

Incidence Calculus

Alan Bundy

DAI Research Paper No.

April 26, 2004

Commissioned by the Encyclopedia of Artificial
Intelligence

Department of Artificial Intelligence
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN
Scotland

© Alan Bundy

Incidence Calculus *

Alan Bundy

Abstract

We describe *incidence calculus*, a logic for probabilistic reasoning. In incidence calculus, probabilities are not directly associated with formulae. Rather sets of possible worlds are directly associated with formulae and probabilities are calculated from these. This enables incidence calculus to be *truth functional*, which a logic based on a purely numeric uncertainty measure cannot be. This, in turn, enables tighter probability intervals to be calculated for theorems of an incidence calculus theory than is possible in a purely numeric uncertainty theory.

1 Introduction

The *incidence calculus* is a logic for probabilistic reasoning. Given upper and lower bounds on the probabilities of the axioms of a logical theory it formalises the derivation of upper and lower bounds on the remaining formulae of the theory. It differs from other logics for reasoning under uncertainty in two respects:

- Probabilities are not directly associated with any formulae. Rather *incidences* are directly associated with some formulae and probabilities are calculated from these incidences. An *incidence* is a set of possible worlds, each with an associated probability. The intended meaning of the incidence of a formula is the set of possible worlds in which the formula is true.
- This indirect encoding enables the incidence calculus to be *truth functional*, that is the incidence of a compound formula can be calculated directly from its parts. Any logic in which probabilities, or any numeric uncertainty values, are associated directly with formulae cannot be truth functional, [Bundy 85]. Truth functionality is an important property of any logic since it enables the calculation of tight upper and lower bounds on the uncertainty values of formulae.

2 A Simple Example

To see how this works in practice, consider the following simple example.

Let there be two propositions, *rainy* and *windy*, and seven possible worlds, *sun*, *mon*, *tues*, *wed*, *thurs*, *fri*, *sat*. Let each possible world be equally probable, *i.e.* occur $\frac{1}{7}$ of the time.

Let $i(\phi)$ denote the incidence of the formula ϕ and let the incidences of our two propositions be:

$$\begin{aligned}i(\textit{rainy}) &= \{\textit{fri}, \textit{sat}, \textit{sun}, \textit{mon}\} \\i(\textit{windy}) &= \{\textit{mon}, \textit{wed}, \textit{fri}\}\end{aligned}$$

*The research reported in this paper was supported by SERC/Alvey grant number GR/C/20826 and an SERC Senior Fellowship to the author. I thank Flavio Correa da Silva and Joe Halpern for recent conversations about the incidence calculus.

From this we can calculate the incidences of $\neg\text{rainy}$, $\text{rainy} \wedge \neg\text{rainy}$ and $\text{rainy} \wedge \text{windy}$, which are:

$$\begin{aligned}
i(\neg\text{rainy}) &= \{\text{mon}, \text{tues}, \text{wed}, \text{thurs}, \text{fri}, \text{sat}, \text{sun}\} \setminus \{\text{fri}, \text{sat}, \text{sun}, \text{mon}\} \\
&= \{\text{tues}, \text{wed}, \text{thurs}\} \\
i(\text{rainy} \wedge \neg\text{rainy}) &= \{\text{fri}, \text{sat}, \text{sun}, \text{mon}\} \cap \{\text{tues}, \text{wed}, \text{thurs}\} \\
&= \{\} \\
i(\text{rainy} \wedge \text{windy}) &= \{\text{fri}, \text{sat}, \text{sun}, \text{mon}\} \cap \{\text{mon}, \text{wed}, \text{fri}\} \\
&= \{\text{fri}, \text{mon}\}
\end{aligned}$$

Let $p(\phi)$ denote the probability of the formula ϕ . We can calculate the probability of each the above formulae by adding up the number of possible worlds in its incidence.

$$\begin{aligned}
p(\text{rainy}) &= \frac{4}{7} & p(\text{rainy} \wedge \neg\text{rainy}) &= 0 \\
p(\text{windy}) &= \frac{3}{7} & p(\text{rainy} \wedge \text{windy}) &= \frac{2}{7} \\
p(\neg\text{rainy}) &= \frac{3}{7}
\end{aligned}$$

Note that although windy and $\neg\text{rainy}$ have the same probability, $\frac{3}{7}$, their conjunctions with rainy have different probabilities, $\frac{2}{7}$ and 0, respectively. If the probability of compound formulae were calculated *solely* from the probabilities of their parts then the probabilities of these two conjunctions would have to be the same. This shows that a probabilistic calculus cannot be truth functional. This is because the *correlation* between two formulae cannot be encoded in their two probabilities, whereas it is encoded in their two incidences.

