Reasoning under Uncertainty with Abstract Argumentation Frameworks

Matthias Thimm

Institute for Web Science and Technologies, University of Koblenz-Landau, Germany

April 26, 2017

This talk contains joint works with Anthony Hunter

Motivation

Scenario:
Decision-making under incomplete/contradictory information

- John is accused of murdering Frank
- Arguments of the court case:
  - John is innocent as long as his guilt is not proven beyond reasonable doubt ($\mathcal{I}$)
  - John is guilty as he supposedly did not like Frank ($\mathcal{G}$)
  - CCTV footage gives evidence that a person looking like John (with uncertainty $p \in [0, 1]$) was present at the time of the crime, giving a reason that John is not innocent ($\mathcal{S}_1$)
  - Other CCTV footage gives evidence that a person looking like John (with uncertainty $p' \in [0, 1]$) was not present at the time of the crime, giving a reason that John is not guilty ($\mathcal{S}_2$)
- Observations:
  - Decision-making needs to involve argumentative reasoning and
  - ... reasoning about quantitative uncertainty

Abstract Argumentation

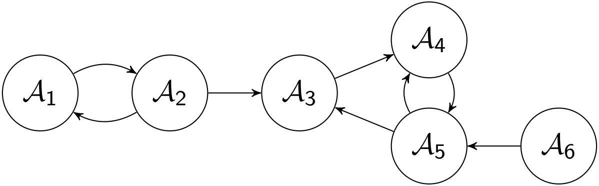

Overview

Abstract Argumentation + Probabilities

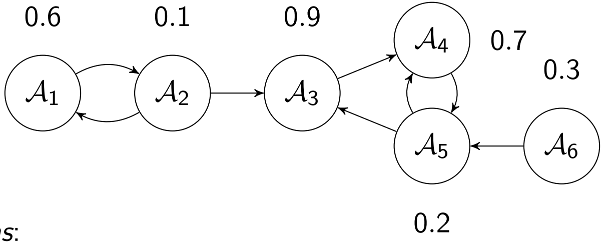

Questions:

- What are the relationships between qualitative uncertainty expressed by attacks and quantitative uncertainty expressed by probabilities?
- Given partial probabilistic information, what should the other probabilities look like?
- Given contradictory probabilistic information, what should the other probabilities look like?

Outline

Outline

Abstract Argumentation

Definition (Abstract Argumentation Framework)

An *abstract argumentation framework* AF is a tuple $\mathsf{AF} = (\mathsf{Arg}, \rightarrow)$ with arguments Arg and an attack relation $\rightarrow \subseteq \mathsf{Arg} \times \mathsf{Arg}$ [Dung,1995].

A labelling L is a function $L : \mathsf{Arg} \rightarrow \{\mathsf{in}, \mathsf{out}, \mathsf{undec}\}$ [Caminada,2006].

Definition

L is *admissible* iff for all $\mathcal{A} \in \mathsf{Arg}$

1. $L(\mathcal{A}) = \mathsf{out} \Longrightarrow \exists \mathcal{B} \in \mathsf{Arg} : L(\mathcal{B}) = \mathsf{in} \wedge \mathcal{B} \rightarrow \mathcal{A}$ and
2. $L(\mathcal{A}) = \mathsf{in} \Longrightarrow \forall \mathcal{B} \in \mathsf{Arg} : \mathcal{B} \rightarrow \mathcal{A} \Rightarrow L(\mathcal{B}) = \mathsf{out}$,

and it is *complete* if additionally $L(\mathcal{A}) = \mathsf{undec}$ implies both

3. $\neg \exists \mathcal{B} \in \mathsf{Arg} : \mathcal{B} \rightarrow \mathcal{A} \wedge L(\mathcal{B}) = \mathsf{in}$ and
4. $\exists \mathcal{B}' \in \mathsf{Arg} : \mathcal{B}' \rightarrow \mathcal{A} \wedge L(\mathcal{B}') \neq \mathsf{out}$.

Semantics

Definition

- L is *grounded* if and only if $\text{in}(L)$ is minimal.
- L is *preferred* if and only if $\text{in}(L)$ is maximal.
- L is *stable* if and only if $\text{undec}(L) = \emptyset$.
- L is *semi-stable* if and only if $\text{undec}(L)$ is minimal.

