

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

Findings of ACL 2022

Dublin, Ireland
22 – 27 May 2022

Volume 1 of 6

ISBN: 978-1-7138-6726-5

Printed from e-media with permission by:

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



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

Copyright© (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>Is Whole Word Masking Always Better for Chinese BERT?: Probing on Chinese Grammatical Error Correction</i>	
Yong Dai, Linyang Li, Cong Zhou, Zhangyin Feng, Enbo Zhao, Xipeng Qiu, Piji Li and Duyu Tang	1
<i>Compilable Neural Code Generation with Compiler Feedback</i>	
Xin Wang, Yasheng Wang, Yao Wan, Fei Mi, Yitong Li, Pingyi Zhou, Jin Liu, Hao Wu, Xin Jiang and Qun Liu	9
<i>Towards Unifying the Label Space for Aspect- and Sentence-based Sentiment Analysis</i>	
Yiming Zhang, Min Zhang, Sai Wu and Junbo Zhao	20
<i>Input-specific Attention Subnetworks for Adversarial Detection</i>	
Emil Biju, Anirudh Sriram, Pratyush Kumar and Mitesh M Khapra	31
<i>RelationPrompt: Leveraging Prompts to Generate Synthetic Data for Zero-Shot Relation Triplet Extraction</i>	
Yew Ken Chia, Lidong Bing, Soujanya Poria and Luo Si	45
<i>Pre-Trained Multilingual Sequence-to-Sequence Models: A Hope for Low-Resource Language Translation?</i>	
En-Shiun Annie Lee, Sarubi Thillainathan, Shravan Nayak, Surangika Ranathunga, David Ifeoluwa Adelani, Ruisi Su and Arya D. McCarthy	58
<i>Multi-Scale Distribution Deep Variational Autoencoder for Explanation Generation</i>	
ZeFeng Cai, Linlin Wang, Gerard de Melo, Fei Sun and Liang He	68
<i>Dual Context-Guided Continuous Prompt Tuning for Few-Shot Learning</i>	
Jie Zhou, Le Tian, Houjin Yu, Zhou Xiao, Hui Su and Jie Zhou	79
<i>Extract-Select: A Span Selection Framework for Nested Named Entity Recognition with Generative Adversarial Training</i>	
Peixin Huang, Xiang Zhao, Minghao Hu, Yang Fang, Xinyi Li and Weidong Xiao	85
<i>Controlled Text Generation Using Dictionary Prior in Variational Autoencoders</i>	
Xianghong Fang, Jian Li, Lifeng Shang, Xin Jiang, Qun Liu and Dit-Yan Yeung	97
<i>Challenges to Open-Domain Constituency Parsing</i>	
Sen Yang, Leyang Cui, Ruoxi Ning, Di Wu and Yue Zhang	112
<i>Going Deeper: Structured Sememe Prediction via Transformer with Tree Attention</i>	
Yining Ye, Fanchao Qi, Zhiyuan Liu and Maosong Sun	128
<i>Table-based Fact Verification with Self-adaptive Mixture of Experts</i>	
Yuxuan Zhou, Xien Liu, Kaiyin Zhou and Ji Wu	139
<i>Investigating Data Variance in Evaluations of Automatic Machine Translation Metrics</i>	
Jiannan Xiang, Huayang Li, Yahui Liu, Lemao Liu, Guoping Huang, Defu Lian and Shuming Shi	150
<i>Sememe Prediction for BabelNet Synsets using Multilingual and Multimodal Information</i>	
Fanchao Qi, Chuancheng Lv, Zhiyuan Liu, Xiaojun Meng, Maosong Sun and Hai-Tao Zheng	158

<i>Query and Extract: Refining Event Extraction as Type-oriented Binary Decoding</i> Sijia Wang, Mo Yu, Shiyu Chang, Lichao Sun and Lifu Huang	169
<i>LEVEN: A Large-Scale Chinese Legal Event Detection Dataset</i> Feng Yao, Chaojun Xiao, Xiaozhi Wang, Zhiyuan Liu, Lei Hou, Cunchao Tu, Juanzi Li, Yun Liu, Weixing Shen and Maosong Sun	183
<i>Analyzing Dynamic Adversarial Training Data in the Limit</i> Eric Wallace, Adina Williams, Robin Jia and Douwe Kiela	202
<i>AbductionRules: Training Transformers to Explain Unexpected Inputs</i> Nathan Young, Qiming Bao, Joshua Bensemann and Michael J. Witbrock	218
<i>On the Importance of Data Size in Probing Fine-tuned Models</i> Houman Mehrafarin, Sara Rajaei and Mohammad Taher Pilehvar	228
<i>RuCCoN: Clinical Concept Normalization in Russian</i> Alexandr Nesterov, Galina Zubkova, Zulfat Miftahutdinov, Vladimir Kokh, Elena Tutubalina, Artem Shelmanov, Anton M. Alekseev, Manvel Avetisian, Andrey Chertok and Sergey Nikolenko .	239
<i>A Sentence is Worth 128 Pseudo Tokens: A Semantic-Aware Contrastive Learning Framework for Sen- tence Embeddings</i> Haochen Tan, Wei Shao, Han Wu, Ke Yang and Linqi Song	246
<i>Eider: Empowering Document-level Relation Extraction with Efficient Evidence Extraction and Inference- stage Fusion</i> Yiqing Xie, Jiaming Shen, Sha Li, Yuning Mao and Jiawei Han	257
<i>Meta-X_NLG: A Meta-Learning Approach Based on Language Clustering for Zero-Shot Cross-Lingual Transfer and Generation</i> Kaushal Kumar Maurya and Maunendra Sankar Desarkar	269
<i>MR-P: A Parallel Decoding Algorithm for Iterative Refinement Non-Autoregressive Translation</i> Hao Cheng and Zhihua Zhang	285
<i>Open Relation Modeling: Learning to Define Relations between Entities</i> Jie Huang, Kevin Chang, Jinjun Xiong and Wen-mei Hwu	297
<i>A Slot Is Not Built in One Utterance: Spoken Language Dialogs with Sub-Slots</i> Sai Zhang, Yuwei Hu, Yuchuan Wu, Jiaman Wu, Yongbin Li, Jian Sun, Caixia Yuan and Xiaojie Wang	309
<i>Towards Transparent Interactive Semantic Parsing via Step-by-Step Correction</i> Lingbo Mo, Ashley Lewis, Huan Sun and Michael White	322
<i>MINER: Multi-Interest Matching Network for News Recommendation</i> Jian Li, Jieming Zhu, Qiwei Bi, Guohao Cai, Lifeng Shang, Zhenhua Dong, Xin Jiang and Qun Liu	343
<i>KSAM: Infusing Multi-Source Knowledge into Dialogue Generation via Knowledge Source Aware Multi-Head Decoding</i> Sixing Wu, Ying Li, Dawei Zhang and Zhonghai Wu	353
<i>Towards Responsible Natural Language Annotation for the Varieties of Arabic</i> A. Stevie Bergman and Mona T. Diab	364
<i>Dynamically Refined Regularization for Improving Cross-corpora Hate Speech Detection</i> Tulika Bose, Nikolaos Aletras, Irina Illina and Dominique Fohr	372

