

60th Annual Meeting of the Association for Computational Linguistics (ACL 2022)

Long Papers

Dublin, Ireland
22 – 27 May 2022

Volume 1 of 12

ISBN: 978-1-7138-6721-0

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© (2022) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2023)

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>AdapLeR: Speeding up Inference by Adaptive Length Reduction</i> Ali Modarressi, Hosein Mohebbi and Mohammad Taher Pilehvar	1
<i>Quantified Reproducibility Assessment of NLP Results</i> Anya Belz, Maja Popovic and Simon Mille	16
<i>Rare Tokens Degenerate All Tokens: Improving Neural Text Generation via Adaptive Gradient Gating for Rare Token Embeddings</i> Sangwon Yu, Jongyoon Song, Heeseung Kim, Seongmin Lee, Woo-Jong Ryu and Sungroh Yoon	29
<i>AlphBERT: Language Model Pre-training and Evaluation from Sub-Word to Sentence Level</i> Amit Seker, Elron Bandel, Dan Bareket, Idan Brusilovsky, Refael Shaked Greenfeld and Reut Tsarfaty	46
<i>Learning to Imagine: Integrating Counterfactual Thinking in Neural Discrete Reasoning</i> Moxin Li, Fuli Feng, Hanwang Zhang, Xiangnan He, Fengbin Zhu and Tat-Seng Chua	57
<i>Domain Adaptation in Multilingual and Multi-Domain Monolingual Settings for Complex Word Identification</i> George-Eduard Zaharia, Răzvan-Alexandru Smădu, Dumitru Clementin Cercel and Mihai Dasca-lu	70
<i>JointCL: A Joint Contrastive Learning Framework for Zero-Shot Stance Detection</i> Bin Liang, Qinglin Zhu, Xiang Li, Min Yang, Lin Gui, Yulan He and Ruifeng Xu	81
<i>[CASPI] Causal-aware Safe Policy Improvement for Task-oriented Dialogue</i> Govardana Sachithanandam Ramachandran, Kazuma Hashimoto and Caiming Xiong	92
<i>UniTranSeR: A Unified Transformer Semantic Representation Framework for Multimodal Task-Oriented Dialog System</i> Zhiyuan Ma, Jianjun Li, Guohui Li and Yongjing Cheng	103
<i>Dynamic Schema Graph Fusion Network for Multi-Domain Dialogue State Tracking</i> Yue Feng, Aldo Lipani, Fanghua Ye, Qiang Zhang and Emine Yilmaz	115
<i>Attention Temperature Matters in Abstractive Summarization Distillation</i> Shengqiang Zhang, Xingxing Zhang, Hangbo Bao and Furu Wei	127
<i>Towards Making the Most of Cross-Lingual Transfer for Zero-Shot Neural Machine Translation</i> Guanhua Chen, Shuming Ma, Yun Chen, Dongdong Zhang, Jia Pan, Wenping Wang and Furu Wei	142
<i>TopWORDS-Seg: Simultaneous Text Segmentation and Word Discovery for Open-Domain Chinese Texts via Bayesian Inference</i> Changzai Pan, Maosong Sun and Ke Deng	158
<i>An Unsupervised Multiple-Task and Multiple-Teacher Model for Cross-lingual Named Entity Recognition</i> Zhuoran Li, Chunming Hu, Xiaohui Guo, Junfan Chen, Wenyi Qin and Richong Zhang	170
<i>Discriminative Marginalized Probabilistic Neural Method for Multi-Document Summarization of Medical Literature</i> Gianluca Moro, Luca Ragazzi, Lorenzo Valgimigli and Davide Freddi	180

<i>Sparse Progressive Distillation: Resolving Overfitting under Pretrain-and-Finetune Paradigm</i>	
Shaoyi Huang, Dongkuan Xu, Ian En-Hsu Yen, Yijue Wang, Sung-En Chang, Bingbing Li, Shiyang Chen, Mimi Xie, Sanguthevar Rajasekaran, Hang Liu and Caiwen Ding	190
<i>CipherDAug: Ciphertext based Data Augmentation for Neural Machine Translation</i>	
Nishant Kambhatla, Logan Born and Anoop Sarkar	201
<i>Overlap-based Vocabulary Generation Improves Cross-lingual Transfer Among Related Languages</i>	
Vaidehi Patil, Partha Talukdar and Sunita Sarawagi	219
<i>Long-range Sequence Modeling with Predictable Sparse Attention</i>	
Yimeng Zhuang, Jing Zhang and Mei Tu	234
<i>Improving Personalized Explanation Generation through Visualization</i>	
Shijie Geng, Zuohui Fu, Yingqiang Ge, Lei Li, Gerard de Melo and Yongfeng Zhang	244
<i>New Intent Discovery with Pre-training and Contrastive Learning</i>	
Yuwei Zhang, Haode Zhang, Li-Ming Zhan, Xiao-Ming Wu and Albert Y.s. Lam	256
<i>Modeling U.S. State-Level Policies by Extracting Winners and Losers from Legislative Texts</i>	
Maryam Davoodi, Eric Waltenburg and Dan Goldwasser	270
<i>Structural Characterization for Dialogue Disentanglement</i>	
Xinbei Ma, Zhuosheng Zhang and Hai Zhao	285
<i>Multi-Party Empathetic Dialogue Generation: A New Task for Dialog Systems</i>	
Ling.Yu Zhu, Zhengkun Zhang, Jun Wang, Hongbin Wang, Haiying wu and Zhenglu Yang . .	298
<i>MISC: A Mixed Strategy-Aware Model integrating COMET for Emotional Support Conversation</i>	
Quan Tu, Yanran Li, Jianwei Cui, Bin Wang, Ji-Rong Wen and Rui Yan	308
<i>GLM: General Language Model Pretraining with Autoregressive Blank Infilling</i>	
Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang and Jie Tang . .	320
<i>QuoteR: A Benchmark of Quote Recommendation for Writing</i>	
Fanchao Qi, Yanhui Yang, Jing Yi, Zhili Cheng, Zhiyuan Liu and Maosong Sun	336
<i>Towards Comprehensive Patent Approval Predictions: Beyond Traditional Document Classification</i>	
Xiaochen Kev Gao, Zhaoyi Hou, Yifei Ning, Kewen Zhao, Beilei He, Jingbo Shang and Vish Krishnan	349
<i>Hypergraph Transformer: Weakly-Supervised Multi-hop Reasoning for Knowledge-based Visual Question Answering</i>	
Yu-Jung Heo, Eun-Sol Kim, Woo Suk Choi and Byoung-Tak Zhang	373
<i>Cross-Utterance Conditioned VAE for Non-Autoregressive Text-to-Speech</i>	
Yang Li, Cheng Yu, Guangzhi Sun, Hua Jiang, Fanglei Sun, Weiqin Zu, Ying Wen, Yang Yang and Jun Wang	391
<i>Mix and Match: Learning-free Controllable Text Generation using Energy Language Models</i>	
Fatemehsadat Mireshghallah, Kartik Goyal and Taylor Berg-Kirkpatrick	401
<i>So Different Yet So Alike! Constrained Unsupervised Text Style Transfer</i>	
Abhinav Ramesh Kashyap, Devamanyu Hazarika, Min-Yen Kan, Roger Zimmermann and Soujanya Poria	416
<i>e-CARE: a New Dataset for Exploring Explainable Causal Reasoning</i>	
Li Du, Xiao Ding, Kai Xiong, Ting Liu and Bing Qin	432

<i>Fantastic Questions and Where to Find Them: FairytaleQA – An Authentic Dataset for Narrative Comprehension</i>	
Ying Xu, Dakuo Wang, Mo Yu, Daniel Ritchie, Bingsheng Yao, Tongshuang Wu, Zheng Zhang, Toby Jia-Jun Li, Nora Bradford, Branda Sun, Tran Bao Hoang, Yisi Sang, Yufang Hou, Xiaojuan Ma, Diyi Yang, Nanyun Peng, Zhou Yu and Mark Warschauer	447
<i>KaFSP: Knowledge-Aware Fuzzy Semantic Parsing for Conversational Question Answering over a Large-Scale Knowledge Base</i>	
Junzhuo Li and Deyi Xiong	461
<i>Multilingual Knowledge Graph Completion with Self-Supervised Adaptive Graph Alignment</i>	
Zijie Huang, Zheng li, Haoming Jiang, Tianyu Cao, Hanqing Lu, Bing Yin, Karthik Subbian, Yizhou Sun and Wei Wang	474
<i>Modeling Hierarchical Syntax Structure with Triplet Position for Source Code Summarization</i>	
Juncai Guo, Jin Liu, Yao Wan, Li Li and Pingyi Zhou	486
<i>FewNLU: Benchmarking State-of-the-Art Methods for Few-Shot Natural Language Understanding</i>	
Yanan Zheng, Jing Zhou, Yujie Qian, Ming Ding, Chonghua Liao, Li Jian, Ruslan Salakhutdinov, Jie Tang, Sebastian Ruder and Zhilin Yang	501
<i>Learn to Adapt for Generalized Zero-Shot Text Classification</i>	
Yiwen Zhang, Caixia Yuan, Xiaojie Wang, Ziwei Bai and Yongbin Liu	517
<i>TableFormer: Robust Transformer Modeling for Table-Text Encoding</i>	
Jingfeng Yang, Aditya Gupta, Shyam Upadhyay, Luheng He, Rahul Goel and Shachi Paul	528
<i>Perceiving the World: Question-guided Reinforcement Learning for Text-based Games</i>	
Yunqiu Xu, Meng Fang, Ling Chen, Yali Du, Joey Tianyi Zhou and Chengqi Zhang	538
<i>Neural Label Search for Zero-Shot Multi-Lingual Extractive Summarization</i>	
Ruipeng Jia, Xingxing Zhang, Yanan Cao, Zheng Lin, Shi Wang and Furu Wei	561
<i>Few-Shot Class-Incremental Learning for Named Entity Recognition</i>	
Rui Wang, Tong Yu, Handong Zhao, Sungchul Kim, Subrata Mitra, Ruiyi Zhang and Ricardo Henao	571
<i>Improving Meta-learning for Low-resource Text Classification and Generation via Memory Imitation</i>	
Yingxiu Zhao, Zhiliang Tian, Huaxiu Yao, Yinhe Zheng, Dongkyu Lee, Yiping Song, Jian Sun and Nevin Zhang	583
<i>Quality Controlled Paraphrase Generation</i>	
Elron Bandel, Ranit Aharonov, Michal Shmueli-Scheuer, Ilya Shnayderman, Noam Slonim and Liat Ein-Dor	596
<i>Controllable Dictionary Example Generation: Generating Example Sentences for Specific Targeted Audiences</i>	
Xingwei He and Siu Ming Yiu	610
<i>AraT5: Text-to-Text Transformers for Arabic Language Generation</i>	
El Moatez Billah Nagoudi, AbdelRahim A. Elmadany and Muhammad Abdul-Mageed	628
<i>Legal Judgment Prediction via Event Extraction with Constraints</i>	
Yi Feng, Chuanyi Li and Vincent Ng	648
<i>Answer-level Calibration for Free-form Multiple Choice Question Answering</i>	
Sawan Kumar	665

<i>Learning When to Translate for Streaming Speech</i>	
Qian Qian Dong, Yaoming Zhu, Mingxuan Wang and Lei Li	680
<i>Compact Token Representations with Contextual Quantization for Efficient Document Re-ranking</i>	
Yingrui Yang, Yifan Qiao and Tao Yang	695
<i>Early Stopping Based on Unlabeled Samples in Text Classification</i>	
HongSeok Choi, Dongha Choi and Hyunju Lee	708
<i>Meta-learning via Language Model In-context Tuning</i>	
Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis and He He	719
<i>It is AI's Turn to Ask Humans a Question: Question-Answer Pair Generation for Children's Story Books</i>	
Bingsheng Yao, Dakuo Wang, Tongshuang Wu, Zheng Zhang, Toby Jia-Jun Li, Mo Yu and Ying Xu	731
<i>Prompt-Based Rule Discovery and Boosting for Interactive Weakly-Supervised Learning</i>	
Rongzhi Zhang, Yue Yu, Pranav Shetty, Le Song and Chao Zhang	745
<i>Constrained Multi-Task Learning for Bridging Resolution</i>	
Hideo Kobayashi, Yufang Hou and Vincent Ng	759
<i>DEAM: Dialogue Coherence Evaluation using AMR-based Semantic Manipulations</i>	
Sarik Ghazarian, Nuan Wen, Aram Galstyan and Nanyun Peng	771
<i>HIBRIDS: Attention with Hierarchical Biases for Structure-aware Long Document Summarization</i>	
Shuyang Cao and Lu Wang	786
<i>De-Bias for Generative Extraction in Unified NER Task</i>	
Shuai Zhang, Yongliang Shen, Zeqi Tan, Yiquan Wu and Weiming Lu	808
<i>An Information-theoretic Approach to Prompt Engineering Without Ground Truth Labels</i>	
Taylor Sorensen, Joshua Robinson, Christopher Michael Rytting, Alexander Glenn Shaw, Kyle Jeffrey Rogers, Alexia Pauline Delorey, Mahmoud Khalil, Nancy Fulda and David Wingate	819
<i>Expanding Pretrained Models to Thousands More Languages via Lexicon-based Adaptation</i>	
Xinyi Wang, Sebastian Ruder and Graham Neubig	863
<i>Language-agnostic BERT Sentence Embedding</i>	
Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan and Wei Wang	878
<i>Nested Named Entity Recognition with Span-level Graphs</i>	
Juncheng Wan, Dongyu Ru, Weinan Zhang and Yong Yu	892
<i>CogTaskonomy: Cognitively Inspired Task Taxonomy Is Beneficial to Transfer Learning in NLP</i>	
Yifei Luo, Minghui Xu and Deyi Xiong	904
<i>RoCBert: Robust Chinese Bert with Multimodal Contrastive Pretraining</i>	
Hui Su, Weiwei Shi, Xiaoyu Shen, Zhou Xiao, Tuo Ji, Jiarui Fang and Jie Zhou	921
<i>Premise-based Multimodal Reasoning: Conditional Inference on Joint Textual and Visual Clues</i>	
Qingxiu Dong, Ziwei Qin, Heming Xia, Tian Feng, Shoujie Tong, Haoran Meng, Lin Xu, Zhongyu Wei, Weidong Zhan, Baobao Chang, Sujian Li, Tianyu Liu and Zhifang Sui	932
<i>Parallel Instance Query Network for Named Entity Recognition</i>	
Yongliang Shen, Xiaobin Wang, Zeqi Tan, Guangwei Xu, Pengjun Xie, Fei Huang, Weiming Lu and Yueting Zhuang	947

