YAO-HUA TAN

IS DEFAULT LOGIC A REINVENTION OF INDUCTIVE-STATISTICAL REASONING?

ABSTRACT. Currently there is hardly any connection between philosophy of science and Artificial Intelligence research. We argue that both fields can benefit from each other. As an example of this mutual benefit we discuss the relation between *Inductive-Statistical Reasoning* and *Default Logic*. One of the main topics in AI research is the study of common-sense reasoning with incomplete information. Default logic is especially developed to formalise this type of reasoning. We show that there is a striking resemblance between inductive-statistical reasoning and default logic. A central theme in the logical positivist study of inductive-statistical reasoning such as Hempel's *Criterion of Maximal Specificity* turns out to be equally important in default logic. We also discuss to what extent the relevance of the results of Logical Positivism to AI research could contribute to a reevaluation of Logical Positivism in general.

1. INTRODUCTION

The research of *Artificial Intelligence* (AI) is a flourishing business. Numerous researchers are studying the problem how computers can be programmed such that they can perform cognitive tasks that used to be the sole domain of humans. The problems that arise in the field of AI are similar in many respects to classical epistemological and methodological questions that have been studied by philosophers for decades. In particular, there is a striking similarity between the issues that are studied in AI research and the so-called *Logical Positivism* tradition, which was the predominant research tradition in the philosophy of science from the thirties until the sixties. In this article we will concentrate on a parallel between AI and philosophy of science that has not been thoroughly investigated before; namely the relation between default logic and inductive-statistical reasoning. We will see that the similarities between these two formalisms are so numerous that the question arises whether default logic is not just a reinvention of inductive-statistical reasoning.

The basic assumption of Logical Positivism is that it is possible to give a logical analysis of scientific knowledge in general, and in particular of notions such as causality, explanation, lawlikeness, the distinction between observation and theoretical terms and inductive logic. These are

Synthese **110:** 357–379, 1997.

also important issues in AI research. Taking this similarity into account, one would expect that AI research is a kind of continuation of Logical Positivism by other means; computers instead of logic. But, strangely enough, there is no such link. Whereas AI research flourishes, Logical Positivism is declining. Interest is rapidly decreasing, and currently there are but a few research groups in the world that work in the logical positivist tradition. The difficult state of affairs Logical Positivism is in can be witnessed from survey books on the results of Logical Positivism (see e.g. Suppe (1977), Stegmüller (1983) and Salmon (1990)). For example, in Stegmüller (1983) the author is very pessimistic about the future of Logical Positivism. According to his opinion all the fundamental questions of Logical Positivism are still unsolved, and no satisfactory logical analysis has been given of typical logical positivist questions such as causality, explanation, lawlikeness, the distinction between observation and theoretical terms or inductive logic. In the preface to the second edition of his book *Erklärung, Begründung, Kausalität* Stegmüller even suggests abandoning the idea that a logical analysis of scientific knowledge is at all possible (see Stegmüller (1983)). According to him the basic assumption of Logical Positivism, that a logical analysis of scientific knowledge is feasible, is a dogma that should be discarded. Most philosophers of science seem to agree with him. Hence, during the last two decades the logically oriented philosophy of science has gradually been overshadowed by the philosophy of science of *historical case studies* and the *social-constructivist* philosophy of science, pioneered by philosophers of science such as Kuhn and Feyerabend. Initially, Logical Positivism looked a very promising research programme, but with hindsight many people have become increasingly sceptical about the objectives of Logical Positivism. In our opinion, this pessimistic view of Logical Positivism is unwarranted. Even if Logical Positivism has not fullfilled all of its promises, a wealth of knowledge and expertise has been acquired regarding problems which are relevant for AI research. In this article we will point out several parallels between Logical Positivism and AI research, and by doing so hope to show that Logical Positivism is still a fruitful research programme.

In recent years, a few publications have appeared in which parallels between Logical Positivism and AI research have been studied. We will first briefly discuss some of these publications that concern themes such as causality, explanation, lawlikeness and inductive logic.

Langley and his collaborators have been developing an expert system, called BACON, which can discover physical laws from raw data, i.e. it can perform a kind of inductive reasoning. For example, Langley claims that BACON has discovered the Ideal Gas Law from data about the pressure, volume and temperature of a gas in a container (see Langley et al. (1987) and Langley and Zytkow (1989)). However, they admit that BACON can not replace the human scientist. One of the limitations of BACON is that it requires a rather sophisticated pre-processing of the data. The system is highly sensitive to the way the data are presented, and it cannot suggest interesting experiments or fact-gathering based on vague intuitions such as humans can. Actually, it might be more appropriate to say that, BACON only discovers elementary mathematical relations between carefully pre-processed data.

Causality is another issue that has been addressed to by some AI researchers. For example, in Shoham (1988) causality is studied from the viewpoint of AI. In this book Shoham provides a new logical analysis of causality which is based on a non-monotonic logic (this notion will be explained later). Shoham shows in detail to what extent his analysis diverts from the traditional logical positivist analysis of causality. Non-monotonic logic is also relevant for the analysis of notions such as 'lawlikeness' and 'disposition'. Already in the forties Carnap pointed out that these notions are closely connected to the notion of so-called *counterfactual* sentences (see Carnap (1948)). Lewis developed the so-called *Conditional Logic* to model the logic of these counterfactual sentences (for references see Lewis (1973)). Van Benthem showed (in Benthem (1989)) that there are very strong parallels between conditional logic and non-monotonic logic.[1]

Recently, the logical analysis of explanation has become an actual theme in AI research again. In the field of non-monotonic logics a variety of alternative inference systems have been studied, which are very suitable to model pragmatic aspects of explanation (see e.g. Benthem (1990)). In Tan (1988) and Tan (1991) a new logical model of explanation and prediction was introduced, which can be viewed as an improvement of Hempel's *Covering Law Model* with respect to explanations that are based on incomplete information. This logical model is based on the non-monotonic logic from Shoham (1988). Using a simple electric circuit it is shown that in case of incomplete information explanation and prediction are not always symmetrical. This is an anomaly for Hempel's model, but our analysis can account for this phenomenon due to the use of non-monotonic logic. This new logical model was applied in Janssen and Tan (1991) to analyse explanations in economics.

Another interesting new view on explanation was introduced in Reiter (1987). In Hempel's *Covering Law Model* of explanation the main objective is to show that a particular fact was to be expected given certain laws. In Reiter's theory about fault diagnosis the objective is to explain the occurrence of *un*expected faults in a system, using certain assumptions about faulty components in this system. Reiter shows that these explanations of faulty behaviour of a system are based on a kind of non-monotonic reasoning. In Janssen and Tan (1992), Reiter's theory about diagnostic reasoning was applied to explanations in economic science. In particular it is shown in this article that Milton Friedman's so-called '*As-if*'-methodology can be viewed as a kind of diagnostic reasoning. This is a remarkable result, because until then no satisfactory logical analysis of Friedman's 'As-if'-methodology had been provided.

