

Shift from Forward to Backward Deliberation in Search of Reconciliation

Hiroyuki Kido¹ and Federico Cerutti²

¹ School of Engineering, The University of Tokyo, Japan
kido@sys.t.u-tokyo.ac.jp

² Department of Computing Science, The University of Aberdeen, UK
f.cerutti@abdn.ac.uk

Abstract. Desire conflicts arise in several real-world contexts. In this paper we propose a mixed deliberation dialogue for reconciliation. A mixed deliberation dialogue is defined as a combination of *forward* and *backward* deliberation dialogues whose goals are *subordinate* and *superordinate* desires of a given desire, respectively. This research and the introduction of mixed deliberation dialogue have been motivated by Kowalski and Toni's reconciliatory scenario: indeed we show that an instantiation of a mixed deliberation dialogue implements key parts of Kowalski and Toni's reconciliatory solution. We also proved the correctness of the mixed deliberation dialogues.

1 Introduction

Dialogue theory encompasses various types of descriptive and formal studies, aimed at various purposes, on the structure of dialogues (van Eemeren et al., 1996). For example, Hintikka's game-theoretic semantics (Hintikka, 1968) and Lorenzen's dialogue logic (Lorenzen, 1961) explore semantics of language, and on the other hand, Hamblin's formal dialectics (Hamblin, 1970) explores descriptive or formal dialogue systems. Since formal dialogue systems can give agents rational interaction and computation mechanisms under uncertain, incomplete, inconsistent, subjective and distributed information, it has received attention from researchers working on formal argumentation (Prakken, 2006; Kok et al., 2010; Prakken, 2005; Fan and Toni, 2012; Wells and Reed, 2006). However, little work has been done for dialogue systems for reconciling conflict by not only searching for means of satisfying either all or parts of given desires, but also searching for means for satisfying their underlying desires behind the given ones. To the best of our knowledge, Kowalski and Toni (Kowalski and Toni, 1994) first argued for the need for reconciliation in the context of argumentation.

Example 1 ((Kowalski and Toni, 1994)). In a recent head-of-sections committee meeting in our Department, we discussed the composition of a new resources committee. Two conflicting arguments were put forward. The Director of Administration argued that, in the interests of efficiency, the members of the new committee should consist of himself and the other principal administrative officers of the Department. The Director of Research argued, in opposition to him, that, in the interests of democracy, the committee should also contain members elected by the Department.

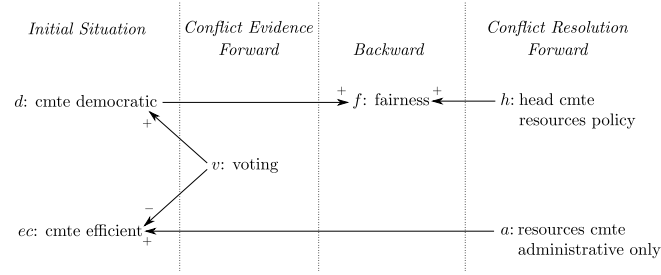


Fig. 1. Rough causal relations behind Kowalski and Toni’s reconciliatory scenario

During the course of the discussion it became clear that the two sides were focusing on different assumptions about the purpose of the new committee: the Director of Administration on its purely administrative function, and the Director of Research on its presumed policy making nature. These two assumptions could be viewed as conflicting solutions to the more general goals of deciding, on the one hand, which group should administer resources, and on the other hand, which group should make policy about resources.

By focusing on the more general goals, it was possible to identify a new solution which was acceptable to both parties: the resources committee will administer resources, whereas the head-of-sections committee will make policy about resources. In the interests of efficiency, the members of the resources committee will consist of administrative officers only. In the interests of democracy, the head-of-sections committee will represent the views and interests of the various Department sections on matters concerning policy about the allocation of resources.

What is interesting about their scenario is that neither the generalized goal nor reconciliatory solution can be obtained by just choosing one of the given alternatives based on utilities or preferences. An open question is how one should have a dialogue to reach the generalized goal and the reconciliatory solution.

