

Argument-based Multi-Issue Negotiation

Thalya Fossey¹, Jean-Guy Mailly² and Pavlos Moraitis^{1,3}

¹Université Paris Cité, LIPADE, F-75006 Paris, France

²Université Toulouse Capitole, IRT, Toulouse, France

³Argument Theory, Paris, France

{thalya.fossey, pavlos.moraitis}@u-paris.fr, jean-guy.mailly@irit.fr

Abstract

Automated negotiation aims at finding agreements between agents with conflicting goals. Existing utility-based approaches guarantee agents satisfaction with negotiation outcomes, especially in multi-issue negotiations where concession mechanisms lead to win-win results. However, they lack explainability and do not consider agents' beliefs. On the other hand, argument-based approaches provide reasons for accepting or rejecting offers but do not include utility modeling for offers or enable concession mechanisms in multi-issue settings. We propose a novel hybrid approach combining both types of approaches. The utility-based component enables agents to make concessions on complex negotiation objects to achieve win-win outcomes, while the argumentation component ensures that accepted offers align with the agents' personal argumentation theories. These theories represent their beliefs, encoding various profiles, ethical considerations, social norms, or legal principles.

1 Introduction

Automated negotiation offers a promising solution for addressing real-life conflicts with improved efficiency (e.g., [Luo *et al.*, 2024]). The goal of a negotiation dialogue is to enable interacting agents to resolve conflicts and arrive at a mutually acceptable agreement. These conflicts revolve around the object of negotiation (e.g., a product such as a car), which may be defined by a single issue, such as price (single-issue negotiations), or by multiple issues, such as price, color, brand, etc. (multi-issue negotiations).

Over the past twenty-five years, computational argumentation has been widely applied to model negotiation dialogues, particularly in the context of single-issue negotiations (see, e.g., [Rahwan *et al.*, 2003; Dimopoulos and Moraitis, 2014] for a survey).

In argumentation-based negotiation (ABN), agents choose offers likely to be accepted and exchange arguments supporting them, based on their theories (e.g., [Amgoud *et al.*, 2007; Amgoud and Kaci, 2006; Kakas and Moraitis, 2006; Dung *et al.*, 2008; Parsons *et al.*, 1998; Hadidi *et al.*, 2010;

Casali *et al.*, 2016]), on their opponent's profile (e.g., [Hadidi *et al.*, 2012; Pilotti *et al.*, 2015; Bonzon *et al.*, 2012]), or both as in [Dimopoulos *et al.*, 2019; Dimopoulos *et al.*, 2021], where an agent first uses its theory to choose the best offer and then uses its incomplete theory about its opponent to find supporting arguments.

Opponent modeling is crucial for negotiation (see e.g. [Baarslag *et al.*, 2016b]) as it enhances the capacity to achieve win-win agreements, especially in cooperative settings.

Learning an opponent's profile involves understanding its acceptance and bidding strategies, deadlines, and preferences. Most proposed works use learning techniques (e.g. [Baarslag *et al.*, 2016a]) better suited for game-theoretic (or utility-based) negotiations. Several methods exist for learning the opponent's preference profile, categorized as Bayesian filters (e.g. [Buffett and Spencer, 2005; Buffett and Spencer, 2007]), perceptron models (e.g. [Zafari and Nassiri-Mofakham, 2016; Pocola, 2022]), and frequency models (e.g. [van Galen Last, 2012; van Krimpen *et al.*, 2013; Tunali *et al.*, 2017; Hosokawa and Fujita, 2020; Pocola, 2022]).

In argumentation-based negotiations two recent works [Dimopoulos *et al.*, 2019; Dimopoulos *et al.*, 2021] propose a different approach to represent the uncertain knowledge an agent has on its opponent by using the Control Argumentation Frameworks (CAFs) from [Dimopoulos *et al.*, 2018].

Utility-based negotiation approaches [Zhan *et al.*, 2024] are well-suited for multi-issue negotiations as they allow for trade-offs, increasing the ability to reach an agreement because the importance of the issues is generally not the same between the negotiating agents. However, they lack information exchange, leading agents to refuse offers they could have accepted if they had information justifying them.

Argumentation-based negotiation (ABN) approaches [Rahwan *et al.*, 2003; Dimopoulos and Moraitis, 2014] overcome the drawback of not allowing agents to justify offers by enabling them to exchange arguments and counterarguments, potentially convincing opponents to accept offers. However, finding a trade-off between multiple issues in multi-issue negotiations is challenging. This is because arguments supporting different issues may express compelling reasons for proposed values, making it difficult to find the best compromise using only qualitative information. Consequently, there is currently no argumentative approach

for multi-issue negotiations based solely on arguments for the choice of offers (i.e. the choice of the values of an issue).

This work advances automated negotiation by combining the advantages of previous approaches. It uses a utility-based part based on frequency models to learn the opponent's profile regarding issues weights and values utilities, seeking trade-offs to increase compromises and win-win situations. This part considers both the proponent and opponent's profiles. It also offers an argumentation part based on [Dimopoulos *et al.*, 2019; Dimopoulos *et al.*, 2021] to represent the opponent's uncertain profile in relation to the practical arguments supporting the values of the issues and the epistemic arguments encoding its beliefs. It also follows this work by choosing the arguments supporting the offer calculated by the utility-based part in the opponent's theory, which has been proven to increase the number of agreements.

However, this argumentation part is used differently from the way it is used in those works and more generally in the other traditional approaches. Indeed, as mentioned above, in traditional approaches, arguments are used by the agents to choose the offers by using their own theories, those of the opponents or both, but also to defend the arguments which support these offers against the arguments which attack them. In our approach, the choice of offers (i.e. the values of the different issues to be negotiated) is handled by the utility-based part. Arguments serve to guarantee that the chosen offers are consistent with the beliefs of the agents, which may encode different profiles, ethical behaviors, social or legal norms, etc., while serving to support the offers and to convince the opponent (as in traditional approaches) in case an inconsistency is detected. This is done by ensuring that the supporting arguments whose role in our approach is to justify the values of the issues (and not to choose them), are acceptable arguments in the theories of negotiating agents.

2 Frequency Models for Multi-Issue Negotiations

Among the different approaches mentioned above concerning opponent learning, we decided in this work to use a frequency model. Indeed, as we work on negotiations which are not executed several times, using the perceptron models (requiring training examples) is not suitable. Frequency and Bayesian models are two commonly used methods, but we decided to go with a frequency method. However, our approach is modular and someone else could integrate and use a Bayesian model.

Frequency models [van Galen Last, 2012; van Krimpen *et al.*, 2013; Tunalı *et al.*, 2017; Hosokawa and Fujita, 2020; Pocola, 2022] allow the agent to learn both the weight of the issues and the utility of the values based on the frequency of their opponent's bids. It relies on the hypothesis that an agent will make its best bids in the beginning of the negotiation and will make concessions over its least preferred issues first. The idea is that an agent will change first the value of the less preferred issue [van Krimpen *et al.*, 2013]. So, the weight of the issues that have the same value from one offer to the next will be increased by a coefficient ϵ , while the weight of the issues for which the value changes will be decreased, such

that the sum of the weights of the issues is still equal to 1.