In this article we concentrate on the relation between default logic and inductive-statistical explanations. We show that there is a remarkable resemblance between these two research traditions. Both research traditions have the same research objective; to develop formalisms for reasoning with incomplete information. The strange thing is that in spite of this common goal, there is hardly any cross-over from the one research tradition to the other. With the exception of Shoham almost no AI researcher considered it worthwhile to study the results that were obtained earlier by logical positivist philosophers. Conversely, even the modern logical positivists seem to be totally unaware of the fact that the formal analysis of explanation is also an important issue in AI. Even in a survey book like Salmon's *Four Decades of Scientific Explanation* that was published as recently as 1990, there is no reference to AI research at all (see Salmon (1990))! One possibility could be that although these two research traditions study the same subject, they are interested in quite different aspects of the same problem. However, our detailed comparison shows that there is no evidence for such a discrepancy. We will see that in both research traditions the crucial problem that had to be dealt with is the problem of *Specificity*, i.e. when two arguments conflict with each other the most specific argument has to be preferred to the less specific argument. This criterion of specificity that was proposed in AI research appears to be very similar to the criterion of maximal specificity that was suggested by Hempel in the early sixties. However, it is not our sole purpose to demonstrate this parallel between inductive-statistical reasoning and default logic, but also to show that these research traditions could benefit from each other. Though they have much in common, there are also differences. For example in AI research the objective to find an implementation of reasoning with incomplete knowledge provides constraints on research that generates new and fascinating problems. On the other hand logical positivists have been very sharp in pointing out the most minute problematic details which might look rather abstract for a computer scientist, but which lurk in the subconsciousness of his expert systems nevertheless.

The organization of this article is as follows. In Section 2 we discuss inductive-statistical reasoning. In particular we discuss Hempel's model of inductive-statistical explanation and the well-known problem of statistical ambiguity. In Section 3 we introduce default logic and we compare it with inductive-statistical reasoning. We also show the parallel between the debate about the criterion of specificity in default logic and the debate about statistical ambiguity in inductive-statistical reasoning. Finally, in Section 4 we draw some conclusions.

2. INDUCTIVE-STATISTICAL EXPLANATIONS

One of the major results of Logical Positivism is the so-called *Covering Law Model* that was introduced by Hempel in the early sixties in his famous article 'Aspects of Scientific Explanation' (see Hempel (1965), Stegmüller (1983) and Salmon (1990) for an historical overview). The basic idea of this covering law model is that a fact or an event is explained by subsumption under a so-called *covering law,* i.e. the task of an explanation is to show that a fact or an event can be seen as an instantiation of a law. In the covering law model two types of explanation are distinguished: *Deductive-Nomological* explanations (DN-explanations) and *Inductive-Statistical* explanations (IS-explanations). In DN-explanations the law is *deterministic*, whereas in IS-explanations the law is *statistical.* An example of a deterministic law is $\forall x(F(x) \rightarrow G(x))$, which says that every object with the property F also has the property G. Hence, if we know that a certain object a has the property F, then we can use this law to derive that it also has the property G. In this sense the fact $G(a)$ is said to be subsumed by the law $\forall x(F(x) \rightarrow G(x))$ and the initial condition $F(a)$. According to Hempel a DN-explanation is an argument with the following logical structure.

$$\begin{array}{l} \forall x(F(x) \rightarrow G(x)) \\ F(a) \\ \hline G(a) \end{array}$$

The line indicates that $G(a)$ is a deductive consequence of the law $\forall x(F(x) \rightarrow G(x))$ and the initial condition $F(a)$. Therefore, this type of explanation is called deductive. The two premises are called the *explanans*, and the conclusion is called the *explanandum*. In IS-explanations the law is of a statistical nature. An example of a *statistical law* is $p(G, F) = r$ with $0 \leq r \leq 1$, where r denotes the probability that an object from the set F is also a member of the set G. The set F is called the *reference class* of the statistical law $p(G, F) = r$. Hempel introduced the idea that an IS-explanation has a structure which is analogous to the deductive structure of the DN-explanations, and which is as follows.

$$\begin{array}{ll} p(G,F) = r & \\ F(a) & \\ \hline\hline G(a) & [r] \end{array}$$

The double line in this argument indicates that the explanandum $G(a)$ is not a deductive consequence of the explanans $p(G, F) = r$ and $F(a)$. The symbol $[r]$ indicates the probability that the explanandum $G(a)$ occurs, given that $F(a)$ occurred. This IS-argument explains the fact that object a has property G in the sense that it was to be expected with probability r, given the statistical law $p(G, F) = r$, and the fact that a had the property F. Hempel has formulated several minimal conditions for IS-explanations. One of these conditions is that an IS-argument only has explanatory value if r is close to 1. The fundamental problem of IS-explanations is that until now no one has ever succeeded in developing a satisfactory logical method, the so-called *inductive logic*, to compute the value of r of a statistical law $p(G, F)$, given a finite number of observations about the sets G and F. Carnap especially did a lot of research in the forties and fifties on inductive logic, but the sought logical method has never been found.[2] Another minimal condition that Hempel stated is that statistical laws have to satisfy the axioms of mathematical probability theory. We will not discuss these axioms here in detail. The only axiom which is relevant for the sequel is the following one.

(Axiom 1) $p(G, F) + p(\neg G, F) = 1$

The symbol $\neg G$ denotes the complement of the set G. This axiom is intuitively plausible. Because, for example, if this axiom is not satisfied, and we have two statistical laws $p(\neg G, F) = 0.9$ and $p(G, F) = 0.9$, and object a has property F, then due to these two laws it is very probable that a has and has not the property G at the same time.

Right from the beginning it was clear to Hempel that two IS-explanations can yield contradictory conclusions even when they satisfy the axioms of probability theory. He called this phenomenon the *statistical ambiguity* of IS-explanations (see Hempel (1965) and Hempel (1968)). Hempel used the following example to illustrate this ambiguity. Let $I(x)$ be the predicate 'x is infected with malaria plasmodium', $M(x)$ be the predicate 'x has malaria', $S(x)$ be the predicate 'x is protected by a genetic property against malaria' and j denote the person John. Furthermore, assume that

$p(M, I) = 0.9$ and $p(\neg M, S) = 0.9$. Now suppose that John is infected with malaria plasmodium, but also has the property P. In this case both explanada of the following two IS-arguments are true, while they yield contradictory conclusions.

