

A General Framework for Sound Assumption-based Argumentation Dialogues

Xiuyi Fan, Francesca Toni

Imperial College London, United Kingdom

Abstract

We propose a formal model for argumentation-based dialogues between agents, using assumption-based argumentation (ABA) as the underlying argumentation framework. Thus, the dialogues amount to conducting an argumentation process in ABA. The model is given in terms of ABA-specific utterances, debate trees and forests implicitly built during and drawn from dialogues, legal-move functions (amounting to protocols) and outcome functions. The model is generic in that it is not restricted to any specific dialogue types and can be used to support a wide range thereof. We prove a formal connection between dialogues and three well-known argumentation semantics (i.e. grounded, admissible and ideal extensions), by giving soundness results for our dialogue models with respect to these semantics. Thus, our dialogues can be seen as a distributed mechanism for successfully determining acceptability of claims (with respect to the semantics considered), while constructing argumentation frameworks and arguments for these claims.

Keywords: Argumentation, Dialogue, Agents, Semantics

1. Introduction

Argumentation-based dialogue systems have attracted considerable research interest in recent years (e.g. see [33, 27]), largely due to the need for agents to communicate and agree in multi-agent systems. Indeed, argumentation is a powerful reasoning abstraction where conflicting positions or opinions are evaluated against one another in order to resolve conflicts. Argumentation has been used quite extensively in AI in the last two decades to support a number of applications and address a number of problems (e.g. see [4, 5, 38] for an overview). To support this line of research, several argumentation frameworks have been proposed through the years, including [11, 7] and many more (see for example [1]

for a recent overview of some approaches). The modern study of formal dialogue systems for argumentation can be deemed to have started with Charles Hamblin’s work [24]. The topic was initially studied within philosophical logic and argumentation theory [26, 43]. Subsequently, researchers from the field of AI & law [23, 31] and multi-agent systems [3, 30] have looked into dialogue systems as well.

This paper presents a two-agent argumentation-based dialogue model, using Assumption-based Argumentation (ABA) [7, 12, 40, 42] for the representation of arguments and attacks, and for determining “success” of dialogues. ABA is well-suited as a foundation for argumentation-based dialogues for a number of reasons. It is a general purpose argumentation framework with several applications (e.g. see [12, 40, 42]) including applications requiring dialogues between agents [20]. It is a structured argumentation frameworks, so that a dialogue model based on it can allow the collective construction of arguments and attacks, and a distributed evaluation of “success”, rather than forcing, as when using abstract argumentation [11], for example as in [36], that arguments and attacks are determined and/or constructed individually by agents or collectively but prior to dialogues. At the same time, it is an instance of abstract argumentation [15, 40] and it admits abstract argumentation as an instance [40], thus allowing our dialogue model to accommodate, as a special case, the communication and evaluation of abstract arguments as well. There are several other structured argumentation frameworks available as a basis for argumentation-based dialogues, notably logic-based argumentation [5], DeLP [22] and, more recently, ASPIC+ [34] (see [1] for a recent survey of all these structured argumentation frameworks). Of these, only ASPIC+ is an instance of abstract argumentation and allows abstract argumentation dialogues to be generated, if required by applications. ASPIC+ is a generalisation of ABA and has been designed to admit ABA as an instance [34]. Thus, our ABA-based dialogue model can also be seen as a dialogue model based on ASPIC+, and extending the functionalities and properties of dialogue models based on precursors of ASPIC+ (e.g. [32, 33], as we will discuss later). Another essential feature of ABA, for the purposes of this article, is that it is equipped with provably correct computational mechanisms with respect to several semantics [14, 15, 41]. We rely upon aspects of these mechanisms, as well as their soundness, in order to prove our formal soundness results.

An ABA framework consists of rules, assumptions, and contraries, specified in a logical language. Informally, rules and assumptions form deductions (arguments); contraries of assumptions provide means of specifying counter-arguments (attacks) against arguments. Within an ABA framework, sets of arguments are

deemed “acceptable” if they fulfil certain properties, e.g., under the semantics of admissible extensions [7, 12], a set of arguments does not attack itself and attacks all arguments that attack it. Then, claims are deemed “acceptable” if they are supported by (are conclusions of) arguments that belong to “acceptable” sets.

Our dialogue model makes use of the same building blocks as ABA, in that a dialogue is composed of utterances whose *content* may be a rule, an assumption, a contrary, or a claim whose “acceptability” needs to be ascertained. In addition, the content of utterances may be a pass, amounting to the agent contributing no information to the dialogue at the time of the utterance. Dialogues start with an agent putting forward a claim. Our dialogue model is generic in that it does not focus on any particular dialogue type, e.g. information seeking, persuasion or negotiation [43], but can be used to support several such dialogue types [17, 18, 19, 20].

Through dialogues, the participating agents construct a “joint knowledge base” by pooling all disclosed information to form an *ABA framework*. The ABA framework drawn from a dialogue δ , referred to as \mathcal{F}_δ , contains all information that the two agents have uttered in the dialogue and gives the context for examining the acceptability of the claim of the dialogue. Conceptually, a dialogue is “successful” if its claim is “acceptable” in \mathcal{F}_δ . Note that the claim of a dialogue may be a belief, and acceptability thereof an indication that the agents may legitimately uphold the belief, or a course of actions, and acceptability thereof an indication that the agents may legitimately choose to adhere to it. Indeed, acceptability has so far shown to be an important criterion for assessing the outcome of various types of dialogues [17, 18, 19, 20], and thus “successful” dialogues can be seen as building blocks of a widely deployable framework for distributed interactions in multi-agent systems. We focus here on three forms of “acceptability” and “success”, with respect to three well-known argumentation semantics.

Rather than checking “success” retrospectively, we define *legal-move functions* guaranteed to generate “successful” dialogues if a limited form of retrospective checking by means of *outcome functions* succeeds. Given a dialogue, a legal-move function returns a set of allowed utterances that can be uttered to extend the dialogue. Legal-move functions can thus be viewed as dialogue protocols. Outcome functions are mappings from dialogues to true / false. Given a dialogue, an outcome function returns true if a certain property holds for that dialogue.

In summary, the main contributions of this work are: (1) a generic formal model for ABA-based dialogues; and (2) the link between this model and standard argumentation semantics (of grounded, admissible and ideal extensions) to define success of these dialogues. We focus on these three semantics as they allow to cap-

ture general forms of credulous reasoning (admissible) and two well-understood forms of sceptical reasoning (grounded and ideal), and are thus suitable for a wide range of problems. Our soundness results are obtained by mapping the *debate tree/forest* generated from a dialogue onto an *abstract dispute tree* [14] that is known to sanction the “acceptability” of the claim [14, 15]. These debate tree/forest can be seen as a *commitment store* [43] holding information that agents disclose and share using the dialogue.

The paper generalises and extends the initial proposal of ABA-based dialogues in [16] in several ways. Firstly, this paper shows soundness results with respect to grounded, admissible and ideal extensions, rather than just admissible extensions as in [16]. Secondly, [16] uses “dialectical trees”, which are mapped onto the *concrete dispute trees* of [14] whereas this work uses a new notion of *debate trees* (see Definition 8.1), which are mapped onto the *abstract dispute trees* of [14], directly allowing to use soundness results from [15]. Moreover, [16] defines dialectical trees constructively, whereas this work defines debate trees declaratively, allowing to prove some novel results (e.g. Lemma 11.1). Thirdly, in this paper we define *debate forests* and use them to study *unrestricted* dialogues, completely absent from [16], which studies *focused* dialogues only.

The article is structured as follows. Section 2 presents background on ABA. Section 3 sets the foundation of our dialogue framework and Section 4 introduces generic notions of legal-move and outcome functions. Section 5 defines specific kinds of these functions to generate special kinds of dialogues, guaranteed to draw ABA frameworks. Section 6 defines the three notions of “successful” dialogue we are after, in a non-constructive way. Section 7 starts refining the dialogue framework by introducing new legal-move and outcome functions that enforce core properties of “successful” dialogues, constructively. Section 8 presents *debate trees* that are then used to define legal-move and outcome functions, in Section 9, allowing to construct “successful” dialogues and prove our soundness results in Sections 10 and 11, for *focused* and *unrestricted* dialogues, respectively. Section 12 discusses related works and Section 13 concludes.

The proposed dialogue model relies upon several notions and formal definitions: the most important amongst these are summarised in a glossary in Appendix A, to aid readability.

2. Background: Assumption-based Argumentation (ABA)

An *ABA framework* is a tuple $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ where

- $\langle \mathcal{L}, \mathcal{R} \rangle$ is a deductive system, with a *language* \mathcal{L} and a set of inference *rules* \mathcal{R} of the form $\beta_0 \leftarrow \beta_1, \dots, \beta_m$ ($m > 0$) or $\beta_0 \leftarrow$ with, for $i = 0, \dots, m$, $\beta_i \in \mathcal{L}$, and, if $m > 1$, then $\beta_i \neq \beta_j$ for $i \neq j, 1 \leq i, j \leq m$;¹
- $\mathcal{A} \subseteq \mathcal{L}$ is a (non-empty) set, whose elements are referred to as *assumptions*,
- \mathcal{C} is a total mapping from \mathcal{A} into $2^{\mathcal{L}} - \{\{\}\}$, where each $c \in \mathcal{C}(\alpha)$ is a *contrary* of α .²

Basically, ABA frameworks can be defined for any logic specified by means of inference rules. Some of the sentences in the underlying language are assumptions.

Given $\beta_0 \leftarrow \beta_1, \dots, \beta_m$ or $\beta_0 \leftarrow$, β_0 is referred as the *head* and β_1, \dots, β_m or the empty sequence, respectively, as the *body*. We will use the following notation: $Head(\beta_0 \leftarrow \beta_1, \dots, \beta_m) = \beta_0$; $Body(\beta_0 \leftarrow \beta_1, \dots, \beta_m) = \beta_1, \dots, \beta_m$, $Head(\beta_0 \leftarrow) = \beta_0$, and $Body(\beta_0 \leftarrow)$ is empty. An ABA framework is *flat* iff no assumption is the head of a rule. We will focus on flat ABA frameworks, so as to be able to use results valid for them (see later).

Example 2.1. Let $\mathcal{AF}_1 = \langle \mathcal{L}_1, \mathcal{R}_1, \mathcal{A}_1, \mathcal{C}_1 \rangle$ be as follows:

- $\mathcal{L}_1 = \{a, b, c, p, q, r, s\}$,
- $\mathcal{R}_1 = \{p \leftarrow a, q \leftarrow b, s, r \leftarrow c, s \leftarrow\}$,
- $\mathcal{A}_1 = \{a, b, c\}$,
- $\mathcal{C}_1(a) = \{q\}$, $\mathcal{C}_1(b) = \{p\}$, $\mathcal{C}_1(c) = \{r, q\}$.

\mathcal{AF}_1 is a flat ABA framework. Let $\mathcal{AF}_2 = \langle \mathcal{L}_2, \mathcal{R}_2, \mathcal{A}_2, \mathcal{C}_2 \rangle$ be as follows:

- $\mathcal{L}_2 = \{a, p, q\}$,
- $\mathcal{R}_2 = \{p \leftarrow a\}$,
- $\mathcal{A}_2 = \{p, a\}$,
- $\mathcal{C}_2(a) = \{q\}$, $\mathcal{C}_2(p) = \{q\}$.

\mathcal{AF}_2 is a non-flat ABA framework, as p is both the head of a rule ($p \leftarrow a$) and an assumption.

¹Standard ABA does not require $\beta_i \neq \beta_j$, but this can be imposed without loss of generality. As an example, consider two ABA frameworks, identical except for rule $p \leftarrow q, q$ in the first and rule $p \leftarrow q$ in the second. Then, trivially, p (or any other sentence in the underlying language) is admissible (or complete, grounded, preferred, ideal, defined later) in the first iff it is in the second.

²Here, as in [21], we (equivalently) define the contrary of an assumption as a total mapping from an assumption to a (non-empty) set of sentences, instead of a mapping from an assumption to a sentence as in the original ABA. This lends itself better to a dialogical setting, as agents may hold different sentences as contrary to the same assumption.

In ABA, informally, arguments are deductions of claims supported by sets of assumptions, and attacks against arguments are directed at the assumptions in their supports, and are given by arguments with an assumption’s contrary as their claim. Formally, given an ABA framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$, an *argument* for $\beta \in \mathcal{L}$ supported by $A \subseteq \mathcal{A}$ is a tree with nodes labelled by sentences in \mathcal{L} or by τ^3 , such that

1. the root is labelled by β
2. for every node N
 - if N is a leaf then N is labelled either by an assumption in A or by τ ;
 - if N is not a leaf and β_0 is the label of N , then there is an inference rule $\beta_0 \leftarrow \beta_1, \dots, \beta_m (m \geq 0)$ and either $m = 0$ and the child of N is τ or $m > 0$ and N has m children, labelled by β_1, \dots, β_m (respectively)
3. A is the set of all assumptions labelling the leaves.

The shorthand $A \vdash \beta$ is used to denote an argument for β supported by A . Given argument $A \vdash \beta$, A is referred to as the *support* and β as the *claim*. Figure 1 illustrates this notion of argument for \mathcal{AF}_1 in Example 2.1.




Figure 1: Arguments in Example 2.1, as trees (above) and using standard shorthand (below).

With the notion of arguments and contrary of assumption, *attack* in an ABA framework is defined as follows:

- an argument $A_1 \vdash \beta_1$ attacks an argument $A_2 \vdash \beta_2$ iff the claim β_1 of the first argument is a contrary of one of the assumptions in the support A_2 of the second argument (i.e., $\exists \alpha \in A_2$ such that $\beta_1 \in \mathcal{C}(\alpha)$);

³The symbol τ is such that $\tau \notin \mathcal{L}$. τ stands for “true” and intuitively represents the empty body of rules.

- a set of arguments As attacks a set of arguments Bs if some argument in As attacks some argument in Bs ;
- a set of assumptions A_1 attacks a set of assumptions A_2 iff an argument supported by a subset of A_1 attacks an argument supported by a subset of A_2 .

Example 2.2. (Continuation of Example 2.1) In \mathcal{AF}_1 , we have that $\{a\} \vdash p$ attacks $\{b\} \vdash q$, $\{b\} \vdash q$ attacks $\{a\} \vdash p$, $\{b\} \vdash q$ attacks $\{c\} \vdash r$, and $\{c\} \vdash r$ attacks $\{a\} \vdash p$. Also, $\{\{a\} \vdash p\}$ attacks $\{\{b\} \vdash q, \{c\} \vdash r\}$ and $\{a\}$ attacks $\{b\}$, $\{b\}$ attacks $\{a, c\}$, $\{b\}$ attacks $\{c\}$, and $\{c\}$ attacks any set of assumptions containing c .

With argument and attack defined, all argumentation semantics introduced in abstract argumentation [11] can be applied in ABA. These semantics can be defined for assumptions, as in [8, 7], or, equivalently, for arguments [15, 40]. Formally, given $\mathcal{F} = \langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$, sets of assumptions can be deemed to be acceptable according to the following semantics⁴:

- *a set of assumptions is admissible* (in \mathcal{F}) iff it does not attack itself and it attacks all $A \subseteq \mathcal{A}$ that attack it;
- *a set of assumptions is complete* (in \mathcal{F}) iff it is admissible (in \mathcal{F}) and contains all assumptions it *defends* (in \mathcal{F}), where a set of assumptions A *defends* another set of assumptions A' (in \mathcal{F}) iff A attacks all sets of assumptions that attack A' ;
- *a set of assumptions is grounded* (in \mathcal{F}) iff it the least set (with respect to \subseteq) that is complete (in \mathcal{F});
- *a set of assumptions is preferred* (in \mathcal{F}) iff it is maximally (with respect to \subseteq) admissible (in \mathcal{F});
- *a set of assumptions is ideal* (in \mathcal{F}) iff it is an admissible set (in \mathcal{F}) contained in all preferred sets (in \mathcal{F});

In this paper, we will also use the following notions of individual arguments and sentences being acceptable according to different argumentation semantics:

⁴This is not an exhaustive list, but includes all semantics relevant to this paper.

- an argument $A \vdash \beta$ is admissible (complete, grounded, preferred, ideal) (in \mathcal{F}) supported by $A' \subseteq \mathcal{A}$ iff $A \subseteq A'$ and A' is admissible (complete, grounded, preferred, ideal, respectively) (in \mathcal{F});
- a sentence is admissible (complete, grounded, preferred, ideal) (in \mathcal{F}) iff it is the claim of an argument that is admissible (complete, grounded, preferred, ideal, respectively) (in \mathcal{F}) supported by some $A \subseteq \mathcal{A}$.

In this paper we will focus on the admissible, grounded and ideal semantics.

Example 2.3. (Continuation of Example 2.2) The sets of assumptions $\{a\}$, $\{b\}$, $\{\}$ are all admissible (and no other set is admissible). Amongst these, $\{\}$ is the only grounded as well as the only ideal set, and $\{a\}$ and $\{b\}$ are the only preferred sets.

The sets of arguments $\{\{a\} \vdash p\}$, $\{\{b\} \vdash q\}$, $\{\{a\} \vdash a\}$, $\{\{b\} \vdash b\}$, $\{\{a\} \vdash p, \{a\} \vdash a\}$, $\{\{b\} \vdash q, \{b\} \vdash b\}$, and $\{\}$ are all admissible (and no other set is admissible). Amongst these, $\{\}$ is the only grounded as well as the only ideal set, and $\{\{a\} \vdash p, \{a\} \vdash a\}$ and $\{\{b\} \vdash q, \{b\} \vdash b\}$ are the only preferred sets. Finally, p and q are both admissible, whereas no sentence is grounded or ideal.

We will use the *abstract dispute trees* of [14] to prove some of our results later, where an *abstract dispute tree* for an argument A (in an ABA framework \mathcal{F}) is a (possibly infinite) tree \mathcal{T}^a such that:⁵

1. every node of \mathcal{T}^a holds an argument B (in \mathcal{F}) and is labelled by either *proponent* (P) or *opponent* (O), but not both, denoted by $L : B$, for $L \in \{P, O\}$;
2. the root of \mathcal{T}^a is a P node holding A;
3. for every P node N holding an argument B, and for every argument C that attacks B (in \mathcal{F}), there exists a child of N , which is an O node holding C;
4. for every O node N holding an argument B, there exists exactly one child of N which is a P node holding an argument which attacks (in \mathcal{F}) some assumption α in the support of B.⁶ α is said to be the *culprit* in B;
5. there are no other nodes in \mathcal{T}^a except those given by 1-4 above.

⁵Here, a stands for 'abstract'. Also, proponent and opponent are roles/fictitious players rather than actual agents, and abstract dispute trees are not to be confused with a dialogical model. We will return to this later.

⁶An argument attacks an assumption if the argument supports a contrary of the assumption.




Figure 2: Admissible abstract dispute tree for Example 2.4.

The set of all assumptions in (the support of arguments hold of) the P nodes in \mathcal{T}^a is called the *defence set* of \mathcal{T}^a .

Abstract dispute trees can be used as the basis for computing various argumentation semantics as follows, given an ABA framework \mathcal{F} :

- An abstract dispute tree \mathcal{T}^a is *admissible* iff no culprit in the argument of an O labelling node in \mathcal{T}^a belongs to the defence set of \mathcal{T}^a . The defence set of an admissible abstract dispute tree for an argument A (in \mathcal{F}) is admissible (Theorem 5.1 in [14]), and thus A as well as the sentence in the root node of A is admissible (in \mathcal{F}).
- An abstract dispute tree is *grounded* iff it is admissible and finite. The defence set of a grounded abstract dispute tree \mathcal{T}^a for an argument A (in \mathcal{F}) is contained in the grounded set [12], and thus A as well as the sentence in the root node of A is grounded (in \mathcal{F}).
- An abstract dispute tree is *ideal* iff for no O node N in the tree there exists an admissible abstract dispute tree with root N (in \mathcal{F}). The defence set of an ideal abstract dispute tree for an argument is ideal (Theorem 3.4 in [15]), and thus A as well as the sentence in the root node of A is ideal (in \mathcal{F}).

