A Dialogue Game to Agree to Disagree about Inconsistent Information

Henk-Jan Lebbink
Institute of Information and
Computing Sciences
Utrecht University
`henkjan@cs.uu.nl`

Cilia Witteman
Diagnostic Decision Making
Faculty of Social Sciences
University of Nijmegen
`c.witteman@ped.kun.nl`

John-Jules Meyer
Institute of Information and
Computing Sciences
Utrecht University
`jj@cs.uu.nl`

Abstract

This paper proposes a dialogue game in which coherent conversational sequences at the speech act level are described of agents that become aware they have a disagreement and settle the dispute by agreeing to disagree when they believe insufficient propositions to resolve the situation. A dialogue game is formulated in which agents can offer information possibly resulting in non-reconcilable, mutually inconsistent belief states. These states are handled by a dialogue rule that defines when to 'agree to disagree'.

1 Introduction

If we are to understand conversations we may need to carefully model the underlying principles that drive them, but we would probably be just as satisfied if we could build computational models that generate useful conversations. In general conversations, participants may have autonomy over their cognitive states but they may also have intentions to change those of others. This autonomy and intention may result in discussions about non-reconcilable beliefs. How to cope with these disputes and how to devise a computational model that identifies them?

A dialogue game by Beun (2001) describes speech acts between agents and identifies three structures with accompanying properties that form a dialogue game that agents need to play to communicate in a sensible way. Agents need to have a cognitive state, dialogue rules to generate speech acts to convey information to other agents and update rules to process incoming information. In Lebbink et al. (2003), a multi-valued logic is introduced to describe inconsistent and biased information in dialogue games without forcing agents to perform belief revision. In the same vein as the FIPA work on agent communication languages (Labrou, 2001), we formulate semantics of the speech acts to offer information and to agree to disagree.

What is lacking in Beun (2001) and Lebbink et al. (2003), is the possibility for agents to recognize irresolvable disagreements and based on this recognition utter an agreement to disagree making the disagreement common belief. This common belief can motivate dialogues on the definition of the terms used in the disagreement or dialogues to weaken arguments in order to retract propositions to resolve the disagreement.

Agents may be motivated to persuade others to accept to believe certain propositions. Consequently, agents may accept to believe proposed propositions, but they may also deny to accept to believe propositions when accepting them would result in inconsistent beliefs. The objective is to present a dialogue game in which cognitive agents become aware of non-reconcilable beliefs and manifest this awareness to others while preserving their private beliefs. We use a formal method to represent classical and inconsistent information in the agents' cognitive states, enabling us to devise dialogue rules that can cope with these types of information. The resulting formalism allows for embedded dialogues and verification of existing dialogues, but notably it is computational which makes it possible to generate all valid dialogues in a dialogue game.

In the following (fictitious) dialogue, Tv (from Sesame Street) tries to insure his new car; he rings

an insurance company and explains his desire to an insurance agent (Ia). The Ia wants to sell Tv an expensive insurance policy because Sesame Street puppets are prone to fast and dangerous driving. The Ia wants Tv to accept that its car is not safe justifying an expensive policy. For the sake of argument, agents do not accept information that renders their beliefs inconsistent. This dialogue is mathematically checked as a valid sequence of utterances of our dialogue game. We omit proofs and conclude that this dialogue is one of the 177 different valid dialogues.

Dialogue 1 (Car insurance in Sesame Street)

1. *Tv to Ia*   ‘My car is a Ferrari.’
2. *Ia to Tv*   ‘Ok.’
3. *Tv to Ia*   ‘(and) My car is safe.’
4. *Ia to Tv*   ‘I don’t believe that.’
5. *Ia to Tv*   ‘(actually) I think your car is not safe.’
6. *Tv to Ia*   ‘I don’t accept that my car is not safe.’
7. *Ia to Tv*   ‘Do you accept that if a car is a Ferrari then it is not a safe car?’
8. *Tv to Ia*   ‘(no) I don’t want to accept that.’
9. *Ia to Tv*   ‘Lets agree to disagree whether your car is safe or not.’
10. *Tv to Ia* ‘Ok.’

