Probability in Argumentation

Nir Oren
n.oren@abdn.ac.uk

University of Aberdeen

5 Sep 2014

Overview

- Why probabilities?
- Approaches
- Applications

Argumentation

- Argumentation formalisms facilitate reasoning in the presence of conflicting information.

Alice has black hair.
Alice has red hair.

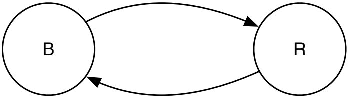

Argumentation

- Argumentation formalisms facilitate reasoning in the presence of conflicting information.

Bob says that Alice has black hair.
Robert says that Alice has red hair.

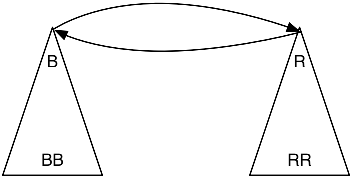

Argumentation

- Argumentation formalisms facilitate reasoning in the presence of conflicting information.

Bob says that Alice has black hair.
Robert says that Alice has red hair.
Charlie says he's pretty sure that Bob hasn't met Alice, but is almost certain that Robert has met her.

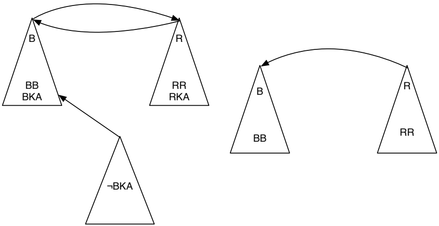

- We can use preferences to obtain extensions.
- But these do not capture the "pretty sure" and "almost" — extensions are not different.

Uncertainty

- Uncertainty is a common facet in everyday life, and is related to, but independent from conflicting information.
- We need to be able to reason about
  - Uncertainty about arguments (i.e. there is a 0.7 likelihood of this set of facts holding).
  - Arguments about uncertainty (the likelihood of this fact is 0.6 because of ...).
- There are still many open questions in both of these areas.

Probabilistic Argumentation Frameworks (PrAFs)

- PrAFs (Li, TAFA-11) are a recent, very popular approach to modelling uncertainty in abstract argumentation.
- PrAFs extend a standard DAF with probabilistic concepts.

$$\langle A, D \rangle$$

Probabilistic Argumentation Frameworks (PrAFs)

- PrAFs (Li, TAFA-11) are a recent, very popular approach to modelling uncertainty in abstract argumentation.
- PrAFs extend a standard DAF with probabilistic concepts.

$$\langle A, D, P_A, P_D \rangle$$

- P_A,P_D encodes the likelihood of an argument or attack.

Interpreting PrAFs

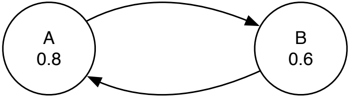

- We can interpret PrAFs via a frequentist approach to probability: $P_A(A) = 0.8$ means that in 8 out of 10 possible worlds (or Argument Frameworks), A exists.

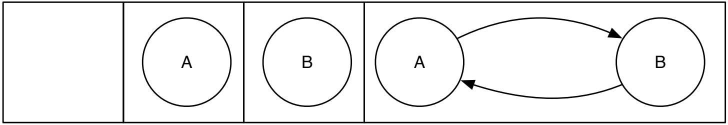

Likelihoods of Argument Frameworks

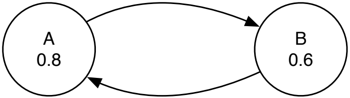

- $P(\emptyset, \emptyset) = ?$
- $P(\{A\}, \emptyset) = ?$
- $P(\{B\}, \emptyset) = ?$
- $P(\{A, B\}, \{(A, B), (B, A)\}) = ?$

Likelihoods of Argument Frameworks

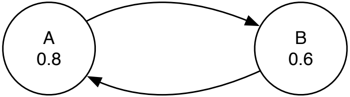

- $P(\emptyset, \emptyset) = 0.08$
- $P(\{A\}, \emptyset) = 0.32$
- $P(\{B\}, \emptyset) = 0.12$
- $P(\{A, B\}, \{(A, B), (B, A)\}) = 0.48$
- Each of these DAFs are *induced* from the original PrAF.

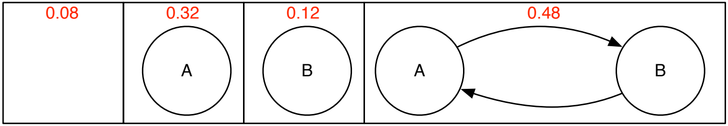

Semantics

- Unlike traditional frameworks, extensions are probabilistic, indicating the likelihood that a set of arguments appears within some extension.
- This probability is computed as the sum of probabilities of the AFs where the argument appears in the Dung extension.

$$P(\emptyset, \emptyset) = 0.08 \qquad P(\{A\}, \emptyset) = 0.32$$

$$P(\{B\}, \emptyset) = 0.12 \qquad P(\{A, B\}, \{(A, B), (B, A)\}) = 0.48$$

- $P(\{A\} \in \text{Grounded}) = 0.32$
- $P(\{A\} \in \text{Preferred(credulous)}) = 0.8$
- $P(\{A\} \in \text{Preferred(skeptical)}) = 0.32$

Computing Probabilistic Justification

Require: A PrAF $\langle A, D, P_A, P_D \rangle$
Require: A set of arguments $T \subseteq A$
Require: X, a semantics

1: $p \leftarrow 0$
2: **for all** $A' \subseteq A, D' \subseteq D$ $(a, b) \in D'$, only if $a, b \in A'$ **do**
3: **if** T is in the X extension of (A', D') **then**
4: $p \leftarrow p + \prod_{a \in A'} P_A(a) \prod_{a \in A \setminus A'} (1 - P_A(a)) \prod_{d \in D'} P_D(d) \prod_{d\ inD \setminus D'} (1 - P_D(d))$
5: **end if**
6: **end for**
7: **return** p

Computing Probabilistic Justification

- The naive algorithm considers all combinations of arguments and attacks. Exponential computational complexity.
- Can we do better?
  - Special cases (Fazzinga et al., IJCAI-13).
  - Approximate solutions

Admissible semantics

- An admissible set Ad is
  - Conflict free (i.e. there is no $a, b \in Ad$ s.t. $(a, b) \in D$
  - contains arguments acceptable w.r.t Ad (i.e. if b attacks $a \in Ad$ then there is a $c \in Ad$ which attacks b).
- So what?

