

Missing Phases of Deliberation Dialogue for Real Applications

Douglas Walton¹, Alice Toniolo², and Timothy J. Norman²

¹ University of Windsor

Centre for Research in Argumentation, Reasoning and Rhetoric

`dwalton@uwindsor.ca`

² Department of Computing Science, University of Aberdeen, Scotland, UK

`{a.toniolo,t.j.norman}@abdn.ac.uk`

Abstract. Models of deliberative dialogue are fundamental for developing autonomous systems that support human practical reasoning. However, we must consider whether these existing models are able to capture the complexity and richness of natural deliberation for developing real applications. In real contexts, circumstances relevant to the decision can change rapidly. In this paper, we introduce an extension to today’s leading model of deliberation dialogue to capture dynamic changes of circumstances during dialogue. Moreover, in natural deliberation, a dialogue may be successful even if a decision on what to do has not been made. A set of criteria is proposed to address the problem of when to close off the practical reasoning phase of dialogue. We then discuss some initial efforts to introduce those characteristics within an existing model of deliberation for agent collaboration. We believe that our extended model of dialogue may represent that richness of natural deliberative dialogue that is yet to be addressed in existing models of agent deliberation.

Keywords: Deliberative dialogue, practical reasoning

1 Introduction

Practical reasoning is the inferential process of arriving at a conclusion to take action. Deliberative dialogue serves as a rational method of decision on what to do, thus rational deliberation cannot be understood without the fabric of practical reasoning that holds it together. Deliberation is often taken to be a solitary process in which an individual arrives at an intelligent conclusion on how to act by forming a goal and collecting data. But even an individual deliberation can be seen as a process in which one “thinks aloud” by asking questions and then answers them himself. By answering such questions in a dialogue format, one can clarify his goals in light of his present situation and its potential future consequences. Practical reasoning in this context is seen as a form of argumentation. Group deliberations are also common as typically an individual is not acting in isolation, where the goal of participants is to decide what to do in collaboration.

Practical reasoning is foundational to current research initiatives in computing, especially in multi-agent systems; e.g., the project of designing systems for

electronic democracy [3]. The capability of the user, or of the agents within the systems, to pose critical questions and reply to counter-arguments for a proposed plan of action is vital for this application, as such deliberations are only useful if weak points in a proposal can be questioned. Often there is a gap between agent-based and human reasoning, however. Formal models of argumentation-based dialogue exploit natural ways for humans to represent justifications and conflicts in making decisions [1, 4]. These models, however, often detach from the natural argumentation to focus mainly on logical formalisms since there is the need to identify tractable models to be developed for real applications. There is also a lack of rigorous evaluation of the effectiveness of these systems, in fact this research typically proves some formal properties, and then illustrates its applications through the use of examples (cf. [7]). In order for argumentation models to be successfully deployed in real-world practical applications where humans and autonomous agents collaborate in making decision, these systems must resemble natural language argumentation and their benefits must be evaluated. However, while formal systems are developed from models of human deliberation, there is little effort in the literature to verify to what extent the richness of human practical reasoning can be represented by those systems. The question that we address in this paper is: are existing models of autonomous deliberative dialogue able to capture the complexity and richness of realistic natural deliberation?

We discuss the leading model of deliberation dialogue in artificial intelligence [5]. We argue that this is the right type of framework to represent how practical reasoning is used in realistic cases, but at least two important additional phases should be integrated. In real world scenarios, the circumstances relevant to the decision can change rapidly. Participants may also share individual knowledge during dialogue that changes the circumstances. It is vitally important for the agent who is to deliberate intelligently to be aware of these changes, thus we propose an extension to the information-seeking phase of the dialogue that takes changes into account during the course of the deliberation itself. Moreover, in many instances, a deliberative dialogue can be very successful educationally in revealing the arguments and positions on both sides, even though the dialogue did not succeed in determining what to do. This raises the problem of when practical reasoning should be closed off and how to establish the success of a dialogue. We propose a solution to this closure problem by presenting ten criteria that determine when a deliberation has been successful in realistic cases.

In addressing our question, we show that these phases of natural dialogue are to a certain extent addressed in the argumentation-based deliberation system presented in [7]. We show that this model is able to capture the information sharing that changes circumstances during the dialogue. We demonstrate through empirical evaluation that this is fundamental for obtaining more successful outcomes in a situation where new information is continually streaming in. We then propose some results on the evaluation of three criteria proposed for successful dialogue. We argue that this dialogue system is a first attempt to fill some of the gaps between human and autonomous deliberation identified in this research.

2 Limitations of Today's Deliberative Dialogue Protocols

In this section we present the leading model of deliberation dialogue by McBurney, Hitchcock and Parsons [5]. We discuss the important phases and we discuss some missing phases required to represent a more realistic deliberation process. For convenience, hereafter we refer to the McBurney, Hitchcock and Parsons deliberation dialogue model [5] as the MHP model.

Deliberation in the MHP dialogue is seen as a resource-bounded procedure that starts from an initial situation where a choice has to be made, and then goes through several other stages as it moves towards a closing stage where the decision is arrived at on the basis of pro and con arguments that have been put forward in a middle stage [9]. According to the MHP dialogue model, deliberation is a formal procedure that goes through eight stages. The dialogue has an opening stage where the question is raised about what is to be done, and a closing stage where the sequence of deliberation is ended. This special question raised at the opening stage is called the *governing question*, meaning that this single question governs the whole dialogue, including the opening and closing stages.

