

17th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2023)

Dubrovnik, Croatia
2-6 May 2023

Volume 1 of 5

ISBN: 978-1-7138-7375-4

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

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

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

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

Phone: 1-570-476-8006
Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

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

Table of Contents

<i>PiC: A Phrase-in-Context Dataset for Phrase Understanding and Semantic Search</i> Thang Pham, Seunghyun Yoon, Trung Bui and Anh Nguyen	1
<i>Enhancing Dialogue Summarization with Topic-Aware Global- and Local- Level Centrality</i> Xinnian Liang, Shuangzhi Wu, Chenhao Cui, Jiaqi Bai, Chao Bian and Zhoujun Li	27
<i>Exploiting Summarization Data to Help Text Simplification</i> Renliang Sun, Zhixian Yang and Xiaojun Wan	39
<i>Shironaam: Bengali News Headline Generation using Auxiliary Information</i> Abu Ubaida Akash, Mir Tafseer Nayeem, Faisal Tareque Shohan and Tanvir Islam	52
<i>PCC: Paraphrasing with Bottom-k Sampling and Cyclic Learning for Curriculum Data Augmentation</i> Hongyuan Lu and Wai Lam	68
<i>A Two-Sided Discussion of Preregistration of NLP Research</i> Anders Søgaard, Daniel Hershcovich and Miryam de Lhoneux	83
<i>WinoDict: Probing language models for in-context word acquisition</i> Julian Martin Eisenschlos, Jeremy R. Cole, Fangyu Liu and William W. Cohen	94
<i>Sentiment as an Ordinal Latent Variable</i> Niklas Stoehr, Ryan Cotterell and Aaron Schein	103
<i>Nationality Bias in Text Generation</i> Pranav Narayanan Venkit, Sanjana Gautam, Ruchi Panchanadikar, Ting-Hao Huang and Shomir Wilson	116
<i>Investigating data partitioning strategies for crosslinguistic low-resource ASR evaluation</i> Zoey Liu, Justin Spence and Emily Prud'hommeaux	123
<i>Shortcomings of Question Answering Based Factuality Frameworks for Error Localization</i> Ryo Kamoi, Tanya Goyal and Greg Durrett	132
<i>Socratic Question Generation: A Novel Dataset, Models, and Evaluation</i> Beng Heng Ang, Sujatha Das Gollapalli and See-Kiong Ng	147
<i>Do we need Label Regularization to Fine-tune Pre-trained Language Models?</i> Ivan Kobzyev, Aref Jafari, Mehdi Rezagholizadeh, Tianda Li, Alan Do-Omri, Peng Lu, Pascal Poupart and Ali Ghodsi	166
<i>COVID-VTS: Fact Extraction and Verification on Short Video Platforms</i> Fuxiao Liu, Yaser Yacoob and Abhinav Shrivastava	178
<i>Multimodal Graph Transformer for Multimodal Question Answering</i> Xuehai He and Xin Wang	189
<i>Retrieval Enhanced Data Augmentation for Question Answering on Privacy Policies</i> Md Rizwan Parvez, Jianfeng Chi, Wasii Uddin Ahmad, Yuan Tian and Kai-Wei Chang	201
<i>FastKASSIM: A Fast Tree Kernel-Based Syntactic Similarity Metric</i> Maximillian Chen, Caitlyn Chen, Xiao Yu and Zhou Yu	211
<i>Friend-training: Learning from Models of Different but Related Tasks</i> Mian Zhang, Lifeng Jin, Linfeng Song, Haitao Mi, Xiabing Zhou and Dong Yu	232

<i>Understanding Transformer Memorization Recall Through Idioms</i>	
Adi Haviv, Ido Cohen, Jacob Gidron, Roei Schuster, Yoav Goldberg and Mor Geva	248
<i>A Discerning Several Thousand Judgments: GPT-3 Rates the Article + Adjective + Numeral + Noun Construction</i>	
Kyle Mahowald	265
<i>Triple-Hybrid Energy-based Model Makes Better Calibrated Natural Language Understanding Models</i>	
Haotian Xu and Yingying Zhang	274
<i>A weakly supervised textual entailment approach to zero-shot text classification</i>	
Marc Pàmies, Joan Llop, Francesco Multari, Nicolau Duran-Silva, César Parra-Rojas, Aitor Gonzalez-Agirre, Francesco Alessandro Massucci and Marta Villegas	286
<i>Fair Enough: Standardizing Evaluation and Model Selection for Fairness Research in NLP</i>	
Xudong Han, Timothy Baldwin and Trevor Cohn	297
<i>CHARD: Clinical Health-Aware Reasoning Across Dimensions for Text Generation Models</i>	
Steven Y. Feng, Vivek Khetan, Bogdan Sacaleanu, Anatole Gershman and Eduard Hovy	313
<i>Prompt Tuning with Contradictory Intentions for Sarcasm Recognition</i>	
Yiyi Liu, Ruqing Zhang, Yixing Fan, Jiafeng Guo and Xueqi Cheng	328
<i>COMBO: A Complete Benchmark for Open KG Canonicalization</i>	
Chengyue Jiang, Yong Jiang, Weiqi Wu, Yuting Zheng, Pengjun Xie and Kewei Tu	340
<i>UScore: An Effective Approach to Fully Unsupervised Evaluation Metrics for Machine Translation</i>	
Jonas Belouadi and Steffen Eger	358
<i>Assistive Recipe Editing through Critiquing</i>	
Diego Antognini, Shuyang Li, Boi Faltings and Julian McAuley	375
<i>DiTTO: A Feature Representation Imitation Approach for Improving Cross-Lingual Transfer</i>	
Shanu Kumar, Soujanya Abbaraju, Sandipan Dandapat, Sunayana Sitaram and Monojit Choudhury	385
<i>John is 50 years old, can his son be 65? Evaluating NLP Models' Understanding of Feasibility</i>	
Himanshu Gupta, Neeraj Varshney, Swaroop Mishra, Kuntal Kumar Pal, Saurabh Arjun Sawant, Kevin Scaria, Siddharth Goyal and Chitta Baral	407
<i>Efficient Encoders for Streaming Sequence Tagging</i>	
Ayush Kaushal, Aditya Gupta, Shyam Upadhyay and Manaal Faruqui	418
<i>Retrieve-and-Fill for Scenario-based Task-Oriented Semantic Parsing</i>	
Akshat Shrivastava, Shrey Desai, Anchit Gupta, Ali Elkahky, Aleksandr Livshits, Alexander Zotov and Ahmed Aly	430
<i>Document Flattening: Beyond Concatenating Context for Document-Level Neural Machine Translation</i>	
Minghao Wu, George Foster, Lizhen Qu and Gholamreza Haffari	448
<i>Scaling Back-Translation with Domain Text Generation for Sign Language Gloss Translation</i>	
Jinhui Ye, Wenxiang Jiao, Xing Wang and Zhaopeng Tu	463
<i>Realistic Conversational Question Answering with Answer Selection based on Calibrated Confidence and Uncertainty Measurement</i>	
Soyeong Jeong, Jinheon Baek, Sung Ju Hwang and Jong Park	477