Admissible semantics

- An admissible set *Ad* is
  - Conflict free (i.e. there is no $a, b \in Ad$ s.t. $(a, b) \in D$
  - contains arguments acceptable w.r.t *Ad* (i.e. if *b* attacks $a \in Ad$ then there is a $c \in Ad$ which attacks *b*).
- The probability that a set $S \subseteq A$ is admissible is the product of
- the likelihood that *S* appears;
- the likelihood that it is conflict free (i.e. that no attacks between its elements appear);
- the likelihood that
  - No argument that might attack an element of *S* exists; or
  - If such an element does exist, it does not attack an element of *S*; or
  - If it does attack an element of *S*, an attack agains the attacker exist from *S*.

Admissible Semantics

Lemma 1 *Given a PrAF* $\mathcal{F} = \langle A, P_A, D, P_D \rangle$ *and a set* $S \subseteq A$, $Pr^{ad}_{\mathcal{F}}(S) = P_1(S) \cdot P_2(S) \cdot P_3(S)$, *where*[3]*:*

$P_1(S) = \prod_{a \in S} P_A(a)$,

$P_2(S) = \prod_{\substack{\langle a,b \rangle \in D \\ \wedge a \in S \\ \wedge b \in S}} \left(1 - P_D(\langle a,b \rangle)\right)$, *and*

$P_3(S) = \prod_{d \in A \setminus S} \left(P_{31}(S,d) + P_{32}(S,d) + P_{33}(S,d)\right)$, *where:*

$P_{31}(S,d) = 1 - P_A(d)$,

$P_{32}(S,d) = P_A(d) \times \prod_{\substack{\langle d,b \rangle \in D \\ \wedge b \in S}} \left(1 - P_D(\langle d,b \rangle)\right)$,

$$P_{33}(S,d) = P_A(d) \times \left(1 - \prod_{\substack{\langle d,b \rangle \in D \\ \wedge b \in S}} \left(1 - P_D(\langle d,b \rangle)\right)\right) \times \left(1 - \prod_{\substack{\langle a,d \rangle \in D \\ \wedge a \in S}} \left(1 - P_D(\langle a,d \rangle)\right)\right).$$

Stable Semantics

- A similar line of reasoning can be used to the stable semantics.
- The main change is that arguments outside the stable set *must* be attacked.

Lemma 2 *Given a PrAF* $\mathcal{F} = \langle A, P_A, D, P_D \rangle$, *and a set* $S \subseteq A$, $Pr_{\mathcal{F}}^{st}(S) = P_1(S) \cdot P_2(S) \cdot P_3(S)$, *where* $P_1(S)$ *and* $P_2(S)$ *are defined as in Lemma 1 and* $P_3(S) = \prod_{d \in A \setminus S} \left(P_{31}(S,d) + P_{32}(S,d) \right)$, *where*

$$P_{31}(S,d) = 1 - P_A(d), \text{ and}$$

$$P_{32}(S,d) = P_A(d) \times \left(1 - \prod_{\substack{\langle a,d \rangle \in D \\ \wedge a \in S}} (1 - P_D(\langle a,d \rangle))\right).$$

Approximating Solutions

- Other semantics are computationally hard ($FP^{\#P} - complete$).
- To address this we utilise Monte-Carlo sampling.
- Algorithm overview:
  - We randomly generate a DAF based on the underlying probability distribution of the PrAF
  - Compute the extension of the generated DAF and record if the set we're testing (X) for is within the extension.
  - Repeat the above n times, assume that the set was within the extension m times.
  - $P'_{PrAF}(X) = m/n$.
- As $n \to \infty$, $P_{PrAF}(X) - P'_{PrAF}(X) \to 0$

When to stop?

- If we're willing to specify some confidence interval ϵ, we can stop when we're (nearly) sure that our solution lies within this confidence interval, i.e. when $|P_{PrAF}(X) - P'_{PrAF}(X)| < \epsilon$
- From basic statistics, if $z_{1-\alpha/2}$ to 1.96)

$$n > \frac{p'(1-p')}{\epsilon^2}(1.96)^2$$

- But, when we run the first trial, $p' = 0$ or $p' = 1$.
- We therefore use the Agresti-Coull method which perturbs the true number of trials and successes slightly to obtain

$$n > \frac{4p'(1-p')}{\epsilon^2} - 4$$

for a tweaked version of p'.

When to stop?