1. The Opening Stage. The governing question is raised.
2. The Inform Stage. Next there is a discussion of goals, any constraints on the actions being considered, and any external facts relevant to the discussion.
3. The Propose Stage. At this stage proposals for possible action are brought forward.
4. The Consider Stage. Comments are made on the proposals that have been brought forward. At this stage, arguments for and against proposals are considered.
5. The Revise Stage. At this stage, the goals, the actions that have been proposed, and the facts that have been collected and that are relevant to the deliberation may be revised. Another type of dialogue, for example a persuasion dialogue, may be embedded in the deliberation dialogue.
6. The Recommend Stage. At this stage participants can recommend a particular action and the other participants can either accept or reject this option.
7. The Confirm Stage. At this stage, all participants must confirm their acceptance of one particular option for action in order for the dialogue to terminate.
8. The Close Stage. At this point the dialogue terminates.

In the MHP dialogue, a formal protocol describes what moves can be made by each party at which stages of the dialogue. They define the speech act of propose, *propose()*, as a valid instance of the type of speech act containing a goal, a fact and an action, including some other elements called constraint, evaluation, and perspective. They also have speech acts for making an assertion, *assert()*, preferring a particular option for action, *prefer()*, asking the other party to justify an assertion, *ask_justify()*, pronouncing on whether a proposal for action should be accepted or rejected, *assert(action)* and *reject()*, retracting a previous location, *retract()*, and withdrawing from the deliberation dialogue *withdraw()*.




Fig. 1. Outline of the McBurney, Hitchcock and Parsons (MHP) Model

The structure of the MHP dialogue is shown in Figure 1. In the opening stage the goal of the dialogue is defined. The goal is for the participants to decide what is the best decision. Once the opening stage is set in place, the dialogue moves forward to a more complex middle stage that has several components. For convenience sake, we call this middle stage the argumentation stage. Whether the case is one of a single agent or group deliberation, the first step in this middle stage is to inform the agent(s) what the circumstances are. Thus, there needs to be an information-seeking dialogue embedded into a deliberation dialogue. For the purposes of modeling an abstract normative model of rational deliberation, the information-seeking part needs to be seen as a distinctive element, and because of its importance it is placed at the beginning of the middle stage.

Once the information is in, the second part of the middle stage is putting forward proposals by all parties to the deliberation dialogue. For this purpose, there is a special speech act of making a proposal. When the agents in the dialogue have formulated some proposals, the deliberation dialogue moves to the third part of the middle stage. This stage corresponds to the consider stage and the revise stage. During this stage each party probes into the proposals put forward by other parties by raising questions about possible problems with the proposal and even by attacking the proposal with counterarguments. In response to an attack on its proposal, the agent may concede, and retract its commitment or refine the proposal. This consider-revise part of the middle stage forms one stage in its own right, because considering and revising naturally go together.

Example 1. The MHP protocol represents many examples of realistic argumentation in deliberation dialogue. We present here a dialogue between Alice and Bob [10]. The governing question and the inform stage are motivated by the fact that they would like to find a suitable house in Windsor. At the propose stage




Fig. 2. Carneades [3] Argument Map of Arguments Supporting the Condo Choice

they narrowed their choices down to three: a condominium, a two-story house, and a bungalow. Both Alice and Bob place a high value on health, costs, and on the environment. A segment of the consider stage of dialogue follows:

Bob: The condo and the bungalow are in the same area, where I can ride to work on my bike in 35 minutes. The bike path goes right along the river straight to my office.

Alice: The problem with the two story house is that it is twice as far away as the two other homes. It would take over an hour for you to ride the bike in.

Bob: If we lived there, I would have to drive the car to work most days.

Alice: Riding the bike to work is more environmentally friendly. Also, you really like riding your bike to work, and it is good exercise.

Bob: Yes, and there is also the factor of the cost of gas. The cost of driving to work adds up to a significant amount over a year. Also, neither of us likes spending a lot of time in the car. It is wasted time, and getting exercise is a good way to spend that time. When you spend your whole day working on a computer, it is really important to get some exercise to break up your day.

Alice: That is offset by the lower taxes of the two story house, because it is outside the city, even though it is a larger house.

Bob: I say let's make an offer on the condo.

The arguments for the condo are shown in Figure 2. Alice and Bob arrive at a decision by listing of all the positive and negative factors of all three options, then weighing each factor that is more or less important to each of them. They combine all this information to determine which option has the most positive and the least negative factors attached to it.

The MHP protocol formally represents the stages of the deliberation between Alice and Bob. However, if the circumstances change, and more information is

needed before making a decision, the protocol may not be able to capture such characteristics. In fact, there is an assumption that all the participants have already gathered full knowledge of the circumstances either prior to the dialogue or during the information-seeking stage. New knowledge can only be added to the commitment store using speech acts *assert(fact)*, exclusively during the *Inform* stage. This ensures that no other information shared during the dialogue will change the view of the circumstances. The deliberation dialogue always arises out of the question of what to do in any given set of circumstances. The information relevant to the circumstances is vital to arriving at a rational decision, therefore, the collection of data relevant to the decision needs to be seen as an important requirement for intelligent deliberation. In any real case, however, new information will need to be introduced during all sequences of argumentation prior to the closing stage and the dialogue model must support the new changes in the shared knowledge base that is used during the information-seeking phase.

