

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

Short Papers

Dublin, Ireland
22 – 27 May 2022

ISBN: 978-1-7138-6722-7

Printed from e-media with permission by:

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



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

Copyright© (2022) by the Association for Computational Linguistics
All rights reserved.

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

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

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

Phone: 1-570-476-8006

Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

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

Table of Contents

<i>BitFit: Simple Parameter-efficient Fine-tuning for Transformer-based Masked Language-models</i> Elad Ben Zaken, Yoav Goldberg and Shauli Ravfogel	1
<i>Are Shortest Rationales the Best Explanations for Human Understanding?</i> Hua Shen, Tongshuang Wu, Wenbo Guo and Ting-Hao Huang	10
<i>Analyzing Wrap-Up Effects through an Information-Theoretic Lens</i> Clara Isabel Meister, Tiago Pimentel, Thomas Hikaru Clark, Ryan D Cotterell and Roger P. Levy 20	
<i>Have my arguments been replied to? Argument Pair Extraction as Machine Reading Comprehension</i> Jianzhu Bao, Jingyi Sun, Qinglin Zhu and Ruifeng Xu	29
<i>High probability or low information? The probability–quality paradox in language generation</i> Clara Isabel Meister, Gian Wiher, Tiago Pimentel and Ryan D Cotterell	36
<i>Disentangled Knowledge Transfer for OOD Intent Discovery with Unified Contrastive Learning</i> Yutao Mou, Keqing He, Yanan Wu, Zhiyuan Zeng, Hong Xu, Huixing Jiang, Wei Wu and Weiran Xu	46
<i>Voxel-informed Language Grounding</i> Rodolfo Corona, Shizhan Zhu, Dan Klein and Trevor Darrell	54
<i>P-Tuning: Prompt Tuning Can Be Comparable to Fine-tuning Across Scales and Tasks</i> Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang and Jie Tang .	61
<i>On Efficiently Acquiring Annotations for Multilingual Models</i> Joel Ruben Antony Moniz, Barun Patra and Matthew R. Gormley	69
<i>Automatic Detection of Entity-Manipulated Text using Factual Knowledge</i> Ganesh Jawahar, Muhammad Abdul-Mageed and Laks V. S. Lakshmanan	86
<i>Does BERT Know that the IS-A Relation Is Transitive?</i> Ruixi Lin and Hwee Tou Ng	94
<i>Buy Tesla, Sell Ford: Assessing Implicit Stock Market Preference in Pre-trained Language Models</i> Chengyu Chuang and Yi Yang	100
<i>Pixie: Preference in Implicit and Explicit Comparisons</i> Amanul Haque, Vaibhav Garg, Hui Guo and Munindar P. Singh	106
<i>Counterfactual Explanations for Natural Language Interfaces</i> George Tolkachev, Stephen Mell, Stephan Zdancewic and Osbert Bastani	113
<i>Predicting Difficulty and Discrimination of Natural Language Questions</i> Matthew Alexander Byrd and Shashank Srivastava	119
<i>How does the pre-training objective affect what large language models learn about linguistic properties?</i> Ahmed Alajrami and Nikolaos Aletras	131
<i>The Power of Prompt Tuning for Low-Resource Semantic Parsing</i> Nathan Schucher, Siva Reddy and Harm de Vries	148

