



An unsupervised framework for tracing textual sources of moral change

Aida Ramezani, Zining Zhu, Frank Rudzicz, Yang Xu

University of Toronto

Findings of the Association for Computational Linguistics: EMNLP 2021

Motivation

Morality plays an important role in our well-being, but our perception of morality changes over time.



1997, Bill Clinton
elected for his
second term
presidency.

1998, Bill Clinton
impeached for
scandal.



2001, 58% of
Americans think
Bill Clinton is
NOT honest and
trustworthy*!



Language is known to be an important factor to construct our moral concerns, and change them.

<http://www.cnn.com/2001/ALLPOLITICS/stories/01/10/cnn.poll.clinton/> (CNN/USA Today/Gallup poll)

Related work


NLP-based approaches to moral inference from text:

- Classification of moral sentiments in short text
 - Moral sentiments in social media posts, and political speeches over time (Garten et al., 2016)
 - Using background knowledge (Lin et al., 2018)
 - Moralization in social media and violence in protests (Mooijman et al., 2018)
- Moral sentiment changes in different moral concepts over time (e.g., *slavery*) (Xie et al., 2019)


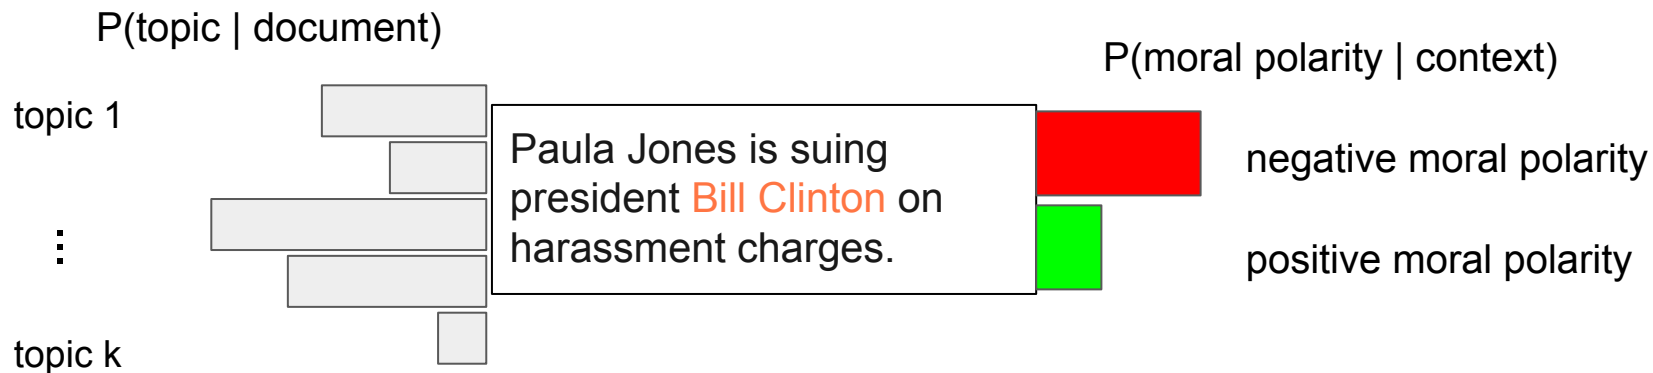
Research problem

What are the textual sources of moral sentiment change, e.g., for entities like political leader?

Example: Why did the moral sentiment toward Bill Clinton become increasingly negative during 1998?
Source: The Clinton-Lewinsky scandal.



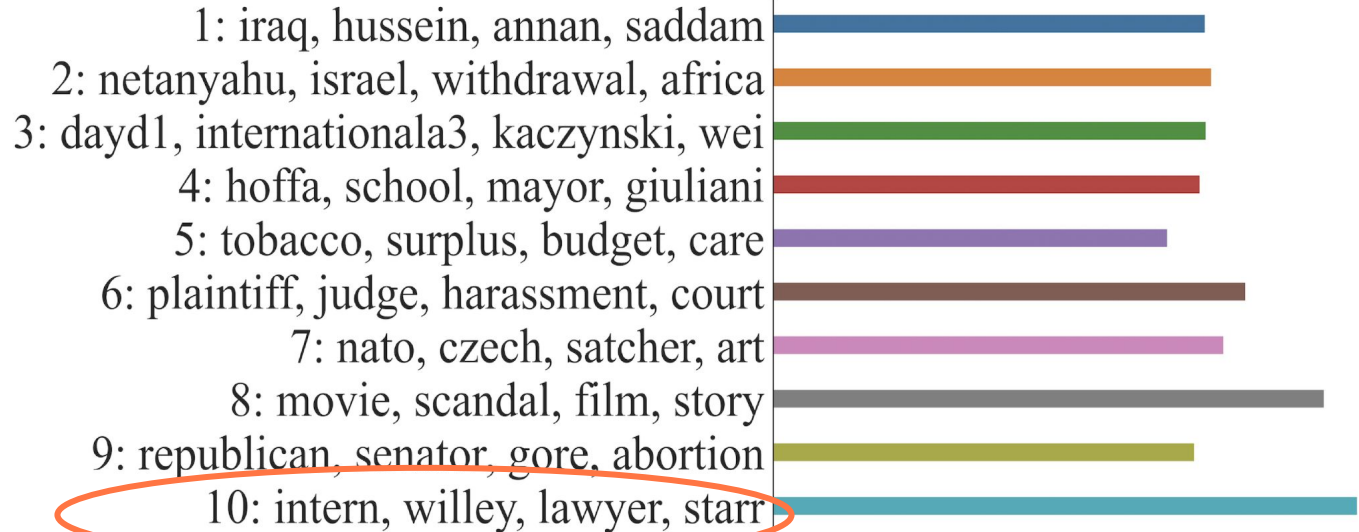
Our framework




Highlights of our results

Identifying salient topics as sources of moral sentiment change for an entity (*Bill Clinton*)

Topics



Application to COVID-19 pandemic (entity = *Donald Trump*)



Conclusion

We developed an unsupervised topic-based approach to identify sources of entities' moral sentiment change in text, and showed the utility of our framework in interpreting moral changes perceived in news.

Thank you!

Contact: armzn@cs.toronto.edu

An unsupervised framework for tracing textual sources of moral change

Aida Ramezani^{1*}, Zining Zhu^{1,2*}, Frank Rudzicz^{1,2,3}, Yang Xu^{1,2,4}

¹ Department of Computer Science, University of Toronto, Toronto, Canada

² Vector Institute for Artificial Intelligence, Toronto, Canada

³ St. Michael's Hospital, Toronto, Canada

⁴ Cognitive Science Program, University of Toronto, Toronto, Canada

{armzn, zining, frank, yangxu}@cs.toronto.edu

Published in Findings of the Association for
Computational Linguistics: EMNLP 2021