RHETORiC: an Audience Conversation Tool that Restores Civility in News Comment Sections

Emilie Bossens
Meaningful Interactions Lab, Institute for Media Studies, KU Leuven
emilie.bossens@kuleuven.be

David Geerts
Digital Society Institute, KU Leuven
david.geerts@kuleuven.be

Elias Storms
University of Antwerp
elias.storms@uantwerpen.be

Michiel Nuytemans
Tree Company
michiel@treecompany.be

Jan Boesman
Institute for Media Studies, KU Leuven, and Tree Company
jan.boesman@kuleuven.be

ABSTRACT

We present a case study where we developed RHETORiC: an audience conversation tool that promotes civil participation in online news comments. By following a human-centered design process, we created and evaluated a novel tool with an interface that supports people in carefully formulating their opinions and arguments, so they can constructively contribute to online discussions about news. Results of a large-scale field study show that our tool succeeds in increasing the level of civility, argumentation and proficiency of comments in comparison with those on social media. On top, both users and journalists report high satisfaction with the RHETORiC tool. In our paper we reflect upon lessons learned about the design process as well as about the tool itself, which contributes to the fields of both Human-Computer Interaction and Digital Journalism.

CCS CONCEPTS

• **Human-centered computing** → Collaborative and social computing; Empirical studies in collaborative and social computing.

KEYWORDS

digital journalism, news comments, online debate, audience engagement

ACM Reference Format:
Emilie Bossens, David Geerts, Elias Storms, Michiel Nuytemans, and Jan Boesman. 2022. RHETORiC: an Audience Conversation Tool that Restores Civility in News Comment Sections. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '22 Extended Abstracts), April 29–May 05, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3491101.3503560

1 INTRODUCTION

Enabling audience participation in debate on news and current affairs increases trust in the media and is a necessary condition for a healthy democracy [7]. Commenting on news is an important form of public engagement, which can result in more civic engagement with public issues [2]. Readers use news comments to shape their idea of what other people think about current topics. For media outlets, on the other hand, audience participation can provide journalistic opportunities in terms of feedback or input for journalistic purposes [7]. However, in recent years, disrespectful and offensive language has come to dominate online comment sections. This has led to a pessimistic view on news comments, both from an audience and a media perspective. Such uncivil commentary on news media platforms can negatively affect readers and news producers, and lead to a diminished perception of an article's quality [11] and news publishers being urged to close down their online comment sections because the effort to manage the comments is too high [7]. Nonetheless, the closure of comment sections in digital journalism has not stopped the desire of people to comment on news articles, and these audience reactions have moved to social media such as Facebook and Twitter, where the online debate is even more polarized and toxic. Since news publishers believe in the value of readers' opinions and its potential input for journalistic editorials, they prefer to host constructive discussions about their content on their own platforms. Consequently, several attempts at enabling more respectful dialogue in news comments sections have been undertaken. For example, NRKbeta, a Norwegian media website, experimented with a quiz that readers needed to pass before they were allowed to leave a comment [12]. However, it is unclear if it really led to more civilized comments. The discrepancy between closed comment sections of online news publishers and an audience that desires to engage with news, shows that there is a clear need for a space where readers can safely share their opinion about news, and thus for solutions to improve the quality of user comments. As reactive interventions do not seem to have the desired effect, design-based proactive interventions could offer better chances of success. In this paper, we present the RHETORiC[1] tool, which guides commenters through the commenting process using conversational prompts, stance taking, live feedback, and clustering of similar comments, in order to improve civility in online news comment sections.

CHI '22 Extended Abstracts, April 29–May 05, 2022, New Orleans, LA, USA

ACM ISBN 978-1-4503-9156-6/22/04.
https://doi.org/10.1145/3491101.3503560

[1] An acronym for Reducing Hate through Editorial Tools for Online Reactions and Comments

