Arguing coalitions in abstract argumentation

Lisha Qiao[1,2] Yiqi Shen[4] Liuwen Yu[1,2,3] Beishui Liao[1,4]
Leendert van der Torre[1,4]

[1]University of Luxembourg, Luxembourg
[2]University of Bologna, Italy
[3]University of Turin, Italy
[4]Zhejiang University, China

Abstract

In this paper, we are interested in different ways in which agents can collaborate in abstract agent argumentation. First, if arguments are accepted when they are put forward by more than one agent, then agents can put forward arguments from other agents of the coalition. Second, agents can put forward arguments to defend arguments from other agents of the coalition. For example, in expert opinion, a domain expert can put forward an argument defending an argument made by a politician, even when the politician cannot judge the correctness of the argument. Third, agents from a coalition can collectively defend an argument they share, without being able to defend the argument individually. In this paper, we formalize the different kinds of collaboration in abstract agent argumentation, and we illustrate the coalition formation with a case study in political debate.

Keywords: abstract argumentation, abstract agent argumentation, coalition formation

1 Introduction

Dung [14] defines an argumentation framework as a set of arguments and a binary relation between them. Agent argumentation frameworks [24] extend Dung's theory with agents or the sources of the arguments. Most approaches to agent argumentation are inspired by social choice and voting theory, and prefer arguments or attacks if they belong to more than one agent [5,6,10,17,20]. Moreover, in this setting, several authors [1,2,3,18] have considered the role of coalitions, for example in the setting of coalition formation [2]. Building on this line of work, in this paper we are interested in the following research question:

How to study coalition formation in abstract agent argumentation?

To answer this question, we adopt the minimal formal framework we introduced earlier for our principle-based analysis of agent argumentation semantics [24]. Agent-based extensions typically introduce various aspects such as knowledge, uncertainty, support, trust and so on. We use a minimal extension of Dung's theory [14] as a common core in those approaches. We limit this paper to an abstract set of agents, associate arguments with agents, and a partitioning of the set of agents to represent coalitions. Our research question breaks down into the following sub-questions:

(1) Which kinds of collaboration among arguing agents can be distinguished?

(2) How can Dung's notion of defense [14] be adapted to incorporate coalitions?

(3) How can we use this theory to reason about coalition formation?

To answer the first sub-question, we distinguish between three reasons for forming coalitions:

(1) Arguments can be accepted when they are put forward by more than one agent.

(2) Arguments can be defended by arguments from other agents.

(3) An argument shared by several agents can be collectively defended by a group without being defended by an individual agent.

For example, in expert opinion, a domain expert can put forward an argument that defends an argument made by a politician, even when the politician cannot judge the correctness of the argument.

To answer the second sub-question, we introduce the notion of coalition defense: arguments can only be defended by arguments from the same coalition. Everything else stays the same.

To answer the third sub-question, we use a running example from coalition formation in politics. A libertarian, a collectivist, an anarchist and a political expert form different coalitions in order to have as many of their arguments accepted as possible.

The potential application of this paper is to see how we can maximize agreement when we want to form a coalition in multi-agent systems (MASs). In the real world, we can also use this coalition framework to distinguish between different reasons for forming coalitions. Our investigation highlights various properties which can be further studied in this formal setting.

The layout of this paper is as follows. Section 2 introduces a case study in politics, and we informally describe the different kinds of collaboration in the example. Section 3 introduces the notion of coalition argumentation framework. Section 4 formalizes coalition defense as well as coalition semantics, and uses the running example to further explain these notions. Section 5 discusses several properties. We discuss related work in section 6, and future work in section 7. Section 8 concludes the paper.

2 Coalition formation in politics

In this section, we introduce a running example from political debate.

2.1 Political views

Let us suppose that there are four kinds of people with different political tendencies: Libertarian (L, for short), Collectivist (O, for short), Anarchist (N, for short), and Political expert(E, for short). They hold different points of view on whether government and laws are necessary for society.

L's point of view involves improving the functions of government and the law. We need government and laws. What government and laws do first of all is protect individual freedom.

O holds that we should improve the functions of government and law. Our government and laws protect the collective interest, so everyone can gain more interest. Sometimes, we need to limit individual freedom to comply with the law.

N argues that we don't need any government or law to protect individual freedom. We only need a government that serves people's interests. The government and laws that serves the elite firstly protect the individual interests of that minority. So, we don't need the current government and law.

E has two arguments: 1) today's elite-led government is committed to realizing collective interests, and 2) we need to limit individual freedom to comply with the law.

2.2 Political conflict

Obviously, there are conflicts in the views of these politicians. Even one politician's views will be self-contradictory.