<i>Data Contamination: From Memorization to Exploitation</i> Inbal Magar and Roy Schwartz	157
<i>Detecting Annotation Errors in Morphological Data with the Transformer</i> Ling Liu and Mans Hulden	166
<i>Estimating the Entropy of Linguistic Distributions</i> Aryaman Arora, Clara Isabel Meister and Ryan D Cotterell	175
<i>Morphological Reinflection with Multiple Arguments: An Extended Annotation schema and a Georgian Case Study</i> David Guriel, Omer Goldman and Reut Tsarfaty	196
<i>DQ-BART: Efficient Sequence-to-Sequence Model via Joint Distillation and Quantization</i> Zheng Li, Zijian Wang, Ming Tan, Ramesh Nallapati, Parminder Bhatia, Andrew Arnold, Bing Xiang and Dan Roth	203
<i>Learning-by-Narrating: Narrative Pre-Training for Zero-Shot Dialogue Comprehension</i> Chao Zhao, Wenlin Yao, Dian Yu, Kaiqiang Song, Dong Yu and Jianshu Chen	212
<i>Kronecker Decomposition for GPT Compression</i> Ali Edalati, Marzieh S. Tahaei, Ahmad Rashid, Vahid Partovi Nia, James J. Clark and Mehdi Rezagholizadeh	219
<i>Simple and Effective Knowledge-Driven Query Expansion for QA-Based Product Attribute Extraction</i> Keiji Shinzato, Naoki Yoshinaga, Yandi Xia and Wei-Te Chen	227
<i>Event-Event Relation Extraction using Probabilistic Box Embedding</i> EunJeong Hwang, Jay-Yoon Lee, Tianyi Yang, Dhruvesh Patel, Dongxu Zhang and Andrew McCallum	235
<i>Sample, Translate, Recombine: Leveraging Audio Alignments for Data Augmentation in End-to-end Speech Translation</i> Tsz Kin Lam, Shigehiko Schamoni and Stefan Riezler	245
<i>Predicting Sentence Deletions for Text Simplification Using a Functional Discourse Structure</i> Bohan Zhang, Prafulla Kumar Choubey and Ruihong Huang	255
<i>Multilingual Pre-training with Language and Task Adaptation for Multilingual Text Style Transfer</i> Huiyuan Lai, Antonio Toral and Malvina Nissim	262
<i>When to Use Multi-Task Learning vs Intermediate Fine-Tuning for Pre-Trained Encoder Transfer Learning</i> Orion Weller, Kevin Seppi and Matt Gardner	272
<i>Leveraging Explicit Lexico-logical Alignments in Text-to-SQL Parsing</i> Runxin Sun, Shizhu He, Chong Zhu, Yaohan He, Jinlong Li, Jun Zhao and Kang Liu	283
<i>Complex Evolutional Pattern Learning for Temporal Knowledge Graph Reasoning</i> Zixuan Li, Saiping Guan, Xiaolong Jin, Weihua Peng, Yajuan Lyu, Yong Zhu, Long Bai, Wei Li, Jiafeng Guo and Xueqi Cheng	290
<i>Mismatch between Multi-turn Dialogue and its Evaluation Metric in Dialogue State Tracking</i> Takyong Kim, Hoonsang Yoon, Yukyung Lee, Pilsung Kang and Misuk Kim	297
<i>LM-BFF-MS: Improving Few-Shot Fine-tuning of Language Models based on Multiple Soft Demonstration Memory</i> Eunhwan Park, Donghyeon Jeon, Seonhoon Kim, Inho Kang and Seung-Hoon Na	310

<i>Towards Fair Evaluation of Dialogue State Tracking by Flexible Incorporation of Turn-level Performances</i>	
Suvodip Dey, Ramamohan Kummara and Maunendra Sankar Desarkar	318
<i>Exploiting Language Model Prompts Using Similarity Measures: A Case Study on the Word-in-Context Task</i>	
Mohsen Tabasi, Kiamehr Rezaee and Mohammad Taher Pilehvar	325
<i>Hierarchical Curriculum Learning for AMR Parsing</i>	
Peiyi Wang, Liang Chen, Tianyu Liu, Damai Dai, Yunbo Cao, Baobao Chang and Zhifang Sui	333
<i>PARE: A Simple and Strong Baseline for Monolingual and Multilingual Distantly Supervised Relation Extraction</i>	
Vipul Kumar Rathore, Kartikeya Badola, Parag Singla and Mausam	340
<i>To Find Waldo You Need Contextual Cues: Debiasing Who's Waldo</i>	
Yiran Luo, Pratyay Banerjee, Tejas Gokhale, Yezhou Yang and Chitta Baral	355
<i>Translate-Train Embracing Translationese Artifacts</i>	
Sicheng Yu, Qianru Sun, Hao Zhang and Jing Jiang	362
<i>C-MORE: Pretraining to Answer Open-Domain Questions by Consulting Millions of References</i>	
Xiang Yue, Xiaoman Pan, Wenlin Yao, Dian Yu, Dong Yu and Jianshu Chen	371
<i>k-Rater Reliability: The Correct Unit of Reliability for Aggregated Human Annotations</i>	
Ka Wong and Praveen Paritosh	378
<i>An Embarrassingly Simple Method to Mitigate Undesirable Properties of Pretrained Language Model Tokenizers</i>	
Valentin Hofmann, Hinrich Schuetze and Janet B. Pierrehumbert	385
<i>SCD: Self-Contrastive Decorrelation of Sentence Embeddings</i>	
Tassilo Klein and Moin Nabi	394
<i>Problems with Cosine as a Measure of Embedding Similarity for High Frequency Words</i>	
Kaitlyn Zhou, Kawin Ethayarajh, Dallas Card and Dan Jurafsky	401
<i>Revisiting the Compositional Generalization Abilities of Neural Sequence Models</i>	
Arkil Patel, Satwik Bhattamishra, Phil Blunsom and Navin Goyal	424
<i>A Copy-Augmented Generative Model for Open-Domain Question Answering</i>	
Shuang Liu, Dong Wang, Xiaoguang Li, Minghui Huang and Meizhen Ding	435
<i>Augmenting Document Representations for Dense Retrieval with Interpolation and Perturbation</i>	
Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang and Jong C. Park	442
<i>WLASL-LEX: a Dataset for Recognising Phonological Properties in American Sign Language</i>	
Federico Tavella, Viktor Schlegel, Marta Romeo, Aphrodite Galata and Angelo Cangelosi ..	453
<i>Investigating person-specific errors in chat-oriented dialogue systems</i>	
Koh Mitsuda, Ryuichiro Higashinaka, Tingxuan Li and Sen Yoshida	464
<i>Direct parsing to sentiment graphs</i>	
David Samuel, Jeremy Barnes, Robin Kurtz, Stephan Oepen, Lilja Øvrelid and Erik Velldal .	470
<i>XDBERT: Distilling Visual Information to BERT from Cross-Modal Systems to Improve Language Understanding</i>	
Chan-Jan Hsu, Hung-yi Lee and Yu Tsao	479