2 BACKGROUND

Techniques to improve the quality of user comments in digital journalism have been extensively studied in the past years. These techniques can be distinguished in terms of reactive and pro-active approaches. Among reactive approaches, the most prominent technique is content moderation. Such moderation can take the form of pre-moderation or post-moderation, and it can be carried out by humans or by AI [7]. Examples of reactive interventions are automatically detecting and removing toxic comments or involving the online community in removing 'bad actors' from a discussion [4]. Pro-active attempts, by contrast, focus more on design-based interventions that aim to prevent hateful comments and encourage commenters to contribute more qualitatively. Techniques that have proven to be successful for a more constructive debate are altering the structure of the comment section or the 'discourse architecture' [6][10], providing a short summary to highlight relevant facts or contrasting information [10], using specific CAPTCHAs to prime positive emotions [14], including conversational prompts [15][19] and presenting contextual cues [8][9]. As the focus of our current research is primarily on stimulating a civil debate, we stick to an overview of relevant work on pro-active interventions.

When a discussion topic is precisely defined, there is a higher chance that the comments will be more constructive [18]. Using a conversational prompt preceding a discussion, such as posing specific questions, the better targeted the response will be [3]. This is confirmed by a more recent study based on a content analysis of 11,218 user comments on Facebook pages of established German news media outlets. The researchers found that conversational prompts in news article posts on Facebook led to more civil and rational comments [19]. Another, more substantial, design-based technique to improve comment quality on news sites is through moving away from a traditional comment section, i.e. a text field only and comments organized chronologically, to a more visually structured design. So-called debate tools, such as Kialo and consider.it, use a format that enables people to engage in thoughtful discussions on controversial topics, to understand different points of view and to visualize pro and contra arguments. Such structured discussion formats can be helpful in encouraging users to base their opinion on thoughtful reasoning and may serve a framing role, affecting the likelihood of extensive reasoning [16]. Prior studies found that when comments are shown in two [16] or three [10] columns with contrasting viewpoints, the comments were of higher deliberative quality compared to comments in a one-column format. The interface Opinion Space goes even a step further, by clustering and projecting opinions in a two-dimensional environment. Points close by correspond to similar opinions, points far away correspond to very different views. A user test revealed that participants agreed more with and had more respect for comments they read via Opinion Space compared to what they read via the chronological list-based interface [5]. In conclusion, well-designed interactions, conversational prompts and discourse architecture can shape discussions for the better, an approach we also take with our RHETORiC tool.

3 METHODS

We followed a human-centered design process, involving various stakeholders in iterative cycles of analyzing user needs, designing and testing prototypes, and ending with a large-scale evaluation in a real-life context.

3.1 Analyzing user needs

First, we organized a multi-stakeholder ideation workshop involving 17 participants, including media professionals, HCI experts and social media users, to come up with ideas for interface elements to increase the civility on online discussions. The workshop took place in November 2019 and lasted for four hours. During the workshop, participants were introduced to the topics of polarization, cognitive biases and argumentation skills, using sensitizing exercises. Afterwards, they were distributed over four groups and asked to generate ideas for interface elements that stimulate a constructive debate using a structured template. After generating the ideas, two groups each were paired and asked to present their ideas to the other group which the other group critiqued in the form of 'roses' (positive elements), 'buds' (elements with potential) or 'thorns' (drawbacks or pitfalls).

In parallel with the ideation workshop, we conducted diary studies followed by interviews with 8 participants (4 male, 4 female, average 39 years old) to study the experiences and reflections from commenters who actively participate in online discussions. We were mainly interested in their motivations and inhibitions for reading and writing comments. In preparation for the interview, we asked participants to keep two diaries for the duration of one week: one for comments they read, another for comments they wrote. For each diary, participants were asked to screenshot comments and answer short contextual questions, which served as material for the in-depth, semi-structured interviews. We probed the reasons why they had selected comments they had encountered, and asked about their decisions and motivations to write comments and what they expect to get out of their participation. The interviews generally lasted between an hour and an hour and a half. All interviews were transcribed and analysed using affinity diagramming.

Lastly, with the goal of defining engagement strategies, we conducted 9 interviews with engagement editors, community moderators, and heads of innovations from media organizations in Flanders and the Netherlands. The interviews lasted between half an hour and two hours, and were analyzed using thematic analysis. In addition, we reviewed relevant literature and studied good practices worldwide. Examples of the analyzed cases are Coral, Hearken, De Correspondent, and City Bureau. Based on the insights from the interviews, literature and best practices, we defined various strategies to engage a broad audience in online discussions.