<i>PANCETTA: Phoneme Aware Neural Completion to Elicit Tongue Twisters Automatically</i> Sedrick Scott Keh, Steven Y. Feng, Varun Gangal, Malihe Alikhani and Eduard Hovy	491
<i>A User-Centered, Interactive, Human-in-the-Loop Topic Modelling System</i> Zheng Fang, Lama Alqazlan, Du Liu, Yulan He and Rob Procter	505
<i>A Survey of Methods for Addressing Class Imbalance in Deep-Learning Based Natural Language Processing</i> Sophie Henning, William Beluch, Alexander Fraser and Annemarie Friedrich	523
<i>Extracting or Guessing? Improving Faithfulness of Event Temporal Relation Extraction</i> Haoyu Wang, Hongming Zhang, Yuqian Deng, Jacob Gardner, Dan Roth and Muhao Chen	541
<i>LoFT: Enhancing Faithfulness and Diversity for Table-to-Text Generation via Logic Form Control</i> Yilun Zhao, Zhenting Qi, Linyong Nan, Lorenzo Jaime Flores and Dragomir Radev	554
<i>PromptDA: Label-guided Data Augmentation for Prompt-based Few Shot Learners</i> Canyu Chen and Kai Shu	562
<i>Incorporating Question Answering-Based Signals into Abstractive Summarization via Salient Span Selection</i> Daniel Deutsch and Dan Roth	575
<i>Patient Outcome and Zero-shot Diagnosis Prediction with Hypernetwork-guided Multitask Learning</i> Shaoyong Ji and Pekka Marttinen	589
<i>A Kind Introduction to Lexical and Grammatical Aspect, with a Survey of Computational Approaches</i> Annemarie Friedrich, Nianwen Xue and Alexis Palmer	599
<i>Incorporating Context into Subword Vocabularies</i> Shaked Yehezkel and Yuval Pinter	623
<i>LoRaLay: A Multilingual and Multimodal Dataset for Long Range and Layout-Aware Summarization</i> Laura Nguyen, Thomas Scialom, Benjamin Piwowarski and Jacopo Staiano	636
<i>ViHOS: Hate Speech Spans Detection for Vietnamese</i> Phu Gia Hoang, Canh Duc Luu, Khanh Quoc Tran, Kiet Van Nguyen and Ngan Luu-Thuy Nguyen	652
<i>Vote'n'Rank: Revision of Benchmarking with Social Choice Theory</i> Mark Rofin, Vladislav Mikhailov, Mikhail Florinsky, Andrey Kravchenko, Tatiana Shavrina, Elena Tutubalina, Daniel Karabekyan and Ekaterina Artemova	670
<i>Combining Parameter-efficient Modules for Task-level Generalisation</i> Edoardo Maria Ponti, Alessandro Sordani, Yoshua Bengio and Siva Reddy	687
<i>Self-imitation Learning for Action Generation in Text-based Games</i> Zijing Shi, Yunqiu Xu, Meng Fang and Ling Chen	703
<i>Investigating the Effect of Relative Positional Embeddings on AMR-to-Text Generation with Structural Adapters</i> Sebastien Montella, Alexis Nasr, Johannes Heinecke, Frederic Bechet and Lina M. Rojas Barahona	727
<i>On the Intersection of Context-Free and Regular Languages</i> Clemente Pasti, Andreas Opedal, Tiago Pimentel, Tim Vieira, Jason Eisner and Ryan Cotterell	737

<i>Social Influence Dialogue Systems: A Survey of Datasets and Models For Social Influence Tasks</i>	
Kushal Chawla, Weiyan Shi, Jingwen Zhang, Gale Lucas, Zhou Yu and Jonathan Gratch	750
<i>Aggregating Crowdsourced and Automatic Judgments to Scale Up a Corpus of Anaphoric Reference for Fiction and Wikipedia Texts</i>	
Juntao Yu, Silviu Paun, Maris Camilleri, Paloma Garcia, Jon Chamberlain, Udo Kruschwitz and Massimo Poesio	767
<i>What Makes Sentences Semantically Related? A Textual Relatedness Dataset and Empirical Study</i>	
Mohamed Abdalla, Krishnapriya Vishnubhotla and Saif Mohammad	782
<i>RevUp: Revise and Update Information Bottleneck for Event Representation</i>	
Mehdi Rezaee and Francis Ferraro	797
<i>NusaX: Multilingual Parallel Sentiment Dataset for 10 Indonesian Local Languages</i>	
Genta Indra Winata, Alham Fikri Aji, Samuel Cahyawijaya, Rahmad Mahendra, Fajri Koto, Ade Romadhony, Kemal Kurniawan, David Moeljadi, Radityo Eko Prasajo, Pascale Fung, Timothy Baldwin, Jey Han Lau, Rico Sennrich and Sebastian Ruder	815
<i>The Functional Relevance of Probed Information: A Case Study</i>	
Michael Hanna, Roberto Zamparelli and David Mareček	835
<i>Do Pretrained Contextual Language Models Distinguish between Hebrew Homograph Analyses?</i>	
Avi Shmidman, Cheyn Shmuel Shmidman, Dan Bareket, Moshe Koppel and Reut Tsarfaty . .	849
<i>Parameter-Efficient Tuning with Special Token Adaptation</i>	
Xiacong Yang, James Y. Huang, Wenxuan Zhou and Muhao Chen	865
<i>Probing Power by Prompting: Harnessing Pre-trained Language Models for Power Connotation Framing</i>	
Shima Khanehzar, Trevor Cohn, Gosia Mikolajczak and Lea Frermann	873
<i>Zero and Few-Shot Localization of Task-Oriented Dialogue Agents with a Distilled Representation</i>	
Mehrad Moradshahi, Sina Semnani and Monica Lam	886
<i>Contextual Semantic Parsing for Multilingual Task-Oriented Dialogues</i>	
Mehrad Moradshahi, Victoria Tsai, Giovanni Campagna and Monica Lam	902
<i>Teacher Intervention: Improving Convergence of Quantization Aware Training for Ultra-Low Precision Transformers</i>	
Minsoo Kim, Kyuhong Shim, Seongmin Park, Wonyong Sung and Jungwook Choi	916
<i>Generative Replay Inspired by Hippocampal Memory Indexing for Continual Language Learning</i>	
Aru Maekawa, Hidetaka Kamigaito, Kotaro Funakoshi and Manabu Okumura	930
<i>A Survey of Multi-task Learning in Natural Language Processing: Regarding Task Relatedness and Training Methods</i>	
Zhihan Zhang, Wenhao Yu, Mengxia Yu, Zhichun Guo and Meng Jiang	943
<i>Conclusion-based Counter-Argument Generation</i>	
Milad Alshomary and Henning Wachsmuth	957
<i>Question-Answer Sentence Graph for Joint Modeling Answer Selection</i>	
Roshni Iyer, Thuy Vu, Alessandro Moschitti and Yizhou Sun	968
<i>Evaluating and Improving the Coreference Capabilities of Machine Translation Models</i>	
Asaf Yehudai, Arie Cattan, Omri Abend and Gabriel Stanovsky	980