<i>Towards Large-Scale Interpretable Knowledge Graph Reasoning for Dialogue Systems</i> Yi-Lin Tuan, Sajjad Beygi, Maryam Fazel-Zarandi, Qiaozi Gao, Alessandra Cervone and William Yang Wang	383
<i>MDERank: A Masked Document Embedding Rank Approach for Unsupervised Keyphrase Extraction</i> Linhan Zhang, Qian Chen, Wen Wang, Chong Deng, ShiLiang Zhang, Bing Li, Wei Wang and Xin Cao	396
<i>Visualizing the Relationship Between Encoded Linguistic Information and Task Performance</i> Jiannan Xiang, Huayang Li, Defu Lian, Guoping Huang, Taro Watanabe and Lemao Liu	410
<i>Efficient Argument Structure Extraction with Transfer Learning and Active Learning</i> Xinyu Hua and Lu Wang	423
<i>Plug-and-Play Adaptation for Continuously-updated QA</i> Kyungjae Lee, Wookje Han, Seung-won Hwang, Hwaran Lee, Joonsuk Park and Sang-Woo Lee	438
<i>Reinforced Cross-modal Alignment for Radiology Report Generation</i> Han Qin and Yan Song	448
<i>What Works and Doesn't Work, A Deep Decoder for Neural Machine Translation</i> Zuchao Li, Yiran Wang, Masao Utiyama, Eiichiro Sumita, Hai Zhao and Taro Watanabe	459
<i>SyMCoM - Syntactic Measure of Code Mixing A Study Of English-Hindi Code-Mixing</i> Prashant Kodali, Anmol Goel, Monojit Choudhury, Manish Shrivastava and Ponnurangam Kumaraguru	472
<i>HybriDialogue: An Information-Seeking Dialogue Dataset Grounded on Tabular and Textual Data</i> Kai Nakamura, Sharon Levy, Yi-Lin Tuan, Wenhui Chen and William Yang Wang	481
<i>NEWS: A Corpus for News Topic-Focused Summarization</i> Seyed Ali Bahrainian, Sheridan Feucht and Carsten Eickhoff	493
<i>Classification without (Proper) Representation: Political Heterogeneity in Social Media and Its Implications for Classification and Behavioral Analysis</i> Kenan Alkiek, Bohan Zhang and David Jurgens	504
<i>Toward More Meaningful Resources for Lower-resourced Languages</i> Constantine Lignos, Nolan Holley, Chester Palen-Michel and Jonne Sälevä	523
<i>Better Quality Estimation for Low Resource Corpus Mining</i> Muhammed Yusuf Kocyigit, Jiho Lee and Derry Wijaya	533
<i>End-to-End Segmentation-based News Summarization</i> Yang Liu, Chenguang Zhu and Michael Zeng	544
<i>Fast Nearest Neighbor Machine Translation</i> Yuxian Meng, Xiaoya Li, Xiayu Zheng, Fei Wu, Xiaofei Sun, Tianwei Zhang and Jiwei Li ..	555
<i>Extracting Latent Steering Vectors from Pretrained Language Models</i> Nishant Subramani, Nivedita Suresh and Matthew E Peters	566
<i>Domain Generalisation of NMT: Fusing Adapters with Leave-One-Domain-Out Training</i> Thuy-Trang Vu, Shahram Khadivi, Dinh Phung and Gholamreza Haffari	582
<i>Reframing Instructional Prompts to GPTk's Language</i> Daniel Khashabi, Chitta Baral, Yejin Choi and Hannaneh Hajishirzi	589

<i>Read Top News First: A Document Reordering Approach for Multi-Document News Summarization</i> Chao Zhao, Tenghao Huang, Somnath Basu Roy Chowdhury, Muthu Kumar Chandrasekaran, Kathleen McKeown and Snigdha Chaturvedi	613
<i>Human Language Modeling</i> Nikita Soni, Matthew Matero, Niranjana Balasubramanian and H. Schwartz	622
<i>Inverse is Better! Fast and Accurate Prompt for Few-shot Slot Tagging</i> Yutai Hou, Cheng Chen, Xianzhen Luo, Bohan Li and Wanxiang Che	637
<i>Cross-Modal Cloze Task: A New Task to Brain-to-Word Decoding</i> Shuxian Zou, Shaonan Wang, Jiajun Zhang and Chengqing Zong	648
<i>Mitigating Gender Bias in Distilled Language Models via Counterfactual Role Reversal</i> Umang Gupta, Jwala Dhamala, Varun Kumar, Apurv Verma, Yada Pruksachatkun, Satyapriya Krishna, Rahul Gupta, Kai-Wei Chang, Greg Ver Steeg and Aram Galstyan	658
<i>Domain Representative Keywords Selection: A Probabilistic Approach</i> Pritom Saha Akash, Jie Huang, Kevin Chang, Yunyao Li, Lucian Popa and ChengXiang Zhai	679
<i>Hierarchical Inductive Transfer for Continual Dialogue Learning</i> Shaoxiong Feng, Xuancheng Ren, Kan Li and Xu Sun	693
<i>Why Exposure Bias Matters: An Imitation Learning Perspective of Error Accumulation in Language Generation</i> Kushal Arora, Layla El Asri, Hareesh Bahuleyan and Jackie CK Cheung	700
<i>Question Answering Infused Pre-training of General-Purpose Contextualized Representations</i> Robin Jia, Mike Lewis and Luke Zettlemoyer	711
<i>Automatic Song Translation for Tonal Languages</i> Fenfei Guo, Chen Zhang, Zhirui Zhang, Qixin He, Kejun Zhang, Jun Xie and Jordan Lee Boyd- Graber	729
<i>Read before Generate! Faithful Long Form Question Answering with Machine Reading</i> Dan SU, Xiaoguang Li, Jindi Zhang, Lifeng Shang, Xin Jiang, Qun Liu and Pascale Fung . . .	744
<i>A Simple yet Effective Relation Information Guided Approach for Few-Shot Relation Extraction</i> Yang Liu, Jinpeng Hu, Xiang Wan and Tsung-Hui Chang	757
<i>MIMICause: Representation and automatic extraction of causal relation types from clinical notes</i> Vivek Khetan, Md Imbesat Hassan Rizvi, Jessica Huber, Paige Bartusiak, Bogdan Eugen Sacaleanu and Andrew Fano	764
<i>Compressing Sentence Representation for Semantic Retrieval via Homomorphic Projective Distillation</i> Xuandong Zhao, Zhiguo Yu, Ming Wu and Lei Li	774
<i>Debiasing Event Understanding for Visual Commonsense Tasks</i> Minji Seo, YeonJoon Jung, Seungtaek Choi, Seung-won Hwang and Bei Liu	782
<i>Fact-Tree Reasoning for N-ary Question Answering over Knowledge Graphs</i> Yao Zhang, Peiyao Li, Hongru Liang, Adam Jatowt and Zhenglu Yang	788
<i>DeepStruct: Pretraining of Language Models for Structure Prediction</i> Chenguang Wang, Xiao Liu, Zui Chen, Haoyun Hong, Jie Tang and Dawn Song	803
<i>The Change that Matters in Discourse Parsing: Estimating the Impact of Domain Shift on Parser Error</i> Katherine Atwell, Anthony Sicilia, Seong Jae Hwang and Malihe Alikhani	824

<i>Mukayese: Turkish NLP Strikes Back</i>	
Ali Safaya, Emirhan Kurtuluş, Arda Goktogan and Deniz Yuret	846
<i>Virtual Augmentation Supported Contrastive Learning of Sentence Representations</i>	
Dejiao Zhang, Wei Xiao, Henghui Zhu, Xiaofei Ma and Andrew Arnold	864
<i>MoEfication: Transformer Feed-forward Layers are Mixtures of Experts</i>	
Zhengyan Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun and Jie Zhou	877
<i>DS-TOD: Efficient Domain Specialization for Task-Oriented Dialog</i>	
Chia-Chien Hung, Anne Lauscher, Simone Paolo Ponzetto and Goran Glavaš	891
<i>Distinguishing Non-natural from Natural Adversarial Samples for More Robust Pre-trained Language Model</i>	
Jiayi Wang, Rongzhou Bao, Zhuosheng Zhang and Hai Zhao	905
<i>Learning Adaptive Axis Attentions in Fine-tuning: Beyond Fixed Sparse Attention Patterns</i>	
Zihan Wang, Jiuxiang Gu, Jason Kuen, Handong Zhao, Vlad I Morariu, Ruiyi Zhang, Ani Nenkova, Tong Sun and Jingbo Shang	916
<i>Using Interactive Feedback to Improve the Accuracy and Explainability of Question Answering Systems Post-Deployment</i>	
Zichao Li, Prakhar Sharma, Xing Han Lu, Jackie CK Cheung and Siva Reddy	926
<i>To be or not to be an Integer? Encoding Variables for Mathematical Text</i>	
Deborah Ferreira, Mokanarangan Thayaparan, Marco Valentino, Julia Rozanova and Andre Freitas	938
<i>GRS: Combining Generation and Revision in Unsupervised Sentence Simplification</i>	
Mohammad Dehghan, Dhruv Kumar and Lukasz Golab	949
<i>BPE vs. Morphological Segmentation: A Case Study on Machine Translation of Four Polysynthetic Languages</i>	
Manuel Mager, Arturo Oncevay, Elisabeth Mager, Katharina Kann and Thang Vu	961
<i>Distributed NLI: Learning to Predict Human Opinion Distributions for Language Reasoning</i>	
Xiang Zhou, Yixin Nie and Mohit Bansal	972
<i>Morphological Processing of Low-Resource Languages: Where We Are and What's Next</i>	
Adam Wiemerslage, Miikka Silfverberg, Changbing Yang, Arya D. McCarthy, Garrett Nicolai, Eliana Colunga and Katharina Kann	988
<i>Learning and Evaluating Character Representations in Novels</i>	
Naoya Inoue, Charuta Pethe, Allen Kim and Steven Skiena	1008
<i>Answer Uncertainty and Unanswerability in Multiple-Choice Machine Reading Comprehension</i>	
Vatsal Raina and Mark Gales	1020
<i>Measuring the Language of Self-Disclosure across Corpora</i>	
Ann-Katrin Reuel, Sebastian Peralta, João Sedoc, Garrick Sherman and Lyle Ungar	1035
<i>When Chosen Wisely, More Data Is What You Need: A Universal Sample-Efficient Strategy For Data Augmentation</i>	
Ehsan Kamalloo, Mehdi Rezagholizadeh and Ali Ghodsi	1048
<i>Explaining Classes through Stable Word Attributions</i>	
Samuel Rönnqvist, Aki-Juhani Kyröläinen, Amanda Myntti, Filip Ginter and Veronika Laippala	1063

