

MOOD: Massive Open Online Deliberation Platform-A Practical Application

Verdiesen, E.P.; Dignum, Virginia; van den Hoven, Jeroen; Cligge, Martijn; Timmermans, Jan; Segers, Lennard

Publication date

2016

Document Version

Final published version

Published in

Proceedings of 22nd European Conference on Artificial Intelligence

Citation (APA)

Verdiesen, E. P., Dignum, V., van den Hoven, J., Cligge, M., Timmermans, J., & Segers, L. (2016). MOOD: Massive Open Online Deliberation Platform-A Practical Application. In G. A. Kaminka, M. Fox, P. Bouquet, E. Hüllermeier, V. Dignum, F. Dignum, & F. van Harmelen (Eds.), *Proceedings of 22nd European Conference on Artificial Intelligence* IOS Press.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

MOOD: Massive Open Online Deliberation Platform

A practical application

Ilse Verdiesen and Martijn Cligge and Jan Timmermans and Lennard Segers
and Virginia Dignum and Jeroen van den Hoven¹

Abstract. Nowadays, public debates often take place on social media platforms like Facebook or Twitter and can be characterized as asynchronous, protracted and ill-structured. The Massive Open Online Deliberation (MOOD) platform aims to structure these debates. Essential is that the platform can differentiate between the moral acceptability and the social acceptance of a debate outcome. We briefly describe the e-deliberation process and look at two existing debate platforms, *Liquidfeedback* and *Debatehub*. We design and build a prototype that mainly focuses on: (1) a method to differentiate and validate *facts* and *opinions*, and (2) a mechanism that maps both the *social acceptance* and the *moral acceptability* of debate outcomes. We research these ethical concepts more in depth and implement several techniques, such as a voting mechanism, in a working prototype that supports a four stage deliberation process. In future applications, machine learning techniques can be integrated in the platform to perform sentiment analysis on a debate.

1 INTRODUCTION

Public deliberation is an important component of decision-making in a democracy. Deliberation can result in an increased likelihood of justifiable policies, can help to identify incompatible moral values and can help people to get a broader perspective on policy questions [9]. The internet could be a valuable medium for public deliberation, because it can be a tool for information dissemination and long distance communication [20]. It allows citizens to share their opinion more easily. However, the debates that are currently held on the internet often take place on social media platforms like Facebook or Twitter and can therefore be characterized as asynchronous, protracted and ill-structured. E-deliberation platforms aim to structure these debates and their respective outcomes. These outcomes can be used by policy makers to make better decisions. In the field of ethics, the differentiation between social acceptance and moral acceptability is essential for the judgment on policies. Furthermore, public debates can be marginally ethical, as they occasionally contain discriminating content, and have statements that can be accepted, or not, by a majority of the crowd [21]. An example of this is a debate on banning polluting vehicles in the city center. This proposal can be accepted by local residents, but unaccepted by downtown business owners. Also, one could question if it is morally acceptable to prohibit access to city centers for potential customers and suppliers of businesses. On the other hand, for local residents the air quality is very important. E-deliberation platforms facilitate debates which should take the views

of both the majority as the minority into account, and therefore strive to be ethically just [21]. However, existing platforms often lack the ability to do so. In this paper, we propose our vision of a refined e-deliberation platform that takes into account the shortcomings of existing platforms by proposing a conceptual design and working prototype.

The paper is structured as follows: in section 2 we introduce the theoretical concepts underlying our design, describe related work in the field of deliberation processes and we analyze some existing platforms that support these processes. Section 3 shows the design choices and the methodologies used for our prototype. In section 4 we demonstrate the implementation and give insight in the framework we used to develop the platform. In the final section we discuss the limitations of our work and provide direction for further research.

2 RELATED WORK

In this section we describe the differentiation between facts and values, the concept of moral acceptability and social acceptance, and the e-deliberation process in general. We also look at two existing platforms that support this process. We analyze their shortcomings and based on these, we state the aspects we have focused on in the design of our prototype.

