

Personalized Models for Fake News Detection

Master/Bachelor Thesis

Motivation

Detecting fake news online is an important and timely issue. We are interested in investigating personal differences in how fake news is spread and developing models to detect fake news. Our group has collected data for this task and for thousands of Reddit users. This can be used to test models for personalization from embeddings, language models, priming, or others. These models have been not been used for many downstream applications and we are interested to find out if these can be more widely applied and how performance differs on downstream tasks. Each method requires a different amount of computation and memory so there are trade-offs to consider.

Difficulty

Analysis



Programming



Literature



Task Description



Our goal is to show that fake news can be more accurately detected when using personalized representations and we expect to see that the performance increases proportionally to the level of detail in the representation. If we can find other tasks where we expect personalization to help and also have a dataset with a reasonable volume of data per user, we can run experiments with those as well (e.g. hate speech, sarcasm). Work can be submitted for publication upon completion.

Contact

Prof. Dr. Lucie Flek
Dr. Charlie Welch

✉ lucie.flek@gmail.com

✉ cwelch@mathematik.uni-marburg.de

🌐 caisa-lab.github.io

☎ 06421 / 28 - 21586

📞 06421 / 28 - 28902

🏠 CAISA Lab FB 12 |

Mathematik und Informatik

📍 Hans-Meerwein-Straße 6,
35043, Marburg

Room: 04C21 (Staircase C,
Level 4)

References

- [1] Charles Welch, Jonathan K. Kummerfeld, Verónica Pérez-Rosas, and Rada Mihalcea. Compositional demographic word embeddings. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, November 2020.
- [2] Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. Misinformation has high perplexity. *CoRR*, abs/2006.04666, 2020.