

fp11
14417

一般不. 刘增

8962804

MACHINE INTELLIGENCE 11

Logic and the acquisition of knowledge

edited by

J. E. HAYES

*Research Associate,
Turing Institute*

D. MICHIE

*Chief Scientist,
Turing Institute*

and

J. RICHARDS

*Head of Industrial Studies,
Turing Institute*



E8962804

CLARENDON PRESS · OXFORD
1988

Oxford University Press, Walton Street, Oxford OX2 6DP

Oxford New York Toronto

Delhi Bombay Calcutta Madras Karachi

Petaling Jaya Singapore Hong Kong Tokyo

Nairobi Dar es Salaam Cape Town

Melbourne Auckland

and associated companies in

Berlin Ibadan

Oxford is a trade mark of Oxford University Press

Published in the United States

by Oxford University Press, New York

© J. E. Hayes, D. Michie, and J. Richards 1988

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of Oxford University Press

British Library Cataloguing in Publication Data

Machine Intelligence.

1. Artificial Intelligence

I. Title II. Michie, Donald III. Richard J.

006.3

ISBN 0-19-853718-2

Library of Congress Cataloging in Publication Data

Data available

Typeset and printed in Northern Ireland at The Universities Press (Belfast) Ltd.

MACHINE INTELLIGENCE 11

MACHINE INTELLIGENCE

- Machine Intelligence 1* (1967) (eds N. Collins and D. Michie) Oliver & Boyd, Edinburgh
- Machine Intelligence 2* (1968) (eds E. Dale and D. Michie) Oliver & Boyd, Edinburgh
- (1 and 2 published as one volume in 1971 by *Edinburgh University Press*) (eds N. Collins, E. Dale, and D. Michie).
- Machine Intelligence 3* (1968) (ed. D. Michie) Edinburgh University Press, Edinburgh
- Machine Intelligence 4* (1969) (eds B. Meltzer and D. Michie) Edinburgh University Press, Edinburgh
- Machine Intelligence 5* (1970) (eds B. Meltzer and D. Michie) Edinburgh University Press, Edinburgh
- Machine Intelligence 6* (1971) (eds B. Meltzer and D. Michie) Edinburgh University Press, Edinburgh
- Machine Intelligence 7* (1972) (eds B. Meltzer and D. Michie) Edinburgh University Press, Edinburgh
- Machine Intelligence 8* (1977) (eds E. W. Elcock and D. Michie) Ellis Horwood, Chichester/Halsted, New York
- Machine Intelligence 9* (1979) (eds J. E. Hayes, D. Michie, and L. Mikulich) Ellis Horwood, Chichester/Halsted, New York
- Machine Intelligence 10* (1982) (eds J. E. Hayes, D. Michie, and Y.-H. Pao) Ellis Horwood, Chichester/Halsted, New York
- Machine Intelligence 11* (1988) (eds J. E. Hayes, D. Michie, and J. Richards) Oxford University Press, Oxford

PREFACE

Held at intervals in Scotland, the first seven International Machine Intelligence Workshops spanning the period of 1965–71 were involved in developing the new subject internationally—in those early days mainly as a mid-Atlantic phenomenon. Japan and continental Europe had yet to enter in strength. Also in the wings was the ill-famed ‘Lighthill report’ which in 1973 stigmatized machine intelligence as a mirage and in the UK demolished its local infrastructure.

Two and a half millennia ago, the historian Thucydides observed that it is not fortifications which make a city but people. In spite of dispersion, the AI culture under challenge evinced both hardiness and solidarity. Included in the exodus from Britain’s ‘AI winter’ were the MI Workshops themselves. Successively they found hospitality in Santa Cruz, USA (1975), Repino, USSR (1977), and Cleveland, USA (1981), by which time the distant tidings of Japan’s Fifth Generation presaged the coming thaw. Preparations were begun to found a new UK centre, the Turing Institute at Glasgow. By 1985 sufficient critical mass existed for the new Institute to be able to host a return after fourteen years to the series’ land of origin. With additional support from the University of Strathclyde, the eleventh Workshop took place at the University’s study centre at Ross Priory near the banks of Loch Lomond.