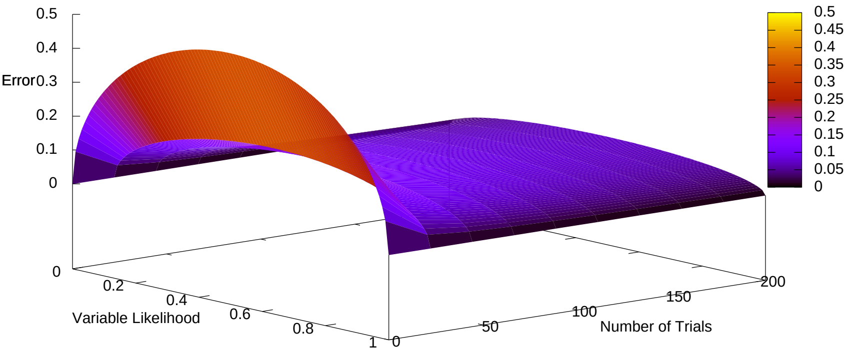

Error
0.5
0.4
0.3
0.2
0.1
0
0
0.2
0.4
0.6
0.8
1
Variable Likelihood
0
50
100
150
200
Number of Trials
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0

Evaluation

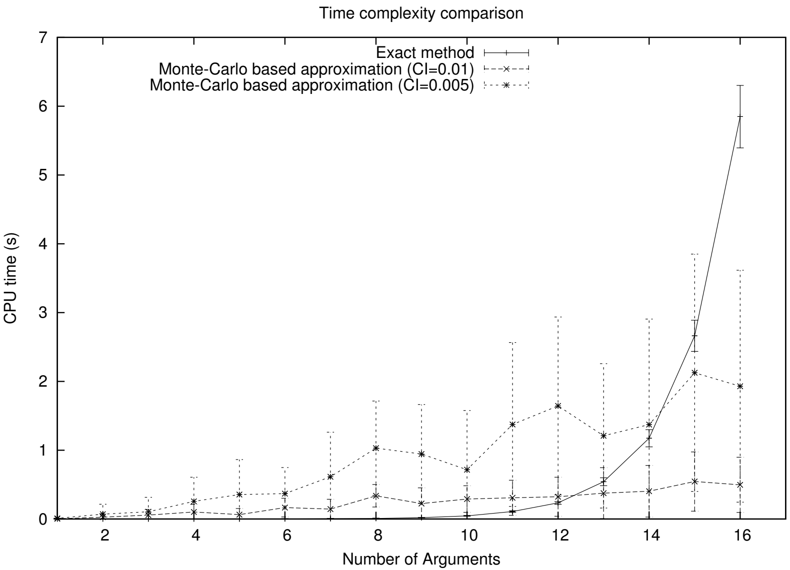
Time complexity comparison
Exact method
Monte-Carlo based approximation (CI=0.01)
Monte-Carlo based approximation (CI=0.005)
CPU time (s)
0
1
2
3
4
5
6
7
2
4
6
8
10
12
14
16
Number of Arguments

Observations

- Monte-Carlo evaluation time increases linearly (due to the longer time taken to evaluate the extension).
- As the permitted error shrinks, the standard deviation increases.

Where are we?

- Given a PrAF and a set of arguments, we can determine the probability that these arguments appear in an extension.
- We have assumed that the likelihood of arguments appearing in the induced DAF is independent of other arguments appearing (and that defeats are conditional only on the source and target arguments appearing).
- This is not a realistic assumption — arguments are (usually) composed of sub arguments; the parent argument can only appear if the child argument is present.
- Given that *a* is a parent of *b*, *a* can only be present in an extension if *b* is present.
- In other words, we need to encode $P(a|b)$.

Removing the Independence Assumption

- A sub-argument *supports* its parent.
- A bipolar argumentation framework is needed to encode support.
- We extend Evidential Argument Frameworks (EAFs) into PrEAFs.

EAFs, Features

- Support.
- Collective attack and support.
- A special argument, η whose support is, in a sense, required for further support or attack to be successful.
- Semantics are very similar to those of Dung.

Evidential Argument Frameworks (EAFs)

- $\langle Args, R_a, R_e \rangle$, $R_a, R_e : 2^{Args} \backslash \emptyset \times Args$
- $\eta \in Args$ such that $(\_, \eta) \notin R_a, R_e$ and $(\{\ldots, \eta\}, \_) \notin R_a$.
- A set S provides evidential support to an argument a iff $a = \eta$ or there is a non-empty $S' \subseteq S$ s.t. $(S', a) \in R_e$ and $\forall x \in S'$, x has evidential support form $S \backslash \{a\}$.
- S carries out an evidence supported attack on a iff $(S', a) \in R_a$, for $S' \subseteq S$ and any $s \in S'$ has evidential support from S.

Other concepts

- An argument a is acceptable w.r.t S iff
  - it is e-supported by S; and
  - if T carries out a minimal e-supported attack against a, S carries out an e-supported attack against T.
- Self-supporting — all arguments in S are e-supported by S.
- conflict free — no $a \in S, B \subseteq S$ s.t. $(B, a) \in R_a$.
- admissible — conflict free and acceptable.
- preferred — maximal w.r.t set inclusion admissible.
- ....

EAFs

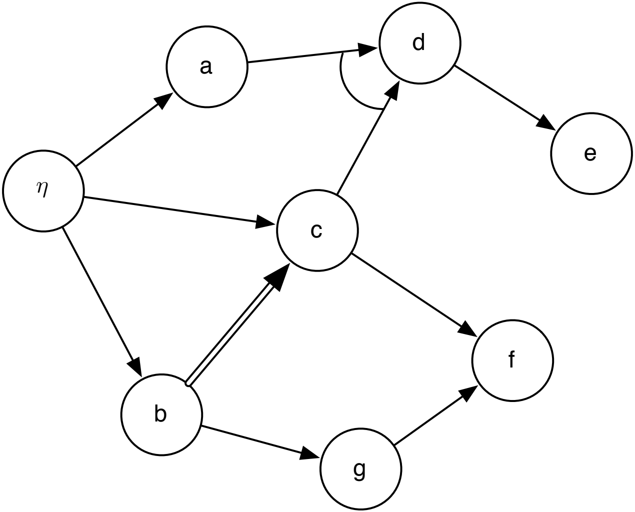

EAFs, η

- For an argument to be considered for inclusion within an extension, or to successfully attack some other argument there must be a *chain of support* from η to the argument.
- η
  - eliminates the possibility of self supporting cycles.
  - Helps ensure consistency of the system.
  - Allows for a simple representation of argument schemes?
- $\eta \to a$ can model a strict or default fact within the system.
- EAFs are, in a sense, underspecified structured argument frameworks.

