Computational Models of Argument
S. Modgil et al. (Eds.)

doi:10.3233/978-1-61499-906-5-325

D-BAS - A Dialog-Based Online Argumentation System

Tobias KRAUTHOFF [a,1], Christian METER [a], Michael BAURMANN [b], Gregor BETZ [c] and Martin MAUVE [a]

[a] *Department of Computer Science, University of Düsseldorf, Germany*
[b] *Department of Sociology, University of Düsseldorf, Germany*
[c] *Department of Philosophy, Karlsruhe Institute of Technology, Germany*

Abstract. In this paper, we present D-BAS, a dialog-based online argumentation system, tailored to support e-participation processes. The main idea of D-BAS is to let users exchange proposals and arguments with each other in the form of a time-shifted dialog where arguments are presented and acted upon one-at-a-time. We highlight the key research challenges that needed to be addressed in order to realize such a system, provide solutions for those challenges, report on a full scale implementation of D-BAS and summarize the findings from a real world e-participation process, where D-BAS provided the infrastructure for online argumentation.

Keywords. online argumentation, dialog-based argumentation, online participation

1. Introduction

E-participation, such as urban planning or participatory budgeting, is an important application area for online argumentation. In these processes citizens that will be affected by future decisions are invited to participate in the decision making process by proposing actions and discussing them with their peers. The results of the discussion, i.e., the proposals and the arguments, are then incorporated in the decision making process.

E-participation is a challenging application area for online argumentation, since the participants might be experts in the problem domain, but they are not experts in argumentation. Additionally, they often have a significant stake in the topic that is being discussed. Therefore they typically want to convey their point of view rather than engage in evidence-based deliberation. Furthermore, the number of individual contributions can be very large in these processes, while at the same time, the available resources are often rather limited. As a consequence it is frequently not an option to have experts in argumentation involved in the process in order to take the input of the participants and then structure it in an appropriate way.

At the same time, however, the result of the online argumentation is not the final outcome in an e-participation process. Instead, it is taken as an input by those that finally make a decision, such as elected representatives. Thus there is a layer above the online

[1]Corresponding Author: Tobias Krauthoff, Department of Computer Science, Universitätsstraße 1, 40225 Düsseldorf, Germany; E-mail: krauthoff@cs.uni-duesseldorf.de

argumentation that is able to interpret and weight the individual proposals and arguments. It is therefore not necessary that online argumentation in e-participation processes arrives at a certain conclusion such as a consensus or any other form of decision. Instead, a well structured set of proposals and (interrelated) arguments is a perfectly acceptable result.

So far, online argumentation in e-participation processes mainly relies on forum-based systems. This leads to well known problems such as limited scalability and a lack of structure [1]. As a consequence there have been several attempts to provide better support for online argumentation. However, so far, none of them has had really significant practical impact. One important reason for this may be that forum-based systems offer something that other systems do not: they allow for a highly complex exchange of arguments and counter-arguments with an intuitive statement-reply-scheme. Other approaches to online argumentation either do not capture the full complexity of argumentation (e.g., pro/con lists) or they require that the user is trained in operating a rather complex technical tool (e.g., the cooperative creation of an argument map).

In this paper we present D-BAS, a dialog-based online argumentation system, that does not require any prior knowledge or training from the user and avoids the shortcomings of forum-based systems while still allowing for complex argumentation. The main application scenario we have in mind for this approach is e-participation, while we do believe, that it might be applicable to other areas as well.

The key idea of our approach is to guide participants through the arguments provided by other users so that they perform a time-shifted dialog with those that have participated before them. The system is driven by a formal data structure capturing the full complexity of argumentation. User interactions, however, have the structure of a regular dialog as it is performed in everyday life. It is the task of system – and not of the participants – to translate between those two views. We call this approach *dialog-based online argumentation*. The output of dialog-based online argumentation is a set of proposals and interrelated arguments, both provided by the participants. While it might be possible to extend dialog-based online argumentation to include group decisions, this is not part of the work described here.