When N only chooses to accept one of his/her own ideas, N prefers the former viewpoint that we don't need any government or law to protect individual freedom because even good government and laws can protect individual interests, including freedom; they may also limit and even deprive individuals of their personal freedom.

L advocates that we need government and laws whose first aim is to protect individual freedom, and objects to N's view that we don't need any government or law. L's point is more robust than N's because government and law are historical choices. And L's argument also contradicts E's argument that we need to limit individual freedom to comply with the law because individual freedom is not absolute but relative; if we obey the law, we don't need to restrict our freedom.

E's argument also attacks L's view when recognizing that freedom is absolute and needs to be limited. The conflict between N and L is clear. We cannot accept both L's view and N's view. N's two points are both stronger than L's view. Because we regard freedom as the most important thing, and any government and law can limit our freedom, we accept N's argument.

A conflict between N and O is also apparent. O is of the same opinion as L that we should improve the functions of government and the law. We already know that N's argument is more robust from this viewpoint. When it comes to individual interests, N's view is that the current government and law prioritize protecting the interests of minority people (the elite), and that we don't need the present government and law. O holds the opposite view.

Intuitively, political experts have professional political knowledge, and their viewpoints are often accepted. But N strongly contradicts the experts' statement. We cannot deny the existence of elite interest groups who are committed to realizing the interests of minority group members firstly. Political experts also admit this. So, N's proposal that we don't need the current government or law proposed is accepted.

2.3 Political coalition structure

There are contradictions in the views of every politician, and to make their ideas accepted, they can choose whether or not to form a coalition. These four politicians can create 15 coalition structures through partitioning. They can work in a grand coalition or work independently. Their coalitions can also be two against two (2:2), two against one against one (2:1:1), and three against one(3:1). For example, L and N form a coalition, O and E form another coalition. Alternatively, L, O and N form a coalition, and E works independently.

2.4 Working in a political coalition

If they can work in a grand coalition, O and E both support the view that we need to limit individual freedom to comply with the law. Intuitively, the more politicians who support this point of view, the more this view will be accepted. So, it is possible that the member of this coalition can accept this point of view in a coalition.

E has two arguments, and these arguments are attacked by L and N respectively. If (s)he doesn't collaborate with other politicians, (s)he cannot let other politicians accept his/her perspective. Suppose (s)he collaborates with other politicians to help argue that the current government and laws give priority to protecting the personal interests of the minority (the elite), and we don't need our current government and laws. In that case, his/her view that the current government is committed to realizing collective interests is likely to be accepted. Thus, E can collaborate with O or N. Although N attacks E, another view by N can help E refute N. Without the coalition, E does not know whose view can help him/her refute N.

L and O have a shared view that we should improve government and legal functions, but N has two arguments for attacking this view. If they work in the same coalition, they know that they can attack N's two points of view respectively such that their shared view can be accepted.

3 Coalition argumentation framework

This section introduces coalition argumentation frameworks. Coalition argumentation frameworks generalize argumentation frameworks studied by Dung [14], which are directed graphs, where the nodes are arguments, and the arrows correspond to the attack relation.

A coalition argumentation framework extends an argumentation framework with a set of agents, a relation associating arguments with agents, and a partitioning of the agents called a coalition structure. An argument can belong to no agent, one agent or multiple agents [24]. Each agent belongs to exactly one coalition.

We write $a \sqsubset \alpha$ to represent that argument a belongs to agent α, or that agent α has argument a. We write $C = \{\alpha\beta\}$ if agents α and β are in the same coalition, e.g. they work together in a political coalition.

Definition 3.1 [Coalition argumentation framework] A *coalition argumentation framework* (CAF) is a 5-tuple $\langle \mathcal{A}, \rightarrow, \mathcal{S}, \sqsubset, C \rangle$ where $\mathcal{A}$ is a set of arguments, $\rightarrow \subseteq \mathcal{A} \times \mathcal{A}$ is a binary relation over $\mathcal{A}$ called attack, $\mathcal{S}$ is a set of agents or sources, $\sqsubset \subseteq \mathcal{A} \times \mathcal{S}$ is a binary relation associating arguments with agents, and $C = \{\mathcal{S}_1, \mathcal{S}_2, ...\mathcal{S}_n\}$ is a set of disjoint subsets of $\mathcal{S}$ such that every agent of $\mathcal{S}$ occurs in exactly one of these subsets, representing the coalition structure.

The following example illustrates the formalization of the running example as a coalition argumentation framework.

Example 3.2 [Coalitions] Consider the running example in section 2. There are four agents L, E, O and N. Their arguments are as follows.
L has two arguments i and a.

i: We should improve the functions of government and the law.
a: We need government and laws.