<i>ProphetChat: Enhancing Dialogue Generation with Simulation of Future Conversation</i>	
Chang Liu, Xu Tan, Chongyang Tao, Zhenxin Fu, Dongyan Zhao, Tie-Yan Liu and Rui Yan .	962
<i>Modeling Multi-hop Question Answering as Single Sequence Prediction</i>	
Semih Yavuz, Kazuma Hashimoto, Yingbo Zhou, Nitish Shirish Keskar and Caiming Xiong .	974
<i>Learning Disentangled Semantic Representations for Zero-Shot Cross-Lingual Transfer in Multilingual Machine Reading Comprehension</i>	
Linjuan Wu, Shaojuan Wu, Xiaowang Zhang, Deyi Xiong, Shizhan Chen, Zhiqiang Zhuang and Zhiyong Feng .	991
<i>Multi-Granularity Structural Knowledge Distillation for Language Model Compression</i>	
Chang Liu, Chongyang Tao, Jiazhan Feng and Dongyan Zhao .	1001
<i>Auto-Debias: Debiasing Masked Language Models with Automated Biased Prompts</i>	
Yue Guo, Yi Yang and Ahmed Abbasi .	1012
<i>Where to Go for the Holidays: Towards Mixed-Type Dialogs for Clarification of User Goals</i>	
Zeming Liu, Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu and Hua Wu .	1024
<i>Semi-supervised Domain Adaptation for Dependency Parsing with Dynamic Matching Network</i>	
Ying Li, Shuaike Li and Min Zhang .	1035
<i>A Closer Look at How Fine-tuning Changes BERT</i>	
Yichu Zhou and Vivek Srikumar .	1046
<i>Sentence-aware Contrastive Learning for Open-Domain Passage Retrieval</i>	
Wu Bo Hong, Zhuosheng Zhang, Jinyuan Wang and Hai Zhao .	1062
<i>FaiRR: Faithful and Robust Deductive Reasoning over Natural Language</i>	
Soumya Sanyal, Harman Singh and Xiang Ren .	1075
<i>HiTab: A Hierarchical Table Dataset for Question Answering and Natural Language Generation</i>	
Zhoujun Cheng, Haoyu Dong, Zhiruo Wang, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou and Dongmei Zhang .	1094
<i>Doctor Recommendation in Online Health Forums via Expertise Learning</i>	
Xiaoxin Lu, Yubo Zhang, Jing Li and Shi Zong .	1111
<i>Continual Prompt Tuning for Dialog State Tracking</i>	
Qi Zhu, Bing Li, Fei Mi, Xiaoyan Zhu and Minlie Huang .	1124
<i>There's a Time and Place for Reasoning Beyond the Image</i>	
Xingyu Fu, Ben Zhou, Ishaan Preetam Chandratreya, Carl Vondrick and Dan Roth .	1138
<i>FORTAP: Using Formulas for Numerical-Reasoning-Aware Table Pretraining</i>	
Zhoujun Cheng, Haoyu Dong, Ran Jia, Pengfei wu, Shi Han, Fan Cheng and Dongmei Zhang	1150
<i>Multimodal fusion via cortical network inspired losses</i>	
Shiv Shankar .	1167
<i>Modeling Temporal-Modal Entity Graph for Procedural Multimodal Machine Comprehension</i>	
Huibin Zhang, Zhengkun Zhang, Yao Zhang, Jun Wang, Yufan Li, Ning Jiang, Xin Wei and Zhenglu Yang .	1179

<i>Explanation Graph Generation via Pre-trained Language Models: An Empirical Study with Contrastive Learning</i>	
Swarnadeep Saha, Prateek Yadav and Mohit Bansal	1190
<i>Unsupervised Extractive Opinion Summarization Using Sparse Coding</i>	
Somnath Basu Roy Chowdhury, Chao Zhao and Snigdha Chaturvedi	1209
<i>LexSubCon: Integrating Knowledge from Lexical Resources into Contextual Embeddings for Lexical Substitution</i>	
George Michalopoulos, Ian McKillop, Alexander Wong and Helen Hong Chen	1226
<i>Think Before You Speak: Explicitly Generating Implicit Commonsense Knowledge for Response Generation</i>	
Pei Zhou, Karthik Gopalakrishnan, Behnam Hedayatnia, Seokhwan Kim, Jay Pujara, Xiang Ren, Yang Liu and Dilek Hakkani-Tur	1237
<i>Flow-Adapter Architecture for Unsupervised Machine Translation</i>	
Yihong Liu, Haris Jabbar and Hinrich Schuetze	1253
<i>Efficient Unsupervised Sentence Compression by Fine-tuning Transformers with Reinforcement Learning</i>	
Demian Gholipour Ghalandari, Chris Hokamp and Georgiana Ifrim	1267
<i>Tracing Origins: Coreference-aware Machine Reading Comprehension</i>	
Zhuosheng Zhang and Hai Zhao	1281
<i>WatClaimCheck: A new Dataset for Claim Entailment and Inference</i>	
Kashif Khan, Ruizhe Wang and Pascal Poupart	1293
<i>FrugalScore: Learning Cheaper, Lighter and Faster Evaluation Metrics for Automatic Text Generation</i>	
Moussa Kamal Eddine, Guokan Shang, Antoine J.-p. Tixier and Michalis Vazirgiannis	1305
<i>A Well-Composed Text is Half Done! Composition Sampling for Diverse Conditional Generation</i>	
Shashi Narayan, Gonçalo Simões, Yao Zhao, Joshua Maynez, Dipanjan Das, Michael Collins and Mirella Lapata	1319
<i>Synthetic Question Value Estimation for Domain Adaptation of Question Answering</i>	
Xiang Yue, Ziyu Yao and Huan Sun	1340
<i>Better Language Model with Hypernym Class Prediction</i>	
He Bai, Tong Wang, Alessandro Sordoni and Peng Shi	1352
<i>Tackling Fake News Detection by Continually Improving Social Context Representations using Graph Neural Networks</i>	
Nikhil Mehta, Maria Leonor Pacheco and Dan Goldwasser	1363
<i>Understanding Gender Bias in Knowledge Base Embeddings</i>	
Yupei Du, Qi Zheng, Yuanbin Wu, Man Lan, Yan Yang and Meirong Ma	1381
<i>Computational Historical Linguistics and Language Diversity in South Asia</i>	
Aryaman Arora, Adam Farris, Samopriya Basu and Suresh Kolichala	1396
<i>Faithful or Extractive? On Mitigating the Faithfulness-Abstractiveness Trade-off in Abstractive Summarization</i>	
Faisal Ladhak, Esin Durmus, He He, Claire Cardie and Kathleen McKeown	1410
<i>Slangvolution: A Causal Analysis of Semantic Change and Frequency Dynamics in Slang</i>	
Daphna Keidar, Andreas Opedal, Zhijing Jin and Mrinmaya Sachan	1422

<i>Spurious Correlations in Reference-Free Evaluation of Text Generation</i>	
Esin Durmus, Faisal Ladhak and Tatsunori Hashimoto	1443
<i>On The Ingredients of an Effective Zero-shot Semantic Parser</i>	
Pengcheng Yin, John Frederick Wieting, Avirup Sil and Graham Neubig	1455
<i>Bias Mitigation in Machine Translation Quality Estimation</i>	
Hanna Behnke, Marina Fomicheva and Lucia Specia	1475
<i>Unified Speech-Text Pre-training for Speech Translation and Recognition</i>	
Yun Tang, Hongyu Gong, Ning Dong, Changhan Wang, Wei-Ning Hsu, Jiatao Gu, Alexei Baevski, Xian Li, Abdelrahman Mohamed, Michael Auli and Juan Pino	1488
<i>Match the Script, Adapt if Multilingual: Analyzing the Effect of Multilingual Pretraining on Cross-lingual Transferability</i>	
Yoshinari Fujinuma, Jordan Lee Boyd-Graber and Katharina Kann	1500
<i>Structured Pruning Learns Compact and Accurate Models</i>	
Mengzhou Xia, Zexuan Zhong and Danqi Chen	1513
<i>How can NLP Help Revitalize Endangered Languages? A Case Study and Roadmap for the Cherokee Language</i>	
Shiyue Zhang, Ben Frey and Mohit Bansal	1529
<i>Differentiable Multi-Agent Actor-Critic for Multi-Step Radiology Report Summarization</i>	
Sanjeev Kumar Karn, Ning Liu, Hinrich Schuetze and Oladimeji Farri	1542
<i>Online Semantic Parsing for Latency Reduction in Task-Oriented Dialogue</i>	
Jiawei Zhou, Jason Eisner, Michael Newman, Emmanouil Antonios Platanios and Sam Thomson	1554
<i>Few-Shot Tabular Data Enrichment Using Fine-Tuned Transformer Architectures</i>	
Asaf Harari and Gilad Katz	1577
<i>Summ^N: A Multi-Stage Summarization Framework for Long Input Dialogues and Documents</i>	
Yusen Zhang, Ansong Ni, Ziming Mao, Chen Henry Wu, Chenguang Zhu, Budhaditya Deb, Ahmed Hassan Awadallah, Dragomir Radev and Rui Zhang	1592
<i>Open Domain Question Answering with A Unified Knowledge Interface</i>	
Kaixin Ma, Hao Cheng, Xiaodong Liu, Eric Nyberg and Jianfeng Gao	1605
<i>Principled Paraphrase Generation with Parallel Corpora</i>	
Aitor Ormazabal, Mikel Artetxe, Aitor Soroa, Gorka Labaka and Eneko Agirre	1621
<i>GlobalWoZ: Globalizing MultiWoZ to Develop Multilingual Task-Oriented Dialogue Systems</i>	
Bosheng Ding, Junjie Hu, Lidong Bing, Mahani Aljunied, Shafiq Joty, Luo Si and Chunyan Miao	1639
<i>Domain Knowledge Transferring for Pre-trained Language Model via Calibrated Activation Boundary Distillation</i>	
Dongha Choi, HongSeok Choi and Hyunju Lee	1658
<i>Retrieval-guided Counterfactual Generation for QA</i>	
Bhargavi Paranjape, Matthew Lamm and Ian Tenney	1670
<i>DYLE: Dynamic Latent Extraction for Abstractive Long-Input Summarization</i>	
Ziming Mao, Chen Henry Wu, Ansong Ni, Yusen Zhang, Rui Zhang, Tao Yu, Budhaditya Deb, Chenguang Zhu, Ahmed Hassan Awadallah and Dragomir Radev	1687

<i>Searching for fingerspelled content in American Sign Language</i> Bowen Shi, Diane Brentari, Greg Shakhnarovich and Karen Livescu	1699
<i>Skill Induction and Planning with Latent Language</i> Pratyusha Sharma, Antonio Torralba and Jacob Andreas	1713
<i>Fully-Semantic Parsing and Generation: the BabelNet Meaning Representation</i> Abelardo Carlos Martinez Lorenzo, Marco Maru and Roberto Navigli	1727
<i>Leveraging Similar Users for Personalized Language Modeling with Limited Data</i> Charles Welch, Chenxi Gu, Jonathan K Kummerfeld, Veronica Perez-Rosas and Rada Mihalcea 1742	
<i>DEEP: DEnoising Entity Pre-training for Neural Machine Translation</i> Junjie Hu, Hiroaki Hayashi, Kyunghyun Cho and Graham Neubig	1753
<i>Multi-Modal Sarcasm Detection via Cross-Modal Graph Convolutional Network</i> Bin Liang, Chenwei Lou, Xiang Li, Min Yang, Lin Gui, Yulan He, Wenjie Pei and Ruifeng Xu 1767	
<i>Composable Sparse Fine-Tuning for Cross-Lingual Transfer</i> Alan Ansell, Edoardo Ponti, Anna Korhonen and Ivan Vulić	1778
<i>Toward Annotator Group Bias in Crowdsourcing</i> Haochen Liu, Joseph Davis Thekinen, Sinem Mollaoglu, Da Tang, Ji Yang, Youlong Cheng, Hui Liu and Jiliang Tang	1797
<i>Under the Morphosyntactic Lens: A Multifaceted Evaluation of Gender Bias in Speech Translation</i> Beatrice Savoldi, Marco Gaido, Luisa Bentivogli, Matteo Negri and Marco Turchi	1807
<i>Answering Open-Domain Multi-Answer Questions via a Recall-then-Verify Framework</i> Zhihong Shao and Minlie Huang	1825
<i>Probing as Quantifying Inductive Bias</i> Alexander Immer, Lucas Torroba Hennigen, Vincent Fortuin and Ryan D Cotterell	1839
<i>Probing Structured Pruning on Multilingual Pre-trained Models: Settings, Algorithms, and Efficiency</i> Yanyang Li, Fuli Luo, Runxin Xu, Songfang Huang, Fei Huang and Liwei Wang	1852
<i>GPT-D: Inducing Dementia-related Linguistic Anomalies by Deliberate Degradation of Artificial Neu- ral Language Models</i> Changye Li, David Knopman, Weizhe Xu, Trevor Cohen and Serguei V. S. Pakhomov	1866
<i>An Empirical Survey of the Effectiveness of Debiasing Techniques for Pre-trained Language Models</i> Nicholas Meade, Elinor Poole-Dayana and Siva Reddy	1878
<i>Exploring and Adapting Chinese GPT to Pinyin Input Method</i> Minghuan Tan, Yong Dai, Duyu Tang, Zhangyin Feng, Guoping Huang, Jing Jiang, Jiwei Li and Shuming Shi	1899
<i>Enhancing Cross-lingual Natural Language Inference by Prompt-learning from Cross-lingual Templa- tes</i> Kunxun Qi, Hai Wan, Jianfeng Du and Haolan Chen	1910
<i>Sense Embeddings are also Biased – Evaluating Social Biases in Static and Contextualised Sense Em- beddings</i> Yi Zhou, Masahiro Kaneko and Danushka Bollegala	1924

<i>Hybrid Semantics for Goal-Directed Natural Language Generation</i> Connor Baumler and Soumya Ray	1936
<i>Predicting Intervention Approval in Clinical Trials through Multi-Document Summarization</i> Georgios Katsimpras and Georgios Paliouras	1947
<i>BiTIIMT: A Bilingual Text-infilling Method for Interactive Machine Translation</i> Yanling Xiao, Lemao Liu, Guoping Huang, Qu Cui, Shujian Huang, Shuming Shi and Jiajun Chen	1958
<i>Distributionally Robust Finetuning BERT for Covariate Drift in Spoken Language Understanding</i> Samuel Broscheit, Quynh Do and Judith Gaspers	1970
<i>Enhancing Chinese Pre-trained Language Model via Heterogeneous Linguistics Graph</i> Yanzeng Li, Jiangxia Cao, Xin Cong, Zhenyu Zhang, Bowen Yu, Hongsong Zhu and Tingwen Liu	1986
<i>Divide and Denoise: Learning from Noisy Labels in Fine-Grained Entity Typing with Cluster-Wise Loss Correction</i> Kunyu Pang, Haoyu Zhang, Jie Zhou and Ting Wang	1997
<i>Towards Robustness of Text-to-SQL Models Against Natural and Realistic Adversarial Table Perturbation</i> Xinyu Pi, Bing Wang, Yan Gao, Jiaqi Guo, Zhoujun Li and Jian-Guang Lou	2007
<i>Overcoming Catastrophic Forgetting beyond Continual Learning: Balanced Training for Neural Machine Translation</i> Chenze Shao and Yang Feng	2023
<i>Metaphors in Pre-Trained Language Models: Probing and Generalization Across Datasets and Languages</i> Ehsan Aghazadeh, Mohsen Fayyaz and Yadollah Yaghoobzadeh	2037
<i>Discrete Opinion Tree Induction for Aspect-based Sentiment Analysis</i> Chenhua Chen, Zhiyang Teng, Zhongqing Wang and Yue Zhang	2051
<i>Investigating Non-local Features for Neural Constituency Parsing</i> Leyang Cui, Sen Yang and Yue Zhang	2065
<i>Learning from Sibling Mentions with Scalable Graph Inference in Fine-Grained Entity Typing</i> Yi Chen, Jiayang Cheng, Haiyun Jiang, Lemao Liu, Haisong Zhang, Shuming Shi and Ruifeng Xu	2076
<i>A Variational Hierarchical Model for Neural Cross-Lingual Summarization</i> Yunlong Liang, Fandong Meng, Chulun Zhou, Jinan Xu, Yufeng Chen, Jinsong Su and Jie Zhou	2088
<i>On the Robustness of Question Rewriting Systems to Questions of Varying Hardness</i> Hai Ye, Hwee Tou Ng and Wenjuan Han	2100
<i>OpenHands: Making Sign Language Recognition Accessible with Pose-based Pretrained Models across Languages</i> Prem Selvaraj, Gokul NC, Pratyush Kumar and Mitesh M Khapra	2114
<i>bert2BERT: Towards Reusable Pretrained Language Models</i> Cheng Chen, Yichun Yin, Lifeng Shang, Xin Jiang, Yujia Qin, Fengyu Wang, Zhi Wang, Xiao Chen, Zhiyuan Liu and Qun Liu	2134

