

Modeling Partially Reliable Information Sources”

Rolf Haenni



Reasoning under **UN**certainty Group
Institute of Computer Science and Applied Mathematics
University of Berne, Switzerland

Formal Epistemology Workshop, Austin TX

May 25-29, 2005

Outline

- 1 Introduction
- 2 Probabilistic Argumentation
- 3 Modeling Partially Reliable Information Sources
- 4 General Solution
- 5 Case Studies
- 6 Conclusion

Outline

- 1 Introduction
- 2 Probabilistic Argumentation
- 3 Modeling Partially Reliable Information Sources
- 4 General Solution
- 5 Case Studies
- 6 Conclusion

Introduction

The Problem of Partially Reliable Information Sources

- YOU want to know whether a *hypothesis* H is true or false
- YOU get *reports* about H from different information sources
- The information sources are not fully reliable
- Problem: combine the reports, i.e.
 - ▶ make YOUR own judgment/opinion
 - ▶ possibly revise YOUR initial judgment/opinion
- Examples:
 - ▶ testimonial reports in court
 - ▶ review reports on accepting/rejecting a paper
 - ▶ expert opinions
 - ▶ sensor signals

Introduction

Basic Assumptions

- Only *two* alternative hypotheses H and $\neg H$
 - ▶ positive reports Rep_i
 - ▶ negative reports $\neg Rep_i$
- All sources are *structurally identical*, i.e.
 - ▶ represented by the same model
 - ▶ characterized by the same set of numerical parameters
- The sources are *conditionally independent* (given the hypothesis)
- $N = n + m$ is the total number of reports
 - ▶ n positive reports
 - ▶ m negative reports

Introduction

The Role of the Model

- The *true* behavior of the sources is unknown
- A *model* represents YOUR knowledge about the sources
- Possible valid models:
 - ▶ to know nothing
 - ▶ to know the true behavior
 - ▶ anything in between
- In general, a model is *incomplete*, *imprecise*, and/or *incorrect*
- The validity of the resulting overall judgement is always with respect to the correctness of the respective model
- Choosing an appropriate model is crucial

Outline

- 1 Introduction
- 2 Probabilistic Argumentation**
- 3 Modeling Partially Reliable Information Sources
- 4 General Solution
- 5 Case Studies
- 6 Conclusion

Probabilistic Argumentation

Basic Idea

- Probabilistic argumentation means to
 - ▶ find *arguments* and *counter-arguments* for a hypothesis
 - ▶ determine respective *probabilistic weights*
- This yields two non-additive measures:
 - ▶ *Degree of Support*: $dsp(H) \in [0, 1]$
 - ▶ *Degree of Possibility*: $dps(H) = 1 - dsp(\neg H) \in [0, 1]$
- The spirit is similar to:
 - ▶ Dempster-Shafer Theory of evidence [Sha76]
 - ▶ Walley's Imprecise Probabilities [Wal91]
 - ▶ Jøsang's Subjective Logic [Jøs01]
 - ▶ Ruspini's Theory of Evidential Reasoning [Rus86]

Probabilistic Argumentation

Basic Idea (cont.)

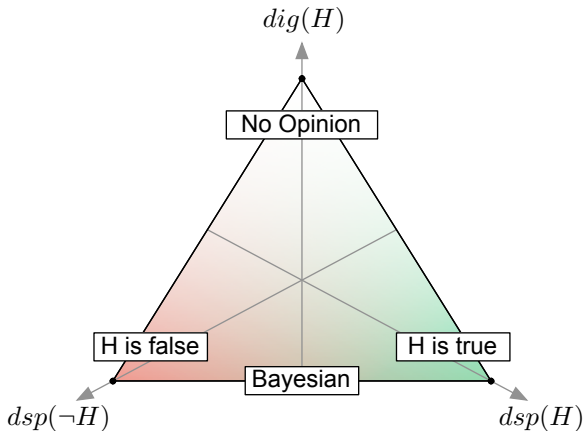
- The gap between $dps(H)$ and $dsp(H)$ measures the amount of available (missing) knowledge
- *Degree of Ignorance*: $dig(H) = dps(H) - dsp(H)$
- An *opinion* is a triple $\omega_H = (dsp(H), dsp(\neg H), dig(H))$
- Example: $\omega_H = (0.5, 0.3, 0.2)$



- Extreme opinions:
 - ▶ No opinion: $\omega_H = (0, 0, 1)$
 - ▶ Absolute opinions: $\omega_H = (1, 0, 0)$, $\omega_H = (0, 1, 0)$
 - ▶ Bayesian opinion: $\omega_H = (p, 1 - p, 0)$


Probabilistic Argumentation

Basic Idea (cont.)



