

Third Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)

Toronto, Canada
14 July 2023

ISBN: 978-1-7138-8236-7

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© (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>Towards Faithful Explanations for Text Classification with Robustness Improvement and Explanation Guided Training</i> | |
| Dongfang Li, Baotian Hu, Qingcai Chen and Shan He | 1 |
| <i>Driving Context into Text-to-Text Privatization</i> | |
| Stefan Arnold, Dilara Yesilbas and Sven Weinzierl | 15 |
| <i>Automated Ableism: An Exploration of Explicit Disability Biases in Sentiment and Toxicity Analysis Models</i> | |
| Pranav Narayanan Venkit, Mukund Srinath and Shomir Wilson | 26 |
| <i>Pay Attention to the Robustness of Chinese Minority Language Models! Syllable-level Textual Adversarial Attack on Tibetan Script</i> | |
| Xi Cao, Dolma Dawa, Nuo Qun and Trashy Nyima | 35 |
| <i>Can we trust the evaluation on ChatGPT?</i> | |
| Rachith Aiyappa, Jisun An, Haewoon Kwak and Yong-yeol Ahn | 47 |
| <i>Improving Factuality of Abstractive Summarization via Contrastive Reward Learning</i> | |
| I-chun Chern, Zhiruo Wang, Sanjan Das, Bhavuk Sharma, Pengfei Liu and Graham Neubig | 55 |
| <i>Examining the Causal Impact of First Names on Language Models: The Case of Social Commonsense Reasoning</i> | |
| Sullam Jeoung, Jana Diesner and Halil Kilicoglu | 61 |
| <i>Reliability Check: An Analysis of GPT-3’s Response to Sensitive Topics and Prompt Wording</i> | |
| Aisha Khatun and Daniel Brown | 73 |
| <i>Sample Attackability in Natural Language Adversarial Attacks</i> | |
| Vyas Raina and Mark Gales | 96 |
| <i>A Keyword Based Approach to Understanding the Overpenalization of Marginalized Groups by English Marginal Abuse Models on Twitter</i> | |
| Kyra Yee, Alice Schoenauer Sebag, Olivia Redfield, Matthias Eck, Emily Sheng and Luca Belli | 108 |
| <i>An Empirical Study of Metrics to Measure Representational Harms in Pre-Trained Language Models</i> | |
| Saghar Hosseini, Hamid Palangi and Ahmed Hassan Awadallah | 121 |
| <i>Linguistic Properties of Truthful Response</i> | |
| Bruce W. Lee, Benedict Florance Arockiaraj and Helen Jin | 135 |
| <i>Debunking Biases in Attention</i> | |
| Shijing Chen, Usman Naseem and Imran Razzak | 141 |
| <i>Guiding Text-to-Text Privatization by Syntax</i> | |
| Stefan Arnold, Dilara Yesilbas and Sven Weinzierl | 151 |
| <i>Are fairness metric scores enough to assess discrimination biases in machine learning?</i> | |
| Fanny Jourdan, Laurent Risser, Jean-michel Loubes and Nicholas Asher | 163 |
| <i>DEPTH+: An Enhanced Depth Metric for Wikipedia Corpora Quality</i> | |
| Saied Alshahrani, Norah Alshahrani and Jeanna Matthews | 175 |

| | |
|--|-----|
| <i>Distinguishing Fact from Fiction: A Benchmark Dataset for Identifying Machine-Generated Scientific Papers in the LLM Era.</i> | |
| Edoardo Mosca, Mohamed Hesham Ibrahim Abdalla, Paolo Basso, Margherita Musumeci and Georg Groh | 190 |
| <i>Detecting Personal Information in Training Corpora: an Analysis</i> | |
| Nishant Subramani, Sasha Luccioni, Jesse Dodge and Margaret Mitchell | 208 |
| <i>Enhancing textual counterfactual explanation intelligibility through Counterfactual Feature Importance</i> | |
| Milan Bhan, Jean-noel Vittaut, Nicolas Chesneau and Marie-jeanne Lesot | 221 |
| <i>Privacy- and Utility-Preserving NLP with Anonymized data: A case study of Pseudonymization</i> | |
| Oleksandr Yermilov, Vipul Raheja and Artem Chernodub | 232 |
| <i>GPTs Don't Keep Secrets: Searching for Backdoor Watermark Triggers in Autoregressive Language Models</i> | |
| Evan Lucas and Timothy Havens | 242 |
| <i>Make Text Unlearnable: Exploiting Effective Patterns to Protect Personal Data</i> | |
| Xinzhe Li and Ming Liu | 249 |
| <i>Training Data Extraction From Pre-trained Language Models: A Survey</i> | |
| Shotaro Ishihara | 260 |
| <i>Expanding Scope: Adapting English Adversarial Attacks to Chinese</i> | |
| Hanyu Liu, Chengyuan Cai and Yanjun Qi | 276 |
| <i>IMBERT: Making BERT Immune to Insertion-based Backdoor Attacks</i> | |
| Xuanli He, Jun Wang, Benjamin Rubinstein and Trevor Cohn | 287 |
| <i>On The Real-world Performance of Machine Translation: Exploring Social Media Post-authors' Perspectives</i> | |
| Ananya Gupta, Jae Takeuchi and Bart Knijnenburg | 302 |
| <i>Enabling Classifiers to Make Judgements Explicitly Aligned with Human Values</i> | |
| Yejin Bang, Tiezheng Yu, Andrea Madotto, Zhaojiang Lin, Mona Diab and Pascale Fung ... | 311 |
| <i>Strength in Numbers: Estimating Confidence of Large Language Models by Prompt Agreement</i> | |
| Gwenyth Portillo Wightman, Alexandra Delucia and Mark Dredze | 326 |