Question Answering</i>	
Fangyuan Xu, Yixiao Song, Mohit Iyyer and Eunsol Choi	3225
<i>HyPe: Better Pre-trained Language Model Fine-tuning with Hidden Representation Perturbation</i>	
Hongyi Yuan, Zheng Yuan, Chuanqi Tan, Fei Huang and Songfang Huang	3246

<i>Generating User-Engaging News Headlines</i>	
Pengshan Cai, Kaiqiang Song, Sangwoo Cho, Hongwei Wang, Xiaoyang Wang, hong yu, Fei Liu and Dong Yu	3265
<i>Word sense extension</i>	
Lei Yu and Yang Xu	3281
<i>PVGRU: Generating Diverse and Relevant Dialogue Responses via Pseudo-Variational Mechanism</i>	
Yongkang Liu, Shi Feng, Daling Wang, Yifei Zhang and Hinrich Schütze	3295
<i>Decoding Symbolism in Language Models</i>	
Meiqi Guo, Rebecca Hwa and Adriana Kovashka	3311
<i>A Survey on Zero Pronoun Translation</i>	
Longyue Wang, Siyou Liu, Mingzhou Xu, Linfeng Song, Shuming Shi and Zhaopeng Tu ..	3325
<i>We Understand Elliptical Sentences, and Language Models should Too: A New Dataset for Studying Ellipsis and its Interaction with Thematic Fit</i>	
Davide Testa, Emmanuele Chersoni and Alessandro Lenci	3340
<i>MPCHAT: Towards Multimodal Persona-Grounded Conversation</i>	
Jaewoo Ahn, Yeda Song, Sangdoo Yun and Gunhee Kim	3354
<i>DOC: Improving Long Story Coherence With Detailed Outline Control</i>	
Kevin Yang, Dan Klein, Nanyun Peng and Yuandong Tian	3378
<i>Dual-Alignment Pre-training for Cross-lingual Sentence Embedding</i>	
Ziheng Li, Shaohan Huang, Zihan Zhang, Zhi-Hong Deng, Qiang Lou, Haizhen Huang, Jian Jiao, Furu Wei, Weiwei Deng and Qi Zhang	3466
<i>Exploring Better Text Image Translation with Multimodal Codebook</i>	
Zhibin Lan, Jiawei Yu, Xiang Li, Wen Zhang, Jian Luan, Bin Wang, Degen Huang and Jinsong Su	3479
<i>FEDLEGAL: The First Real-World Federated Learning Benchmark for Legal NLP</i>	
Zhuo Zhang, Xiangjing Hu, Jingyuan Zhang, Yating Zhang, Hui Wang, Lizhen Qu and Zenglin Xu	3492
<i>A Gradient Control Method for Backdoor Attacks on Parameter-Efficient Tuning</i>	
Naibin Gu, Peng Fu, Xiyu Liu, Zhengxiao Liu, Zheng Lin and Weiping Wang	3508
<i>History Semantic Graph Enhanced Conversational KBQA with Temporal Information Modeling</i>	
Hao Sun, Yang Li, Liwei Deng, Bowen Li, Binyuan Hui, Binhua Li, Yunshi Lan, Yan Zhang and Yongbin Li	3521
<i>From the One, Judge of the Whole: Typed Entailment Graph Construction with Predicate Generation</i>	
Zhibin Chen, Yansong Feng and Dongyan Zhao	3534
<i>Alleviating Over-smoothing for Unsupervised Sentence Representation</i>	
Nuo Chen, Linjun Shou, Jian Pei, Ming Gong, Bowen Cao, Jianhui Chang, Jia Li and Daxin Jiang	3552
<i>Memory-efficient NLLB-200: Language-specific Expert Pruning of a Massively Multilingual Machine Translation Model</i>	
Yeskendir Koishkenov, Alexandre Berard and Vassilina Nikoulina	3567

<i>DAMP: Doubly Aligned Multilingual Parser for Task-Oriented Dialogue</i> William Held, Christopher Hidey, Fei Liu, Eric Y Zhu, Rahul Goel, Diyi Yang and Rushin Shah	3586
<i>From Characters to Words: Hierarchical Pre-trained Language Model for Open-vocabulary Language Understanding</i> Li Sun, Florian Luisier, Kayhan Batmanghelich, Dinei Florencio and Cha Zhang	3605
<i>MatSci-NLP: Evaluating Scientific Language Models on Materials Science Language Tasks Using Text-to-Schema Modeling</i> Yu Song, Santiago Miret and Bang Liu	3621
<i>Code4Struct: Code Generation for Few-Shot Event Structure Prediction</i> Xingyao Wang, Sha Li and Heng Ji	3640
<i>GENEVA: Benchmarking Generalizability for Event Argument Extraction with Hundreds of Event Types and Argument Roles</i> Tanmay Parekh, I-Hung Hsu, Kuan-Hao Huang, Kai-Wei Chang and Nanyun Peng	3664
<i>Efficient Semiring-Weighted Earley Parsing</i> Andreas Opedal, Ran Zmigrod, Tim Vieira, Ryan Cotterell and Jason Eisner	3687
<i>Tree-Based Representation and Generation of Natural and Mathematical Language</i> Alexander Scarlatos and Andrew Lan	3714
<i>ParaLS: Lexical Substitution via Pretrained Paraphraser</i> Jipeng Qiang, Kang Liu, Yun Li, Yunhao Yuan and Yi Zhu	3731
<i>Peer-Label Assisted Hierarchical Text Classification</i> Junru Song, Feifei Wang and Yang Yang	3747
<i>Free Lunch for Efficient Textual Commonsense Integration in Language Models</i> Wanyun Cui and Xingran Chen	3759
<i>A Probabilistic Framework for Discovering New Intents</i> Yunhua Zhou, Guofeng Quan and Xipeng Qiu	3771
<i>MultiTACRED: A Multilingual Version of the TAC Relation Extraction Dataset</i> Leonhard Hennig, Philippe Thomas and Sebastian Möller	3785
<i>Towards Higher Pareto Frontier in Multilingual Machine Translation</i> yichong huang, Xiaocheng Feng, Xinwei Geng, Baohang Li and Bing Qin	3802
<i>Small Pre-trained Language Models Can be Fine-tuned as Large Models via Over-Parameterization</i> Ze-Feng Gao, Kun Zhou, Peiyu Liu, Wayne Xin Zhao and Ji-Rong Wen	3819
<i>Entity Tracking in Language Models</i> Najoung Kim and Sebastian Schuster	3835
<i>A Textual Dataset for Situated Proactive Response Selection</i> Naoki Otani, Jun Araki, HyeongSik Kim and Eduard H Hovy	3856
<i>DiffusionNER: Boundary Diffusion for Named Entity Recognition</i> Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu and Yueting Zhuang	3875
<i>WACO: Word-Aligned Contrastive Learning for Speech Translation</i> Siqi Ouyang, Rong Ye and Lei Li	3891

<i>Cross-lingual Continual Learning</i>	
Meryem M’hamdi, Xiang Ren and Jonathan May	3908
<i>Faithful Question Answering with Monte-Carlo Planning</i>	
Ruixin Hong, Hongming Zhang, Hong Zhao, Dong Yu and Changshui Zhang	3944
<i>Unbalanced Optimal Transport for Unbalanced Word Alignment</i>	
Yuki Arase, Han Bao and Sho Yokoi	3966
<i>Guiding Computational Stance Detection with Expanded Stance Triangle Framework</i>	
Zhengyuan Liu, Yong Keong Yap, Hai Leong Chieu and Nancy Chen	3987
<i>Analyzing and Reducing the Performance Gap in Cross-Lingual Transfer with Fine-tuning Slow and Fast</i>	
Yiduo Guo, Yaobo Liang, Dongyan Zhao, Bing Liu and Nan Duan	4002
<i>Improving Self-training for Cross-lingual Named Entity Recognition with Contrastive and Prototype Learning</i>	
Ran Zhou, Xin Li, Lidong Bing, Erik Cambria and Chunyan Miao	4018
<i>MM-SHAP: A Performance-agnostic Metric for Measuring Multimodal Contributions in Vision and Language Models & Tasks</i>	
Letitia Parcalabescu and Anette Frank	4032
<i>Towards Boosting the Open-Domain Chatbot with Human Feedback</i>	
Hua Lu, Siqi Bao, Huang He, Fan Wang, Hua Wu and Haifeng Wang	4060
<i>Knowledge-enhanced Mixed-initiative Dialogue System for Emotional Support Conversations</i>	
Yang Deng, Wenxuan Zhang, Yifei Yuan and Wai Lam	4079
<i>UTC-IE: A Unified Token-pair Classification Architecture for Information Extraction</i>	
Hang Yan, Yu Sun, Xiaonan Li, Yunhua Zhou, Xuanjing Huang and Xipeng Qiu	4096
<i>Social-Group-Agnostic Bias Mitigation via the Stereotype Content Model</i>	
Ali Omrani, Alireza Salkhordeh Ziabari, Charles Yu, Preni Golazizian, Brendan Kennedy, Mohammad Atari, Heng Ji and Morteza Dehghani	4123
<i>Revisiting the Gold Standard: Grounding Summarization Evaluation with Robust Human Evaluation</i>	
Yixin Liu, Alex Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Jason Wu, Caiming Xiong and Dragomir Radev	4140
<i>FIREBALL: A Dataset of Dungeons and Dragons Actual-Play with Structured Game State Information</i>	
Andrew Zhu, Karmanya Aggarwal, Alexander H Feng, Lara J. Martin and Chris Callison-Burch	4171
<i>A fine-grained comparison of pragmatic language understanding in humans and language models</i>	
Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina Fedorenko and Edward Gibson	4194
<i>Counterfactual Multihop QA: A Cause-Effect Approach for Reducing Disconnected Reasoning</i>	
Wangzhen Guo, Qinkang Gong, Yanghui Rao and Hanjiang Lai	4214
<i>Causal-Debias: Unifying Debiasing in Pretrained Language Models and Fine-tuning via Causal Invariant Learning</i>	
Fan Zhou, Yuzhou Mao, Liu Yu, Yi Yang and Ting Zhong	4227
<i>Parameter-Efficient Fine-Tuning without Introducing New Latency</i>	
Baohao Liao, Yan Meng and Christof Monz	4242

<i>MANNER: A Variational Memory-Augmented Model for Cross Domain Few-Shot Named Entity Recognition</i>	
Jinyuan Fang, Xiaobin Wang, Zaiqiao Meng, Pengjun Xie, Fei Huang and Yong Jiang	4261
<i>MASSIVE: A 1M-Example Multilingual Natural Language Understanding Dataset with 51 Typologically-Diverse Languages</i>	
Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur and Prem Natarajan	4277
<i>Distilling Script Knowledge from Large Language Models for Constrained Language Planning</i>	
Siyu Yuan, Jiangjie Chen, Ziquan Fu, Xuyang Ge, Soham Pranav Shah, Charles Jankowski, Yanghua Xiao and Deqing Yang	4303
<i>RED^{FM}: a Filtered and Multilingual Relation Extraction Dataset</i>	
Pere Lluís Huguet Cabot, Simone Tedeschi, Axel-Cyrille Ngonga Ngomo and Roberto Navigli	4326
<i>Modeling Appropriate Language in Argumentation</i>	
Timon Ziegenbein, Shahbaz Syed, Felix Lange, Martin Potthast and Henning Wachsmuth . .	4344
<i>CELDA: Leveraging Black-box Language Model as Enhanced Classifier without Labels</i>	
Hyunsoo Cho, Youna Kim and Sang-goo Lee	4364
<i>MvP: Multi-view Prompting Improves Aspect Sentiment Tuple Prediction</i>	
Zhibin Gou, qingyan guo and Yujiu Yang	4380
<i>ACCENT: An Automatic Event Commonsense Evaluation Metric for Open-Domain Dialogue Systems</i>	
Sarik Ghazarian, Yijia Shao, Rujun Han, Aram Galstyan and Nanyun Peng	4398
<i>Explanation-based Finetuning Makes Models More Robust to Spurious Cues</i>	
Josh Magnus Ludan, Yixuan Meng, Tai Nguyen, Saurabh Shah, Qing Lyu, Marianna Apidianaki and Chris Callison-Burch	4420
<i>CAME: Confidence-guided Adaptive Memory Efficient Optimization</i>	
Yang Luo, Xiaozhe REN, Zangwei Zheng, ZHUO JIANG, Xin Jiang and Yang You	4442
<i>On Second Thought, Let's Not Think Step by Step! Bias and Toxicity in Zero-Shot Reasoning</i>	
Omar Shaikh, Hongxin Zhang, William Held, Michael S Bernstein and Diyi Yang	4454
<i>Solving Math Word Problems via Cooperative Reasoning induced Language Models</i>	
Xinyu Zhu, Junjie Wang, Lin Zhang, Yuxiang Zhang, Yongfeng Huang, ruyi gan, Jiaying Zhang and Yujiu Yang	4471
<i>Exploiting Biased Models to De-bias Text: A Gender-Fair Rewriting Model</i>	
Chantal Amrhein, Florian Schottmann, Rico Sennrich and Samuel Lüubli	4486
<i>Early Discovery of Disappearing Entities in Microblogs</i>	
Satoshi Akasaki, Naoki Yoshinaga and Masashi Toyoda	4507
<i>DiffusionBERT: Improving Generative Masked Language Models with Diffusion Models</i>	
Zhengfu He, Tianxiang Sun, Qiong Tang, Kuanning Wang, Xuanjing Huang and Xipeng Qiu	4521
<i>Lifting the Curse of Capacity Gap in Distilling Language Models</i>	
Chen Zhang, Yang Yang, Jiahao Liu, Jingang Wang, Yunsen Xian, Benyou Wang and Dawei Song	4535

<i>Towards Faithful Dialogues via Focus Learning</i>	
Yifan Deng, Xingsheng Zhang, Heyan Huang and Yue Hu	4554
<i>Back Translation for Speech-to-text Translation Without Transcripts</i>	
Qingkai Fang and Yang Feng	4567
<i>Prompter: Zero-shot Adaptive Prefixes for Dialogue State Tracking Domain Adaptation</i>	
Ibrahim Taha Aksu, Min-Yen Kan and Nancy Chen	4588
<i>Enhancing Dialogue Generation via Dynamic Graph Knowledge Aggregation</i>	
Chen Tang, Hongbo Zhang, Tyler Loakman, Chenghua Lin and Frank Guerin	4604
<i>Multi-modal Action Chain Abductive Reasoning</i>	
Mengze Li, Tianbao Wang, Jiahe Xu, Kairong Han, Shengyu Zhang, Zhou Zhao, Jiaxu Miao, wenqiao zhang, Shiliang Pu and Fei Wu	4617
<i>Exploring the Capacity of Pretrained Language Models for Reasoning about Actions and Change</i>	
Weinan He, Canming Huang, Zhanhao Xiao and Yongmei Liu	4629
<i>Unified Demonstration Retriever for In-Context Learning</i>	
Xiaonan Li, Kai Lv, Hang Yan, Tianyang Lin, Wei Zhu, Yuan Ni, GUOTONG XIE, Xiaoling Wang and Xipeng Qiu	4644
<i>Movie101: A New Movie Understanding Benchmark</i>	
Zihao Yue, Qi Zhang, Anwen Hu, Liang Zhang, Ziheng Wang and Qin Jin	4669
<i>Enhancing Language Representation with Constructional Information for Natural Language Understanding</i>	
Lvxiaowei Xu, Jianwang Wu, Jiawei Peng, Zhilin Gong, Ming Cai and Tianxiang Wang	4685
<i>Query Structure Modeling for Inductive Logical Reasoning Over Knowledge Graphs</i>	
Siyuan Wang, Zhongyu Wei, meng han, Zhihao Fan, Haijun Shan, Qi Zhang and Xuanjing Huang	4706
<i>DimonGen: Diversified Generative Commonsense Reasoning for Explaining Concept Relationships</i>	
Chenzhenyi Liu, Jie Huang, Kerui Zhu and Kevin Chen-Chuan Chang	4719
<i>Incorporating Attribution Importance for Improving Faithfulness Metrics</i>	
Zhixue Zhao and Nikolaos Aletras	4732
<i>Reward Gaming in Conditional Text Generation</i>	
Richard Yuanzhe Pang, Vishakh Padmakumar, Thibault Sellam, Ankur Parikh and He He ..	4746
<i>Hidden Schema Networks</i>	
Ramses J Sanchez, Lukas Alexander Conrads, Pascal Welke, Kostadin Cvejovski and Cesar Ali Ojeda Marin	4764
<i>Towards Robust Low-Resource Fine-Tuning with Multi-View Compressed Representations</i>	
Linlin Liu, Xingxuan Li, Megh Thakkar, Xin Li, Shafiq Joty, Luo Si and Lidong Bing	4799
<i>An Ordinal Latent Variable Model of Conflict Intensity</i>	
Niklas Stoehr, Lucas Torroba Hennigen, Josef Valvoda, Robert West, Ryan Cotterell and Aaron Schein	4817
<i>Multilingual Conceptual Coverage in Text-to-Image Models</i>	
Michael S Saxon and William Yang Wang	4831

<i>Pre-Training to Learn in Context</i>	
Yuxian Gu, Li Dong, Furu Wei and Minlie Huang	4849
<i>Ethical Considerations for Machine Translation of Indigenous Languages: Giving a Voice to the Speakers</i>	
Manuel Mager, Elisabeth Albine Mager, Katharina Kann and Ngoc Thang Vu	4871
<i>Revisiting non-English Text Simplification: A Unified Multilingual Benchmark</i>	
Michael Joseph Ryan, Tarek Naous and Wei Xu	4898
<i>Don't Generate, Discriminate: A Proposal for Grounding Language Models to Real-World Environments</i>	
Yu Gu, Xiang Deng and Yu Su	4928
<i>Privacy-Preserving Domain Adaptation of Semantic Parsers</i>	
Fatemehsadat Mireshghallah, Yu Su, Tatsunori Hashimoto, Jason Eisner and Richard Shin	4950
<i>Guide the Many-to-One Assignment: Open Information Extraction via IoU-aware Optimal Transport</i>	
Kaiwen Wei, Yiran Yang, li jin, Xian Sun, Zequn Zhang, Jingyuan Zhang, Xiao yu Li, Linhao Zhang, Jintao Liu and Guo Zhi	4971
<i>Actively Supervised Clustering for Open Relation Extraction</i>	
Jun Zhao, Yongxin Zhang, Qi Zhang, Tao Gui, Zhongyu Wei, Minlong Peng and Mingming Sun	4985
<i>ConvGQR: Generative Query Reformulation for Conversational Search</i>	
Fengran Mo, Kelong Mao, Yutao Zhu, Yihong Wu, Kaiyu Huang and Jian-Yun Nie	4998
<i>KILM: Knowledge Injection into Encoder-Decoder Language Models</i>	
Yan Xu, Mahdi Namazifar, Devamanyu Hazarika, Aishwarya Padmakumar, Yang Liu and Dilek Hakkani-Tur	5013
<i>VSTAR: A Video-grounded Dialogue Dataset for Situated Semantic Understanding with Scene and Topic Transitions</i>	
Yuxuan Wang, Zilong Zheng, Xueliang Zhao, Jinpeng Li, Yueqian Wang and Dongyan Zhao	5036
<i>NLPeer: A Unified Resource for the Computational Study of Peer Review</i>	
Nils Dycke, Ilia Kuznetsov and Iryna Gurevych	5049
<i>IM-TQA: A Chinese Table Question Answering Dataset with Implicit and Multi-type Table Structures</i>	
Mingyu Zheng, Yang Hao, Wenbin Jiang, Zheng Lin, Yajuan Lyu, QiaoQiao She and Weiping Wang	5074
<i>Z-Code++: A Pre-trained Language Model Optimized for Abstractive Summarization</i>	
Pengcheng He, Baolin Peng, Song Wang, Yang Liu, Ruochen Xu, Hany Hassan, Yu Shi, Chengguang Zhu, Wayne Xiong, Michael Zeng, Jianfeng Gao and Xuedong Huang	5095
<i>Mixture-of-Domain-Adapters: Decoupling and Injecting Domain Knowledge to Pre-trained Language Models' Memories</i>	
Shizhe Diao, Tianyang Xu, Ruijia Xu, Jiawei Wang and Tong Zhang	5113
<i>Unsupervised Graph-Text Mutual Conversion with a Unified Pretrained Language Model</i>	
Yi Xu, Shuqian Sheng, Jiexing Qi, Luoyi Fu, Zhouhan Lin, Xinbing Wang and Chenghu Zhou	5130
<i>Randomized Smoothing with Masked Inference for Adversarially Robust Text Classifications</i>	
Han Cheol Moon, Shafiq Joty, Ruochen Zhao, Megh Thakkar and Chi Xu	5145