3.2 Design and testing

In the design phase, we first conducted an online experiment with 255 participants to empirically test whether two novel interface design elements, namely a discussion statement and a way to indicate an opinion, can improve the quality of online comments. By formulating a specific statement in the comment section, we aimed to encourage users to comment more on-topic and to discourage them to expand on controversial details. The opinion element also acts as

an anchor that guides participants through the commenting process. We explored the impact of both elements in a between-subjects design with four conditions. The control condition was similar to a traditional comment section design, i.e., a text field only. In the three experimental conditions a discussion statement was added, either alone or in combination with agree-disagree buttons or an agree-disagree scale. Participants were requested to comment on a (controversial) news article and to complete a series of questions, including the short user experience questionnaire [13]. To measure the quality of participants' comments and study variations between the conditions, we coded all comments on five dimensions: argumentation, civility, relevance, orthography and length. For more details about this experiment, see [1].

Next, during a structured workshop with HCI and media professionals using Miro, we co-designed the main use case of our audience conversation tool, describing the different steps in the interaction between the user and the system. This formed the input for the first prototype which we created in Figma. Via individual video calls using the walk-through approach, the design was evaluated with five participants. The interviewer asked the participant questions about the buttons and features while going through the prototype. Participants were explicitly asked to think out loud throughout the walk-through. Additionally, the interviewer asked questions such as "what should you do here?", "where would you click on' and 'what do you expect to happen next?'.

3.3 Large-scale evaluation in a live environment

In the third and final research phase, the iterative development of the RHETORiC tool concluded with a small- and large-scale test in a live environment. We kicked off with a pilot study, where we placed the tool in a closed environment and shared the link with family and friends. In total, 58 contributions were collected. Via an online questionnaire, participants' feedback about the tool was gathered. Based on the feedback, we made small adjustments to the tool. Subsequently, we conducted a long-term, real-life test aimed at a final evaluation of the design with participation of a broad audience. The live test ran on the website of Het Nieuwsblad, a national publisher, during six weeks in the summer of 2021. Journalists added the RHETORiC component to nine online articles about a range of themes. In total, 1733 contributions from readers were collected. Seven of the nine articles with the audience conversation tool were shared on the Facebook page of the news publisher. Both the comments submitted via the new tool and Facebook were analyzed using a text analytics API in terms of argumentation, proficiency, toxicity and sentiment. To get more insight into journalists' and users' experiences, we conducted respectively a double interview with two journalists that tested the tool and an online survey. A total of 317 participants (partially) filled in the questionnaire. We asked them about their general impression of the tool, how they experienced each different step in the flow, how they experienced the new tool in comparison with other ways of commenting (traditional comments and social media), whether they would participate in the future, and what they liked and disliked about the tool.

4 RESULTS

In this section, we will describe those results that were effectively used as input for the design of the RHETORiC tool. We will discuss the impact of our choices and some lessons learned in a separate section.

4.1 Ideas for interface elements: writing help, taking position and substantiating statements

During the ideation workshop, 34 unique ideas for interface elements were generated. After clustering the ideas based on their most salient features, we arrived at ten categories of interface elements that could improve the quality of user comments. Three of those categories were selected as they were the most promising in terms of potential impact, and feasibility of implementing it by the project partners.

The first category we selected consisted of ideas conceived as a writing help tool, meant to be used while a person is writing a comment. The common element of all ideas in this category was that through continuous analysis of the comment that is being written, the system detects argument strength, and in various ways presents this to the commenters so they can decide to adjust their comment if they want to. The second category of ideas we continued with, was the idea of letting commenters take in a clear position before writing a comment. Some ideas let commenters take in a position on either side in a staged discussion (whether or not they agree with it), while others suggest asking people to first indicate what their own position is, before letting them argue in favour of it. Finally, we also considered the idea of stimulating commenters to substantiate their statement(s). The common idea is that commenters would have to be more conscious of the statements they make, and make sure they have supporting arguments or evidence to back them up.