2 A Dialogue Game to Offer Information

In accordance with Wittgenstein’s dictum “meaning is use”, we propose a dialogue game that gives meaning to the conception of offering information. To understand a word or sentence is to know how to use it. And to be able to use a word or a sentence is to be capable of recognizing the contexts in which it is appropriate to utter it (Ellenbogen, 2003). This activity of speaking is described in a normative way, governed by dialogue rules that dictate correct and incorrect use of communicative acts. We could say that an agent understands a word when it can distinguish between correct and incorrect uses.

Following the approach taken in Beun (2001) and Lebbink et al. (2003), a dialogue game is a set of dialogue rules that describe which communicative acts an agent may utter given its current cognitive state. A dialogue game also has a set of update rules that describe the changes of the cognitive state given an uttered communicative act.

We assume that information can only accumulate in the participants’ cognitive states and cannot be retracted. In this information-monotonic approach additions may introduce inconsistent beliefs. Although the dialogue rules we are to present will prohibit inconsistencies, agents use the possibility of inconsistencies in a look-ahead fashion when deciding to believe propositions. Agents are considered omnipotent and use equal consequence relations. In addition, agents can only speak to one agent at a time via an ideal half-duplex communication channel which means that no information is lost and that information can only flow in one direction at a time. No restrictions are made on the number of participants.

2.1 Ordering of Information: Bilattices

Whereas in classical logic terms are assigned a truth-value true or false, in multi-valued logic (MVL) new truth-values are introduced to represent uncertain, non-determined or other epistemic attitudes (Rescher, 1969). In previous work (Lebbink et al., 2003), truth-values from a bilattice structure (Ginsberg, 1988; Fitting, 1991) are used to define a MVL, and theories of this MVL are used to represent the agent’s cognitive state. Theories of MVL are sets of multi-valued propositions which are terms taken from some ontology with an assigned truth-value from a bilattice structure. Next to these propositions, an implication operation for four truth-values $\mathsf{imp_4}$ is added with a reading similar to the one from classical logic: if the antecedent of the implication is part of a theory, then the consequent is also present.

Two terms are used to denote the information of the example dialogue: $\mathsf{is\_a(this\_car, ferrari)}$ and $\mathsf{is\_a(this\_car, safe)}$ stating that it is true that this car denoted by $\mathsf{this\_car}$ is a Ferrari and a safe one respectively. The multi-valued proposition $\mathsf{is\_a(this\_car, ferrari){:}f}$ is read as ‘term $\mathsf{is\_a(this\_car, ferrari)}$ has truth-value f’. If this proposition is part of a theory T that represents the beliefs of some agent, we say that the agent believes that ‘this car is not a Ferrari’. An implication between the fact that if some car is a Ferrari then that car is not safe, is represented by the proposition $\mathsf{imp_4(is\_a(X, ferrari){:}t, is\_a(X, safe){:}f){:}t}$.

A bilattice is an algebraic structure that formal-

Figure 1: Bilattice representing partial and inconsistent information.

izes an intuitive space of generalized truth-values with two lattice orderings. In Figure 1 the bilattice for a four-valued logic (Belnap Jr., 1977) is graphically depicted; t and f stand for the classical truth-values true and false respectively; non-orthodox truth-value u and i represent a complete lack of information (unknown) and the inconsistent information state. Truth-values are ordered by the amount of truth $\leq_t$ and the amount of information $\leq_k$; only the latter is of interest to us. For instance, u has less information than t and f, denoted by u $\leq_k$ t and u $\leq_k$ f, but t and f are unrelated to one another in the k-order, that is, t $\not\leq_k$ f and f $\not\leq_k$ t. Bilattices with more truth-values and even a continuum of truth-values can be used to represent biased information or probabilities (Ginsberg, 1988); we use only the four from Figure 1.

2.2 Communicative Acts

The communicative act of offering information and its two corresponding answers to accept and reject information are defined next. The act of an *offer* $[x, y, p]^!$ is uttered by a speaker (x) directed to a listener (y) and is read as *'Are you (y) willing to accept to believe proposition p?'* In the first line of the example dialogue, Tv states that its car is a Ferrari, which we consider equal to the phrase *'Are you, Ia, willing to accept to believe that it is true that my car is a Ferrari?'*. In answer to this, in the second line, the Ia grants Tv's offer. The act of *granting an offer* $[x, y, p]^{!+}$ is read as *'I (x) am willing to accept to believe p.'* An agent may also deny an offer, which is done in line four when the Ia denies to believe that Tv's car is safe. The act of *denying an offer* $[x, y, p]^{!-}$ is read as *'I (x) am not willing to accept to believe p.'* The act of agreeing to disagree $[x, y, (p, q)]^\Delta$ is read as *'Are you (y) willing to agree that we are in disagreement over proposition p and q.'* Precise contexts for correct use of these acts are defined in Section 2.6.