Example 2.4. (Continuation of Example 2.3) Figure 2 gives an example of an (infinite) admissible dispute tree, with (admissible) defence set $\{a\}$.

3. Basic Dialogue Notions

This section presents the basis of our model of argumentation-based dialogues between (two) agents, a_1 and a_2 . In our dialogue model, agents can utter claims (to be debated), rules, assumptions and contraries, or pass. Thus, dialogues allow

agents to “build” a shared repository of components of ABA frameworks. Note that a_1, a_2 may or may not be equipped with ABA frameworks. ABA is used as a lingua franca in the spirit of the Argument Interchange Format (AIF) [9, 37], in the sense that, internally, agents can use knowledge representation formalisms different from ABA, but, while communicating, agents convert their internal representation into ABA. Even if a_1, a_2 are equipped with ABA frameworks, the agents may or may not be truthful, in that, for example, an agent may utter “made-up” rules which do not exist in its ABA framework.

We assume the agents share a common logical language \mathcal{L} and a (non-empty, possibly infinite) set \mathcal{ID} that:

- is totally ordered, with the ordering given by $<$;
- contains a special element ID_0 which is the least element with respect to $<$.

Example 3.1. We can choose \mathcal{ID} to be the set of non-negative integers, $\mathbb{N} \cup \{0\}$. The total order relation $<$ is the standard $<$ relation, such that for $a, b \in \mathcal{ID}$, $a < b$ iff there exists some $c \in \mathbb{N}$ such that $a + c = b$. For this choice of \mathcal{ID} , ID_0 is 0.

For simplicity, in all our examples \mathcal{ID} and ID_0 are as in Example 3.1. However, several other choices are possible, including a closed interval of real numbers and a set of strings (with $<$ the lexicographic order). Utterances are defined as follows:

Definition 3.1. An *utterance* from agent a_i to agent a_j ($i, j = 1, 2, i \neq j$) is a tuple $\langle a_i, a_j, T, C, ID \rangle$ where:

- C (the *content*) is of one of the following forms:
 - $claim(\chi)$ for some $\chi \in \mathcal{L}$ (a *claim*),
 - $rl(\beta_0 \leftarrow \beta_1, \dots, \beta_m)$ for some $\beta_0, \dots, \beta_m \in \mathcal{L}$ with $m \geq 0$ (a *rule*),
 - $asm(\alpha)$ for some $\alpha \in \mathcal{L}$ (an *assumption*),
 - $ctr(\alpha, \beta)$ for some $\alpha, \beta \in \mathcal{L}$ (a *contrary*),
 - a *pass* sentence π , such that $\pi \notin \mathcal{L}$.
- $ID \in \mathcal{ID} \setminus \{ID_0\}$ (the *identifier*).
- $T \in \mathcal{ID}$ (the *target*); we impose that $T < ID$.

We refer to an utterance with content π as a *pass utterance*, to an utterance with content $claim(-)$ as a *claim utterance*⁷ and to an utterance with content other than π or $claim(-)$ as *regular*. We say that an utterance from a_i is *made by* a_i .

Note that our utterances are generic and not tailored to any specific types of dialogues, but can be used to provide specialised semantics for specific dialogue types (such as persuasion and inquiry [43]), e.g. following the lines of [17, 18, 19, 20].

Notation 3.1. We use \mathcal{U} to denote the set of all possible utterances as in Definition 3.1, given \mathcal{L} , and \mathcal{U}^i to denote all utterances from a_i in \mathcal{U} , i.e. of the form $\langle a_i, -, -, - \rangle$.

Intuitively, when the content of an utterance is π , the utterance indicates that the agent making it does not have or want to contribute any information (i.e. claim, rule, assumption, contrary) that can be added to the dialogue at that point, either because no such information is in the agent’s possession or because the agent chooses not to disclose such information. The target of an utterance is the identifier of some earlier utterance in the dialogue, as we will see next.

Definition 3.2. A dialogue $\mathcal{D}_{a_j}^{a_i}(\chi)$ (between a_i and a_j , $i, j \in \{1, 2\}$, $i \neq j$, for $\chi \in \mathcal{L}$), is a sequence $\langle u_1, \dots, u_n \rangle$, $n \geq 0$, where each u_l , $l = 1, \dots, n$, is in \mathcal{U} , and:

1. $u_1 = \langle a_i, a_j, -, - \rangle$;
2. the content of u_l is $claim(\chi)$ iff $l = 1$;
3. the target of pass and claim utterances is ID_0 ; the target of regular utterances is not ID_0 ;
4. for every utterance $u_l = \langle -, -, T, -, - \rangle$, such that $l > 1$ and $T \neq ID_0$, there exists some $u_k = \langle -, -, C, T \rangle$, such that $C \neq \pi$ and $k < l$;
5. for $0 \leq k < l \leq n$, if ID_l is the identifier of u_l and ID_k is the identifier of u_k then $ID_k < ID_l$.

Given a dialogue $\delta = \langle u_1, \dots, u_n \rangle$ and an utterance u , we define $\delta \circ u = \langle u_1, \dots, u_n, u \rangle$. Given $\delta = \langle u_1, \dots, u_n \rangle$, we say that each u_i , $1 \leq i \leq n$, is *in* δ .

Notation 3.2. We use \mathcal{D} to denote the set of all dialogues as in Definition 3.2.

⁷Throughout, $-$ stands for an anonymous variable as in Prolog.

Example 3.2. Given $\mathcal{L} = \{s, a, b, c, d, g, q, r\}$, and $\mathcal{ID} = \mathbb{N} \cup \{0\}$ with $ID_0 = 0$, a possible dialogue $\mathcal{D}_{a_2}^{a_1}(s)$ is as follows:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 0, \pi, 2 \rangle$
$\langle a_1, a_2, 1, \text{rl}(s \leftarrow a), 3 \rangle$	$\langle a_2, a_1, 3, \text{asm}(a), 4 \rangle$
$\langle a_1, a_2, 4, \text{ctr}(a, q), 5 \rangle$	$\langle a_2, a_1, 5, \text{rl}(q \leftarrow b), 6 \rangle$
$\langle a_1, a_2, 6, \text{asm}(b), 7 \rangle$	$\langle a_2, a_1, 7, \text{ctr}(b, c), 8 \rangle$
$\langle a_1, a_2, 8, \text{asm}(c), 9 \rangle$	$\langle a_2, a_1, 9, \text{ctr}(c, r), 10 \rangle$
$\langle a_1, a_2, 0, \pi, 11 \rangle$	$\langle a_2, a_1, 0, \pi, 12 \rangle$

Our dialogue model allows pass utterances being uttered at any moment throughout a dialogue. Given any dialogue, the π -pruned sequence obtained from the dialogue consists of all regular utterances in the dialogue.

Example 3.3. The π -pruned sequence obtained from the dialogue δ in Example 3.2 is δ' as follows:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	
$\langle a_1, a_2, 1, \text{rl}(s \leftarrow a), 3 \rangle$	$\langle a_2, a_1, 3, \text{asm}(a), 4 \rangle$
$\langle a_1, a_2, 4, \text{ctr}(a, q), 5 \rangle$	$\langle a_2, a_1, 5, \text{rl}(q \leftarrow b), 6 \rangle$
$\langle a_1, a_2, 6, \text{asm}(b), 7 \rangle$	$\langle a_2, a_1, 7, \text{ctr}(b, c), 8 \rangle$
$\langle a_1, a_2, 8, \text{asm}(c), 9 \rangle$	$\langle a_2, a_1, 9, \text{ctr}(c, r), 10 \rangle$

Here $\delta' = \langle u'_1, \dots, u'_9 \rangle$ where $u'_1 = u_1, u'_2 = u_3, u'_3 = u_4, u'_4 = u_5, u'_5 = u_6, u'_6 = u_7, u'_7 = u_8, u'_8 = u_{10}$, for $u_i = \langle -, -, -, i \rangle$ in the original δ .

Note that, since no regular utterance has a pass utterance as its target (by condition 3 in Definition 3.2), the target of every utterance in a π -pruned sequence is guaranteed to be in this sequence. Also, since the first utterance in a dialogue can never be a pass utterance (by condition 2 of Definition 3.2), the first utterance in a π -pruned sequence is always the same as the first utterance in the original dialogue. Moreover, it is trivially true that, for any utterance u , if u is not in a dialogue δ , u is not in the π -pruned sequence of δ .

By means of dialogues agents “build” a shared repository of components of an ABA framework, defined as follows:

Definition 3.3. The (argumentation) framework drawn from a dialogue $\delta = \langle u_1, \dots, u_n \rangle \in \mathcal{D}$ is $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ where:

- $\mathcal{R}_\delta = \{\rho | rl(\rho) \text{ is the content of some } u_i \text{ in } \delta\}$;
- $\mathcal{A}_\delta = \{\alpha | asm(\alpha) \text{ is the content of some } u_i \text{ in } \delta\}$;
- for any $\alpha \in \mathcal{A}_\delta$, $\mathcal{C}_\delta(\alpha) = \{\beta | ctr(\alpha, \beta) \text{ is the content of some } u_i \text{ in } \delta\}$.

Clearly, the framework drawn from a dialogue represents all information that has been disclosed by the two agents in the dialogue.

Example 3.4. The framework drawn from the dialogue δ in Example 3.2 is $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$, in which

- $\mathcal{A}_\delta = \{a, b, c\}$;
- $\mathcal{R}_\delta = \{s \leftarrow a, q \leftarrow b\}$;
- $\mathcal{C}_\delta(a) = \{q\}, \mathcal{C}_\delta(b) = \{c\}, \mathcal{C}_\delta(c) = \{r\}$.

Note that the framework drawn from the dialogues in Example 3.4 is an ABA framework, but, in general, the framework drawn from a dialogue may or may not be an ABA framework, as we will discuss later in Section 5.

3.1. Illustration: Twelve Angry Men

Throughout the paper, we will illustrate our model in the context of the following example, adapted from the movie *Twelve Angry Men*, an example of argumentation-based collaborative reasoning [2]. Here, twelve jurors need to decide whether to condemn a boy, accused of murder, or acquit him, after a trial where two witnesses have provided evidence against the boy. According to the law, the jurors should acquit the boy if they do not believe that the trial has proven him guilty convincingly. We focus on the reasoning of two of the jurors: juror 8, played by Henry Fonda (a_1), and juror 9, played by Joseph Sweeney (a_2). We can model their exchanges of opinions as the two-agents dialogue $\delta = \mathcal{D}_{a_2}^{a_1}(boy_innocent)$ in Table 1. Here, a_1 starts the dialogue by putting forward the claim that the boy is innocent (*boy_innocent*). Then both agents contribute rules, assumptions and contraries for and against the claim (directly or indirectly). The content of utterances in this dialogue should be self-explanatory. For example, *boy_innocent* \leftarrow *boy_not_proven_guilty* says that the boy should be deemed to be innocent if he cannot be proven guilty. A natural language reading of this dialogue is in Table 2. The framework drawn from this dialogue δ is $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ where

- $\mathcal{A}_\delta = \{boy_not_proven_guilty, w1_is_believable, w2_is_believable\}$;
- $\mathcal{R}_\delta = \{boy_innocent \leftarrow boy_not_proven_guilty, boy_proven_guilty \leftarrow w1_is_believable,$

Table 1: Dialogue δ for Example 3.1.

$\langle a_1, a_2, 0, \text{claim}(\text{boy_innocent}), 1 \rangle$
$\langle a_2, a_1, 1, \text{rl}(\text{boy_innocent} \leftarrow \text{boy_not_proven_guilty}), 2 \rangle$
$\langle a_1, a_2, 2, \text{asm}(\text{boy_not_proven_guilty}), 3 \rangle$
$\langle a_2, a_1, 3, \text{ctr}(\text{boy_not_proven_guilty}, \text{boy_proven_guilty}), 4 \rangle$
$\langle a_1, a_2, 4, \text{rl}(\text{boy_proven_guilty} \leftarrow \text{w1_is_believable}), 5 \rangle$
$\langle a_2, a_1, 5, \text{asm}(\text{w1_is_believable}), 6 \rangle$
$\langle a_1, a_2, 6, \text{ctr}(\text{w1_is_believable}, \text{w1_not_believable}), 7 \rangle$
$\langle a_1, a_2, 7, \text{rl}(\text{w1_not_believable} \leftarrow \text{w1_contradicted_by_w2}), 8 \rangle$
$\langle a_1, a_2, 8, \text{rl}(\text{w1_contradicted_by_w2} \leftarrow), 9 \rangle$
$\langle a_2, a_1, 4, \text{rl}(\text{boy_proven_guilty} \leftarrow \text{w2_is_believable}), 10 \rangle$
$\langle a_1, a_2, 10, \text{asm}(\text{w2_is_believable}), 11 \rangle$
$\langle a_2, a_1, 11, \text{ctr}(\text{w2_is_believable}, \text{w2_not_believable}), 12 \rangle$
$\langle a_1, a_2, 12, \text{rl}(\text{w2_not_believable} \leftarrow \text{w2_has_poor_eyesight}), 13 \rangle$
$\langle a_2, a_1, 13, \text{rl}(\text{w2_has_poor_eyesight} \leftarrow), 14 \rangle$
$\langle a_1, a_2, 0, \pi, 15 \rangle$
$\langle a_2, a_1, 0, \pi, 16 \rangle$

Table 2: A possible natural language reading of the dialogue in Table 1.

a_1 :	The boy is innocent.
a_2 :	The boy is innocent if he is not proven guilty.
a_1 :	We assume the boy is not proven guilty.
a_2 :	Proving that the boy is guilty disagrees with this assumption.
a_1 :	The boy is guilty if witness 1 is believable.
a_2 :	We assume witness 1 is believable.
a_1 :	Showing witness 1 cannot be believed disagrees with this assumption.
a_1 :	Witness 1 cannot be believed if it is contradicted by witness 2.
a_1 :	Witness 1 is indeed contradicted by witness 2.
a_2 :	The boy is guilty if witness 2 is believable.
a_1 :	We assume witness 2 is believable.
a_2 :	Showing witness 2 cannot be believed disagrees with this assumption.
a_1 :	Witness 2 cannot be believed if it has a poor eyesight.
a_2 :	Witness 2 indeed has a poor eyesight.
a_1 :	OK.
a_2 :	OK.

$$\begin{aligned}
& \text{boy_proven_guilty} \leftarrow \text{w2_is_believable}, \\
& \text{w1_not_believable} \leftarrow \text{w1_contradicted_by_w2}, \\
& \text{w1_contradicted_by_w2} \leftarrow, \\
& \text{w2_not_believable} \leftarrow \text{w2_has_poor_eyesight}, \\
& \text{w2_has_poor_eyesight} \leftarrow \}
\end{aligned}$$

- $\mathcal{C}_\delta(\text{boy_not_proven_guilty}) = \{\text{boy_proven_guilty}\},$
 $\mathcal{C}_\delta(\text{w1_is_believable}) = \{\text{w1_is_not_believable}\},$
 $\mathcal{C}_\delta(\text{w2_is_believable}) = \{\text{w2_is_not_believable}\}.$

Note that this is an ABA framework.

4. Constructing Dialogues

This section gives the basic notions used in our model to allow agents to construct dialogues. As in other dialogue models, notably that of [33], these notions amount to *turn-making functions*, to determine which agent should make an utterance at any point in a dialogue, *legal-move functions*, to determine which utterances agents can make during dialogues, and *outcome functions*, to determine whether dialogues have desirable properties.

In Example 3.2, the two agents take turns in making utterances: a strict interleaving is enforced between these two agents. In general, such requirement is unnecessary, i.e., an agent may be allowed to make a few consecutive utterances before the other agent makes any. We define a turn-making function to specify which agent makes the next utterance in a dialogue.

Definition 4.1. A *turn-making function* is a mapping $\gamma : \mathcal{D} \mapsto \{a_1, a_2\}$ such that, given $\delta = \mathcal{D}_{a_j}^{a_i}(\chi) = \langle u_1, \dots, u_n \rangle, i, j \in \{1, 2\}, i \neq j$:

$$\gamma(\delta) = \begin{cases} a_i & \text{if } n = 0, \\ a_x \in \{a_i, a_j\} & \text{otherwise.} \end{cases}$$

A dialogue $\langle u_1, \dots, u_n \rangle$ ($n > 0$) is *compatible with a turn-making function* γ iff for each $l = 1, \dots, n$, if $u_l = \langle a_x, -, -, - \rangle$ then $\gamma(\langle u_1, \dots, u_{l-1} \rangle) = a_x$.

Our definition of turn-making function is very liberal, in that it states that a_i starts $\mathcal{D}_{a_j}^{a_i}(\chi)$ and all subsequent utterances are made by any of the agents as dictated by γ . As observed earlier, $\gamma(\delta)$ for $\delta = \mathcal{D}_{a_2}^{a_1}(\chi)$ in Example 3.2 forces a strict interleaving between a_1 and a_2 , whereas the dialogue in Table 3.2 does not.

Definition 4.2. A *legal-move function* (with respect to a turn-taking function γ) is a mapping $\lambda : \mathcal{D} \mapsto 2^{\mathcal{U}}$ such that, given $\delta = \langle u_1, \dots, u_n \rangle \in \mathcal{D}$, for all $u \in \lambda(\delta)$:

1. $\delta \circ u$ is a dialogue;
2. $\delta \circ u$ is compatible with γ if δ is;
3. if $u = \langle -, -, T, C, - \rangle$, then there is no $i, 1 \leq i \leq n$, such that $u_i = \langle -, -, T, C, - \rangle$.

Given $\delta = \langle u_1, \dots, u_n \rangle$, if $u_{m+1} \in \lambda(\langle u_1, \dots, u_m \rangle)$ for all m such that $0 \leq m < n$, we say that δ is *compatible* with λ .

Here condition 3 regulates that there is no repeated utterance to the same target in a dialogue. However, the definition of legal-move function does not impose any “mentalistic” requirement on agents, such as that they utter information they hold true, similarly to communication protocols in multi-agent systems [44].

Notation 4.1. We use Λ to denote the set of all legal-move functions.

Different types of dialogues will require a different turn-making functions. For example, negotiation typically requires a strict interleaving, whereas inquiry does not (although strict interleaving may be used for inquiry too, as in [6]). As our dialogue framework is generic, from now on we will assume as given a generic turn-making function γ and we will omit to mention it in definitions, assuming implicitly that all our dialogues are compatible with this γ . In particular, we will omit the turn-making function when giving legal-move functions.

We conclude this section by defining outcome functions, to determine whether dialogues have desirable properties.

Definition 4.3. An *outcome function* is a mapping $\omega : \mathcal{D} \times \Lambda \mapsto \{true, false\}$.

Notation 4.2. We use Ω to denote the set of all outcome functions.

In the remainder we give several concrete legal-move and outcome functions.

5. ABA Dialogues

In this section, we refine our dialogues so that the frameworks drawn from them are guaranteed to be flat ABA frameworks. The resulting dialogues are referred to as *ABA dialogues*. This refinement builds upon specific kinds of legal-move function and outcome function to restrict the kind of allowed utterances.

As mentioned earlier, in general, the framework drawn from a dialogue may not be an ABA framework, since the contrary of some assumption may be empty:

Example 5.1. Given $\mathcal{L} = \{s, a\}$, let δ be:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 1, \text{rl}(s \leftarrow a), 2 \rangle$
$\langle a_1, a_2, 2, \text{asm}(a), 3 \rangle$	

The framework drawn from this dialogue, $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$, has:

- $\mathcal{A}_\delta = \{a\}$;
- $\mathcal{R}_\delta = \{s \leftarrow a\}$;
- $\mathcal{C}_\delta(a) = \{\}$.