<i>SESCORE2: Learning Text Generation Evaluation via Synthesizing Realistic Mistakes</i> Wenda Xu, Xian Qian, Mingxuan Wang, Lei Li and William Yang Wang	5166
<i>Tokenization and the Noiseless Channel</i> Vilém Zouhar, Clara Meister, Juan Luis Gastaldi, Li Du, Mrinmaya Sachan and Ryan Cotterell	5184
<i>Contextual Distortion Reveals Constituency: Masked Language Models are Implicit Parsers</i> Jiayi Li and Wei Lu	5208
<i>MetaAdapt: Domain Adaptive Few-Shot Misinformation Detection via Meta Learning</i> Zhenrui Yue, Huimin Zeng, Yang Zhang, Lanyu Shang and Dong Wang	5223
<i>Tackling Modality Heterogeneity with Multi-View Calibration Network for Multimodal Sentiment Detection</i> yiwei wei, Shaozu Yuan, Ruosong Yang, Lei Shen, zhangmeizhi li, Longbiao Wang and Meng Chen.....	5240
<i>COLA: Contextualized Commonsense Causal Reasoning from the Causal Inference Perspective</i> Zhaowei Wang, Quyet V. Do, Hongming Zhang, Jiayao Zhang, Weiqi Wang, Tianqing Fang, Yangqiu Song, Ginny Y. Wong and Simon See	5253
<i>MEMEX: Detecting Explanatory Evidence for Memes via Knowledge-Enriched Contextualization</i> Shivam Sharma, Ramaneswaran S, Udit Arora, Md. Shad Akhtar and Tanmoy Chakraborty	5272
<i>WikiHowQA: A Comprehensive Benchmark for Multi-Document Non-Factoid Question Answering</i> Valeriia Bolotova-Baranova, Vladislav Blinov, Sofya Filippova, Falk Scholer and Mark Sanderson	5291
<i>Making Language Models Better Reasoners with Step-Aware Verifier</i> Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang LOU and Weizhu Chen	5315
<i>Distributed Marker Representation for Ambiguous Discourse Markers and Entangled Relations</i> Dongyu Ru, Lin Qiu, Xipeng Qiu, Yue Zhang and Zheng Zhang	5334
<i>MISGENDERED: Limits of Large Language Models in Understanding Pronouns</i> Tamanna Hossain, Sunipa Dev and Sameer Singh	5352
<i>Reasoning with Language Model Prompting: A Survey</i> Shuofei Qiao, Yixin Ou, Ningyu Zhang, Xiang Chen, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang and Huajun Chen	5368
<i>Tackling Ambiguity with Images: Improved Multimodal Machine Translation and Contrastive Evaluation</i> Matthieu Futral, Cordelia Schmid, Ivan Laptev, Benoît Sagot and Rachel Bawden	5394
<i>Hybrid Knowledge Transfer for Improved Cross-Lingual Event Detection via Hierarchical Sample Selection</i> Luis Fernando Guzman Nateras, Franck Deroncourt and Thien Huu Nguyen	5414
<i>BLEURT Has Universal Translations: An Analysis of Automatic Metrics by Minimum Risk Training</i> Yiming Yan, Tao Wang, Chengqi Zhao, Shujian Huang, Jiajun CHEN and Mingxuan Wang	5428
<i>Cross-modal Attention Congruence Regularization for Vision-Language Relation Alignment</i> Rohan Pandey, Rulin Shao, Paul Pu Liang, Ruslan Salakhutdinov and Louis-Philippe Morency	5444

<i>Enhancing Personalized Dialogue Generation with Contrastive Latent Variables: Combining Sparse and Dense Persona</i>	
Yihong Tang, Bo Wang, Miao Fang, Dongming Zhao, Kun Huang, Ruifang He and Yuexian Hou	
	5456
<i>Can LMs Learn New Entities from Descriptions? Challenges in Propagating Injected Knowledge</i>	
Yasumasa Onoe, Michael J.Q. Zhang, Shankar Padmanabhan, Greg Durrett and Eunsol Choi	
	5469
<i>Explaining How Transformers Use Context to Build Predictions</i>	
Javier Ferrando, Gerard I. Gállego, Ioannis Tsiamas and Marta R. Costa-jussà	
	5486
<i>DISCO: Distilling Counterfactuals with Large Language Models</i>	
Zeming Chen, Qiyue Gao, Antoine Bosselut, Ashish Sabharwal and Kyle Richardson	
	5514
<i>Non-Sequential Graph Script Induction via Multimedia Grounding</i>	
Yu Zhou, Sha Li, Manling Li, Xudong Lin, Shih-Fu Chang, Mohit Bansal and Heng Ji	
	5529
<i>SCOTT: Self-Consistent Chain-of-Thought Distillation</i>	
Peifeng Wang, Zhengyang Wang, Zheng Li, Yifan Gao, Bing Yin and Xiang Ren	
	5546
<i>Clinical Note Owns its Hierarchy: Multi-Level Hypergraph Neural Networks for Patient-Level Representation Learning</i>	
Nayeon Kim, Yinhua Piao and Sun Kim	
	5559
<i>Incorporating Distributions of Discourse Structure for Long Document Abstractive Summarization</i>	
Dongqi Pu, Yifan Wang and Vera Demberg	
	5574
<i>Evaluating Open-Domain Question Answering in the Era of Large Language Models</i>	
Ehsan Kamaloo, Nouha Dziri, Charles Clarke and Davood Rafiei	
	5591
<i>No clues good clues: out of context Lexical Relation Classification</i>	
Lucia Pitarch, Jordi Bernad, Lacramioara Dranca, Carlos Bobed Lisboa and Jorge Gracia	
	5607
<i>Won't Get Fooled Again: Answering Questions with False Premises</i>	
Shengding Hu, Yifan Luo, Huadong Wang, Xingyi Cheng, Zhiyuan Liu and Maosong Sun	
	5626
<i>What the DAAM: Interpreting Stable Diffusion Using Cross Attention</i>	
Raphael Tang, Linqing Liu, Akshat Pandey, Zhiying Jiang, Gefei Yang, Karun Kumar, Pontus Stenetorp, Jimmy Lin and Ferhan Ture	
	5644
<i>Zero-shot Faithful Factual Error Correction</i>	
Kung-Hsiang Huang, Hou Pong Chan and Heng Ji	
	5660
<i>Open-Domain Hierarchical Event Schema Induction by Incremental Prompting and Verification</i>	
Sha Li, Ruining Zhao, Manling Li, Heng Ji, Chris Callison-Burch and Jiawei Han	
	5677
<i>Zero-shot Approach to Overcome Perturbation Sensitivity of Prompts</i>	
Mohna Chakraborty, Adithya Kulkarni and Qi Li	
	5698
<i>Free Lunch: Robust Cross-Lingual Transfer via Model Checkpoint Averaging</i>	
Fabian David Schmidt, Ivan Vulić and Goran Glavaš	
	5712
<i>Cross-View Language Modeling: Towards Unified Cross-Lingual Cross-Modal Pre-training</i>	
Yan Zeng, Wangchunshu Zhou, Ao Luo, Ziming Cheng and Xinsong Zhang	
	5731
<i>Unsupervised Discontinuous Constituency Parsing with Mildly Context-Sensitive Grammars</i>	
Songlin Yang, Roger Levy and Yoon Kim	
	5747

<i>Simplicity Bias in Transformers and their Ability to Learn Sparse Boolean Functions</i> Satwik Bhattamishra, Arkil Patel, Varun Kanade and Phil Blunsom	5767
<i>Counterspeeches up my sleeve! Intent Distribution Learning and Persistent Fusion for Intent-Conditioned Counterspeech Generation</i> Rishabh Gupta, Shaily Desai, Manvi Goel, Anil Bandhakavi, Tanmoy Chakraborty and Md. Shad Akhtar	5792
<i>DITTO: Data-efficient and Fair Targeted Subset Selection for ASR Accent Adaptation</i> Suraj N Kothawade, Anmol Reddy Mekala, D.Chandra Sekhara SS Hetha Havya, Mayank Kothari, Rishabh K Iyer, Ganesh Ramakrishnan and Preethi Jyothi	5810
<i>Verify-and-Edit: A Knowledge-Enhanced Chain-of-Thought Framework</i> Ruochen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin and Lidong Bing	5823
<i>Bridging the Domain Gaps in Context Representations for k-Nearest Neighbor Neural Machine Translation</i> Zhiwei Cao, Baosong Yang, Huan Lin, Suhang Wu, Xiangpeng Wei, Dayiheng Liu, Jun Xie, Min Zhang and Jinsong Su	5841
<i>Node Placement in Argument Maps: Modeling Unidirectional Relations in High & Low-Resource Scenarios</i> Iman Jundi, Neele Falk, Eva Maria Vecchi and Gabriella Lapesa	5854
<i>Towards a Common Understanding of Contributing Factors for Cross-Lingual Transfer in Multilingual Language Models: A Review</i> Fred Philippy, Siwen Guo and Shohreh Haddadan	5877
<i>Toward Human-Like Evaluation for Natural Language Generation with Error Analysis</i> Qingyu Lu, Liang Ding, Liping Xie, Kanjian Zhang, Derek F. Wong and Dacheng Tao	5892
<i>Connective Prediction for Implicit Discourse Relation Recognition via Knowledge Distillation</i> Hongyi Wu, Hao Zhou, Man Lan, Yuanbin Wu and Yadong Zhang	5908
<i>What is the best recipe for character-level encoder-only modelling?</i> Kris Cao	5924
<i>Unifying Cross-Lingual and Cross-Modal Modeling Towards Weakly Supervised Multilingual Vision-Language Pre-training</i> Zejun Li, Zhihao Fan, Jingjing Chen, Qi Zhang, Xuanjing Huang and Zhongyu Wei	5939
<i>Learning OHelps for Learning More: Handling the Unlabeled Entity Problem for Class-incremental NER</i> Ruotian Ma, xuanting chen, zhang lin, Xin Zhou, Junzhe Wang, Tao Gui, Qi Zhang, Xiang Gao and Yun Wen Chen	5959
<i>Scene Graph as Pivoting: Inference-time Image-free Unsupervised Multimodal Machine Translation with Visual Scene Hallucination</i> Hao Fei, Qian Liu, Meishan Zhang, Min Zhang and Tat-Seng Chua	5980
<i>CoLaDa: A Collaborative Label Denoising Framework for Cross-lingual Named Entity Recognition</i> Tingting Ma, Qianhui Wu, Huiqiang Jiang, Börje F. Karlsson, Tiejun Zhao and Chin-Yew Lin	5995
<i>Dialect-robust Evaluation of Generated Text</i> Jiao Sun, Thibault Sellam, Elizabeth Clark, Tu Vu, Timothy Dozat, Dan Garrette, Aditya Siddhant, Jacob Eisenstein and Sebastian Gehrmann	6010

<i>Understanding and Improving the Robustness of Terminology Constraints in Neural Machine Translation</i>	
Huaao Zhang, Qiang Wang, Bo - Qin, Zelin Shi, haibo wang and MING CHEN	6029
<i>Language model acceptability judgements are not always robust to context</i>	
Koustuv Sinha, Jon Gauthier, Aaron Mueller, Kanishka Misra, Keren Fuentes, Roger Levy and Adina Williams	6043
<i>RobuT: A Systematic Study of Table QA Robustness Against Human-Annotated Adversarial Perturbations</i>	
Yilun Zhao, Chen Zhao, Linyong Nan, Zhenting Qi, Wenlin Zhang, Xiangru Tang, Boyu Mi and Dragomir Radev	6064
<i>Morphological Inflection: A Reality Check</i>	
Jordan Kodner, Sarah Payne, Salam Khalifa and Zoey Liu	6082
<i>TOME: A Two-stage Approach for Model-based Retrieval</i>	
Ruiyang Ren, Wayne Xin Zhao, Jing Liu, Hua Wu, Ji-Rong Wen and Haifeng Wang	6102
<i>Using Neural Machine Translation for Generating Diverse Challenging Exercises for Language Learner</i>	
Frank Palma Gomez, Subhadarshi Panda, Michael M Flor and Alla Rozovskaya	6115
<i>Similarity-weighted Construction of Contextualized Commonsense Knowledge Graphs for Knowledge-intensive Argumentation Tasks</i>	
Moritz Plen, Juri Opitz, Philipp Heinisch, Philipp Cimiano and Anette Frank	6130
<i>miCSE: Mutual Information Contrastive Learning for Low-shot Sentence Embeddings</i>	
Tassilo Klein and Moin Nabi	6159
<i>Learning Non-linguistic Skills without Sacrificing Linguistic Proficiency</i>	
Mandar Sharma, Nikhil Muralidhar and Naren Ramakrishnan	6178
<i>Forgotten Knowledge: Examining the Citational Amnesia in NLP</i>	
Janvijay Singh, Mukund Rungta, Diyi Yang and Saif M. Mohammad	6192
<i>Measuring the Instability of Fine-Tuning</i>	
Yupei Du and Dong Nguyen	6209
<i>FairPrism: Evaluating Fairness-Related Harms in Text Generation</i>	
Eve Fleisig, Aubrie N Amstutz, Chad Atalla, Su Lin Blodgett, Hal Daumé III, Alexandra Olteanu, Emily Sheng, Dan Vann and Hanna Wallach	6231
<i>Factually Consistent Summarization via Reinforcement Learning with Textual Entailment Feedback</i>	
Paul Roit, Johan Ferret, Lior Shani, Roei Aharoni, Geoffrey Cideron, Robert Dadashi, Matthieu Geist, Sertan Girgin, Leonard Hussenot, Orgad Keller, Nikola Momchev, Sabela Ramos Garea, Piotr Stanczyk, Nino Vieillard, Olivier Bachem, Gal Elidan, Avinatan Hassidim, Olivier Pietquin and Idan Szpektor	6252
<i>SIMMC-VR: A Task-oriented Multimodal Dialog Dataset with Situated and Immersive VR Streams</i>	
Te-Lin Wu, Satwik Kottur, Andrea Madotto, Mahmoud Azab, Pedro Rodriguez, Babak Damavandi, Nanyun Peng and Seungwhan Moon	6273
<i>Multilingual LLMs are Better Cross-lingual In-context Learners with Alignment</i>	
Eshaan Tanwar, Subhabrata Dutta, Manish Borthakur and Tanmoy Chakraborty	6292
<i>APOLLO: A Simple Approach for Adaptive Pretraining of Language Models for Logical Reasoning</i>	
Soumya Sanyal, Yichong Xu, Shuohang Wang, Ziyi Yang, Reid Pryzant, Wenhao Yu, Chenguang Zhu and Xiang Ren	6308

<i>MultiTabQA: Generating Tabular Answers for Multi-Table Question Answering</i>	
Vaishali Pal, Andrew Yates, Evangelos Kanoulas and Maarten de Rijke	6322
<i>To Copy Rather Than Memorize: A Vertical Learning Paradigm for Knowledge Graph Completion</i>	
Rui Li, Xu Chen, Chaozhuo Li, Yanming Shen, Jianan Zhao, Yujing Wang, Weihao Han, Hao Sun, Weiwei Deng, Qi Zhang and Xing Xie	6335
<i>CoAD: Automatic Diagnosis through Symptom and Disease Collaborative Generation</i>	
Huimin Wang, Wai Chung Kwan, Kam-Fai Wong and Yefeng Zheng	6348
<i>Long-Tailed Question Answering in an Open World</i>	
Yi Dai, Hao Lang, Yinhe Zheng, Fei Huang and Yongbin Li	6362
<i>Parallel Context Windows for Large Language Models</i>	
Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Dov Karpas, Amnon Shashua, Kevin Leyton-Brown and Yoav Shoham	6383
<i>Efficient Transformers with Dynamic Token Pooling</i>	
Piotr Nawrot, Jan Chorowski, Adrian Lancucki and Edoardo Maria Ponti	6403
<i>Did the Models Understand Documents? Benchmarking Models for Language Understanding in Document-Level Relation Extraction</i>	
Haotian Chen, Bingsheng Chen and Xiangdong Zhou	6418
<i>ContraCLM: Contrastive Learning For Causal Language Model</i>	
Nihal Jain, Dejiao Zhang, Wasi Uddin Ahmad, Zijian Wang, Feng Nan, Xiaopeng Li, Ming Tan, Ramesh Nallapati, Baishakhi Ray, Parminder Bhatia, Xiaofei Ma and Bing Xiang	6436
<i>Advancing Multi-Criteria Chinese Word Segmentation Through Criterion Classification and Denoising</i>	
Tzu Hsuan Chou, Chun-Yi Lin and Hung-Yu Kao	6460
<i>Infusing Hierarchical Guidance into Prompt Tuning: A Parameter-Efficient Framework for Multi-level Implicit Discourse Relation Recognition</i>	
Haodong Zhao, Ruifang He, Mengnan Xiao and Jing Xu	6477
<i>Contrastive Learning with Adversarial Examples for Alleviating Pathology of Language Model</i>	
Pengwei Zhan, Jing Yang, Xiao Huang, Chunlei Jing, Jingying Li and Liming Wang	6493
<i>Are Fairy Tales Fair? Analyzing Gender Bias in Temporal Narrative Event Chains of Children’s Fairy Tales</i>	
Paulina Toro Isaza, Guangxuan Xu, Toye Oloko, Yufang Hou, Nanyun Peng and Dakuo Wang	6509
<i>FutureTOD: Teaching Future Knowledge to Pre-trained Language Model for Task-Oriented Dialogue</i>	
Weihao Zeng, Keqing He, Yejie Wang, Chen Zeng, Jingang Wang, Yunsen Xian and Weiran Xu	6532
<i>LAMBADA: Backward Chaining for Automated Reasoning in Natural Language</i>	
Mehran Kazemi, Najoung Kim, Deepti Bhatia, Xin Xu and Deepak Ramachandran	6547
<i>PeaCoK: Persona Commonsense Knowledge for Consistent and Engaging Narratives</i>	
Silin Gao, Beatriz Borges, Soyoun Oh, Deniz Bayazit, Saya Kanno, Hiromi Wakaki, Yuki Mitsufuji and Antoine Bosselut	6569
<i>OpenSR: Open-Modality Speech Recognition via Maintaining Multi-Modality Alignment</i>	
Xize Cheng, Tao Jin, Linjun Li, Wang Lin, Xinyu Duan and Zhou Zhao	6592