These observations motivate us to formalize reconciliatory dialogues as consisting of *forward* and *backward* deliberation dialogues. We show how shifting between *forward* and *backward* deliberation dialogue highlights the linkage with more general goals. Indeed, on one hand, the Director of Administration desires *efficiency*; on the other hand, the Director of Research wants to promote *democracy* via *voting* although this demotes *efficiency* — *forward deliberation*. Then, since one of the outcomes of adopting *democracy* is *fairness*, and since there is no reason in favour of not desiring *fairness*, we can consider it a “desirable” outcome of the *democracy* — *backward deliberation*. Finally, shifting back to a *forward* deliberation dialogue, the two more general goals become evident, because *efficiency* will be improved if the resources committee is composed by administrative officers only, and *fairness* will be ensured if the head committee will take care of making policies about resources.

Example 2 (continued). According to Fig. 1:

- *d* represents “the resources committee is democratic”;
- *ec* represents “the resources committee is efficient”;

- v represents “the resources committee is elected”;
- f represents “the policy is fair”;
- h represents “the head committee makes policy about resources”;
- a represents “the resources committee is composed by administrative officers only”.

This paper contributes to the state of the art of studies on formal dialogue and argumentation. In fact, to the best of our knowledge, this paper first handles the processes leading from conflict detection to justification of reconciliation in terms of a series of dialogues. In particular, first of all, this paper gives underlying dialogue and inference principles behind reconciliation, and second of all, we address Kowalski and Toni’s academically-acknowledged scenario that cannot be solved by simply taking advantage of utilities or preferences that many formal dialogue-based and argumentation-based approaches assume.

2 Preliminaries

Dung’s theory of acceptability semantics (Dung, 1995) reformulates consequence notions of various kinds of nonmonotonic logics. The semantics is defined on a pair $AF = \langle AR, attacks \rangle$, called an abstract argumentation framework, where AR is a set of arguments and $attacks$ is a binary relation on AR , i.e., $attacks \subseteq AR \times AR$.

Definition 1 (Grounded extension). (Dung, 1995) Let $AF = \langle AR, attacks \rangle$ be an abstract argumentation framework, $S \subseteq AR$ and $a \in AR$.

- S is conflict-free iff there are no $a, b \in S$ such that $a attacks b$, i.e., $(a, b) \notin attacks$.
- a is acceptable with respect to S iff, for all $b \in AR$, if $b attacks a$ then there is $c \in S$ such that $c attacks b$.
- The characteristic function, $F_{AF} : \mathfrak{P}(AR) \rightarrow \mathfrak{P}(AR)$, is defined as $F_{AF}(S) = \{a \in AR \mid a \text{ is acceptable with respect to } S\}$.
- S is the grounded extension iff it is the least fixed point of F_{AF} .

Given an argumentation framework $AF = \langle AR, attacks \rangle$, an argument $a \in AR$ is *justified* in AF iff a is in the grounded extension of AF , and *overruled* otherwise.

Logical languages and rules of inference instantiate abstract argumentation frameworks. \mathcal{L}_0 is a language of modal propositional logic with single modal operator D representing the operator “it is desirable that”. \mathcal{L}_0 is closed under truth-functional operations. \mathcal{L}_0 conforms to the axiomatic system KD in which $D(a \rightarrow b) \rightarrow Da \rightarrow Db$ and $Da \rightarrow \neg D\neg a$ are axioms. \mathcal{L}_1 is a language defined on \mathcal{L}_0 . \mathcal{L}_1 consists of so-called defeasible conditionals, or defaults. They commonly have the forms “ $a \Rightarrow b$ ” where a is a conjunction of literals, i.e., atomic propositions or their negation, in \mathcal{L}_0 and b is a literal in \mathcal{L}_0 , and mean that if a is the case then b is normally the case. Operator D is assumed not to appear in defeasible conditionals. We assume a fixed, but arbitrary theory $T \subseteq \mathcal{L}_0 \cup \mathcal{L}_1$.