<i>Document-Level Planning for Text Simplification</i> Liam Cripwell, Joël Legrand and Claire Gardent	993
<i>Efficient Hybrid Generation Framework for Aspect-Based Sentiment Analysis</i> Haoran Lv, Junyi Liu, Henan Wang, Yaoming Wang, Jixiang Luo and Yaxiao Liu	1007
<i>What's New? Summarizing Contributions in Scientific Literature</i> Hiroaki Hayashi, Wojciech Kryscinski, Bryan McCann, Nazneen Rajani and Caiming Xiong	1019
<i>Find Parent then Label Children: A Two-stage Taxonomy Completion Method with Pre-trained Language Model</i> Fei Xia, Yixuan Weng, Shizhu He, Kang Liu and Jun Zhao	1032
<i>Meta Self-Refinement for Robust Learning with Weak Supervision</i> Dawei Zhu, Xiaoyu Shen, Michael Hedderich and Dietrich Klakow	1043
<i>Looking for a Needle in a Haystack: A Comprehensive Study of Hallucinations in Neural Machine Translation</i> Nuno M. Guerreiro, Elena Voita and André Martins	1059
<i>Investigating UD Treebanks via Dataset Difficulty Measures</i> Artur Kulmizev and Joakim Nivre	1076
<i>On Robustness of Prompt-based Semantic Parsing with Large Pre-trained Language Model: An Empirical Study on Codex</i> Terry Yue Zhuo, Zhuang Li, Yujin Huang, Fatemeh Shiri, Weiqing Wang, Gholamreza Haffari and Yuan-Fang Li	1090
<i>Leveraging Task Dependency and Contrastive Learning for Case Outcome Classification on European Court of Human Rights Cases</i> Santosh T.Y.S.S, Marcel Perez San Blas, Phillip Kemper and Matthias Grabmair	1103
<i>Semi-supervised Relation Extraction via Data Augmentation and Consistency-training</i> Komal Teru	1112
<i>Event Temporal Relation Extraction with Bayesian Translational Model</i> Xingwei Tan, Gabriele Pergola and Yulan He	1125
<i>Persona Expansion with Commonsense Knowledge for Diverse and Consistent Response Generation</i> Donghyun Kim, Youbin Ahn, Wongyu Kim, Chanhee Lee, Kyungchan Lee, Kyong-Ho Lee, Jeonguk Kim, Donghoon Shin and Yeonsoo Lee	1139
<i>UnifEE: Unified Evidence Extraction for Fact Verification</i> Nan Hu, Zirui Wu, Yuxuan Lai, Chen Zhang and Yansong Feng	1150
<i>MiniALBERT: Model Distillation via Parameter-Efficient Recursive Transformers</i> Mohammadmahdi Nouriborji, Omid Rohanian, Samaneh Kouchaki and David A. Clifton ..	1161
<i>Multilingual Normalization of Temporal Expressions with Masked Language Models</i> Lukas Lange, Jannik Strötgen, Heike Adel and Dietrich Klakow	1174
<i>K-hop neighbourhood regularization for few-shot learning on graphs: A case study of text classification</i> Niels van der Heijden, Ekaterina Shutova and Helen Yannakoudakis	1187
<i>What Clued the AI Doctor In? On the Influence of Data Source and Quality for Transformer-Based Medical Self-Disclosure Detection</i> Mina Valizadeh, Xing Qian, Pardis Ranjbar-Noiey, Cornelia Caragea and Natalie Parde . . .	1201

<i>Improving Visual-Semantic Embedding with Adaptive Pooling and Optimization Objective</i> Zijian Zhang, Chang Shu, Ya Xiao, Yuan Shen, Di Zhu, Youxin Chen, Jing Xiao, Jey Han Lau, Qian Zhang and Zheng Lu	1217
<i>Policy-based Reinforcement Learning for Generalisation in Interactive Text-based Environments</i> Edan Toledo, Jan Buys and Jonathan Shock	1230
<i>Logic Against Bias: Textual Entailment Mitigates Stereotypical Sentence Reasoning</i> Hongyin Luo and James Glass	1243
<i>Entity Tracking via Effective Use of Multi-Task Learning Model and Mention-guided Decoding</i> Janvijay Singh, Fan Bai and Zhen Wang	1255
<i>Conversational Tree Search: A New Hybrid Dialog Task</i> Dirk Vath, Lindsey Vanderlyn and Ngoc Thang Vu	1264
<i>A Human Subject Study of Named Entity Recognition in Conversational Music Recommendation Queries</i> Elena Epure and Romain Hennequin	1281
<i>Entity Disambiguation with Entity Definitions</i> Luigi Procopio, Simone Conia, Edoardo Barba and Roberto Navigli	1297
<i>Exploring Paracrawl for Document-level Neural Machine Translation</i> Yusser Al Ghussin, Jingyi Zhang and Josef van Genabith	1304
<i>Poor Man’s Quality Estimation: Predicting Reference-Based MT Metrics Without the Reference</i> Vilem Zouhar, Shehzaad Dhuliawala, Wangchunshu Zhou, Nico Daheim, Tom Kocmi, Yuchen Eleanor Jiang and Mrinmaya Sachan	1311
<i>Integrating Translation Memories into Non-Autoregressive Machine Translation</i> Jitao Xu, Josep Crego and Franois Yvon	1326
<i>Shorten the Long Tail for Rare Entity and Event Extraction</i> Pengfei Yu and Heng Ji	1339
<i>Do Deep Neural Networks Capture Compositionality in Arithmetic Reasoning?</i> Keito Kudo, Yoichi Aoki, Tatsuki Kuribayashi, Ana Brassard, Masashi Yoshikawa, Keisuke Sa- kaguchi and Kentaro Inui	1351
<i>BLM-AgrF: A New French Benchmark to Investigate Generalization of Agreement in Neural Networks</i> Aixiu An, Chunyang Jiang, Maria A. Rodriguez, Vivi Nastase and Paola Merlo	1363
<i>Robustification of Multilingual Language Models to Real-world Noise in Crosslingual Zero-shot Set- tings with Robust Contrastive Pretraining</i> Asa Cooper Stickland, Sailik Sengupta, Jason Krone, Saab Mansour and He He	1375
<i>Unsupervised Anomaly Detection in Multi-Topic Short-Text Corpora</i> Mira Ait-Saada and Mohamed Nadif	1392
<i>Metaphor Detection with Effective Context Denoising</i> Shun Wang, Yucheng Li, Chenghua Lin, Loic Barrault and Frank Guerin	1404
<i>Low-Resource Compositional Semantic Parsing with Concept Pretraining</i> Subendhu Rongali, Mukund Sridhar, Haidar Khan, Konstantine Arkoudas, Wael Hamza and An- drew McCallum	1410

