

# **Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT 2022)**

Online  
10-15 July 2022

Volume 1 of 8

ISBN: 978-1-7138-5618-4

**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. (2022)

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](mailto: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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

## Table of Contents

<i>Social Norms Guide Reference Resolution</i> Mitchell Abrams and Matthias Scheutz .....	1
<i>Learning Natural Language Generation with Truncated Reinforcement Learning</i> Alice Martin, Guillaume Quispe, Charles Ollion, Sylvain Le Corff, Florian Strub and Olivier Pietquin .....	12
<i>Language Model Augmented Monotonic Attention for Simultaneous Translation</i> Sathish Reddy Indurthi, Mohd Abbas Zaidi, Beomseok Lee, Nikhil Kumar Lakumarapu and Sangha Kim .....	38
<i>What Makes a Good and Useful Summary? Incorporating Users in Automatic Summarization Research</i> Maartje Ter Hoeve, Julia Kiseleva and Maarten de Rijke .....	46
<i>ErAConD: Error Annotated Conversational Dialog Dataset for Grammatical Error Correction</i> Xun Yuan, Derek Pham, Sam Davidson and Zhou Yu .....	76
<i>Semantic Diversity in Dialogue with Natural Language Inference</i> Katherine Stasaski and Marti Hearst .....	85
<i>LEA: Meta Knowledge-Driven Self-Attentive Document Embedding for Few-Shot Text Classification</i> S. K. Hong and Tae Young Jang .....	99
<i>Enhancing Self-Attention with Knowledge-Assisted Attention Maps</i> Jiangang Bai, Yujing Wang, Hong Sun, Ruonan Wu, Tianmeng Yang, Pengfei Tang, Defu Cao, Mingliang Zhang1, Yunhai Tong, Yaming Yang, Jing Bai, Ruofei Zhang, Hao Sun and Wei Shen ..	107
<i>Batch-Softmax Contrastive Loss for Pairwise Sentence Scoring Tasks</i> Anton Chernyavskiy, Dmitry Ilvovsky, Pavel Kalinin and Preslav Nakov .....	116
<i>NewsEdits: A News Article Revision Dataset and a Novel Document-Level Reasoning Challenge</i> Alexander Spangher, Xiang Ren, Jonathan May and Nanyun Peng .....	127
<i>Putting the Con in Context: Identifying Deceptive Actors in the Game of Mafia</i> Samee Omotayo Ibraheem, Gaoyue Zhou and John DeNero .....	158
<i>SUBS: Subtree Substitution for Compositional Semantic Parsing</i> Jingfeng Yang, Le Zhang and Diyi Yang .....	169
<i>Two Contrasting Data Annotation Paradigms for Subjective NLP Tasks</i> Paul Rottger, Bertie Vidgen, Dirk Hovy and Janet B. Pierrehumbert .....	N/A
<i>Do Deep Neural Nets Display Human-like Attention in Short Answer Scoring?</i> Zijie Zeng, Xinyu Li, Dragan Gasevic and Guanliang Chen .....	191
<i>Knowledge-Grounded Dialogue Generation with a Unified Knowledge Representation</i> Yu Li, Baolin Peng, Yelong Shen, Yi Mao, Lars Liden, Zhou Yu and Jianfeng Gao .....	206
<i>CERES: Pretraining of Graph-Conditioned Transformer for Semi-Structured Session Data</i> Rui Feng, Chen Luo, Qingyu Yin, Bing Yin, Tuo Zhao and Chao Zhang .....	219
<i>Political Ideology and Polarization: A Multi-dimensional Approach</i> Barea Sinno, Bernardo Oviedo, Katherine Atwell, Malihe Alikhani and Junyi Jessy Li .....	231

<i>Cooperative Self-training of Machine Reading Comprehension</i>	
Hongyin Luo, Shang-Wen Li, Mingye Gao, Seunghak Yu and James R. Glass . . . . .	244
<i>GlobEnc: Quantifying Global Token Attribution by Incorporating the Whole Encoder Layer in Transformers</i>	
Ali Modarressi, Mohsen Fayyaz, Yadollah Yaghoobzadeh and Mohammad Taher Pilehvar . . .	258
<i>A Robustly Optimized BMRC for Aspect Sentiment Triplet Extraction</i>	
Shu Liu, Kaiwen Li and Zuhe Li . . . . .	272
<i>Seed-Guided Topic Discovery with Out-of-Vocabulary Seeds</i>	
Yu Zhang, Yu Meng, Xuan Wang, Sheng Wang and Jiawei Han . . . . .	279
<i>Towards Process-Oriented, Modular, and Versatile Question Generation that Meets Educational Needs</i>	
Xu Wang, Simin Fan, Jessica Houghton and Lu Wang . . . . .	291
<i>SwahBERT: Language Model of Swahili</i>	
Gati L Martin, Medard Edmund Mswahili, Young-Seob Jeong and Jeong Young-Seob . . . . .	303
<i>Deconstructing NLG Evaluation: Evaluation Practices, Assumptions, and Their Implications</i>	
Kaitlyn Zhou, Su Lin Blodgett, Adam Trischler, Hal Daumé III, Kaheer Suleman and Alexandra Olteanu . . . . .	N/A
<i>TSTR: Too Short to Represent, Summarize with Details! Intro-Guided Extended Summary Generation</i>	
Sajad Sotudeh and Nazli Goharian . . . . .	325
<i>Empathic Machines: Using Intermediate Features as Levers to Emulate Emotions in Text-To-Speech Systems</i>	
Saiteja Kosgi, Sarath Sivaprasad, Niranjan Pedanekar, Anil Kumar Nelakanti and Vineet Gandhi	336
<i>The Why and The How: A Survey on Natural Language Interaction in Visualization</i>	
Henrik Voigt, Ozge Alacam, Monique Meuschke, Kai Lawonn and Sina Zarriß . . . . .	348
<i>Understand before Answer: Improve Temporal Reading Comprehension via Precise Question Understanding</i>	
Hao Huang, Xiubo Geng, Guodong Long and Daxin Jiang . . . . .	375
<i>User-Driven Research of Medical Note Generation Software</i>	
Tom Knoll, Francesco Moramarco, Alex Papadopoulos Korfiatis, Rachel Young, Claudia Ruffini, Mark Perera, Christian Perstl, Ehud Reiter, Anya Belz and Aleksandar Savkov . . . . .	385
<i>Ask Me Anything in Your Native Language</i>	
Nikita Sorokin, Dmitry Abulkhanov, Irina Piontkovskaya and Valentin Malykh . . . . .	395
<i>Diversifying Neural Dialogue Generation via Negative Distillation</i>	
Yiwei Li, Shaoxiong Feng, Bin Sun and Kan Li . . . . .	407
<i>On Synthetic Data for Back Translation</i>	
Jiahao Xu, Yubin Ruan, Wei Bi, Guoping Huang, Shuming Shi, Lihui Chen and Lemao Liu .	419
<i>Mapping the Design Space of Human-AI Interaction in Text Summarization</i>	
Ruijia Cheng, Alison Smith-Renner, Ke Zhang, Joel R. Tetreault and Alejandro Jaimes-Larrarte	431
<i>Towards Robust and Semantically Organised Latent Representations for Unsupervised Text Style Transfer</i>	
Sharan Narasimhan, Suvodip Dey and Maunendra Sankar Desarkar . . . . .	456

<i>An Exploration of Post-Editing Effectiveness in Text Summarization</i>	
Vivian Lai, Alison Smith-Renner, Ke Zhang, Ruijia Cheng, Wenjuan Zhang, Joel R. Tetreault and Alejandro Jaimes-Larrarte .....	475
<i>Automatic Correction of Human Translations</i>	
Jessy Lin, Geza Kovacs, Aditya Shastry, Joern Wuebker and John DeNero .....	494
<i>On the Robustness of Reading Comprehension Models to Entity Renaming</i>	
Jun Yan, Yang Xiao, Sagnik Mukherjee, Bill Yuchen Lin, Robin Jia and Xiang Ren .....	508
<i>Explaining Why: How Instructions and User Interfaces Impact Annotator Rationales When Labeling Text Data</i>	
Cynthia Sullivan, William Brackenbury, Andrew Michael McNut, Kevin Bryson, kbyllofficial@gmail.com, kbyllofficial@gmail.com, Yuxin Chen, Michael Littman, Chenhao Tan and Blase Ur .....	N/A
<i>Fine-tuning Pre-trained Language Models for Few-shot Intent Detection: Supervised Pre-training and Isotropization</i>	
Haode Zhang, Haowen Liang, Yuwei Zhang, Li-Ming Zhan, Xiao-Ming Wu, Xiaolei Lu and Albert Y.s. Lam .....	532
<i>Cross-document Misinformation Detection based on Event Graph Reasoning</i>	
Xueqing Wu, Kung-Hsiang Huang, Yi Fung and Heng Ji .....	543
<i>Disentangled Action Recognition with Knowledge Bases</i>	
Zhekun Luo, Shalini Ghosh, Devin Guillory, Keizo Kato, Trevor Darrell and Huijuan Xu ...	559
<i>Machine-in-the-Loop Rewriting for Creative Image Captioning</i>	
Vishakh Padmakumar and He He .....	573
<i>A Word is Worth A Thousand Dollars: Adversarial Attack on Tweets Fools Stock Prediction</i>	
Yong Xie, Dakuo Wang, Pin-Yu Chen, Jinjun Xiong, Sijia Liu and Oluwasanmi O Koyejo ...	587
<i>Building Multilingual Machine Translation Systems That Serve Arbitrary XY Translations</i>	
Akiko Eriguchi, Shufang Xie, Tao Qin and Hany Hassan .....	600
<i>Non-Autoregressive Neural Machine Translation with Consistency Regularization Optimized Variational Framework</i>	
Minghao Zhu, Junli Wang and Chungang Yan .....	607
<i>User-Centric Gender Rewriting</i>	
Bashar Alhafni, Nizar Habash and Houda Bouamor .....	618
<i>Reframing Human-AI Collaboration for Generating Free-Text Explanations</i>	
Sarah Wiegrefe, Jack Hessel, Swabha Swayamdipta, Mark Riedl and Yejin Choi .....	632
<i>EmRel: Joint Representation of Entities and Embedded Relations for Multi-triple Extraction</i>	
Benfeng Xu, Quan Wang, Yajuan Lyu, Yabing Shi, Yong Zhu, Jie Gao and Zhendong Mao ..	659
<i>Meta Learning for Natural Language Processing: A Survey</i>	
Hung-yi Lee, Shang-Wen Li and Thang Vu .....	666
<i>Analyzing Modality Robustness in Multimodal Sentiment Analysis</i>	
Devamanyu Hazarika, Yingting Li, Bo Cheng, Shuai Zhao, Roger Zimmermann and Soujanya Poria .....	685
<i>Fuse It More Deeply! A Variational Transformer with Layer-Wise Latent Variable Inference for Text Generation</i>	
Jinyi Hu, Xiaoyuan Yi, Wenhao Li, Maosong Sun and Xing Xie .....	697

