

2

AD-A250 617



DTIC  
ELECTE  
MAY 28 1992  
S C D

# On Probability Distributions Over Possible Worlds

Fahiem Bacchus\*  
Department of Philosophy &  
Department of Computer Science  
University of Rochester  
Rochester, New York  
14627

March 8, 1988

## Abstract

In *Probabilistic Logic* Nilsson uses the device of a probability distribution over the set of possible worlds to assign probabilities to the sentences of a logical language. In his paper Nilsson concentrated on inference and associated computational issues. This paper, on the other hand, examines the probabilistic semantics in more detail, particularly for the case of first order languages, and attempts to explain some of the features and limitations of this form of probability logic. It is pointed out that the device of assigning probabilities to logical sentences has certain expressive limitations. In particular, statistical assertions are not easily expressed by such a device. This leads to certain difficulties with attempts to give probabilistic semantics to default reasoning using probabilities assigned to logical sentences.

\*This research was supported by a Post-Doctoral fellowship funded by the U.S. Army Signals Warfare Laboratory.

DISTRIBUTION STATEMENT A  
Approved for public release;  
Distribution Unlimited

92-13691

92 0 22 011

# REPORT DOCUMENTATION PAGE

Form Approved  
OPM No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Information and Regulatory Affairs, Office of Management and Budget, Washington, DC 20503.

|  |  |   |   |  |
|--|--|---|---|--|
| 1. AGENCY USE ONLY (Leave Blank)   |  | 2. REPORT DATE<br><br>1988                                  | 3. REPORT TYPE AND DATES COVERED<br><br>Unknown                   |  |
| 4. TITLE AND SUBTITLE<br><br>On Probability Distributions Over Possible Worlds   |  |   | 5. FUNDING NUMBERS<br><br>DAAB10-86-C-0567                        |  |
| 6. AUTHOR(S)<br><br>Fahiem Bacchus   |  |   |   |  |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)<br>University of Rochester<br>Department of Philosophy<br>Rochester, NY 14627   |  |   | 8. PERFORMING ORGANIZATION REPORT NUMBER                          |  |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)<br>U.S. Army CECOM Signals Warfare Directorate<br>Vint Hill Farms Station<br>Warrenton, VA 22186-5100  |  |   | 10. SPONSORING/MONITORING AGENCY REPORT NUMBER<br><br>92-TRF-0021 |  |
| 1. SUPPLEMENTARY NOTES   |  |   |   |  |
| 12a. DISTRIBUTION/AVAILABILITY STATEMENT<br><br>Statement A; Approved for public release; distribution unlimited.  |  |   | 12b. DISTRIBUTION CODE  |  |
| 13. ABSTRACT (Maximum 200 words) In Probabilistic logic Nilsson used the device of a probability distribution over the set of possible worlds to assign probabilities to the sentences of a logical language. In his paper Nilsson concentrated on inference and associated computational issues. This paper, on the other hand, examines the probabilistic semantics in more detail, particularly for the case of first order languages, and attempts to explain some of the features and limitation of this form of probability logic. It is pointed out that the device of assigning probabilities to logical sentences has certain expressive limitations. In particular, statistical assertions are not easily expressed by such a device. This leads to certain difficulties with attempts to give probabilistic semantics to default reasoning using probabilities assigned to logical sentences. |  |   |   |  |
| 14. SUBJECT TERMS<br><br>Artificial Intelligence, Data Fusion, Probability, Probability Distributions  |  |   | 15. NUMBER OF PAGES<br><br>11                                     |  |
|  |  |   | 16. PRICE CODE  |  |
| 17. SECURITY CLASSIFICATION OF REPORT<br><br>UNCLASSIFIED  | 18. SECURITY CLASSIFICATION OF THIS PAGE<br><br>UNCLASSIFIED | 19. SECURITY CLASSIFICATION OF ABSTRACT<br><br>UNCLASSIFIED | 20. LIMITATION OF ABSTRACT<br><br>UL                              |  |