Example 2. Brian had a problem with his printer. Whenever he scanned a document using the automatic document feeder, a black line appeared down the middle of the page. To solve the problem he searches for a troubleshooting guide for its printer. This guide gave a series of instructions. If this problem happens when you print and copy/scan, the problem is in the print engine itself, hence refer to a certain website for help troubleshooting. However in Brian's case, the defect only occurred when scanning from the automatic document feeder. Following the instruction in the troubleshooting guide, first Brian opened the scanner cover checked for debris. There was none. Next he located the small strip of glass at the left of the main glass area. Next he carefully cleaned this strip of glass with a soft cloth, and then scanned a document to see if this fixed the problem. It did not. Then he looked at the small strip of glass and he could see that it was covered by a thin plastic piece. He then managed to pull the plastic piece out. He found a small black mark in the middle on the bottom of the clear plastic. He tried cleaning the plastic, but it did not work, as he found by scanning a test document. He then showed the plastic piece to his wife Anna, and asked her if there was some way it might be possible to clean it to remove the small black mark. She used a soft cleaning pad and managed to remove the black mark. Brian went through the scanning procedure as a test. Success! There was no longer a black line down the middle of the pages that were scanned. Brian's problem was solved. The steps of the decision are shown in Figure 3.

One of the most important aspects of this case is that the deliberation was based on a prior process of information-seeking dialogue. Moreover, the information-seeking dialogue was interwoven with the steps taken during the sequence of practical reasoning in the deliberation dialogue. The MHP dialogue permits the representation of Brian's monologue in order to make a decision on what to do to fix the printer, considering the search for information as the troubleshooting instructions. However, if an agent were to act on Brian's behalf, it would have to add a new piece of information to its knowledge base; i.e., "The scanner is not working even if the instructions were followed". Since this information changes the circumstances, it should also be asserted in the dialogue.




Fig. 3. Sequences of Problem Solving in the Printer Example

Otherwise the agent would continue to propose the same decision; i.e., “Follow the instruction in the guide”. As previously discussed, the MHP model does not consider these changes in the knowledge-base during the dialogue.

Continuing with our analysis of the MHP dialogue, we see that unanimity of the participants is required for a decision on the course of action to be taken, but this assumption is made for purposes of simplicity in the initial model of deliberation. It is also possible, however, that in real instances of deliberation, there may be no agreement on the best proposal for action to solve the problem posed by the governing question. In this case, in the MHP model, presumably the deliberation dialogue must be considered a failure, because the governing question has not been answered. Nevertheless, the deliberation could have educational value, in that the consideration of the pro and con arguments might have shown deficiencies in some of the proposals. This revelation might have deepened the understanding of the participants. Note that the only stages which must occur in every dialogue that terminates normally are the opening and the closing stage. It appears that we must have the closing stage in every deliberation dialogue, or at least in every successful deliberation dialogue. Hence, in the MHP model, a dialogue cannot really be a deliberation dialogue unless it terminates in a closing stage where a decision is arrived at on the best course of action to take. This remark is a revealing comment on the general problem of formulating goals in formal systems of dialogue. It also affects the persuasion type of dialogue. The goal of a persuasion dialogue is to resolve the conflict of opinions agreed upon as the issue at the opening stage. However, in many instances, a persuasion dialogue can be very successful educationally in revealing the arguments and positions on both sides. Determining the criteria of success of a dialogue in meeting its goal is a general problem for formal dialogue systems.

For the above reasons the MHP model of deliberation dialogue needs some revisions to bring it in line with an open knowledge-base and with a closing stage that considers when the dialogue ends and when it is to be considered successful.

3 An Extended Dialogue Model for Open Knowledge

In this section, we propose an extension to the MHP deliberation model that makes the role of the knowledge base of the agents more explicit and takes into account changes of the circumstances of deliberation.

One of the most important factors in intelligent deliberation is that in real scenarios the circumstances of the world can change, and it is vitally important for the agent who is to deliberate intelligently to be aware of these changes and to take them into account during the course of the deliberation itself. Intelligent deliberation needs to be both informed and flexible. The agent needs not only to be aware of the relevant circumstances, but also to be aware of relevant changes in them. In fact, the knowledge base that agents have may be incomplete. Some of the agents may know things that the others do not know, and it is important that the knowledge base is left open during the argumentation stage so that new, relevant information that might affect proposals and commitments can come in. The MHP model has led to the development of a number of deliberation protocols employed in multiagent systems where agents have complete shared knowledge of the circumstances [1, 4]. One of the worst errors in a rational deliberation, however, is for the agent to become inflexible by failing to take new relevant developments into account in arriving at an informed decision on what to do.

For this reason, deliberation as a framework for rational argumentation needs to be extended from the structure described above. Instead of the knowledge base being fixed at the opening stage, propositions need to be added to it and deleted from it as the sequence of deliberation proceeds through the argumentation stage. In other words, deliberation is not merely the putting forward of proposals and the acceptance or rejection of those proposals based on arguments for and against them without considering changes of circumstances. It must also comprise incoming knowledge of the circumstances during the argumentation stage that may affect how these proposals are to be evaluated.

In the revised model, outlined in Figure 4, there needs to be a knowledge base set in place at the opening stage and agreed to by all parties in the deliberation as representing the circumstances of the situation in which the governing question is framed. This knowledge base is fixed in place at the opening stage and is part of what defines the choice that is to be made as stated by the governing question of the deliberation. However, once we reach the argumentation stage, this knowledge base needs to be opened so that if the situation changes, and information about these changes becomes available to the agents, this knowledge is taken into account during this stage where proposals are put forward and considered. Hence, as shown in Figure 4, the opening stage includes a knowledge base, but when this knowledge base is opened, new information can come in.

