Foundations of Knowledge Representation

Lecture 10: Abstract Argumentation

Sebastian Rudolph

based on slides of
Sarah Gaggl and
Stefan Woltran

Introduction

Argumentation:

The study of processes "concerned with how assertions are proposed, discussed, and resolved in the context of issues upon which several diverging opinions may be held".

[Bench-Capon and Dunne, Argumentation in AI, AIJ 171:619-641, 2007]

Formal Models of Argumentation are concerned with

- representation of an argument (i.e. an expression of opinion)
- representation of the relationship between arguments
- solving conflicts between the arguments ("acceptability")

Overall Process

The overall process of managing argumentation frameworks consists of the following steps:

1. Starting point: knowledge-base
2. Form arguments
3. Identify conflicts
4. Abstract from internal structure
5. Resolve conflicts
6. Draw conclusions

Overall Process – Form Arguments

Consider the following knowledge base:

Example

$$\Delta = \{s, r, w, s \rightarrow \neg r, r \rightarrow \neg w, w \rightarrow \neg s\}$$

From this, form arguments:

$$\langle\{w, w\rightarrow\neg s\}, \neg s\rangle$$

$$\langle\{s, s\rightarrow\neg r\}, \neg r\rangle \qquad \langle\{r, r\rightarrow\neg w\}, \neg w\rangle$$

Overall Process – Identify Conflicts

Example

$$\Delta = \{s, r, w, s \to \neg r, r \to \neg w, w \to \neg s\}$$

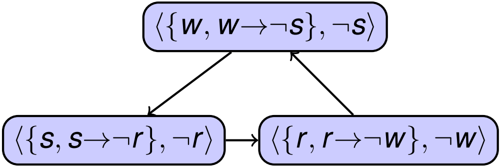

Overall Process – Abstract from Internal Structure

Example (Knowledge Base)

$$\Delta = \{s, r, w, s \rightarrow \neg r, r \rightarrow \neg w, w \rightarrow \neg s\}$$

F_Δ:

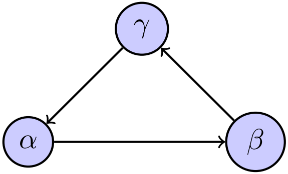

Overall Process – Resolve Conflicts

Example (Knowledge Base)

$$\Delta = \{s, r, w, s \to \neg r, r \to \neg w, w \to \neg s\}$$

F_Δ:

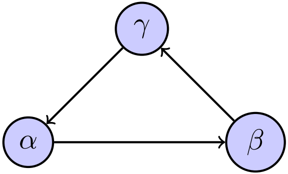

$pref(F_\Delta) = \{\emptyset\}$
$stage(F_\Delta) = \{\{\alpha\}, \{\beta\}, \{\gamma\}\}$

Overall Process – Draw Conclusions

Example (Knowledge Base)

$$\Delta = \{s, r, w, s \to \neg r, r \to \neg w, w \to \neg s\}$$

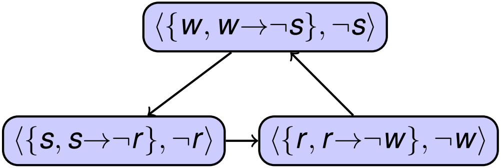

$Cn_{pref}(F_\Delta) = Cn(\top)$
$Cn_{stage}(F_\Delta) = Cn(\neg r \vee \neg w \vee \neg s)$

The Overall Process (ctd.)

Some Remarks

- Main idea dates back to Dung [1995]; has then been refined by several authors (Prakken, Gordon, Caminada, etc.)
- Separation between logical (forming arguments) and nonmonotonic reasoning ("abstract argumentation frameworks")
- Abstraction allows to compare several KR formalisms on a conceptual level ("calculus of conflict")

The Overall Process (ctd.)

Main Challenge

- All steps in the argumentation process are, in general, intractable.
- This calls for:
  - careful complexity analysis (identification of tractable fragments)
  - re-use of established tools for implementations (reduction method)

Approaches to Form Arguments

Classical Arguments [Besnard & Hunter, 2001]

- Given is a KB (a set of propositions) Δ
- an argument is a pair (Φ, α), such that $\Phi \subseteq \Delta$ is consistent, $\Phi \models \alpha$ and for no $\Psi \subset \Phi$, $\Psi \models \alpha$
- conflicts between arguments (Φ, α) and (Φ', α') arise if Φ and α' are contradicting.