# 1 Introduction

In [1] Nilsson describes a method of assigning probabilities to the sentences of a logic through a probability distribution over the set of possible worlds. Each possible world is a unique and consistent set of truth value assignments for all of the sentences in the logic. Although this approach is unproblematic when applied to propositional languages, certain difficulties arise when dealing with first order languages. By taking a different tack these difficulties can be overcome, and indeed, it has already been demonstrated that probabilities can be coherently assigned to the sentences of any first order language (Gaifman [2], Scott and Krauss [3]).

While the method of assigning probabilities to logical formulas is capable of representing probabilistic degrees of belief, it is incapable of effectively representing statistical assertions. It is argued that many types of defaults have a natural statistical interpretation, but cannot be represented by probabilities over logical formulas, because of this limitation. Some authors have attempted to represent defaults in exactly this way (Geffner and Pearl [4], Pearl [5]), and the difficulties with their systems can be demonstrated.

It is pointed out that although probabilities over logical formulas fails to do the job, statistical assertions can be efficiently represented in other types of probability logics, logics which go beyond the simple extension of first order logic offered by Nilsson's probabilistic logic.

# 2 The Propositional Case

A natural semantic model for a propositional language is simply a subset of the set of atomic symbols (Chang and Keisler [6]). This subset is the set of atomic symbols which are assigned the truth value true ( $t$ ). Hence, in the propositional case Nilsson's concept of possible worlds, i.e., a set of consistent truth value assignments, has a natural correspondence with the set of semantic models. Each possible world is completely determined by its truth value assignments to the atomic symbols of the language, and the assignments to the atomic symbols can be viewed as being the characteristic function of a semantic model (with  $t = 1$ ,  $f = 0$ ).

For example, in a propositional language with two atomic symbols  $\{A, B\}$  there are four possible worlds with corresponding semantic models ( $\sigma$  is used

Accession For  
DTIC GRA&I  
PREC TAB  
Availability  
Classification

|                     |                         |
|---------------------|-------------------------|
| By _____            |                         |
| Distribution/ _____ |                         |
| Availability Codes  |                         |
| Dist                | Avail and/or<br>Special |
| A-1                 |                         |

to indicate the truth evaluation function).

1.  $\{A^\sigma = t, B^\sigma = t\}$  or  $\{A, B\}$ .
2.  $\{A^\sigma = t, B^\sigma = f\}$  or  $\{A\}$ .
3.  $\{A^\sigma = f, B^\sigma = t\}$  or  $\{B\}$ .
4.  $\{A^\sigma = f, B^\sigma = f\}$  or  $\{\}$ , i.e., the empty set.

An equivalent way of looking at things, which will turn out to be more useful when we move to first order languages, is to consider the *atoms*<sup>1</sup> of the language. When the language has a finite number of atomic symbols each possible world can be represented as a single sentence: a sentence formed by conjoining each atomic symbol or its negation, such a sentence is called an atom. Corresponding to the four cases above we have the four atoms  $A \wedge B$ ,  $A \wedge \neg B$ ,  $\neg A \wedge B$ , and  $\neg A \wedge \neg B$ .

Given a probability distribution over the set of possible worlds it is possible to assign a probability to each sentence of the language. Each sentence is given a probability equal to the probability of the set of worlds in which it is true.

Equivalently, a probability distribution can be placed over the sentences of the logic, more precisely over the Lindenbaum-Tarski algebra of the language. This algebra is generated by grouping the sentences into equivalence classes. Two sentences,  $\alpha$  and  $\beta$ , are in the same equivalence class iff  $\vdash_0 \alpha \leftrightarrow \beta$ , where  $\vdash_0$  indicates deducible from the propositional axioms.