A rendition of the dialogue from Section 1 is used in the remainder of this paper.

Dialogue 2 (Car insurance in Sesame Street)

1. $[\text{tv}, \text{ia}, \text{is\_a}(\text{this\_car}, \text{ferrari}){:}\text{t}]^!$
2. $[\text{ia}, \text{tv}, \text{is\_a}(\text{this\_car}, \text{ferrari}){:}\text{t}]^{!+}$
3. $[\text{tv}, \text{ia}, \text{is\_a}(\text{this\_car}, \text{safe}){:}\text{t}]^!$
4. $[\text{ia}, \text{tv}, \text{is\_a}(\text{this\_car}, \text{safe}){:}\text{t}]^{!-}$
5. $[\text{ia}, \text{tv}, \text{is\_a}(\text{this\_car}, \text{safe}){:}\text{f}]^!$
6. $[\text{tv}, \text{ia}, \text{is\_a}(\text{this\_car}, \text{safe}){:}\text{f}]^{!-}$
7. $[\text{ia}, \text{tv}, \text{imp}(\text{is\_a}(\text{X}, \text{ferrari}){:}\text{t}, \text{is\_a}(\text{X}, \text{safe}){:}\text{f}){:}\text{t}]^!$
8. $[\text{tv}, \text{ia}, \text{imp}(\text{is\_a}(\text{X}, \text{ferrari}){:}\text{t}, \text{is\_a}(\text{X}, \text{safe}){:}\text{f}){:}\text{t}]^{!-}$
9. $[\text{ia}, \text{tv}, (\text{is\_a}(\text{this\_car}, \text{safe}){:}\text{f}, \text{is\_a}(\text{this\_car}, \text{safe}){:}\text{t})]^\Delta$
10. $[\text{tv}, \text{ia}, (\text{is\_a}(\text{this\_car}, \text{safe}){:}\text{f}, \text{is\_a}(\text{this\_car}, \text{safe}){:}\text{t})]^\Delta$

2.3 The Agent's Cognitive State

An agent's cognitive state consists of a number of mental constructs which are theories of MVL. An agent's belief state is probably the most important construct next to its desire state. A proposition p is said to be believed by an agent x if p is part of mental construct B_x, that is, $p \in B_x$. Analogously, an agent x desires that an agent y believes a proposition p if p is part of mental construct $D_x^{B_y}$, that is, $p \in D_x^{B_y}$. For example, Tv desires that the Ia believes that it is true that this car is safe, that is, $\text{is\_a}(\text{this\_car}, \text{safe}){:}\text{t} \in D_{\text{tv}}^{B_{\text{ia}}}$.

Agents keep record of all other agents' explicitly communicated beliefs, desires and accompanying consequences. The mental construct for manifested beliefs $M_x B_y$ represents agent y's beliefs that agent x is aware of. For instance, $\psi{:}\text{t} \in M_x B_y$ states that x is aware that y believes that ψ has at least truth-value t. An agent also records other agents' communicated desires, this is done in $M_x D_y^{B_z}$. For instance, Tv is aware that the Ia desires that Tv believes that this car is not safe, that is, $\text{is\_a}(\text{this\_car}, \text{safe}){:}\text{f} \in M_{\text{tv}} D_{\text{ia}}^{B_{\text{tv}}}$. Also, agents need to keep record of explicitly stated ignorance of other agents; the third type of mental construct is the manifested ignorance state. $p \in B_x I_y$ states that agent x is aware that agent y is ignorant towards p.