<i>Easy Adaptation to Mitigate Gender Bias in Multilingual Text Classification</i> Xiaolei Huang .....	717
<i>On the Use of External Data for Spoken Named Entity Recognition</i> Ankita Pasad, Felix Wu, Suwon Shon, Karen Livescu and Kyu Han .....	724
<i>Long-term Control for Dialogue Generation: Methods and Evaluation</i> Ramya Ramakrishnan, Hashan Buddhika Narangodage, Mauro Schilman, Kilian Q Weinberger and Ryan McDonald .....	738
<i>Learning Dialogue Representations from Consecutive Utterances</i> Zhihan Zhou, Dejiao Zhang, Wei Xiao, Nicholas Dingwall, Xiaofei Ma, Andrew Arnold and Bing Xiang .....	754
<i>On the Machine Learning of Ethical Judgments from Natural Language</i> Zeerak Talat, Hagen Blix, Josef Valvoda, Maya Indira Ganesh, Ryan Cotterell and Adina Williams 769	
<i>NeuroLogic A*esque Decoding: Constrained Text Generation with Lookahead Heuristics</i> Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, Jungo Kasai, Daniel Khashabi, Ronan Le Bras, Lianhui Qin, Youngjae Yu, Rowan Zellers, Noah Smith and Yejin Choi .....	780
<i>PARADISE: Exploiting Parallel Data for Multilingual Sequence-to-Sequence Pretraining</i> Machel Reid and Mikel Artetxe .....	800
<i>Explaining Toxic Text via Knowledge Enhanced Text Generation</i> Rohit Sridhar and Diyi Yang .....	811
<i>Teaching BERT to Wait: Balancing Accuracy and Latency for Streaming Disfluency Detection</i> Angelica Chen, Vicky Zayats, Daniel David Walker and Dirk Padfield .....	827
<i>GRAM: Fast Fine-tuning of Pre-trained Language Models for Content-based Collaborative Filtering</i> Yoonseok Yang, Kyu Seok Kim, Minsam Kim and Juneyoung Park .....	839
<i>Generating Repetitions with Appropriate Repeated Words</i> Toshiki Kawamoto, Hidetaka Kamigaito, Kotaro Funakoshi and Manabu Okumura .....	852
<i>Textless Speech-to-Speech Translation on Real Data</i> Ann Lee, Hongyu Gong, Paul-Ambroise Duquenne, Holger Schwenk, Peng-Jen Chen, Changhan Wang, Sravya Popuri, Yossi Adi, Juan Pino, Jiatao Gu and Wei-Ning Hsu .....	860
<i>WALNUT: A Benchmark on Semi-weakly Supervised Learning for Natural Language Understanding</i> Guoqing Zheng, Giannis Karamanolakis, Kai Shu and Ahmed Hassan Awadallah .....	873
<i>CompactIE: Compact Facts in Open Information Extraction</i> Farima Fatahi Bayat, Nikita Bhutani and H. Jagadish .....	900
<i>CoSim: Commonsense Reasoning for Counterfactual Scene Imagination</i> Hyoungun Kim, Abhay Zala and Mohit Bansal .....	911
<i>Abstraction not Memory: BERT and the English Article System</i> Harish Tayyar Madabushi, Dagmar Divjak and Petar Milin .....	924
<i>OmniTab: Pretraining with Natural and Synthetic Data for Few-shot Table-based Question Answering</i> Zhengbao Jiang, Yi Mao, Pengcheng He, Graham Neubig and Weizhu Chen .....	932
<i>Provably Confidential Language Modelling</i> Xuandong Zhao, Lei Li and Yu-Xiang Wang .....	943

<i>KAT: A Knowledge Augmented Transformer for Vision-and-Language</i> Liangke Gui, Borui Wang, Qiuyuan Huang, Alexander G Hauptmann, Yonatan Bisk and Jianfeng Gao .....	956
<i>When a sentence does not introduce a discourse entity, Transformer-based models still sometimes refer to it</i> Sebastian Schuster and Tal Linzen .....	969
<i>On Curriculum Learning for Commonsense Reasoning</i> Adyasha Maharana and Mohit Bansal .....	983
<i>DocTime: A Document-level Temporal Dependency Graph Parser</i> Puneet Mathur, Vlad I Morariu, Verena Kaynig-Fittkau, Jiuxiang Gu, Franck Dernoncourt, Quan Hung Tran, Ani Nenkova, Dinesh Manocha and Rajiv Jain .....	993
<i>FactPEGASUS: Factuality-Aware Pre-training and Fine-tuning for Abstractive Summarization</i> David Wan and Mohit Bansal .....	1010
<i>ScAN: Suicide Attempt and Ideation Events Dataset</i> Bhanu Pratap Singh Rawat, Samuel Kovaly, Hong Yu and Wilfred Pigeon .....	1029
<i>Socially Aware Bias Measurements for Hindi Language Representations</i> Vijit Malik, Sunipa Dev, Akihiro Nishi, Nanyun Peng and Kai-Wei Chang .....	1041
<i>AmbiPun: Generating Humorous Puns with Ambiguous Context</i> Anirudh Mittal, Yufei Tian and Nanyun Peng .....	1053
<i>EmpHi: Generating Empathetic Responses with Human-like Intents</i> Mao Yan Chen, Siheng Li and Yujiu Yang .....	1063
<i>Yes, No or IDK: The Challenge of Unanswerable Yes/No Questions</i> Elior Sulem, Jamaal Hay and Dan Roth .....	1075
<i>Inducing and Using Alignments for Transition-based AMR Parsing</i> Andrew Drozdov, Jiawei Zhou, Radu Florian, Andrew McCallum, Tahira Naseem, Yoon Kim and Ramón Fernandez Astudillo .....	1086
<i>Masked Part-Of-Speech Model: Does Modeling Long Context Help Unsupervised POS-tagging?</i> Xiang Zhou, Shiyue Zhang and Mohit Bansal .....	1099
<i>DREAM: Improving Situational QA by First Elaborating the Situation</i> Yuling Gu, Bhavana Dalvi and Peter Clark .....	1115
<i>CoSe-Co: Text Conditioned Generative CommonSense Contextualizer</i> Rachit Bansal, Milan Aggarwal, Sumit Bhatia, Jivat Neet Kaur and Balaji Krishnamurthy ..	1128
<i>Probing via Prompting</i> Jiaoda Li, Ryan Cotterell and Mrinmaya Sachan .....	1144
<i>Database Search Results Disambiguation for Task-Oriented Dialog Systems</i> Kun Qian, Satwik Kottur, Ahmad Beirami, Shahin Shayandeh, Paul A. Crook, Alborz Geramifard, Zhou Yu and Chinnadhurai Sankar .....	1158
<i>Unsupervised Slot Schema Induction for Task-oriented Dialog</i> Dian Yu, Mingqiu Wang, Yuan Cao, Izhak Shafran, Laurent El Shafey and Hagen Soltau ...	1174

<i>Towards a Progression-Aware Autonomous Dialogue Agent</i>	
Abraham Sanders, Tomek Strzalkowski, Mei Si, Albert Chang, Deepanshu Dey, Jonas Braasch and Dakuo Wang .....	1194
<i>Cross-Domain Detection of GPT-2-Generated Technical Text</i>	
Juan Diego Rodriguez, Todd Hay, David Gros, Zain Shamsi and Ravi Srinivasan .....	1213
<i>DISAPERE: A Dataset for Discourse Structure in Peer Review Discussions</i>	
Neha Nayak Kennard, Tim O’Gorman, Rajarshi Das, Akshay Sharma, Chhandak Bagchi, Matthew Clinton, Pranay Kumar Yelugam, Hamed Zamani and Andrew McCallum .....	1234
<i>MultiSpanQA: A Dataset for Multi-Span Question Answering</i>	
Haonan Li, Martin Tomko, Maria Vasardani and Timothy Baldwin .....	1250
<i>Context-Aware Abbreviation Expansion Using Large Language Models</i>	
Shanqing Cai, Subhashini Venugopalan, Katrin Tomanek, Ajit Narayanan, Meredith Ringel Morris and Michael Brenner .....	1261
<i>Theory-Grounded Measurement of U.S. Social Stereotypes in English Language Models</i>	
Yang Trista Cao, Anna Sotnikova, Hal Daumé III, Rachel Rudinger and Linda Zou .....	1276
<i>Sort by Structure: Language Model Ranking as Dependency Probing</i>	
Max Müller-Eberstein, Rob Van Der Goot and Barbara Plank .....	1296
<i>Quantifying Synthesis and Fusion and their Impact on Machine Translation</i>	
Arturo Oncevay, Duygu Ataman, Niels Van Berkel, Barry Haddow, Alexandra Birch and Johannes Bjerva .....	1308
<i>Commonsense and Named Entity Aware Knowledge Grounded Dialogue Generation</i>	
Deeksha Varshney, Akshara Prabhakar and Asif Ekbal .....	1322
<i>Efficient Hierarchical Domain Adaptation for Pretrained Language Models</i>	
Alexandra Chronopoulou, Matthew E Peters and Jesse Dodge .....	1336
<i>Hatemoji: A Test Suite and Adversarially-Generated Dataset for Benchmarking and Detecting Emoji-Based Hate</i>	
Hannah Rose Kirk, Bertie Vidgen, Paul Rottger, Tristan Thrush and Scott A. Hale .....	1352
<i>On the Economics of Multilingual Few-shot Learning: Modeling the Cost-Performance Trade-offs of Machine Translated and Manual Data</i>	
Kabir Ahuja, Monojit Choudhury and Sandipan Dandapat .....	1369
<i>Learning to Selectively Learn for Weakly Supervised Paraphrase Generation with Model-based Reinforcement Learning</i>	
Haiyan Yin, Dingcheng Li and Ping Li .....	1385
<i>Quality-Aware Decoding for Neural Machine Translation</i>	
Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. De Souza, Perez Ogayo, Graham Neubig and Andre Martins .....	1396
<i>Pretrained Models for Multilingual Federated Learning</i>	
Orion Weller, Marc Marone, Vladimir Braverman, Dawn Lawrie and Benjamin Van Durme .....	1413
<i>AcTune: Uncertainty-Based Active Self-Training for Active Fine-Tuning of Pretrained Language Models</i>	
Yue Yu, Ling kai Kong, Jieyu Zhang, Rongzhi Zhang and Chao Zhang .....	1422