Example 3 (continued). To illustrate our proposal, let us consider the following additional propositions:

- p represents “the resources committee is dictatorial”;
- i represents “the policy is effective”;

4 H. Kido, F. Cerutti

- er represents “the resources is efficient”;
- m represents “the resources committee makes policy about resources”;
- r represents “the head of sections committee represents various views of interests”

Moreover let us consider the following rules — which provide the casual relationships depicted in Figure 1:

$$\begin{aligned} v \Rightarrow d; \quad v \Rightarrow \neg ec; \quad a \Rightarrow ec; \quad d \wedge m \Rightarrow f; \\ d \wedge m \Rightarrow \neg i; \quad p \wedge m \Rightarrow i; \quad h \wedge r \Rightarrow f. \end{aligned}$$

Any rules of inference that are not valid — w.r.t. modal logic KD — are called defeasible inference rules and they are represented by \rightsquigarrow . We use letters a, b, c, \dots of the alphabet to represent literals in $\mathcal{L}_0 \cup \mathcal{L}_1$, lower-case greek letters $\alpha, \beta, \gamma, \dots$ to represent metavariables of them and A, B, C, \dots to represent sequences of them.

Definition 2 (Forward practical syllogisms). (Bench-Capon and Prakken, 2006) *Positive forward practical syllogism, denoted by PFPS, and negative forward practical syllogism, denoted by NFPS, are defined as follows, respectively*³.

$$\begin{aligned} \text{PFPS} : D\alpha, \gamma, \beta \wedge \gamma \Rightarrow \alpha \rightsquigarrow D\beta \\ \text{NFPS} : D\alpha, \gamma, \beta \wedge \gamma \Rightarrow \neg\alpha \rightsquigarrow D\neg\beta \end{aligned}$$

The positive one intuitively states that if one believes that α is desirable, γ is the case and if β is realized under γ is the case then α is satisfied, then one defeasibly concludes that β is desirable. The negative one, however, concludes that β is undesirable from a different premise stating that if β is realized under γ is the case then α is frustrated. We describe them as *FPS* without distinction. We often silently use abbreviated forms $D\alpha, \beta \Rightarrow \alpha \rightsquigarrow D\beta$ and $D\alpha, \beta \Rightarrow \neg\alpha \rightsquigarrow D\neg\beta$.

Definition 3 (Argument). Let $\Sigma \subseteq T$, $p_i \in \mathcal{L}_0 \cup \mathcal{L}_1$. A sequence $A = p_1, p_2, \dots, p_n$ is an argument from Σ to p_n iff (1), for all $i(1 \leq i \leq n)$, $p_i \in \Sigma$ (base case) or p_i is derived from preceding formulae $p_j(j < i)$ by applying a rule of inference, and (2), for all $i(1 \leq i \leq n)$, the sequence obtained by eliminating p_i from A does not satisfy (1).

The first condition states that an argument is a deduction from Σ . The second assures that an argument consists of a minimal number of formulae. In this paper, $\Sigma \vdash_x p_n$ denotes that there is an argument from Σ to p_n where rules of inference are restricted to only x . For example, $\Sigma \vdash_{\text{PFPS}} Da$ represents that Da is derived from Σ by applying only positive or negative forward practical syllogisms zero or more time.

3 Formalizing Goals of Dialogues

This section aims to give formal definitions of dialogue goals. Given a desirable outcome or desire, we use term *superordinate desire* to refer to a desire such that, once it is assumed, then a given desire is actually derived as desirable, but no other desire is derived as undesirable. This idea is embodied by using the forward practical syllogisms.

³ The authors originally call them positive and negative practical syllogisms, respectively, and the conclusion of the negative one has $\neg D\beta$.