<i>As Little as Possible, as Much as Necessary: Detecting Over- and Undertranslations with Contrastive Conditioning</i>	
Jannis Vamvas and Rico Sennrich	490
<i>How Distributed are Distributed Representations? An Observation on the Locality of Syntactic Information in Verb Agreement Tasks</i>	
Bingzhi Li, Guillaume Wisniewski and Benoit Crabbé	501
<i>Machine Translation for Livonian: Catering to 20 Speakers</i>	
Matīss Rikters, Marili Tomingas, Tuuli Tuisk, Valts Ernštreits and Mark Fishel	508
<i>Fire Burns, Sword Cuts: Commonsense Inductive Bias for Exploration in Text-based Games</i>	
Dongwon Kelvin Ryu, Ehsan Shareghi, Meng Fang, Yunqiu Xu, Shirui Pan and Reza Haf ...	515
<i>A Simple but Effective Pluggable Entity Lookup Table for Pre-trained Language Models</i>	
Deming Ye, Yankai Lin, Peng Li, Maosong Sun and Zhiyuan Liu	523
<i>S⁴-Tuning: A Simple Cross-lingual Sub-network Tuning Method</i>	
Runxin Xu, Fuli Luo, Baobao Chang, Songfang Huang and Fei Huang	530
<i>Region-dependent temperature scaling for certainty calibration and application to class-imbalanced token classification</i>	
Hillary Dawkins and Isar Nejadgholi	538
<i>Developmental Negation Processing in Transformer Language Models</i>	
Antonio Laverghetta Jr. and John Licato	545
<i>Canary Extraction in Natural Language Understanding Models</i>	
Rahil Parikh, Christophe Dupuy and Rahul Gupta	552
<i>On the Intrinsic and Extrinsic Fairness Evaluation Metrics for Contextualized Language Representations</i>	
Yang Trista Cao, Yada Pruksachatkun, Kai-Wei Chang, Rahul Gupta, Varun Kumar, Jwala Dhama and Aram Galstyan	561
<i>Sequence-to-sequence AMR Parsing with Ancestor Information</i>	
Chen Yu and Daniel Gildea	571
<i>Zero-Shot Dependency Parsing with Worst-Case Aware Automated Curriculum Learning</i>	
Miryam De Lhoneux, Sheng Zhang and Anders Søgaard	578
<i>PriMock57: A Dataset Of Primary Care Mock Consultations</i>	
Alex Papadopoulos Korfiatis, Francesco Moramarco, Radmila Sarac and Aleksandar Savkov	588
<i>UniGDD: A Unified Generative Framework for Goal-Oriented Document-Grounded Dialogue</i>	
Chang Gao, Wenxuan Zhang and Wai Lam	599
<i>DMix: Adaptive Distance-aware Interpolative Mixup</i>	
Ramit Sawhney, Megh Thakkar, Shrey Pandit, Ritesh Singh Soun, Di Jin, Diyi Yang and Lucie Flek	606
<i>Sub-Word Alignment is Still Useful: A Vest-Pocket Method for Enhancing Low-Resource Machine Translation</i>	
Minhan Xu and Yu Hong	613
<i>HYPHEN: Hyperbolic Hawkes Attention For Text Streams</i>	
Shivam Agarwal, Ramit Sawhney, Sanchit Ahuja, Ritesh Singh Soun and Sudheer Chava ...	620