<i>Label Anchored Contrastive Learning for Language Understanding</i>	
Zhenyu Zhang, Yuming Zhao, Meng Chen and Xiaodong He .....	1437
<i>Go Back in Time: Generating Flashbacks in Stories with Event Temporal Prompts</i>	
Rujun Han, Hong Chen, Yufei Tian and Nanyun Peng .....	1450
<i>Forecasting COVID-19 Caseloads Using Unsupervised Embedding Clusters of Social Media Posts</i>	
Felix Drinkall, Stefan Zohren and Janet B. Pierrehumbert .....	1471
<i>Many Hands Make Light Work: Using Essay Traits to Automatically Score Essays</i>	
Rahul Kumar, Sandeep Mathias, Sriparna Saha and Pushpak Bhattacharyya .....	1485
<i>Natural Language Inference with Self-Attention for Veracity Assessment of Pandemic Claims</i>	
Miguel Arana-Catania, Elena Kochkina, Arkaitz Zubiaga, Maria Liakata, Robert Procter and Yulan He .....	1496
<i>Beyond Emotion: A Multi-Modal Dataset for Human Desire Understanding</i>	
Ao Jia, Yu He, Yazhou Zhang, Sagar Uprety, Dawei Song and Christina Lioma .....	1512
<i>Relation-Specific Attentions over Entity Mentions for Enhanced Document-Level Relation Extraction</i>	
Jiaxin Yu, Deqing Yang and Shuyu Tian .....	1523
<i>Twitter-COMMs: Detecting Climate, COVID, and Military Multimodal Misinformation</i>	
Giscard Biamby, Grace Luo, Trevor Darrell and Anna Rohrbach .....	1530
<i>BlonDe: An Automatic Evaluation Metric for Document-level Machine Translation</i>	
Yuchen Eleanor Jiang, Tianyu Liu, Shuming Ma, Dongdong Zhang, Jian Yang, Haoyang Huang, Rico Sennrich, Ryan Cotterell, Mrinmaya Sachan and Ming Zhou .....	1550
<i>Disentangled Learning of Stance and Aspect Topics for Vaccine Attitude Detection in Social Media</i>	
Lixing Zhu, Zheng Fang, Gabriele Pergola, Robert Procter and Yulan He .....	1566
<i>SKILL: Structured Knowledge Infusion for Large Language Models</i>	
Fedor Moiseev, Zhe Dong, Enrique Alfonseca and Martin Jaggi .....	1581
<i>Same Neurons, Different Languages: Probing Morphosyntax in Multilingual Pre-trained Models</i>	
Karolina Stanczak, Edoardo Ponti, Lucas Torroba Hennigen, Ryan Cotterell and Isabelle Augenstein .....	1589
<i>Aspect Is Not You Need: No-aspect Differential Sentiment Framework for Aspect-based Sentiment Analysis</i>	
Jiahao Cao, Rui Liu, Huailiang Peng, Lei Jiang and Xu Bai .....	1599
<i>MoEBERT: from BERT to Mixture-of-Experts via Importance-Guided Adaptation</i>	
Simiao Zuo, Qingru Zhang, Chen Liang, Pengcheng He, Tuo Zhao and Weizhu Chen .....	1610
<i>Implicit n-grams Induced by Recurrence</i>	
Xiaobing Sun and Wei Lu .....	1624
<i>Guiding Visual Question Generation</i>	
Nihir Vedd, Zixu Wang, Marek Rei, Yishu Miao and Lucia Specia .....	1640
<i>OPERA: Operation-Pivoted Discrete Reasoning over Text</i>	
Yongwei Zhou, Junwei Bao, Chaoqun Duan, Haipeng Sun, Jiahui Liang, Yifan Wang, Jing Zhao, Youzheng Wu, Xiaodong He and Tiejun Zhao .....	1655
<i>Improving Multi-Document Summarization through Referenced Flexible Extraction with Credit-Awareness</i>	
Yun-Zhu Song, Yi-Syuan Chen and Hong-Han Shuai .....	1667

<i>Improving Constituent Representation with Hypertree Neural Networks</i> Hao Zhou, Gongshen Liu and Kewei Tu .....	1682
<i>Measuring Fairness with Biased Rulers: A Comparative Study on Bias Metrics for Pre-trained Language Models</i> Pieter Delobelle, Ewoenam Kwaku Tokpo, Toon Calders and Bettina Berendt .....	1693
<i>MuCPAD: A Multi-Domain Chinese Predicate-Argument Dataset</i> Yahui Liu, Haoping Yang, Chen Gong, Qingrong Xia, Zhenghua Li and Min Zhang .....	1707
<i>Representation Learning for Conversational Data using Discourse Mutual Information Maximization</i> Bishal Santra, Sumegh Roychowdhury, Aishik Mandal, Vasu Gurram, Atharva Naik, Manish Gupta and Pawan Goyal .....	1718
<i>ValCAT: Variable-Length Contextualized Adversarial Transformations Using Encoder-Decoder Language Model</i> Chuyun Deng, Mingxuan Liu, Yue Qin, Jia Zhang, Hai-Xin Duan and Donghong Sun .....	1735
<i>A Study of Syntactic Multi-Modality in Non-Autoregressive Machine Translation</i> Kexun Zhang, Rui Wang, Xu Tan, Junliang Guo, Yi Ren, Tao Qin and Tie-Yan Liu .....	1747
<i>CIAug: Equipping Interpolative Augmentation with Curriculum Learning</i> Ramit Sawhney, Ritesh Singh Soun, Shrey Pandit, Megh Thakkar, Sarvagya Malaviya and Yuval Pinter .....	1758
<i>Proposition-Level Clustering for Multi-Document Summarization</i> Ori Ernst, Avi Caciularu, Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Jacob Goldberger and Ido Dagan .....	1765
<i>Non-Autoregressive Machine Translation: It's Not as Fast as it Seems</i> Jindřich Helcl, Barry Haddow and Alexandra Birch .....	1780
<i>BAD-X: Bilingual Adapters Improve Zero-Shot Cross-Lingual Transfer</i> Marinela Parović, Goran Glavaš, Ivan Vulić and Anna Korhonen .....	1791
<i>Combining Humor and Sarcasm for Improving Political Parody Detection</i> Xiao Ao, Danae Sanchez Villegas, Daniel Preotiu-Pietro and Nikolaos Aletras .....	1800
<i>TIE: Topological Information Enhanced Structural Reading Comprehension on Web Pages</i> Zihan Zhao, Lu Chen, Ruisheng Cao, Hongshen Xu, Xingyu Chen and Kai Yu .....	1808
<i>RSTGen: Imbuing Fine-Grained Interpretable Control into Long-Form Text Generators</i> Rilwan Akanni Adewoyin, Ritabrata Dutta and Yulan He .....	1822
<i>Intent Detection and Discovery from User Logs via Deep Semi-Supervised Contrastive Clustering</i> Rajat Kumar, Mayur Patidar, Vaibhav Varshney, Lovekesh Vig and Gautam Shroff .....	1836
<i>Extending Multi-Text Sentence Fusion Resources via Pyramid Annotations</i> Daniela Brook Weiss, Paul Roit, Ori Ernst and Ido Dagan .....	1854
<i>The Devil is in the Details: On the Pitfalls of Vocabulary Selection in Neural Machine Translation</i> Tobias Domhan, Eva Hasler, Ke Tran, Sony Trenous, Bill Byrne and Felix Hieber .....	1861
<i>MultiCite: Modeling realistic citations requires moving beyond the single-sentence single-label setting</i> Anne Lauscher, Brandon Ko, Bailey Kuehl, Sophie Johnson, Arman Cohan, David Jurgens and Kyle Lo .....	1875

<i>DEGREE: A Data-Efficient Generation-Based Event Extraction Model</i>	
I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang and Nanyun Peng .....	1890
<i>Bridging the Gap between Language Models and Cross-Lingual Sequence Labeling</i>	
Nuo Chen, Linjun Shou, Ming Gong, Jian Pei and Daxin Jiang .....	1909
<i>Hero-Gang Neural Model For Named Entity Recognition</i>	
Jinpeng Hu, Yaling Shen, Yang Liu, Xiang Wan and Tsung-Hui Chang .....	1924
<i>MGIMN: Multi-Grained Interactive Matching Network for Few-shot Text Classification</i>	
Jianhai Zhang, Mieradilijiang Maimaiti, Gao Xing, Yuanhang Zheng and Ji Zhang .....	1937
<i>All You May Need for VQA are Image Captions</i>	
Soravit Changpinyo, Doron Kukliansy, Idan Szepktor, Xi Chen, Nan Ding and Radu Soricut	1947
<i>Frustratingly Easy System Combination for Grammatical Error Correction</i>	
Muhammad Reza Qorib, Seung-Hoon Na and Hwee Tou Ng .....	1964
<i>Simple Local Attentions Remain Competitive for Long-Context Tasks</i>	
Wenhan Xiong, Barlas Oguz, Anchit Gupta, Xilun Chen, Diana Liskovich, Omer Levy, Scott Yih and Yashar Mehdad .....	1975
<i>Even the Simplest Baseline Needs Careful Re-investigation: A Case Study on XML-CNN</i>	
Si-An Chen, Jie-jyun Liu, Tsung-Han Yang, Hsuan-Tien Lin and Chih-Jen Lin .....	1987
<i>Multi-Relational Graph Transformer for Automatic Short Answer Grading</i>	
Rajat Agarwal, Varun Khurana, Karish Grover, Mukesh Mohania and Vikram Goyal .....	2001
<i>Event Schema Induction with Double Graph Autoencoders</i>	
Xiaomeng Jin, Manling Li and Heng Ji .....	2013
<i>CSIQA: A Dataset for Assisting Code-based Question Answering in an Introductory Programming Course</i>	
Changyoon Lee, Yeon Seonwoo and Alice Oh .....	2026
<i>Unsupervised Cross-Lingual Transfer of Structured Predictors without Source Data</i>	
Kemal Kurniawan, Lea Frermann, Philip Schulz and Trevor Cohn .....	2041
<i>Don't Take It Literally: An Edit-Invariant Sequence Loss for Text Generation</i>	
Guangyi Liu, Zichao Yang, Tianhua Tao, Xiaodan Liang, Junwei Bao, Zhen Li, Xiaodong He, Shuguang Cui and Zhiting Hu .....	2055
<i>Modeling Exemplification in Long-form Question Answering via Retrieval</i>	
Shufan Wang, Fangyuan Xu, Laure Thompson, Eunsol Choi and Mohit Iyyer .....	2079
<i>D2U: Distance-to-Uniform Learning for Out-of-Scope Detection</i>	
Eyup Halit Yilmaz and Cagri Toraman .....	2093
<i>Reference-free Summarization Evaluation via Semantic Correlation and Compression Ratio</i>	
Yizhu Liu, Qi Jia and Kenny Q. Zhu .....	2109
<i>KroneckerBERT: Significant Compression of Pre-trained Language Models Through Kronecker Decomposition and Knowledge Distillation</i>	
Marzieh S. Tahaei, Ella Charlaix, Vahid Partovi Nia, Ali Ghodsi and Mehdi Rezagholizadeh	2116

<i>Building a Role Specified Open-Domain Dialogue System Leveraging Large-Scale Language Models</i> Sanghwan Bae, Donghyun Kwak, Sungdong Kim, Donghoon Ham, Soyoung Kang, Sang-Woo Lee and Woomyoung Park .....	2128
<i>Sentence-Level Resampling for Named Entity Recognition</i> Xiaochen Wang and Yue Wang .....	2151
<i>Word Tour: One-dimensional Word Embeddings via the Traveling Salesman Problem</i> Ryoma Sato .....	2166
<i>On the Diversity and Limits of Human Explanations</i> Chenhao Tan .....	2173
<i>Locally Aggregated Feature Attribution on Natural Language Model Understanding</i> Sheng Zhang, Jin Wang, Haitao Jiang and Rui Song .....	2189
<i>Generic and Trend-aware Curriculum Learning for Relation Extraction</i> Nidhi Vakil and Hadi Amiri .....	2202
<i>On Systematic Style Differences between Unsupervised and Supervised MT and an Application for High-Resource Machine Translation</i> Kelly Marchisio, Markus Freitag and David Grangier .....	2214
<i>Evidentiality-guided Generation for Knowledge-Intensive NLP Tasks</i> Akari Asai, Matt Gardner and Hannaneh Hajishirzi .....	2226
<i>Modularized Transfer Learning with Multiple Knowledge Graphs for Zero-shot Commonsense Reasoning</i> Yu Jin Kim, Beong-woo Kwak, Youngwook Kim, Reinald Kim Amplayo, Seung-won Hwang and Jinyoung Yeo .....	2244
<i>Learning to Express in Knowledge-Grounded Conversation</i> Xueliang Zhao, Tingchen Fu, Chongyang Tao, Wei Wu, Dongyan Zhao and Rui Yan .....	2258
<i>End-to-End Chinese Speaker Identification</i> Dian Yu, Ben Zhou and Dong Yu .....	2274
<i>MINION: a Large-Scale and Diverse Dataset for Multilingual Event Detection</i> Amir Pouran Ben Veyseh, Minh Van Nguyen, Franck Dernoncourt and Thien Huu Nguyen .....	2286
<i>Do Prompt-Based Models Really Understand the Meaning of Their Prompts?</i> Albert Webson and Ellie Pavlick .....	2300
<i>GPL: Generative Pseudo Labeling for Unsupervised Domain Adaptation of Dense Retrieval</i> Kexin Wang, Nandan Thakur, Nils Reimers and Iryna Gurevych .....	2345
<i>Sparse Distillation: Speeding Up Text Classification by Using Bigger Student Models</i> Qinyuan Ye, Madian Khabsa, Mike Lewis, Sinong Wang, Xiang Ren and Aaron Jaech .....	2361
<i>Towards Understanding Large-Scale Discourse Structures in Pre-Trained and Fine-Tuned Language Models</i> Patrick Huber and Giuseppe Carenini .....	2376
<i>SAIS: Supervising and Augmenting Intermediate Steps for Document-Level Relation Extraction</i> Yuxin Xiao, Zecheng Zhang, Yuning Mao, Carl Yang and Jiawei Han .....	2395