Definition 4 (Superordinate desires). Let $Dg, Dh \in \mathcal{L}_0$. Dh is a superordinate desire of Dg in T iff there is $\Sigma \subseteq T$ such that $\Sigma \not\vdash_{FPS} Dg$ and $\Sigma \cup \{Dh\} \vdash_{FPS} Dg$, and $T \cup \{Dh\} \not\vdash_{FPS} D\neg t$, for all $Dt \in T \cup \{Dg\}$.

Example 4 (continued). Df is a superordinate desire of Dd in $T_1 = \{d \wedge m \Rightarrow f, d \wedge m \Rightarrow \neg i, p \wedge m \Rightarrow i, Dp, m\}$. However, $D\neg i$ is not because, given $\Sigma_1 = \{d \wedge m \Rightarrow \neg i, m\} \subseteq T_1$, although $\Sigma_1 \cup \{D\neg i\} \vdash_{FPS} Dd$ is the case, but $T_1 \cup \{D\neg i\} \vdash_{FPS} D\neg p$ is the case.

Given a desire, we use the term *subordinate desire* to refer to a desire such that if it is realized then the given desire is satisfied, but no other desires are frustrated. Namely, a subordinate desire of a desire is desirable as a means of satisfying the desire, but not as a means of frustrating any other desires.

Definition 5 (Subordinate desires). Let $Dg, Dh \in \mathcal{L}_0$. Dg is a subordinate desire of Dh in T iff there is $\Sigma \subseteq T$ such that $\Sigma \not\vdash_{FPS} Dg$ and $\Sigma \cup \{Dh\} \vdash_{FPS} Dg$, and $T \cup \{Dh\} \not\vdash_{FPS} D\neg g$.

Example 5 (continued). Da is a subordinate desire of Dec in $T_2 = \{v \Rightarrow d, v \Rightarrow \neg ec, a \Rightarrow ec, Dv\}$. However, $D\neg v$ is not because, given $\Sigma_2 = \{v \Rightarrow \neg ec\}$, although $\Sigma_2 \cup \{Dec\} \vdash_{FPS} D\neg v$, but $T_2 \cup \{Dec\} \vdash_{FPS} Dv$.

Finally, we introduce the notion of *reconciliatory desires* defined by combining superordinate desires and subordinate desires. As we will see in the next section, reconciliatory desires have a significant role in reconciliation.

Definition 6 (Reconciliatory desires). Let $Dg, Dh \in \mathcal{L}_0$. Dg is a reconciliatory desire of Dh in T iff there is $Di \in \mathcal{L}_0$ such that Di is a superordinate desire of Dh in T and Dg is a subordinate desire of Di in T .

Example 6 (continued). $T = \{d \wedge m \Rightarrow f, d \wedge m \Rightarrow \neg i, p \wedge m \Rightarrow i, m, Dp, h \wedge r \Rightarrow f, r\}$. Dh is a reconciliatory desire of Dd in T because Df is a superordinate desire of Dd and Dh is a subordinate desire of Df .

4 Formalizing Deliberation Dialogues

4.1 Backward Practical Syllogisms

We introduce a *backward* version of the practical syllogism. This inference pattern, differently from the traditional forward practical syllogism, represents the result of a critical thinking approach. Indeed, let us suppose that α is a desirable outcome ($D\alpha$), and it is true that $\alpha \Rightarrow \beta$. Therefore, if α is accepted as desirable, β , as it is a material implication deriving from α , should be considered. What the backward practical syllogism is doing here is to suggest to explore the world by defeasibly assuming that $D\beta$ holds too. Clearly, this is the case only if there is no evidence of the contrary — i.e. that $D\neg\beta$ holds. Hereafter, \sim represents *negation as failure*: i.e. $\sim\chi$ is true either $\neg\chi$ or if each attempt to make χ true fails.