Probabilistic Argumentation

Basic Idea (cont.)



Probabilistic Argumentation

Formal Model

- Suppose to have two sets of variables V and $W \subseteq V$
- Let \mathcal{L}_V be a formal language over V
- The *knowledge base* is a set $\Sigma \subseteq \mathcal{L}_V$ of sentences
- The elements of W are called *probabilistic variables*
- The elements $\mathbf{x} \in \Theta_W$ are *probabilistic states* (scenarios)
- $\mathbf{p}_W = \mathbf{p}(\Theta_W)$ denotes a prior probability distribution over W

Definition (Probabilistic Argumentation System)

$$\mathcal{A} = (V, \mathcal{L}_V, \Sigma, W, \mathbf{p}_W)$$

Probabilistic Argumentation

Formal Model (cont.)

- Subsets $E \subseteq \Theta_W$ are *probabilistic events*: $P(E) = \sum_{\mathbf{x} \in E} p(\mathbf{x})$
- Arguments for H : $Args(H) = \{\mathbf{x} \in \Theta_W : \Sigma_{\mathbf{x}} \models H\}$
- Conflicts: $Args(\perp) = \{\mathbf{x} \in \Theta_W : \Sigma_{\mathbf{x}} \models \perp\}$

Definition (Degree of Support/Possibility)

$$dsp(H) = P(Args(H) | Args(\perp)^c) = \frac{P(Args(H)) - P(Args(\perp))}{1 - P(Args(\perp))}$$

$$dps(H) = 1 - dsp(H)$$

Probabilistic Argumentation


Example

- Barbara promises to organize a barbecue (B) in case of a sunny day (s)
- The probability of sunshine is 60%
- Nothing is known about possible bad-weather alternatives
- $V = \{B, s\}$, $W = \{s\}$, $p(s) = 0.6$, $\Sigma = \{s \rightarrow B\}$
- Arguments: $Args(B) = \{s\}$
- Counter-Arguments: $Args(\neg B) = \{\}$
- Conflicts: $Args(\perp) = \{\}$
- $dsp(B) = 0.6$, $dsp(\neg B) = 0$, $dps(B) = 1$, $dig(B) = 0.4$
- $\omega_B = (0.6, 0, 0.4)$

Probabilistic Argumentation

Remarks

- Degree of support is the probability of the event that H is a logical consequence of our knowledge Σ
 - ⇒ *probability of provability*
 - ⇒ *epistemological probability*
- Special cases:
 - ▶ Logical reasoning: $W = \emptyset$
 - ▶ Probabilistic reasoning: $W = V$
- Probabilistic reasoning unifies logic and probability theory



Outline

- 1 Introduction
- 2 Probabilistic Argumentation
- 3 Modeling Partially Reliable Information Sources**
- 4 General Solution
- 5 Case Studies
- 6 Conclusion

Modeling Partially Reliable Information Sources

Example 1

- Nothing is known about the sources
- n positive, m negative reports
- no prior information about H
- $V = \{H, Rep_1, \dots, Rep_N\}$, $W = \emptyset$
- $\Sigma = \{Rep_1, \dots, Rep_n, \neg Rep_{n+1}, \dots, \neg Rep_N\}$
- $dsp(H) = 0$, $dps(H) = 1$, $dig(H) = 1$
- $\omega_H = (0, 0, 1)$