<i>What to Learn, and How: Toward Effective Learning from Rationales</i> Samuel Carton, Surya Kanoria and Chenhao Tan	1075
<i>Listening to Affected Communities to Define Extreme Speech: Dataset and Experiments</i> Antonis Maronikolakis, Axel Wisioerek, Leah Nann, Haris Jabbar, Sahana Udupa and Hinrich Schuetze	1089
<i>Entropy-based Attention Regularization Frees Unintended Bias Mitigation from Lists</i> Giuseppe Attanasio, Debora Nozza, Dirk Hovy and Elena Baralis	1105
<i>From BERT's Point of View: Revealing the Prevailing Contextual Differences</i> Carolin M. Schuster and Simon Hegelich	1120
<i>Learning Bias-reduced Word Embeddings Using Dictionary Definitions</i> Haozhe An, Xiaojiang Liu and Donald Zhang	1139
<i>Knowledge Graph Embedding by Adaptive Limit Scoring Loss Using Dynamic Weighting Strategy</i> Jinfa Yang, Xianghua Ying, Yongjie Shi, Xin Tong, Ruibin Wang, Taiyan Chen and Bowei Xing	1153
<i>OCR Improves Machine Translation for Low-Resource Languages</i> Oana Ignat, Jean Maillard, Vishrav Chaudhary and Francisco Guzmán	1164
<i>CoCoLM: Complex Commonsense Enhanced Language Model with Discourse Relations</i> Changlong Yu, Hongming Zhang, Yangqiu Song and Wilfred Ng	1175
<i>Learning to Robustly Aggregate Labeling Functions for Semi-supervised Data Programming</i> Ayush Maheshwari, Krishnateja Killamsetty, Ganesh Ramakrishnan, Rishabh K Iyer, Marina Danilevsky and Lucian Popa	1188
<i>Multi-Granularity Semantic Aware Graph Model for Reducing Position Bias in Emotion Cause Pair Extraction</i> Yinan Bao, Qianwen Ma, Lingwei Wei, Wei Zhou and Songlin Hu	1203
<i>Cross-lingual Inference with A Chinese Entailment Graph</i> Tianyi Li, Sabine Weber, Mohammad Javad Hosseini, Liane Guillou and Mark Steedman ..	1214
<i>Multi-task Learning for Paraphrase Generation With Keyword and Part-of-Speech Reconstruction</i> Xuhang Xie, Xuesong Lu and Bei Chen	1234
<i>MDCSpell: A Multi-task Detector-Corrector Framework for Chinese Spelling Correction</i> Chenxi Zhu, Ziqiang Ying, Boyu Zhang and Feng Mao	1244
<i>S²SQL: Injecting Syntax to Question-Schema Interaction Graph Encoder for Text-to-SQL Parsers</i> Binyuan Hui, Ruiying Geng, Lihan Wang, Bowen Qin, Yanyang Li, Bowen Li, Jian Sun and Yongbin Li	1254
<i>Constructing Open Cloze Tests Using Generation and Discrimination Capabilities of Transformers</i> Mariano Felice, Shiva Taslimipoor and Paula Buttery	1263
<i>Co-training an Unsupervised Constituency Parser with Weak Supervision</i> Nickil Maveli and Shay B Cohen	1274
<i>HiStruct+: Improving Extractive Text Summarization with Hierarchical Structure Information</i> Qian Ruan, Malte Ostendorff and Georg Rehm	1292
<i>An Isotropy Analysis in the Multilingual BERT Embedding Space</i> Sara Rajaei and Mohammad Taher Pilehvar	1309

<i>Multi-Stage Prompting for Knowledgeable Dialogue Generation</i>	
Zihan Liu, Mostofa Patwary, Ryan Prenger, Shrimai Prabhumoye, Wei Ping, Mohammad Shoeybi and Bryan Catanzaro	1317
<i>DuReader_vis: A Chinese Dataset for Open-domain Document Visual Question Answering</i>	
Le Qi, Shangwen Lv, Hongyu Li, Jing Liu, Yu Zhang, Qiaoqiao She, Hua Wu, Haifeng Wang and Ting Liu	1338
<i>Coloring the Blank Slate: Pre-training Imparts a Hierarchical Inductive Bias to Sequence-to-sequence Models</i>	
Aaron Mueller, Robert Frank, Tal Linzen, Luheng Wang and Sebastian Schuster	1352
<i>C³KG: A Chinese Commonsense Conversation Knowledge Graph</i>	
Dawei Li, Yanran Li, Jiayi Zhang, Ke Li, Chen Wei, Jianwei Cui and Bin Wang	1369
<i>Graph Neural Networks for Multiparallel Word Alignment</i>	
Ayyoob Imani, Lütfi Kerem Senel, Masoud Jalili Sabet, François Yvon and Hinrich Schuetze	1384
<i>Sentiment Word Aware Multimodal Refinement for Multimodal Sentiment Analysis with ASR Errors</i>	
Yang Wu, Yanyan Zhao, Hao Yang, Song Chen, Bing Qin, Xiaohuan Cao and Wenting Zhao	1397
<i>A Novel Framework Based on Medical Concept Driven Attention for Explainable Medical Code Prediction via External Knowledge</i>	
Tao Wang, Linhai Zhang, Chenchen Ye, Junxi Liu and Deyu Zhou	1407
<i>Effective Unsupervised Constrained Text Generation based on Perturbed Masking</i>	
Yingwen Fu, Wenjie Ou, Zhou Yu and Yue Lin	1417
<i>Combining (Second-Order) Graph-Based and Headed-Span-Based Projective Dependency Parsing</i>	
Songlin Yang and Kewei Tu	1428
<i>End-to-End Speech Translation for Code Switched Speech</i>	
Orion Weller, Matthias Sperber, Telmo Pires, Hendra Setiawan, Christian Gollan, Dominic C Telaar and Matthias Paulik	1435
<i>A Transformational Biencoder with In-Domain Negative Sampling for Zero-Shot Entity Linking</i>	
Kai Sun, Richong Zhang, Samuel Mensah, Yongyi Mao and Xudong Liu	1449
<i>Finding the Dominant Winning Ticket in Pre-Trained Language Models</i>	
Zhuocheng Gong, Di He, Yelong Shen, Tie-Yan Liu, Weizhu Chen, Dongyan Zhao, Ji-Rong Wen and Rui Yan	1459
<i>Thai Nested Named Entity Recognition Corpus</i>	
Weerayut Buaphet, Can Udomcharoenchaikit, Peerat Limkonchotiwat, Attapol Rutherford and Sarana Nutanong	1473
<i>Two-Step Question Retrieval for Open-Domain QA</i>	
Yeon Seonwoo, Juhee Son, Jiho Jin, Sang-Woo Lee, Ji-Hoon Kim, Jung-Woo Ha and Alice Oh	1487
<i>Semantically Distributed Robust Optimization for Vision-and-Language Inference</i>	
Tejas Gokhale, Abhishek Chaudhary, Pratyay Banerjee, Chitta Baral and Yezhou Yang	1493
<i>Learning from Missing Relations: Contrastive Learning with Commonsense Knowledge Graphs for Commonsense Inference</i>	
Yong-Ho Jung, Jun-Hyung Park, Joon-Young Choi, Mingyu Lee, Junho Kim, Kang-Min Kim and SangKeun Lee	1514