This technique is not limited to languages with a finite number of atomic symbols. When the language is finite, however, the atoms will be sentences of the language, and the probability distribution will be completely specified by the probabilities of the atoms (the e-classes of). Any sentence can be written as a disjunction of a unique set of atoms, and its probability will be the sum of the probabilities of these atoms. For example, if we specify the probabilities  $\{A \wedge B = .5, A \wedge \neg B = .1, \neg A \wedge B = .2, \neg A \wedge \neg B = .2\}$ , then the sentence  $A \vee B$  will have probability 0.8 as it can be written as  $(A \wedge B) \vee (A \wedge \neg B) \vee (\neg A \wedge B)$ .

---

<sup>1</sup>An atom in a Boolean algebra is a minimal non-zero element (Bell and Machover, [7]).

### 3 First Order Languages

When the move is made to first order languages certain problems arise. The first problem is that we lose the nice correspondence between possible worlds and semantic models. The normal semantic model for a first order language is considerably more complex than the model for a propositional language, and the truth value of the sentences in a first order language is determined both by the model and by an interpretation (i.e., the mapping from the symbols to the semantic entities). For a given truth value assignment to the sentences (possible world) there will be many different (in fact an infinite number) of model/interpretation pairs which will yield the same truth values. Hence, the semantic structure of the possible worlds is unclear.

Another difficulty, which Nilsson is aware of, is that Nilsson's techniques depend on being able to generate consistent truth value assignments for a set of sentences. These are used as 0/1 column vectors in his  $V$  matrix. This technique is limited to languages in which the consistency of a finite set of sentences can be established. The consistency of a set of first order sentences is not decidable, except in special cases (see Ackermann [8] for an interesting survey).

These difficulties can be avoided if instead of probability distributions over possible worlds we consider probability distributions over the Lindenbaum-Tarski (L-T) algebra of the language. It has already been demonstrated by Gaifman [2] that a probability measure can be defined over the L-T algebra of sentences of a first order language. Every sentence in the language will have a probability equal to the probability of its equivalence class, and furthermore, the probabilities will satisfy the condition

$$\text{If } \vdash \neg(\alpha \wedge \beta) \text{ then } p[\alpha \vee \beta] = p[\alpha] + p[\beta],$$

where  $\vdash$  indicates deducible from the first order axioms. This means that the probability measure preserves the partial order of the algebra. In this partial order we have  $\alpha < \beta$  iff  $\alpha \wedge \beta = \alpha$ ; hence,  $p[\beta] = p[\beta \wedge \alpha] + p[\beta \wedge \neg\alpha]$  (by the above condition), and  $p[\beta] = p[\alpha] + p[\beta \wedge \neg\alpha] \geq p[\alpha]$ . Under this partial order the conjunction and disjunction operators generate the greatest lower bound (infimum) and least upper bound (supremum).

To examine what happens to quantified sentences under such a probability measure it is sufficient to note that for L-T algebras we have that

(\*)

$$|\exists x\alpha| = \bigvee_{t \in T} |\alpha(x/t)|,$$

where  $| \circ |$  indicates the e-class of the formula, and  $T$  is the set of terms of the language. What this means is that each existentially quantified sentence is equal to the supremum of all its instantiations. This implies that the probability of any existentially quantified sentence must be greater than or equal to the probability of any of its instances. Similarly, the probability of any universally quantified sentence must be less than or equal to any of its instances.

This interpretation also makes sense in terms of Nilsson's possible worlds. In any possible world the existential must be true if any of its instantiations are. Hence, the set of possible worlds in which the existential is true includes the set of possible worlds in which any instantiation is true, and the existential must have a probability greater than or equal to the probability of any of its instances.