The titles of the twenty papers which now emerge are indicative of a continuing trend towards unity of approach. Logical models of deductive and inductive reasoning become ever more central and find a common frame in interactive environments for practical problem solving. We also see the first demonstrations that the fruits of past solutions can be systematically digested by an automated solver and built into incremental bodies of new, human-type, knowledge.

The long expected maturation of machine intelligence is evidently at last occurring apace. An adolescent’s elders not uncommonly warn, as elders of the physical sciences have of AI, that the youth may have outgrown his own strength. Has the maturation of machine intelligence been of this kind? With some confidence we commit this eleventh volume to the hands of its readers and invite them to pursue the question to their own conclusions.

February 1988

Donald Michie
Editor in Chief

CONTENTS

COMPUTATION AND LOGIC

1. Partial models and non-monotonic inference 3
K. KONOLIGE
2. Equational programming 21
N. DERSHOWITZ and D. A. PLAISTED
3. Beyond LOGLISP: combining functional and relational programming in a reduction setting 57
J. A. ROBINSON
4. Concurrent computer architecture for unification operations 69
J. V. OLDFIELD and C. D. STORMON

DEDUCTIVE PROBLEM-SOLVING AND PROOF

5. Integrating decision procedures into heuristic theorem provers: a case study of linear arithmetic 83
R. S. BOYER and J. S. MOORE
6. A problem simplification approach that generates heuristics for constraint-satisfaction problems 125
R. DECHTER and J. PEARL
7. The relation between programming and specification languages with particular reference to Anna 157
A. D. MCGETTRICK and J. G. STELL

LOGIC PROGRAMMING TOOLS AND APPLICATIONS

8. YAPES: yet another PROLOG expert system 167
T. B. NIBLETT
9. LogiCalc: a PROLOG spreadsheet 193
F. KRIWACZEK
10. Representing legislation as logic programs 209
M. SERGOT

MACHINE LEARNING: METHODS AND INSTRUMENTS

11. Incremental learning of concept descriptions: a method and experimental results 263
R. E. REINKE and R. S. MICHALSKI

CONTENTS

12. Generating expert rules from examples in PROLOG B. ARBAB and D. MICHIE	289
13. Decision trees and multi-valued attributes J. R. QUINLAN	305
14. RuleFactory: a new inductive learning shell S. RENNER	319
15. Intelligence architecture and inference: VLSI generalized associative memory devices D. R. MCGREGOR and J. R. MALONE	333

AUTOMATING THE ACQUISITION OF KNOWLEDGE FOR COMPLEX DOMAINS

16. Expert against oracle A. J. ROYCROFT	347
17. Inductive acquisition of chess strategies S. H. MUGGLETON	375
18. Validation of a weather forecasting expert system S. ZUBRICK	391
19. Comparison of ACLS and classical linear methods in a biological application B. SHEPHERD, J. PIPER, and D. RUTOVITZ	423
20. Automatic synthesis and compression of cardiological knowledge I. BRATKO, I. MOZETIČ, and N. LAVRAČ	435

INDEX	455
-------	-----

COMPUTATION AND LOGIC

Partial Models and Non-monotonic Inference

K. Konolige

Artificial Intelligence Center, SRI International and
Center for the Study of Language and Information, Stanford University, USA

Abstract

The non-monotonic character of common-sense reasoning is well recognized, as we often jump to conclusions that are not strictly justified by our partial knowledge of a situation. Most formalizations of this idea are best described as syntactic transformations on theories, with little or no semantic underpinnings. In this paper we develop a method of non-monotonic reasoning from a strictly semantic viewpoint, namely, as conjectures about how the missing information in a partial model should be filled in. The advantages of this approach are a natural and intuitively satisfying formalization of diverse types of non-monotonic reasoning, among them domain closure, the unique names hypothesis, and default reasoning.

1. INTRODUCTION

The importance of non-monotonic reasoning for common-sense domains is widely recognized in the field of Artificial Intelligence (AI). In this paper we will be concerned with such reasoning in its most general form, that is, in inferences that are *defeasible*: given more information, we may retract them.