<i>LITE: Intent-based Task Representation Learning Using Weak Supervision</i> Naoki Otani, Michael Gamon, Sujay Kumar Jauhar, Mei Yang, Sri Raghu Malireddi and Oriana Riva .....	2410
<i>Does Summary Evaluation Survive Translation to Other Languages?</i> Spencer Braun, Oleg Vasilyev, Neslihan Iskender and John Bohannon .....	2425
<i>A Shoulder to Cry on: Towards A Motivational Virtual Assistant for Assuaging Mental Agony</i> Tulika Saha, Saichethan Miriyala Reddy, Anindya Sundar Das, Sriparna Saha and Pushpak Bhat- tacharyya .....	2436
<i>SueNes: A Weakly Supervised Approach to Evaluating Single-Document Summarization via Negative Sampling</i> Forrest Sheng Bao, Ge Luo, Hebi Li, Minghui Qiu, Yinfei Yang, Youbiao He and Cen Chen .....	2450
<i>Combating the Curse of Multilinguality in Cross-Lingual WSD by Aligning Sparse Contextualized Word Representations</i> Gábor Berend .....	2459
<i>Cheat Codes to Quantify Missing Source Information in Neural Machine Translation</i> Proyag Pal and Kenneth Heafield .....	2472
<i>WiC = TSV = WSD: On the Equivalence of Three Semantic Tasks</i> Bradley Hauer and Grzegorz Kondrak .....	2478
<i>What do tokens know about their characters and how do they know it?</i> Ayush Kaushal and Kyle Mahowald .....	2487
<i>AnswerSumm: A Manually-Curated Dataset and Pipeline for Answer Summarization</i> Alexander Fabbri, Xiaojian Wu, Srini Iyer, Haoran Li and Mona T. Diab .....	2508
<i>Paragraph-based Transformer Pre-training for Multi-Sentence Inference</i> Luca Di Liello, Siddhant Garg, Luca Soldaini and Alessandro Moschitti .....	2521
<i>Text Style Transfer via Optimal Transport</i> Nasim Nouri .....	2532
<i>Exploring the Role of Task Transferability in Large-Scale Multi-Task Learning</i> Vishakh Padmakumar, Leonard Lausen, Miguel Ballesteros, Sheng Zha, He He and George Karypis .....	2542
<i>Interactive Query-Assisted Summarization via Deep Reinforcement Learning</i> Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Ido Dagan and Yael Amsterdamer .....	2551
<i>Data Augmentation with Dual Training for Offensive Span Detection</i> Nasim Nouri .....	2569
<i>Training Mixed-Domain Translation Models via Federated Learning</i> Peyman Passban, Tanya Roosta, Rahul Gupta, Ankit Chadha and Clement Chung .....	2576
<i>QAFactEval: Improved QA-Based Factual Consistency Evaluation for Summarization</i> Alexander Fabbri, Chien-Sheng Wu, Wenhao Liu and Caiming Xiong .....	2587
<i>How Gender Debiasing Affects Internal Model Representations, and Why It Matters</i> Hadas Orgad, Seraphina Goldfarb-Tarrant and Yonatan Belinkov .....	2602
<i>A Structured Span Selector</i> Tianyu Liu, Yuchen Eleanor Jiang, Ryan Cotterell and Mrinmaya Sachan .....	2629

<i>Unified Semantic Typing with Meaningful Label Inference</i>	
James Y. Huang, Bangzheng Li, Jiashu Xu and Muhao Chen . . . . .	2642
<i>Learning To Retrieve Prompts for In-Context Learning</i>	
Ohad Rubin, Jonathan Herzig and Jonathan Berant . . . . .	2655
<i>Necessity and Sufficiency for Explaining Text Classifiers: A Case Study in Hate Speech Detection</i>	
Esma Balkir, Isar Nejadgholi, Kathleen C. Fraser and Svetlana Kiritchenko . . . . .	2672
<i>Learning to Retrieve Passages without Supervision</i>	
Ori Ram, Gal Shachaf, Omer Levy, Jonathan Berant and Amir Globerson . . . . .	2687
<i>Re2G: Retrieve, Rerank, Generate</i>	
Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Ankita Naik, Pengshan Cai and Alfio Gliozzo . . . . .	2701
<i>Don't sweat the small stuff, classify the rest: Sample Shielding to protect text classifiers against adversarial attacks</i>	
Jonathan Rusert and Padmini Srinivasan . . . . .	2716
<i>Federated Learning with Noisy User Feedback</i>	
Rahul Sharma, Anil Ramakrishna, Ansel MacLaughlin, Anna Rumshisky, Jimit Majmudar, Clement Chung, Salman Avestimehr and Rahul Gupta . . . . .	2726
<i>Gender Bias in Masked Language Models for Multiple Languages</i>	
Masahiro Kaneko, Aizhan Imankulova, Danushka Bollegala and Naoaki Okazaki . . . . .	2740
<i>Multi-Domain Targeted Sentiment Analysis</i>	
Orith Toledo-Ronen, Matan Orbach, Yoav Katz and Noam Slonim . . . . .	2751
<i>Falsesum: Generating Document-level NLI Examples for Recognizing Factual Inconsistency in Summarization</i>	
Prasetya Ajie Utama, Joshua Bambrick, Nafise Sadat Moosavi and Iryna Gurevych . . . . .	2763
<i>Dynamic Gazetteer Integration in Multilingual Models for Cross-Lingual and Cross-Domain Named Entity Recognition</i>	
Besnik Fetahu, Anjie Fang, Oleg Rokhlenko and Shervin Malmasi . . . . .	2777
<i>MetaICL: Learning to Learn In Context</i>	
Sewon Min, Mike Lewis, Luke Zettlemoyer and Hannaneh Hajishirzi . . . . .	2791
<i>Enhancing Knowledge Selection for Grounded Dialogues via Document Semantic Graphs</i>	
Sha Li, Mahdi Namazifar, Di Jin, Mohit Bansal, Heng Ji, Yang Liu and Dilek Hakkani-Tur . . . . .	2810
<i>Using Natural Sentence Prompts for Understanding Biases in Language Models</i>	
Sarah Alnegheimish, Alicia Guo and Yi Sun . . . . .	2824
<i>Robust Conversational Agents against Imperceptible Toxicity Triggers</i>	
Ninareh Mehrabi, Ahmad Beirami, Fred Morstatter and Aram Galstyan . . . . .	2831
<i>Selective Differential Privacy for Language Modeling</i>	
Weiyang Shi, Aiqi Cui, Evan Li, Ruoxi Jia and Zhou Yu . . . . .	2848
<i>Do Trajectories Encode Verb Meaning?</i>	
Dylan Ebert, Chen Sun and Ellie Pavlick . . . . .	2860
<i>Long Context Question Answering via Supervised Contrastive Learning</i>	
Avi Caciularu, Ido Dagan, Jacob Goldberger and Arman Cohan . . . . .	2872

<i>The USMLE® Step 2 Clinical Skills Patient Note Corpus</i>	
Victoria Yaneva, Janet Mee, Le An Ha, Polina Harik, Michael Jodoin and Alex J Mechaber	2880
<i>Learning to Borrow– Relation Representation for Without-Mention Entity-Pairs for Knowledge Graph Completion</i>	
Huda Hakami, Mona Hakami, Angrosh Mandya and Danushka Bollegala	2887
<i>Improving Entity Disambiguation by Reasoning over a Knowledge Base</i>	
Tom Ayoola, Joseph Fisher and Andrea Pierleoni	2899
<i>Modal Dependency Parsing via Language Model Priming</i>	
Jiarui Yao, Nianwen Xue and Bonan Min	2913
<i>Document-Level Relation Extraction with Sentences Importance Estimation and Focusing</i>	
Wang Xu, Kehai Chen, Lili Mou and Tiejun Zhao	2920
<i>Are All the Datasets in Benchmark Necessary? A Pilot Study of Dataset Evaluation for Text Classification</i>	
Yang Xiao, Jinlan Fu, See-Kiong Ng and Pengfei Liu	2930
<i>Triggerless Backdoor Attack for NLP Tasks with Clean Labels</i>	
Leilei Gan, Jiwei Li, Tianwei Zhang, Xiaoya Li, Yuxian Meng, Fei Wu, Yi Yang, Shangwei Guo and Chun Fan	2942
<i>PPL-MCTS: Constrained Textual Generation Through Discriminator-Guided MCTS Decoding</i>	
Antoine Chaffin, Vincent Claveau and Ewa Kijak	2953
<i>Interpretable Proof Generation via Iterative Backward Reasoning</i>	
Hanhao Qu, Yu Cao, Jun Gao, Liang Ding and Ruifeng Xu	2968
<i>Domain Confused Contrastive Learning for Unsupervised Domain Adaptation</i>	
Quanyu Long, Tianze Luo, Wenya Wang and Sinno Pan	2982
<i>Incorporating Centering Theory into Neural Coreference Resolution</i>	
Haixia Chai and Michael Strube	2996
<i>Progressive Class Semantic Matching for Semi-supervised Text Classification</i>	
Haiming Xu, Lingqiao Liu and Ehsan M Abbasnejad	3003
<i>Low Resource Style Transfer via Domain Adaptive Meta Learning</i>	
Xiangyang Li, Xiang Long, Yu Xia and Sujian Li	3014
<i>Features or Spurious Artifacts? Data-centric Baselines for Fair and Robust Hate Speech Detection</i>	
Alan Ramponi and Sara Tonelli	3027
<i>Document-Level Event Argument Extraction by Leveraging Redundant Information and Closed Boundary Loss</i>	
Hanzhang Zhou and Kezhi Mao	3041
<i>A Few Thousand Translations Go a Long Way! Leveraging Pre-trained Models for African News Translation</i>	
David Ifeoluwa Adelani, Jesujoba Oluwadara Alabi, Angela Fan, Julia Kreutzer, Xiaoyu Shen, Machel Reid, Dana Ruiter, Dietrich Klakow, Peter Nabende, Ernie Chang, Tajuddeen Gwadabe, Freshia Sackey, Bonaventure F. P. Dossou, Chris Chinenye Emezue, Colin Leong, Michael Beukman, Shamsuddeen Hassan Muhammad, Guyo Dub Jarso, Oreen Yousuf, Andre Niyongabo Rubungo, Gilles Hacheme, Eric Peter Wairagala, Muhammad Umair Nasir, Benjamin Ayoade Ajibade, Tunde Oluwaseyi Ajayi, Yvonne Wambui Gitau, Jade Abbott, Mohamed Ahmed, Millicent Ochieng, Anuoluwapo	