This paper is structured as follows. In Section 2 we give a brief overview of related work in the area of online argumentation. Section 3 explains the model of online argumentation used by D-BAS. The general idea of dialog-based online argumentation is outlined in Section 4. The solutions to key challenges are presented in Sections 5, 6 and 7. In Section 8 we describe the implementation of D-BAS as a fully functional dialog-based online argumentation system. We show that the idea of dialog-based online argumentation is viable by summarizing the findings from a deployment of D-BAS in a real world setting in Section 9. The paper is then concluded by a summary in Section 10.

2. Related Work

The general idea, key challenges of dialog-based online argumentation and details about our argumentation framework were already given in [2]. We have analyzed the details of a real world deployment and evaluation of D-BAS in a paper presented at AI ^{3} 2017 [3]. The current submission, on the other hand, only briefly summarizes the key findings of the real world deployment in order to show that the idea of dialog-based online argumentation is actually viable and focus developing a fully fledged dialog-based online argumentation system called D-BAS, targeting e-participation processes.

Current approaches for online argumentation in e-participation processes can be roughly separated into three main groups: forum-based approaches, pro and contra lists and argumentation maps.

The first group, also called asynchronous threaded discussions, allows participants to exchange arguments by means of a sequence of text contributions. Those approaches have encountered much criticism in the past, because they are believed to lead to a high degree of redundancy [4] and polarization [5,6] while scaling poorly with the number of involved participants and supporting non-collaborativness [7]. However, in practice they are, by far, the most commonly used approach to support online argumentation.

It has been suggested, e.g. ConsiderIt [8], to use online pro and contra lists to aid collective decision making processes. These lists work very well for evaluating a given proposal. However, they are not suitable to deal with more general positions and alternatives since they do not support the exchange of arguments and counter arguments.

Online systems for argument mapping enable participants to structure their arguments and the relation between them in an argument map. Examples are Carneades [9] and Deliberatorium [10]. While those systems do avoid the shortcomings of forum-based approaches, they require the users to become familiar with their notations and the semantics of formal argumentation. Therefore, in practice, they are used by experts or students who want to learn about the logic of argumentation rather than by average users that want to take part in an e-participation process.

The idea of engaging in a formalized dialog to exchange arguments is used by so-called dialog games, which follow a set of rules to react to the statements of each other, see [11]. In contrast to our work, dialog games focus on the real-time interaction between users in order to learn something about a subject at hand. They do not seek to provide better instruments for online argumentation.

In addition to general work on online argumentation there are three individual systems that are related to our work. The first one is the *Structured Consultation Tool* (SCT) [12]. Its primary goal is to allow a government agency to elaborate and present a justification for a given action. While the SCT explicitly seeks feedback on the arguments provided by the government agency, it does so in a questionnaire kind of way. This is valid for gathering feedback on government proposals, but it is unsuitable for an online argumentation, where the dynamic exchange of arguments is the main focus.

The Carneades Opinion Formation and Polling Tool [13] is part of the Carneades argumentation mapping system. It allows participants to provide structured, questionnaire-style feedback on a given argumentation consisting of multiple arguments and positions put forward by – potentially – many agents. This tool can be regarded as a generalization of the SCT. As with the SCT the questionnaire-style feedback is well suited for an evaluation of government activities by citizens but it does not fit the idea of an online argumentation amongst peers.

The third system that is related to our work is Arvina [14] and its predecessor MAgtALO [15]. Both systems allow a user to conduct a dialog between robots and humans. As a basis, they use an existing discussion where the positions and arguments of some real-world persons are marked. In contrast to D-BAS Arvina and MAgtALO are driven by the questions of the users. Thus there is no need for the users to react to replies from the system by providing their own arguments.

3. The System Perspective

In the following we assume that every online argumentation is identified by a *topic* that describes what the argumentation is about. *Statements* are the most basic primitives used in an online argumentation. Individual participants might consider a given statement to be true or false. The negation of a statement is itself a statement. A *position* is a prescriptive statement, i.e., a statement which recommends or demands that a certain action can be taken.