O has three arguments i, g, l.

i: We should improve the functions of government and the law.
g: Our government and laws guarantee the collective interest, so everyone can gain more interest.
l: Sometimes, we need to limit individual freedom to comply with the law.

N has two arguments n and o.

n: We don't need any government and laws to protect individual freedom.
o: We only need the government and laws that serve people's interests.

E has two arguments c and l.

c: Today's elite-led government is committed to the collective interest.
l: We need to limit individual freedom to comply with the law.

Consider the coalition argumentation framework of the running example depicted in Figure 1, which contains $\mathcal{A} = \{i, n, o, a, g, c, l\}$, $\rightarrow = \{a \rightarrow n, n \rightarrow i, n \rightarrow o, g \rightarrow o, o \rightarrow g, o \rightarrow i, o \rightarrow c, l \rightarrow a, a \rightarrow l\}$, $\mathcal{S}$ = {L,E,O,N,}, $\sqsubset = \{(a, L), (i, L), (n, N), (o, N), (i, O), (g, O), (c, E), (l, O), (l, E)\}$.

An example of a coalition structure is C = {LON,E}, in which L, O and N work together and E works independently. We abbreviate this to C = {LON,E}. Using this abbreviation, all possible coalition structures are as follows. C_1 = {LEON}, C_2 = {L,O,N,E}, C_3 = {LON,E}, C_4 = {LEO,N}, C_5 = {ONE,L}, C_6 = {LEN,O}, C_7 = {LO,NE}, C_8 = {LN,OE}, C_9 = {ON,LE}, C_{10} = {LO,N,E}, C_{11} = {LN,O,E}, C_{12} = {LE,O,N}, C_{13} = {ON,L,E}, C_{14} = {OE,L,N}, C_{15} = {NE,L,O}.

Note that in this example, all arguments belong to at least one agent, and arguments i and l belong to two agents. Moreover, note that the arguments of most agents do not conflict with their own other arguments, only with those of the other agents. An exception is arguments n and o which belong to agent N. The attack from n to o means that without any interaction with other agents, agent N accepts n, but if this argument is rejected due to attacks from other agents, (s)he accepts argument o.

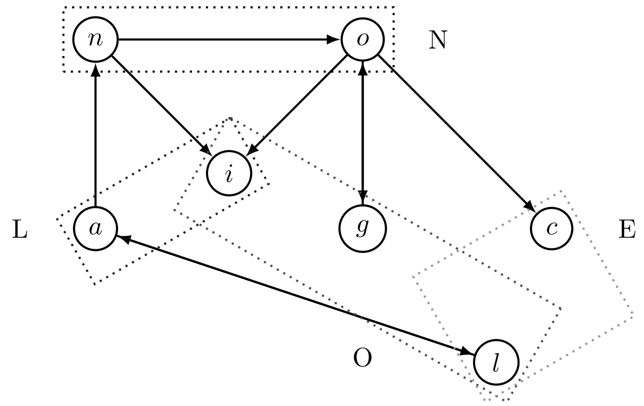

Fig. 1. Abstract framework of the political debate without coalitions

4 Coalition defense semantics

We now introduce a new kind of defense for coalition argumentation frameworks, which we call coalition defense. We adapt Dung's notions of defense and admissibility for coalition argumentation frameworks. Dung [14] defines conflict-freeness as a situation where attacking and attacked arguments cannot be accepted at the same time. A set of arguments defends an argument when the former attacks all attackers of the latter. A set is admissible if it is conflict-free and it defends all its elements. In our theory, roughly, if an agent puts forward an argument, it can only be coalition defended by arguments belonging to agents from the same coalition. Coalition admissibility means that a set of arguments is conflict-free and the coalition defends all its elements.

- **Definition 4.1** [Coalition admissible] Let $\langle \mathcal{A}, \rightarrow, \mathcal{S}, \sqsubset, C \rangle$:
  - $E \subseteq \mathcal{A}$ is *conflict-free* iff there are no arguments a and b in E such that a attacks b.
  - $E \subseteq \mathcal{A}$ coalition defends c iff there is a group of agents G in C, there exists an agent α in G having argument c, and for all arguments b in $\mathcal{A}$ attacking c, there exists an argument a in E such that a attacks b, and there is an agent β in G having a.
  - $E \subseteq \mathcal{A}$ is *coalition admissible* iff it is conflict-free and a coalition defends all its elements.

Example 4.2 [Political coalition, continued from Example 3.2] Consider the running example depicted in Figure 1. It illustrates several arguments that can be defended by arguments from other agents in the same coalition. For example, g defends c iff O and E are in the same coalition, and l defends n iff O and N are in the same coalition. Note also that n defends g iff N and O are in the same coalition, even when g also defends itself.