<i>Retrieval-free Knowledge Injection through Multi-Document Traversal for Dialogue Models</i>	
Rui Wang, Jianzhu Bao, Fei Mi, Yi Chen, Hongru Wang, Yasheng Wang, Yitong Li, Lifeng Shang, Kam-Fai Wong and Ruifeng Xu	6608
<i>BERM: Training the Balanced and Extractable Representation for Matching to Improve Generalization Ability of Dense Retrieval</i>	
Shicheng Xu, Liang Pang, Huawei Shen and Xueqi Cheng	6620
<i>Multiview Identifiers Enhanced Generative Retrieval</i>	
Yongqi Li, Nan Yang, Liang Wang, Furu Wei and Wenjie Li	6636
<i>Prompting Language Models for Linguistic Structure</i>	
Terra Blevins, Hila Gonen and Luke Zettlemoyer	6649
<i>Trillion Dollar Words: A New Financial Dataset, Task & Market Analysis</i>	
Agam Shah, Suvan Satya Paturi and Sudheer Chava	6664
<i>RE-Matching: A Fine-Grained Semantic Matching Method for Zero-Shot Relation Extraction</i>	
Jun Zhao, WenYu Zhan, Xin Zhao, Qi Zhang, Tao Gui, Zhongyu Wei, Junzhe Wang, Minlong Peng and Mingming Sun	6680
<i>SQuARe: A Large-Scale Dataset of Sensitive Questions and Acceptable Responses Created through Human-Machine Collaboration</i>	
Hwaran Lee, Seokhee Hong, Joonsuk Park, Takyong Kim, Meeyoung Cha, Yejin Choi, BYOUNGPIL KIM, Gunhee Kim, Eun-Ju Lee, Yong Lim, Alice Oh, Sangchul Park and Jung-Woo Ha ...	6692
<i>Towards standardizing Korean Grammatical Error Correction: Datasets and Annotation</i>	
Soyoung Yoon, Sungjoon Park, Gyuwan Kim, Junhee Cho, Kihyo Park, Gyu Tae Kim, Minjoon Seo and Alice Oh	6713
<i>FLaME: Few-shot Learning from Natural Language Explanations</i>	
Yangqiaoyu Zhou, Yiming Zhang and Chenhao Tan	6743
<i>Learning Symbolic Rules over Abstract Meaning Representations for Textual Reinforcement Learning</i>	
Subhajit Chaudhury, Sarathkrishna Swaminathan, Daiki Kimura, Prithviraj Sen, Keerthiram Murgesan, Rosario Uceda-Sosa, Michiaki Tatsubori, Achille Fokoue, Pavan Kapanipathi, Asim Munawar and Alexander Gray	6764
<i>Counterfactual Debiasing for Fact Verification</i>	
Weizhi Xu, Qiang Liu, Shu Wu and Liang Wang	6777
<i>What social attitudes about gender does BERT encode? Leveraging insights from psycholinguistics</i>	
Julia Watson, Barend Beekhuizen and Suzanne Stevenson	6790
<i>Rethinking Multimodal Entity and Relation Extraction from a Translation Point of View</i>	
Changmeng Zheng, Junhao Feng, Yi Cai, Xiaoyong Wei and Qing Li	6810
<i>Annotating and Detecting Fine-grained Factual Errors for Dialogue Summarization</i>	
Rongxin Zhu, Jianzhong Qi and Jey Han Lau	6825
<i>Improving the Robustness of Summarization Systems with Dual Augmentation</i>	
Xiuying Chen, Guodong Long, Chongyang Tao, Mingzhe Li, Xin Gao, Chengqi Zhang and Xiangliang Zhang	6846
<i>Interpretable Math Word Problem Solution Generation via Step-by-step Planning</i>	
mengxue zhang, Zichao Wang, Zhichao Yang, weiqi feng and Andrew Lan	6858

<i>TemplateGEC: Improving Grammatical Error Correction with Detection Template</i> Yinghao Li, Xuebo Liu, Shuo Wang, Peiyuan Gong, Derek F. Wong, Yang Gao, Heyan Huang and Min Zhang	6878
<i>Deep Model Compression Also Helps Models Capture Ambiguity</i> Hancheol Park and Jong Park	6893
<i>Are Experts Needed? On Human Evaluation of Counselling Reflection Generation</i> Zixiu Wu, Simone Balloccu, Ehud Reiter, Rim Helaoui, Diego Reforgiato Recupero and Daniele Riboni	6906
<i>PairSpanBERT: An Enhanced Language Model for Bridging Resolution</i> Hideo Kobayashi, Yufang Hou and Vincent Ng	6931
<i>Compounding Geometric Operations for Knowledge Graph Completion</i> Xiou Ge, Yun Cheng Wang, Bin Wang and C.-C. Jay Kuo	6947
<i>Few-shot In-context Learning on Knowledge Base Question Answering</i> Tianle Li, Xueguang Ma, Alex Zhuang, Yu Gu, Yu Su and Wenhui Chen	6966
<i>Fact-Checking Complex Claims with Program-Guided Reasoning</i> Liangming Pan, Xiaobao Wu, Xinyuan Lu, Anh Tuan Luu, William Yang Wang, Min-Yen Kan and Preslav Nakov	6981
<i>Patton: Language Model Pretraining on Text-Rich Networks</i> Bowen Jin, Wentao Zhang, Yu Zhang, Yu Meng, Xinyang Zhang, Qi Zhu and Jiawei Han ..	7005
<i>Soft Language Clustering for Multilingual Model Pre-training</i> Jiali Zeng, Yufan Jiang, Yongjing Yin, Yi Jing, Fandong Meng, Binghui Lin, Yunbo Cao and Jie Zhou	7021
<i>Curriculum Learning for Graph Neural Networks: A Multiview Competence-based Approach</i> Nidhi Vakil and Hadi Amiri	7036
<i>When and how to paraphrase for named entity recognition?</i> Saket Sharma, Aviral Joshi, Yiyun Zhao, Namrata Mukhija, Hanoz Bhatena, Prateek Singh and Sashank Santhanam	7052
<i>UniEvent: Unified Generative Model with Multi-Dimensional Prefix for Zero-Shot Event-Relational Reasoning</i> Zhengwei Tao, Zhi Jin, Haiyan Zhao, Chengfeng Dou, yongqiang zhao, Tao Shen and Chongyang Tao	7088
<i>Are Machine Rationales (Not) Useful to Humans? Measuring and Improving Human Utility of Free-text Rationales</i> Brihi Joshi, Ziyi Liu, Sahana Ramnath, Aaron Chan, Zhewei Tong, Shaoliang Nie, Qifan Wang, Yejin Choi and Xiang Ren	7103
<i>Automatic Annotation of Direct Speech in Written French Narratives</i> Noé Durandard, Viet Anh TRAN, Gaspard Michel and Elena V. Epure	7129
<i>Automatic Creation of Named Entity Recognition Datasets by Querying Phrase Representations</i> Hyunjae Kim, jaehyo yoo, Seunghyun Yoon and Jaewoo Kang	7148
<i>Dynamic Transformers Provide a False Sense of Efficiency</i> Yiming Chen, Simin Chen, Zexin Li, Wei Yang, Cong Liu, Robby Tan and Haizhou Li	7164

<i>Empowering Cross-lingual Behavioral Testing of NLP Models with Typological Features</i> Ester Hlavnova and Sebastian Ruder	7181
<i>Local Byte Fusion for Neural Machine Translation</i> Makesh Narsimhan Sreedhar, Xiangpeng Wan, Yu Cheng and Junjie Hu	7199
<i>Where's the Point? Self-Supervised Multilingual Punctuation-Agnostic Sentence Segmentation</i> Benjamin Minixhofer, Jonas Pfeiffer and Ivan Vulić	7215
<i>Multi-target Backdoor Attacks for Code Pre-trained Models</i> Yanzhou Li, Shangqing Liu, Kangjie Chen, Xiaofei Xie, Tianwei Zhang and Yang Liu	7236
<i>Learning Better Masking for Better Language Model Pre-training</i> Dongjie Yang, Zhuosheng Zhang and Hai Zhao	7255
<i>VisText: A Benchmark for Semantically Rich Chart Captioning</i> Benny J. Tang, Angie Boggust and Arvind Satyanarayan	7268
<i>Byte-Level Grammatical Error Correction Using Synthetic and Curated Corpora</i> Svanhvít Lilja Ingólfssdóttir, Petur Orri Ragnarsson, Haukur Páll Jónsson, Haukur Barri Simonarson, Vilhjalmur Thorsteinsson and Vésteinn Snæbjarnarson	7299
<i>Multi-Level Knowledge Distillation for Out-of-Distribution Detection in Text</i> Qianhui Wu, Huiqiang Jiang, Haonan Yin, Börje F. Karlsson and Chin-Yew Lin	7317
<i>Peeking inside the black box: A Commonsense-aware Generative Framework for Explainable Complaint Detection</i> Apoorva Singh, Raghav Jain, Prince Jha and Sriparna Saha	7333
<i>MMDialog: A Large-scale Multi-turn Dialogue Dataset Towards Multi-modal Open-domain Conversation</i> Jiazhan Feng, Qingfeng Sun, Can Xu, Pu Zhao, Yaming Yang, Chongyang Tao, Dongyan Zhao and Qingwei Lin	7348
<i>ByGPT5: End-to-End Style-conditioned Poetry Generation with Token-free Language Models</i> Jonas Belouadi and Steffen Eger	7364
<i>Envisioning Future from the Past: Hierarchical Duality Learning for Multi-Turn Dialogue Generation</i> Ang Lv, Jinpeng Li, Shufang Xie and Rui Yan	7382
<i>DualGATs: Dual Graph Attention Networks for Emotion Recognition in Conversations</i> Duzhen Zhang, Feilong Chen and Xiuyi Chen	7395
<i>Consistent Prototype Learning for Few-Shot Continual Relation Extraction</i> Xiudi Chen, Hui Wu and xiaodong shi	7409
<i>Matching Pairs: Attributing Fine-Tuned Models to their Pre-Trained Large Language Models</i> Myles Foley, Amrith Rawat, Taesung Lee, Yufang Hou, Gabriele Picco and Giulio Zizzo ..	7423
<i>Large Language Models Meet NL2Code: A Survey</i> Daoguang Zan, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Wang Yongji and Jian-Guang LOU	7443
<i>When Does Aggregating Multiple Skills with Multi-Task Learning Work? A Case Study in Financial NLP</i> Jingwei Ni, Zhijing Jin, QIAN WANG, Mrinmaya Sachan and Markus Leippold	7465

<i>Enhancing Grammatical Error Correction Systems with Explanations</i>	
Yuejiao Fei, Leyang Cui, Sen Yang, Wai Lam, Zhenzhong Lan and Shuming Shi	7489
<i>Linguistic representations for fewer-shot relation extraction across domains</i>	
Sireesh Gururaja, Ritam Dutt, Tinglong Liao and Carolyn Rosé	7502
<i>DarkBERT: A Language Model for the Dark Side of the Internet</i>	
Youngjin Jin, Eugene Jang, Jian Cui, Jin-Woo Chung, Yongjae Lee and Seungwon Shin	7515
<i>MDACE: MIMIC Documents Annotated with Code Evidence</i>	
Hua Cheng, Rana Jafari, April D Russell, Russell Klopfer, Edmond Lu, Benjamin R Striner and Matthew R. Gormley	7534
<i>Towards Zero-Shot Multilingual Transfer for Code-Switched Responses</i>	
Ting-Wei Wu, Changsheng Zhao, Ernie Chang, Yangyang Shi, Pierce I-Jen Chuang, Vikas Chandra and Biing Juang	7551
<i>One Network, Many Masks: Towards More Parameter-Efficient Transfer Learning</i>	
Guangtao Zeng, Peiyuan Zhang and Wei Lu	7564
<i>Can Language Models Make Fun? A Case Study in Chinese Comical Crosstalk</i>	
jianquan li, XiangBo Wu, Xiaokang Liu, Qianqian Xie, Prayag Tiwari and Benyou Wang	7581
<i>Convergence and Diversity in the Control Hierarchy</i>	
Alexandra Cristina Butoi, Ryan Cotterell and David Chiang	7597
<i>ConFEDE: Contrastive Feature Decomposition for Multimodal Sentiment Analysis</i>	
Jiuding Yang, Yakun Yu, Di Niu, Weidong Guo and Yu Xu	7617
<i>Using Domain Knowledge to Guide Dialog Structure Induction via Neural Probabilistic Soft Logic</i>	
Connor F Pryor, Quan Yuan, Jeremiah Liu, Mehran Kazemi, Deepak Ramachandran, Tania Bedrax-Weiss and Lise Getoor	7631
<i>Are You Copying My Model? Protecting the Copyright of Large Language Models for EaaS via Backdoor Watermark</i>	
Wenjun Peng, Jingwei Yi, Fangzhao Wu, Shangxi Wu, Bin Benjamin Bin Zhu, Lingjuan Lyu, Binxing Jiao, Tong Xu, Guangzhong Sun and Xing Xie	7653
<i>Answering Ambiguous Questions via Iterative Prompting</i>	
Weiwei Sun, Hengyi Cai, Hongshen Chen, Pengjie Ren, Zhumin CHEN, Maarten de Rijke and Zhaochun Ren	7669
<i>A Dataset of Argumentative Dialogues on Scientific Papers</i>	
Federico Ruggeri, Mohsen Mesgar and Iryna Gurevych	7684
<i>Massively Multilingual Lexical Specialization of Multilingual Transformers</i>	
Tommaso Green, Simone Paolo Ponzetto and Goran Glavaš	7700
<i>RL4F: Generating Natural Language Feedback with Reinforcement Learning for Repairing Model Outputs</i>	
Afra Feyza Akyurek, Ekin Akyurek, Ashwin Kalyan, Peter Clark, Derry Tanti Wijaya and Niket Tandon	7716
<i>WebIE: Faithful and Robust Information Extraction on the Web</i>	
Chenxi Whitehouse, Clara Vania, Alham Fikri Aji, Christos Christodoulopoulos and Andrea Pierleoni	7734

<i>NormBank: A Knowledge Bank of Situational Social Norms</i>	
Caleb Ziems, Jane Dwivedi-Yu, Yi-Chia Wang, Alon Halevy and Diyi Yang	7756
<i>DIP: Dead code Insertion based Black-box Attack for Programming Language Model</i>	
CheolWon Na, YunSeok Choi and Jee-Hyong Lee	7777
<i>Modeling Structural Similarities between Documents for Coherence Assessment with Graph Convolutional Networks</i>	
Wei Liu, Xiyan Fu and Michael Strube	7792
<i>HiTIN: Hierarchy-aware Tree Isomorphism Network for Hierarchical Text Classification</i>	
He Zhu, Chong Zhang, Junjie Huang, Junran Wu and Ke Xu	7809
<i>Contextual Knowledge Learning for Dialogue Generation</i>	
Wen Zheng, Natasa Milic-Frayling and Ke Zhou	7822
<i>Easy Guided Decoding in Providing Suggestions for Interactive Machine Translation</i>	
Ke Wang, Xin Ge, Jiayi Wang, Yuqi Zhang and Yu Zhao	7840
<i>Discourse-Centric Evaluation of Document-level Machine Translation with a New Densely Annotated Parallel Corpus of Novels</i>	
Yuchen Eleanor Jiang, Tianyu Liu, Shuming Ma, Dongdong Zhang, Mrinmaya Sachan and Ryan Cotterell	7853
<i>CMOT: Cross-modal Mixup via Optimal Transport for Speech Translation</i>	
Yan Zhou, Qingkai Fang and Yang Feng	7873
<i>On the Evaluation of Neural Selective Prediction Methods for Natural Language Processing</i>	
Zhengyao Gu and Mark Hopkins	7888
<i>Speech-Text Pre-training for Spoken Dialog Understanding with Explicit Cross-Modal Alignment</i>	
Tianshu Yu, haoyu gao, Ting-En Lin, Min Yang, Yuchuan Wu, Wentao Ma, chao wang, Fei Huang and Yongbin Li	7900
<i>Text Style Transfer with Contrastive Transfer Pattern Mining</i>	
Jingxuan Han, Quan Wang, Licheng Zhang, Weidong Chen, Yan Song and Zhendong Mao ..	7914
<i>Zero- and Few-Shot Event Detection via Prompt-Based Meta Learning</i>	
Zhenrui Yue, Huimin Zeng, Mengfei Lan, Heng Ji and Dong Wang	7928
<i>Text Style Transfer Back-Translation</i>	
Daimeng Wei, Zhanglin Wu, Hengchao Shang, Zongyao Li, Minghan Wang, Jiaxin GUO, Xiaoyu Chen, Zhengzhe YU and Hao Yang	7944
<i>Generating Visual Spatial Description via Holistic 3D Scene Understanding</i>	
Yu Zhao, Hao Fei, Wei Ji, Jianguo Wei, Meishan Zhang, Min Zhang and Tat-Seng Chua ...	7960
<i>Continual Knowledge Distillation for Neural Machine Translation</i>	
Yuanchi Zhang, Peng Li, Maosong Sun and Yang Liu	7978
<i>Query Refinement Prompts for Closed-Book Long-Form QA</i>	
Reinald Kim Amplayo, Kellie Webster, Michael Collins, Dipanjan Das and Shashi Narayan	7997
<i>CONE: An Efficient COarse-to-fiNE Alignment Framework for Long Video Temporal Grounding</i>	
Zhijian Hou, Wanjun Zhong, Lei Ji, DIFEI GAO, Kun Yan, W.K. Chan, Chong-Wah Ngo, Mike Zheng Shou and Nan Duan	8013

<i>Few-Shot Document-Level Event Argument Extraction</i> Xianjun Yang, Yujie Lu and Linda R Petzold	8029
<i>ParaAMR: A Large-Scale Syntactically Diverse Paraphrase Dataset by AMR Back-Translation</i> Kuan-Hao Huang, Varun Iyer, I-Hung Hsu, Anoop Kumar, Kai-Wei Chang and Aram Galstyan	8047
<i>Towards Understanding and Improving Knowledge Distillation for Neural Machine Translation</i> Songming Zhang, Yunlong Liang, Shuaibo Wang, Yufeng Chen, Wenjuan Han, Jian Liu and Jinan Xu	8062
<i>Multi-Row, Multi-Span Distant Supervision For Table+Text Question Answering</i> Wishwajeet Kumar, Yash Gupta, Saneem Ahmed Chemmengath, Jaydeep Sen, Soumen Chakrabarti, Samarath Bharadwaj and Feifei Pan	8080
<i>HAHE: Hierarchical Attention for Hyper-Relational Knowledge Graphs in Global and Local Level</i> Haoran Luo, Haihong E, Yuhao Yang, Yikai Guo, Mingzhi Sun, Tianyu Yao, Zichen Tang, Kaiyang Wan, Meina Song and Wei Lin	8095
<i>ORGAN: Observation-Guided Radiology Report Generation via Tree Reasoning</i> Wenjun Hou, Kaishuai Xu, Yi Cheng, Wenjie Li and Jiang Liu	8108
<i>Data Curation Alone Can Stabilize In-context Learning</i> Ting-Yun Chang and Robin Jia	8123
<i>MidMed: Towards Mixed-Type Dialogues for Medical Consultation</i> Xiaoming Shi, Zeming Liu, Chuan Wang, Haitao Leng, Kui Xue, Xiaofan Zhang and Shaoting Zhang	8145
<i>FiD-ICL: A Fusion-in-Decoder Approach for Efficient In-Context Learning</i> Qinyuan Ye, Iz Beltagy, Matthew Peters, Xiang Ren and Hannaneh Hajishirzi	8158
<i>S2ynRE: Two-stage Self-training with Synthetic data for Low-resource Relation Extraction</i> Benfeng Xu, Quan Wang, Yajuan Lyu, Dai Dai, Yongdong Zhang and Zhendong Mao	8186
<i>DSEE: Dually Sparsity-embedded Efficient Tuning of Pre-trained Language Models</i> Xuxi Chen, Tianlong Chen, Weizhu Chen, Ahmed Hassan Awadallah, Zhangyang Wang and Yu Cheng	8208
<i>CASE: Aligning Coarse-to-Fine Cognition and Affection for Empathetic Response Generation</i> Jinfeng Zhou, Chujie Zheng, Bo Wang, Zheng Zhang and Minlie Huang	8223
<i>Comparative evaluation of boundary-relaxed annotation for Entity Linking performance</i> Gabriel Herman Bernardim Andrade, Shuntaro Yada and Eiji ARAMAKI	8238
<i>Do CoNLL-2003 Named Entity Taggers Still Work Well in 2023?</i> Shuheng Liu and Alan Ritter	8254
<i>READIN: A Chinese Multi-Task Benchmark with Realistic and Diverse Input Noises</i> Chenglei Si, Zhengyan Zhang, Yingfa Chen, Xiaozhi Wang, Zhiyuan Liu and Maosong Sun	8272
<i>MAD-TSC: A Multilingual Aligned News Dataset for Target-dependent Sentiment Classification</i> Evan Dufraisse, Adrian Popescu, Julien Tourille, Armelle Brun and Jerome Deshayes	8286
<i>A New Dataset and Empirical Study for Sentence Simplification in Chinese</i> Shiping Yang, Renliang Sun and Xiaojun Wan	8306