We distinguish between first-order and second-order *argumentation*. First of all, there is argumentation for or against statements. Here, some statement (the premise) is said to be a reason for another statement (the conclusion). This leads to a first-order argument, consisting in a premise-statement, a conclusion-statement and a reason relation between both. With this structure it is straightforward to represent undermines and rebuttals. A first-order argument A attacks another first-order argument B iff A's conclusion is the negation of a premise of B; and A is a rebuttal of B iff A and B have contradictory conclusions.

Still, we must not presuppose that untrained participants in dialog-based online argumentation advance only deductive reasons and valid arguments. The reason-relations claimed by users might be more or less cogent – and more or less evident for other users, which may trigger a discussion about the strengths of reason-relations. That is what we call second-order argumentation. Accordingly, we allow reasoning not only about statements but also about whether one statement really supports (attacks) another statement. A second-order argument consists of a statement (the premise) that is cited as a reason for why another reason-relation does not hold. Second-order arguments allow us to express undercutting attacks, namely as arguments against reason-relations pertained in other arguments.

As a consequence we use the following definition: an argument consists of one or more statements, which form the premise(s); one statement or the second-order claim that a certain reason-relation does or does not hold, which forms the conclusion; and the reason-relation between premise and conclusion. Together, the arguments of a debate form a (partially connected) *web of reasons* (WoR).

We would like to stress that our data structure and the distinction between first-order and second-order argumentation in a user dialog does not commit us to a specific argumentation-theoretic framework. On the contrary, the dialogs we model can be interpreted in quite different ways:

- *Deductive argumentation:* The arguments we model can be understood as enthymemes, i.e. incomplete arguments, that can in principle be reconstructed as deductively valid arguments if all implicit assumptions are made explicit. On this view, second-order argumentation would actually be argumentation about the plausibility of those implicit assumptions.
- *Probabilistic reasons:* The reason-relations can be explicated in probabilistic terms. On this view, a second-order argumentation undermines or establishes the probabilistic reason-relation maintained in another argument.
- *Defeasible reasons:* The arguments we describe can be understood as defeasible reasons. On this view, a second-order argumentation defeats another argument (or attacks a defeater).

We conceive this theoretical openness of our argumentative dialog model as a major strength. Note that the data we generate can also be used to check how well the alternative paradigms of rational argumentation can cope with the discussions we protocol.

4. The User Perspective

The foundation of dialog-based online argumentation is a novel way to navigate an existing set of arguments pertaining to a given subject. Instead of presenting many arguments at once – in maps or lists of arguments – the user is shown only a single argument at a time. It is then possible for the user to respond to that statement, either by selecting a statement provided by another user or by entering a new statement. Based on this response and, possibly, the data gathered from the responses of other participants, the system selects the next argument that is shown to the user. In this way the user and the system perform a dialog where the system selects arguments that are likely to be of interest to the user and the user provides feedback on those arguments.

Both, the user and the system, profit from the dialog. The user is efficiently guided towards those arguments that are particularly relevant for her. This also reduces redundancy, polarization and the occurrence of logical fallacies. The system, on the other hand, will increase its knowledge base with every response from a participant. This can then be used to improve the selection of arguments for the next user and to provide a summary of the online argumentation at hand.

5. The First Challenge: Feedback

The most basic building block of dialog-based online argumentation is gathering feedback from a user regarding a given argument. This is done by asking a question derived from the statements pertaining to the argument in the WoR. For example, if we have a premise-conclusion structure, the question generated by the system would be “What do you think about the following argument: ...?”. The system then offers a set of answers from which the user can choose. This set has to be constructed in a way that enables an untrained user to provide precise feedback on the argument. A simple choice between: “I agree with this argument” and “I do not agree with this argument” could undoubtedly be made by an untrained participant. However, both statements are not precise and have little significance. For example “I do not agree with this argument” might refer to several distinct scenarios: the user might disagree with the premise, the user might think that the conclusion is not supported by the premise or the user might consider this to be a valid argument, but at the same time she might consider, that it is weaker than other arguments supporting the negation of its conclusion.