Figure 1 also illustrates that an argument shared by several arguments can be attacked by several arguments, where each argument cannot be defended against all these attacks by an individual agent, but can be defended when the agents work together. For example, a and g defend i iff L and O are in the same coalition.

Dung [14] presents admissibility-based semantics, which refers to an evaluation standard for the acceptability of argument sets and is used to select an acceptable argument set from a group of conflicting arguments. Coalition semantics is an evaluation criterion for the acceptability of argument sets, which is used to select acceptable argument sets from a set of contradictory arguments. It is represented as coalition extensions which are sets of acceptable arguments. And an empty set is considered to defend itself. The difference between coalition semantics and Dung's semantics is that in coalition semantics, arguments can only be defended by arguments from the coalition.

Each extension can be regarded as a set of collectively acceptable arguments. There are four kinds of coalition extensions: coalition complete extensions, coalition grounded extensions, coalition preferred extensions and coalition stable extensions. A coalition complete extension is conflict-free, and it can defend all elements and contains all the arguments it defends, so it's a set of closure. Coalition grounded extensions are the smallest coalition complete extensions, which minimizes the set of compatible arguments. Coalition preferred extensions are the largest coalition complete extensions, which maximizes the set of compatible arguments. Coalition stable extensions are conflict-free and attack all arguments that do not belong to the coalition.

Definition 4.3 [Coalition extensions] Let $\langle \mathcal{A}, \rightarrow, \mathcal{S}, \sqsubset, C \rangle$:

- $E \subseteq \mathcal{A}$ is a *coalition complete extension* iff E is coalition admissible and it contains all the arguments the coalition defends, $E = \{a | E$ coalition defends a$\}$.
- $E \subseteq \mathcal{A}$ is a *coalition grounded extension* iff it is the smallest (for set inclusion) coalition complete extension.
- $E \subseteq \mathcal{A}$ is a *coalition preferred extension* iff it is the largest (for set inclusion) coalition complete extension.
- $E \subseteq \mathcal{A}$ is a *coalition stable extension* iff it is conflict-free and it attacks all the arguments in $\mathcal{A} \backslash E$.

The following example illustrates coalition extensions.

Example 4.4 [Political coalition, continued from Example 4.2] Reconsidering the running example depicted in Figure 1, the extensions of each coalition are

Table 1
The coalition extension of all possible coalition structures. We abbreviate Coalition Grounded as $\mathbb{CG}$.