<i>Vision-Language Pre-Training for Multimodal Aspect-Based Sentiment Analysis</i>	
Yan Ling, Jianfei Yu and Rui Xia	2149
<i>You might think about slightly revising the title: Identifying Hedges in Peer-tutoring Interactions</i>	
Yann Raphalen, Chloé Clavel and Justine Cassell	2160
<i>Efficient Cluster-Based k-Nearest-Neighbor Machine Translation</i>	
Dexin Wang, Kai Fan, Boxing Chen and Deyi Xiong	2175
<i>Headed-Span-Based Projective Dependency Parsing</i>	
Songlin Yang and Kewei Tu	2188
<i>Decoding Part-of-Speech from Human EEG Signals</i>	
Alex Graeme Murphy, Bernd Bohnet, Ryan McDonald and Uta Noppeneay	2201
<i>Robust Lottery Tickets for Pre-trained Language Models</i>	
Rui Zheng, Bao Rong, Yuhao Zhou, Di Liang, Sirui Wang, Wei Wu, Tao Gui, Qi Zhang and Xuanjing Huang	2211
<i>Knowledgeable Prompt-tuning: Incorporating Knowledge into Prompt Verbalizer for Text Classification</i>	
Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu and Maosong Sun	2225
<i>Cross-Lingual Contrastive Learning for Fine-Grained Entity Typing for Low-Resource Languages</i>	
Xu Han, Yuqi Luo, Weize Chen, Zhiyuan Liu, Maosong Sun, Zhou Botong, Hao Fei and Suncong Zheng	2241
<i>MELM: Data Augmentation with Masked Entity Language Modeling for Low-Resource NER</i>	
Ran Zhou, Xin Li, Ruidan He, Lidong Bing, Erik Cambria, Luo Si and Chunyan Miao	2251
<i>Word2Box: Capturing Set-Theoretic Semantics of Words using Box Embeddings</i>	
Shib Sankar Dasgupta, Michael Boratko, Siddhartha Mishra, Shriya Atmakuri, Dhruvesh Patel, Xiang Lorraine Li and Andrew McCallum	2263
<i>IAM: A Comprehensive and Large-Scale Dataset for Integrated Argument Mining Tasks</i>	
Liyang Cheng, Lidong Bing, Ruidan He, Qian Yu, Yan Zhang and Luo Si	2277
<i>PLANET: Dynamic Content Planning in Autoregressive Transformers for Long-form Text Generation</i>	
Zhe Hu, Hou Pong Chan, Jiachen Liu, Xinyan Xiao, Hua Wu and Lifu Huang	2288
<i>CTRLEval: An Unsupervised Reference-Free Metric for Evaluating Controlled Text Generation</i>	
Pei Ke, Hao Zhou, Yankai Lin, Peng Li, Jie Zhou, Xiaoyan Zhu and Minlie Huang	2306
<i>Beyond the Granularity: Multi-Perspective Dialogue Collaborative Selection for Dialogue State Tracking</i>	
Jinyu Guo, Kai Shuang, Jijie Li, Zihan Wang and Yixuan Liu	2320
<i>Are Prompt-based Models Clueless?</i>	
Pride Kavumba, Ryo Takahashi and Yusuke Oda	2333
<i>Learning Confidence for Transformer-based Neural Machine Translation</i>	
Yu Lu, Jiali Zeng, Jiajun Zhang, Shuangzhi Wu and Mu Li	2353
<i>Things not Written in Text: Exploring Spatial Commonsense from Visual Signals</i>	
Xiao Liu, Da Yin, Yansong Feng and Dongyan Zhao	2365

<i>Conditional Bilingual Mutual Information Based Adaptive Training for Neural Machine Translation</i> Songming Zhang, Yijin Liu, Fandong Meng, Yufeng Chen, Jinan Xu, Jian Liu and Jie Zhou	2377
<i>ClusterFormer: Neural Clustering Attention for Efficient and Effective Transformer</i> Ningning Wang, Guobing Gan, Peng Zhang, Shuai Zhang, Junqiu Wei, Qun Liu and Xin Jiang	2390
<i>Bottom-Up Constituency Parsing and Nested Named Entity Recognition with Pointer Networks</i> Songlin Yang and Kewei Tu	2403
<i>Redistributing Low-Frequency Words: Making the Most of Monolingual Data in Non-Autoregressive Translation</i> Liang Ding, Longyue Wang, Shuming Shi, Dacheng Tao and Zhaopeng Tu	2417
<i>Dependency Parsing as MRC-based Span-Span Prediction</i> Leilei Gan, Yuxian Meng, Kun Kuang, Xiaofei Sun, Chun Fan, Fei Wu and Jiwei Li	2427
<i>Adversarial Soft Prompt Tuning for Cross-Domain Sentiment Analysis</i> Hui Wu and Xiaodong Shi	2438
<i>Generating Scientific Claims for Zero-Shot Scientific Fact Checking</i> Dustin Wright, David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Isabelle Augenstein and Lucy Lu Wang	2448
<i>Modeling Dual Read/Write Paths for Simultaneous Machine Translation</i> Shaolei Zhang and Yang Feng	2461
<i>ExtEnD: Extractive Entity Disambiguation</i> Edoardo Barba, Luigi Procopio and Roberto Navigli	2478
<i>Hierarchical Sketch Induction for Paraphrase Generation</i> Tom Hosking, Hao Tang and Mirella Lapata	2489
<i>Alignment-Augmented Consistent Translation for Multilingual Open Information Extraction</i> Keshav Kolluru, Muqeeth Mohammed, Shubham Mittal, Soumen Chakrabarti and Mausam	2502
<i>Text-to-Table: A New Way of Information Extraction</i> Xueqing Wu, Jiacheng Zhang and Hang Li	2518
<i>Accelerating Code Search with Deep Hashing and Code Classification</i> Wenchao Gu, Yanlin Wang, Lun Du, Hongyu Zhang, Shi Han, Dongmei Zhang and Michael Lyu	2534
<i>Other Roles Matter! Enhancing Role-Oriented Dialogue Summarization via Role Interactions</i> Haitao Lin, Junnan Zhu, Lu Xiang, Yu Zhou, Jiajun Zhang and Chengqing Zong	2545
<i>ClarET: Pre-training a Correlation-Aware Context-To-Event Transformer for Event-Centric Generation and Classification</i> Yucheng Zhou, Tao Shen, Xiubo Geng, Guodong Long and Daxin Jiang	2559
<i>Measuring and Mitigating Name Biases in Neural Machine Translation</i> Jun Wang, Benjamin I. P. Rubinstein and Trevor Cohn	2576
<i>Understanding and Improving Sequence-to-Sequence Pretraining for Neural Machine Translation</i> Wenxuan Wang, Wenxiang Jiao, Yongchang Hao, Xing Wang, Shuming Shi, Zhaopeng Tu and Michael Lyu	2591

<i>MSCTD: A Multimodal Sentiment Chat Translation Dataset</i>	
Yunlong Liang, Fandong Meng, Jinan Xu, Yufeng Chen and Jie Zhou	2601
<i>Learning Disentangled Textual Representations via Statistical Measures of Similarity</i>	
Pierre Colombo, Guillaume Staerman, Nathan Noiry and Pablo Piantanida	2614
<i>On the Sensitivity and Stability of Model Interpretations in NLP</i>	
Fan Yin, Zhouxing Shi, Cho-Jui Hsieh and Kai-Wei Chang	2631
<i>Down and Across: Introducing Crossword-Solving as a New NLP Benchmark</i>	
Saurabh Kulshreshtha, Olga Kovaleva, Namrata Shivgunde and Anna Rumshisky	2648
<i>Generating Data to Mitigate Spurious Correlations in Natural Language Inference Datasets</i>	
Yuxiang Wu, Matt Gardner, Pontus Stenetorp and Pradeep Dasigi	2660
<i>GL-CLeF: A Global-Local Contrastive Learning Framework for Cross-lingual Spoken Language Understanding</i>	
Libo Qin, Qiguang Chen, Tianbao Xie, Qixin Li, Jian-Guang Lou, Wanxiang Che and Min-Yen Kan	2677
<i>Good Examples Make A Faster Learner: Simple Demonstration-based Learning for Low-resource NER</i>	
Dong-Ho Lee, Akshen Kadakia, Kangmin Tan, Mahak Agarwal, Xinyu Feng, Takashi Shibuya, Ryosuke Mitani, Toshiyuki Sekiya, Jay Pujara and Xiang Ren	2687
<i>Contextual Representation Learning beyond Masked Language Modeling</i>	
Zhiyi Fu, Wangchunshu Zhou, Jingjing Xu, Hao Zhou and Lei Li	2701
<i>Efficient Hyper-parameter Search for Knowledge Graph Embedding</i>	
Yongqi Zhang, Zhanke Zhou, Quanming Yao and Yong Li	2715
<i>A Meta-framework for Spatiotemporal Quantity Extraction from Text</i>	
Qiang Ning, Ben Zhou, Hao Wu, Haoruo Peng, Chuchu Fan and Matt Gardner	2736
<i>Leveraging Visual Knowledge in Language Tasks: An Empirical Study on Intermediate Pre-training for Cross-Modal Knowledge Transfer</i>	
Woojeong Jin, Dong-Ho Lee, Chenguang Zhu, Jay Pujara and Xiang Ren	2750
<i>A Good Prompt Is Worth Millions of Parameters: Low-resource Prompt-based Learning for Vision-Language Models</i>	
Woojeong Jin, Yu Cheng, Yelong Shen, Weizhu Chen and Xiang Ren	2763
<i>Continual Few-shot Relation Learning via Embedding Space Regularization and Data Augmentation</i>	
Chengwei Qin and Shafiq Joty	2776
<i>Variational Graph Autoencoding as Cheap Supervision for AMR Coreference Resolution</i>	
Irene Li, Linfeng Song, Kun Xu and Dong Yu	2790
<i>Identifying Chinese Opinion Expressions with Extremely-Noisy Crowdsourcing Annotations</i>	
Xin Zhang, Guangwei Xu, Yueheng Sun, Meishan Zhang, Xiaobin Wang and Min Zhang	2801
<i>Sequence-to-Sequence Knowledge Graph Completion and Question Answering</i>	
Apoorv Umang Saxena, Adrian Kochsiek and Rainer Gemulla	2814
<i>Learning to Mediate Disparities Towards Pragmatic Communication</i>	
Yuwei Bao, Sayan Ghosh and Joyce Chai	2829
<i>Unsupervised Corpus Aware Language Model Pre-training for Dense Passage Retrieval</i>	
Luyu Gao and Jamie Callan	2843

<i>Multimodal Dialogue Response Generation</i>	
Qingfeng Sun, Yujing Wang, Can Xu, Kai Zheng, Yaming Yang, Huang Hu, Fei Xu, Jessica Zhang, Xiubo Geng and Daxin Jiang	2854
<i>CAKE: A Scalable Commonsense-Aware Framework For Multi-View Knowledge Graph Completion</i>	
Guanglin Niu, Bo Li, Yongfei Zhang and Shiliang Pu	2867
<i>Confidence Based Bidirectional Global Context Aware Training Framework for Neural Machine Translation</i>	
Chulun Zhou, Fandong Meng, Jie Zhou, Min Zhang, Hongji Wang and Jinsong Su	2878
<i>BRIO: Bringing Order to Abstractive Summarization</i>	
Yixin Liu, Pengfei Liu, Dragomir Radev and Graham Neubig	2890
<i>Leveraging Relaxed Equilibrium by Lazy Transition for Sequence Modeling</i>	
Xi Ai and Bin Fang	2904
<i>FIBER: Fill-in-the-Blanks as a Challenging Video Understanding Evaluation Framework</i>	
Santiago Castro, Ruoyao Wang, Pingxuan Huang, Ian Stewart, Oana Ignat, Nan Liu, Jonathan C. Stroud and Rada Mihalcea	2925
<i>KenMeSH: Knowledge-enhanced End-to-end Biomedical Text Labelling</i>	
Xindi Wang, Robert Mercer and Frank Rudzicz	2941
<i>A Taxonomy of Empathetic Questions in Social Dialogs</i>	
Ekaterina Svikhnushina, Iuliana Voinea, Anuradha Welivita and Pearl Pu	2952
<i>Enhanced Multi-Channel Graph Convolutional Network for Aspect Sentiment Triplet Extraction</i>	
Hao Chen, Zepeng Zhai, Fangxiang Feng, Ruifan Li and Xiaojie Wang	2974
<i>ProtoTEX: Explaining Model Decisions with Prototype Tensors</i>	
Anubrata Das, Chitrang Gupta, Venelin Kovatchev, Matthew Lease and Junyi Jessy Li	2986
<i>Show Me More Details: Discovering Hierarchies of Procedures from Semi-structured Web Data</i>	
Shuyan Zhou, Li Zhang, Yue Yang, Qing Lyu, Pengcheng Yin, Chris Callison-Burch and Graham Neubig	2998
<i>Cross-Modal Discrete Representation Learning</i>	
Alexander H. Liu, SouYoung Jin, Cheng-I Lai, Andrew Rouditchenko, Aude Oliva and James R. Glass	3013
<i>Improving Event Representation via Simultaneous Weakly Supervised Contrastive Learning and Clustering</i>	
Jun Gao, Wei Wang, Changlong Yu, Huan Zhao, Wilfred Ng and Ruifeng Xu	3036
<i>Contrastive Visual Semantic Pretraining Magnifies the Semantics of Natural Language Representations</i>	
Robert Wolfe and Aylin Caliskan	3050
<i>ConTinTin: Continual Learning from Task Instructions</i>	
Wenpeng Yin, Jia Li and Caiming Xiong	3062
<i>Automated Crossword Solving</i>	
Eric Wallace, Nicholas Tomlin, Albert Xu, Kevin Yang, Eshaan Pathak, Matthew L. Ginsberg and Dan Klein	3073
<i>Learned Incremental Representations for Parsing</i>	
Nikita Kitaev, Thomas Lu and Dan Klein	3086