In addition to these categories, we formulated recommendations based on the feedback on the ideas. Firstly, our participants suggested that it is important to find the right tone when designing audience conversation tools, and that we should avoid becoming too patronizing in the way we provide feedback or advice. Secondly, we would have to find the right balance between imposing or forcing certain actions on users, which might be seen as too controlling, or merely suggesting certain actions, which can be easily ignored. Thirdly, the tools should prevent as much as possible that commenters can 'game the system' or provide counter measures if they do.

4.2 Motivations for online participation: impact, appreciation and anonymity

The diary study and interviews done with people who comment online provided us with several insights regarding motivations and barriers to participate in online debates that led to concrete recommendations for designing our RHETORiC tool.

First of all, it was clear that a sense of impact motivates people to participate, and absence of impact can lead to drop-out. Sometimes, this was even conceived as a sense of 'civic duty'. Our participants said they were less likely to participate when they did not think their contribution had an impact both in terms of reach and in

terms of effect on others. This means that emphasising the impact a contribution has, can help overcome this obstacle. Another reason for not participating in online debates is when someone doesn't think they have anything to say. This can be mitigated by showing in advance what has already been said, giving both an indication of what type of contributions are useful, as well as making it easier for people to add something new to the discussion. Vice-versa, receiving appreciation (in the form of replies or likes) is an important motivator for continued participation. Contributions need to be valued and appreciated, so it's important to provide possibilities to contribute other than commenting, e.g. likes on other comments, polls, a position indicator, etc.

Finally, anonymity also plays an important role, although this is something that news providers have conflicting opinions about (see further). Writing a comment means exposing oneself, so some participants were reluctant to do so under their real name, especially when they fear either backlash or want to hide their opinion from people they know. Since the exposure resulting from participation can be an obstacle, the possibility of posting anonymous comments therefore needs to be carefully considered.

4.3 Engaging users to participate: impact, transparency, feedback and reward

As final step in the user needs analysis, based on the interviews with media professionals, we defined requirements for increasing the engagement of commenters with online discussions. When discussing reader participation with professionals, it became clear that they take a broad view on engaging their readers in online debates. While commenting is the best known and most widely used way for making people participate in online debates, news providers that practice participatory journalism employ various other forms of participation to involve readers, such as direct calls, crowdsourcing, conversations/debates, polls, live chats, reader questions, opportunities to report news or submit audiovisual material. Our interviewees confirmed that some of these approaches could also work when transferred to the world of online debates.

According to the interviewed media professionals, the key aspects to engaging audiences are impact, transparency, feedback and reward. The first aspect resonates with the motivation of the commenters we interviewed. Our interviewees stressed that you cannot expect input without creating impact. Moreover, it is important to be transparent up front about what will be done with the user's input and how the user will be visible in their input. It is also necessary to be transparent about what is expected from a good contribution, the time commitment that is being asked from users, and the length of time in which they can provide input. The last two important aspects about engaging users to participate are giving feedback on the obtained input (so they know it is actually being read or used) and reward (some) participants by putting contributions in the picture. An additional recommendation from our participants is to be selective, and not ask for input on everything, e.g. by letting people respond to a part of the article, asking specific questions or limit the input in time.

Finally, and in contrast with commenters, media professionals do not prefer anonymity in online discussions, as they believe it is one of the reasons for the increased incivility in online debates. Nevertheless, they do agree that it's important that the threshold to participate is not too high, and that it is necessary to find a good balance. We will discuss further on how we handled anonymity during the large-scale evaluation, and what we learned.

4.4 Experimental results: higher civility and relevance

As one of the ideas from the ideation workshop included taking a clear position, we wanted to test if this could indeed increase the civility of comments, among other things. The results of the controlled online experiment we set up for this revealed that the civility and relevance of participants' comments in the three novel interfaces (a discussion statement, a discussion statement with an 'agree' and 'disagree' button and a discussion statement with an opinion scale) are significantly higher compared to a traditional interface (we describe the results in more detail in [1]). This suggests that when readers have a specific statement to comment on, in particular in the context of controversial news, they are more likely to behave civil and stay on-topic than when the focus of the comment section is more open. Although our results clearly illustrate that the statement has a positive impact on civility and relevance, we found no influence on argumentation, comment length and orthography. Regarding user experience, however, we found that participants perceived the design with the statement and the opinion scale as less clear, efficient, easy and supportive compared to the traditional Interface. This could be the result of the higher degree of interaction required from the participant to move the slider, compared to writing a traditional comment. The results of the experiment motivated us to continue with the statement and incorporate this along with the 'agree' and 'disagree' buttons in the final design of our tool.