Example

$$\langle\{s, s \rightarrow \neg r\}, \neg r\rangle \longrightarrow \langle\{r, r \rightarrow \neg w\}, \neg w\rangle$$

Approaches to Form Arguments

Other Approaches

- arguments are trees of statements
- claims are obtained via strict and defeasible rules
- different notions of conflict: rebuttal, undercut, etc.

Dung's Abstract Argumentation Frameworks

Example

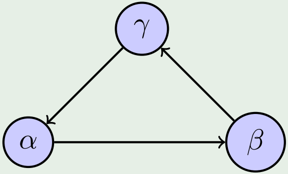

Main Properties

- Abstract from the concrete content of arguments but only consider the relation between them
- Semantics select subsets of arguments respecting certain criteria
- Simple, yet powerful, formalism
- Most active research area in the field of argumentation.

Dung's Abstract Argumentation Frameworks

Definition

An argumentation framework (AF) is a pair (A, R) where

- A is a set of arguments,
- $R \subseteq A \times A$ is a relation representing the conflicts ("attacks").

Example

F=({a,b,c,d,e} , {(a,b),(c,b),(c,d),(d,c),(d,e),(e,e)})

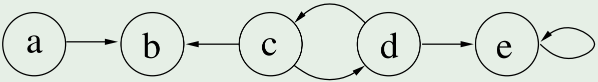

Basic Properties

Conflict-Free Sets

Given an AF $F = (A, R)$.
A set $S \subseteq A$ is conflict-free in F, if, for each $a, b \in S$, $(a, b) \notin R$.

Basic Properties

Conflict-Free Sets

Given an AF $F = (A, R)$.
A set $S \subseteq A$ is conflict-free in F, if, for each $a, b \in S$, $(a, b) \notin R$.

Example

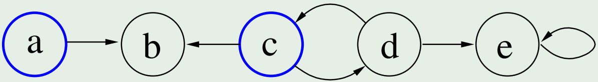

$cf(F) = \{\{a, c\},$

Basic Properties

Conflict-Free Sets

Given an AF $F = (A, R)$.
A set $S \subseteq A$ is conflict-free in F, if, for each $a, b \in S$, $(a, b) \notin R$.

Example

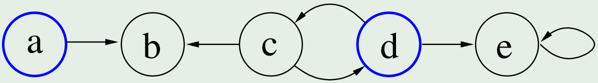

$cf(F) = \{\{a, c\}, \{a, d\},$

Basic Properties

Conflict-Free Sets

Given an AF $F = (A, R)$.
A set $S \subseteq A$ is conflict-free in F, if, for each $a, b \in S$, $(a, b) \notin R$.

Example

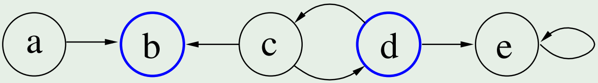

$cf(F) = \{\{a, c\}, \{a, d\}, \{b, d\},$

Basic Properties

Conflict-Free Sets

Given an AF $F = (A, R)$.
A set $S \subseteq A$ is conflict-free in F, if, for each $a, b \in S$, $(a, b) \notin R$.

Example

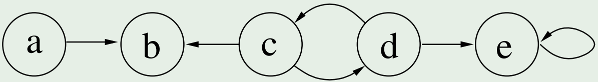

$cf(F) = \{\{a, c\}, \{a, d\}, \{b, d\}, \{a\}, \{b\}, \{c\}, \{d\}, \emptyset\}$

Basic Properties (ctd.)

Admissible Sets [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is admissible in F, if

- S is conflict-free in F
- each $a \in S$ is defended by S in F
  - $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Basic Properties (ctd.)

