Preferences and Constraints in Abstract Argumentation

Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, Irina Trubitsyna

Department of Informatics, Modeling, Electronics and System Engineering (DIMES)
University of Calabria
Italy
{g.alfano, greco, fparisi, i.trubitsyna}@dimes.unical.it

32nd International Joint Conference on Artificial Intelligence
19th-25th August 2023
Macao, S.A.R.

Argumentation and Domain Knowledge

- A general way for representing arguments and relationships (rebuttals) between them
- It allows representing dialogues, making decisions, and handling inconsistency and uncertainty

Abstract Argumentation Framework (AF) [Dung1995]: arguments are abstract entities (no attention is paid to their internal structure) that may attack and/or be attacked by other arguments

Example (AF describing what a person is going to have for lunch)

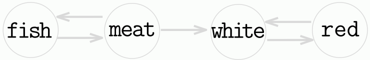

(S)he will have either `fish` or `meat`, and will drink either `white` wine or `red` wine. However, if (s)he will have `meat`, then (s)he will not drink `white` wine. Every solution (extension according to a semantics) represents a menu.

- However, in some cases it is difficult to accurately model domain knowledge in a natural and easy-to-understand way

Argumentation and Domain Knowledge

- A general way for representing arguments and relationships (rebuttals) between them
- It allows representing dialogues, making decisions, and handling inconsistency and uncertainty

Abstract Argumentation Framework (AF) [Dung1995]: arguments are abstract entities (no attention is paid to their internal structure) that may attack and/or be attacked by other arguments

Example (AF describing what a person is going to have for lunch)

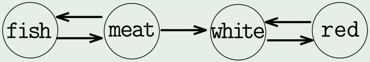

(S)he will have either `fish` or `meat`, and will drink either `white` wine or `red` wine. However, if (s)he will have `meat`, then (s)he will not drink `white` wine. Every solution (extension according to a semantics) represents a menu.

- However, in some cases it is difficult to accurately model domain knowledge in a natural and easy-to-understand way

Argumentation and Domain Knowledge

- A general way for representing arguments and relationships (rebuttals) between them
- It allows representing dialogues, making decisions, and handling inconsistency and uncertainty

Abstract Argumentation Framework (AF) [Dung1995]: arguments are abstract entities (no attention is paid to their internal structure) that may attack and/or be attacked by other arguments

Example (AF describing what a person is going to have for lunch)

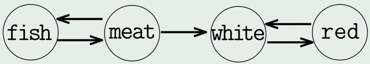

(S)he will have either `fish` or `meat`, and will drink either `white` wine or `red` wine. However, if (s)he will have `meat`, then (s)he will not drink `white` wine. Every solution (extension according to a semantics) represents a menu.

- However, in some cases it is difficult to accurately model domain knowledge in a natural and easy-to-understand way

Preferences and Constraints in AF

- Several proposals have been made to extend AF with the aim of better modeling the knowledge to be represented, e.g. Preference-based AF (PAF) and Constrained AF (CAF)
- We introduce three new frameworks generalizing PAF and CAF:

1) *extended Preference-based AF* (ePAF), an extension of (P)AF where preferences are 3-valued (e.g. $\texttt{red}^{\mathbf{t}} \succ \texttt{red}^{\mathbf{u}}$)
2) *extended Preference-based Constrained AF* (ePCAF), combining the features of CAF and ePAF
3) *multi-agent* ePCAF (mPCAF), dealing with multiple agents sharing the same AF and having different constraints and preferences

- Complexity of verification, credulous/skeptical acceptance problems

	AF			CAF			PAF			ePAF / ePCAF / mPCAF		
σ	Ver_σ	CA_σ	SA_σ	Ver_σ	CA_σ	SA_σ	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}
co	P	*NP*-c	P	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	P	co*NP*-c	Σ_2^P-c	Π_2^P-c
st	P	*NP*-c	co*NP*-c	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c
pr	co*NP*-c	*NP*-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P
ss	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P

Preferences and Constraints in AF

- Several proposals have been made to extend AF with the aim of better modeling the knowledge to be represented, e.g. Preference-based AF (PAF) and Constrained AF (CAF)
- We introduce three new frameworks generalizing PAF and CAF:

1) *extended Preference-based AF* (ePAF), an extension of (P)AF where preferences are 3-valued (e.g. $\texttt{red}^{\mathbf{t}} \succ \texttt{red}^{\mathbf{u}}$)
2) *extended Preference-based Constrained AF* (ePCAF), combining the features of CAF and ePAF
3) *multi-agent* ePCAF (mPCAF), dealing with multiple agents sharing the same AF and having different constraints and preferences