3 Formal Definitions

The incidence calculus is defined formally as follows.

Definition 1 *Propositional Language*

$\mathcal{L}(P)$ is the propositional language formed from P , where P is a finite set of propositions.

$\mathcal{L}(P)$ is the smallest set containing the truth values and the members of P and closed under the operations of negation, disjunction, conjunction and implication, *i.e.*

$$\begin{aligned}
&\text{true}, \text{false} \in \mathcal{L}(P), \\
&\text{if } p \in P \text{ then } p \in \mathcal{L}(P), \text{ and} \\
&\text{if } \phi, \psi \in \mathcal{L}(P) \text{ then } \neg\phi \in \mathcal{L}(P), \phi \vee \psi \in \mathcal{L}(P), \phi \wedge \psi \in \mathcal{L}(P) \text{ and} \\
&\phi \rightarrow \psi \in \mathcal{L}(P).
\end{aligned}$$

Definition 2 *Possible Worlds*

A *possible world* is a primitive, *i.e.* undefined, object of incidence calculus, but can be thought of as a partial interpretation of some logical formulae.

The *probability* that a possible world, w , is the real world is represented by a function, ϱ , from possible worlds to real numbers between 0 and 1.

If I is a set of possible worlds then $wp(I)$ is called the *weighted probability* of I , and is defined to be:

$$wp(I) = \sum_{w \in I} \varrho(w) \tag{1}$$

Definition 3 *Incidence Calculus*

An incidence calculus theory is a quintuple $\langle \mathcal{W}, \varrho, \mathcal{P}, \mathcal{A}, i \rangle$, where:

\mathcal{W} is a finite set of possible worlds.

For all $w \in \mathcal{W}$, $\varrho(w)$ is the probability of w and $wp(\mathcal{W}) = 1$.

\mathcal{P} is a set of propositions. $\mathcal{L}(\mathcal{P})$ is the language of the theory.

\mathcal{A} is a distinguished set of formulae in $\mathcal{L}(\mathcal{P})$ called the *axioms* of the theory.

i is a function from the axioms, \mathcal{A} , to $2^{\mathcal{W}}$, the set of subsets of \mathcal{W} . $i(\phi)$ is called the *incidence* of ϕ . $i(\phi)$ is to be thought of as the set of possible worlds in \mathcal{W} in which ϕ is true, *i.e.* $i(\phi) = \{w \in \mathcal{W} | w \models \phi\}$.

i is extended to a function from $\mathcal{L}(\mathcal{A})$ to $2^{\mathcal{W}}$ by the following *defining equations of incidence*.

$$i(\text{true}) = \mathcal{W} \quad (2)$$

$$i(\text{false}) = \{\} \quad (3)$$

$$i(\neg\phi) = \mathcal{W} \setminus i(\phi) \quad (4)$$

$$i(\phi \wedge \psi) = i(\phi) \cap i(\psi) \quad (5)$$

$$i(\phi \vee \psi) = i(\phi) \cup i(\psi) \quad (6)$$

$$i(\phi \rightarrow \psi) = (\mathcal{W} \setminus i(\phi)) \cup i(\psi) \quad (7)$$

Note that the equations (2) to (7) define i *truth functionally*, that is the incidence of a compound formula can be calculated solely from the incidences of its constituent sub-formulae.

In particular, it is easy to calculate that the incidence of any propositional tautology is \mathcal{W} and that of any propositional contradiction is $\{\}$.