In order to get precise and meaningful feedback from the user, the system has to differentiate between the scenarios by means of a set of answers that the user can choose from. Experiments with a prototype system that allowed users to react to arguments of a pre-constructed online-argumentation led us to one key observation: giving feedback on an argument is a two step process. The first step requires just a single click from the user to determine her initial reaction to the argument, e.g. the user rejects the premise of the argument. As a second step the user can then provide a justification for her choice either by selecting an existing statement or by providing a new one. Separating the two steps facilitates very fast feedback and a clean design of the user interface.

Next, we constructed the set of reactions that would be offered to the user in order to respond to a given argument. As a basis we used attacks as they are defined in argumentation theory and added two options that are frequently used in informal argumentation. This led to the following set: (1) Reject the premise. (2) Accept the premise and, as a consequence, the conclusion. (3) Accept the premise but disagree that this leads to accepting the conclusion. (4) Accept the premise but state that there is a stronger argument that leads to rejecting the conclusion. (5) Do not care about the argument.

Once the user has selected an answer and provided a reason, the system uses this to update the internal information of the WoR and to select the next argument that is presented to the user.

5.1. Optimizing the Representation of Questions and Answers

In an early implementation of our system we simply repeated the full *premise* and *conclusion* for each individual option that the user could choose. First tests showed that this leads to very long feedback options where some text parts were repeated several times. Participants in those tests indicated that this was not acceptable since they lost their focus when reading all the feedback options. As a solution to this problem, we use terms like "my point of view", "their statement" or "their point of view" instead of repeating the actual premise and conclusion of the argument. In order to make sure that the participants can easily determine what those terms refer to, both the terms and the premise or conclusion they refer to are colored in the same way. An example is shown in Figure 1

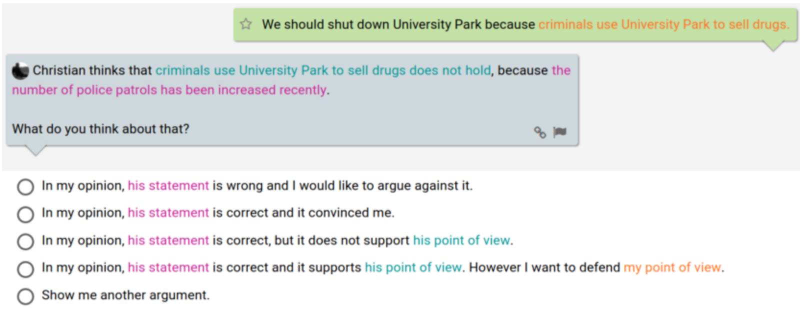

Figure 1. Challenging the user's argument and getting feedback from the user.

5.2. User Evaluation

We conducted two experiments with 18 participants each. The goal of the experiments was to verify whether we have included all relevant feedback options. For the first experiment we used the initial feedback options without substituting the premises and conclusions. The second experiment then employed those substitutions including coloring and highlighting. In both experiments the number of male and female participants was about the same and the age of the participants covered a wide range.

Both experiments were designed as follows: the participant was shown two arguments. A first argument represented the opinion of the participant and a second argument was an attack on the first argument. The participant was then asked: “What is your opinion regarding the second argument?”. The participant would then write down her answer in a simple text field. Afterwards, the participant was shown the feedback options described above and asked the question: “Which option would you choose?”. Comparing the text entered with the chosen option allowed us to determine whether the user is able to indicate the desired reply by using the feedback options.

The first, and possibly most important, result of the experiment was that every single reply given in the text field could be matched to one of the feedback options we described above. This indicates that our feedback options are complete. We then investigated whether the participants provided consistent answers and thereby showing that the feedback options were presented in a way that the user would understand.

In general, the user answers could be mapped to the appropriate feedback option with a chance of 72.2% for the first experiment. Especially in the second experiment there were less differences between the free text and the chosen feedback option, because the mapping increased to 83.3%. We believe that the improvement is caused by a reduction in the complexity of the feedback options when using color-coded replies.