(2a)
$$\begin{array}{l} p(M, I) = 0.9 \\ \underline{I(j) \qquad\qquad} \\ \hline M(j) \end{array} [0.9]$$

(2b)
$$\begin{array}{l} p(\neg M, S) = 0.9 \\ \underline{S(j) \qquad\qquad} \\ \hline \neg M(j) \end{array} [0.9]$$

Note that the statistical laws in (2a) and (2b) do not violate Axiom 1, because these two laws do not have the same reference class. Hempel mentions that there is also a more specific statistical law which says that if somebody has the property S and is actually infected, then it is still highly probable that this person will not get malaria, i.e. $p(\neg M, S \wedge I) = 0.9$.[3] This more specific statistical law yields the following argument.

(2c)
$$\begin{array}{l} p(\neg M, S \wedge I) = 0.9 \\ \underline{S(j) \qquad\qquad\qquad} \\ \hline \neg M(j) \end{array} [0.9]$$

Comparing these three arguments (2a), (2b) and (2c), most people will agree that (2c) is the best explanation. When one knows that John is infected and also has the property S, then one expects that he will not get malaria. Hempel used this example to argue that it is of crucial importance to always use the most specific statistical law in IS-arguments. He calls this the *Requirement of Maximal Specificity* (RMS). The qualification 'the most specific' is of course relative. If one only knows that John is infected, but one does not know whether he has property P, then one considers the law $p(M, I) = 0.9$ the most specific one. In that case (2a) is the best explanation. Only if one knows that John is infected *and* that he has the property S, then one considers the law $p(\neg M, S \wedge I) = 0.9$ the most specific one, and (2c) the best explanation. Hence, the qualification 'the most specific' is relative to the knowledge state of a particular person. Let K denote the knowledge state of a person. Hempel assumes that this state K is deductively closed, i.e. K docs not only contain all the facts and laws that a person knows, but also all logical consequences of this knowledge. Furthermore, Hempel also assumes that K contains the axioms of probability theory. In Hempel (1968) the requirement of maximal specificity is defined as follows. An IS-argument of the form:

$$\begin{array}{l} p(G,F) = r \\ \underline{F(a) \qquad\qquad\quad} [r] \\ \hline G(a) \end{array}$$

is an *acceptable* IS-explanation with respect to a knowledge state K, if the following requirement of maximal specificity is satisfied.

RMS: For any class H for which the following two sentences are contained in K

(i) $\forall x(H(x) \rightarrow F(x))$,

(ii) $H(a)$,

there is a statistical law $p(G,H) = r'$ in K such that $r = r'$, unless this law is a theorem of probability theory.

The condition (i) implies that the class H is a subset of the class F. The basic idea of RMS is that if F and H both contain the object a, and H is a subset of F, then H provides more specific information about a than F, and therefore the law $p(G,H)$ should be preferred over the law $p(G,F)$. The ‘unless’-clause is needed to avoid the following problem. If we substitute the class $F \wedge G$ for H, i.e. $H = F \wedge G$, then it is obvious that H satisfies (i) and (ii). Since $p(G, F \wedge G) = 1$ is a theorem of probability theory, it follows with RMS that $p(G,F)$ must also have the value 1. But this implies that RMS without the ‘unless’-clause would exclude every IS-argument that contains a statistical law with a value smaller than 1, which is of course not what we want.

One can easily check that RMS rules out argument (2a) in the ‘malaria’ example. Let us assume that the knowledge state K contains all the premises of the arguments (2a)–(2c). The first premise of (2a) is the law $p(M,I) = 0.9$. Due to Axiom 1 it follows immediately that $p(\neg M, I) = 0.1$. However, this law is excluded by RMS, since (i) $S \wedge I$ is a subset of I, and (ii) John is an element of $S \wedge I$, while $p(\neg M, I) = 0.1$ and $p(\neg M, S \wedge I) = 0.9$. This violates RMS, because, given (i) and (ii), we should have $p(\neg M, I) = p(\neg M, S \wedge I)$. Argument (2b) does satisfy RMS, because (i) $S^* \wedge I$ is the only subset of S in K that satisfies the condition (ii) that it contains John as an element, and $p(\neg M, S) = p(\neg M, S \wedge I)$. Argument (3c) satisfies RMS for the trivial reason that K does not contain information about more specific subsets than $S \wedge I$. There are of course more specific subsets than $S \wedge I$, e.g. the subset $S \wedge I \wedge A$, where $A(x)$ is the predicate 'x lives in Amsterdam', but this information is not contained in K. This shows that RMS selects the right IS-explanations. Before we argued that, given the knowledge state K, the arguments (2b) and (2c) are intuitively acceptable IS-explanations, while (2a) is not.

Several authors have criticized Hempel's definition of RMS. We will discuss one of these critical comments that is also relevant for the comparison of IS-explanations with default arguments that we will discuss later.

In Humphreys (1968) Humphreys argues that RMS is too restrictive. There are cases in which RMS forces you to reject intuitively acceptable IS-explanations. Humphreys gives the following example. Let K contain the premises of the arguments (2a) and (2c). Furthermore, assume that K also contains the information that John is an element of the finite subset $S' \subseteq S \wedge I$, which consists of exactly four elements. Since S' contains four elements, it is obvious according to Humphreys, that the value of the statistical law $p(\neg M, S')$ is either 0, 1/4, 1/2, 3/4 or 1. However, this implies due to RMS that the argument (3c) is unacceptable, because (i) $S^1 \subseteq S \wedge I$, and (ii) John is an element of S', but $p(\neg M, S \wedge I) = 0.9$, which is not equivalent to either 0, 1/4, 1/2, 3/4 or 1. Although Hempel questions Humphreys claim that the value of $p(\neg M, S')$ is either 0, 1/4, 1/2, 3/4 or 1, he does admit that RMS has to be adapted, therefore he adds the following extra condition N_1 to RMS.

(N_1) The classes G and F in a statistical law $p(G, F) = r$ should not be finite on purely logical grounds.

By the criterion 'not finite on purely logical grounds' Hempel means that G and F should not be logically equivalent with a finite conjunction such as for example $G(a_1) \wedge G(a_2) \wedge \cdots \wedge G(a_n)$. The underlying idea is that G and F should refer to real natural sorts, which are infinite. According to Hempel Humphreys' example is excluded by N_1, because S' does not satisfy N_1.

