

61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)

Short Papers

Toronto, Canada
9 – 14 July 2023

Volume 1 of 3

ISBN: 978-1-7138-8174-2

Printed from e-media with permission by:

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



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

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

Printed with permission by Curran Associates, Inc. (2024)

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>Should you marginalize over possible tokenizations?</i> Nadezhda Chirkova, Germán Kruszewski, Jos Rozen and Marc Dymetman	1
<i>Back to Patterns: Efficient Japanese Morphological Analysis with Feature-Sequence Trie</i> Naoki Yoshinaga	13
<i>Transformed Protoform Reconstruction</i> Young Min Kim, Kalvin Chang, Chenxuan Cui and David R. Mortensen	24
<i>Ellipsis-Dependent Reasoning: a New Challenge for Large Language Models</i> Daniel Hardt	39
<i>Bootstrapping Neural Relation and Explanation Classifiers</i> Zheng Tang and Mihai Surdeanu	48
<i>A Fast Algorithm for Computing Prefix Probabilities</i> Franz Nowak and Ryan Cotterell	57
<i>Analyzing Text Representations by Measuring Task Alignment</i> Cesar Gonzalez-Gutierrez, Audi Primadhanty, Francesco Cazzaro and Ariadna Julieta Quattoni 70	
<i>Tracing Linguistic Markers of Influence in a Large Online Organisation</i> Prashant Khare, Ravi Shekhar, Mladen Karan, Stephen McQuistin, Colin Perkins, Ignacio Castro, Gareth Tyson, Patrick G.T. Healey and Matthew Purver	82
<i>Metaphor Detection via Explicit Basic Meanings Modelling</i> Yucheng Li, Shun Wang, Chenghua Lin and Frank Guerin	91
<i>xSIM++: An Improved Proxy to Bitext Mining Performance for Low-Resource Languages</i> Mingda Chen, Kevin Heffernan, Onur Çelebi, Alexandre Mourachko and Holger Schwenk . .	101
<i>Graph Propagation based Data Augmentation for Named Entity Recognition</i> Jiong Cai, Shen Huang, Yong Jiang, Zeqi Tan, Pengjun Xie and Kewei Tu	110
<i>Dataset Distillation with Attention Labels for Fine-tuning BERT</i> Aru Maekawa, Naoki Kobayashi, Kotaro Funakoshi and Manabu Okumura	119
<i>Multi-Document Summarization with Centroid-Based Pretraining</i> Ratish Surendran Puduppully, Parag Jain, Nancy Chen and Mark Steedman	128
<i>Scaling in Cognitive Modelling: a Multilingual Approach to Human Reading Times</i> Andrea Gregor de Varda and Marco Marelli	139
<i>Improving Generalization in Language Model-based Text-to-SQL Semantic Parsing: Two Simple Se- mantic Boundary-based Techniques</i> Daking Rai, Bailin Wang, Yilun Zhou and Ziyu Yao	150
<i>HiPool: Modeling Long Documents Using Graph Neural Networks</i> Irene Li, Aosong Feng, Dragomir Radev and Rex Ying	161
<i>A Weakly Supervised Classifier and Dataset of White Supremacist Language</i> Michael Miller Yoder, Ahmad Diab, David West Brown and Kathleen M Carley	172