It is not usually possible to infer the incidence of formulae in $\mathcal{L}(\mathcal{P}) \setminus \mathcal{L}(\mathcal{A})$. What we can do is define upper and lower bounds on the incidence using the functions i^* and i_* , respectively. For all $\phi \in \mathcal{L}(\mathcal{P})$ these are defined as follows.

$$i^*(\phi) = \bigcap_{\psi \in \mathcal{L}(\mathcal{A})} \{i(\psi) | i(\phi \rightarrow \psi) = \mathcal{W}\} \quad (8)$$

$$i_*(\phi) = \bigcup_{\psi \in \mathcal{L}(\mathcal{A})} \{i(\psi) | i(\psi \rightarrow \phi) = \mathcal{W}\} \quad (9)$$

It is easy to see that $i_*(\phi) = i(\phi) = i^*(\phi)$ for all $\phi \in \mathcal{L}(\mathcal{A})$ and $i_*(\phi) \subseteq i^*(\phi)$ for all $\phi \in \mathcal{L}(\mathcal{P})$

Definition 4 *Probability*

The *probability* that a formula, ϕ , is true is represented using the partial function p from formulae to real numbers in the interval 0 to 1. When $i(\phi)$ is defined, $p(\phi)$ is defined as:

$$p(\phi) = wp(i(\phi)) \quad (10)$$

For all $\phi \in \mathcal{L}(\mathcal{P})$ we define upper and lower bounds on the probability of ϕ using the functions p^* and p_* , respectively, as follows.

$$p^*(\phi) = wp(i^*(\phi))$$

$$p_*(\phi) = wp(i_*(\phi))$$

The *conditional probability* that a formula, ϕ , is true given that another formula ψ is true is represented using the partial binary function p from pairs of formulae to real numbers in the interval 0 to 1. For all $\phi, \psi \in \mathcal{L}(\mathcal{P})$ it is defined as:

$$p(\phi|\psi) = \frac{p(\phi \wedge \psi)}{p(\psi)} \quad (11)$$

The *correlation* between two formulae, ϕ and ψ is represented using the partial binary function c from pairs of formulae to real numbers in the interval -1 to 1. For all $\phi, \psi \in \mathcal{L}(\mathcal{P})$ it is defined as:

$$c(\phi, \psi) = \frac{p(\phi \wedge \psi) - p(\phi).p(\psi)}{\sqrt{p(\phi).p(\neg\phi).p(\psi).p(\neg\psi)}} \quad (12)$$

The correlation between two formulae is a measure of their degree of dependence. $c(\phi, \psi) = 1$ means ϕ and ψ always co-occur, $c(\phi, \neg\phi) = -1$ means they never co-occur and $c(\phi, \psi) = 0$ means they are independent, co-occurring in a random way. Note that $p(\neg\phi)$ and $p(\neg\psi)$ are needed to define $c(\phi, \psi)$; it cannot be defined in terms of $p(\phi)$ and $p(\psi)$ alone.

From the definitions of i , (2) to (7), wp , (1), p , (10), and c , (12), we can derive the following rules of probability.

$$\begin{aligned} p(\text{true}) &= 1 \\ p(\text{false}) &= 0 \\ p(\neg\phi) &= 1 - p(\phi) \\ p(\phi \vee \psi) &= p(\phi) + p(\psi) - p(\phi \wedge \psi) \\ p(\phi \rightarrow \psi) &= p(\neg\phi) + p(\psi) - p(\neg\phi \wedge \psi) \\ p(\phi \wedge \psi) &= p(\phi).p(\psi) + c(\phi, \psi).\sqrt{p(\phi).p(\neg\phi).p(\psi).p(\neg\psi)} \end{aligned} \quad (13)$$

Note that these equations do not define p truth functionally. In particular, the probability of $\phi \wedge \psi$ is given not just in terms of the probability of ϕ and ψ but also in terms of $c(\phi, \psi)$. This deficiency is not just an artifact of these equations, but is an endemic property of a probabilistic logic, indeed of any purely arithmetic uncertainty calculus. Similar properties hold for conditional probability, correlation and any combination. That is, it is not possible to calculate the conditional probability or correlation of a compound formula solely from the conditional probabilities, correlations and probabilities of its parts, [Bundy 85].