<i>Knowledge Enhanced Reflection Generation for Counseling Dialogues</i>	
Siqi Shen, Veronica Perez-Rosas, Charles Welch, Soujanya Poria and Rada Mihalcea	3096
<i>Misinfo Reaction Frames: Reasoning about Readers' Reactions to News Headlines</i>	
Saadia Gabriel, Skyler Hallinan, Maarten Sap, Pemi Nguyen, Franziska Roesner, Eunsol Choi and Yejin Choi	3108
<i>On Continual Model Refinement in Out-of-Distribution Data Streams</i>	
Bill Yuchen Lin, Sida Wang, Xi Victoria Lin, Robin Jia, Lin Xiao, Xiang Ren and Scott Yih	3128
<i>Achieving Conversational Goals with Unsupervised Post-hoc Knowledge Injection</i>	
Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick and Julian McAuley	3140
<i>Generated Knowledge Prompting for Commonsense Reasoning</i>	
Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi and Hannaneh Hajishirzi	3154
<i>Training Data is More Valuable than You Think: A Simple and Effective Method by Retrieving from Training Data</i>	
Shuohang Wang, Yichong Xu, Yuwei Fang, Yang Liu, Siqi Sun, Ruochen Xu, Chenguang Zhu and Michael Zeng	3170
<i>Life after BERT: What do Other Muppets Understand about Language?</i>	
Vladislav Lialin, Kevin Zhao, Namrata Shivagunde and Anna Rumshisky	3180
<i>Tailor: Generating and Perturbing Text with Semantic Controls</i>	
Alexis Ross, Tongshuang Wu, Hao Peng, Matthew E Peters and Matt Gardner	3194
<i>TruthfulQA: Measuring How Models Mimic Human Falsehoods</i>	
Stephanie Lin, Jacob Hilton and Owain Evans	3214
<i>Adaptive Testing and Debugging of NLP Models</i>	
Marco Tulio Ribeiro and Scott M Lundberg	3253
<i>Right for the Right Reason: Evidence Extraction for Trustworthy Tabular Reasoning</i>	
Vivek Gupta, Shuo Zhang, Alakananda Vempala, Yujie He, Temma Choji and Vivek Srikumar	3268
<i>Interactive Word Completion for Plains Cree</i>	
William Abbott Lane, Atticus Galvin Harrigan and Antti Arppe	3284
<i>LAGr: Label Aligned Graphs for Better Systematic Generalization in Semantic Parsing</i>	
Dora Jambor and Dzmitry Bahdanau	3295
<i>ToxiGen: A Large-Scale Machine-Generated Dataset for Adversarial and Implicit Hate Speech Detection</i>	
Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray and Ece Kamar	3309
<i>Direct Speech-to-Speech Translation With Discrete Units</i>	
Ann Lee, Peng-Jen Chen, Changan Wang, Jiatao Gu, Sravya Popuri, Xutai Ma, Adam Polyak, Yossi Adi, Qing He, Yun Tang, Juan Pino and Wei-Ning Hsu	3327
<i>Hallucinated but Factual! Inspecting the Factuality of Hallucinations in Abstractive Summarization</i>	
Meng Cao, Yue Dong and Jackie CK Cheung	3340

<i>EntSUM: A Data Set for Entity-Centric Extractive Summarization</i>	
Mounica Maddela, Mayank Kulkarni and Daniel Preotiuc-Pietro	3355
<i>Sentence-level Privacy for Document Embeddings</i>	
Casey Meehan, Khalil Mrini and Kamalika Chaudhuri	3367
<i>Dataset Geography: Mapping Language Data to Language Users</i>	
Fahim Faisal, Yinkai Wang and Antonios Anastasopoulos	3381
<i>ILDAE: Instance-Level Difficulty Analysis of Evaluation Data</i>	
Neeraj Varshney, Swaroop Mishra and Chitta Baral	3412
<i>Image Retrieval from Contextual Descriptions</i>	
Benno Krojer, Vaibhav Adlakha, Vibhav Vineet, Yash Goyal, Edoardo Ponti and Siva Reddy	3426
<i>Multilingual Molecular Representation Learning via Contrastive Pre-training</i>	
Zhihui Guo, Pramod Sharma, Andy Daniel Martinez, Liang Du and Robin Abraham	3441
<i>Investigating Failures of Automatic Translation in the Case of Unambiguous Gender</i>	
Adi Renduchintala and Adina Williams	3454
<i>Cross-Task Generalization via Natural Language Crowdsourcing Instructions</i>	
Swaroop Mishra, Daniel Khashabi, Chitta Baral and Hannaneh Hajishirzi	3470
<i>Imputing Out-of-Vocabulary Embeddings with LOVE Makes Language Models Robust with Little Cost</i>	
Lihu Chen, Gael Varoquaux and Fabian M. Suchanek	3488
<i>NumGLUE: A Suite of Fundamental yet Challenging Mathematical Reasoning Tasks</i>	
Swaroop Mishra, Arindam Mitra, Neeraj Varshney, Bhavdeep Singh Sachdeva, Peter Clark, Chitta Baral and Ashwin Kalyan	3505
<i>Upstream Mitigation Is Not All You Need: Testing the Bias Transfer Hypothesis in Pre-Trained Language Models</i>	
Ryan Steed, Swetasudha Panda, Ari Kobren and Michael L Wick	3524
<i>Improving Multi-label Malevolence Detection in Dialogues through Multi-faceted Label Correlation Enhancement</i>	
Yangjun Zhang, Pengjie Ren, Wentao Deng, Zhumin Chen and Maarten de Rijke	3543
<i>How Do We Answer Complex Questions: Discourse Structure of Long-form Answers</i>	
Fangyuan Xu, Junyi Jessy Li and Eunsol Choi	3556
<i>Understanding Iterative Revision from Human-Written Text</i>	
Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez and Dongyeop Kang	3573
<i>Making Transformers Solve Compositional Tasks</i>	
Santiago Ontanon, Joshua Ainslie, Zachary Fisher and Vaclav Cvicek	3591
<i>Can Transformer be Too Compositional? Analysing Idiom Processing in Neural Machine Translation</i>	
Verna Dankers, Christopher G. Lucas and Ivan Titov	3608
<i>ConditionalQA: A Complex Reading Comprehension Dataset with Conditional Answers</i>	
Haitian Sun, William W. Cohen and Ruslan Salakhutdinov	3627

<i>Prompt-free and Efficient Few-shot Learning with Language Models</i> Rabeeh Karimi Mahabadi, Luke Zettlemoyer, James Henderson, Lambert Mathias, Marzieh Saeidi, Veselin Stoyanov and Majid Yazdani	3638
<i>Continual Sequence Generation with Adaptive Compositional Modules</i> Yanzhe Zhang, Xuezhi Wang and Diyi Yang	3653
<i>An Investigation of the (In)effectiveness of Counterfactually Augmented Data</i> Nitish Joshi and He He	3668
<i>Inducing Positive Perspectives with Text Reframing</i> Caleb Ziems, Minzhi Li, Anthony Zhang and Diyi Yang	3682
<i>VALUE: Understanding Dialect Disparity in NLU</i> Caleb Ziems, Jiaao Chen, Camille Harris, Jessica Anderson and Diyi Yang	3701
<i>From the Detection of Toxic Spans in Online Discussions to the Analysis of Toxic-to-Civil Transfer</i> John Pavlopoulos, Leo Laugier, Alexandros Xenos, Jeffrey S. Sorensen and Ion Androutsopoulos	3721
<i>FormNet: Structural Encoding beyond Sequential Modeling in Form Document Information Extraction</i> Chen-Yu Lee, Chun-Liang Li, Timothy Dozat, Vincent Perot, Guolong Su, Nan Hua, Joshua Ainslie, Renshen Wang, Yasuhisa Fujii and Tomas Pfister	3735
<i>The Moral Integrity Corpus: A Benchmark for Ethical Dialogue Systems</i> Caleb Ziems, Jane A. Yu, Yi-Chia Wang, Alon Y. Halevy and Diyi Yang	3755
<i>Token Dropping for Efficient BERT Pretraining</i> Le Hou, Richard Yuanzhe Pang, Tianyi Zhou, Yuexin Wu, Xinying Song, Xiaodan Song and Denny Zhou	3774
<i>DialFact: A Benchmark for Fact-Checking in Dialogue</i> Prakhar Gupta, Chien-Sheng Wu, Wenhao Liu and Caiming Xiong	3785
<i>The Trade-offs of Domain Adaptation for Neural Language Models</i> David Grangier and Dan Iter	3802
<i>Towards Afrocentric NLP for African Languages: Where We Are and Where We Can Go</i> Ife Adebara and Muhammad Abdul-Mageed	3814
<i>Ensembling and Knowledge Distilling of Large Sequence Taggers for Grammatical Error Correction</i> Maksym Tarnavskiy, Artem Chernodub and Kostiantyn Omelianchuk	3842
<i>Speaker Information Can Guide Models to Better Inductive Biases: A Case Study On Predicting Code-Switching</i> Alissa Ostapenko, Shuly Wintner, Melinda Fricke and Yulia Tsvetkov	3853
<i>Detecting Unassimilated Borrowings in Spanish: An Annotated Corpus and Approaches to Modeling</i> Elena Álvarez-Mellado and Constantine Lignos	3868
<i>Is Attention Explanation? An Introduction to the Debate</i> Adrien Bibal, Rémi Cardon, David Alfter, Rodrigo Wilkens, Xiaou Wang, Thomas François and Patrick Watrin	3889
<i>There Are a Thousand Hamlets in a Thousand People’s Eyes: Enhancing Knowledge-grounded Dialogue with Personal Memory</i> Tingchen Fu, Xueliang Zhao, Chongyang Tao, Ji-Rong Wen and Rui Yan	3901

<i>Neural Pipeline for Zero-Shot Data-to-Text Generation</i> Zdeněk Kasner and Ondrej Dusek	3914
<i>Not always about you: Prioritizing community needs when developing endangered language technology</i> Zoey Liu, Crystal Richardson, Richard John Hatcher and Emily Prud'hommeaux	3933
<i>Automatic Identification and Classification of Bragging in Social Media</i> Mali Jin, Daniel Preotiuc-Pietro, A. Seza Dođruöz and Nikolaos Aletras	3945
<i>Automatic Error Analysis for Document-level Information Extraction</i> Aliva Das, Xinya Du, Barry Wang, Kejian Shi, Jiayuan Gu, Thomas Porter and Claire Cardie	3960
<i>Learning Functional Distributional Semantics with Visual Data</i> Yinhong Liu and Guy Emerson	3976
<i>ePiC: Employing Proverbs in Context as a Benchmark for Abstract Language Understanding</i> Sayan Ghosh and Shashank Srivastava	3989
<i>Chart-to-Text: A Large-Scale Benchmark for Chart Summarization</i> Shankar Kantharaj, Rixie Tiffany Ko Leong, Xiang Lin, Ahmed Masry, Megh Thakkar, Enamul Hoque and Shafiq Joty	4005
<i>Characterizing Idioms: Conventionality and Contingency</i> Michaela Socolof, Jackie CK Cheung, Michael Wagner and Timothy J. O'Donnell	4024
<i>Bag-of-Words vs. Graph vs. Sequence in Text Classification: Questioning the Necessity of Text-Graphs and the Surprising Strength of a Wide MLP</i> Lukas Galke and Ansgar Scherp	4038
<i>Generative Pretraining for Paraphrase Evaluation</i> Jack Weston, Raphael Lenain, Udepa Meepegama and Emil Fristed	4052
<i>Incorporating Stock Market Signals for Twitter Stance Detection</i> Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd and Nigel Collier	4074
<i>Multilingual Mix: Example Interpolation Improves Multilingual Neural Machine Translation</i> Yong Cheng, Ankur Bapna, Orhan Firat, Yuan Cao, Pidong Wang and Wolfgang Macherey .	4092
<i>Word Segmentation as Unsupervised Constituency Parsing</i> Raquel G. Alhama	4103
<i>SafetyKit: First Aid for Measuring Safety in Open-domain Conversational Systems</i> Emily Dinan, Gavin Abercrombie, A. Stevie Bergman, Shannon L. Spruit, Dirk Hovy, Y-Lan Boureau and Verena Rieser	4113
<i>Zero-Shot Cross-lingual Semantic Parsing</i> Tom Sherborne and Mirella Lapata	4134
<i>The Paradox of the Compositionality of Natural Language: A Neural Machine Translation Case Study</i> Verna Dankers, Elia Bruni and Dieuwke Hupkes	4154
<i>Multilingual Document-Level Translation Enables Zero-Shot Transfer From Sentences to Documents</i> Biao Zhang, Ankur Bapna, Melvin Johnson, Ali Dabirmoghaddam, Naveen Arivazhagan and Orhan Firat	4176

<i>Cross-Lingual Phrase Retrieval</i>	
Heqi Zheng, Xiao Zhang, Zewen Chi, Heyan Huang, Yan Tan, Tian Lan, Wei Wei and Xian-Ling Mao	4193
<i>Improving Compositional Generalization with Self-Training for Data-to-Text Generation</i>	
Sanket Vaibhav Mehta, Jinfeng Rao, Yi Tay, Mihir Kale, Ankur P Parikh and Emma Strubell	4205
<i>MMCoQA: Conversational Question Answering over Text, Tables, and Images</i>	
Yongqi Li, Wenjie Li and Liqiang Nie	4220
<i>Effective Token Graph Modeling using a Novel Labeling Strategy for Structured Sentiment Analysis</i>	
Wenxuan Shi, Fei Li, Jingye Li, Hao Fei and Donghong Ji	4232
<i>PromDA: Prompt-based Data Augmentation for Low-Resource NLU Tasks</i>	
Yufei Wang, Can Xu, Qingfeng Sun, Huang Hu, Chongyang Tao, Xiubo Geng and Daxin Jiang	4242
<i>Disentangled Sequence to Sequence Learning for Compositional Generalization</i>	
Hao Zheng and Mirella Lapata	4256
<i>RST Discourse Parsing with Second-Stage EDU-Level Pre-training</i>	
Nan Yu, Meishan Zhang, Guohong Fu and Min Zhang	4269
<i>SimKGC: Simple Contrastive Knowledge Graph Completion with Pre-trained Language Models</i>	
Liang Wang, Wei Zhao, Zhuoyu Wei and Jingming Liu	4281
<i>Do Transformer Models Show Similar Attention Patterns to Task-Specific Human Gaze?</i>	
Oliver Eberle, Stephanie Brandl, Jonas Pilot and Anders Søgaard	4295
<i>LexGLUE: A Benchmark Dataset for Legal Language Understanding in English</i>	
Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Martin Katz and Nikolaos Aletras	4310
<i>DiBiMT: A Novel Benchmark for Measuring Word Sense Disambiguation Biases in Machine Translation</i>	
Niccolò Campolungo, Federico Martelli, Francesco Saina and Roberto Navigli	4331
<i>Improving Word Translation via Two-Stage Contrastive Learning</i>	
Yaoyiran Li, Fangyu Liu, Nigel Collier, Anna Korhonen and Ivan Vulić	4353
<i>Scheduled Multi-task Learning for Neural Chat Translation</i>	
Yunlong Liang, Fandong Meng, Jinan Xu, Yufeng Chen and Jie Zhou	4375
<i>FairLex: A Multilingual Benchmark for Evaluating Fairness in Legal Text Processing</i>	
Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Felix Schwemer and Anders Søgaard	4389
<i>Towards Abstractive Grounded Summarization of Podcast Transcripts</i>	
Kaiqiang Song, Chen Li, Xiaoyang Wang, Dong Yu and Fei Liu	4407
<i>FiNER: Financial Numeric Entity Recognition for XBRL Tagging</i>	
Lefteris Loukas, Manos Fergadiotis, Ilias Chalkidis, Eirini Spyropoulou, Prodromos Malakasiotis, Ion Androutsopoulos and Georgios Paliouras	4419
<i>Keywords and Instances: A Hierarchical Contrastive Learning Framework Unifying Hybrid Granularities for Text Generation</i>	
Mingzhe Li, XieXiong Lin, Xiuying Chen, Jinxiong Chang, Qishen Zhang, Feng Wang, Taifeng Wang, Zhongyi Liu, Wei Chu, Dongyan Zhao and Rui Yan	4432