- Complexity of verification, credulous/skeptical acceptance problems

	AF			CAF			PAF			ePAF / ePCAF / mPCAF		
σ	Ver_σ	CA_σ	SA_σ	Ver_σ	CA_σ	SA_σ	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}
co	P	*NP*-c	P	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	P	co*NP*-c	Σ_2^P-c	Π_2^P-c
st	P	*NP*-c	co*NP*-c	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c
pr	co*NP*-c	*NP*-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P
ss	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P

Outline

Argumentation Semantics and Decision Problems

- Several semantics have been proposed to identify “reasonable” sets of arguments (called *extensions*)

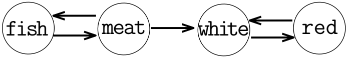

Semantic σ	Set of $\sigma-$extensions of AF Λ
complete (co)	$\{E_0 = \emptyset, E_1 = \{\texttt{fish}, \texttt{white}\}, E_2 = \{\texttt{fish}, \texttt{red}\}, E_3 = \{\texttt{meat}, \texttt{red}\}, E_4 = \{\texttt{fish}\}, E_5 = \{\texttt{red}\}\}$
preferred (pr)	$\{E_1, E_2, E_3\}$
semi-stable (sst)	$\{E_1, E_2, E_3\}$
stable (st)	$\{E_1, E_2, E_3\}$
grounded (gr)	$\{E_0\}$

- Verification problem: Ver_σ is the problem of checking whether a given set of arguments is a σ-extension
- Credulous (resp. Skeptical) acceptance problem: for a goal argument g, CA_σ (resp. SA_σ) consists in deciding whether g belongs to at least one (resp. every) σ-extension

Argumentation Semantics and Decision Problems

- Several semantics have been proposed to identify "reasonable" sets of arguments (called *extensions*)

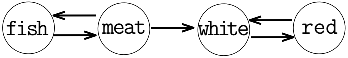

Semantic σ	Set of $\sigma-$extensions of AF Λ
complete ($\mathtt{co}$)	$\{E_0 = \emptyset, E_1 = \{\mathtt{fish}, \mathtt{white}\}, E_2 = \{\mathtt{fish}, \mathtt{red}\}, E_3 = \{\mathtt{meat}, \mathtt{red}\}, E_4 = \{\mathtt{fish}\}, E_5 = \{\mathtt{red}\}\}$
preferred ($\mathtt{pr}$)	$\{E_1, E_2, E_3\}$
semi-stable ($\mathtt{sst}$)	$\{E_1, E_2, E_3\}$
stable ($\mathtt{st}$)	$\{E_1, E_2, E_3\}$
grounded ($\mathtt{gr}$)	$\{E_0\}$

- Verification problem: Ver_σ is the problem of checking whether a given set of arguments is a σ-extension
- Credulous (resp. Skeptical) acceptance problem: for a goal argument g, CA_σ (resp. SA_σ) consists in deciding whether g belongs to at least one (resp. every) σ-extension

Preference-based AF

- A *Preference-based AF (PAF)* is an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ augmented with a preference relation $>$, that is a strict partial order over the set of arguments, e.g. $\texttt{meat} > \texttt{fish}$
- "best extensions" semantics: the preference relation $>$ is used to obtain a preference relation $\sqsupseteq$ over the extensions of the underlying AF, and then the *best* extensions w.r.t. $\sqsupseteq$ are selected
- Different proposals to define the best extensions, we focus on KVT: $E \sqsupseteq F$ if $\forall a, b \in \mathcal{A}$ the relation $a > b$ with $a \in F \setminus E$ and $b \in E \setminus F$ does not hold ($E \sqsupset F$, if $E \sqsupseteq F$ and $F \not\sqsupseteq E$)
- Given a PAF $\Delta = \langle \mathcal{A}, \mathcal{R}, > \rangle$, the best σ-extensions of Δ are the extensions $E \in \sigma(\langle \mathcal{A}, \mathcal{R} \rangle)$ such that there is no $F \in \sigma(\langle \mathcal{A}, \mathcal{R} \rangle)$ with $F \sqsupset E$

Example (PAF built from our example AF by adding preference $\texttt{meat} > \texttt{fish}$)

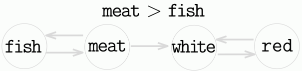

preferred extensions: $\{E_1 = \{\texttt{fish}, \texttt{white}\}, E_2 = \{\texttt{fish}, \texttt{red}\}, E_3 = \{\texttt{meat}, \texttt{red}\}\}$
$E_3 \sqsupset E_1$ and $E_3 \sqsupset E_2$