<i>Factual or Contextual? Disentangling Error Types in Entity Description Generation</i> Navita Goyal, Ani Nenkova and Hal Daumé III	8322
<i>Weakly Supervised Vision-and-Language Pre-training with Relative Representations</i> Chi Chen, Peng Li, Maosong Sun and Yang Liu	8341
<i>HermEs: Interactive Spreadsheet Formula Prediction via Hierarchical Formulet Expansion</i> Wanrong He, Haoyu Dong, Yihuai Gao, zhichao fan, Xingzhuo Guo, Zhitao Hou, Xiao Lv, Ran Jia, Shi Han and Dongmei Zhang	8356
<i>ArgU: A Controllable Factual Argument Generator</i> Sougata Saha and Rohini K Srihari	8373
<i>Learning Answer Generation using Supervision from Automatic Question Answering Evaluators</i> Matteo Gabburo, Siddhant Garg, Rik Koncel-Kedziorski and Alessandro Moschitti	8389
<i>RECAP: Retrieval-Enhanced Context-Aware Prefix Encoder for Personalized Dialogue Response Generation</i> Shuai Liu, Hyundong J Cho, Marjorie Freedman, Xuezhe Ma and Jonathan May	8404
<i>Don't Parse, Choose Spans! Continuous and Discontinuous Constituency Parsing via Autoregressive Span Selection</i> Songlin Yang and Kewei Tu	8420
<i>Laziness Is a Virtue When It Comes to Compositionality in Neural Semantic Parsing</i> Maxwell Crouse, Pavan Kapanipathi, Subhajt Chaudhury, Tahira Naseem, Ramon Fernandez Astudillo, Achille Fokoue and Tim Klinger	8434
<i>AD-KD: Attribution-Driven Knowledge Distillation for Language Model Compression</i> Siyue Wu, Hongzhan Chen, Xiaojun Quan, Qifan Wang and Rui Wang	8449
<i>(QA)²: Question Answering with Questionable Assumptions</i> Najoung Kim, Phu Mon Htut, Samuel R. Bowman and Jackson Petty	8466
<i>Attributable and Scalable Opinion Summarization</i> Tom Hosking, Hao Tang and Mirella Lapata	8488
<i>Targeted Data Generation: Finding and Fixing Model Weaknesses</i> Zexue He, Marco Tulio Ribeiro and Fereshte Khani	8506
<i>HiFi: High-Information Attention Heads Hold for Parameter-Efficient Model Adaptation</i> Anchun Gui and Han Xiao	8521
<i>CFSum Coarse-to-Fine Contribution Network for Multimodal Summarization</i> Min Xiao, Junnan Zhu, Haitao Lin, Yu Zhou and Chengqing Zong	8538
<i>On Scientific Debt in NLP: A Case for More Rigour in Language Model Pre-Training Research</i> Made Nindyatama Nityasya, Haryo Akbarianto Wibowo, Alham Fikri Aji, Genta Indra Winata, Radityo Eko Prasajo, Phil Blunsom and Adhiguna Kuncoro	8554
<i>End-to-end Knowledge Retrieval with Multi-modal Queries</i> Man Luo, Zhiyuan Fang, Tejas Gokhale, Yezhou Yang and Chitta Baral	8573
<i>AV-TranSpeech: Audio-Visual Robust Speech-to-Speech Translation</i> Rongjie Huang, Huadai Liu, Xize Cheng, Yi Ren, Linjun Li, Zhenhui Ye, Jinzheng He, Lichao Zhang, Jinglin Liu, Xiang Yin and Zhou Zhao	8590

<i>Dual Class Knowledge Propagation Network for Multi-label Few-shot Intent Detection</i> Feng Zhang, Wei Chen, Fei Ding and Tengjiao Wang	8605
<i>VendorLink: An NLP approach for Identifying & Linking Vendor Migrants & Potential Aliases on Darknet Markets</i> Vageesh Kumar Saxena, Nils Rethmeier, Gijs van Dijck and Gerasimos Spanakis	8619
<i>Element-aware Summarization with Large Language Models: Expert-aligned Evaluation and Chain-of-Thought Method</i> Yiming Wang, Zhuosheng Zhang and Rui Wang	8640
<i>Efficient Shapley Values Estimation by Amortization for Text Classification</i> Chenghao Yang, Fan Yin, He He, Kai-Wei Chang, Xiaofei Ma and Bing Xiang.....	8666
<i>PeerDA: Data Augmentation via Modeling Peer Relation for Span Identification Tasks</i> Weiwen Xu, Xin Li, Yang Deng, Wai Lam and Lidong Bing.....	8681
<i>Dynamic Regularization in UDA for Transformers in Multimodal Classification</i> Ivonne Monter-Aldana, Adrian Pastor Lopez Monroy and Fernando Sanchez-Vega.....	8700
<i>Conflicts, Villains, Resolutions: Towards models of Narrative Media Framing</i> Lea Frermann, Jiatong Li, Shima Khanehzar and Gosia Mikolajczak	8712
<i>bgGLUE: A Bulgarian General Language Understanding Evaluation Benchmark</i> Momchil Hardalov, Pepa Atanasova, Todor Mihaylov, Galia Angelova, Kiril Simov, Petya Osenova, Veselin Stoyanov, Ivan K. Koychev, Preslav Nakov and Dragomir Radev	8733
<i>DuNST: Dual Noisy Self Training for Semi-Supervised Controllable Text Generation</i> Yuxi Feng, Xiaoyuan Yi, Xiting Wang, Laks Lakshmanan, V.S. and Xing Xie	8760
<i>What does the Failure to Reason with Respectively” in Zero/Few-Shot Settings Tell Us about Language Models?</i> Ruixiang Cui, Seolhwa Lee, Daniel Hershcovich and Anders Søgaard	8786
<i>BLIND: Bias Removal With No Demographics</i> Hadas Orgad and Yonatan Belinkov	8801
<i>How do humans perceive adversarial text? A reality check on the validity and naturalness of word-based adversarial attacks</i> Salijona Dyrnishi, Salah GHAMIZI and Maxime Cordy.....	8822
<i>Soft Alignment Objectives for Robust Adaptation of Language Generation</i> Michal Štefánik, Marek Kadlcik and Petr Sojka	8837
<i>The CRINGE Loss: Learning what language not to model</i> Leonard Adolphs, Tianyu Gao, Jing Xu, Kurt Shuster, Sainbayar Sukhbaatar and Jason Weston	8854
<i>Modeling User Satisfaction Dynamics in Dialogue via Hawkes Process</i> Fanghua Ye, zhiyuan hu and Emine Yilmaz.....	8875
<i>Towards Identifying Fine-Grained Depression Symptoms from Memes</i> Shweta Yadav, Cornelia Caragea, Chenye Zhao, Naincy Kumari, Marvin A Solberg and Tanmay Sharma	8890
<i>SLUE Phase-2: A Benchmark Suite of Diverse Spoken Language Understanding Tasks</i> Suwon Shon, Siddhant Arora, Chyi-Jiunn Lin, Ankita Pasad, Felix Wu, Roshan S Sharma, Wei-Lun Wu, Hung-yi Lee, Karen Livescu and Shinji Watanabe.....	8906

<i>My side, your side and the evidence: Discovering aligned actor groups and the narratives they weave</i>	
Pavan Holur, David Chong, Timothy R Tangherlini and Vwani Roychowdhury	8938
<i>Characterizing and Measuring Linguistic Dataset Drift</i>	
Tyler A Chang, Kishalay Halder, Neha Anna John, Yogarshi Vyas, Yassine Benajiba, Miguel Ballesteros and Dan Roth	8953
<i>WebCPM: Interactive Web Search for Chinese Long-form Question Answering</i>	
Yujia Qin, Zihan Cai, Dian Jin, Lan Yan, Shihao Liang, Kunlun Zhu, Yankai Lin, Xu Han, Ning Ding, Huadong Wang, Ruobing Xie, Fanchao Qi, Zhiyuan Liu, Maosong Sun and Jie Zhou	8968
<i>Synthesize, Prompt and Transfer: Zero-shot Conversational Question Generation with Pre-trained Language Model</i>	
Hongwei Zeng, Bifan Wei, Jun Liu and Weiping Fu	8989
<i>FormNetV2: Multimodal Graph Contrastive Learning for Form Document Information Extraction</i>	
Chen-Yu Lee, Chun-Liang Li, Hao Zhang, Timothy Dozat, Vincent Perot, Guolong Su, Xiang Zhang, Kihyuk Sohn, NIKOLAY GLUSHNEV, Renshen Wang, Joshua Ainslie, Shangbang Long, Siyang Qin, Yasuhisa Fujii, Nan Hua and Tomas Pfister	9011
<i>MixCE: Training Autoregressive Language Models by Mixing Forward and Reverse Cross-Entropies</i>	
Shiyue Zhang, Shijie Wu, Ozan Irsoy, Steven Lu, Mohit Bansal, Mark Dredze and David Rosenberg	9027
<i>Knowledgeable Parameter Efficient Tuning Network for Commonsense Question Answering</i>	
Ziwan Zhao, Linmei Hu, Hanyu Zhao, Yingxia Shao and Yequan Wang	9051
<i>BLASER: A Text-Free Speech-to-Speech Translation Evaluation Metric</i>	
Mingda Chen, Paul-Ambroise Augustin Duquenne, Pierre Y Andrews, Justine T Kao, Alexandre Mourachko, Holger Schwenk and Marta R. Costa-jussà	9064
<i>NLPositionality: Characterizing Design Biases of Datasets and Models</i>	
Sebastin Santy, Jenny T Liang, Ronan Le Bras, Katharina Reinecke and Maarten Sap	9080
<i>Backpack Language Models</i>	
John Hewitt, John Thickstun, Christopher D. Manning and Percy Liang	9103
<i>WinoQueer: A Community-in-the-Loop Benchmark for Anti-LGBTQ+ Bias in Large Language Models</i>	
Virginia K. Felkner, Ho-Chun Herbert Chang, Eugene Jang and Jonathan May	9126
<i>Grounded Multimodal Named Entity Recognition on Social Media</i>	
Jianfei Yu, Ziyang Li, Jieming Wang and Rui Xia	9141
<i>Preserving Commonsense Knowledge from Pre-trained Language Models via Causal Inference</i>	
Junhao Zheng, Qianli Ma, Shengjie Qiu, Yue Wu, Peitian Ma, Junlong Liu, Huawen Feng, Xichen Shang and Haibin Chen	9155
<i>Translation-Enhanced Multilingual Text-to-Image Generation</i>	
Yaoyiran Li, Ching-Yun Chang, Stephen Rawls, Ivan Vulić and Anna Korhonen	9174
<i>Benchmarking Large Language Model Capabilities for Conditional Generation</i>	
Joshua Maynez, Priyanka Agrawal and Sebastian Gehrmann	9194
<i>lilGym: Natural Language Visual Reasoning with Reinforcement Learning</i>	
Anne Wu, Kianté Brantley, Noriyuki Kojima and Yoav Artzi	9214

<i>Unsupervised Melody-to-Lyrics Generation</i>	
Yufei Tian, Anjali Narayan-Chen, Shereen Oraby, Alessandra Cervone, Gunnar A Sigurdsson, Chenyang Tao, Wenbo Zhao, Tagyoung Chung, Jing Huang and Nanyun Peng	9235
<i>Causality-aware Concept Extraction based on Knowledge-guided Prompting</i>	
Siyu Yuan, Deqing Yang, Jinxi Liu, Shuyu Tian, Jiaqing Liang, Yanghua Xiao and Rui Xie .	9255
<i>Span-level Aspect-based Sentiment Analysis via Table Filling</i>	
Mao Zhang, Yongxin Zhu, Zhen Liu, Zhimin Bao, Yunfei Wu, Xing Sun and Linli Xu	9273
<i>Limitations of Language Models in Arithmetic and Symbolic Induction</i>	
Jing Qian, Hong Wang, Zekun Li, Shiyang Li and Xifeng Yan	9285
<i>EEL: Efficiently Encoding Lattices for Reranking</i>	
Prasann Singhal, Jiacheng Xu, Xi Ye and Greg Durrett	9299
<i>CLAPSpeech: Learning Prosody from Text Context with Contrastive Language-Audio Pre-Training</i>	
Zhenhui Ye, Rongjie Huang, Yi Ren, Ziyue Jiang, Jinglin Liu, Jinzheng He, Xiang Yin and Zhou Zhao	9317
<i>Revisiting Cross-Lingual Summarization: A Corpus-based Study and A New Benchmark with Improved Annotation</i>	
Yulong Chen, Huajian Zhang, Yijie Zhou, Xuefeng Bai, Yueguan Wang, Ming Zhong, Jianhao Yan, Yafu Li, Judy Li, Xianchao Zhu and Yue Zhang	9332
<i>Learning Dynamic Contextualised Word Embeddings via Template-based Temporal Adaptation</i>	
Xiaohang Tang, Yi Zhou and Danushka Bollegala	9352
<i>How poor is the stimulus? Evaluating hierarchical generalization in neural networks trained on child-directed speech</i>	
Aditya Yedetore, Tal Linzen, Robert Frank and R. Thomas McCoy	9370
<i>GanLM: Encoder-Decoder Pre-training with an Auxiliary Discriminator</i>	
Jian Yang, Shuming Ma, Li Dong, Shaohan Huang, Haoyang Huang, Yuwei Yin, Dongdong Zhang, Liqun Yang, Furu Wei and Zhoujun Li	9394
<i>Linear Guardedness and its Implications</i>	
Shauli Ravfogel, Yoav Goldberg and Ryan Cotterell	9413
<i>Searching for Needles in a Haystack: On the Role of Incidental Bilingualism in PaLM’s Translation Capability</i>	
Eleftheria Briakou, Colin Cherry and George Foster	9432
<i>Open Set Relation Extraction via Unknown-Aware Training</i>	
Jun Zhao, Xin Zhao, WenYu Zhan, Qi Zhang, Tao Gui, Zhongyu Wei, Yun Wen Chen, Xiang Gao and Xuanjing Huang	9453
<i>Learning to Imagine: Visually-Augmented Natural Language Generation</i>	
Tianyi Tang, Yushuo Chen, Yifan Du, Junyi Li, Wayne Xin Zhao and Ji-Rong Wen	9468
<i>Generating Hashtags for Short-form Videos with Guided Signals</i>	
Tiezheng Yu, Hanchao Yu, Davis Liang, Yuning Mao, Shaoliang Nie, Po-Yao Huang, Madian Khabsa, Pascale Fung and Yi-Chia Wang	9482
<i>NEUROSTRUCTURAL DECODING: Neural Text Generation with Structural Constraints</i>	
Mohaddeseh Bastan, Mihai Surdeanu and Niranjana Balasubramanian	9496

<i>The Best of Both Worlds: Combining Human and Machine Translations for Multilingual Semantic Parsing with Active Learning</i>	
Zhuang Li, Lizhen Qu, Philip Cohen, Raj V Tumuluri and Gholamreza Haffari	9511
<i>Ideology Prediction from Scarce and Biased Supervision: Learn to Disregard the What” and Focus on the How”!</i>	
Chen Chen, Dylan Walker and Venkatesh Saligrama	9529
<i>Unsupervised Extractive Summarization of Emotion Triggers</i>	
Tiberiu Sosea, Hongli Zhan, Junyi Jessy Li and Cornelia Caragea	9550
<i>Document-Level Event Argument Extraction With a Chain Reasoning Paradigm</i>	
Jian Liu, Chen Liang, Jinan Xu, Haoyan Liu and Zhe Zhao	9570
<i>Pre-training Multi-party Dialogue Models with Latent Discourse Inference</i>	
Yiyang Li, Xinting Huang, Wei Bi and Hai Zhao	9584
<i>Interpreting Positional Information in Perspective of Word Order</i>	
Zhang Xilong, Liu Ruochen, Liu Jin and Liang Xuefeng	9600
<i>I2D2: Inductive Knowledge Distillation with NeuroLogic and Self-Imitation</i>	
Chandra Bhagavatula, Jena D. Hwang, Doug Downey, Ronan Le Bras, Ximing Lu, Lianhui Qin, Keisuke Sakaguchi, Swabha Swayamdipta, Peter West and Yejin Choi	9614
<i>More than Classification: A Unified Framework for Event Temporal Relation Extraction</i>	
Quzhe Huang, Yutong Hu, Shengqi Zhu, Yansong Feng, Chang Liu and Dongyan Zhao	9631
<i>Multi-Source Test-Time Adaptation as Dueling Bandits for Extractive Question Answering</i>	
Hai Ye, Qizhe Xie and Hwee Tou Ng	9647
<i>Decoupling Pseudo Label Disambiguation and Representation Learning for Generalized Intent Discovery</i>	
Yutao Mou, Xiaoshuai Song, Keqing He, Chen Zeng, Pei Wang, Jingang Wang, Yunsen Xian and Weiran Xu	9661
<i>DecompEval: Evaluating Generated Texts as Unsupervised Decomposed Question Answering</i>	
Pei Ke, Fei Huang, Fei Mi, Yasheng Wang, Qun Liu, Xiaoyan Zhu and Minlie Huang	9676
<i>Backdoor Neural Code Search</i>	
Weisong Sun, Yuchen Chen, Guan hong Tao, Chunrong Fang, Xiangyu Zhang, Qianjun Zhang and Bin Luo	9692
<i>Concise Answers to Complex Questions: Summarization of Long-form Answers</i>	
Abhilash C Potluri, Fangyuan Xu and Eunsol Choi	9709
<i>Towards Better Entity Linking with Multi-View Enhanced Distillation</i>	
Yi Liu, Yuan Tian, Jianxun Lian, xinlong wang, Yanan Cao, Fang Fang, Wen Zhang, Haizhen Huang, Weiwei Deng and Qi Zhang	9729
<i>A Measure-Theoretic Characterization of Tight Language Models</i>	
Li Du, Lucas Torroba Hennigen, Tiago Pimentel, Clara Meister, Jason Eisner and Ryan Cotterell	9744
<i>PAED: Zero-Shot Persona Attribute Extraction in Dialogues</i>	
Luyao Zhu, Wei Li, Rui Mao, Vlad Pandealea and Erik Cambria	9771
<i>PromptRank: Unsupervised Keyphrase Extraction Using Prompt</i>	
Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun and Xiaoyan Bai	9788

