

Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT 2021)

Online
6 – 11 June 2021

Volume 1 of 8

ISBN: 978-1-7138-3009-2

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© (2021) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2021)

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>Knowledge Router: Learning Disentangled Representations for Knowledge Graphs</i> Shuai Zhang, Xi Rao, Yi Tay and Ce Zhang	1
<i>Distantly Supervised Relation Extraction with Sentence Reconstruction and Knowledge Base Priors</i> Fenia Christopoulou, Makoto Miwa and Sophia Ananiadou	11
<i>Cross-Task Instance Representation Interactions and Label Dependencies for Joint Information Extraction with Graph Convolutional Networks</i> Minh Van Nguyen, Viet Lai and Thien Huu Nguyen	27
<i>Abstract Meaning Representation Guided Graph Encoding and Decoding for Joint Information Extraction</i> Zixuan Zhang and Heng Ji	39
<i>A Frustratingly Easy Approach for Entity and Relation Extraction</i> Zexuan Zhong and Danqi Chen	50
<i>Event Time Extraction and Propagation via Graph Attention Networks</i> Haoyang Wen, Yanru Qu, Heng Ji, Qiang Ning, Jiawei Han, Avi Sil, Hanghang Tong and Dan Roth	62
<i>Probing Word Translations in the Transformer and Trading Decoder for Encoder Layers</i> Hongfei Xu, Josef van Genabith, Qiuhui Liu and Deyi Xiong	74
<i>Mediators in Determining what Processing BERT Performs First</i> Aviv Slobodkin, Leshem Choshen and Omri Abend	86
<i>Automatic Generation of Contrast Sets from Scene Graphs: Probing the Compositional Consistency of GQA</i> Yonatan Bitton, Gabriel Stanovsky, Roy Schwartz and Michael Elhadad	94
<i>Multilingual Language Models Predict Human Reading Behavior</i> Nora Hollenstein, Federico Pirovano, Ce Zhang, Lena Jäger and Lisa Beinborn	106
<i>Do Syntactic Probes Probe Syntax? Experiments with Jabberwocky Probing</i> Rowan Hall Maudslay and Ryan Cotterell	124
<i>A Non-Linear Structural Probe</i> Jennifer C. White, Tiago Pimentel, Naomi Saphra and Ryan Cotterell	132
<i>Concealed Data Poisoning Attacks on NLP Models</i> Eric Wallace, Tony Zhao, Shi Feng and Sameer Singh	139
<i>Backtranslation Feedback Improves User Confidence in MT, Not Quality</i> Vilém Zouhar, Michal Novák, Matúš Žilinec, Ondřej Bojar, Mateo Obregón, Robin L. Hill, Frédéric Blain, Marina Fomicheva, Lucia Specia and Lisa Yankovskaya	151
<i>Data Filtering using Cross-Lingual Word Embeddings</i> Christian Herold, Jan Rosendahl, Joris Vanvinckenroye and Hermann Ney	162

<i>Improving the Lexical Ability of Pretrained Language Models for Unsupervised Neural Machine Translation</i>	
Alexandra Chronopoulou, Dario Stojanovski and Alexander Fraser	173
<i>Neural Machine Translation without Embeddings</i>	
Uri Shaham and Omer Levy	181
<i>Counterfactual Data Augmentation for Neural Machine Translation</i>	
Qi Liu, Matt Kusner and Phil Blunsom	187
<i>Cultural and Geographical Influences on Image Translatability of Words across Languages</i>	
Nikzad Khani, Isidora Tourni, Mohammad Sadegh Rasooli, Chris Callison-Burch and Derry Tanti Wijaya	198
<i>Multilingual BERT Post-Pretraining Alignment</i>	
Lin Pan, Chung-Wei Hang, Haode Qi, Abhishek Shah, Saloni Potdar and Mo Yu	210
<i>A Million Tweets Are Worth a Few Points: Tuning Transformers for Customer Service Tasks</i>	
Amir Hadifar, Sofie Labat, Veronique Hoste, Chris Develder and Thomas Demeester	220
<i>Paragraph-level Rationale Extraction through Regularization: A case study on European Court of Human Rights Cases</i>	
Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapatsanis, Nikolaos Aletras, Ion Androutsopoulos and Prodromos Malakasiotis	226
<i>Answering Product-Questions by Utilizing Questions from Other Contextually Similar Products</i>	
Ohad Rozen, David Carmel, Avihai Mejer, Vitaly Mirkis and Yftah Ziser	242
<i>EnSidNet: Enhanced Hybrid Siamese-Deep Network for grouping clinical trials into drug-development pathways</i>	
Lucia Pagani	254
<i>DATE: Detecting Anomalies in Text via Self-Supervision of Transformers</i>	
Andrei Manolache, Florin Brad and Elena Burceanu	267
<i>A Simple Approach for Handling Out-of-Vocabulary Identifiers in Deep Learning for Source Code</i>	
Nadezhda Chirkova and Sergey Troshin	278
<i>Fast and Scalable Dialogue State Tracking with Explicit Modular Decomposition</i>	
Dingmin Wang, Chenghua Lin, Qi Liu and Kam-Fai Wong	289
<i>Augmented SBERT: Data Augmentation Method for Improving Bi-Encoders for Pairwise Sentence Scoring Tasks</i>	
Nandan Thakur, Nils Reimers, Johannes Daxenberger and Iryna Gurevych	296
<i>SmBoP: Semi-autoregressive Bottom-up Semantic Parsing</i>	
Ohad Rubin and Jonathan Berant	311
<i>SGL: Speaking the Graph Languages of Semantic Parsing via Multilingual Translation</i>	
Luigi Procopio, Rocco Tripodi and Roberto Navigli	325
<i>Unifying Cross-Lingual Semantic Role Labeling with Heterogeneous Linguistic Resources</i>	
Simone Conia, Andrea Bacciu and Roberto Navigli	338

<i>Fool Me Twice: Entailment from Wikipedia Gamification</i>	
Julian Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger and Jordan Boyd-Graber	352
<i>Meta-Learning for Domain Generalization in Semantic Parsing</i>	
Bailin Wang, Mirella Lapata and Ivan Titov	366
<i>Aspect-Controlled Neural Argument Generation</i>	
Benjamin Schiller, Johannes Daxenberger and Iryna Gurevych	380
<i>Text Generation from Discourse Representation Structures</i>	
Jiangming Liu, Shay B. Cohen and Mirella Lapata	397
<i>APo-VAE: Text Generation in Hyperbolic Space</i>	
Shuyang Dai, Zhe Gan, Yu Cheng, Chenyang Tao, Lawrence Carin and Jingjing Liu	416
<i>DART: Open-Domain Structured Data Record to Text Generation</i>	
Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher and Nazneen Fatema Rajani	432
<i>When Being Unseen from mBERT is just the Beginning: Handling New Languages With Multilingual Language Models</i>	
Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot and Djamé Seddah	448
<i>Multi-Adversarial Learning for Cross-Lingual Word Embeddings</i>	
Haozhou Wang, James Henderson and Paola Merlo	463
<i>Multi-view Subword Regularization</i>	
Xinyi Wang, Sebastian Ruder and Graham Neubig	473
<i>mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer</i>	
Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua and Colin Raffel	483
<i>MetaXL: Meta Representation Transformation for Low-resource Cross-lingual Learning</i>	
Mengzhou Xia, Guoqing Zheng, Subhabrata Mukherjee, Milad Shokouhi, Graham Neubig and Ahmed Hassan Awadallah	499
<i>Open Domain Question Answering over Tables via Dense Retrieval</i>	
Jonathan Herzig, Thomas Müller, Syrine Krichene and Julian Eisenschlos	512
<i>Open-Domain Question Answering Goes Conversational via Question Rewriting</i>	
Raviteja Anantha, Svitlana Vakulenko, Zhucheng Tu, Shayne Longpre, Stephen Pulman and Srinivas Chappidi	520
<i>QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering</i>	
Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang and Jure Leskovec	535
<i>XOR QA: Cross-lingual Open-Retrieval Question Answering</i>	
Akari Asai, Jungo Kasai, Jonathan Clark, Kenton Lee, Eunsol Choi and Hannaneh Hajishirzi ..	547
<i>SPARTA: Efficient Open-Domain Question Answering via Sparse Transformer Matching Retrieval</i>	
Tiancheng Zhao, Xiaopeng Lu and Kyusong Lee	565

<i>Implicitly Abusive Language – What does it actually look like and why are we not getting there?</i>	
Michael Wiegand, Josef Ruppenhofer and Elisabeth Eder	576
<i>The Importance of Modeling Social Factors of Language: Theory and Practice</i>	
Dirk Hovy and Diyi Yang	588
<i>On learning and representing social meaning in NLP: a sociolinguistic perspective</i>	
Dong Nguyen, Laura Rosseel and Jack Grieve	603
<i>Preregistering NLP research</i>	
Emiel van Miltenburg, Chris van der Lee and Emiel Krahmer	613
<i>Get Your Vitamin C! Robust Fact Verification with Contrastive Evidence</i>	
Tal Schuster, Adam Fisch and Regina Barzilay	624
<i>Representing Numbers in NLP: a Survey and a Vision</i>	
Avijit Thawani, Jay Pujara, Filip Ilievski and Pedro Szekely	644
<i>Extending Multi-Document Summarization Evaluation to the Interactive Setting</i>	
Ori Shapira, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, Yael Amsterdamer and Ido Dagan	657
<i>Identifying Helpful Sentences in Product Reviews</i>	
Iftah Gamzu, Hila Gonen, Gilad Kutiel, Ran Levy and Eugene Agichtein	678
<i>Noisy Self-Knowledge Distillation for Text Summarization</i>	
Yang Liu, Sheng Shen and Mirella Lapata	692
<i>Improving Zero and Few-Shot Abstractive Summarization with Intermediate Fine-tuning and Data Augmentation</i>	
Alexander Fabbri, Simeng Han, Haoyuan Li, Haoran Li, Marjan Ghazvininejad, Shafiq Joty, Dragomir Radev and Yashar Mehdad	704
<i>Enhancing Factual Consistency of Abstractive Summarization</i>	
Chenguang Zhu, William Hinthorn, Ruochen Xu, Qingkai Zeng, Michael Zeng, Xuedong Huang and Meng Jiang	718
<i>Few-shot Intent Classification and Slot Filling with Retrieved Examples</i>	
Dian Yu, Luheng He, Yuan Zhang, Xinya Du, Panupong Pasupat and Qi Li	734
<i>"Nice Try, Kiddo": Investigating Ad Hominems in Dialogue Responses</i>	
Emily Sheng, Kai-Wei Chang, Prem Natarajan and Nanyun Peng	750
<i>Human-like informative conversations: Better acknowledgements using conditional mutual information</i>	
Ashwin Paranjape and Christopher Manning	768
<i>A Comparative Study on Schema-Guided Dialogue State Tracking</i>	
Jie Cao and Yi Zhang	782
<i>Spoken Language Understanding for Task-oriented Dialogue Systems with Augmented Memory Networks</i>	
Jie Wu, Ian Harris and Hongzhi Zhao	797
<i>How to Motivate Your Dragon: Teaching Goal-Driven Agents to Speak and Act in Fantasy Worlds</i>	
Prithviraj Ammanabrolu, Jack Urbanek, Margaret Li, Arthur Szlam, Tim Rocktäschel and Jason Weston	807