In addition, higher-order manifested mental constructs are needed for agents to remember to whom they stated their desires and beliefs. This information is needed to prevent them from uttering offers more than once; this is addressed by dialogue rules (Section 2.6). Mind-bending mental

constructs are needed to encode this information; the construct $M_x M_y B_x$ states that x is aware of y's awareness that x believes p. From a usage perspective, if an agent has answered an offer regarding some proposition then this proposition is part of this construct; this is addressed by update rules (Section 2.7). Likewise, if an agent has proposed an agreement to disagree it is not allowed to utter the same agreement again. Therefore, a record needs to be kept of these agreements. For instance, $M_{\mathsf{tv}} M_{\mathsf{ia}} \Delta(\mathsf{tv}, \mathsf{ia})$ states that Tv is aware of the Ia's awareness of their disagreement, this situation is described in Section 3.2.

A dialogue game defines a space of different dialogues that unfold by applying dialogue rules and update rules. Given a (initial) cognitive state of all agents participating in the dialogue (hereafter collective state), all sequences of communicative acts can be generated that are considered valid in the game. For our example dialogue the following initial collective state is used. The Ia has the desire that Tv believes its car is not safe, and Tv has the desire that the Ia believes that its car is a Ferrari and safe. The Ia believes that if a car is a Ferrari then the car is not safe, and Tv believes that its car is a Ferrari and a safe one. Formally,
is_a(this_car, safe):f $\in D_{\mathsf{ia}}^{B_{\mathsf{tv}}}$,
is_a(this_car, ferrari):t $\in D_{\mathsf{tv}}^{B_{\mathsf{ia}}} \cap B_{\mathsf{tv}}$,
imp$_4$(is_a(X, ferrari):t, is_a(X, safe):f):t $\in B_{\mathsf{ia}}$, and
is_a(this_car, safe):t $\in D_{\mathsf{tv}}^{B_{\mathsf{ia}}} \cap B_{\mathsf{tv}}$.

2.4 Cognitive Processes

In our formalism, agents can perform two cognitive processes: deciding to believe offered information and deducing consequences of newly accepted beliefs. Other central concepts, such as choice between permissible communicative acts, or which strategy to use to persuade others, will not be included in the descriptions presented here.

From a mathematical perspective, an agent is persuaded to believe a proposition when its cognitive state changes from not believing the proposition to believing it (Walton and Krabbe, 1995); the proposition is set-theoretically added to the belief state. Agents can have different criteria for accepting to believe something. For simplicity, we assume agents to be very credulous: they accept to believe offered propositions that are consistent with their current beliefs.

Reasoning is restricted to the agents' capacity to draw conclusions based on believed implication rules and antecedents. If an implication rule is believed by the Ia, imp$_4$(is_a(X, ferrari):t, is_a(X, safe):f):t $\in B_x$ and the Ia is persuaded to believe an antecedent that the car is a Ferrari, is_a(this_car, ferrari):t $\in B_{\mathsf{ia}}$, then it also concludes to believe the consequent that the car is not safe, is_a(this_car, safe):f $\in B_{\mathsf{ia}}$. However, if the Ia already believes is_a(this_car, safe):t, then its belief state becomes inconsistent, that is, is_a(this_car, safe):i $\in B_{\mathsf{ia}}$. The closure $cl(B_{\mathsf{ia}} \cup \{p\})$ corresponds with the set of propositions including Ia's beliefs plus p with consequents.

2.5 Motivations to Communicate

In Beun (2001) an agent's motivation to utter a question is to balance its belief and desire state. Stated in our terminology, an agent may pose a question regarding some proposition if it has the desire to be in a belief state in which it believes the proposition and it currently does not (yet) believe the proposition. We give a similar motivation to offer information. An agent x may offer information to an agent y regarding some proposition p if x has the desire that y is to believe the proposition, and x is not aware that y already believes the proposition. Stated differently, an agent's motivation to utter an offer is to balance its desire and manifested belief state; an agent x is motivated to offer p to y if $p \in D_x^{B_y}$ and $p \notin M_x B_y$.

According to the Gricean maxims of cooperation, speakers are forbidden to ask anything they already believe (Grice, 1975). Analogously, speakers are forbidden to put forward information if they are aware the listener already believes the information. Next to giving restrictions, these maxims also provide motivations for granting and denying questions and offers: both should always be answered. In the next paragraph, the motivations and restrictions are combined to form the preconditions for 'correct' usage of our communicative acts; these preconditions provide the semantics of the acts and give communicative acts meaning in the context of a dialogue game.