Following the same principle, if an agent picks the same value for an issue many times, this value has a high utility. This utility is computed using a function based on the frequency with which the value has been proposed by the opponent. A widely adopted function for computing the utility of values, as outlined in [Tunalı *et al.*, 2017], is defined as follows:

$$u_i(j) = \frac{(1 + \sum_{o \in \mathcal{O}_{1 \rightarrow t}} \delta_i(j, o))^\gamma}{\max_{k \in \mathcal{V}_i} (1 + \sum_{o \in \mathcal{O}_{1 \rightarrow t}} \delta_i(k, o))^\gamma} \quad (1)$$

where $\mathcal{O}_{1 \rightarrow t}$ represents the set of offers made by the opponent up to time t and $\delta_i(j, o)$ equals 1 if value j appears in offer o , and 0 otherwise. γ is an exponential filter applied to moderate the growth of unbalanced values when the opponent repeatedly sends the same offer.

3 Background on Abstract Argumentation

We briefly recall basics of abstract argumentation [Dung, 1995]. A (finite) *abstract Argumentation Framework* (AF) is a directed graph $\langle A, R \rangle$ with A the set of arguments and $R \subseteq A \times A$ the attack relation. Standard reasoning with AFs is based on *extensions*, i.e. sets of collectively acceptable arguments. Here, we focus on *stable* extensions, i.e. sets of arguments $S \subseteq A$ s.t. $\forall a, b \in S, (a, b) \notin R$, and $\forall b \in A \setminus S, \exists a \in S$ s.t. $(a, b) \in R$. We write $\text{st}(\langle A, R \rangle)$ for the set of stable extensions of the AF $\langle A, R \rangle$. Our approach could be adapted for other semantics in the literature; see e.g. [Dung, 1995; Baroni *et al.*, 2011] for more details on these semantics.

Our approach relies on Control Argumentation Frameworks (CAFs) [Dimopoulos *et al.*, 2018], which generalize AFs by integrating several kinds of uncertainties in the framework (about the existence of an argument, the existence of an attack, or the direction of an attack), and by adding a set of *control arguments and attacks* allowing an agent to influence the acceptability of some target arguments, whatever the actual state of the uncertain parts of the framework.¹ More details are provided in the supplementary material.¹

4 Negotiation Framework

4.1 Negotiating Agent

We study bilateral agent negotiation over objects characterized by several issues. We formally define the constructs of this negotiation.

Definition 1 (Negotiation object). *A negotiation object x is a finite set $\mathcal{I}(x)$ of issues (or attributes) characterizing x . Each issue $i(x) \in \mathcal{I}(x)$ takes values in a finite domain $\mathcal{V}_{i(x)} = \{v_1^i, \dots, v_m^i\}$.*

Notice that for two issues $i(x) \neq j(x)$ the value of m may differ. If there is no risk of ambiguity about the negotiation object, we simply write i and \mathcal{V}_i instead of $i(x)$ and $\mathcal{V}_{i(x)}$.

¹The supplementary material is available here: <https://hal.science/hal-05064307>.

Example 1. We consider a negotiation between a car seller and a buyer, with the issues $\mathcal{I}(\text{car}) = \{\text{col}, \text{pr}, \text{del}\}$, corresponding respectively to the color, the price and the delivery date of the car. The values domains of the issues are $\mathcal{V}_{\text{col}} = \{\text{blue}, \text{red}, \text{green}\}$, $\mathcal{V}_{\text{pr}} = \{11000, 10500, 10000\}$ (in \$), and $\mathcal{V}_{\text{del}} = \{1, \dots, 5\}$ (in months).

In [Amgoud *et al.*, 2008], $A^{ag} = A_p^{ag} \cup A_e^{ag}$ is the set of arguments used in the negotiation by agent ag , where A_p^{ag} is the set of practical arguments that support the offers (modeling practical reasoning for making decisions) and A_e^{ag} is the set of epistemic arguments that represent the agent’s beliefs (modeling epistemic reasoning about the agents’ beliefs regarding the world and other agents). In our context, the arguments in A_p^{ag} support the values of the issues. Each value can be supported by several arguments. Formally, $A_p^{ag} = \bigcup_{i \in \mathcal{I}(x)} A_{p,i}^{ag}$, where $A_{p,i}^{ag}$ is the set of practical arguments associated with the issue i . Each $A_{p,i}^{ag} = \bigcup_{k=1}^m A_{p,v_k^i}^{ag}$, where $A_{p,v_k^i}^{ag}$ is the set of practical arguments supporting the value v_k^i in \mathcal{V}_i of the issue i . The sets of arguments supporting two different values are disjoint. Formally, $\forall v_k^i \neq v_l^i, A_{p,v_k^i}^{ag} \cap A_{p,v_l^i}^{ag} = \emptyset$.

We assume that agents have preferences over the possible values of the issues.

Definition 2 (Preferences on an issue’s value domain). Let a negotiation object x , i an issue and \mathcal{V}_i its value domain. We define a weak (preference) order [Fishburn, 1968] for an agent ag , \succeq_{ag}^i on \mathcal{V}_i s.t.:

$$\forall v_k^i, v_l^i \in \mathcal{V}_i, v_k^i \succeq_{ag}^i v_l^i \text{ or } v_l^i \succeq_{ag}^i v_k^i \text{ (possibly both)}$$

Each value is also associated with a utility.

Definition 3 (Utility of a value). Given i an issue, \mathcal{V}_i its domain and \succeq_{ag}^i the weak order on \mathcal{V}_i . The utility function $u_i^{ag} : \mathcal{V}_i \rightarrow [0, 1]$ s.t. $\forall i \in \mathcal{I}(x), \forall v_k^i \in \mathcal{V}_i, u_i^{ag}(v_k^i) \in [0, 1]$, i.e. u_i^{ag} maps the values to their utilities. The utilities must be consistent with the preferences, i.e. $u_i^{ag}(v_k^i) \geq u_i^{ag}(v_l^i)$ iff $v_k^i \succeq_{ag}^i v_l^i$.

For defining the utility of an offer (for an agent), we use (as usually in the literature) a Linear Additive Utility function, as we consider there is no interdependence between the issues.

Definition 4 (Utility of an offer). Let $\mathcal{I}(x)$ the set of issues, $\mathcal{W}^{ag} = \{w_i^{ag} \mid i \in \mathcal{I}(x)\}$ the set of weights of the issues, $0 \leq w_i^{ag} \leq 1$, $\sum_{i \in \mathcal{I}(x)} w_i^{ag} = 1$, and u_i^{ag} the values’ utility function. Then the utility of the offer o is defined as:

$$U^{ag}(o) = \sum_{i \in \mathcal{I}(x)} [w_i^{ag} * u_i^{ag}(v_k^i)]$$

Details on how to adapt to our multi-issue setting the definition of the agents argumentation theories originally proposed by [Dimopoulos *et al.*, 2021] for single-issue negotiation are provided in the supplementary material.