PrEAFs

$$PrEAF = \langle A, R_d, R_s, P_s \rangle$$

- $P_s : R_s \rightarrow (0, 1]$ captures the probability of support.
- We assume that if SR_sa, then it is not the case that $S' \subseteq S$ such that $S'R_sa$.
- We also assume no *support loops*.

PrEAFs

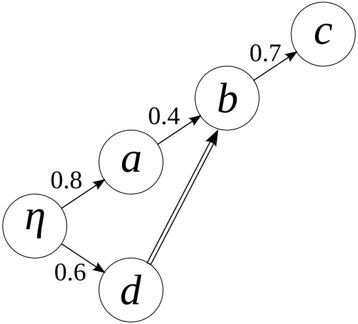

Inducible EAF

- An EAF $I = \langle A^I, R^I_d, R^I_s \rangle$ can be induced from a PrEAF $P = \langle A, R_d, R_s, P_s \rangle$ iff all of the following hold.
  - $A^I \subseteq A$ and $\eta \in A^I$
  - $R^I_s \subseteq R_s$
  - $R^I_d = R_d \cap (A^I \times A^I)$
  - $\forall SR_s a \in R_s$ s.t. $S \subseteq A^I$ and $P_s(SR_s a) = 1$, $SR_s a \in R^I_s$
  - $\forall a \in A \backslash \{\eta\}$ s.t. $\exists SR_s a \in R^I_s$, $a \in A^I$

$$P_{PrEAF}(X) = \sum_{EAF \in I(PrEAF)} P^I_{PrEAF}(EAF)\xi^S(X, EAF)$$

- We need a way to compute $P^I_{PrEAF}(EAF)$.

Computing Inducible EAF Probabilities

- We begin by considering whether an induced EAF is valid, i.e., can be induced from some given PrEAF.
- A naïve approach will consider every subset of 2^A. If this set of arguments does not include any unsupported arguments, then it is a valid inducible EAF.
- Given a valid inducible EAF, its probability of being induced is the joint probability that
  - Its arguments appear; and
  - Those arguments not in the EAF do not appear.
- These depend on the probabilities that an argument's supporting arguments are also present, all the way up to η.
- We can therefore identify inducible EAFs by beginning at η, and identifying which EAFs can be induced, repeating the process for these induced EAFs until no more EAFs can be induced.

Inducing EAFs

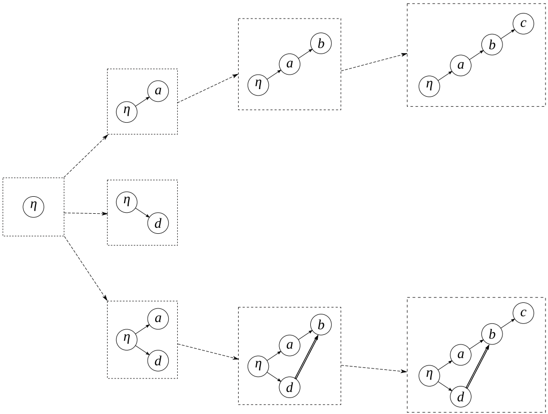

- There's a small complication when creating this tree — due to arguments supporting multiple other arguments, there may be two ways of inducing the same EAF. Our process must prevent double-counting of EAFs.

Induced EAF probabilities

- The probability that an EAF $\langle A^I, R_s^I, R_d^I \rangle$ is induced is the probability that its arguments do appear, and arguments not within it do not appear (given the arguments that are within it).

$$P(A^I) \times P(\overline{A \backslash A^I} | A^I)$$

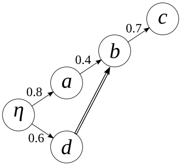

- $P^I_{PrEAF}(\{\eta, a, d\}) = P(\{\eta, a, d\}) \times P(\overline{\{b, c\}} | \{\eta, a, d\})$

$P(\overline{A \backslash A^l} | A^l)$

- Consider the probability that all arguments except those in A^l are not present within the induced EAF (given that A^l is present).
- We have two possible situations
  - An argument not in A^l supports an argument in A^l.
  - An argument in A^l supports an argument not in A^l.
- So we must consider the support links

$$\{r_s | r_s \in R_s \backslash R_s^l \text{ where } r_s = (S, a) \text{ and } S \subseteq A^l\}$$

- $P(\overline{A \backslash A^l} | A^l)$ is equivalent to the likelihood that each of these support links is not induced, i.e.

$$P(\overline{A \backslash A^l} | A^l) = \prod_{\{r_s | r_s \in R_s \backslash R_s^l \text{ and } Src(r_s) \subseteq A^l\}} (1 - P_s(r_s))$$

$P(A^I)$

- We can compute $P(A^I)$ recursively, based on the probability of its parent EAF being induced.

$$P(A^I) = \begin{cases} 1 & \text{if } I = \langle\{\eta\}, \emptyset, \emptyset\rangle \\ P(A^F) \times P((A^I \backslash A^F)|A^F) & \text{otherwise} \end{cases}$$

- We know $P(A^F)$ — the probability of the parent EAF being induced — by using our tree expansion.
- $P((A^I \backslash A^F)|A^F)$ is the conditional probability that an argument is present in the induced EAF but not the parent, given the arguments already present in the parent EAF.