<i>BOLT: Fast Energy-based Controlled Text Generation with Tunable Biases</i> Xin Liu, Muhammad Khalifa and Lu Wang	186
<i>mOKB6: A Multilingual Open Knowledge Base Completion Benchmark</i> Shubham Mittal, Keshav Kolluru, Soumen Chakrabarti and Mausam -	201
<i>Covering Uncommon Ground: Gap-Focused Question Generation for Answer Assessment</i> Roni Rabin, Alexandre Djerbetian, Roeel Engelberg, Lidan Hackmon, Gal Elidan, Reut Tsarfaty and Amir Globerson	215
<i>Detoxifying Text with MaRCO: Controllable Revision with Experts and Anti-Experts</i> Skyler R Hallinan, Alisa Liu, Yejin Choi and Maarten Sap	228
<i>A Natural Bias for Language Generation Models</i> Clara Meister, Wojciech Jan Stokowiec, Tiago Pimentel, Lei Yu, Laura Rimell and Adhiguna Kuncoro	243
<i>Simple Augmentations of Logical Rules for Neuro-Symbolic Knowledge Graph Completion</i> Ananjan Nandi, Navdeep Kaur, Parag Singla and Mausam -	256
<i>Parameter-efficient Weight Ensembling Facilitates Task-level Knowledge Transfer</i> Xingtai Lv, Ning Ding, Yujia Qin, Zhiyuan Liu and Maosong Sun	270
<i>Faithfulness Tests for Natural Language Explanations</i> Pepa Atanasova, Oana-Maria Camburu, Christina Lioma, Thomas Lukasiewicz, Jakob Grue Si- monsens and Isabelle Augenstein	283
<i>COGEN: Abductive Commonsense Language Generation</i> rohola zandie, Diwanshu Shekhar and Mohammad Mahoor	295
<i>Multimodal Relation Extraction with Cross-Modal Retrieval and Synthesis</i> Xuming Hu, Zhijiang Guo, ZHIYANG TENG, Irwin King and Philip S. Yu	303
<i>Characterization of Stigmatizing Language in Medical Records</i> Keith Harrigan, Ayah Zirikly, Brant Chee, Alya Ahmad, Anne R Links, Somnath Saha, Mary Catherine Beach and Mark Dredze	312
<i>Abstractive Summarizers are Excellent Extractive Summarizers</i> Daniel Varab and Yumo Xu	330
<i>Language Models Get a Gender Makeover: Mitigating Gender Bias with Few-Shot Data Interventions</i> Himanshu Thakur, Atishay Jain, Praneetha Vaddamanu, Paul Pu Liang and Louis-Philippe Mo- rency	340
<i>PLUE: Language Understanding Evaluation Benchmark for Privacy Policies in English</i> Jianfeng Chi, Wasi Uddin Ahmad, Yuan Tian and Kai-Wei Chang	352
<i>Stop Pre-Training: Adapt Visual-Language Models to Unseen Languages</i> Yasmine Karoui, Rémi Lebret, Negar Foroutan Eghlidi and Karl Aberer	366
<i>BUCA: A Binary Classification Approach to Unsupervised Commonsense Question Answering</i> Jie He, Simon Chi Lok U, Victor Gutierrez-Basulto and Jeff Z. Pan	376
<i>Nichelle and Nancy: The Influence of Demographic Attributes and Tokenization Length on First Name Biases</i> Haozhe An and Rachel Rudinger	388