Although RMS solves many conflicts between IS-arguments, it does not solve all conflicts. In Stegmüller (1983, 785) a well-known example is discussed of a conflict between two IS-arguments that is not solved by RMS. The example is as follows. Let $P(x)$ be the predicate 'x is a philosopher', $M(x)$ be the predicate 'x is a millionaire', $D(x)$ be the predicate 'x owns a diamond mine' and j again denote John. Furthermore, assume that 99% of all philosophers are not millionaires, whereas 99% of all diamond mine owners are millionaires; i.e. we have the two statistical laws $p(\neg M, P) = 0.99$ and $p(M, D) = 0.99$, respectively. Now suppose that John is a philosopher as well as a diamond mine owner, then we have the following two IS-arguments.

$$\text{(3a)} \quad \begin{array}{l} p(\neg M, P) = 0.99 \\ P(j) \\ \hline\hline \neg M(j) \end{array} [0.99]$$

$$\text{(3b)} \quad \begin{array}{l} p(M, D) = 0.99 \\ D(j) \\ \hline\hline M(j) \end{array} [0.99]$$

If the knowledge state K is the deductive closure of the premises from (3a) and (3b), then both arguments violate RMS. Argument (3a) is not an acceptable IS-explanation, because (i) $P \wedge D$ is a subset of P, and (ii) $P \wedge D$ contains j, but K does not contain a statistical law $p(\neg M, P \wedge D) = r$. An analogous argument holds for (3b). The basic problem is that the class P is not a subset of the class D, and neither is D a subset of P. Therefore, neither of these two conflicting statistical laws is preferred over the other. If, for example, D was a subset of P, then with RMS the law $p(M, D) = 0.99$ would be preferred over the law $p(\neg M, P) = 0.99$, and hence (3b) would be the only acceptable IS-explanation. From these observations Stegmüller draws the conclusion that if there are two conflicting statistical laws for which there is no subset relation between their respective reference classes, then no conclusion should be drawn at all. In the next section we will see that analogous unsolvable conflicts occur in default logic.

3. DEFAULT LOGIC

One of the major topics in AI research is the study of reasoning with incomplete information in expert systems. To formalise this type of reasoning new logics were developed, the so-called *non-monotonic logics.* For a concise survey of non-monotonic logics the reader is referred to Ginsburg (1987), Besnard (1989), Lukaszewicz (1990) and Marek and Truszcsynski (1993). Non-monotonic logics can be viewed as a kind of probabilistic reasoning without numbers. One of the best-known non-monotonic logics is *Default Logic* which was developed by Reiter (see Reiter (1980)). In this section we will discuss default logic and compare it with inductive-statistical reasoning.

The classical example of non-monotonic reasoning is the so-called ‘Tweety’ example. Usually, birds can fly, unless they are abnormal. Penguins are an example of abnormal birds, because they cannot fly. If one only knows that Tweety is a bird, then one draws the conclusion that Tweety can fly. However, if one knows not only that Tweety is a bird but also that he is a penguin, then one draws the opposite conclusion that Tweety cannot fly. Hence, in this example the addition of new information does invalidate earlier conclusions. Default logic was introduced to formalise these types of non-monotonic reasoning. Essentially, default logic is an ordinary first-order predicate logic, say $\mathcal{L}$, extended with extra inference rules that are called default rules. The logical form of a *default rule* is as follows:

$$(\alpha(x)\ :\ \beta_1(x),\ldots,\beta_n(x)/\omega(x)).$$

The subformulas $\alpha(x)$, $\beta_i(x)$, and $\omega(x)$ are predicate logical formulas from $\mathcal{L}$ with the free variable x. The subformula $\alpha(x)$ is called the *prerequisite*, $\beta_i(x)$ are the *justifications* and $\omega(x)$ is the *consequent* of the default rule. The intuitive interpretation of a default rule is as follows: if the prerequisite $\alpha(x)$ is true, and all justifications $\beta_i(x)$ are consistent with the available information (i.e. $\neg\beta_i(x)$ is not derivable from the available information), then one can assume that the consequent $\omega(x)$ is true. A *default theory*, which is denoted by $\Delta = \langle W, D\rangle$, consists of a set D of default rules, and a set W of predicate logical formulas. W can be viewed as the knowledge of a person he is certain of. The defaults generate defeasible conclusions from W. A set of conclusions that are derivable from a default theory Δ is called an extension of Δ. Let Th(S) denote the deductive closure in the logic $\mathcal{L}$ of a set S of $\mathcal{L}$-formulas, i.e. $\{\Delta | S \vdash_{\mathcal{L}} \Delta\}$. A set of $\mathcal{L}$-formulas E is an *extension* of the default theory $\Delta = \langle W, D\rangle$, if E is the smallest set such that

(i) $W \subseteq E$,

(ii) $E = \text{Th}(E)$,

(iii) For each default rule $(\alpha(x)\text{: } \beta_1(x),\ldots,\beta_n(x)/\omega(x)) \in D$, and each term t: if $\alpha(t) \in E$, and $\neg\beta_1(t),\ldots,\neg\beta_n(t) \notin E$, then $\omega(t) \in E$.

We explain this definition with a formalisation of the Tweety example in default logic. Let $B(x)$ be the predicate 'x is a bird', $P(x)$ the predicate 'x is a penguin' and $F(x)$ the predicate 'x can fly' and t denote Tweety. Rules of the form 'usually, A's are B's, unless they are C's are translated into default rules. The default rule 'usually, birds can fly, unless they are penguins' is translated into the following default rule.

$$(B(x) : \neg P(x)/F(x))$$

If one only knows that Tweety is a bird, then the default theory is $\Delta = \langle W, D\rangle$ with

$$W = \{B(t)\} \quad \text{and} \quad D = \{(B(x) : \neg P(x)/F(x))\}.$$

One can easily check that the formula $F(t)$ can be derived with the default rule, hence it is in the extension E of Δ. The crucial observation is that the formula $P(t)$ is not in E, because it is not in W, and there is no default rule in D with a consequent of the form $P(x)$, hence $P(t)$ is not derivable. Since $P(t) \notin E$, and hence also $\neg\neg P(t) \notin E$, the justification of the default rule is satisfied. The prerequisite $B(t)$ is satisfied, since $B(t) \in W$, and W is a subset of E. From $B(t) \in E$ and $\neg\neg P(t) \notin E$ it follows with the default rule that $F(t) \in E$. Hence Δ yields the conclusion that Tweety can fly. This conclusion is invalidated if we add more information. If we know not only that Tweety is a bird but also that he is a penguin, then the new default theory $\Delta' = \langle W', D\rangle$ is

$$W' = \{B(t), P(t)\} \quad \text{and} \quad D = \{(B(x) : \neg P(x)/F(x))\}.$$

Since the formula $P(t)$ is now in the extension E' of Δ', because W' is a subset of E', the justification of the default rule is no longer satisfied, hence $F(t)$ is not derivable any more, and therefore $F(t) \notin E'$. The conclusion that Tweety can fly does not follow from Δ'.

