Computational Models of Argument
S. Modgil et al. (Eds.)

doi:10.3233/978-1-61499-906-5-345

Choosing Appropriate Arguments from Trustworthy Sources

Alison R. PANISSON [a,1], Simon PARSONS [b], Peter MCBURNEY [b], and Rafael H. BORDINI [a]

[a] *School of Technology, PUCRS – Porto Alegre, Brazil*
[b] *Department of Informatics, KCL – London, UK*

Abstract. Recently, argumentation frameworks have been extended in order to consider trust when defining preferences between arguments, given that arguments (or information that supports the arguments) from more trustworthy sources may be preferred to arguments from less trustworthy sources. Although such literature presents interesting results on argumentation-based reasoning and how agents define preferences between arguments, there is little work taking into account agent strategies for argumentation-based dialogues using such information. In this work, we propose an argumentation framework in which agents consider how much the recipient of an argument trusts others in order to choose the most suitable argument for that particular recipient, i.e., arguments constructed using information from those sources that the recipient trusts. Our approach aims to allow agents to construct more effective arguments, depending on the recipients and on their views on the trustworthiness of potential sources.

Keywords. Multi-Agent Systems, Argumentation, Trust and Reputation

1. Introduction

Argumentation-based techniques have been applied to Multi-Agent Systems (MAS) in order to allow agents to reason in conditions of uncertainty and to have richer communication frameworks. Also, inspired by human societies, given the adoption of an *open society* paradigm to MAS, trust and reputation models have been applied to MAS in order to take into account different kinds of social relationship between agents [1]. Following [1,2], trust can arise from two views: (i) the first is a subjective property assessed particularly by each individual, in which an agent directly or indirectly undertakes interactions with other agents. This point of view describes the *trust* of an individual x from the point of view of an individual y; (ii) the second is a societal view of trust, which consists of observations by the society of the past behaviour of agents (which is called *reputation*) which are made available to agents who have not directly interacted with them previously. This point of view can be seen as common-sense knowledge, in the sense that the society agrees about the reputation of each individual in that society [1].

In this paper, we consider both views of trust in MAS, the individual and societal views, i.e., *trust* and *reputation*. In particular, we put forward strategies for agents to

[1] Corresponding Author: Alison R. Panisson – e-mail: alison.panisson@acad.pucrs.br

choose more effective arguments depending on the intended recipients, particularly for recipients who consider trust to determine preferences between arguments and, consequently, to decide the acceptability of arguments. Also, it will not always be the case that agents have access to others' individual views on trust; therefore, the social view of trust (i.e., reputation) also plays an important role here. In our approach, an agent will be able to construct the "best" arguments for a recipient when it has access to the recipient's individual views on the trustworthiness of others and, when that is not the case, it will be able to construct good arguments based on reputations, considering that agents with good social reputation tend to be trustworthy from the views of individuals too.

2. Trust and Reputation in Multi-Agent Systems

We assume that agents have an internal representation of how much they trust other agents, as well as having access to the reputation of each agent in that society. This can be based on approaches that combine *trust* and *reputation* in MAS, for example [2].

We use the following notation in this paper: given an agent Ag, Π^{Ag} represents Ag's private knowledge regarding trust, and Δ^{Ag} represents all of Ag's private knowledge about the other agents and about the world (environment, organisation, etc.), i.e., $\Pi^{Ag} \subset \Delta^{Ag}$. The reputation knowledge is denoted by Γ. Particularly interesting for our approach is that trust/reputation systems allow agents to have a (partial) model about other agents private knowledge regarding trust, which we will represent here as $\Delta^{Ag_i}_{\Pi'^{Ag_j}}$, meaning that part of Ag_i's knowledge base contains knowledge about the Ag_j's private knowledge base regarding trust (note that $\Pi'^{Ag_j} \subseteq \Pi^{Ag_j}$).