<i>Linking Entities to Unseen Knowledge Bases with Arbitrary Schemas</i> Yogarshi Vyas and Miguel Ballesteros	834
<i>Self-Training with Weak Supervision</i> Giannis Karamanolakis, Subhabrata Mukherjee, Guoqing Zheng and Ahmed Hassan Awadallah	845
<i>Neural Language Modeling for Contextualized Temporal Graph Generation</i> Aman Madaan and Yiming Yang	864
<i>Probabilistic Box Embeddings for Uncertain Knowledge Graph Reasoning</i> Xuelu Chen, Michael Boratko, Muhao Chen, Shib Sankar Dasgupta, Xiang Lorraine Li and Andrew McCallum	882
<i>Document-Level Event Argument Extraction by Conditional Generation</i> Sha Li, Heng Ji and Jiawei Han	894
<i>Template Filling with Generative Transformers</i> Xinya Du, Alexander Rush and Claire Cardie	909
<i>Towards Interpreting and Mitigating Shortcut Learning Behavior of NLU models</i> Mengnan Du, Varun Manjunatha, Rajiv Jain, Ruchi Deshpande, Franck Dernoncourt, Jiuxiang Gu, Tong Sun and Xia Hu	915
<i>On Attention Redundancy: A Comprehensive Study</i> Yuchen Bian, Jiaji Huang, Xingyu Cai, Jiahong Yuan and Kenneth Church	930
<i>Does BERT Pretrained on Clinical Notes Reveal Sensitive Data?</i> Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg and Byron Wallace	946
<i>Low-Complexity Probing via Finding Subnetworks</i> Victor Sanh and Alexander Rush	960
<i>An Empirical Comparison of Instance Attribution Methods for NLP</i> Pouya Pezeshkpour, Sarthak Jain, Byron Wallace and Sameer Singh	967
<i>Generalization in Instruction Following Systems</i> Soham Dan, Michael Zhou and Dan Roth	976
<i>LightningDOT: Pre-training Visual-Semantic Embeddings for Real-Time Image-Text Retrieval</i> Siqi Sun, Yen-Chun Chen, Linjie Li, Shuohang Wang, Yuwei Fang and Jingjing Liu	982
<i>Measuring Social Biases in Grounded Vision and Language Embeddings</i> Candace Ross, Boris Katz and Andrei Barbu	998
<i>MTAG: Modal-Temporal Attention Graph for Unaligned Human Multimodal Language Sequences</i> Jianing Yang, Yongxin Wang, Ruitao Yi, Yuying Zhu, Azaan Rehman, Amir Zadeh, Soujanya Poria and Louis-Philippe Morency	1009
<i>Grounding Open-Domain Instructions to Automate Web Support Tasks</i> Nancy Xu, Sam Masling, Michael Du, Giovanni Campagna, Larry Heck, James Landay and Monica Lam	1022
<i>Modular Networks for Compositional Instruction Following</i> Rodolfo Corona, Daniel Fried, Coline Devin, Dan Klein and trevor darrell	1033

<i>Improving Cross-Modal Alignment in Vision Language Navigation via Syntactic Information</i> Jialu Li, Hao Tan and Mohit Bansal	1041
<i>Improving Pretrained Models for Zero-shot Multi-label Text Classification through Reinforced Label Hierarchy Reasoning</i> Hui Liu, Danqing Zhang, Bing Yin and Xiaodan Zhu	1051
<i>Fine-Tuning Pre-trained Language Model with Weak Supervision: A Contrastive-Regularized Self-Training Approach</i> Yue Yu, Simiao Zuo, Haoming Jiang, Wendi Ren, Tuo Zhao and Chao Zhang	1063
<i>Posterior Differential Regularization with f-divergence for Improving Model Robustness</i> Hao Cheng, Xiaodong Liu, Lis Pereira, Yaoliang Yu and Jianfeng Gao	1078
<i>Understanding Hard Negatives in Noise Contrastive Estimation</i> Wenzheng Zhang and Karl Stratos	1090
<i>Certified Robustness to Word Substitution Attack with Differential Privacy</i> Wenjie Wang, Pengfei Tang, Jian Lou and Li Xiong	1102
<i>DReCa: A General Task Augmentation Strategy for Few-Shot Natural Language Inference</i> Shikhar Murty, Tatsunori Hashimoto and Christopher Manning	1113
<i>Harnessing Multilinguality in Unsupervised Machine Translation for Rare Languages</i> Xavier Garcia, Aditya Siddhant, Orhan Firat and Ankur Parikh	1126
<i>Macro-Average: Rare Types Are Important Too</i> Thamme Gowda, Weiqiu You, Constantine Lignos and Jonathan May	1138
<i>Assessing Reference-Free Peer Evaluation for Machine Translation</i> Sweta Agrawal, George Foster, Markus Freitag and Colin Cherry	1158
<i>The Curious Case of Hallucinations in Neural Machine Translation</i> Vikas Raunak, Arul Menezes and Marcin Junczys-Dowmunt	1172
<i>Towards Continual Learning for Multilingual Machine Translation via Vocabulary Substitution</i> Xavier Garcia, Noah Constant, Ankur Parikh and Orhan Firat	1184
<i>Towards Modeling the Style of Translators in Neural Machine Translation</i> Yue Wang, Cuong Hoang and Marcello Federico	1193
<i>Self-Supervised Test-Time Learning for Reading Comprehension</i> Pratyay Banerjee, Tejas Gokhale and Chitta Baral	1200
<i>Capturing Row and Column Semantics in Transformer Based Question Answering over Tables</i> Michael Glass, Mustafa Canim, Alfio Gliozzo, Saneem Chemmengath, Vishwajeet Kumar, Rishav Chakravarti, Avi Sil, Feifei Pan, Samarth Bharadwaj and Nicolas Rodolfo Fauceglia	1212
<i>Explainable Multi-hop Verbal Reasoning Through Internal Monologue</i> Zhengzhong Liang, Steven Bethard and Mihai Surdeanu	1225
<i>Robust Question Answering Through Sub-part Alignment</i> Jifan Chen and Greg Durrett	1251
<i>Text Modular Networks: Learning to Decompose Tasks in the Language of Existing Models</i> Tushar Khot, Daniel Khachabi, Kyle Richardson, Peter Clark and Ashish Sabharwal	1264

<i>RECONSIDER: Improved Re-Ranking using Span-Focused Cross-Attention for Open Domain Question Answering</i>	
Srinivasan Iyer, Sewon Min, Yashar Mehdad and Wen-tau Yih	1280
<i>On the Transferability of Minimal Prediction Preserving Inputs in Question Answering</i>	
Shayne Longpre, Yi Lu and Chris DuBois	1288
<i>Understanding by Understanding Not: Modeling Negation in Language Models</i>	
Arian Hosseini, Siva Reddy, Dzmitry Bahdanau, R Devon Hjelm, Alessandro Sordoni and Aaron Courville	1301
<i>DuoRAT: Towards Simpler Text-to-SQL Models</i>	
Torsten Scholak, Raymond Li, Dzmitry Bahdanau, Harm de Vries and Chris Pal	1313
<i>Looking Beyond Sentence-Level Natural Language Inference for Question Answering and Text Summarization</i>	
Anshuman Mishra, Dhruvesh Patel, Aparna Vijayakumar, Xiang Lorraine Li, Pavan Kapanipathi and Kartik Talamadupula	1322
<i>Structure-Grounded Pretraining for Text-to-SQL</i>	
Xiang Deng, Ahmed Hassan Awadallah, Christopher Meek, Oleksandr Polozov, Huan Sun and Matthew Richardson	1337
<i>Incremental Few-shot Text Classification with Multi-round New Classes: Formulation, Dataset and System</i>	
Congying Xia, Wenpeng Yin, Yihao Feng and Philip Yu	1351
<i>Temporal Reasoning on Implicit Events from Distant Supervision</i>	
Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal and Dan Roth	1361
<i>Disentangling Semantics and Syntax in Sentence Embeddings with Pre-trained Language Models</i>	
James Y. Huang, Kuan-Hao Huang and Kai-Wei Chang	1372
<i>Structure-Aware Abstractive Conversation Summarization via Discourse and Action Graphs</i>	
Jiaao Chen and Diyi Yang	1380
<i>A New Approach to Overgenerating and Scoring Abstractive Summaries</i>	
Kaiqiang Song, Bingqing Wang, Zhe Feng and Fei Liu	1392
<i>D2S: Document-to-Slide Generation Via Query-Based Text Summarization</i>	
Edward Sun, Yufang Hou, Dakuo Wang, Yunfeng Zhang and Nancy X. R. Wang	1405
<i>Efficient Attentions for Long Document Summarization</i>	
Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji and Lu Wang	1419
<i>RefSum: Refactoring Neural Summarization</i>	
Yixin Liu, Zi-Yi Dou and Pengfei Liu	1437
<i>Annotating and Modeling Fine-grained Factuality in Summarization</i>	
Tanya Goyal and Greg Durrett	1449
<i>Larger-Context Tagging: When and Why Does It Work?</i>	
Jinlan Fu, Liangjing Feng, Qi Zhang, Xuanjing Huang and Pengfei Liu	1463
<i>Neural Sequence Segmentation as Determining the Leftmost Segments</i>	
Yangming Li, Lemaou Liu and Kaisheng Yao	1476