<i>Capture Human Disagreement Distributions by Calibrated Networks for Natural Language Inference</i> Yuxia Wang, Minghan Wang, Yimeng Chen, Shimin Tao, Jiaxin Guo, Chang Su, Min Zhang and Hao Yang	1524
<i>Efficient, Uncertainty-based Moderation of Neural Networks Text Classifiers</i> Jakob Smedegaard Andersen and Walid Maalej	1536
<i>Revisiting Automatic Evaluation of Extractive Summarization Task: Can We Do Better than ROUGE?</i> Mousumi Akter, Naman Bansal and Shubhra Kanti Karmaker	1547
<i>Open Vocabulary Extreme Classification Using Generative Models</i> Daniel Simig, Fabio Petroni, Pouya Yanki, Kashyap Popat, Christina Du, Sebastian Riedel and Majid Yazdani	1561
<i>Decomposed Meta-Learning for Few-Shot Named Entity Recognition</i> Tingting Ma, Huiqiang Jiang, Qianhui Wu, Tiejun Zhao and Chin-Yew Lin	1584
<i>TegTok: Augmenting Text Generation via Task-specific and Open-world Knowledge</i> Chao-Hong Tan, Jia-Chen Gu, Chongyang Tao, Zhen-Hua Ling, Can Xu, Huang Hu, Xiubo Geng and Daxin Jiang	1597
<i>EmoCaps: Emotion Capsule based Model for Conversational Emotion Recognition</i> Zaijing Li, Fengxiao Tang, Ming Zhao and Yusen Zhu	1610
<i>Logic-Driven Context Extension and Data Augmentation for Logical Reasoning of Text</i> Siyuan Wang, Wanjun Zhong, Duyu Tang, Zhongyu Wei, Zhihao Fan, Daxin Jiang, Ming Zhou and Nan Duan	1619
<i>Transfer Learning and Prediction Consistency for Detecting Offensive Spans of Text</i> Amir Pouran Ben Veysch, Ning Xu, Quan Hung Tran, Varun Manjunatha, Franck Deroncourt and Thien Huu Nguyen	1630
<i>Learning Reasoning Patterns for Relational Triple Extraction with Mutual Generation of Text and Graph</i> Yubo Chen, Yunqi Zhang and Yongfeng Huang	1638
<i>Document-Level Event Argument Extraction via Optimal Transport</i> Amir Pouran Ben Veysch, Minh Van Nguyen, Franck Deroncourt, Bonan Min and Thien Huu Nguyen	1648
<i>N-Shot Learning for Augmenting Task-Oriented Dialogue State Tracking</i> Ibrahim Taha Aksu, Zhengyuan Liu, Min-Yen Kan and Nancy F. Chen	1659
<i>Document-Level Relation Extraction with Adaptive Focal Loss and Knowledge Distillation</i> Qingyu Tan, Ruidan He, Lidong Bing and Hwee Tou Ng	1672
<i>Calibration of Machine Reading Systems at Scale</i> Shehzaad Zuzar Dhuliawala, Leonard Adolphs, Rajarshi Das and Mrinmaya Sachan	1682
<i>Towards Adversarially Robust Text Classifiers by Learning to Reweight Clean Examples</i> Jianhan Xu, Cenyuan Zhang, Xiaoqing Zheng, Linyang Li, Cho-Jui Hsieh, Kai-Wei Chang and Xuanjing Huang	1694
<i>Morphosyntactic Tagging with Pre-trained Language Models for Arabic and its Dialects</i> Go Inoue, Salam Khalifa and Nizar Habash	1708

<i>How Pre-trained Language Models Capture Factual Knowledge? A Causal-Inspired Analysis</i> Shaobo Li, Xiaoguang Li, Lifeng Shang, Zhenhua Dong, Chengjie Sun, Bingquan Liu, Zhenzhou Ji, Xin Jiang and Qun Liu	1720
<i>Metadata Shaping: A Simple Approach for Knowledge-Enhanced Language Models</i> Simran Arora, Sen Wu, Enci Liu and Christopher Re	1733
<i>Enhancing Natural Language Representation with Large-Scale Out-of-Domain Commonsense</i> Wanyun Cui and Xingran Chen	1746
<i>Weighted self Distillation for Chinese word segmentation</i> Rian He, Shubin Cai, Zhong Ming and Jiale Zhang	1757
<i>Sibylvariant Transformations for Robust Text Classification</i> Fabrice Y Harel-Canada, Muhammad Ali Gulzar, Nanyun Peng and Miryung Kim	1771
<i>DaLC: Domain Adaptation Learning Curve Prediction for Neural Machine Translation</i> Cheonbok Park, Hantae Kim, Ioan Calapodescu, Hyun Chang Cho and Vassilina Nikoulina	1789
<i>Hey AI, Can You Solve Complex Tasks by Talking to Agents?</i> Tushar Khot, Kyle Richardson, Daniel Khashabi and Ashish Sabharwal	1808
<i>Modality-specific Learning Rates for Effective Multimodal Additive Late-fusion</i> Yiqun Yao and Rada Mihalcea	1824
<i>BiSyn-GAT+: Bi-Syntax Aware Graph Attention Network for Aspect-based Sentiment Analysis</i> Shuo Liang, Wei Wei, Xian-Ling Mao, Fei Wang and Zhiyong He	1835
<i>IndicBART: A Pre-trained Model for Indic Natural Language Generation</i> Raj Dabre, Himani Shrotriya, Anoop Kunchukuttan, Ratish Puduppully, Mitesh M Khapra and Pratyush Kumar	1849
<i>Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models</i> Jianmo Ni, Gustavo Hernandez Abrego, Noah Constant, Ji Ma, Keith Hall, Daniel Cer and Yinfei Yang	1864
<i>Improving Relation Extraction through Syntax-induced Pre-training with Dependency Masking</i> Yuanhe Tian, Yan Song and Fei Xia	1875
<i>Striking a Balance: Alleviating Inconsistency in Pre-trained Models for Symmetric Classification Tasks</i> Ashutosh Kumar and Aditya Joshi	1887
<i>Diversifying Content Generation for Commonsense Reasoning with Mixture of Knowledge Graph Experts</i> Wenhao Yu, Chenguang Zhu, Lianhui Qin, Zhihan Zhang, Tong Zhao and Meng Jiang	1896
<i>Dict-BERT: Enhancing Language Model Pre-training with Dictionary</i> Wenhao Yu, Chenguang Zhu, Yuwei Fang, Donghan Yu, Shuohang Wang, Yichong Xu, Michael Zeng and Meng Jiang	1907
<i>A Feasibility Study of Answer-Unaware Question Generation for Education</i> Liam Dugan, Eleni Miltsakaki, Shriyash Kaustubh Upadhyay, Etan Jacob Ginsberg, Hannah Gonzalez, DaHyeon Choi, Chuning Yuan and Chris Callison-Burch	1919
<i>Relevant CommonSense Subgraphs for What if... Procedural Reasoning</i> Chen Zheng and Parisa Kordjamshidi	1927

<i>Combining Feature and Instance Attribution to Detect Artifacts</i>	
Pouya Pezeshkpour, Sarthak Jain, Sameer Singh and Byron C Wallace	1934
<i>Leveraging Expert Guided Adversarial Augmentation For Improving Generalization in Named Entity Recognition</i>	
Aaron Reich, Jiaao Chen, Aastha Agrawal, Yanzhe Zhang and Diyi Yang	1947
<i>Label Semantics for Few Shot Named Entity Recognition</i>	
Jie Ma, Miguel Ballesteros, Srikanth Doss, Rishita Anubhai, Sunil Mallya, Yaser Al-Onaizan and Dan Roth	1956
<i>Detection, Disambiguation, Re-ranking: Autoregressive Entity Linking as a Multi-Task Problem</i>	
Khalil Mrini, Shaoliang Nie, Jiatao Gu, Sinong Wang, Maziar Sanjabi and Hamed Firooz ..	1972
<i>VISITRON: Visual Semantics-Aligned Interactively Trained Object-Navigator</i>	
Ayush Shrivastava, Karthik Gopalakrishnan, Yang Liu, Robinson Piramuthu, Gokhan Tur, Devi Parikh and Dilek Hakkani-Tur	1984
<i>Investigating Selective Prediction Approaches Across Several Tasks in IID, OOD, and Adversarial Settings</i>	
Neeraj Varshney, Swaroop Mishra and Chitta Baral	1995
<i>Unsupervised Natural Language Inference Using PHL Triplet Generation</i>	
Neeraj Varshney, Pratyay Banerjee, Tejas Gokhale and Chitta Baral	2003
<i>Data Augmentation and Learned Layer Aggregation for Improved Multilingual Language Understanding in Dialogue</i>	
Evgeniia Razumovskaia, Ivan Vulić and Anna Korhonen	2017
<i>Ranking-Constrained Learning with Rationales for Text Classification</i>	
Juanyan Wang, Manali Sharma and Mustafa Bilgic	2034
<i>CaM-Gen: Causally Aware Metric-Guided Text Generation</i>	
Navita Goyal, Roodram Paneri, Ayush Agarwal, Udit Kalani, Abhilasha Sancheti and Niyati Chhaya	2047
<i>Training Dynamics for Text Summarization Models</i>	
Tanya Goyal, Jiacheng Xu, Junyi Jessy Li and Greg Durrett	2061
<i>Richer Countries and Richer Representations</i>	
Kaitlyn Zhou, Kawin Ethayarajh and Dan Jurafsky	2074
<i>BBQ: A hand-built bias benchmark for question answering</i>	
Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut and Samuel R. Bowman	2086
<i>Zero-shot Learning for Grapheme to Phoneme Conversion with Language Ensemble</i>	
Xinjian Li, Florian Metze, David R Mortensen, Shinji Watanabe and Alan Black	2106
<i>Dim Wihl Gat Tun: The Case for Linguistic Expertise in NLP for Under-Documented Languages</i>	
Clarissa Forbes, Farhan Samir, Bruce Harold Oliver, Changbing Yang, Edith Coates, Garrett Nicolai and Miikka Silfverberg	2116
<i>Question Generation for Reading Comprehension Assessment by Modeling How and What to Ask</i>	
Bilal Ghanem, Lauren Lutz Coleman, Julia Rivard Dexter, Spencer McIntosh von der Ohe and Alona Fyshe	2131