$P(A^I \backslash A^F | A^F)$

- If we have only one possible support link from the parent EAF to all introduced arguments, then the probability of these new arguments appearing is simply the product of all support link likelihoods.

$$P(A^I \backslash A^F | A^F) = \prod_{\{a | a \in A^I \backslash A^F \text{ and } Tgt(r_s)=a\}} P_s(r_s)$$

- If multiple ways exist of supporting a newly appearing argument, then we must consider the likelihood that at least one of these exist.

$$P(a|A^F) = 1 - \prod_{r_s \in Sups(a,I,A^F,A^{F'})} (1 - P_s(r_s))$$

$P(A^I \backslash A^F | A^F)$

- Extending this to all arguments, we obtain

$$P(A^I \backslash A^F | A^F) = \prod_{a \in A^I \backslash A^F} \left(1 - \left(\prod_{r_s \in Sups\left(a, I, A^F, A^{F'}\right)} (1 - P_s(r_s)) \right) \right)$$

- Here, *Sups* function identifies the newly appearing support links for an argument *a* given an EAF, its parent and grandparent.

Where are we?

- We can now expand an induced EAF tree, starting at an EAF containing only the argument η, and compute the probability of a specific node being induced in this tree based on its parent's probability of being induced.
- By computing the extension of this induced EAF, and summing over all other EAFs, we can compute the likelihood of a set of arguments appearing within an extension.
- Complexity is dependent on the number of possible induced EAFs.
- If everything is supported by η (i.e., reduced to a PrAF), we have exponential complexity.
- Termination of our algorithms are guaranteed.

What do probabilities mean?

1. Likelihood of an argument being considered justified (Hunter, COMMA-12)
2. Likelihood that an argument is known by an agent (Li et al, TAFA-11,COMMA-12,ArgMAS-13)
3. Likelihood that an agent believes an argument (Thimm, ECAI-12, ECAI-14, Hunter, IJAR-13, ArXiv-14)

What do probabilities mean?

1. Likelihood of an argument being considered justified (Hunter, COMMA-12)
2. Likelihood that an argument is known by an agent (Li et al, TAFA-11,COMMA-12,ArgMAS-13)
3. Likelihood that an agent believes an argument (Thimm, ECAI-12, ECAI-14, Hunter, IJAR-13, ArXiv-14)

- Structural uncertainty - uncertainty about the structure of the argument graph (1 and 2).
- Epistemic uncertainty - uncertainty about agent beliefs (3).

Epistemic Extensions (taken from Hunter, ArXIV-14)

- A probability function maps sets of arguments to a probability value $P : 2^A \rightarrow [0, 1]$, s.t. $\sum'_A \subseteq AP(A) = 1$

$$P(a) = \sum_{a \in E \subseteq A} P(E)$$

- Arguments are labelled based on the probability associated with them: a is `in` if $(P(a) > 0.5)$, `out` if $P(a) < 0.5$ and `undec` otherwise.
- *What constraints can be placed on the probability function?*

Some Constraints

COH For every $a, b \in A$, if $a \rightarrow b$, then $P(a) \leq 1 - P(b)$

SFOU If $P(a) \geq 0.5$ for every $a \in A$ which is not attacked.

FOU If $P(a) = 1$ for every $a \in A$ which is not attacked.

SOPT If $P(a) \geq 1 - \sum_{b \text{ s.t. } b \rightarrow a} P(b)$ whenever an attack against a exists.

OPT If $P(a) \geq 1 - \sum_{b \text{ s.t. } b \rightarrow a} P(b)$.

JUS If **COH** and **OPT**

TER If $P(a) \in \{0, 0.5, 1\}$ for any $a \in A$

Properties

- $P_{JUS} \subset P_{COH} \subset P$
- $P_{OPT} = P_{SOPT} \cap P_{FOU}$
- $P_{FOU} \subset P_{SFOU}$
- $\emptyset \subset P_{TER} \subset P$
- If a labelling is admissible, then $P \in P_{SFOU}$
- A complete probability function is one s.t. $P \in P_{COH} \cap P_{FOU} \cap P_{TER}$

Classical Extensions

- Given a complete probability function, the following association between restrictions and classical extensions exists.

No restriction	Complete
No a s.t. $P(a) = 0.5$	Stable
Maximal arguments s.t. $P(a) = 1$	Preferred
Maximal arguments s.t. $P(a) = 0$	Preferred
Maximal arguments s.t. $P(a) = 0.5$	Grounded
Minimal arguments s.t. $P(a) = 1$	Grounded
Minimal arguments s.t. $P(a) = 0$	Grounded
Minimal arguments s.t. $P(a) = 0.5$	Stable

Non-standard Extensions

- Additional properties can be introduced.

  RAT If $a \rightarrow b$ then $P(a) > 0.5$ implies $P(b) \leq 0.5$
  NEU P(a)=0.5 for all $a \in A$
  MAX P(a)=1 for all $a \in A$
  MIN P(a)=0 for all $a \in A$
  INV If $a \rightarrow b$ then $P(a) = 1 - P(b)$

Non-standard Extensions

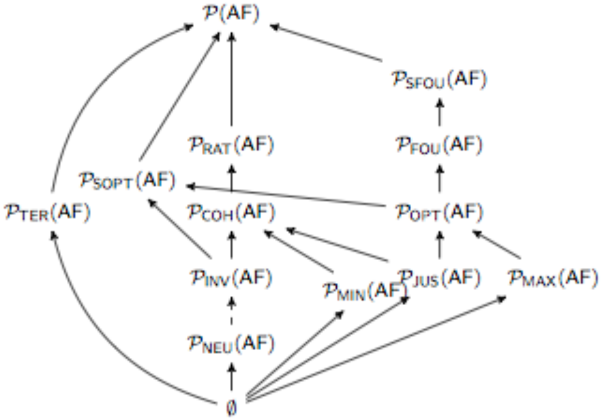

More results

- For an admissible labelling, there is an associated RAT probability function, whose in labellings are conflict-free.
- If an AF contains an odd cycle, then any INV probability function is neutral.
- All attackers of A have the same probability value if an INV function is used.
- If attacks exist between arguments, no function exists which satisfies RAT and MAX.