<i>PCFGs Can Do Better: Inducing Probabilistic Context-Free Grammars with Many Symbols</i> Songlin Yang, Yanpeng Zhao and Kewei Tu	1487
<i>GEMNET: Effective Gated Gazetteer Representations for Recognizing Complex Entities in Low-context Input</i> Tao Meng, Anjie Fang, Oleg Rokhlenko and Shervin Malmasi	1499
<i>Video-aided Unsupervised Grammar Induction</i> Songyang Zhang, LINFENG Song, LIFENG Jin, Kun Xu, Dong Yu and Jiebo Luo	1513
<i>Generating Negative Samples by Manipulating Golden Responses for Unsupervised Learning of a Response Evaluation Model</i> ChaeHun Park, Eugene Jang, Wonsuk Yang and Jong Park	1525
<i>How Robust are Fact Checking Systems on Colloquial Claims?</i> Byeongchang Kim, Hyunwoo Kim, Seokhee Hong and Gunhee Kim	1535
<i>Fine-grained Post-training for Improving Retrieval-based Dialogue Systems</i> Janghoon Han, Taesuk Hong, Byoungjae Kim, Youngjoong Ko and Jungyun Seo	1549
<i>Put Chatbot into Its Interlocutor's Shoes: New Framework to Learn Chatbot Responding with Intention</i> Hsuan Su, Jiun-Hao Jhan, Fan-yun Sun, Saurav Sahay and Hung-yi Lee	1559
<i>Adding Chit-Chat to Enhance Task-Oriented Dialogues</i> Kai Sun, Seungwhan Moon, Paul Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho and Claire Cardie	1570
<i>Incorporating Syntax and Semantics in Coreference Resolution with Heterogeneous Graph Attention Network</i> Fan Jiang and Trevor Cohn	1584
<i>Context Tracking Network: Graph-based Context Modeling for Implicit Discourse Relation Recognition</i> Yingxue Zhang, Fandong Meng, Peng Li, Ping Jian and Jie Zhou	1592
<i>Improving Neural RST Parsing Model with Silver Agreement Subtrees</i> Naoki Kobayashi, Tsutomu Hirao, Hidetaka Kamigaito, Manabu Okumura and Masaaki Nagata	1600
<i>RST Parsing from Scratch</i> Thanh-Tung Nguyen, Xuan-Phi Nguyen, Shafiq Joty and Xiaoli Li	1613
<i>Did they answer? Subjective acts and intents in conversational discourse</i> Elisa Ferracane, Greg Durrett, Junyi Jessy Li and Katrin Erk	1626
<i>Evaluating the Impact of a Hierarchical Discourse Representation on Entity Coreference Resolution Performance</i> Sopan Khosla, James Fiacco and Carolyn Rosé	1645
<i>Bridging Resolution: Making Sense of the State of the Art</i> Hideo Kobayashi and Vincent Ng	1652
<i>Explicitly Modeling Syntax in Language Models with Incremental Parsing and a Dynamic Oracle</i> Yikang Shen, Shawn Tan, Alessandro Sordoni, Siva Reddy and Aaron Courville	1660
<i>Revisiting the Weaknesses of Reinforcement Learning for Neural Machine Translation</i> Samuel Kiegeand and Julia Kreutzer	1673

<i>Learning to Organize a Bag of Words into Sentences with Neural Networks: An Empirical Study</i>	
Chongyang Tao, Shen Gao, Juntao Li, Yansong Feng, Dongyan Zhao and Rui Yan	1682
<i>Mask Attention Networks: Rethinking and Strengthen Transformer</i>	
Zhihao Fan, Yeyun Gong, Dayiheng Liu, Zhongyu Wei, Siyuan Wang, Jian Jiao, Nan Duan, Ruofei Zhang and Xuanjing Huang	1692
<i>ERNIE-Gram: Pre-Training with Explicitly N-Gram Masked Language Modeling for Natural Language Understanding</i>	
Dongling Xiao, Yu-Kun Li, Han Zhang, Yu Sun, Hao Tian, Hua Wu and Haifeng Wang	1702
<i>Lattice-BERT: Leveraging Multi-Granularity Representations in Chinese Pre-trained Language Models</i>	
Yuxuan Lai, Yijia Liu, Yansong Feng, Songfang Huang and Dongyan Zhao	1716
<i>Modeling Event Plausibility with Consistent Conceptual Abstraction</i>	
Ian Porada, Kaheer Suleman, Adam Trischler and Jackie Chi Kit Cheung	1732
<i>UmlsBERT: Clinical Domain Knowledge Augmentation of Contextual Embeddings Using the Unified Medical Language System Metathesaurus</i>	
George Michalopoulos, Yuanxin Wang, Hussam Kaka, Helen Chen and Alexander Wong	1744
<i>Field Embedding: A Unified Grain-Based Framework for Word Representation</i>	
Junjie Luo, Xi Chen, Jichao Sun, Yuejia Xiang, Ningyu Zhang and Xiang Wan	1754
<i>MelBERT: Metaphor Detection via Contextualized Late Interaction using Metaphorical Identification Theories</i>	
Minjin Choi, Sunkyung Lee, Eunseong Choi, Heesoo Park, Junhyuk Lee, Dongwon Lee and Jongwuk Lee	1763
<i>Non-Parametric Few-Shot Learning for Word Sense Disambiguation</i>	
Howard Chen, Mengzhou Xia and Danqi Chen	1774
<i>Why Do Document-Level Polarity Classifiers Fail?</i>	
Karen Martins, Pedro O.S Vaz-de-Melo and Rodrygo Santos	1782
<i>A Unified Span-Based Approach for Opinion Mining with Syntactic Constituents</i>	
Qingrong Xia, Bo Zhang, Rui Wang, Zhenghua Li, Yue Zhang, Fei Huang, Luo Si and Min Zhang	1795
<i>Target-specified Sequence Labeling with Multi-head Self-attention for Target-oriented Opinion Words Extraction</i>	
Yuhao Feng, Yanghui Rao, Yuyao Tang, Ninghua Wang and He Liu	1805
<i>Does syntax matter? A strong baseline for Aspect-based Sentiment Analysis with RoBERTa</i>	
Junqi Dai, Hang Yan, Tianxiang Sun, Pengfei Liu and Xipeng Qiu	1816
<i>Domain Divergences: A Survey and Empirical Analysis</i>	
Abhinav Ramesh Kashyap, Devamanyu Hazarika, Min-Yen Kan and Roger Zimmermann	1830
<i>Target-Aware Data Augmentation for Stance Detection</i>	
Yingjie Li and Cornelia Caragea	1850
<i>End-to-end ASR to jointly predict transcriptions and linguistic annotations</i>	
Motoi Omachi, Yuya Fujita, Shinji Watanabe and Matthew Wiesner	1861

<i>Source and Target Bidirectional Knowledge Distillation for End-to-end Speech Translation</i>	
Hirofumi Inaguma, Tatsuya Kawahara and Shinji Watanabe	1872
<i>Searchable Hidden Intermediates for End-to-End Models of Decomposable Sequence Tasks</i>	
Siddharth Dalmia, Brian Yan, Vikas Raunak, Florian Metze and Shinji Watanabe	1882
<i>SPLAT: Speech-Language Joint Pre-Training for Spoken Language Understanding</i>	
Yu-An Chung, Chenguang Zhu and Michael Zeng	1897
<i>Worldly Wise (WoW) - Cross-Lingual Knowledge Fusion for Fact-based Visual Spoken-Question Answering</i>	
Kiran Ramnath, Leda Sari, Mark Hasegawa-Johnson and Chang Yoo	1908
<i>Align-Refine: Non-Autoregressive Speech Recognition via Iterative Realignment</i>	
Ethan A. Chi, Julian Salazar and Katrin Kirchhoff	1920
<i>Everything Has a Cause: Leveraging Causal Inference in Legal Text Analysis</i>	
Xiao Liu, Da Yin, Yansong Feng, Yuting Wu and Dongyan Zhao	1928
<i>Counterfactual Supporting Facts Extraction for Explainable Medical Record Based Diagnosis with Graph Network</i>	
Haoran Wu, Wei Chen, Shuang Xu and Bo Xu	1942
<i>Personalized Response Generation via Generative Split Memory Network</i>	
Yuwei Wu, Xuezhe Ma and Diyi Yang	1956
<i>Towards Few-shot Fact-Checking via Perplexity</i>	
Nayeon Lee, Yejin Bang, Andrea Madotto and Pascale Fung	1971
<i>Active² Learning: Actively reducing redundancies in Active Learning methods for Sequence Tagging and Machine Translation</i>	
Rishi Hazra, Parag Dutta, Shubham Gupta, Mohammed Abdul Qaathir and Ambedkar Dukkipati	1982
<i>Generating An Optimal Interview Question Plan Using A Knowledge Graph And Integer Linear Programming</i>	
Soham Datta, Prabir Mallick, Sangameshwar Patil, Indrajit Bhattacharya and Girish Palshikar	1996
<i>Model Extraction and Adversarial Transferability, Your BERT is Vulnerable!</i>	
Xuanli He, Lingjuan Lyu, Lichao Sun and Qiongkai Xu	2006
<i>A Global Past-Future Early Exit Method for Accelerating Inference of Pre-trained Language Models</i>	
Kaiyuan Liao, Yi Zhang, Xuancheng Ren, Qi Su, Xu Sun and Bin He	2013
<i>Masked Conditional Random Fields for Sequence Labeling</i>	
Tianwen Wei, Jianwei Qi, Shenghuan He and Songtao Sun	2024
<i>Heterogeneous Graph Neural Networks for Concept Prerequisite Relation Learning in Educational Data</i>	
Chenghao Jia, Yongliang Shen, Yechun Tang, Lu Sun and Weiming Lu	2036
<i>Be Careful about Poisoned Word Embeddings: Exploring the Vulnerability of the Embedding Layers in NLP Models</i>	
Wenkai Yang, Lei Li, Zhiyuan Zhang, Xuancheng Ren, Xu Sun and Bin He	2048
<i>DA-Transformer: Distance-aware Transformer</i>	
Chuhan Wu, Fangzhao Wu and Yongfeng Huang	2059

<i>ASAP: A Chinese Review Dataset Towards Aspect Category Sentiment Analysis and Rating Prediction</i> Jiahao Bu, Lei Ren, Shuang Zheng, Yang Yang, Jingang Wang, Fuzheng Zhang and Wei Wu .	2069
<i>Are NLP Models really able to Solve Simple Math Word Problems?</i> Arkil Patel, Satwik Bhattamishra and Navin Goyal	2080
<i>WRIME: A New Dataset for Emotional Intensity Estimation with Subjective and Objective Annotations</i> Tomoyuki Kajiwara, Chenhui Chu, Noriko Takemura, Yuta Nakashima and Hajime Nagahara	2095
<i>KPQA: A Metric for Generative Question Answering Using Keyphrase Weights</i> Hwanhee Lee, Seunghyun Yoon, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Joongbo Shin and Kyomin Jung	2105
<i>StylePTB: A Compositional Benchmark for Fine-grained Controllable Text Style Transfer</i> Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard Hovy, Barnabás Póczos, Ruslan Salakhutdinov and Louis-Philippe Morency	2116
<i>Blow the Dog Whistle: A Chinese Dataset for Cant Understanding with Common Sense and World Knowledge</i> Canwen Xu, Wangchunshu Zhou, Tao Ge, Ke Xu, Julian McAuley and Furu Wei	2139
<i>COVID-19 Named Entity Recognition for Vietnamese</i> Thinh Hung Truong, Mai Hoang Dao and Dat Quoc Nguyen	2146
<i>Framing Unpacked: A Semi-Supervised Interpretable Multi-View Model of Media Frames</i> Shima Khanehzar, Trevor Cohn, Gosia Mikolajczak, Andrew Turpin and Lea Frermann	2154
<i>Automatic Classification of Neutralization Techniques in the Narrative of Climate Change Scepticism</i> Shraey Bhatia, Jey Han Lau and Timothy Baldwin	2167
<i>Suicide Ideation Detection via Social and Temporal User Representations using Hyperbolic Learning</i> Ramit Sawhney, Harshit Joshi, Rajiv Ratn Shah and Lucie Flek	2176
<i>WikiTalkEdit: A Dataset for modeling Editors' behaviors on Wikipedia</i> Kokil Jaidka, Andrea Ceolin, Iknoor Singh, Niyati Chhaya and Lyle Ungar	2191
<i>The structure of online social networks modulates the rate of lexical change</i> Jian Zhu and David Jurgen	2201
<i>Modeling Framing in Immigration Discourse on Social Media</i> Julia Mendelsohn, Ceren Budak and David Jurgen	2219
<i>Modeling the Severity of Complaints in Social Media</i> Mali Jin and Nikolaos Aletras	2264
<i>What About the Precedent: An Information-Theoretic Analysis of Common Law</i> Josef Valvoda, Tiago Pimentel, Niklas Stoehr, Ryan Cotterell and Simone Teufel	2275
<i>Introducing CAD: the Contextual Abuse Dataset</i> Bertie Vidgen, Dong Nguyen, Helen Margetts, Patricia Rossini and Rebekah Tromble	2289
<i>Lifelong Learning of Hate Speech Classification on Social Media</i> Jing Qian, Hong Wang, Mai ElSherief and Xifeng Yan	2304