4.5 Use case and prototype

Based on the insights from the user analysis and the experiment, we developed a use case which was translated into the RHETORiC prototype in Figma. We describe the step-by-step flow based on the screenshots of Figure 1

First, users are presented with a question or discussion statement and buttons via which they can indicate their opinion (a). Next, users are requested to substantiate their opinion by writing one argument. During writing, users receive real-time feedback about the quality and language of their contribution (b). After submitting their argument, we present, based on an algorithm, similar comments to the user and ask them to indicate whether the argument is similar to one's own argument or not. This step supports the algorithm in clustering arguments in the final overview (c). On the next screen, the user is exposed to a series of divergent arguments on the topic and participants can "like" other contributions when they agree. However, participants cannot respond to each other directly (d). On the last screen, after the users enter a name (which is not mandatory), all arguments appear clustered by content similarity. Important to note, but out of scope for this paper, is that in the back-end uncivil contributions automatically go into pre-moderation. This prototype, developed in Figma, was evaluated with end-users. Based on the walk-through evaluations with the prototype, users reported two critical issues: (1) there should be a

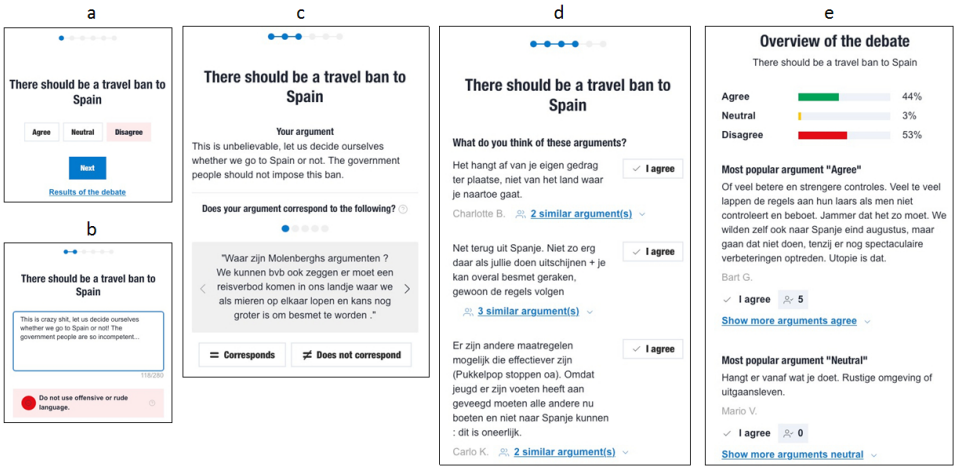

Figure 1: Step-by-step flow of the RHETORiC tool (main interface elements and example comment translated to English)

more nuanced choice in addition to the agree and disagree buttons, and (2) the detection of similar comments requires additional explanation. Promisingly, all participants expressed interest in using the tool, especially when it would manage to counter the hostile atmosphere of many current online discussions.

4.6 Results from the large-scale evaluation: increased civility, better engagement

Before reporting our results, it is important to note that most respondents participated only once via the RHETORiC tool. Our results thus reflect a first impression of this new environment by those who were willing to fill in the questionnaire, and do not include experiences of repeated use.

As our tool aims to increase civility in online discussions, arguably the most important measure is the quality of comments written using our tool in comparison with social media. We used Mann Whitney U-tests to compare differences in comment quality and, across all topics, comments in the RHETORiC tool were significantly better argued ($U=53019.5$, $p<.001$), more skillfully written ($U=55452.5$, $p<.001$), less toxic ($U=61628.5$, $p<.05$), and more positively worded ($U=59341.5$, $p<.05$) than comments on the same news article submitted via the Facebook page of the news medium. This means our tool indeed is doing what it was designed for.