2.1 Facts and values

The distinction between facts and values is a much-debated concept in the world of ethics. Many philosophers have had their thoughts on how to filter descriptive statements from normative statements. Descriptive statements, also referred to as factual statements, describe factual matters and can be used to assert, deny or communicate about facts [13]. Normative statements, which can also be viewed as value judgments, deal with how people judge human decisions and conduct [16]. They are concerned with how people value factual matters and circumstances. We adhere to this distinction in developing our prototype.

2.2 Moral acceptability

Morality is concerned with the distinction between right and wrong and contains principles for good and bad behavior. These principles depend on the political, cultural and religious context they are defined in [6]. They govern our thoughts, emotions and behavior and can be viewed at a personal, interpersonal or collective level [4]. Morality can also be studied on a system level from a more functional approach and can be described as: *'Moral systems are interlocking*

¹ Delft University of Technology, The Netherlands, email authors: ep.verdiesen@student.tudelft.nl, m.e.cligge@student.tudelft.nl, l.h.j.segers@student.tudelft.nl, j.timmermans-1@student.tudelft.nl

sets of values, virtues, norms, practices, identities, institutions, technologies, and evolved psychological mechanisms that work together to suppress or regulate selfishness and make social life possible.' [8, p. 368]. This systematic approach resulted in the Moral Foundations Theory which uses a moral reasoning model based on the principles of harm, fairness, liberty, loyalty, authority, and purity [21]. We use these principles to define the moral acceptability of the alternatives proposed in the debate process.

2.3 Social acceptance

Social acceptance is a combination of individual feelings, perceived benefits and risks and above all, it is a social process in which people are influenced by various types of interactions. Views and available information are important for social acceptance [10]. Research shows that indicators for social acceptance are knowledge, fear and perceptions of the public [1]. We found that literature on measuring social acceptance is scarce. We turned to the field of ethics and looked at the Social Choice theory which provides a theoretical framework to reach a collective decision on social welfare. This theory is based on combining individual opinions, preferences and interests of people and links welfare economics and voting theory to aggregate preferences and behaviors of individuals. We define social acceptance as the collective decision on the preferences of individuals.

2.4 (E)-deliberation

In this paper, we define deliberation as a critical examination of a certain issue where the examination is based on the weighting of pro and con arguments for that issue. A deliberative process allows multiple participants to receive and exchange information, to critically examine this information, to form a collective judgment (based on the provided information) about a certain issue, which determines the decision-making on a certain issue [7]. E-deliberation platforms are platforms that make use of the modern online communication technologies to support such a deliberation process. The platforms capture collective judgments regarding complex social and political issues, such as decision-making over referendums, trade treaties and the use of killer-robots. These platforms intend to overcome legitimacy problems that may arise in public debates and public decision-making in controversial and adversarial arenas. E-deliberation platforms can be used to structure these deliberation processes by providing logic to support reasoning, voting procedures and reputation mechanisms [21]. E-deliberation platforms can be used by decision makers and citizens, to receive the opinions and information from debate participants on certain topics. For example, a decision maker might use it to introduce legislative proposals to citizens and to subsequently see how citizens evaluate these proposals via the collective judgment of the crowd.

2.5 Analysis of existing e-deliberation platforms

In order to get an understanding of the characteristics of the available e-deliberation platforms and to see if these platforms can be refined, we analyzed two existing platforms; LiquidFeedback and Debate Hub. We choose these two platforms because, in our opinion, these are two of the most investigated platforms and we were constrained by a limited amount of research time. In this analysis we mainly focused on how the deliberative process is implemented, how the collective judgments of the crowd are formed and how facts and values are differentiated and evaluated in order to identify gaps in the existing platforms which we use as input for our prototype.