<i>Learning to Recognize Dialect Features</i>	
Dorottya Demszky, Devyani Sharma, Jonathan Clark, Vinodkumar Prabhakaran and Jacob Eisenstein	2315
<i>It's Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners</i>	
Timo Schick and Hinrich Schütze	2339
<i>Static Embeddings as Efficient Knowledge Bases?</i>	
Philipp Dufter, Nora Kassner and Hinrich Schütze	2353
<i>Highly Efficient Knowledge Graph Embedding Learning with Orthogonal Procrustes Analysis</i>	
Xutan Peng, Guanyi Chen, Chenghua Lin and Mark Stevenson	2364
<i>Rethinking Network Pruning – under the Pre-train and Fine-tune Paradigm</i>	
Dongkuan Xu, Ian En-Hsu Yen, Jinxi Zhao and Zhibin Xiao	2376
<i>Towards a Comprehensive Understanding and Accurate Evaluation of Societal Biases in Pre-Trained Transformers</i>	
Andrew Silva, Pradyumna Tambwekar and Matthew Gombolay	2383
<i>Detoxifying Language Models Risks Marginalizing Minority Voices</i>	
Albert Xu, Eshaan Pathak, Eric Wallace, Suchin Gururangan, Maarten Sap and Dan Klein ...	2390
<i>HONEST: Measuring Hurtful Sentence Completion in Language Models</i>	
Debora Nozza, Federico Bianchi and Dirk Hovy	2398
<i>EaSe: A Diagnostic Tool for VQA based on Answer Diversity</i>	
Shailza Jolly, Sandro Pezzelle and Moin Nabi	2407
<i>DeCEMBERT: Learning from Noisy Instructional Videos via Dense Captions and Entropy Minimization</i>	
Zineng Tang, Jie Lei and Mohit Bansal	2415
<i>Improving Generation and Evaluation of Visual Stories via Semantic Consistency</i>	
Adyasha Maharana, Darryl Hannan and Mohit Bansal	2427
<i>Multilingual Multimodal Pre-training for Zero-Shot Cross-Lingual Transfer of Vision-Language Models</i>	
Po-Yao Huang, Mandela Patrick, Junjie Hu, Graham Neubig, Florian Metze and Alexander Hauptmann	2443
<i>Video Question Answering with Phrases via Semantic Roles</i>	
Arka Sadhu, Kan Chen and Ram Nevatia	2460
<i>From Masked Language Modeling to Translation: Non-English Auxiliary Tasks Improve Zero-shot Spoken Language Understanding</i>	
Rob van der Goot, Ibrahim Sharaf, Aizhan Imankulova, Ahmet Üstün, Marija Stepanović, Alan Ramponi, Siti Oryza Khairunnisa, Mamoru Komachi and Barbara Plank	2479
<i>WEC: Deriving a Large-scale Cross-document Event Coreference dataset from Wikipedia</i>	
Alon Eirew, Arie Cattan and Ido Dagan	2498
<i>Challenging distributional models with a conceptual network of philosophical terms</i>	
Yvette Oortwijn, Jelke Bloem, Pia Sommerauer, Francois Meyer, Wei Zhou and Antske Fokkens	2511

<i>KILT: a Benchmark for Knowledge Intensive Language Tasks</i>	
Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel and Sebastian Riedel	2523
<i>A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios</i>	
Michael A. Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen and Dietrich Klakow	2545
<i>Temporal Knowledge Graph Completion using a Linear Temporal Regularizer and Multivector Embeddings</i>	
Chengjin Xu, Yung-Yu Chen, Mojtaba Nayyeri and Jens Lehmann	2569
<i>UDALM: Unsupervised Domain Adaptation through Language Modeling</i>	
Constantinos Karouzou, Georgios Paraskevopoulos and Alexandros Potamianos	2579
<i>Beyond Black & White: Leveraging Annotator Disagreement via Soft-Label Multi-Task Learning</i>	
Tommaso Fornaciari, Alexandra Uma, Silviu Paun, Barbara Plank, Dirk Hovy and Massimo Poesio	2591
<i>Clustering-based Inference for Biomedical Entity Linking</i>	
Rico Angell, Nicholas Monath, Sunil Mohan, Nishant Yadav and Andrew McCallum	2598
<i>Variance-reduced First-order Meta-learning for Natural Language Processing Tasks</i>	
Lingxiao Wang, Kevin Huang, Tengyu Ma, Quanquan Gu and Jing Huang	2609
<i>Diversity-Aware Batch Active Learning for Dependency Parsing</i>	
Tianze Shi, Adrian Benton, Igor Malioutov and Ozan Irsoy	2616
<i>How many data points is a prompt worth?</i>	
Teven Le Scao and Alexander Rush	2627
<i>Can Latent Alignments Improve Autoregressive Machine Translation?</i>	
Adi Haviv, Lior Vassertail and Omer Levy	2637
<i>Smoothing and Shrinking the Sparse Seq2Seq Search Space</i>	
Ben Peters and André F. T. Martins	2642
<i>Unified Pre-training for Program Understanding and Generation</i>	
Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray and Kai-Wei Chang	2655
<i>Hyperparameter-free Continuous Learning for Domain Classification in Natural Language Understanding</i>	
Ting Hua, Yilin Shen, Changsheng Zhao, Yen-Chang Hsu and Hongxia Jin	2669
<i>On the Embeddings of Variables in Recurrent Neural Networks for Source Code</i>	
Nadezhda Chirkova	2679
<i>Cross-Lingual Word Embedding Refinement by ℓ_1 Norm Optimisation</i>	
Xutan Peng, Chenghua Lin and Mark Stevenson	2690
<i>Semantic Frame Forecast</i>	
Chieh-Yang Huang and Ting-Hao Huang	2702
<i>MUSER: MULTImodal Stress detection using Emotion Recognition as an Auxiliary Task</i>	
Yiqun Yao, Michalis Papakostas, Mihai Burzo, Mohamed Abouelenien and Rada Mihalcea ..	2714

<i>Learning to Decompose and Organize Complex Tasks</i>	
Yi Zhang, Sujay Kumar Jauhar, Julia Kiseleva, Ryen White and Dan Roth	2726
<i>Continual Learning for Text Classification with Information Disentanglement Based Regularization</i>	
Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang and Diyi Yang	2736
<i>Learning from Executions for Semantic Parsing</i>	
Bailin Wang, Mirella Lapata and Ivan Titov	2747
<i>Learning to Synthesize Data for Semantic Parsing</i>	
Bailin Wang, Wenpeng Yin, Xi Victoria Lin and Caiming Xiong	2760
<i>Edge: Enriching Knowledge Graph Embeddings with External Text</i>	
Saed Rezayi, Handong Zhao, Sungchul Kim, Ryan Rossi, Nedim Lipka and Sheng Li	2767
<i>FLIN: A Flexible Natural Language Interface for Web Navigation</i>	
Sahisnu Mazumder and Oriana Riva	2777
<i>Game-theoretic Vocabulary Selection via the Shapley Value and Banzhaf Index</i>	
Roma Patel, Marta Garnelo, Ian Gemp, Chris Dyer and Yoram Bachrach	2789
<i>Incorporating External Knowledge to Enhance Tabular Reasoning</i>	
J. Neeraja, Vivek Gupta and Vivek Srikumar	2799
<i>Compositional Generalization for Neural Semantic Parsing via Span-level Supervised Attention</i>	
Pengcheng Yin, Hao Fang, Graham Neubig, Adam Pauls, Emmanouil Antonios Platanios, Yu Su, Sam Thomson and Jacob Andreas	2810
<i>Domain Adaptation for Arabic Cross-Domain and Cross-Dialect Sentiment Analysis from Contextualized Word Embedding</i>	
Abdellah El Mekki, Abdelkader El Mahdaouy, Ismail Berrada and Ahmed Khoumsi	2824
<i>Multi-task Learning of Negation and Speculation for Targeted Sentiment Classification</i>	
Andrew Moore and Jeremy Barnes	2838
<i>A Disentangled Adversarial Neural Topic Model for Separating Opinions from Plots in User Reviews</i>	
Gabriele Pergola, Lin Gui and Yulan He	2870
<i>Graph Ensemble Learning over Multiple Dependency Trees for Aspect-level Sentiment Classification</i>	
Xiaochen Hou, Peng Qi, Guangtao Wang, Rex Ying, Jing Huang, Xiaodong He and Bowen Zhou	2884
<i>Emotion-Infused Models for Explainable Psychological Stress Detection</i>	
Elsbeth Turcan, Smaranda Muresan and Kathleen McKeown	2895
<i>Aspect-based Sentiment Analysis with Type-aware Graph Convolutional Networks and Layer Ensemble</i>	
Yuanhe Tian, Guimin Chen and Yan Song	2910
<i>Supertagging-based Parsing with Linear Context-free Rewriting Systems</i>	
Thomas Ruprecht and Richard Mörbitz	2923
<i>Outside Computation with Superior Functions</i>	
Parker Riley and Daniel Gildea	2936
<i>Learning Syntax from Naturally-Occurring Bracketings</i>	
Tianze Shi, Ozan İrsoy, Igor Malioutov and Lillian Lee	2941