This phenomenon that conclusions are invalidated by expansions of the premise set is called *non-monotonicity*. A logic is called monotonic if conclusions (φ) are preserved under expansions (ψ) of the premise set (Φ), i.e. if $\Phi \vdash \varphi$, then $\Phi, \psi \vdash \varphi$. The Tweety example shows that default logic is non-monotonic. From Δ follows that Tweety can fly. If we expand Δ with $P(t)$ to Δ', then this conclusion is not preserved.

Hempel's RMS produces also non-monotonic effects in inductive-statistical reasoning. The malaria example is non-monotonic in the following sense. In the previous section it was observed that the conflict between argument (2a) on the one hand and the arguments (2b) and (2c) on the other hand depends on the knowledge state K. If K contains besides the three statistical laws only the information that John is infected, then RMS determines that (2a) is the best explanation. In that case K implies the conclusion that John will get malaria. However, if K is expanded with the premise $S(j)$, i.e. the information that John also has the property S, to the new knowledge state K', then RMS determines that (2b) and (2c) are the best explanations. In that case K' implies that John will not get malaria. Hence, the conclusion that John will get malaria is not preserved under expansion of K to K'.

In default logic the same type of conflicts occur as in inductive-statistical reasoning. In Etherington (1987) a conflict between default rules is discussed which is a classical example in the literature about default logic, and which is analogous to the 'malaria' example. This example is as follows. Consider the following rules.

"*Usually, adults are employed*"

"*Usually, high school dropouts are adults*"

"*Usually, high school dropouts are unemployed*"

Let $A(x)$ be the predicate 'x is an adult', $E(x)$ be the predicate 'x is employed', $D(x)$ be the predicate 'x is a high school dropout' and j denote again John. Etherington translates these rules in the following default rules.

$$D = \{(A(x) : E(x)/E(x)),\ (D(x) : A(x)/A(x)),$$
$$(D(x) : \neg E(x)/\neg E(x))\}$$

Suppose that W contains the information that John is a dropout, i.e. $W = \{D(j)\}$. Now the question is whether John is employed or not. In other words, is the conclusion $E(j)$ or the conclusion $\neg E(j)$ in the extension of the default theory $\Delta = \langle W, D\rangle$? Intuitively, the second conclusion that John is unemployed is the right one. However, we can derive both conclusions from this default theory. The default theory Δ has two extensions; one extension E_1 that contains $E(j)$, and another extension E_2 that contains $\neg E(j)$. In E_2 the conclusion $\neg E(j)$ is derived with the third default rule. In E_1 the conclusion $E(j)$ is derived by first deriving the formula $A(j)$ with the second default rule, and subsequently deriving $E(j)$ with the first default rule. Note that the first and third default rule cannot be applied simultaneously, because the consequent of the one rule violates the justification of the other rule. Since E_2 contains the right conclusion, we should prefer it over E_1. This conflict between multiple extensions of one default theory corresponds to the different conflicting IS-arguments that could be constructed with the same information. Hempel proposed his RMS to solve some of these conflicts. Etherington proposes a similar criterion to discriminate between conflicting extensions.

Common sense reasoning usually prefers one of the competing default by virtue of its prerequisite being more specific. (Etherington (1987, 47))[4]

According to this criterion the default rule with the 'most specific' prerequisite is preferred in case of conflicts. Let $A(x)$ and $B(x)$ be the prerequisites of the default rules d_1 and d_2. The prerequisite $A(x)$ is *more* *specific* than $B(x)$ if the set that the predicate A refers to is a subset of the set that B refers to, i.e. if the sentence $\forall x(A(x) \rightarrow B(x))$ is true. It is obvious that this criterion can be considered as the analogue of RMS in default logic. Hence, we reformulate this criterion of maximal specificity in default logic as follows.

RMS*: If a default theory has multiple conflicting extensions, then the extension is preferred which is generated by the most specific defaults.

Stegmüller's 'diamond mine owner' example showed that RMS does not solve all conflicts between IS-arguments. Likewise in default logic there are conflicts that are not solved by RMS* either. The classical example of an unsolvable conflict in default logic is the so-called 'Quaker-Republican' example. Usually, republicans are not pacifists. Usually, quakers are pacifists. Former president Nixon was a republican as well as a quaker. Was he a pacifist or not? The information is represented in default logic as follows. Let $Q(x)$ be the predicate 'x is a quaker', $R(x)$ be the predicate 'x is a republican', $P(x)$ be the predicate 'x is a pacifist' and n denote Nixon. The default theory is $\Delta = \langle W, D \rangle$ with

$$W = \{R(n), Q(n)\} \text{ and } D = \{(R(x) : \neg P(x)/\neg P(x)), (Q(x) : P(x)/P(x))\}.$$

This default theory also has two extensions. The one extension E_1 contains the conclusion $P(n)$, i.e. Nixon is a pacifist, while the other extension E_2 contains the conclusion $\neg P(n)$. This conflict cannot be solved by RMS*, because we cannot determine which of the two default rules is the most specific one. We do not know whether R is a subset of Q or vice versa, and it might very well be the case that neither of them is a subset of the other. This example of an unsolvable conflict in default logic is analogous to the 'diamond mine owner' example of Stegmüller. His proposal was that in such cases we should refrain from drawing any conclusion. In default logic there are two approaches. Some AI researchers argue that only those conclusions from a default theory Δ are acceptable that are contained in *every* extension of Δ (see for example Horty et al. (1987)). In this approach neither $P(n)$ nor $\neg P(n)$ follow from Δ, because none of these formulas are contained in the intersection of E_1 and E_2. This approach is in line with Stegmüller's point of view that in these conflicts no conclusion should be drawn. However, most AI researchers are more liberal than Horty and Stegmüller. They are of the opinion that different extensions should be considered and studied as different perspectives on the same problem. In that case there is not just one set of conclusions of a default theory, but a theory might have several different sets of conclusions that are correct relative to a certain extension. The first approach is called the *sceptical approach*, and the second is called the *credulous approach* to default reasoning. That this choice between these approaches is not just a matter of taste becomes clear from the discussion about so-called *cascaded ambiguities*. These are examples in which two or more default conflicts interfere in a chain of inferences. It appears that in cascaded ambiguities the sceptical and credulous approaches produce different conclusions (see Touretzky et al. (1987)). It could also be useful to make this distinction between sceptical and credulous approaches to IS-arguments, and to study cascaded ambiguities in inductive-statistical reasoning.