Definition 7 (Backward practical syllogisms). *Positive backward practical syllogism, denoted by PBPS, and negative backward practical syllogism, denoted by NBPS, are defined as follows, respectively:*

$$\begin{aligned} \text{PBPS} &: D\alpha, \sim D\neg\beta, \gamma, \alpha \wedge \gamma \Rightarrow \beta \rightsquigarrow D\beta \\ \text{NBPS} &: D\alpha, \sim D\beta, \gamma, \neg\alpha \wedge \gamma \Rightarrow \beta \rightsquigarrow D\neg\beta \end{aligned}$$

The positive backward practical syllogism intuitively states that if one believes that α is desirable, γ is the case and if α is realized under γ is the case then β is realized, then one defeasibly concludes that β might be desirable unless there are reasons for believing otherwise. The negative one, however, concludes that β is undesirable — unless proven otherwise — from a different premise stating that if α is not realized under γ is the case then β is realized. We describe them as *BPS* without distinction.

4.2 General Elements of Dialogues

Players exchange moves during a dialogue. We consider two players, the *proponent* — who chooses the *topic* of the dialogue — and the *opponent*. Each player can exchange a single move during its own turn (*turntaking rule*). Each move consists of a *speech act* — the content of the move — and a *type* of dialogue — the context in which moves are put forward.

Definition 8 (Move). *Let $a \in \mathcal{L}_0 \cup \mathcal{L}_1$: a move is a tuple $\langle \text{speech act}, \text{type} \rangle$ where $\text{speech act} \in \{\text{claim}(D\alpha), \text{why}(\alpha), \text{since}(\Phi \rightsquigarrow \alpha), \text{fact}(\alpha)\}$ and $\text{type} \in \{B, F\}$, i.e., backward deliberation dialogue or forward deliberation dialogue.*

Let us define the set of allowed replies to a move.

Definition 9 (Allowed replies). *Let M be a set of moves and $X \in \{B, F\}$: the following table depicts the allowed reply to the locutions.*

Locutions	Replies
$\langle \text{claim}(D\alpha), X \rangle$	$\langle \text{why}(D\alpha), X \rangle, \langle \text{claim}(D\neg\alpha), X \rangle$
$\langle \text{why}(\alpha), X \rangle$	$\langle \text{since}(\Phi \rightsquigarrow \alpha), X \rangle, \langle \text{fact}(\alpha), X \rangle$
$\langle \text{since}(\Phi = \{\dots, \beta, \dots\} \rightsquigarrow \alpha), X \rangle$	$\langle \text{why}(\beta), X \rangle$
$\langle \text{since}(\Phi = \{\dots, \sim\beta, \dots\} \rightsquigarrow \alpha), X \rangle$	$\langle \text{since}(\Psi \rightsquigarrow \beta), X \rangle, \langle \text{fact}(\beta), X \rangle$
$\langle \text{fact}(\alpha), X \rangle$	

In the following, if $m \in M$ is a reply to $n \in M$, then we will say that m *attacks* n or $m \longrightarrow n$. In particular, if $n = \langle \text{claim}(D\alpha), X \rangle$ and $m = \langle \text{claim}(D\neg\alpha), X \rangle$, then $m \longrightarrow n$ and $n \longrightarrow m$. Note that no move attacks another whose types are different. A dialogue framework is defined as an instantiation of an abstract argumentation framework whose arguments and attacks are instantiated by moves and attacks on the set of moves, respectively.

Definition 10 (Dialogue frameworks). *A dialogue framework is a pair $DF = \langle M, \text{attacks} \rangle$ where M is a set of moves and $\text{attacks} = \{\langle m, n \rangle \mid m, n \in M, m \longrightarrow n\}$.*

We use Dung's acceptability semantics to evaluate acceptable moves in dialogue frameworks because it is rational to think that moves successfully replying to critical questions are worthy of acceptance. In general, dialogue frameworks are constructed by multiple agents who freely participate and make moves from their private knowledge bases. This knowledge is invisible to others and they can only see what they said during a dialogue.