2.6 Dialogue Rules

In group decisions, different experts make decisions as a group by agreeing on the assumptions they need to make to come to one common decision. This means that the assumptions (which are beliefs) should be non-conflicting. A motivation for someone who facilitates group decision making is to introduce dialogue rules that enable experts to offer information with the objective that experts become aware of other agent's attitudes towards their beliefs. One way to check whether assumptions are conflicting is to offer these to others and conclude from their responses whether they agree to believe these. In this section, the dialogue rules are defined allowing an agent to utter communicative acts given its cognitive state.

An offer $[x, y, p]^!$ is defined applicable when the speaker (x) is motivated to utter an offer. As stated before, the speaker desires the listener (y) to believe proposition p and the speaker is not aware that the listener already believes p, that is, $p \in D_x^{B_y}$ and $p \notin M_x B_y$. Of course, a dialogue game can be conceived in which meaning is given to offering information even when the speaker is aware the listener already believes this. Such a different game is played when the speaker wants to convey that it was *not* aware that the listener believed the proposition, although it was. Furthermore, it is assumed that in this dialogue game, agents are not allowed to propose information that they do not believe themselves. Due to the ideal communication channel, agents are also not allowed to pose a communicative act more than once: after uttering an offer, they are aware that the listener is aware of the speaker's desire to induce a belief change. This is represented by the mental construct that the speaker is aware that the listener is aware of the speaker's desire to induce a belief change associated with p in the listener. Also, proposition p should not be part of this mental construct, that is, $p \notin M_x M_y D_x^{B_y}$ which can only be true if speaker has already offered p.

Definition 1 (offer) *If $p \in D_x^{B_y}$, $p \notin M_x B_y$, $p \in B_x$ and $p \notin M_x M_y D_x^{B_y}$ then an offer $[x, y, p]^!$ is applicable.*

The communicative act of granting an offer $[x, y, p]^{!+}$ is applicable when the (granting) speaker (x) believes the proposition p. Whether agent x is persuaded to believe the proposition as a result of the offer x is responding to, does not matter. Obviously, a speaker may only grant an offer if it is aware the listener has the desire to make the speaker believe p, that is, $p \in M_x D_y^{B_x}$. In this dialogue game, an agent x can only be aware of this if the other agent y has uttered an offer. To ensure that the information represented by granting an offer is not superfluous, e.g. stated more than once, the speaker may not be aware that the listener is aware that the speaker believes the proposition it granted, that is, $p \notin M_x M_y B_x$.

Definition 2 (granting an offer) *If $p \in B_x$, $p \in M_x D_y^{B_x}$ and $p \notin M_x M_y B_x$ then granting an offer $[x, y, p]^{!+}$ is applicable.*

An offer of information can be answered by granting or denying to accept to believe the proposition. The act of denying an offer $[x, y, p]^{!-}$ is similar to the act of granting with the difference that the speaker had not been persuaded to believe p, that is, $p \notin B_x$. Equal to the act of granting an offer, the speaker must be aware that the listener has the desire to induce a cognitive state change in the speaker, that is, $p \in M_x D_y^{B_x}$. To prevent that the denial is not superfluous, the speaker may not be aware that the listener is already aware that it has explicitly stated that it does not believe the proposition, that is, $p \notin M_x M_y I_x$.

Definition 3 (denying an offer) *If $p \notin B_x$, $p \in M_x D_y^{B_x}$ and $p \notin M_x M_y I_x$ then denying an offer $[x, y, p]^{!-}$ is applicable.*

A follow-up offer is an offer that substantiates some claim to believe another proposition. This offer is syntactically indistinguishable from the offer defined in Definition 1. However, the follow-up offer is a different speech act from a semantic perspective: it has the following preconditions. The speaker has the desire that the listener believes some proposition p, but the listener does (not yet) believe p and the speaker has already offered him p. The speaker may utter a follow-up offer regarding proposition q if q added to the listeners belief state would make him accept to believe p. Formally, if q is added set theoretically to $M_x B_y$, then p becomes part of the manifested belief state, that is, $p \in cl(M_x B_y \cup \{q\})$. The proposition q may in

this case be offered when the listener does not believe it and the speaker has not proposed it before.