## 4 The Representation of Statistical Knowledge

Probabilities attached to logical sentences can be interpreted as degrees of belief in those sentences. Instead of either asserting a sentence or its negation, as in ordinary logic, one can attach some intermediate degree to it, a degree of belief. So, for example, one could represent a degree of belief of greater than 0.9 in the assertion "Tweety can fly" by assigning the sentence  $Fly(Tweety)$  a probability  $> 0.9$ . However, it is not easy to represent statistical information, for example, the assertion "More than 90% of all birds fly."<sup>2</sup>

---

<sup>2</sup>It is the case that first order logic is in some sense universally expressive. That is, set theory can be constructed in first order logic, and thus, sufficient mathematics can be built up inside the language to represent statements of this form. This is not, however, an efficient representation, nor is there any direct reflection in the semantics of the statistical information. Since probabilities attached to logical formulas generalizes ordinary logic,

First, propositional languages do not seem to possess sufficient power to represent these kinds of statements. This particular statistical statement is an assertion which indicates some relationship between the properties of being a bird and being able to fly, but it is not an assertion about any particular bird. This indicates that some sort of variable is required, i.e., this statement cannot be encode as a statement about any particular bird. Propositional languages do not have variables, and so are inadequate for this task even when they are generalized to take on probabilities instead of just 1/0 truth values.

When we move to first order languages we do get access to variables, variables which can range over the set of individuals. A seemingly reasonable way to represent this statement is to consider the probabilistic generalization of the universal sentence  $\forall x Bird(x) \rightarrow Fly(x)$ . The universal in 1/0 first order logic says that all birds fly, so if we attach a probability of  $> 0.9$  perhaps we will get what we need. Unfortunately, this does not work. If there is single bird who is thought to be unable to fly, this universal will be forced to have a probability close to zero. That is, the probability of this universal must be  $1 - p[\exists x Bird(x) \wedge \neg Fly(x)]$ . Hence, if one believes to degree greater than 0.1 that a non-flying bird exists, then the probability of the universal must be  $< .9$ .

Since universal quantification or its dual existential quantification are the only ones available in a first order language, it does not seem that moving to first order languages allows us to represent statistical assertions. There is, however, one more avenue available: conditional probabilities. We have probabilities attached to sentences hence with two sentences we can form conditional probabilities. It has been suggested (Cheeseman [9]) that meta-quantified statements of the following form can be used to capture statistical statements, in particular for the statement about birds:

$$\forall(x)p[Fly(x)|Bird(x)] > 0.9.$$

In this assertion quantification is occurring at a level outside of the language (at a meta-level): " $p[Fly(x)|Bird(x)]$ " is not a formula of the first order language. This statement is intended to assert that for every term,  $t$ , in the object language the conditional probability of the sentence  $Fly(x/t)$ , with the variable  $x$  substituted by the term  $t$ , given  $Bird(x/t)$  is  $> 0.9$ .

---

statistical information could be represented in this manner; however, I am concerned here with efficient representations.

However, this formulation also falls prey to any known exception. Say that there is some individual, denoted by the constant  $c$ , who is thought to be a bird, i.e.,  $p[Bird(c)]$  is high, and for some reason or the other is also believed to be unable to fly, i.e.,  $p[Fly(c)]$  is low, then clearly this statement cannot be true for the instance when  $x$  is  $c$ ; hence, the meta-level universal statement cannot be true. It is important to note that it does not matter what other things are known about the individual  $c$ . For example, it does not matter if  $c$  is known to be an ostrich; it is still the case that the conditional probability of  $Fly(c)$  given  $Bird(c)$  will not be  $> 0.9$ . Hence, there is no way that the statistical statement "More than 90% of all birds fly" can be represented by the assertion that the conditional probability is greater than 0.9 for all substitutions of  $x$ : this assertion will be false for certain substitutions. The problem here is that the statistical statement implies that  $p[Fly(x)|Bird(x)] > 0.9$  for a *random*  $x$ , but a universally quantified  $x$  is not the same as a random variable  $x$ ; furthermore, the simple device of assigning probabilities to sentences of a logical language does not give you access to random variables. This point has also been raised by Schubert [10].