As pointed out in the MHP model, there will need to be dialectical shifts during the argumentation stage from the deliberation to an information-seeking dialogue in which new knowledge of the circumstances comes in. If new circumstances relevant to the choice to be made in the governing question are found, or come to be known by information-seeking, there is a shift to the next phase where proposals are put forward, considered and evaluated. The proposals may




Fig. 4. Outline of a Revised Model of Deliberation Dialogue

need to be modified or evaluated differently once this new information comes to be known. During this stage, some agents will have knowledge that others lack, and thus one important type of speech act is that of asserting a proposition that represents factual knowledge, *assert(fact)*, to inform the other agents involved in the deliberation process. In contrast with the MHP model, this *fact* may be a new piece of information that was not known at previous stages of the dialogue. Thus, during the argumentation stage, new information that comes in may alter the structure of the choice being made, by dynamically revising each participant's knowledge of the world, and this leads to suggesting a new option.

In the new model, there is a cycle of proposing, considering and revising as new information comes in, and evaluating the proposals in light of this new information. During the argumentation stage there can be a cyclical flow of argumentation as new knowledge comes in that requires rethinking and re-evaluating proposals. There is a danger that this cyclical process can continue infinitely, stalling the deliberation. Hence, it is necessary to have a closing stage so that the argumentation stage can be terminated once it has been judged that enough information on the circumstances of the decision has been taken into account. Further intake and evaluation of new information would mean delay of a kind that would interfere with the making of a timely choice necessary for properly answering the governing question. At some point a decision may have to be made on determining which is the best proposal given the knowledge of the circumstances that has been obtained by the research carried out so far.

In this section, we discussed the need for the MHP model to be revised, especially in relation to the cases where the agents taking part in the deliberation have partial views of the circumstances of the world. The revised version offered

here provides for the sharing of information between agents while deliberating so that better proposal can be identified. In the next section we discuss in more detail the problem highlighted above of the closing stage of deliberation.

4 Criteria for Closure of Deliberative Dialogue

In Section 2, we discussed the dialogue structure in terms of an opening, an argumentation and a closing stage. The closure problem is the problem of determining when practical reasoning ends. When can the searching for knowledge about a case be closed off so that the premises of the practical reasoning provide an evidential base sufficient to prove the conclusion? The closure problem of a deliberation dialogue is a problem of determining the conditions for closure of practical reasoning. A deliberation may have to be closed off and a decision taken based on the pro and con arguments put forward for practical reasons, typically time and money. A decision by majority vote may have to be taken to meet the practical demand for closing of the discussion. However, the depth, comprehensiveness and thoroughness of the pro and con arguments brought for and against the proposals is the most valuable feature of a deliberation leading to an supported conclusion offering a well-reasoned decision for action. There are also some issues remaining about how to represent the closing stage. What if it cannot be determined, on the evidence, which is the best proposal?

In a case of deliberation, like that between Alice and Bob, the dialogue can be closed when they have collected enough information about what is available on the real estate market, and they have discussed the matter thoroughly enough to critically examine all the pro and con arguments on both sides of every available option. At that point, even though the deliberation could be reopened before they make an offer and are committed to it, the dialogue has reached the closing point. If one party has put forward a particular proposal, and nobody else has any objections, then, at least temporarily, the deliberation is closed. They have found a proposed course of action that they all agree on, or at least that nobody disagrees with, or is willing to contend further. Hence the deliberation is over. Based on the arguments for and against brought forward and criticized during the dialogue sequence, a reasoned decision can now be arrived at on what to do. There are four options. The original proposal can be accepted or the original proposal can be rejected, in which cases the deliberation stops. The proposal can be modified, to consider the objections made during the dialogue, or a counterproposal can be brought forward. In the latter two cases, the procedure of formulating a proposal goes around the cycle again, except that this time a new proposal has been formulated in place of the one that was not accepted.

The decision of when the deliberation has reached its closing stage cannot be made by any current model of deliberation. In an emergency, the closing stage may have to be reached quickly. In other cases, determining when the closing stage has been reached may depend on costs, or on how much time there is for discussion. In theory, the closing stage should only be reached when the arguments and proposals considered on all sides have been thoroughly enough

discussed so that all the relevant factors have been discussed. Then, an evaluation of all the argumentation that has taken place through the argumentation stage should be evaluated thoroughly enough so that all the potentially usable proposals have been discussed to decide which is the best. In practice, a decision in the circumstances may have to be made within time constraints, and so a determination of which proposal is best may still be subject to pro and contra argumentation. In such cases, the argumentation stage should be closed off and some means taken to arrive at a decision, for example taking a vote. However, in this context, can we determine whether a dialogue has been successful?

In this framework, we define ten criteria that determine the extent to which a deliberation has been successful:

1. Whether the proposals that were discussed represent all the proposals that should be considered, or whether some proposals that should have been discussed were not.
2. The accuracy and completeness of the information regarding the circumstances of the case made available to the agents during the opening stage.
3. How well arguments were critically questioned or attacked by counterarguments.
4. How well the agents followed the procedural rules by allowing the other agents to present their proposals and arguments openly, and how they responded to proposals.
5. How thoroughly each of the proposals that were put forward during the deliberation were engaged by supporting or attacking arguments.
6. Whether any arguments that should have been considered were not given due consideration.
7. How good the arguments were supporting or attacking each of the proposals, depending on the validity of the arguments and the factual accuracy of their premises.
8. Whether the argumentation avoided personal attacks, or was unduly influenced by opinion leaders or personalities who dominated discussion during the argumentation stage.
9. The relevance of the arguments put forward during the argumentation stage.
10. The taking into account of the values of the group of agents engaged in the deliberation.

Relevance of an argument is determined by how it fits into a sequence of argumentation that connects to the problem or choice of action set as the issue of the deliberation at the opening stage. To declare that the closing stage has been reached, the participants in the deliberation must reach a consensus on whether the deliberation has been successful, based on the ten criteria above, so that one proposal has been shown to be superior to the others. However, if for practical reasons, such as time limits, this cannot be shown, then they can vote on the proposal using majority rule.