2.5.1 LiquidFeedback

LiquidFeedback is designed and built by the Public Software Group of Berlin. The deliberation process consists of four phases; the admission phase, the discussion phase, the verification phase and the voting phase, where each phase has a fixed amount of time. Users of the platform can initiate a debate by proposing a certain issue, for example *'What should the town council do in order to improve the air quality in the city center?'*. Proposing of issues takes place in the admission phase, where users can support certain issues by voting. In the next step of the admission phase participants can provide alternatives to the proposed issues. An example of an alternative for the earlier described issue could be *'Polluting vehicles should be banned from the city center in the weekend'*. A discussion on a topic follows after a issue reached a certain quorum of votes in the admission phase. A discussion consists of the earlier mentioned alternatives and suggestions provided by discussants to improve the proposed alternatives. Users who provided issues and alternatives can choose to update their draft versions, based on the provided suggestions. After the discussion phase, discussants enter the verification phase. In the verification phase it is not possible anymore to change the draft alternatives, although new alternatives can still be added to the list of alternatives. At the end of the verification phase, users need to vote again on the list of alternatives. Only the alternatives that reached a certain quorum enter the next phase, which is the voting phase. This second quorum reduces the workload for participants in the voting phase. In the voting phase, participants can vote against or in favor of remaining alternatives which have passed the second quorum [2]. The voting mechanism for this last phase is conform the Schulze method, which will be explained in section 3.4 of this paper. An advantage of the Schulze method is that it takes minorities into account, so that alternatives that have a low amount of votes still have chance to reach the quorum.

LiquidFeedback is a well substantiated e-deliberation platform. However, we found that it could be improved in some areas. Firstly, LiquidFeedback does not elaborate on the differentiation of facts and values. If someone provides an alternative in the first three phases of the deliberation process, where is this alternative based on? Is it based on an opinion of someone, or is it based on a fact with corresponding literature? The platform does not explain how facts and opinions (values) are differentiated and how facts and corresponding sources are evaluated. Secondly, the platform does not differentiate in the outcome between social acceptance and moral acceptability. Social acceptance and moral acceptability often differ and that differentiation is important for decision-making and judgment [21]. The exact differences will be defined in section 3.2 and 3.3. Thirdly, in our opinion is LiquidFeedback a platform where participants can only provide alternatives for certain issues and subsequently modify these alternatives when participants do not support them. We miss a debate structure which is more focused on providing pro and con arguments with facts and corresponding literature, just like is done during a "real world" offline debate. These aspects are in our opinion crucial for a well-structured deliberation process, because requiring participants to add literature could result in a deliberation process of higher quality.

2.5.2 Debate Hub

The second existing platform we analyzed is Debate Hub. This platform is an initiative from the Open University's Knowledge Management Institute. The platform consists of debates where people can

provide debate topics, alternatives, and arguments. It does not have a well-defined deliberation process with different phases and fixed amounts of time as LiquidFeedback has, however, it has some sequence which users have to follow. The first step is initiating a debate topic or issue, such as the example provided in section 2.2.1; ‘*What should the town council do in order to improve the air quality in the city center?*’. After that, participants can add alternatives, like; ‘*Polluting vehicles should be banned from the city center in the weekend*’. Consequently, participants can add pro-or con arguments to these alternatives. The structure of the argument form allows participants to add literature to their arguments. Participants can vote on alternatives and arguments, but there is no voting mechanism that filters out the most accepted alternatives or arguments like LiquidFeedback has.

After analyzing Debate Hub, we found that Debate Hub has a very different setup compared to LiquidFeedback, since it does not have a deliberation process with distinctive phases and fixed times. The debates pages are more like forms to structure an online debate. In our opinion, the following aspects could be improved; firstly, there is no quorum for initiating a debate. By not implementing a quorum, there will be many debates without any participants. Secondly, although there is some distinction between facts and values, the facts are not validated. Thirdly, there is no distinction between social acceptance and moral acceptability. Users only can show their support for certain alternatives or arguments, but it is not clear how users evaluate the moral acceptability of certain alternatives or arguments. Lastly, there is no voting method that takes minorities into account.