\mathcal{F}_δ , is not an ABA framework as the contrary of a in this framework is empty.

Even when the framework drawn from a dialogue is an ABA framework, it may not be flat, as the agents may disagree on the assumptions:

Example 5.2. Given $\mathcal{L} = \{s, a, q, r\}$, let δ be:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 1, \text{rl}(s \leftarrow a), 2 \rangle$
$\langle a_1, a_2, 2, \text{asm}(a), 3 \rangle$	$\langle a_2, a_1, 2, \text{rl}(a \leftarrow q), 4 \rangle$
$\langle a_1, a_2, 3, \text{ctr}(a, r), 5 \rangle$	

The framework $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ has:

- $\mathcal{A}_\delta = \{a\}$;
- $\mathcal{R}_\delta = \{s \leftarrow a, a \leftarrow q\}$;
- $\mathcal{C}_\delta(a) = \{r\}$.

Here a is both an assumption and the head of a rule, hence \mathcal{F}_δ is not flat.

We ensure flatness of the framework drawn from a dialogue using a specific kind of legal-move function, as follows.

Definition 5.1. A *flat legal-move function* is a legal-move function $\lambda \in \Lambda$ such that, given a dialogue $\delta = \langle u_1, \dots, u_n \rangle \in \mathcal{D}$, for all $u = \langle -, -, -, C, - \rangle \in \lambda(\delta)$, then

- $C = \text{asm}(\alpha)$ only if there exists no u_i for $1 \leq i \leq n$ with content $\text{rl}(\rho)$ and $\text{Head}(\rho) = \alpha$;
- $C = \text{rl}(\rho)$ only if there exists no u_i for $1 \leq i \leq n$ with content $\text{asm}(\alpha)$ and $\text{Head}(\rho) = \alpha$.

δ in Example 5.2 is not compatible with a flat legal-move function.

Lemma 5.1. Given a dialogue $\delta \in \mathcal{D}$, if δ is compatible with a flat legal-move function and the framework $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ drawn from δ is an ABA framework, then \mathcal{F}_δ is a flat ABA framework.

The proof of this and all results in this article can be found in Appendix B.

In order to guarantee that the framework drawn from a dialogue is an ABA framework, we will use a specific kind of outcome function imposing that each assumption has a non-empty set as its contrary:

Definition 5.2. The *ABA outcome function* ω_{ABA} is such that given a dialogue $\delta \in \mathcal{D}$ and a legal-move function $\lambda \in \Lambda$, $\omega_{ABA}(\delta, \lambda) = true$ iff δ is compatible with λ and the framework $\langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ drawn from δ is such that for all $\alpha \in \mathcal{A}_\delta$, $\mathcal{C}_\delta(\alpha) \neq \{\}$. A *λ -ABA dialogue* is a dialogue $\delta \in \mathcal{D}$ with $\lambda \in \Lambda$ flat such that $\omega_{ABA}(\delta, \lambda) = true$. An *ABA dialogue* is a λ -ABA dialogue for some $\lambda \in \Lambda$.

Lemma 5.2. The framework drawn from an ABA dialogue is a flat ABA framework.

In the remainder, all dialogues will be λ -ABA dialogues, for some (flat) λ . In the remainder of the paper we will impose further restrictions on λ to obtain specialised forms of dialogues. First, however, let us give a concrete example of ABA dialogue.

5.1. Illustration: Twelve Angry Men

The dialogue δ in Table 1 is an ABA dialogue. Indeed, it is compatible with a flat legal-move function λ_{fl} since there is no β such that both $asm(\beta)$ and $rl(\beta \leftarrow \dots)$ are contents of utterances in δ . As a result, $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ is a (flat) ABA framework and indeed $\omega_{ABA}(\delta, \lambda_{fl}) = true$, since for all assumptions $\alpha \in \mathcal{A}_\delta = \{boy_not_proven_guilty, w1_is_believable, w2_is_believable\}$ the contrary of α is not empty.

6. Successful Dialogues

The definitions so far jointly establish the foundations of a generic dialogue framework for determining acceptability of claims (in the flat ABA framework drawn from ABA dialogues). In this section, we define three notions of *successful* dialogues, intuitively amounting to sanctioning the initial claims as acceptable according to the admissible, grounded and ideal semantics, respectively.

Definition 6.1. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta$, let \mathcal{F}_δ be the ABA framework drawn from δ . Then δ is:

- *a-successful* iff χ is admissible in \mathcal{F}_δ ;
- *g-successful* iff χ is grounded in \mathcal{F}_δ ;
- *i-successful* iff χ is ideal in \mathcal{F}_δ .

These notions equate success of a dialogue with determining whether its claim is semantically acceptable, according to the three semantics we focus on, in the ABA framework drawn from the dialogue. Thus, successful dialogues can be seen as a distributed mechanism for determining acceptability while also building a shared knowledge base (the ABA framework) to determine acceptability. As discussed earlier, this shared knowledge base may or may not reflect the agents' individual knowledge bases, depending on whether they are truthful or not.

Proposition 6.1. Given a dialogue $\delta \in \mathcal{D}$, if δ is g-successful, then δ is also a-successful and i-successful.

Proposition 6.2. Given a dialogue $\delta \in \mathcal{D}$, if δ is i-successful, then δ is a-successful.

Example 6.1. Let the dialogue δ be:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(a), 1 \rangle$	$\langle a_2, a_1, 1, \text{asm}(a), 2 \rangle$
$\langle a_1, a_2, 2, \text{ctr}(a, c), 3 \rangle$	$\langle a_2, a_1, 3, \text{asm}(c), 4 \rangle$
$\langle a_1, a_2, 4, \text{ctr}(c, a), 5 \rangle$	$\langle a_2, a_1, 5, \text{asm}(a), 6 \rangle$
$\langle a_1, a_2, 0, \pi, 7 \rangle$	$\langle a_2, a_1, 0, \pi, 8 \rangle$

The ABA framework, $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$, drawn from this dialogue has:

- $\mathcal{A}_\delta = \{a, c\}$;
- $\mathcal{R}_\delta = \{\}$;
- $\mathcal{C}_\delta(a) = \{c\}, \mathcal{C}_\delta(c) = \{a\}$.

Here, argument $\{a\} \vdash a$ attacks argument $\{c\} \vdash c$ and vice versa. Then $\{a\} \vdash a$ defends itself and, hence, it is admissible. Thus a is admissible. Since $\{c\} \vdash c$ is also admissible, a is not grounded nor ideal. Since a is admissible, but not grounded, nor ideal, in \mathcal{F}_δ , δ is a-successful, but not g- or i-successful.

Example 6.2. Let the dialogue δ be:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(a), 1 \rangle$	$\langle a_2, a_1, 1, \text{asm}(a), 2 \rangle$
$\langle a_1, a_2, 2, \text{ctr}(a, c), 3 \rangle$	$\langle a_2, a_1, 3, \text{asm}(c), 4 \rangle$
$\langle a_1, a_2, 4, \text{ctr}(c, b), 5 \rangle$	$\langle a_2, a_1, 5, \text{rl}(b \leftarrow), 6 \rangle$
$\langle a_1, a_2, 0, \pi, 7 \rangle$	$\langle a_2, a_1, 0, \pi, 8 \rangle$

The ABA framework, $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$, drawn from this dialogue has:

- $\mathcal{A}_\delta = \{a, c\}$;
- $\mathcal{R}_\delta = \{b \leftarrow\}$;
- $\mathcal{C}_\delta(a) = \{c\}, \mathcal{C}_\delta(c) = \{b\}$.

Here, we have arguments $\{a\} \vdash a$, $\{c\} \vdash c$ and $\{\} \vdash b$. $\{c\} \vdash c$ attacks $\{a\} \vdash a$ and $\{\} \vdash b$ defends $\{a\} \vdash a$ by attacking $\{c\} \vdash c$. No argument attacks $\{\} \vdash b$. Hence, a is admissible, grounded and ideal in \mathcal{F}_δ , and consequently δ is a-, g- and i-successful.

Example 6.3. Let the dialogue δ be:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(a), 1 \rangle$	$\langle a_2, a_1, 1, \text{asm}(a), 2 \rangle$
$\langle a_1, a_2, 2, \text{ctr}(a, b), 3 \rangle$	$\langle a_2, a_1, 3, \text{rl}(b \leftarrow c, d), 4 \rangle$
$\langle a_1, a_2, 4, \text{asm}(c), 5 \rangle$	$\langle a_2, a_1, 4, \text{asm}(d), 6 \rangle$
$\langle a_1, a_2, 5, \text{ctr}(c, a), 7 \rangle$	$\langle a_2, a_1, 6, \text{ctr}(d, c), 8 \rangle$
$\langle a_1, a_2, 8, \text{asm}(c), 9 \rangle$	$\langle a_2, a_1, 7, \text{asm}(a), 10 \rangle$
$\langle a_1, a_2, 10, \text{ctr}(a, b), 11 \rangle$	$\langle a_2, a_1, 9, \text{ctr}(c, a), 12 \rangle$
$\langle a_1, a_2, 0, \pi, 13 \rangle$	$\langle a_2, a_1, 0, \pi, 14 \rangle$

The ABA framework drawn from this dialogue, $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$, has:

- $\mathcal{A}_\delta = \{a, c, d\}$;
- $\mathcal{R}_\delta = \{b \leftarrow c, d\}$;
- $\mathcal{C}_\delta(a) = \{b\}, \mathcal{C}_\delta(c) = \{a\}, \mathcal{C}_\delta(d) = \{c\}$.

The set of assumptions $\{a, d\}$ is the only preferred set, $\{a\}$ is admissible and ideal, and $\{\}$ is grounded. Thus, a is ideal (and admissible), but not grounded, in \mathcal{F}_δ . Hence, the dialogue δ is i-successful (and a-successful) but not g-successful.

Note that we could define a notion of *p-successful* dialogue, with preferred claim in the ABA framework drawn from the dialogue. Since every admissible set (of assumptions/arguments) is necessarily contained in a preferred set (see [11, 7]), and every preferred set is admissible by definition, trivially a dialogue is *p-successful* iff it is a-successful. Similarly, we could define a notion of *c-successful*

dialogue, with claim in the ABA framework drawn from the dialogue. However, since every preferred set is complete (see [11, 7]) and thus every admissible set is contained in a complete set, trivially a dialogue is c-successful iff it is a-successful. Therefore, we can focus on g-, a- and i-successful dialogues.

6.1. Illustration: Twelve Angry Men

The claim, *boy_innocent*, of the dialogue in Table 1 is supported by:

$$A = \{boy_not_proven_guilty\} \vdash boy_innocent,$$

attacked by:

$$B = \{w1_is_believable\} \vdash boy_proven_guilty \text{ and}$$

$$C = \{w2_is_believable\} \vdash boy_proven_guilty.$$

B and C are in turn attacked by:

$$D = \{\} \vdash w1_is_not_believable \text{ and } E = \{\} \vdash w2_is_not_believable,$$

respectively. Since D and E are supported by the empty set, they cannot be attacked. Hence, *boy_innocent* is admissible, grounded, and ideal in \mathcal{F}_δ , and this dialogue is a-, g- and i-successful.

In this section we have given non-constructive notions of successful dialogues, sanctioning the acceptability of their claims. In the remainder of the paper we show how successful dialogues can be constructed by deploying appropriate legal-move and outcome functions.

7. Related and Exhaustive Dialogues

In this section, we refine ABA dialogues, by using a specialised class of legal-move functions to avoid “disconnected” dialogues, and by using a new outcome function to avoid “incomplete” dialogues. The resulting refined ABA dialogues bring us closer to constructing successful dialogues.

The following is an example of “disconnected” dialogue:

Example 7.1. Consider dialogue δ :

a_1	a_2
$\langle a_1, a_2, 0, claim(s), 1 \rangle$	$\langle a_2, a_1, 1, rl(p \leftarrow q), 2 \rangle$
$\langle a_1, a_2, 1, asm(a), 3 \rangle$	

Here, there is no connection between contents of utterances and their targets, e.g., there is no connection between the rule $p \leftarrow q$ and the sentence s , and between the assumption a and the sentence s . The dialogue is not a-successful (and thus not g- and i-successful) as indeed s is not even supported by an argument in \mathcal{F}_δ .

Note that the second and third utterances in the dialogue in Example 7.1 are connected to their target in the sense of the following definition of related-ness:

Definition 7.1. For any two utterances $u_i, u_j \in \mathcal{U}$, $u_i \neq u_j$, u_j is *related to* u_i iff $u_i = \langle -, -, -, -, \text{ID} \rangle$, $u_j = \langle -, -, \text{ID}, -, - \rangle$.

Indeed, condition 4 in Definition 3.2 enforces that all regular utterances in a dialogue are related according to this basic, syntactic notion. The following is a more purposeful notion of related-ness:

Definition 7.2. For any two utterances $u_i, u_j \in \mathcal{U}$, $u_i = \langle -, -, -, C_i, - \rangle$, $u_j = \langle -, -, -, C_j, - \rangle$, u_j is *top-down related to* u_i iff u_j is related to u_i (as in Definition 7.1) and one of the following cases holds:

1. $C_j = rl(\rho_j)$, $Head(\rho_j) = \beta$ and either $C_i = rl(\rho_i)$ with $\beta \in Body(\rho_i)$, or $C_i = ctr(-, \beta)$, or $C_i = claim(\beta)$;
2. $C_j = asm(\alpha)$ and either $C_i = rl(\rho)$ with $\alpha \in Body(\rho)$, or $C_i = ctr(-, \alpha)$, or $C_i = claim(\alpha)$;
3. $C_j = ctr(\alpha, -)$ and $C_i = asm(\alpha)$.

Intuitively, an utterance is top-down related to another if it contributes to expanding an argument (case 1), identifies an assumption in the support of an argument (case 2) or starts the construction of a counter-argument (case 3). Note that an utterance may be top-down related to an utterance from the same agent or not. Also, no pass utterance can be top-down related to an utterance and vice versa.

We use the notion of top-down related-ness to define a new class of legal-move functions and corresponding connected dialogues:

Definition 7.3. A *top-down related legal-move function* $\lambda \in \Lambda$ is such that for $\delta = \langle u_1, \dots, u_n \rangle$, $\delta \in \mathcal{D}$, and for all $u \in \lambda(\delta)$ such that u is a regular utterance, then u is top-down related to some u_m in δ ($1 \leq m \leq n$). A *top-down related dialogue* is a λ -ABA dialogue for a top-down related $\lambda \in \Lambda$.

For simplicity, in the remainder, we use “related” to mean “top-down related”. The dialogue in Example 3.2 is related; the dialogue in Example 7.1 is not. Note that a dialogue can be compatible with a legal-move function that fulfils several purposes, e.g., being flat and related, as is the case in Example 3.2. In the remainder, all dialogues will be related.

The following is an example of an “incomplete” dialogue:

Example 7.2. The following dialogue δ :

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 1, \text{rl}(s \leftarrow a, b), 2 \rangle$
$\langle a_1, a_2, 2, \text{rl}(a \leftarrow c), 3 \rangle$	$\langle a_2, a_1, 3, \text{rl}(c \leftarrow), 4 \rangle$
$\langle a_1, a_2, 2, \text{rl}(b \leftarrow c), 5 \rangle$	

is “incomplete”, as $\langle a_2, a_1, 5, \text{rl}(c \leftarrow), 6 \rangle$ can be uttered as $\text{rl}(c \leftarrow)$ is in \mathcal{F}_δ (since it is the content of utterance with identifier 4).

The next notion we introduce prevents this form of “incompleteness” and gives a sense of *exhaustiveness*. Exhaustiveness is introduced to ensure a dialogue is “complete” in the sense that there are no “unsaid” utterances in such dialogues that would explicitly bring about important arguments for determining success. This feature will ease the proof of soundness results later.

Definition 7.4. The *exhaustive outcome function* ω_{ex} is such that, given $\delta \in \mathcal{D}$ and $\lambda \in \Lambda$, $\omega_{ex}(\delta, \lambda) = \text{true}$ iff $\omega_{ABA}(\delta, \lambda) = \text{true}$ and, given $\langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ the ABA framework drawn from δ , there exists no $u' \in \lambda(\delta)$ with content either:

- $asm(\alpha)$, for $\alpha \in \mathcal{A}_\delta$, or
- $\text{rl}(\rho)$, for $\rho \in \mathcal{R}_\delta$, or
- $\text{ctr}(\alpha, \beta)$, for $\beta \in \mathcal{C}_\delta(\alpha)$,

such that $\omega_{ABA}(\delta \circ u', \lambda) = \text{true}$.

A λ -*exhaustive dialogue* is a λ -ABA dialogue δ such that $\omega_{ex}(\delta, \lambda) = \text{true}$ and λ is related. An *exhaustive dialogue* is a λ -exhaustive dialogue for some $\lambda \in \Lambda$.

Note that exhaustiveness does not force agents to contribute to dialogues all relevant information they hold. Rather, it enforces that if an utterance u with content equal to an assumption, a rule or contrary has been made compatibly with a certain λ , then another utterance u' with the same content as u should be made in δ , if such u' is allowed by λ . The dialogue in Example 7.2 is not exhaustive, whereas the dialogue in Example 3.2 is.

7.1. Illustration: Twelve Angry Men

The dialogue δ shown in Table 1 is related and λ -exhaustive, for λ related. Indeed, in addition to being an ABA dialogue (as we have seen in Section 5.1), δ is compatible with a (top-down) related legal-move function λ as all regular utterances are related; and it is λ -exhaustive as there is no utterance u' related to any existing utterance with content an assumption, rule or contrary already in $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ (see Section 5.1) that can be added to δ .

8. Debate Trees

In this section we define *debate trees* to keep track of information that has been disclosed in dialogues. These are then used to define new legal-move functions and outcome functions guaranteeing the construction of successful dialogues (see later in Section 10). The new legal-move functions allow to decide utterances on the basis of debate trees drawn from dialogues. The new outcome functions determine the outcome of dialogues on the basis of conditions satisfied by the corresponding debate trees. Informally (see Definition 8.1 below for the formal definition), each node of debate trees

1. contains one sentence,
2. is *tagged* as either *unmarked* (*um*), *marked-rule* (*mr*) or *marked-assumption* (*ma*),
3. is *labelled* either *proponent* (*P*) or *opponent* (*O*) as in the abstract dispute trees of [14], and
4. has an ID taken from an utterance in the dialogue.

Intuitively, the sentence of each node in a debate tree represents an argument's claim or element of the argument's support. A node is tagged *unmarked* if its sentence is only mentioned in the claim or the body of a rule or contrary of an assumption, but without any further examination, *marked-rule* if its sentence is the head of an uttered rule, and *marked-assumption* if its sentence has been explicitly uttered as an assumption. A node is labelled *P* (*O*) if it is (directly or indirectly) for (against, respectively) the claim of the dialogue. The ID is used to identify the node's corresponding utterance in the dialogue.

As an illustration, possible nodes from utterances in Example 3.2 are:

$$\begin{array}{ll} (s, um : P[1]), & (s, mr : P[1]), \\ (a, um : P[3]), & (a, ma : P[4]), \\ (q, um : O[5]). & \end{array}$$

Note that, since debate trees can be viewed as constructed in steps as dialogues proceed, not all nodes in the above list are in a debate tree at every step. Namely, certain nodes tagged as *um*, such as $(s, um : P[1])$ and $(q, um : O[5])$ are replaced by nodes with “updated” information about the same sentences, i.e., $(s, mr : P[1])$ and $(q, mr : O[5])$, respectively, after new utterances are inserted into the dialogue. For instance, the content of utterance 5 in the dialogue in Example 3.2 is $ctr(a, q)$, so when this utterance occurs, q is only mentioned as a contrary of a , it is uncertain whether q is an assumption or the head of a rule, hence the tag in



Figure 3: Nodes in a debate tree connected as in arguments (left) or attacks between them (right).