4 Inferring New Incidences from Old

Given an assignment of incidences to a set of axioms of a theory, \mathcal{A} , we want to derive the incidences of the remaining formulae of the language $\mathcal{L}(\mathcal{P})$. Unfortunately, it is generally only possible to derive upper and lower bounds for the incidences of these formulae. For instance, consider the *modus ponens* rule of inference:

$$\frac{\phi \rightarrow \psi, \quad \phi}{\psi}$$

It is not possible to calculate $i(\psi)$ from $i(\phi \rightarrow \psi)$ and $i(\phi)$ alone. Suppose $\mathcal{W} = \{a, b\}$, $i(\phi \rightarrow \psi) = \{a, b\}$ and $i(\phi) = \{a\}$ then all we can say about $i(\psi)$ is that it lies between $\{a\}$ and $\{a, b\}$. In fact, $i_*(\psi) = \{a\}$ and $i^*(\psi) = \{a, b\}$.

In general, therefore, the best we can hope for is an inference mechanism that, given the upper and lower bounds on the incidences of the hypotheses of an inference step, calculates the tightest upper and lower bounds of the incidence of the

conclusion. We have realised this as a process of constraint propagation. Given a finite set of formulae and some initial assignment of upper and lower bounds, new assignments are calculated from old assignments and the process iterates until it terminates. At termination the tightest bound have been calculated.

The set of formulae on which this constraint propagation is performed (the constraint set) must be defined with some care. For completeness it must be the language of: the axioms, the formulae whose incidences we are interested in calculating, and all sub-formulae of all these formulae. Unfortunately, this is an infinite set, but constraint propagation must be performed on a finite set if it is to terminate. The solution is to let each formula of the language be represented by a canonical form, where there are only a finite number of these canonical forms.

Definition 5 *Canonical Form*

A formula in $\mathcal{L}(P)$ is in canonical form if it has the form $\bigwedge_{i=1}^m (\neg \bigwedge_{j=1}^n l_i^j)$ where $l_i^j = p$ or $l_i^j = \neg p$ for some $p \in P$.

Note that for each $\phi \in \mathcal{L}(P)$ there is an equivalent formula in canonical form. We first transform ϕ into conjunctive normal form and then use de Morgan's law to turn each conjunct $\bigvee_{j=1}^m l_i^j$ into $\neg \bigwedge_{j=1}^m \neg l_i^j$ and then cancel all double negations.

Let $\mathcal{L}'(P)$ be the subset of formulae in $\mathcal{L}(P)$ that are in canonical form.

For readability we will sometimes abbreviate a formula in canonical form by one of its non-canonical equivalents, *e.g.* write $\phi \rightarrow \psi$ to stand for $\neg(\phi \wedge \neg\psi)$.

Definition 6 *Constraint Set*

Let \mathcal{A} be the set of axioms of a theory and \mathcal{T} be the set of formulae whose incidence we are interested in calculating (the 'theorems').

Let $sf(F)$ be the set of subformulae of the formulae F , *i.e.*

- if $\phi \in F$ then $\phi \in sf(F)$;
- if $\neg\phi \in sf(F)$ then $\phi \in sf(F)$;
- if $\phi \wedge \psi \in sf(F)$ then $\phi, \psi \in sf(F)$;
- if $\phi \vee \psi \in sf(F)$ then $\phi, \psi \in sf(F)$;
- if $\phi \rightarrow \psi \in sf(F)$ then $\phi, \psi \in sf(F)$.

The constraint set is $sf(\mathcal{L}'(\mathcal{A} \cup \mathcal{T}))$.

The constraint set is organised as a network in which each formula has pointers to each of its sub-formulae and super-formulae. The initial assignment consists of upper and lower bounds equal to the incidence for the axioms, and the default upper and lower bounds \mathcal{W} and $\{\}$, respectively, for all other formulae. The idea of the constraint propagation is to improve on these default assignments by replacing them with the values of i^* and i_* for each formula.

To describe this constraint propagation process we must first be more precise about the representation of an assignment of upper and lower bounds.

Definition 7 *Assignments*

Let F be an *assignment* of upper and lower incidence bounds to a set of formulae.

sup_F is a function from formulae to sets of possible worlds which defines the current assignment of upper bounds.

inf_F is a function from formulae to sets of possible worlds which defines the current assignment of lower bounds.

