Novel Algorithms for Abstract Dialectical Frameworks based on Complexity Analysis of Subclasses and SAT Solving

Thomas Linsbichler[1] Marco Maratea[2] Andreas Niskanen[3]
Johannes P. Wallner[1] Stefan Woltran[1]

[1] Institute of Logic and Computation, TU Wien, Austria
[2] DIBRIS, University of Genova, Italy
[3] HIIT, Department of Computer Science, University of Helsinki, Finland

July 18, 2018 @ IJCAI-ECAI 2018, Stockholm, Sweden

Motivation

Argumentation in Artificial Intelligence (AI)

- An active area of modern AI research
- Applications in law, medicine, eGovernment, debating technologies
- Central formalism: Dung's argumentation frameworks (AFs)
  - Arguments as nodes and attacks as edges in a directed graph
  - Complexity-sensitive procedures for reasoning in AFs implemented

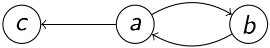

Motivation

Argumentation in Artificial Intelligence (AI)

- An active area of modern AI research
- Applications in law, medicine, eGovernment, debating technologies
- Central formalism: Dung's argumentation frameworks (AFs)
  - Arguments as nodes and attacks as edges in a directed graph
  - Complexity-sensitive procedures for reasoning in AFs implemented

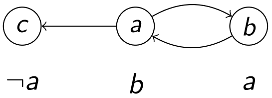

Abstract Dialectical Frameworks (ADFs)

- Powerful generalization of AFs: each argument equipped with an acceptance condition (a propositional formula)
- Expressive power comes with a price: higher computational complexity

Contributions

- Complexity analysis of ADF subclasses
  - Investigate two new subclasses: acyclic and concise ADFs
  - Constant distance to a subclass: k-bipolar, k-acyclic and k-concise
- Algorithms for argument acceptance problems in ADFs
  - Make use of input ADF being k-bipolar for a sufficiently low value of k
  - Based on incremental SAT solving
- Experimental evaluation of the resulting system
  - Capable of outperforming the state-of-the-art

Syntax of Abstract Dialectical Frameworks

Abstract Dialectical Framework (ADF)

A tuple $D = (A, L, C)$, where

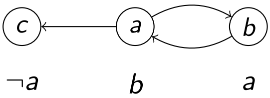

- A is a finite set of **arguments**
- $L \subseteq A \times A$ is a set of **links**
- $C = \{\varphi_a\}_{a \in A}$ is a set of **acceptance conditions**
  - each φ_a is a propositional formula over the parents of a

Interpretations

- An interpretation I maps each argument to a truth value in $\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$
- J is at least as informative as I, $I \leq_i J$, if all arguments that I maps to $\mathbf{t}$ or $\mathbf{f}$ are mapped likewise by J

Syntax of Abstract Dialectical Frameworks

Abstract Dialectical Framework (ADF)

A tuple $D = (A, L, C)$, where

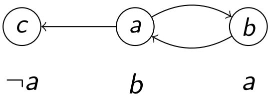

- A is a finite set of **arguments**
- $L \subseteq A \times A$ is a set of **links**
- $C = \{\varphi_a\}_{a \in A}$ is a set of **acceptance conditions**
  - each φ_a is a propositional formula over the parents of a

Interpretations

- An interpretation I maps each argument to a truth value in $\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$
- J is at least as informative as I, $I \leq_i J$, if all arguments that I maps to **t** or **f** are mapped likewise by J

Semantics of Abstract Dialectical Frameworks

- Semantics σ identify interpretations that are meaningful in the context of argument acceptance
  - Map an ADF D to a set $\sigma(D)$ of σ-interpretations
- Standard AF semantics can be generalized to ADFs

Preferred semantics

Given an ADF D, an interpretation I is preferred, $I \in prf(D)$, if I is admissible and $\leq_i$-maximal.

Semantics of Abstract Dialectical Frameworks

- Semantics σ identify interpretations that are meaningful in the context of argument acceptance
  - Map an ADF D to a set $\sigma(D)$ of σ-interpretations
- Standard AF semantics can be generalized to ADFs

Preferred semantics

Given an ADF D, an interpretation I is preferred, $I \in prf(D)$, if I is admissible and $\leq_i$-maximal.

ADF Reasoning Tasks

Let σ be an ADF semantics.