So What?

- Can we use these properties to assign probabilities to arguments?
- Assume a partial function $\pi : A \to [0, 1]$
- What are the “best” probabilities to assign to arguments not in the domain of π?

The Idea

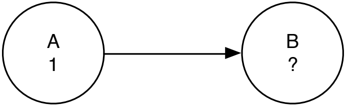

The Idea

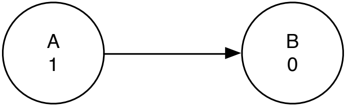

The Idea

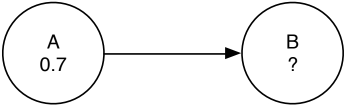

The Idea

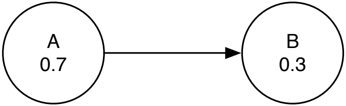

The Idea

- What if we want COH (If $a \rightarrow b$ then $P(a) \leq 1 - P(b)$)?

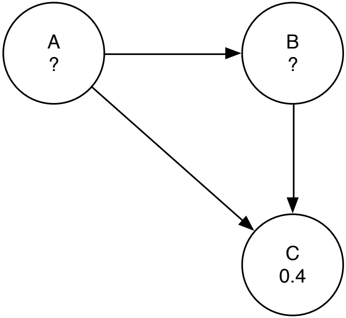

The Idea

- What if we want COH (If $a \to b$ then $P(a) \leq 1 - P(b)$)?

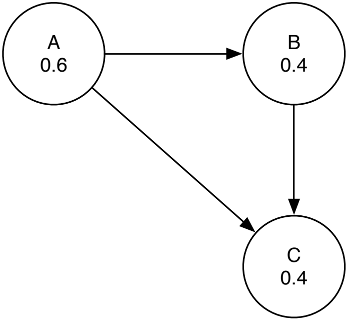

The Idea

- What if we want COH (If $a \rightarrow b$ then $P(a) \leq 1 - P(b)$)?

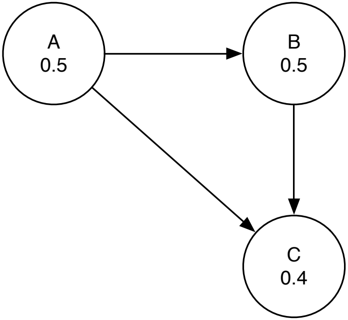

- Multiple probability functions can satisfy the coherence here.

Selecting Probability Functions

- Maximise entropy ($-\sum_{E \subseteq Arg} P(E) log P(E)$).
- This requires that the sets be
  - convex (if $x_1, x_2 \in X$ then $\delta x_1 + (1 - \delta)x_2 \in X$ for $\delta \in [0, 1]$).
  - closed (for any converging sequence $x_1, x_2, \ldots$ where $x_i \in X$, $lim_{i \rightarrow \infty} x_i \in X$)
- All sets satisfying properties (except for P_{RAT}) are convex and closed.

P_{RAT}

- Consider $a \rightarrow b$
- $P_1(a) = 0.5$, $P_1(b) = 0.4$
- $P_2(a) = 0.4$, $P_2(b) = 0.8$
- What convex combination is not in P_{RAT} (RAT: If $a \rightarrow b$ then $P(a) = 1 - P(b)$)?

P_{RAT}

- Consider $a \rightarrow b$
- $P_1(a) = 0.5$, $P_1(b) = 0.4$
- $P_2(a) = 0.4$, $P_2(b) = 0.8$
- What convex combination is not in P_{RAT} (RAT: If $a \rightarrow b$ then $P(a) = 1 - P(b)$)?
- For $P = 0.5P_1 + 0.5P_2$ we have $P(a) = 0.7, P(b) = 0.6$

Open Questions

- What should one do if the coherence condition is not satisfied (e.g. P(b)=0.7, P(c)=0.6)?
- One selects probabilities "as close as possible" to ones that meet the condition, or perhaps selects probability functions that meet the condition "as close as possible" to the original assignment.
- How does one define "as close as possible"?

Applications

- Reasoning about uncertain knowledge
- Persuasion
- Opponent Modelling

The Problem

The wealthy should pay more tax.

The Problem

The Problem

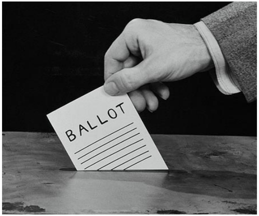

What should he say to be convincing?

The Problem

Why?

Problem Abstraction

- The speaker would like to see some action pursued.
- Using a variant of Atkinson's scheme for practical reasoning (without CQ10[1]), we can identify attacks between the speaker's universe of possible arguments.

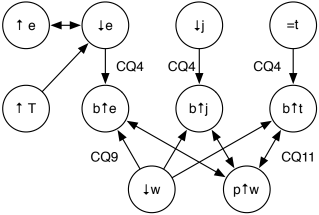

- Separate epistemic and practical critical questions; the former are resolved outside our framework through an appeal to truth.