E_3 is the only best preferred extension (under KTV criterion)

Preference-based AF

- A *Preference-based AF (PAF)* is an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ augmented with a preference relation $>$, that is a strict partial order over the set of arguments, e.g. $\texttt{meat} > \texttt{fish}$
- "best extensions" semantics: the preference relation $>$ is used to obtain a preference relation $\sqsupseteq$ over the extensions of the underlying AF, and then the *best* extensions w.r.t. $\sqsupseteq$ are selected
- Different proposals to define the best extensions, we focus on KVT: $E \sqsupseteq F$ if $\forall a, b \in \mathcal{A}$ the relation $a > b$ with $a \in F \setminus E$ and $b \in E \setminus F$ does not hold ($E \sqsupset F$, if $E \sqsupseteq F$ and $F \not\sqsupseteq E$)
- Given a PAF $\Delta = \langle \mathcal{A}, \mathcal{R}, > \rangle$, the best σ-extensions of Δ are the extensions $E \in \sigma(\langle \mathcal{A}, \mathcal{R} \rangle)$ such that there is no $F \in \sigma(\langle \mathcal{A}, \mathcal{R} \rangle)$ with $F \sqsupset E$

Example (PAF built from our example AF by adding preference $\texttt{meat} > \texttt{fish}$)

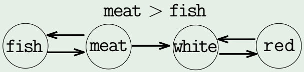

preferred extensions: $\{E_1 = \{\texttt{fish}, \texttt{white}\}$, $E_2 = \{\texttt{fish}, \texttt{red}\}$, $E_3 = \{\texttt{meat}, \texttt{red}\}\}$
$E_3 \sqsupset E_1$ and $E_3 \sqsupset E_2$

E_3 is the only best preferred extension (under KTV criterion)

Constrained AF

- A *Constrained AF (CAF)* is an AF $\langle\mathcal{A},\mathcal{R}\rangle$ augmented with a set $\mathcal{C}$ of integrity constraints, that is a set of propositional formulae built from $\mathcal{L}_\mathcal{A}$
- $\mathcal{L}_\mathcal{A}$ denotes the propositional language defined from a set of arguments $\mathcal{A}$ and the connectives $\wedge$, $\vee$, $\Rightarrow$, $\neg$
- A *constraint* is a formula of one of the following forms: (*i*) $\varphi \Rightarrow v$, or (*ii*) $v \Rightarrow \varphi$, where φ is a propositional formula in $\mathcal{L}_\mathcal{A}$ and $v \in \{\mathbf{f},\mathbf{u},\mathbf{t}\}$
- Constraints are interpreted under Lukasiewicz's logic
- Extensions (of the AF) not satisfying the constraints are filtered out
- Given a CAF $\langle\mathcal{A},\mathcal{R},\mathcal{C}\rangle$, a set $S \subseteq \mathcal{A}$ is a σ-extension for $\langle\mathcal{A},\mathcal{R},\mathcal{C}\rangle$ if S is a σ-extension for $\langle\mathcal{A},\mathcal{R}\rangle$ and $S \models \mathcal{C}$

Example (CAF built from our example AF by adding $\mathtt{meat} \Rightarrow \mathbf{f}$(alse))

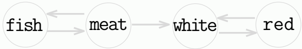

preferred extensions: $\{E_1 = \{\mathtt{fish},\mathtt{white}\}$, $E_2 = \{\mathtt{fish},\mathtt{red}\}$, $E_3 = \{\mathtt{meat},\mathtt{red}\}\}$

$E_3 \not\models \{\mathtt{meat} \Rightarrow \mathbf{f}\}$

E_1 and E_2 are the only best preferred extensions

Constrained AF

- A *Constrained AF (CAF)* is an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ augmented with a set $\mathcal{C}$ of integrity constraints, that is a set of propositional formulae built from $\mathcal{L}_{\mathcal{A}}$
- $\mathcal{L}_{\mathcal{A}}$ denotes the propositional language defined from a set of arguments $\mathcal{A}$ and the connectives $\wedge$, $\vee$, $\Rightarrow$, $\neg$
- A *constraint* is a formula of one of the following forms: (i) $\varphi \Rightarrow v$, or (ii) $v \Rightarrow \varphi$, where φ is a propositional formula in $\mathcal{L}_{\mathcal{A}}$ and $v \in \{\mathbf{f}, \mathbf{u}, \mathbf{t}\}$
- Constraints are interpreted under Lukasiewicz's logic
- Extensions (of the AF) not satisfying the constraints are filtered out
- Given a CAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$, a set $S \subseteq \mathcal{A}$ is a σ-extension for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ if S is a σ-extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ and $S \models \mathcal{C}$