As for subjective responses of the people that used the tool, more than 8 out of 10 respondents reported (very) positive feelings after using the RHETORiC tool ($M=3.94/5$). More than 90% found it worthwhile to participate in the debate using our tool and would also do so in the future. The different steps in the user flow (real-time feedback, indication of similarity of arguments and overview of the debate) were considered as clear, useful, helpful and simple. More than 90% of respondents indicated that they found the new environment (much) better for responding to news than social media such as Facebook. On our question to rank three ways of commenting on online news (the RHETORiC tool, traditional comments and social media) according to their preference, nearly 3 out of 4 participants put the RHETORiC tool first. Journalists who tested the tool in the back-end also reported positive experiences. According to them, the tool succeeds in doing what seemed impossible to them on social media: exchanging qualitative arguments and increase reader engagement.

Via open questions, participants shared what they liked and disliked about the RHETORiC tool. Firstly, participants perceived the tool as easy, clear and well-organized. Both the step-by-step flow and the structured result of the debate were appreciated. Some respondents realized that the step-by-step flow encourages reflection and a more nuanced view: "*I was almost forced to think better and deeper about the issue*" and "*You are more or less led to explain your answer.*" Other users acknowledged the added value of exposing them to diverse opinions, as illustrated by this quote: "*Because for me this was innovative and well-organized. But above all the direct link to other, and not unimportant "meaningful", opinions was great because it made me think and look at things from a different perspective. You can call this mind-blowing!*". Secondly, participants felt that the comments were more civil and on-topic compared to comments in a traditional comment section or on social media. They reported having less fear of getting "attacked" by other commenters after sharing an opinion. This goes hand in hand with the fact that readers were not able to respond to each other, which was something that several participants liked. The following quote illustrates the contrast between the new tool and other ways of commenting: "*This debate environment keeps the debate close to the essence and has a high chance to avoid completely meaningless and unsubstantiated criticism and nonsense, where on social media or classic comments, the biggest mob has the most prominent opinion. (...)*". Thirdly, another

element that often came up among respondents and which may lay at the basis of the increased civility and relevance, is the more closed environment and the greater sense of anonymity compared to social media. Users felt that the target group is smaller and consists mainly of interested people that want to say something about the topic at hand.

As expected by journalists, several users reported feeling more involved in the news topic after having participated in the debate via the RHETORiC tool. One participant noted: "*I used to respond to some articles already via the "correct" button. Whether anyone took it into account was a mystery to me. Now I feel more involved in the topic.*" Moreover, our results indicate that the RHETORiC tool can motivate the silent majority to participate in an online debate. 45% of our respondents (almost) never comments on online platforms, but they did via the RHETORiC tool. One participant wrote "*This way (of commenting) really struck me, it was the first time I commented on an article (. . .)*".

Even though the vast majority of the respondents evaluated the tool positively, some more critical feedback was shared as well, albeit to a much lesser extent. First, during the step of indicating similarity of arguments, participants reported having difficulties when their argument only partially matched a proposed argument. "*No response is ever completely identical*", noted one respondent. Second, at some points the automatic real-time feedback provided unjustified criticism, which elicited feelings of frustration among a few users.

5 LESSONS LEARNED

5.1 Process

First of all, we underestimated the sensitivity of the research topic (online news comments), and the uncertainty of the outcome, on the process itself. During the project, we had two media partners (a national newspaper and a national public broadcaster) on board, which was very instrumental in getting useful input and feedback, as well as for live-testing the RHETORiC tool. However, the broadcaster was apprehensive of using the tool on their own website, out of fear of negative reactions, especially as they have a societal responsibility as public institution. Whereas comments under news articles on Facebook could be still seen as less their responsibility (despite moderating those comments as well), if something would go wrong with the tool on their own website, and offensive comments about sensitive topics would get through, this could have a negative impact on their reputation. As a result, the public broadcaster did not allow the tool to be tested with a real-life audience, which is why it was only tested on the website of the national newspaper who was willing to take this calculated risk. However, although a large-scale A/B testing with a traditional comment section on the website itself would yield better comparable results, both media partners did not want to allow unmoderated comments on their website.