The judgment of the success of the deliberation can be evaluated in two ways. As indicated just above, the participants can reach a consensus that one proposal has been shown to be superior. This is called an internal evaluation. The other

way is an external evaluation carried out at a metalevel. Once a deliberation has been carried out in a given case by a group of agents, another group of agents can then keep a record of the argumentation in the deliberation, for example by keeping a transcript of the discussions. Then they can analyze and evaluate the argumentation in the deliberation dialogue, and arrive at an evaluation of how successful the dialogue was according to the ten criteria set out above.

Satisfying these criteria implies that enough proposals have been thoroughly discussed and evaluated (as instances of practical reasoning) so that further discussion is unlikely to turn up anything new or useful for making a good decision or solving the problem. The dialogue may be re-opened if it turns out the proposed solution does not work once it has been tried. However, because of time or cost limitations that do not permit the collecting of more evidence, the deliberation may have to be closed off prematurely. This decision has to be made on a basis of practical reasoning. More discussion might mean delay, and that might have negative consequences that interfere with the goal of the deliberation dialogue. This solution to the closure problem, while not yet implemented in a computational system, is useful for the project of improving the current models of deliberation and may also give useful insights on how to deal with the closure problem in other types of dialogue such as persuasion or information-seeking.

5 Effects of Open Knowledge-Base in Agent Applications

Our objective is to understand whether existing models of deliberative dialogue capture the richness of human deliberation. We argued that an extended model is necessary to capture the changes of circumstances during dialogue. In this section, we analyze a model of deliberation for agents discussing interdependent plans with individual goals, presented in our previous research [6, 7]. We show how this model permits agents to gather information about circumstances, not only at the initial stage of the dialogue but during its course. We empirically demonstrate that this is necessary for agents to identify an agreement in situations where agents do not have a shared knowledge about each other's plans.

The problem addressed by Toniolo et al. [7] is one of coordination of goals and actions between agents where a decision made by one of them may interfere with a decision made by others due to differences in goals. Each agent's plan is internally consistent with regard to its knowledge of the circumstances and of how actions can be carried out. However, in a typical case there may be conflicts between independent plans. The model of dialogue uses argumentation schemes to deal with problems caused by conflicting goals, scheduling constraints and norms. These schemes are reasoning patterns formed by a set of premises in favour of a conclusion and a set of critical questions that can be put forward against a stated argument [8]. A key feature of this model is that it provides for the sharing of information between agents during the deliberation process.

The example of the deliberation used is one where agents are discussing the repair of the water supply in a particular location where a disaster has occurred. One agent proposes to stop the water supply in the location while the other

argues about the need for water to perform activities that include the building of a field hospital to deal with casualties. The first agent, however, has a goal of stopping the water supply that is contaminated. Water that is contaminated is not safe, and that constitutes an argument against allowing a water supply in the location. The two agents then discuss other options, for example building the field hospital first but not using the water supply until after it can be guaranteed that the water is safe. They discuss other options as well, for example building the field hospital in a different location. By communicating the circumstances to each other, and by constructing arguments to respond to arguments put forward by the other side, for example the argument that building water supply at this location would not be safe, the dialogue can proceed in an orderly way.

We consider two agents, each with individual goals and norms, that form interdependent plans to achieve their goals. An agent proposes an action A_{dk} for discussion that may affect the other agent's plan. We assume that there are a number of actions A_{dl} to be discussed in each agent's plan. The dialogue proceeds in a turn-taking fashion following a simplified version of the dialogue protocol of Kok et al. [4]. This protocol comprises the deliberation stages and the speech acts of the MHP model. After the opening stage, an action A_{dk} is proposed by an agent. Then, agents exchange pro or con arguments for the adoption of A_{dk} . When an agent fails to defend its claims, it must pass and the dialogue pauses. Agents may replan considering new information acquired during the dialogue. If the issue has not been solved, and an agent finds a suitable alternative, a subsequent argumentation phase is initiated to debate the new proposal. The dialogue terminates when an agreement is found or no other alternatives exist. In the latter case, if the proponent failed to defend A_{dk} , the action and the goal that that action would contribute to achieve are dropped. If the opponent failed, it is forced to rearrange its plan to include A_{dk} dropping some of its goals if necessary. If there are other actions A_{dl} to be discussed a new dialogue will be opened. When the agents have no further issues, they will inform each other that they are satisfied and the process ends. The deliberation is shown in Figure 5.

The characteristic of this dialogue structure is that it allows the agents to construct alternatives dynamically as they acquire information during the sequence of dialogue. This exchange of information enables them to reduce the number of conflicts in their interdependent plans related to scheduling or norm violations. When such conflicts arise, they can be dealt with by two means: (1) argument exchanges, using an argument that fits the argumentation scheme for argument from negative consequences, and (2) by exchanging information about the circumstances of the case as new circumstances come to be known by one party. The argument scheme from negative consequences is [8]:

- Major Premise: If A_{dk} is brought about, then consequences C will occur.
- Minor Premise: Consequences C are bad.
- Conclusion: Therefore A_{dk} should not be brought about.

The critical questions include "How strong is the likelihood that the cited consequences will occur?" and "What evidence supports the claim that the cited consequences will occur?". This type of arguments describes what conflicts do




Fig. 5. Overview of agent deliberation process

not permit an action to be performed, or a state of the world to be achieved. We refer to this type of arguments as Arg_{def} . Information about new circumstances can be exchanged by offering support to previously stated claims. These statements explain part of the plans or new obligations. This information asserted by an agent represents a change of the circumstances known by the other agent because it was not part of their common knowledge. In fact, new obligations may become active or new actions may become possible only given the new knowledge about a state of the world that was shared during the dialogue. We refer to these statements as defeasible supporting arguments Arg_{sup} , since circumstances of the agents' plans may change or the knowledge of an agent may contradict that of others. A defeat is defined for an argument that contradicts the premises or the conclusions of another argument.