<i>When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories</i>	
Alex Troy Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi and Hannaneh Hajishirzi	9802
<i>infoVerse: A Universal Framework for Dataset Characterization with Multidimensional Meta-information</i>	
Jaehyung Kim, Yekyung Kim, Karin Johanna Denton de Langis, Jinwoo Shin and Dongyeop Kang	9823
<i>SeeGULL: A Stereotype Benchmark with Broad Geo-Cultural Coverage Leveraging Generative Models</i>	
Akshita Jha, Aida Mostafazadeh Davani, Chandan K Reddy, Shachi Dave, Vinodkumar Prabhakaran and Sunipa Dev	9851
<i>Automated Metrics for Medical Multi-Document Summarization Disagree with Human Evaluations</i>	
Lucy Lu Wang, Yulia Otmakhova, Jay DeYoung, Thinh Hung Truong, Bailey E. Kuehl, Erin A Bransom and Byron C. Wallace.....	9871
<i>Say What You Mean! Large Language Models Speak Too Positively about Negative Commonsense Knowledge</i>	
Jiangjie Chen, Wei Shi, Ziquan Fu, Sijie Cheng, Lei Li and Yanghua Xiao	9890
<i>An Inner Table Retriever for Robust Table Question Answering</i>	
Weizhe Lin, Rexhina Blloshmi, Bill Byrne, Adria de Gispert and Gonzalo Iglesias	9909
<i>SIMSUM: Document-level Text Simplification via Simultaneous Summarization</i>	
Sofia Blinova, Xinyu Zhou, Martin Jaggi, Carsten Eickhoff and Seyed Ali Bahrainian	9927
<i>SimOAP: Improve Coherence and Consistency in Persona-based Dialogue Generation via Over-sampling and Post-evaluation</i>	
Junkai Zhou, Liang Pang, Huawei Shen and Xueqi Cheng	9945
<i>NatLogAttack: A Framework for Attacking Natural Language Inference Models with Natural Logic</i>	
Zi'ou Zheng and Xiaodan Zhu	9960
<i>Cognitive Reframing of Negative Thoughts through Human-Language Model Interaction</i>	
Ashish Sharma, Kevin Rushton, Inna Wanyin Lin, David Wadden, Khendra G Lucas, Adam Miner, Theresa Nguyen and Tim Althoff	9977
<i>Dating Greek Papyri with Text Regression</i>	
John Pavlopoulos, Maria Konstantinidou, Isabelle Marthot-Santaniello, Holger Essler and Asimina Paparigopoulou	10001
<i>Interleaving Retrieval with Chain-of-Thought Reasoning for Knowledge-Intensive Multi-Step Questions</i>	
Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot and Ashish Sabharwal	10014
<i>Direct Fact Retrieval from Knowledge Graphs without Entity Linking</i>	
Jinheon Baek, Alham Fikri Aji, Jens Lehmann and Sung Ju Hwang	10038
<i>DisentQA: Disentangling Parametric and Contextual Knowledge with Counterfactual Question Answering</i>	
Ella Neeman, Roei Aharoni, Or Honovich, Leshem Choshen, Idan Szpektor and Omri Abend	10056
<i>A New Direction in Stance Detection: Target-Stance Extraction in the Wild</i>	
Yingjie Li, Krishna K Garg and Cornelia Caragea	10071

<i>Improved Instruction Ordering in Recipe-Grounded Conversation</i>	
Duong Minh Le, Ruohao Guo, Wei Xu and Alan Ritter	10086
<i>Token-wise Decomposition of Autoregressive Language Model Hidden States for Analyzing Model Predictions</i>	
Byung-Doh Oh and William Schuler	10105
<i>Document-Level Multi-Event Extraction with Event Proxy Nodes and Hausdorff Distance Minimization</i>	
Xinyu Wang, Lin Gui and Yulan He	10118
<i>Dialog-Post: Multi-Level Self-Supervised Objectives and Hierarchical Model for Dialogue Post-Training</i>	
Zhenyu Zhang, Lei Shen, Yuming Zhao, Meng Chen and Xiaodong He	10134
<i>Language Detoxification with Attribute-Discriminative Latent Space</i>	
Jin Myung Kwak, Minseon Kim and Sung Ju Hwang	10149
<i>Just Like a Human Would, Direct Access to Sarcasm Augmented with Potential Result and Reaction</i>	
Changrong Min, Ximing Li, Liang Yang, Zhilin Wang, Bo Xu and Hongfei LIN	10172
<i>Adaptive and Personalized Exercise Generation for Online Language Learning</i>	
Peng Cui and Mrinmaya Sachan	10184
<i>NLP Reproducibility For All: Understanding Experiences of Beginners</i>	
Shane Storks, Keunwoo Peter Yu, Ziqiao Ma and Joyce Chai	10199
<i>Why Did the Chicken Cross the Road? Rephrasing and Analyzing Ambiguous Questions in VQA</i>	
Elias Stengel-Eskin, Jimena Guallar-Blasco, Yi Zhou and Benjamin Van Durme	10220
<i>UMRSpell: Unifying the Detection and Correction Parts of Pre-trained Models towards Chinese Missing, Redundant, and Spelling Correction</i>	
Zheyu He, Yujin Zhu, Linlin Wang and Liang Xu	10238
<i>LAIT: Efficient Multi-Segment Encoding in Transformers with Layer-Adjustable Interaction</i>	
Jeremiah Milbauer, Annie Louis, Mohammad Javad Hosseini, Alex Fabrikant, Donald Metzler and Tal Schuster	10251
<i>Local Interpretation of Transformer Based on Linear Decomposition</i>	
Sen Yang, Shujian Huang, wei zou, Jianbing Zhang, Xinyu Dai and Jiajun CHEN	10270
<i>DataFinder: Scientific Dataset Recommendation from Natural Language Descriptions</i>	
Vijay Viswanathan, Luyu Gao, Tongshuang Wu, Pengfei Liu and Graham Neubig	10288
<i>Multilingual Event Extraction from Historical Newspaper Adverts</i>	
Nadav Borenstein, Natália da Silva Perez and Isabelle Augenstein	10304
<i>BIC: Twitter Bot Detection with Text-Graph Interaction and Semantic Consistency</i>	
Zhenyu Lei, Herun Wan, Wenqian Zhang, Shangbin Feng, Zilong Chen, Jundong Li, Qinghua Zheng and Minnan Luo	10326
<i>Do I have the Knowledge to Answer? Investigating Answerability of Knowledge Base Questions</i>	
Mayur Patidar, Prayushi Faldu, Avinash Kumar Singh, Lovekesh Vig, Indrajit Bhattacharya and Mausam -	10341
<i>Understanding Client Reactions in Online Mental Health Counseling</i>	
Anqi Li, Lizhi Ma, Yaling Mei, Hongliang He, Shuai Zhang, Huachuan Qiu and Zhenzhong Lan	10358

<i>Nonlinear Structural Equation Model Guided Gaussian Mixture Hierarchical Topic Modeling</i>	
HeGang Chen, Pengbo Mao, Yuyin Lu and Yanghui Rao	10377
<i>Revisiting Token Dropping Strategy in Efficient BERT Pretraining</i>	
Qihuang Zhong, Liang Ding, Juhua Liu, Xuebo Liu, Min Zhang, Bo Du and Dacheng Tao	10391
<i>The Benefits of Bad Advice: Autocontrastive Decoding across Model Layers</i>	
Ariel Gera, Roni Friedman, Ofir Arviv, Chulaka Gunasekara, Benjamin Sznajder, Noam Slonim and Eyal Shnarch	10406
<i>FACTIFY-5WQA: 5W Aspect-based Fact Verification through Question Answering</i>	
Anku Rani, S.M Towhidul Islam Tonmoy, Dwip D Dalal, Shreya Gautam, Megha Chakraborty, Aman Chadha, Amit Sheth and Amitava Das	10421
<i>Naamapadam: A Large-Scale Named Entity Annotated Data for Indic Languages</i>	
Arnav Anil Mhaske, Harshit Kedia, Sumanth Doddapaneni, Mitesh M. Khapra, Pratyush Kumar, Rudra Murthy and Anoop Kunchukuttan	10441
<i>CREPE: Open-Domain Question Answering with False Presuppositions</i>	
Xinyan Velocity Yu, Sewon Min, Luke Zettlemoyer and Hannaneh Hajishirzi	10457
<i>Joint Document-Level Event Extraction via Token-Token Bidirectional Event Completed Graph</i>	
Qizhi Wan, Changxuan Wan, Keli Xiao, Dexi Liu, Chenliang Li, Bolong Zheng, Xiping Liu and Rong Hu	10481
<i>Robust Representation Learning with Reliable Pseudo-labels Generation via Self-Adaptive Optimal Transport for Short Text Clustering</i>	
Xiaolin Zheng, Mengling Hu, Weiming Liu, Chaochao Chen and Xinting Liao	10493
<i>Multilingual Knowledge Graph Completion with Language-Sensitive Multi-Graph Attention</i>	
Rongchuan Tang, Yang Zhao, Chengqing Zong and Yu Zhou	10508
<i>What are the Desired Characteristics of Calibration Sets? Identifying Correlates on Long Form Scientific Summarization</i>	
Griffin Adams, Bichlien H Nguyen, Jake Allen Smith, Yingce Xia, Shufang Xie, Anna Ostropolets, Budhaditya Deb, Yuan-Jyue Chen, Tristan Naumann and Noémie Elhadad	10520
<i>Annotating Mentions Alone Enables Efficient Domain Adaptation for Coreference Resolution</i>	
Nupoor Gandhi, Anjalie Field and Emma Strubell	10543
<i>A Universal Discriminator for Zero-Shot Generalization</i>	
Haike Xu, Zongyu Lin, Jing Zhou, Yanan Zheng and Zhilin Yang	10559
<i>Syntax and Geometry of Information</i>	
Raphaël Bailly, Laurent Leblond and Kata Gábor	10576
<i>GreenKGC: A Lightweight Knowledge Graph Completion Method</i>	
Yun Cheng Wang, Xiou Ge, Bin Wang and C.-C. Jay Kuo	10596
<i>Unsupervised Open-domain Keyphrase Generation</i>	
Lam Thanh Do, Pritom Saha Akash and Kevin Chen-Chuan Chang	10614
<i>A Cognitive Stimulation Dialogue System with Multi-source Knowledge Fusion for Elders with Cognitive Impairment</i>	
Jiyue Jiang, Sheng Wang, Qintong Li, Lingpeng Kong and Chuan Wu	10628

<i>Plug-and-Play Knowledge Injection for Pre-trained Language Models</i> Zhengyan Zhang, Zhiyuan Zeng, Yankai Lin, Huadong Wang, Deming Ye, Chaojun Xiao, Xu Han, Zhiyuan Liu, Peng Li, Maosong Sun and Jie Zhou	10641
<i>Two Birds One Stone: Dynamic Ensemble for OOD Intent Classification</i> Yunhua Zhou, Jianqiang Yang, Pengyu Wang and Xipeng Qiu	10659
<i>SWiPE: A Dataset for Document-Level Simplification of Wikipedia Pages</i> Philippe Laban, Jesse Vig, Wojciech Kryscinski, Shafiq Joty, Caiming Xiong and Chien-Sheng Jason Wu	10674
<i>Are Message Passing Neural Networks Really Helpful for Knowledge Graph Completion?</i> Juanhui Li, Harry Aaron Shomer, Jiayuan Ding, Yiqi Wang, Yao Ma, Neil Shah, Jiliang Tang and Dawei Yin	10696
<i>A dynamic programming algorithm for span-based nested named-entity recognition in $O(n^2)$</i> Caio Corro	10712
<i>Target-Side Augmentation for Document-Level Machine Translation</i> Guangsheng Bao, ZHIYANG TENG and Yue Zhang	10725
<i>Rethinking Masked Language Modeling for Chinese Spelling Correction</i> Hongqiu Wu, Shaohua Zhang, Yuchen Zhang and Hai Zhao	10743
<i>A Multi-Modal Context Reasoning Approach for Conditional Inference on Joint Textual and Visual Clues</i> Yunxin Li, Baotian Hu, Chen Xinyu, Yuxin Ding, Lin Ma and Min Zhang	10757
<i>Simple and Effective Unsupervised Speech Translation</i> Changhan Wang, Hirofumi Inaguma, Peng-Jen Chen, Iliia Kulikov, Yun Tang, Wei-Ning Hsu, Michael Auli and Juan Pino	10771
<i>Modeling What-to-ask and How-to-ask for Answer-unaware Conversational Question Generation</i> Xuan Long Do, BOWEI ZOU, Shafiq Joty, Tran Anh Tai, Liangming Pan, Nancy Chen and Ai Ti Aw	10785
<i>CHEER: Centrality-aware High-order Event Reasoning Network for Document-level Event Causality Identification</i> Meiqi Chen, Yixin Cao, Yan Zhang and Zhiwei Liu	10804
<i>f-Divergence Minimization for Sequence-Level Knowledge Distillation</i> Yuqiao Wen, Zichao Li, Wenyu Du and Lili Mou	10817
<i>Supervised Adversarial Contrastive Learning for Emotion Recognition in Conversations</i> Dou Hu, Yinan Bao, Lingwei Wei, Wei Zhou and Songlin Hu	10835
<i>A Novel Table-to-Graph Generation Approach for Document-Level Joint Entity and Relation Extraction</i> Ruoyu Zhang, Yanzeng Li and Lei Zou	10853
<i>A Synthetic Data Generation Framework for Grounded Dialogues</i> Jianzhu Bao, Rui Wang, Yasheng Wang, Aixin Sun, Yitong Li, Fei Mi and Ruifeng Xu	10866
<i>MasakhaPOS: Part-of-Speech Tagging for Typologically Diverse African languages</i> Cheikh M. Bamba Dione, David Ifeoluwa Adelani, Peter Nabende, Jesujoba Alabi, Thapelo Andrew Sindane, Happy Buzaaba, Shamsuddeen Hassan Muhammad, Chris Chinenye Emezue, Perez Ogayo, Anuoluwapo Aremu, Catherine Gitau, Derguene Mbaye, Jonathan Mukiibi, Blessing K Sibanda, Bonaventure F. P. Dossou, Andiswa Bukula, Rooweither Mabuya, Allahsera Auguste Tapo, Edwin	

Munkoh-Buabeng, victoire Memdjokam Koagne, Fatoumata Ouoba Kabore, Amelia Taylor, Godson K KALIPE, Tebogo Macucwa, Vukosi Marivate, Tajuddeen Gwadabe, Mboning Tchiazee Elvis, Ikechukwu Ekene Onyenwe, Gratien G. Atindogbe, Tolulope Anu Adelani, Idris Akinade, Olanrewaju Samuel, Marien NAHIMANA, Théogène MUSABEYEZU, Emile Niyomutabazi, Ester Chimhenga, Kudzai Gotosa, Patrick Mizha, Apelete AGBOLO, SEYDOU TRAORE, Chinedu Uchechukwu, Aliyu Yakubu Yusuf, Muhammad Sulaiman Abdullahi and Dietrich Klakow	10883
<i>Semantic Structure Enhanced Event Causality Identification</i>	
Zhilei Hu, Zixuan Li, Xiaolong Jin, Long Bai, Saiping Guan, Jiafeng Guo and Xueqi Cheng	10901
<i>Weakly-Supervised Spoken Video Grounding via Semantic Interaction Learning</i>	
Ye Wang, Wang Lin, Shengyu Zhang, Tao Jin, Linjun Li, Xize Cheng and Zhou Zhao	10914
<i>Rehearsal-free Continual Language Learning via Efficient Parameter Isolation</i>	
Zhicheng Wang, Yufang Liu, Tao Ji, Xiaoling Wang, Yuanbin Wu, Congcong Jiang, Ye Chao, Zhencong Han, Ling Wang, Xu Shao and Wenqiu Zeng	10933
<i>Label-Aware Hyperbolic Embeddings for Fine-grained Emotion Classification</i>	
Chih Yao Chen, Tun Min Hung, Yi-Li Hsu and Lun-Wei Ku	10947
<i>Combo of Thinking and Observing for Outside-Knowledge VQA</i>	
Qingyi Si, Yuchen Mo, Zheng Lin, HUIZHAN JI and Weiping Wang	10959
<i>AMPERE: AMR-Aware Prefix for Generation-Based Event Argument Extraction Model</i>	
I-Hung Hsu, Zhiyu Xie, Kuan-Hao Huang, Prem Natarajan and Nanyun Peng	10976
<i>Your spouse needs professional help: Determining the Contextual Appropriateness of Messages through Modeling Social Relationships</i>	
David Jurgens, Agrima Seth, Jackson Sargent, Athena Aghighi and Michael Geraci	10994
<i>TART: Improved Few-shot Text Classification Using Task-Adaptive Reference Transformation</i>	
Shuo Lei, Xuchao Zhang, Jianfeng He, Fanglan Chen and Chang-Tien Lu	11014
<i>How Do In-Context Examples Affect Compositional Generalization?</i>	
Shengnan An, Zeqi Lin, Qiang Fu, Bei Chen, Nanning Zheng, Jian-Guang LOU and Dongmei Zhang	11027
<i>Attractive Storyteller: Stylized Visual Storytelling with Unpaired Text</i>	
Dingyi Yang and Qin Jin	11053
<i>Multitask Pretraining with Structured Knowledge for Text-to-SQL Generation</i>	
Robert Giaquinto, Dejiao Zhang, Benjamin Kleiner, Yang Li, Ming Tan, Parminder Bhatia, Ramesh Nallapati and Xiaofei Ma	11067
<i>WSPAlign: Word Alignment Pre-training via Large-Scale Weakly Supervised Span Prediction</i>	
Qiyu Wu, Masaaki Nagata and Yoshimasa Tsuruoka	11084
<i>Distill or Annotate? Cost-Efficient Fine-Tuning of Compact Models</i>	
Junmo Kang, Wei Xu and Alan Ritter	11100
<i>OD-RTE: A One-Stage Object Detection Framework for Relational Triple Extraction</i>	
Jinzhong Ning, Zhihao Yang, Yuanyuan Sun, Zhizheng Wang and Hongfei LIN	11120
<i>I Cast Detect Thoughts: Learning to Converse and Guide with Intents and Theory-of-Mind in Dungeons and Dragons</i>	
Pei Zhou, Andrew Zhu, Jennifer Hu, Jay Pujara, Xiang Ren, Chris Callison-Burch, Yejin Choi and Prithviraj Ammanabrolu	11136

<i>Multitask Pre-training of Modular Prompt for Chinese Few-Shot Learning</i>	
Tianxiang Sun, Zhengfu He, Qin Zhu, Xipeng Qiu and Xuanjing Huang	11156
<i>Is GPT-3 a Good Data Annotator?</i>	
BOSHENG DING, Chengwei Qin, Linlin Liu, Yew Ken Chia, Boyang Li, Shafiq Joty and Lidong Bing	11173
<i>Multi-Grained Knowledge Retrieval for End-to-End Task-Oriented Dialog</i>	
Fanqi Wan, Weizhou Shen, Ke Yang, Xiaojun Quan and Wei Bi	11196
<i>Few-shot Event Detection: An Empirical Study and a Unified View</i>	
Yubo Ma, Zehao Wang, Yixin Cao and Aixin Sun	11211
<i>How to Plant Trees in Language Models: Data and Architectural Effects on the Emergence of Syntactic Inductive Biases</i>	
Aaron Mueller and Tal Linzen	11237
<i>ClarifyDelphi: Reinforced Clarification Questions with Defeasibility Rewards for Social and Moral Situations</i>	
Valentina Pyatkin, Jena D. Hwang, Vivek Srikumar, Ximing Lu, Liwei Jiang, Yejin Choi and Chandra Bhagavatula	11253
<i>HINT: Hypernetwork Instruction Tuning for Efficient Zero- and Few-Shot Generalisation</i>	
Hamish Ivison, Akshita Bhagia, Yizhong Wang, Hannaneh Hajishirzi and Matthew Peters	11272
<i>Measuring Inductive Biases of In-Context Learning with Underspecified Demonstrations</i>	
Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen and He He	11289
<i>An Inclusive Notion of Text</i>	
Iliia Kuznetsov and Iryna Gurevych	11311
<i>AlignScore: Evaluating Factual Consistency with A Unified Alignment Function</i>	
Yuheng Zha, Yichi Yang, Ruichen Li and Zhiting Hu	11328
<i>Multi-source Semantic Graph-based Multimodal Sarcasm Explanation Generation</i>	
Liqiang Jing, Xuemeng Song, Kun Ouyang, Mengzhao Jia and Liqiang Nie	11349
<i>Counterfactual Active Learning for Out-of-Distribution Generalization</i>	
Xun Deng, Wenjie Wang, Fuli Feng, Hanwang Zhang, Xiangnan He and Yong Liao	11362
<i>Multi-granularity Temporal Question Answering over Knowledge Graphs</i>	
Ziyang Chen, Jinzhi Liao and Xiang Zhao	11378
<i>A New Aligned Simple German Corpus</i>	
Vanessa Toborek, Moritz Busch, Malte Boßert, Christian Bauckhage and Pascal Welke	11393
<i>Introducing Semantics into Speech Encoders</i>	
Derek Q Xu, Shuyan Annie Dong, Changhan Wang, Suyoun Kim, Zhaojiang Lin, Bing Liu, Akshat Shrivastava, Shang-Wen Li, Liang-Hsuan Tseng, Guan-Ting Lin, Alexei Baeovski, Hung-yi Lee, Yizhou Sun and Wei Wang	11413
<i>Constrained Tuple Extraction with Interaction-Aware Network</i>	
Xiaojun Xue, Chunxia Zhang, Tianxiang Xu and Zhendong Niu	11430
<i>MultiInstruct: Improving Multi-Modal Zero-Shot Learning via Instruction Tuning</i>	
Zhiyang Xu, Ying Shen and Lifu Huang	11445