In MAS, *trust* can be formalised as an asymmetric relation between agents $tr : Ags \times Ags \mapsto \mathbb{R}$ [3,4,5], which returns either a value between 0 and 1, representing how much an agent trusts another, or $null$ in case the agents have no acquaintance. Both cases can be seen in the trust network in Figure 1, where we have $tr(Ag4, Ag5) = 0$ and $tr(Ag1, Ag4) = null$. Furthermore, trust can be a transitive relation, so an agent Ag_i can evaluate its private trust on Ag_j directly or indirectly, that is, based on direct interaction or, for example, the observations of another agent, e.g., $tr(Ag1, Ag3) = 0.7$ using $Ag2$ as witness, in Figure 1.

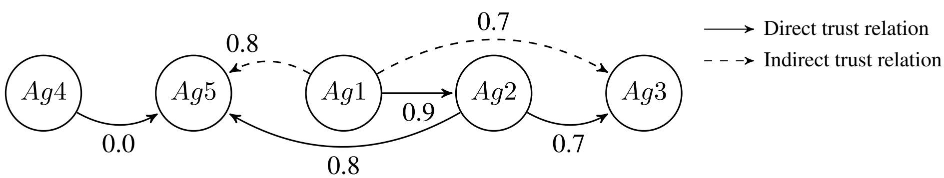

Figure 1. Trust Network Example [5]

When evaluating trust indirectly, an agent uses the observations of other agents. Such evaluation is based on the particularities of each agent's profile and it considers the different paths and levels of depth in its trust network [3,4,5]. During this process of considering the experience of other agents, an agent is able to model the other agents' private views on trust. Considering Figure 1, Ag_1 is able to model the trust Ag_2 has on Ag_3 and Ag_5, i.e., $\Delta^{Ag_1}_{\Pi'^{Ag_2}} \models tr(Ag_2, Ag_3) = 0.7$, and $\Delta^{Ag_1}_{\Pi'^{Ag_2}} \models tr(Ag_2, Ag_5) = 0.8$.

When an agent Ag_i has not evaluated its private trust on another agent Ag_j, so that $\Pi^{Ag_i} \not\models tr(Ag_i, Ag_j)$, the only immediate information of how much Ag_i can trust agent Ag_j is the information from the reputation system, that is $\Gamma \models rp(Ag_j)$, with $rp : Ags \mapsto \mathbb{R}$. In cases where an agent Ag_i has previously evaluated its private trust on another agent Ag_j, directly or indirectly, both pieces of information will be available to it. In this work, we will assume that a rational agent will always consider first its private evaluation, and when it is unable to access such information, it considers the reputation of that agent[2].

$$tr(Ag_i, Ag_j) = \begin{cases} tr(Ag_i, Ag_j) & \text{if } \Pi^{Ag_i} \models tr(Ag_i, Ag_j) \\ rp(Ag_j) & \text{otherwise, with } \Gamma \models rp(Ag_j) \end{cases} \quad (1)$$

3. Argumentation-Based Reasoning with Trust

Interesting work on combining trust and argumentation can be found in the literature. Most of that work focuses on how agents consider the trust on other agents in order to define preferences between arguments and to define which ones are acceptable [3,5], i.e., argumentation-based reasoning with trust. Based on that work, in this section, we describe how agents may reason about the acceptability of arguments considering trust.

In [5,6], arguments are constructed using defeasible inference rules. Similarly to other structured argumentation frameworks, arguments are represented as tuples $\langle S, c\rangle$, with S (the support) a set of premises and inference rules used to draw c (the conclusion of the argument). Considering the arguments' structure, conflict between arguments are of two types (we use "$\neg$" for strong negation and a general operator for contradictory information $\overline{\varphi} \equiv \neg\varphi$): (i) an argument $\langle S_1, c_1\rangle$ *rebuts* another argument $\langle S_2, c_2\rangle$ iff $c_1 \equiv \overline{c_2}$; and (ii) an argument $\langle S_1, c_1\rangle$ *undercuts* another argument $\langle S_2, c_2\rangle$ iff $c_1 \equiv \overline{\varphi}$ for some $\varphi \in S_2$.