We now define the theories and profiles of the agents. Consider two agents α and β negotiating over the object x with the issues $\mathcal{I}(x)$. Their value domains \mathcal{V}_i (with $i \in \mathcal{I}(x)$) are the same for both agents, but each agent may assign different weights to the issues, different utilities to the values of the issues, may have different practical arguments supporting the

values and different epistemic arguments. Each agent has an initial profile composed of its own issues’ weights, the preferences and utilities for each value of each issue and its personal theory on the one hand; and its beliefs about the opponent’s profile (i.e. weights, preferences and utilities) before the beginning of the negotiation and about its argumentation theory on the other hand. Formally, the profile Pr_α of an agent α is:

Definition 5 (Agents’ profiles). Let α, β be two agents.

The profile of α is $Pr_\alpha := \langle \mathcal{I}(x), \mathcal{V}, u_\alpha, W_\alpha, u_\beta^?, W_\beta^?, T^\alpha, CAF^{\alpha,\beta} \rangle$, where $\mathcal{I}(x), \mathcal{V} = \{\mathcal{V}_i \mid i \in \mathcal{I}(x)\}$ are the issues and their domains, u_α, W_α are the personal utilities and weights of α , $u_\beta^?, W_\beta^?$ are the utilities and weights that α believes β has, T^α is α ’s argumentation theory, $CAF^{\alpha,\beta}$ is the argumentation theory α believes β has.

By symmetry, the profile of β is $Pr_\beta := \langle \mathcal{I}(x), \mathcal{V}, u_\beta, W_\beta, u_\alpha^?, W_\alpha^?, T^\beta, CAF^{\beta,\alpha} \rangle$.

4.2 Opponent Profile Modeling

In order to model the opponent’s profile, we use two methods. First, we learn the opponent’s weights and utilities, using a frequency model [van Krimpen *et al.*, 2013; Tunali *et al.*, 2017]. We define a function *updateProfile* computing the new profile of the opponent. This function increases the weight of the issues for which the value remained the same from one offer to the next by a coefficient ϵ [van Krimpen *et al.*, 2013]. To update the utilities of the values, we use formula 1 (Section 2) [Tunali *et al.*, 2017]. The updated profile is used to choose the next offer to send to the opponent. Then, we learn the opponent argumentation theory by using the Control Argumentation Frameworks (CAFs) [Dimopoulos *et al.*, 2018] like in [Dimopoulos *et al.*, 2019; Dimopoulos *et al.*, 2021]. Actually, the CAF is updated when agents exchange practical arguments supporting their offers along with epistemic and control arguments defending them.

4.3 Utility Concessions

The agent uses a utility concession profile [Baarslag *et al.*, 2015] which follows a time dependent formula to compute the utility U_{target} of the offers it can propose or accept at time t . Here, the time corresponds to the steps in the negotiation (i.e. the number of offers sent by the agent). For keeping the time t in $[0, 1]$, we set $t = \text{step} \div \text{maximumSteps}$. The formula is the following :

$$U_{target}(t) = 1 - (1 - U_{min}) \times t^{1/e} \quad (2)$$

with U_{min} and e parameters depending on the chosen profile.

4.4 Utility-based Values Selection

In this section we explain how agent α selects for each issue the value it will send to agent β at a step t .

An offer is represented as $o = \langle \sigma, S, R \rangle$, where $\sigma = \{\langle i, v_k^i, \theta^i \rangle \mid i \in \mathcal{I}(x)\}$ is a set where each issue in the negotiation is assigned a value along with a practical argument supporting that value, S is an extension that contains all the arguments θ^i appearing in σ and R is the set of attacks defending the arguments in S .

At the beginning of the negotiation, we compute the set Ω^α containing all the possible sets of pairs $\omega = \{\langle i, v_k^i \rangle \mid i \in$

$\mathcal{I}(x)\}$, with $v_k^i \in \mathcal{V}_i$ s.t. each issue is assigned exactly one value from its domain. At each step t , we compute the set $\Omega^{U_{target}} \subseteq \Omega^\alpha$ such that $\forall \omega_l \in \Omega^{U_{target}}, U^\alpha(\omega_l)$ (calculated with the formula from definition 4, see [Coeboorn and Jennings, 2004]) equals the utility target $U_{target}^\alpha(t)$ (see formula 2, and [Hosokawa and Fujita, 2020]). More formally:

$$\Omega^{U_{target}} \leftarrow \{\omega_l \mid \omega_l \in \Omega^\alpha, U^\alpha(\omega_l) = U_{target}^\alpha(t)\} \quad (3)$$

As in [Coeboorn and Jennings, 2004], agent α wants to choose the offer it thinks is best for agent β . In order to do so, agent α computes the utility $U^{\alpha,\beta}(\omega_l)$ of each $\omega_l \in \Omega^{U_{target}}$ (see definition 4), based on agent α 's current beliefs over agent β weights and utilities. This is done with the function *compute_omega_set* ($U_{target}^\alpha(t), \Omega^\alpha$) which returns the set $\omega \in \Omega^{U_{target}}$ that has the best utility for β (according to α 's beliefs), i.e. $\omega = \text{argmax}_{\Omega^{U_{target}}} U^{\alpha,\beta}$. When $t > 1$, before computing the new U_{target}^α , the function returns the sets $\omega_l \in \Omega^\alpha$ that were previously returned (i.e. have a higher utility than the new utility target) but the arguments supporting the offers were not acceptable in the current state of the agents' argumentation theories (see Section 4.5). However, through the exchange of arguments during the negotiation process (see Sections 4.6, 4.7 and 4.8), these previously not acceptable arguments may now be considered acceptable, enabling an offer to be sent when it was previously not feasible. This ensures that an offer, which could lead to an agreement from a utilitarian perspective, is disregarded only if all its supporting arguments remain not acceptable throughout the entire negotiation process.