<i>Single Sequence Prediction over Reasoning Graphs for Multi-hop QA</i>	
Gowtham Ramesh, Makesh Narsimhan Sreedhar and Junjie Hu	11466
<i>Contrastive Error Attribution for Finetuned Language Models</i>	
Faisal Ladhak, Esin Durmus and Tatsunori Hashimoto	11482
<i>DARE: Towards Robust Text Explanations in Biomedical and Healthcare Applications</i>	
Adam Daniel Ivankay, Mattia Rigotti and Pascal Frossard	11499
<i>Neural Machine Translation for Mathematical Formulae</i>	
Felix Petersen, Moritz Schubotz, Andre Greiner-Petter and Bela Gipp	11534
<i>Query-Efficient Black-Box Red Teaming via Bayesian Optimization</i>	
Deokjae Lee, JunYeong Lee, Jung-Woo Ha, Jin-Hwa Kim, Sang-Woo Lee, Hwaran Lee and Hyun Oh Song	11551
<i>SSD-LM: Semi-autoregressive Simplex-based Diffusion Language Model for Text Generation and Modular Control</i>	
Xiaochuang Han, Sachin Kumar and Yulia Tsvetkov	11575
<i>Recall, Expand, and Multi-Candidate Cross-Encode: Fast and Accurate Ultra-Fine Entity Typing</i>	
Chengyue Jiang, Wenyang Hui, Yong Jiang, Xiaobin Wang, Pengjun Xie and Kewei Tu ...	11597
<i>MIR-GAN: Refining Frame-Level Modality-Invariant Representations with Adversarial Network for Audio-Visual Speech Recognition</i>	
Yuchen Hu, Chen Chen, Ruizhe Li, Heqing Zou and Eng Siong Chng	11610
<i>Understanding Factual Errors in Summarization: Errors, Summarizers, Datasets, Error Detectors</i>	
Liyan Tang, Tanya Goyal, Alex Fabbri, Philippe Laban, Jiacheng Xu, Semih Yavuz, Wojciech Kryscinski, Justin F Rousseau and Greg Durrett	11626
<i>GIFT: Graph-Induced Fine-Tuning for Multi-Party Conversation Understanding</i>	
Jia-Chen Gu, Zhenhua Ling, Quan Liu, Cong Liu and Guoping Hu	11645
<i>Hybrid Uncertainty Quantification for Selective Text Classification in Ambiguous Tasks</i>	
Artem Vazhentsev, Gleb Kuzmin, Akim Tsvigun, Alexander Panchenko, Maxim Panov, Mikhail Burtsev and Artem Shelmanov	11659
<i>BLOOM+1: Adding Language Support to BLOOM for Zero-Shot Prompting</i>	
Zheng Xin Yong, Hailey Schoelkopf, Niklas Muennighoff, Alham Fikri Aji, David Ifeoluwa Adelani, KHALID ALMUBARAK, M Saiful Bari, Lintang Sutawika, Jungo Kasai, Ahmed Baruwa, Genta Indra Winata, Stella Biderman, Edward Raff, Dragomir Radev and Vassilina Nikoulina	11682
<i>Logic-driven Indirect Supervision: An Application to Crisis Counseling</i>	
Mattia Medina Grespan, Meghan Broadbent, Xinyao Zhang, Katherine E Axford, Brent Kious, Zac Imel and Vivek Srikumar	11704
<i>Grounding Characters and Places in Narrative Text</i>	
Sandeep Soni, Amanpreet Sihra, Elizabeth F. Evans, Matthew Wilkens and David Bamman	11723
<i>From Pretraining Data to Language Models to Downstream Tasks: Tracking the Trails of Political Biases Leading to Unfair NLP Models</i>	
Shangbin Feng, Chan Young Park, Yuhan Liu and Yulia Tsvetkov	11737
<i>SLABERT Talk Pretty One Day: Modeling Second Language Acquisition with BERT</i>	
Aditya Yadavalli, Alekhya Yadavalli and Vera Tobin	11763

<i>Contrastive Novelty-Augmented Learning: Anticipating Outliers with Large Language Models</i> Albert Xu, Xiang Ren and Robin Jia	11778
<i>Learning to Initialize: Can Meta Learning Improve Cross-task Generalization in Prompt Tuning?</i> Chengwei Qin, Shafiq Joty, Qian Li and Ruochen Zhao	11802
<i>Rethinking the Role of Scale for In-Context Learning: An Interpretability-based Case Study at 66 Billion Scale</i> Hritik Bansal, Karthik Gopalakrishnan, Saket Dingliwal, Sravan Bodapati, Katrin Kirchhoff and Dan Roth	11833
<i>Question-Answering in a Low-resourced Language: Benchmark Dataset and Models for Tigrinya</i> Fitsum Gaim, Wonsuk Yang, Hancheol Park and Jong Park	11857
<i>ESCOXLM-R: Multilingual Taxonomy-driven Pre-training for the Job Market Domain</i> Mike Zhang, Rob van der Goot and Barbara Plank	11871
<i>CITADEL: Conditional Token Interaction via Dynamic Lexical Routing for Efficient and Effective Multi-Vector Retrieval</i> Minghan Li, Sheng-Chieh Lin, Barlas Oguz, Asish Ghoshal, Jimmy Lin, Yashar Mehdad, Wen-tau Yih and Xilun Chen	11891
<i>MultiCapCLIP: Auto-Encoding Prompts for Zero-Shot Multilingual Visual Captioning</i> Bang Yang, Fenglin Liu, Xian Wu, Yaowei Wang, Xu Sun and Yuexian Zou	11908
<i>Transfer and Active Learning for Dissonance Detection: Addressing the Rare-Class Challenge</i> Vasudha Varadarajan, Swanie Juhng, Syeda Mahwish, Xiaoran Liu, Jonah Keith Luby, Christian C. Luhmann and H. Andrew Schwartz	11923
<i>In-sample Curriculum Learning by Sequence Completion for Natural Language Generation</i> Qi Jia, Yizhu Liu, Haifeng Tang and Kenny Q. Zhu	11937
<i>Product Question Answering in E-Commerce: A Survey</i> Yang Deng, Wenxuan Zhang, Qian Yu and Wai Lam	11951
<i>Towards Domain-Agnostic and Domain-Adaptive Dementia Detection from Spoken Language</i> Shahla Farzana and Natalie Parde	11965
<i>Generalizing Backpropagation for Gradient-Based Interpretability</i> Kevin Du, Lucas Torroba Hennigen, Niklas Stoehr, Alex Warstadt and Ryan Cotterell	11979
<i>UPPAM: A Unified Pre-training Architecture for Political Actor Modeling based on Language</i> Xinyi Mou, Zhongyu Wei, Qi Zhang and Xuanjing Huang	11996
<i>Generic Temporal Reasoning with Differential Analysis and Explanation</i> Yu Feng, Ben Zhou, Haoyu Wang, Helen Jin and Dan Roth	12013
<i>Model-Based Simulation for Optimising Smart Reply</i> Benjamin Towle and Ke Zhou	12030
<i>Beyond Contrastive Learning: A Variational Generative Model for Multilingual Retrieval</i> John Wieting, Jonathan Clark, William Cohen, Graham Neubig and Taylor Berg-Kirkpatrick	12044
<i>On the Blind Spots of Model-Based Evaluation Metrics for Text Generation</i> Tianxing He, Jingyu Zhang, Tianle Wang, Sachin Kumar, Kyunghyun Cho, James Glass and Yulia Tsvetkov	12067

<i>Dealing with Semantic Underspecification in Multimodal NLP</i>	
Sandro Pezzelle	12098
<i>Trigger Warning Assignment as a Multi-Label Document Classification Problem</i>	
Matti Wiegmann, Magdalena Wolska, Christopher Schröder, Ole Borchardt, Benno Stein and Martin Potthast	12113
<i>WhitenedCSE: Whitening-based Contrastive Learning of Sentence Embeddings</i>	
Wenjie Zhuo, Yifan Sun, Xiaohan Wang, Linchao Zhu and Yi Yang	12135
<i>Federated Learning for Semantic Parsing: Task Formulation, Evaluation Setup, New Algorithms</i>	
Tianshu Zhang, Changchang Liu, Wei-Han Lee, Yu Su and Huan Sun	12149
<i>Causality-Guided Multi-Memory Interaction Network for Multivariate Stock Price Movement Prediction</i>	
Di Luo, Weiheng Liao, Shuqi Li, Xin Cheng and Rui Yan	12164
<i>DSRM: Boost Textual Adversarial Training with Distribution Shift Risk Minimization</i>	
SongYang Gao, Shihan Dou, Yan Liu, Xiao Wang, Qi Zhang, Zhongyu Wei, Jin Ma and Ying Shan	12177
<i>A Simple and Flexible Modeling for Mental Disorder Detection by Learning from Clinical Questionnaires</i>	
Hoyun Song, Jisu Shin, Huije Lee and Jong Park	12190
<i>Downstream Datasets Make Surprisingly Good Pretraining Corpora</i>	
Kundan Krishna, Saurabh Garg, Jeffrey P. Bigham and Zachary Lipton	12207
<i>Towards Open-World Product Attribute Mining: A Lightly-Supervised Approach</i>	
Liyan Xu, Chenwei Zhang, Xian Li, Jingbo Shang and Jinho D. Choi	12223
<i>XDailyDialog: A Multilingual Parallel Dialogue Corpus</i>	
Zeming Liu, Ping Nie, Jie Cai, Haifeng Wang, Zheng-Yu Niu, PENG ZHANG, Mrinmaya Sachan and Kaiping Peng	12240
<i>PAL to Lend a Helping Hand: Towards Building an Emotion Adaptive Polite and Empathetic Counseling Conversational Agent</i>	
Kshitij Mishra, Priyanshu Priya and Asif Ekbal	12254
<i>Bidirectional Generative Framework for Cross-domain Aspect-based Sentiment Analysis</i>	
Yue Deng, Wenxuan Zhang, Sinno Jialin Pan and Lidong Bing	12272
<i>Contrastive Decoding: Open-ended Text Generation as Optimization</i>	
Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer and Mike Lewis	12286
<i>Resolving Indirect Referring Expressions for Entity Selection</i>	
Mohammad Javad Hosseini, Filip Radlinski, Silvia Pareti and Annie Louis	12313
<i>Accelerating Transformer Inference for Translation via Parallel Decoding</i>	
Andrea Santilli, Silvio Severino, Emilian Postolache, Valentino Maiorca, Michele Mancusi, Riccardo Marin and Emanuele Rodola	12336
<i>Hard Sample Aware Prompt-Tuning</i>	
Yuanjian Xu, Qi An, Jiahuan Zhang, Peng Li and Zaiqing Nie	12356
<i>WikiBio: a Semantic Resource for the Intersectional Analysis of Biographical Events</i>	
Marco Antonio Stranisci, Rossana Damiano, Enrico Mensa, Viviana Patti, Daniele P. Radicioni and Tommaso Caselli	12370

<i>Best-k Search Algorithm for Neural Text Generation</i>	
Jiacheng Xu, Caiming Xiong, silvio savarese and Yingbo Zhou	12385
<i>Towards Leaving No Indic Language Behind: Building Monolingual Corpora, Benchmark and Models for Indic Languages</i>	
Sumanth Doddapaneni, Rahul Aralikkatte, Gowtham Ramesh, Shreya Goyal, Mitesh M. Khapra, Anoop Kunchukuttan and Pratyush Kumar	12402
<i>Transforming Visual Scene Graphs to Image Captions</i>	
Xu Yang, Jiawei Peng, Zihua Wang, Haiyang Xu, Qinghao Ye, Chenliang Li, Songfang Huang, Fei Huang, Zhangzikang Li and Yu Zhang	12427
<i>Hybrid Transducer and Attention based Encoder-Decoder Modeling for Speech-to-Text Tasks</i>	
Yun Tang, Anna Y Sun, Hirofumi Inaguma, Xinyue Chen, Ning Dong, Xutai Ma, Paden D Tomassello and Juan Pino	12441
<i>Improving Domain Generalization for Prompt-Aware Essay Scoring via Disentangled Representation Learning</i>	
Zhiwei Jiang, Tianyi Gao, Yafeng Yin, Meng Liu, Hua Yu, Zifeng Cheng and Qing Gu	12456
<i>What's the Meaning of Superhuman Performance in Today's NLU?</i>	
Simone Tedeschi, Johan Bos, Thierry Declerck, Jan Hajič, Daniel Hershcovich, Eduard H Hovy, Alexander Koller, Simon Krek, Steven Schockaert, Rico Sennrich, Ekaterina Shutova and Roberto Navigli	12471
<i>PromptNER: Prompt Locating and Typing for Named Entity Recognition</i>	
Yongliang Shen, Zeqi Tan, Shuhui Wu, Wenqi Zhang, Rongsheng Zhang, Yadong Xi, Weiming Lu and Yueting Zhuang	12492
<i>Hints on the data for language modeling of synthetic languages with transformers</i>	
Rodolfo Joel Zevallos and Nuria Bel	12508
<i>Neural Machine Translation Methods for Translating Text to Sign Language Glosses</i>	
Dele Zhu, Vera Czehmann and Eleftherios Avramidis	12523
<i>Revisiting Event Argument Extraction: Can EAE Models Learn Better When Being Aware of Event Co-occurrences?</i>	
Yuxin He, Jingyue Hu and Buzhou Tang	12542
<i>HAUSER: Towards Holistic and Automatic Evaluation of Simile Generation</i>	
Qianyu He, Yikai Zhang, Jiaqing Liang, Yuncheng Huang, Yanghua Xiao and Yunwen Chen	12557
<i>Large-scale Lifelong Learning of In-context Instructions and How to Tackle It</i>	
Jisoo Mok, Jaeyoung Do, Sungjin Lee, Tara Taghavi, Seunghak Yu and Sungroh Yoon	12573
<i>Controllable Text Generation via Probability Density Estimation in the Latent Space</i>	
Yuxuan Gu, Xiaocheng Feng, Sicheng Ma, Lingyuan Lingyuan Zhang, Heng Gong, Weihong Zhong and Bing Qin	12590
<i>Learning Latent Relations for Temporal Knowledge Graph Reasoning</i>	
Mengqi Zhang, Yuwei Xia, Qiang Liu, Shu Wu and Liang Wang	12617
<i>DT-Solver: Automated Theorem Proving with Dynamic-Tree Sampling Guided by Proof-level Value Function</i>	
Haiming Wang, Ye Yuan, Zhengying Liu, Jianhao Shen, Yichun Yin, Jing Xiong, Enze Xie, Han Shi, Yujun Li, lin li, Jian Yin, Zhenguo Li and Xiaodan Liang	12632

<i>Unsupervised Selective Rationalization with Noise Injection</i>	
Adam Storek, Melanie Subbiah and Kathleen McKeown	12647
<i>Understanding In-Context Learning via Supportive Pretraining Data</i>	
Xiaochuang Han, Daniel Simig, Todor Mihaylov, Yulia Tsvetkov, Asli Celikyilmaz and Tianlu Wang	12660
<i>ETHICIST: Targeted Training Data Extraction Through Loss Smoothed Soft Prompting and Calibrated Confidence Estimation</i>	
Zhexin Zhang, Jiaxin Wen and Minlie Huang	12674
<i>Effective Contrastive Weighting for Dense Query Expansion</i>	
Xiao Wang, Sean MacAvaney, Craig Macdonald and Iadh Ounis	12688
<i>Improving the Detection of Multilingual Online Attacks with Rich Social Media Data from Singapore</i>	
Janosch Haber, Bertie Vidgen, Matthew S Chapman, Vibhor Agarwal, Roy Ka-Wei Lee, Yong Keong Yap and Paul Röttger	12705
<i>Reanalyzing L2 Preposition Learning with Bayesian Mixed Effects and a Pretrained Language Model</i>	
Jakob Prange and Man Ho Ivy Wong	12722
<i>Socratic Pretraining: Question-Driven Pretraining for Controllable Summarization</i>	
Artidoro Pagnoni, Alex Fabbri, Wojciech Kryscinski and Chien-Sheng Jason Wu	12737
<i>MatCha: Enhancing Visual Language Pretraining with Math Reasoning and Chart Derendering</i>	
Fangyu Liu, Francesco Piccinno, Syrine Krichene, Chenxi Pang, Kenton Lee, Mandar Joshi, Yasemin Altun, Nigel Collier and Julian Martin Eisenschlos	12756
<i>MGR: Multi-generator Based Rationalization</i>	
Wei Liu, Haozhao Wang, Jun Wang, Ruixuan Li, Xinyang Li, YuanKai Zhang and Yang Qiu	12771
<i>BUMP: A Benchmark of Unfaithful Minimal Pairs for Meta-Evaluation of Faithfulness Metrics</i>	
Liang Ma, Shuyang Cao, Robert L Logan IV, Di Lu, Shihao Ran, Ke Zhang, Joel Tetreault and Alejandro Jaimes	12788
<i>Is Fine-tuning Needed? Pre-trained Language Models Are Near Perfect for Out-of-Domain Detection</i>	
Rheeya Uppaal, Junjie Hu and Yixuan Li	12813
<i>UniSumm and SummZoo: Unified Model and Diverse Benchmark for Few-Shot Summarization</i>	
Yulong Chen, Yang Liu, Ruochen Xu, Ziyi Yang, Chenguang Zhu, Michael Zeng and Yue Zhang	12833
<i>RADE: Reference-Assisted Dialogue Evaluation for Open-Domain Dialogue</i>	
Zhengliang Shi, Weiwei Sun, Shuo Zhang, Zhen Zhang, Pengjie Ren and Zhaochun Ren ..	12856
<i>An AMR-based Link Prediction Approach for Document-level Event Argument Extraction</i>	
Yuqing Yang, Qipeng Guo, Xiangkun Hu, Yue Zhang, Xipeng Qiu and Zheng Zhang	12876
<i>PuMer: Pruning and Merging Tokens for Efficient Vision Language Models</i>	
Qingqing Cao, Bhargavi Paranjape and Hannaneh Hajishirzi	12890
<i>Gloss-Free End-to-End Sign Language Translation</i>	
Kezhou Lin, Xiaohan Wang, Linchao Zhu, Ke Sun, bang zhang and Yi Yang	12904
<i>TAGPRIME: A Unified Framework for Relational Structure Extraction</i>	
I-Hung Hsu, Kuan-Hao Huang, Shuning Zhang, Wenxin Cheng, Prem Natarajan, Kai-Wei Chang and Nanyun Peng	12917