$(q, um : O[5])$ is um . However, after utterance 6, which has content $rl(q \leftarrow b)$, it is known that q is the head of the rule $q \leftarrow b$, hence $(q, um : O[5])$ is replaced by $(q, mr : O[5])$. Thus nodes tagged um may be replaced by other nodes during the construction of a debate tree. Also, when “updating” a node, its ID is not always changed, as seen in the example, where $(s, um : P[1])$ and $(s, mr : P[1])$ have the same ID (1). This is because that ID is used for book keeping purposes to decide where to insert new nodes. The rules for updating IDs are given later.

In a debate tree, nodes are connected in two cases:

1. they belong to the same argument, and
2. they form attacks between two arguments.

Figure 3 (left) gives an example of the first case, for the dialogue in Example 3.2. Here, we see the connected nodes in a debate tree on the far left and, next to them, the corresponding argument (represented as a tree). In the second case, two nodes $N_1 = (\alpha, ma : L[id])$ and $N_2 = (\beta, - : L'[-])$, $L, L' \in \{P, O\}, L' \neq L$, are connected if there is an utterance $u = \langle -, -, id, ctr(\alpha, \beta), - \rangle$ in the dialogue (e.g., see Figure 3, right). Hence, the two nodes are connected if the parent node contains an assumption and the child node contains a contrary of that assumption.

We give the formal definition of a debate tree as follows.

Definition 8.1. Given a dialogue $\mathcal{D}_{a_j}^{\alpha_i}(\chi) = \delta$, the *debate tree drawn from δ* is a tree $\mathcal{T}(\delta)$, where:

1. nodes of $\mathcal{T}(\delta)$ can be characterised as follows:
 - (a) nodes are tuples $(S, F : L[U])$ where
 - S is a sentence in \mathcal{L} ,
 - F (the tag) is either um (*unmarked*), mr (*marked-rule*) or ma (*marked-assumption*),
 - L (the label) is either P (*proponent*) or O (*opponent*),
 - $U \in \mathcal{ID}$ (the ID);
 - (b) $(\beta, - : -[-])$ is a node in $\mathcal{T}(\delta)$ iff there is an utterance $\langle -, -, -, C, - \rangle$ in δ such that C is either:

- i. $claim(\beta)$, or
 - ii. $rl(s_0 \leftarrow s_1, \dots, s_m)$ where $\beta = s_i$ for some $i, 0 < i \leq m$, or
 - iii. $asm(\beta)$, or
 - iv. $ctr(-, \beta)$;
- (c) $(\beta, um : _ [t])$ is a node in $\mathcal{T}(\delta)$ iff there is no utterance $\langle -, -, t, rl(\beta \leftarrow \dots), - \rangle$ or $\langle -, -, t, asm(\beta), - \rangle$ in δ ;
2. edges of $\mathcal{T}(\delta)$ are of 6 types (below $L, L' \in \{\mathbf{P}, \mathbf{O}\}, L \neq L'$):
- (a) for $N = (s_0, mr : L[t])$ and $N' = (s_i, um : L[id])$, N is the parent of N' iff there is one utterance in δ of the form $\langle -, -, t, rl(s_0 \leftarrow \dots, s_i, \dots), id \rangle$ and there is no utterance in δ of the forms $\langle -, -, id, rl(s_i \leftarrow \dots), - \rangle$ or $\langle -, -, id, asm(s_i), - \rangle$;
 - (b) for $N = (s_0, mr : L[t])$ and $N' = (s_i, mr : L[id])$, N is the parent of N' iff there are two utterances in δ of the forms $\langle -, -, t, rl(s_0 \leftarrow \dots, s_i, \dots), id \rangle$ and $\langle -, -, id, rl(s_i \leftarrow \dots), - \rangle$, respectively;
 - (c) for $N = (s_0, mr : L[t])$ and $N' = (s_i, ma : L[id])$, N is the parent of N' iff there are two utterances in δ of the forms $\langle -, -, t, rl(s_0 \leftarrow \dots, s_i, \dots), t' \rangle$ and $\langle -, -, t', asm(s_i), id \rangle$, respectively;
 - (d) for $N = (s_0, ma : L[t])$ and $N' = (s_i, um : L'[id])$, N is the parent of N' iff there is one utterance in δ of the form $\langle -, -, t, ctr(s_0, s_i), id \rangle$, and there is no utterance in δ of the forms $\langle -, -, id, rl(s_i \leftarrow \dots), - \rangle$ or $\langle -, -, id, asm(s_i), - \rangle$;
 - (e) for $N = (s_0, ma : L[t])$ and $N' = (s_i, mr : L'[id])$, N is the parent of N' iff there are two utterances in δ of the forms $\langle -, -, t, ctr(s_0, s_i), id \rangle$ and $\langle -, -, id, rl(s_i \leftarrow \dots), - \rangle$, respectively;
 - (f) for $N = (s_0, ma : L[t])$ and $N' = (s_i, ma : L'[id])$, N is the parent of N' iff there are two utterances in δ of the forms $\langle -, -, t, ctr(s_0, s_i), t' \rangle$ and $\langle -, -, t', asm(s_i), id \rangle$, respectively.

We say that a dialogue δ *draws the debate tree* $\mathcal{T}(\delta)$.

Notation 8.1. For convenience, we call $(-, mr : _ [-])$, $(-, ma : _ [-])$, and $(-, um : _ [-])$, a *rule*, *assumption*, and *unmarked node*, respectively.

The debate tree $\mathcal{T}(\delta)$ for δ in Example 3.2 is shown in Figure 4.

Definition 8.1 gives the characteristics of debate trees. Debate trees are constructed as dialogues progress. The construction of debate trees is given below.




Figure 4: The debate tree drawn from the dialogue in Example 3.2.

Lemma 8.1. Given a dialogue $\mathcal{D}_{a_j}^{\alpha_i}(\chi) = \delta$, the debate tree $\mathcal{T}(\delta)$ is $\mathcal{T}^m(\delta)$ in the sequence $\mathcal{T}^0(\delta), \mathcal{T}^1(\delta), \dots, \mathcal{T}^m(\delta)$ constructed inductively from the π -pruned sequence $\delta' = \langle u'_1 \dots, u'_m \rangle$ obtained from δ , as follows ($L, L' \in \{P, O\}, L \neq L'$):

1. $\mathcal{T}^0(\delta)$ is empty;
2. $\mathcal{T}^1(\delta)$ contains a single node:

$$(\chi, um : P[id_1]),$$
 where id_1 is the identifier of $u'_1 = u_1$;
3. let $\mathcal{T}^i(\delta)$ be the i -th tree, $0 < i < m$, let $u'_{i+1} = \langle -, -, t, C, id \rangle$ and $\langle -, -, -, C_t, t \rangle$ be the target utterance of u'_{i+1} ; then $\mathcal{T}^{i+1}(\delta)$ is obtained as follows:
 - (a) if $C = rl(\beta_0 \leftarrow \beta_1, \dots, \beta_l)$ then $\mathcal{T}^{i+1}(\delta)$ is $\mathcal{T}^i(\delta)$ with additional nodes:

$$(\beta_1, um : L[id]), \dots, (\beta_l, um : L[id])$$
 as children of the node $(\beta_0, - : L[t])$, and the node $(\beta_0, - : L[t])$ is replaced by $(\beta_0, mr : L[t])$;
 - (b) if $C = asm(\alpha)$ then $\mathcal{T}^{i+1}(\delta)$ is $\mathcal{T}^i(\delta)$ with the node $(\alpha, um : L[t])$ replaced by $(\alpha, ma : L[id])$;
 - (c) if $C = ctr(\alpha, \beta)$ then $\mathcal{T}^{i+1}(\delta)$ is $\mathcal{T}^i(\delta)$ with an additional node: $(\beta, um : L[id])$ child of $(\alpha, ma : L[t])$.

The construction of the debate tree in Figure 4 is shown in Figure 5.

Arguments can be drawn from a debate tree, as follows:

Definition 8.2. An argument drawn from a debate tree $\mathcal{T}(\delta)$ is a sub-tree \mathcal{T} of $\mathcal{T}(\delta)$ such that:

1. all nodes in \mathcal{T} have the same label (either P or O);

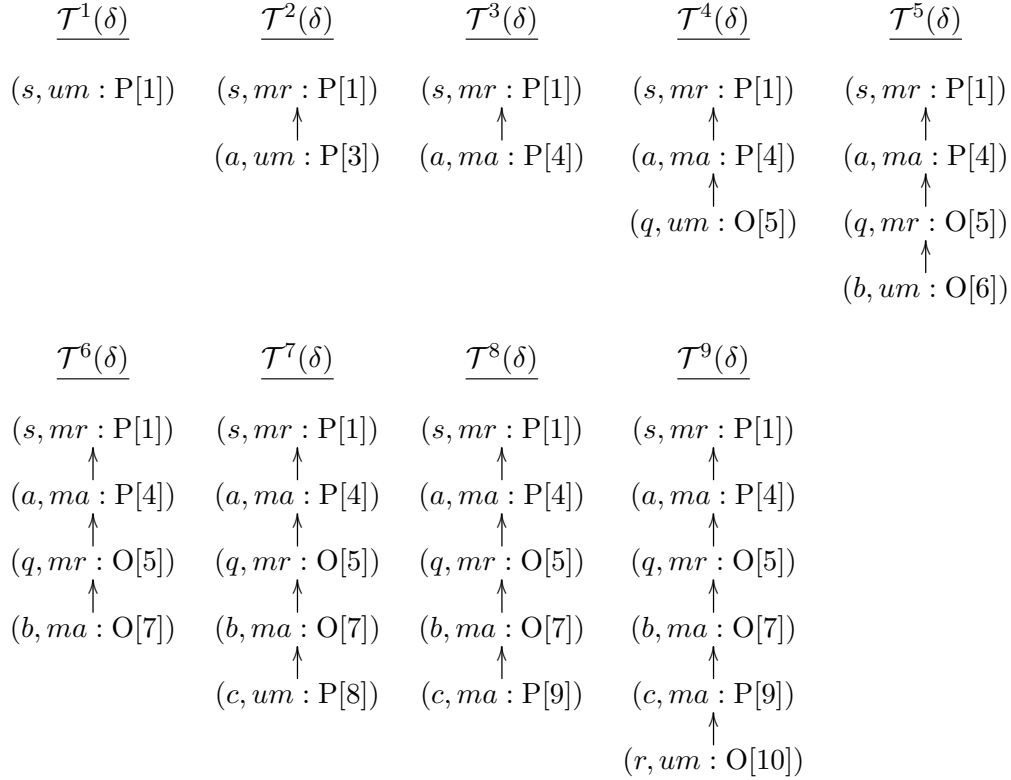


Figure 5: Construction of the debate tree in Figure 4 (using utterances u'_1, \dots, u'_9 in Example 3.3).

2. if there is an utterance $\langle -, -, -, rl(\beta_0 \leftarrow \beta_1, \dots, \beta_m), t \rangle$ in δ and $(\beta_0, mr : L[t])$ is in \mathcal{T} , then $(\beta_1, - : L[-]), \dots, (\beta_m, - : L[-])$ are in \mathcal{T} ;
3. there does not exist a node N' in $\mathcal{T}(\delta)$ such that N' is parent or child of some node N_i in \mathcal{T} , N' is not in \mathcal{T} and N_i, N' have the same label.

The sentence χ in the root of an argument \mathcal{T} drawn from a debate tree is the *claim* of \mathcal{T} . An argument \mathcal{T} drawn from a debate tree is *actual* if for all nodes $(-, F : -[-])$ in \mathcal{T} , F is either mr or ma ; if there is at least one node of the form $(-, um : -[-])$ in \mathcal{T} , then \mathcal{T} is *potential*. An argument drawn from a debate tree is a *proponent (opponent) argument* if its nodes are labelled P (O, respectively).

Example 8.1. Given the debate tree in Figure 4, we can draw three actual arguments and one potential argument, as shown in Figure 6.



Figure 6: The actual arguments (the first three starting from left) and potential argument (right most) drawn from the debate tree in Figure 4.

Note that, in debate trees, as in abstract dispute trees [14], proponent/opponent are roles/fictitious players rather than actual agents and agents can play separate or both such roles in a dialogue. For example, in the case of inquiry dialogues [43], agents may want to play both roles interchangeably, whereas in persuasion dialogues [43] the agent putting forward the claim may play solely the role of proponent and the other agent the role of opponent.

Notation 8.2. We refer to an actual or potential argument drawn from a debate tree simply as *actual* or *potential argument* (respectively).

We refer to arguments in ABA (see Section 2) as *ABA arguments*, to distinguish them from arguments drawn from debate trees.

Given a debate tree $\mathcal{T}(\delta)$, we say that *a node is in an argument* (in $\mathcal{T}(\delta)$) iff it is a node in some argument drawn from $\mathcal{T}(\delta)$.

If \mathcal{T} is a potential argument and χ is its claim, \mathcal{T} is written as $S_a, S_r \vdash^t \chi$, where $S_a = \{\alpha | (\alpha, ma : -[-]) \text{ is a node in } \mathcal{T}\}$ and $S_r = \{\beta | (\beta, um : -[-]) \text{ is a node in } \mathcal{T}\}$. If \mathcal{T} is an actual argument and χ is its claim, \mathcal{T} is written as $S \vdash^t \chi$, where $S = \{\alpha | (\alpha, ma : -[-]) \text{ is a node in } \mathcal{T}\}$.

Definition 8.2 gives a means of talking about arguments in the context of debate trees. The following lemma sanctions that actual arguments can be mapped to equivalent ABA arguments.

Lemma 8.2. Given a dialogue $\delta \in \mathcal{D}$, for each actual argument $S \vdash^t \beta$ drawn from the debate tree $\mathcal{T}(\delta)$, there exists an ABA argument $S \vdash \beta$ in the ABA framework drawn from δ .

The ABA arguments corresponding to actual arguments drawn from the debate tree drawn from the dialogue in Example 3.2 (see Figures 4, 6) are $\{a\} \vdash s$, $\{b\} \vdash q$, and $\{c\} \vdash c$.

For exhaustive dialogues, the other direction of Lemma 8.2 holds as well:

Lemma 8.3. Given an exhaustive dialogue $\delta \in \mathcal{D}$, for each ABA argument $S \vdash \beta$ in the ABA framework drawn from δ , $S \vdash^t \beta$ can be drawn from $\mathcal{T}(\delta)$.

Lemma 8.3 does not hold if a dialogue is not exhaustive, as shown next.

Example 8.2. Let a dialogue $\delta \in \mathcal{D}$ be:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(a), 1 \rangle$	$\langle a_2, a_1, 1, \text{rl}(a \leftarrow b, c), 2 \rangle$
$\langle a_1, a_2, 2, \text{rl}(b \leftarrow), 3 \rangle$	$\langle a_2, a_1, 2, \text{rl}(c \leftarrow b), 4 \rangle$

The ABA framework \mathcal{F}_δ drawn from δ has:

- $\mathcal{R}_\delta = \{a \leftarrow b, c, b \leftarrow, c \leftarrow b\}$,
- $\mathcal{A}_\delta = \{\}$.

Clearly, $\{\} \vdash a$ is an argument in \mathcal{F}_δ . The debate tree $\mathcal{T}(\delta)$ drawn from δ is in Figure 7. We can see that no actual argument can be drawn from this tree, as one leaf node, $(b, um : P[4])$, of this debate tree is an unmarked node.




Figure 7: The debate tree drawn from the dialogue in Example 8.2.

As in [14], we introduce the defence set and the culprits of a debate tree. We use these two concepts for our soundness result in the next sections.

Definition 8.3. Given a debate tree $\mathcal{T}(\delta)$,

- the *defence set* $\mathcal{DEF}(\mathcal{T}(\delta))$ is the union of all assumptions α in proponent nodes of the form $N = (\alpha, ma : P[-])$, such that N is in an actual argument;
- the *culprits* $\mathcal{CUL}(\mathcal{T}(\delta))$ is the set of all assumptions α in opponent nodes $N = (\alpha, ma : O[-])$ such that the child of N in $\mathcal{T}(\delta)$ is $N' = (-, - : P[-])$ and both N and N' are in actual arguments.

8.1. Illustration: Twelve Angry Men

The debate tree drawn from the dialogue δ shown in Figure 1 is $\mathcal{T}(\delta)$ shown in Figure 8. Here, the defence set containing the only assumption in proponent nodes is $\{\text{boy_not_proven_guilty}\}$ and the culprits set containing the two assumptions in opponent nodes is $\{w1_is_believable, w2_is_believable\}$.




Figure 8: The debate tree drawn from the dialogue in Table 1.

9. Legal-Move and Outcome Functions from Debate Trees

In this section we use debate trees to provide new, refined legal-move and outcome functions. The definition thereof requires agents to consult the current debate tree before making utterances or deciding whether the dialogue has achieved the desired outcome. In particular, when an agent decides what to utter, it needs to take the current debate tree into account and make sure that its new utterance will keep the tree fulfilling desired properties. Thus, the debate tree drawn from a dialogue can be seen as a *commitment store* [43] holding information that agents disclose and share using the dialogue.

The new notions of legal-move and outcome functions will then be used, in Sections 10 and 11, to generate provably successful dialogues, in the sense of Section 6. Note that the first such notion (of patient legal-move function, see Section 9.1) is not necessary to guarantee successful dialogues, but it is intuitive and considerably simplifies the proof of results.

9.1. Patient Legal-Move Functions

During dialogues, agents may in general choose to start attacking arguments while these arguments are still under construction. In patient dialogues, defined below, this is not allowed, as, in this type of dialogues, arguments are fully expanded (cf. actual) before being attacked.

Definition 9.1. A debate tree $\mathcal{T}(\delta)$ is *patient* iff for all nodes $N = (-, ma : -[-])$ in $\mathcal{T}(\delta)$ such that N has a child, then N is in an actual argument drawn from $\mathcal{T}(\delta)$.

A legal-move function $\lambda \in \Lambda$ is *patient* iff for every $\delta \in \mathcal{D}$ such that $\mathcal{T}(\delta)$ is patient, for every $u \in \lambda(\delta)$, $\mathcal{T}(\delta \circ u)$ is still patient.

The debate tree in Figure 4 is patient, whereas the one in the following example is not.

Example 9.1. Let dialogue δ be as follows:

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 1, \text{rl}(s \leftarrow a, b), 2 \rangle$
$\langle a_1, a_2, 2, \text{asm}(a), 3 \rangle$	$\langle a_2, a_1, 3, \text{ctr}(a, c), 4 \rangle$

$\mathcal{T}(\delta)$, given in Figure 9, is not patient as the argument $\{a\}, \{b\} \vdash^t s$ is not actual, since its node $(b, um : P[2])$ has the tag um , yet the assumption node $(a, ma : P[3])$ already has a child $(c, um : O[4])$.




Figure 9: The non-patient debate tree drawn from the dialogue in Example 9.1.

In the remainder, dialogues are compatible with a patient legal-move function.

9.2. Last-Word Outcome Function and Defensive Dialogues

We refine the notion of exhaustive outcome function guaranteeing that the (fictitious) proponent has the last-word in the dialogue, namely all leaves in the debate tree are proponent nodes or “dead-end” opponent nodes (not corresponding to any un-attacked actual arguments). We first define the notion of a node being *properly attacked* in a debate tree:

Definition 9.2. Given a debate tree $\mathcal{T}(\delta)$, a node N in $\mathcal{T}(\delta)$ is *properly attacked* iff N is of the form $(-, ma : -[-])$ and N has a child N' that is in an actual argument.