## 5 The Representation of Defaults

There are many different defaults which have a natural statistical justification, the famous example of "Birds fly" being one of them. A natural reason for assuming by default that a particular bird can fly is simply the fact that, in a statistical sense, most birds do fly. This is not to say that all defaults have a statistical interpretation: there are many different notions of typicality which do not have a straightforward statistical interpretation, e.g., "Dogs give live birth" (Carlson [11], Nutter [12], also see Brachman [13] for a discussion of different notions of typicality).

Since probabilities attached to the sentences of a logic do not offer any easy way of representing statistical assertions, it is not surprising that attempts to use this formalism to give meaning to defaults leads to certain difficulties.

Recently Geffner and Pearl [4] have proposed giving semantics to defaults through meta-quantified conditional probability statements (also Pearl [5]<sup>3</sup>). For example, the default "Birds fly" is given meaning through

---

<sup>3</sup>Pearl uses a slightly different notion of probabilities within  $\epsilon$  of one. The technical



the meta-quantified statement  $\forall x p[Fly(x)|Bird(x)] \approx 1$ . In order to allow penguins to be non-flying birds they have the separate default rule:  $\forall x p[\neg Fly(x)|Penguin(x)] \approx 1$ . They also have universal statements like  $\forall x Penguin(x) \rightarrow Bird(x)$ . The probability of these universals is one; thus, as discussed above, every instantiation must also have probability one.

To examine the difficulties which arise from this approach consider the following example. Say that we have a logical language with the predicates *Bird*, *Fly*, and *Penguin*, some set of terms  $\{t_i\}$ , and a probability distribution over the sentences of the language which satisfies the default rules, i.e., for all terms  $t_i$ ,  $p[Fly(t_i)|Bird(t_i)] \approx 1$  and  $p[\neg Fly(t_i)|Penguin(t_i)] \approx 1$ , and in which the universal  $\forall x Penguin(x) \rightarrow Bird(x)$ , has probability one. Some simple facts which follow from the universal having probability one are that for all terms  $t_i$ ,  $p[Bird(t_i)] \geq p[Penguin(t_i)]$ , and  $p[Bird(t_i) \wedge Penguin(t_i)] = p[Penguin(t_i)]$ . Consider the following derivation:

$$\begin{aligned}
1 &\approx p[Fly(t_i)|Bird(t_i)] \\
&= \frac{p[Fly(t_i) \wedge Bird(t_i) \wedge \neg Peng(t_i)]}{p[Bird(t_i)]} + \frac{p[Fly(t_i) \wedge Bird(t_i) \wedge Peng(t_i)]}{p[Bird(t_i)]} \\
&\leq \frac{p[Fly(t_i) \wedge Bird(t_i) \wedge \neg Peng(t_i)]}{p[Penguin(t_i)]} + \frac{p[Fly(t_i) \wedge Bird(t_i) \wedge Peng(t_i)]}{p[Penguin(t_i)]} \\
&\leq \frac{p[\neg Penguin(t_i)]}{p[Penguin(t_i)]} + \frac{p[Fly(t_i) \wedge Penguin(t_i)]}{p[Penguin(t_i)]} \\
&= Odds[\neg Penguin(t_i)] + \approx 0
\end{aligned}$$

That is, the constraints imply that for any term  $t_i$  the odds that  $t_i$  is not a penguin must be at least  $\approx 1$ ; hence  $p[Penguin(t_i)]$  cannot be much greater than 0.5. Since  $\approx 0.5$  is an upper bound on the probability of all of these sentences, it must also be the case that it is an upper bound on the probability of the sentence  $\exists x Penguin(x)$ , by equation  $\star$ .

That is, if we accept the defaults we are must reject any sort of high level of belief in the *existence* of penguins.