We can now give the rules of inference which propagate the lower and lower bounds around the network of formulae.

Definition 8 *The Rules of Inference*

A rule of inference is a mapping from assignments to assignments. Let F be the assignment before the rule fires and G be the assignment afterwards. In each case, G is the same as F except for the changes, on some particular formula, defined below for each rule.

$$\begin{array}{ll}
\text{Not1:} & sup_G(\phi) = (\mathcal{W} \setminus inf_F(\neg\phi)) \cap sup_F(\phi) \\
\text{Not2:} & inf_G(\phi) = (\mathcal{W} \setminus sup_F(\neg\phi)) \cup inf_F(\phi) \\
\text{Not3:} & sup_G(\neg\phi) = (\mathcal{W} \setminus inf_F(\phi)) \cap sup_F(\neg\phi) \\
\text{Not4:} & inf_G(\neg\phi) = (\mathcal{W} \setminus sup_F(\phi)) \cup inf_F(\neg\phi) \\
\text{And1:} & sup_G(\phi) = (sup_F(\phi \wedge \psi) \cup (\mathcal{W} \setminus inf_F(\psi))) \cap sup_F(\phi) \\
\text{And2:} & inf_G(\phi) = inf_F(\phi \wedge \psi) \cup inf_F(\phi) \\
\text{And3:} & sup_G(\psi) = (sup_F(\phi \wedge \psi) \cup (\mathcal{W} \setminus inf_F(\phi))) \cap sup_F(\psi) \\
\text{And4:} & inf_G(\psi) = inf_F(\phi \wedge \psi) \cup inf_F(\psi) \\
\text{And5:} & sup_G(\phi \wedge \psi) = sup_F(\phi) \cap sup_F(\psi) \cap sup_F(\phi \wedge \psi) \\
\text{And6:} & inf_G(\phi \wedge \psi) = (inf_F(\phi) \cap inf_F(\psi)) \cup inf_F(\phi \wedge \psi)
\end{array}$$

Note that it is only necessary to give rules for negation and conjunction because these rules only operate on formulae in canonical form.

The exhaustive application of these rules of inference will terminate, [Bundy 86]. Suppose the final assignment is F , then $sup_F(\phi) = i^*(\phi)$ and $inf_F(\phi) = i_*(\phi)$, [Correa daSilva & Bundy 91].

This inference mechanism can be quite efficiently implemented if sets of possible worlds are represented by bit vectors, with one bit for each possible world in \mathcal{W} . \cap , \cup and \setminus are then represented by logical and, or and not, respectively.

5 An Example of Inference

We can illustrate the constraint propagation process with the *modus ponens* example with which we started §4. That is: $\mathcal{A} = \{\phi \rightarrow \psi, \phi\}$, $\mathcal{T} = \{\psi\}$, $\mathcal{W} = \{a, b\}$, $i(\phi \rightarrow \psi) = \{a, b\}$ and $i(\phi) = \{a\}$. We want to calculate upper and lower bounds on $i(\psi)$.

We must first translate $\phi \rightarrow \psi$ into its logically equivalent, canonical form: $\neg(\phi \wedge \neg\psi)$. The initial assignment of upper and lower bounds is given and explained in table 1. We have omitted from the table irrelevant members of $sf(\mathcal{L}'(\mathcal{A} \cup \mathcal{T}))$.

Ignoring rule applications which have no effect on the assignment, the inference process proceeds as follows.

$$\begin{array}{ll}
sup_2(\phi \wedge \neg\psi) & = (\mathcal{W} \setminus inf_1(\neg(\phi \wedge \neg\psi))) \cap sup_1(\phi \wedge \neg\psi) & \text{by rule Not1} \\
& = (\{a, b\} \setminus \{a, b\}) \cap \{a, b\} \\
& = \{\} \\
sup_3(\neg\psi) & = (sup_2(\phi \wedge \neg\psi) \cup (\mathcal{W} \setminus inf_2(\phi))) \cap sup_2(\neg\psi) & \text{by rule And3} \\
& = (\{\} \cup (\{a, b\} \setminus \{a\})) \cap \{a, b\} \\
& = \{b\} \\
inf_4(\psi) & = (\mathcal{W} \setminus sup_3(\neg\psi)) \cup inf_3(\psi) & \text{by rule Not2} \\
& = (\{a, b\} \setminus \{b\}) \cup \{a\} \\
& = \{a\}
\end{array}$$

After this none of the rules can change the assignment of these formulae and the process terminates. The final assignment is given in table 2. As required, this gives $i_*(\psi) = \{a\}$ and $i^*(\psi) = \{a, b\}$.