<i>Bot-Adversarial Dialogue for Safe Conversational Agents</i>	
Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston and Emily Dinan	2950
<i>Non-Autoregressive Semantic Parsing for Compositional Task-Oriented Dialog</i>	
Arun Babu, Akshat Shrivastava, Armen Aghajanyan, Ahmed Aly, Angela Fan and Marjan Ghazvininejad	2969
<i>Example-Driven Intent Prediction with Observers</i>	
Shikib Mehri and Mihail Eric	2979
<i>Imperfect also Deserves Reward: Multi-Level and Sequential Reward Modeling for Better Dialog Management</i>	
Zhengxu Hou, Bang Liu, Ruihui Zhao, Zijing Ou, Yafei Liu, Xi Chen and Yefeng Zheng	2993
<i>Action-Based Conversations Dataset: A Corpus for Building More In-Depth Task-Oriented Dialogue Systems</i>	
Derek Chen, Howard Chen, Yi Yang, Alexander Lin and Zhou Yu	3002
<i>Controlling Dialogue Generation with Semantic Exemplars</i>	
Prakhar Gupta, Jeffrey Bigham, Yulia Tsvetkov and Amy Pavel	3018
<i>COIL: Revisit Exact Lexical Match in Information Retrieval with Contextualized Inverted List</i>	
Luyu Gao, Zhu Yun Dai and Jamie Callan	3030
<i>X-Class: Text Classification with Extremely Weak Supervision</i>	
Zihan Wang, Dheeraj Mekala and Jingbo Shang	3043
<i>Fine-tuning Encoders for Improved Monolingual and Zero-shot Polylingual Neural Topic Modeling</i>	
Aaron Mueller and Mark Dredze	3054
<i>Exploring the Relationship Between Algorithm Performance, Vocabulary, and Run-Time in Text Classification</i>	
Wilson Fearn, Orion Weller and Kevin Seppi	3069
<i>Faithfully Explainable Recommendation via Neural Logic Reasoning</i>	
Yaxin Zhu, Yikun Xian, Zuohui Fu, Gerard de Melo and Yongfeng Zhang	3083
<i>You Sound Like Someone Who Watches Drama Movies: Towards Predicting Movie Preferences from Conversational Interactions</i>	
Sergey Volokhin, Joyce Ho, Oleg Rokhlenko and Eugene Agichtein	3091
<i>Reading and Acting while Blindfolded: The Need for Semantics in Text Game Agents</i>	
Shunyu Yao, Karthik Narasimhan and Matthew Hausknecht	3097
<i>SOrT-ing VQA Models : Contrastive Gradient Learning for Improved Consistency</i>	
Sameer Dharur, Purva Tendulkar, Dhruv Batra, Devi Parikh and Ramprasaath R. Selvaraju	3103
<i>Semi-Supervised Policy Initialization for Playing Games with Language Hints</i>	
Tsu-Jui Fu and William Yang Wang	3112
<i>Revisiting Document Representations for Large-Scale Zero-Shot Learning</i>	
Jihyung Kil and Wei-Lun Chao	3117
<i>Negative language transfer in learner English: A new dataset</i>	
Leticia Farias Wanderley, Nicole Zhao and Carrie Demmans Epp	3129

<i>SentSim: Crosslingual Semantic Evaluation of Machine Translation</i> Yurun Song, Junchen Zhao and Lucia Specia	3143
<i>Quality Estimation for Image Captions Based on Large-scale Human Evaluations</i> Tomer Levinboim, Ashish V. Thapliyal, Piyush Sharma and Radu Soricut	3157
<i>CaSiNo: A Corpus of Campsite Negotiation Dialogues for Automatic Negotiation Systems</i> Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May and Jonathan Gratch	3167
<i>News Headline Grouping as a Challenging NLU Task</i> Philippe Laban, Lucas Bandarkar and Marti A. Hearst	3186
<i>Olá, Bonjour, Salve! XFORMAL: A Benchmark for Multilingual Formality Style Transfer</i> Eleftheria Briakou, Di Lu, Ke Zhang and Joel Tetreault	3199
<i>Grouping Words with Semantic Diversity</i> Karine Chubarian, Abdul Rafae Khan, Anastasios Sidiropoulos and Jia Xu	3217
<i>Noise Stability Regularization for Improving BERT Fine-tuning</i> Hang Hua, Xingjian Li, Dejing Dou, Chengzhong Xu and Jiebo Luo	3229
<i>FlowPrior: Learning Expressive Priors for Latent Variable Sentence Models</i> Xiaoan Ding and Kevin Gimpel	3242
<i>HTCInfoMax: A Global Model for Hierarchical Text Classification via Information Maximization</i> Zhongfen Deng, Hao Peng, Dongxiao He, Jianxin Li and Philip Yu	3259
<i>Knowledge Guided Metric Learning for Few-Shot Text Classification</i> Dianbo Sui, Yubo Chen, Binjie Mao, Delai Qiu, Kang Liu and Jun Zhao	3266
<i>Ensemble of MRR and NDCG models for Visual Dialog</i> Idan Schwartz	3272
<i>Supervised Neural Clustering via Latent Structured Output Learning: Application to Question Intents</i> Iryna Haponchyk and Alessandro Moschitti	3364
<i>ConVEx: Data-Efficient and Few-Shot Slot Labeling</i> Matthew Henderson and Ivan Vulić	3375
<i>CREAD: Combined Resolution of Ellipses and Anaphora in Dialogues</i> Bo-Hsiang Tseng, Shruti Bhargava, Jiarui Lu, Joel Ruben Antony Moniz, Dhivya Piraviperumal, Lin Li and Hong Yu	3390
<i>Knowledge-Driven Slot Constraints for Goal-Oriented Dialogue Systems</i> Piyawat Lertvittayakumjorn, Daniele Bonadiman and Saab Mansour	3407
<i>Clipping Loops for Sample-Efficient Dialogue Policy Optimisation</i> Yen-Chen Wu and Carl Edward Rasmussen	3420
<i>Integrating Lexical Information into Entity Neighbourhood Representations for Relation Prediction</i> Ian Wood, Mark Johnson and Stephen Wan	3429
<i>Noisy-Labeled NER with Confidence Estimation</i> Kun Liu, Yao Fu, Chuanqi Tan, Moshu Chen, Ningyu Zhang, Songfang Huang and Sheng Gao	3437

<i>TABBIE: Pretrained Representations of Tabular Data</i>	
Hiroshi Iida, Dung Thai, Varun Manjunatha and Mohit Iyyer	3446
<i>Better Feature Integration for Named Entity Recognition</i>	
Lu Xu, Zhanming Jie, Wei Lu and Lidong Bing	3457
<i>ZS-BERT: Towards Zero-Shot Relation Extraction with Attribute Representation Learning</i>	
Chih-Yao Chen and Cheng-Te Li	3470
<i>Graph Convolutional Networks for Event Causality Identification with Rich Document-level Structures</i>	
Minh Tran Phu and Thien Huu Nguyen	3480
<i>A Context-Dependent Gated Module for Incorporating Symbolic Semantics into Event Coreference Resolution</i>	
Tuan Lai, Heng Ji, Trung Bui, Quan Hung Tran, Franck Dernoncourt and Walter Chang	3491
<i>Multi-Style Transfer with Discriminative Feedback on Disjoint Corpus</i>	
Navita Goyal, Balaji Vasan Srinivasan, Anandhavelu N and Abhilasha Sancheti	3500
<i>FUDGE: Controlled Text Generation With Future Discriminators</i>	
Kevin Yang and Dan Klein	3511
<i>Controllable Text Simplification with Explicit Paraphrasing</i>	
Mounica Maddela, Fernando Alva-Manchego and Wei Xu	3536
<i>Knowledge Graph Based Synthetic Corpus Generation for Knowledge-Enhanced Language Model Pre-training</i>	
Oshin Agarwal, Heming Ge, Siamak Shakeri and Rami Al-Rfou	3554
<i>Choose Your Own Adventure: Paired Suggestions in Collaborative Writing for Evaluating Story Generation Models</i>	
Elizabeth Clark and Noah A. Smith	3566
<i>InfoXLM: An Information-Theoretic Framework for Cross-Lingual Language Model Pre-Training</i>	
Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang and Ming Zhou	3576
<i>Context-Interactive Pre-Training for Document Machine Translation</i>	
Pengcheng Yang, Pei Zhang, Boxing Chen, Jun Xie and Weihua Luo	3589
<i>Code-Mixing on Sesame Street: Dawn of the Adversarial Polyglots</i>	
Samson Tan and Shafiq Joty	3596
<i>X-METRA-ADA: Cross-lingual Meta-Transfer learning Adaptation to Natural Language Understanding and Question Answering</i>	
Meryem M’hamdi, Doo Soon Kim, Franck Dernoncourt, Trung Bui, Xiang Ren and Jonathan May	3617
<i>Explicit Alignment Objectives for Multilingual Bidirectional Encoders</i>	
Junjie Hu, Melvin Johnson, Orhan Firat, Aditya Siddhant and Graham Neubig	3633
<i>Cross-lingual Cross-modal Pretraining for Multimodal Retrieval</i>	
Hongliang Fei, Tan Yu and Ping Li	3644

<i>Wikipedia Entities as Rendezvous across Languages: Grounding Multilingual Language Models by Predicting Wikipedia Hyperlinks</i>	
Iacer Calixto, Alessandro Raganato and Tommaso Pasini	3651
<i>multiPROver: Generating Multiple Proofs for Improved Interpretability in Rule Reasoning</i>	
Swarnadeep Saha, Prateek Yadav and Mohit Bansal	3662
<i>Adaptable and Interpretable Neural Memory Over Symbolic Knowledge</i>	
Pat Verga, Haitian Sun, Livio Baldini Soares and William Cohen	3678
<i>CLEVR_HYP: A Challenge Dataset and Baselines for Visual Question Answering with Hypothetical Actions over Images</i>	
Shailaja Keyur Sampat, Akshay Kumar, Yezhou Yang and Chitta Baral	3692
<i>Refining Targeted Syntactic Evaluation of Language Models</i>	
Benjamin Newman, Kai-Siang Ang, Julia Gong and John Hewitt	3710
<i>Universal Adversarial Attacks with Natural Triggers for Text Classification</i>	
Liwei Song, Xinwei Yu, Hsuan-Tung Peng and Karthik Narasimhan	3724
<i>QuadrupletBERT: An Efficient Model For Embedding-Based Large-Scale Retrieval</i>	
Peiyang Liu, Sen Wang, Xi Wang, Wei Ye and Shikun Zhang	3734
<i>Dynamically Disentangling Social Bias from Task-Oriented Representations with Adversarial Attack</i>	
Liwen Wang, Yuanmeng Yan, Keqing He, Yanan Wu and Weiran Xu	3740
<i>An Empirical Investigation of Bias in the Multimodal Analysis of Financial Earnings Calls</i>	
Ramit Sawhney, Arshiya Aggarwal and Rajiv Ratn Shah	3751
<i>Beyond Fair Pay: Ethical Implications of NLP Crowdsourcing</i>	
Boaz Shmueli, Jan Fell, Soumya Ray and Lun-Wei Ku	3758
<i>On Transferability of Bias Mitigation Effects in Language Model Fine-Tuning</i>	
Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves and Xiang Ren	3770
<i>Case Study: Deontological Ethics in NLP</i>	
Shrimai Prabhumoye, Brendon Boldt, Ruslan Salakhutdinov and Alan W Black	3784
<i>Privacy Regularization: Joint Privacy-Utility Optimization in Language Models</i>	
Fatemehsadat Miresghallah, Huseyin Inan, Marcello Hasegawa, Victor Rühle, Taylor Berg-Kirkpatrick and Robert Sim	3799
<i>On the Impact of Random Seeds on the Fairness of Clinical Classifiers</i>	
Silvio Amir, Jan-Willem van de Meent and Byron Wallace	3808
<i>Topic Model or Topic Twaddle? Re-evaluating Semantic Interpretability Measures</i>	
Caitlin Doogan and Wray Buntine	3824
<i>Discourse Probing of Pretrained Language Models</i>	
Fajri Koto, Jey Han Lau and Timothy Baldwin	3849
<i>UniDrop: A Simple yet Effective Technique to Improve Transformer without Extra Cost</i>	
Zhen Wu, Lijun Wu, Qi Meng, Yingce Xia, Shufang Xie, Tao Qin, Xinyu Dai and Tie-Yan Liu	3865