Definition 4 (follow-up offer) *If* $p \in D_x^{B_y}$*,* $p \notin M_x B_y$*,* $p \in M_x M_y D_x^{B_y}$*,* $p \in cl(M_x B_y \cup \{q\})$*,* $q \notin M_x B_y$ *and* $q \notin M_x M_y D_x^{B_y}$ *then offer* $[x, y, q]^!$ *is applicable.*

2.7 Update Rules

Update rules define the agent's change of cognitive state given a communicative act directed at the agent.

After a speaker (x) has offered proposition p to a listener (y), that is, after $[x, y, p]^!$, the following properties for the cognitive states hold. The listener is aware that the speaker desires the listener to believe proposition p, that is, $p \in M_y D_x^{B_y}$, and the speaker is aware that the listener is aware of this, that is, $p \in M_x M_y D_x^{B_y}$. In addition, after an offer the listener is aware that the speaker believes the proposition, that is, $p \in M_y B_x$, and the speaker is aware that the listener is aware of this, $p \in M_x M_y B_x$. Offering a proposition may have the effect that the listener is persuaded to believe the proposition ($p \in B_y$). Note that being persuaded is not encoded in update rules but in the agent's cognitive processes (Section 2.4).

Definition 5 (offer) *after the update of an offer* $[x, y, p]^!$ *holds that* $p \in M_y D_x^{B_y}$*,* $p \in M_x M_y D_x^{B_y}$*,* $p \in M_y B_x$ *and* $p \in M_x M_y B_x$ *.*

After a speaker (x) has granted an offer regarding proposition p to a listener (y), that is $[x, y, p]^{!+}$, the following properties for the cognitive state hold. The listener is aware the speaker believes proposition p, that is, $p \in M_y B_x$, and the speaker is aware that the listener is aware of this, that is, $p \in M_x M_y B_x$. Remember, this mental construct is used to represent that the speaker has granted the offer.

Definition 6 (granting an offer) *after the update of granting an offer* $[x, y, p]^{!+}$ *holds that* $p \in M_y B_x$ *and* $p \in M_x M_y B_x$*.*

After a speaker (x) has denied an offer regarding proposition p to a listener (y), that is $[x, y, p]^{!-}$, the following properties for the cognitive state hold. Similar to granting an offer, after denying

an offer, the listener is aware that the speaker does not believe the proposition, that is, it is ignorant towards it, $p \in M_y I_x$. Also, the speaker is aware of this, which is represented by $p \in M_x M_y I_x$.

Definition 7 (denying an offer) *after the update of denying an offer* $[x, y, p]^{!-}$ *holds that* $p \in M_y I_x$ *and* $p \in M_x M_y I_x$*.*

3 Agree to Disagree

If a group of experts are unable to agree on a decision when for example two experts disagree on some propositions that are needed to agree, a persuasion dialogue may resolve the disagreement by adding information to the expert's belief state (Walton and Krabbe, 1995). If all methods to persuade have become exhausted, the experts can conclude that they disagree on a specific subject and that they both agree on this. This agreement may trigger the meta dialogue in which for example a coin flipping method is proposed to resolve the problem. In this section, a dialogue rule for agreeing to disagree is proposed, making the disagreement common belief.

3.1 Disagreement

Two pieces of information are in disagreement when they are not subsumed under each other in the information order. Stated differently, two pieces of information disagree when the truth-values representing the information are not related with respect to the information order $\leq_k$. For example, t disagrees with f because t $\not\leq_k$ f and f $\not\leq_k$ t, but u agrees with t because u $\leq_k$ t. An agent believing a proposition ψ : u and another agent believing ψ : t do not disagree about ψ, the latter agent is just more informed than the former naive agent. Equally, t agrees with i because t $\leq_k$ i, see also Figure 1.

A *disagreement* between two agents x and y about the truth-value of term ψ exists if (and only if) x believes a proposition $\psi : \theta_1$ and y believes proposition $\psi : \theta_2$, and the truth-values disagree. In line 1 of the example dialogue, Tv states that the car is a Ferrari. After the update of cognitive states, the Ia believes this and he concludes that this car is not a safe car, but he is not yet aware that Tv believes that this car is safe.