Thus, a node is properly attacked iff it is an assumption node and there is an actual argument attacking the assumption in the node. Intuitively, this notion allows to distinguish assumptions against which actual attacks have been constructed (and thus ‘properly attacked’) from assumptions currently under attack

but for which no actual attack has yet been finalised. In what we call ‘last-word’ dialogues, the debate tree is such that all its opponent arguments have an assumption node which is properly attacked:

Definition 9.3. A debate tree $\mathcal{T}(\delta)$ is *last-word* iff

1. for all leaf nodes N in $\mathcal{T}(\delta)$, N is either $(-, mr : P[-])$ or $(-, ma : P[-])$, and
2. if a leaf node N is of the form $(-, - : O[-])$, then either
 - (a) N is in a potential argument, or
 - (b) N is in an actual argument that contains one node N' of the form $(\alpha, ma : O[-])$ such that there is another node N'' in $\mathcal{T}(\delta)$ of the form $(\alpha, ma : O[-])$, $N'' \neq N$, and N'' is properly attacked.

The *last-word outcome function* ω_{lw} is such that, given $\delta \in \mathcal{D}$ and $\lambda \in \Lambda$, $\omega_{lw}(\delta, \lambda) = true$ iff $\omega_{ex}(\delta, \lambda) = true$ and $\mathcal{T}(\delta)$ is last-word.

A *λ -defensive dialogue* is a λ -exhaustive dialogue $\delta \in \mathcal{D}$ such that $\omega_{lw}(\delta, \lambda) = true$. A *defensive dialogue* is a λ -defensive dialogue for some $\lambda \in \Lambda$.

Intuitively, the last-word outcome function specifies a winning condition for the (fictitious) proponent: either the proponent finishes the dialogue with rules and un-attacked assumptions (condition 1), or the (fictitious) opponent does not pose any valid attacks via actual argument (condition 2a), or any valid attacks posed by the opponent has been answered with valid counter attacks (condition 2b).

The dialogue in Example 3.2 is defensive, as the debate tree in Figure 4 shows. The next example shows a dialogue that is not.

Example 9.2. The following dialogue δ

a_1	a_2
$\langle a_1, a_2, 0, claim(s), 1 \rangle$	$\langle a_2, a_1, 1, rl(s \leftarrow a, b), 2 \rangle$
$\langle a_1, a_2, 0, \pi, 3 \rangle$	$\langle a_2, a_1, 0, \pi, 4 \rangle$

is not defensive. Indeed, the debate tree $\mathcal{T}(\delta)$ (see Figure 10) is not last-word, as there are two unmarked leaf node in $\mathcal{T}(\delta)$: $(a, um : P[2])$ and $(b, um : P[2])$.

9.3. Conflict-Free Outcome Function and Conflict-Free Dialogues

As in the case of abstract dispute trees, the defence set of a debate tree may attack itself. We refine the notion of outcome function to avoid this, as follows:




Figure 10: A non-defensive dialogue in Example 9.2.

Definition 9.4. The *conflict-free outcome function* ω_{cf} is such that, given $\delta \in \mathcal{D}$ and a legal-move function λ , $\omega_{cf}(\delta, \lambda) = true$ iff $\omega_{lw}(\delta, \lambda) = true$ and $\mathcal{DEF}(\mathcal{T}(\delta)) \cap \mathcal{CUL}(\mathcal{T}(\delta)) = \{\}$.

A λ -*conflict-free dialogue* is a λ -defensive dialogue $\delta \in \mathcal{D}$ such that $\omega_{cf}(\delta, \lambda) = true$. A *conflict-free dialogue* is a λ -conflict-free dialogue for some $\lambda \in \Lambda$.

Lemma 9.1. Let δ be a conflict-free dialogue and let \mathcal{F}_δ be the ABA framework drawn from δ . Then, $\mathcal{DEF}(\mathcal{T}(\delta))$ does not attack itself in \mathcal{F}_δ .

9.4. Filtered Legal-Move Functions

Another legal-move function that relies on the debate tree as a commitment store is the filtered legal-move function, defined in terms of filtered debate trees. These are trees where the same assumptions are not attacked repeatedly, unnecessarily. Instead, in a filtered debate tree, if an assumption has already been marked, then it does not need to be “dealt with” again. Formally:

Definition 9.5. A debate tree $\mathcal{T}(\delta)$ is *filtered* iff for any two nodes $\mathcal{T}(\delta)$: $N_1 = (\beta, ma : L[id_1])$, $N_2 = (\beta, ma : L[id_2])$, $L \in \{\mathbf{P}, \mathbf{O}\}$, $N_1 \neq N_2$, if N_1 has a child N'_1 in an actual argument \mathcal{T}_1 and N_2 has a child N'_2 in an actual argument \mathcal{T}_2 , then $\mathcal{T}_1 \neq \mathcal{T}_2$.⁸ A legal-move function $\lambda \in \Lambda$ is *filtered* iff for every $\delta \in \mathcal{D}$ such that $\mathcal{T}(\delta)$ is filtered, for every $u \in \lambda(\delta)$, $\mathcal{T}(\delta \circ u)$ is still filtered.

Example 9.3. The following dialogue (from which the debate tree in Figure 11 (left) is drawn) is compatible with a filtered legal-move function:

a_1	a_2
$\langle a_1, a_2, 0, claim(a), 1 \rangle$	$\langle a_2, a_1, 1, asm(a), 2 \rangle$
$\langle a_1, a_2, 2, ctr(a, b), 3 \rangle$	$\langle a_2, a_1, 3, asm(b), 4 \rangle$
$\langle a_1, a_2, 4, ctr(b, a), 5 \rangle$	$\langle a_2, a_1, 5, asm(a), 6 \rangle$
$\langle a_1, a_2, 6, ctr(a, b), 7 \rangle$	$\langle a_2, a_1, 0, \pi, 8 \rangle$

⁸Note that, when comparing two arguments for equality, we only consider the sentence and the tag in each argument, while ignoring the label and the ID.



Figure 11: Filtered (left) and non-filtered (right) debate trees for Example 9.3.

Here the argument $\{a\} \vdash a$ is attacked by $\{b\} \vdash b$ once and only once, and, in the filtered debate tree (Figure 11 left), since the node $(a, ma : P[2])$ has a child $(b, ma : O[4])$, the node $(a, ma : P[6])$ cannot be attacked by the same argument $\{b\} \vdash b$, hence $(b, um : O[7])$ is an *unmarked* node, rather than a marked assumption node. Instead, the dialogue below

a_1	a_2
$\langle a_1, a_2, 0, claim(a), 1 \rangle$	$\langle a_2, a_1, 1, asm(a), 2 \rangle$
$\langle a_1, a_2, 2, ctr(a, b), 3 \rangle$	$\langle a_2, a_1, 3, asm(b), 4 \rangle$
$\langle a_1, a_2, 4, ctr(b, a), 5 \rangle$	$\langle a_2, a_1, 5, asm(a), 6 \rangle$
$\langle a_1, a_2, 6, ctr(a, b), 7 \rangle$	$\langle a_2, a_1, 7, asm(b), 8 \rangle$

(from which the debate tree in Figure 11 (right) is drawn) is not compatible with a filtered legal-move function. Here, the argument $\{a\} \vdash a$ is attacked by $\{b\} \vdash b$ twice.

Thus, filtered legal-move functions bring efficiency to dialogues, as, in dialogues that are compatible with a filtered legal-move function, any assumption is attacked at most once by the same argument.

9.5. Illustration: Twelve Angry Men

The debate tree $\mathcal{T}(\delta)$ given in Figure 8 is patient and filtered. Hence, δ in Table 1 is compatible with a patient and filtered (as well as flat and related) legal-move function λ . The two leaf nodes in Figure 8 are:

$(w1_contradicted_by_w2, mr : P[8])$ and $(w2_has_poor_eyesight, mr : P[13])$,

both are proponent nodes with marked rules, hence of the form $(-, mr : P[-])$. By Definition 9.3, $\mathcal{T}(\delta)$ is last-word and $\omega_{lw}(\delta, \lambda) = true$. Moreover, $\mathcal{DEF}(\mathcal{T}(\delta)) =$

$\{boy_not_proven_guilty\}$ and $\mathcal{CUL}(\mathcal{T}(\delta)) = \{w1_is_believable, w2_is_believable\}$. Thus, by Definition 9.4, $\omega_{cf}(\delta, \lambda) = true$ and δ is λ -conflict-free.

Overall, in this section we have introduced the basic ingredients to formally link dialogues with argumentation semantics and show how successful dialogues can be constructed. In the next sections, we will do so for different categories of dialogues, that vary in how broadly or narrowly they debate the given claim.

10. Focused Dialogues

In this section we link dialogues and argumentation semantics for dialogues that *focus* on constructing a single abstract dispute tree for the claim, where the (fictitious) proponent only needs to build single (supporting and defending) arguments, whereas the (fictitious) opponent may build several (attacking) arguments.

Definition 10.1. A debate tree $\mathcal{T}(\delta)$ is *focused* iff

1. for all arguments A drawn from $\mathcal{T}(\delta)$, if A contains a node $(-, ma : O[-])$, then there is at most one node N of the form $(-, ma : O[-])$ in A such that N has any child, and such node N has a single child;
2. for all nodes of the form $(\beta_0, mr : P[t])$ with children $(\beta_1, - : P[-])$, \dots , $(\beta_n, - : P[-])$ there must be an utterance in δ of the form $\langle -, -, t, rl(\beta_0 \leftarrow \beta_1, \dots, \beta_n), - \rangle$.

A legal-move function $\lambda \in \Lambda$ is *focused* iff for every $\delta \in \mathcal{D}$ such that $\mathcal{T}(\delta)$ is focused, for every $u \in \lambda(\delta)$, $\mathcal{T}(\delta \circ u)$ is still focused. A *focused dialogue* is a dialogue compatible with a focused legal-move function.

In focused trees, no alternative ways to support or defend the claim are considered simultaneously, e.g. an opponent argument is only attacked by a single proponent argument, whereas a proponent argument can be attacked in as many ways as the number of its assumptions and ways to prove their contraries. Moreover, the claim is supported by a single proponent argument. The tree in Figure 4 is focused. Figures 12 and 13 show two non-focused debate trees drawn from the dialogues in Tables 3 and 4, respectively, and violating condition 1 and 2 in Definition 10.1, respectively. In these (unfocused) dialogues the agents stick to the original claim but they shift the way they carry out the debate. For example, in the dialogue in Table 3 two different ways to defeat an opponent argument are explored concurrently (by building arguments against a and b), although one suffices to determine acceptability of the claim.

Table 3: A dialogue that draws a non-focused debate tree.

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(\chi), 1 \rangle$	$\langle a_2, a_1, 1, \text{asm}(\chi), 2 \rangle$
$\langle a_1, a_2, 2, \text{ctr}(\chi, s), 3 \rangle$	$\langle a_2, a_1, 3, \text{rl}(s \leftarrow a, b), 4 \rangle$
$\langle a_1, a_2, 4, \text{asm}(a), 5 \rangle$	$\langle a_2, a_1, 4, \text{asm}(b), 6 \rangle$
$\langle a_1, a_2, 5, \text{ctr}(a, c), 7 \rangle$	$\langle a_2, a_1, 6, \text{ctr}(b, c), 8 \rangle$

Table 4: Another dialogue that draws a non-focused debate tree.

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 1, \text{rl}(s \leftarrow a), 2 \rangle$
$\langle a_1, a_2, 1, \text{rl}(s \leftarrow b), 3 \rangle$	

λ -defensive dialogues for λ focused give debate trees corresponding to *abstract dispute trees* [14] with the same defence set and culprits. Formally:

Lemma 10.1. Given a λ -defensive dialogue $\delta \in \mathcal{D}$ for χ , for λ -focused, let $\mathcal{T}(\delta)$ be the debate tree drawn from δ and $\mathcal{F}_\delta = \langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$ be the ABA framework drawn from δ . Then there is an abstract dispute tree \mathcal{T}^a for $S \vdash \chi$ for some $S \subseteq \mathcal{A}_\delta$, such that $\mathcal{DEF}(\mathcal{T}(\delta)) = \mathcal{DEF}(\mathcal{T}^a)$ and $\mathcal{CUL}(\mathcal{T}(\delta)) = \mathcal{CUL}(\mathcal{T}^a)$.

Lemma 10.1 shows the connection between the debate trees drawn from λ -defensive for λ -related and focused dialogues and abstract dispute trees. An analogous result can be obtained for λ -defensive and filtered dialogues. Indeed, since filtered legal-move functions perform a form of “filtering” resulting in pruned trees, the result relies upon inserting nodes that have been filtered back into the debate tree, resulting in *expanded debate trees*, defined as follows:

Definition 10.2. Given a debate tree $\mathcal{T}(\delta)$ where δ is λ -defensive for λ -focused, we construct a (possibly infinite) sequence $\mathcal{T}_0^E, \dots, \mathcal{T}_n^E, \dots$ of trees as follows:

- Delete all nodes N from $\mathcal{T}(\delta)$ where N is in a potential argument. Let $\mathcal{T}'(\delta)$ be the result.
- $\mathcal{T}_0^E = \mathcal{T}'(\delta)$.




Figure 12: A non-focused tree (drawn from the dialogue in Table 3). Here, the opponent argument $\{a, b\} \vdash^t s$ has two proponent nodes as its children: $(c, um : P[7])$ and $(c, um : P[8])$.




Figure 13: A non-focused tree (drawn from the dialogue in Table 4). Here, a proponent node (the root) has two proponent children from two different utterances.

- Suppose \mathcal{T}_i^E , for $i \geq 0$, has been constructed; then \mathcal{T}_{i+1}^E is obtained by adding arguments $\mathcal{T}_1, \dots, \mathcal{T}_k$ simultaneously to leaf nodes N_1, \dots, N_k of \mathcal{T}_i^E , as children, respectively, where \mathcal{T}_j and N_j , $1 \leq j \leq k$ are such that:
 1. N_j is of the form $(\alpha, ma : L_0[-])$ ($L_0 \in \{\mathbf{P}, \mathbf{O}\}$),
 2. there is N_j^0 of the form $(\alpha, ma : -[-])$ in $\mathcal{T}'(\delta)$, $N_j^0 \neq N_j$,
 3. N_j^0 has a child N_j' in $\mathcal{T}'(\delta)$, and N_j' is in an actual argument \mathcal{T}_j'' in $\mathcal{T}'(\delta)$.
 4. modify all nodes $(\beta, F : -[id])$ in \mathcal{T}_j'' to $(\beta, F : P[id])$ if $L_0 = \mathbf{O}$, or $(\beta, F : O[id])$ if $L_0 = \mathbf{P}$; let the result be \mathcal{T}_j .

The *expanded debate tree* $\mathcal{T}^E(\delta)$ of $\mathcal{T}(\delta)$ is the limit⁹ of this sequence.

Note that $\mathcal{T}^E(\delta)$ of $\mathcal{T}(\delta)$ is the last element of this sequence, if the sequence is finite (when no leaf node is an assumption node).

Example 10.1. We illustrate the notion of expanded debate tree with the debate tree in Figure 14 for the dialogue in Example 6.3. Here, filtering has been applied

⁹The limit of a sequence of trees is a (possibly infinite) tree \mathcal{T} such that every tree in the sequence is a top-portion of \mathcal{T} , and every finite top-portion of \mathcal{T} is a top-portion of some tree in the sequence. A top-portion of a tree is a subtree that contains the root of the tree.

to both: $(b, um : O[11])$ and $(a, um : O[12])$. Hence, there is no repeated attack against these two assumptions. The expanded debate tree is shown in Figure 15.




Figure 14: The debate tree drawn from the dialogue in Example 6.3.




Figure 15: The expanded debate tree of the debate tree shown in Figure 14.

Lemma 10.2. Given a λ -defensive dialogue δ for χ , for λ -focused and filtered, let $\mathcal{T}(\delta)$ be the debate tree drawn from δ . There is an abstract dispute tree \mathcal{T}^a for $S \vdash \chi$ for some S , such that $\mathcal{DEF}(\mathcal{T}(\delta)) = \mathcal{DEF}(\mathcal{T}^a)$ and $\mathcal{CUL}(\mathcal{T}(\delta)) = \mathcal{CUL}(\mathcal{T}^a)$.

In order to guarantee that dialogues are (g-/a-/i-)successful, we restrict them to be conflict-free (as given in Definition 9.4):

Theorem 10.1. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta \in \mathcal{D}$, if δ is λ -conflict-free for λ -focused, then δ is g-successful and χ is grounded in the ABA framework \mathcal{F}_δ drawn from δ (supported by $\mathcal{DEF}(\mathcal{T}(\delta))$).

Theorem 10.1 gives a “recipe” for generating a- and i-successful dialogues, by virtue of proposition 6.1. The following Theorems 10.2 and 10.3 characterise a larger class of dialogues that can be proven to be a- and i-successful, respectively.

Theorem 10.2. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta \in \mathcal{D}$, if δ is λ -conflict-free for λ -focused and filtered, then δ is a-successful and χ is admissible in the ABA framework \mathcal{F}_δ drawn from δ (supported by $\mathcal{DEF}(\mathcal{T}(\delta))$).

Definition 10.3. A debate tree $\mathcal{T}(\delta)$ is *ideal* iff none of the opponent arguments drawn from $\mathcal{T}(\delta)$ belongs to an admissible set of arguments in the ABA framework \mathcal{F}_δ drawn from δ .

The *ideal outcome function* ω_i is such that, given $\delta \in \mathcal{D}$ and $\lambda \in \Lambda$, $\omega_i(\delta, \lambda) = \text{true}$ iff $\omega_{cf}(\delta, \lambda) = \text{true}$ and $\mathcal{T}(\delta)$ is ideal.

A λ -*ideal dialogue* is a λ -conflict-free dialogue $\delta \in \mathcal{D}$ such that $\omega_i(\delta, \lambda) = \text{true}$. An *ideal dialogue* is a λ -ideal dialogue for some $\lambda \in \Lambda$.

Theorem 10.3. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta \in \mathcal{D}$, if δ is λ -ideal for λ -focused and filtered, then δ is i-successful and χ is ideal in the ABA framework \mathcal{F}_δ drawn from δ (supported by $\mathcal{DEF}(\mathcal{T}(\delta))$).

10.1. Illustration: Twelve Angry Men

The debate tree $\mathcal{T}(\delta)$ given in Figure 8 is focused. Hence, δ in Table 1 is focused. Then, by Theorems 10.2, 10.1 and Proposition 6.1, δ is a-/g-/i-successful, as we noted in Section 6.1. The techniques presented in this section give a more constructive way to guarantee success.

11. Debate Forests and Unrestricted Dialogues

Focused dialogues are restricted, in that they force a single way of supporting or defending claims. Instead, in realistic dialogues (of various types) shifts inter- and intra-topic are likely to occur. In this section we consider unrestricted dialogues, that may not be focused by allowing multiple ways of supporting or defending claims (while sticking however to the initial claim). These may be useful to support brain storming and explore different alternative ways to determine acceptability of claims. We will show that analogous versions of Theorems 10.1, 10.2 and 10.3 hold for such unrestricted dialogues, that thus can be still be used to determine acceptability of claims.




Figure 16: The (non-focused) debate tree drawn from the (non-focused) dialogue in Example 11.1.