Finally we will discuss to what extent the critical comments about Hempel's RMS also apply to the specificity criterion RMS* in default logic. Humphreys' critical comment about RMS is interesting for the comparison of inductive-statistical reasoning with default logic, because in this respect default logic is better than inductive-statistical reasoning. Humphreys' criticism is directly related to the *quantitative* character of statistical laws. Since numbers do not play a role in default logic, Humphreys' argument does not hold for default logic. Actually, in the literature about non-monotonic logics the non-quantitative character of these logics is often mentioned as a reason to prefer these logics over probabilistic formalisms such as inductive-statistical reasoning or bayesian statistics (see for example McCarthy (1980)). In the early years of AI research there were some attempts to use probabilistic reasoning in expert systems, but these attempts were not very successful. For example, in the development of the famous medical expert system MYCIN Buchanan and Shortliffe initially intended to implement a probabilistic reasoner in this system. This attempt failed, however, due to complexity problems. Therefore, they introduced their own reasoner which is based on, as they call it, *certainty factors* (see Buchanan and Shortliffe (1984)). For simple tasks this reasoning system suffices, but it does violate some of the basic axioms of probability theory. Currently, the controversy between 'quantitative' approaches to reasoning with incomplete information such as bayesian statistics and 'qualitative' approaches such as default logic is still undecided. Critical comments like Humphreys' might support the conclusion that the qualitative approaches are better than the quantitative ones.

In Section 2 it was said that the fundamental problem of inductive-statistical reasoning is that until now no adequate inductive logic has been found yet, i.e. there is no method to compute the value r of a statistical law $p(G, F) = r$ on the basis of a finite number of observations about the classes F and G. Although these numbers do not appear in default logic, it is clear that a similar problem also occurs in default logic. The underlying problem is how to represent our everyday knowledge in statistical laws or default rules. Which sentences do we translate into statistical laws or default rules, and which sentences do we simply translate in classical logical formulas? This might look like a trivial question, and in toy problems it usually is a trivial question, but in developing large real-life expert systems it can become a very difficult problem (as every knowledge engineer can tell you!). Although AI researchers do suffer from these application problems, they do not study it in a very systematic way (a notable exception is Veltman (1996)). A systematic approach would be to look for minimal conditions for the application of a default rule.

Let us first consider the simplest default rules of the following form.

$$(A(x) : B(x)/B(x))$$

Such a default rule, of which the consequent is identical to the justification, is called a *normal* default rule. Are there any reasonable minimal conditions for normal default rules? Let us compare normal default rules with statistical laws of the form $p(B, A) = r$. Hempel proposed several minimal conditions for statistical laws. Do they make sense for default rules?

For example, initially Hempel required that the value r of a statistical law should be larger than 0.5. Is it reasonable to require that a normal default rule is only acceptable if the corresponding statistical law has a value r larger than 0.5? This seems a reasonable condition. The expression "usually A's are B's" seems to express that *most* A's are B's. So, we could require that only if we know that most A's are B's, then one can represent this information by a default rule. However, some researchers argue that default rules do not have this probabilistic character. For example, Veltman argues against this interpretation with the example that it is said that 'usually, crocodiles live thirty years', while more than seventy percent of the crocodiles actually die before they are three weeks old.[5] Another counter example against the probabilistic interpretation of default rules can be taken from the legal practice. In court the accused is by default assumed to be innocent. This is not because judges think that most of the accused are innocent, but because it is one of the fundamental principles of righteous law that the accused has the benefit of the doubt. Clearly, the question to what extent default rules have a probabilistic character deserves further investigation.

Another minimal condition that Hempel proposed was that statistical laws should satisfy the axioms of probability theory. Is it not reasonable to require the same for default rules? For example, statistical laws have to satisfy the axiom $p(G, F) + p(\neg G, F) = 1$. If we do require of normal default rules that the corresponding statistical law has a value larger than 0.5, then this axiom would imply that the two defaults

$$(A(x) : B(x)/B(x)) \quad \text{and} \quad (A(x) : \neg B(x)/\neg B(x))$$

should never be contained in the same default set D. This seems a reasonable condition that has not been proposed before as far as we know. Of course this is just a first start. The default logical equivalents of other theorems of probability theory still have to be investigated more closely. Some questions that are important in this context are (1) does the *non-additivity* of inductive-statistical reasoning also holds for default logic, and (2) if the notorious Lottery Paradox in inductive-statistical reasoning can also be formulated in default logic. Actually, the AI researcher Poole reported in Poole (1990) that lottery paradoxes do occur in default logic, and that one should be very careful to avoid these paradoxes in expert systems. The interesting thing was that Poole reports that he first discovered the lottery-paradox in the study of large expert systems, subsequently he searched the literature for an adequate analysis of this problem, and then he arrived at the logical positivist analysis of this paradox. This is a clear example how the results of Logical Positivism can be relevant to very practical problems in AI research.

With respect to the general default rules it is much harder to formulate reasonable minimal conditions, because it is not immediately clear what the justification of a default rule corresponds to in a statistical law. One proposal could be that the justifications $B_i(x)$ in the general default rule

$$(A(x) : B_1(x), B_2(x), \ldots, B_n(x)/C(x))$$

is represented as in the following statistical law

$$p(C, A \wedge B_1 \wedge \cdots \wedge B_n) = r$$

This proposal has the disadvantage that the syntactic distinction between the prerequisite and the justification of a default rule is blurred. For example, this translation has the result that in the default $(B(x) : \neg P(x)/F(x))$ about flying birds the conclusion that a bird can fly gets the same probability as the assumption that it is not a penguin, which is clearly counter-intuitive. Hence, this is also an issue that has to be further investigated.

Up till now we mainly discussed to what extent the logical positivist perspective could be relevant for AI research. We will finish this section with a few words about the opposite influence. To what extent is the AI perspective relevant for Logical Positivism? In the introduction we already mentioned the relevance of non-monotonic logics for the formal analysis of causality and explanation. A concise survey of recent developments in default logic can be found in Lukaszewicz (1990) and Marek and Truszczynski (1993). We briefly discuss just a few of these developments.

Besides default logic there are the two alternative important non-monotonic logics; namely *Circumscription* (see McCarthy (1980)) and *Auto-epistemic Logic* (see Moore (1985)). One of the main questions was whether these logics and default logic are intertranslatable. By now it has become clear that they are inter-translatable (see Konolige (1988) and Lukaszewicz (1990)). Circumscription is most suitable for fundamental research of non-mononotonic logic because of its so-called *preferential semantics*. The preferential semantics has become a new trend in logic (see Shoham (1988), Benthem (1989) and Kraus et al. (1990)). For implementation purposes, however, default logic is more suitable than either circumscription or auto-epistemic logic. Auto-epistemic logic has given a new impetus to the logic of knowledge and belief. For example, Halpern, has extended the classical epistemic logic, that started with Hintikka (1962), with new modal operators that contributed to the further development of auto-epistemic logic (see Halpern (1986)). One of these new operators is the so-called *common knowledge* operator. With this operator it is possible to logically represent the situation ‘I know, that you know, that I know . . . etc.’. Lewis argued already in the sixties that this type of knowledge is essential for maintaining social conventions (see Lewis (1969)). This is of course an important philosophical issue.