Admissible Sets [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is admissible in F, if

- S is conflict-free in F
- each $a \in S$ is defended by S in F
  - $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

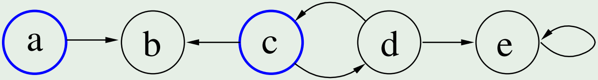

$adm(F) = \{\{a, c\},$

Basic Properties (ctd.)

Admissible Sets [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is admissible in F, if

- S is conflict-free in F
- each $a \in S$ is defended by S in F
  - $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

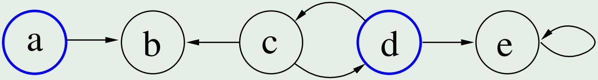

$adm(F) = \{\{a, c\}, \{a, d\},$

Basic Properties (ctd.)

Admissible Sets [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is admissible in F, if

- S is conflict-free in F
- each $a \in S$ is defended by S in F
  - $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

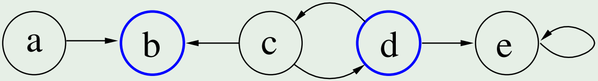

$adm(F) = \{\{a, c\}, \{a, d\},$ ~~$\{b, d\}$~~,

Basic Properties (ctd.)

Admissible Sets [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is admissible in F, if

- S is conflict-free in F
- each $a \in S$ is defended by S in F
  - $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

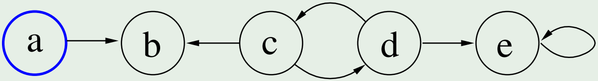

$adm(F) = \{\{a, c\}, \{a, d\}, \sout{\{b, d\}}, \{a\}, \sout{\{b\}}, \{c\}, \{d\}, \emptyset\}$

Basic Properties (ctd.)

Dung's Fundamental Lemma

Let S be admissible in an AF F and a, a' arguments in F defended by S in F. Then,

1. $S' = S \cup \{a\}$ is admissible in F
2. a' is defended by S' in F

Semantics

Naive Extensions

Given an AF $F = (A, R)$. A set $S \subseteq A$ is naive in F, if

- S is conflict-free in F
- for each $T \subseteq A$ conflict-free in F, $S \not\subset T$

Semantics

Naive Extensions

Given an AF $F = (A, R)$. A set $S \subseteq A$ is naive in F, if

- S is conflict-free in F
- for each $T \subseteq A$ conflict-free in F, $S \not\subset T$

Example

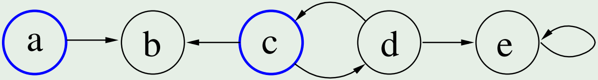

$naive(F) = \{\{a, c\},$

Semantics

Naive Extensions

Given an AF $F = (A, R)$. A set $S \subseteq A$ is naive in F, if

- S is conflict-free in F
- for each $T \subseteq A$ conflict-free in F, $S \not\subset T$

Example

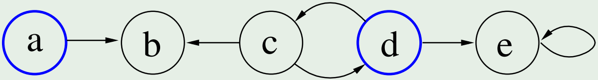

$naive(F) = \{\{a, c\}, \{a, d\},$

Semantics

Naive Extensions

Given an AF $F = (A, R)$. A set $S \subseteq A$ is naive in F, if

- S is conflict-free in F
- for each $T \subseteq A$ conflict-free in F, $S \not\subset T$

Example

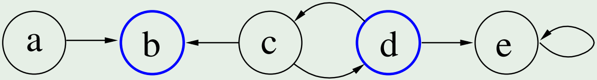

$naive(F) = \{\{a, c\}, \{a, d\}, \{b, d\},$

Semantics

Naive Extensions

Given an AF $F = (A, R)$. A set $S \subseteq A$ is naive in F, if

- S is conflict-free in F
- for each $T \subseteq A$ conflict-free in F, $S \not\subset T$

Example

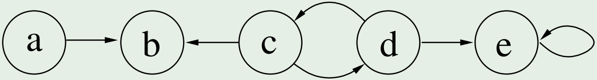

$naive(F) = \{\{a, c\}, \{a, d\}, \{b, d\},$ ~~$\{a\}$~~, ~~$\{b\}$~~, ~~$\{c\}$~~, ~~$\{d\}$~~, ~~$\emptyset$~~$\}$

Grounded Extension [Dung, 1995]

Given an AF $F = (A, R)$. The unique grounded extension of F is defined as the outcome S (initially empty) of the following "algorithm":

1. put each argument $a \in A$ which is not attacked in F into S; if no such argument exists, return S;
2. remove from F all (new) arguments in S and all arguments attacked by them (together with all adjacent attacks); and continue with Step 1.