Formula	inf_1	sup_1
ϕ	$\{a\}$	$\{a\}$
ψ	$\{\}$	$\{a, b\}$
$\neg\psi$	$\{\}$	$\{a, b\}$
$\phi \wedge \neg\psi$	$\{\}$	$\{a, b\}$
$\neg(\phi \wedge \neg\psi)$	$\{a, b\}$	$\{a, b\}$

We will number the consecutive assignments 1, 2, etc. The two axioms ϕ and $\neg(\phi \wedge \neg\psi)$ are assigned upper and lower bounds equal to their incidences. The other formulae are all the proper sub-formulae of $\neg(\phi \wedge \neg\psi)$. All non-axioms are assigned the default upper and lower bounds of $\{a, b\}$ and $\{\}$, respectively.

Table 1: Initial Assignment of Upper and Lower Bounds

Formula	inf_4	sup_4
ϕ	$\{a\}$	$\{a\}$
ψ	$\{a\}$ (Not2)	$\{a, b\}$
$\neg\psi$	$\{\}$	$\{b\}$ (And3)
$\phi \wedge \neg\psi$	$\{\}$	$\{\}$ (Not1)
$\neg(\phi \wedge \neg\psi)$	$\{a, b\}$	$\{a, b\}$

Those assignments that have been changed from table 1 have been labelled with the rule of inference that made the change.

Table 2: Final Assignment of Upper and Lower Bounds

6 Assigning Incidences to Formulae

The incidence calculus inference mechanism assumes that incidences have been assigned to the axioms. Unfortunately, an uncertainty inference problem may not always be presented in this form. Often the uncertainty of formulae will be presented initially as an assignment of probabilities to the axioms, with or without some record of the correlations between them. It is then necessary to assign incidences to the axioms which respect the assignment of probabilities and correlations.

In making this assignment there are several variables: the size of \mathcal{W} , the definition of ϱ and the definition of i . We will consider only one simple scheme. We will take \mathcal{W} to be fixed, *e.g.* 100 possible worlds. The greater this size the greater the sensitivity of the calculations, but the greater their computational expense. Low sensitivity is often enough in artificial intelligence applications. We will also take ϱ to give the same value for each possible world, *e.g.* $\frac{1}{100}$. The probability assignment must now be realised solely by the assignment of equi-probable possible worlds to axioms using the function i .

As one example of how this can be done we describe the Monte Carlo method due to [Corlett & Todd 85], interleaved with the inference mechanism in order to take account of the logical structure of the axioms. In this method formulae are assumed to be assigned probabilities and incidences are assigned which respect these probabilities. Correlations are assumed not to be available. Some correlations are forced by the logical structure, but otherwise axioms are assumed to be pairwise independent.

Definition 9 Monte Carlo Incidence Assignment Method

We assume that $p(\phi)$ is known for all $\phi \in \mathcal{A}$.

Recall that $p^*(\phi) = wp(i^*(\phi))$ and $p_*(\phi) = wp(i_*(\phi))$.

Define p' as: $p'(\phi) = \frac{p(\phi) - p_*(\phi)}{wp(i^*(\phi) \setminus i_*(\phi))} \quad (= 1 - \frac{p^*(\phi) - p(\phi)}{wp(i^*(\phi) \setminus i_*(\phi))})$.

Initialise $i_*(\phi) = \{\}$ and $i^*(\phi) = \mathcal{W}$ for all $\phi \in sf(\mathcal{L}'(\mathcal{A}))$.

For each $\phi \in \mathcal{A}$.