Example 11.1. The following dialogue δ

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 1, \text{rl}(s \leftarrow p), 2 \rangle$
$\langle a_1, a_2, 1, \text{rl}(s \leftarrow q), 3 \rangle$	$\langle a_2, a_1, 3, \text{rl}(q \leftarrow a), 4 \rangle$
$\langle a_1, a_2, 2, \text{rl}(p \leftarrow b, c), 5 \rangle$	$\langle a_2, a_1, 4, \text{asm}(a), 6 \rangle$
$\langle a_1, a_2, 6, \text{ctr}(a, r), 7 \rangle$	$\langle a_2, a_1, 7, \text{rl}(r \leftarrow), 8 \rangle$
$\langle a_1, a_2, 5, \text{asm}(b), 9 \rangle$	$\langle a_2, a_1, 5, \text{rl}(c \leftarrow), 10 \rangle$
$\langle a_1, a_2, 9, \text{ctr}(b, k), 11 \rangle$	$\langle a_2, a_1, 0, \pi, 12 \rangle$

is not focused as both utterances $\langle a_2, a_1, 1, \text{rl}(s \leftarrow p), 2 \rangle$ and $\langle a_1, a_2, 1, \text{rl}(s \leftarrow q), 3 \rangle$ expand s . The (non-focused) debate tree $\mathcal{T}(\delta)$ is shown in Figure 16. Nonetheless, δ is a- and g-successful, as s is admissible and grounded in the framework \mathcal{F}_δ drawn from δ , since there is an argument $\{b\} \vdash s$ but no attack against b in \mathcal{F}_δ .

To study the acceptability of claims in unrestricted dialogues, we first introduce the notion of arguments being *attacked* as follows.

Definition 11.1. Given a debate tree \mathcal{T} and an argument A drawn from \mathcal{T} , A is *attacked* in \mathcal{T} iff there is a node $N = (-, ma : _[-])$ in A such that N has a child node m in \mathcal{T} . We say that the sub-tree rooted at m in \mathcal{T} is an *attacker* of A in \mathcal{T} .

Example 11.2. Consider the dialogue δ

a_1	a_2
$\langle a_1, a_2, 0, \text{claim}(s), 1 \rangle$	$\langle a_2, a_1, 0, \pi, 2 \rangle$
$\langle a_1, a_2, 1, \text{asm}(s), 3 \rangle$	$\langle a_2, a_1, 3, \text{ctr}(s, a), 4 \rangle$
$\langle a_1, a_2, 4, \text{rl}(a \leftarrow q), 5 \rangle$	$\langle a_2, a_1, 4, \text{rl}(a \leftarrow p), 6 \rangle$
$\langle a_1, a_2, 5, \text{asm}(q), 7 \rangle$	$\langle a_2, a_1, 6, \text{asm}(p), 8 \rangle$
$\langle a_1, a_2, 7, \text{ctr}(q, b), 9 \rangle$	$\langle a_2, a_1, 8, \text{ctr}(p, c), 10 \rangle$
$\langle a_1, a_2, 9, \text{rl}(b \leftarrow), 11 \rangle$	$\langle a_2, a_1, 10, \text{rl}(c \leftarrow), 12 \rangle$
$\langle a_1, a_2, 0, \pi, 13 \rangle$	$\langle a_2, a_1, 0, \pi, 14 \rangle$

The debate tree drawn from this dialogue is in Figure 17. Here, the argument $\{s\} \vdash^t s$ is attacked; and the attacker of $\{s\} \vdash^t s$ is shown in Figure 18 (right).




Figure 17: The debate tree drawn from the dialogue in Example 11.2.




Figure 18: The argument $\{s\} \vdash^t s$ (left) and its *attacker* (right) for the debate tree in Figure 17.

The notion of *attacked* differs from the notion of *properly attacked* given in Definition 9.2 in two ways:

- Definition 11.1 defines *attacked* with respect to an argument, or a sub-tree, of a debate tree, whereas Definition 9.2 defines *properly attacked* with respect to a single (assumption) node in a debate tree;

- Definition 11.1 states that an argument is *attacked* as long as there are *some* nodes in the debate tree that are children of the argument, whereas Definition 9.2 defines an assumption node being *properly attacked* iff there is an actual argument that hangs from that node.

Debate trees are insufficient to help determine the acceptability of claims of non-focused dialogues. We introduce the notion of *debate forest*, composed of (debate) trees and defined in terms of the notions of *attacked* and *attacker* just given.

Definition 11.2. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta = \langle u_1, \dots, u_n \rangle$, the *debate forest* $\mathcal{F}(\delta)$ drawn from δ is the set of trees $\mathcal{F}^m(\delta)$ in the sequence $\mathcal{F}^0(\delta), \mathcal{F}^1(\delta), \dots, \mathcal{F}^m(\delta)$ constructed inductively from the π -pruned sequence $\delta' = \langle u'_1 \dots, u'_m \rangle$ obtained from δ , as follows (below, $L, L' \in \{\mathbf{P}, \mathbf{O}\}, L \neq L'$):

1. $\mathcal{F}^0(\delta)$ is empty;
2. $\mathcal{F}^1(\delta)$ consists of a single debate tree $\mathcal{T}_1^1(\delta)$, which consists of a single node $(\chi, um : \mathbf{P}[id_1])$, where id_1 is the identifier of $u_1 = u'_1$;
3. Let $\mathcal{F}^i(\delta)$ be the i -th forest consisting of trees $\mathcal{T}_1^i(\delta), \dots, \mathcal{T}_l^i(\delta)$, let $u'_{i+1} = \langle -, -, t, C, id \rangle$, and let $u'_i = \langle -, -, -, C_t, t \rangle$ be the target utterance of u'_{i+1} ; then $\mathcal{F}^{i+1}(\delta)$ is obtained as follows.
 - (a) if $C = rl(\beta_0 \leftarrow \beta_1, \dots, \beta_x)$ then $\mathcal{F}^{i+1}(\delta)$ is obtained in one of the following two cases:
 - i. if there is no debate tree in $\mathcal{F}^i(\delta)$ that contains $(\beta_0, mr : \mathbf{P}[t])$, then $\mathcal{F}^{i+1}(\delta)$ is $\mathcal{F}^i(\delta)$ updated as follows:
for each $\mathcal{T}_j^i(\delta), 0 < j \leq l$ in $\mathcal{F}^i(\delta)$ such that $\mathcal{T}_j^i(\delta)$ contains $(\beta_0, - : \mathbf{L}[t])$, then

$$(\beta_1, um : \mathbf{L}[id]), \dots, (\beta_x, um : \mathbf{L}[id])$$
are added to $\mathcal{T}_j^i(\delta)$ as children of $(\beta_0, - : \mathbf{L}[t])$; and if $(\beta_0, - : \mathbf{L}[t])$ is $(\beta_0, um : \mathbf{L}[t])$, then it is replaced by $(\beta_0, mr : \mathbf{L}[t])$ in each $\mathcal{T}_j^i(\delta)$ that contains $(\beta_0, um : \mathbf{L}[t])$;
 - ii. else, $\mathcal{F}^{i+1}(\delta)$ is $\mathcal{F}^i(\delta)$ with k additional debate trees, where k is the number of debate trees in $\mathcal{F}^i(\delta)$ containing $(\beta_0, mr : \mathbf{P}[t])$; for each $\mathcal{T}_j^i(\delta)$ containing $(\beta_0, mr : \mathbf{P}[t])$, a new debate tree is created by (1) copying $\mathcal{T}_j^i(\delta)$ and (2) replacing all children and sub-trees rooted at these children of $(\beta_0, mr : \mathbf{P}[t])$ with new children

$$(\beta_1, um : \mathbf{P}[id]), \dots, (\beta_x, um : \mathbf{P}[id]);$$
 - (b) if $C = asm(\alpha)$ then $\mathcal{F}^{i+1}(\delta)$ is $\mathcal{F}^i(\delta)$ with all $\mathcal{T}_j^i(\delta)$ that contain $(\alpha, um : \mathbf{L}[t])$ with $(\alpha, um : \mathbf{L}[t])$ replaced by $(\alpha, ma : \mathbf{L}[id])$;
 - (c) if $C = ctr(\alpha, \beta)$ then $\mathcal{F}^{i+1}(\delta)$ is obtained as follows:

- i. if there exists no $\mathcal{T}_j^i(\delta)$ in $\mathcal{F}^i(\delta)$ such that $\mathcal{T}_j^i(\delta)$ contains a node $N = (\alpha, ma : O[t])$ where N is in an argument A such that A is *attacked*, then $\mathcal{F}^{i+1}(\delta)$ is $\mathcal{F}^i(\delta)$ with all $\mathcal{T}_j^i(\delta)$, $0 < j \leq l$, that contain $(\alpha, ma : L[t])$ each having a new node $(\beta, um : L'[id])$ as a child of $(\alpha, ma : L[t])$;
- ii. otherwise, $\mathcal{F}^{i+1}(\delta)$ is $\mathcal{F}^i(\delta)$ with k additional debate trees, where k is the number of debate trees in $\mathcal{F}^i(\delta)$ that contain $(\alpha, ma : O[t])$. For each $\mathcal{T}_j^i(\delta)$ that contains $(\alpha, ma : O[t])$, a new debate tree is created by (1) copying $\mathcal{T}_j^i(\delta)$, (2) removing from it the *attacker* of the argument that contains $(\alpha, ma : O[t])$, and (3) adding $(\beta, um : P[id])$ as a child of $(\alpha, ma : O[t])$.

The construction of the debate forest drawn from the dialogue in Example 11.1 is shown in Figures 19, 20 and 21.

Conceptually, we construct debate forests using the same procedure for constructing debate trees, i.e. inserting new nodes and locating parents, etc. However, since the main purpose of having a debate forest is to ensure that every individual debate tree in the forest is focused, if we encounter a new utterance u such that adding u will yield a non-focused debate tree (as in cases 3(a)ii and 3(c)ii), we (I) duplicate all existing trees in the forest that contain the target of u , (II) remove all sub-trees of the target of u from the duplications, and (III) “add u ” to the duplicated trees. Note that cases 3(a)i, 3(b), and 3(c)i describe “normal” cases where adding u to the forest will not yield non-focused debate trees.

In this process, it is possible that more trees are duplicated than needed but the forest is a set, so only a single version of identical trees is kept. Hence, at the end, we obtain a set of unique trees in a debate forest.

Note that each such tree in the debate forest drawn from a dialogue is a top-portion of the dialogue tree drawn from the dialogue.

In order to study semantic properties of debate forests, we introduce the concept of sub-dialogues.

Definition 11.3. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta$, δ' is a *sub-dialogue* of δ iff it is a dialogue for χ between a_i and a_j and, for all utterances u in δ' , u is in δ . We say that δ is the *full-dialogue* of δ' .

Example 11.3. Two sub-dialogues of the dialogue in Example 11.1 are in tables 5 and 6. The first sub-dialogue (in Table 5) is neither a- nor g-successful as the proponent fails to defend a , an assumption in the argument $\{a\} \vdash s$. The second sub-dialogue (in Table 6) is both a- and g-successful as the proponent is able to construct the argument $\{b\} \vdash s$ and defend it (since b is not attacked).




Figure 19: Construction of the debate forest in Example 11.1 (Part 1).




Figure 20: Construction of the debate forest in Example 11.1 (Part 2).




Figure 21: Construction of the debate forest in Example 11.1 (Part 3).

Table 5: A sub-dialogue of the dialogue in Example 11.1 (part 1).

a_1	a_2
$\langle a_1, a_2, 0, claim(s), 1 \rangle$	
$\langle a_1, a_2, 1, rl(s \leftarrow q), 3 \rangle$	$\langle a_2, a_1, 3, rl(q \leftarrow a), 4 \rangle$
	$\langle a_2, a_1, 4, asm(a), 6 \rangle$
$\langle a_1, a_2, 6, ctr(a, r), 7 \rangle$	$\langle a_2, a_1, 7, rl(r \leftarrow), 8 \rangle$

Table 6: A sub-dialogue of the dialogue in Example 11.1 (part 2).

a_1	a_2
$\langle a_1, a_2, 0, claim(s), 1 \rangle$	$\langle a_2, a_1, 1, rl(s \leftarrow p), 2 \rangle$
$\langle a_1, a_2, 2, rl(p \leftarrow b, c), 5 \rangle$	
$\langle a_1, a_2, 5, asm(b), 9 \rangle$	$\langle a_2, a_1, 5, rl(c \leftarrow), 10 \rangle$
$\langle a_1, a_2, 9, ctr(b, k), 11 \rangle$	

Note that both example sub-dialogues in Example 11.3 are exhaustive, however Definition 11.3 does not impose that sub-dialogues are exhaustive.

Sub-dialogues can be used to prove properties of unrestricted dialogues in that, intuitively, each unrestricted (but exhaustive) dialogue can be understood as the collection of several independent focused sub-dialogues. Hence, each sub-dialogue draws a tree in the debate forest drawn from the full (unrestricted) dialogue. Thus, as soon as one debate tree is found to be “good”, then the particular sub-dialogue drawing the debate tree becomes successful, and the full non-focused dialogue is successful.

Lemma 11.1. Given a debate forest $\mathcal{F}(\delta)$ drawn from a dialogue δ , each tree in $\mathcal{F}(\delta)$ is a debate tree drawn from a focused sub-dialogue δ_i of δ .

For example, the two debate trees $\mathcal{T}_1^{11}(\delta)$ and $\mathcal{T}_2^{11}(\delta)$ in Figure 20 and 21 are drawn from the sub-dialogues in Tables 5 and 6, respectively.

Note that, in the case of unrestricted dialogues, the notion of filtered legal-move functions is overly strong, in that it imposes filtering across trees in the forest, thus ruling out successful dialogues unnecessarily. In order to prove semantic properties for a wider class of unrestricted dialogues, we relax the definition of filtered legal-move functions by redefining it with respect to debate forests:

Definition 11.4. A debate forest $\mathcal{F}(\delta)$ is *filtered* iff all debate trees in $\mathcal{F}(\delta)$ are filtered. A legal-move function $\lambda \in \Lambda$ is *filtered* iff for every $\delta \in \mathcal{D}$ such that $\mathcal{F}(\delta)$ is filtered, for every $u \in \lambda(\delta)$, $\mathcal{F}(\delta \circ u)$ is still filtered.

This notion of legal-move function is a generalised version of the notion in Definition 9.5, since a debate tree is a special case of a debate forest that consists of a single tree. From now on we will assume that filtered dialogues are dialogues compatible with this generalised notion of filtered legal-move function.

Similarly, patient legal-move functions can be redefined with respect to debate forests as follows. (Again, we assume that dialogues in later discussion are compatible with the following generalised notion of patient legal-move function).

Definition 11.5. A debate forest $\mathcal{F}(\delta)$ is *patient* iff all debate trees in $\mathcal{F}(\delta)$ are patient debate trees. A legal-move function $\lambda \in \Lambda$ is *patient* iff for every $\delta \in \mathcal{D}$ such that $\mathcal{F}(\delta)$ is patient, for every $u \in \lambda(\delta)$, $\mathcal{F}(\delta \circ u)$ is still patient.

We can obtain similar results as for focused dialogues, by using the same combinations of legal move-functions and outcome functions, as follows.

Theorem 11.1. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta \in \mathcal{D}$ such that δ is λ -exhaustive, if there is a debate tree $\mathcal{T}(\delta_i)$ drawn from a sub-dialogue δ_i of δ such that δ_i is λ -conflict-free, then δ is g-successful and χ is grounded in the ABA framework \mathcal{F}_δ drawn from δ (supported by $\mathcal{DEF}(\mathcal{T}(\delta_i))$).

Theorem 11.2. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta \in \mathcal{D}$ such that δ is λ -exhaustive for λ filtered, if there is a debate tree $\mathcal{T}(\delta_i)$ drawn from a sub-dialogue δ_i of δ such that δ_i is λ -conflict-free, then δ is a-successful and χ is admissible in the ABA framework \mathcal{F}_δ drawn from δ (supported by $\mathcal{DEF}(\mathcal{T}(\delta_i))$).

Theorems 11.1 and 11.2 do not refer to debate forests drawn from dialogues directly, rather, they specify conditions of sub-dialogues thereof. This is partly because ω_{cf} is defined with respect to properties of debate trees rather than debate forest (see Definition 9.4). Though we could overwrite Definition 9.4 and redefine it with respect to debate forests, we choose not to as the current version of Theorems 11.1 and 11.2 clearly indicates that testing the acceptability of the claim of an unrestricted dialogue can be reduced to testing the acceptability of the claim of sub-dialogues that draw debate trees and exhibit desirable properties.

To lift the results for the ideal semantics to unrestricted dialogues, we modify the notion of ideal outcome function so that it focuses on a sub-dialogue but uses the ABA framework drawn from the full dialogue to check opponent nodes:

Definition 11.6. Let $\delta \in \mathcal{D}$, $\mathcal{F}(\delta)$ be the debate forest drawn from δ and δ_i be a sub-dialogue of δ . $\mathcal{T}(\delta_i)$ in $\mathcal{F}(\delta)$ is *ideal* with respect to δ iff none of the opponent arguments drawn from $\mathcal{T}(\delta_i)$ belongs to an admissible set of arguments in the ABA framework \mathcal{F}_δ drawn from δ . The *ideal outcome function* ω_i is such that, given $\delta \in \mathcal{D}$ and $\lambda \in \Lambda$, $\omega_i(\delta, \lambda) = true$ iff there exists a sub-dialogue δ_i of δ such that $\omega_{cf}(\delta_i, \lambda) = true$ and $\mathcal{T}(\delta_i)$ is ideal with respect to δ . A λ -*ideal dialogue* is a dialogue $\delta \in \mathcal{D}$ such that $\omega_i(\delta, \lambda) = true$.

The following example illustrates the computational advantages of this definition in the case of unrestricted dialogues.

Example 11.4. Given the dialogue δ :

a_1	a_2
$\langle a_1, a_2, 0, claim(s), 1 \rangle$	$\langle a_2, a_1, 0, rl(s \leftarrow a), 2 \rangle$
$\langle a_1, a_2, 1, rl(s \leftarrow b), 3 \rangle$	$\langle a_2, a_1, 2, asm(a), 4 \rangle$
$\langle a_1, a_2, 4, ctr(a, q), 5 \rangle$	$\langle a_2, a_1, 5, asm(q), 6 \rangle$
$\langle a_1, a_2, 6, ctr(q, r), 7 \rangle$	$\langle a_2, a_1, 7, rl(r \leftarrow), 8 \rangle$
$\langle a_1, a_2, 3, rl(b \leftarrow a), 9 \rangle$	




Figure 22: The debate forest for the dialogue δ in Example 11.4.

the forest $\mathcal{F}(\delta)$ drawn from δ is shown in Figure 22, where δ_1 and δ_2 are the two sub-dialogues that draw the individual debate trees. We can see that δ_1 is ideal in $\mathcal{F}(\delta)$ as the opponent argument, $\{q\} \vdash q$, drawn from $\mathcal{T}(\delta_1)$, is not in an admissible set in \mathcal{F}_δ , even though δ_2 is not “complete” yet as $\langle -, -, 9, ctr(a, q), 10 \rangle$ can still be uttered. Such “early termination” is possible because of the redefined ideal outcome function.

Theorem 11.3. Given a dialogue $\mathcal{D}_{a_j}^{a_i}(\chi) = \delta \in \mathcal{D}$, if δ is λ -ideal for λ -focused and filtered, then δ is i -successful and χ is ideal in the ABA framework \mathcal{F}_δ drawn from δ (supported by $\mathcal{DEF}(\mathcal{T}(\delta))$).

11.1. Illustration: Twelve Angry Men

We modify the dialogue shown in Table 1/Table 2 to illustrate non-focused dialogues. In Table 2, after a_1 making the utterance:

Witness 2 cannot be believed if it has a poor eyesight.

a_1 makes another utterance:

Witness 2 cannot be believed if it has a conflict of interest.