Sem.	Coalition Complete	$\mathbb{CG}$	Coalition Preferred	Coalition Stable
C_1 = {LEON}	$\{\emptyset, \{a\}, \{g,c\}, \{a,o\}, \{a,g,c,i\}, \{l,n,g,c\}\}$	$\{\emptyset\}$	$\{\{a,o\}, \{a,g,c,i\}, \{l,n,g,c\}\}$	$\{\{a,o\}, \{a,g,c,i\}, \{l,n,g,c\}\}$
C_2 = {L,E,N,O}	$\{\emptyset, \{a\}, \{g\}, \{l\}, \{a,g\}, \{l,g\}\}$	$\{\emptyset\}$	$\{\{a,g\}, \{l,g\}\}$	$\times$
C_3 = {LNO,E}	$\{\emptyset, \{a\}, \{g\}, \{a,o\}, \{a,g,i\}, \{l,n,g\}\}$	$\{\emptyset\}$	$\{\{a,o\}, \{a,g,i\}, \{l,n,g\}\}$	$\{\{a,o\}\}$
C_4 = {LEO,N}	$\{\emptyset, \{a\}, \{l\}, \{g,c\}, \{l,g,c\}, \{a,g,c,i\}\}$	$\{\emptyset\}$	$\{\{l,g,c,\}, \{a,g,c,i\}\}$	$\{\{a,g,c,i\}\}$
C_5 = {ENO,L}	$\{\emptyset, \{a\}, \{g,c\}, \{a,g,c\}, \{l,n,g,c\}\}$	$\{\emptyset\}$	$\{\{a,g,c,\}, \{l,n,g,c\}\}$	$\{\{l,n,g,c\}\}$
C_6 = {ELN,O}	$\{\emptyset, \{a\}, \{l\}, \{g\}, \{a,o\}, \{a,g\}, \{l,g\}\}$	$\{\emptyset\}$	$\{\{a,o\}, \{a,g\}, \{l,g\}\}$	$\{\{a,o\}\}$
C_7 = {LO,EN}	$\{\emptyset, \{a\}, \{l\}, \{g\}, \{l,g\}, \{a,g,i\}\}$	$\{\emptyset\}$	$\{\{l,g\}, \{a,g,i\}\}$	$\times$
C_8 = {LN,EO}	$\{\emptyset, \{a\}, \{l\}, \{g,c\}, \{a,o\}, \{a,g,c\}, \{l,g,c\}\}$	$\{\emptyset\}$	$\{\{a,o\}, \{a,g,c\}, \{l,g,c\}\}$	$\{\{a,o\}\}$
C_9 = {NO,EL}	$\{\emptyset, \{a\}, \{g\}, \{l,n\}, \{a,g\}, \{l,n,g\}\}$	$\{\emptyset\}$	$\{\{a,g\}, \{l,n,g\}\}$	$\times$
C_{10} = {LO,E,N}	$\{\emptyset, \{a\}, \{l\}, \{g\}, \{l,g\}, \{a,g,i\}\}$	$\{\emptyset\}$	$\{\{l,g\}, \{a,g,i\}\}$	$\times$
C_{11} = {LN,E,O}	$\{\emptyset, \{a\}, \{g\}, \{l\}, \{a,o\}, \{a,g\}, \{l,g\}\}$	$\{\emptyset\}$	$\{\{a,o\}, \{a,g\}, \{l,g\}\}$	$\{\{a,o\}\}$
C_{12} = {EL,N,O}	$\{\emptyset, \{a\}, \{l\}, \{g\}, \{a,g\}, \{l,g\}\}$	$\{\emptyset\}$	$\{\{a,g\}, \{l,g\}\}$	$\times$
C_{13} = {ON,E,L}	$\{\emptyset, \{a\}, \{g\}, \{l,n\}, \{a,g\}, \{l,n,g\}\}$	$\{\emptyset\}$	$\{\{a,g\}, \{l,n,g\}\}$	$\times$
C_{14} = {EO,L,N}	$\{\emptyset, \{a\}, \{l\}, \{g,c\}, \{a,g,c\}, \{l,g,c\}\}$	$\{\emptyset\}$	$\{\{a,g,c\}, \{l,g,c\}\}$	$\times$
C_{15} = {EN,L,O}	$\{\emptyset, \{a\}, \{l\}, \{g\}, \{a,g\}, \{l,g\}\}$	$\{\emptyset\}$	$\{\{a,g\}, \{l,g\}\}$	$\times$

listed in Table 1.

If all agents are in a grand coalition, it is the same as Dung's graph. If all agents work independently, a and g, or l and g can be accepted. a and g cannot coalition defend i, because these arguments do not pertain to agents from the same coalition. i is not accepted. i is accepted when L and O are in the same coalition, for example $C_3 = \{\text{LNO,E}\}$. i is not accepted when L and O are not in the same coalition, such as $C_2 = \{\text{L,E,N,O}\}$. An argument shared by several agents can be attacked by several arguments, where each argument cannot be defended against all these attacks by an individual agent, but can be defended when the agents work together.

c is not accepted when E and O are not in the same coalition, like in $C_3 = \{\text{LNO,E}\}$. We accept c when O and E are in the same coalition or N and E are in the same coalition, such as $C_6 = \{\text{LEN,O}\}$. We can say that arguments can be defended by arguments from other agents in the same coalition.

Example 4.5 [Political coalition, continued from Example 4.4] Reconsider the running example depicted in Figure 1. We now briefly explain the idea of social semantics that prefers arguments that belong to more than one agent. Since argument l belongs to two agents and argument a belongs to only one agent, we can say that argument l is preferred to argument a. This can be represented by removing the attack from a to l. The effect on the extensions in Table 1 is that l is now in all extensions, and all extensions with a can be removed from Table 1.

Definition 4.6 [Restricted coalition defend] Let $\langle \mathcal{A}, \rightarrow, \mathcal{S}, \sqsubset, C \rangle$:
E restricted coalition defends $a \iff$

- E coalition defends a (e.g. based on coalition C_i)
- $\forall$ agent $\alpha \in C_i$, $\forall$ argument $\beta \sqsubset \alpha$, a and β do not attack each other.

This definition makes sure that a coalition cannot be formed randomly. Agents who hold contradictory arguments will not tend to form coalitions in most cases, except when they have no conflicting interests. By this definition, agents who holds contradictory arguments can still form a coalition, but the cost is that they cannot defend the arguments in conflict.

Also, we can define restricted coalition admissible and restricted coalition semantics as we show above.

Definition 4.7 [Self-organizing coalition] A coalition C_i is a self-organizing coalition about argument $a \iff \forall$ agent $\alpha \in C_i$, $\forall$ argument $\beta \sqsubset \alpha$, a and β do not attack each other.