2.6 Discussion related work

Based on the previous section we can conclude that the two analyzed platforms are complete, but have drawbacks in some areas. LiquidFeedback has a deliberation process with distinctive phases in which results of the most accepted alternatives are listed, while Debate Hub has a very clear way of structuring the debate itself by letting users provide debate topics or issues, alternatives and pro-and con arguments (just like in “real world” offline debates). We built a prototype that focuses on the one hand on combining the best of both platforms (by using parts of the debate page structure of Debate Hub and by using parts of the deliberation process of LiquidFeedback) and on the other hand on aspects of both platforms that could be improved. We defined a design objective for our prototype which is based on the earlier described analysis. Our design objective mainly focuses on the following aspects: (1) a method to differentiate and validate *facts* and *opinions*, and (2) a mechanism that supports both the *social acceptance* and the *moral acceptability* of debate outcomes.

3 METHODOLOGY

In this section we describe the methodologies we used in our deliberation process design and we state which methods we implemented in our platform.

3.1 Facts and values

The goal of differentiating between facts and values for our system is to have a clear discussion that is based on facts, and let participants have a discussion over values which are derived from those facts. We think that by keeping the structure of the debate page of Debate Hub, we are able to structure the debate in such a way that participants have to provide a fact with the corresponding source for every argument they make. The structure of the page where people

can add an argument with facts requires users to add a URL which supports their facts. This will be explained in section 4.1 in more detail. To validate the facts and sources provided by participants, we use the methodology of online encyclopedia Wikipedia. Wikipedia implemented crowd-sourcing technology, where users (the crowd or editors) have the responsibility of (1) adding content to the encyclopedia and (2) validating all of the content. This is done by panels of experts. The composition of these panels is formed throughout the existence of the website. Groups of active editors are specialized in certain topics, and if false content on certain pages exists, they will correct this content [18]. We incorporate this concept in our platform, by letting users report on facts they think are not correct. If a fact reaches a certain amount of so-called report votes, a group of users will be notified to check this fact. This group of users is randomly selected and they have the responsibility to validate the reported fact and/ or source. If they are not able to judge if a fact is correct or incorrect, they can inform a group of users which are expert in the field of where the source comes from. We propose a two step procedure with a randomly selected panel and an expert panel to limit the workload for the expert panel. In other words, the validation of facts in this methodology relies on the wisdom of the crowd. We realize that this methodology might be vulnerable for groupthink and strategic behavior, but we think that Wikipedia proves that the wisdom of the crowd works, if implemented correctly.

3.2 Moral acceptability

To survey the moral acceptability of the alternatives we use the Moral Foundations Questionnaire (MFQ) that was developed [8] based on the Moral Foundation Theory. The MFQ can be used to measure a broad range of moral concerns. The MFQ consists of two parts, one about moral relevance and the other one is about moral judgment. We intended to use the fifteen questions of the first part as an instrument to assess the moral acceptability of the proposed alternatives in the debates. We performed a small test to check the understandability of the questions. It turned out that the questions in their original form were hard to understand by the testers and did not fit the way we want to measure the alternatives. Therefore we decided to adjust the MFQ questions slightly to make them more applicable to our design of the debate process and understandable for the user. An example of this modification is the rephrasing the statement *Whether or not some people were treated differently than others* into the question: *Do you think that as a result of the alternative above: Someone is treated differently from others?* We realize that this impacts the validity of this instrument which means that research is needed to validate the modified questions. Since our prototype is merely a proof of concept we chose not to test this validity at this moment.

3.3 Social acceptance

As described in paragraph 2.3, the Social Choice theory takes the preferences of individuals into account, therefore we regard it as a suitable means to measure social acceptance. We studied several voting mechanisms that are being used in Social Choice Theory and chose to implement one to determine the social acceptance of the alternatives of the debates. These voting mechanisms are described in the next paragraph.

3.4 Voting mechanisms

Voting is a popular method to reach a joint decision based on aggregated preferences of individuals. One of the most used voting mech-

anisms in elections is the Schulze method which is used by Ubuntu, several Pirate Party political parties, OpenStack and LiquidFeedback [17]. This preferential voting method satisfies among others the criteria of anonymity, the Condorcet criterion and independence of clones [19]. Voters can list their preferences anonymously which is an important prerequisite for elections. The Condorcet criterion selects a single winner by majority rule in pairwise comparisons over every other candidates. Clone independence is a criterion that prevents certain types of strategic behavior in the voting process which means that it is impossible to be insincere about a voter's real preferences in order to secure a more favorable outcome. In the Schulze method every voter submits an ordered preference list for the candidates presented to the voter. All candidates are compared pairwise and a directed graph with the strongest path is created based on all votes and pair-wised comparisons. The output can be determined by looking which candidate defeated all other candidates and this one is declared the winner [17].