	Input	Decision
$Cred_\sigma$	ADF D, argument $a \in A$	$\exists I \in \sigma(D), I(a) = \mathbf{t}$?
$Skept_\sigma$	ADF D, argument $a \in A$	$\forall I \in \sigma(D), I(a) = \mathbf{t}$?
$Exists_\sigma^{>}$	ADF D, interpretation I	$\exists J \in \sigma(D), J >_i I$?
Ver_σ	ADF D, interpretation I	$I \in \sigma(D)$?

Example

Now $\{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}\}$ and $\{a \mapsto \mathbf{f}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}\}$ are preferred in D, so a is credulously but not skeptically accepted under preferred.

ADF Reasoning Tasks

Let σ be an ADF semantics.

	Input	Decision
$Cred_\sigma$	ADF D, argument $a \in A$	$\exists I \in \sigma(D), I(a) = \mathbf{t}$?
$Skept_\sigma$	ADF D, argument $a \in A$	$\forall I \in \sigma(D), I(a) = \mathbf{t}$?
$Exists_\sigma^{>}$	ADF D, interpretation I	$\exists J \in \sigma(D), J >_i I$?
Ver_σ	ADF D, interpretation I	$I \in \sigma(D)$?

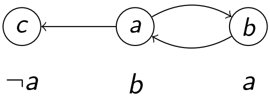

Example

Now $\{a \mapsto \mathbf{t}, b \mapsto \mathbf{t}, c \mapsto \mathbf{f}\}$ and $\{a \mapsto \mathbf{f}, b \mapsto \mathbf{f}, c \mapsto \mathbf{t}\}$ are preferred in D, so a is credulously but not skeptically accepted under preferred.

ADF Subclasses

Subclasses

An ADF $D = (A, L, C)$ is

- bipolar, if every link $(a, b) \in L$ is attacking or supporting,
- acyclic, if the directed graph (A, L) is acyclic,
- concise for a semantics σ, if there is exactly one σ-interpretation.

Distance to Subclasses

Let $k \geq 1$. An ADF $D = (A, L, C)$ is

- k-bipolar, if every argument has at most k non-bipolar incoming links,
- k-acyclic, if removing links from parents of k arguments results in an acyclic ADF,
- k-concise for a semantics σ, if there are at most k σ-interpretations.

ADF Subclasses

Subclasses

An ADF $D = (A, L, C)$ is

- bipolar, if every link $(a, b) \in L$ is attacking or supporting,
- acyclic, if the directed graph (A, L) is acyclic,
- concise for a semantics σ, if there is exactly one σ-interpretation.

Distance to Subclasses

Let $k \geq 1$. An ADF $D = (A, L, C)$ is

- k-bipolar, if every argument has at most k non-bipolar incoming links,
- k-acyclic, if removing links from parents of k arguments results in an acyclic ADF,
- k-concise for a semantics σ, if there are at most k σ-interpretations.

Complexity of ADFs and ADF Subclasses

σ	$Cred_\sigma$	$Skept_\sigma$	$Exists_\sigma$	Ver_σ
cf	NP-c	trivial	NP-c	NP-c
nai	NP-c	Π_2^P-c	NP-c	DP-c
adm	Σ_2^P-c	trivial	Σ_2^P-c	coNP-c
grd	coNP-c	coNP-c	coNP-c	DP-c
com	Σ_2^P-c	coNP-c	Σ_2^P-c	DP-c
prf	Σ_2^P-c	Π_3^P-c	Σ_2^P-c	Π_2^P-c

Table: Complexity of general ADFs [Strass and Wallner, 2015].

Complexity of ADFs and ADF Subclasses

σ	$Cred_\sigma$	$Skept_\sigma$	$Exists_\sigma$	Ver_σ
cf	in P	trivial	in P	in P
nai	in P	coNP-c	in P	in P
adm	NP-c	trivial	NP-c	in P
grd	in P	in P	in P	in P
com	NP-c	in P	NP-c	in P
prf	NP-c	Π_2^P-c	NP-c	coNP-c

Table: Complexity of bipolar ADFs [Strass and Wallner, 2015].

Complexity of ADFs and ADF Subclasses

σ	$Cred_\sigma$	$Skept_\sigma$	$Exists_\sigma$	Ver_σ
cf	in P	trivial	in P	in P
nai	in P	coNP-c	in P	in P
adm	NP-c	trivial	NP-c	in P
grd	in P	in P	in P	in P
com	NP-c	in P	NP-c	in P
prf	NP-c	Π_2^P-c	NP-c	coNP-c

Table: Complexity of k-bipolar ADFs (this paper).