Example

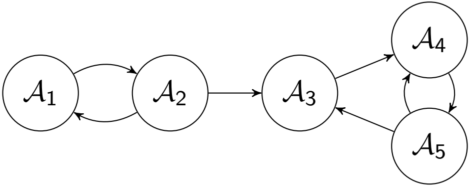

$L(\mathcal{A}_1) = \text{in}$ $\quad L(\mathcal{A}_2) = \text{out}$ $\quad L(\mathcal{A}_3) = \text{out}$
$L(\mathcal{A}_4) = \text{out}$ $\quad L(\mathcal{A}_5) = \text{in}$

L is admissible, complete, preferred, stable, and semi-stable.

$L'(\mathcal{A}_1) = \text{undec}$ $\quad L'(\mathcal{A}_2) = \text{undec}$ $\quad L'(\mathcal{A}_3) = \text{undec}$
$L'(\mathcal{A}_4) = \text{undec}$ $\quad L'(\mathcal{A}_5) = \text{undec}$

L' is admissible, complete, and grounded.

Outline

Probability Functions

Idea: Exchange "labelling" by "probability function"

Definition

Let AF $= (\text{Arg}, \rightarrow)$ be an AF. A *probability function P* on AF is a function $P : 2^{\text{Arg}} \rightarrow [0, 1]$ with

$$\sum_{E \subseteq \text{Arg}} P(X) = 1$$

and we define the probability of an argument $\mathcal{A} \in \text{Arg}$ as

$$P(\mathcal{A}) = \sum_{\mathcal{A} \in E \subseteq \text{Arg}} P(E)$$

Epistemic Extensions

Let P be any probability function.

Definition

The labelling $L_P : \text{Arg} \rightarrow \{\text{in}, \text{out}, \text{undec}\}$ defined via

- $L_P(\mathcal{A}) = \text{in} \quad \text{iff} \quad P(\mathcal{A}) > 0.5$
- $L_P(\mathcal{A}) = \text{out} \quad \text{iff} \quad P(\mathcal{A}) < 0.5$
- $L_P(\mathcal{A}) = \text{undec} \quad \text{iff} \quad P(\mathcal{A}) = 0.5$

is called the *epistemic labelling* of P. The set

$$E_P = \{\mathcal{A} \mid L_P(\mathcal{A}) = \text{in}\}$$

is called the *epistemic extension* of P.

Example

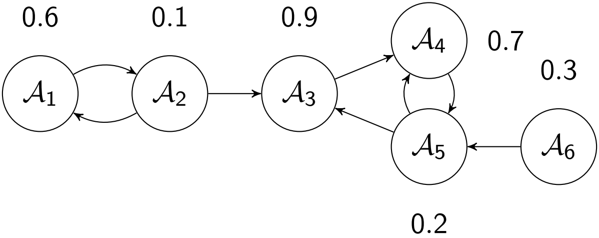

$\rightarrow E_P = \{\mathcal{A}_1, \mathcal{A}_3, \mathcal{A}_4\}$

Questions:

- When does a probability function P adhere to the structure of AF?
- When is an epistemic extension "meaningful" in some sense?
- What are the probabilistic versions of admissibility, completeness, . . . ?

Coherence

COH P is *coherent* if $\mathcal{A} \rightarrow \mathcal{B}$ implies $P(\mathcal{A}) \leq 1 - P(\mathcal{B})$

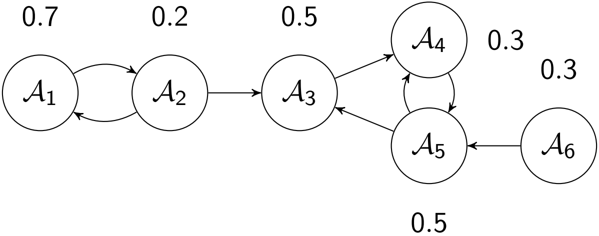

Foundation

SFOU P is *semi-founded* if $P(\mathcal{A}) \geq 0.5$ for every unattacked $\mathcal{A}$
FOU P is *founded* if $P(\mathcal{A}) = 1$ for every unattacked $\mathcal{A}$