Definition 11 (Collaborative theory). Let DF be a dialogue framework. A collaborative theory built from DF , denoted by $T(DF)$, is the set $\{a \in \mathcal{L}_0 \cup \mathcal{L}_1 \mid \text{There is a move } m \text{ in } DF \text{ whose speech act is } \text{fact}(a)\}$.

4.3 Forward, Backward and Mixed Deliberation Dialogues

A forward (resp. backward) deliberation protocol is defined using FPS (resp. BPS).

Definition 12 (Forward (resp. backward) deliberation protocol). Let M be a set of moves and \mathcal{DF} be a set of dialogue frameworks and $Dh \in \mathcal{L}_0$. A forward (resp. backward) deliberation dialogue protocol is a function $P_F : \mathcal{DF} \rightarrow 2^M$ (resp. $P_B : \mathcal{DF} \rightarrow 2^M$) where $m \in P_F(DF)$ (resp. $m \in P_B(DF)$) if and only if:

- $m = \langle \text{claim}(Dh), F \rangle$ (resp. $m = \langle \text{claim}(Dh), B \rangle$), if $DF = \emptyset$;
- $m \notin DF$ and $\exists n \in DF$ such that $m \longrightarrow n$, if $DF \neq \emptyset$.

Moreover, if the speech act of m is $\text{since}(A \rightsquigarrow a)$ then $A \rightsquigarrow a$ is an application of forward (resp. backward) practical syllogisms.

A dialogue framework DF is a forward (resp. backward) deliberation dialogue iff DF is constructed by the forward (resp. backward) deliberation protocol. The first move of DF is called the *subject* of DF .

Moreover, let us notice that Definition 12, in conjunction with Definitions 9 and 10, makes redundant the defeasible premise in the backward practical syllogism (Def. 7). Indeed, let us suppose there is a move $m = \langle \text{claim}(D\beta), B \rangle$ which is justified using an instance of the BPS , namely $D\alpha, \sim D\neg\beta, \gamma, \alpha \wedge \gamma \Rightarrow \beta \rightsquigarrow D\beta$. According to Def. 12, to prove $D\neg\beta$, there must be a move $n = \langle \text{claim}(D\neg\beta), B \rangle$: from Defs. 9 and 10, $m \longrightarrow n$ and $n \longrightarrow m$. According to the Grounded semantics (cf. Def. 1), if m (resp. n) is justified, then n (resp. m) is not. This is equivalent to saying that $D\alpha$ is an acceptable desire if we prove that $D\neg\alpha$ is not. This is a stronger requirement than defeasible premise of Definition 7, therefore hereafter we silently drop the defeasible premise (i.e. $\sim D\neg\beta$) in BPS .

We now define mixed dialogues as a combination of forward and backward deliberation dialogues where agents search for reconciliatory desires and justify them. We give claim moves in forward deliberation dialogues a role to shift dialogue types from backward to forward deliberation dialogues. Formally, a move $\langle \text{claim}(D\alpha), B \rangle$ is allowed to attack $\langle \text{why}(D\alpha), F \rangle$ in mixed dialogues.

Definition 13 (Mixed deliberation protocol). Let N be a set of moves and \mathcal{DF} be a set of dialogue frameworks. A mixed deliberation protocol is a function $P_M : \mathcal{DF} \rightarrow 2^N$ where $m \in P_M(DF)$ iff $m \in P_F(DF) \cup P_B(DF)$ or, $m = \langle \text{claim}(Da), B \rangle$ and $\langle \text{why}(Da), F \rangle$ is in DF .




Fig. 2. A mixed deliberation dialogue with the justified subject $\langle \text{claim}(Dh), F \rangle$

A dialogue framework DF is a *mixed deliberation dialogue* iff DF is constructed by the mixed deliberation protocol.

Example 7 (continued). Figure 2 shows an example of a mixed deliberation dialogue DF where the subject $\text{claim}(Dh)$ is justified. The collaborative theory built from DF is $T(DF) = \{Dd, m, d \wedge m \Rightarrow f, r, h \wedge r \Rightarrow f\}$.

