LREC 2020 Workshop
Language Resources and Evaluation Conference
11–16 May 2020

First International Workshop on Social Threats in Online Conversations: Understanding and Management

PROCEEDINGS

Editors:
Archna Bhatia and Samira Shaikh

Proceedings of the LREC 2020 First International Workshop on Social Threats in Online Conversations: Understanding and Management

Edited by: Archna Bhatia and Samira Shaikh

ISBN: 979-10-95546-39-9
EAN: 9791095546399

For more information:
European Language Resources Association (ELRA)
9 rue des Cordelières
75013, Paris
France
http://www.elra.info
Email: lrec@elda.org

Introduction

Welcome to the LREC 2020 Workshop on Social Threats in Online Conversations (STOC): Understanding and Management. The First STOC workshop was accepted to be held in conjunction with LREC 2020 (The 12th Edition of Language Resources and Evaluation Conference). Motivated by the need of using natural language processing (NLP) and computational sociolinguistics techniques in conjunction with metadata analysis, that can provide a better means for detecting and countering social engineering (SE) attacks and disinformation campaigns in a wide variety of online, conversational contexts, the main goal of the workshop was to glean actions and intentions of adversaries in online conversations, from the adversaries' language use coupled with communication content.

The organizing committee consisted of Archna Bhatia, Adam Dalton, Bonnie J. Dorr, Samira Shaikh and Tomek Strzałkowski. We solicited papers from a wide range of disciplines, including cybersecurity, NLP, computational sociolinguistics, Human-Computer Interaction and Psychology. We received a total of nine papers and accepted eight of these. Each paper was reviewed by at least three reviewers, two of the papers were reviewed by four reviewers to arrive at a decision.

The eight accepted papers dealt with a wide range of topics related to Social Threats online. In "The Panacea Threat Intelligence and Active Defense Platform", Dalton et al. describe a system that supports NLP components, including Ask and Framing detection, Named Entity Recognition, Dialogue Engineering and Stylometry for active defenses against SE attacks. The novelty of this system is in engaging the SE attacker using bots to waste the attacker's time and resources.

Bhatia et al. develop a paradigm for extensible lexical development based on Lexical Conceptual Structure in "Adaptation of a Lexical Organization for Social Engineering Detection and Response Generation". The paradigm supports resource extensions for new applications such as SE detection and response generation. Authors demonstrate that their proposed lexical organization refinements improve ask/framing detection and top ask identification, and yield qualitative improvements in response generation for defense from SE.

Kim et al. describe an automated approach to determine if a participant in online conversation is a cyberpredator in "Analysis of Online Conversations to Detect Cyberpredators Using Recurrent Neural Networks". Their experiments using recurrent neural networks on two datasets demonstrate that their approach provides better recall than prior, similar approaches.

Abeywardana and Thayasivam present their results to apply traditional structured privacy preserving techniques on unstructured data, including Twitter messages in their paper titled "A Privacy Preserving Data Publishing Middleware for Unstructured, Textual Social Media Data". They also make available a corpus of tweets that have been annotated for privacy related attributes.

In the paper "Information Space Dashboard", the authors Krumbiegel, Pritzkau and Schmitz created a dashboard that supports an analyst in generating a common operational picture of the information space, link it with an operational picture of the physical space and, thus, contribute to overarching situational awareness. They demonstrate their method on the analysis of historical data regarding violent anti-migrant protests and respective counter-protests that took place in Chemnitz in 2018.

Blackburn et al. in the paper "Corpus Development for Studying Online Disinformation Campaign: A Narrative + Stance Approach" discuss an end-to-end effort towards developing a corpus for studying disinformation campaigns across platforms, focusing on the Syrian White Helmets case study. They focused on the most challenging annotation tasks and an exploration of automated methods to extract narrative elements from microblogs.

Pascucci et al. describe their web service platform for disinformation detection in hotel reviews written in English in “Is this hotel review truthful or deceptive? A platform for disinformation detection through computational stylometry”. Using the corpus of Deceptive Opinion Spam corpus, consisting of hotel reviews in four categories positive truthful, negative truthful, positive deceptive and negative deceptive reviews, the authors investigated four classifiers and demonstrated that Logistic Regression is the best performance algorithm for disinformation detection.

In “Email Threat Detection Using Distinct Neural Network Approaches”, Castillo et al. use neural networks to detect malicious content in email interactions. Their goal is to obtain highly accurate detection of malicious emails based on email text alone. Their results show that back-propagation both with and without recurrent neural layers outperforms current state of the art techniques that include supervised learning algorithms with stylometric elements of texts as features.

The STOC workshop at LREC 2020 is cancelled due to Covid-19 pandemic, but these proceedings touch upon the research being conducted to study various dimensions of social threats in online conversations through techniques involving NLP, machine learning, computational sociolinguistics, and stylometry. We hope that this will provide a ground for future discussions and follow up workshops on topics related to social threats and SE.

Organizers:

Archna Bhatia, Institute for Human and Machine Cognition
Adam Dalton, Institute for Human and Machine Cognition
Bonnie J. Dorr, Institute for Human and Machine Cognition
Samira Shaikh, University of North Carolina at Charlotte
Tomek Strzalkowski, Rensselaer Polytechnic Institute

Program Committee:

Ehab Al-Shaer, UNCC
Genevieve Bartlett, USC-ISI
Emily Bender, U Washington
Larry Bunch, IHMC
Esteban Castillo, RPI
Dave DeAngelis, USC-ISI
Mona Diab, GWU/Google
Sreekar Dhaduvai, SUNY Albany
Min Du, UC Berkeley
Maxine Eskenazi, CMU
William Ferguson, Raytheon
Mark Finlayson, FIU
Marjorie Freedman, USC-ISI
Bryanna Hebenstreit, SUNY Albany
Christopher Hidey, Columbia
Scott Langevin, Uncharted
Christian Lebiere, CMU
Kristina Lerman, USC/ISI
Fei Liu, UCF
Amir Masoumzadeh, SUNY Albany
Kathleen McKeown, Columbia
Alex Memory, Leidos
Chris Miller, SIFT
Mark Orr, University of Virginia
Ian Perera, IHMC
Alan Ritter, OSU
Emily Grace Saldanha, PNNL
Sashank Santhanam, UNCC
Sonja Schmer-Galunder, SIFT
Svitlana Volkova, PNNL
Ning Yu, Leidos
Zhou Yu, UC Davis
Alan Zemel, SUNY Albany

Invited Speakers:

Rosanna E. Guadagno, Stanford University
Ian Harris, University of California Irvine

Table of Contents