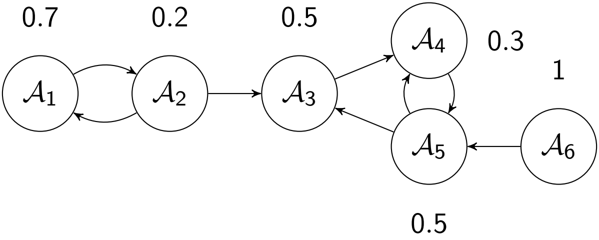

Optimism

SOPT P is *semi-optimistic* if $P(\mathcal{A}) \geq 1 - \sum_{\mathcal{B} \in \mathsf{Att}_{\mathsf{AF}}(\mathcal{A})} P(\mathcal{B})$ for every $\mathcal{A} \in \mathsf{Arg}$ with at least one attacker

OPT P is *optimistic* if $P(\mathcal{A}) \geq 1 - \sum_{\mathcal{B} \in \mathsf{Att}_{\mathsf{AF}}(\mathcal{A})} P(\mathcal{B})$ for every $\mathcal{A} \in \mathsf{Arg}$

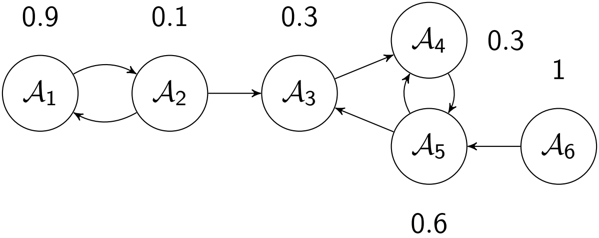

Justifiability/Ternary

<u>JUS</u> P is *justifiable* if P is coherent and optimistic
TER P is *ternary* if $P(\mathcal{A}) \in \{0, 0.5, 1\}$ for every $\mathcal{A} \in \text{Arg}$

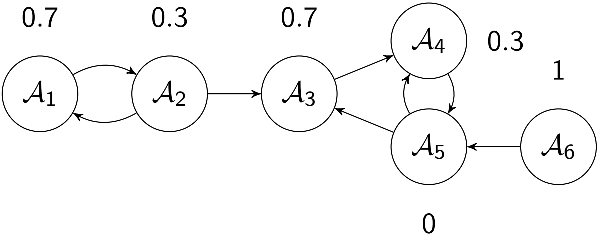

Rationality

RAT P is *rational* if $\mathcal{A} \to \mathcal{B}$ then $P(\mathcal{A}) > 0.5$ implies $P(\mathcal{B}) \leq 0.5$

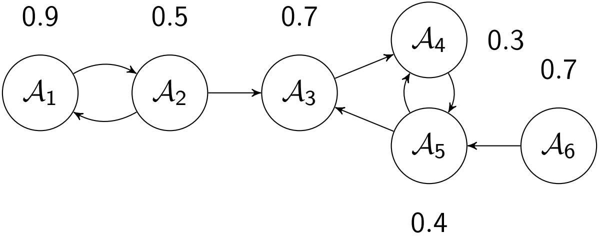

Involution

INV P is *involutary* if $\mathcal{A} \to \mathcal{B}$ implies $P(\mathcal{A}) = 1 - P(\mathcal{B})$

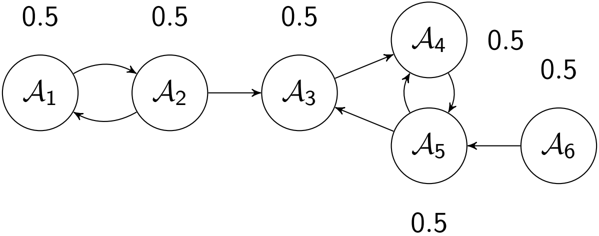

Neutrality/Maximality/Minimality

NEU P is *neutral* if $P(\mathcal{A}) = 0.5$ for every $\mathcal{A} \in \text{Arg}$
<u>MAX</u> P is *maximal* if $P(\mathcal{A}) = 1$ for every $\mathcal{A} \in \text{Arg}$
MIN P is *minimal* if $P(\mathcal{A}) = 0$ for every $\mathcal{A} \in \text{Arg}$