A second research topic which has attracted a lot of attention is the question which default theories have exactly one extension. This is an important property for implementation purposes, because having at most one extension reduces the complexity of the reasoning process considerably. On the other hand there are default theories, the so-called *incoherent* ones, which do not have an extension at all. An example of an incoherent default theory is the following one:

$$W = \{A(j)\} \quad \text{and} \quad D = \{(A(x) : B(x)/\neg B(x))\}$$

A variety of syntactical constraints on default rules to exclude incoherency have been proposed. For example, Etherington has shown that a default theory that contains only normal default rules is always coherent (see Etherington (1987)). Normal default rules are rules of the form ($\alpha : \beta/\beta$), i.e. the justification is identical to the conclusion.

The third and perhaps most important research topic is the implementation of default logic in expert systems. This is not a trivial matter, because default logic is not decidable. This is related to the fact that the truth of a justification of a default rule depends on the non-derivability of the negation of its justification, and this is undecidable. By now reasonably efficient implementations are discovered for large fragments of default logic (see Marek and Truszcsynski (1993), Risch and Schwind (1994) and Schaerf and Cadoli (1995)).

4. CONCLUSIONS

The aim of this article was two-fold. First to show that there is a striking similarity between AI research and Logical Positivism. A similarity which is usually neglected in the AI literature. Secondly, to show that both research traditions can benefit from each other.

Regarding our first aim we think that our detailed comparison of inductive-statistical reasoning with default logic has clearly shown that both formalisms are similar in two respects. First of all both formalisms were developed to model the same kind of reasoning with incomplete information. Secondly, in both formalisms the fundamental research topic that had to be solved was how to handle ambiguity in conflicting arguments. It turned out that Hempel's RMS could solve some of these conflicts, but only if there is a subset relation between the reference classes of a pair of conflicting statistical laws. Stegmüller's example of the diamond mine owner showed that without such a subset relation the conflict cannot be solved. We showed in some detail that the very same issues can be recognized in AI research. It appeared that Etherington's criterion of specificity RMS* is the analogue of RMS in default logic. And also in default logic this criterion of specificity can only be applied if there is a subset relation between the prerequisites of a pair of conflicting default rules. In the 'quaker-republican' example there was no such subset relation, and hence there were two extensions. This example is almost an alphabetical variant of the 'diamond mine owner' example, but we also pointed out that there are differences between the two formalisms. Hence, although both formalism were clearly developed to model the same phenomena, it is too simple to say that default logic is just a reinvention of inductive-statistical reasoning.

With respect to the second aim, the mutual benefit of AI research and Logical Positivism, our conclusions have a more tentative character. We discussed the relevance of Humphrey's critical comment on Hempel's RMS to Etherington's criterion of maximal specificity. Humphrey's comment was interesting, because in this respect the qualitative approach of reasoning with incomplete information was better off than the quantitative approach. Furthermore we made a first attempt to investigate whether Hempel's minimal conditions for statistical laws made sense for default rules. It appeared that some of these conditions do make sense. Finally, we pointed out that some of the well-known paradoxes of inductive-statistical reasoning such as for example the lottery paradox also play a role in default logic. The main benefit for Logical Positivism from AI research is of course that eventually they can run their ideas on the computer. A good example of a first attempt towards such a new logical positivism is Thagard (1988).

Does this comparison mean that a new future for Logical Positivism is in sight? the crucial question is what one considers to be the essential characteristic of Logical Positivism; the *problems* or the *methods*? If we identify Logical Positivism with the search of finding the perfect logical analysis of causality, explanation or inductive logic, then AI research is of little help. AI researchers have their own problems they are interested in. This distinction in 'problem focus' is related to difference in research goals of these two research traditions. The traditional logical positivists were mainly interested in analysing the cognitive activities of scientists. The goal of AI research is to program a computer such that it can reason with incomplete information. Occasionally they strike on deep philosophical problems, but they do not look for such problems. When they can circumvent these problems with an ingenious programming trick, they do not hesitate to act accordingly. The magic of the human mind is replaced by the turbo-power of the computer. The main question is no longer "How works the human mind?", but "How can we make the computer do it?". The same transition happened in aviation. Leonardo da Vinci claimed that humans have to learn how to fly from the birds. So he designed a machine that gave wings to a human being. This machine was a disaster. The modern jet-plane does not look like a bird at all, and yet it is hundred times faster than a bird. Imitating nature does not always produce the best results. Currently, in AI research the development of a new programming language like LISP or PROLOG has a much greater impact than a new theory about the human mind (see e.g. Sterling and Shapiro (1986)). Fascinating though these new programming techniques are, there still remains the problem how to apply these techniques. In other words, what are the conditions for the appropriate use of default rules. Philosophers of science can contribute to the solution of this problem. In this sense the methods and results of Logical Positivism can be applied fruitfully in AI research.

We finish with one last remark about Stegmüller. In his despair he suggested that we should discard the *Third Dogma* of empiricism; i.e. the assumption that philosophical problems can be solved by logical analysis. This looks like the behaviour of a rejected lover. From the fact that not all aims of Logical Positivism are achieved, one should not jump to the conclusion that the logically oriented philosophy of science is a complete failure. Instead we should accept our limitations, and realize that there are all kinds of less ambitious research goals that can be solved by logical techniques. This is the big challenge for the formal philosophy of science, and we hope that this article contributes to a revival of this approach in philosophy of science.

NOTES

[1] Roughly speaking, the parallel between these two logics is that the preference relation in the semantics of a non-monotonic logic is analogous to the accessibility relation in the possible world semantics of a conditional logic.
[2] For a concise overview of the developments in the study of inductive logic see Kuipers (1978).
[3] We make a slight abuse of notation here, because we use the conjunction symbol also to denote the union of the two sets S and I.
[4] Actually, specificity is no longer viewed as the only principle to resolve conflicts between defaults. Especially in the case of modelling of legal reasoning, where laws are often modelled as default rules, there are also other principles such as for example *Lex Posterior*, which says that the earlier law is preferred to the later one (see e.g. Prakken (1993)).
[5] Personal communication.

REFERENCES

Benthem, J. F. A. K. van: 1989, 'Semantic Parallels in Natural Language and Computation', in: H-D. Ebbinghaus, J. Fernandez-Prida, M. Carrido, D. Lascar and A. Rodriquez Artalejo (eds.), *Logic Colloquium Granada 1987*, North-Holland, Amsterdam, pp. 331–375.