4.5 Practical Arguments Selection

In this section we discuss Algorithm 1, which describes how α selects the practical arguments θ^i that support the values in ω and allow to complete the pairs $\langle i, v_k^i \rangle$ into triplets $\langle i, v_k^i, \theta^i \rangle$ in σ . The key idea behind this algorithm is as follows: an agent begins by checking its own theory to see if an offer, selected based on utility, is supported by an acceptable argument. Next, it examines the uncertain argumentation theory it holds about its opponent to identify whether the offer is supported by an acceptable argument or if it can introduce into the negotiation control arguments to make the supporting argument acceptable when it initially is not. More particularly, first, agent α needs to check if there is a set of arguments $\Phi = \{\phi^{i_1}, \dots, \phi^{i_n}\}$, that is credulously accepted in its theory T^α . In other words, it needs to check if the set ω that represents a utilitarian point of view is in adequacy with its beliefs that can represent ethical, social or legal points of view. This combination of views is necessary when negotiations concern not only economical transactions but also resolution of political or geostrategic conflicts. Since there can be several arguments supporting each value, we use a function *comp_phi*($\omega, \text{testedPhi}$) (line 1) that returns each time a different combination of arguments $\phi^i \in A_{p,v_k^i}^\alpha$, such that there is exactly one argument ϕ^i supporting each value v_k^i in ω . When such a set Φ is computed by *comp_phi*($\omega, \text{testedPhi}$), it is stored in the set *testedPhi* in

order to prevent this function from recalculating it (lines 2 and 30). Then we call a function *credulously_accepted*(Φ, T^α) (line 4) that returns a Boolean *true* if all practical arguments $\phi^i \in \Phi$ are together in at least one extension, and *false* otherwise. We compute new combinations of arguments Φ using *comp_phi*($\omega, \text{testedPhi}$) (line 29) until we find a set Φ that is credulously accepted in T^α or until all the combinations have been tested unsuccessfully (i.e. *comp_phi*($\omega, \text{testedPhi}$) returns \emptyset). If there is a set Φ credulously accepted in T^α , it means that ω is in adequacy with agent α 's beliefs.

In order to convince agent β that the set ω is the best one, and since agents do not necessarily have the same arguments in their personal theories, agent α needs to find in $CAF^{\alpha,\beta}$ a set of practical arguments $\Theta = \{\theta^{i_1}, \dots, \theta^{i_n}\}$ supporting the values in ω , which are credulously accepted in $CAF^{\alpha,\beta}$. $CAF^{\alpha,\beta}$ represents agent α 's current knowledge of agent β 's beliefs. Similarly as in T^α , there can be several practical arguments in $CAF^{\alpha,\beta}$ that support the same value. The function *comp_theta*($\omega, \text{testedTheta}$) (line 5) returns each time a new combination of arguments in the set Θ , such that $\theta^i \in A_{p,v_k^i}^{\alpha,\beta}$ and there is exactly one argument θ^i supporting each value v_k^i in ω . As with the sets Φ , each time a set Θ is computed it is stored in *testedTheta* (lines 6, 13, 17 and 26) so that we prevent its recalculation. Then, we check if Θ is credulously accepted in $A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}$. The function *credulously_accepted*($\Theta, A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}$) (line 8) returns *true* if there is at least one extension in each completion (see the supplementary material for more information on completions) that contains all arguments $\theta^i \in \Theta$, *false* otherwise. We compute a new set Θ (line 12) until we find one that is credulously accepted in $A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}$ or until all the combinations have been tested unsuccessfully (i.e. *comp_theta*($\omega, \text{testedTheta}$) returns \emptyset). If such a set Θ is found, Algorithm 1 returns the offer $\langle \sigma, \emptyset, \emptyset \rangle$, where $\sigma = \{\langle i, v_k^i, \theta^i \rangle \mid i \in \mathcal{I}(x)\}$, with $\langle i, v_k^i \rangle \in \omega$, and $\theta^i \in \Theta$. The function *comp_sigma*(ω, Θ) (lines 9 and 22) computes the set σ based on the set ω of couples $\langle i, v_k^i \rangle$ and the set Θ . If no combination Θ is credulously accepted in $A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}$, we search for a configuration of $CAF^{\alpha,\beta}$, using control arguments, where a set Θ is credulously accepted. This means there is at least one extension in each completion that includes all arguments in Θ , achieved using the function *comp_contr_conf*($\omega, \Theta, CAF^{\alpha,\beta}$) (line 19). This function returns an extension S such that $\Theta \subseteq S$ (if it exists), and \emptyset otherwise. We continue computing a new set Θ (line 25) until the function *comp_contr_conf*($\omega, \Theta, CAF^{\alpha,\beta}$) returns an extension S or no new sets remain. If *comp_contr_conf*($\omega, \Theta, CAF^{\alpha,\beta}$) returns an extension S , we compute the set of attacks R (line 21) that defends the arguments in S . Then, Algorithm 1 returns the offer $\langle \sigma, S, R \rangle$.

If no set Φ is credulously accepted in T^α , it means the set ω is not in adequacy with agent α 's beliefs and Algorithm 1 returns $\langle \emptyset, \emptyset, \emptyset \rangle$. Similarly, if there is no set Θ credulously accepted in $A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}$ or $CAF^{\alpha,\beta}$, then the set ω is not in adequacy with agent α 's knowledge of agent β 's beliefs and Algorithm 1 returns $\langle \emptyset, \emptyset, \emptyset \rangle$.

Algorithm 1 *choose_practical_arguments*($\omega, T^\alpha, CAF^{\alpha,\beta}$)

```

1:  $\Phi \leftarrow \text{comp\_phi}(\omega, \text{testedPhi})$ 
2:  $\text{testedPhi} \leftarrow \{\Phi\}$ 
3: while  $\Phi \neq \emptyset$  do
4:   if  $\text{credulously\_accepted}(\Phi, T^\alpha) = \text{true}$  then
5:      $\Theta \leftarrow \text{comp\_theta}(\omega, \text{testedTheta})$ 
6:      $\text{testedTheta} \leftarrow \{\Theta\}$ 
7:     while  $\Theta \neq \emptyset$  do
8:       if  $\text{credulously\_accepted}(\Theta, A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}) = \text{true}$  then
9:          $\sigma \leftarrow \text{comp\_sigma}(\omega, \Theta)$ 
10:        return  $\langle \sigma, \emptyset, \emptyset \rangle$ 
11:       end if
12:        $\Theta \leftarrow \text{comp\_theta}(\omega, \text{testedTheta})$ 
13:        $\text{testedTheta} \leftarrow \text{testedTheta} \cup \{\Theta\}$ 
14:     end while
15:      $\text{testedTheta} \leftarrow \emptyset$ 
16:      $\Theta \leftarrow \text{comp\_theta}(\omega, \text{testedTheta})$ 
17:      $\text{testedTheta} \leftarrow \{\Theta\}$ 
18:     while  $\Theta \neq \emptyset$  do
19:        $S \leftarrow \text{comp\_contr\_conf}(\Theta, CAF^{\alpha,\beta})$ 
20:       if  $S \neq \emptyset$  then
21:          $R \leftarrow \{(a, b) \mid a \in S, b \in A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}\}$ 
22:          $\sigma \leftarrow \text{comp\_sigma}(\omega, \Theta)$ 
23:         return  $\langle \sigma, S, R \rangle$ 
24:       end if
25:        $\Theta \leftarrow \text{comp\_theta}(\omega, \text{testedTheta})$ 
26:        $\text{testedTheta} \leftarrow \text{testedTheta} \cup \{\Theta\}$ 
27:     end while
28:   end if
29:    $\Phi \leftarrow \text{comp\_phi}(\omega, \text{testedPhi})$ 
30:    $\text{testedPhi} \leftarrow \text{testedPhi} \cup \{\Phi\}$ 
31: end while
32: return  $\langle \emptyset, \emptyset, \emptyset \rangle$ 