<i>Improving Syntactic Probing Correctness and Robustness with Control Tasks</i> Weicheng Ma, Brian C Wang, Hefan Zhang, Lili Wang, Rolando Coto-Solano, Saeed Hassanpour and Soroush Vosoughi	402
<i>Split-NER: Named Entity Recognition via Two Question-Answering-based Classifications</i> Jatin Arora and Youngja Park	416
<i>Credible without Credit: Domain Experts Assess Generative Language Models</i> Denis Peskoff and Brandon Stewart	427
<i>Grokking of Hierarchical Structure in Vanilla Transformers</i> Shikhar Murty, Pratyusha Sharma, Jacob Andreas and Christopher D. Manning	439
<i>Zero-shot Cross-lingual Transfer With Learned Projections Using Unlabeled Target-Language Data</i> Ujan Deb, Ridayesh Ramesh Parab and Preethi Jyothi	449
<i>Context-Aware Transformer Pre-Training for Answer Sentence Selection</i> Luca Di Liello, Siddhant Garg and Alessandro Moschitti	458
<i>Toward Expanding the Scope of Radiology Report Summarization to Multiple Anatomies and Modalities</i> Zhihong Chen, Maya Varma, Xiang Wan, Curtis Langlotz and Jean-Benoit Delbrouck	469
<i>Efficient Diagnosis Assignment Using Unstructured Clinical Notes</i> Louis Blankemeier, Jason Fries, Robert Tinn, Joseph S Preston, Nigam Shah and Akshay Chaud- hari	485
<i>MetaVL: Transferring In-Context Learning Ability From Language Models to Vision-Language Models</i> Masoud Monajatipoor, Liunian Harold Li, Mozhddeh Rouhsedaghat, Lin Yang and Kai-Wei Chang 495	
<i>On the Interpretability and Significance of Bias Metrics in Texts: a PMI-based Approach</i> Francisco Valentini, Germán Federico Rosati, Damián Blasi, Diego Fernandez Slezak and Edgar Altszyler	509
<i>Surface-Based Retrieval Reduces Perplexity of Retrieval-Augmented Language Models</i> Ehsan Doostmohammadi, Tobias Norlund, Marco Kuhlmann and Richard Johansson	521
<i>MIReAD: Simple Method for Learning High-quality Representations from Scientific Documents</i> Anastasiia Razdaibiedina and Aleksandr V. Brechalov	530
<i>KNOW How to Make Up Your Mind! Adversarially Detecting and Alleviating Inconsistencies in Natural Language Explanations</i> Myeongjun Jang, Bodhisattwa Prasad Majumder, Julian McAuley, Thomas Lukasiewicz and Oana-Maria Camburu	540
<i>Measuring the Effect of Influential Messages on Varying Personas</i> Chenkai Sun, Jinning Li, Hou Pong Chan, ChengXiang Zhai and Heng Ji	554
<i>Going Beyond Sentence Embeddings: A Token-Level Matching Algorithm for Calculating Semantic Textual Similarity</i> Hongwei Wang and Dong Yu	563
<i>Robust Learning for Multi-party Addressee Recognition with Discrete Addressee Codebook</i> Pengcheng Zhu, Wei Zhou, Kuncai Zhang, Yuankai Ma and Haiqing Chen	571
<i>TwistList: Resources and Baselines for Tongue Twister Generation</i> Tyler Loakman, Chen Tang and Chenghua Lin	579

<i>Substitution-based Semantic Change Detection using Contextual Embeddings</i>	
Dallas Card	590
<i>Probing Physical Reasoning with Counter-Commonsense Context</i>	
Kazushi Kondo, Saku Sugawara and Akiko Aizawa	603
<i>Morphological Inflection with Phonological Features</i>	
David Guriel, Omer Goldman and Reut Tsarfaty	613
<i>A Holistic Approach to Reference-Free Evaluation of Machine Translation</i>	
Hanming Wu, Wenjuan Han, Hui Di, Yufeng Chen and Jinan Xu	623
<i>Balancing Lexical and Semantic Quality in Abstractive Summarization</i>	
Jeewoo Sul and Yong Suk Choi	637
<i>Learning Neuro-Symbolic World Models with Conversational Proprioception</i>	
Don Joven Agravante, Daiki Kimura, Michiaki Tsubori, Asim Munawar and Alexander Gray	648
<i>In and Out-of-Domain Text Adversarial Robustness via Label Smoothing</i>	
Yahan Yang, Soham Dan, Dan Roth and Insup Lee	657
<i>LM-CPPF: Paraphrasing-Guided Data Augmentation for Contrastive Prompt-Based Few-Shot Fine-Tuning</i>	
Amirhossein Abaskohi, Sascha Rothe and Yadollah Yaghoobzadeh	670
<i>Considerations for meaningful sign language machine translation based on glosses</i>	
Mathias Müller, Zifan Jiang, Amit Moryossef, Annette Rios and Sarah Ebling	682
<i>Detecting Contradictory COVID-19 Drug Efficacy Claims from Biomedical Literature</i>	
Daniel N Sosa, Malavika Suresh, Christopher Potts and Russ B Altman	694
<i>The Role of Global and Local Context in Named Entity Recognition</i>	
Arthur Amalvy, Vincent Labatut and Richard Dufour	714
<i>Joint End-to-end Semantic Proto-role Labeling</i>	
Elizabeth Spaulding, Gary Kazantsev and Mark Dredze	723
<i>Improving Automatic Quotation Attribution in Literary Novels</i>	
Krishnapriya Vishnubhotla, Frank Rudzicz, Graeme Hirst and Adam Hammond	737
<i>Modular Visual Question Answering via Code Generation</i>	
Sanjay Subramanian, Medhini Narasimhan, Kushal M Khangaonkar, Kevin Yang, Arsha Nagrani, Cordelia Schmid, Andy Zeng, Trevor Darrell and Dan Klein	747
<i>Target-Based Offensive Language Identification</i>	
Marcos Zampieri, Skye Morgan, Kai North, Tharindu Ranasinghe, Austin Simmons, Paridhi Khandelwal, Sara Rosenthal and Preslav Nakov	762
<i>Unsupervised Subtitle Segmentation with Masked Language Models</i>	
David Ponce, Thierry Etchegoyhen and Victor Ruiz	771
<i>Exploring Continual Learning for Code Generation Models</i>	
Prateek Yadav, Qing Sun, Hantian Ding, Xiaopeng Li, Dejiao Zhang, Ming Tan, Parminder Bhatia, Xiaofei Ma, Ramesh Nallapati, Murali Krishna Ramanathan, Mohit Bansal and Bing Xiang . . .	782
<i>Deep Active Learning for Morphophonological Processing</i>	
Seyed Morteza Mirbostani, Yasaman Boreshban, Salam Khalifa, SeyedAbolghasem Mirroshandel and Owen Rambow	793