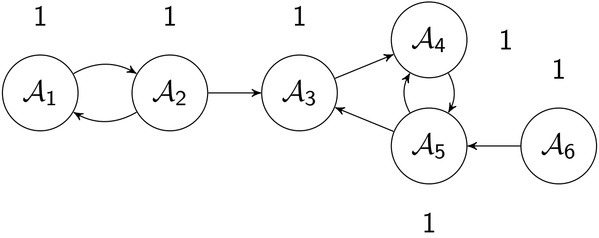

Relationships between probabilistic notions 1/2

COH	P is *coherent* if $\mathcal{A} \to \mathcal{B}$ implies $P(\mathcal{A}) \leq 1 - P(\mathcal{B})$
SOPT	P is *semi-optimistic* if $P(\mathcal{A}) \geq 1 - \sum_{\mathcal{B} \in \mathrm{Att}_{\mathrm{AF}}(\mathcal{A})} P(\mathcal{B})$ for every $\mathcal{A} \in \mathrm{Arg}$ with at least one attacker
OPT	P is *optimistic* if $P(\mathcal{A}) \geq 1 - \sum_{\mathcal{B} \in \mathrm{Att}_{\mathrm{AF}}(\mathcal{A})} P(\mathcal{B})$ for every $\mathcal{A} \in \mathrm{Arg}$
FOU	P is *founded* if $P(\mathcal{A}) = 1$ for every unattacked $\mathcal{A}$
JUS	P is *justifiable* if P is coherent and optimistic
RAT	P is *rational* if $\mathcal{A} \to \mathcal{B}$ then $P(\mathcal{A}) > 0.5$ implies $P(\mathcal{B}) \leq 0.5$
INV	P is *involutary* if $\mathcal{A} \to \mathcal{B}$ implies $P(\mathcal{A}) = 1 - P(\mathcal{B})$

Observations

OPT = SOPT + FOU

JUS $\Rightarrow$ COH

COH $\Rightarrow$ RAT

INV $\Rightarrow$ COH

INV $\Rightarrow$ SOPT

Relationships between probabilistic notions 2/2

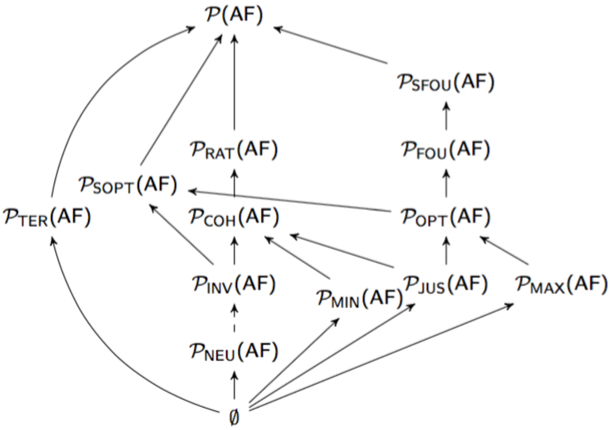

Relationships to classical abstract argumentation

How do these probabilistic concepts relate to concepts from abstract argumentation?

Observations

- $P \in \mathcal{P}_{\mathsf{COH}}(\mathsf{AF}) \cap \mathcal{P}_{\mathsf{FOU}}(\mathsf{AF}) \cap \mathcal{P}_{\mathsf{TER}}(\mathsf{AF})$ if and only if L_P is a complete labelling.
- If L_P is admissible then P is justifiable
- The grounded labelling corresponds to the justifiable probability function with maximum entropy
- Stable labellings correspond to justifiable probability functions with minimum entropy
- If P is rational then E_P is conflict-free

Outline

Decision-making with Partial Information

Assume an agent's knowledge consists of an AF and partial probabilistic information:

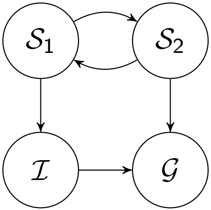

$+ \qquad P(\mathcal{S}_1) = 0.7$

Question: What should be reasonably inferred for $P(\mathcal{S}_2)$, $P(\mathcal{I})$, and $P(\mathcal{G})$?