<i>Made of Steel? Learning Plausible Materials for Components in the Vehicle Repair Domain</i> Annerose Eichel, Helena Schlipf and Sabine Schulte im Walde	1420
<i>Self-Adapted Utterance Selection for Suicidal Ideation Detection in Lifeline Conversations</i> Zhong-Ling Wang, Po-Hsien Huang, Wen-Yau Hsu and Hen-Hsen Huang	1436
<i>Can Pretrained Language Models (Yet) Reason Deductively?</i> Zhangdie Yuan, Songbo Hu, Ivan Vulić, Anna Korhonen and Zaiqiao Meng	1447
<i>Selective In-Context Data Augmentation for Intent Detection using Pointwise V-Information</i> Yen-Ting Lin, Alexandros Papangelis, Seokhwan Kim, Sungjin Lee, Devamanyu Hazarika, Mahdi Namazifar, Di Jin, Yang Liu and Dilek Hakkani-Tur	1463
<i>Multilingual Representation Distillation with Contrastive Learning</i> Weiting Tan, Kevin Heffernan, Holger Schwenk and Philipp Koehn	1477
<i>On the inconsistency of separable losses for structured prediction</i> Caio Corro	1491
<i>A Systematic Search for Compound Semantics in Pretrained BERT Architectures</i> Filip Miletic and Sabine Schulte im Walde	1499
<i>Efficiently Upgrading Multilingual Machine Translation Models to Support More Languages</i> Simeng Sun, Maha Elbayad, Anna Sun and James Cross	1513
<i>Summarize and Generate to Back-translate: Unsupervised Translation of Programming Languages</i> Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray and Kai-Wei Chang	1528
<i>The Impacts of Unanswerable Questions on the Robustness of Machine Reading Comprehension Models</i> Son Quoc Tran, Phong Nguyen-Thuan Do, Uyen Le and Matt Kretchmar	1543
<i>FrameBERT: Conceptual Metaphor Detection with Frame Embedding Learning</i> Yucheng Li, Shun Wang, Chenghua Lin, Frank Guerin and Loic Barrault	1558
<i>Towards More Efficient Insertion Transformer with Fractional Positional Encoding</i> Zhisong Zhang, Yizhe Zhang and Bill Dolan	1564
<i>SODAPOP: Open-Ended Discovery of Social Biases in Social Commonsense Reasoning Models</i> Haozhe An, Zongxia Li, Jieyu Zhao and Rachel Rudinger	1573
<i>Augmenting Pre-trained Language Models with QA-Memory for Open-Domain Question Answering</i> Wenhu Chen, Pat Verga, Michiel de Jong, John Wieting and William W. Cohen	1597
<i>Gold Doesn't Always Glitter: Spectral Removal of Linear and Nonlinear Guarded Attribute Information</i> Shun Shao, Yftah Ziser and Shay B. Cohen	1611
<i>CTC Alignments Improve Autoregressive Translation</i> Brian Yan, Siddharth Dalmia, Yosuke Higuchi, Graham Neubig, Florian Metze, Alan W Black and Shinji Watanabe	1623
<i>Modelling Temporal Document Sequences for Clinical ICD Coding</i> Boon Liang Clarence Ng, Diogo Santos and Marek Rei	1640
<i>LongEval: Guidelines for Human Evaluation of Faithfulness in Long-form Summarization</i> Kalpesh Krishna, Erin Bransom, Bailey Kuehl, Mohit Iyyer, Pradeep Dasigi, Arman Cohan and Kyle Lo	1650

<i>Cluster-Guided Label Generation in Extreme Multi-Label Classification</i> Taehee Jung, Joo-kyung Kim, Sungjin Lee and Dongyeop Kang	1670
<i>Empathy Identification Systems are not Accurately Accounting for Context</i> Andrew Lee, Jonathan Kummerfeld, Larry An and Rada Mihalcea	1686
<i>Enhancing Multi-Document Summarization with Cross-Document Graph-based Information Extraction</i> Zixuan Zhang, Heba Elfardy, Markus Dreyer, Kevin Small, Heng Ji and Mohit Bansal	1696
<i>What happens before and after: Multi-Event Commonsense in Event Coreference Resolution</i> Sahithya Ravi, Chris Tanner, Raymond Ng and Vered Shwartz	1708
<i>Multi-Modal Bias: Introducing a Framework for Stereotypical Bias Assessment beyond Gender and Race in Vision–Language Models</i> Sepehr Janghorbani and Gerard De Melo	1725
<i>CyLE: Cylinder Embeddings for Multi-hop Reasoning over Knowledge Graphs</i> Chau Duc Minh Nguyen, Tim French, Wei Liu and Michael Stewart	1736
<i>Fiction-Writing Mode: An Effective Control for Human-Machine Collaborative Writing</i> Wenjie Zhong, Jason Naradowsky, Hiroya Takamura, Ichiro Kobayashi and Yusuke Miyao .	1752
<i>Robustness Challenges in Model Distillation and Pruning for Natural Language Understanding</i> Mengnan Du, Subhabrata Mukherjee, Yu Cheng, Milad Shokouhi, Xia Hu and Ahmed Hassan Awadallah	1766
<i>Don't Blame the Annotator: Bias Already Starts in the Annotation Instructions</i> Mihir Parmar, Swaroop Mishra, Mor Geva and Chitta Baral	1779
<i>Performance Prediction via Bayesian Matrix Factorisation for Multilingual Natural Language Processing Tasks</i> Viktoria Schram, Daniel Beck and Trevor Cohn	1790
<i>Unified Neural Topic Model via Contrastive Learning and Term Weighting</i> Sungwon Han, Mingi Shin, Sungkyu Park, Changwook Jung and Meeyoung Cha	1802
<i>Don't Mess with Mister-in-Between: Improved Negative Search for Knowledge Graph Completion</i> Fan Jiang, Tom Drummond and Trevor Cohn	1818
<i>Semantic Frame Induction with Deep Metric Learning</i> Kosuke Yamada, Ryohei Sasano and Koichi Takeda	1833
<i>The Devil is in the Details: On Models and Training Regimes for Few-Shot Intent Classification</i> Mohsen Mesgar, Thy Thy Tran, Goran Glavaš and Iryna Gurevych	1846
<i>Iterative Document-level Information Extraction via Imitation Learning</i> Yunmo Chen, William Gantt, Weiwei Gu, Tongfei Chen, Aaron White and Benjamin Van Durme	1858
<i>CLICK: Contrastive Learning for Injecting Contextual Knowledge to Conversational Recommender System</i> Hyeongjun Yang, Heesoo Won, Youbin Ahn and Kyong-Ho Lee	1875
<i>LEALLA: Learning Lightweight Language-agnostic Sentence Embeddings with Knowledge Distillation</i> Zhuoyuan Mao and Tetsuji Nakagawa	1886
<i>Synthesizing Human Gaze Feedback for Improved NLP Performance</i> Varun Khurana, Yaman Kumar, Nora Hollenstein, Rajesh Kumar and Balaji Krishnamurthy	1895