6. The Second Challenge: Navigating the Web of Reasons

The second major challenge for dialog-based online argumentation is how the system should select the arguments that are presented to the participant. We address this challange through three mechanisms: (1) bootstrapping the dialog by identifying the first argument that should be presented to a new user; (2) selecting the next argument based on the prior actions of the user and (3) an optional lateral entry into an ongoing argumentation.

6.1. Bootstrapping

The first thing that the system usually needs to do when a new user wants to participate is to choose an initial argument to present to the user. This is challenging since the system has no information on the user, yet. We solve this problem by asking the participant which initial position she is interested in. This position is selected as the starting point in the WoR. The user is then invited to indicate her attitude towards this position: she can support or attack the position.

After the supporting or the attacking option is chosen, the user is asked to select or provide a statement explaining her choice. This statement is used as the premise, whereby the position (or its negation) is the conclusion. This completes the first argument and ends bootstrapping.

6.2. Selecting the Next Argument

The selection of the next argument that is presented to the participant can be based on different selection strategies. We have chosen a simple antagonistic strategy that mimics typical human behaviour in an argumentation: we look at the participation history of a user to identify the most recent argument that she selected. Then we search the WoR for an argument of prior users which attacks (undermine, rebut or undercut) that argument. This argument is shown to the user who then has the opportunity to react to it. This process continues until the WoR contains no counter argument to the argument of the current user. The overall intention is to simulate a real discussion where participants challenge the arguments of other participants.

6.3. The Quick Lateral Entry

When a user has already particpated in a discussion she should be able to navigate directly to an argument instead of starting from scratch. To this end, we provide two additional means of navigation: First, the user can search for any statement via a query mask. As outlined above she is then invited to indicate her attitude towards this statement, completing the alternative bootstrapping process. Second, the underlying data structure can be viewed as interactive argumentation map. After the selection of a statement, the system displays all arguments connected to this statement. The user is then invited to select any of these arguments.

7. The Third Challenge: Accepting and Maintaining Arguments

The key to incorporating new arguments in dialog-based online argumentation is to nudge the users to provide arguments in an appropriate way. Currently, we use three mechanisms for this purpose.

First, users can enter their own statements only within the dialog, so that given statements automatically are connected to the WoR in an appropriate fashion. Second, we apply sentence openers to frame the statements of the users. In this way the user is guided towards making structured and well-formed statements. Third, we automatically match the text entered by a user with existing statements in the WoR by means of Elasticsearch[2]. This reduces redundancy if the user chooses to use an existing statement instead of providing a new statement. An example of statement proposals during the users input as well as the sentence opener is given in Figure 2.

While those mechanisms improve the quality of the arguments provided by the users, they can not prevent that a given user input is incorrect. To address this problem we use a decentralized moderation system, so that the every participant is able to review statements and propose improvements by means of review queues. If one user flags a statements due a specific reason, other users can go through those queues and vote on the action to be taken. Once a sufficiently clear-cut collective opinion has been reached, the appropriate action is taken, e.g. the statement might be replaced or deleted or the flagging might be discarded. Based on our experience with dialog-based online argumentation we suggest the following review queues:

- *Delete:* Statements, which have been flagged as harmful, abusive or offtopic, will be deleted, if positive collective consensus is reached.
- *Duplicate:* Statements which have been identified as a duplicates will be merged, if a positive collective consensus is reached.
- *Edit:* Proposals for updating already existing statements. If positive collective consensus is reached, these statements will be replaced with the proposed version.