3.2 Awareness of Disagreements

An agent (x) is aware of a disagreement with another agent (y) if and only if x believes a proposition $\psi{:}\theta_1$ and x is aware that y believes proposition $\psi{:}\theta_2$ and θ_1 disagrees with θ_2. In line 3 of the example dialogue Tv states that its car is safe. After the cognitive state update of this act the Ia is aware that Tv believes that the car is safe; the Ia is now aware of a disagreement because it believes that the car is safe. After line 4, Tv is also aware of the disagreement.

Under the assumption that the dialogue only results in additions of the agents' cognitive states, it can be proven that disagreement awareness always implies the existence of a disagreement and that it is not possible that agents are incorrectly aware of a disagreement. Note that if agents have a disagreement, they need not be aware of this.

A second-order disagreement awareness exists when an agent (x) is aware that another agent (y) believes a proposition $\psi{:}\theta_1$ and x is aware that y is aware that x believes proposition $\psi{:}\theta_2$ and θ_1 disagrees with θ_2. In line 5 the Ia states that it believes that the car is not safe, after the update, the Ia is aware of a second order disagreement. After line 6, Tv is also aware of this disagreement.

3.3 Resolving Disagreements

The minimal piece of information that is needed to resolve a disagreement between two agents about a term ψ is represented by proposition $\psi{:}\xi$ that if added to one of the agent's belief state resolves the disagreement. Remember that in the current dialogue game only additions of information are possible and consequently, resolving disagreement can only take place by adding sufficient information to one of the two agent's, rendering it possibly inconsistent.

3.4 Dialogue Rule to Agree to Disagree

The situation in which participants may utter an agreement to disagree can now be equated in a dialogue rule. The speaker (x) may propose to agree to disagree about term ψ to the listener (y) if:

1. The speaker is aware that it has a disagreement about term ψ with the listener, that is, $\psi{:}\theta_1 \in B_x$, $\psi{:}\theta_2 \in M_x B_y$ and θ_1 disagrees with θ_2.

2. The speaker is aware that the listener is also aware of the disagreement, that is, $\psi : \theta_3 \in M_x M_y B_x$ and θ_3 disagrees with θ_2.

3. The speaker is *not* aware of a set of propositions that it has not offered to the listener before that could have resolved the disagreement if the listener had accepted to believe them. Suppose the proposition $\psi{:}\xi_1$ represents the minimal amount of information that if added to the listener's belief state resolves the disagreement. If for a set of propositions Φ that is believed by the speaker ($\Phi \subseteq B_x$) holds that if Φ were added to the listener's belief state then the disagreement would have been resolved, that is, $\psi{:}\xi_1 \in cl(M_x B_y \cup \Phi)$. Furthermore, if a set of proposition $\Phi \subseteq B_x$ has been offered to the listener then holds that $\Phi \subseteq M_x M_y D_x^{B_y}$. We can now state that the speaker has no methods (sets of propositions Φ) left to persuade the listener by $(\forall \Phi \subseteq B_x)(\psi{:}\xi_1 \in cl(M_x B_y \cup \Phi) \Rightarrow \Phi \subseteq M_x M_y D_x^{B_y})$.

4. According to the speaker, the listener is *not* aware of a set of proposition not offered to the speaker before that can resolve the disagreement if the speaker accepts to believe them. That is, for all sets of propositions Ψ believed by the listener according to the speaker ($\Psi \subseteq M_x B_y$) holds that if these propositions Ψ were believed by the speaker then the speaker would have believed $\psi{:}\xi_2$ representing the minimal information that is needed to resolve the disagreement (according to the listeners that the speaker is aware of), that is, $\psi{:}\xi_2 \in cl(M_x M_y B_x \cup \Psi)$. If all these propositions Ψ have been offered and apparently this did not resolve the disagreement, that is $\Psi \subseteq M_x M_y I_x$, the speaker is aware that the listener has no methods left to resolve the situation. Formally, $(\forall \Psi \subseteq M_x B_y)(\psi : \xi_2 \in cl(M_x M_y B_x \cup \Psi) \Rightarrow \Psi \subseteq M_x M_y I_x)$.