<i>EPT-X: An Expression-Pointer Transformer model that generates eXplanations for numbers</i> Bugeun Kim, Kyung Seo Ki, Sangkyu Rhim and Gahgene Gweon	4442
<i>Identifying the Human Values behind Arguments</i> Johannes Kiesel, Milad Alshomary, Nicolas Handke, Xiaoni Cai, Henning Wachsmuth and Benno Stein	4459
<i>BenchIE: A Framework for Multi-Faceted Fact-Based Open Information Extraction Evaluation</i> Kiril Gashteovski, Mingying Yu, Bhushan Kotnis, Carolin Lawrence, Mathias Niepert and Goran Glavaš	4472
<i>Leveraging Unimodal Self-Supervised Learning for Multimodal Audio-Visual Speech Recognition</i> Xichen Pan, Peiyu Chen, Yichen Gong, Helong Zhou, Xinbing Wang and Zhouhan Lin	4491
<i>SummaReranker: A Multi-Task Mixture-of-Experts Re-ranking Framework for Abstractive Summarization</i> Mathieu Ravaut, Shafiq Joty and Nancy F. Chen	4504
<i>Understanding Multimodal Procedural Knowledge by Sequencing Multimodal Instructional Manuals</i> Te-Lin Wu, Alex Spangher, Pegah Alipoormolabashi, Marjorie Freedman, Ralph M. Weischedel and Nanyun Peng	4525
<i>Zoom Out and Observe: News Environment Perception for Fake News Detection</i> Qiang Sheng, Juan Cao, Xueyao Zhang, Rundong Li, Danding Wang and Yongchun Zhu . . .	4543
<i>Divide and Rule: Effective Pre-Training for Context-Aware Multi-Encoder Translation Models</i> Lorenzo Lupo, Marco Dinarelli and Laurent Besacier	4557
<i>Saliency as Evidence: Event Detection with Trigger Saliency Attribution</i> Jian Liu, Yufeng Chen and Jinan Xu	4573
<i>SRLAE – Semantic Role Labeling for Emotions: A Unified Evaluation Framework</i> Cesare Campagnano, Simone Conia and Roberto Navigli	4586
<i>Context Matters: A Pragmatic Study of PLMs’ Negation Understanding</i> Reto Gubelmann and Siegfried Handschuh	4602
<i>Probing for Predicate Argument Structures in Pretrained Language Models</i> Simone Conia and Roberto Navigli	4622
<i>Multilingual Generative Language Models for Zero-Shot Cross-Lingual Event Argument Extraction</i> Kuan-Hao Huang, I-Hung Hsu, Prem Natarajan, Kai-Wei Chang and Nanyun Peng	4633
<i>Identifying Moments of Change from Longitudinal User Text</i> Adam Tsakalidis, Federico Nanni, Anthony Hills, Jenny Chim, Jiayu Song and Maria Liakata	4647
<i>Multi-Task Pre-Training for Plug-and-Play Task-Oriented Dialogue System</i> Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai and Yi Zhang . .	4661
<i>Graph Enhanced Contrastive Learning for Radiology Findings Summarization</i> Jinpeng Hu, Zhuo Li, Zhihong Chen, Zhen Li, Xiang Wan and Tsung-Hui Chang	4677
<i>Semi-Supervised Formality Style Transfer with Consistency Training</i> Ao Liu, An Wang and Naoaki Okazaki	4689
<i>Cross-Lingual Ability of Multilingual Masked Language Models: A Study of Language Structure</i> Yuan Chai, Yaobo Liang and Nan Duan	4702

<i>Rare and Zero-shot Word Sense Disambiguation using Z-Reweighting</i>	
Ying Su, Hongming Zhang, Yangqiu Song and Tong Zhang	4713
<i>Nibbling at the Hard Core of Word Sense Disambiguation</i>	
Marco Maru, Simone Conia, Michele Bevilacqua and Roberto Navigli	4724
<i>Large Scale Substitution-based Word Sense Induction</i>	
Matan Eyal, Shoval Sadde, Hillel Taub-Tabib and Yoav Goldberg	4738
<i>Can Synthetic Translations Improve Bitext Quality?</i>	
Eleftheria Briakou and Marine Carpuat	4753
<i>Unsupervised Dependency Graph Network</i>	
Yikang Shen, Shawn Tan, Alessandro Sordoni, Peng Li, Jie Zhou and Aaron Courville	4767
<i>WikiDiverse: A Multimodal Entity Linking Dataset with Diversified Contextual Topics and Entity Types</i>	
Xuwu Wang, Junfeng Tian, Min Gui, Zhixu Li, Rui Wang, Ming Yan, Lihan Chen and Yanghua Xiao	4785
<i>Rewire-then-Probe: A Contrastive Recipe for Probing Biomedical Knowledge of Pre-trained Language Models</i>	
Zaiqiao Meng, Fangyu Liu, Ehsan Shareghi, Yixuan Su, Charlotte Anne Collins and Nigel Collier	4798
<i>Fine- and Coarse-Granularity Hybrid Self-Attention for Efficient BERT</i>	
Jing Zhao, Yifan Wang, Junwei Bao, Youzheng Wu and Xiaodong He	4811
<i>Compression of Generative Pre-trained Language Models via Quantization</i>	
Chaofan Tao, Lu Hou, Wei Zhang, Lifeng Shang, Xin Jiang, Qun Liu, Ping Luo and Ngai Wong	4821
<i>Visual-Language Navigation Pretraining via Prompt-based Environmental Self-exploration</i>	
Xiwen Liang, Fengda Zhu, Li Lingling, Hang Xu and Xiaodan Liang	4837
<i>DialogVED: A Pre-trained Latent Variable Encoder-Decoder Model for Dialog Response Generation</i>	
Wei Chen, Yeyun Gong, Song Wang, Bolun Yao, Weizhen Qi, Zhongyu Wei, Xiaowu Hu, Bartuer Zhou, Yi Mao, Weizhu Chen, Biao Cheng and Nan Duan	4852
<i>Contextual Fine-to-Coarse Distillation for Coarse-grained Response Selection in Open-Domain Conversations</i>	
Wei Chen, Yeyun Gong, Can Xu, Huang Hu, Bolun Yao, Zhongyu Wei, Zhihao Fan, Xiaowu Hu, Bartuer Zhou, Biao Cheng, Daxin Jiang and Nan Duan	4865
<i>Textomics: A Dataset for Genomics Data Summary Generation</i>	
Mu-Chun Wang, Zixuan Liu and Sheng Wang	4878
<i>A Contrastive Framework for Learning Sentence Representations from Pairwise and Triple-wise Perspective in Angular Space</i>	
Yuhao Zhang, Hongji Zhu, Yongliang Wang, Nan Xu, Xiaobo LI and Binqiang Zhao	4892
<i>Packed Levitated Marker for Entity and Relation Extraction</i>	
Deming Ye, Yankai Lin, Peng Li and Maosong Sun	4904
<i>An Interpretable Neuro-Symbolic Reasoning Framework for Task-Oriented Dialogue Generation</i>	
Shiquan Yang, Rui Zhang, Sarah Monazam Erfani and Jey Han Lau	4918
<i>Impact of Evaluation Methodologies on Code Summarization</i>	
Pengyu Nie, Jiyang Zhang, Junyi Jessy Li, Ray Mooney and Milos Gligoric	4936

<i>KG-FiD: Infusing Knowledge Graph in Fusion-in-Decoder for Open-Domain Question Answering</i> Donghan Yu, Chenguang Zhu, Yuwei Fang, Wenhao Yu, Shuohang Wang, Yichong Xu, Xiang Ren, Yiming Yang and Michael Zeng	4961
<i>Which side are you on? Insider-Outsider classification in conspiracy-theoretic social media</i> Pavan S Holur, Tianyi Wang, Shadi Shabsavari, Timothy R Tangherlini and Vwani Roychowdhury	4975
<i>Learning From Failure: Data Capture in an Australian Aboriginal Community</i> Eric Le Ferrand, Steven Bird and Laurent Besacier	4988
<i>Deep Inductive Logic Reasoning for Multi-Hop Reading Comprehension</i> Wenya Wang and Sinno Pan	4999
<i>CICERO: A Dataset for Contextualized Commonsense Inference in Dialogues</i> Deepanway Ghosal, Siqi Shen, Navonil Majumder, Rada Mihalcea and Soujanya Poria	5010
<i>A Comparative Study of Faithfulness Metrics for Model Interpretability Methods</i> Chun Sik Chan, Huanqi Kong and Liang Guanqing	5029
<i>SPoT: Better Frozen Model Adaptation through Soft Prompt Transfer</i> Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou' and Daniel Cer	5039
<i>Pass off Fish Eyes for Pearls: Attacking Model Selection of Pre-trained Models</i> Biru Zhu, Yujia Qin, Fanchao Qi, Yangdong Deng, Zhiyuan Liu, Maosong Sun and Ming Gu	5060
<i>Educational Question Generation of Children Storybooks via Question Type Distribution Learning and Event-centric Summarization</i> Zhenjie Zhao, Yufang Hou, Dakuo Wang, Mo Yu, Chengzhong Liu and Xiaojuan Ma	5073
<i>HeterMPC: A Heterogeneous Graph Neural Network for Response Generation in Multi-Party Conversations</i> Jia-Chen Gu, Chao-Hong Tan, Chongyang Tao, Zhen-Hua Ling, Huang Hu, Xiubo Geng and Daxin Jiang	5086
<i>The patient is more dead than alive: exploring the current state of the multi-document summarisation of the biomedical literature</i> Yulia Otmakhova, Karin Verspoor, Timothy Baldwin and Jey Han Lau	5098
<i>A Multi-Document Coverage Reward for RELAXed Multi-Document Summarization</i> Jacob Parnell, Inigo Jauregi Unanue and Massimo Piccardi	5112
<i>KNN-Contrastive Learning for Out-of-Domain Intent Classification</i> Yunhua Zhou, Peiju Liu and Xipeng Qiu	5129
<i>A Neural Network Architecture for Program Understanding Inspired by Human Behaviors</i> Renyu Zhu, Lei Yuan, Xiang Li, Ming Gao and Wenyan Cai	5142
<i>FaVIQ: FAct Verification from Information-seeking Questions</i> Jungsoo Park, Sewon Min, Jaewoo Kang, Luke Zettlemoyer and Hannaneh Hajishirzi	5154
<i>Simulating Bandit Learning from User Feedback for Extractive Question Answering</i> Ge Gao, Eunsol Choi and Yoav Artzi	5167
<i>Beyond Goldfish Memory: Long-Term Open-Domain Conversation</i> Jing Xu, Arthur Szlam and Jason E Weston	5180

<i>ReCLIP: A Strong Zero-Shot Baseline for Referring Expression Comprehension</i> Sanjay Subramanian, William Merrill, Trevor Darrell, Matt Gardner, Sameer Singh and Anna Rohrbach	5198
<i>Dynamic Prefix-Tuning for Generative Template-based Event Extraction</i> Xiao Liu, Heyan Huang, Ge Shi and Bo Wang	5216
<i>E-LANG: Energy-Based Joint Inferencing of Super and Swift Language Models</i> Mohammad Akbari, Amin Banitalebi-Dehkordi and Yong Zhang	5229
<i>PRIMERA: Pyramid-based Masked Sentence Pre-training for Multi-document Summarization</i> Wen Xiao, Iz Beltagy, Giuseppe Carenini and Arman Cohan	5245
<i>Dynamic Global Memory for Document-level Argument Extraction</i> Xinya Du, Sha Li and Heng Ji	5264
<i>Measuring the Impact of (Psycho-)Linguistic and Readability Features and Their Spill Over Effects on the Prediction of Eye Movement Patterns</i> Daniel Wiechmann and Elma Kerz	5276
<i>Alternative Input Signals Ease Transfer in Multilingual Machine Translation</i> Simeng Sun, Angela Fan, James Cross, Vishrav Chaudhary, Chau Tran, Philipp Koehn and Francisco Guzmán	5291
<i>Phone-ing it in: Towards Flexible Multi-Modal Language Model Training by Phonetic Representations of Data</i> Colin Leong and Daniel Lee Whitenack	5306
<i>Noisy Channel Language Model Prompting for Few-Shot Text Classification</i> Sewon Min, Mike Lewis, Hannaneh Hajishirzi and Luke Zettlemoyer	5316
<i>Multilingual unsupervised sequence segmentation transfers to extremely low-resource languages</i> C. M. Downey, Shannon Drizin, Levon Haroutunian and Shivin Thukral	5331
<i>KinyaBERT: a Morphology-aware Kinyarwanda Language Model</i> Antoine Nzeyimana and Andre Niyongabo Rubungo	5347
<i>On the Calibration of Pre-trained Language Models using Mixup Guided by Area Under the Margin and Saliency</i> Seo Yeon Park and Cornelia Caragea	5364
<i>IMPLI: Investigating NLI Models' Performance on Figurative Language</i> Kevin Stowe, Prasetya Ajie Utama and Iryna Gurevych	5375
<i>QAConv: Question Answering on Informative Conversations</i> Chien-Sheng Wu, Andrea Madotto, Wenhao Liu, Pascale Fung and Caiming Xiong	5389
<i>Prix-LM: Pretraining for Multilingual Knowledge Base Construction</i> Wenxuan Zhou, Fangyu Liu, Ivan Vulić, Nigel Collier and Muhao Chen	5412
<i>Semantic Composition with PSHRG for Derivation Tree Reconstruction from Graph-Based Meaning Representations</i> Chun Hei Lo, Wai Lam and Hong Cheng	5425
<i>HOLM: Hallucinating Objects with Language Models for Referring Expression Recognition in Partially-Observed Scenes</i> Volkan Cirik, Louis-Philippe Morency and Taylor Berg-Kirkpatrick	5440