1. If $p(\phi) < p_*(\phi)$ or $p(\phi) > p^*(\phi)$ then stop with failure.
2. Else randomly assign each of the possible worlds in $i^*(\phi) \setminus i_*(\phi)$ to $i(\phi)$ with a probability of $p'(\phi)$.
Let $i^*(\phi) = i_*(\phi) = i(\phi)$.
3. Run the inference mechanism on the formulae in $sf(\mathcal{L}'(\mathcal{A}))$ with initial assignment $inf_1(\psi) = i_*(\psi)$ and $sup_1(\psi) = i^*(\psi)$ for all $\psi \in sf(\mathcal{L}'(\mathcal{A}))$. Suppose F is the assignment on termination.
Let $i_*(\psi) = inf_F(\psi)$ and $i^*(\psi) = sup_F(\psi)$ for all $\psi \in sf(\mathcal{L}'(\mathcal{A}))$.

For instance, let there be two propositions, *sunny* and *dry*, and ten, equi-probable, possible worlds named 1 to 10, *i.e.* $\mathcal{W} = \{1, \dots, 10\}$, $\varrho(w) = 0 \cdot 1$ for all $w \in \mathcal{W}$. Let the given assignment of probabilities be $p(\textit{sunny}) = 0 \cdot 4$, $p(\textit{sunny} \rightarrow \textit{dry}) = 0 \cdot 8$.

Suppose we choose to assign an incidence to *sunny* first. $i_*(\textit{sunny}) = \{\}$ and $i^*(\textit{sunny}) = \{1, \dots, 10\}$, so $p'(\textit{sunny}) = 0 \cdot 4$. Suppose the random assignment of possible worlds to $i(\textit{sunny})$ yields $\{1, 2, 4, 7\}$. When the inference mechanism is run it will reassign $\{3, 5, 6, 8, 9, 10\}$ to $i_*(\textit{sunny} \rightarrow \textit{dry})$.

Now we must assign an incidence to $\textit{sunny} \rightarrow \textit{dry}$. $i_*(\textit{sunny} \rightarrow \textit{dry}) = \{3, 5, 6, 8, 9, 10\}$ and $i^*(\textit{sunny} \rightarrow \textit{dry}) = \{1, \dots, 10\}$, so $p'(\textit{sunny} \rightarrow \textit{dry}) = 0 \cdot 5$.

The random assignment of possible worlds to $i(\text{sunny} \rightarrow \text{dry})$ might then yield $\{1, 2, 3, 5, 6, 8, 9, 10\}$.

As a side effect this process gives the following bounds on the probability of dry : $0 \cdot 2 \leq p(\text{dry}) \leq 1$.

Note that the correlation between the two axioms is $c(\text{sunny}, \text{sunny} \rightarrow \text{dry}) \simeq 0 \cdot 61$, so they are not independent.

With such a small \mathcal{W} this assignment is necessarily crude, but greater sensitivity can be gained by using larger \mathcal{W} .

7 Comparisons with Other Logics

Incidence calculus is similar to Nilsson's *probabilistic logic*, [Nilsson 84]. The main difference is that instead of taking the concept of possible world as a primitive, probabilistic logic uses Herbrand models of $\mathcal{L}(\mathcal{P})$. A Herbrand model is a model of a logical theory, in the sense of Tarski, which is formed using the symbols of the logic. In this way it is possible to form an exhaustive and disjoint set of models. The number of Herbrand models and which formulae are true in each one is fixed by this process. Thus the incidence of each formula is also known, in contrast with incidence calculus in which only upper and lower incidence bounds may be known for some formulae.

Given some desired assignment of probabilities to the formulae of \mathcal{A} , probabilistic logic calculates an assignment of probabilities to the Herbrand models which will yield the desired assignment to the formulae. In general, this is a prohibitively expensive process. Nilsson suggests various ways to circumvent it in special cases. Incidence calculus avoids this problem by using uninterpreted possible worlds rather than Herbrand models. This permits the cruder, but much less expensive probability assignment mechanism described in §6.

Incidence calculus is also similar to Dempster-Shafer theory, [Shafer 76], especially as formalised in [Fagin & Halpern 89]. Both systems permit only partial definition of the probabilities of some formulae. Dempster-Shafer theory achieves this by defining the incidence of all formulae, but not defining the probabilities of all the possible worlds, *i.e.* i is a total function, but ϱ is a partial function. Since p , the probability of a formula, is defined using both i and ϱ (see equations (1) and (10) above), it too is only partially defined. Incidence calculus achieves a similar effect the other way round, *i.e.* ϱ is total but i is partial. In most cases these two alternative solutions are equivalent, [Correa daSilva & Bundy 90].