Modeling Partially Reliable Information Sources

Example 2

- All reports are generated at random with $P(r_i) = 0.5$
- $(r_i \rightarrow Rep_i) \wedge (\neg r_i \rightarrow \neg Rep_i) \equiv r_i \leftrightarrow Rep_i$
- no prior information about H
- $V = \{H, Rep_1, \dots, Rep_N, r_1, \dots, r_N\}$
- $W = \{r_1, \dots, r_N\}, P(r_i) = 0.5$
- $\Sigma = \left\{ \begin{array}{l} r_1 \leftrightarrow Rep_1, \dots, r_N \leftrightarrow Rep_N, \\ Rep_1, \dots, Rep_n, \neg Rep_{n+1}, \dots, \neg Rep_N \end{array} \right\}$
- $dsp(H) = 0, dps(H) = 1, dig(H) = 1$
- $\omega_H = (0, 0, 1)$

Modeling Partially Reliable Information Sources

Example 3

- All sources are aware of the true state of H
- Reliable sources tell the truth, unreliable sources lie:
 $[rel_i \rightarrow (H \leftrightarrow Rep_i)] \wedge [\neg rel_i \rightarrow (H \leftrightarrow \neg Rep_i)]$
- 90% of the sources are reliable: $P(rel_i) = 0.9$
- no prior information about H
- $V = \{H, Rep_1, \dots, Rep_N, rel_1, \dots, rel_N\}$
- $W = \{rel_1, \dots, rel_N\}, P(rel_i) = 0.9$
- $\Sigma = \left\{ \begin{array}{l} rel_1 \leftrightarrow (H \leftrightarrow Rep_1), \dots, rel_N \leftrightarrow (H \leftrightarrow Rep_N), \\ Rep_1, \dots, Rep_n, \neg Rep_{n+1}, \dots, \neg Rep_N \end{array} \right\}$
- $dsp(H) = dps(H) = \frac{0.9^{n-m}}{0.9^{n-m} + 0.1^{n-m}} = \frac{1}{1 + \frac{1}{9^{n-m}}}, dig(H) = 0$


Modeling Partially Reliable Information Sources

Complete and Incomplete Models

- $W_i \subseteq W$ is the set of probabilistic variables relevant to source i
- $V_i = W_i \cup \{H\}$ is the set of all variables relevant to source i
- *Complete model*:
⇒ report Rep_i is unambiguously determined for all $\mathbf{x} \in \Theta_{V_i}$
- $s = |V_i|$
- Total number of possible models: $m(s) = 3^{2^s}$
- Total number of complete models: $c(s) = 2^{2^s}$
- Total number of incomplete models: $i(s) = m(s) - c(s)$
- Many models are equivalent up to symmetry

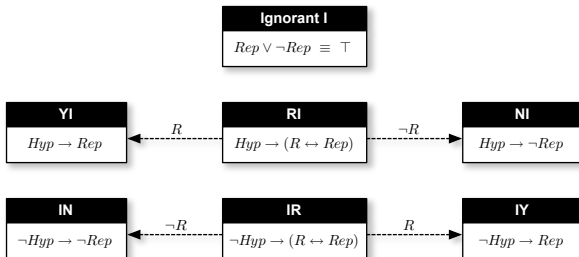
Modeling Partially Reliable Information Sources

Taxonomy of Complete Models



Modeling Partially Reliable Information Sources

Taxonomy of Incomplete Models



Modeling Partially Reliable Information Sources

Reliability-Based Models

- All sources are aware of the true state of H
- *Reliable* sources tell the truth (behave perfectly):
 $rel_i \rightarrow (H \leftrightarrow Rep_i)$
- *Unreliable* sources may do anything else: $\neg rel_i \rightarrow (M)$
- (M) denotes any complete or incomplete model except (P)

$$\Sigma = \left\{ \begin{array}{l} rel_1 \rightarrow (H \leftrightarrow Rep_1), \dots, rel_N \leftrightarrow (H \leftrightarrow Rep_N), \\ \neg rel_1 \rightarrow (M), \dots, \neg rel_N \rightarrow (M), \\ Rep_1, \dots, Rep_n, \neg Rep_{n+1}, \dots, \neg Rep_N \end{array} \right\}$$
- The probability of source i being reliable is $P(rel_i) = \rho$
- New models: (PR), (PD), (PI), ...