<i>Multi Task Learning For Zero Shot Performance Prediction of Multilingual Models</i>	
Kabir Ahuja, Shanu Kumar, Sandipan Dandapat and Monojit Choudhury	5454
<i>∞-former: Infinite Memory Transformer</i>	
Pedro Henrique Martins, Zita Marinho and Andre Martins	5468
<i>Systematic Inequalities in Language Technology Performance across the World's Languages</i>	
Damian E Blasi, Antonios Anastasopoulos and Graham Neubig	5486
<i>CaMEL: Case Marker Extraction without Labels</i>	
Leonie Weissweiler, Valentin Hofmann, Masoud Jalili Sabet and Hinrich Schuetze	5506
<i>Improving Generalizability in Implicitly Abusive Language Detection with Concept Activation Vectors</i>	
Isar Nejadgholi, Kathleen C. Fraser and Svetlana Kiritchenko	5517
<i>Reports of personal experiences and stories in argumentation: datasets and analysis</i>	
Neele Falk and Gabriella Lapesa	5530
<i>Non-neural Models Matter: a Re-evaluation of Neural Referring Expression Generation Systems</i>	
Fahime Same, Guanyi Chen and Kees Van Deemter	5554
<i>Bridging the Generalization Gap in Text-to-SQL Parsing with Schema Expansion</i>	
Chen Zhao, Yu Su, Adam Pauls and Emmanouil Antonios Platanios	5568
<i>Predicate-Argument Based Bi-Encoder for Paraphrase Identification</i>	
Qiwei Peng, David Weir, Julie Weeds and Yekun Chai	5579
<i>MINER: Improving Out-of-Vocabulary Named Entity Recognition from an Information Theoretic Perspective</i>	
Xiao Wang, Shihan Dou, Limao Xiong, Yicheng Zou, Qi Zhang, Tao Gui, Liang Qiao, Zhanzhan Cheng and Xuanjing Huang	5590
<i>Leveraging Wikipedia article evolution for promotional tone detection</i>	
Christine De Kock and Andreas Vlachos	5601
<i>From text to talk: Harnessing conversational corpora for humane and diversity-aware language technology</i>	
Mark Dingemanse and Andreas Liesenfeld	5614
<i>Flooding-X: Improving BERT's Resistance to Adversarial Attacks via Loss-Restricted Fine-Tuning</i>	
Qin Liu, Rui Zheng, Bao Rong, Jingyi Liu, ZhiHua Liu, Zhanzhan Cheng, Liang Qiao, Tao Gui, Qi Zhang and Xuanjing Huang	5634
<i>RoMe: A Robust Metric for Evaluating Natural Language Generation</i>	
Md Rashad Al Hasan Rony, Liubov Kovriguina, Debanjan Chaudhuri, Ricardo Usbeck and Jens Lehmann	5645
<i>Finding Structural Knowledge in Multimodal-BERT</i>	
Victor Siemen Janusz Milewski, Miryam de Lhoneux and Marie-Francine Moens	5658
<i>Fully Hyperbolic Neural Networks</i>	
Weize Chen, Xu Han, Yankai Lin, Hexu Zhao, Zhiyuan Liu, Peng Li, Maosong Sun and Jie Zhou	5672
<i>Neural Machine Translation with Phrase-Level Universal Visual Representations</i>	
Qingkai Fang and Yang Feng	5687

<i>M3ED: Multi-modal Multi-scene Multi-label Emotional Dialogue Database</i>	
Jinming Zhao, Tenggao Zhang, Jingwen Hu, Yuchen Liu, Qin Jin, Xinchao Wang and Haizhou Li	5699
<i>Few-shot Named Entity Recognition with Self-describing Networks</i>	
Jiawei Chen, Qing Liu, Hongyu Lin, Xianpei Han and Le Sun	5711
<i>SpeechT5: Unified-Modal Encoder-Decoder Pre-Training for Spoken Language Processing</i>	
Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie Liu, Tom Ko, Qing Li, Yu Zhang, Zhihua Wei, Yao Qian, Jinyu Li and Furu Wei	5723
<i>Human Evaluation and Correlation with Automatic Metrics in Consultation Note Generation</i>	
Francesco Moramarco, Alex Papadopoulos Korfiatis, Mark Perera, Damir Juric, Jack Flann, Ehud Reiter, Anya Belz and Aleksandar Savkov	5739
<i>Unified Structure Generation for Universal Information Extraction</i>	
Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu Lin, Xianpei Han, Le Sun and Hua Wu	5755
<i>Subgraph Retrieval Enhanced Model for Multi-hop Knowledge Base Question Answering</i>	
Jing Zhang, Xiaokang Zhang, Jifan Yu, Jian Tang, Jie Tang, Cuiping Li and Hong Chen	5773
<i>Pre-training to Match for Unified Low-shot Relation Extraction</i>	
Fangchao Liu, Hongyu Lin, Xianpei Han, Boxi Cao and Le Sun	5785
<i>Can Prompt Probe Pretrained Language Models? Understanding the Invisible Risks from a Causal View</i>	
Boxi Cao, Hongyu Lin, Xianpei Han, Fangchao Liu and Le Sun	5796
<i>Evaluating Extreme Hierarchical Multi-label Classification</i>	
Enrique Amigo and Agustín D. Delgado	5809
<i>What does the sea say to the shore? A BERT based DST style approach for speaker to dialogue attribution in novels</i>	
Carolina Cuesta-Lazaro, Animesh Prasad and Trevor Wood	5820
<i>Measuring Fairness of Text Classifiers via Prediction Sensitivity</i>	
Satyapriya Krishna, Rahul Gupta, Apurv Verma, Jwala Dhamala, Yada Pruksachatkun and Kai-Wei Chang	5830
<i>RotateQVS: Representing Temporal Information as Rotations in Quaternion Vector Space for Temporal Knowledge Graph Completion</i>	
Kai Chen, Ye Wang, Yitong Li and Aiping Li	5843
<i>Feeding What You Need by Understanding What You Learned</i>	
Xiaoqiang Wang, Bang Liu, Fangli Xu, Bo Long, Siliang Tang and Lingfei Wu	5858
<i>Probing Simile Knowledge from Pre-trained Language Models</i>	
Weijie Chen, Yongzhu Chang, Rongsheng Zhang, Jiashu Pu, Guandan Chen, Le Zhang, Yadong Xi, Yijiang Chen and Chang su	5875
<i>An Effective and Efficient Entity Alignment Decoding Algorithm via Third-Order Tensor Isomorphism</i>	
Xin Mao, Meirong Ma, Hao Yuan, Jianchao Zhu, Zongyu Wang, Rui Xie, Wei Wu and Man Lan	5888
<i>Entailment Graph Learning with Textual Entailment and Soft Transitivity</i>	
Zhibin Chen, Yansong Feng and Dongyan Zhao	5899

<i>Logic Traps in Evaluating Attribution Scores</i>	
Yiming Ju, Yuanzhe Zhang, Zhao Yang, Zhongtao Jiang, Kang Liu and Jun Zhao	5911
<i>Continual Pre-training of Language Models for Math Problem Understanding with Syntax-Aware Memory Network</i>	
Zheng Gong, Kun Zhou, Xin Zhao, Jing Sha, Shijin Wang and Ji-Rong Wen	5923
<i>Multitasking Framework for Unsupervised Simple Definition Generation</i>	
Cunliang Kong, Yun Chen, Hengyuan Zhang, Liner Yang and Erhong Yang	5934
<i>Learning to Reason Deductively: Math Word Problem Solving as Complex Relation Extraction</i>	
Zhanming Jie, Jierui Li and Wei Lu	5944
<i>When did you become so smart, oh wise one?! Sarcasm Explanation in Multi-modal Multi-party Dialogues</i>	
Shivani Kumar, Atharva Kulkarni, Md Shad Akhtar and Tanmoy Chakraborty	5956
<i>Toward Interpretable Semantic Textual Similarity via Optimal Transport-based Contrastive Sentence Learning</i>	
Seonghyeon Lee, Dongha Lee, Seongbo Jang and Hwanjo Yu	5969
<i>Pre-training and Fine-tuning Neural Topic Model: A Simple yet Effective Approach to Incorporating External Knowledge</i>	
Linhai Zhang, Xuemeng Hu, Boyu Wang, Deyu Zhou, Qian-Wen Zhang and Yunbo Cao	5980
<i>Multi-View Document Representation Learning for Open-Domain Dense Retrieval</i>	
Shunyu Zhang, Yaobo Liang, Ming Gong, Daxin Jiang and Nan Duan	5990
<i>Graph Pre-training for AMR Parsing and Generation</i>	
Xuefeng Bai, Yulong Chen and Yue Zhang	6001
<i>Turning Tables: Generating Examples from Semi-structured Tables for Endowing Language Models with Reasoning Skills</i>	
Ori Yoran, Alon Talmor and Jonathan Berant	6016
<i>RNG-KBQA: Generation Augmented Iterative Ranking for Knowledge Base Question Answering</i>	
Xi Ye, Semih Yavuz, Kazuma Hashimoto, Yingbo Zhou and Caiming Xiong	6032
<i>Rethinking Self-Supervision Objectives for Generalizable Coherence Modeling</i>	
Prathyusha Jwalapuram, Shafiq Joty and Xiang Lin	6044
<i>Just Rank: Rethinking Evaluation with Word and Sentence Similarities</i>	
Bin Wang, C.-c. Jay Kuo and Haizhou Li	6060
<i>MarkupLM: Pre-training of Text and Markup Language for Visually Rich Document Understanding</i>	
Junlong Li, Yiheng Xu, Lei Cui and Furu Wei	6078
<i>CLIP Models are Few-Shot Learners: Empirical Studies on VQA and Visual Entailment</i>	
Haoyu Song, Li Dong, Weinan Zhang, Ting Liu and Furu Wei	6088
<i>KQA Pro: A Dataset with Explicit Compositional Programs for Complex Question Answering over Knowledge Base</i>	
Shulin Cao, Jiaxin Shi, Liangming Pan, Lunyiu Nie, Yutong Xiang, Lei Hou, Juanzi Li, Bin He and Hanwang Zhang	6101
<i>Debiased Contrastive Learning of Unsupervised Sentence Representations</i>	
Kun Zhou, Beichen Zhang, Xin Zhao and Ji-Rong Wen	6120

<i>MSP: Multi-Stage Prompting for Making Pre-trained Language Models Better Translators</i> Zhixing Tan, Xiangwen Zhang, Shuo Wang and Yang Liu	6131
<i>SalesBot: Transitioning from Chit-Chat to Task-Oriented Dialogues</i> Ssu Chiu, Maolin Li, Yen-Ting Lin and Yun-Nung Chen	6143
<i>UCTopic: Unsupervised Contrastive Learning for Phrase Representations and Topic Mining</i> Jiacheng Li, Jingbo Shang and Julian McAuley	6159
<i>XLM-E: Cross-lingual Language Model Pre-training via ELECTRA</i> Zewen Chi, Shaohan Huang, Li Dong, Shuming Ma, Bo Zheng, Saksham Singhal, Payal Bajaj, Xia Song, Xian-Ling Mao, Heyan Huang and Furu Wei	6170
<i>Nested Named Entity Recognition as Latent Lexicalized Constituency Parsing</i> Chao Lou, Songlin Yang and Kewei Tu	6183
<i>Can Explanations Be Useful for Calibrating Black Box Models?</i> Xi Ye and Greg Durrett	6199
<i>OIE@OIA: an Adaptable and Efficient Open Information Extraction Framework</i> Xin Wang, Minlong Peng, Mingming Sun and Ping Li	6213
<i>ReACC: A Retrieval-Augmented Code Completion Framework</i> Shuai Lu, Nan Duan, Hojae Han, Daya Guo, Seung-won Hwang and Alexey Svyatkovskiy .	6227
<i>Does Recommend-Revise Produce Reliable Annotations? An Analysis on Missing Instances in DocRED</i> Quzhe Huang, Shibo Hao, Yuan Ye, Shengqi Zhu, Yansong Feng and Dongyan Zhao	6241
<i>UniPELT: A Unified Framework for Parameter-Efficient Language Model Tuning</i> Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Scott Yih and Madian Khabsa	6253
<i>An Empirical Study of Memorization in NLP</i> Xiaosen Zheng and Jing Jiang	6265
<i>AmericasNLI: Evaluating Zero-shot Natural Language Understanding of Pretrained Multilingual Models in Truly Low-resource Languages</i> Abteen Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios, Ivan Vladimir Meza Ruiz, Gustavo A. Giménez-Lugo, Elisabeth Mager, Graham Neubig, Alexis Palmer, Rolando Coto-Solano, Thang Vu and Katharina Kann	6279
<i>Towards Learning (Dis)-Similarity of Source Code from Program Contrasts</i> Yangruibo Ding, Luca Buratti, Saurabh Pujar, Alessandro Morari, Baishakhi Ray and Saikat Chakraborty	6300
<i>Guided Attention Multimodal Multitask Financial Forecasting with Inter-Company Relationships and Global and Local News</i> Gary Ang and Ee-Peng Lim	6313
<i>On Vision Features in Multimodal Machine Translation</i> Bei Li, Chuanhao Lv, Zefan Zhou, Tao Zhou, Tong Xiao, Anxiang Ma and JingBo Zhu	6327
<i>CONTaiNER: Few-Shot Named Entity Recognition via Contrastive Learning</i> Sarkar Snigdha Sarathi Das, Arzoo Katiyar, Rebecca J. Passonneau and Rui Zhang	6338

