


The language of online commentary: Extracting information from news comments

SFU

Discourse Processing Lab

Luca Cavasso • Jennifer Fest • Emilie Francis • Varada Kolhatkar • Maite Taboada • Hanhan Wu



Online news comments

News comments

- Readers' reactions posted after a news article
- Often informative, part of a dialogue – **constructive**
- Sometimes offensive, abusive, derogatory – **toxic**

Media and technology

- Toronto Star decides to close comments as of December 2015 [1]
- CBC News does not allow comments for stories on indigenous people [2]
- New Yorker: psychological toll on human monitors [3]
- Google: Jigsaw, devoted to helping monitor comments [4]

theguardian

The dark side of Guardian comments

As part of a series on the rising global phenomenon of online harassment, the Guardian commissioned research into the 70m comments left on its site since 2006 and discovered that of the 10 most abused writers eight are women, and the two men are black. Hear from three of those writers, explore the data and help us host better conversations online

SOCC: The SFU Opinion and Comments Corpus

The Globe and Mail

- All opinion articles + all comments
- 5 years: January 2012 – December 2016
- 10,339 articles – 1,280,454 comments
- 6,895,696 words (articles) + 77,238,179 words (comments)



Annotation

1. Appraisal [5]

Attitude { Affect
Appreciation
Judgement

Polarity { Positive
Negative

The shallowness of this debate on government finances is depressing

Affect, Negative

Great piece! Thanks.

Appreciation, Positive

Why does the Globe publish this kind of trash?

Judgement, Negative

Clinton lost, not because she's a woman but because she was perceived to be an establishment candidate.

2. Negation & scope [6]

not <she's a woman>

3. Nonveridicality [7]

perceived to be

4. Coherence relations [8]

Cause (because); Contrast (but)

Data and annotations to be made publicly available

Current work: Constructiveness & toxicity

Which comment is constructive?



MARIE HENEIN

Thank you, Hillary. Now women know retreat is not an option

MARIE HENEIN

Special to The Globe and Mail
Published Thursday, Nov. 10, 2016 4:46PM EST
Last updated Friday, Nov. 11, 2016 1:02PM EST

Constructive comments

- Survey of 100 people through SurveyMonkey
- What does 'constructive' mean?

- Composite of the answers: Constructive comments intend to **create a civil dialogue** through remarks that are **relevant** to the article and not intended to merely provoke an emotional response. Typically targeted to **specific points** and supported by **appropriate evidence**

Operationalizing constructiveness

- Argumentation [9] Constructive comments more likely to
 - Contain modal verbs and argumentative connectives
- Rhetorical relations [10]: Constructive comments more likely to contain persuasive and presentational relations
 - Cause, Comparison, Condition, Contrast, Explanation

Crowd annotation

- CrowdFlower. 1,200 comments from 10 articles. Labels:
 - Constructive vs. Non-constructive
 - **Not toxic** – **Mildly toxic** – **Toxic** – **Very toxic**
- Results: **no difference** in toxicity between constructive and non-constructive comments

1

I have 3 daughters, and I told them that Mrs. Clinton lost because she did not have a platform. The only message that I got from her was that Mr. Trump is not fit to be in office and that she wanted to be the first female President. I honestly believe that she lost because she offered no hope, or direction, to the average American. Mr. Trump, with all his shortcomings, at least offered change and some hope.