Next to the Schulze method we considered to implement the Ranked pairs algorithm, because this method is even more robust to strategic behavior [19] and it satisfies most of the same criteria as the Schulze method. Both are Condorcet methods, but they produce a different order of winners due to the fact that the Schulze algorithm reverses a larger majority than the Ranked Pairs algorithm for the majorities on which the two orders of finish disagree [17]. We found that there is less information available about the Ranked pairs algorithm than about the Schulze method. Ranked pairs is also harder to understand which negatively impacts the transparency of the voting mechanism. Therefore, we chose to implement the Schulze method in our prototype and used the PHP library of Julien Boudry that was available on GitHub [3]. We analyzed and tested the implementation of this algorithm with voting example to determine if the open-source algorithm was correct, which it turned out to be.

4 IMPLEMENTATION

In this section we describe the techniques we implemented in our prototype that we developed in the ten weeks of our Information Architecture design project course. We explain our choices for the framework we used and sketch our plan to test our application.

4.1 MOOD deliberation process

In our prototype we implemented the actual e-deliberation process based on the methods described in the previous section. We built a deliberation process consisting four stages: (1) proposal and initiation of a debate, (2) the actual debate in which user can cast votes to support an alternative, (3) the selection of alternatives via preference voting and measuring the moral acceptability of the alternatives and (4) reporting of the results. These stages are depicted in figure 1 which are translated to the application in the overview of the debate page in figure 2.

In stage one, a user can initiate a debate by submitting a proposal to the MOOD platform. This proposal needs to be described in a generic way and should preferably be posed as an open question. The initiator has four weeks to raise support for the debate and to reach a voting threshold. We set the threshold with an initial value of ten votes, but we will have to test if this value proves to be correct. The threshold procedure resembles the procedure for citizen initiatives in The Netherlands [5]. After reaching the voting threshold the proposal enters stage two of the debate. Once the threshold is reached, an initiator cannot withdraw his proposed debate, because this would

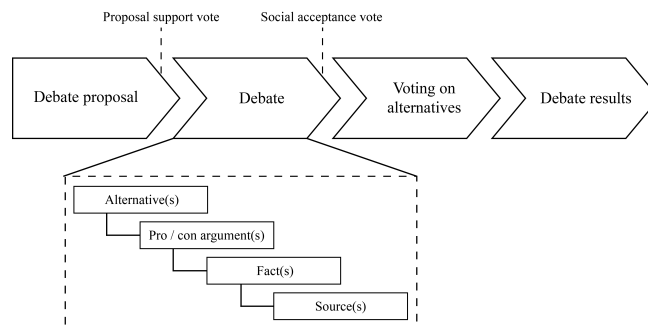


Figure 1. MOOD deliberation process

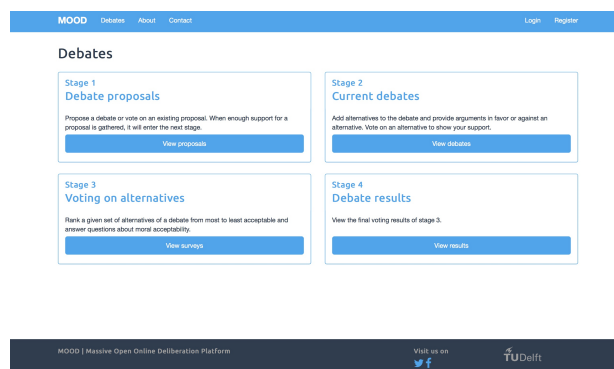


Figure 2. Screenshot debate page

mean that all aspects of a certain debate, like arguments, sources and facts, will be deleted and to our opinion valuable information will be lost.