Grounded Extension [Dung, 1995]

Given an AF $F = (A, R)$. The unique grounded extension of F is defined as the outcome S (initially empty) of the following "algorithm":

1. put each argument $a \in A$ which is not attacked in F into S; if no such argument exists, return S;
2. remove from F all (new) arguments in S and all arguments attacked by them (together with all adjacent attacks); and continue with Step 1.

Example

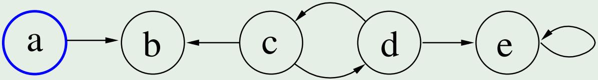

$ground(F) = \{\{a\}\}$

Semantics (ctd.)

Complete Extension [Dung, 1995]

Given an AF (A, R). A set $S \subseteq A$ is complete in F, if

- S is admissible in F
- each $a \in A$ defended by S in F is contained in S
  - Recall: $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Semantics (ctd.)

Complete Extension [Dung, 1995]

Given an AF (A, R). A set $S \subseteq A$ is complete in F, if

- S is admissible in F
- each $a \in A$ defended by S in F is contained in S
  - Recall: $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

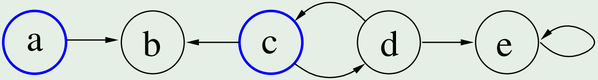

$comp(F) = \{\{a, c\},$

Semantics (ctd.)

Complete Extension [Dung, 1995]

Given an AF (A, R). A set $S \subseteq A$ is complete in F, if

- S is admissible in F
- each $a \in A$ defended by S in F is contained in S
  - Recall: $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

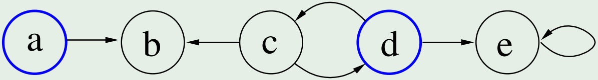

$comp(F) = \{\{a, c\}, \{a, d\},$

Semantics (ctd.)

Complete Extension [Dung, 1995]

Given an AF (A, R). A set $S \subseteq A$ is complete in F, if

- S is admissible in F
- each $a \in A$ defended by S in F is contained in S
  - Recall: $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

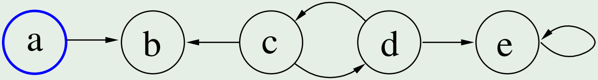

$comp(F) = \{\{a, c\}, \{a, d\}, \{a\},$

Semantics (ctd.)

Complete Extension [Dung, 1995]

Given an AF (A, R). A set $S \subseteq A$ is complete in F, if

- S is admissible in F
- each $a \in A$ defended by S in F is contained in S
  - Recall: $a \in A$ is defended by S in F, if for each $b \in A$ with $(b, a) \in R$, there exists a $c \in S$, such that $(c, b) \in R$.

Example

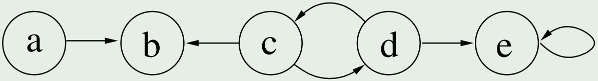

$comp(F) = \{\{a, c\}, \{a, d\}, \{a\},$ ~~$\{c\}, \{d\}, \emptyset$~~ $\}$

Semantics (ctd.)

Properties of the Grounded Extension

For any AF F, the grounded extension of F is the subset-minimal complete extension of F.

Semantics (ctd.)

Properties of the Grounded Extension

For any AF F, the grounded extension of F is the subset-minimal complete extension of F.

Remark

Since there exists exactly one grounded extension for each AF F, we often write $\mathit{ground}(F) = S$ instead of $\mathit{ground}(F) = \{S\}$.

Semantics (ctd.)

Preferred Extensions [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a preferred extension of F, if

- S is admissible in F
- for each $T \subseteq A$ admissible in F, $S \not\subset T$

Semantics (ctd.)