<i>Memory-efficient Temporal Moment Localization in Long Videos</i> Cristian Rodriguez-Opazo, Edison Marrese-Taylor, Basura Fernando, Hiroya Takamura and Qi Wu	1909
<i>Extracting Victim Counts from Text</i> Mian Zhong, Shehzaad Dhuliawala and Niklas Stoehr	1925
<i>ConEntail: An Entailment-based Framework for Universal Zero and Few Shot Classification with Supervised Contrastive Pretraining</i> Ranran Haoran Zhang, Aysa Xuemo Fan and Rui Zhang	1941
<i>Guide the Learner: Controlling Product of Experts Debiasing Method Based on Token Attribution Similarities</i> Ali Modarressi, Hossein Amirkhani and Mohammad Taher Pilehvar	1954
<i>Task and Sentiment Adaptation for Appraisal Tagging</i> Lin Tian, Xiuzhen Zhang, Myung Hee Kim and Jennifer Biggs	1960
<i>DREEM: Guiding Attention with Evidence for Improving Document-Level Relation Extraction</i> Youmi Ma, An Wang and Naoaki Okazaki	1971
<i>Span-based Named Entity Recognition by Generating and Compressing Information</i> Nhung T. H. Nguyen, Makoto Miwa and Sophia Ananiadou	1984
<i>An In-depth Analysis of Implicit and Subtle Hate Speech Messages</i> Nicolas Ocampo, Ekaterina Sviridova, Elena Cabrio and Serena Villata	1997
<i>MTEB: Massive Text Embedding Benchmark</i> Niklas Muennighoff, Nouamane Tazi, Loic Magne and Nils Reimers	2014
<i>Step by Step Loss Goes Very Far: Multi-Step Quantization for Adversarial Text Attacks</i> Piotr Gaiński and Klaudia Bałazy	2038
<i>TwirGCN: Temporally Weighted Graph Convolution for Question Answering over Temporal Knowledge Graphs</i> Aditya Sharma, Apoorv Saxena, Chitrang Gupta, Mehran Kazemi, Partha Talukdar and Soumen Chakrabarti	2049
<i>ZELDA: A Comprehensive Benchmark for Supervised Entity Disambiguation</i> Marcel Milich and Alan Akbik	2061
<i>GLADIS: A General and Large Acronym Disambiguation Benchmark</i> Lihu Chen, Gael Varoquaux and Fabian M. Suchanek	2073
<i>Probing Cross-Lingual Lexical Knowledge from Multilingual Sentence Encoders</i> Ivan Vulić, Goran Glavaš, Fangyu Liu, Nigel Collier, Edoardo Maria Ponti and Anna Korhonen	2089
<i>Pento-DIARef: A Diagnostic Dataset for Learning the Incremental Algorithm for Referring Expression Generation from Examples</i> Philipp Sadler and David Schlangen	2106
<i>Mitigating Exposure Bias in Grammatical Error Correction with Data Augmentation and Reweighting</i> Hannan Cao, Wenmian Yang and Hwee Tou Ng	2123
<i>Plausible May Not Be Faithful: Probing Object Hallucination in Vision-Language Pre-training</i> Wenliang Dai, Zihan Liu, Ziwei Ji, Dan Su and Pascale Fung	2136

<i>Characterizing the Entities in Harmful Memes: Who is the Hero, the Villain, the Victim?</i>	
Shivam Sharma, Atharva Kulkarni, Tharun Suresh, Himanshi Mathur, Preslav Nakov, Md. Shad Akhtar and Tanmoy Chakraborty	2149
<i>Systematic Investigation of Strategies Tailored for Low-Resource Settings for Low-Resource Dependency Parsing</i>	
Jivnesh Sandhan, Laxmidhar Behera and Pawan Goyal	2164
<i>Compositional Generalisation with Structured Reordering and Fertility Layers</i>	
Matthias Lindemann, Alexander Koller and Ivan Titov	2172
<i>Investigating Multi-source Active Learning for Natural Language Inference</i>	
Ard Snijders, Douwe Kiela and Katerina Margatina	2187
<i>Towards a Unified Multi-Domain Multilingual Named Entity Recognition Model</i>	
Mayank Kulkarni, Daniel Preotiuc-Pietro, Karthik Radhakrishnan, Genta Indra Winata, Shijie Wu, Lingjue Xie and Shaohua Yang	2210
<i>Do Neural Topic Models Really Need Dropout? Analysis of the Effect of Dropout in Topic Modeling</i>	
Suman Adhya, Avishek Lahiri and Debarshi Kumar Sanyal	2220
<i>A Psycholinguistic Analysis of BERT's Representations of Compounds</i>	
Lars Buijtelaar and Sandro Pezzelle	2230
<i>Measuring Normative and Descriptive Biases in Language Models Using Census Data</i>	
Samia Touileb, Lilja Øvrelid and Erik Velldal	2242
<i>UDAPTER - Efficient Domain Adaptation Using Adapters</i>	
Bhavivyva Malik, Abhinav Ramesh Kashyap, Min-Yen Kan and Soujanya Poria	2249
<i>Efficient CTC Regularization via Coarse Labels for End-to-End Speech Translation</i>	
Biao Zhang, Barry Haddow and Rico Sennrich	2264
<i>Exploring Category Structure with Contextual Language Models and Lexical Semantic Networks</i>	
Joseph Renner, Pascal Denis, Remi Gilleron and Angèle Brunellière	2277
<i>An Empirical Study of Clinical Note Generation from Doctor-Patient Encounters</i>	
Asma Ben Abacha, Wen-wai Yim, Yadan Fan and Thomas Lin	2291
<i>Instruction Clarification Requests in Multimodal Collaborative Dialogue Games: Tasks, and an Analysis of the CoDraw Dataset</i>	
Brielen Madureira and David Schlangen	2303
<i>Can Synthetic Text Help Clinical Named Entity Recognition? A Study of Electronic Health Records in French</i>	
Nicolas Hiebel, Olivier Ferret, Karen Fort and Aurélie Névéal	2320
<i>IRMA: the 335-million-word Italian coRpus for studying MisinformAtion</i>	
Fabio Carrella, Alessandro Miani and Stephan Lewandowsky	2339
<i>Parameter-Efficient Korean Character-Level Language Modeling</i>	
Marco Cогnetta, Sangwhan Moon, Lawrence Wolf-sonkin and Naoaki Okazaki	2350
<i>Opportunities and Challenges in Neural Dialog Tutoring</i>	
Jakub Macina, Nico Daheim, Lingzhi Wang, Tanmay Sinha, Manu Kapur, Iryna Gurevych and Mrinmaya Sachan	2357