Due to limitations of space, we only show the relationships between acceptability status of mixed deliberation dialogues and reconciliatory desires defined on a collaborative theory. It corresponds to show correctness of these dialogues in terms of whether status of subjects is consistent with these desires. We say a dialogue framework is finite if the number of moves in it is finite.

Definition 14 (Well dialogue framework). Let DF be a finite dialogue framework and P be a protocol. DF is well iff there is no $m \in P(DF)$ such that the speech act of m is $\text{why}(a)$ or $\text{since}(A \rightsquigarrow a)$ where, for all $c \in A$, there is a move n in DF such that the speech act of n is $\text{since}(B \rightsquigarrow b)$ and $c \in B$.

We show that mixed deliberation dialogues are *sound* in the sense that the justified subjects are necessarily reconciliatory desires in the collaborative theories and *complete* in the sense that subordinate desires in the collaborative theories are necessarily the justified subjects.

Theorem 1. Let $X \in \{F, B\}$ and DF be a well mixed deliberation dialogue whose subject is $\langle \text{claim}(Dh), X \rangle$. There is $Dg \in T(DF)$ such that Dh is a reconciliatory desire of Dg in $T(DF)$ iff $\langle \text{claim}(Dh), X \rangle$ is justified in DF .

5 Conclusions and Future Work

We proposed a mixed deliberation dialogue for reconciliatory desires. A mixed deliberation dialogue is defined as a combination of forward and backward deliberation whose goals are subordinate and superordinate desires, respectively. We showed correctness of dialogues based on the fact that the subject of well mixed deliberation dialogue is justified iff the subject is a reconciliatory desire of a desire in the collaborative theory. Weakness of our formalization is, first, that it restricts inference mechanisms to only forward and backward practical syllogisms, and, second, that it does not allow agents to challenge or rebut facts put forward in dialogues. We think that our formalization can be served as a foundation for overcoming these limitations.

Bibliography

- T. J. M. Bench-Capon and H. Prakken. Justifying actions by accruing arguments. In *Proc. of the 1st International Conference on Computational Models of Argument*, pages 247–258, 2006.
- P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n -person games. *Artificial Intelligence*, 77: 321–357, 1995.
- X. Fan and F. Toni. Argumentation dialogues for two-agent conflict resolution. In *Proc. of the 4th International Conference on Computational Models of Argument*, pages 249–260, 2012.
- C. L. Hamblin. *Fallacies*. Methuen, 1970.
- J. Hintikka. *Language-Games for Quantifiers*, pages 46–72. Oxford: Blackwell, in american philosophical quarterly monograph series 2: studies in logical theory edition, 1968.
- E. M. Kok, J. J. C. Meyer, H. Prakken, and G. A. W. Vreeswijk. A formal argumentation framework for deliberation dialogues. In *Proc. of the 7th International Workshop on Argumentation in Multi-Agent Systems*, pages 31–48, 2010.
- R. A. Kowalski and F. Toni. Argument and reconciliation. In *Proc. of the 5th Generation Computer Systems Workshop on Application of Logic Programming to Legal Reasoning*, pages 9–16, 1994.
- P. Lorenzen. *Ein dialogisches konstruktivitätskriterium*, pages 193–200. Oxford: Pergamon, in infinitistic methods edition, 1961.
- H. Prakken. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation*, 15:1009–1040, 2005.
- H. Prakken. Formal systems for persuasion dialogue. *The Knowledge Engineering Review*, 21(2):163–188, 2006.
- F. H. van Eemeren, R. Grootendorst, and F. S. Henkemans. *Fundamentals of Argumentation Theory: A Handbook of Historical Backgrounds and Contemporary Developments*. Routledge, 1996.
- S. Wells and C. Reed. Knowing when to bargain. In *Proc. of the 1st International Conference on Computational Models of Argument*, pages 235–246, 2006.