In our dialogue protocol, we consider two ways to exchange supporting statements. An agent may ask “why?”, (e.g., “Why do you want to perform this action?”) and the other agent counter-argues this by explaining some circumstances, $argue(Arg_{sup})$. More importantly, an agent may take the initiative to exchange new information about circumstances at any point of the dialogue. We identify the intentional act of exchanging this information with $disclose(Arg_{sup})$. Agents are able to modify their plans dynamically and add new pieces of information into their knowledge-base throughout the dialogue. However, it is fundamental for agents to share the information about new circumstances with other agents, in order to establish a more favorable agreement; i.e., identify adequate alternatives to solve conflicts within interdependent plans. We argue that by introducing the $disclose(Arg_{sup})$ act in our dialogue protocol we are able to represent the type of extension proposed in this paper that permits agents to consider changes of circumstances during the course of the deliberation.

Evaluating the effects of an open knowledge-base. In the dialogue, introducing $disclose(Arg_{sup})$ is necessary for agents to find better alternatives for plan conflicts. Here, we assess the effects of this on the identification of an agreement.

Table 1. Speech acts of each protocol $\mathcal{P}$.

Speech Act	Attacks	Surrenders
$propose(A_{dk})$	- $[\mathcal{P}_{ncir}, \mathcal{P}_{cir}]reject(A_{dk})$ - $[\mathcal{P}_{ncir}, \mathcal{P}_{cir}]why(A_{dk})$	$accept(A_{dk})$
$reject(A_{dk})$	- $[\mathcal{P}_{cir}]disclose(Arg_{sup})$ where Arg_{sup} explains A_{dk} - $[\mathcal{P}_{ncir}, \mathcal{P}_{cir}]why(\neg A_{dk})$	$withdraw(A_{dk})$
$withdraw(A_{dk})$		
$accept(A_{dk})$		
$argue(Arg_{def_i})$	- $[\mathcal{P}_{ncir}, \mathcal{P}_{cir}]why(\ell)$ and $\ell \in Arg_{def_i}$ - $[\mathcal{P}_{ncir}, \mathcal{P}_{cir}]argue(Arg_{def_j})$ and Arg_{def_j} defeats Arg_{def_i} - $[\mathcal{P}_{cir}]disclose(Arg_{sup})$ and Arg_{sup} explains $\ell \in Arg_{def_i}$	$withdraw(A_{dk})$ or $accept(A_{dk})$
$why(\ell)$	- $[\mathcal{P}_{ncir}, \mathcal{P}_{cir}]argue(Arg_{sup})$ and Arg_{sup} explains ℓ	$withdraw(A_{dk})$ or $accept(A_{dk})$
$disclose(Arg_{sup})$ $[\mathcal{P}_{cir}]$	- $[\mathcal{P}_{cir}]argue(Arg_{def})$ and Arg_{def} defeats Arg_{sup} - $[\mathcal{P}_{cir}]why(\ell)$ and $\ell \in Arg_{sup}$	$withdraw(A_{dk})$ or $accept(A_{dk})$

In order to show the difference introduced by $disclose(Arg_{sup})$ we define two protocols, $\mathcal{P}_{cir}$ and $\mathcal{P}_{ncir}$, to be used by agents in the argumentation phase. In [6] those are respectively $\mathcal{P}_{sym}$ and $\mathcal{P}_{asym}$. Both protocols use speech acts similar to the MHP model. However, agents may only perform $disclose(Arg_{sup})$ within protocol $\mathcal{P}_{cir}$, which represents a more flexible protocol, while with the use of $\mathcal{P}_{ncir}$ agents are constrained to argue only against some negative consequences in adopting a new action. The protocols are presented in Table 1, each speech act in the right and middle columns may respond to one act in the left column previously performed by the other agent.

Here we report important results from the empirical evaluation of a system employing our model of deliberation. Full details of the implementation of this system can be found in [6]. We intend to show that using $\mathcal{P}_{cir}$ agents are able to identify better agreements. Two agents start with individual interdependent plans and discuss about some dependencies using $\mathcal{P}_{cir}$ and $\mathcal{P}_{ncir}$. Individual plans have an average of 5 objectives and 50 actions. Collaboration between agents is enforced for a 5%-10% of the total number of actions in each plan. We ran experiments on 275 pairs of plans with conflicts related to scheduling or norm violations in a range between 20 to 80. This permits us to study different aspects of robustness of the argumentation system as the complexity of the problem increases. The complexity is given by the total number of conflicts that exist between two initial plans (TOT). Our hypothesis for the experiments is:

Hypothesis 1 *The use of a protocol where agents are able to share information about their view of the circumstances of the decision during the course of deliberation increases the number of successful outcomes between agents.*

Here a successful outcome must be defined. Note that each action A_{dk} proposed for discussion is part of a chain of actions that enables an agent to achieve a goal in its own plan. If no alternatives can be found, the goal must be dropped.




Fig. 6. Percentage of dialogues with successful outcome as total conflicts increases.

Definition 1. A successful outcome of deliberation is achieved if an agreement on the action proposed for discussion A_{dk} was reached or if a valid alternative was accepted without the need for any of the agents to drop their goals.