<i>Evaluating the Robustness of Discrete Prompts</i>	Yoichi Ishibashi, Danushka Bollegala, Katsuhito Sudoh and Satoshi Nakamura	2373
<i>Assessing Out-of-Domain Language Model Performance from Few Examples</i>	Prasann Singhal, Jarad Forristal, Xi Ye and Greg Durrett	2385
<i>Mind the Labels: Describing Relations in Knowledge Graphs With Pretrained Models</i>	Zdeněk Kasner, Ioannis Konstas and Ondrej Dusek	2398
<i>Shapley Head Pruning: Identifying and Removing Interference in Multilingual Transformers</i>	William Held and Diyi Yang	2416
<i>Why Don't You Do It Right? Analysing Annotators' Disagreement in Subjective Tasks</i>	Marta Sandri, Elisa Leonardelli, Sara Tonelli and Elisabetta Jezek	2428
<i>Analyzing Challenges in Neural Machine Translation for Software Localization</i>	Sai Koneru, Matthias Huck, Miriam Exel and Jan Niehues	2442
<i>Bootstrapping Multilingual Semantic Parsers using Large Language Models</i>	Abhijeet Awasthi, Nitish Gupta, Bidisha Samanta, Shachi Dave, Sunita Sarawagi and Partha Talukdar	2455
<i>Modeling Complex Event Scenarios via Simple Entity-focused Questions</i>	Mahnaz Koupaee, Greg Durrett, Nathanael Chambers and Niranjan Balasubramanian	2468
<i>Uncovering Implicit Inferences for Improved Relational Argument Mining</i>	Ameer Saadat-Yazdi, Jeff Z. Pan and Nadin Kokciyan	2484
<i>How people talk about each other: Modeling Generalized Intergroup Bias and Emotion</i>	Venkata Subrahmanyam Govindarajan, Katherine Atwell, Barea Sinno, Malihe Alikhani, David I. Beaver and Junyi Jessy Li	2496
<i>Semantic Parsing for Conversational Question Answering over Knowledge Graphs</i>	Laura Perez-Beltrachini, Parag Jain, Emilio Monti and Mirella Lapata	2507
<i>MAPL: Parameter-Efficient Adaptation of Unimodal Pre-Trained Models for Vision-Language Few-Shot Prompting</i>	Oscar Mañas, Pau Rodriguez Lopez, Saba Ahmadi, Aida Nematzadeh, Yash Goyal and Aishwarya Agrawal	2523
<i>ComSearch: Equation Searching with Combinatorial Strategy for Solving Math Word Problems with Weak Supervision</i>	Qianying Liu, Wenyu Guan, Jianhao Shen, Fei Cheng and Sadao Kurohashi	2549
<i>Towards preserving word order importance through Forced Invalidation</i>	Hadeel Al-Negheimish, Pranava Madhyastha and Alessandra Russo	2563
<i>How Many and Which Training Points Would Need to be Removed to Flip this Prediction?</i>	Jinghan Yang, Sarthak Jain and Byron C. Wallace	2571
<i>Reinforced Sequence Training based Subjective Bias Correction</i>	Karthic Madanagopal and James Caverlee	2585
<i>Detecting Lexical Borrowings from Dominant Languages in Multilingual Wordlists</i>	John E. Miller and Johann-Mattis List	2599
<i>Towards Integration of Discriminability and Robustness for Document-Level Relation Extraction</i>	Jia Guo, Stanley Kok and Lidong Bing	2606

<i>Penguins Don't Fly: Reasoning about Generics through Instantiations and Exceptions</i> Emily Allaway, Jena D. Hwang, Chandra Bhagavatula, Kathleen McKeown, Doug Downey and Yejin Choi	2618
<i>Adding Instructions during Pretraining: Effective way of Controlling Toxicity in Language Models</i> Shrimai Prabhunoye, Mostofa Patwary, Mohammad Shoeybi and Bryan Catanzaro	2636
<i>Multi2Claim: Generating Scientific Claims from Multi-Choice Questions for Scientific Fact-Checking</i> Neset Tan, Trung Nguyen, Josh Bensemann, Alex Peng, Qiming Bao, Yang Chen, Mark Gahegan and Michael Witbrock	2652
<i>On Evaluation of Document Classifiers using RVL-CDIP</i> Stefan Larson, Gordon Lim and Kevin Leach	2665
<i>Event Linking: Grounding Event Mentions to Wikipedia</i> Xiaodong Yu, Wenpeng Yin, Nitish Gupta and Dan Roth	2679
<i>SwitchPrompt: Learning Domain-Specific Gated Soft Prompts for Classification in Low-Resource Domains</i> Koustava Goswami, Lukas Lange, Jun Araki and Heike Adel	2689
<i>Do dialogue representations align with perception? An empirical study</i> Sarenne Wallbridge, Peter Bell and Catherine Lai	2696
<i>Methods for Measuring, Updating, and Visualizing Factual Beliefs in Language Models</i> Peter Hase, Mona Diab, Asli Celikyilmaz, Xian Li, Zornitsa Kozareva, Veselin Stoyanov, Mohit Bansal and Srinivasan Iyer	2714
<i>Improving Sign Recognition with Phonology</i> Lee Kezar, Jesse Thomason and Zed Sehyr	2732
<i>Parameter-efficient Modularised Bias Mitigation via AdapterFusion</i> Deepak Kumar, Oleg Lesota, George Zerveas, Daniel Cohen, Carsten Eickhoff, Markus Schedl and Navid Rekabsaz	2738
<i>LingMess: Linguistically Informed Multi Expert Scorers for Coreference Resolution</i> Shon Otmazgin, Arie Cattan and Yoav Goldberg	2752
<i>Finding the Law: Enhancing Statutory Article Retrieval via Graph Neural Networks</i> Antoine Louis, Gijs van Dijck and Gerasimos Spanakis	2761
<i>Behavior Cloned Transformers are Neurosymbolic Reasoners</i> Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté and Prithviraj Ammanabrolu	2777
<i>Bridging the Gap Between BabelNet and HowNet: Unsupervised Sense Alignment and Sememe Prediction</i> Xiang Zhang, Ning Shi, Bradley Hauer and Grzegorz Kondrak	2789
<i>The StatCan Dialogue Dataset: Retrieving Data Tables through Conversations with Genuine Intent</i> Xing Han Lu, Siva Reddy and Harm de Vries	2799
<i>Question Generation Using Sequence-to-Sequence Model with Semantic Role Labels</i> Alireza Naeiji, Aijun An, Heidar Davoudi, Marjan Delpisheh and Muath Alzghool	2830
<i>StyLEX: Explaining Style Using Human Lexical Annotations</i> Shirley Anugrah Hayati, Kyumin Park, Dheeraj Rajagopal, Lyle Ungar and Dongyeop Kang	2843