<i>Cree Corpus: A Collection of nêhiyawêwin Resources</i> Daniela Teodorescu, Josie Matalski, Delaney Alexa Lothian, Denilson Barbosa and Carrie Demmans Epp	6354
<i>Learning to Rank Visual Stories From Human Ranking Data</i> Chi-Yang Hsu, Yun-Wei Chu, Vincent Chen, Kuan-Chieh Lo, Chacha Chen, Ting-Hao Huang and Lun-Wei Ku	6365
<i>Universal Conditional Masked Language Pre-training for Neural Machine Translation</i> Pengfei Li, Liangyou Li, Meng Zhang, Minghao Wu and Qun Liu	6379
<i>CARETS: A Consistency And Robustness Evaluative Test Suite for VQA</i> Carlos E Jimenez, Olga Russakovsky and Karthik R Narasimhan	6392
<i>Phrase-aware Unsupervised Constituency Parsing</i> Xiaotao Gu, Yikang Shen, Jiaming Shen, Jingbo Shang and Jiawei Han	6406
<i>Achieving Reliable Human Assessment of Open-Domain Dialogue Systems</i> Tianbo Ji, Yvette Graham, Gareth J. F. Jones, Chenyang Lyu and Qun Liu	6416
<i>Updated Headline Generation: Creating Updated Summaries for Evolving News Stories</i> Sheena Panthaplackel, Adrian Benton and Mark Dredze	6438
<i>SaFeRDialogues: Taking Feedback Gracefully after Conversational Safety Failures</i> Megan Ung, Jing Xu and Y-Lan Boureau	6462
<i>Compositional Generalization in Dependency Parsing</i> Emily Goodwin, Siva Reddy, Timothy J. O'Donnell and Dzmitry Bahdanau	6482
<i>ASPECTNEWS: Aspect-Oriented Summarization of News Documents</i> Ojas Ahuja, Jiacheng Xu, Akshay Kumar Gupta, Kevin Horecka and Greg Durrett	6494
<i>MemSum: Extractive Summarization of Long Documents Using Multi-Step Episodic Markov Decision Processes</i> Nianlong Gu, Elliott Ash and Richard Hahnloser	6507
<i>CLUES: A Benchmark for Learning Classifiers using Natural Language Explanations</i> Rakesh R Menon, Sayan Ghosh and Shashank Srivastava	6523
<i>Substructure Distribution Projection for Zero-Shot Cross-Lingual Dependency Parsing</i> Freda Shi, Kevin Gimpel and Karen Livescu	6547
<i>Multilingual Detection of Personal Employment Status on Twitter</i> Manuel Tonneau, Dhaval Adjudah, Joao Palotti, Nir Grinberg and Samuel Fraiberger	6564
<i>MultiHiertt: Numerical Reasoning over Multi Hierarchical Tabular and Textual Data</i> Yilun Zhao, Yunxiang Li, Chenying Li and Rui Zhang	6588
<i>Transformers in the loop: Polarity in neural models of language</i> Lisa Bylinina and Alexey Tikhonov	6601
<i>Bridging the Data Gap between Training and Inference for Unsupervised Neural Machine Translation</i> Zhiwei He, Xing Wang, Rui Wang, Shuming Shi and Zhaopeng Tu	6611
<i>SDR: Efficient Neural Re-ranking using Succinct Document Representation</i> Nachshon Cohen, Amit Portnoy, Besnik Fetahu and Amir Ingber	6624

<i>The AI Doctor Is In: A Survey of Task-Oriented Dialogue Systems for Healthcare Applications</i> Mina Valizadeh and Natalie Parde	6638
<i>SHIELD: Defending Textual Neural Networks against Multiple Black-Box Adversarial Attacks with Stochastic Multi-Expert Patcher</i> Thai Le, Noseong Park and Dongwon Lee	6661
<i>Accurate Online Posterior Alignments for Principled Lexically-Constrained Decoding</i> Soumya Chatterjee, Sunita Sarawagi and Preethi Jyothi	6675
<i>Leveraging Task Transferability to Meta-learning for Clinical Section Classification with Limited Data</i> Zhuohao Chen, Jangwon Kim, Ram Bhakta and Mustafa Y Sir	6690
<i>Reinforcement Guided Multi-Task Learning Framework for Low-Resource Stereotype Detection</i> Rajkumar Pujari, Erik Oveson, Priyanka Kulkarni and Elnaz Nouri	6703
<i>Letters From the Past: Modeling Historical Sound Change Through Diachronic Character Embeddings</i> Sidsel Boldsen and Patrizia Paggio	6713
<i>A Token-level Reference-free Hallucination Detection Benchmark for Free-form Text Generation</i> Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen and Bill Dolan	6723
<i>Low-Rank Softmax Can Have Unargmaxable Classes in Theory but Rarely in Practice</i> Andreas Grivas, Nikolay Bogoychev and Adam Lopez	6738
<i>Prompt for Extraction? PAIE: Prompting Argument Interaction for Event Argument Extraction</i> Yubo Ma, Zehao Wang, Yixin Cao, Mukai Li, Meiqi Chen, Kun Wang and Jing Shao	6759
<i>Reducing Position Bias in Simultaneous Machine Translation with Length-Aware Framework</i> Shaolei Zhang and Yang Feng	6775
<i>A Statutory Article Retrieval Dataset in French</i> Antoine Louis and Gerasimos Spanakis	6789
<i>ParaDetox: Detoxification with Parallel Data</i> Varvara Logacheva, Daryna Dementieva, Sergey Ustyantsev, Daniil Moskovskiy, David Dale, Irina Krotova, Nikita Semenov and Alexander Panchenko	6804
<i>Interpreting Character Embeddings With Perceptual Representations: The Case of Shape, Sound, and Color</i> Sidsel Boldsen, Manex Agirrezabal and Nora Hollenstein	6819
<i>Fine-Grained Controllable Text Generation Using Non-Residual Prompting</i> Fredrik Carlsson, Joey Öhman, Fangyu Liu, Severine Verlinden, Joakim Nivre and Magnus Sahlgren	6837
<i>Language-Agnostic Meta-Learning for Low-Resource Text-to-Speech with Articulatory Features</i> Florian Lux and Thang Vu	6858
<i>TwittIrish: A Universal Dependencies Treebank of Tweets in Modern Irish</i> Lauren Cassidy, Teresa Lynn, James Barry and Jennifer Foster	6869
<i>Length Control in Abstractive Summarization by Pretraining Information Selection</i> Yizhu Liu, Qi Jia and Kenny Q. Zhu	6885
<i>CQG: A Simple and Effective Controlled Generation Framework for Multi-hop Question Generation</i> Zichu Fei, Qi Zhang, Tao Gui, Di Liang, Sirui Wang, Wei Wu and Xuanjing Huang	6896

<i>Word Order Does Matter and Shuffled Language Models Know It</i>	
Mostafa Abdou, Vinit Ravishankar, Artur Kulmizev and Anders Søgaard	6907
<i>An Empirical Study on Explanations in Out-of-Domain Settings</i>	
George Chrysostomou and Nikolaos Aletras	6920
<i>MILIE: Modular & Iterative Multilingual Open Information Extraction</i>	
Bhushan Kotnis, Kiril Gashteovski, Daniel Onoro Rubio, Ammar Shaker, Vanesa Rodriguez-Tembras, Makoto Takamoto, Mathias Niepert and Carolin Lawrence	6939
<i>What Makes Reading Comprehension Questions Difficult?</i>	
Saku Sugawara, Nikita Nangia, Alex Warstadt and Samuel R. Bowman	6951
<i>From Simultaneous to Streaming Machine Translation by Leveraging Streaming History</i>	
Javier Iranzo Sanchez, Jorge Civera and Alfons Juan-Císcar	6972
<i>A Rationale-Centric Framework for Human-in-the-loop Machine Learning</i>	
Jinghui Lu, Linyi Yang, Brian Mac Namee and Yue Zhang	6986
<i>Challenges and Strategies in Cross-Cultural NLP</i>	
Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust and Anders Søgaard	6997
<i>Prototypical Verbalizer for Prompt-based Few-shot Tuning</i>	
Ganqu Cui, Shengding Hu, Ning Ding, Longtao Huang and Zhiyuan Liu	7014
<i>Clickbait Spoiling via Question Answering and Passage Retrieval</i>	
Matthias Hagen, Maik Fröbe, Artur Jurk and Martin Potthast	7025
<i>BERT Learns to Teach: Knowledge Distillation with Meta Learning</i>	
Wangchunshu Zhou, Canwen Xu and Julian McAuley	7037
<i>STEMM: Self-learning with Speech-text Manifold Mixup for Speech Translation</i>	
Qingkai Fang, Rong Ye, Lei Li, Yang Feng and Mingxuan Wang	7050
<i>Integrating Vectorized Lexical Constraints for Neural Machine Translation</i>	
Shuo Wang, Zhixing Tan and Yang Liu	7063
<i>MPII: Multi-Level Mutual Promotion for Inference and Interpretation</i>	
Yan Liu, Sanyuan Chen, Yazheng Yang and Qi Dai	7074
<i>StableMoE: Stable Routing Strategy for Mixture of Experts</i>	
Damai Dai, Li Dong, Shuming Ma, Bo Zheng, Zhifang Sui, Baobao Chang and Furu Wei ..	7085
<i>Boundary Smoothing for Named Entity Recognition</i>	
Enwei Zhu and Jinpeng Li	7096
<i>Incorporating Hierarchy into Text Encoder: a Contrastive Learning Approach for Hierarchical Text Classification</i>	
Zihan Wang, Peiyi Wang, Lianzhe Huang, Xin Sun and Houfeng Wang	7109
<i>Signal in Noise: Exploring Meaning Encoded in Random Character Sequences with Character-Aware Language Models</i>	
Mark Bo Chu, Bhargav Srinivasa Desikan, Ethan O Nadler, Donald Ruggiero Lo Sardo, Elise Darragh-Ford and Douglas Richard Guilbeault	7120

<i>Hyperlink-induced Pre-training for Passage Retrieval in Open-domain Question Answering</i>	
Jiawei Zhou, Xiaoguang Li, Lifeng Shang, Lan Luo, Ke Zhan, Enrui Hu, Xinyu Zhang, Hao Jiang, Zhao Cao, Fan Yu, Xin Jiang, Qun Liu and Lei Chen	7135
<i>AdaLoGN: Adaptive Logic Graph Network for Reasoning-Based Machine Reading Comprehension</i>	
Xiao Li, Gong Cheng, Ziheng Chen, Yawei Sun and Yuzhong Qu	7147
<i>CAMERO: Consistency Regularized Ensemble of Perturbed Language Models with Weight Sharing</i>	
Chen Liang, Pengcheng He, Yelong Shen, Weizhu Chen and Tuo Zhao	7162
<i>Interpretability for Language Learners Using Example-Based Grammatical Error Correction</i>	
Masahiro Kaneko, Sho Takase, Ayana Niwa and Naoaki Okazaki	7176
<i>Rethinking Negative Sampling for Handling Missing Entity Annotations</i>	
Yangming Li, Lemao Liu and Shuming Shi	7188
<i>Distantly Supervised Named Entity Recognition via Confidence-Based Multi-Class Positive and Unlabeled Learning</i>	
Kang Zhou, Yuepei Li and Qi Li	7198
<i>UniXcoder: Unified Cross-Modal Pre-training for Code Representation</i>	
Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou and Jian Yin	7212
<i>One Country, 700+ Languages: NLP Challenges for Underrepresented Languages and Dialects in Indonesia</i>	
Alham Fikri Aji, Genta Indra Winata, Fajri Koto, Samuel Cahyawijaya, Ade Romadhony, Rahmad Mahendra, Kemal Kurniawan, David Moeljadi, Radityo Eko Prasajo, Timothy Baldwin, Jey Han Lau and Sebastian Ruder	7226
<i>Is GPT-3 Text Indistinguishable from Human Text? Scarecrow: A Framework for Scrutinizing Machine Text</i>	
Yao Dou, Maxwell Forbes, Rik Koncel-Kedziorski, Noah Smith and Yejin Choi	7250
<i>Transkimmer: Transformer Learns to Layer-wise Skim</i>	
Yue Guan, Zhengyi Li, Jingwen Leng, Zhouhan Lin and Minyi Guo	7275
<i>SkipBERT: Efficient Inference with Shallow Layer Skipping</i>	
Jue Wang, Ke Chen, Gang Chen, Lidan Shou and Julian McAuley	7287
<i>Pretraining with Artificial Language: Studying Transferable Knowledge in Language Models</i>	
Ryokan Ri and Yoshimasa Tsuruoka	7302
<i>mLUKE: The Power of Entity Representations in Multilingual Pretrained Language Models</i>	
Ryokan Ri, Ikuya Yamada and Yoshimasa Tsuruoka	7316
<i>Evaluating Factuality in Text Simplification</i>	
Ashwin Devaraj, William Berkeley Sheffield, Byron C Wallace and Junyi Jessie Li	7331
<i>Requirements and Motivations of Low-Resource Speech Synthesis for Language Revitalization</i>	
Aidan Pine, Dan Wells, Nathan Brinklow, Patrick William Littell and Korin Richmond	7346
<i>Sharpness-Aware Minimization Improves Language Model Generalization</i>	
Dara Bahri, Hossein Mobahi and Yi Tay	7360
<i>Adversarial Authorship Attribution for Deobfuscation</i>	
Wanyue Zhai, Jonathan Rusert, Zubair Shafiq and Padmini Srinivasan	7372

<i>Weakly Supervised Word Segmentation for Computational Language Documentation</i>	
Shu Okabe, Laurent Besacier and François Yvon	7385
<i>SciNLI: A Corpus for Natural Language Inference on Scientific Text</i>	
Mobashir Sadat and Cornelia Caragea	7399
<i>Neural reality of argument structure constructions</i>	
Bai Li, Zining Zhu, Guillaume Thomas, Frank Rudzicz and Yang Xu	7410
<i>On the Robustness of Offensive Language Classifiers</i>	
Jonathan Rusert, Zubair Shafiq and Padmini Srinivasan	7424
<i>Few-shot Controllable Style Transfer for Low-Resource Multilingual Settings</i>	
Kalpesh Krishna, Deepak Nathani, Xavier Garcia, Bidisha Samanta and Partha Talukdar ...	7439
<i>ABC: Attention with Bounded-memory Control</i>	
Hao Peng, Jungo Kasai, Nikolaos Pappas, Dani Yogatama, Zhaofeng Wu, Lingpeng Kong, Roy Schwartz and Noah Smith	7469
<i>The Dangers of Underclaiming: Reasons for Caution When Reporting How NLP Systems Fail</i>	
Samuel R. Bowman	7484
<i>RELiC: Retrieving Evidence for Literary Claims</i>	
Katherine Thai, Yapei Chang, Kalpesh Krishna and Mohit Iyyer	7500
<i>Analyzing Generalization of Vision and Language Navigation to Unseen Outdoor Areas</i>	
Raphael Schumann and Stefan Riezler	7519
<i>Adapting Coreference Resolution Models through Active Learning</i>	
Michelle Yuan, Patrick Xia, Chandler May, Benjamin Van Durme and Jordan Lee Boyd-Graber	7533
<i>An Imitation Learning Curriculum for Text Editing with Non-Autoregressive Models</i>	
Sweta Agrawal and Marine Carpuat	7550
<i>Memorisation versus Generalisation in Pre-trained Language Models</i>	
Michael Tänzler, Sebastian Ruder and Marek Rei	7564
<i>ChatMatch: Evaluating Chatbots by Autonomous Chat Tournaments</i>	
Ruolan Yang, Zitong Li, Haifeng Tang and Kenny Q. Zhu	7579
<i>Do self-supervised speech models develop human-like perception biases?</i>	
Juliette Millet and Ewan Dunbar	7591
<i>Vision-and-Language Navigation: A Survey of Tasks, Methods, and Future Directions</i>	
Jing Gu, Eliana Stefani, Qi Wu, Jesse Thomason and Xin Eric Wang	7606
<i>Learning to Generate Programs for Table Fact Verification via Structure-Aware Semantic Parsing</i>	
Suixin Ou and Yongmei Liu	7624
<i>Cluster & Tune: Boost Cold Start Performance in Text Classification</i>	
Eyal Shnarch, Ariel Gera, Alon Halfon, Lena Dankin, Leshem Choshen, Ranit Aharonov and Noam Slonim	7639
<i>Overcoming a Theoretical Limitation of Self-Attention</i>	
David Chiang and Peter Cholak	7654