When two arguments are in conflict, i.e., the arguments attack each other, this does not necessarily mean that an argument defeats the other. Defeat is a "successful" attack, and it considers the set of arguments that defend each other, including preferences between the conflicting arguments [7]. Thus, considering that the approach we presented in [6] is not able to decide some conflicts in which arguments rebut each other, we later extended it to consider different meta-information commonly available in MAS platforms, in order to define preferences between arguments. Among the meta-information considered, in [5] we gave special attention to trust as meta-information.

Trust can be applied to beliefs, based on the trust value applied to the sources of these beliefs [5]. We consider that both trust and reputation of an agent can be available. Thus, considering a belief φ of the agent Ag_1, φ can come from multiple sources and in order to know how much Ag_1 trusts φ, the *tr* value associated with each source of φ for Ag_1 must be considered. Thus, a function $trb_i : \varphi \mapsto \mathbb{R}$ can be defined so that $trb_i(\varphi)$ returns the trust value that Ag_i has for belief φ based on the trust level Ag_i has on the sources of the information φ. The operation that calculates $trb_i(\varphi)$ varies according to agent profiles, corresponding to different attitudes towards one's sources of information. In this paper we consider three different agents profiles: *optimistic*, *pessimistic*, and *fair*. An *optimistic*[3] agent considers the most trustworthy source (equation (2)), a *pessimistic*

[2]Note that $tr(Ag_i, Ag_j)$ aggregates the evaluation of direct and indirect trust presented in Section 2.

[3]Our *optimistic* profile corresponds to the *credulous* profile described in [5].

agent considers the least trustworthy source (equation (3)), and a *fair* agent considers the average of the trust it has in the sources (equation (4)).

$$trb_i(\varphi) = max\{tr(Ag_i, Ag_j), ..., tr(Ag_i, Ag_n)\} \quad (2)$$
$$trb_i(\varphi) = min\{tr(Ag_i, Ag_j), ..., tr(Ag_i, Ag_n)\} \quad (3)$$
$$trb_i(\varphi) = avg\{tr(Ag_i, Ag_j), ..., tr(Ag_i, Ag_n)\} \quad (4)$$

where $\{Ag_j, ..., Ag_n\}$ is the set of sources that informed φ to Ag_i. After calculating how much agents trust their beliefs, they are able to calculate how much they trust arguments.

The trust value on an argument depends on the values of each element in its support. Again, how to combine the trust values of beliefs to calculate a trust value for arguments depends on the desired behaviour of the agents. An *optimistic* agent will consider the most trustworthy piece of information in order to assign a trust value for that argument (equation (5)), a *pessimistic*[4] agent the least trustworthy piece of information (equation (6)), and a *fair* agent the average of all pieces of information (equation (7)).

$$tra_i(\langle S, c\rangle) = \max\{trb_i(\varphi_1), \ldots, trb_i(\varphi_n)\}. \quad (5)$$
$$tra_i(\langle S, c\rangle) = \min\{trb_i(\varphi_1), \ldots, trb_i(\varphi_n)\}. \quad (6)$$
$$tra_i(\langle S, c\rangle) = avg\{trb_i(\varphi_1), \ldots, trb_i(\varphi_n)\}. \quad (7)$$

with $tra_i : \langle S, c\rangle \mapsto \mathbb{R}$ and $S = \{\varphi_1, \ldots, \varphi_n\}$. When an agent has multiple arguments for the same conclusion c, for example, the argument $\langle S_1, c\rangle$ and $\langle S_2, c\rangle$, it can opt for the argument that has the highest trust value: $\max\{tra_i(\langle S_1, c\rangle), \ldots, tra_i(\langle S_n, c\rangle)\}$.