However, most implementations of Dempster-Shafer theory work directly with the upper and lower probability bounds on formulae, and not via incidences. This means that their inference mechanism inevitably calculates looser bounds than the incidence calculus mechanism described in §4 above. This is because a probability calculus is not truth functional, whereas an incidence calculus is. The following example illustrates this.

Let $p(\text{rainy}) = \frac{4}{7}$ and $p(\text{windy}) = \frac{3}{7}$. If we try to calculate $p(\text{rainy} \wedge \text{windy})$ solely from the probabilities of its sub-formulae, then the best we can do is place it in a rather wide interval. If it is never both windy and rainy then $p(\text{rainy} \wedge \text{windy}) = 0$ and if it is always windy when it is rainy then $p(\text{rainy} \wedge \text{windy}) = \frac{3}{7}$. Every probability between these extremes is possible, thus:

$$0 \leq p(\text{rainy} \wedge \text{windy}) \leq \frac{3}{7}$$

Note that using equation (13) gives an even wider interval with upper bound $\frac{24}{49}$!

On the other hand, we saw in §2 that if the incidences of rainy and windy are known then we can calculate an exact value for the incidence of $\text{rainy} \wedge \text{windy}$ and

hence of $p(\text{rainy} \wedge \text{windy})$. This is because incidence calculus is truth functional, which no purely probabilistic calculus could be.

8 Conclusion

The incidence calculus is a logic for uncertain reasoning in which the uncertainty value is represented by the set of possible worlds in which it is true rather than by a numeric value. A numeric value, namely the probability, is readily recovered from this set. The advantage of this arrangement is that the incidence calculus is truth functional — leading to the calculation of tighter bounds on the probability of a formula than is possible with a calculus working with probabilities alone.

Incidence bounds on new formulae can be inferred from bounds on old formulae by a process of constraint propagation using rules of inference.

References

- [Bundy 85] Alan Bundy. Incidence calculus: A mechanism for probabilistic reasoning. *Journal of Automated Reasoning*, 1(3):263–284, 1985. Earlier version in Proceedings of FGCS-84 and in Proceedings of the Workshop on Uncertainty and Probability. Also available from Edinburgh as DAI Research Paper No 216.
- [Bundy 86] Alan Bundy. Correctness criteria of some algorithms for uncertain reasoning using incidence calculus. *Journal of Automated Reasoning*, 2(2):109–126, 1986. Also available from Edinburgh as DAI Research Paper No. 259.
- [Corlett & Todd 85] R. A. Corlett and S. J. Todd. A Monte-Carlo approach to uncertain inference. In P. Ross, editor, *Proceedings of AISB-85*, pages 28–34. Society for the Study of Artificial Intelligence and Simulation of Behaviour, 1985.
- [Correa daSilva & Bundy 90] F. S. Correa da Silva and A. Bundy. On some equivalence relations between incidence calculus and Dempster-Shafer theory of evidence. In P. P. Henrion, M. & Bonnissonne, editor, *Procs. of the Sixth Conference on Uncertainty in Artificial Intelligence*, pages 378–383. GE Corp., 1990. Also available from Edinburgh as DAI Research Paper 470.
- [Correa daSilva & Bundy 91] F. S. Correa da Silva and A. Bundy. A rational reconstruction of incidence calculus. Research Paper 517, Dept. of Artificial Intelligence, University of Edinburgh, 1991.
- [Fagin & Halpern 89] R. Fagin and J. Y. Halpern. Uncertainty, belief, and probability. In N. S. Sridharan, editor, *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1161–1167. Morgan Kaufmann, 1989. Also available as IBM Research Report RJ-6191.

[Nilsson 84]

N. Nilsson. Probabilistic logic. Technical Note 321, SRI International, 1984.

[Shafer 76]

G. Shafer. *A Mathematical Theory of Evidence*. Princeton Univ. Press, 1976.