<i>Counterfactual reasoning: Testing language models’ understanding of hypothetical scenarios</i> Jiaxuan Li, Lang Yu and Allyson Ettinger	804
<i>Bhasa-Abhijnaanam: Native-script and romanized Language Identification for 22 Indic languages</i> Yash H. Madhani, Mitesh M. Khapra and Anoop Kunchukuttan	816
<i>Using contradictions improves question answering systems</i> Etienne Fortier-Dubois and Domenic Anthony Rosati	827
<i>Token-Level Self-Evolution Training for Sequence-to-Sequence Learning</i> Keqin Peng, Liang Ding, Qihuang Zhong, Yuanxin Ouyang, Wenge Rong, Zhang Xiong and Dacheng Tao	841
<i>Gradient Ascent Post-training Enhances Language Model Generalization</i> Dongkeun Yoon, Joel Jang, Sungdong Kim and Minjoon Seo	851
<i>An Open Dataset and Model for Language Identification</i> Laurie V Burchell, Alexandra Birch, Nikolay Bogoychev and Kenneth Heafield	865
<i>Evaluating Paraphrastic Robustness in Textual Entailment Models</i> Dhruv Verma, Yash Kumar Lal, Shreyashee Sinha, Benjamin Van Durme and Adam Poliak ..	880
<i>Are Pre-trained Language Models Useful for Model Ensemble in Chinese Grammatical Error Correction?</i> Chenming Tang, Xiuyu Wu and Yunfang Wu	893
<i>Improving Factuality of Abstractive Summarization without Sacrificing Summary Quality</i> Tanay Dixit, Fei Wang and Muhao Chen	902
<i>With a Little Push, NLI Models can Robustly and Efficiently Predict Faithfulness</i> Julius Steen, Juri Opitz, Anette Frank and Katja Markert	914
<i>A Better Way to Do Masked Language Model Scoring</i> Carina Kauf and Anna Ivanova	925
<i>ChatGPT for Zero-shot Dialogue State Tracking: A Solution or an Opportunity?</i> Michael Heck, Nurul Lubis, Benjamin Matthias Ruppik, Renato Vukovic, Shutong Feng, Christian Geishauser, Hsien-chin Lin, Carel van Niekerk and Milica Gasic	936
<i>Controllable Mixed-Initiative Dialogue Generation through Prompting</i> Maximillian Chen, Xiao Yu, Weiyang Shi, Urvi Awasthi and Zhou Yu	951
<i>Enhancing Event Causality Identification with Counterfactual Reasoning</i> Feiteng Mu and Wenjie Li	967
<i>Contrastive Bootstrapping for Label Refinement</i> Shudi Hou, Yu Xia, Muhao Chen and Sujian Li	976
<i>NollySenti: Leveraging Transfer Learning and Machine Translation for Nigerian Movie Sentiment Classification</i> Iyanuoluwa Adeola Shode, David Ifeoluwa Adelani, Jing Peng and Anna Feldman	986
<i>Trading Syntax Trees for Wordpieces: Target-oriented Opinion Words Extraction with Wordpieces and Aspect Enhancement</i> Samuel Mensah, Kai Sun and Nikolaos Aletras	999
<i>An (unhelpful) guide to selecting the best ASR architecture for your under-resourced language</i> Robert Jimerson, Zoey Liu and Emily Prud’hommeaux	1008