Preferred Extensions [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a preferred extension of F, if

- S is admissible in F
- for each $T \subseteq A$ admissible in F, $S \not\subset T$

Example

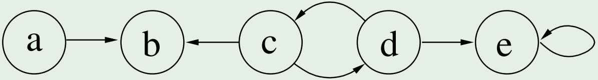

$pref(F) = \{\{a, c\}, \{a, d\},$ ~~$\{a\}, \{c\}, \{d\}, \emptyset$~~ $\}$

Semantics (ctd.)

Stable Extensions [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a stable extension of F, if

- S is conflict-free in F
- for each $a \in A \setminus S$, there exists a $b \in S$, such that $(b, a) \in R$

Semantics (ctd.)

Stable Extensions [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a stable extension of F, if

- S is conflict-free in F
- for each $a \in A \setminus S$, there exists a $b \in S$, such that $(b, a) \in R$

Example

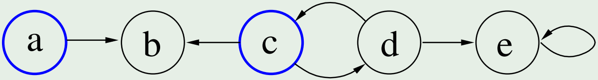

$\mathit{stable}(F) = \{$ ~~$\{a, c\}$~~ $\}$

Semantics (ctd.)

Stable Extensions [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a stable extension of F, if

- S is conflict-free in F
- for each $a \in A \setminus S$, there exists a $b \in S$, such that $(b, a) \in R$

Example

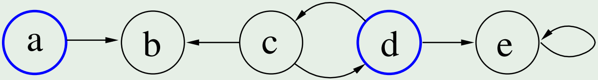

$stable(F) = \{$~~$\{a, c\}$~~$, \{a, d\},$

Semantics (ctd.)

Stable Extensions [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a stable extension of F, if

- S is conflict-free in F
- for each $a \in A \setminus S$, there exists a $b \in S$, such that $(b, a) \in R$

Example

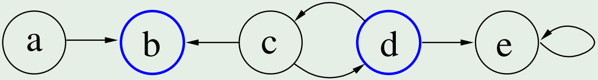

$stable(F) = \{$ ~~$\{a, c\}$~~, $\{a, d\}$, ~~$\{b, d\}$~~,

Semantics (ctd.)

Stable Extensions [Dung, 1995]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a stable extension of F, if

- S is conflict-free in F
- for each $a \in A \setminus S$, there exists a $b \in S$, such that $(b, a) \in R$

Example

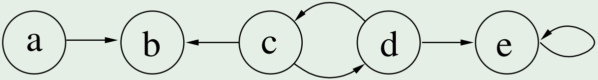

stable$(F) = \{$ ~~$\{a, c\}$~~, $\{a, d\}$, ~~$\{b, d\}$~~, ~~$\{a\}$~~, ~~$\{b\}$~~, ~~$\{c\}$~~, ~~$\{d\}$~~, ~~$\emptyset$~~, $\}$

Semantics (ctd.)

Some Relations

For any AF F the following relations hold:

1. Each stable extension of F is admissible in F
2. Each stable extension of F is also a preferred one
3. Each preferred extension of F is also a complete one

Semantics (ctd.)

Semi-Stable Extensions [Caminada, 2006]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a semi-stable extension of F, if

- S is admissible in F
- for each $T \subseteq A$ admissible in F, $S^+ \not\subset T^+$
    - for $S \subseteq A$, define $S^+ = S \cup \{a \mid \exists b \in S \text{ with } (b, a) \in R\}$

Semantics (ctd.)

Semi-Stable Extensions [Caminada, 2006]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a semi-stable extension of F, if

- S is admissible in F
- for each $T \subseteq A$ admissible in F, $S^+ \not\subset T^+$
  - for $S \subseteq A$, define $S^+ = S \cup \{a \mid \exists b \in S \text{ with } (b, a) \in R\}$

Example

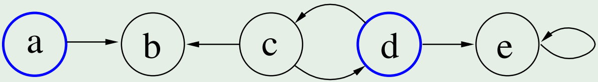

$semi(F) = \{$ ~~$\{a, c\}$~~, $\{a, d\}$, ~~$\{a\}, \{c\}, \{d\}, \emptyset$~~ $\}$

Semantics (ctd.)