```

4.6 The Biding Strategy

In this section we discuss Algorithm 2, describing agent α 's biding strategy. The key idea behind it is as follows: an agent seeks to learn its opponent's profile to propose a utility-based offer that is likely to be accepted. It then applies Algorithm 1 to identify acceptable arguments supporting this offer, iterating the process as long as it can find offers with utility greater than its reservation value, when previously selected offers cannot be supported by acceptable arguments in the theories of both agents. More particularly, first, agent α updates its opponent's profile by using the frequency models from Section 2 (line 1). Then it computes the target utility $U_{\text{target}}^\alpha(t)$ at step t . Once this is done, it calls the function $\text{compute_omega_set}(U_{\text{target}}^\alpha(t), \Omega^\alpha)$ (formula 3) to choose the best set ω at this step. While no offer has been chosen, $o(\alpha, \beta) = \langle \emptyset, \emptyset, \emptyset \rangle$ and $\text{compute_omega_set}(U_{\text{target}}^\alpha(t), \Omega^\alpha)$ does not return \emptyset , it calls Algorithm 1 (line 8) which returns an offer $o(\alpha, \beta) = \langle \sigma, S, R \rangle$ if there is a set of practical arguments supporting the values in ω credulously accepted in T^α and a set of practical arguments supporting the values in ω credulously accepted in

$CAF^{\alpha,\beta}$. Otherwise, it returns $\langle \emptyset, \emptyset, \emptyset \rangle$. In this case (line 9), we remove ω from Ω^α (line 10) and choose a new ω for this step t (line 11). If Algorithm 1 returns $\langle \emptyset, \emptyset, \emptyset \rangle$ for all the sets ω at step t , we compute U_{target}^α for the next step t (line 15). If $U_{\text{target}}^\alpha > \text{reservationValue}$, we stay in the loop. Once we leave the loop, if $o(\alpha, \beta) = \langle \emptyset, \emptyset, \emptyset \rangle$, then the agent does not have any new offers to send to its opponent. If the opponent's last message was *no_offer* it means the opponent does not have new offers either, and agent α quits the negotiation (line 20). Otherwise, the agent sends *no_offer* (line 18); if an offer is set, it is sent to agent β .

Algorithm 2 *makeBid*(reservationValue , $\text{profile}^{\alpha,\beta}$, Ω^α , step , $\text{message}(\beta, \alpha)$, T^α , $CAF^{\alpha,\beta}$)

```

1: call  $\text{updateProfile}(\text{profile}^{\alpha,\beta})$ 
2:  $t \leftarrow \text{step}$ 
3: compute  $U_{\text{target}}^\alpha(t)$ 
4:  $o(\alpha, \beta) \leftarrow \langle \emptyset, \emptyset, \emptyset \rangle$ 
5: while  $o(\alpha, \beta) = \langle \emptyset, \emptyset, \emptyset \rangle$  and  $U_{\text{target}}^\alpha(t) \geq \text{reservationValue}$  do
6:    $\omega \leftarrow \text{compute\_omega\_set}(U_{\text{target}}^\alpha(t), \Omega^\alpha)$ 
7:   while  $o(\alpha, \beta) = \langle \emptyset, \emptyset, \emptyset \rangle$  and  $\omega \neq \emptyset$  do
8:      $o(\alpha, \beta) \leftarrow \text{chose\_pract\_arg}(\omega, T^\alpha, CAF^{\alpha,\beta})$ 
9:     if  $o(\alpha, \beta) = \langle \emptyset, \emptyset, \emptyset \rangle$  then
10:       $\Omega^\alpha \leftarrow \Omega^\alpha \setminus \{\omega\}$ 
11:       $\omega \leftarrow \text{compute\_omega\_set}(U_{\text{target}}^\alpha(t), \Omega^\alpha)$ 
12:    end if
13:   end while
14:    $t \leftarrow t + 1$ 
15:   compute  $U_{\text{target}}^\alpha(t)$ 
16: end while
17: if  $o(\alpha, \beta) = \langle \emptyset, \emptyset, \emptyset \rangle$  and  $\text{message}(\beta, \alpha) \neq \text{no\_offer}$  then
18:    $\text{message}(\alpha, \beta) \leftarrow \text{no\_offer}$ 
19: else if  $o(\alpha, \beta) = \langle \emptyset, \emptyset, \emptyset \rangle$  then
20:    $\text{message}(\alpha, \beta) \leftarrow \text{quit}$ 
21: else
22:    $\text{message}(\alpha, \beta) \leftarrow \text{propose}(o(\alpha, \beta), \emptyset, \emptyset)$ 
23: end if
24: send( $\text{message}(\alpha, \beta)$ )

```

4.7 The Acceptance Strategy

In this section we describe the acceptance strategy of the agents. The key idea behind this algorithm is as follows: the agent evaluates whether it can accept the proposed offer after updating its argumentation theory with the information provided alongside the offer. If it cannot, it communicates the reasons for its inability to accept. More particularly, when agent α receives an offer $o(\beta, \alpha)$ if the extension S is not null it updates its own theory T^α and $CAF^{\alpha,\beta}$ by integrating the arguments in S and the attacks in R to them (lines 3-6).

For updating the theories T^α and $CAF^{\alpha,\beta}$ we use the operator \oplus defined in [Dimopoulos *et al.*, 2019; Dimopoulos *et al.*, 2021] as follows: given three sets A_1, A_2, A_3 , $(A_1, A_2) \oplus A_3$ is the pair (A'_1, A'_2) such that $A'_1 = A_1 \setminus (A_1 \cap A_3)$ and $A'_2 = A_2 \cup (A_1 \cap A_3)$. This operator is used in Algorithms 3.

Then, agent α accepts an offer (line 17) if the offer's utility $U^\alpha(o(\beta, \alpha))$ is above its current $U_{target}^\alpha(t)$, and the supporting arguments in the offer are credulously accepted in T^α .

If $U^\alpha(o(\beta, \alpha)) < U_{target}^\alpha$ agent α sends a reject message (line 9). If $U^\alpha(o(\beta, \alpha)) > U_{target}^\alpha$, agent α calls a function $get_pract_arg(\sigma)$ that extracts the set Φ of practical arguments supporting the values of the offer in its theory T^α . Then, it checks if Φ is credulously accepted in T^α . If it is not the case, agent α computes the set of arguments Q , and the set of attacks *Reasons* explaining this rejection and sends the rejection message (lines 12-15). If it rejects the offer, it makes a counteroffer by calling Algorithm 2.