Correspondingly, in Table 1, after the utterance

$\langle a_1, a_2, 12, rl(w2_not_believable \leftarrow w2_has_poor_eyesight), 13 \rangle$,

another utterance

$\langle a_1, a_2, 12, rl(w2_not_believable \leftarrow w2_has_conflict_of_interest), 14 \rangle$

is added to the dialogue. (Utterances with ID 14, 15 and 16 in the original dialogue are then given with ID 15, 16 and 17, respectively.) This modified dialogue is not focused, as the new utterance leads to another way of supporting *w2_not_believable*, which defends the claim *boy_innocent*.

The resulting debate forest contains two debate trees: $\mathcal{T}_1(\delta)$ and $\mathcal{T}_2(\delta)$. $\mathcal{T}_1(\delta)$ is the debate tree shown in Figure 8. $\mathcal{T}_2(\delta)$ is $\mathcal{T}_1(\delta)$ with node

$$(w2_has_poor_eyesight, mr : P[13])$$

replaced by

$$(w2_has_conflict_of_interest, um : P[14]).$$

12. Related Work

Several proposals for argumentation based dialogues exist in the literature. We briefly review here those that are most closely related to our work.

McBurney and Parsons [27] give an overview of dialogue games for argumentation relaying upon syntax and semantics of dialogue protocols. Our work can be seen as providing a novel dialogue game for ABA, hence the syntax of our dialogue is based on ABA and the semantics of our dialogues are standard ABA semantics. These argumentation semantics can be used to define specialised semantics for specific dialogue types (such as persuasion and inquiry [43]), e.g. following the lines of [17, 18, 19, 20].

Black and Hunter [6] present a formal system for inquiry dialogues based on DeLP [22] as the underlying argumentation framework. Our work differs from theirs as (1) it defines a mechanism for any type of dialogue whereas they focus on inquiry dialogues; (2) it uses ABA whereas they use DeLP; (3) it does not force an agent to disclose all knowledge whereas they force full disclosure of all relevant knowledge for the purpose of inquiry; (4) it does not force a strict interleaving whereas they do. In addition to being more general, our model benefits from using ABA as it makes use of formal soundness results thereof to prove soundness of dialogues. Moreover, since abstract argumentation is an instance of ABA, our model also provides a model for abstract argumentation dialogues too.

Parsons et al [29] examine three notions of relevance in dialogues where utterances are arguments and attacks:

- R1 (every new utterance has a direct impact on the claim),
- R2 (every new utterance directly or indirectly impacts the claim), and
- R3 (every new utterance has a direct impact on the previous one).

Our patient legal-move function and focused legal-move function jointly resemble

their R1 related-ness. Our related legal-move function has some of the features of their R2 related-ness in that all related utterances have impact on the claim. However, our utterances are at a finer granularity level, as they correspond to rules, assumptions, and contraries, whereas in [29] utterances are at the *argument level*, i.e., each utterance contains an argument. Thus, there is no direct mapping between our work and their relevance.

Prakken [32] defines a formal system for persuasion. The main differences with our work are: (1) since that work focuses on persuasion dialogues, proponent and opponent roles are pre-assigned to agents before the dialogue whereas in our work agents can play both roles within the same dialogue; (2) Prakken focuses on the grounded semantics, whereas we allow admissibility, grounded, and ideal semantics; (3) his set of utterances refers to arguments and attacks, as in the case of [29]; (4) he forces the support of arguments to be minimal, whereas we do not, in the spirit of [13]; (5) he does not allow agents to jointly construct arguments whereas we do. Roughly speaking, in terms of obtaining the soundness result, i.e., a “successful” dialogue means that the claim of the dialogue is “acceptable”, there are two main differences between his approach and ours. Firstly, we rely on mapping our dialogues to abstract dispute trees whereas he uses a form of labelling. Hence, in our case, “successful” dialogues can be mapped to “good” abstract dispute trees; whereas, in Prakken’s case, “successful” dialogues have winning arguments that are labelled *in*. Secondly, to support non-focused dialogues, we use debate forests, and show that individual trees in a forest can be mapped to abstract dispute trees, whereas Prakken puts all arguments into a single tree and then identifies a “winning strategy” that represents a sub part of the single tree. Overall, our definitions allow to leverage on formal results for ABA (and in particular the connection between abstract dispute trees in ABA and our debate trees and forests) to formally prove soundness of our dialogue model with respect to a range of argumentation semantics. It is not clear whether and how this could be done in the case of Prakken’s trees, as the labels already carry a semantic connotation and are intrinsically linked to a specific argumentation semantics.

Thang et.al [39] propose a dialogue system for persuasion. Similar to Prakken’s work, they pre-assign the roles of proponent and opponent to agents participating in a dialogue. Moreover, they use an annotation scheme similar to Prakken’s IN/OUT labelling to compute the acceptability of the dialogue. Unlike Prakken’s work, they have used ABA as their underlying representation and define an ABA-based dialogue system. Yet, unlike our generic framework composed of various components (legal-move functions and outcome functions), their work has a limited focus on persuasion. They point out computational inefficiency of our notion

of exhaustive dialogue, also present in our earlier work [16], as some utterances might be repeated in exhaustive dialogues without adding information to the ABA framework drawn from these dialogues. We believe that this potential loss of efficiency is a small price to pay if enforcing exhaustiveness helps to define appropriate legal-move and outcome functions to establish the correspondence between debate trees and abstract dispute trees and formally prove soundness. Moreover, the potential inefficiency can be mitigated in implementations of our dialogues, where such utterances can be brought about automatically.

Prakken et.al [35] introduce an argumentation dialogue game for analysing IT security risk assessment, with *ASPIC*⁺ [34] as the underlying argumentation framework, under the grounded semantics. As [32], this work pre-assigns the roles of proponent and opponent to the two participating agents and uses a version of the IN/OUT labelling to compute acceptability. Instead of using explicitly defined structures such as our legal-move functions, this work relies on a single notion of *legal dialogue* which, not unlike our outcome functions, serves the purpose of validating the dialogue by checking that it exhibits desirable properties.

13. Conclusions

We have presented a formal dialogue model for assumption-based argumentation (ABA) [7, 12, 40]. We use ABA for the representation of arguments and attacks and for determining “success” of dialogues. ABA is well-suited as a foundation for argumentation-based dialogues since it is a general purpose, structured argumentation framework with several applications (e.g. see [12, 40]), well-understood theoretically (e.g. as an instance of other argumentation frameworks, notably abstract argumentation [15, 40]) and ASPIC+ [34], and equipped with provably correct computational mechanisms with respect to several semantics [14, 15, 41] that we rely upon in order to prove our formal soundness results.

We proved that our model is sound by connecting it with the admissibility, grounded and ideal argumentation semantics for ABA. Thus, our dialogues can be seen as a distributed mechanism for computing admissible, grounded and ideal extensions (supporting claims) of argumentation frameworks while these are generated/jointly constructed. Our dialogue model has two main advantages:

1. Our dialogues meet soundness criteria with respect to different argumentation semantics, obtained by lifting soundness results for the underlying argumentation framework, ABA;
2. Our model is not tailored to any dialogue type. Indeed, it is generic and flexible as it is composed of many loosely coupled components (i.e., legal-move

Table 7: Summary of types of dialogues defined/studied in this paper. Here, λ_{fl} , λ_{rt} , λ_f , λ_p , λ_{fi} denote flat, related, focused, patient, and filtered legal move functions, respectively.

Outcome functions	Legal-move functions	Dialogues	Results
ω_{ABA}	λ_{fl}	ABA	Lemma 5.2
ω_{ex} (extending ω_{ABA})	$\lambda_{fl}, \lambda_{rt}$	exhaustive	Lemmas 8.2, 8.3
ω_{lw} (extending ω_{ex})	$\lambda_{fl}, \lambda_{rt}$	defensive	
ω_{cf} (extending ω_{lw})	$\lambda_{fl}, \lambda_{rt}$	conflict-free	Lemma 9.1
”	$\lambda_{fl}, \lambda_{rt}, \lambda_p, \lambda_f,$	g-successful	Theorem 10.1
”	$\lambda_{fl}, \lambda_{rt}, \lambda_p$	g-successful	Theorem 11.1
”	$\lambda_{fl}, \lambda_{rt}, \lambda_p, \lambda_{fi}, \lambda_f$	a-successful	Theorem 10.2
”	$\lambda_{fl}, \lambda_{rt}, \lambda_p, \lambda_{fi}$	a-successful	Theorem 11.2
ω_i (extending ω_{cf})	$\lambda_{fl}, \lambda_{rt}, \lambda_p, \lambda_{fi}, \lambda_f$	i-successful	Theorem 10.3
”	$\lambda_{fl}, \lambda_{rt}, \lambda_p, \lambda_{fi}$	i-successful	Theorem 11.3

and outcome functions) that can be freely combined to construct dialogues with different aims and of different types.

We have assumed that agents exchange their views in ABA-format, namely ABA serves as a standard for the exchange of information between agents. However, agents may adopt an internal representation different from ABA.

Our dialogue model relies upon several legal-move functions and outcome functions. Table 7 summarises the main kinds of dialogues we have defined with these legal-move and outcome functions, as well as the core properties they enjoy.

Our work opens several avenues for future research, as follows.

We have focused on soundness results only, and ignored the important issue of completeness. Completeness results for dispute derivations for ABA in [15, 41] are a useful starting point for studying the completeness of our dialogues.

Our model for argumentation-based dialogues has very clear formal properties, thanks to the use of ABA and leveraging formal results for its computational counterparts. It would be interesting to see whether the model could be modified and/or extended to support argumentation dialogue with respect to other types of structured argumentation. This may require the need for additional utterances, e.g., in the case of ASPIC+ [34], conveying preferences by an agent of one rule over another.

We have only considered two agents, for simplicity. We believe that our model

naturally generalises to several agents. Indeed, debate trees and forests are completely neutral as to which agent plays the (fictitious) role of proponent and opponent, and focus instead on the content of utterances. This generalisation would require extending the turn-making function to accommodate several agents, allowing, in particular, for agents to play different roles within the same dialogue (e.g. in the case of inquiry [43]) or to stick to a specific role (e.g. in the case of persuasion [43]) and form coalitions (for and against acceptability of the dialogue claim). We plan to study this generalisation in the future, taking into account the challenges raised in the literature [10].

We have focused on three semantics for ABA (although some other semantics, of preferred and complete extensions, are accommodated for free, as we discuss at the end of Section 6). It would be interesting to consider other semantics, e.g. that of stable extensions [11, 7].

For simplicity we have restricted dialogues to be patient, but non-patient dialogues could also be used to guarantee successful dialogues.

The soundness results for i-successful dialogues are obtained in a semi-constructive manner, by checking that opponent arguments in debate trees do not belong to admissible sets in the ABA framework drawn from the dialogue. It would be useful to be able to check this condition using dialogues in turn, e.g. inspired by the notion of Fail-dispute derivation in [15].

We have focused on dialogues that generate *flat* ABA frameworks. It would be interesting to see whether the model could be generalised to the case of non-flat ABA frameworks, and in particular to deal with settings where agents may disagree on whether a sentence is an assumption or not.

Future work also includes applying our dialogue framework to various dialogue types, e.g., negotiation and deliberation, following the lines of [17, 18, 19, 20], and extending our framework to model sequences of dialogues. Although these initial attempts show promise in the deployment of our general model, much work is required to explore this challenge in full. An interesting question, for example, is whether all forms of deliberation as discussed in [28] can indeed be supported and/or accommodated.

Finally, it would be interesting to see how our framework can be applied in human-agent dialogical interactions.

Acknowledgements

This work was supported by the EPSRC project TRaDAr (Transparent Rational Decisions by Argumentation): EP/J020915/1. We are grateful to K. Apt

for pointing out [2], from which we have drawn several illustrations. We thank A. Hunter and A. Edalat, as well as anonymous reviewers, for very useful comments.

Appendix A. Glossary

- *Utterance* (**Definition 3.1**) is the basic building block in our dialogue model.
- *Dialogue* (**Definition 3.2**) is a sequence of utterances.
- *Framework drawn from a Dialogue* (**Definition 3.3**) contains all rules, assumptions and contraries disclosed in a dialogue.
- *Turn-making Function* (**Definition 4.1**) decides the agent making the next utterance.
- *Legal-move Function* (**Definition 4.2**) specifies how to continue a dialogue.
- *Outcome Function* (**Definition 4.3**) allows to determine whether dialogues have desirable properties.
- *Flat Legal-move Function* (**Definition 5.1**) and ABA outcome function (**Definition 5.2**) ensure that the framework drawn from a dialogue is a flat ABA framework.
- *A-/G-/I-successful Dialogues* (**Definition 6.1**) are dialogues with their claims being admissible, grounded, and ideal, respectively, in the ABA frameworks drawn from them.
- *Related Utterances* (**Definition 7.1**) are “syntactically” related utterances.
- *Top-down Related Utterances* (**Definition 7.2**) are “semantically” related utterances.
- *Top-down Related Legal-move Function* (**Definition 7.3**) ensures “semantically” related relations between utterances.
- *Exhaustive Outcome Function* (**Definition 7.4**) ensures that all utterances compatible with the used legal-move functions are uttered in the dialogue.
- *Debate Tree drawn from Dialogue* (**Definition 8.1**) keeps track of dialogue progress, for book-keeping purpose.

- *Argument drawn from a Debate Tree* (**Definition 8.2**) defines how to extract ABA arguments from debate trees.
- *Defence Set and Culprits* (**Definition 8.3**) are proponent and attacked opponent assumptions from a debate tree.
- In *Patient Debate Trees* (**Definition 9.1**) arguments are not attacked until fully constructed.
- *Patient Legal-move Function* (**Definition 9.1**) constructs dialogues drawing patient debate trees.
- In *Last-word Debate Tree* (**Definition 9.3**) all leaves are proponent nodes.
- *Defensive Dialogues* (**Definition 9.3**) draw last-word debate trees.
- *Conflict-free Outcome Function* (**Definition 9.4**) enforces that there is no overlap between defence set and culprits.
- *Conflict-free Dialogues* (**Definition 9.4**) are defensive dialogues for which the conflict-free outcome function returns true.
- In *Filtered Debate Trees* (**Definition 9.5**) there is no repeated attack between arguments. In *Filtered Debate Forests* (**Definition 11.4**) all debate trees are filtered.
- *Filtered Legal-move Functions* (**Definition 9.5, 11.4**) construct dialogue drawing filtered debate trees/forests.
- *Properly Attacked Node* (**Definition 9.2**) is an assumption node which has been attacked by some actual argument.
- *Focused Debate Trees* (**Definition 10.1**) contain a single way of defending their root.
- *Focused Legal-move Function* (**Definition 10.1**) constructs dialogues drawing focused debate trees.
- *Ideal Debate Tree* (**Definition 10.3**) has the property that none of the opponent arguments drawn from the tree is admissible.
- *Ideal Outcome Function* (**Definition 10.3, 11.6**) returns true iff the debate tree / forest drawn from a dialogue is ideal.

- *Ideal Dialogue* (**Definition 10.3, 11.6**) is a conflict-free dialogue for which the ideal outcome function returns true.
- *Attacked Argument* (**Definition 11.1**) is an argument in a debate tree with an assumption node that has a child.
- *Debate Forest* (**Definition 11.2**) is a set of debate trees.
- *Sub-dialogue* (**Definition 11.3**) is a set of utterances extracted from a dialogue.
- *Patient Debate Forest* (**Definition 11.5**) is a debate forest containing only patient debate trees.
- *Ideal Debate Forest* (**Definition 11.6**) a debate forest containing only ideal debate trees.

Appendix B. Proofs

Proof of Lemma 5.1. By Definition 5.1, for any $\beta \in \mathcal{L}$, if $\beta \in \mathcal{A}_\delta$ then there exists no ρ , such that $\rho \in \mathcal{R}_\delta$ and $Head(\rho) = \beta$; and if there exists ρ , such that $\rho \in \mathcal{R}_\delta$ then $\beta \notin \mathcal{A}_\delta$. Therefore, \mathcal{F}_δ is flat.

Proof of Lemma 5.2. Trivial.

Proof of Proposition 6.1. Since δ is g-successful, χ is grounded. By definition of grounded semantics (see Section 2), χ is also complete and admissible. Hence g-successful implies a-successful. By Theorem 2.1 (iii) of [15], which states that an ideal set of arguments is a superset of the grounded set of arguments, we have that g-successful implies i-successful.

Proof of Proposition 6.2. Trivial by definition of ideal semantics (see Section 2).

Proof of Lemma 8.1. It is easy to see that debate trees for related dialogues are always guaranteed to be well-formed, in that each non-root node in them has exactly one parent. We show, given a dialogue δ , that this inductive process constructs a debate tree $\mathcal{T}(\delta)$ as in Definition 8.1.

Condition 1(a) and 1(b) in Definition 8.1 are trivially true since nodes in $\mathcal{T}^0(\delta), \dots, \mathcal{T}^m(\delta)$ are nodes inserted in accordance with utterances in δ hence, if a sentence β is in a node, then β must be in some utterance in the dialogue. Condition 1(c) in Definition 8.1 is true as 3(a) and 3(b) in the lemma specify how

nodes tagged *um* are replaced by nodes tagged *ma* and *mr*. Condition 2(a) through 2(f) in Definition 8.1 are jointly met by 3(a) through 3(c) in the lemma.

Hence, this inductive process yields a debate tree as per Definition 8.1.

Proof of Lemma 8.2. This lemma is trivially true as a node in an actual argument can be mapped to a node in an ABA argument by 1) dropping the tag and the ID, and 2) adding nodes τ as children of leaf nodes of the form $(-, mr : -[-])$ (as each of these node represents a rule with an empty body).

Proof of Lemma 8.3. To show $S \vdash^t \beta$ can be drawn from $\mathcal{T}(\delta)$ is to show all sentences in $S \vdash \beta$ have corresponding nodes in $\mathcal{T}(\delta)$; and these nodes are connected in the right order with the same label. Let \mathcal{F}_δ be $\langle \mathcal{L}, \mathcal{R}_\delta, \mathcal{A}_\delta, \mathcal{C}_\delta \rangle$. Since δ is λ -exhaustive, for all sentences $\beta_0, \beta_1, \dots, \beta_n \in \mathcal{L}$, if $\beta_0 \leftarrow \beta_1, \dots, \beta_n \in \mathcal{R}_\delta$, then nodes $n_0 = (\beta_0, - : L[-])$, $n_1 = (\beta_1, - : L[-])$, \dots , $n_2 = (\beta_2, - : L[-])$.

Proof of Lemma 9.1. Assume otherwise, let $D = \mathcal{DEF}(\mathcal{T}(\delta))$, $A, A' \subseteq D$ and $A \vdash c$ attacking $A' \vdash c'$ in \mathcal{F}_δ . By definition of defence set, all assumptions in D are labelled P. Since δ is λ -exhaustive, by Lemma 8.3, we know that $A \vdash^t c$ and $A' \vdash^t c'$ are in $\mathcal{T}(\delta)$. Since A and A' are in two arguments attacking each other, they cannot have the same labelling. So either A is labelled P or A' is labelled P, but not both: contradiction, since both A and A' are subsets of D .

Proof of Lemma 10.1. We can transform debate trees into abstract dispute trees. Given a debate tree $\mathcal{T}(\delta)$, its equivalent abstract dispute tree \mathcal{T}^a is constructed as follows.