Stage Extensions [Verheij, 1996]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is a stage extension of F, if

- S is conflict-free in F
- for each $T \subseteq A$ conflict-free in F, $S^+ \not\subset T^+$
  - recall $S^+ = S \cup \{a \mid \exists b \in S \text{ with } (b, a) \in R\}$

Semantics (ctd.)

Ideal Extension [Dung, Mancarella & Toni 2007]

Given an AF $F = (A, R)$. A set $S \subseteq A$ is an ideal extension of F, if

- S is admissible in F and contained in each preferred extension of F
- there is no $T \supset S$ admissible in F and contained in each of $\mathit{pref}(F)$

Properties of Ideal Extensions

For any AF F the following observations hold:

1. there exists exactly one ideal extension of F
2. the ideal extension of F is also a complete one

Relations between Semantics

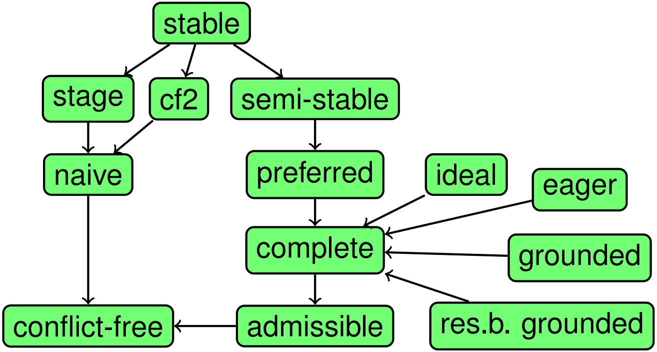

Figure: An arrow from semantics σ to semantics τ encodes that each σ-extension is also a τ-extension.

Characteristics of Argumentation Semantics

Example

$pref(F) = \{\{a, d, e\}, \{b, c, e\}, \{a, b\}\}$
$naive(F) = \{\{a, d, e\}, \{b, c, e\}, \{a, b, e\}\}$

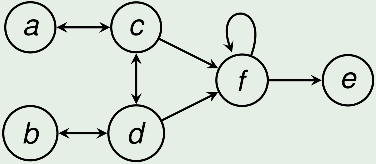

Characteristics of Argumentation Semantics

Natural Questions

- How to change the AF if we want $\{a, b, e\}$ instead of $\{a, b\}$ in $pref(F)$?
- How to change the AF if we want $\{a, b, d\}$ instead of $\{a, b\}$ in $pref(F)$?
- Can we have equivalent AFs without argument f?

$\Rightarrow$ Realizability

Some Properties . . .

Theorem

For any AFs F and G, we have

- $adm(F) = adm(G) \Longrightarrow \sigma(F) = \sigma(G)$, for $\sigma \in \{pref, ideal\}$;
- $comp(F) = comp(G) \Longrightarrow \theta(F) = \theta(G)$, for $\theta \in \{pref, ideal, ground\}$;
- no other such relation between the different semantics (*adm*, *pref*, *ideal*, *semi*, *ground*, *comp*, *stable*) in terms of standard equivalence holds.

Decision Problems on AFs

Credulous Acceptance

Cred_σ: Given AF $F = (A, R)$ and $a \in A$; is a contained in at least one σ-extension of F?

Skeptical Acceptance

Skept_σ: Given AF $F = (A, R)$ and $a \in A$; is a contained in every σ-extension of F?

If no extension exists then all arguments are skeptically accepted and no argument is credulously accepted[1].

[1] This is only relevant for stable semantics.

Decision Problems on AFs

Hence we are also interested in the following problem:

Skeptically and Credulously accepted

Skept'_σ: Given AF $F = (A, R)$ and $a \in A$; is a contained in every and at least one σ-extension of F?

Further Decision Problems

Verifying an extension

Ver_σ: Given AF $F = (A, R)$ and $S \subseteq A$; is S a σ-extension of F?

Does there exist an extension?

Exists_σ: Given AF $F = (A, R)$; Does there exist a σ-extension for F?

Does there exist a nonempty extensions?

$\text{Exists}_\sigma^{\neg\emptyset}$: Does there exist a non-empty σ-extension for F?