A second issue we ran into, was the operationalization of comment quality. In other words, what separates a high-quality comment from a low-quality one. In literature, the quality of online comments has been approached as a multifaceted and complex construct, and various efforts have been undertaken to formulate a definition and define a set of normative criteria. Moreover, the definition may vary according to the perspective, for example a user, journalist or moderator. Therefore, we recommend to approach comment quality from various stakeholder perspectives and come to better understand the normative construct.

Thirdly, we wanted to set up an experiment for assessing what the impact was of some design elements, which needed to be controlled in order to provide valid results (i.e. ensure internal validity). However, there are also many external factors that influence how people comment, such as the selected article, other comments, the branding of platform, etc. Moreover, the fact that other people can see and react to comments is essential, and not easy to replicate in an experiment. We therefore took various measures to mimic a real-life news comment section (i.e. ensuring a certain level of external validity), and made the participants believe that other participants would read their comment and be able to react to it. A pre-test showed that we largely succeeded, as participants rated the tool as being realistic, and believing that other would read their comment.

5.2 Tool

We also gained valuable insights about news commenting tools in general. First of all, throughout the whole study participants were concerned that tools like this could create an 'echo chamber', and only expose them to like-minded people and comments. It is thus important to expose readers and commenters to a wide range of (counter)arguments. This is not just an interface design issue but has implications on how the algorithms themselves are designed. We highly recommend working in an interdisciplinary team that considers the ethical implications of the technology used.

A similar concern arose regarding the live feedback on the comments that were written, indicating whether or not something was offensive or well-argued. Some participants were frustrated that the machine learning tool indicated something as not well-argued, while they thought it was. We believe more work is needed on 'explainable AI', e.g. by showing how reliable an automatic rating is, as well as giving more control to users e.g. to indicate why they think it is not accurate. This would remove frustration and give control back to users.

Finally, the issue of anonymity is a continuous source of debate. Our results arguably point to (the possibility of) anonymity of commenters leading to more participation and that it is positive for increasing civility. This seems in conflict with what is called the "online disinhibition effect", meaning that we lose our inhibitions when discussing in online fora because of being anonymous [17]. Although this could mean being more offensive or uncivil, it could also mean less fear of negative reactions and more honest conversations. We believe that a tool like ours, that combines anonymity with other measures such as live comment feedback and more structured interaction, could lead to more civil and constructive online conversations.

6 CONCLUSION

To overcome the drawbacks of a traditional news comment section, we gathered new knowledge on how to promote civil participation in digital journalism. This was done by designing and evaluating an

interface that includes design interventions aimed at increasing the added value of audience participation for both readers and editors. We followed a human-centered design process to create interface elements that support people in carefully formulating their opinions and arguments so they can learn how to constructively contribute to online discussions about news. The result is the RHETORiC tool with a step-by-step flow that guides commenters through the commenting process based on conversational prompts, stance taking, live feedback, and clustering of similar comments. Results from a large-scale live test on the website of a national news publisher are promising. The tool succeeds in doing what seemed impossible on social media: exchanging qualitative arguments and increasing news reader engagement.

ACKNOWLEDGMENTS

We thank all participants for their time. This work was supported by the RHETORiC project, realized in collaboration with imec, with project support from VLAIO (Flanders Innovation and Entrepreneurship). Project partners are VRT, Mediahuis, Tree Company, Wieni, Textgain, ID Lab (UGent) and Mintlab (KU Leuven).

REFERENCES

[1] Emilie Bossens, Elias Storms, and David Geerts. 2021. Improving the Debate: Interface Elements that Enhance Civility and Relevance in Online News Comments. In Human-Computer Interaction – INTERACT 2021, 433–450.