<i>Prediction Difference Regularization against Perturbation for Neural Machine Translation</i>	
Dengji Guo, Zhengrui Ma, Min Zhang and Yang Feng	7665
<i>Make the Best of Cross-lingual Transfer: Evidence from POS Tagging with over 100 Languages</i>	
Wietse de Vries, Martijn Wieling and Malvina Nissim	7676
<i>Should a Chatbot be Sarcastic? Understanding User Preferences Towards Sarcasm Generation</i>	
Silviu Vlad Oprea, Steven R. Wilson and Walid Magdy	7686
<i>How Do Seq2Seq Models Perform on End-to-End Data-to-Text Generation?</i>	
Xunjian Yin and Xiaojun Wan	7701
<i>Probing for Labeled Dependency Trees</i>	
Max Müller-Eberstein, Rob Van Der Goot and Barbara Plank	7711
<i>DoCoGen: Domain Counterfactual Generation for Low Resource Domain Adaptation</i>	
Nitay Calderon, Eyal Ben-David, Amir Feder and Roi Reichart	7727
<i>LiLT: A Simple yet Effective Language-Independent Layout Transformer for Structured Document Understanding</i>	
Jiapeng Wang, Lianwen Jin and Kai Ding	7747
<i>Dependency-based Mixture Language Models</i>	
Zhixian Yang and Xiaojun Wan	7758
<i>Can Unsupervised Knowledge Transfer from Social Discussions Help Argument Mining?</i>	
Subhabrata Dutta, Jeevesh Juneja, Dipankar Das and Tanmoy Chakraborty	7774
<i>Entity-based Neural Local Coherence Modeling</i>	
Sungho Jeon and Michael Strube	7787
<i>That Is a Suspicious Reaction!: Interpreting Logits Variation to Detect NLP Adversarial Attacks</i>	
Edoardo Mosca, Shreyash Agarwal, Javier Rando Ramírez and Georg Groh	7806
<i>Local Languages, Third Spaces, and other High-Resource Scenarios</i>	
Steven Bird	7817
<i>That Sleepen Al the Nyght with Open Ye! Cross-era Sequence Segmentation with Switch-memory</i>	
Xuemei Tang and Qi Su	7830
<i>Fair and Argumentative Language Modeling for Computational Argumentation</i>	
Carolin Holtermann, Anne Lauscher and Simone Paolo Ponzetto	7841
<i>Learning Adaptive Segmentation Policy for End-to-End Simultaneous Translation</i>	
Ruiqing Zhang, Zhongjun He, Hua Wu and Haifeng Wang	7862
<i>Can Pre-trained Language Models Interpret Similes as Smart as Human?</i>	
Qianyu He, Sijie Cheng, Zhixu Li, Rui Xie and Yanghua Xiao	7875
<i>CBLUE: A Chinese Biomedical Language Understanding Evaluation Benchmark</i>	
Ningyu Zhang, Moshu Chen, Zhen Bi, Xiaozhuan Liang, Lei Li, Xin Shang, Kangping Yin, Chuanqi Tan, Jian Xu, Fei Huang, Luo Si, Yuan Ni, Guotong Xie, Zhifang Sui, Baobao Chang, Hui Zong, Zheng Yuan, Linfeng Li, Jun Yan, Hongying Zan, Kunli Zhang, Buzhou Tang and Qingcai Chen	7888
<i>Learning Non-Autoregressive Models from Search for Unsupervised Sentence Summarization</i>	
Puyuan Liu, Chenyang Huang and Lili Mou	7916

<i>Learning to Generalize to More: Continuous Semantic Augmentation for Neural Machine Translation</i> Xiangpeng Wei, Heng Yu, Yue Hu, Rongxiang Weng, Weihua Luo and Rong Jin	7930
<i>Lexical Knowledge Internalization for Neural Dialog Generation</i> Zhiyong Wu, Wei Bi, Xiang Li, Lingpeng Kong and Ben Kao	7945
<i>Modeling Syntactic-Semantic Dependency Correlations in Semantic Role Labeling Using Mixture Models</i> Junjie Chen, Xiangheng He and Yusuke Miyao	7959
<i>Learning the Beauty in Songs: Neural Singing Voice Beautifier</i> Jinglin Liu, Chengxi Li, Yi Ren, Zhiying Zhu and Zhou Zhao	7970
<i>A Model-agnostic Data Manipulation Method for Persona-based Dialogue Generation</i> Yu Cao, Wei Bi, Meng Fang, Shuming Shi and Dacheng Tao	7984
<i>LinkBERT: Pretraining Language Models with Document Links</i> Michihiro Yasunaga, Jure Leskovec and Percy Liang	8003
<i>Improving Time Sensitivity for Question Answering over Temporal Knowledge Graphs</i> Chao Shang, Guangtao Wang, Peng Qi and Jing Huang	8017
<i>Self-supervised Semantic-driven Phoneme Discovery for Zero-resource Speech Recognition</i> Liming Wang, Siyuan Feng, Mark A. Hasegawa-Johnson and Chang D. Yoo	8027
<i>Softmax Bottleneck Makes Language Models Unable to Represent Multi-mode Word Distributions</i> Haw-Shiuan Chang and Andrew McCallum	8048
<i>Ditch the Gold Standard: Re-evaluating Conversational Question Answering</i> Huihan Li, Tianyu Gao, Manan Goenka and Danqi Chen	8074
<i>Fantastically Ordered Prompts and Where to Find Them: Overcoming Few-Shot Prompt Order Sensitivity</i> Yao Lu, Max Bartolo, Alastair Philip Moore, Sebastian Riedel and Pontus Stenetorp	8086
<i>Situated Dialogue Learning through Procedural Environment Generation</i> Prithviraj Ammanabrolu, Renee Jia and Mark Riedl	8099
<i>UniTE: Unified Translation Evaluation</i> Yu Wan, Dayiheng Liu, Baosong Yang, Haibo Zhang, Boxing Chen, Derek F. Wong and Lidia S. Chao	8117
<i>Program Transfer for Answering Complex Questions over Knowledge Bases</i> Shulin Cao, Jiaxin Shi, Zijun Yao, Xin Lv, Jifan Yu, Lei Hou, Juanzi Li, Zhiyuan Liu and Jinghui Xiao	8128
<i>EAG: Extract and Generate Multi-way Aligned Corpus for Complete Multi-lingual Neural Machine Translation</i> Yulin Xu, Zhen Yang, Fandong Meng and Jie Zhou	8141
<i>Using Context-to-Vector with Graph Retrofitting to Improve Word Embeddings</i> Jiangbin Zheng, Yile Wang, Ge Wang, Jun Xia, Yufei Huang, Guojiang Zhao, Yue Zhang and Stan Z. Li	8154
<i>Multimodal Sarcasm Target Identification in Tweets</i> Jiquan Wang, Lin Sun, Yi Liu, Meizhi Shao and Zengwei Zheng	8164

<i>Flexible Generation from Fragmentary Linguistic Input</i> Peng Qian and Roger P. Levy	8176
<i>Revisiting Over-Smoothness in Text to Speech</i> Yi Ren, Xu Tan, Tao Qin, Zhou Zhao and Tie-Yan Liu	8197
<i>Coherence boosting: When your pretrained language model is not paying enough attention</i> Nikolay Malkin, Zhen Wang and Nebojsa Jojic	8214
<i>Uncertainty Estimation of Transformer Predictions for Misclassification Detection</i> Artem Vazhentsev, Gleb Kuzmin, Artem Shelmanov, Akim Tsvigun, Evgenii Tsymbalov, Kirill Fedyanin, Maxim Panov, Alexander Panchenko, Gleb Gennadjevich Gusev, Mikhail Burtsev, Manvel Avetisian and Leonid Zhukov	8237
<i>VALSE: A Task-Independent Benchmark for Vision and Language Models Centered on Linguistic Phenomena</i> Letitia Parcalabescu, Michele Cafagna, Lilitta Muradjan, Anette Frank, Iacer Calixto and Albert Gatt	8253
<i>The Grammar-Learning Trajectories of Neural Language Models</i> Leshem Choshen, Guy Hacohen, Daphna Weinshall and Omri Abend	8281
<i>Generating Scientific Definitions with Controllable Complexity</i> Tal August, Katharina Reinecke and Noah Smith	8298
<i>Label Semantic Aware Pre-training for Few-shot Text Classification</i> Aaron Mueller, Jason Krone, Salvatore Romeo, Saab Mansour, Elman Mansimov, Yi Zhang and Dan Roth	8318
<i>ODE Transformer: An Ordinary Differential Equation-Inspired Model for Sequence Generation</i> Bei Li, Quan Du, Tao Zhou, Yi Jing, Shuhan Zhou, Xin Zeng, Tong Xiao, JingBo Zhu, Xuebo Liu and Min Zhang	8335
<i>A Comparison of Strategies for Source-Free Domain Adaptation</i> Xin Su, Yiyun Zhao and Steven Bethard	8352
<i>Ethics Sheets for AI Tasks</i> Saif M. Mohammad	8368
<i>Learning Disentangled Representations of Negation and Uncertainty</i> Jake A Vasilakes, Chrysoula Zerva, Makoto Miwa and Sophia Ananiadou	8380
<i>latent-GLAT: Glancing at Latent Variables for Parallel Text Generation</i> Yu Bao, Hao Zhou, Shujian Huang, Dongqi Wang, Lihua Qian, Xinyu Dai, Jiajun Chen and Lei Li	8398
<i>PPT: Pre-trained Prompt Tuning for Few-shot Learning</i> Yuxian Gu, Xu Han, Zhiyuan Liu and Minlie Huang	8410
<i>Deduplicating Training Data Makes Language Models Better</i> Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch and Nicholas Carlini	8424
<i>Improving the Generalizability of Depression Detection by Leveraging Clinical Questionnaires</i> Thong Nguyen, Andrew Yates, Ayah Zirikly, Bart Desmet and Arman Cohan	8446
<i>Internet-Augmented Dialogue Generation</i> Mojtaba Komeili, Kurt Shuster and Jason E Weston	8460

<i>SUPERB-SG: Enhanced Speech processing Universal PERFORMANCE Benchmark for Semantic and Generative Capabilities</i>	
Hsiang-Sheng Tsai, Heng-Jui Chang, Wen-Chin Huang, Zili Huang, Kushal Lakhotia, Shu-wen Yang, Shuyan Dong, Andy T. Liu, Cheng-I Lai, Jiatong Shi, Xuankai Chang, Phil Hall, Hsuan-Jui Chen, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed and Hung-yi Lee	8479
<i>Knowledge Neurons in Pretrained Transformers</i>	
Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang and Furu Wei	8493
<i>Meta-Learning for Fast Cross-Lingual Adaptation in Dependency Parsing</i>	
Anna Langedijk, Verna Dankers, Phillip Lippe, Sander Bos, Bryan Cardenas Guevara, Helen Yannakoudakis and Ekaterina Shutova	8503
<i>French CrowS-Pairs: Extending a challenge dataset for measuring social bias in masked language models to a language other than English</i>	
Aurélie Névéol, Yoann Dupont, Julien Bezançon and Karën Fort.	8521
<i>Few-Shot Learning with Siamese Networks and Label Tuning</i>	
Thomas Müller, Guillermo Pérez-Torró and Marc Franco-Salvador	8532
<i>Inferring Rewards from Language in Context</i>	
Jessy Lin, Daniel Fried, Dan Klein and Anca Dragan	8546
<i>Generating Biographies on Wikipedia: The Impact of Gender Bias on the Retrieval-Based Generation of Women Biographies</i>	
Angela Fan and Claire Gardent	8561
<i>Your Answer is Incorrect... Would you like to know why? Introducing a Bilingual Short Answer Feedback Dataset</i>	
Anna Filighera, Siddharth Singh Parihar, Tim Steuer, Tobias Meuser and Sebastian Ochs	8577
<i>Towards Better Characterization of Paraphrases</i>	
Timothy Liu and De Wen Soh	8592
<i>SummScreen: A Dataset for Abstractive Screenplay Summarization</i>	
Mingda Chen, Zewei Chu, Sam Wiseman and Kevin Gimpel	8602
<i>Sparsifying Transformer Models with Trainable Representation Pooling</i>	
Michał Pietruszka, Łukasz Borchmann and Łukasz Garncarek	8616
<i>Uncertainty Determines the Adequacy of the Mode and the Tractability of Decoding in Sequence-to-Sequence Models</i>	
Felix Stahlberg, Ilia Kulikov and Shankar Kumar	8634
<i>FlipDA: Effective and Robust Data Augmentation for Few-Shot Learning</i>	
Jing Zhou, Yanan Zheng, Jie Tang, Li Jian and Zhilin Yang	8646
<i>Text-Free Prosody-Aware Generative Spoken Language Modeling</i>	
Eugene Kharitonov, Ann Lee, Adam Polyak, Yossi Adi, Jade Copet, Kushal Lakhotia, Tu Anh Nguyen, Morgane Riviere, Abdelrahman Mohamed, Emmanuel Dupoux and Wei-Ning Hsu	8666
<i>Lite Unified Modeling for Discriminative Reading Comprehension</i>	
Yilin Zhao, Hai Zhao, Libin Shen and Yinggong Zhao	8682
<i>Bilingual alignment transfers to multilingual alignment for unsupervised parallel text mining</i>	
Chih-chan Tien and Shane Steinert-Threlkeld	8696

<i>End-to-End Modeling via Information Tree for One-Shot Natural Language Spatial Video Grounding</i>	
Mengze Li, Tianbao Wang, Haoyu Zhang, Shengyu Zhang, Zhou Zhao, Jiaxu Miao, Wenqiao Zhang, Wenming Tan, Jin Wang, Peng Wang, Shiliang Pu and Fei Wu	8707
<i>RNSum: A Large-Scale Dataset for Automatic Release Note Generation via Commit Logs Summarization</i>	
Hisashi Kamezawa, Noriki Nishida, Nobuyuki Shimizu, Takashi Miyazaki and Hideki Nakayama	8718
<i>Improving Machine Reading Comprehension with Contextualized Commonsense Knowledge</i>	
Kai Sun, Dian Yu, Jianshu Chen, Dong Yu and Claire Cardie	8736
<i>Modeling Persuasive Discourse to Adaptively Support Students' Argumentative Writing</i>	
Thiemo Wambsganss and Christina Niklaus	8748
<i>Active Evaluation: Efficient NLG Evaluation with Few Pairwise Comparisons</i>	
Akash Kumar Mohankumar and Mitesh M Khapra	8761
<i>The Moral Debater: A Study on the Computational Generation of Morally Framed Arguments</i>	
Milad Alshomary, Roxanne El Baff, Timon Gurcke and Henning Wachsmuth	8782
<i>Pyramid-BERT: Reducing Complexity via Successive Core-set based Token Selection</i>	
Xin Huang, Ashish Khetan, Rene Bidart and Zohar Karnin	8798
<i>Probing for the Usage of Grammatical Number</i>	
Karim Lasri, Tiago Pimentel, Alessandro Lenci, Thierry Poibeau and Ryan D Cotterell	8818