Algorithm 3 *accept_offer*($U_{target}^\alpha(t), o(\beta, \alpha), T^\alpha, CAF^{\alpha, \beta}$)

```

1:  $\langle \sigma, S, R \rangle = o(\beta, \alpha)$ 
2: if  $S \neq \emptyset$  then
3:    $T^\alpha = \langle A^\alpha \cup S, \rightarrow_\alpha \cup R \rangle$ 
4:    $(A_U^{\alpha, \beta}, A_F^{\alpha, \beta}) = (A_U^{\alpha, \beta}, A_F^{\alpha, \beta}) \oplus S$ 
5:    $(\neg \rightarrow_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) = (\neg \rightarrow_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) \oplus R$ 
6:    $(\rightleftharpoons_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) = (\rightleftharpoons_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) \oplus R$ 
7: end if
8: if  $U^\alpha(o(\beta, \alpha)) < U_{target}^\alpha(t)$  then
9:    $message(\alpha, \beta) \leftarrow reject(o(\beta, \alpha), \emptyset, \emptyset)$ 
10: else
11:    $\Phi \leftarrow get\_pract\_arg(\sigma)$ 
12:   if not credulously_accepted( $\Phi, T^\alpha$ ) then
13:     Compute  $Q \subseteq \xi$  where  $\xi$  is an extension and  $Q$  is
       the set of arguments from which each  $\phi_j \in \Phi$  are
       reachable
14:      $Reasons \leftarrow \{(r, \phi_j) | (r, \phi_j) \in \rightarrow_\alpha, r \in Q, \phi_j \in \Phi\}$ 
15:      $message(\alpha, \beta) \leftarrow reject(o(\beta, \alpha), Q, Reasons)$ 
16:   else
17:      $message(\alpha, \beta) \leftarrow accept(o(\beta, \alpha), \emptyset, \emptyset)$ 
18:   end if
19: end if
20:  $send(message(\alpha, \beta))$ 
21: if ( $message(\alpha, \beta) \leftarrow reject(o(\beta, \alpha), Q, Reasons)$ )
   then
22:   call makeBid(reservationValue, profile $^{\alpha, \beta}$ ,  $\Omega^\alpha$ ,
     step, message( $\beta, \alpha$ ),  $T^\alpha$ ,  $CAF^{\alpha, \beta}$ )
23: end if

```

4.8 The Negotiation Protocol

This section discusses the overall negotiation protocol (Algorithm 4). This algorithm enables observing how an agent's theory is updated after receiving a reject message. While the negotiation is ongoing, agents α and β exchange messages that can either be an offer, the absence of an offer, the rejection of an offer, the acceptance of an offer or the agent informing its opponent that it quits the negotiation. If the opponent message is a rejection and the set of arguments Q explaining the rejection in the theory is not \emptyset , the agent updates its theories T^α and $CAF^{\alpha, \beta}$ (lines 10-13). Then the agent waits for its opponent to make a new bid. If the message is an acceptance or if its opponent quits (the opponent no longer wishes to negotiate) the negotiation ends. If the message is an of-

fer, then the agent checks if it can accept the offer by calling Algorithm 3. If agent β did not have an offer, agent α calls Algorithm 2 to find its next offer to send to agent β . Algorithm 4 describes this protocol for agent α . A comprehensive example illustrating the execution of the protocol is provided in the supplementary material.

Algorithm 4 *negotiationProtocol*(*reservationValue*, *profile* $^{\alpha, \beta}$, Ω^α , *step*, *message*(β, α), T^α , $CAF^{\alpha, \beta}$)

```

1: if agent goes first then
2:   makeBid(reservationValue, profile $^{\alpha, \beta}$ ,  $\Omega^\alpha$ , step,
     message( $\beta, \alpha$ ),  $T^\alpha$ ,  $CAF^{\alpha, \beta}$ )
3: end if
4: negotiation  $\leftarrow true$ 
5: while negotiation do
6:   get message( $\beta, \alpha$ )
7:   switch message( $\beta, \alpha$ ) do
8:     case reject( $o(\alpha, \beta), Q, Reasons$ ):
9:       if  $Q \neq \emptyset$  then
10:         $T^\alpha = \langle A^\alpha \cup Q, \rightarrow_\alpha \cup Reasons \rangle$ 
11:         $(A_U^{\alpha, \beta}, A_F^{\alpha, \beta}) = (A_U^{\alpha, \beta}, A_F^{\alpha, \beta}) \oplus Q$ 
12:         $(\neg \rightarrow_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) = (\neg \rightarrow_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) \oplus Reasons$ 
13:         $(\rightleftharpoons_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) = (\rightleftharpoons_{\alpha, \beta}, \rightarrow_{\alpha, \beta}) \oplus Reasons$ 
14:       end if
15:       get message( $\beta, \alpha$ )
16:       case accept( $o(\alpha, \beta), \emptyset, \emptyset$ ):
17:         negotiation  $\leftarrow false$ 
18:       case quit:
19:         negotiation  $\leftarrow false$ 
20:       case propose( $o(\beta, \alpha), \emptyset, \emptyset$ ):
21:         accept_offer( $U_{target}^\alpha(t), o(\beta, \alpha), T^\alpha, CAF^{\alpha, \beta}$ )
22:       case no_offer:
23:         makeBid(reservationValue, profile $^{\alpha, \beta}$ ,  $\Omega^\alpha$ ,
           step, message( $\beta, \alpha$ ),  $T^\alpha$ ,  $CAF^{\alpha, \beta}$ )
24:       end switch
25: end while

```

5 Experimental Results

Our negotiation framework was implemented in Java using the Jade platform (<https://jade.tilab.com/>). The acceptability of arguments in CAFs was computed with the solver taeydenae [Niskanen *et al.*, 2020b; Niskanen *et al.*, 2020a].

Our experiments include 100 multi-issue negotiations with 2 issues and 100 negotiations with 3 issues, with the number of values for each issue included between $min = 3$ and $max = 5$. More details on the benchmark generation are given in the supplementary material.

Our results are summarized in Figure 1. We first compare the percentage of agreements between our argument-based multi-issue negotiation protocol and a multi-issue negotiation protocol without argumentation. The negotiations without argumentation were run on the genius platform [Lin *et al.*, 2014] using the Atlas3 [Mori and Ito, 2017] party (i.e. agent model) for both agents. We selected this agent model because our framework currently employs the same utility target formula. However, since our framework is modular, any utility




Figure 1: Percentage of agreements in multi-issue negotiations: without and with argumentation, with 2 and 3 issues.

