

Workshop on Natural Language Processing for Internet Freedom (NLP4IF 2019)

Censorship, Disinformation, and
Propaganda

Hong Kong, China
4 November 2019

ISBN: 978-1-7138-0086-6

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© (2019) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2020)

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>Assessing Post Deletion in Sina Weibo: Multi-modal Classification of Hot Topics</i> Meisam Navaki Arefi, Rajkumar Pandi, Michael Carl Tschantz, Jedidiah R. Crandall, King-wa Fu, Dahlia Qiu Shi and Miao Sha	1
<i>Detecting context abusiveness using hierarchical deep learning</i> Ju-Hyoung Lee, Jun-U Park, Jeong-Won Cha and Yo-Sub Han	10
<i>How Many Users Are Enough? Exploring Semi-Supervision and Stylometric Features to Uncover a Russian Troll Farm</i> Nayeema Nasrin, Kim-Kwang Raymond Choo, Myung Ko and Anthony Rios	20
<i>Identifying Nuances in Fake News vs. Satire: Using Semantic and Linguistic Cues</i> Or Levi, Pedram Hosseini, Mona Diab and David Broniatowski	31
<i>Calls to Action on Social Media: Detection, Social Impact, and Censorship Potential</i> Anna Rogers, Olga Kovaleva and Anna Rumshisky	36
<i>Mapping (Dis-)Information Flow about the MH17 Plane Crash</i> Mareike Hartmann, Yevgeniy Golovchenko and Isabelle Augenstein	45
<i>Generating Sentential Arguments from Diverse Perspectives on Controversial Topic</i> ChaeHun Park, Wonsuk Yang and Jong Park	56
<i>Rumor Detection on Social Media: Datasets, Methods and Opportunities</i> Quanzhi Li, Qiong Zhang, Luo Si and Yingchi Liu	66
<i>Unraveling the Search Space of Abusive Language in Wikipedia with Dynamic Lexicon Acquisition</i> Wei-Fan Chen, Khalid Al Khatib, Matthias Hagen, Henning Wachsmuth and Benno Stein	76
<i>CAUnLP at NLP4IF 2019 Shared Task: Context-Dependent BERT for Sentence-Level Propaganda Detection</i> Wenjun Hou and Ying Chen	83
<i>Fine-Grained Propaganda Detection with Fine-Tuned BERT</i> Shehel Yoosuf and Yin Yang	87
<i>Neural Architectures for Fine-Grained Propaganda Detection in News</i> Pankaj Gupta, Khushbu Saxena, Usama Yaseen, Thomas Runkler and Hinrich Schütze	92
<i>Fine-Tuned Neural Models for Propaganda Detection at the Sentence and Fragment levels</i> Tariq Alhindi, Jonas Pfeiffer and Smaranda Muresan	98
<i>Divisive Language and Propaganda Detection using Multi-head Attention Transformers with Deep Learning BERT-based Language Models for Binary Classification</i> Norman Mapes, Anna White, Radhika Medury and Sumeet Dua	103
<i>On Sentence Representations for Propaganda Detection: From Handcrafted Features to Word Embeddings</i> André Ferreira Cruz, Gil Rocha and Henrique Lopes Cardoso	107

<i>JUSTDeep at NLP4IF 2019 Task 1: Propaganda Detection using Ensemble Deep Learning Models</i> Hani Al-Omari, Malak Abdullah, Ola AlTiti and Samira Shaikh	113
<i>Detection of Propaganda Using Logistic Regression</i> Jinfen Li, Zhihao Ye and Lu Xiao	119
<i>Cost-Sensitive BERT for Generalisable Sentence Classification on Imbalanced Data</i> Harish Tayyar Madabushi, Elena Kochkina and Michael Castelle	125
<i>Understanding BERT performance in propaganda analysis</i> Yiqing Hua	135
<i>Pretrained Ensemble Learning for Fine-Grained Propaganda Detection</i> Ali Fadel, Ibrahim Tuffaha and Mahmoud Al-Ayyoub	139
<i>NSIT@NLP4IF-2019: Propaganda Detection from News Articles using Transfer Learning</i> Kartik aggarwal and Anubhav Sadana	143
<i>Sentence-Level Propaganda Detection in News Articles with Transfer Learning and BERT-BiLSTM-Capsule Model</i> George-Alexandru Vlad, Mircea-Adrian Tanase, Cristian Onose and Dumitru-Clementin Cercel	150
<i>Synthetic Propaganda Embeddings To Train A Linear Projection</i> Adam Ek and Mehdi Ghanimifard	157
<i>Findings of the NLP4IF-2019 Shared Task On Fine-grained Propaganda Detection</i> Giovanni Da San Martino, Alberto Barrón-Cedeño and Preslav Nakov	162