Example (CAF built from our example AF by adding $\texttt{meat} \Rightarrow \mathbf{f}$(alse))

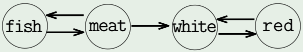

preferred extensions: $\{E_1 = \{\texttt{fish}, \texttt{white}\}$, $E_2 = \{\texttt{fish}, \texttt{red}\}$, $E_3 = \{\texttt{meat}, \texttt{red}\}\}$

$E_3 \not\models \{\texttt{meat} \Rightarrow \mathbf{f}\}$

E_1 and E_2 are the only best preferred extensions

Outline

Outline

ePAF Syntax

- Most of AF semantics are 3-valued (arguments can be either *accepted*, *defeated*, or *undecided*)
- We introduce a form of preferences which are 3-valued
- The extended Preference-based AF (ePAF) is an extension of AF (and PAF under KTV criterion) where preferences are 3-valued

(Extended preferences)

Let $\mathcal{A}$ be a set of arguments, an (extended) preference relation, denoted as $\succ$, is a strict partial order (i.e. an irreflexive, asymmetric, and transitive relation) over $\mathcal{A}^V = \{a^v \mid a \in \mathcal{A} \wedge v \in \{\mathbf{f}, \mathbf{u}, \mathbf{t}\}\}$ of the form $a^{v_1} \succ b^{v_2}$.

- An extended preference compares arguments' statuses in two extensions

Example

- $\texttt{red}^{\mathbf{t}} \succ \texttt{red}^{\mathbf{u}}$ means that we prefer menus containing red wine w.r.t. menus where red wine is undecided
- $\texttt{fish}^{\mathbf{t}} \succ \texttt{red}^{\mathbf{f}}$ states that we prefer menus containing $\texttt{fish}$ w.r.t. menus where $\texttt{red}$ is false (i.e. defeated)

ePAF Syntax

- Most of AF semantics are 3-valued (arguments can be either *accepted*, *defeated*, or *undecided*)
- We introduce a form of preferences which are 3-valued
- The extended Preference-based AF (ePAF) is an extension of AF (and PAF under KTV criterion) where preferences are 3-valued

(Extended preferences)

Let $\mathcal{A}$ be a set of arguments, an (extended) preference relation, denoted as $\succ$, is a strict partial order (i.e. an irreflexive, asymmetric, and transitive relation) over $\mathcal{A}^V = \{a^v \mid a \in \mathcal{A} \wedge v \in \{\mathbf{f}, \mathbf{u}, \mathbf{t}\}\}$ of the form $a^{v_1} \succ b^{v_2}$.

- An extended preference compares arguments' statuses in two extensions

Example

- $\mathtt{red}^{\mathbf{t}} \succ \mathtt{red}^{\mathbf{u}}$ means that we prefer menus containing red wine w.r.t. menus where red wine is undecided
- $\mathtt{fish}^{\mathbf{t}} \succ \mathtt{red}^{\mathbf{f}}$ states that we prefer menus containing `fish` w.r.t. menus where `red` is false (i.e. defeated)

Semantics of ePAF

- An extended PAF (ePAF) is an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ augmented with an extended preference relation $\succ$

(ePAF Semantics)

Given an ePAF $\Delta = \langle \mathcal{A}, \mathcal{R}, \succ \rangle$ and two distinct sets of arguments $E, F \subseteq \mathcal{A}$, we have that $E \sqsupseteq F$ if $\nexists\, a^{v_1} \succ b^{v_2}$ such that $a \in v_1(F) \setminus v_1(E)$, $b \in v_2(E) \setminus v_2(F)$ holds (where $v_1, v_2 \in \{\mathbf{f}, \mathbf{u}, \mathbf{t}\}$).
Moreover, $E \sqsupset F$, if $E \sqsupseteq F$ and $F \not\sqsupseteq E$.