<i>A Risk-Averse Mechanism for Suicidality Assessment on Social Media</i>	
Ramit Sawhney, Atula Tejaswi Neerkaje and Manas Gaur	628
<i>When classifying grammatical role, BERT doesn't care about word order... except when it matters</i>	
Isabel Papadimitriou, Richard Futrell and Kyle Mahowald	636
<i>Triangular Transfer: Freezing the Pivot for Triangular Machine Translation</i>	
Meng Zhang, Liangyou Li and Qun Liu	644
<i>Can Visual Dialogue Models Do Scorekeeping? Exploring How Dialogue Representations Incrementally Encode Shared Knowledge</i>	
Brielen Madureira and David Schlangen	651
<i>Focus on the Target's Vocabulary: Masked Label Smoothing for Machine Translation</i>	
Liang Chen, Runxin Xu and Baobao Chang	665
<i>Contrastive Learning-Enhanced Nearest Neighbor Mechanism for Multi-Label Text Classification</i>	
Xi'ao Su, Ran Wang and Xinyu Dai	672
<i>NoisyTune: A Little Noise Can Help You Finetune Pretrained Language Models Better</i>	
Chuhan Wu, Fangzhao Wu, Tao Qi and Yongfeng Huang	680
<i>Adjusting the Precision-Recall Trade-Off with Align-and-Predict Decoding for Grammatical Error Correction</i>	
Xin Sun and Houfeng Wang	686
<i>On the Effect of Isotropy on VAE Representations of Text</i>	
Lan Zhang, Wray Buntine and Ehsan Shareghi	694
<i>Efficient Classification of Long Documents Using Transformers</i>	
Hyunji Hayley Park, Yogarshi Vyas and Kashif Shah	702
<i>Rewarding Semantic Similarity under Optimized Alignments for AMR-to-Text Generation</i>	
Lisa Jin and Daniel Gildea	710
<i>An Analysis of Negation in Natural Language Understanding Corpora</i>	
Md Mosharaf Hossain, Dhivya Chinnappa and Eduardo Blanco	716
<i>Primum Non Nocere: Before working with Indigenous data, the ACL must confront ongoing colonialism</i>	
Lane Schwartz	724
<i>Unsupervised multiple-choice question generation for out-of-domain Q&A fine-tuning</i>	
Guillaume Le Berre, Christophe Cerisara, Philippe Langlais and Guy Lapalme	732
<i>Can a Transformer Pass the Wug Test? Tuning Copying Bias in Neural Morphological Inflection Models</i>	
Ling Liu and Mans Hulden	739
<i>Probing the Robustness of Trained Metrics for Conversational Dialogue Systems</i>	
Jan Milan Deriu, Don Tuggener, Pius Von Däniken and Mark Cieliebak	750
<i>Rethinking and Refining the Distinct Metric</i>	
Siyang Liu, Sahand Sabour, Yinhe Zheng, Pei Ke, Xiaoyan Zhu and Minlie Huang	762
<i>How reparametrization trick broke differentially-private text representation learning</i>	
Ivan Habernal	771
<i>Towards Consistent Document-level Entity Linking: Joint Models for Entity Linking and Coreference Resolution</i>	
Klim Zaporojets, Johannes Deleu, Yiwei Jiang, Thomas Demeester and Chris Develder	778

<i>A Flexible Multi-Task Model for BERT Serving</i>	
Tianwen Wei, Jianwei Qi and Shenghuan He	785
<i>Understanding Game-Playing Agents with Natural Language Annotations</i>	
Nicholas Tomlin, Andre Wang He and Dan Klein	797
<i>Code Synonyms Do Matter: Multiple Synonyms Matching Network for Automatic ICD Coding</i>	
Zheng Yuan, Chuanqi Tan and Songfang Huang	808
<i>CoDA21: Evaluating Language Understanding Capabilities of NLP Models With Context-Definition Alignment</i>	
Lütfi Kerem Senel, Timo Schick and Hinrich Schuetze	815
<i>On the Importance of Effectively Adapting Pretrained Language Models for Active Learning</i>	
Katerina Margatina, Loic Barrault and Nikolaos Aletras	825
<i>A Recipe for Arbitrary Text Style Transfer with Large Language Models</i>	
Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch and Jason Wei	837
<i>DiS-ReX: A Multilingual Dataset for Distantly Supervised Relation Extraction</i>	
Abhyuday Bhartiya, Kartikeya Badola and Mausam	849
<i>(Un)solving Morphological Inflection: Lemma Overlap Artificially Inflates Models' Performance</i>	
Omer Goldman, David Guriel and Reut Tsarfaty	864
<i>Text Smoothing: Enhance Various Data Augmentation Methods on Text Classification Tasks</i>	
Xing Wu, Chaochen Gao, Meng Lin, Liangjun Zang and Songlin Hu	871