Considering that agents are able to calculate how much they trust arguments, based on how much they trust the information that composes that argument, agents are able to define preferences over arguments. Using such preferences, agents are able to resolve conflicts that they could not resolve previously [5]. Considering two arguments, $\langle S_1, c_1\rangle$ *defeats* $\langle S_2, c_2\rangle$ iff either $\langle S_1, c_1\rangle$ rebuts $\langle S_2, c_2\rangle$ and $tra(\langle S_1, c_1\rangle) > tra(\langle S_2, c_2\rangle)$ or $\langle S_1, c_1\rangle$ undercuts $\langle S_2, c_2\rangle$ at $\varphi \in S_2$, and there is no $\langle S_3, \varphi\rangle$ such that $tra(\langle S_1, c_1\rangle) < tra(\langle S_3, \varphi\rangle)$. An argument $\langle S, c\rangle$ is *acceptable* to an agent Ag iff for all other arguments that attacks $\langle S, c\rangle$, either $\langle S, c\rangle$ defeats such arguments, or there are other acceptable arguments that defeat those arguments attacking $\langle S, c\rangle$.

4. Dialogue Strategies Using Information about Trust

In this work, we propose a strategy for argumentation-based dialogues in which agents consider trust in order to define preferences between arguments. As discussed above, different choices regarding how agents consider trust will provide different agent profiles. Therefore, in this work, on one hand we propose a general approach in terms of such choices, modularising our choices as replaceable components, and on the other hand, the strategy we propose takes into consideration the profiles previously defined that the target of an argument could adopt; that is, the receiver could be *optimistic*, *pessimistic*, or *fair*.

Imagine that an agent Ag_i needs to construct an argument supporting a claim c for another agent Ag_j. When this happens, Ag_i needs to consider which arguments it is able to construct from the information available to it, and who are the sources of

[4]This is similar to the *sceptical* profile defined in [5], and it implements the weakest link principle.

information that Ag_j trusts. In order to describe this process, we introduce Algorithm 1, which describes the reasoning an agent may carry out in order to construct the best argument it can, considering the information it has. That is, it considers both the different arguments supporting c it is able to construct as well as the information about how much the recipient trusts the available sources of information.

In Algorithm 1, the agent starts looking for each argument it is able to construct from its knowledge base, depending on its profile, i.e., $\Delta^{Ag_i} \models^{pr} \langle S, c\rangle$. Different profiles have been defined in the argumentation literature for MAS; for example, [8] introduces different *assertion attitudes* defining different behaviours that agents are able to adopt towards asserting arguments, in particular the following agent profiles are defined: (i) a *confident* agent asserts any proposition for which it can construct an argument; and (ii) a *thoughtful* agent asserts any proposition for which it can construct an acceptable argument. That is, while a *confident* agent just needs to be able to construct an argument in order to put it forward during a dialogue, a *thoughtful* agent needs to be able to construct an acceptable argument, which means considering the set of arguments it is able to construct and the attack relations between them.

Algorithm 1 Argument Selection Algorithm

1: **procedure** ARGUMENTSELECTION(c, Ag_j)
2:     bestArg $\leftarrow$ null;
3:     argVal $\leftarrow$ 0;
4:     **for** each $\langle S, c\rangle$ which $\Delta^{Ag_i} \models^{pr} \langle S, c\rangle$ **do**
5:         $S' \leftarrow$ GETBESTSOURCES(Ag_j, S)
6:         tempArgVal $\leftarrow$ GETARGVAL(Ag_j,S')
7:         **if** argVal $\leq$ tempArgVal **then**
8:             bestArg $\leftarrow \langle S', c\rangle$
9:             argVal $\leftarrow$ tempArgVal
    return bestArg

$$\Delta^{Ag_i} \models^{pr} \langle S, c\rangle = \begin{cases} \texttt{true} & \text{if } \langle S, c\rangle \text{ is acceptable in } \Delta^{Ag_i} \\ \texttt{false} & \text{if } \langle S, c\rangle \text{ is not acceptable in } \Delta^{Ag_i} \end{cases} \tag{8}$$

In order to allow different attitudes to be considered, we define this choice in equation (8), in which we state that Ag_i will use only arguments that are acceptable to it, i.e., the *thoughtful* agent assertion attitude from [8].