- The best extensions are obtained as for PAF but using the above-defined criterion to compare extensions

Example (ePAF built from our AF by adding preference $\mathtt{fish}^{\mathbf{t}} \succ \mathtt{meat}^{\mathbf{t}}$)

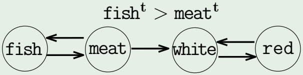

preferred extensions: $\{E_1 = \{\mathtt{fish}, \mathtt{white}\}$, $E_2 = \{\mathtt{fish}, \mathtt{red}\}$, $E_3 = \{\mathtt{meat}, \mathtt{red}\}\}$
$E_1 \sqsupset E_3$ and $E_2 \sqsupset E_3$

E_1 and E_2 are the best preferred extensions

Outline

Combining Preferences with Constraints

- We extend CAF with (extended) preferences to express several kinds of desiderata among extensions
- The resulting framework is called *extended Preference-based Constrained Argumentation Framework* (ePCAF)
- An ePCAF is a CAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ augmented with an (extended) preference relation $\succ$
- The semantics of an ePCAF is given by the best extensions selected among those that satisfy the constraints

(Semantics)

Given an ePCAF $\Delta = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \succ \rangle$, a σ-extension E for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is a best σ-extension for Δ if there is no σ-extension F for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ such that $F \sqsupset E$.

Combining Preferences with Constraints

- We extend CAF with (extended) preferences to express several kinds of desiderata among extensions
- The resulting framework is called *extended Preference-based Constrained Argumentation Framework* (ePCAF)
- An ePCAF is a CAF $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ augmented with an (extended) preference relation $\succ$
- The semantics of an ePCAF is given by the best extensions selected among those that satisfy the constraints

(Semantics)

Given an ePCAF $\Delta = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \succ \rangle$, a σ-extension E for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is a best σ-extension for Δ if there is no σ-extension F for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ such that $F \sqsupset E$.

Combining Preferences with Constraints

- The semantics of an ePCAF is given by the best extensions selected among those that satisfy the constraints

Example

Consider the ePCAF $\Delta = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \succ \rangle$ where:

- $\mathcal{C} = \{\mathtt{white} \Rightarrow \mathbf{f}\}$
- $\succ = \{\mathtt{fish}^{\mathbf{t}} > \mathtt{meat}^{\mathbf{t}}\}$ and
- the set of the preferred extension of the underlying AF is $\mathtt{pr}(\langle \mathcal{A}, \mathcal{R} \rangle) = \{E_1 = \{\mathtt{fish}, \mathtt{white}\}, E_2 = \{\mathtt{fish}, \mathtt{red}\}, E_3 = \{\mathtt{meat}, \mathtt{red}\}\}$
- $\mathtt{pr}(\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle) = \{E_2, E_3\}$

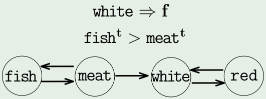

$\sqsupseteq$	E_2	E_3
E_2	✓	✓
E_3		✓

$\sqsupset$	E_2	E_3
E_2		✓
E_3		

As $\mathtt{white}$ must be false, there are only two preferred extensions satisfying the constraint: E_2 and E_3. Then, the only best preferred extension is E_2

Outline

Dealing with Multiple Agents

- Multiple agents sharing the same AF and having different constraints and preferences (represented by different ePCAFs)
- A multi-agent ePCAF (mPCAF) is a set of ePCAFs $\{\langle\mathcal{A},\mathcal{R},\mathcal{C}_1,\succ_1\rangle,\langle\mathcal{A},\mathcal{R},\mathcal{C}_2,\succ_2\rangle,\ldots,\langle\mathcal{A},\mathcal{R},\mathcal{C}_n,\succ_n\rangle\}$, one for each agent
- Each agent i has ePCAF $\Delta_i = \langle\mathcal{A},\mathcal{R},\mathcal{C}_i,\succ_i\rangle$ with its set of best σ-extensions in $\sigma(\Delta_i)$
- A set of arguments $S \subseteq \mathcal{A}$ is said to be a *possible* (resp. *necessary*) *best σ-extension* of Δ iff $S \in \sigma(\Delta_i)$ for some (resp. every) $i \in [1, n]$
- Two variants of the verification problem for mPCAF: the *possible* (resp. *necessary*) *verification* problem, is the problem of deciding whether S is possible (resp. necessary) best σ-extension of Δ
- Two variants of the credulous (resp. skeptical) acceptance problem: a goal argument can be possibly/necessarily credulously (resp. skeptically) accepted