Aremu, Perez Ogayo, Jonathan Mukiibi, Fatoumata Ouoba Kabore, Godson Koffi Kalipe, Derguene Mbaye, Allahsera Auguste Tapo, Victoire Memdjokam Koagne, Edwin Munkoh-Buabeng, Valencia Wagner, Idris Abdulmumin, Ayodele Awokoya, Happy Buzaaba, Blessing Sibanda, Andiswa Bukula and Sam Manthalu .....	3053
<i>Should We Rely on Entity Mentions for Relation Extraction? Debiasing Relation Extraction with Counterfactual Analysis</i>	
Yiwei Wang, Muhao Chen, Wenxuan Zhou, Yujun Cai, Yuxuan Liang, Dayiheng Liu, Baosong Yang, Juncheng Liu and Bryan Hooi .....	3071
<i>Analyzing Encoded Concepts in Transformer Language Models</i>	
Hassan Sajjad, Nadir Durrani, Fahim Dalvi, Firoj Alam, Abdul Rafae Khan and Jia Xu ....	3082
<i>Boosted Dense Retriever</i>	
Patrick Lewis, Barlas Oguz, Wenhan Xiong, Fabio Petroni, Scott Yih and Sebastian Riedel .	3102
<i>MuCGEC: a Multi-Reference Multi-Source Evaluation Dataset for Chinese Grammatical Error Correction</i>	
Yue Zhang, Zhenghua Li, Zuyi Bao, Jiacheng Li, Bo Zhang, Chen Li, Fei Huang and Min Zhang	3118
<i>NeuS: Neutral Multi-News Summarization for Mitigating Framing Bias</i>	
Nayeon Lee, Yejin Bang, Tiezheng YU, Andrea Madotto and Pascale Fung .....	3131
<i>Enhance Incomplete Utterance Restoration by Joint Learning Token Extraction and Text Generation</i>	
Shumpei Inoue, Tsungwei Liu, Son Hong Nguyen and Minh-Tien Nguyen.....	3149
<i>Efficient Constituency Tree based Encoding for Natural Language to Bash Translation</i>	
Shikhar Bharadwaj and Shirish Shevade.....	3159
<i>Privacy-Preserving Text Classification on BERT Embeddings with Homomorphic Encryption</i>	
Garam Lee, Minsoo Kim, Jai Hyun Park, Seung-won Hwang and Jung Hee Cheon .....	3169
<i>ITA: Image-Text Alignments for Multi-Modal Named Entity Recognition</i>	
Xinyu Wang, Min Gui, Yong Jiang, Zixia Jia, Nguyen Bach, Tao Wang, Zhongqiang Huang and Kewei Tu .....	3176
<i>A Dataset for N-ary Relation Extraction of Drug Combinations</i>	
Aryeh Tiktinsky, Vijay Viswanathan, Danna Niezni, Dana Meron Azagury, Yosi Shamay, Hillel Taub-Tabib, Tom Hope and Yoav Goldberg .....	3190
<i>Curriculum: A Broad-Coverage Benchmark for Linguistic Phenomena in Natural Language Understanding</i>	
Zeming Chen and Qiyue Gao .....	3204
<i>Neural Language Taskonomy: Which NLP Tasks are the most Predictive of fMRI Brain Activity?</i>	
Subba Reddy Oota, Jashn Arora, Veeral Agarwal, Mounika Marreddy, Manish Gupta and Bapi Raju Surampudi .....	3220
<i>FactGraph: Evaluating Factuality in Summarization with Semantic Graph Representations</i>	
Leonardo F. R. Ribeiro, Mengwen Liu, Iryna Gurevych, Markus Dreyer and Mohit Bansal .	3238
<i>Unsupervised Paraphrasability Prediction for Compound Nominalizations</i>	
John Sie Yuen Lee, Ho Hung Lim and Carol Carol Webster .....	3254
<i>Global Entity Disambiguation with BERT</i>	
Ikuya Yamada, Koki Washio, Hiroyuki Shindo and Yuji Matsumoto.....	3264



<i>Clues Before Answers: Generation-Enhanced Multiple-Choice QA</i>	
Zixian Huang, Ao Wu, Jiaying Zhou, Yu Gu, Yue Zhao and Gong Cheng .....	3272
<i>Towards Efficient NLP: A Standard Evaluation and A Strong Baseline</i>	
Xiangyang Liu, Tianxiang Sun, Junliang He, Jiawen Wu, Lingling Wu, Xinyu Zhang, Hao Jiang, Zhao Cao, Xuanjing Huang and Xipeng Qiu .....	3288
<i>Stylized Knowledge-Grounded Dialogue Generation via Disentangled Template Rewriting</i>	
Qingfeng Sun, Can Xu, Huang Hu, Yujing Wang, Jian Miao, Xiubo Geng, Yining Chen, Fei Xu and Daxin Jiang .....	3304
<i>LUNA: Learning Slot-Turn Alignment for Dialogue State Tracking</i>	
Yifan Wang, Jing Zhao, Junwei Bao, Chaoqun Duan, Youzheng Wu and Xiaodong He .....	3319
<i>Crossroads, Buildings and Neighborhoods: A Dataset for Fine-grained Location Recognition</i>	
Pei Chen, Haotian Xu, Cheng Zhang and Ruihong Huang .....	3329
<i>Tricks for Training Sparse Translation Models</i>	
Dheeru Dua, Shruti Bhosale, Vedanuj Goswami, James Cross, Mike Lewis and Angela Fan .....	3340
<i>Persona-Guided Planning for Controlling the Protagonist’s Persona in Story Generation</i>	
Zhexin Zhang, Jiaxin Wen, Jian Guan and Minlie Huang .....	3346
<i>CHEF: A Pilot Chinese Dataset for Evidence-Based Fact-Checking</i>	
Xuming Hu, Zhijiang Guo, GuanYu Wu, Aiwei Liu, Lijie Wen and Philip S. Yu .....	3362
<i>VGNMN: Video-grounded Neural Module Networks for Video-Grounded Dialogue Systems</i>	
Hung Le, Nancy F. Chen and Steven Hoi .....	3377
<i>Multimodal Dialogue State Tracking</i>	
Hung Le, Nancy F. Chen and Steven Hoi .....	3394
<i>On the Use of Bert for Automated Essay Scoring: Joint Learning of Multi-Scale Essay Representation</i>	
Yongjie Wang, Chuang Wang, Ruobing Li and Hui Lin .....	3416
<i>Recognition of They/Them as Singular Personal Pronouns in Coreference Resolution</i>	
Connor Baumler and Rachel Rudinger .....	3426
<i>TWEETSPIN: Fine-grained Propaganda Detection in Social Media Using Multi-View Representations</i>	
Prashanth Vijayaraghavan and Soroush Vosoughi .....	3433
<i>UserIdentifier: Implicit User Representations for Simple and Effective Personalized Sentiment Analysis</i>	
Fatemehsadat Mireshghallah, Vaishnavi Shrivastava, Milad Shokouhi, Taylor Berg-Kirkpatrick, Robert Sim and Dimitrios Dimitriadis .....	3449
<i>Improving Neural Models for Radiology Report Retrieval with Lexicon-based Automated Annotation</i>	
Luyao Shi, Tanveer Syeda-mahmood and Tyler Baldwin .....	3457
<i>Transparent Human Evaluation for Image Captioning</i>	
Jungo Kasai, Keisuke Sakaguchi, Lavinia Dunagan, Jacob Daniel Morrison, Ronan Le Bras, Yejin Choi and Noah Smith .....	3464
<i>Lifting the Curse of Multilinguality by Pre-training Modular Transformers</i>	
Jonas Pfeiffer, Naman Goyal, Xi Victoria Lin, Xian Li, James Cross, Sebastian Riedel and Mikel Artetxe .....	3479

<i>DocAMR: Multi-Sentence AMR Representation and Evaluation</i>	
Tahira Naseem, Austin Blodgett, Sadhana Kumaravel, Tim O’Gorman, Young-Suk Lee, Jeffrey Flanigan, Ramón Fernandez Astudillo, Radu Florian, Salim Roukos and Nathan Schneider . . . . .	3496
<i>Learning to Transfer Prompts for Text Generation</i>	
Junyi Li, Tianyi Tang, Jian-Yun Nie, Ji-Rong Wen and Xin Zhao . . . . .	3506
<i>ElitePLM: An Empirical Study on General Language Ability Evaluation of Pretrained Language Models</i>	
Junyi Li, Tianyi Tang, Zheng Gong, Lixin Yang, Zhuohao Yu, Zhipeng Chen, Jingyuan Wang, Xin Zhao and Ji-Rong Wen . . . . .	3519
<i>Bidimensional Leaderboards: Generate and Evaluate Language Hand in Hand</i>	
Jungo Kasai, Keisuke Sakaguchi, Ronan Le Bras, Lavinia Dunagan, Jacob Daniel Morrison, Alexander Fabbri, Yejin Choi and Noah Smith . . . . .	3540
<i>Improving In-Context Few-Shot Learning via Self-Supervised Training</i>	
Mingda Chen, Jingfei Du, Ramakanth Pasunuru, Todor Mihaylov, Srini Iyer, Veselin Stoyanov and Zornitsa Kozareva . . . . .	3558
<i>Exposing the Limits of Video-Text Models through Contrast Sets</i>	
Jae Sung Park, Sheng Shen, Ali Farhadi, Trevor Darrell, Yejin Choi and Anna Rohrbach . . . . .	3574
<i>Zero-shot Sonnet Generation with Discourse-level Planning and Aesthetics Features</i>	
Yufei Tian and Nanyun Peng . . . . .	3587
<i>Benchmarking Intersectional Biases in NLP</i>	
John P. Lalor, Yi Yang, Kendall Smith, Nicole Forsgren and Ahmed Abbasi . . . . .	3598
<i>When is BERT Multilingual? Isolating Crucial Ingredients for Cross-lingual Transfer</i>	
Ameet Deshpande, Partha Talukdar and Karthik R Narasimhan . . . . .	3610
<i>How Conservative are Language Models? Adapting to the Introduction of Gender-Neutral Pronouns</i>	
Stephanie Brandl, Ruixiang Cui and Anders Søgaard . . . . .	3624
<i>Prompt Waywardness: The Curious Case of Discretized Interpretation of Continuous Prompts</i>	
Daniel Khashabi, Xinxu Lyu, Sewon Min, Lianhui Qin, Kyle Richardson, Sean Welleck, Hannaneh Hajishirzi, Tushar Khot, Ashish Sabharwal, Sameer Singh and Yejin Choi . . . . .	3631
<i>Contrastive Representation Learning for Cross-Document Coreference Resolution of Events and Entities</i>	
Benjamin Hsu and Graham Horwood . . . . .	3644
<i>Learning the Ordering of Coordinate Compounds and Elaborate Expressions in Hmong, Lahu, and Chinese</i>	
Chenxuan Cui, Katherine J. Zhang and David R Mortensen . . . . .	3656
<i>FRUIT: Faithfully Reflecting Updated Information in Text</i>	
Robert L. Logan IV, Alexandre Tachard Passos, Sameer Singh and Ming-Wei Chang . . . . .	3670
<i>Multi2WOZ: A Robust Multilingual Dataset and Conversational Pretraining for Task-Oriented Dialog</i>	
Chia-Chien Hung, Anne Lauscher, Ivan Vulić, Simone Paolo Ponzetto and Goran Glavaš . . . . .	3687
<i>ChapterBreak: A Challenge Dataset for Long-Range Language Models</i>	
Simeng Sun, Katherine Thai and Mohit Iyyer . . . . .	3704
<i>ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction</i>	
Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts and Matei Zaharia . . . . .	3715