target formula can be implemented, allowing other agents to be used for negotiations without argumentation. As we can see, our protocol (red bars) has a percentage of agreement of 67% for 2 issues and 64% for three issues, whereas the multi-issue negotiation protocol without argumentation (blue bars) has an overall percentage of agreements of 94% for 2 issues and 100% for 3 issues. However, simply comparing the overall percentage of agreements is insufficient for real world situations. To assess the added value of our protocol, it is also essential to evaluate the percentage of agreements that align with the agents’ beliefs (i.e., their argumentation theories) for the negotiation protocol without argumentation. By using arguments to verify that offers acceptable to both agents from a utility perspective also align with their beliefs (which may represent societal, ethical or legal principles, among other principles), our protocol ensures the exclusion of offers that violate these principles and would likely be rejected in real-life situations. In contrast, negotiations conducted using the multi-issue negotiation protocol without argumentation relied solely on the utility of the offers to reach agreements. This explains the difference in the number of agreements (for 2 and 3 issues) between the two negotiation protocols, with the protocol without argumentation being favored when arguments are not taken into account. When considering the arguments for verifying compliance with beliefs, the negotiation protocol without arguments (represented by blue striped bars) shows a significant drop in the percentage of consistent agreements, reaching only 40% for two issues and 29% for three issues. This occurs because the values of the issues in the other agreements (i.e. that are lost) are supported by arguments that are not acceptable in the agents’ theories (i.e. they conflict with their beliefs). The difference in the percentages of agreements reached between our protocol and the protocol without argumentation—this time favoring our protocol when considering the arguments supporting the offers—arises from a fundamental distinction. In the latter, if a supporting argument for an offer is not accepted in the agents’ theories, the agreement on that offer—reached purely based on utilities—is disregarded, as no argumentation protocol is available to resolve such conflicts. In contrast, our protocol allows for the possibility of exchanging arguments, enabling initially rejected supporting arguments to potentially become accepted in the agents’ theories, thus facilitating agreement. Therefore, we show that our approach is better suited for handling multi-issue negotiations in real-world settings.

6 Discussion and Future Work

In this paper, we presented a novel hybrid approach that combines the advantages of the two main traditional approaches, namely utility-based and argumentation-based approaches, enabling effective adaptation of argumentation-based approaches in multi-issue negotiations. Our hybrid approach leverages traditional (single-issue) argumentation approaches (e.g. [Dimopoulos *et al.*, 2019]), because it uses the most effective techniques (i.e., frequency models) in the literature to seek trade-offs across multiple issues while accounting for both the proponent’s and the opponent’s profiles. This allows for an increase in the number of agreements, similar to traditional utility-based negotiations. At the same time, our approach offers a significant advantage over traditional utility-based multi-issue negotiations [Zhan *et al.*, 2024]. By incorporating argumentation, it enables an agent to explain the rationale behind each offer and persuade its opponent that the offer not only satisfies utilitarian expectations but also aligns with its profile, ethical principles, legal requirements and societal norms, among other principles. This capability is particularly critical for addressing automated negotiations, not only in business contexts but also in more complex domains such as political or geostrategic negotiations. The way in which the argumentation is used is a major difference with the work proposed in [Xue-jie *et al.*, 2013] which is, to our knowledge, the only one close to our approach in the sense that it uses utilities to represent the preferences of agents and arguments to support the offers. Indeed this work uses argumentation in a traditional way (i.e. only to support offers like the practical arguments of our work). Furthermore, this work uses a simple mechanism to calculate the utilities of offers based exclusively on the information exchanged between the agents (i.e. without possibility of learning the profile of the opponent) and considering that the agents will reveal their utilities when a conflict is detected, which is an unrealistic assumption in many real-world applications. Regarding computational aspects, our work relies on well-known techniques, like checking the credulous acceptability of a set of arguments or the controllability of a CAF, which are respectively NP-complete [Dvorák and Dunne, 2017] and Σ_3^P -complete [Niskanen *et al.*, 2020a] for the stable semantics.

To conclude, we highlight ideas for future research. A first one consists in studying the impact of using other models for learning the opponents utilities (e.g. Bayesian filters [Buffett and Spencer, 2005; Buffett and Spencer, 2007]). Another interesting track consists in using other forms of controllability [Mailly, 2020; Gaignier *et al.*, 2021], notably in presence of probabilities, that could allow agents to find agreements in situations where no offer is acceptable with respect to *all* the completions but may be acceptable with respect to the most probable ones, or when negotiation are subject to time constraints. Finally, a promising direction for future work is the creation of an NLP-powered platform that allows users to define agents’ utilities, preferences and argumentation theories using natural language.

Acknowledgments

The second author is funded by the French National Research Agency under grants ANR-22-CE23-0005 (AGGREEY) and ANR-22-CPJ1-0061-01 (AIDAL).

References

- [Amgoud and Kaci, 2006] Leila Amgoud and Souhila Kaci. On the study of negotiation strategies. In *Agent Communication II, International Workshops on Agent Communication, AC 2005 and AC 2006, Utrecht, Netherlands, July 25, 2005 and Hakodate, Japan, May 9, 2006, Selected and Revised Papers*, pages 150–163, 2006.
- [Amgoud et al., 2007] Leila Amgoud, Yannis Dimopoulos, and Pavlos Moraitis. A unified and general framework for argumentation-based negotiation. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems, (AAMAS07)*, pages 963–970, 2007.
- [Amgoud et al., 2008] Leila Amgoud, Yannis Dimopoulos, and Pavlos Moraitis. Making decisions through preference-based argumentation. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008*, pages 113–123. AAAI Press, 2008.
- [Baarslag et al., 2015] Tim Baarslag, Reyhan Aydoğan, Koen V. Hindriks, Katsuhide Fujita, Takayuki Ito, and Catholijn M. Jonker. The automated negotiating agents competition, 2010–2015. *AI Magazine*, 36(4):115–118, 2015.
- [Baarslag et al., 2016a] Tim Baarslag, Mark J. C. Hendriks, Koen V. Hindriks, and Catholijn M. Jonker. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, 30(5):849–898, 2016.
- [Baarslag et al., 2016b] Tim Baarslag, Mark J. C. Hendriks, Koen V. Hindriks, and Catholijn M. Jonker. A survey of opponent modeling techniques in automated negotiation. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, pages 575–576. ACM, 2016.
- [Baroni et al., 2011] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowl. Eng. Rev.*, 26(4):365–410, 2011.
- [Bonzon et al., 2012] Elise Bonzon, Yannis Dimopoulos, and Pavlos Moraitis. Knowing each other in argumentation-based negotiation. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*, pages 1413–1414, 2012.
- [Buffett and Spencer, 2005] Scott Buffett and Bruce Spencer. Learning opponents’ preferences in multi-object automated negotiation. In *Proceedings of the 7th International Conference on Electronic Commerce, ICEC ’05*, page 300–305, New York, NY, USA, 2005. Association for Computing Machinery.
- [Buffett and Spencer, 2007] Scott Buffett and Bruce Spencer. A bayesian classifier for learning opponents’ preferences in multi-object automated negotiation. *Electron. Commer. Rec. Appl.*, 6(3):274–284, oct 2007.
- [Casali et al., 2016] Ana Casali, Pablo Pilotti, and Carlos Iván Chesñevar. Assessing communication strategies in argumentation-based negotiation agents equipped with belief revision. *Argument Comput.*, 7(2-3):175–200, 2016.
- [Coehoorn and Jennings, 2004] Robert M. Coehoorn and Nicholas R. Jennings. Learning on opponent’s preferences to make effective multi-issue negotiation trade-offs. In *Proceedings of the 6th International Conference on Electronic Commerce, ICEC ’04*, page 59–68, New York, NY, USA, 2004. Association for Computing Machinery.
- [Dimopoulos and Moraitis, 2014] Yannis Dimopoulos and Pavlos Moraitis. Advances in argumentation-based negotiation. In *Negotiation and Argumentation in Multi-Agent Systems: Fundamentals, Theories, Systems and Applications*, pages 82–125, 2014.
- [Dimopoulos et al., 2018] Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. Control argumentation frameworks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4678–4685. AAAI Press, 2018.
- [Dimopoulos et al., 2019] Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. Argumentation-based negotiation with incomplete opponent profiles. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’19)*, pages 1252–1260, may 2019.
- [Dimopoulos et al., 2021] Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. Arguing and negotiating using incomplete negotiators profiles. *Auton. Agents Multi Agent Syst.*, 35(2):18, 2021.
- [Dung et al., 2008] Phan Minh Dung, Phan Minh Thang, and Francesca Toni. Towards argumentation-based contract negotiation. In *Computational Models of Argument: Proceedings of COMMA 2008, Toulouse, France, May 28-30, 2008.*, pages 134–146, 2008.
- [Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.
- [Dvorák and Dunne, 2017] Wolfgang Dvorák and Paul E. Dunne. Computational problems in formal argumentation and their complexity. *FLAP*, 4(8), 2017.
- [Fishburn, 1968] Peter C. Fishburn. Utility theory. *Management Science*, 14(5):335–378., 1968.
- [Gaignier et al., 2021] Fabrice Gaignier, Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. Probabilistic control argumentation frameworks. In *AAMAS ’21: 20th*