[2] https://www.elastic.co/

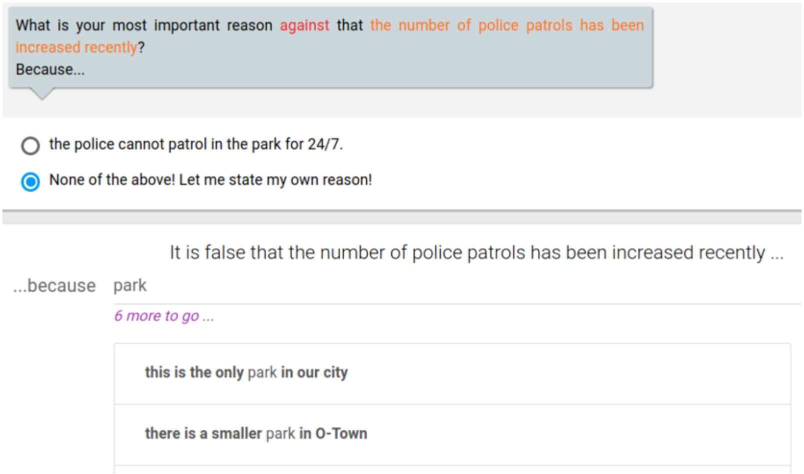

Figure 2. Example of statement proposals during the users input as well as the sentence opener.

- *Optimization Required:* Sometimes a user is not able or willing to provide an update for a statement that requires a revision. In this case the statement can be flagged and other users can provide an update while going through this queue. The update will then be submitted to the edit queue and the respective entry will be deleted from the optimization required queue once positive consensus is reached.
- *Split:* Even though we use framing, on rare occasions users include both premise and conclusion or multiple distinct premises in a single text contribution. This can be flagged by other users. If a positive collective consensus is reached, the flagged statement will split appropriately.

In order to motivate users to participate by providing statements or by taking part in the review system, they gain reputation by helpful actions and in order to deter them from abusing the system, they loose reputation if their actions are considered unhelpful. The actions that a user can take, in particular which review queue she can use, depends on the reputation of the user.

8. D-BAS: Implementing a Dialog-Based Online Argumentation System

We have developed an application for dialog-based online argumentation called *Dialog-Based Argumentation System* (D-BAS) which implements the ideas described above. It is available both as a web-based service[3] and as open source software[4].

D-BAS' backend is written in *Python 3* with usage of the *Pylons Pyramid web framework*. We use *nginx* as proxy, *uwsgi* as webserver and *Chameleon* as HTML template engine. Additionally we use *Node.js* with *Socket.IO* for asynchronous and bidirectional communications, e.g. notifications about specific events. D-BAS' data structure is managed by Pythons SQL toolkit *SQLAlchemy*. D-BAS' frontend is built upon a number of established technologies like *HTML*, *JavaScript* with *jQuery*, *Bootstrap* and *SASS*.

[3] https://dbas.cs.hhu.de
[4] https://github.com/hhucn/dbas

To allow future applications to use the functionality of D-BAS we implemented an *application programming interface* (API), which provides the possibility to access D-BAS' backend. This abstraction of the core argumentation functionality of D-BAS can then be used to enable dialog-based discussions in an arbitrary context.

Furthermore, it is possible to query the contents of the database in a very flexible way, where a developer can specify which data she wants to access. Most contents of our discussions are freely queryable except for private information, e.g. the user's password-hash. For this we implemented *GraphQL*, which allows flexible queries for the required data. Additionally, we offer *GraphiQL*[5] as an in-browser IDE for exploring GraphQL. Querying data or giving the transferred statements a structure is the main goal of the *Argument Interchange Format*, an universal format for the exchange of arguments by [16], but we decided to make it more flexible for the developers to query data from our database. All necessary information to use the API and more examples are available in the documentation of D-BAS[6].

9. Findings from a real-world online participation process based on D-BAS

D-BAS was used in a real-world online participation process where all students of our computer science department were invited to propose and discuss improvements to the computer science studies program. The main issue was how to deal with an increased number of students. The number of enrolled students has more than doubled in the past three years, leading to numerous problems such as overcrowded lectures.

A full report on this process an its results can be found in [3]. Here, we only seek to answer the question: "Does D-BAS enable a large group of untrained participants to make proposals and discuss then in such a way, that the resulting WoR is reasonably well-formed and helpful for those that have to evaluate the proposals?"