Outline

- 1 Introduction
- 2 Probabilistic Argumentation
- 3 Modeling Partially Reliable Information Sources
- 4 General Solution**
- 5 Case Studies
- 6 Conclusion

General Solution

Combining Opinions

- Suppose YOU get a single report from source i
 $\Rightarrow \omega_H^i = (dsp_i(H), dsp_i(\neg H), dig_i(H))$
- Suppose ω_H^0 is YOUR *initial opinion* (prior knowledge)
- Combined opinion: $\omega_H = \omega_H^0 \otimes \omega_H^1 \otimes \dots \otimes \omega_H^N$
- \otimes = Dempster's rule of combination
- *Commonality* function:
 - $x_i := dsp_i(H) + dig_i(H)$
 - $y_i := dsp_i(\neg H) + dig_i(H)$
 - $z_i := dig_i(H)$
- (x_i, y_i, z_i) represents the information obtained from source i

General Solution

Theorem (General Solution)

$$dsp(H) = \frac{\prod_{i=0}^N x_i - \prod_{i=0}^N z_i}{\prod_{i=0}^N x_i + \prod_{i=0}^N y_i - \prod_{i=0}^N z_i}$$

$$dps(H) = \frac{\prod_{i=0}^N x_i}{\prod_{i=0}^N x_i + \prod_{i=0}^N y_i - \prod_{i=0}^N z_i}$$

General Solution

Identical Sources

- Suppose all source are structurally identical
 - ⇒ Initial opinion: $(X, Y, Z) = (x_0, y_0, z_0)$
 - ⇒ Positive reports: $(X_1, Y_1, Z_1) = (x_i, y_i, z_i), i = 1 \dots n$
 - ⇒ Negative reports: $(X_2, Y_2, Z_2) = (x_i, y_i, z_i), i = n+1 \dots N$

Theorem (Identical Sources)

$$dsp(H) = \frac{XX_1^n X_2^m - ZZ_1^n Z_2^m}{XX_1^n X_2^m + YY_1^n Y_2^m - ZZ_1^n Z_2^m}$$

$$dps(hyp) = \frac{XX_1^n X_2^m}{XX_1^n X_2^m + YY_1^n Y_2^m - ZZ_1^n Z_2^m}$$

General Solution

Special Cases of Prior Knowledge

- Suppose a prior distribution $h = P(H)$ is given
 - $\Rightarrow X = h$
 - $\Rightarrow Y = 1 - h = \bar{h}$
 - $\Rightarrow Z = 0$

Theorem (Prior Distribution)

$$dsp(H) = dps(\{H\}) = \frac{h}{h + \bar{h} \left(\frac{Y_1}{X_1}\right)^n \left(\frac{Y_2}{X_2}\right)^m}$$

General Solution

Special Cases of Prior Knowledge (cont.)

- Suppose no prior knowledge is given
 - $\Rightarrow X = 1$
 - $\Rightarrow Y = 1$
 - $\Rightarrow Z = 1$

Theorem (No Prior Knowledge)

$$dsp(H) = \frac{X_1^n X_2^m - Z_1^n Z_2^m}{X_1^n X_2^m + Y_1^n Y_2^m - Z_1^n Z_2^m}$$

$$dps(hyp) = \frac{X_1^n X_2^m}{X_1^n X_2^m + Y_1^n Y_2^m - Z_1^n Z_2^m}$$

General Solution

Model	Positive Report			Negative Report		
	X_1	Y_1	Z_1	X_2	Y_2	Z_2
(Y)	1	1	1	0	0	0
(N)	0	0	0	1	1	1
(P)	1	0	0	0	1	0
(A)	0	1	0	1	0	0
(R)	r	r	r	\bar{r}	\bar{r}	\bar{r}
(YR)	1	r	r	0	\bar{r}	0
(RY)	r	1	r	\bar{r}	0	0
(NR)	0	r	0	1	\bar{r}	\bar{r}
(RN)	r	0	0	\bar{r}	1	\bar{r}
(L)	r	\bar{r}	0	\bar{r}	r	0
(D)	p	q	pq	\bar{p}	\bar{q}	$\bar{p}\bar{q}$