<i>Quantifying Language Variation Acoustically with Few Resources</i>	
Martijn Bartelds and Martijn Wieling .....	3735
<i>Adaptable Adapters</i>	
Nafise Sadat Moosavi, Quentin Delfosse, Kristian Kersting and Iryna Gurevych .....	3742
<i>Models in the Loop: Aiding Crowdworkers with Generative Annotation Assistants</i>	
Max Bartolo, Tristan Thrush, Sebastian Riedel, Pontus Stenetorp, Robin Jia and Douwe Kiela	3754
<i>GMN: Generative Multi-modal Network for Practical Document Information Extraction</i>	
Haoyu Cao, Jiefeng Ma, Antai Guo, Yiqing Hu, Hao Liu, Deqiang Jiang, Yinsong Liu and Bo Ren .....	3768
<i>One Reference Is Not Enough: Diverse Distillation with Reference Selection for Non-Autoregressive Translation</i>	
Chenze Shao, Xuanfu Wu and Yang Feng .....	3779
<i>Can Rationalization Improve Robustness?</i>	
Howard Chen, Jacqueline He, Karthik R Narasimhan and Danqi Chen .....	3792
<i>On the Effectiveness of Sentence Encoding for Intent Detection Meta-Learning</i>	
Tingting Ma, Qianhui Wu, Zhiwei Yu, Tiejun Zhao and Chin-Yew Lin .....	3806
<i>A Computational Acquisition Model for Multimodal Word Categorization</i>	
Uri Berger, Gabriel Stanovsky, Omri Abend and Lea Freermann .....	3819
<i>Residue-Based Natural Language Adversarial Attack Detection</i>	
Vyas Raina and Mark Gales .....	3836
<i>Does it Really Generalize Well on Unseen Data? Systematic Evaluation of Relational Triple Extraction Methods</i>	
Juhyuk Lee, Min-Joong Lee, June Yong Yang and Eunho Yang .....	3849
<i>From spoken dialogue to formal summary: An utterance rewriting for dialogue summarization</i>	
Yue Fang, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Bo Long, Yanyan Lan and Yanquan Zhou .....	3859
<i>EASE: Entity-Aware Contrastive Learning of Sentence Embedding</i>	
Sosuke Nishikawa, Ryokan Ri, Ikuya Yamada, Yoshimasa Tsuruoka and Isao Echizen .....	3870
<i>Is Neural Topic Modelling Better than Clustering? An Empirical Study on Clustering with Contextual Embeddings for Topics</i>	
Zihan Zhang, Meng Fang, Ling Chen and Mohammad Reza Namazi Rad .....	3886
<i>Dynamic Multistep Reasoning based on Video Scene Graph for Video Question Answering</i>	
Jianguo Mao, Wenbin Jiang, Xiangdong Wang, Zhifan Feng, Yajuan Lyu, Hong Liu and Yong Zhu .....	3894
<i>TRUE: Re-evaluating Factual Consistency Evaluation</i>	
Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim and Yossi Matias .....	3905
<i>Knowledge Inheritance for Pre-trained Language Models</i>	
Yujia Qin, Yankai Lin, Jing Yi, Jiajie Zhang, Xu Han, Zhengyan Zhang, Yusheng Su, Zhiyuan Liu, Peng Li, Maosong Sun and Jie Zhou .....	3921

<i>Bi-SimCut: A Simple Strategy for Boosting Neural Machine Translation</i> Pengzhi Gao, Zhongjun He, Hua Wu and Haifeng Wang .....	3938
<i>On Transferability of Prompt Tuning for Natural Language Processing</i> Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Huadong Wang, Kaiyue Wen, Zhiyuan Liu, Peng Li, Juanzi Li, Lei Hou, Maosong Sun and Jie Zhou .....	3949
<i>DocEE: A Large-Scale and Fine-grained Benchmark for Document-level Event Extraction</i> MeiHan Tong, Bin Xu, Shuai Wang, Meihuan Han, Yixin Cao, Jiangqi Zhu, Siyu Chen, Lei Hou and Juanzi Li .....	3970
<i>Towards Debiasing Translation Artifacts</i> Koel Dutta Chowdhury, Richa Jalota, Cristina España-Bonet and Josef Van Genabith .....	3983
<i>WECHSEL: Effective initialization of subword embeddings for cross-lingual transfer of monolingual language models</i> Benjamin Minixhofer, Fabian Paischer and Navid Rekasaz .....	3992
<i>A New Concept of Knowledge based Question Answering (KBQA) System for Multi-hop Reasoning</i> Yu Wang, v.srinivasan@samsung.com v.srinivasan@samsung.com and Hongxia Jin .....	4007
<i>Bilingual Tabular Inference: A Case Study on Indic Languages</i> Chaitanya Agarwal, Vivek Gupta, Anoop Kunchukuttan and Manish Shrivastava .....	4018
<i>Generative Biomedical Entity Linking via Knowledge Base-Guided Pre-training and Synonyms-Aware Fine-tuning</i> Hongyi Yuan, Zheng Yuan and Sheng Yu .....	4038
<i>Robust Self-Augmentation for Named Entity Recognition with Meta Reweighting</i> Linzhi Wu, Pengjun Xie, Jie Zhou, Meishan Zhang, Ma Chunping, Guangwei Xu and Min Zhang 4049	
<i>Unsupervised Stem-based Cross-lingual Part-of-Speech Tagging for Morphologically Rich Low-Resource Languages</i> Ramy Eskander, Cass Lowry, Sujay Khandagale, Judith Lynn Klavans, Maria Polinsky and Smaranda Muresan .....	4061
<i>Optimising Equal Opportunity Fairness in Model Training</i> Aili Shen, Xudong Han, Trevor Cohn, Timothy Baldwin and Lea Frermann .....	4073
<i>Leaner and Faster: Two-Stage Model Compression for Lightweight Text-Image Retrieval</i> Siyu Ren and Kenny Q. Zhu .....	4085
<i>Joint Learning-based Heterogeneous Graph Attention Network for Timeline Summarization</i> Jingyi You, Dongyuan Li, Hidetaka Kamigaito, Kotaro Funakoshi and Manabu Okumura ..	4091
<i>Early Rumor Detection Using Neural Hawkes Process with a New Benchmark Dataset</i> Fengzhu Zeng and Wei Gao .....	4105
<i>Emp-RFT: Empathetic Response Generation via Recognizing Feature Transitions between Utterances</i> Wongyu Kim, Youbin Ahn, Donghyun Kim and Kyong-Ho Lee .....	4118
<i>KCD: Knowledge Walks and Textual Cues Enhanced Political Perspective Detection in News Media</i> Wenqian Zhang, Shangbin Feng, Zilong Chen, Zhenyu Lei, Jundong Li and Minnan Luo ...	4129
<i>Collective Relevance Labeling for Passage Retrieval</i> Jihyuk Kim, Minsoo Kim and Seung-won Hwang .....	4141

<i>COGMEN: COntextualized GNN based Multimodal Emotion recognition</i>	
Abhinav Joshi, Ashwani Bhat, Ayush Jain, Atin Vikram Singh and Ashutosh Modi . . . . .	4148
<i>Revisit Overconfidence for OOD Detection: Reassigned Contrastive Learning with Adaptive Class-dependent Threshold</i>	
Yanan Wu, Keqing He, Yuanmeng Yan, QiXiang Gao, Zhiyuan Zeng, Fujia Zheng, Lulu Zhao, Huixing Jiang, Wei Wu and Weiran Xu . . . . .	4165
<i>AISFG: Abundant Information Slot Filling Generator</i>	
Yang Yan, Junda Ye, Zhongbao Zhang and Liwen Wang . . . . .	4180
<i>Improving negation detection with negation-focused pre-training</i>	
Thinh Hung Truong, Timothy Baldwin, Trevor Cohn and Karin Verspoor . . . . .	4188
<i>Practice Makes a Solver Perfect: Data Augmentation for Math Word Problem Solvers</i>	
Vivek Kumar, Rishabh Maheshwary and Vikram Pudi . . . . .	4194
<i>DiffCSE: Difference-based Contrastive Learning for Sentence Embeddings</i>	
Yung-Sung Chuang, Rumen Dangovski, Hongyin Luo, Yang Zhang, Shiyu Chang, Marin Soljagic, Shang-Wen Li, Scott Yih, Yoon Kim and James R. Glass . . . . .	4207
<i>Generative Cross-Domain Data Augmentation for Aspect and Opinion Co-Extraction</i>	
Junjie Li, Jianfei Yu and Rui Xia . . . . .	4219
<i>ProQA: Structural Prompt-based Pre-training for Unified Question Answering</i>	
Wanjun Zhong, Yifan Gao, Ning Ding, Yujia Qin, Zhiyuan Liu, Ming Zhou, Jiahai Wang, Jian Yin and Nan Duan . . . . .	4230
<i>A Data Cartography based MixUp for Pre-trained Language Models</i>	
Seo Yeon Park and Cornelia Caragea . . . . .	4244
<i>Grapheme-to-Phoneme Conversion for Thai using Neural Regression Models</i>	
Tomohiro Yamasaki . . . . .	4251
<i>Generating Authentic Adversarial Examples beyond Meaning-preserving with Doubly Round-trip Translation</i>	
Siyu Lai, Zhen Yang, Fandong Meng, Xue Zhang, Yufeng Chen, Jinan Xu and Jie Zhou . . .	4256
<i>TVShowGuess: Character Comprehension in Stories as Speaker Guessing</i>	
Yisi Sang, Xiangyang Mou, Mo Yu, Shunyu Yao, Jing Li and Jeffrey Stanton . . . . .	4267
<i>Causal Distillation for Language Models</i>	
Zhengxuan Wu, Atticus Geiger, Joshua Rozner, Elisa Kreiss, Hanson Lu, Thomas Icard, Christopher Potts and Noah Goodman . . . . .	4288
<i>FNet: Mixing Tokens with Fourier Transforms</i>	
James Lee-Thorp, Joshua Ainslie, Ilya Eckstein and Santiago Ontanon . . . . .	4296
<i>Answer Consolidation: Formulation and Benchmarking</i>	
Wenxuan Zhou, Qiang Ning, Heba Elfardy, Kevin Small and Muhao Chen . . . . .	4314
<i>Informativeness and Invariance: Two Perspectives on Spurious Correlations in Natural Language</i>	
Jacob Eisenstein . . . . .	4326
<i>FOAM: A Follower-aware Speaker Model For Vision-and-Language Navigation</i>	
Zi-Yi Dou and Nanyun Peng . . . . .	4332