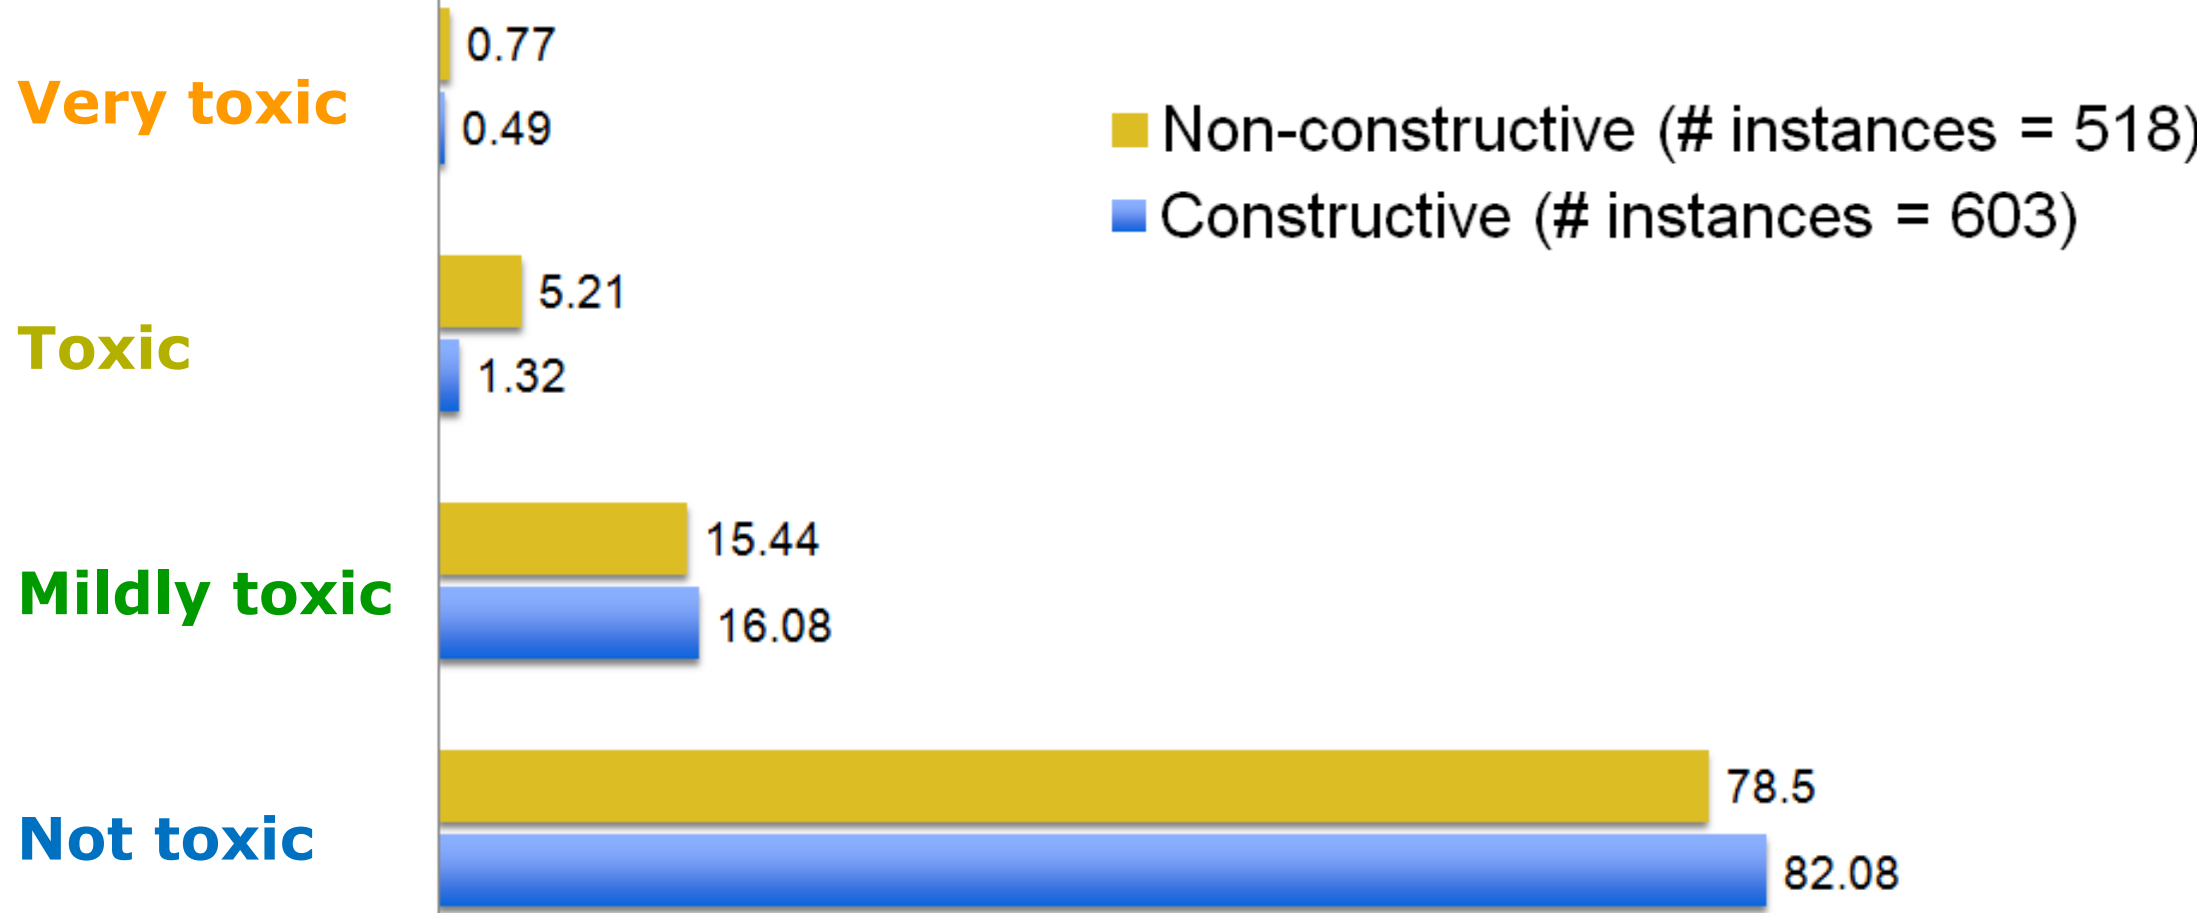
2

This article was a big disappointment. Thank you Ms Henein. Now women know that wasting their time reading your emotion-based opinion piece is not an option.



Operationalizing toxicity

- Toxic comments are likely to offend or cause distress
- Personal attacks, harassment, profanity, abuse, hate speech [11, 12]
- Sentiment analysis [13]: toxic comments contain (extreme) negative sentiment



Future work

- Register analysis [14]:
 - Are comments more like conversations or monologues?
- Topic modelling in articles and comments [15]:
 - What is more frequently discussed in articles?
 - Which article topics are discussed in the comments?
 - Which topics have more/fewer replies?
 - Link to sentiment

Applications

- Comment organization: sort by topic and sentiment
- Filtering: delete toxic comments
- Public opinion: what topics are discussed the most?

References

- [1] Cooke, M. (2015). Toronto Star closes commenting on thestar.com. *Toronto Star*, December 16, 2015. <http://www.thestar.com/news/2015/12/16/toronto-star-closing-commenting-on-thestarcom.html>
- [2] Uncivil dialogue: Commenting and stories about indigenous people. (2015). *CBC News*, November 30, 2015. <http://www.cbc.ca/newsblogs/community/editorsblog/2015/11/uncivil-dialogue-commenting-and-stories-about-indigenous-people.html>
- [3] Chen, A. (2017). The human toll of protecting the Internet from the worst of humanity. *The New Yorker*, January 28, 2017. <http://www.newyorker.com/tech/elements/the-human-toll-of-protecting-the-internet-from-the-worst-of-humanity>
- [4] <https://jigsaw.google.com/>
- [5] Martin, J. R., & White, P. R. R. (2005). *The Language of Evaluation*. New York: Palgrave.
- [6] Blanco, E., & Moldovan, D. (2013). Retrieving implicit positive meaning from negated statements. *Natural Language Engineering*, 20(4), 501-535.