<i>tWT-WT: A Dataset to Assert the Role of Target Entities for Detecting Stance of Tweets</i> Ayush Kaushal, Avirup Saha and Niloy Ganguly	3879
<i>Learning to Learn to be Right for the Right Reasons</i> Pride Kavumba, Benjamin Heinzerling, Ana Brassard and Kentaro Inui	3890
<i>Double Perturbation: On the Robustness of Robustness and Counterfactual Bias Evaluation</i> Chong Zhang, Jieyu Zhao, Huan Zhang, Kai-Wei Chang and Cho-Jui Hsieh	3899
<i>Explaining Neural Network Predictions on Sentence Pairs via Learning Word-Group Masks</i> Hanjie Chen, Song Feng, Jatin Ganhotra, Hui Wan, Chulaka Gunasekara, Sachindra Joshi and Yangfeng Ji	3917
<i>Almost Free Semantic Draft for Neural Machine Translation</i> Xi Ai and Bin Fang	3931
<i>Pruning-then-Expanding Model for Domain Adaptation of Neural Machine Translation</i> Shuhao Gu, Yang Feng and Wanying Xie	3942
<i>Multi-Hop Transformer for Document-Level Machine Translation</i> Long Zhang, Tong Zhang, Haibo Zhang, Baosong Yang, Wei Ye and Shikun Zhang	3953
<i>Continual Learning for Neural Machine Translation</i> Yue Cao, Hao-Ran Wei, Boxing Chen and Xiaojun Wan	3964
<i>Self-Training for Unsupervised Neural Machine Translation in Unbalanced Training Data Scenarios</i> Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita and Tiejun Zhao	3975
<i>Smart-Start Decoding for Neural Machine Translation</i> Jian Yang, Shuming Ma, Dongdong Zhang, Juncheng Wan, Zhoujun Li and Ming Zhou	3982
<i>Multi-Task Learning with Shared Encoder for Non-Autoregressive Machine Translation</i> Yongchang Hao, Shilin He, Wenxiang Jiao, Zhaopeng Tu, Michael Lyu and Xing Wang	3989
<i>ER-AE: Differentially Private Text Generation for Authorship Anonymization</i> Haohan Bo, Steven H. H. Ding, Benjamin C. M. Fung and Farkhund Iqbal	3997
<i>Distantly Supervised Transformers For E-Commerce Product QA</i> Happy Mittal, Aniket Chakrabarti, Belhassen Bayar, Animesh Anant Sharma and Nikhil Rasiwasia	4008
<i>Quantitative Day Trading from Natural Language using Reinforcement Learning</i> Ramit Sawhney, Arnav Wadhwa, Shivam Agarwal and Rajiv Ratn Shah	4018
<i>Restoring and Mining the Records of the Joseon Dynasty via Neural Language Modeling and Machine Translation</i> Kyeongpil Kang, Kyohoon Jin, Soyoung Yang, Soojin Jang, Jaegul Choo and Youngbin Kim	4031
<i>Modeling Diagnostic Label Correlation for Automatic ICD Coding</i> Shang-Chi Tsai, Chao-Wei Huang and Yun-Nung Chen	4043
<i>Self-Supervised Contrastive Learning for Efficient User Satisfaction Prediction in Conversational Agents</i> Mohammad Kachuee, Hao Yuan, Young-Bum Kim and Sungjin Lee	4053
<i>A recipe for annotating grounded clarifications</i> Luciana Benotti and Patrick Blackburn	4065

<i>Grey-box Adversarial Attack And Defence For Sentiment Classification</i>	
Ying Xu, Xu Zhong, Antonio Jimeno Yepes and Jey Han Lau	4078
<i>How low is too low? A monolingual take on lemmatisation in Indian languages</i>	
Kumar Saunack, Kumar Saurav and Pushpak Bhattacharyya	4088
<i>Causal Effects of Linguistic Properties</i>	
Reid Pryzant, Dallas Card, Dan Jurafsky, Victor Veitch and Dhanya Sridhar	4095
<i>Dynabench: Rethinking Benchmarking in NLP</i>	
Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts and Adina Williams	4110
<i>Translational NLP: A New Paradigm and General Principles for Natural Language Processing Research</i>	
Denis Newman-Griffis, Jill Fain Lehman, Carolyn Rosé and Harry Hochheiser	4125
<i>Predicting Discourse Trees from Transformer-based Neural Summarizers</i>	
Wen Xiao, Patrick Huber and Giuseppe Carenini	4139
<i>Probing for Bridging Inference in Transformer Language Models</i>	
Onkar Pandit and Yufang Hou	4153
<i>Is Incoherence Surprising? Targeted Evaluation of Coherence Prediction from Language Models</i>	
Anne Beyer, Sharid Loáiciga and David Schlangen	4164
<i>Stay Together: A System for Single and Split-antecedent Anaphora Resolution</i>	
Juntao Yu, Nafise Sadat Moosavi, Silviu Paun and Massimo Poesio	4174
<i>Redefining Absent Keyphrases and their Effect on Retrieval Effectiveness</i>	
Florian Boudin and Ygor Gallina	4185
<i>CoRT: Complementary Rankings from Transformers</i>	
Marco Wrzalik and Dirk Krechel	4194
<i>Multi-source Neural Topic Modeling in Multi-view Embedding Spaces</i>	
Pankaj Gupta, Yatin Chaudhary and Hinrich Schütze	4205
<i>Inductive Topic Variational Graph Auto-Encoder for Text Classification</i>	
Qianqian Xie, Jimin Huang, Pan Du, Min Peng and Jian-Yun Nie	4218
<i>Self-Alignment Pretraining for Biomedical Entity Representations</i>	
Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella and Nigel Collier	4228
<i>TaxoClass: Hierarchical Multi-Label Text Classification Using Only Class Names</i>	
Jiaming Shen, Wenda Qiu, Yu Meng, Jingbo Shang, Xiang Ren and Jiawei Han	4239
<i>MERMAID: Metaphor Generation with Symbolism and Discriminative Decoding</i>	
Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan and Nanyun Peng	4250
<i>On Learning Text Style Transfer with Direct Rewards</i>	
Yixin Liu, Graham Neubig and John Wieting	4262
<i>Focused Attention Improves Document-Grounded Generation</i>	
Shrimai Prabhumoye, Kazuma Hashimoto, Yingbo Zhou, Alan W Black and Ruslan Salakhutdinov	4274

<i>NeuroLogic Decoding: (Un)supervised Neural Text Generation with Predicate Logic Constraints</i> Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula and Yejin Choi	4288
<i>Ask what’s missing and what’s useful: Improving Clarification Question Generation using Global Knowledge</i> Bodhisattwa Prasad Majumder, Sudha Rao, Michel Galley and Julian McAuley	4300
<i>Progressive Generation of Long Text with Pretrained Language Models</i> Bowen Tan, Zichao Yang, Maruan Al-Shedivat, Eric Xing and Zhiting Hu	4313
<i>SOCCKER: An Information-Sparse Discourse State Tracking Collection in the Sports Commentary Domain</i> Ruochen Zhang and Carsten Eickhoff	4325
<i>Plot-guided Adversarial Example Construction for Evaluating Open-domain Story Generation</i> Sarik Ghazarian, Zixi Liu, Akash S M, Ralph Weischedel, Aram Galstyan and Nanyun Peng	4334
<i>MultiOpEd: A Corpus of Multi-Perspective News Editorials</i> Siyi Liu, Sihao Chen, Xander Uyttendaele and Dan Roth	4345
<i>Swords: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality</i> Mina Lee, Chris Donahue, Robin Jia, Alexander Iyabor and Percy Liang	4362
<i>"I'm Not Mad": Commonsense Implications of Negation and Contradiction</i> Liwei Jiang, Antoine Bosselut, Chandra Bhagavatula and Yejin Choi	4380
<i>Identifying Medical Self-Disclosure in Online Communities</i> Mina Valizadeh, Pardis Ranjbar-Noiey, Cornelia Caragea and Natalie Parde	4398
<i>Language in a (Search) Box: Grounding Language Learning in Real-World Human-Machine Interaction</i> Federico Bianchi, Ciro Greco and Jacopo Tagliabue	4409
<i>Finding Concept-specific Biases in Form–Meaning Associations</i> Tiago Pimentel, Brian Roark, Søren Wichmann, Ryan Cotterell and Damián Blasi	4416
<i>How (Non-)Optimal is the Lexicon?</i> Tiago Pimentel, Irene Nikkarinen, Kyle Mahowald, Ryan Cotterell and Damián Blasi	4426
<i>Word Complexity is in the Eye of the Beholder</i> Sian Gooding, Ekaterina Kochmar, Seid Muhie Yimam and Chris Biemann	4439
<i>Linguistic Complexity Loss in Text-Based Therapy</i> Jason Wei, Kelly Finn, Emma Templeton, Thalia Wheatley and Soroush Vosoughi	4450
<i>Ab Antiquo: Neural Proto-language Reconstruction</i> Carlo Meloni, Shauli Ravfogel and Yoav Goldberg	4460
<i>On Biasing Transformer Attention Towards Monotonicity</i> Annette Rios, Chantal Amrhein, Noëmi Aepli and Rico Sennrich	4474
<i>Extracting a Knowledge Base of Mechanisms from COVID-19 Papers</i> Tom Hope, Aida Amini, David Wadden, Madeleine van Zuylen, Sravanthi Parasa, Eric Horvitz, Daniel Weld, Roy Schwartz and Hannaneh Hajishirzi	4489
<i>Constrained Multi-Task Learning for Event Coreference Resolution</i> Jing Lu and Vincent Ng	4504