Formalization

- β : Arg $\rightarrow [0,1]$ partial function, *partial probability assignment*
- Probability function $P \in \mathcal{P}(\mathsf{AF})$ is β-*compliant* if for every $\mathcal{A} \in \operatorname{dom} \beta$ we have $\beta(\mathcal{A}) = P(\mathcal{A})$; let $\mathcal{P}^{\beta}(\mathsf{AF}) \subseteq \mathcal{P}(\mathsf{AF})$ be the set of all such functions
- $T \subseteq \{\mathsf{RAT}, \mathsf{COH}, \mathsf{SFOU}, \mathsf{FOU}, \mathsf{OPT}, \mathsf{SOPT}, \mathsf{JUS}\}$
- Define

$$\mathcal{P}^{\beta}_{T}(\mathsf{AF}) = \mathcal{P}_{T}(\mathsf{AF}) \cap \mathcal{P}^{\beta}(\mathsf{AF})$$

- Assume $\mathcal{P}^{\beta}_{T}(\mathsf{AF}) \neq \emptyset$
- *Possible probabilities* of $\mathcal{A}$ under constraints of β

$$\mathsf{p}^{\beta}_{T,\mathsf{AF}}(\mathcal{A}) = \{P(\mathcal{A}) \mid P \in \mathcal{P}^{\beta}_{T}(\mathsf{AF})\}$$

Example

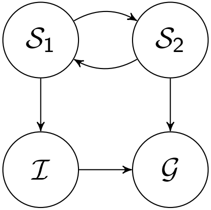

$+ \quad P(\mathcal{S}_1) = 0.7$

- Assume $T_1 = \{\text{COH}\}$
- Then

$$\mathsf{p}^{\beta_1}_{T_1,\mathsf{AF}}(\mathcal{S}_2) = [0, 0.3]$$
$$\mathsf{p}^{\beta_1}_{T_1,\mathsf{AF}}(\mathcal{I}) = [0, 0.3]$$
$$\mathsf{p}^{\beta_1}_{T_1,\mathsf{AF}}(\mathcal{G}) = [0, 0.7]$$

Results

$T \subseteq \{\text{COH,SFOU,FOU, OPT,SOPT,JUS}\}$

1. For all $\beta : \text{Arg} \rightarrow [0,1]$, $\mathcal{P}^{\beta}(\text{AF}) \neq \emptyset$
2. $\mathcal{P}^{\beta}(\text{AF})$ is connected, convex, and closed.
3. $\mathcal{P}_{T}(\text{AF})$ and $\mathcal{P}_{T}^{\beta}(\text{AF})$ are connected, convex, and closed.
4. $\text{p}_{T,\text{AF}}^{\beta}(\mathcal{A})$ is connected, convex, and closed.
5. Deciding $p \in \text{p}_{T,\text{AF}}^{\beta}(\mathcal{A})$ for some $p \in [0,1]$ is NP-complete.
6. Deciding $[l,u] = \text{p}_{T,\text{AF}}^{\beta}(\mathcal{A})$ for some $l,u \in [0,1]$ is D^{P}-complete.
7. Computing $l,u \in [0,1]$ such that $[l,u] = \text{p}_{T,\text{AF}}^{\beta}(\mathcal{A})$ is FP^{NP}-complete.

Outline

Decision-making with Contradictory Information

Assume an agent's knowledge consists of an AF and partial probabilistic information:

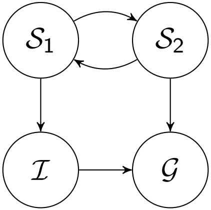

$+ \qquad P(\mathcal{S}_1) = 0.7,\ P(\mathcal{S}_2) = 0.4$

Question: What should be reasonably inferred for $P(\mathcal{I})$ and $P(\mathcal{G})$?

- Recall
  - $\mathcal{P}_T(\mathsf{AF})$ = Probability functions satisfying T
  - $\mathcal{P}^\beta(\mathsf{AF})$ = Probability functions compatible with β
  - $\mathcal{P}_T^\beta(\mathsf{AF}) = \mathcal{P}_T(\mathsf{AF}) \cap \mathcal{P}^\beta(\mathsf{AF})$
- We now allow for $\mathcal{P}_T^\beta(\mathsf{AF}) = \emptyset$

We use *inconsistency measures* [Thimm 2013; De Bona and Finger 2015; Grant and Hunter 2013] as an analytical tool

Let d be some distance (e. g. p-norm distance)

Definition

$\mathcal{I}_T^d(\beta, \mathsf{AF}) = d(\mathcal{P}^\beta(\mathsf{AF}), \mathcal{P}_T(\mathsf{AF}))$