Complexity of ADFs and ADF Subclasses

σ	$Cred_\sigma$	$Skept_\sigma$	$Exists_\sigma$	Ver_σ
cf	in P	trivial	in P	in P
nai	in P	coNP-c	in P	in P
adm	NP-c	trivial	NP-c	in P
grd	in P	in P	in P	in P
com	NP-c	in P	NP-c	in P
prf	NP-c	Π_2^P-c	NP-c	coNP-c

Table: Complexity of *k*-bipolar ADFs (this paper).

Complexity results for other subclasses, e.g.:

- acyclic ADFs: most decision problems tractable
- *k*-acyclic ADFs: no observed drops in complexity

Results on concise and *k*-concise and more details in paper!

Algorithms for Acceptance in ADFs

Skeptical acceptance under preferred via SAT solving

- Π_3^P-complete in general, and Π_2^P-complete for *k*-bipolar ADFs
- Goal: delegate suitable NP fragments to SAT solvers
- Complexity of $Exists^{>}_{adm}$ is NP-complete for *k*-bipolar ADFs
- Provide encoding of $Exists^{>}_{adm}$ as an instance of SAT
  - bipolar ADFs: polynomial encoding
  - *k*-bipolar ADFs: polynomial encoding, but exponential in *k*
- Complexity-sensitive: detect when input ADF is *k*-bipolar for low *k*

Skeptical Acceptance under Preferred for *k*-bipolar ADFs

Given an ADF D and an argument α.

- Form the encoding φ for $Exists^{>}_{adm}(D, I_{\mathbf{u}})$.
- If φ is unsatisfiable, reject.
- While there exists a truth assignment to φ:
  - Extract the corresponding admissible interpretation I.
  - Iteratively search for a preferred interpretation:
    - Similarly solve the problem $Exists^{>}_{adm}(D, I)$ via SAT.
    - If a solution exists, set I as the corresponding interpretation.
  - If $I(\alpha) \neq \mathbf{t}$, reject.
  - Otherwise, exclude all $J \leq_i I$ from the search space by refining φ.
- Accept.

Implementation and Empirical Evaluation

k++ADF: SAT-based system for reasoning in ADFs

- Implements the encodings and algorithms
- Includes MiniSAT 2.2.0 as the underlying SAT solver

Experimental setup

- Benchmark ADFs generated from ICCMA 2017 AFs
- 1800 second timeout for each instance
- Compare to existing systems for ADFs: QADF, YADF, goDiamond

Implementation and Empirical Evaluation

k++ADF: SAT-based system for reasoning in ADFs

- Implements the encodings and algorithms
- Includes MiniSAT 2.2.0 as the underlying SAT solver

Experimental setup

- Benchmark ADFs generated from ICCMA 2017 AFs
- 1800 second timeout for each instance
- Compare to existing systems for ADFs: QADF, YADF, goDiamond

Skeptical acceptance under preferred

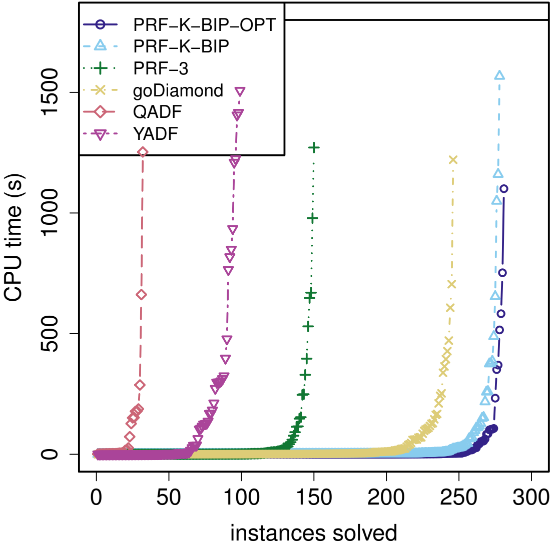

Paper Summary

Contributions

- Complexity analysis of ADF subclasses
- Algorithms for credulous and skeptical acceptance under preferred semantics based on incremental SAT solving
- Empirical evaluation of the system k++ADF, available in open source:

  `http://www.cs.helsinki.fi/group/coreo/k++adf/`
- More in paper: complexity results for further subclasses, details on encodings and algorithms, additional experiments, ...
- Future work: sharper complexity bounds, extending the system