<i>Empirical Evaluation of Pre-trained Transformers for Human-Level NLP: The Role of Sample Size and Dimensionality</i>	
Adithya V Ganesan, Matthew Matero, Aravind Reddy Ravula, Huy Vu and H. Andrew Schwartz	4515
<i>Leveraging Deep Representations of Radiology Reports in Survival Analysis for Predicting Heart Failure Patient Mortality</i>	
Hyun Gi Lee, Evan Sholle, Ashley Beecy, Subhi Al’Aref and Yifan Peng	4533
<i>On the Use of Context for Predicting Citation Worthiness of Sentences in Scholarly Articles</i>	
Rakesh Gosangi, Ravneet Arora, Mohsen Gheisarieha, Debanjan Mahata and Haimin Zhang	4539
<i>Data and Model Distillation as a Solution for Domain-transferable Fact Verification</i>	
Mitch Paul Mithun, Sandeep Suntwal and Mihai Surdeanu	4546
<i>Adapting Coreference Resolution for Processing Violent Death Narratives</i>	
Ankith Uppunda, Susan Cochran, Jacob Foster, Alina Arseniev-Koehler, Vickie Mays and Kai-Wei Chang	4553
<i>Time-Stamped Language Model: Teaching Language Models to Understand The Flow of Events</i>	
Hossein Rajaby Faghihi and Parisa Kordjamshidi	4560
<i>If You Want to Go Far Go Together: Unsupervised Joint Candidate Evidence Retrieval for Multi-hop Question Answering</i>	
Vikas Yadav, Steven Bethard and Mihai Surdeanu	4571
<i>SPARTQA: A Textual Question Answering Benchmark for Spatial Reasoning</i>	
Roshanak Mirzaee, Hossein Rajaby Faghihi, Qiang Ning and Parisa Kordjamshidi	4582
<i>A Dataset of Information-Seeking Questions and Answers Anchored in Research Papers</i>	
Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith and Matt Gardner	4599
<i>Differentiable Open-Ended Commonsense Reasoning</i>	
Bill Yuchen Lin, Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Xiang Ren and William Cohen	4611
<i>Does Structure Matter? Encoding Documents for Machine Reading Comprehension</i>	
Hui Wan, Song Feng, Chulaka Gunasekara, Siva Sankalp Patel, Sachindra Joshi and Luis Lastras	4626
<i>Multi-Step Reasoning Over Unstructured Text with Beam Dense Retrieval</i>	
Chen Zhao, Chenyan Xiong, Jordan Boyd-Graber and Hal Daumé III	4635
<i>Scalable and Interpretable Semantic Change Detection</i>	
Syrielle Montariol, Matej Martinc and Lidia Pivovarova	4642
<i>Scalar Adjective Identification and Multilingual Ranking</i>	
Aina Garí Soler and Marianna Apidianaki	4653
<i>ESC: Redesigning WSD with Extractive Sense Comprehension</i>	
Edoardo Barba, Tommaso Pasini and Roberto Navigli	4661
<i>Recent advances in neural metaphor processing: A linguistic, cognitive and social perspective</i>	
Xiaoyu Tong, Ekaterina Shutova and Martha Lewis	4673

<i>Constructing Taxonomies from Pretrained Language Models</i> Catherine Chen, Kevin Lin and Dan Klein	4687
<i>Event Representation with Sequential, Semi-Supervised Discrete Variables</i> Mehdi Rezaee and Francis Ferraro	4701
<i>Seq2Emo: A Sequence to Multi-Label Emotion Classification Model</i> Chenyang Huang, Amine Trabelsi, Xuebin Qin, Nawshad Farruque, Lili Mou and Osmar Zaiane 4717	
<i>Knowledge Enhanced Masked Language Model for Stance Detection</i> Kornraphop Kawintiranon and Lisa Singh	4725
<i>Learning Paralinguistic Features from Audiobooks through Style Voice Conversion</i> Zakaria Aldeneh, Matthew Perez and Emily Mower Provost	4736
<i>Adapting BERT for Continual Learning of a Sequence of Aspect Sentiment Classification Tasks</i> Zixuan Ke, Hu Xu and Bing Liu	4746
<i>Adversarial Learning for Zero-Shot Stance Detection on Social Media</i> Emily Allaway, Malavika Srikanth and Kathleen McKeown	4756
<i>Efficiently Summarizing Text and Graph Encodings of Multi-Document Clusters</i> Ramakanth Pasunuru, Mengwen Liu, Mohit Bansal, Sujith Ravi and Markus Dreyer	4768
<i>Enriching Transformers with Structured Tensor-Product Representations for Abstractive Summarization</i> Yichen Jiang, Asli Celikyilmaz, Paul Smolensky, Paul Soulos, Sudha Rao, Hamid Palangi, Roland Fernandez, Caitlin Smith, Mohit Bansal and Jianfeng Gao	4780
<i>What's in a Summary? Laying the Groundwork for Advances in Hospital-Course Summarization</i> Griffin Adams, Emily Alsentzer, Mert Ketenci, Jason Zucker and Noémie Elhadad	4794
<i>Understanding Factuality in Abstractive Summarization with FRANK: A Benchmark for Factuality Metrics</i> Artidoro Pagnoni, Vidhisha Balachandran and Yulia Tsvetkov	4812
<i>GSum: A General Framework for Guided Neural Abstractive Summarization</i> Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang and Graham Neubig	4830
<i>What Will it Take to Fix Benchmarking in Natural Language Understanding?</i> Samuel R. Bowman and George Dahl	4843
<i>TuringAdvice: A Generative and Dynamic Evaluation of Language Use</i> Rowan Zellers, Ari Holtzman, Elizabeth Clark, Lianhui Qin, Ali Farhadi and Yejin Choi	4856
<i>Multitask Learning for Emotionally Analyzing Sexual Abuse Disclosures</i> Ramit Sawhney, Puneet Mathur, Taru Jain, Akash Kumar Gautam and Rajiv Ratn Shah	4881
<i>Self Promotion in US Congressional Tweets</i> Jun Wang, Kelly Cui and Bei Yu	4893
<i>Profiling of Intertextuality in Latin Literature Using Word Embeddings</i> Patrick J. Burns, James Brofos, Kyle Li, Pramit Chaudhuri and Joseph P. Dexter	4900
<i>Identifying inherent disagreement in natural language inference</i> Xinliang Frederick Zhang and Marie-Catherine de Marneffe	4908

<i>Modeling Human Mental States with an Entity-based Narrative Graph</i>	
I-Ta Lee, Maria Leonor Pacheco and Dan Goldwasser	4916
<i>A Simple and Efficient Multi-Task Learning Approach for Conditioned Dialogue Generation</i>	
Yan Zeng and Jian-Yun Nie	4927
<i>Hurdles to Progress in Long-form Question Answering</i>	
Kalpesh Krishna, Aurko Roy and Mohit Iyyer	4940
<i>ENTRUST: Argument Reframing with Language Models and Entailment</i>	
Tuhin Chakrabarty, Christopher Hidey and Smaranda Muresan	4958
<i>Paragraph-level Simplification of Medical Texts</i>	
Ashwin Devaraj, Iain Marshall, Byron Wallace and Junyi Jessy Li	4972
<i>An Empirical Study on Neural Keyphrase Generation</i>	
Rui Meng, Xingdi Yuan, Tong Wang, Sanqiang Zhao, Adam Trischler and Daqing He	4985
<i>Attention Head Masking for Inference Time Content Selection in Abstractive Summarization</i>	
Shuyang Cao and Lu Wang	5008
<i>Factual Probing Is [MASK]: Learning vs. Learning to Recall</i>	
Zexuan Zhong, Dan Friedman and Danqi Chen	5017
<i>Evaluating Saliency Methods for Neural Language Models</i>	
Shuoyang Ding and Philipp Koehn	5034
<i>Contextualized Perturbation for Textual Adversarial Attack</i>	
Dianqi Li, Yizhe Zhang, Hao Peng, Liqun Chen, Chris Brockett, Ming-Ting Sun and Bill Dolan	5053
<i>DirectProbe: Studying Representations without Classifiers</i>	
Yichu Zhou and Vivek Srikumar	5070
<i>Evaluating the Values of Sources in Transfer Learning</i>	
Md Rizwan Parvez and Kai-Wei Chang	5084
<i>Too Much in Common: Shifting of Embeddings in Transformer Language Models and its Implications</i>	
Daniel Biś, Maksim Podkorytov and Xiuwen Liu	5117
<i>On the Inductive Bias of Masked Language Modeling: From Statistical to Syntactic Dependencies</i>	
Tianyi Zhang and Tatsunori Hashimoto	5131
<i>Limitations of Autoregressive Models and Their Alternatives</i>	
Chu-Cheng Lin, Aaron Jaech, Xin Li, Matthew R. Gormley and Jason Eisner	5147
<i>On the Transformer Growth for Progressive BERT Training</i>	
Xiaotao Gu, Liyuan Liu, Hongkun Yu, Jing Li, Chen Chen and Jiawei Han	5174
<i>Revisiting Simple Neural Probabilistic Language Models</i>	
Simeng Sun and Mohit Iyyer	5181
<i>ReadTwice: Reading Very Large Documents with Memories</i>	
Yury Zemlyanskiy, Joshua Ainslie, Michiel de Jong, Philip Pham, Ilya Eckstein and Fei Sha .	5189

<i>SCRIPT: Self-Critic PreTraining of Transformers</i> Erik Nijkamp, Bo Pang, Ying Nian Wu and Caiming Xiong	5196
<i>Learning How to Ask: Querying LMs with Mixtures of Soft Prompts</i> Guanghui Qin and Jason Eisner	5203
<i>Nutri-bullets Hybrid: Consensual Multi-document Summarization</i> Darsh Shah, Lili Yu, Tao Lei and Regina Barzilay	5213
<i>AVA: an Automatic eValuation Approach for Question Answering Systems</i> Thuy Vu and Alessandro Moschitti	5223
<i>SpanPredict: Extraction of Predictive Document Spans with Neural Attention</i> Vivek Subramanian, Matthew Engelhard, Sam Berchuck, Liqun Chen, Ricardo Henao and Lawrence Carin	5234
<i>Text Editing by Command</i> Felix Faltings, Michel Galley, Gerold Hintz, Chris Brockett, Chris Quirk, Jianfeng Gao and Bill Dolan	5259
<i>A Deep Metric Learning Approach to Account Linking</i> Aleem Khan, Elizabeth Fleming, Noah Schofield, Marcus Bishop and Nicholas Andrews	5275
<i>Improving Factual Completeness and Consistency of Image-to-Text Radiology Report Generation</i> Yasuhide Miura, Yuhao Zhang, Emily Tsai, Curtis Langlotz and Dan Jurafsky	5288
<i>Multimodal End-to-End Sparse Model for Emotion Recognition</i> Wenliang Dai, Samuel Cahyawijaya, Zihan Liu and Pascale Fung	5305
<i>MIMOQA: Multimodal Input Multimodal Output Question Answering</i> Hrituraj Singh, Anshul Nasery, Denil Mehta, Aishwarya Agarwal, Jatin Lamba and Balaji Vasani Srinivasan	5317
<i>OCID-Ref: A 3D Robotic Dataset With Embodied Language For Clutter Scene Grounding</i> Ke-Jyun Wang, Yun-Hsuan Liu, Hung-Ting Su, Jen-Wei Wang, Yu-Siang Wang, Winston Hsu and Wen-Chin Chen	5333
<i>Unsupervised Vision-and-Language Pre-training Without Parallel Images and Captions</i> Liunian Harold Li, Haoxuan You, Zhecan Wang, Alireza Zareian, Shih-Fu Chang and Kai-Wei Chang	5339
<i>Multitasking Inhibits Semantic Drift</i> Athul Paul Jacob, Mike Lewis and Jacob Andreas	5351
<i>Probing Contextual Language Models for Common Ground with Visual Representations</i> Gabriel Ilharco, Rowan Zellers, Ali Farhadi and Hannaneh Hajishirzi	5367
<i>BBAEG: Towards BERT-based Biomedical Adversarial Example Generation for Text Classification</i> Ishani Mondal	5378
<i>Targeted Adversarial Training for Natural Language Understanding</i> Lis Pereira, Xiaodong Liu, Hao Cheng, Hoifung Poon, Jianfeng Gao and Ichiro Kobayashi ..	5385
<i>Latent-Optimized Adversarial Neural Transfer for Sarcasm Detection</i> Xu Guo, Boyang Li, Han Yu and Chunyan Miao	5394