<i>Improving Compositional Generalization with Latent Structure and Data Augmentation</i>	
Linlu Qiu, Peter Shaw, Panupong Pasupat, Pawel Krzysztof Nowak, Tal Linzen, Fei Sha and Kristina Toutanova .....	4341
<i>Joint Extraction of Entities, Relations, and Events via Modeling Inter-Instance and Inter-Label Dependencies</i>	
Minh Van Nguyen, Bonan Min, Franck Dernoncourt and Thien Huu Nguyen .....	4363
<i>Linguistic Frameworks Go Toe-to-Toe at Neuro-Symbolic Language Modeling</i>	
Jakob Prange, Nathan Schneider and Lingpeng Kong .....	4375
<i>Imagination-Augmented Natural Language Understanding</i>	
Yujie Lu, Wanrong Zhu, Xin Eric Wang, Miguel Eckstein and William Yang Wang .....	4392
<i>What company do words keep? Revisiting the distributional semantics of J.R. Firth &amp; Zellig Harris</i>	
Mikael Brunila and Jack LaViolette .....	4403
<i>Compositional Task-Oriented Parsing as Abstractive Question Answering</i>	
Wenting Zhao, Konstantine Arkoudas, Weiqi Sun and Claire Cardie .....	4418
<i>Learning Cross-Lingual IR from an English Retriever</i>	
Yulong Li, Martin Franz, Md Arafat Sultan, Bhavani Iyer, Young-Suk Lee and Avirup Sil ..	4428
<i>Testing the Ability of Language Models to Interpret Figurative Language</i>	
Emmy Liu, Chenxuan Cui, Kenneth Zheng and Graham Neubig .....	4437
<i>Multi-Vector Models with Textual Guidance for Fine-Grained Scientific Document Similarity</i>	
Sheshera Mysore, Arman Cohan and Tom Hope .....	4453
<i>CHAI: A CHatbot AI for Task-Oriented Dialogue with Offline Reinforcement Learning</i>	
Siddharth Verma, Justin Fu, Sherry Yang and Sergey Levine .....	4471
<i>Connecting the Dots between Audio and Text without Parallel Data through Visual Knowledge Transfer</i>	
Yanpeng Zhao, Jack Hessel, Youngjae Yu, Ximing Lu, Rowan Zellers and Yejin Choi .....	4492
<i>SURF: Semantic-level Unsupervised Reward Function for Machine Translation</i>	
Atijit Anuchitanukul and Julia Ive .....	4508
<i>Disentangling Categorization in Multi-agent Emergent Communication</i>	
Washington Garcia, Hamilton Scott Clouse and Kevin R. B. Butler .....	4523
<i>Show, Don't Tell: Demonstrations Outperform Descriptions for Schema-Guided Task-Oriented Dialogue</i>	
Raghav Gupta, Harrison Lee, Jeffrey Zhao, Yuan Cao, Abhinav Rastogi and Yonghui Wu ..	4541
<i>Does Pre-training Induce Systematic Inference? How Masked Language Models Acquire Commonsense Knowledge</i>	
Ian Porada, Alessandro Sordoni and Jackie CK Cheung .....	4550
<i>Using Paraphrases to Study Properties of Contextual Embeddings</i>	
Laura Burdick, Jonathan K Kummerfeld and Rada Mihalcea .....	4558
<i>Measure and Improve Robustness in NLP Models: A Survey</i>	
Xuezhi Wang, Haohan Wang and Diyi Yang .....	4569
<i>Learning to Generate Examples for Semantic Processing Tasks</i>	
Danilo Croce, Simone Filice, Giuseppe Castellucci and Roberto Basili .....	4587

<i>Symbolic Knowledge Distillation: from General Language Models to Commonsense Models</i>	
Peter West, Chandra Bhagavatula, Jack Hessel, Jena D. Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck and Yejin Choi . . . . .	4602
<i>GenIE: Generative Information Extraction</i>	
Martin Josifoski, Nicola De Cao, Maxime Peyrard, Fabio Petroni and Robert West . . . . .	4626
<i>Entity Linking via Explicit Mention-Mention Coreference Modeling</i>	
Dhruv Agarwal, Rico Angell, Nicholas Monath and Andrew McCallum . . . . .	4644
<i>Massive-scale Decoding for Text Generation using Lattices</i>	
Jiacheng Xu, Siddhartha Jonnalagadda and Greg Durrett . . . . .	4659
<i>Disentangling Indirect Answers to Yes-No Questions in Real Conversations</i>	
Krishna Chaitanya Sanagavarapu, Jathin Pranav Singaraju, Anusha Kakileti, Anirudh Kaza, Aaron Abraham Mathews, Helen Li, Nathan Raul Brito and Eduardo Blanco . . . . .	4677
<i>Quantifying Adaptability in Pre-trained Language Models with 500 Tasks</i>	
Belinda Z. Li, Jane A. Yu, Madian Khabsa, Luke Zettlemoyer, Alon Y. Halevy and Jacob Andreas	4696
<i>Counterfactually Augmented Data and Unintended Bias: The Case of Sexism and Hate Speech Detection</i>	
Indira Sen, Mattia Samory, Claudia Wagner and Isabelle Augenstein . . . . .	4716
<i>A Study of the Attention Abnormality in Trojaned BERTs</i>	
Weimin Lyu, Songzhu Zheng, Tengfei Ma and Chao Chen . . . . .	4727
<i>EPiDA: An Easy Plug-in Data Augmentation Framework for High Performance Text Classification</i>	
Minyi Zhao, Lu Zhang, Yi Xu, Jiandong Ding, Jihong Guan and Shuigeng Zhou . . . . .	4742
<i>Partial-input baselines show that NLI models can ignore context, but they don't.</i>	
Neha Srikanth and Rachel Rudinger . . . . .	4753
<i>Lifelong Pretraining: Continually Adapting Language Models to Emerging Corpora</i>	
Xisen Jin, Dejiao Zhang, Henghui Zhu, Wei Xiao, Shang-Wen Li, Xiaokai Wei, Andrew Arnold and Xiang Ren . . . . .	4764
<i>Learning as Conversation: Dialogue Systems Reinforced for Information Acquisition</i>	
Pengshan Cai, Hui Wan, Fei Liu, Mo Yu, Hong yu and Sachindra Joshi . . . . .	4781
<i>Dynamic Programming in Rank Space: Scaling Structured Inference with Low-Rank HMMs and PC-FGs</i>	
Songlin Yang, Wei Liu and Kewei Tu . . . . .	4797
<i>What Factors Should Paper-Reviewer Assignments Rely On? Community Perspectives on Issues and Ideals in Conference Peer-Review</i>	
Terne Sasha Thorn Jakobsen and Anna Rogers . . . . .	4810
<i>Reducing Disambiguation Biases in NMT by Leveraging Explicit Word Sense Information</i>	
Niccolò Campolungo, Tommaso Pasini, Denis Emelin and Roberto Navigli . . . . .	4824
<i>Mining Clues from Incomplete Utterance: A Query-enhanced Network for Incomplete Utterance Rewriting</i>	
Shuzheng Si, Shuang Zeng and Baobao Chang . . . . .	4839

<i>Domain-Oriented Prefix-Tuning: Towards Efficient and Generalizable Fine-tuning for Zero-Shot Dialogue Summarization</i>	
Lulu Zhao, Fujia Zheng, Weihao Zeng, Keqing He, Weiran Xu, Huixing Jiang, Wei Wu and Yanan Wu .....	4848
<i>Interactive Symbol Grounding with Complex Referential Expressions</i>	
Rimvydas Rubavicius and Alex Lascarides .....	4863
<i>Generalized Quantifiers as a Source of Error in Multilingual NLU Benchmarks</i>	
Ruixiang Cui, Daniel Herscovich and Anders Søgaard .....	4875
<i>Exact Paired-Permutation Testing for Structured Test Statistics</i>	
Ran Zmigrod, Tim Vieira and Ryan Cotterell .....	4894
<i>A Balanced Data Approach for Evaluating Cross-Lingual Transfer: Mapping the Linguistic Blood Bank</i>	
Dan Malkin, Tomasz Limisiewicz and Gabriel Stanovsky .....	4903
<i>SSEGCN: Syntactic and Semantic Enhanced Graph Convolutional Network for Aspect-based Sentiment Analysis</i>	
Zheng Zhang, Zili Zhou and Yanna Wang .....	4916
<i>Mitigating Toxic Degeneration with Empathetic Data: Exploring the Relationship Between Toxicity and Empathy</i>	
Allison Lahkala, Charles Welch, Béla Neuendorf and Lucie Flek .....	4926
<i>DUCK: Rumour Detection on Social Media by Modelling User and Comment Propagation Networks</i>	
Lin Tian, Xiuzhen Zhang and Jey Han Lau .....	4939
<i>Jam or Cream First? Modeling Ambiguity in Neural Machine Translation with SCONES</i>	
Felix Stahlberg and Shankar Kumar .....	4950
<i>SkillSpan: Hard and Soft Skill Extraction from English Job Postings</i>	
Mike Zhang, Kristian Nørgaard Jensen, Sif Dam Sonniks and Barbara Plank .....	4962
<i>RAAT: Relation-Augmented Attention Transformer for Relation Modeling in Document-Level Event Extraction</i>	
Yuan Liang, Zhuoxuan Jiang, di Yin and Bo Ren .....	4985
<i>A Double-Graph Based Framework for Frame Semantic Parsing</i>	
Ce Zheng, Xudong Chen, Runxin Xu and Baobao Chang .....	4998
<i>An Enhanced Span-based Decomposition Method for Few-Shot Sequence Labeling</i>	
Peiyi Wang, Runxin Xu, Tianyu Liu, Qingyu Zhou, Yunbo Cao, Baobao Chang and Zhifang Sui	5012
<i>A Two-Stream AMR-enhanced Model for Document-level Event Argument Extraction</i>	
Runxin Xu, Peiyi Wang, Tianyu Liu, Shuang Zeng, Baobao Chang and Zhifang Sui .....	5025
<i>Robust (Controlled) Table-to-Text Generation with Structure-Aware Equivariance Learning</i>	
Fei Wang, Zhewei Xu, Pedro Szekely and Muhao Chen .....	5037
<i>JointLK: Joint Reasoning with Language Models and Knowledge Graphs for Commonsense Question Answering</i>	
Yueqing Sun, Qi Shi, Le Qi and Yu Zhang .....	5049
<i>Models In a Spelling Bee: Language Models Implicitly Learn the Character Composition of Tokens</i>	
Itay Itzhak and Omer Levy .....	5061



<i>A Corpus for Understanding and Generating Moral Stories</i> Jian Guan, Ziqi Liu and Minlie Huang .....	5069
<i>Modeling Multi-Granularity Hierarchical Features for Relation Extraction</i> Xinnian Liang, Shuangzhi Wu, Mu Li and Zhoujun Li .....	5088
<i>Cross-modal Contrastive Learning for Speech Translation</i> Rong Ye, Mingxuan Wang and Lei Li .....	5099
<i>Meet Your Favorite Character: Open-domain Chatbot Mimicking Fictional Characters with only a Few Utterances</i> Seungju Han, Beomsu Kim, Jin Yong Yoo, Seokjun Seo, Sangbum Kim, Enkhbayar Erdenee and Buru Chang .....	5114
<i>DynamicTOC: Persona-based Table of Contents for Consumption of Long Documents</i> Himanshu Maheshwari, Nethraa Sivakumar, Shelly Jain, Tanvi Karandikar, Vinay Aggarwal, Navita Goyal and Sumit Shekhar .....	5133
<i>KALA: Knowledge-Augmented Language Model Adaptation</i> Minki Kang, Jinheon Baek and Sung Ju Hwang .....	5144
<i>On the Effect of Pretraining Corpora on In-context Learning by a Large-scale Language Model</i> Seongjin Shin, Sang-Woo Lee, Hwijee Ahn, Sungdong Kim, HyoungSeok Kim, Boseop Kim, Kyunghyun Cho, Gichang Lee, Woomyoung Park, Jung-Woo Ha and Nako Sung .....	5168
<i>Sketching as a Tool for Understanding and Accelerating Self-attention for Long Sequences</i> Yifan Chen, Qi Zeng, Dilek Hakkani-Tur, Di Jin, Heng Ji and Yun Yang .....	5187
<i>Partner Personas Generation for Dialogue Response Generation</i> Hongyuan Lu, Wai Lam, Hong Cheng and Helen M. Meng .....	5200
<i>Semantically Informed Slang Interpretation</i> Zhewei Sun, Richard Zemel and Yang Xu .....	5213
<i>Dual-Channel Evidence Fusion for Fact Verification over Texts and Tables</i> Nan Hu, Zirui Wu, Yuxuan Lai, Xiao Liu and Yansong Feng .....	5232
<i>TreeMix: Compositional Constituency-based Data Augmentation for Natural Language Understanding</i> Le Zhang, Zichao Yang and Diyi Yang .....	5243
<i>Syn2Vec: Synset Colexification Graphs for Lexical Semantic Similarity</i> John Harvill, Roxana Girju and Mark A. Hasegawa-Johnson .....	5259
<i>On the Origin of Hallucinations in Conversational Models: Is it the Datasets or the Models?</i> Nouha Dziri, Sivan Milton, Mo Yu, Osmar Zaiane and Siva Reddy .....	5271
<i>Is My Favorite New Movie My Favorite Movie? Probing the Understanding of Recursive Noun Phrases</i> Qing Lyu, Zheng Hua, Daoxin Li, Li Zhang, Marianna Apidianaki and Chris Callison-Burch	5286
<i>Original or Translated? A Causal Analysis of the Impact of Translationese on Machine Translation Performance</i> Jingwei Ni, Zhijing Jin, Markus Freitag, Mrinmaya Sachan and Bernhard Schölkopf .....	5303
<i>Visual Commonsense in Pretrained Unimodal and Multimodal Models</i> Chenyu Zhang, Benjamin Van Durme, Zhuowan Li and Elias Stengel-Eskin .....	5321