This definition shows which agents will tend to form a coalition to defend a specific argument.

In particular, we notice that: if coalition E defends a based on a self-organizing coalition about a, then E restricted coalition defends a.

5 Properties

Given a coalition argumentation framework $CAF = \langle \mathcal{A}, \rightarrow, \mathcal{S}, \sqsubset, C \rangle$, the coalition demonstrates different properties under different semantics as follows:

First, under stable semantics, since every argument that does not belong to an extension should be attacked by the extension, each possible coalition that makes non-empty extensions should be able to involve all arguments. There might be more than one coalition that will enforce an extension or a set of extensions. So, we may define a partial order over a set of coalitions in terms of some criteria, such as their size, authorities and values, etc.

Second, under grounded semantics, if $\langle \mathcal{A}, \rightarrow \rangle$ has no non-empty grounded extension, it is impossible to form a coalition that may enforce a non-empty grounded extension. Under preferred semantics, there exists at least one extension. So, there exists at least one coalition to enforce each preferred extension of $\langle \mathcal{A}, \rightarrow \rangle$. In this case, we may also define a partial order of coalitions.

Third, if we replace all the agents of a coalition by one agent, then we derive the same extensions under agent defense semantics.

6 Related work

Given a set of agents and a set of arguments where each agent may have several arguments and an argument may belong to several agents, we are interested in how agents have the ability to defend arguments in the form of coalitions. There are other variants of semantics that adapt these notions, such as weak defense for weak admissibility semantics [7], but that is not based on the agent metaphor. Arisaka and Satoh [3] adapt the notion of conflict-free to conflict-eliminability and then apply their four new coalition formability semantics to a Japanese political example, while in our paper, we use a concrete and realistic running example to show coalition defense and the corresponding new semantics. Kontarinis and Toni [19] analyse the identification of the malicious behavior of agents in the form of a bipolar argumentation framework which, together with the work of Panisson et al. [21], may inspire work of agent reduction semantics based on trustfulness. Our own recent work [24] involved a complete analysis of four types of semantics of agent argumentation.

There is a lot of early work on how to generate coalitions of agents with different mechanisms [1,13,18]. Boella et al. [8] developed social viewpoints for arguing about coalitions, considering attacks on attacks in the context of reasoning about coalitions. Amgoud [2] discusses task allocation via coalition formation. She points out that agents need to form a coalition to help each other in order to fulfill joint tasks better, and proposes a framework where several coalition structures can be generated and then evaluated by agents based on their preferences. She presents a proof theory that agents do not need to test the whole structure in order to check whether or not the given coalition is good. In our work, we pre-define the coalition of agents (only agents in the same coalition have the ability to defend each other's arguments) and then we evaluate the semantics of the framework through the new defense. Arisaka et al. [4] extend agent argumentation frameworks with coalitions among agents. Rienstra et al. [22] consider the case where agents may have different semantics, for example one agent uses grounded semantics and another agent uses preferred semantics. Bulling et al. [9] use Alternating-time temporal logic (ATL) as the technical basis for modelling coalition formation. They extend ATL for modelling coalitions through argumentation and reasoning about the abilities of coalitions of agents and the formation of coalitions, where coalition formation is part of the logical language. In our work, we model coalition formation in a different way by extending Dung's abstract argumentation, associating arguments with agents, and partitioning of the agents to represent coalitions. They define defense and conflict-free based on defeat rather than attack, and put forward a valid coalition in which members in a coalition are undefeated. In our work, the arguments proposed by agents in the same coalition can conflict.

Cayrol & Lagasquie-Schiex [11] propose a coalition of arguments instead of a coalition of agents in the setting of bipolar argumentation. They define a meta-framework consisting of a set of meta-arguments as well as their conflict relation such that an attack exists only at the meta-level. Support relation is used to relate members in a meta-argument. Their work differs from ours in how the coalition of arguments is conflict-free. Whereas they argue that arguments represent agents and only the ones who want to cooperate come together, in our work, we allow arguments belonging to agents who are in a coalition conflict. Another difference is that they follow Dung's methodology for defining semantics with the construction of a meta-argumentation framework, while our work defines coalition defense and coalition semantics.

To meet the maximum agreement, Leite & Martins [20] introduce an abstract model of argumentation where agents can vote for or against an issue. Another related work is the aggregation of individual views into collective acceptability, which in general is split into two directions: semantics aggregation and structured aggregation. Some authors [5,10] capture the notion that individual members need to defend collective decisions in order to reach a compatible outcome, and propose to address judgement aggregation by combining different individual evaluations (semantics) of the situation represented by an argumentation framework. On the other hand, Chen et al. [12] evaluate how to aggregate abstract argumentation frameworks with the preservation of semantics. Hunter et al. [16] take an epistemic approach to probabilistic argumentation where arguments are believed or not believed by varying degrees, providing an alternative to the subtle standard in Dung's framework. Due to space limitations, we focus on the form of coalitions, while we also inflect the maximum agreement through social semantics that prefers arguments that belong to more than one agent.