<i>Self-training Improves Pre-training for Natural Language Understanding</i>	
Jingfei Du, Edouard Grave, Beliz Gunel, Vishrav Chaudhary, Onur Celebi, Michael Auli, Veselin Stoyanov and Alexis Conneau	5408
<i>Supporting Clustering with Contrastive Learning</i>	
Dejiao Zhang, Feng Nan, Xiaokai Wei, Shang-Wen Li, Henghui Zhu, Kathleen McKeown, Ramesh Nallapati, Andrew O. Arnold and Bing Xiang	5419
<i>TITA: A Two-stage Interaction and Topic-Aware Text Matching Model</i>	
Xingwu Sun, Yanling Cui, Hongyin Tang, Qiuyu Zhu, Fuzheng Zhang and Beihong Jin	5431
<i>Neural Quality Estimation with Multiple Hypotheses for Grammatical Error Correction</i>	
Zhenghao Liu, Xiaoyuan Yi, Maosong Sun, Liner Yang and Tat-Seng Chua	5441
<i>Neural Network Surgery: Injecting Data Patterns into Pre-trained Models with Minimal Instance-wise Side Effects</i>	
Zhiyuan Zhang, Xuancheng Ren, Qi Su, Xu Sun and Bin He	5453
<i>Discrete Argument Representation Learning for Interactive Argument Pair Identification</i>	
Lu Ji, Zhongyu Wei, Jing Li, Qi Zhang and Xuanjing Huang	5467
<i>On Unifying Misinformation Detection</i>	
Nayeon Lee, Belinda Z. Li, Sinong Wang, Pascale Fung, Hao Ma, Wen-tau Yih and Madian Khabsa	5479
<i>Frustratingly Easy Edit-based Linguistic Steganography with a Masked Language Model</i>	
Honai Ueoka, Yugo Murawaki and Sadao Kurohashi	5486
<i>Few-Shot Text Classification with Triplet Networks, Data Augmentation, and Curriculum Learning</i>	
Jason Wei, Chengyu Huang, Soroush Vosoughi, Yu Cheng and Shiqi Xu	5493
<i>Do RNN States Encode Abstract Phonological Alternations?</i>	
Miikka Silfverberg, Francis Tyers, Garrett Nicolai and Mans Hulden	5501
<i>Pre-training with Meta Learning for Chinese Word Segmentation</i>	
Zhen Ke, Liang Shi, Songtao Sun, Erli Meng, Bin Wang and Xipeng Qiu	5514
<i>Decompose, Fuse and Generate: A Formation-Informed Method for Chinese Definition Generation</i>	
Hua Zheng, Damai Dai, Lei Li, Tianyu Liu, Zhifang Sui, Baobao Chang and Yang Liu	5524
<i>User-Generated Text Corpus for Evaluating Japanese Morphological Analysis and Lexical Normalization</i>	
Shohei Higashiyama, Masao Utiyama, Taro Watanabe and Eiichiro Sumita	5532
<i>GPT Perdetry Test: Generating new meanings for new words</i>	
Nikolay Malkin, Sameera Lanka, Pranav Goel, Sudha Rao and Nebojsa Jojic	5542
<i>Universal Semantic Tagging for English and Mandarin Chinese</i>	
Wenxi Li, Yiyang Hou, Yajie Ye, Li Liang and Weiwei Sun	5554
<i>ShadowGNN: Graph Projection Neural Network for Text-to-SQL Parser</i>	
Zhi Chen, Lu Chen, Yanbin Zhao, Ruisheng Cao, Zihan Xu, Su Zhu and Kai Yu	5567
<i>Contextualized and Generalized Sentence Representations by Contrastive Self-Supervised Learning: A Case Study on Discourse Relation Analysis</i>	
Hirokazu Kiyomaru and Sadao Kurohashi	5578

<i>AMR Parsing with Action-Pointer Transformer</i>	
Jiawei Zhou, Tahira Naseem, Ramón Fernandez Astudillo and Radu Florian	5585
<i>NL-EDIT: Correcting Semantic Parse Errors through Natural Language Interaction</i>	
Ahmed Elgohary, Christopher Meek, Matthew Richardson, Adam Fourney, Gonzalo Ramos and Ahmed Hassan Awadallah	5599
<i>Unsupervised Concept Representation Learning for Length-Varying Text Similarity</i>	
Xuchao Zhang, Bo Zong, Wei Cheng, Jingchao Ni, Yanchi Liu and Haifeng Chen	5611
<i>Augmenting Knowledge-grounded Conversations with Sequential Knowledge Transition</i>	
Haolan Zhan, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Yongjun Bao and Yanyan Lan . .	5621
<i>Adversarial Self-Supervised Learning for Out-of-Domain Detection</i>	
Zhiyuan Zeng, Keqing He, Yuanmeng Yan, Hong Xu and Weiran Xu	5631
<i>Leveraging Slot Descriptions for Zero-Shot Cross-Domain Dialogue StateTracking</i>	
Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho and Rajen Subba	5640
<i>Hierarchical Transformer for Task Oriented Dialog Systems</i>	
Bishal Santra, Potnuru Anusha and Pawan Goyal	5649
<i>Measuring the ‘I don’t know’ Problem through the Lens of Gricean Quantity</i>	
Huda Khayrallah and João Sedoc	5659
<i>RTFE: A Recursive Temporal Fact Embedding Framework for Temporal Knowledge Graph Completion</i>	
Youri Xu, Haihong E, Meina Song, wenyu song, Xiaodong Lv, wang haotian and yang jinrui .	5671
<i>Open Hierarchical Relation Extraction</i>	
Kai Zhang, Yuan Yao, Ruobing Xie, Xu Han, Zhiyuan Liu, Fen Lin, Leyu Lin and Maosong Sun	5682
<i>Jointly Extracting Explicit and Implicit Relational Triples with Reasoning Pattern Enhanced Binary Pointer Network</i>	
Yubo Chen, Yunqi Zhang, Changran Hu and Yongfeng Huang	5694
<i>Multi-Grained Knowledge Distillation for Named Entity Recognition</i>	
Xuan Zhou, Xiao Zhang, Chenyang Tao, Junya Chen, Bing Xu, Wei Wang and Jing Xiao	5704
<i>SGG: Learning to Select, Guide, and Generate for Keyphrase Generation</i>	
Jing Zhao, Junwei Bao, Yifan Wang, Youzheng Wu, Xiaodong He and Bowen Zhou	5717
<i>Towards Sentiment and Emotion aided Multi-modal Speech Act Classification in Twitter</i>	
Tulika Saha, Apoorva Upadhyaya, Sriparna Saha and Pushpak Bhattacharyya	5727
<i>Generative Imagination Elevates Machine Translation</i>	
Quanyu Long, Mingxuan Wang and Lei Li	5738
<i>Non-Autoregressive Translation by Learning Target Categorical Codes</i>	
Yu Bao, Shujian Huang, Tong Xiao, Dongqi Wang, Xinyu Dai and Jiajun CHEN	5749
<i>Training Data Augmentation for Code-Mixed Translation</i>	
Abhirut Gupta, Aditya Vavre and Sunita Sarawagi	5760

<i>Rethinking Perturbations in Encoder-Decoders for Fast Training</i> Sho Takase and Shun Kiyono	5767
<i>Context-aware Decoder for Neural Machine Translation using a Target-side Document-Level Language Model</i> Amane Sugiyama and Naoki Yoshinaga	5781
<i>Machine Translated Text Detection Through Text Similarity with Round-Trip Translation</i> Hoang-Quoc Nguyen-Son, Tran Thao, Seira Hidano, Ishita Gupta and Shinsaku Kiyomoto ..	5792
<i>TR-BERT: Dynamic Token Reduction for Accelerating BERT Inference</i> Deming Ye, Yankai Lin, Yufei Huang and Maosong Sun	5798
<i>Breadth First Reasoning Graph for Multi-hop Question Answering</i> Yongjie Huang and Meng Yang	5810
<i>Improving Zero-Shot Cross-lingual Transfer for Multilingual Question Answering over Knowledge Graph</i> Yucheng Zhou, Xiubo Geng, Tao Shen, Wenqiang Zhang and Daxin Jiang	5822
<i>RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering</i> Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu and Haifeng Wang	5835
<i>DAGN: Discourse-Aware Graph Network for Logical Reasoning</i> Yinya Huang, Meng Fang, Yu Cao, Liwei Wang and Xiaodan Liang	5848
<i>Designing a Minimal Retrieve-and-Read System for Open-Domain Question Answering</i> Sohee Yang and Minjoon Seo	5856
<i>Unsupervised Multi-hop Question Answering by Question Generation</i> Liangming Pan, Wenhui Chen, Wenhan Xiong, Min-Yen Kan and William Yang Wang	5866
<i>Sliding Selector Network with Dynamic Memory for Extractive Summarization of Long Documents</i> Peng Cui and Le Hu	5881
<i>AdaptSum: Towards Low-Resource Domain Adaptation for Abstractive Summarization</i> Tiezheng Yu, Zihan Liu and Pascale Fung	5892
<i>QMSum: A New Benchmark for Query-based Multi-domain Meeting Summarization</i> Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu and Dragomir Radev	5905
<i>MM-AVS: A Full-Scale Dataset for Multi-modal Summarization</i> Xiyang Fu, Jun Wang and Zhenglu Yang	5922
<i>MediaSum: A Large-scale Media Interview Dataset for Dialogue Summarization</i> Chenguang Zhu, Yang Liu, Jie Mei and Michael Zeng	5927
<i>Improving Faithfulness in Abstractive Summarization with Contrast Candidate Generation and Selection</i> Sihao Chen, Fan Zhang, Kazuo Sone and Dan Roth	5935
<i>Inference Time Style Control for Summarization</i> Shuyang Cao and Lu Wang	5942