- [7] Giannakidou, A. (2002). Licensing and sensitivity in polarity items: From downward entailment to nonveridicality. In M. Andronis, A. Pycha & K. Yoshimura (Eds.), *Papers from the 38th Meeting of the Chicago Linguistics Society*. Chicago: CLS.
- [8] Taboada, M., & Mann, W. (2006). Rhetorical Structure Theory: Looking back and moving ahead. *Discourse Studies*, 8(3), 423-459.
- [9] Stede, M. (2016). Towards assessing depth of argumentation. *Proceedings of the 26th International Conference on Computational Linguistics* (pp. 3308-3317). Osaka.
- [10] Peldszus, A., & Stede, M. (2016). Rhetorical structure and argumentation structure in monologue text. *Proceedings of the 3rd Workshop on Argument Mining, ACL* (pp. 103-112). Berlin.
- [11] Nobata, C., Tetreault, J., Thomas, A., Mehdad, Y., & Chang, Y. (2016). Abusive language detection in online user content. *Proceedings of the 25th International Conference on World Wide Web* (pp. 145-153). Montréal.
- [12] Warner, W., & Hirschberg, J. (2012). Detecting hate speech on the World Wide Web. *Proceedings of the Second Workshop on Language in Social Media* (pp. 19-26). Montréal.
- [13] Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2), 267-307.
- [14] Biber, D. (1988). *Variation across Speech and Writing*. Cambridge: Cambridge University Press.
- [15] Blei, D. M. (2012). Probabilistic topic models. *Communications of the ACM*, 55(4), 77-84.

Acknowledgements

Thanks to Emma Mileva, Rada Trnavac, Silvia Vázquez and Kazuki Yabe for their help in collecting data and testing crowd questions. Thanks to Yue Wang for the poster template.

Funding: SSHRC Insight Grant ("Opinion, subjectivity and point of view in language"); NSERC Discovery Grant ("A computational treatment of negation and speculation in natural language").

Future work