Dealing with Multiple Agents

- Multiple agents sharing the same AF and having different constraints and preferences (represented by different ePCAFs)
- A multi-agent ePCAF (mPCAF) is a set of ePCAFs $\{\langle \mathcal{A}, \mathcal{R}, \mathcal{C}_1, \succ_1 \rangle, \langle \mathcal{A}, \mathcal{R}, \mathcal{C}_2, \succ_2 \rangle, ..., \langle \mathcal{A}, \mathcal{R}, \mathcal{C}_n, \succ_n \rangle\}$, one for each agent
- Each agent i has ePCAF $\Delta_i = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}_i, \succ_i \rangle$ with its set of best σ-extensions in $\sigma(\Delta_i)$
- A set of arguments $S \subseteq \mathcal{A}$ is said to be a *possible* (resp. *necessary*) *best* σ*-extension* of Δ iff $S \in \sigma(\Delta_i)$ for some (resp. every) $i \in [1, n]$
- Two variants of the verification problem for mPCAF: the *possible* (resp. *necessary*) *verification* problem, is the problem of deciding whether S is possible (resp. necessary) best σ-extension of Δ
- Two variants of the credulous (resp. skeptical) acceptance problem: a goal argument can be possibly/necessarily credulously (resp. skeptically) accepted

Outline

Outline

Complexity Results

- *Verification* problem (Ver_σ): deciding whether a set S of arguments is a σ-extension of ePAF/ePCAF/mPCAF
- *Credulous (resp. Skeptical) Acceptance* problem (CA_σ and SA_σ): deciding whether a goal argument g belongs to any (resp. all) σ-extension of ePAF/ePCAF/mPCAF

	AF			CAF			PAF			ePAF / ePCAF / mPCAF		
σ	Ver_σ	CA_σ	SA_σ	Ver_σ	CA_σ	SA_σ	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}
co	P	*NP*-c	P	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	P	co*NP*-c	Σ_2^P-c	Π_2^P-c
st	P	*NP*-c	co*NP*-c	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c
pr	co*NP*-c	*NP*-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P
ss	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P

- The complexity bounds for ePAF do not increase w.r.t. those of PAF, except for $SA_{\texttt{co}}$ that becomes Π_2^P-complete
- ePCAF is more expressive than CAF, particularly if we consider Ver_σ
- ePCAF is more expressive than both CAF and PAF, though the complexity bounds do not increase w.r.t. that of ePAF
- In the multiple agents scenario, no increase in complexity w.r.t. eP(C)AF

Complexity Results

- *Verification* problem (Ver_σ): deciding whether a set S of arguments is a σ-extension of ePAF/ePCAF/mPCAF
- *Credulous (resp. Skeptical) Acceptance* problem (CA_σ and SA_σ): deciding whether a goal argument g belongs to any (resp. all) σ-extension of ePAF/ePCAF/mPCAF

	AF			CAF			PAF			ePAF / ePCAF / mPCAF		
σ	Ver_σ	CA_σ	SA_σ	Ver_σ	CA_σ	SA_σ	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}	Ver_{σ_k}	CA_{σ_k}	SA_{σ_k}
co	P	*NP*-c	P	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	P	co*NP*-c	Σ_2^P-c	Π_2^P-c
st	P	*NP*-c	co*NP*-c	P	*NP*-c	co*NP*-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c
pr	co*NP*-c	*NP*-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P
ss	co*NP*-c	Σ_2^P-c	Π_2^P-c	co*NP*-c	Σ_2^P-c	Π_2^P-c	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P

- The complexity bounds for ePAF do not increase w.r.t. those of PAF, except for SA_{co} that becomes Π_2^P-complete
- ePCAF is more expressive than CAF, particularly if we consider Ver_σ
- ePCAF is more expressive than both CAF and PAF, though the complexity bounds do not increase w.r.t. that of ePAF
- In the multiple agents scenario, no increase in complexity w.r.t. eP(C)AF

Outline

Conclusions and future work

- We have introduced novel frameworks extending PAF and CAF with (3-valued) preferences and constraints
- Extended preferences and (3-valued) constraints as well as our complexity results can carry over to other AF-based frameworks, such as AFN and ASAF, that can be rewritten in AF
- FW1: investigate other criteria to define the best extensions, e.g. (variants of) democratic and elitist approaches
- FW2: investigate other forms of constraints such as weak and epistemic constraints
- FW3: investigate preferences and constraints in other frameworks extending AF (e.g. incomplete and probabilistic AFs)

Conclusions and future work

- We have introduced novel frameworks extending PAF and CAF with (3-valued) preferences and constraints
- Extended preferences and (3-valued) constraints as well as our complexity results can carry over to other AF-based frameworks, such as AFN and ASAF, that can be rewritten in AF
- FW1: investigate other criteria to define the best extensions, e.g. (variants of) democratic and elitist approaches
- FW2: investigate other forms of constraints such as weak and epistemic constraints
- FW3: investigate preferences and constraints in other frameworks extending AF (e.g. incomplete and probabilistic AFs)

Thank you for your attention!