- International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, pages 519–527. ACM, 2021.
- [Hadidi *et al.*, 2010] Nabila Hadidi, Yannis Dimopoulos, and Pavlos Moraitis. Argumentative alternating offers. In *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Toronto, Canada, May 10-14, 2010, Volume 1-3, pages 441–448, 2010.
- [Hadidi *et al.*, 2012] Nabila Hadidi, Yannis Dimopoulos, and Pavlos Moraitis. Tactics and concessions for argumentation-based negotiation. In *Computational Models of Argument - Proceedings of COMMA 2012*, Vienna, Austria, September 10-12, 2012, pages 285–296, 2012.
- [Hosokawa and Fujita, 2020] Yuta Hosokawa and Katsuhide Fujita. Opponent’s preference estimation considering their offer transition in multi-issue closed negotiations. *IEICE Transactions on Information and Systems*, E103.D:2531–2539, 12 2020.
- [Kakas and Moraitis, 2006] Antonis C. Kakas and Pavlos Moraitis. Adaptive agent negotiation via argumentation. In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, Hakodate, Japan, May 8-12, 2006, pages 384–391, 2006.
- [Lin *et al.*, 2014] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014.
- [Luo *et al.*, 2024] Xudong Luo, Yanling Li, Qiaojuan Huang, and Jieyu Zhan. A survey of automated negotiation: Human factor, learning, and application. *Comput. Sci. Rev.*, 54:100683, 2024.
- [Mailly, 2020] Jean-Guy Mailly. Possible controllability of control argumentation frameworks. In *Computational Models of Argument - Proceedings of COMMA 2020*, Perugia, Italy, September 4-11, 2020, volume 326 of *Frontiers in Artificial Intelligence and Applications*, pages 283–294. IOS Press, 2020.
- [Mori and Ito, 2017] Akiyuki Mori and Takayuki Ito. *Atlas3: A Negotiating Agent Based on Expecting Lower Limit of Concession Function*, volume 674, pages 169–173. 04 2017.
- [Niskanen *et al.*, 2020a] Andreas Niskanen, Daniel Neugebauer, and Matti Järvisalo. Controllability of control argumentation frameworks. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 1855–1861. ijcai.org, 2020.
- [Niskanen *et al.*, 2020b] Andreas Niskanen, Daniel Neugebauer, Matti Järvisalo, and Jörg Rothe. Deciding acceptance in incomplete argumentation frameworks. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI 2020)*, pages 2942–2949, United States, 2020. AAAI Press.
- [Parsons *et al.*, 1998] Simon Parsons, Carles Sierra, and Nicholas R. Jennings. Agents that reason and negotiate by arguing. *J. Log. Comput.*, 8(3):261–292, 1998.
- [Pilotti *et al.*, 2015] Pablo Pilotti, Ana Casali, and Carlos Iván Chesñevar. A belief revision approach for argumentation-based negotiation agents. *Applied Mathematics and Computer Science*, 25(3):455–470, 2015.
- [Pocola, 2022] Octav Pocola. *Opponent Modeling in Automated Bilateral Negotiation: Can Machine Learning Techniques Outperform State-of-the-Art Heuristic Techniques?* PhD thesis, TU Delft Electrical Engineering, Mathematics and Computer Science, 2022.
- [Rahwan *et al.*, 2003] Iyad Rahwan, Sarvapali D. Ramchurn, Nicholas R. Jennings, Peter McBurney, Simon Parsons, and Liz Sonenberg. Argumentation-based negotiation. *Knowledge Eng. Review*, 18(4):343–375, 2003.
- [Tunali *et al.*, 2017] Okan Tunali, Reyhan Aydogan, and Víctor Sánchez-Anguix. Rethinking frequency opponent modeling in automated negotiation. In *PRIMA 2017: Principles and Practice of Multi-Agent Systems - 20th International Conference, Nice, France, October 30 - November 3, 2017, Proceedings*, pages 263–279, 2017.
- [van Galen Last, 2012] Niels van Galen Last. Agent Smith: Opponent model estimation in bilateral multi-issue negotiation. *Studies in Computational Intelligence*, 383:167–174, 2012.
- [van Krimpen *et al.*, 2013] Thijs van Krimpen, Daphne Looije, and Siamak Hajizadeh. Hardheaded. *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pages 223–227, 2013.
- [Xue-jie *et al.*, 2013] Dong Xue-jie, Chen Jian, Hu Ying-lan, Jiang Guo-rui, and Huang Ti-yun. Multi-attribute negotiation model based on internal factors argumentation. In *2013 International Conference on Management Science and Engineering 20th Annual Conference Proceedings*, pages 20–27, 2013.
- [Zafari and Nassiri-Mofakham, 2016] Farhad Zafari and Faria Nassiri-Mofakham. Popponent: Highly accurate, individually and socially efficient opponent preference model in bilateral multi issue negotiations. *Artificial Intelligence*, 237:59–91, 2016.
- [Zhan *et al.*, 2024] Haolan Zhan, Yufei Wang, Zhuang Li, Tao Feng, Yuncheng Hua, Suraj Sharma, Lizhen Qu, Zhaleh Semnani-Azad, Ingrid Zukerman, and Reza Haf-fari. Let’s negotiate! A survey of negotiation dialogue systems. In *Findings of the Association for Computational Linguistics: EACL 2024, St. Julian’s, Malta, March 17-22, 2024*, pages 2019–2031. Association for Computational Linguistics, 2024.