<i>The Ecological Fallacy in Annotation: Modeling Human Label Variation goes beyond Sociodemographics</i>	
Matthias Orlikowski, Paul Röttger, Philipp Cimiano and Dirk Hovy	1017
<i>Decomposed scoring of CCG dependencies</i>	
Aditya Bhargava and Gerald Penn	1030
<i>Do GPTs Produce Less Literal Translations?</i>	
Vikas Raunak, Arul Menezes, Matt Post and Hany Hassan	1041
<i>Environmental Claim Detection</i>	
Dominik Stammach, Nicolas Webersinke, Julia Anna Bingler, Mathias Kraus and Markus Leopold	1051
<i>Black-box language model explanation by context length probing</i>	
Ondřej Cifka and Antoine Liutkus	1067
<i>Let Me Check the Examples: Enhancing Demonstration Learning via Explicit Imitation</i>	
Sirui Wang, Kaiwen Wei, Hongzhi Zhang, Yuntao Li and Wei Wu	1080
<i>The Inside Story: Towards Better Understanding of Machine Translation Neural Evaluation Metrics</i>	
Ricardo Rei, Nuno M. Guerreiro, Marcos Treviso, Luisa Coheur, Alon Lavie and André Martins	1089
<i>Typo-Robust Representation Learning for Dense Retrieval</i>	
Panuthep Tasawong, Wuttikorn Ponwitayarat, Peerat Limkonchotiwat, Can Udomcharoenchaikit, Ekapol Chuangsuwanich and Sarana Nutanong	1106
<i>Focused Prefix Tuning for Controllable Text Generation</i>	
Congda Ma, Tianyu Zhao, Makoto Shing, Kei Sawada and Manabu Okumura	1116
<i>ReAugKD: Retrieval-Augmented Knowledge Distillation For Pre-trained Language Models</i>	
Jianyi Zhang, Aashiq Muhamed, Aditya Anantharaman, Guoyin Wang, Changyou Chen, Kai Zhong, Qingjun Cui, Yi Xu, Belinda Zeng, Trishul Chilimbi and Yiran Chen	1128
<i>Debiasing Generative Named Entity Recognition by Calibrating Sequence Likelihood</i>	
Yu Xia, Yongwei Zhao, Wenhao Wu and Sujian Li	1137
<i>Deriving Language Models from Masked Language Models</i>	
Lucas Torroba Hennigen and Yoon Kim	1149
<i>UniTRec: A Unified Text-to-Text Transformer and Joint Contrastive Learning Framework for Text-based Recommendation</i>	
Zhiming Mao, Huimin Wang, Yiming Du and Kam-Fai Wong	1160
<i>Reasoning Implicit Sentiment with Chain-of-Thought Prompting</i>	
Hao Fei, Bobo Li, Qian Liu, Lidong Bing, Fei Li and Tat-Seng Chua	1171
<i>Latent Positional Information is in the Self-Attention Variance of Transformer Language Models Without Positional Embeddings</i>	
Ta-Chung Chi, Ting-Han Fan, Li-Wei Chen, alexander rudnicky and Peter J Ramadge	1183
<i>Is Anisotropy Truly Harmful? A Case Study on Text Clustering</i>	
Mira Ait-Saada and Mohamed Nadif	1194
<i>Class based Influence Functions for Error Detection</i>	
Thang Nguyen-Duc, Hoang Thanh-Tung, Quan Hung Tran, Dang Huu-Tien, Hieu Ngoc Nguyen, Anh T. V. Dau and Nghi D. Q. Bui	1204