In order to prove our hypothesis, we counted the number of successful dialogues in the 550 experiment runs for the two protocols. The percentage of successful dialogues on average is significantly higher by about 25% in protocol $\mathcal{P}_{cir}$ (64.51%) than in protocol $\mathcal{P}_{ncir}$ (40.69%) with $p \ll 0.001$. This shows that the dialogue is more likely to be successful with a protocol where agents are able to intentionally exchange information about circumstances. We studied the robustness of successful outcomes of the dialogue when the complexity of the problems increases. Figure 6 presents the results. There is a significant improvement of the success of the outcome in $\mathcal{P}_{cir}$ ($p \ll 0.001$). The successful outcome decreases when the complexity increases, probably caused by the difficulty of finding an agreement when the problems become harder. However, $\mathcal{P}_{cir}$ is shown to be the most robust for increasing problem complexity, with an average of 60% of successful outcomes even in the most challenging problems. The results prove that by permitting agents to disclose information about individual circumstances we obtain a significant improvement on the number of agreements established.

These results showed that the introduction of a flexible knowledge-base where agents can add pieces of information about circumstances while deliberating is fundamental for establishing more successful agreements. Moreover, we show a first attempt to introduce that richness of natural deliberative dialogue that was argued to be missing in dialogue protocols such as the MHP model.

6 Measures for Criteria of Closure

In Section 4 we formulated ten criteria for assessing whether a deliberative dialogue has been successful. In this section, we discuss some measures for external evaluation of success carried out by looking at the record of deliberation.

The criteria proposed are necessary in human deliberation to determine whether enough proposals have been discussed and whether they are particularly useful to determine the educational value of deliberation in situations when an open knowledge base is introduced. However, existing systems for evaluation (e.g., [2]) use dialogue protocols based on the MHP model, thus their methods for evaluation are limited to assessing whether the agreement was found or not. We argue that if agents were to follow a more natural deliberation dialogue our criteria are necessary to assess the educational benefits of protocols for agent deliberation. However, while the criteria function as guideline for evaluation, we must define specific measures to compute these benefits. We propose a measure for evaluating the following criteria within the system presented in Section 5:

- 2) The accuracy and completeness of the information regarding the circumstances of the case;
- 5) How thoroughly each of the proposals that were put forward during the deliberation were engaged by supporting or attacking arguments;
- 6) Whether any arguments that should have been considered were not given due consideration.

In our system, we want to assess the educational role of the argumentation phase in terms of how well agents are able to align their knowledge about the circumstances. The goal of the agents is not only to find an agreement but to identify one that resolves as many conflicts between interdependent plans as possible through dialogue. This establishes how beneficial or successful the deliberative dialogue has been. The measure of a successful outcome in Def.1 is not enough to determine such a result. We define a beneficial outcome as:

Definition 2. *A feasible plan is one that an agent is able to enact without impeding another's goals. A dialogue about two interdependent plans is beneficial when, at the end of the dialogue, the feasibility of the plans has increased. This is determined by an increase of the number of conflicts solved between the plans.*

We now discuss further empirical results of our model to show the benefits of a more flexible protocol, $\mathcal{P}_{cir}$, able to share information about new circumstances, in comparison with a more restrictive protocol $\mathcal{P}_{ncir}$ within a deliberation dialogue. In addition, we introduce a condition $\mathcal{G}_{prior}$, based on protocol $\mathcal{P}_{cir}$ where agents strategically select arguments Arg_{def} to be exchanged. The selection is based upon the idea that agents aim to solve as many conflicts as possible during dialogue. In order to do so, agents prefer arguments that subsume other arguments in terms of their conflicts. For example, an agent would prefer to state that stopping the water supply damages both building and running an hospital, rather than stating that it simply impedes the running of the hospital. We propose the following hypotheses for evaluation of the criteria:

Hypothesis 2 *The use of a protocol where agents can share additional information about circumstances $\mathcal{P}_{cir}$ during dialogue is more effective in conveying information and thus, increases the number of conflicts resolved (Criterion 2).*




Fig. 7. Conflicts solved and Arguments as total number of conflicts increases.

Hypothesis 3 *The use of a protocol where agents can share additional information about circumstances $\mathcal{P}_{cir}$ leads to an increase in the number of arguments exchanged (Criterion 5).*

Hypothesis 4 *The use of a strategy for selecting arguments $\mathcal{G}_{prior}$ is more effective in conveying the appropriate information about conflicts without increasing the number of arguments exchanged (Criterion 6).*

The results are presented in Figure 7. Side A represents the data and linear regressions of the number of conflicts solved ($SOLV$), while side B represents the number of arguments exchanged ($ARGS$) in both cases, when the total number of conflicts increases (TOT). Results are statistically significant at $p \ll 0.001$.

Figure 7-A provides evidence for Hyp.2. The graph shows that the number of conflicts solved is significantly higher when agents employ $\mathcal{P}_{cir}$ than using $\mathcal{P}_{ncir}$. This is a measure for testing the accuracy and completeness of the information regarding the circumstances of the case. We showed that many conflicts may only be solved through the exchange of information about the circumstances.

Our second hypothesis is verified in Figure 7-B. The top line corresponds to the number of arguments moved in $\mathcal{P}_{cir}$ and the bottom line is $\mathcal{P}_{ncir}$. In a protocol where more information about circumstances is exchanged, there is more possibility for agents to discuss their proposal and claims more thoroughly. The results prove the claim that the arguments exchange $ARGS$ in $\mathcal{P}_{cir}$ are significantly higher than with $\mathcal{P}_{ncir}$. We attribute the increase in the number of conflicts solved showed for Hyp.2 to the information shared within the two protocols, since the experiments were performed on the same pairs of plans.