<i>QuALITY: Question Answering with Long Input Texts, Yes!</i>	
Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny L Ma, Jana Thompson, He He and Samuel R. Bowman . . . . .	5336
<i>ExSum: From Local Explanations to Model Understanding</i>	
Yilun Zhou, Marco Tulio Ribeiro and Julie Shah . . . . .	5359
<i>Maximum Bayes Smatch Ensemble Distillation for AMR Parsing</i>	
Young-Suk Lee, Ramón Fernandez Astudillo, Hoang Thanh Lam, Tahira Naseem, Radu Florian and Salim Roukos . . . . .	5379
<i>When Does Syntax Mediate Neural Language Model Performance? Evidence from Dropout Probes</i>	
Mycal Tucker, Tiwalayo Eisape, Peng Qian, Roger P. Levy and Julie Shah . . . . .	5393
<i>Modeling Task Interactions in Document-Level Joint Entity and Relation Extraction</i>	
Liyang Xu and Jinho D. Choi . . . . .	5409
<i>Few-Shot Semantic Parsing with Language Models Trained on Code</i>	
Richard Shin and Benjamin Van Durme . . . . .	5417
<i>CORWA: A Citation-Oriented Related Work Annotation Dataset</i>	
Xiangci Li, Biswadip Mandal and Jessica Ouyang . . . . .	5426
<i>Overcoming Catastrophic Forgetting During Domain Adaptation of Seq2seq Language Generation</i>	
Dingcheng Li, Zheng Chen, Eunah Cho, Jie Hao, Xiaohu Liu, Fan Xing, Chenlei Guo and Yang Liu . . . . .	5441
<i>Extreme Zero-Shot Learning for Extreme Text Classification</i>	
Yuanhao Xiong, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu and Inderjit S Dhillon . . . . .	5455
<i>ConflIBERT: A Pre-trained Language Model for Political Conflict and Violence</i>	
Yibo Hu, MohammadSaleh Hosseini, Erick Skorupa Parolin, Javier Osorio, Latifur Khan, Patrick Brandt and Vito D’Orazio . . . . .	5469
<i>Automatic Multi-Label Prompting: Simple and Interpretable Few-Shot Classification</i>	
Han Wang, Canwen Xu and Julian McAuley . . . . .	5483
<i>Few-shot Subgoal Planning with Language Models</i>	
Lajanugen Logeswaran, Yao Fu, Moontae Lee and Honglak Lee . . . . .	5493
<i>IDPG: An Instance-Dependent Prompt Generation Method</i>	
Zhuofeng Wu, Sinong Wang, Jiatao Gu, Rui Hou, Yuxiao Dong, V.G.Vinod Vydiswaran and Hao Ma . . . . .	5507
<i>Embedding Hallucination for Few-shot Language Fine-tuning</i>	
Yiren Jian, Chongyang Gao and Soroush Vosoughi . . . . .	5522
<i>Cryptocurrency Bubble Detection: A New Stock Market Dataset, Financial Task &amp; Hyperbolic Models</i>	
Ramit Sawhney, Shivam Agarwal, Vivek Mittal, Paolo Rosso, Vikram Nanda and Sudheer Chava	5531
<i>Nearest Neighbor Knowledge Distillation for Neural Machine Translation</i>	
Zhixian Yang, Renliang Sun and Xiaojun Wan . . . . .	5546
<i>DEMIX Layers: Disentangling Domains for Modular Language Modeling</i>	
Suchin Gururangan, Mike Lewis, Ari Holtzman, Noah Smith and Luke Zettlemoyer . . . . .	5557

<i>Contrastive Learning for Prompt-based Few-shot Language Learners</i>	
Yiren Jian, Chongyang Gao and Soroush Vosoughi .....	5577
<i>Cross-Lingual Event Detection via Optimized Adversarial Training</i>	
Luis Fernando Guzman-Nateras, Minh Van Nguyen and Thien Huu Nguyen .....	5588
<i>Identifying Implicitly Abusive Remarks about Identity Groups using a Linguistically Informed Approach</i>	
Michael Wiegand, Elisabeth Eder and Josef Ruppenhofer .....	5600
<i>Label Definitions Improve Semantic Role Labeling</i>	
Li Zhang, Ishan Jindal and Yunyao Li .....	5613
<i>Shedding New Light on the Language of the Dark Web</i>	
Youngjin Jin, Eugene Jang, Yongjae Lee, Seungwon Shin and Jin-Woo Chung .....	5621
<i>Conceptualizing Treatment Leakage in Text-based Causal Inference</i>	
Adel Daoud, Connor Thomas Jerzak and Richard Johansson .....	5638
<i>Consistency Training with Virtual Adversarial Discrete Perturbation</i>	
Jungsoo Park, Gyuwan Kim and Jaewoo Kang .....	5646
<i>CONFIT: Toward Faithful Dialogue Summarization with Linguistically-Informed Contrastive Fine-tuning</i>	
Xiangru Tang, Arjun Nair, Borui Wang, Bingyao Wang, Jai Amit Desai, Aaron Wade, Haoran Li, Asli Celikyilmaz, Yashar Mehdad and Dragomir Radev .....	5657
<i>CoMPM: Context Modeling with Speaker's Pre-trained Memory Tracking for Emotion Recognition in Conversation</i>	
Joosung Lee and Woon Lee .....	5669
<i>Investigating Crowdsourcing Protocols for Evaluating the Factual Consistency of Summaries</i>	
Xiangru Tang, Alexander Fabbri, Haoran Li, Ziming Mao, Griffin Thomas Adams, Borui Wang, Asli Celikyilmaz, Yashar Mehdad and Dragomir Radev .....	5680
<i>DialSummEval: Revisiting Summarization Evaluation for Dialogues</i>	
Mingqi Gao and Xiaojun Wan .....	5693
<i>Hyperbolic Relevance Matching for Neural Keyphrase Extraction</i>	
Mingyang Song, Yi Feng and Liping Jing .....	5710
<i>Template-free Prompt Tuning for Few-shot NER</i>	
Ruotian Ma, Xin Zhou, Tao Gui, Yiding Tan, Linyang Li, Qi Zhang and Xuanjing Huang ..	5721
<i>Few-Shot Document-Level Relation Extraction</i>	
Nicholas Popovic and Michael Färber .....	5733
<i>LaMemo: Language Modeling with Look-Ahead Memory</i>	
Haozhe Ji, Rongsheng Zhang, Zhenyu Yang, Zhipeng Hu and Minlie Huang .....	5747
<i>Exploiting Inductive Bias in Transformers for Unsupervised Disentanglement of Syntax and Semantics with VAEs</i>	
Ghazi Felhi, Joseph Le Roux and Djamel Seddah .....	5763
<i>Neighbors Are Not Strangers: Improving Non-Autoregressive Translation under Low-Frequency Lexical Constraints</i>	
Chun Zeng, Jiangjie Chen, Tianyi Zhuang, Rui Xu, Hao Yang, Qin Ying, Shimin Tao and Yanghua Xiao .....	5777

<i>What do Toothbrushes do in the Kitchen? How Transformers Think our World is Structured</i> Alexander Henlein and Alexander Mehler .....	5791
<i>Less is More: Learning to Refine Dialogue History for Personalized Dialogue Generation</i> Hanxun Zhong, Zhicheng Dou, Yutao Zhu, Hongjin Qian and Ji-Rong Wen .....	5808
<i>A Holistic Framework for Analyzing the COVID-19 Vaccine Debate</i> Maria Leonor Pacheco, Tunazzina Islam, Monal Mahajan, Andrey Shor, Ming Yin, Lyle Ungar and Dan Goldwasser .....	5821
<i>Learning to Win Lottery Tickets in BERT Transfer via Task-agnostic Mask Training</i> Yuanxin Liu, Fandong Meng, Zheng Lin, Peng Fu, Yanan Cao, Weiping Wang and Jie Zhou	5840
<i>You Don't Know My Favorite Color: Preventing Dialogue Representations from Revealing Speakers' Private Personas</i> Haoran Li, Yangqiu Song and Lixin Fan .....	5858
<i>Explaining Dialogue Evaluation Metrics using Adversarial Behavioral Analysis</i> Baber Khalid and Sungjin Lee .....	5871
<i>Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection</i> Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi and Noah Smith	5884
<i>Non-Autoregressive Chinese ASR Error Correction with Phonological Training</i> Zheng Fang, Ruiqing Zhang, Zhongjun He, Hua Wu and Yanan Cao .....	5907
<i>Hate Speech and Counter Speech Detection: Conversational Context Does Matter</i> Xinchen Yu, Eduardo Blanco and Lingzi Hong .....	5918
<i>DACSA: A large-scale Dataset for Automatic summarization of Catalan and Spanish newspaper Articles</i> Encarnación Segarra Soriano, Vicent Ahuir, Lluís-F. Hurtado and José Ángel González . . .	5931
<i>Time Waits for No One! Analysis and Challenges of Temporal Misalignment</i> Kelvin Luu, Daniel Khashabi, Suchin Gururangan, Karishma Mandyam and Noah Smith . . .	5944
<i>MCSE: Multimodal Contrastive Learning of Sentence Embeddings</i> Miaoran Zhang, Marius Mosbach, David Ifeoluwa Adelani, Michael A. Hedderich and Dietrich Klakow .....	5959
<i>HiURE: Hierarchical Exemplar Contrastive Learning for Unsupervised Relation Extraction</i> Shuliang Liu, Xuming Hu, Chenwei Zhang, Shu'ang Li, Lijie Wen and Philip S. Yu .....	5970
<i>Diagnosing Vision-and-Language Navigation: What Really Matters</i> Wanrong Zhu, Yuankai Qi, Pradyumna Narayana, Kazoo Sone, Sugato Basu, Xin Eric Wang, Qi Wu, Miguel Eckstein and William Yang Wang .....	5981
<i>Aligning to Social Norms and Values in Interactive Narratives</i> Prithviraj Ammanabrolu, Liwei Jiang, Maarten Sap, Hannaneh Hajishirzi and Yejin Choi . .	5994
<i>MOVER: Mask, Over-generate and Rank for Hyperbole Generation</i> Yunxiang Zhang and Xiaojun Wan .....	6018
<i>Embarrassingly Simple Performance Prediction for Abductive Natural Language Inference</i> Emils Kadiķis, Vaibhav Srivastav and Roman Klinger .....	6031
<i>Re-Examining System-Level Correlations of Automatic Summarization Evaluation Metrics</i> Daniel Deutsch, Rotem Dror and Dan Roth .....	6038