The online participation process took place from may, 9^{th} until may, 28^{th} 2017 and all students of our computer science department were invited via e-mail. In total there were 318 unique visitors who added 22 positions and 255 statements. The resulting argumentation map is shown in Figure 3a. The typical depth of a sequence of arguments varies between three and four. This clearly shows that there was (time-shifted) interaction between the participants.

In order to allow others to analyze the discussion, we summarized the main facts online[7] and we offer a dump of our database, which is licensed under the Creative Commons License CC BY-NC-SA[8].

As a first step we investigated the quality of the resulting argumentaion by taking a look at the proposed positions. The students added 22 positions to the argumentation, where each position lead to further reactions, indicating that they were of interest to others. Furthermore, no position was a duplicate of another position and all proposed positions where reasonable. While it is not possible to prove that no other means of online argumentation might lead to more or better positions, the absolute number indicates that the argumentation was extremely successful at gathering meaningful positions.

[5] https://dbas.cs.hhu.de/graphiql
[6] https://dbas.cs.hhu.de/docs/api/v2.html
[7] https://dbas.cs.hhu.de/fieldexperiment
[8] https://creativecommons.org/licenses/by-nc-sa/3.0/

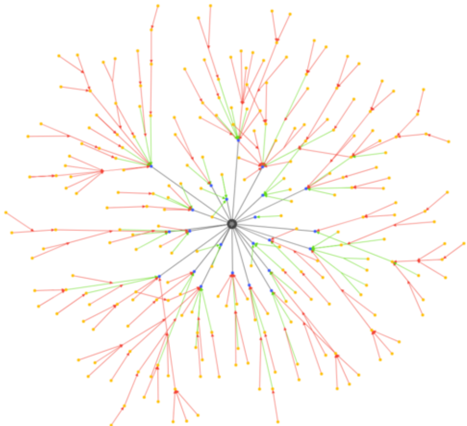

(a) Complete view of the discussion. The grey dot is the root of the discussion, the blue dots are positions and the yellow dots are statements that are not positions. Green arrows denote supporting arguments and red arrows denote attacking arguments.

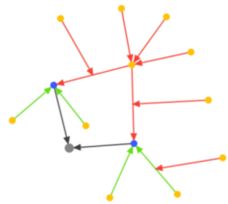

(b) Enlarged view of a connected subgraph during a discussion.

Figure 3. Argumentation graph created by participants in D-BAS.

Second, we examined how the participants interacted during the online argumentation. The discussion consists of 265 statements in total. In order to examine the participant's interactivity, it is important to understand how the results of the argumentation look like. Essentially, each position is the start of a sub-graph of arguments. Since statements can be reused, the sub-graphs of the positions are interconnected, as shown in Figure 3b. The size of the subgraphs was between 2 and 44 with an average of 13. This shows that for all proposals there was a significant exchange of arguments.

Third, we analyzed the selected feedback options. Users selected 200 undermines, 44 supports, 137 undercuts, 56 rebuts; 19 times they wanted to see another attacking argument and 104 times they went a step back. We manually investigated if those reactions were used appropriately, that is, if the reaction made sense in relation to the argument it was a reaction to. This holds true for every single reaction. This is surprising, since at least the undercut is a challenging type of reaction. While we were very pleased with this result, it should be noted that the participants were all computer science students. It is not certain that this result would remain unchanged with a different set of participants.

Summarizing, the field experiment indicates that it is possible to lead a high quality online argumentation by using dialog-based online argumentation. It demonstrates in a real-world setting that participants with no background in formal argumentation are able to collectively argue about a topic in such a way that the resulting formal argumentation map is reasonable and non degenerated.

10. Conclusion

In this paper we have presented D-BAS, a system for dialog-based online argumentation. We have identified and solved three main challenges: providing feedback on existing arguments, selecting the next argument that should be presented to the user and incorporating user input. The resulting system was fully implemented as an open source, web-based service with a well-defined API that can be used by other applications. Further, in a real world deployment we have shown that untrained participants of an e-participation process are able to use D-BAS and that the resulting WoR is reasonably well-formed.