<i>Leveraging Prefix Transfer for Multi-Intent Text Revision</i>	
Ruining Chong, Cunliang Kong, Liu Wu, Zhenghao Liu, Ziyi Jin, Liner Yang, Yange Fan, Hanghang Fan and Erhong Yang	1219
<i>Learning Multi-Step Reasoning by Solving Arithmetic Tasks</i>	
Tianduo Wang and Wei Lu	1229
<i>Towards Adaptive Prefix Tuning for Parameter-Efficient Language Model Fine-tuning</i>	
Zhen-Ru Zhang, Chuanqi Tan, Haiyang Xu, Chengyu Wang, Jun Huang and Songfang Huang	1239
<i>Improving Gender Fairness of Pre-Trained Language Models without Catastrophic Forgetting</i>	
Zahra Fatemi, Chen Xing, Wenhao Liu and Caiming Xiong	1249
<i>Class-Incremental Learning based on Label Generation</i>	
Yijia Shao, Yiduo Guo, Dongyan Zhao and Bing Liu	1263
<i>Evaluating pragmatic abilities of image captioners on A3DS</i>	
Polina Tsvilodub and Michael Franke	1277
<i>The Art of Prompting: Event Detection based on Type Specific Prompts</i>	
Sijia Wang, Mo Yu and Lifu Huang	1286
<i>Exploring the Impact of Layer Normalization for Zero-shot Neural Machine Translation</i>	
Zhuoyuan Mao, Raj Dabre, Qianying Liu, Haiyue Song, Chenhui Chu and Sadao Kurohashi	1300
<i>Do Models Really Learn to Follow Instructions? An Empirical Study of Instruction Tuning</i>	
Po-Nien Kung and Nanyun Peng	1317
<i>Self-Distilled Quantization: Achieving High Compression Rates in Transformer-Based Language Models</i>	
James O’Neill and Sourav Dutta	1329
<i>Modality Adaption or Regularization? A Case Study on End-to-End Speech Translation</i>	
Yuchen Han, Chen Xu, Tong Xiao and Jingbo Zhu	1340
<i>Uncertainty-Aware Bootstrap Learning for Joint Extraction on Distantly-Supervised Data</i>	
Yufe Li, Xiao Yu, Yanchi Liu, Haifeng Chen and Cong Liu	1349
<i>Text-to-SQL Error Correction with Language Models of Code</i>	
Ziru Chen, Shijie Chen, Michael White, Raymond Mooney, Ali Payani, Jayanth Srinivasa, Yu Su and Huan Sun	1359
<i>The Tail Wagging the Dog: Dataset Construction Biases of Social Bias Benchmarks</i>	
Nikil Rooshan Selvam, Sunipa Dev, Daniel Khoshdel, Tushar Khot and Kai-Wei Chang	1373
<i>Summarizing, Simplifying, and Synthesizing Medical Evidence using GPT-3 (with Varying Success)</i>	
Chantal Shaib, Millicent L Li, Sebastian A Joseph, Iain Marshall, Junyi Jessy Li and Byron C. Wallace	1387
<i>Prefix Propagation: Parameter-Efficient Tuning for Long Sequences</i>	
Jonathan X Li, Will Aitken, Rohan Bhambhoria and Xiaodan Zhu	1408
<i>Listener Model for the PhotoBook Referential Game with CLIP Scores as Implicit Reference Chain</i>	
Shih-Lun Wu, Yi-Hui Chou and Liangze Li	1420
<i>Bring More Attention to Syntactic Symmetry for Automatic Postediting of High-Quality Machine Translations</i>	
Baikjin Jung, Myungji Lee, Jong-Hyeok Lee and Yunsu Kim	1433