Algorithm 2 Source Selection Algorithm

function GETBESTSOURCES(Ag_j,S)
    finalArg $\leftarrow$ { }
    for each $\varphi_{[Src]} \in$ S **do**
        bestSrc = BESTSOURCE(Ag_j,$Src \cup \{Ag_i\}$)
        finalArg $\leftarrow$ finalArg $\cup\ \varphi_{[\text{bestSrc}]}$
    return finalArg

Algorithm 3 Argument Value Algorithm

function GETARGVAL(Ag_j,S)
    Sources $\leftarrow$ { }
    for each $\varphi_{[\text{Src}]} \in$ S **do**
        Sources $\leftarrow$ Sources $\cup$ Src
    return ARGVAL(Ag_j,Sources)

After selecting one of the arguments Ag_i is able to construct in Algorithm 1, on line 5 it makes the selection of the best sources for the information it used on the support of that argument, considering how much the recipient of the argument trusts those sources. Thus, in Algorithm 2, for each information used, φ, Ag_i chooses the best sources from $[Src]$, according to how much the recipient trusts that sources (see equation (10)). Considering the three profiles previously introduced, the best general strategy is choosing only one source for each piece of information, which the recipient trust most (see equation (9)). Thus, the agent avoids a lower trust value when facing a *pessimistic* or *fair* agent.

$$\textsc{bestSource}(Ag_j, Src) = \arg\max_{s \in Src} \textsc{trustA}(Ag_j, s) \tag{9}$$

$$\textsc{trustA}(Ag_j, s) = \begin{cases} tr(Ag_j, s) & \text{if } \Delta^{Ag_i}_{\Pi'^{Ag_j}} \models tr(Ag_j, s) \\ rp(s) & \text{if } rp(s) \geq tr(Ag_i, s) \\ & \text{ with } \Delta^{Ag_i} \models tr(Ag_i, s) \\ tr(Ag_i, s) & \text{otherwise,} \\ & \text{with } \Delta^{Ag_i} \models tr(Ag_i, s) \end{cases} \tag{10}$$

Both equations (9) and (10) can be replaced by other choices depending on the application domain and the agent profiles being considered. Here, while equation (9) is part of the general strategy we are proposing in this paper, equation (10) describes how Ag_i considers how much others agents trust that information. Note that Ag_i considers first the recipient's private evaluation of trust, when this information is available, otherwise it considers the reputation of other agents. Also, Ag_i can provide its own observations of how much it trusts the others agents when such trust is superior to the other agents reputation (the last cases in equation (10)).

$$\textsc{argVal}(Ag_j, Src) = \min_{s \in Src} \textsc{trustA}(Ag_j, s) \tag{11}$$

After selecting the best source for each piece of information used in that argument, from the perspective of the recipient, using Algorithm 3, Ag_i aggregates the sources and calculates the trust value for that argument based on equation[5] (11). This equation can be easily replaced in order to model other strategies considering other agent profiles. Finally, comparing the strength of each argument, Ag_i finds the best argument it can construct based on the information it has. It is important to emphasise that, on one hand the *thoughtful* agent assertion attitude is used as the profile of the agent who is putting the argument forward as part of a dialogue, and on the other hand the choices in equations (9), (10), and (11) represent the considerations by this agent on how the other agents will evaluate that argument[6].

5. Example

In order to show our approach, we consider a scenario in which an agent Ag needs to construct an argument supporting c, for the recipient Ag_1, aiming to persuade Ag_1 about c. The state of Ag's knowledge base is as follows[7]:

$$\Delta^{Ag} = \left\{ \begin{array}{ll} x_{[s(Ag_4), s(Ag_2)]} & a_{[s(Ag_4), s(Ag_3), s(Ag_5)]} \\ y_{[s(Ag_4)]} & b_{[s(Ag_2), s(Ag_3), s(Ag_4)]} \\ (x, y \Rightarrow c)_{[s(Ag_2), s(Ag_5)]} & (a, b \Rightarrow c)_{[s(Ag_3), s(Ag_5)]} \end{array} \right\}$$

[5]Equation (11) implements the weakest link principle in which the trust on an argument will be equal to the minimum trust value found in its support.