... see you at the poster session!

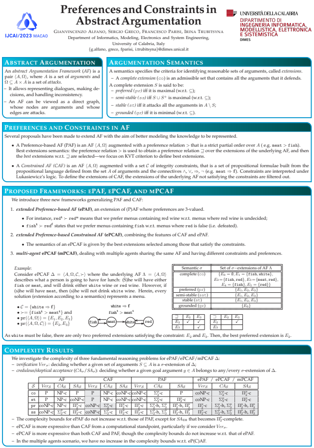

Preferences and Constraints in Abstract Argumentation

IJCAI/2023 MACAO
Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, Irina Trubitsyna
Department of Informatics, Modeling, Electronics and System Engineering,
University of Calabria, Italy
{g.alfano, greco, fparisi, i.trubitsyna}@dimes.unical.it
Università della Calabria
Dipartimento di Ingegneria Informatica, Modellistica, Elettronica e Sistemistica
DIMES

Abstract Argumentation

An *abstract Argumentation Framework* (*AF*) is a pair $\langle A, \Omega \rangle$, where A is a set of *arguments* and $\Omega \subseteq A \times A$ is a set of *attacks*.
- It allows representing dialogues, making decisions, and handling inconsistency.
- An AF can be viewed as a direct graph, whose nodes are arguments and whose edges are attacks.

Argumentation Semantics

A semantics specifies the criteria for identifying reasonable sets of arguments, called *extensions*.
- A *complete extension* ($\mathtt{co}$) is an admissible set that contains all the arguments that it defends.

A complete extension S is said to be:
- *preferred* ($\mathtt{pr}$) iff it is maximal (w.r.t. $\subseteq$);
- *semi-stable* ($\mathtt{ss}$) iff $S \cup S^+$ is maximal (w.r.t. $\subseteq$);
- *stable* ($\mathtt{st}$) iff it attacks all the arguments in $A \setminus S$;
- *grounded* ($\mathtt{gr}$) iff it is minimal (w.r.t. $\subseteq$).

Preferences and Constraints in AF

Several proposals have been made to extend AF with the aim of better modeling the knowledge to be represented.

- A Preference-based AF (PAF) is an AF $\langle A, \Omega \rangle$ augmented with a preference relation $>$ that is a strict partial order over A (e.g. $\mathtt{meat} > \mathtt{fish}$). Best extensions semantics: the preference relation $>$ is used to obtain a preference relation $\sqsupseteq$ over the extensions of the underlying AF, and then the *best* extensions w.r.t. $\sqsupseteq$ are selected—we focus on KVT criterion to define best extensions.
- A *Constrained AF (CAF)* is an AF $\langle A, \Omega \rangle$ augmented with a set $\mathcal{C}$ of integrity constraints, that is a set of propositional formulae built from the propositional language defined from the set A of arguments and the connectives $\wedge, \vee, \Rightarrow, \neg$ (e.g. $\mathtt{meat} \Rightarrow \mathtt{f}$). Constraints are interpreted under Lukasiewicz's logic. To define the extensions of CAF, the extensions of the underlying AF not satisfying the constraints are filtered out.

Proposed Frameworks: ePAF, ePCAF, and mPCAF

We introduce three new frameworks generalizing PAF and CAF:

1. ***extended Preference-based AF* (ePAF)**, an extension of (P)AF where preferences are 3-valued.
   - For instance, $\mathtt{red}^{\mathbf{t}} \succ \mathtt{red}^{\mathbf{u}}$ means that we prefer menus containing red wine w.r.t. menus where red wine is undecided;
   - $\mathtt{fish}^{\mathbf{t}} \succ \mathtt{red}^{\mathbf{f}}$ states that we prefer menus containing $\mathtt{fish}$ w.r.t. menus where $\mathtt{red}$ is false (i.e. defeated).
2. ***extended Preference-based Constrained AF* (ePCAF)**, combining the features of CAF and ePAF.
   - The semantics of an ePCAF is given by the best extensions selected among those that satisfy the constraints.
3. ***multi-agent* ePCAF (mPCAF)**, dealing with multiple agents sharing the same AF and having different constraints and preferences.