$\mathcal{I}_T^d(\beta, \mathsf{AF})$ = degree of inconsistency of β wrt. AF

Example

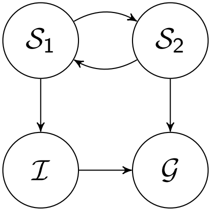

- β_1 defined by $\beta_1(\mathcal{S}_1) = 0.7$ and $\beta_1(\mathcal{S}_2) = 0.4$ ($T_1 = \{\text{COH}\}$)

$$\mathcal{I}^{d_1}_{T_1}(\beta_1, \text{AF}) = 0.1 \qquad (d_1 = \text{ Manhattan distance})$$
$$\mathcal{I}^{d_2}_{T_1}(\beta_1, \text{AF}) = 0.037 \qquad (d_2 = \text{ Euclidean distance})$$

- β_2 defined by $\beta_2(\mathcal{S}_1) = 0.8$ and $\beta_2(\mathcal{S}_2) = 0.9$:

$$\mathcal{I}^{d_1}_{T_1}(\beta_2, \text{AF}) = 0.7$$
$$\mathcal{I}^{d_2}_{T_1}(\beta_2, \text{AF}) \approx 0.403$$

Inconsistent-tolerant Reasoning

$\mathcal{I}^d_T(\beta, \mathsf{AF}) = d(\mathcal{P}^\beta(\mathsf{AF}), \mathcal{P}_T(\mathsf{AF}))$

Idea:

- Select those probability functions of reasoning that minimize the above distance
- Apply the same mechanism to those functions as in the case $\mathcal{P}_T(\mathsf{AF}) \cap \mathcal{P}^\beta(\mathsf{AF}) \neq \emptyset$

Definition

$\Pi_{T,d,\mathsf{AF}}(\beta) = \{P \in \mathcal{P}^\beta(\mathsf{AF}) \mid d(P, \mathcal{P}_T(\mathsf{AF})) \text{ minimal}\}$

$\pi^{\beta,d}_{T,\mathsf{AF}}(\mathcal{A}) = \{P(\mathcal{A}) \mid P \in \Pi_{T,d,\mathsf{AF}}(\beta)\}$

Example

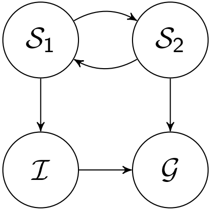

- β_1 defined by $\beta_1(\mathcal{S}_1) = 0.7$ and $\beta_1(\mathcal{S}_2) = 0.4$

$$\pi^{\beta_1, d_2}_{T_1, \mathsf{AF}}(\mathcal{I}) \approx [0.0284, 0.383]$$
$$\pi^{\beta_1, d_2}_{T_1, \mathsf{AF}}(\mathcal{G}) \approx [0.0270, 0.682]$$

Some technical Results

1. If $\mathcal{P}^{\beta}_{T}(\mathsf{AF}) \neq \emptyset$ then $\Pi_{T,d,\mathsf{AF}}(\beta) = \mathcal{P}^{\beta}_{T}(\mathsf{AF})$ for every pre-metrical distance measure d
2. $\Pi_{T,d,\mathsf{AF}}(\beta) \neq \emptyset$
3. $\Pi_{T,d,\mathsf{AF}}(\beta)$ is connected, convex, and closed
4. Given the value of $\mathcal{I}^{d}_{T}(\beta, \mathsf{AF})$ the computational complexity of reasoning tasks is not harder as in the case of partial assignments

Outline

Summary

- We proposed a probabilistic setting on top of abstract argumentation
- Several properties (=semantics) can be extended to the probabilistic setting
- We applied our probabilistic framework to the problem of decision-making with incomplete probabilistic information
- Completing incomplete and "Repairing" contradictory probabilistic information

Thank you for your attention

References:

[Thimm. A Probabilistic Semantics for Abstract Argumentation. ECAI 2012]
[Thimm. Inconsistency Measures for Probabilistic Logics. In AIJ 197:1–24, 2013]
[Hunter, Thimm. Probabilistic Argumentation with Incomplete Information. ECAI 2014]
[Potyka, Thimm. Probabilistic Reasoning with Inconsistent Beliefs using Inconsistency Measures. IJCAI 2015]
[Hunter, Thimm. On Partial Information and Contradictions in Probabilistic Abstract Argumentation. KR 2016]