<i>An Embarrassingly Easy but Strong Baseline for Nested Named Entity Recognition</i>	
Hang Yan, Yu Sun, Xiaonan Li and Xipeng Qiu	1442
<i>Hexatagging: Projective Dependency Parsing as Tagging</i>	
Afra Amini, Tianyu Liu and Ryan Cotterell	1453
<i>Understanding Demonstration-based Learning from a Causal Perspective</i>	
Ruiyi Zhang and Tong Yu	1465
<i>RAMP: Retrieval and Attribute-Marking Enhanced Prompting for Attribute-Controlled Translation</i>	
Gabriele Sarti, Phu Mon Htut, Xing Niu, Benjamin Hsu, Anna Currey, Georgiana Dinu and Maria Nadejde	1476
<i>Zero-Shot and Few-Shot Stance Detection on Varied Topics via Conditional Generation</i>	
Haoyang Wen and Alexander Hauptmann	1491
<i>Discourse-Level Representations can Improve Prediction of Degree of Anxiety</i>	
Swanie Juhng, Matthew Matero, Vasudha Varadarajan, Johannes C Eichstaedt, Adithya V Ganesan and H. Andrew Schwartz	1500
<i>Controlling the Extraction of Memorized Data from Large Language Models via Prompt-Tuning</i>	
Mustafa Safa Ozdayi, Charith Peris, Jack FitzGerald, Christophe Dupuy, Jimit Majmudar, Haidar Khan, Rahil Parikh and Rahul Gupta	1512
<i>MultiTool-CoT: GPT-3 Can Use Multiple External Tools with Chain of Thought Prompting</i>	
Tatsuro Inaba, Hirokazu Kiyomaru, Fei Cheng and Sadao Kurohashi	1522
<i>mPMR: A Multilingual Pre-trained Machine Reader at Scale</i>	
Weiwen Xu, Xin Li, Wai Lam and Lidong Bing	1533
<i>MOSPC: MOS Prediction Based on Pairwise Comparison</i>	
Kexin Wang, Yunlong Zhao, Qianqian Dong, Tom Ko and Mingxuan Wang	1547
<i>LI-RAGE: Late Interaction Retrieval Augmented Generation with Explicit Signals for Open-Domain Table Question Answering</i>	
Weizhe Lin, Rexhina Blloshmi, Bill Byrne, Adria de Gispert and Gonzalo Iglesias	1557
<i>How Well Apply Simple MLP to Incomplete Utterance Rewriting?</i>	
Jiang Li, Xiangdong Su, Xinlan Ma and Guanglai Gao	1567
<i>XL-LEXEME: WiC Pretrained Model for Cross-Lingual LEXical sEMantic changE</i>	
Pierluigi Cassotti, Lucia Siciliani, Marco DeGemmis, Giovanni Semeraro and Pierpaolo Basile	1577
<i>Theory-Grounded Computational Text Analysis</i>	
Arya D. McCarthy and Giovanna Maria Dora Dore	1586
<i>AMRs Assemble! Learning to Ensemble with Autoregressive Models for AMR Parsing</i>	
Abelardo Carlos Martínez Lorenzo, Pere Lluís Huguet Cabot and Roberto Navigli	1595
<i>MolXPT: Wrapping Molecules with Text for Generative Pre-training</i>	
Zequn Liu, Wei Zhang, Yingce Xia, Lijun Wu, Shufang Xie, Tao Qin, Ming Zhang and Tie-Yan Liu	1606
<i>A Study on the Efficiency and Generalization of Light Hybrid Retrievers</i>	
Man Luo, Shashank Jain, Anchit Gupta, Arash Einolghozati, Barlas Oguz, Debojeet Chatterjee, Xilun Chen, Chitta Baral and Peyman Heidari	1617