<i>TABi: Type-Aware Bi-Encoders for Open-Domain Entity Retrieval</i>	
Megan Leszczynski, Daniel Y Fu, Mayee F Chen and Christopher Re	2147
<i>Hierarchical Recurrent Aggregative Generation for Few-Shot NLG</i>	
Giulio Zhou, Gerasimos Lampouras and Ignacio Iacobacci	2167
<i>Training Text-to-Text Transformers with Privacy Guarantees</i>	
Natalia Ponomareva, Jasmijn Bastings and Sergei Vassilvitskii	2182
<i>Revisiting Uncertainty-based Query Strategies for Active Learning with Transformers</i>	
Christopher Schröder, Andreas Niekler and Martin Potthast	2194
<i>The impact of lexical and grammatical processing on generating code from natural language</i>	
Nathanaël Beau and Benoit Crabbé	2204
<i>Seq2Path: Generating Sentiment Tuples as Paths of a Tree</i>	
Yue Mao, Yi Shen, Jingchao Yang, Xiaoying Zhu and Longjun Cai	2215
<i>Mitigating the Inconsistency Between Word Saliency and Model Confidence with Pathological Contrastive Training</i>	
Pengwei Zhan, Yang Wu, Shaolei Zhou, Yunjian Zhang and Liming Wang	2226
<i>Your fairness may vary: Pretrained language model fairness in toxic text classification</i>	
Ioana Baldini, Dennis Wei, Karthikeyan Natesan Ramamurthy, Moninder Singh and Mikhail Yurochkin	2245
<i>ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning</i>	
Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty and Enamul Hoque	2263
<i>A Novel Perspective to Look At Attention: Bi-level Attention-based Explainable Topic Modeling for News Classification</i>	
Dairui Liu, Derek Greene and Ruihai Dong	2280
<i>Learn and Review: Enhancing Continual Named Entity Recognition via Reviewing Synthetic Samples</i>	
Yu Xia, Quan Wang, Yajuan Lyu, Yong Zhu, Wenhao Wu, Sujian Li and Dai Dai	2291
<i>Phoneme transcription of endangered languages: an evaluation of recent ASR architectures in the single speaker scenario</i>	
Gilles Boulianne	2301
<i>Does BERT really agree ? Fine-grained Analysis of Lexical Dependence on a Syntactic Task</i>	
Karim Lasri, Alessandro Lenci and Thierry Poibeau	2309
<i>Combining Static and Contextualised Multilingual Embeddings</i>	
Katharina Hämmerl, Jindřich Libovický and Alexander Fraser	2316
<i>An Accurate Unsupervised Method for Joint Entity Alignment and Dangling Entity Detection</i>	
Shengxuan Luo and Sheng Yu	2330
<i>Square One Bias in NLP: Towards a Multi-Dimensional Exploration of the Research Manifold</i>	
Sebastian Ruder, Ivan Vulić and Anders Søgaard	2340
<i>Systematicity, Compositionality and Transitivity of Deep NLP Models: a Metamorphic Testing Perspective</i>	
Edoardo Manino, Julia Rozanova, Danilo Carvalho, Andre Freitas and Lucas Carvalho Cordeiro	2355

<i>Improving Neural Political Statement Classification with Class Hierarchical Information</i> Erenay Dayanik, Andre Blessing, Nico Blokker, Sebastian Haunss, Jonas Kuhn, Gabriella Lapesa and Sebastian Pado	2367
<i>Enabling Multimodal Generation on CLIP via Vision-Language Knowledge Distillation</i> Wenliang Dai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu and Pascale Fung	2383
<i>Co-VQA : Answering by Interactive Sub Question Sequence</i> Ruonan Wang, Yuxi Qian, Fangxiang Feng, Xiaojie Wang and Huixing Jiang	2396
<i>A Simple Hash-Based Early Exiting Approach For Language Understanding and Generation</i> Tianxiang Sun, Xiangyang Liu, Wei Zhu, Zhichao Geng, Lingling Wu, Yilong He, Yuan Ni, Guotong Xie, Xuanjing Huang and Xipeng Qiu	2409
<i>Auxiliary tasks to boost Biaffine Semantic Dependency Parsing</i> Marie Candito	2422
<i>Syntax-guided Contrastive Learning for Pre-trained Language Model</i> Shuai Zhang, Wang Lijie, Xinyan Xiao and Hua Wu	2430
<i>Improved Multi-label Classification under Temporal Concept Drift: Rethinking Group-Robust Algorithms in a Label-Wise Setting</i> Ilias Chalkidis and Anders Søgaard	2441
<i>ASCM: An Answer Space Clustered Prompting Method without Answer Engineering</i> Zhen Wang, Yating Yang, Zhou Xi, Bo Ma, Lei Wang, Rui Dong and Azmat Anwar	2455
<i>Why don't people use character-level machine translation?</i> Jindřich Libovický, Helmut Schmid and Alexander Fraser	2470
<i>Seeking Patterns, Not just Memorizing Procedures: Contrastive Learning for Solving Math Word Problems</i> Zhongli Li, Wenxuan Zhang, Chao Yan, Qingyu Zhou, Chao Li, Hongzhi Liu and Yunbo Cao 2486	
<i>xGQA: Cross-Lingual Visual Question Answering</i> Jonas Pfeiffer, Gregor Geigle, Aishwarya Kamath, Jan-Martin O. Steitz, Stefan Roth, Ivan Vulić and Iryna Gurevych	2497
<i>Automatic Speech Recognition and Query By Example for Creole Languages Documentation</i> Cécile Macaire, Didier Schwab, Benjamin Lecouteux and Emmanuel Schang	2512
<i>MReD: A Meta-Review Dataset for Structure-Controllable Text Generation</i> Chenhui Shen, Liying Cheng, Ran Zhou, Lidong Bing, Yang You and Luo Si	2521
<i>Single Model Ensemble for Subword Regularized Models in Low-Resource Machine Translation</i> Sho Takase, Tatsuya Hiraoka and Naoaki Okazaki	2536
<i>Detecting Various Types of Noise for Neural Machine Translation</i> Christian Herold, Jan Rosendahl, Joris Vanvinckenroye and Hermann Ney	2542
<i>DU-VLG: Unifying Vision-and-Language Generation via Dual Sequence-to-Sequence Pre-training</i> Luyang Huang, Guocheng Niu, Jiachen Liu, Xinyan Xiao and Hua Wu	2552
<i>HiCLRE: A Hierarchical Contrastive Learning Framework for Distantly Supervised Relation Extraction</i> Dongyang Li, Taolin Zhang, Nan Hu, Chengyu Wang and Xiaofeng He	2567

