

Abusive Language Detection: Too much Digital, not enough Humanities?

Sara Tonelli

Fondazione Bruno Kessler
Via Santa Croce, 77
38122 Trento, Italy
satonelli@fbk.eu

Abstract

Social media messages are often written to attack specific groups of users based on their religion, ethnicity or social status, and they can be particularly threatening to vulnerable users such as teenagers. While it is important to develop tools and algorithms that automatically detect abusive messages, so that they can be deleted and hateful accounts can be blocked, it is also important to understand the shared language used by online communities, to what extent offensiveness can be generalized, and what is the role of dataset annotation in avoiding biases in online abuse detection. This talk gives an overview of current state-of-the-art approaches in abusive language detection and discusses ongoing work to better understand this phenomenon, analysing the role of communities, textual context and stakeholders' involvement.

Bio

Sara Tonelli is the head of the Digital Humanities research group at Fondazione Bruno Kessler in Trento, and adjunct professor of Language Interfaces at the Dept. of Psychology and Cognitive Science, University of Trento. She got a PhD in Language Sciences from Università Ca' Foscari in Venice in 2010, after which she joined FBK first as a post-doc in the Natural Language Processing group and then as a tenured researcher leading the newly founded DH group. She has been involved in several national and European projects dealing with historical archives, event and temporal information extraction and political stance detection and, more recently, with social media monitoring and hate speech detection. Her research interests are highly interdisciplinary, trying to apply and adapt advanced text analysis approaches to social sciences and historical investigation.