<i>Model-Generated Pretraining Signals Improves Zero-Shot Generalization of Text-to-Text Transformers</i> Linyuan Gong, Chenyan Xiong, Xiaodong Liu, Payal Bajaj, Yiqing Xie, Alvin Cheung, Jianfeng Gao and Xia Song	12933
<i>BITE: Textual Backdoor Attacks with Iterative Trigger Injection</i> Jun Yan, Vansh Gupta and Xiang Ren	12951
<i>A Crosslingual Investigation of Conceptualization in 1335 Languages</i> Yihong Liu, Haotian Ye, Leonie Weissweiler, Philipp Wicke, Renhao Pei, Robert Zangenfeind and Hinrich Schütze	12969
<i>Exploring and Verbalizing Academic Ideas by Concept Co-occurrence</i> Yi Xu, Shuqian Sheng, Bo Xue, Luoyi Fu, Xinbing Wang and Chenghu Zhou	13001
<i>mCLIP: Multilingual CLIP via Cross-lingual Transfer</i> Guanhua Chen, Lu Hou, Yun Chen, Wenliang Dai, Lifeng Shang, Xin Jiang, Qun Liu, Jia Pan and Wenping Wang	13028
<i>Distantly Supervised Course Concept Extraction in MOOCs with Academic Discipline</i> Mengying Lu, Yuquan Wang, Jifan Yu, Yexing Du, Lei Hou and Juanzi Li	13044
<i>Extrinsic Evaluation of Machine Translation Metrics</i> Nikita Moghe, Tom Sherborne, Mark Steedman and Alexandra Birch	13060
<i>ExplainMeetSum: A Dataset for Explainable Meeting Summarization Aligned with Human Intent</i> Hyun Kim, Minsoo Cho and Seung-Hoon Na	13079
<i>A Cross-Modality Context Fusion and Semantic Refinement Network for Emotion Recognition in Conversation</i> Xiaoheng Zhang and Yang Li	13099
<i>CAT: A Contextualized Conceptualization and Instantiation Framework for Commonsense Reasoning</i> Weiqi Wang, Tianqing Fang, Baixuan Xu, Chun Yi Louis Bo, Yangqiu Song and Lei Chen	13111
<i>The Elephant in the Room: Analyzing the Presence of Big Tech in Natural Language Processing Research</i> Mohamed Abdalla, Jan Philip Wahle, Terry Lima Ruas, Aurélie Névéol, Fanny Ducel, Saif M. Mohammad and Karen Fort	13141
<i>Language of Bargaining</i> Mourad Heddaya, Solomon E Dworkin, Chenhao Tan, Rob Voigt and Alexander K Zentefis	13161
<i>Do Question Answering Modeling Improvements Hold Across Benchmarks?</i> Nelson F. Liu, Tony Lee, Robin Jia and Percy Liang	13186
<i>VLN-Trans: Translator for the Vision and Language Navigation Agent</i> Yue Zhang and Parisa Kordjamshidi	13219
<i>Bridging the Gap between Decision and Logits in Decision-based Knowledge Distillation for Pre-trained Language Models</i> Qinhong Zhou, Zonghan Yang, Peng Li and Yang Liu	13234
<i>Continual Contrastive Finetuning Improves Low-Resource Relation Extraction</i> Wenxuan Zhou, Sheng Zhang, Tristan Naumann, Muhao Chen and Hoifung Poon	13249
<i>KGA: A General Machine Unlearning Framework Based on Knowledge Gap Alignment</i> Lingzhi Wang, Tong Chen, Wei Yuan, Xingshan Zeng, Kam-Fai Wong and Hongzhi Yin	13264

<i>UniCoRN: Unified Cognitive Signal Reconstruction bridging cognitive signals and human language</i> Nuwa Xi, Sendong Zhao, Haochun Wang, Chi Liu, Bing Qin and Ting Liu	13277
<i>Dense-ATOMIC: Towards Densely-connected ATOMIC with High Knowledge Coverage and Massive Multi-hop Paths</i> Xiangqing Shen, Siwei Wu and Rui Xia	13292
<i>Shrinking Embeddings for Hyper-Relational Knowledge Graphs</i> Bo Xiong, Mojtaba Nayyeri, Shirui Pan and Steffen Staab	13306
<i>CTC-based Non-autoregressive Speech Translation</i> Chen Xu, Xiaoqian Liu, Xiaowen Liu, Qingxuan Sun, Yuhao Zhang, Murun Yang, Qianqian Dong, Tom Ko, Mingxuan Wang, Tong Xiao, Anxiang Ma and Jingbo Zhu	13321
<i>Attention as a Guide for Simultaneous Speech Translation</i> Sara Papi, Matteo Negri and Marco Turchi	13340
<i>On Complementarity Objectives for Hybrid Retrieval</i> Dohyeon Lee, Seung-won Hwang, Kyungjae Lee, Seungtaek Choi and Sunghyun Park	13357
<i>C-STANCE: A Large Dataset for Chinese Zero-Shot Stance Detection</i> Chenye Zhao, Yingjie Li and Cornelia Caragea	13369
<i>Wukong-Reader: Multi-modal Pre-training for Fine-grained Visual Document Understanding</i> Haoli Bai, Zhiguang Liu, Xiaojun Meng, li wentao, Shuang Liu, Yifeng LUO, nian xie, Rongfu Zheng, Liangwei Wang, Lu Hou, Jiansheng Wei, Xin Jiang and Qun Liu	13386
<i>PaCE: Unified Multi-modal Dialogue Pre-training with Progressive and Compositional Experts</i> Yunshui Li, Binyuan Hui, ZhiChao Yin, Min Yang, Fei Huang and Yongbin Li	13402
<i>MVP-Tuning: Multi-View Knowledge Retrieval with Prompt Tuning for Commonsense Reasoning</i> Yongfeng Huang, Yanyang Li, Yichong Xu, Lin Zhang, ruyi gan, Jiaying Zhang and Liwei Wang	13417
<i>PEIT: Bridging the Modality Gap with Pre-trained Models for End-to-End Image Translation</i> Shaolin Zhu, Shangjie Li, Yikun Lei and Deyi Xiong	13433
<i>Topic-Guided Sampling For Data-Efficient Multi-Domain Stance Detection</i> Erik Arakelyan, Arnav Arora and Isabelle Augenstein	13448
<i>DiSCoMaT: Distantly Supervised Composition Extraction from Tables in Materials Science Articles</i> Tanishq Gupta, Mohd Zaki, Devanshi Khatsuriya, Kausik Hira, N M Anoop Krishnan and Mausam -	13465
<i>Self-Instruct: Aligning Language Models with Self-Generated Instructions</i> Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi and Hannaneh Hajishirzi	13484
<i>Disentangled Phonetic Representation for Chinese Spelling Correction</i> Zihong Liang, Xiaojun Quan and Qifan Wang	13509
<i>Dissecting Transformer Length Extrapolation via the Lens of Receptive Field Analysis</i> Ta-Chung Chi, Ting-Han Fan, alexander rudnicky and Peter J Ramadge	13522
<i>CHBias: Bias Evaluation and Mitigation of Chinese Conversational Language Models</i> Jiaxu Zhao, Meng Fang, Zijing Shi, Yitong Li, Ling Chen and Mykola Pechenizkiy	13538

<i>Learning New Skills after Deployment: Improving open-domain internet-driven dialogue with human feedback</i>	
Jing Xu, Megan Ung, Mojtaba Komeili, Kushal Arora, Y-Lan Boureau and Jason Weston .	13557
<i>Uncovering and Categorizing Social Biases in Text-to-SQL</i>	
Yan Liu, Yan Gao, Zhe Su, Xiaokang Chen, Elliott Ash and Jian-Guang LOU	13573
<i>On the Compositional Generalization in Versatile Open-domain Dialogue</i>	
Tingchen Fu, Xueliang Zhao, Lema Liu and Rui Yan	13585
<i>What is the Real Intention behind this Question? Dataset Collection and Intention Classification</i>	
Maryam Sadat Mirzaei, Kourosh Meshgi and Satoshi Sekine	13606
<i>Conjunct Resolution in the Face of Verbal Omissions</i>	
Royi Rassin, Yoav Goldberg and Reut Tsarfaty	13623
<i>Training Models to Generate, Recognize, and Reframe Unhelpful Thoughts</i>	
Mounica Maddela, Megan Ung, Jing Xu, Andrea Madotto, Heather M Foran and Y-Lan Boureau	13641
<i>Learning In-context Learning for Named Entity Recognition</i>	
Jiawei Chen, Yaojie Lu, Hongyu Lin, Jie Lou, Wei Jia, Dai Dai, Hua Wu, Boxi Cao, Xianpei Han and Le Sun	13661
<i>Holistic Prediction on a Time-Evolving Attributed Graph</i>	
Shohei Yamasaki, Yuya Sasaki, Panagiotis Karras and Makoto Onizuka	13676
<i>Modeling Instance Interactions for Joint Information Extraction with Neural High-Order Conditional Random Field</i>	
Zixia Jia, Zhaohui Yan, Wenjuan Han, Zilong Zheng and Kewei Tu	13695
<i>Training Trajectories of Language Models Across Scales</i>	
Mengzhou Xia, Mikel Artetxe, Chunting Zhou, Xi Victoria Lin, Ramakanth Pasunuru, Danqi Chen, Luke Zettlemoyer and Veselin Stoyanov	13711
<i>A Diverse Set of Freely Available Linguistic Resources for Turkish</i>	
Duygu ALTINOK	13739
<i>Measuring Consistency in Text-based Financial Forecasting Models</i>	
Linyi Yang, Yingpeng Ma and Yue Zhang	13751
<i>Optimal Transport for Unsupervised Hallucination Detection in Neural Machine Translation</i>	
Nuno M. Guerreiro, Pierre Colombo, Pablo Piantanida and André Martins	13766
<i>RankCSE: Unsupervised Sentence Representations Learning via Learning to Rank</i>	
Jiduan Liu, Jiahao Liu, Qifan Wang, Jingang Wang, Wei Wu, Yunsen Xian, Dongyan Zhao, Kai Chen and Rui Yan	13785
<i>Entailment as Robust Self-Learner</i>	
Jiaxin Ge, Hongyin Luo, Yoon Kim and James Glass	13803
<i>ReCode: Robustness Evaluation of Code Generation Models</i>	
Shiqi Wang, Zheng Li, Haifeng Qian, Chenghao Yang, Zijian Wang, Mingyue Shang, Varun Kumar, Samson Tan, Baishakhi Ray, Parminder Bhatia, Ramesh Nallapati, Murali Krishna Ramanathan, Dan Roth and Bing Xiang	13818

<i>EPIC: Multi-Perspective Annotation of a Corpus of Irony</i> Simona Frenda, Alessandro Pedrani, Valerio Basile, Soda Marem Lo, Alessandra Teresa Cignarella, Raffaella Panizzon, Cristina Marco, Bianca Scarlini, Viviana Patti, Cristina Bosco and Davide Bernardi	13844
<i>Dialogue Summarization with Static-Dynamic Structure Fusion Graph</i> Shen Gao, Xin Cheng, Mingzhe Li, Xiuying Chen, Jinpeng Li, Dongyan Zhao and Rui Yan	13858
<i>Large-Scale Correlation Analysis of Automated Metrics for Topic Models</i> Jia Peng Lim and Hady Lauw	13874
<i>U-CREAT: Unsupervised Case Retrieval using Events extrAcTion</i> Abhinav Joshi, Akshat Sharma, Sai Kiran Tanikella and Ashutosh Modi	13899
<i>ArgAnalysis35K : A large-scale dataset for Argument Quality Analysis</i> Omkar Jayant Joshi, Priya N Pitre and Yashodhara Haribhakta.....	13916
<i>Reference Matters: Benchmarking Factual Error Correction for Dialogue Summarization with Fine-grained Evaluation Framework</i> Mingqi Gao, Xiaojun Wan, Jia Su, Zhefeng Wang and baoxing Huai	13932
<i>Minding Language Models' (Lack of) Theory of Mind: A Plug-and-Play Multi-Character Belief Tracker</i> Melanie Sclar, Sachin Kumar, Peter West, Alane Suhr, Yejin Choi and Yulia Tsvetkov	13960
<i>Don't Retrain, Just Rewrite: Countering Adversarial Perturbations by Rewriting Text</i> Ashim Gupta, Carter Wood Blum, Temma Choji, Yingjie Fei, Shalin Shah, Alakananda Vempala and Vivek Srikumar	13981
<i>Aggregating Multiple Heuristic Signals as Supervision for Unsupervised Automated Essay Scoring</i> Cong Wang, Zhiwei Jiang, Yafeng Yin, Zifeng Cheng, Shiping Ge and Qing Gu	13999
<i>Mitigating Label Biases for In-context Learning</i> Yu Fei, Yifan Hou, Zeming Chen and Antoine Bosselut.....	14014
<i>QUEST: A Retrieval Dataset of Entity-Seeking Queries with Implicit Set Operations</i> Chaitanya Malaviya, Peter Shaw, Ming-Wei Chang, Kenton Lee and Kristina Toutanova ..	14032
<i>Dynamic Heterogeneous-Graph Reasoning with Language Models and Knowledge Representation Learning for Commonsense Question Answering</i> Yujie Wang, Hu Zhang, Jiye Liang and Ru Li	14048
<i>Do You Hear The People Sing? Key Point Analysis via Iterative Clustering and Abstractive Summarization</i> Hao Li, Viktor Schlegel, Riza Batista-Navarro and Goran Nenadic	14064
<i>Ambiguous Learning from Retrieval: Towards Zero-shot Semantic Parsing</i> Shan Wu, Chunlei Xin, Hongyu Lin, Xianpei Han, Cao Liu, Jiansong Chen, Fan Yang, Guanglu Wan and Le Sun	14081
<i>Explicit Syntactic Guidance for Neural Text Generation</i> Yafu Li, Leyang Cui, Jianhao Yan, Yongjing Yin, Wei Bi, Shuming Shi and Yue Zhang ...	14095
<i>What does a Text Classifier Learn about Morality? An Explainable Method for Cross-Domain Comparison of Moral Rhetoric</i> Enrico Liscio, Oscar Araque, Lorenzo Gatti, Ionut L. Constantinescu, Catholijn M Jonker, Kyriaki Kalimeri and Pradeep Kumar Murukannaiah.....	14113

<i>Graph-based Relation Mining for Context-free Out-of-vocabulary Word Embedding Learning</i> Ziran Liang, Yuyin Lu, HeGang Chen and Yanghui Rao	14133
<i>Multimodal Persona Based Generation of Comic Dialogs</i> Harsh Agrawal, Aditya M. Mishra, Manish Gupta and Mausam -	14150
<i>LLM-Blender: Ensembling Large Language Models with Pairwise Ranking and Generative Fusion</i> Dongfu Jiang, Xiang Ren and Bill Yuchen Lin	14165
<i>Seen to Unseen: Exploring Compositional Generalization of Multi-Attribute Controllable Dialogue Generation</i> Weihao Zeng, Lulu Zhao, Keqing He, Ruotong Geng, Jingang Wang, Wei Wu and Weiran Xu	14179
<i>Generating Structured Pseudo Labels for Noise-resistant Zero-shot Video Sentence Localization</i> Minghang Zheng, Shaogang Gong, Hailin Jin, Yuxin Peng and Yang Liu	14197
<i>IndicMT Eval: A Dataset to Meta-Evaluate Machine Translation Metrics for Indian Languages</i> Ananya Sai B, Tanay Dixit, Vignesh Nagarajan, Anoop Kunchukuttan, Pratyush Kumar, Mitesh M. Khapra and Raj Dabre	14210
<i>Weaker Than You Think: A Critical Look at Weakly Supervised Learning</i> Dawei Zhu, Xiaoyu Shen, Marius Mosbach, Andreas Joseph Stephan and Dietrich Klakow	14229
<i>Prompt Tuning Pushes Farther, Contrastive Learning Pulls Closer: A Two-Stage Approach to Mitigate Social Biases</i> Yingji Li, Mengnan Du, Xin Wang and Ying Wang	14254
<i>Towards Understanding Omission in Dialogue Summarization</i> Yicheng Zou, Kaitao Song, Xu Tan, Zhongkai Fu, Qi Zhang, Dongsheng Li and Tao Gui ..	14268
<i>Python Code Generation by Asking Clarification Questions</i> Haau-Sing (Xiaocheng) Li, Mohsen Mesgar, André Martins and Iryna Gurevych	14287
<i>A Compare-and-contrast Multistage Pipeline for Uncovering Financial Signals in Financial Reports</i> Jia-Huei Ju, Yu-Shiang Huang, Cheng-Wei Lin, Che Lin and Chuan-Ju Wang	14307
<i>Improving the robustness of NLI models with minimax training</i> Michalis Korakakis and Andreas Vlachos	14322
<i>USSA: A Unified Table Filling Scheme for Structured Sentiment Analysis</i> Zepeng Zhai, Hao Chen, Ruifan Li and Xiaojie WANG	14340
<i>PAD-Net: An Efficient Framework for Dynamic Networks</i> Shwai He, Liang Ding, Daize Dong, Boan Liu, Fuqiang Yu and Dacheng Tao	14354
<i>Resolving Ambiguities in Text-to-Image Generative Models</i> Ninareh Mehrabi, Palash Goyal, Apurv Verma, Jwala Dhamala, Varun Kumar, Qian Hu, Kai-Wei Chang, Richard Zemel, Aram Galstyan and Rahul Gupta	14367
<i>Knowledge Unlearning for Mitigating Privacy Risks in Language Models</i> Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran and Minjoon Seo	14389
<i>Unnatural Instructions: Tuning Language Models with (Almost) No Human Labor</i> Or Honovich, Thomas Scialom, Omer Levy and Timo Schick	14409

<i>To Adapt or to Annotate: Challenges and Interventions for Domain Adaptation in Open-Domain Question Answering</i>	
Dheeru Dua, Emma Strubell, Sameer Singh and Pat Verga	14429
<i>A Survey for Efficient Open Domain Question Answering</i>	
Qin Zhang, Shangsi Chen, Dongkuan Xu, Qingqing Cao, Xiaojun Chen, Trevor Cohn and Meng Fang	14447
<i>Script Normalization for Unconventional Writing of Under-Resourced Languages in Bilingual Communities</i>	
Sina Ahmadi and Antonios Anastasopoulos	14466
<i>Compositional Generalization without Trees using Multiset Tagging and Latent Permutations</i>	
Matthias Lindemann, Alexander Koller and Ivan Titov	14488
<i>ManagerTower: Aggregating the Insights of Uni-Modal Experts for Vision-Language Representation Learning</i>	
Xiao Xu, Bei Li, Chenfei Wu, Shao-Yen Tseng, Anahita Bhiwandiwalla, Shachar Rosenman, Vasudev Lal, Wanxiang Che and Nan Duan	14507
<i>Finding the Pillars of Strength for Multi-Head Attention</i>	
Jinjie Ni, Rui Mao, Zonglin Yang, Han Lei and Erik Cambria	14526
<i>Jointprop: Joint Semi-supervised Learning for Entity and Relation Extraction with Heterogeneous Graph-based Propagation</i>	
yandan zheng, Anran Hao and Anh Tuan Luu	14541
<i>Reasoning over Hierarchical Question Decomposition Tree for Explainable Question Answering</i>	
Jiajie Zhang, Shulin Cao, Tingjian Zhang, Xin Lv, Juanzi Li, Lei Hou, Jiaxin Shi and Qi Tian	14556
<i>Faking Fake News for Real Fake News Detection: Propaganda-Loaded Training Data Generation</i>	
Kung-Hsiang Huang, Kathleen McKeown, Preslav Nakov, Yejin Choi and Heng Ji	14571
<i>A Length-Extrapolatable Transformer</i>	
Yutao Sun, Li Dong, Barun Patra, Shuming Ma, Shaohan Huang, Alon Benhaim, Vishrav Chaudhary, Xia Song and Furu Wei	14590
<i>A Survey of Deep Learning for Mathematical Reasoning</i>	
Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck and Kai-Wei Chang	14605
<i>A Systematic Study of Knowledge Distillation for Natural Language Generation with Pseudo-Target Training</i>	
Nitay Calderon, Subhabrata Mukherjee, Roi Reichart and Amir Kantor	14632
<i>Vision Language Pre-training by Contrastive Learning with Cross-Modal Similarity Regulation</i>	
Chaoya Jiang, Wei Ye, Haiyang Xu, Songfang Huang, Fei Huang and Shikun Zhang	14660
<i>Tell2Design: A Dataset for Language-Guided Floor Plan Generation</i>	
Sicong Leng, Yang Zhou, Mohammed Haroon Dupty, Wee Sun Lee, Sam C Joyce and Wei Lu	14680
<i>Are Human Explanations Always Helpful? Towards Objective Evaluation of Human Natural Language Explanations</i>	
Bingsheng Yao, Prithviraj Sen, Lucian Popa, James Hendler and Dakuo Wang	14698