[6]Corresponding to the choices agents make based on their profiles, i.e., equations (2) and (5) for an *optimistic* agent, (3) and (6) for a *pessimistic* agent, and (4) and (7) for a *fair* agent.

[7]We annotate the sources of each information in the format of labels $s(Ag)$, with Ag the agent name.

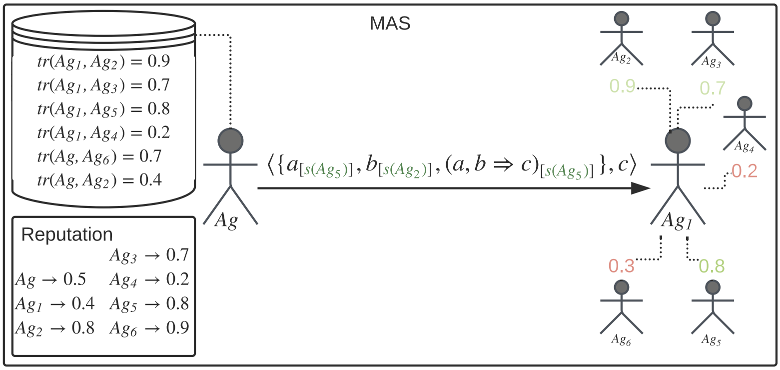

Figure 2. A strong argument for Ag_1, given that Ag_1 considers all sources used in that argument trustworthy.

Also, considering the trust network in Figure 1, and Ag's model about Ag_1 private knowledge regarding trust, shown in Figure 2, Ag is able to construct two arguments supporting c, i.e., $Arg_1 = \langle\{a, b, (a, b \Rightarrow c)\}, c\rangle$ and $Arg_2 = \langle\{x, y, (x, y \Rightarrow c)\}, c\rangle$. In this particular case, both arguments are acceptable for Ag, i.e., $\Delta^{Ag} \models^{pr} Arg_1$ and $\Delta^{Ag} \models^{pr} Arg_2$; therefore, both are considered by Algorithm 1. After that, Algorithm 2 selects the best sources, i.e., $\{a_{[s(Ag_5)]}, b_{[s(Ag_2)]}, (a, b \Rightarrow c)_{[s(Ag_5)]}\}$ for Arg_1 and $\{x_{[s(Ag_2)]}, y_{[s(Ag_4)]}, (x, y \Rightarrow c)_{[s(Ag_2)]}\}$ for Arg_2. Using the best sources for each information, Algorithm 3 returns the following argument values: $Arg_1 = 0.8$ and $Arg_2 = 0.2$. Then, Algorithm 1 returns Arg_1. Note that Arg_1 will have the following values for Ag_1, depending on Ag_1's profile: (i) *optimistic* = 0.9, *pessimistic* = 0.8, and *fair* = 0.86. There is no other source selection that provides a better result for this argument. Also, argument Arg_2 could have the following values for Ag_1, depending on Ag_1's profile: (i) *optimistic* = 0.9, *pessimistic* = 0.2, and *fair* = 0.66.