<i>Prompt-Driven Neural Machine Translation</i>	
Yafu Li, Yongjing Yin, Jing Li and Yue Zhang	2579
<i>On Controlling Fallback Responses for Grounded Dialogue Generation</i>	
Hongyuan Lu, Wai Lam, Hong Cheng and Helen M. Meng	2591
<i>CRAFT: A Benchmark for Causal Reasoning About Forces and Interactions</i>	
Tayfun Ates, M. Şamil Ateşoğlu, Çağatay Yiğit, Ilker Kesen, Mert Kobas, Erkut Erdem, Aykut Erdem, Tilbe Goksun and Deniz Yuret	2602
<i>A Graph Enhanced BERT Model for Event Prediction</i>	
Li Du, Xiao Ding, Yue Zhang, Ting Liu and Bing Qin	2628
<i>Long Time No See! Open-Domain Conversation with Long-Term Persona Memory</i>	
Xinchao Xu, Zhibin Gou, Wenquan Wu, Zheng-Yu Niu, Hua Wu, Haifeng Wang and Shihang Wang	2639
<i>Lacking the Embedding of a Word? Look it up into a Traditional Dictionary</i>	
Elena Sofia Ruzzetti, Leonardo Ranaldi, Michele Mastromattei, Francesca Fallucchi, Noemi Scarpato and Fabio Massimo Zanzotto	2651
<i>MTRec: Multi-Task Learning over BERT for News Recommendation</i>	
Qiwei Bi, Jian Li, Lifeng Shang, Xin Jiang, Qun Liu and Hanfang Yang	2663
<i>Cross-domain Named Entity Recognition via Graph Matching</i>	
Junhao Zheng, Haibin Chen and Qianli Ma	2670
<i>Assessing Multilingual Fairness in Pre-trained Multimodal Representations</i>	
Jialu Wang, Yang Liu and Xin Eric Wang	2681
<i>More Than Words: Collocation Retokenization for Latent Dirichlet Allocation Models</i>	
Jin Cheevaprawatdomrong, Alexandra Schofield and Attapol Rutherford	2696
<i>Generalized but not Robust? Comparing the Effects of Data Modification Methods on Out-of-Domain Generalization and Adversarial Robustness</i>	
Tejas Gokhale, Swaroop Mishra, Man Luo, Bhavdeep Singh Sachdeva and Chitta Baral	2705
<i>ASSIST: Towards Label Noise-Robust Dialogue State Tracking</i>	
Fanghua Ye, Yue Feng and Emine Yilmaz	2719
<i>Graph Refinement for Coreference Resolution</i>	
Lesly Miculicich and James Henderson	2732
<i>ECO v1: Towards Event-Centric Opinion Mining</i>	
Ruoxi Xu, Hongyu Lin, Meng Liao, Xianpei Han, Jin Xu, Wei Tan, Yingfei Sun and Le Sun	2743
<i>Deep Reinforcement Learning for Entity Alignment</i>	
Lingbing Guo, Yuqiang Han, Qiang Zhang and Huajun Chen	2754
<i>Breaking Down Multilingual Machine Translation</i>	
Ting-Rui Chiang, Yi-Pei Chen, Yi-Ting Yeh and Graham Neubig	2766
<i>Mitigating Contradictions in Dialogue Based on Contrastive Learning</i>	
Weizhao Li, Junsheng Kong, Ben Liao and Yi Cai	2781
<i>ELLE: Efficient Lifelong Pre-training for Emerging Data</i>	
Yujia Qin, Jiajie Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun and Jie Zhou	2789

<i>EnCBP: A New Benchmark Dataset for Finer-Grained Cultural Background Prediction in English</i> Weicheng Ma, Samiha Datta, Lili Wang and Soroush Vosoughi	2811
<i>Cutting Down on Prompts and Parameters: Simple Few-Shot Learning with Language Models</i> Robert L. Logan IV, Ivana Balazevic, Eric Wallace, Fabio Petroni, Sameer Singh and Sebastian Riedel	2824
<i>uFACT: Unfaithful Alien-Corpora Training for Semantically Consistent Data-to-Text Generation</i> Tisha Anders, Alexandru Coca and Bill Byrne	2836
<i>Good Night at 4 pm?! Time Expressions in Different Cultures</i> Vered Shwartz	2842
<i>Extracting Person Names from User Generated Text: Named-Entity Recognition for Combating Human Trafficking</i> Yifei Li, Pratheeksha Nair, Kellin Pelrine and Reihaneh Rabbany	2854
<i>OneAligner: Zero-shot Cross-lingual Transfer with One Rich-Resource Language Pair for Low-Resource Sentence Retrieval</i> Tong Niu, Kazuma Hashimoto, Yingbo Zhou and Caiming Xiong	2869
<i>Suum Cuique: Studying Bias in Taboo Detection with a Community Perspective</i> Osama Khalid, Jonathan Rusert and Padmini Srinivasan	2883
<i>Modeling Intensification for Sign Language Generation: A Computational Approach</i> Mert Inan, Yang Zhong, Sabit Hassan, Lorna Quandt and Malihe Alikhani	2897
<i>Controllable Natural Language Generation with Contrastive Prefixes</i> Jing Qian, Li Dong, Yelong Shen, Furu Wei and Weizhu Chen	2912
<i>Revisiting the Effects of Leakage on Dependency Parsing</i> Nathaniel Krasner, Miriam Wanner and Antonios Anastasopoulos	2925
<i>Learning to Describe Solutions for Bug Reports Based on Developer Discussions</i> Sheena Panthaplackel, Junyi Jessy Li, Milos Gligoric and Ray Mooney	2935
<i>Perturbations in the Wild: Leveraging Human-Written Text Perturbations for Realistic Adversarial Attack and Defense</i> Thai Le, Jooyoung Lee, Kevin Yen, Yifan Hu and Dongwon Lee	2953
<i>Improving Chinese Grammatical Error Detection via Data augmentation by Conditional Error Generation</i> Tianchi Yue, Shulin Liu, Huihui Cai, Tao Yang, Shengkang Song and TingHao YU	2966
<i>Modular and Parameter-Efficient Multimodal Fusion with Prompting</i> Sheng Liang, Mengjie Zhao and Hinrich Schuetze	2976
<i>Synchronous Refinement for Neural Machine Translation</i> Kehai Chen, Masao Utiyama, Eiichiro Sumita, Rui Wang and Min Zhang	2986
<i>HIE-SQL: History Information Enhanced Network for Context-Dependent Text-to-SQL Semantic Parsing</i> Yanzhao Zheng, Haibin Wang, Baohua Dong, Xingjun Wang and Changshan Li	2997
<i>CRASpell: A Contextual Typo Robust Approach to Improve Chinese Spelling Correction</i> Shulin Liu, Shengkang Song, Tianchi Yue, Tao Yang, Huihui Cai, TingHao YU and Shengli Sun	3008

<i>Gaussian Multi-head Attention for Simultaneous Machine Translation</i> Shaolei Zhang and Yang Feng	3019
<i>Composing Structure-Aware Batches for Pairwise Sentence Classification</i> Andreas Waldis, Tilman Beck and Iryna Gurevych	3031
<i>Factual Consistency of Multilingual Pretrained Language Models</i> Constanza Fierro and Anders Søgaard	3046
<i>Selecting Stickers in Open-Domain Dialogue through Multitask Learning</i> Zhexin Zhang, Yeshuang Zhu, Zhengcong Fei, Jinchao Zhang and Jie Zhou	3053
<i>ZiNet: Linking Chinese Characters Spanning Three Thousand Years</i> Yang Chi, Fausto Giunchiglia, Daqian Shi, Xiaolei Diao, Chuntao Li and Hao Xu	3061
<i>How Can Cross-lingual Knowledge Contribute Better to Fine-Grained Entity Typing?</i> Hailong Jin, Tiansi Dong, Lei Hou, Juanzi Li, Hui Chen, Zelin Dai and Qu Yincen	3071
<i>AMR-DA: Data Augmentation by Abstract Meaning Representation</i> Ziyi Shou, Yuxin Jiang and Fangzhen Lin	3082
<i>Using Pre-Trained Language Models for Producing Counter Narratives Against Hate Speech: a Comparative Study</i> Serra Sinem Tekiroglu, Helena Bonaldi, Margherita Fanton and Marco Guerini	3099
<i>Improving Robustness of Language Models from a Geometry-aware Perspective</i> Bin Zhu, Zhaoquan Gu, Le Wang, Jinyin Chen and Qi Xuan	3115
<i>Task-guided Disentangled Tuning for Pretrained Language Models</i> Jiali Zeng, Yufan Jiang, Shuangzhi Wu, Yongjing Yin and Mu Li	3126
<i>Exploring the Impact of Negative Samples of Contrastive Learning: A Case Study of Sentence Embedding</i> Rui Cao, Yihao Wang, Yuxin Liang, Ling Gao, Jie Zheng, Jie Ren and Zheng Wang	3138
<i>The Inefficiency of Language Models in Scholarly Retrieval: An Experimental Walk-through</i> Shruti Singh and Mayank Singh	3153
<i>Fusing Heterogeneous Factors with Triaffine Mechanism for Nested Named Entity Recognition</i> Zheng Yuan, Chuanqi Tan, Songfang Huang and Fei Huang	3174
<i>UNIMO-2: End-to-End Unified Vision-Language Grounded Learning</i> Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu and Haifeng Wang	3187
<i>The Past Mistake is the Future Wisdom: Error-driven Contrastive Probability Optimization for Chinese Spell Checking</i> Yinghui Li, Qingyu Zhou, Yangning Li, Zhongli Li, Ruiyang Liu, Rongyi Sun, Zizhen Wang, Chao Li, Yunbo Cao and Hai-Tao Zheng	3202
<i>XFUND: A Benchmark Dataset for Multilingual Visually Rich Form Understanding</i> Yiheng Xu, Tengchao Lv, Lei Cui, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang and Furu Wei	3214
<i>Type-Driven Multi-Turn Corrections for Grammatical Error Correction</i> Shaopeng Lai, Qingyu Zhou, Jiali Zeng, Zhongli Li, Chao Li, Yunbo Cao and Jinsong Su ..	3225