<i>Rethinking Annotation: Can Language Learners Contribute?</i> Haneul Yoo, Rifki Afina Putri, Changyoon Lee, Youngin Lee, So-Yeon Ahn, Dongyeop Kang and Alice Oh	14714
<i>Information Screening whilst Exploiting! Multimodal Relation Extraction with Feature Denoising and Multimodal Topic Modeling</i> Shengqiong Wu, Hao Fei, Yixin Cao, Lidong Bing and Tat-Seng Chua	14734
<i>MultiEMO: An Attention-Based Correlation-Aware Multimodal Fusion Framework for Emotion Recognition in Conversations</i> Tao Shi and Shao-Lun Huang	14752
<i>Learning Language-Specific Layers for Multilingual Machine Translation</i> Telmo Pires, Robin M. Schmidt, Yi-Hsiu Liao and Stephan Peitz	14767
<i>Personality Understanding of Fictional Characters during Book Reading</i> Mo Yu, Jiangnan Li, Shunyu Yao, Wenjie Pang, Xiaochen Zhou, Zhou Xiao, Fandong Meng and Jie Zhou	14784
<i>StoryTrans: Non-Parallel Story Author-Style Transfer with Discourse Representations and Content Enhancing</i> Xuekai Zhu, Jian Guan, Minlie Huang and Juan Liu	14803
<i>Towards Benchmarking and Improving the Temporal Reasoning Capability of Large Language Models</i> Qingyu Tan, Hwee Tou Ng and Lidong Bing	14820
<i>Finding the SWEET Spot: Analysis and Improvement of Adaptive Inference in Low Resource Settings</i> Daniel Rotem, Michael Hassid, Jonathan Mamou and Roy Schwartz	14836
<i>Large Language Models Are Reasoning Teachers</i> Namgyu Ho, Laura Schmid and Se-Young Yun	14852
<i>Abductive Commonsense Reasoning Exploiting Mutually Exclusive Explanations</i> Wenting Zhao, Justin Chiu, Claire Cardie and Alexander Rush	14883
<i>PESCO: Prompt-enhanced Self Contrastive Learning for Zero-shot Text Classification</i> Yau-Shian Wang, Ta-Chung Chi, Ruohong Zhang and YIMING YANG	14897
<i>Visually-augmented pretrained language models for NLP tasks without images</i> Hangyu Guo, Kun Zhou, Wayne Xin Zhao, Qinyu Zhang and Ji-Rong Wen	14912
<i>Using counterfactual contrast to improve compositional generalization for multi-step quantitative reasoning</i> Armineh Nourbakhsh, Sameena Shah and Carolyn Rosé	14930
<i>A Needle in a Haystack: An Analysis of High-Agreement Workers on MTurk for Summarization</i> Lining Zhang, Simon Mille, Yufang Hou, Daniel Deutsch, Elizabeth Clark, Yixin Liu, Saad Mahmood, Sebastian Gehrmann, Miruna Adriana Clinciu, Khyathi Raghavi Chandu and João Sedoc	14944
<i>TAVT: Towards Transferable Audio-Visual Text Generation</i> Wang Lin, Tao Jin, Wenwen Pan, Linjun Li, Xize Cheng, Ye Wang and Zhou Zhao	14983
<i>MeetingQA: Extractive Question-Answering on Meeting Transcripts</i> Archiki Prasad, Trung Bui, Seunghyun Yoon, Hanieh Deilamsalehy, Franck Dernoncourt and Mohit Bansal	15000
<i>FERMAT: An Alternative to Accuracy for Numerical Reasoning</i> Jasivan Alex Sivakumar and Nafise Sadat Moosavi	15026

<i>Don't Forget Your ABC's: Evaluating the State-of-the-Art in Chat-Oriented Dialogue Systems</i> Sarah E. Finch, James D. Finch and Jinho D. Choi	15044
<i>Decoder Tuning: Efficient Language Understanding as Decoding</i> Janqu CUI, Wentao Li, Ning Ding, Longtao Huang, Zhiyuan Liu and Maosong Sun	15072
<i>The KITMUS Test: Evaluating Knowledge Integration from Multiple Sources</i> Akshatha Arodi, Martin Pömsl, Kaheer Suleman, Adam Trischler, Alexandra Olteanu and Jackie Chi Kit Cheung	15088
<i>CREST: A Joint Framework for Rationalization and Counterfactual Text Generation</i> Marcos Treviso, Alexis Ross, Nuno M. Guerreiro and André Martins	15109
<i>Towards Unifying Multi-Lingual and Cross-Lingual Summarization</i> Jiaan Wang, Fandong Meng, Duo Zheng, Yunlong Liang, Zhixu Li, Jianfeng Qu and Jie Zhou	15127
<i>On Improving Summarization Factual Consistency from Natural Language Feedback</i> Yixin Liu, Budhaditya Deb, Milagro Teruel, Aaron L Halfaker, Dragomir Radev and Ahmed Hassan Awadallah	15144
<i>From Dogwhistles to Bullhorns: Unveiling Coded Rhetoric with Language Models</i> Julia Mendelsohn, Ronan Le Bras, Yejin Choi and Maarten Sap	15162
<i>Exploring Large Language Models for Classical Philology</i> Frederick Riemenschneider and Anette Frank	15181
<i>LayoutMask: Enhance Text-Layout Interaction in Multi-modal Pre-training for Document Understanding</i> Yi Tu, Ya Guo, Huan Chen and jinyang tang	15200
<i>Hearing Lips in Noise: Universal Viseme-Phoneme Mapping and Transfer for Robust Audio-Visual Speech Recognition</i> Yuchen Hu, Ruizhe Li, Chen Chen, Chengwei Qin, Qiu-Shi Zhu and Eng Siong Chng	15213
<i>An Extensible Plug-and-Play Method for Multi-Aspect Controllable Text Generation</i> Xuancheng Huang, Zijun Liu, Peng Li, Tao Li, Maosong Sun and Yang Liu	15233
<i>Double-Branch Multi-Attention based Graph Neural Network for Knowledge Graph Completion</i> Hongcai Xu, Junpeng Bao and Wenbo Liu	15257
<i>Dual Cache for Long Document Neural Coreference Resolution</i> Qipeng Guo, Xiangkun Hu, Yue Zhang, Xipeng Qiu and Zheng Zhang	15272
<i>Knowledge Transfer in Incremental Learning for Multilingual Neural Machine Translation</i> Kaiyu Huang, Peng Li, Jin Ma, Ting Yao and Yang Liu	15286
<i>DisorBERT: A Double Domain Adaptation Model for Detecting Signs of Mental Disorders in Social Media</i> Mario Ezra Aragon, Adrian Pastor Lopez Monroy, Luis C Gonzalez, David E. Losada and Manuel Montes	15305
<i>Toward Interactive Dictation</i> Belinda Z. Li, Jason Eisner, Adam Pauls and Sam Thomson	15319
<i>CodeIE: Large Code Generation Models are Better Few-Shot Information Extractors</i> Peng Li, Tianxiang Sun, Qiong Tang, Hang Yan, Yuanbin Wu, Xuanjing Huang and Xipeng Qiu	15339

<i>Beyond English-Centric Bitexts for Better Multilingual Language Representation Learning</i> Barun Patra, Saksham Singhal, Shaohan Huang, Zewen Chi, Li Dong, Furu Wei, Vishrav Chaudhary and Xia Song	15354
<i>Bridging The Gap: Entailment Fused-T5 for Open-retrieval Conversational Machine Reading Comprehension</i> Xiao Zhang, Heyan Huang, Zewen Chi and Xian-Ling Mao	15374
<i>LiveChat: A Large-Scale Personalized Dialogue Dataset Automatically Constructed from Live Streaming</i> Jingsheng Gao, Yixin Lian, Ziyi Zhou, yuzhuo fu and Baoyuan Wang	15387
<i>Prompting PaLM for Translation: Assessing Strategies and Performance</i> David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar and George Foster	15406
<i>Exploring Lottery Prompts for Pre-trained Language Models</i> Yulin Chen, Ning Ding, Xiaobin Wang, Shengding Hu, Haitao Zheng, Zhiyuan Liu and Pengjun Xie	15428
<i>A Facial Expression-Aware Multimodal Multi-task Learning Framework for Emotion Recognition in Multi-party Conversations</i> Wenjie Zheng, Jianfei Yu, Rui Xia and Shijin Wang	15445
<i>TeAST: Temporal Knowledge Graph Embedding via Archimedean Spiral Timeline</i> Jiang Li, Xiangdong Su and Guanglei Gao	15460
<i>Human Inspired Progressive Alignment and Comparative Learning for Grounded Word Acquisition</i> Yuwei Bao, Barrett Martin Lattimer and Joyce Chai	15475
<i>Conjunct Lengths in English, Dependency Length Minimization, and Dependency Structure of Coordination</i> Adam Przepiórkowski and Michał Woźniak	15494
<i>LeXFiles and LegalLAMA: Facilitating English Multinational Legal Language Model Development</i> Ilias Chalkidis, Nicolas Garneau, Catalina Goanta, Daniel Katz and Anders Søgaard	15513
<i>Revisiting Commonsense Reasoning in Machine Translation: Training, Evaluation and Challenge</i> Xuebo Liu, Yutong Wang, Derek F. Wong, Runzhe Zhan, Liangxuan Yu and Min Zhang ..	15536
<i>NOTABLE: Transferable Backdoor Attacks Against Prompt-based NLP Models</i> Kai Mei, Zheng Li, Zhenting Wang, Yang Zhang and Shiqing Ma	15551
<i>Revisiting Relation Extraction in the era of Large Language Models</i> Somin Wadhwa, Silvio Amir and Byron C. Wallace	15566
<i>Pre-trained Language Models Can be Fully Zero-Shot Learners</i> Xuandong Zhao, Siqi Ouyang, Zhiguo Yu, Ming Wu and Lei Li	15590
<i>Can Large Language Models Be an Alternative to Human Evaluations?</i> Cheng-Han Chiang and Hung-yi Lee	15607
<i>HyperMixer: An MLP-based Low Cost Alternative to Transformers</i> Florian Mai, Arnaud Pannatier, Fabio J Fehr, Haolin Chen, Francois Marelli, Francois Fleuret and James Henderson	15632

<i>UnitY: Two-pass Direct Speech-to-speech Translation with Discrete Units</i>	
Hirofumi Inaguma, Sravya Popuri, Ilia Kulikov, Peng-Jen Chen, Changan Wang, Yu-An Chung, Yun Tang, Ann Lee, Shinji Watanabe and Juan Pino	15655
<i>Estimating the Uncertainty in Emotion Attributes using Deep Evidential Regression</i>	
Wen Wu, Chao Zhang and Philip C. Woodland.....	15681
<i>Annotation-Inspired Implicit Discourse Relation Classification with Auxiliary Discourse Connective Generation</i>	
Wei Liu and Michael Strube	15696
<i>Plug-and-Play Document Modules for Pre-trained Models</i>	
Chaojun Xiao, Zhengyan Zhang, Xu Han, Chi-Min Chan, Yankai Lin, Zhiyuan Liu, xiangyang li, Zhonghua Li, Zhao Cao and Maosong Sun	15713
<i>An Empirical Analysis of Parameter-Efficient Methods for Debiasing Pre-Trained Language Models</i>	
Zhongbin Xie and Thomas Lukasiewicz	15730
<i>Two-Stage Fine-Tuning for Improved Bias and Variance for Large Pretrained Language Models</i>	
Lijing Wang, Yingya Li, Timothy A Miller, Steven Bethard and Guergana Savova	15746
<i>A Comparative Study on the Impact of Model Compression Techniques on Fairness in Language Models</i>	
Krithika Ramesh, Arnav Chavan, Shrey Pandit and Sunayana Sitaram	15762
<i>Ranking-Enhanced Unsupervised Sentence Representation Learning</i>	
Yeon Seonwoo, Guoyin Wang, Changmin Seo, Sajal Choudhary, Jiwei Li, Xiang Li, Puyang Xu, Sunghyun Park and Alice Oh.....	15783
<i>To Revise or Not to Revise: Learning to Detect Improvable Claims for Argumentative Writing Support</i>	
Gabriella Skitalinskaya and Henning Wachsmuth	15799
<i>Human-in-the-loop Evaluation for Early Misinformation Detection: A Case Study of COVID-19 Treatments</i>	
Ethan Adrian Mendes, Yang Chen, Wei Xu and Alan Ritter	15817
<i>Composition-contrastive Learning for Sentence Embeddings</i>	
Sachin J. Chanchani and Ruihong Huang	15836
<i>Causes and Cures for Interference in Multilingual Translation</i>	
Uri Shaham, Maha Elbayad, Vedanuj Goswami, Omer Levy and Shruti Bhosale	15849
<i>Understanding and Bridging the Modality Gap for Speech Translation</i>	
Qingkai Fang and Yang Feng.....	15864
<i>Few-shot Reranking for Multi-hop QA via Language Model Prompting</i>	
Muhammad Khalifa, Lajanugen Logeswaran, Moontae Lee, Honglak Lee and Lu Wang... ..	15882
<i>DICE: Data-Efficient Clinical Event Extraction with Generative Models</i>	
Mingyu Derek Ma, Alexander K Taylor, Wei Wang and Nanyun Peng	15898
<i>XSemPLR: Cross-Lingual Semantic Parsing in Multiple Natural Languages and Meaning Representations</i>	
Yusen Zhang, Jun Wang, Zhiguo Wang and Rui Zhang	15918
<i>INK: Injecting kNN Knowledge in Nearest Neighbor Machine Translation</i>	
Wenhao Zhu, Jingjing Xu, Shujian Huang, Lingpeng Kong and Jiajun CHEN	15948

<i>Uncertainty Guided Label Denoising for Document-level Distant Relation Extraction</i>	
Qi Sun, Kun Huang, Xiaocui Yang, Pengfei Hong, Kun Zhang and Soujanya Poria	15960
<i>Cross-Modal Attribute Insertions for Assessing the Robustness of Vision-and-Language Learning</i>	
Shivaen Ramshetty, Gaurav Verma and Srijan Kumar	15974
<i>Crosslingual Generalization through Multitask Finetuning</i>	
Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, KHALID ALMUBARAK, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff and Colin Raffel	15991
<i>Evaluate AMR Graph Similarity via Self-supervised Learning</i>	
Ziyi Shou and Fangzhen Lin	16112
<i>Analyzing Transformers in Embedding Space</i>	
Guy Dar, Mor Geva, Ankit Gupta and Jonathan Berant	16124
<i>Few-Shot Data-to-Text Generation via Unified Representation and Multi-Source Learning</i>	
Alexander Hanbo Li, Mingyue Shang, Evangelia Spiliopoulou, Jie Ma, Patrick Ng, Zhiguo Wang, Bonan Min, William Yang Wang, Kathleen McKeown, Vittorio Castelli, Dan Roth and Bing Xiang	16171
<i>FactKG: Fact Verification via Reasoning on Knowledge Graphs</i>	
Jiho Kim, Sungjin Park, Yeonsu Kwon, Yohan Jo, James Thorne and Edward Choi	16190
<i>DrBERT: A Robust Pre-trained Model in French for Biomedical and Clinical domains</i>	
Yanis Labrak, Adrien Bazoge, Richard Dufour, Mickael Rouvier, Emmanuel Morin, Béatrice Daille and Pierre-Antoine Gourraud	16207
<i>Discriminative Reasoning with Sparse Event Representation for Document-level Event-Event Relation Extraction</i>	
Changsen Yuan, Heyan Huang, Yixin Cao and Yonggang Wen	16222
<i>Facilitating Fine-grained Detection of Chinese Toxic Language: Hierarchical Taxonomy, Resources, and Benchmarks</i>	
Junyu Lu, Bo Xu, Xiaokun Zhang, Changrong Min, Liang Yang and Hongfei LIN	16235
<i>SpeechMatrix: A Large-Scale Mined Corpus of Multilingual Speech-to-Speech Translations</i>	
Paul-Ambroise Augustin Duquenne, Hongyu Gong, Ning Dong, Jingfei Du, Ann Lee, Vedanuj Goswami, Changhan Wang, Juan Pino, Benoît Sagot and Holger Schwenk	16251
<i>Character-Aware Models Improve Visual Text Rendering</i>	
Rosanne Liu, Dan Garrette, Chitwan Saharia, William Chan, Adam Roberts, Sharan Narang, Irina Blok, RJ Mical, Mohammad Norouzi and Noah Constant	16270
<i>IDRISI-RA: The First Arabic Location Mention Recognition Dataset of Disaster Tweets</i>	
Reem Suwaileh, Muhammad Imran and Tamer Elsayed	16298
<i>FSUIE: A Novel Fuzzy Span Mechanism for Universal Information Extraction</i>	
Tianshuo Peng, Zuchao Li, Lefei Zhang, Bo Du and Hai Zhao	16318
<i>What Do NLP Researchers Believe? Results of the NLP Community Metasurvey</i>	
Julian Michael, Ari Holtzman, Alicia Parrish, Aaron Mueller, Alex Wang, Angelica Chen, Divyam Madaan, Nikita Nangia, Richard Yuanzhe Pang, Jason Phang and Samuel R. Bowman	16334

<i>Prototype-Guided Pseudo Labeling for Semi-Supervised Text Classification</i> Weiyi Yang, Richong Zhang, Junfan Chen, Lihong Wang and Jaemin Kim	16369
<i>LENS: A Learnable Evaluation Metric for Text Simplification</i> Mounica Maddela, Yao Dou, David Heineman and Wei Xu	16383
<i>MeetingBank: A Benchmark Dataset for Meeting Summarization</i> Yebowen Hu, Timothy Jeewun Ganter, Hanieh Deilamsalehy, Franck Dernoncourt, Hassan Foroosh and Fei Liu	16409
<i>UniEX: An Effective and Efficient Framework for Unified Information Extraction via a Span-extractive Perspective</i> yang ping, JunYu Lu, ruyi gan, Junjie Wang, Yuxiang Zhang, Pingjian Zhang and Jiaxing Zhang	16424
<i>DEplain: A German Parallel Corpus with Intralingual Translations into Plain Language for Sentence and Document Simplification</i> Regina Stodden, Omar Momen and Laura Kallmeyer	16441
<i>A Neural Divide-and-Conquer Reasoning Framework for Image Retrieval from Linguistically Complex Text</i> Yunxin Li, Baotian Hu, Yuxin Ding, Lin Ma and Min Zhang	16464
<i>RARR: Researching and Revising What Language Models Say, Using Language Models</i> Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Y Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan and Kelvin Guu	16477