5. The speaker is not aware that it proposed to agree to disagree regarding the disagreement before. Agreements to disagree are kept record of similar to manifested beliefs and desires. $M_x M_y \Delta(x,y)$ states that x is aware of the agreement to disagree by x between the agents.

If these five preconditions hold, the speaker may propose an agreement to disagree regarding ψ, denoted $[x, y, (\psi{:}\theta_3, \psi{:}\theta_2)]^\Delta$.

3.5 Generation of the Example Dialogue

The proposed dialogue game gives preconditions to offer information. Combined with the formal rule of the speech acts, the update rules and the (not presented) axioms of theories of MVL, the example dialogue can be proven to follow from the initial collective state. Note that the agreement to disagree is based on information regarding the two agents involved in the disagreement. This disagreement may need to be retracted when new information is gained on how to resolve the situation. This is because the agreement is based on the absence of beliefs of only the two agents that actually have the disagreement. If another agent offers a proposition that resolves the disagreement, the agreement to disagree needs to be retracted.

With the help of software tools, the complete state space of the dialogue game with its initial collective state is generated in one tenth of a second, resulting in a graph with 37 nodes representing the collective states and with 66 edges representing speech act utterances with associated cognitive state updates. This network comprises 177 different dialogues with three different final collective states. One has to remember that an agent accepts to believe a proposition if it is consistent with its current belief state. This makes offering information of crucial importance, resulting in the three different endings.

4 Conclusions

We have given a formal semantics to the act of uttering a proposal to agree to disagree; these semantics are defined by formulating the rules of usage in the context of a computational dialogue game for offering information. We have shown that semantics of communicative acts can be given with a dialogue game in an intuitive manner, and that given the dialogue game a formal system emerges in which sequences of communicative acts can be checked valid dialogues. Also dialogues can be generated from the dialogues and update rules providing the possibility to analyse dialogue games on useful properties like termination or whether the unbalanced desire/belief states are resolved in the terminating collective states.

If two agents utter their agreement to disagree they are mutually aware of the disagreement they have about some proposition, this awareness is vital in group decisions because these decisions should not be based on information that agents disagree upon. Once agents agree to disagree about a proposition it cannot be used in future reasoning, even previous decisions in which this proposition played a role may be compromised.

Future research addresses agents that strategically select which speech acts to utter with the intention to arrive at a collective state in which desirable properties hold. Other research centres around speech acts for retracting information. Retracting information is an act of offering *not* to believe a proposition, that is, an offering to forget. Retractions of information enhance the current dialogue game by to retracting agreements to disagree when new information comes to light.

References

[Belnap Jr.1977] N.D. Belnap Jr. 1977. a useful four-valued logic. In J.M. Dunn and G. Epstein, editors, *Modern Uses of Multiple-Valued Logic*, pages 8–37.

[Beun2001] R.J. Beun. 2001. On the Generation of Coherent Dialogue: A Computational Approach. *Pragmatics & Cognition*, 9(1):37–68.

[Ellenbogen2003] S. Ellenbogen. 2003. *Wittgenstein's Account of Truth*. SUNY series in philosophy. State University of New York Press.

[Fitting1991] M. Fitting. 1991. Bilattices and the semantics of logic programming. *J. of Logic Programming*, 11:91–116.

[Ginsberg1988] M.L. Ginsberg. 1988. Multivalued Logics: A Uniform Approach to Reasoning in Artificial Intelligence. *Comput. Intelligence*, 4:265–316.

[Grice1975] H.P. Grice. 1975. Logic and conversation. In P. Cole and J.L. Morgan, editors, *Speech Acts*, volume 11 of *Syntax and Semantics*, pages 41–58.

[Labrou2001] Y. Labrou. 2001. Standardizing Agent Communication. *LNCS*, 2086:74–98.

[Lebbink et al.2003] H.-J. Lebbink, C.L.M. Witteman, and J.-J.Ch. Meyer. 2003. Dialogue Games for Inconsistent and Biased Information. *presented at LCMAS03 (www.win.tue.nl/ẽvink/lcmas03.html)*.

[Rescher1969] N. Rescher. 1969. *Many-valued Logic*. McGraw-Hill.

[Walton and Krabbe1995] D.N. Walton and E.C.W. Krabbe. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press. Albany, NY, USA.