6. Related Work

The study of argumentation strategies has been considered an under-developed/neglected sub-area of argumentation in MAS [9,10,11]. Among the work that has investigated argumentation strategies, some authors have made unrealistic assumptions for the development of MAS based on the *open* paradigm, such as perfect knowledge. Such work has not been considered here. In contrast, some approaches have proposed strategic argumentation exploiting the *opponent model* [9,10,12,13,14]. Those approaches suggest that agents are able to choose which arguments to put forward in a dialogue, based on expected counterarguments from the opponent, given some search mechanism on the model othe opponent. Our work differs from those, given that, while they propose approaches for agents to choose *which* arguments to put forward in argumentation-based dialogues, our work focuses on how an agent chooses the best argument depending on the recipient. Another important question that arises from the work considering the opponent model is how does an agent obtain such information? In most of the approaches, such information is assumed to be given, which, as described in [11], is an unrealistic assumption in most applied settings. Here, we have detailed how agents can make use of the information from trust mechanisms to model other agents private trust.

7. Conclusion

In this work, we proposed an approach to creating strategies for agents to engage in argumentation-based dialogues in which agents choose the best argument they can depending on the views on trust held by the recipient of that argument. In particular, we focused on strategies for constructing the best arguments for agents who define preferences over arguments based on who are the sources of the information used in those arguments and how much they trust those sources. Differently from most work in strategic argumentation, we show how an agent can acquire the opponents model regarding trust and, after that, how the agent uses that information in order to construct stronger arguments for those recipients.

Considering the *logic*, *dialectic* and *rhetoric* dimensions of argumentation, as described in [15], to the best of our knowledge, our work is the first to consider strategies for the rhetoric dimension. While most work in the literature focus on which arguments an agent could put forward in a dialogue, our work focus on how strong and persuasive those arguments can be, based on the perspective of the recipient. We argue that both approaches are complementary, but we believe the rhetoric dimension might have long-term consequences in terms of inter-agent relationships, which is a subject that deserves further investigation.

References

[1] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial intelligence review*, 24(1):33–60, 2005.
[2] T. D. Huynh, N.R. Jennings, and N.R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *AAMAS*, pgs 119–154, 2006.
[3] S. Parsons, Y. Tang, E. Sklar, P. McBurney, and K. Cai. Argumentation-based reasoning in agents with varying degrees of trust. In *AAMAS*, pgs 879–886, 2011.
[4] Y. Tang, K. Cai, P. McBurney, E. Sklar, and S. Parsons. Using argumentation to reason about trust and belief. *Journal of Logic and Computation*, 22(5):979–1018, 2011.
[5] A. R. Panisson, V. S. Melo, and R. H. Bordini. Using preferences over sources of information in argumentation-based reasoning. In *Brazilian Conference on Intelligent Systems*, pgs 31–36, 2016.
[6] A. R. Panisson, F. Meneguzzi, R. Vieira, and R. H. Bordini. An Approach for Argumentation-based Reasoning Using Defeasible Logic in Multi-Agent Programming Languages. In *ArgMAS*, 2014.
[7] D. Walton, C. Reed, and F. Macagno. *Argumentation Schemes*. Cambridge University Press, 2008.
[8] S. Parsons and P. McBurney. Argumentation-based dialogues for agent co-ordination. *Group Decision and Negotiation*, 12(5):415–439, 2003.
[9] N. Oren and T. J. Norman. Arguing using opponent models. In *ArgMAS*, 2009.
[10] C. Hadjinikolis, Y. Siantos, S. Modgil, E. Black, and P. McBurney. Opponent modelling in persuasion dialogues. In *IJCAI*, pages 164–170, 2013.
[11] M. Thimm. Strategic argumentation in multi-agent systems. *KI-Künstliche Intelligenz*, 28(3):159–168, 2014.
[12] N. Hadidi, Y. Dimopoulos, P. Moraitis, et al. Tactics and concessions for argumentation-based negotiation. In *COMMA*, pages 285–296, 2012.
[13] T. Rienstra, M. Thimm, and N. Oren. Opponent models with uncertainty for strategic argumentation. In *IJCAI*, pages 332–338, 2013.
[14] M. Thimm and A. J. García. Classification and strategical issues of argumentation games on structured argumentation frameworks. In *AAMAS*, pgs 1247–1254, 2010.
[15] D. H. Cohen. Evaluating arguments and making meta-arguments. *Informal Logic*, 21(2), 2001.