1. Delete all nodes N from $\mathcal{T}(\delta)$ where N is in a potential argument. Let $\mathcal{T}'(\delta)$ be the result.
2. Modify $\mathcal{T}'(\delta)$ by appending a new *flag* field $Z = \{0, 1\}$ to each remaining node in $\mathcal{T}(\delta)$ and initialise Z to 0 for all nodes, i.e., a node now looks like $(-, - : -[-])[0]$. Let $\mathcal{T}''(\delta)$ be the result.
3. \mathcal{T}^a is \mathcal{T}_m^a in the sequence $\mathcal{T}_1^a, \dots, \mathcal{T}_m^a$ constructed inductively as follows:
 - \mathcal{T}_0^a is empty;
 - let A be the argument in $\mathcal{T}''(\delta)$ that contains the root of $\mathcal{T}(\delta)$. Set the flags of all nodes in $\mathcal{T}''(\delta)$ that are also in A to 1. Let $\mathcal{T}_1^a(\delta)$ be the result; then \mathcal{T}_1^a contains a single node that holds A and is labelled by P;

- let \mathcal{T}_i^a be the i -th tree, for $0 < i < m$, then \mathcal{T}_{i+1}^a is \mathcal{T}_i^a with an additional node $L : B$, where B is an argument drawn from $\mathcal{T}_i''(\delta)$, child of B' , another argument drawn from $\mathcal{T}_i''(\delta)$, such that:
 - the flag of at least one node in B is 0;
 - the root node of B has a parent node p in $\mathcal{T}_i''(\delta)$ with flag equal to 1 and such that p is in B' ;
 - L is P if the root node of B is a P node, otherwise L is O ;
 - set the flags of all nodes in $\mathcal{T}_i''(\delta)$ that are also in B to 1, let $\mathcal{T}_{i+1}''(\delta)$ be the result.
- 4. m is the smallest index such that there is no node in $\mathcal{T}_m''(\delta)$ with its flag equal to 0.

\mathcal{T}^a constructed above is an abstract dispute tree as follows.

1. Every node of $\mathcal{T}^a = \mathcal{T}_m^a$ contains a single argument as there is no potential argument in $\mathcal{T}'(\delta)$ and $\mathcal{T}''(\delta)$. For each argument, there is a unique node in \mathcal{T}^a . Each node in \mathcal{T}^a is labelled either P or O as arguments drawn from $\mathcal{T}(\delta)$ are labelled either P or O .
2. The root node of \mathcal{T}^a contains the argument that has the claim of the dialogue. The root is labelled P by construction of $\mathcal{T}(\delta)$.
3. By definition of λ -exhaustive outcome function, since δ is λ -exhaustive, every assumption is attacked in as many ways as possible in \mathcal{F}_δ . Hence a P node in \mathcal{T}^a has as many children as its attacks by actual arguments.
4. By definition of patient and focused legal-move function, since δ is patient and focused, there is only one way of attacking an argument labelled by a O node. By the last-word outcome function, since δ is defensive, there is no un-attacked argument labelled by a O node. Therefore, every O node in \mathcal{T}^a with an assumption has one and only one P node as its child.

Since \mathcal{T}^a contains the same actual arguments as $\mathcal{T}(\delta)$ and arguments have the same P/O labelling in both \mathcal{T}^a and $\mathcal{T}(\delta)$, we have $\mathcal{DEF}(\mathcal{T}(\delta)) = \mathcal{DEF}(\mathcal{T}^a)$ and $\mathcal{CUL}(\mathcal{T}(\delta)) = \mathcal{CUL}(\mathcal{T}^a)$.

Proof of Lemma 10.2. We can transform debate trees into abstract dispute trees using the procedure as shown in the proof of Lemma 10.1 with a modification. After deleting all nodes N from $\mathcal{T}(\delta)$ where N is in a potential argument (step 1), we replace $\mathcal{T}(\delta)$ with $\mathcal{T}^E(\delta)$, the expanded debate tree of $\mathcal{T}(\delta)$. The rest of the construction of \mathcal{T}^a remains unchanged.

It is easy to see that \mathcal{T}^a constructed with the modified $\mathcal{T}^E(\delta)$ is an abstract dispute tree and $\mathcal{DEF}(\mathcal{T}(\delta)) = \mathcal{DEF}(\mathcal{T}^a)$ and $\mathcal{CUL}(\mathcal{T}(\delta)) = \mathcal{CUL}(\mathcal{T}^a)$, as shown in the proof of Lemma 10.1.

Proof of Theorem 10.1. Similar to the proof of Lemma 10.1, if $\omega_{cf}(\delta, \lambda) = true$, then there exists an abstract dispute tree \mathcal{T}^a such that $\mathcal{DEF}(\mathcal{T}(\delta)) = \mathcal{DEF}(\mathcal{T}^a)$ and $\mathcal{CUL}(\mathcal{T}(\delta)) = \mathcal{CUL}(\mathcal{T}^a)$. As shown in [12], as a direct consequence of Theorem 3.7 in [25], we obtain that the defence set of a grounded abstract dispute tree is a subset of the grounded set of arguments. Hence δ is g-successful and χ is grounded and supported by $\mathcal{DEF}(\mathcal{T}(\delta))$.

Proof of Theorem 10.2. If $\omega_{cf}(\delta, \lambda) = true$, by Lemma 10.2 there exists an abstract dispute tree \mathcal{T}^a such that $\mathcal{DEF}(\mathcal{T}(\delta)) = \mathcal{DEF}(\mathcal{T}^a)$ and $\mathcal{CUL}(\mathcal{T}(\delta)) = \mathcal{CUL}(\mathcal{T}^a)$. By Theorem 5.1 of [14], the theorem holds.

Proof of Theorem 10.3. Let \mathcal{F}_δ be $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$. From Lemma 10.2, we know that there is an abstract dispute tree \mathcal{T}^a for $S \vdash \chi$ for some $S \subseteq \mathcal{A}$. Since δ is ideal, by Definition 10.3 we know that none of the opponent arguments drawn from $\mathcal{T}(\delta)$ belongs to an admissible set of arguments in \mathcal{F}_δ . Hence, \mathcal{T}^a is ideal. By Theorem 3.4 in [15] the theorem holds.

Proof of Lemma 11.1. We show that (I) the inductive process in Definition 11.2 constructs a set of debate trees (II) each of these is for a sub-dialogue of δ and (III) each such sub-dialogue is focused.

(I) For each tree $\mathcal{T}_j^i(\delta)$ in $\mathcal{F}^i(\delta)$, condition 1(a) and 1(b) in Definition 8.1 are trivially true as all nodes in each tree are inserted in accordance with utterances in δ hence if a sentence β is in a node, β must be in some utterance in the dialogue. Condition 2(a)-2(f) in Definition 8.1 are met by condition 3(a)-3(c) in Definition 11.2. Hence, each tree in a debate forest is a debate tree.

(II) Given that $\mathcal{F}(\delta)$ contains l debate trees $\mathcal{T}_1(\delta), \dots, \mathcal{T}_l(\delta)$, the sub-dialogue $\delta_i, 0 < i \leq l$, that draws the debate tree $\mathcal{T}_i(\delta)$ is constructed as follows.

1. δ_i is initialised to empty;
2. For each node $(\beta, F : _ [id]) = n$ in $\mathcal{T}(\delta_i)$,
 - if $u_{id} = \langle _, _, t, _, id \rangle$ is in δ but not in δ_i , then add u_{id} to δ_i ;
 - let u_t be the utterance in δ such that u_{id} is related to u_t ; if u_t is not in δ_i , then add u_t to δ_i ;
3. Sort δ_i in the order of utterance IDs.

It is easy to see that each δ_i constructed as above is a sub-dialogue of δ and $\mathcal{T}_i(\delta)$ is drawn from δ_i .

(III) Let us consider any sub-dialogue δ_i from part (II). We first show that δ_i is related, and then that it is focused.

δ_i is related: Trivially δ_i is an ABA dialogue, given that δ is. We now show δ_i is compatible with a related $\lambda_i \in \Lambda$ and that $\omega_{ex}(\delta_i, \lambda_i) = true$.

By Definition 8.1, we know that, in a debate tree, there are 6 ways in which two nodes N and N' can be connected (N is the parent of N'):

1. $N = (-, mr : L[id])$ and $N' = (-, mr : L[id'])$, or
2. $N = (-, mr : L[id])$ and $N' = (-, ma : L[id'])$, or
3. $N = (-, mr : L[id])$ and $N' = (-, um : L[id'])$, or
4. $N = (-, ma : L[id])$ and $N' = (-, mr : L'[id'])$, or
5. $N = (-, ma : L[id])$ and $N' = (-, ma : L'[id'])$, or
6. $N = (-, ma : L[id])$ and $N' = (-, um : L'[id'])$,

in which $L, L' \in \{P, O\}, L \neq L', id, id' \in \mathcal{ID} \setminus \{ID_0\}$. By Definition 7.2, the utterance $u_{id'} = \langle -, -, -, id' \rangle$ is related to the utterance $u_{id} = \langle -, -, -, id \rangle$ in cases 1, 3, 4 and 6. In cases 2 and 5, there is an utterance $u_t = \langle -, -, id, asm(-), t \rangle$, such that u_t is related to u_{id} and $u_{id'}$ is related to u_t . It can be seen that the previous construction of δ_i includes all utterances $u_{id}, u_{id'}$ and u_t but no other utterance. Hence δ_i is compatible with a related $\lambda_i \in \Lambda$.

δ_i is focused: Definition 10.1 defines two conditions for debate trees to be focused. Condition (a) in Definition 10.1 is met by 3(c) in Definition 11.2 as this ensures that, in a single tree, a proponent argument is only attacked by a single opponent argument. Condition (b) in Definition 10.1 is met by 3(a) in Definition 11.2 as this ensures that, when a proponent node is expanded such that it has a set of proponent children, then these children must be from a single utterance.

Proof of Theorem 11.1. Let $\mathcal{F}(\delta)$ be the debate forest drawn from δ and $\mathcal{F}(\delta)$ contains debate trees $\mathcal{T}(\delta_1), \dots, \mathcal{T}(\delta_n)$. Given Lemma 11.1, we know all trees in $\mathcal{F}(\delta)$ are debate trees drawn from some sub-dialogue of δ and each sub-dialogue is focused.

Since δ_i is focused, defensive and conflict-free, δ_i is g-successful and χ is grounded in \mathcal{AF}_i , the ABA framework drawn from δ_i (By Theorem 10.1). We need to show χ is also grounded in \mathcal{AF} , the ABA framework drawn from δ .

Since δ_i is g-successful in \mathcal{AF}_i , then it is not the case that there is an argument that attacks $\mathcal{DEF}(\mathcal{T}(\delta_i))$ that has not been countered attacked in the ABA

framework drawn from δ . By Definition 11.2 each debate tree in $\mathcal{F}(\delta)$ represents a set of arguments that support the claim of δ . Hence each tree contains its own set of defence set, i.e., $\mathcal{DEF}(\mathcal{T}(\delta_i))$ is different from $\mathcal{DEF}(\mathcal{T}(\delta_j))$, where $1 \leq i, j \leq n, i \neq j$. Therefore, if a defence set of a tree is grounded in the ABA framework drawn from the sub-dialogue, it is also grounded in the ABA framework drawn from the full-dialogue. Hence δ is g-successful as δ_i is.

Proof of Theorem 11.2. Similar to the proof of Theorem 11.1, we can show that properties of δ_i imply properties of δ .

Since δ is filtered, by Definition 11.4, we know that δ_i is filtered. Also, by Lemma 11.1, δ_i is focused. Hence, by Theorem 10.2, δ_i is a-successful in \mathcal{AF}_i (the ABA framework drawn from δ_i) and χ is admissible in \mathcal{AF}_i . Again, by the construction of debate forest (Definition 11.2), we know that each debate tree in the debate forest contains its own defence set, so arguments in other trees do not affect arguments in \mathcal{AF}_i . Hence, Hence δ is a-successful as δ_i is.

Proof of Theorem 11.3. Let $\mathcal{F}(\delta)$ be the debate forest drawn from δ and let $\mathcal{T}(\delta_1), \dots, \mathcal{T}(\delta_n)$ be the trees in $\mathcal{F}(\delta)$. By Lemma 11.1, all trees in $\mathcal{F}(\delta)$ are debate trees drawn from some sub-dialogue of δ and each sub-dialogue is focused.

Since δ is λ -ideal, by Definition 11.6, there is a debate tree $\mathcal{T}(\delta_i)$ such that none of the opponent arguments drawn from $\mathcal{T}(\delta_i)$ belongs to an admissible set of arguments in \mathcal{F}_δ .

By Lemma 10.2, there is an abstract dispute tree \mathcal{T}^a for $S \vdash \chi$ for some $S \subseteq \mathcal{A}$. Since δ is ideal, \mathcal{T}^a is ideal. By Theorem 3.4 in [15] the theorem holds.

References

- [1] *Argument & Computation, Special Issue: Tutorials on Structured Argumentation*, 5(1), 2014.
- [2] J. Alcolea-Banegas. Teaching argumentation theory and practice: The case of *12 angry men*. In *Tools for Teaching Logic*, pages 1–8. Springer, 2011.
- [3] L. Amgoud, S. Parsons, and N. Maudet. Modelling dialogues using argumentation. In *Proceedings of the 4th International Conference on Multi-Agent Systems*, pages 31–38, Boston, MA, 2000. IEEE Press.
- [4] Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

- [5] P. Besnard and A. Hunter. *Elements of Argumentation*. The MIT Press, 2008.
- [6] E. Black and A. Hunter. An inquiry dialogue system. *Autonomous Agents and Multi-Agent System*, 19:173–209, 2009.
- [7] A. Bondarenko, P. M. Dung, R. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1-2):63–101, 1997.
- [8] A. Bondarenko, F. Toni, and R. A. Kowalski. An assumption-based framework for non-monotonic reasoning. In *Proceedings of the second international workshop on logic programming and non-monotonic reasoning*, pages 171–189, Cambridge, MA, USA, 1993. MIT Press.
- [9] Carlos Iván Chesñevar, Jarred McGinnis, Sanjay Modgil, Iyad Rahwan, Chris Reed, Guillermo Ricardo Simari, Matthew South, Gerard Vreeswijk, and Steven Willmott. Towards an argument interchange format. *Knowledge Engineering Review*, 21(4):293–316, 2006.
- [10] Frank P. M. Dignum and Gerard A. W. Vreeswijk. Towards a testbed for multi-party dialogues. In *Advances in Agent Communication*, volume 2922 of *Lecture Notes in Computer Science*, pages 212–230. Springer, 2004.
- [11] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [12] P. M. Dung, R. A. Kowalski, and F. Toni. Assumption-based argumentation. In *Argumentation in AI*, pages 25–44. Springer, 2009.
- [13] P. M. Dung, F. Toni, and P. Mancarella. Some design guidelines for practical argumentation systems. In *Proceedings of the Third International Conference on Computational Models of Argument*, pages 183–194. IOS Press, 2010.
- [14] P.M. Dung, R.A. Kowalski, and F. Toni. Dialectic proof procedures for assumption-based, admissible argumentation. *Artificial Intelligence*, 170:114–159, 2006.
- [15] P.M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence, Special Issue on Argumentation in Artificial Intelligence*, 171(10–15):642–674, 2007.

- [16] X. Fan and F. Toni. Assumption-based argumentation dialogues. In *Twenty-Second International Joint Conference on Artificial Intelligence*, pages 198–203. IJCAI/AAAI, 2011.
- [17] X. Fan and F. Toni. A first step towards argumentation dialogues for discovery. In *Proceedings of the 1st International Workshop on Theory and Applications of Formal Argumentation*, pages 263–279. Springer, 2011.
- [18] X. Fan and F. Toni. Agent strategies for ABA-based information-seeking and inquiry dialogues. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 324–329. IOS Press, 2012.
- [19] X. Fan and F. Toni. Mechanism design for argumentation-based persuasion. In *Proceedings of the 4th International Conference on Computational Models of Argument*, pages 322–333. IOS Press, 2012.
- [20] X. Fan, F. Toni, A. Mocanu, and M. Williams. Multi-agent decision making with assumption-based argumentation. In *Proceeding of 13th International Conference on Autonomous Agents and Multiagent Systems*, 2014. To appear.
- [21] D. Gaertner and F. Toni. Hybrid argumentation and its properties. In *Proceeding of the 2008 conference on Computational Models of Argument*, pages 183–195. IOS Press, 2008.
- [22] Alejandro J. García and Guillermo R. Simari. Defeasible logic programming: An argumentative approach. *Theory and Practice of Logic Programming*, 4(2):95–138, January 2004.
- [23] T. F. Gordon. The pleadings game: an exercise in computational dialectics. *Artificial Intelligence & Law*, 2(4):239–292, 1994.
- [24] C. L. Hamblin. Mathematical models of dialogue. *Theoria*, 37(2):130–155, 1971.
- [25] A. Kakas and F. Toni. Computing argumentation in logic programming. *Journal of Logic and Computation*, 9:515–562, 1999.
- [26] J. Mackenzie. Four dialogue systems. *Studia Logica*, 49(4):567–583, 1990.

- [27] P. McBurney and S. Parsons. Dialogue games for agent argumentation. In I. Rahwan and G. R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 261–280. Springer, 2009.
- [28] Peter McBurney, David Hitchcock, and Simon Parsons. The eightfold way of deliberation dialogue. *International Journal of Intelligent Systems*, 22(1):95–132, 2007.
- [29] S. Parsons, P. McBurney, E. Sklar, and M. Wooldridge. On the relevance of utterances in formal inter-agent dialogues. In *Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems*, pages 47–62. Springer, 2007.
- [30] S. Parsons, M. Wooldridge, and L. Amgoud. Properties and complexity of some formal inter-agent dialogues. *Journal of Logic and Computation*, 13:347–376, 2003.
- [31] H. Prakken. Modelling reasoning about evidence in legal procedure. In *Proceedings of the 8th international conference on Artificial Intelligence and Law*, pages 119–128, New York, NY, USA, 2001. ACM.
- [32] H. Prakken. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation*, 15:1009–1040, 2005.
- [33] H. Prakken. Formal systems for persuasion dialogue. *Knowledge Engineering Review*, 21(2):163–188, 2006.
- [34] H. Prakken. An abstract framework for argumentation with structured arguments. *Argument & Computation*, 1(2):93–124, 2010.
- [35] H. Prakken, D. Ionita, and R. Wieringa. Risk assessment as an argumentation game. In *Proceedings of the 14th International Workshop on Computational Logic and Multi-Agent Systems*. Springer, 2013.
- [36] I. Rahwan, K. Larson, and F. A. Tohmé. A characterisation of strategy-proofness for grounded argumentation semantics. In *Twenty-First International Joint Conference on Artificial Intelligence*, pages 251–256. IJ-CAI/AAAI, 2009.
- [37] I. Rahwan and C. Reed. The argument interchange format. In Iyad Rahwan and Guillermo R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 383–402. Springer, 2009.

- [38] Iyad Rahwan and Guillermo R. Simari. *Argumentation in Artificial Intelligence*. Springer, 2009.
- [39] P.M. Thang, P.M. Dung, and N.D. Hung. Towards argument-based foundation for sceptical and credulous dialogue games. In *Proceeding of the 2012 conference on Computational Models of Argument*, pages 398–409. IOS Press, 2012.
- [40] F. Toni. Reasoning on the web with assumption-based argumentation. In Thomas Eiter and Thomas Krennwallner, editors, *Reasoning Web. Semantic Technologies for Advanced Query Answering*, volume 7487 of *Lecture Notes in Computer Science*, pages 370–386. Springer, 2012.
- [41] F. Toni. A generalised framework for dispute derivations in assumption-based argumentation. *Artificial Intelligence*, 195:1–43, 2013.
- [42] Francesca Toni. A tutorial on assumption-based argumentation. *Argument & Computation, Special Issue: Tutorials on Structured Argumentation*, 5(1):89–117, 2014.
- [43] D. Walton and E. Krabbe. *Commitment in Dialogue: Basic concept of interpersonal reasoning*. State University of New York Press, 1995.
- [44] M Wooldridge. *An Introduction to Multiagent Systems*. Wiley, Chichester, UK, 2. edition, 2009.