Benthem, J. F. A. K. van: 1990, 'Inference, Methodology and Semantics', in: A. Blynov and P. Bystrov (eds.), *Festschrift for Vladimir Smirnov*, Philosophical Institute, Academy of Science of the USSR, Moskow.

Besnard, P.: 1989, *An Introduction to Default Logic*, Springer Verlag, Berlin.

Buchanan, B. and E. Shortliffe: 1984, *Rule-based Expertsystems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, Reading, MA.

Carnap, R.: 1948, *Meaning and Necessity*, Chicago University Press, Chicago.

Etherington, D. W.: 1987, 'Formalizing Nonmonotonic Reasoning Systems', *Artificial Intelligence* **31**.

Etherington, D. W.: 1988, *Reasoning with Incomplete Information*, Pitman, London.

Ginsberg, M. (ed.): 1987, *Readings in Non-Monotonic Logic*, Morgan Kaufman, Los Altos, CA.

Halpern, J.: 1986, 'Reasoning about Knowledge: an Overview', *Proc. of the First Conference about Theoretical Aspects of Reasoning about Knowledge*, Morgan Kaufman, Los Altos, CA.

Hempel, C. G.: 1965, 'Aspects of Scientific Explanation', in: C. G. Hempel, *Aspects of Scientific Explanation and other Essays in the Philosophy of Science*, The Free Press, New York.

Hempel, C. G.: 1968, 'Maximal Specificity and Lawlikeness in Probabilistic Explanation', *Philosophy of Science* **35**, 116–33.

Hintikka, J.: 1962, *Knowledge and Belief*, Cornell University Press, Ithaca.

Horty, J., R. Thomason, and D. Touretzky: 1987, *A Skeptical Theory of Inheritance in Nonmonotonic Semantic Networks*, Technical Report CMU-CS-87-175, Carnegie Mellon University, Pittsburgh.

Humphreys, W.: 1968, 'Statistical Ambiguity and Maximal Specificity', *Philosophy of Science* **35**, 112–5.

Janssen, M. C. W. and Y. H. Tan: 1991, 'Why Friedman's Non-monotonic Reasoning Defies Hempel's Covering Law Model', *Synthese* **86**, 255–84.

Janssen, M. C. W. and Y. H. Tan: 1992, 'Friedman's Permanent Income Hypothesis as an Example of Diagnostic Reasoning', *Economics and Philosophy* **8**, 23–49.

Konolige, K.: 1988, 'On the Relation Between Default and Autoepistemic Logic', *Artificial Intelligence* **35**, 343–82.

Kraus, S., D. Lehmann and M. Magidor: 1990, 'Nonmonotonic Reasoning, Preferential Models and Cumulative Logics', *Artificial Intelligence* **44**, 167–207.

Kuipers, T. A. F: 1978, *Studies in Inductive Probability and Rational Expectation*, Reidel, Dordrecht.

Langley, P. and J. Zytkow: 1989, 'Datadriven Approaches to Empirical Discovery', *Artificial Intelligence* **13**.

Langley, P., J. Zytkow, H. Simon and G. Bradshaw: 1987, *Scientific Discovery: Computational Explorations of the Creative Processes*, MIT Press, Cambridge MA.

Lewis, D.: 1969, *Convention: A Philosophical Study*, Harvard University Press, Harvard.

Lewis, D.: 1973, *Counterfactuals*, Blackwell, Oxford.

Lukaszewicz, W.: 1990, *Non-Monotonic Reasoning – Formalization of Common Sense Reasoning*, Series in Artificial Intelligence, Ellis Horwood, Chichester.

Marek, W. and M. Truszczynski: 1993, *Nonmonotonic Logic: Context-Dependent Reasoning*, Springer-Verlag, Berlin.

McCarthy, J.: 1980, 'Circumscription – A Form of Non-monotonic Reasoning', *Artificial Intelligence* **13**, 27–39.

Moore, R. C.: 1985, 'Semantical Considerations on Nonmonotonic Logic', *Artiflcial Intelligence* **25**, 75–94.

Poole, D.: 1990, 'The Effect of Knowledge on Belief: Conditioning, Specificity and the Lottery Paradox in Default Reasoning', *Artificial Intelligence* **49**, 281–307.

Prakken, H.: 1993, 'An Argumentation Framework in Default Logic', *Annales of Mathematics and Artificial Intelligence* **9**, 91–131.

Reiter, R.: 1980, 'A Logic for Default Reasoning', *Artificial Intelligence* **13**, 81–132.

Reiter, R.: 1987, 'A Theory of Diagnosis from First Principles', *Artificial Intelligence* **32**, 57–96.

Risch, V. and C. Schwind: 1994, Tableau-based Characterization and Theorem proving for Default Logic, *Journal of Automated Reasoning* **13**, 223–242.

Salmon, W. C.: 1990, *Four Decades of Scientific Explanation*, University of Minnesota Press, Minneapolis.

Schaerf, M. and M. Cadoli: 1995, 'Tractable Reasoning via Approximations', *Artificial Intelligence* **74**, 1–62.

Shoham, Y.: 1988, *Reasoning about Change. Time and Causation from the Standpoint of Artificial Intelligence*, MIT Press, Cambridge MA.

Stegmüller, W.: 1983, *Erklärung, Begründung, Kausalität*, Second edition, Springer Verlag, Berlin.

Sterling, I. and E. Shapiro: 1986, *The Art of Prolog*, MIT Press, Cambridge MA.

Suppe, F.: 1977, *The Structure of Scientific Theories*, second edition, University of Illinois Press, Urbana.

Thagard, P.: 1988, *Computational Philosophy of Science*, MIT Press, Cambridge MA.

Tan, Y. H.: 1988, 'Explanations with Incomplete Information; a Problem for Hempel's Theory about Causal Explanations', in W. Callebaut and P. Mostert (eds.), *Proceedings of the 9th Dutch Philosophy Conference*, Eburon, Delft (in Dutch).

Tan, Y. H.: 1991, 'Non-monotonic Epistemic Aspects of Scientific Explanations', *Logique et Analyse* **133-134**, 197–220.

Touretzky, D., Horty, J. and Thomason, R.: 1987, 'A Clash of Intuitions: The Current State of Nonmonotonic Multiple Inheritance Systems', *Proceedings of the 10th International Joint Conference on Artificial Intelligence, IJCAI'87*, 476–482.

Veltman, F.: 1996, 'Defaults in Update Semantics', to appear in the *Journal of Philosophical Logic*.

Erasmus University Research Institute for Decision and Information Systems (EURIDIS)
Erasmus University Rotterdam
P.O. Box 1738
3000 DR Rotterdam
The Netherlands
ytan@euridis.fbk.eur.nl