<i>Leveraging Knowledge in Multilingual Commonsense Reasoning</i>	
Yuwei Fang, Shuohang Wang, Yichong Xu, Ruochen Xu, Siqi Sun, Chenguang Zhu and Michael Zeng	3237
<i>Encoding and Fusing Semantic Connection and Linguistic Evidence for Implicit Discourse Relation Recognition</i>	
Wei Xiang, Bang Wang, Lu Dai and Yijun Mo	3247
<i>One Agent To Rule Them All: Towards Multi-agent Conversational AI</i>	
Christopher Clarke, Joseph J Peper, Karthik Krishnamurthy, Walter Talamonti, Kevin Leach, Walter Lasecki, Yiping Kang, Lingjia Tang and Jason Mars	3258
<i>Word-level Perturbation Considering Word Length and Compositional Subwords</i>	
Tatsuya Hiraoka, Sho Takase, Kei Uchiumi, Atsushi Keyaki and Naoaki Okazaki	3268
<i>Bridging Pre-trained Language Models and Hand-crafted Features for Unsupervised POS Tagging</i>	
Houquan Zhou, Yang Li, Zhenghua Li and Min Zhang	3276
<i>Controlling the Focus of Pretrained Language Generation Models</i>	
Jiabao Ji, Yoon Kim, James R. Glass and Tianxing He	3291
<i>Comparative Opinion Summarization via Collaborative Decoding</i>	
Hayate Iso, Xiaolan Wang, Stefanos Angelidis and Yoshihiko Suhara	3307
<i>IsoScore: Measuring the Uniformity of Embedding Space Utilization</i>	
William Rudman, Nate Gillman, Taylor Rayne and Carsten Eickhoff	3325
<i>A Natural Diet: Towards Improving Naturalness of Machine Translation Output</i>	
Markus Freitag, David Vilar, David Grangier, Colin Cherry and George Foster	3340
<i>From Stance to Concern: Adaptation of Propositional Analysis to New Tasks and Domains</i>	
Brodie Mather, Bonnie J Dorr, Adam Dalton, William de Beaumont, Owen Rambow and Sonja M. Schmer-Galunder	3354
<i>CUE Vectors: Modular Training of Language Models Conditioned on Diverse Contextual Signals</i>	
Scott Novotney, Sreeparna Mukherjee, Zeeshan Ahmed and Andreas Stolcke	3368
<i>Cross-Lingual UMLS Named Entity Linking using UMLS Dictionary Fine-Tuning</i>	
Rina Galperin, Shachar Schnapp and Michael Elhadad	3380
<i>Aligned Weight Regularizers for Pruning Pretrained Neural Networks</i>	
James O’ Neill, Sourav Dutta and Haytham Assem	3391
<i>Consistent Representation Learning for Continual Relation Extraction</i>	
Kang Zhao, Hua Xu, Jiangong Yang and Kai Gao	3402
<i>Event Transition Planning for Open-ended Text Generation</i>	
Qintong Li, Piji Li, Wei Bi, Zhaochun Ren, Yuxuan Lai and Lingpeng Kong	3412
<i>Comprehensive Multi-Modal Interactions for Referring Image Segmentation</i>	
Kanishk Jain and Vineet Gandhi	3427
<i>MetaWeighting: Learning to Weight Tasks in Multi-Task Learning</i>	
Yuren Mao, Zekai Wang, Weiwei Liu, Xuemin Lin and Pengtao Xie	3436
<i>Improving Controllable Text Generation with Position-Aware Weighted Decoding</i>	
Yuxuan Gu, Xiaocheng Feng, Sicheng Ma, Jiaming Wu, Heng Gong and Bing Qin	3449

<i>Prompt Tuning for Discriminative Pre-trained Language Models</i>	
Yuan Yao, Bowen Dong, Ao Zhang, Zhengyan Zhang, Ruobing Xie, Zhiyuan Liu, Leyu Lin, Maosong Sun and Jianyong Wang	3468
<i>Two Birds with One Stone: Unified Model Learning for Both Recall and Ranking in News Recommendation</i>	
Chuhan Wu, Fangzhao Wu, Tao Qi and Yongfeng Huang	3474
<i>What does it take to bake a cake? The RecipeRef corpus and anaphora resolution in procedural text</i>	
Biaoyan Fang, Timothy Baldwin and Karin Verspoor	3481
<i>MERIt: Meta-Path Guided Contrastive Learning for Logical Reasoning</i>	
Fangkai Jiao, Yangyang Guo, Xueming Song and Liqiang Nie	3496
<i>THE-X: Privacy-Preserving Transformer Inference with Homomorphic Encryption</i>	
Tianyu Chen, Hangbo Bao, Shaohan Huang, Li Dong, Binxing Jiao, Daxin Jiang, Haoyi Zhou, Jianxin Li and Furu Wei	3510
<i>HLDC: Hindi Legal Documents Corpus</i>	
Arnav Kapoor, Mudit Dhawan, Anmol Goel, Arjun T H, Akshala Bhatnagar, Vibhu Agrawal, Amul Agrawal, Arnab Bhattacharya, Ponnurangam Kumaraguru and Ashutosh Modi	3521
<i>Rethinking Document-level Neural Machine Translation</i>	
Zewei Sun, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Shujian Huang, Jiajun Chen and Lei Li	3537
<i>Incremental Intent Detection for Medical Domain with Contrast Replay Networks</i>	
Guirong Bai, Shizhu He, Kang Liu and Jun Zhao	3549
<i>LaPraDoR: Unsupervised Pretrained Dense Retriever for Zero-Shot Text Retrieval</i>	
Canwen Xu, Daya Guo, Nan Duan and Julian McAuley	3557
<i>Do Pre-trained Models Benefit Knowledge Graph Completion? A Reliable Evaluation and a Reasonable Approach</i>	
Xin Lv, Yankai Lin, Yixin Cao, Lei Hou, Juanzi Li, Zhiyuan Liu, Peng Li and Jie Zhou	3570
<i>EICO: Improving Few-Shot Text Classification via Explicit and Implicit Consistency Regularization</i>	
Lei Zhao and Cheng Yao	3582
<i>Improving the Adversarial Robustness of NLP Models by Information Bottleneck</i>	
Senyuan Zhang, Xiang Zhou, Yixin Wan, Xiaoqing Zheng, Kai-Wei Chang and Cho-Jui Hsieh	3588
<i>Incorporating Dynamic Semantics into Pre-Trained Language Model for Aspect-based Sentiment Analysis</i>	
Kai Zhang, Kun Zhang, Mengdi Zhang, Hongke Zhao, Qi Liu, Wei Wu and Enhong Chen ..	3599
<i>DARER: Dual-task Temporal Relational Recurrent Reasoning Network for Joint Dialog Sentiment Classification and Act Recognition</i>	
Bowen Xing and Ivor Tsang	3611
<i>Divide and Conquer: Text Semantic Matching with Disentangled Keywords and Intents</i>	
Yicheng Zou, Hongwei Liu, Tao Gui, Junzhe Wang, Qi Zhang, Meng Tang, Haixiang Li and Daniell Wang	3622
<i>Modular Domain Adaptation</i>	
Junshen K Chen, Dallas Card and Dan Jurafsky	3633