<i>Comparing Intrinsic Gender Bias Evaluation Measures without using Human Annotated Examples</i> Masahiro Kaneko, Danushka Bollegala and Naoaki Okazaki	2857
<i>Faithfulness-Aware Decoding Strategies for Abstractive Summarization</i> David Wan, Mengwen Liu, Kathleen McKeown, Markus Dreyer and Mohit Bansal	2864
<i>Dynamic Benchmarking of Masked Language Models on Temporal Concept Drift with Multiple Views</i> Katerina Margatina, Shuai Wang, Yogarshi Vyas, Neha Anna John, Yassine Benajiba and Miguel Ballesteros	2881
<i>Real-Time Visual Feedback to Guide Benchmark Creation: A Human-and-Metric-in-the-Loop Workflow</i> Anjana Arunkumar, Swaroop Mishra, Bhavdeep Singh Sachdeva, Chitta Baral and Chris Bryan	2899
<i>COMPS: Conceptual Minimal Pair Sentences for testing Robust Property Knowledge and its Inheritance in Pre-trained Language Models</i> Kanishka Misra, Julia Rayz and Allyson Ettinger	2928
<i>Probabilistic Robustness for Data Filtering</i> Yu Yu, Abdul Rafae Khan, Shahram Khadivi and Jia Xu	2950
<i>Unsupervised Improvement of Factual Knowledge in Language Models</i> Nafis Sadeq, Byungkyu Kang, Prarit Lamba and Julian McAuley	2960
<i>Learning to Ignore Adversarial Attacks</i> Yiming Zhang, Yangqiaoyu Zhou, Samuel Carton and Chenhao Tan	2970
<i>Should You Mask 15% in Masked Language Modeling?</i> Alexander Wettig, Tianyu Gao, Zexuan Zhong and Danqi Chen	2985
<i>How do Words Contribute to Sentence Semantics? Revisiting Sentence Embeddings with a Perturbation Method</i> Wenlin Yao, Lifeng Jin, Hongming Zhang, Xiaoman Pan, Kaiqiang Song, Dian Yu, Dong Yu and Jianshu Chen	3001
<i>AutoTriggER: Label-Efficient and Robust Named Entity Recognition with Auxiliary Trigger Extraction</i> Dong-Ho Lee, Ravi Kiran Selvam, Sheikh Muhammad Sarwar, Bill Yuchen Lin, Fred Morstatter, Jay Pujara, Elizabeth Boschee, James Allan and Xiang Ren	3011
<i>Incorporating Task-Specific Concept Knowledge into Script Learning</i> Chenkai Sun, Tie Xu, ChengXiang Zhai and Heng Ji	3026
<i>DeepMaven: Deep Question Answering on Long-Distance Movie/TV Show Videos with Multimedia Knowledge Extraction and Synthesis</i> Yi Fung, Han Wang, Tong Wang, Ali Kebarighotbi, Mohit Bansal, Heng Ji and Prem Natarajan	3041
<i>Salient Span Masking for Temporal Understanding</i> Jeremy R. Cole, Aditi Chaudhary, Bhuwan Dhingra and Partha Talukdar	3052
<i>PECO: Examining Single Sentence Label Leakage in Natural Language Inference Datasets through Progressive Evaluation of Cluster Outliers</i> Michael Saxon, Xinyi Wang, Wenda Xu and William Yang Wang	3061
<i>Weakly-Supervised Questions for Zero-Shot Relation Extraction</i> Saeed Najafi and Alona Fyshe	3075

<i>DiffQG: Generating Questions to Summarize Factual Changes</i> Jeremy R. Cole, Palak Jain, Julian Martin Eisenschlos, Michael J.Q. Zhang, Eunsol Choi and Bhuwan Dhingra	3088
<i>Contextual Dynamic Prompting for Response Generation in Task-oriented Dialog Systems</i> Sandesh Swamy, Narges Tabari, Chacha Chen and Rashmi Gangadharaiah	3102
<i>Why Can't Discourse Parsing Generalize? A Thorough Investigation of the Impact of Data Diversity</i> Yang Janet Liu and Amir Zeldes	3112
<i>Enriching Biomedical Knowledge for Low-resource Language Through Large-scale Translation</i> Long Phan, Tai Dang, Hieu Tran, Trieu H. Trinh, Vy Phan, Lam D. Chau and Minh-Thang Luong	3131
<i>Syntax-guided Neural Module Distillation to Probe Compositionality in Sentence Embeddings</i> Rohan Pandey	3143
<i>Closed-book Question Generation via Contrastive Learning</i> Xiangjue Dong, Jiaying Lu, Jianling Wang and James Caverlee	3150
<i>A Hybrid Detection and Generation Framework with Separate Encoders for Event Extraction</i> Ge Shi, Yunyue Su, Yongliang Ma and Ming Zhou	3163
<i>Path Spuriousness-aware Reinforcement Learning for Multi-Hop Knowledge Graph Reasoning</i> Chunyang Jiang, Tianchen Zhu, Haoyi Zhou, Chang Liu, Ting Deng, Chunming Hu and Jianxin Li	3181
<i>Self-Adaptive Named Entity Recognition by Retrieving Unstructured Knowledge</i> Kosuke Nishida, Naoki Yoshinaga and Kyosuke Nishida	3193
<i>When Do Pre-Training Biases Propagate to Downstream Tasks? A Case Study in Text Summarization</i> Faisal Ladhak, Esin Durmus, Mirac Suzgun, Tianyi Zhang, Dan Jurafsky, Kathleen McKeown and Tatsunori Hashimoto	3206
<i>BERT Shows Garden Path Effects</i> Tovah Irwin, Kyra Wilson and Alec Marantz	3220
<i>Models Teaching Models: Improving Model Accuracy with Slingshot Learning</i> Lachlan O'Neill, Nandini Anantharama, Satya Borgohain and Simon D. Angus	3233
<i>A Federated Approach for Hate Speech Detection</i> Jay Gala, Deep Gandhi, Jash Mehta and Zeerak Talat	3248
<i>Learning the Legibility of Visual Text Perturbations</i> Dev Seth, Rickard Stureborg, Danish Pruthi and Bhuwan Dhingra	3260
<i>DyLoRA: Parameter-Efficient Tuning of Pre-trained Models using Dynamic Search-Free Low-Rank Adaptation</i> Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan Kobyzev and Ali Ghodsi	3274
<i>Conversational Emotion-Cause Pair Extraction with Guided Mixture of Experts</i> DongJin Jeong and JinYeong Bak	3288
<i>Language Generation Models Can Cause Harm: So What Can We Do About It? An Actionable Survey</i> Sachin Kumar, Vidhisha Balachandran, Lucille Njoo, Antonios Anastasopoulos and Yulia Tsvetkov	3299