Another related work is concept accrual, *i.e.* arguments that cannot defeat their target on their own but can succeed together [23] can also increase the acceptability of their inner arguments. For example, let A, B, C be three arguments. A defeats B, A defeats C, neither B nor C can defeat A but B and C accrues to defeat A. Accrual differs from coalition. The former increases the acceptability of arguments through preferences or strength. It mainly concerns the defeat relation between arguments (accruals). The latter increases the acceptability of arguments through defense. It has nothing to do with defeat or preferences. In our future work, it may be valuable to incorporate accrual in CAF. The key points are:

- defining preferences among sets of arguments on the strength of each argument;
- prescribing that arguments can only accrue with arguments in the same coalition;
- identical arguments that belong to different agents can accrue to strengthen themselves.

In our work, we define coalitions based on agents, not arguments, for which we define the new notion of coalition defense. However, the semantics are the same. Our choice can more strongly indicate the meaning of a coalition in the sense used by society, i.e. agency.

7 Future work

The example can be considered in another way. In our example, the coalition helps the agents to jointly make their arguments stronger and increase the acceptability of the arguments they defend. Intuitively, the coalition may also make the arguments it defends weaker or decrease their acceptability.

The idea of persuasion in the work of Anthony Hunter [15], where he talks about persuasion using probability argumentation, may give us some insight. When one agent wants to persuade another, (s)he has uncertain information. We can use a coalition about a set of agents, and probably, we don't know with certainty which argument would be proposed by others, but we know the probability of that. We can do this by forming a coalition that has a higher probability of reaching some goals.

The main goal of collective argumentation is to achieve maximum agreement among a set of agents. Thus, another future work is to discuss how we can reach maximum agreement when we want to form a coalition.

In our paper, the running example is static, but presumably agents not only enter into coalitions but also leave them. This issue of dynamics can be addressed by modularity. When agents join or leave, we just consider the arguments that attack or are attacked by the arguments which belong to the changed agent. In this way, we can get new extensions.

We consider all the possibilities of making up coalitions. And one important thing is to evaluate different coalitions: which is better and what's the criterion? One potential way is to consider the stability of coalitions. If we think of agents taking a principled stand on issues, their positions dictate which coalitions they can enter into and sustain even as the coalitions expand. The strength of arguments can then be measured by the stability of coalitions coming together in support of those arguments and against counter-arguments. It is also related to the dynamics of coalition.

8 Conclusions

We have adapted Dung's notion of defense [14] to incorporate coalitions and proposed a coalition argumentation framework. In this formal framework, we mainly distinguished between two reasons for forming coalitions and find the properties of coalitions.

In this paper we investigated how coalition formation can be analyzed among agents that put forward arguments and try to persuade each other. Coalition formation is typically studied in a game theoretic setting, but we adopted the standard abstract model of argumentation introduced by Dung, we associated arguments with agents, and we introduced new coalition semantics.

We distinguished between three kinds of collaboration among arguments. First, agents may form coalitions to put forward the same arguments, if putting forward the same arguments increases the strength of such arguments in the debate. Second, agents may put forward arguments to defend the arguments of other agents in the coalition. Third, agents may work together to provide defenses for the attackers on their shared arguments.

In our formal approach, we focused on the second and third kinds of collaboration, because the first is already widely studied in social semantics using techniques from voting theory and social choice. We extended Dung's theory of abstract argumentation in two ways. First, we introduced agents and coalition structures in argumentation frameworks. Secondly, we adapted the notion of defense [14] such that arguments can only be defended by arguments from the same coalition. Everything else stayed exactly the same.

We showed how this theory can be used for coalition formation using an extended case study in political debate. With four agents and seven arguments, we showed how coalitions affect the accepted arguments for various kinds of coalition semantics. Considering the political debate example, if L doesn't collaborate with O, argument i will not be accepted.

Our investigation highlighted some properties that can be studied in a formal setting. We leave the formal analysis and the formalization of the formation of coalitions to the journal extension of this paper.

Acknowledgements

This work has received funding from the EU H2020 research and innovation programme under Marie Skłodowska-Curie Actions Innovative Training Networks European Joint Doctorate grant agreement No. 814177 in Law, Science and Technology, a Joint Doctorate on the Rights of Internet of Everything.