<i>Detection of Adversarial Examples in Text Classification: Benchmark and Baseline via Robust Density Estimation</i>	
KiYoon Yoo, Jangho Kim, Jiho Jang and Nojun Kwak	3656
<i>Platt-Bin: Efficient Posterior Calibrated Training for NLP Classifiers</i>	
Rishabh Singh and Shirin Goshtasbpour	3673
<i>Addressing Resource and Privacy Constraints in Semantic Parsing Through Data Augmentation</i>	
Kevin Yang, Olivia Deng, Charles Chen, Richard Shin, Subhro Roy and Benjamin Van Durme	3685
<i>Improving Candidate Retrieval with Entity Profile Generation for Wikidata Entity Linking</i>	
Tuan Lai, Heng Ji and ChengXiang Zhai	3696
<i>Local Structure Matters Most: Perturbation Study in NLU</i>	
Louis Cloutre, Prasanna Parthasarathi, Amal Zouaq and Sarath Chandar	3712
<i>Probing Factually Grounded Content Transfer with Factual Ablation</i>	
Peter West, Chris Quirk, Michel Galley and Yejin Choi	3732
<i>ED2LM: Encoder-Decoder to Language Model for Faster Document Re-ranking Inference</i>	
Kai Hui, Honglei Zhuang, Tao Chen, Zhen Qin, Jing Lu, Dara Bahri, Ji Ma, Jai Gupta, Cicero Nogueira dos Santos, Yi Tay and Donald Metzler	3747
<i>Benchmarking Answer Verification Methods for Question Answering-Based Summarization Evaluation Metrics</i>	
Daniel Deutsch and Dan Roth	3759
<i>Prior Knowledge and Memory Enriched Transformer for Sign Language Translation</i>	
Tao Jin, Zhou Zhao, Meng Zhang and Xingshan Zeng	3766
<i>Discontinuous Constituency and BERT: A Case Study of Dutch</i>	
Konstantinos Kogkalidis and Gijs Wijnholds	3776
<i>Probing Multilingual Cognate Prediction Models</i>	
Clémentine Fourier and Benoît Sagot	3786
<i>A Neural Pairwise Ranking Model for Readability Assessment</i>	
Justin Lee and Sowmya Vajjala	3802
<i>First the Worst: Finding Better Gender Translations During Beam Search</i>	
Danielle Saunders, Rosie Sallis and Bill Byrne	3814
<i>Dialogue Summaries as Dialogue States (DS2), Template-Guided Summarization for Few-shot Dialogue State Tracking</i>	
Jamin Shin, Hangeol Yu, Hyeongdon Moon, Andrea Madotto and Juneyoung Park	3824
<i>Unsupervised Preference-Aware Language Identification</i>	
Xingzhang Ren, Baosong Yang, Dayiheng Liu, Haibo Zhang, Xiaoyu Lv, Liang Yao and Jun Xie	3847
<i>Using NLP to quantify the environmental cost and diversity benefits of in-person NLP conferences</i>	
Piotr Przybyła and Matthew Shardlow	3853
<i>Interpretable Research Replication Prediction via Variational Contextual Consistency Sentence Masking</i>	
Tianyi Luo, Rui Meng, Xin Eric Wang and Yang Liu	3864

<i>Chinese Synesthesia Detection: New Dataset and Models</i>	
Xiaotong Jiang, Qingqing Zhao, Yunfei Long and Zhongqing Wang	3877
<i>Rethinking Offensive Text Detection as a Multi-Hop Reasoning Problem</i>	
Qiang Zhang, Jason Naradowsky and Yusuke Miyao	3888
<i>On the Safety of Conversational Models: Taxonomy, Dataset, and Benchmark</i>	
Hao Sun, Guangxuan Xu, Jiawen Deng, Jiale Cheng, Chujie Zheng, Hao Zhou, Nanyun Peng, Xiaoyan Zhu and Minlie Huang	3906
<i>Word Segmentation by Separation Inference for East Asian Languages</i>	
Yu Tong, Jingzhi Guo, Jizhe Zhou, Ge Chen and Guokai Zheng	3924
<i>Unsupervised Chinese Word Segmentation with BERT Oriented Probing and Transformation</i>	
Wei Li, Yuhan Song, Qi Su and Yanqiu Shao	3935
<i>E-KAR: A Benchmark for Rationalizing Natural Language Analogical Reasoning</i>	
Jiangjie Chen, Rui Xu, Ziquan Fu, Wei Shi, Zhongqiao Li, Xinbo Zhang, Changzhi Sun, Lei Li, Yanghua Xiao and Hao Zhou	3941
<i>Implicit Relation Linking for Question Answering over Knowledge Graph</i>	
Yao Zhao, Jiacheng Huang, Wei Hu, Qijin Chen, XiaoXia Qiu, Chengfu Huo and Weijun Ren	3956
<i>Attention Mechanism with Energy-Friendly Operations</i>	
Yu Wan, Baosong Yang, Dayiheng Liu, Rong Xiao, Derek F. Wong, Haibo Zhang, Boxing Chen and Lidia S. Chao	3969
<i>Probing BERT's priors with serial reproduction chains</i>	
Takateru Yamakoshi, Thomas L. Griffiths and Robert D. Hawkins	3977
<i>Interpreting the Robustness of Neural NLP Models to Textual Perturbations</i>	
Yunxiang Zhang, Liangming Pan, Samson Tan and Min-Yen Kan	3993
<i>Zero-Shot Dense Retrieval with Momentum Adversarial Domain Invariant Representations</i>	
Ji Xin, Chenyan Xiong, Ashwin Srinivasan, Ankita Sharma, Damien Jose and Paul N. Bennett	4008
<i>A Few-Shot Semantic Parser for Wizard-of-Oz Dialogues with the Precise ThingTalk Representation</i>	
Giovanni Campagna, Sina Semnani, Ryan Othniel Kearns, Lucas Jun Koba Sato, Silei Xu and Monica Lam	4021
<i>GCPG: A General Framework for Controllable Paraphrase Generation</i>	
Kexin Yang, Dayiheng Liu, Wenqiang Lei, Baosong Yang, Haibo Zhang, Xue Zhao, Wenqing Yao and Boxing Chen	4035
<i>CrossAligner & Co: Zero-Shot Transfer Methods for Task-Oriented Cross-lingual Natural Language Understanding</i>	
Milan Gritta, Ruoyu Hu and Ignacio Iacobacci	4048
<i>Attention as Grounding: Exploring Textual and Cross-Modal Attention on Entities and Relations in Language-and-Vision Transformer</i>	
Nikolai Ilinykh and Simon Dobnik	4062
<i>Improving Zero-Shot Cross-lingual Transfer Between Closely Related Languages by Injecting Character-Level Noise</i>	
Noëmi Aeppli and Rico Sennrich	4074

<i>Structural Supervision for Word Alignment and Machine Translation</i>	
Lei Li, Kai Fan, Hongjia Li and Chun Yuan	4084
<i>Focus on the Action: Learning to Highlight and Summarize Jointly for Email To-Do Items Summarization</i>	
Kexun Zhang, Jiaao Chen and Diyi Yang	4095
<i>Exploring the Capacity of a Large-scale Masked Language Model to Recognize Grammatical Errors</i>	
Ryo Nagata, Manabu Kimura and Kazuaki Hanawa	4107
<i>Should We Trust This Summary? Bayesian Abstractive Summarization to The Rescue</i>	
Alexios Gidiotis and Grigorios Tsoumakas	4119
<i>On the data requirements of probing</i>	
Zining Zhu, Jixuan Wang, Bai Li and Frank Rudzicz	4132
<i>Translation Error Detection as Rationale Extraction</i>	
Marina Fomicheva, Lucia Specia and Nikolaos Aletras	4148
<i>Towards Collaborative Neural-Symbolic Graph Semantic Parsing via Uncertainty</i>	
Zi Lin, Jeremiah Zhe Liu and Jingbo Shang	4160
<i>Towards Few-shot Entity Recognition in Document Images: A Label-aware Sequence-to-Sequence Framework</i>	
Zilong Wang and Jingbo Shang	4174
<i>On Length Divergence Bias in Textual Matching Models</i>	
Lan Jiang, Tianshu Lyu, Yankai Lin, Meng Chong, Xiaoyong Lyu and Dawei Yin	4187
<i>What is wrong with you?: Leveraging User Sentiment for Automatic Dialog Evaluation</i>	
Sarik Ghazarian, Behnam Hedayatnia, Alexandros Papangelis, Yang Liu and Dilek Hakkani-Tur	4194