This problem is similar to the problem discussed in the previous section; a universally quantified variable is not the same as a random variable, and cannot be used to encode a random variable.

differences between this approach and that of Geffner and Pearl do not make any difference to the following discussion; the anomalies presented also appear in Pearl's system.

## 6 Conclusions

It has been demonstrated that although probabilities can be assigned to the sentences of any first order language, the resulting probability logics are not powerful enough to efficiently represent statistical assertions. It has also been demonstrated that attempts to give defaults a probabilistic semantics using these probability logics leads to certain anomalies.

Statistical facts, it has been argued, give a natural justification to many default inferences. This implies that probabilities might still be useful for giving semantics to default rules and a justification to default inferences. For example, the default rule "Birds fly" could be represented as a statistical assertion that some large percentage of birds fly, and the default inference "Tweety flies" could be given the justification that Tweety probably does fly if to the best of our knowledge Tweety was a randomly selected bird.

Probability logics which accomplish this have already been developed (Bacchus [14], Kyburg [15]), but these logics go beyond the simple device of assigning probabilities to the sentences of a logical language. Bacchus uses a logic which has a probability distribution over the domain of discourse, this logic is capable of expressing statistical information, and possesses a sound and complete proof theory capable of reasoning with statistical facts. Default inferences are handled by an inductive mechanism which forms defeasible conclusions, conclusions which can be defeated by new information. Kyburg uses an object language/meta-language formalism, and has explored the inductive formation of defeasible conclusions in greater detail [16].

## 7 Acknowledgement

The author is thankful to Henry Kyburg for his helpful suggestions.

## References

- [1] Nils J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71-87, 1986.
- [2] H. Gaifman. Concerning measures in first order calculi. *Israel Journal of Mathematics*, 2:1-18, 1964.

- [3] Dana Scott and Krauss Peter. Assigning probabilities to logical formulas. In Jaakko Hintikka and Patrick Suppes, editors, *Aspects of Inductive Logic*, North-Holland, 1966.
- [4] Hector Geffner and Judea Pearl. *Sound Defeasible Inference*. Technical Report CSD870058, Cognitive Systems Laboratory, U.C.L.A., Los Angeles, CA. 90024-1596, U.S.A., 1987.
- [5] Judea Pearl. *Probabilistic Semantics for Inheritance Hierarchies with Exceptions*. Technical Report CSD870052, Cognitive Systems Laboratory, U.C.L.A., Los Angeles, CA. 90024-1596, U.S.A., 1987.
- [6] C. C. Chang and H. J. Keisler. *Model Theory*. North-Holland, Amsterdam, 1973.
- [7] John Bell and Moshé Machover. *A Course in Mathematical Logic*. Elsevier, Netherlands, 1977.
- [8] W. Ackermann. *Solvable Cases of the Decision Problem*. *Studies in Logic and the Foundations of Mathematics*, North-Holland, Amsterdam, 1968.
- [9] Peter Cheeseman. An inquiry into computer understanding. *To appear in Computational Intelligence*, 1987.
- [10] Lenhart K. Schubert. A travesty of truth: a reply to Cheeseman. *To appear in Computational Intelligence*, 1987.
- [11] G. Carlson. *Reference to Kinds in English*. PhD thesis, University of Massachusetts, 1977. unpublished—available from Indiana University Linguistics Club.
- [12] Jane Terry Nutter. Uncertainty and probability. In *Proceedings of the 10th IJCAI*, pages 373-379, 1987.
- [13] Ronald J. Brachman. I lied about the trees. *AI Magazine*, 6(3):80-93, 1985.
- [14] Fahiem Bacchus. Statistically founded degrees of belief. In *CSCSI-88 (Canadian Artificial Intelligence Conference)*, 1988.

- [15] Henry E. Kyburg, Jr. *The Logical Foundations of Statistical Inference*.  
D. Reidel, 1974.
- [16] Henry E. Kyburg, Jr. The reference class. *Philosophy of Science*,  
50(3):374-397, September 1983.