General Solution

Model	Positive Report			Negative Report		
	X_1	Y_1	Z_1	X_2	Y_2	Z_2
(I)	1	1	1	1	1	1
(YI)	1	1	1	0	1	0
(IY)	1	1	1	1	0	0
(NI)	0	1	0	1	1	1
(IN)	1	0	0	1	1	1
(RI)	r	1	r	\bar{r}	1	\bar{r}
(IR)	1	r	r	1	\bar{r}	\bar{r}
(PR)	$1 - \bar{\rho}\bar{r}$	$\bar{\rho}r$	$\bar{\rho}r$	$\bar{\rho}\bar{r}$	$1 - \bar{\rho}r$	$\bar{\rho}\bar{r}$
(PD)	$1 - \bar{\rho}\bar{p}$	$\bar{\rho}q$	$\bar{\rho}pq$	$\bar{\rho}\bar{p}$	$1 - \bar{\rho}q$	$\bar{\rho}\bar{p}\bar{q}$
(PI)	1	$\bar{\rho}$	$\bar{\rho}$	$\bar{\rho}$	1	$\bar{\rho}$

Outline

- 1 Introduction
- 2 Probabilistic Argumentation
- 3 Modeling Partially Reliable Information Sources
- 4 General Solution
- 5 Case Studies**
- 6 Conclusion

Case Studies

The Model (L)

Liar L
$Hyp \leftrightarrow (R \leftrightarrow Rep)$

$$r = Pr(R)$$

- No prior knowledge: $dsp(H) = dps(H) = \frac{1}{1 + (\frac{r}{1-r})^{n-m}}$
 - ▶ Laplace's formula for $m = 0$
 - ▶ Condorcet's Jury Theorem for $n \rightarrow \infty$ and m fixed
- Given prior knowledge: $dsp(H) = dps(H) = \frac{h}{h + \bar{h}(\frac{r}{1-r})^{n-m}}$
 - ▶ Boole's formula for $m = 0$

Case Studies

The Model (D)

Indicator D

$$Hyp \rightarrow (P \leftrightarrow Rep)$$

$$\neg Hyp \rightarrow (Q \leftrightarrow Rep)$$

$$p = Pr(P)$$

$$q = Pr(Q)$$

- No prior knowledge: $dsp(H) = \frac{1 - q^n \bar{q}^m}{1 + \left(\frac{p}{q}\right)^n \left(\frac{\bar{p}}{\bar{q}}\right)^m - q^n \bar{q}^m}$

$$dps(H) = \frac{1}{1 + \left(\frac{p}{q}\right)^n \left(\frac{\bar{p}}{\bar{q}}\right)^m - q^n \bar{q}^m}$$

- Given prior knowledge: $dsp(H) = dps(H) = \frac{h}{h + \bar{h} \left(\frac{q}{p}\right)^n \left(\frac{\bar{q}}{\bar{p}}\right)^m}$
 \Rightarrow 1st Bayesian Network in the paper (e.g. see [BH03])

Outline

- 1 Introduction
- 2 Probabilistic Argumentation
- 3 Modeling Partially Reliable Information Sources
- 4 General Solution
- 5 Case Studies
- 6 Conclusion**

Conclusion

- The general model of partially reliable sources allows to reproduce a number of previously known results
- It explains the relationship between different approaches
- It clarifies the role of prior knowledge
- It demonstrates the elegance of probabilistic argumentation and its generality as a unified theory of logical and probabilistic reasoning
- An important open question is how to model possible dependencies between the sources

Outline

- 1 Introduction
- 2 Probabilistic Argumentation
- 3 Modeling Partially Reliable Information Sources
- 4 General Solution
- 5 Case Studies
- 6 Conclusion

References



L. Bovens and S. Hartmann.

Bayesian Epistemology.

Oxford University Press, 2003.



A. Jøsang.

A logic for uncertain probabilities.

International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 9(3):279–311, 2001.



E. H. Ruspini.

The logical foundations of evidential reasoning.

Technical Report 408, SRI International, AI Center, Menlo Park, USA, 1986.



G. Shafer.

The Mathematical Theory of Evidence.

Princeton University Press, 1976.



P. Walley.

Statistical Reasoning with Imprecise Probabilities.

Monographs on Statistics and Applied Probability 42. Chapman and Hall, London, UK, 1991.