<i>The Mechanical Bard: An Interpretable Machine Learning Approach to Shakespearean Sonnet Generation</i>	
Edwin Agnew, Michelle Qiu, Lily Zhu, Sam Wiseman and Cynthia Rudin	1627
<i>When to Use Efficient Self Attention? Profiling Text, Speech and Image Transformer Variants</i>	
Anuj Diwan, Eunsol Choi and David Harwath	1639
<i>Evaluating Zero-Shot Event Structures: Recommendations for Automatic Content Extraction (ACE) Annotations</i>	
Erica Cai and Brendan O’Connor	1651
<i>Event Extraction as Question Generation and Answering</i>	
Di Lu, Shihao Ran, Joel Tetreault and Alejandro Jaimes	1666
<i>Are Sample-Efficient NLP Models More Robust?</i>	
Nelson F. Liu, Ananya Kumar, Percy Liang and Robin Jia	1689
<i>Diversity-Aware Coherence Loss for Improving Neural Topic Models</i>	
Raymond Li, Felipe Gonzalez-Pizarro, Linzi Xing, Gabriel Murray and Giuseppe Carenini .	1710
<i>NarrowBERT: Accelerating Masked Language Model Pretraining and Inference</i>	
Haoxin Li, Phillip Keung, Daniel Cheng, Jungo Kasai and Noah A. Smith	1723
<i>S3HQA: A Three-Stage Approach for Multi-hop Text-Table Hybrid Question Answering</i>	
Fangyu Lei, Xiang Li, Yifan Wei, Shizhu He, Yiming Huang, Jun Zhao and Kang Liu	1731
<i>Towards Fewer Hallucinations in Knowledge-Grounded Dialogue Generation via Augmentative and Contrastive Knowledge-Dialogue</i>	
Bin Sun, Yitong Li, Fei Mi, fanhu bie, Yiwei Li and Kan Li	1741
<i>AutoConv: Automatically Generating Information-seeking Conversations with Large Language Models</i>	
Siheng Li, Cheng Yang, Yichun Yin, Xinyu Zhu, Zesen Cheng, Lifeng Shang, Xin Jiang, Qun Liu and Yujiu Yang	1751
<i>STT4SG-350: A Speech Corpus for All Swiss German Dialect Regions</i>	
Michel Plüss, Jan Deriu, Yanick Schraner, Claudio Paonessa, Julia Hartmann, Larissa Schmidt, Christian Scheller, Manuela Hürlimann, Tanja Samardžić, Manfred Vogel and Mark Cieliebak . . .	1763
<i>Teaching Small Language Models to Reason</i>	
Lucie Charlotte Magister, Jonathan Mallinson, Jakub Dominik Adamek, Eric Malmi and Aliaksei Severyn	1773
<i>A Simple and Effective Framework for Strict Zero-Shot Hierarchical Classification</i>	
Rohan Bhambhora, Lei Chen and Xiaodan Zhu	1782
<i>A Simple Concatenation can Effectively Improve Speech Translation</i>	
Linlin Zhang, Kai Fan, Boxing Chen and Luo Si	1793
<i>ScoNe: Benchmarking Negation Reasoning in Language Models With Fine-Tuning and In-Context Learning</i>	
Jingyuan S. She, Christopher Potts, Samuel R. Bowman and Atticus Geiger	1803
<i>Revisiting Automated Prompting: Are We Actually Doing Better?</i>	
Yulin Zhou, Yiren Zhao, Iliia Shumailov, Robert Mullins and Yarin Gal	1822
<i>Mind the Gap between the Application Track and the Real World</i>	
Ananya Ganesh, Jie Cao, E. Margaret Perkoff, Rosy Southwell, Martha Palmer and Katharina Kann	1833

<i>How to Distill your BERT: An Empirical Study on the Impact of Weight Initialisation and Distillation Objectives</i>	
Xinpeng Wang, Leonie Weissweiler, Hinrich Schütze and Barbara Plank	1843
<i>ACTC: Active Threshold Calibration for Cold-Start Knowledge Graph Completion</i>	
Anastasiia Sedova and Benjamin Roth	1853
<i>Task-Aware Specialization for Efficient and Robust Dense Retrieval for Open-Domain Question Answering</i>	
Hao Cheng, Hao Fang, Xiaodong Liu and Jianfeng Gao	1864
<i>Linear Classifier: An Often-Forgotten Baseline for Text Classification</i>	
Yu-Chen Lin, Si-An Chen, Jie-Jyun Liu and Chih-Jen Lin	1876
<i>Randomized Positional Encodings Boost Length Generalization of Transformers</i>	
Anian Ruoss, Grégoire Delétang, Tim Genewein, Jordi Grau-Moya, Róbert Csordás, Mehdi Benani, Shane Legg and Joel Veness	1889
<i>Table and Image Generation for Investigating Knowledge of Entities in Pre-trained Vision and Language Models</i>	
Hidetaka Kamigaito, Katsuhiko Hayashi and Taro Watanabe	1904
<i>Improving Grammar-based Sequence-to-Sequence Modeling with Decomposition and Constraints</i>	
Chao Lou and Kewei Tu	1918
<i>TeCS: A Dataset and Benchmark for Tense Consistency of Machine Translation</i>	
Yiming Ai, Zhiwei He, Kai Yu and Rui Wang	1930