Acknowledgements

This work was done in the context of the PhD programme "Online Participation"[9], funded by the Ministry of Innovation, Science and Research in North Rhine Westphalia, Germany. We thank Björn Ebbinghaus and Marc Feger for their assistance.

References

[1] M. Klein, "Using Metrics to Enable Large-Scale Deliberation," in *Collective intelligence in organizations: A workshop of the ACM Group 2010 Conference*, 2010, pp. 103–233.

[2] T. Krauthoff, M. Baurmann, G. Betz, and M. Mauve, "Dialog-Based Online Argumentation," in *Proceedings of the 2016 Conference on Computational Models of Argument*. iOS Press, 2016.

[3] T. Krauthoff, C. Meter, and M. Mauve, "Dialog-Based Online Argumentation: Findings from a Field Experiment," in *Proceedings of the 1st Workshop on Advances in Argumentation in Artificial Intelligence*, vol. 2012. CEUR-WS, 2017, pp. 85–99.

[4] L. Iandoli, I. Quinto, P. Spada, M. Klein, and R. Calabretta, "Supporting argumentation in online political debate: Evidence from an experiment of collective deliberation," *New Media & Society*, 2017.

[5] M. Van Alstyne, E. Brynjolfsson *et al.*, "Electronic Communities: Global Village or Cyberbalkans?" in *Proceedings of the 17th International Conference on Information Systems. New York: Wiley*, 1996.

[6] C. R. Sunstein, "The law of group polarization," *Journal of political philosophy*, vol. 10, no. 2, pp. 175–195, 2002.

[7] M. Klein, "Enabling large-scale deliberation using attention-mediation metrics," *Computer Supported Cooperative Work (CSCW)*, vol. 21, no. 4-5, pp. 449–473, 2012.

[8] T. Kriplean, J. Morgan, D. Freelon, A. Borning, and L. Bennett, "Supporting Reflective Public Thought with ConsiderIt," in *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. ACM Press, 2012, pp. 265–274.

[9] T. F. Gordon and D. Walton, "The Carneades Argumentation Framework – Using Presumptions and Exceptions to Model Critical Questions," in *6th computational models of natural argument workshop (CMNA), European conference on artificial intelligence (ECAI), Italy*, vol. 6, 2006, pp. 5–13.

[10] M. Klein and L. Iandoli, "Supporting Collaborative Deliberation Using a Large-Scale Argumentation System: The MIT Collaboratorium," 2008.

[11] S. Wells, "Supporting Argumentation Schemes in Argumentative Dialogue Games," *Studies in Logic, Grammar and Rhetoric*, vol. 36, no. 1, pp. 171–191, 2014.

[12] T. Bench-Capon, K. Atkinson, and A. Wyner, "Using Argumentation to Structure E-Participation in Policy Making," *Transactions on Large-Scale Data-and Knowledge-Centered Systems XVIII*, vol. 8980, pp. 1–29, 2015.

[13] T. F. Gordon, "Structured Consultation with Argument Graphs," *From Knowledge Representation to Argumentation in AI. A Festschrift in Honour of Trevor Bench-Capon on the Occasion of his 60th Birthday*, pp. 115–133, 2013.

[14] F. Bex, J. Lawrence, M. Snaith, and C. Reed, "Implementing the Argument Web." *Communications of the ACM*, vol. 56, no. 10, pp. 66–73, 2013.

[15] S. Wells and C. Reed, "MAgtALO: an Agent-Based system for persuasive online interaction," in *AISB 2008 Convention Communication, Interaction and Social Intelligence*, vol. 1. Citeseer, 2008, p. 29.

[16] C. Chesñevar, J. McGinnis, S. Modgil, I. Rahwan, C. Reed, G. Simari, M. South, G. Vreeswijk, and S. Willmott, "Towards an Argument Interchange Format," *The Knowledge Engineering Review*, vol. 21, no. 04, pp. 293–316, 2007.

[9] http://fortschrittskolleg.de/en