Detecting Fake News by Leveraging Emotion

Johnny Scanlon

April 2022

Abstract

Fake news has become an increasingly salient issue in the past number of years. The advent of the Internet and subsequent rise of social networks have provided a powerful platform for the rapid sharing of information. With this, the issue of widespread *misinformation* has become a serious problem.

Past work has shown that there is often a distinct emotional component to fake news. This project seeks to exploit this by explicitly leveraging information about the emotional content of news items in order to detect whether the news is "fake". The novel use of deepMoji is proposed as a means of detecting fake news. This architecture classifies emotions in text by returning probabilities of specific emojis being associated with the text. Emojis themselves are a powerful tool for conveying nuanced emotion across digital platforms. This project uses this to detect the differences in emotions conveyed between fake and real news.

Their use is examined in two specific ways. The first is to see how well these emotional features can be used to discriminate between real and fake news. The second is to see if they can be used to improve the results of popular state-of-the-art transformer models like BERT.

Through testing on the LIAR data set and a data set released as part of the AAAI Constraint-2021 shared task, it is found that these emotional features on their own are capable of discriminating between real and fake news. This confirms that there are distinct and measurable differences in the emotions typically conveyed in real and fake news. Positive results were found by combining this with BERT. However, it was not possible to confirm the statistical significance of this improvement.