[2] Marcel Broersma. 2019. Audience Engagement. In Tim P. Vos, Folker Hanusch, Dimitra Dimitrakopoulou, Margaretha Geertsema-Sligh, and Annika Sehl (eds.). *The International Encyclopedia of Journalism Studies* Wiley. https://doi.org/10.1002/9781118841570.iejs0060

[3] James Curran, Shanto Iyengar, Anker Brink Lund, and Inka Salovaara-Moring. 2009. Media System, Public Knowledge and Democracy: A Comparative Study. European Journal of Communication 24, 1: 5–26. https://doi.org/10.1177/0267323108098943

[4] Nicholas Diakopoulos and Mor Naaman. 2011. Towards quality discourse in online news comments. In Proceedings of the ACM 2011 conference on Computer supported cooperative work - CSCW '11, 133. https://doi.org/10.1145/1958824.1958844

[5] Siamak Faridani, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg. 2010. Opinion space: a scalable tool for browsing online comments. In Proceedings of the 28th international conference on Human factors in computing systems - CHI '10, 1175. https://doi.org/10.1145/1753326.1753502

[6] Thomas B. Ksiazek and Nina Springer. 2018. User Comments in Digital Journalism. In The Routledge Handbook of Developments in Digital Journalism Studies (1st ed.), Scott A. Eldridge and Bob Franklin (eds.). Routledge, 475–486. https://doi.org/10.4324/9781315270449-37

[7] Thomas B. Ksiazek and Nina Springer. 2020. User Comments and Moderation in Digital Journalism: Disruptive Engagement. Routledge. https://doi.org/10.4324/9780429276132

[8] Edith Manosevitch. 2009. The Reflective Cue: Prompting Citizens for Greater Consideration of Reasons. International Journal of Public Opinion Research 21, 2: 187–203. https://doi.org/10.1093/ijpor/edp013

[9] Edith Manosevitch, Nili Steinfeld, and Azi Lev-On. 2014. Promoting online deliberation quality: cognitive cues matter. Information, Communication & Society 17, 10: 1177–1195. https://doi.org/10.1080/1369118X.2014.899610

[10] Cynthia Peacock, Joshua M Scacco, and Natalie Jomini Stroud. 2019. The deliberative influence of comment section structure. Journalism 20, 6: 752–771.https://doi.org/10.1177/1464884917711791

[11] Fabian Prochazka, Patrick Weber, and Wolfgang Schweiger. 2018. Effects of Civility and Reasoning in User Comments on Perceived Journalistic Quality. Journalism Studies 19, 1: 62–78. https://doi.org/10.1080/1461670X.2016.1161497

[12] Christine Schmidt. 2017. Remember that Norwegian site that made readers take a quiz before commenting? Here's an update on it. Retrieved from https://www.niemanlab.org/2017/08/remember-that-norwegian-site-that-makes-readers-take-a-quiz-before-commenting-heres-an-update-on-it/

[13] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. 2017. Design and Evaluation of a Short Version of the User Experience Questionnaire (UEQ-S). International Journal of Interactive Multimedia and Artificial Intelligence 4, 6: 103. https://doi.org/10.9781/ijimai.2017.09.001

[14] Joseph Seering, Tianmi Fang, Luca Damasco, Mianhong "Cherie" Chen, Likang Sun, and Geoff Kaufman. 2019. Designing User Interface Elements to Improve the Quality and Civility of Discourse in Online Commenting Behaviors. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 1–14. https://doi.org/10.1145/3290605.3300836

[15] Natalie Jomini Stroud, Joshua M. Scacco, Ashley Muddiman, and Alexander L. Curry. 2015. Changing Deliberative Norms on News Organizations' Facebook Sites. Journal of Computer-Mediated Communication 20, 2: 188–203. https://doi.org/10.1111/jcc4.12104

[16] Mingxiao Sui and Raymond J Pingree. 2016. In Search of Reason-Centered Discussion on China's Twitter: The Effects of Initiating Post and Discussion Format on Reasoning. 16.

[17] John Suler. 2004. The Online Disinhibition Effect. CyberPsychology & Behavior 7, 3: 321–326. https://doi.org/10.1089/1094931041291295

[18] W. Ben Towne and James D. Herbsleb. 2012. Design Considerations for Online Deliberation Systems. Journal of Information Technology & Politics 9, 1: 97–115. https://doi.org/10.1080/19331681.2011.637711

[19] Marc Ziegele, Oliver Quiring, Katharina Esau, and Dennis Friess. 2020. Linking News Value Theory With Online Deliberation: How News Factors and Illustration Factors in News Articles Affect the Deliberative Quality of User Discussions in SNS' Comment Sections. Communication Research 47, 6: 860–890. https://doi.org/10.1177/0093650218797884