In stage two the actual debate is held. Discussants can react to a debate by submitting alternatives which consist of pro- and con arguments (figure 3). It is also possible for users to add pro- or con arguments to an existing alternative. Arguments need to be substantiated by facts and sources to reference these facts to differentiate them from values. Although not built in our prototype yet, these facts will be validated by means of crowd-sourcing. The facts can be contested by other users and if a certain threshold is reached, the administrator will review the fact. If the fact is not valid then it will be marked in the database as rejected and will not be visible to the users. In a future version of the MOOD platform an expert panel will take over this task from the administrator to provide a more independent judgment of a contested fact. A debate will have a pre-set duration which is set by the initiator. In this stage, all users can vote to support an alternative. The five top alternatives will be selected and the debate will enter the next phase.

In the third stage of the debate, a voter can list his or her preferences of alternatives. The preferences are calculated by the Schulze voting mechanism. By this, the social acceptance of the alternatives in a debate is measured. After the voting, a list of alternatives is created ranking the alternatives that received the most votes. Next, the moral acceptability of the alternatives is surveyed with questions that are based on the MFQ for the selected alternatives. Per alternative seven questions will be asked to measure the ethical principles of *harm*, *fairness* and *authority*. The answers will be used to gain insight in the moral acceptability of the different alternatives in a debate.

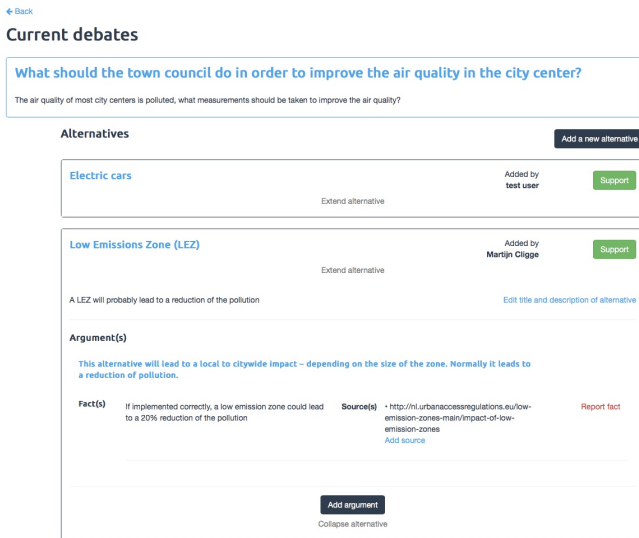


Figure 3. Screenshot alternatives page

In the fourth and final stage the social acceptance and moral acceptability results of the debate will be presented (figure 4). The results will be available to all users which will enhance the transparency of the debate.



Figure 4. Screenshot results page

4.2 Framework

We chose an open-source framework to develop our prototype, because it is easily available and it enhances the transparency and traceability of our platform. We used the free open-source PHP framework *Laravel* to build the prototype. This framework is available on

GitHub and can be used under the terms of a MIT license. According to their official website, it can be used to build elegant web applications that are *'...delivered at warp speed.'* [11]. It is developed via a Model-View-Controller (MVC) architecture. This is a category of software applications that consists of three interconnected parts that separate the internal representation of information from the way the information is presented to the user. The *Model* component handles the data, logic and rules of the application and stores the data it receives from the controller. The *View* shows the output of the application and generates new output when the model changes. The *Controller* accepts and converts the input into commands for the model and the view [15]. Laravel is one of the most popular PHP frameworks at this moment and includes features, such as a Composer architecture for Artisan, Laravel's Command Line Interface, Eloquent Object-Relational-Mapping (ORM) to put constraints on the relationships between database objects and Query builder to program queries automatically [14]. To create the database we used the open-source PHPMyAdmin software that handles MySQL queries for the web [12]. We used *bootstrap* to adjust the layout of the web application dynamically to the (mobile) device of the user. This free and open-source library is hosted on GitHub. Using bootstrap we aim to enhance the user experience for our prototype.