Example:
Consider ePCAF $\Delta = \langle A, \Omega, \mathcal{C}, \succ \rangle$ where the underlying AF $\Lambda = \langle A, \Omega \rangle$ describes what a person is going to have for lunch: (S)he will have either $\mathtt{fish}$ or $\mathtt{meat}$, and will drink either $\mathtt{white}$ wine or $\mathtt{red}$ wine. However, if (s)he will have $\mathtt{meat}$, then (s)he will not drink $\mathtt{white}$ wine. Herein, every solution (extension according to a semantics) represents a menu.

- $\mathcal{C} = \{\mathtt{white} \Rightarrow \mathtt{f}\}$
- $\succ = \{\mathtt{fish}^{\mathbf{t}} > \mathtt{meat}^{\mathbf{t}}\}$ and
- $\mathtt{pr}(\langle A, \Omega \rangle) = \{E_1, E_2, E_3\}$
- $\mathtt{pr}(\langle A, \Omega, \mathcal{C} \rangle) = \{E_2, E_3\}$

$\mathtt{white} \Rightarrow \mathtt{f}$
$\mathtt{fish}^{\mathbf{t}} > \mathtt{meat}^{\mathbf{t}}$
fish meat white red

Semantic σ	Set of σ-extensions of AF Λ
complete ($\mathtt{co}$)	$\{E_0 = \emptyset, E_1 = \{\mathtt{fish}, \mathtt{white}\}, E_2 = \{\mathtt{fish}, \mathtt{red}\}, E_3 = \{\mathtt{meat}, \mathtt{red}\}, E_4 = \{\mathtt{fish}\}, E_5 = \{\mathtt{red}\}\}$
preferred ($\mathtt{pr}$)	$\{E_1, E_2, E_3\}$
semi-stable ($\mathtt{sst}$)	$\{E_1, E_2, E_3\}$
stable ($\mathtt{st}$)	$\{E_1, E_2, E_3\}$
grounded ($\mathtt{gr}$)	$\{E_0\}$

$\sqsupseteq$	E_2	E_3
E_2	✓	✓
E_3		✓

$\sqsupset$	E_2	E_3
E_2		✓
E_3		

As $\mathtt{white}$ must be false, there are only two preferred extensions satisfying the constraint: E_2 and E_3. Then, the best preferred extension is E_2.

Complexity Results

We investigate the complexity of three fundamental reasoning problems for ePAF/ePCAF/mPCAF Δ:
- *verification* Ver_σ: deciding whether a given set of arguments $S \subseteq A$ is a σ-extension of Δ;
- *credulous/skeptical acceptance* (CA_σ / SA_σ): deciding whether a given goal argument $g \in A$ belongs to any/every σ-extension of Δ.

	AF			CAF			PAF			ePAF / ePCAF / mPCAF		
$\mathcal{S}$	$Ver_\mathcal{S}$	$CA_\mathcal{S}$	$SA_\mathcal{S}$	$Ver_\mathcal{S}$	$CA_\mathcal{S}$	$SA_\mathcal{S}$	$Ver_\mathcal{S}$	$CA_\mathcal{S}$	$SA_\mathcal{S}$	$Ver_\mathcal{S}$	$CA_\mathcal{S}$	$SA_\mathcal{S}$
co	P	NP-c	P	P	NP-c	coNP-c	coNP-c	Σ_2^p-c	P	coNP-c	Σ_2^p-c	Π_2^p-c
st	P	NP-c	coNP-c	P	NP-c	coNP-c	coNP-c	Σ_2^p-c	Π_2^p-c	coNP-c	Σ_2^p-c	Π_2^p-c
pr	coNP-c	NP-c	Π_2^p-c	coNP-c	Σ_2^p-c	Π_2^p-c	Π_2^p-c	Σ_2^p-h, Σ_3^p	Π_2^p-h, Π_3^p	Π_2^p-c	Σ_2^p-h, Σ_3^p	Π_2^p-h, Π_3^p
ss	coNP-c	Σ_2^p-c	Π_2^p-c	coNP-c	Σ_2^p-c	Π_2^p-c	Π_2^p-c	Σ_2^p-h, Σ_3^p	Π_2^p-h, Π_3^p	Π_2^p-c	Σ_2^p-h, Σ_3^p	Π_2^p-h, Π_3^p

- The complexity bounds for ePAF do not increase w.r.t. those of PAF, except for $SA_{\mathtt{co}}$ that becomes Π_2^p-complete.
- ePCAF is more expressive than CAF from a computational standpoint, particularly if we consider Ver_σ.
- ePCAF is more expressive than both CAF and PAF, though the complexity bounds do not increase w.r.t. that of ePAF.
- In the multiple agents scenario, we have no increase in the complexity bounds w.r.t. eP(C)AF.