References

[1] Samir Aknine, Suzanne Pinson, and M Shakun. Coalition formation problem: New multi-agent methods with preference models. In *Proceedings of the AAAI workshop on Coalition formation in dynamic multi-agent environments*, pages 51–56, 2002.

[2] Leila Amgoud. An argumentation-based model for reasoning about coalition structures. In *International Workshop on Argumentation in Multi-Agent Systems*, pages 217–228. Springer, 2005.

[3] Ryuta Arisaka and Ken Satoh. Coalition formability semantics with conflict-eliminable sets of arguments. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil*, pages 1469–1471, 2017.

[4] Ryuta Arisaka, Ken Satoh, and Leendert W. N. van der Torre. Anything you say may be used against you in a court of law - abstract agent argumentation (triple-a). volume 10791, pages 427–442. Springer, 2017.

[5] Edmond Awad, Richard Booth, Fernando Tohmé, and Iyad Rahwan. Judgement aggregation in multi-agent argumentation. *J. Log. Comput.*, 27(1):227–259, 2017.

[6] Haris Aziz, Felix Brandt, Edith Elkind, and Piotr Skowron. Computational social choice: The first ten years and beyond. In *Computing and software science*, pages 48–65. Springer, 2019.

[7] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Comparing weak admissibility semantics to their Dung-style counterparts–reduct, modularization, and strong equivalence in abstract argumentation. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 17, pages 79–88, 2020.

[8] Guido Boella, Leendert Van Der Torre, and Serena Villata. Social viewpoints for arguing about coalitions. In *Pacific Rim International Conference on Multi-Agents*, pages 66–77. Springer, 2008.

[9] Nils Bulling, Jürgen Dix, and Carlos Iván Chesnevar. Modelling coalitions: Atl+ argumentation. In *AAMAS (2)*, pages 681–688, 2008.

[10] Martin Caminada and Gabriella Pigozzi. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems*, 22(1):64–102, 2011.

[11] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Coalitions of arguments: A tool for handling bipolar argumentation frameworks. *Int. J. Intell. Syst.*, 25(1):83–109, 2010.

[12] Weiwei Chen and Ulle Endriss. Preservation of semantic properties in collective argumentation: The case of aggregating abstract argumentation frameworks. *Artificial Intelligence*, 269:27–48, 2019.

[13] Viet Dung Dang and Nicholas R. Jennings. Generating coalition structures with finite bound from the optimal guarantees. In *3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, pages 564–571. IEEE Computer Society, 2004.

[14] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357, 1995.

[15] Anthony Hunter. Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In Qiang Yang and Michael J. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 3055–3061. AAAI Press, 2015.

[16] Anthony Hunter, Sylwia Polberg, and Matthias Thimm. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artif. Intell.*, 281:103236, 2020.

[17] Nikos Karanikolas, Pierre Bisquert, and Christos Kaklamanis. A voting argumentation framework: Considering the reasoning behind preferences. In *ICAART 2019-11th International Conference on Agents and Artificial Intelligence*, volume 1, pages 42–53. SCITEPRESS-Science and Technology Publications, 2019.

[18] Matthias Klusch and Andreas Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17(3):42–47, 2002.

[19] Dionysios Kontarinis and Francesca Toni. Identifying malicious behavior in multi-party bipolar argumentation debates. In *Multi-Agent Systems and Agreement Technologies - 13th European Conference*, volume 9571 of *Lecture Notes in Computer Science*, pages 267–278. Springer, 2015.

[20] João Leite and João G. Martins. Social abstract argumentation. In Toby Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 2287–2292. IJCAI/AAAI, 2011.

[21] Alison R. Panisson, Simon Parsons, Peter McBurney, and Rafael H. Bordini. Choosing appropriate arguments from trustworthy sources. In *Computational Models of Argument - Proceedings of COMMA 2018, Warsaw, Poland, 12-14 September 2018*, volume 305, pages 345–352. IOS Press, 2018.

[22] Tjitze Rienstra, Alan Perotti, Serena Villata, Dov M Gabbay, and Leendert van der Torre. Multi-sorted argumentation. In *International Workshop on Theorie and Applications of Formal Argumentation*, pages 215–231. Springer, 2011.

[23] Julien Rossit, Jean-Guy Mailly, Yannis Dimopoulos, and Pavlos Moraitis. United we stand: Accruals in strength-based argumentation. *Argument & Computation*, 12:87–133, 2020.

[24] Liuwen Yu, Dongheng Chen, Lisha Qiao, Yiqi Shen, and Leendert van der Torre. A principle-based analysis of abstract agent argumentation semantics. 2021. under review.