4.3 Testing

At the time that we are writing this paper we did not test our web application yet. Our first test will focus on the usability of our application. We will ask a small group of individuals (3-5 people) to walk through our application via scenario testing. The test scenario focuses on the e-deliberation process of our application. We ask our testers to follow this scenario to see if they understand the different steps in the process and to assess if the application is easy to use. The scenario starts by asking the user to make a new account and subsequently login with this account. After that, our testers will propose a new debate in the first stage of our deliberation process. Next, testers have to work themselves to the different stages, by adding alternatives, arguments, facts and sources in stage 2, by ranking the most social acceptable alternatives in stage 3, by filling in the survey on moral acceptability and by viewing the results in stage 4. We already prepared some debate issues in stage 2, like "No fast food should be sold in the University canteen, because it leads to obesity". We have designed two different kind of setups for our scenario. In the first setup, we will provide users with some explanation and a clear walk-through description which describes every step in the scenario. In the second setup, we ask our testers to follow the same steps as in the first setup, but we give them very minimal explanation and no clear walk-through description. We ask them to think out loud while performing the scenario with the second setup. The results of our test will be available after this paper is drafted, therefore these are not included in this document right now.

5 CONCLUSION AND DISCUSSION

In this paper we gave an overview of the e-deliberation process and existing platforms Liquidfeedback and Debatehub. We built a prototype that focuses on the one hand on combining the best of both platforms (by using parts of the debate page structure of Debatehub and by using parts of the deliberation process of Liquidfeedback) and on the other hand on aspects of both platforms that could be improved. Our design objective mainly focuses on the following

aspects: (1) a method to differentiate and validate *facts* and *opinions*, and (2) a mechanism that supports both the *social acceptance* and the *moral acceptability* of debate outcomes. We researched these concepts more in depth and implemented several techniques to meet these aspects which resulted in a working prototype.

5.1 Limitations

Due to the little available development time, our prototype has several limitations. We focused our research on the topics of the differentiation between facts and values, social acceptance, moral acceptability and voting mechanisms. Time lacked to extensively study these topics and we realize that this scoping can lead to conformation bias, which means that we only used literature that substantiates our ideas and did not consider alternative views. The time constraint also affected the features of our prototype. One of the features that we did not manage to implement, is that of reputation score to distinguish between experts of certain discussion topics and regular users. This distinction is useful to create expert panels to validate the contesting of facts in the stage of the actual debate. Another feature we did not implement is an algorithm that creates a random panel to evaluate a contested fact. In the current application this task is performed by the administrator. Furthermore, a limitation is that we modified the MFQ questionnaire, but we did not study the effect of this instrument. Next to this, we chose to run the application on a small centralised server of the University which limits the amount of users that can simultaneously take part in the debate and impacts the scalability. To accommodate more users, a distributed or cloud server is needed to upscale the application in the future. Finally, we made a trade-off regarding the privacy of users and security of the platform. A limitation of our current design is that an administrator or auditor can trace a vote back to a user who casted it. Although, this violates the anonymity requirement of voting, this information is only visible for an administrator or auditor and not for any other user. More importantly, it enables full traceability, which contributes to more transparency and credibility via audits of the voting results. It is not possible for users to see how often is voted on alternatives in stage two to limit strategic behaviour which could occur when an alternative received many votes and people might want to vote on an alternative that is popular. Nevertheless, strategic behaviour could occur when users register with multiple e-mail addresses in order to be able to cast more votes. We have not been able to implement a counter measure for this in our prototype.

5.2 Future research

These limitations lead to recommendations for future work. We did not manage to study the revised MFQ questions. Its validity and applicability to measure moral acceptability in debates should be researched. We also recommend to extent the literature study for mechanisms to differentiate between facts and values, for social acceptance, moral acceptability and voting mechanisms and find alternate views on these topics. An extension of the voting stage would also be a possible addition to a future version of the application. Adding a second round of preferential voting, after the publication of the results of the moral acceptability survey, would allow people to change their mind and vote for a different alternative than they did the first time. We did not manage to include all features in our prototype that we described in our list of requirements. A mechanism for crowd-sourcing should be added to categorize the facts that are added to the debate. Next to this, it should be possible to forward a contested fact

to an expert panel for an independent judgment. Also, tracking the reputation score of users should be added as a feature to our prototype. These features are crucial to develop the MOOD platform into a more mature application. Additionally, sentiment analysis on content provided by the users could be implemented in the MOOD platform to sense the atmosphere of the debate. On the other hand machine learning techniques can also be used to support the MOOD platform. For example validate facts by means of crowd-sourcing applications or Watson APIs.