<i>TraVLR: Now You See It, Now You Don't! A Bimodal Dataset for Evaluating Visio-Linguistic Reasoning</i> Keng Ji Chow, Samson Tan and Min-Yen Kan	3322
<i>Paraphrase Acquisition from Image Captions</i> Marcel Gohsen, Matthias Hagen, Martin Potthast and Benno Stein	3348
<i>Generation-Based Data Augmentation for Offensive Language Detection: Is It Worth It?</i> Camilla Casula and Sara Tonelli	3359
<i>Quantifying Context Mixing in Transformers</i> Hosein Mohebbi, Willem Zuidema, Grzegorz Chrupała and Afra Alishahi	3378
<i>KGVL-BART: Knowledge Graph Augmented Visual Language BART for Radiology Report Generation</i> Kaveri Kale, Pushpak Bhattacharyya, Milind Gune, Aditya Shetty and Rustom Lawyer	3401
<i>A simple but effective model for attachment in discourse parsing with multi-task learning for relation labeling</i> Zineb Bennis, Julie Hunter and Nicholas Asher	3412
<i>How Far Can It Go? On Intrinsic Gender Bias Mitigation for Text Classification</i> Ewoenam Kwaku Tokpo, Pieter Delobelle, Bettina Berendt and Toon Calders	3418
<i>Multimodal Event Transformer for Image-guided Story Ending Generation</i> Yucheng Zhou and Guodong Long	3434
<i>Improving Cross-modal Alignment for Text-Guided Image Inpainting</i> Yucheng Zhou and Guodong Long	3445
<i>Semantic Specialization for Knowledge-based Word Sense Disambiguation</i> Sakae Mizuki and Naoaki Okazaki	3457
<i>Concept-based Persona Expansion for Improving Diversity of Persona-Grounded Dialogue</i> Donghyun Kim, Youbin Ahn, Chanhee Lee, Wongyu Kim, Kyong-Ho Lee, Donghoon Shin and Yeonsoo Lee	3471
<i>RPTCS: A Reinforced Persona-aware Topic-guiding Conversational System</i> Zishan Ahmad, Kshitij Mishra, Asif Ekbal and Pushpak Bhattacharyya	3482
<i>What Did You Learn To Hate? A Topic-Oriented Analysis of Generalization in Hate Speech Detection</i> Tom Bourgeade, Patricia Chiril, Farah Benamara and Véronique Moriceau	3495
<i>End-to-end Case-Based Reasoning for Commonsense Knowledge Base Completion</i> Zonglin Yang, Xinya Du, Erik Cambria and Claire Cardie	3509
<i>Exploring Segmentation Approaches for Neural Machine Translation of Code-Switched Egyptian Arabic-English Text</i> Marwa Gaser, Manuel Mager, Injy Hamed, Nizar Habash, Slim Abdennadher and Ngoc Thang Vu	3523
<i>Identifying the limits of transformers when performing model-checking with natural language</i> Tharindu Madusanka, Riza Batista-navarro and Ian Pratt-hartmann	3539
<i>Improving the Generalizability of Collaborative Dialogue Analysis With Multi-Feature Embeddings</i> Ayesha Enayet and Gita Sukthankar	3551
<i>MetaQA: Combining Expert Agents for Multi-Skill Question Answering</i> Haritz Puerto, Gözde Şahin and Iryna Gurevych	3566

<i>BERT Is Not The Count: Learning to Match Mathematical Statements with Proofs</i>	
Weixian Waylon Li, Yftah Ziser, Maximin Coavoux and Shay B. Cohen	3581
<i>Lessons Learned from a Citizen Science Project for Natural Language Processing</i>	
Jan-Christoph Klie, Ji-Ung Lee, Kevin Stowe, Gözde Şahin, Nafise Sadat Moosavi, Luke Bates, Dominic Petrak, Richard Eckart De Castilho and Iryna Gurevych	3594
<i>Contrastive Learning with Keyword-based Data Augmentation for Code Search and Code Question Answering</i>	
Shinwoo Park, Youngwook Kim and Yo-Sub Han	3609
<i>Large Scale Multi-Lingual Multi-Modal Summarization Dataset</i>	
Yash Verma, Anubhav Jangra, Raghvendra Verma and Sriparna Saha	3620
<i>External Knowledge Acquisition for End-to-End Document-Oriented Dialog Systems</i>	
Tuan M. Lai, Giuseppe Castellucci, Saar Kuzi, Heng Ji and Oleg Rokhlenko	3633
<i>In-Depth Look at Word Filling Societal Bias Measures</i>	
Matúš Pikuliak, Ivana Beňová and Viktor Bachratý	3648
<i>Retrieval-augmented Image Captioning</i>	
Rita Ramos, Desmond Elliott and Bruno Martins	3666
<i>Automatic Evaluation and Analysis of Idioms in Neural Machine Translation</i>	
Christos Baziotis, Prashant Mathur and Eva Hasler	3682
<i>Representation biases in sentence transformers</i>	
Dmitry Nikolaev and Sebastian Padó	3701
<i>AbLit: A Resource for Analyzing and Generating Abridged Versions of English Literature</i>	
Melissa Roemmele, Kyle Shaffer, Katrina Olsen, Yiyi Wang and Steve DeNeefe	3717
<i>Self-training Reduces Flicker in Retranslation-based Simultaneous Translation</i>	
Sukanta Sen, Rico Sennrich, Biao Zhang and Barry Haddow	3734
<i>Social Commonsense for Explanation and Cultural Bias Discovery</i>	
Lisa Bauer, Hanna Tischer and Mohit Bansal	3745
<i>Counter-GAP: Counterfactual Bias Evaluation through Gendered Ambiguous Pronouns</i>	
Zhongbin Xie, Vid Kocijan, Thomas Lukasiewicz and Oana-Maria Camburu	3761
<i>The NLP Task Effectiveness of Long-Range Transformers</i>	
Guanghui Qin, Yukun Feng and Benjamin Van Durme	3774
<i>Creation and evaluation of timelines for longitudinal user posts</i>	
Anthony Hills, Adam Tsakalidis, Federico Nanni, Ioannis Zachos and Maria Liakata	3791
<i>Semi-supervised New Event Type Induction and Description via Contrastive Loss-Enforced Batch Attention</i>	
Carl Edwards and Heng Ji	3805
<i>Multilingual Content Moderation: A Case Study on Reddit</i>	
Meng Ye, Karan Sikka, Katherine Atwell, Sabit Hassan, Ajay Divakaran and Malihe Alikhani	3828
<i>GrIPS: Gradient-free, Edit-based Instruction Search for Prompting Large Language Models</i>	
Archiki Prasad, Peter Hase, Xiang Zhou and Mohit Bansal	3845

<i>DiscoScore: Evaluating Text Generation with BERT and Discourse Coherence</i> Wei Zhao, Michael Strube and Steffen Eger.....	3865
<i>Know your audience: specializing grounded language models with listener subtraction</i> Aaditya K Singh, David Ding, Andrew Saxe, Felix Hill and Andrew Lampinen.....	3884
<i>Meeting the Needs of Low-Resource Languages: The Value of Automatic Alignments via Pretrained Models</i> Abteem Ebrahimi, Arya D. Mccarthy, Arturo Oncevay, John E. Ortega, Luis Chiruzzo, Gustavo Giménez-Lugo, Rolando Coto-Solano and Katharina Kann.....	3912