[1] Which prevents attacks between identical actions

Problem Abstraction

- We model the system as a VAF $\langle X, A, \nu, \eta \rangle$
- Knowledge of individual listener (audience) α: $x_\alpha \subseteq X$
- Value ordering of α: $v_i \succ_\alpha v_j$

Problem Abstraction

- The speaker is aware of all arguments in the VAF.
- They can introduce a subset of them, which are added to all listener's knowledge bases.

Problem Abstraction

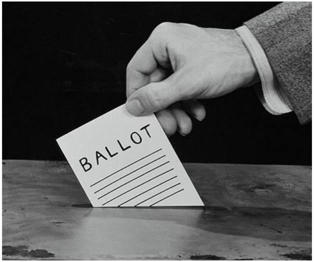

- Listeners compute the extension of their private knowledge base to identify a most preferred action.
- Social choice function aggregates individual preferred actions to global choice.

What should be said? (I)

- If there are
  1. No attacks between practical arguments for the same action (no CQ10);
  2. symmetric attacks among pairs of practical arguments for distinct actions; and
  3. No attacks between practical and epistemic arguments (vice-versa is ok).
- Then the speaker should
  - advance all practical arguments for the goal action; and
  - all epistemic arguments which do not directly/indirectly attack the advanced practical arguments.

Why?

- We relax the assumption of agreement on the value set (CQ16).
- We introduce a *threshold value partition* $T_\alpha : 2^\nu \to \{+, -\}$
- Normally
  1. $\forall v, v \in 2^\nu \setminus \emptyset, T_\alpha(v) = +$. All values should be acted on.
  2. If $T_\alpha(\{v_i\}) = +$ then for any $V \supset \{v_i\}$, $T_\alpha(V) = +$. Accruals of values should be acted on.
  3. $\forall v_1, v_2 \in \nu$, if $v_1 \succ_\alpha v_2$ and $T_\alpha(v_2) = +$ then $T_\alpha v_1 = +$. More preferred values than some minimally preferred value should be acted on.
- If these are not satisfied, then a threshold value partition is *degenerate*.

Dealing with degeneracy

- In the presence of degeneracy, an exhaustive evaluation of the effects of the introduction of a set of arguments is needed.
- We utilise a probabilistic model of listeners through PVAFs, an extension of PrAFs.
  - $P_X : L \times X \rightarrow [0, 1]$ likelihood of individual $i \in L$ knowing an argument.
  - $P_S : L \times S \rightarrow [0, 1]$ likelihood of individual $i \in L$ having some preference ordering.
  - $P_T : L \times S \times T$ likelihood of individual $i \in L$ being audience $\alpha \in S$ and having some threshold preference ordering.

Modelling the listeners

$$P_{VAF}(l, D) = \prod_{x \in X} P_X(l, x) \prod_{x \in X' \setminus X} (1 - P_X(l, x)) \quad (1)$$

$$P_{ext}(l, Y) = \sum_{\alpha \in P_S(l)} \sum_{i \in IV} P_{VAF}(l, i) P_S(l, \alpha) extension(Y, i, \alpha) \quad (2)$$

$$P_{exec}(l, a) = \sum_{t \in T} \sum_{e \in E} P_{ext}(l, e) P_T(l, t|s) \text{ s.t. } a \in e \text{ and } T_l(a) = + \quad (3)$$

$$P_{final}(L, a) = \sum_{l \in 2^L} VF(l) \prod_{l_i \in l} P_{exec}(l_i, a) \prod_{l_j \notin l} 1 - P_{exec}(l_j, a) \quad (4)$$

- $P_{final}(L, a)$ computes the likelihood of an action being selected by the entire group of listeners.
- The introduction of arguments X' by the speaker yields a new PVAF for which $P'_X(l, x) = 1$ if $x \in X'$ and $P_X(l, x)$ otherwise.
- Exponential complexity in the number of arguments.
- Instead, we can utilise Monte-Carlo sampling to obtain approximate solutions quickly.
- However, considering all possible argument combinations is still impractical.

Open Issues

- Intelligent possible argument search, e.g. through genetic algorithm search.
- How should CQ10, CQ16 be dealt with?
- Extend model to dialogue.
- More complex opponent modelling.

Probabilistic Opponent Modelling (Rienstra et al, IJCAI-13)

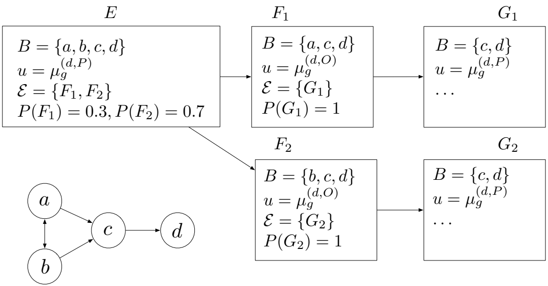

- We can reason about different depths of opponent models - "I believe that you believe that I believe . . . ".
- Better to reason about the likelihood of an opponent knowing an argument than the likelihood of an opponent's entire model. Does opponent agree with a specific argument?