ACKNOWLEDGEMENTS

We would like to thank Dirk Helbing for his thoughts and input to our prototype.

REFERENCES

- [1] Getachew Assefa and Björn Frostell, 'Social sustainability and social acceptance in technology assessment: A case study of energy technologies', *Technology in Society*, **29**(1), 63–78, (2007).
- [2] Jan Behrens, Axel Kistner, Andreas Nitsche, Björn Swierczek, and Björn Swierczek, *The principles of LiquidFeedback*, Interaktive Demokratie e. V., 2014.
- [3] Julien Boudry, Condorcet. <https://github.com/julien-boudry/Condorcet>. Retrieved at: 26-05-2016.
- [4] Taya R Cohen and Lily Morse, 'Moral character: What it is and what it does', *Research in Organizational Behavior*, **34**, 43–61, (2014).
- [5] Tweede Kamer der Staten Generaal n.d. Burgerinitiatief. <https://www.tweedekamer.nl/kamerleden/commissies/verz/burgerinitiatieven>. Retrieved at: 07-07-2016.
- [6] Naomi Ellemers, Stefano Pagliaro, and Manuela Barreto, 'Morality and behavioural regulation in groups: A social identity approach', *European Review of Social Psychology*, **24**(1), 160–193, (2013).
- [7] James D Fearon, 'Deliberation as discussion', *Deliberative democracy*, **44**, 56, (1998).
- [8] Jesse Graham, Brian A Nosek, Jonathan Haidt, Ravi Iyer, Spassena Koleva, and Peter H Ditto, 'Mapping the moral domain.', *Journal of personality and social psychology*, **101**(2), 366, (2011).
- [9] Amy Gutmann and Dennis Thompson, *Democracy and disagreement*, Harvard University Press, 2009.
- [10] Nicole MA Huijts, Cees JH Midden, and Anneloes L Meijnders, 'Social acceptance of carbon dioxide storage', *Energy policy*, **35**(5), 2780–2789, (2007).
- [11] Laravel n.d. <https://laravel.com/>. Retrieved at: 26-05-2016.
- [12] PHPMyAdmin n.d. <http://www.phpmyadmin.net/>. Retrieved at: 26-05-2016.
- [13] Wikipedia n.d. Descriptive statement. https://en.wikipedia.org/wiki/Positive_statement. Retrieved at: 07-07-2016.
- [14] Wikipedia n.d. Laravel. <https://en.wikipedia.org/wiki/Laravel>. Retrieved at: 26-05-2016.
- [15] Wikipedia n.d. Model-view-controller. <https://en.wikipedia.org/wiki/Modelviewcontroller>. Retrieved at: 26-05-2016.
- [16] Wikipedia n.d. Normative statement. https://en.wikipedia.org/wiki/Normative_statement. Retrieved at: 07-07-2016.
- [17] Wikipedia n.d. Schulze method. https://en.wikipedia.org/wiki/schulze_method. Retrieved at: 26-05-2016.
- [18] Wikipedia n.d. Wikipedia:editorial oversight and control. https://en.wikipedia.org/wiki/Wikipedia:Editorial_oversight_and_control. Retrieved at: 26-05-2016.
- [19] David C Parkes and Lirong Xia, 'A complexity-of-strategic-behavior comparison between schulze's rule and ranked pairs', in *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*. American Association for Artificial Intelligence, (2012).
- [20] Rebecca J Romsdahl, 'Political deliberation and e-participation in policy-making', *CLCWeb: Comparative Literature and Culture*, **7**(2), 7, (2005).
- [21] Jeroen van den Hoven and Virginia Dignum, 'Moods: Massive open online deliberation'. draft.