For our third hypothesis, we look at the number of arguments exchanged between $\mathcal{G}_{prior}$ and $\mathcal{P}_{cir}$, and the number of conflicts solved in Figure 7. The results show evidence for the claim that agents selecting arguments strategically solve more conflicts. However, in both conditions $\mathcal{G}_{prior}$ and $\mathcal{P}_{cir}$ the number of arguments exchanged is similar; there were some arguments more important for solving conflicts exchanged in $\mathcal{G}_{prior}$, but they were never moved in $\mathcal{P}_{cir}$. Using these two conditions we can verify criterion 6: we demonstrated that in dialogue $\mathcal{P}_{cir}$, compared to $\mathcal{G}_{prior}$, some arguments were not given enough consideration.

We have presented here an initial simple method for measuring some of the criteria presented for resolving the closure problem. There is the need, however, to develop a more standard method for measuring these criteria in order to determine the educational value of deliberation in agent-based systems. More investigation within computational system is required to establish general conditions under which it is possible to measure these criteria, or to determine whether certain criteria are more useful than others in specific domains.

7 Conclusion

We have explored the problem of whether the current models for agent deliberation capture the richness of human deliberative dialogue. Agent deliberation protocols are based on models of natural deliberation. Black and Atkinson [1], for example, consider agents with different expertise that deliberate over the best action to perform in collaboration. Agents establish what to believe about the surrounding circumstances, and then propose and discuss actions that promote or demote societal values. Similarly, in Kok et al. [4] a team of agents deliberate about the best option to achieve a goal. In the dialogue agents can propose, withdraw, or challenge options. However, little effort has been invested in understanding whether these systems can effectively represent natural deliberation.

In order to address this problem, we discussed the work of McBurney et al. [5], which presents a model of deliberation underpinning many dialogue systems including [1, 4]. We showed that this represents a good model of how deliberation proceeds in real settings, permitting agents to interweave phases of information-seeking and argumentation for practical reasoning. An important characteristic for intelligent rational deliberation is, however, missing: considering the change of circumstances during the course of the dialogue. We proposed an extension to the MHP dialogue in which agents are open to the exchange of new knowledge. We showed that our framework [6, 7] considers, to a certain extent, the situation where agents can willingly share information about new circumstances. We demonstrated that this is necessary to identify more successful outcomes.

New circumstances may initiate an infinite cycle of rethinking and reevaluating proposals. We argued that some criteria must be established to determine when the practical reasoning can be closed off and to declare whether the deliberation has been successful. People engaging in natural deliberation may find dialogue educationally successful even if the debate terminated without an agreement. In the MHP model, however, agents are required to find an agreement to

declare the dialogue successful. How this can be extended so that agent deliberation better reflects human deliberation is a question to which we provided some answers. We proposed ten criteria to determine whether the dialogue has been successful; criteria that we argue should be implemented in agent-based models of deliberation dialogue. Measures for these criteria must also be defined to permit the evaluation of the educational benefits of the dialogue. Here, we demonstrated that, thanks to the sharing of information about new circumstances, agents are able to identify more beneficial agreements.

We proposed two new phases that extend the MHP dialogue to consider dynamic changes of circumstances during dialogue and to address the problem of how to determine success. We believe that our extended model will facilitate the development of applications that are able to represent rich deliberation processes to support human decision-making in a more effective way.

Acknowledgments. This research is partially supported by Social Sciences and Humanities Research Council of Canada Insight Grant 435-2012-0104. This research is also partially supported by the award made by the RCUK Digital Economy program to the dot.rural Digital Economy Hub; award ref.: EP/G066051/1.

References

1. Black, E., Atkinson, K.: Dialogues that account for different perspectives in collaborative argumentation. In: *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems*. pp. 867–874 (2009)
2. Black, E., Bentley, K.: An empirical study of a deliberation dialogue system. In: Modgil, S., Oren, N., Toni, F. (eds.) *Theory and Applications of Formal Argumentation*, *Lecture Notes in Computer Science*, vol. 7132, pp. 132–146. Springer Berlin Heidelberg (2012)
3. Gordon, T.F., Prakken, H., Walton, D.: The carneades model of argument and burden of proof. *Artificial Intelligence* 171(10–15), 875 – 896 (2007)
4. Kok, E.M., Meyer, J.J.C., Prakken, H., Vreeswijk, G.A.W.: A formal argumentation framework for deliberation dialogues. In: McBurney, P., Rahwan, I., Parsons, S. (eds.) *Argumentation in Multi-Agent Systems*, *Lecture Notes in Computer Science*, vol. 6614, pp. 31–48. Springer Berlin Heidelberg (2011)
5. McBurney, P., Hitchcock, D., Parsons, S.: The eightfold way of deliberation dialogue. *International Journal of Intelligent Systems* 22(1), 95–132 (2007)
6. Toniolo, A.: *Models of argument for deliberative dialogue in complex domains*. Ph.D. thesis, University of Aberdeen (2013)
7. Toniolo, A., Norman, T.J., Sycara, K.: An empirical study of argumentation schemes for deliberative dialogue. In: *Proceedings of the Twentieth European Conference on Artificial Intelligence*. vol. 242, pp. 756–761 (2012)
8. Walton, D.: *A pragmatic theory of fallacy*. University of Alabama Press, Tuscaloosa and London (1995)
9. Walton, D.N., Krabbe, E.C.W.: *Commitment in dialogue: Basic concepts of interpersonal reasoning*. State University of New York Press (1995)
10. Walton, D.: An argumentation model of deliberative decision-making. *Technologies for Supporting Reasoning Communities and Collaborative Decision Making: Cooperative Approaches* pp. 1–17 (2011)