Probabilistic Opponent Modelling

$maxEU = u(\pi)$
$bestMoves = \emptyset$
for all $M \in \texttt{legalmoves}(\pi)$ **do**
    $eu = 0$
    for all $E \in \mathcal{E}$ **do**
        $(oUtil, oMoves) = (\texttt{move}_u((\pi, M), E))$
        for all $M' \in oMoves$ **do**
            $(nUtil, nMoves) = \texttt{move}_u((\pi, M, M'), (B, u, \mathcal{E}, P))$
            $eu = eu + nUtil * P(E') * \frac{1}{|oMoves|}$
        end for
    end for
    if $eu > maxEU$ **then**
        $bestMoves = \emptyset$
    end if
    if $eu \geq maxEU$ **then**
        $bestMoves = bestMoves \cup \{M\}$
        $maxEU = eu$
    end if
end for

Arguing about uncertainty

- We've focused on arguing *with* probabilities - what conclusions can we draw from uncertain information.
- We have already seen that there are sometimes several ways of associating probabilities with arguments; how do we pick the right set?
  - Information from a UAV indicates the presence of hostile forces with a likelihood of XX%
  - Information from ground forces indicates hostile forces with a likelihood of YY%
  - What is the likelihood of hostile forces?
- Bayes' Rule; averaging; weighted averaging; maximum entropy; …

Argument Schemes about unertainty (Tang et al. ArgMAS-13

- Tang's work is rooted in Dempster-Shafer theory.
- An argument is a conclusion with evidence for it

$$A_1 = \langle p, \{\neg p : 0.3, p \wedge q : 0.6, q : 0.1\}\rangle$$

- We can compute *belief*, *disbelief* and *uncertainty* values:

$$b(p) = m(p \wedge q) = 0.6$$
$$d(p) = m(\neg p) = 0.3$$
$$u(p) = m(q) = 0.1$$

Argument Schemes about unertainty (Tang et al. ArgMAS-13

- Tang's work is rooted in Dempster-Shafer theory.
- An argument is a conclusion with evidence for it

$$A_2 = \langle p \rightarrow q, \{\neg p \wedge q : 0.5, \neg q : 0.3, \neg p : 0.2\}\rangle$$

- We can compute *belief*, *disbelief* and *uncertainty* values:

$$b(p \rightarrow q) = m(\neg p \wedge q) = 0.5$$
$$d(p \rightarrow q) = 0$$
$$u(p \rightarrow q) = m(p) + m(\neg q) = 0.5$$

Argument Schemes

- We can view rules of combination as argument schemes, taking in evidence arguments
- And outputting another evidence argument.
- Critical questions lead to the acceptance or rejection of applying the scheme.

Example: Dempster's rule of combination

- Combination function:

$$m(E_1 \otimes^D E_2, P) = \frac{\Sigma_{P=B\wedge C} m(E_1, B) m(E_2, C)}{1 - \Sigma_{B\wedge C=\perp} m(E_1, B) m(E_2, C)}$$

- Critical questions
  1. Is the evidence for consonant, consistent or arbitrary focal subsets of the frame of discernment?
  2. Is each piece of evidence equally reliable?
  3. Is each piece of evidence independent?
  4. *Should conflict between evidence be ignored in the mass assignments that result from combination?*
  5. Is there a restricted stochastic process behind the evidence which can be exploited to obtain a more accurate combination?
  6. Is the evidence all informative (i.e. no E_i in the premises contains the totally uncertain focal element Ω with $m(E_i, \Omega) > 0$) ?

Scheme example 2: Mixing-or-average rule

- Rule pattern:

$$\delta = \frac{\langle h_1, E_1 \rangle, ..., \langle h_m, E_m \rangle}{\langle h, E \rangle}$$

- Combination function:

$$m(E_1 \otimes^M ... \otimes^M E_n, P) = \frac{1}{n}\Sigma_{i=1}^{n} w_i \cdot m(E_i, P)$$

  where $P \in E_1 \cup E_2 \cup ... \cup E_n$ and w_i is the weight assigned to reliability of the corresponding piece of evidence.
- Critical questions
  1. Do the assigned weights not reflect the nature of the input evidence?
  2. Is there a restricted stochastic process behind the evidence which can be exploited to obtain a more accurate combination?

Argumentation driven by the Dempster-Shafer schemes

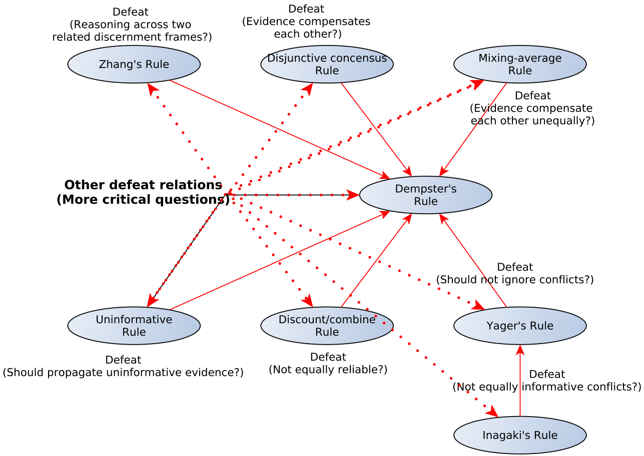

What have we done?

- Reasoned about the likelihood of arguments in the presence of uncertainty
  - Uncertainty about knowledge of argument; uncertainty about justification status: Pr(E)AF.
  - Uncertainty about belief (in base validity) of an argument: Epistemic extensions.
  - Via epistemic extensions, we can constrain what “reasonable” beliefs are.
- Maximised the likelihood of persuasion of others within a dialogue.
- Reasoned about how arguments should be combined in the presence of uncertainty.

Some Open Issues

- What is the relationship between PrAFs and the Epistemic Extension interpretation?
  - What happens if we feed the output of the former into the latter?
  - Can we obtain constraints over “reasonable” PrAF likelihoods?
- How to combine uncertainty with imprecise/fuzzy information (Dondio)
- Older approaches take a BN type approach, with arguments weakening each other. To move to a Dung type framework, better definitions of conflict between uncertain information is required ($p(a) = 0$ conflicts with $p(a) = 1$); does $p(a) = 0.3999$ conflict with $P(a) = 0.4$).

Acknowledgements

I like to thank Anthony Hunter for his invaluable comments regarding earlier versions of these slides.