The purpose of this paper is to introduce a form of non-monotonic inference based on the notion of a *partial model* of the world. We take partial models to reflect our partial knowledge of the true state of affairs. We then define non-monotonic inference as the process of filling in unknown parts of the model with *conjectures*: statements that could turn out to be false, given more complete knowledge. To take a standard example from default reasoning: since most birds can fly, if Tweety is a bird it is reasonable to assume that she can fly, at least in the absence of any information to the contrary. We thus have some justification for filling in our partial picture of the world with this conjecture. If our knowledge includes the fact that Tweety is an ostrich, then no such justification exists, and the conjecture must be retracted.

Of course, there are many different ways to represent partial knowledge of the world; in AI, first-order theories (FOTs) are a widely used method. However, FOTs are in a sense *too* partial for the purpose of non-monotonic inference—it is often difficult to decide just how the ‘partial’ should be filled. For example, consider the sentence

$$\text{Bird}(\text{Tweety}) \vee \text{Bird}(\text{Opus}) \quad (1.1)$$

This sentence gives us partial information about the world, in the sense we know either Tweety or Opus (or both) is a bird; but given just (1.1) it is impossible to conclude that we know Tweety to be a bird, or that we know Opus to be a bird.

Now suppose we are given a default rule stated informally as

$$\text{In the absence of conflicting information, assume that a bird flies.} \quad (1.2)$$

How can this rule be applied to our bird theory (1.1) to make conjectures about the ability of Tweety and Opus to fly? One approach is to relate the application of the default to a consistency condition on the theory, as in the default theories of Reiter (1980). Roughly speaking, our informal rule translates into the following rule for extending an FOT:

$$\text{If in a theory } x \text{ is a bird and it is consistent to assume that } x \text{ can fly, do so.} \quad (1.3)$$

Unfortunately such a default rule yields no new information when applied to (1.1). The disjunction does not permit us to conclude that any particular individual is a bird, and so it is impossible to instantiate the variable x in the antecedent of the default rule.

But clearly our intuitions are that (1.2) tells us something more about the theory (1.1). Suppose we ask what possible partial states of affairs would make (1.1) true. One of the following two is a minimally necessary condition:

1. Tweety is a bird.
2. Opus is a bird.

Now the application of the default rule is straightforward for each case, so we conjecture that either Tweety or Opus can fly.

One conclusion to be drawn from this example is that default reasoning should be based on an analysis of the models that a theory admits. It is the claim of this paper that partial models are an appropriate and natural level of description for the application of default rules, and other types of non-monotonic reasoning as well. In the next section, we support this claim by discussing general principles for implementing non-monotonic reasoning as conjectures on partial models, and by criticizing another model-based framework for non-monotonic inference, McCarthy’s (1980, 1984) circumscription schema, from this point of view. The rest of

the paper is devoted to illustrating the general principles using a particular type of partial model based on Hintikka interpretations, defined in Section 3. Because these models use the constants of a theory as their domain, they admit very natural treatment of assumptions involving equality and the naming of individuals, which are illustrated in Section 4, along with other types of default reasoning, including domain closure and the assumption of disjoint domains.

2. A SEMANTICS FOR NON-MONOTONIC INFERENCE

In this section we consider some general principles of a partial-model approach to non-monotonic inference, and introduce notation to be used throughout the paper. An analysis of circumscription based on these principles is also presented.

2.1. Conjectures on partial models

Any consistent set of sentences (or *theory*) T in a first-order language is satisfied by a set of (first-order) models. To continue the example from the Introduction: let *Tweety* refer to the individual TWEETY, and *Opus* to OPUS, and let BIRD and FLY be the properties of being a bird and flying, respectively. Now consider the models of $Bird(Tweety) \vee Bird(Opus)$:

$$\begin{array}{ll}
 M_1 = BIRD: \{TWEETY\} & FLY: \{\} \\
 M_2 = BIRD: \{TWEETY\} & FLY: \{TWEETY\} \\
 M_3 = BIRD: \{TWEETY, e_1\} & FLY: \{\} \\
 M_4 = BIRD: \{TWEETY, e_1\} & FLY: \{TWEETY\} \\
 M_5 = BIRD: \{TWEETY, e_1\} & FLY: \{TWEETY, e_1\} \\
 M_6 = BIRD: \{TWEETY, e_1, e_2\} & FLY: \{\} \\
 \vdots & \\
 M_i = BIRD: \{OPUS\} & FLY: \{\} \\
 M_{i+1} = BIRD: \{OPUS\} & FLY: \{OPUS\} \\
 M_{i+2} = BIRD: \{OPUS, e_1\} & FLY: \{\} \\
 M_{i+3} = BIRD: \{OPUS, e_1\} & FLY: \{OPUS\} \\
 M_{i+4} = BIRD: \{OPUS, e_1\} & FLY: \{OPUS, e_1\} \\
 \vdots & \\
 M_j = BIRD: \{OPUS, TWEETY\} & FLY: \{\} \\
 \vdots &
 \end{array} \tag{2.1}$$

These models naturally fall into two groups, corresponding to one of the two disjuncts in the theory: either Tweety is a bird, or Opus is (there are models such as M_j in which both these are true; such models fall into both groups). We can represent these groups by using the notion of a *partial model*. A partial model contains only a part of the information necessary in a (complete) model; by *extending* the partial model, we arrive at a set of models. In this example, we could construct two partial models by

specifying just a part of the extension of the BIRD relation:

$$\begin{aligned} m_1 &= \text{TWEETY} \in \text{BIRD} \\ m_2 &= \text{OPUS} \in \text{BIRD}. \end{aligned} \tag{2.2}$$

The extension of m_1 includes $M_1 - M_5$ and M_j ; the extension of m_2 includes $M_i - M_{i+4}$ and M_j . We write $E(m)$ for the set of extensions of a partial model m .

We have in m_1 and m_2 a formal model-theoretic counterpart of the informal reasoning we carried out in the Introduction. We can formulate the default rule 1.2 as the following conjecture:

$$\begin{aligned} &\text{If within a partial model } x \text{ is a bird and it is consistent} \\ &\text{to assume that } x \text{ flies, do so.} \end{aligned} \tag{2.3}$$

Note that this is exactly the default rule (1.3), except that ‘theory’ has been changed to ‘partial model’. A proposition P is *consistent* with a partial model m if there is an extension of m satisfying P . In the case of m_1 there are models in which Tweety flies, and so (2.3) picks out just that subset $\{M_2, M_5, \dots\}$ of $E(m_1)$; similarly, for m_2 we get the subset $\{M_{i+1}, M_{i+4}, \dots\}$. Since the world could be described by either m_1 or m_2 , we take the union $\{M_2, M_5, M_{i+1}, M_{i+4}, \dots\}$ of these models as the result of default reasoning. Obviously, $\text{Fly}(\text{Tweety}) \vee \text{Fly}(\text{Opus})$ is satisfied by each of these models.

To sum up: let T be a theory and α a conjecture on partial models. A conjecture picks out a non-empty subset of the extensions of a partial model. Non-monotonic inference can be viewed as the following process:

1. Let \mathbf{M} be all models of T . Form a set of all partial models \mathbf{m} . Let \mathbf{M}' be $\mathbf{M} - E(\mathbf{m})$, i.e. all models not in the extension of some member of \mathbf{m} .
2. Let $C(\alpha, \mathbf{m})$ be the set of extensions of \mathbf{m} chosen by the conjecture.
3. We say that a set of sentences T' is *inferred by α from T* if every member of T' is satisfied by each of $\mathbf{M}' \cup C(\alpha, \mathbf{m})$. We write this as $T \vdash_\alpha T'$.

Remarks. The general nature of non-monotonic inference here is the pruning of the set of models of a theory. For any given language L , we may have in mind certain types of models, the intended interpretations of L . For example, in studying resolution, we restrict our attention to Herbrand interpretations, in which all terms denote themselves.

We consider some general technical points of this definition. First, the inference operator $T \vdash_\alpha$ can be non-monotonic in T , as is easily shown by example. Let α be the conjecture that picks out only those extensions of a partial model in which P is false. We have:

$$\{Q\} \vdash_\alpha \neg P \tag{2.4}$$

but

$$\{Q, P\} \not\vdash_{\alpha} \neg P. \quad (2.5)$$

A special case, which is monotonic in T , is the conjecture δ that picks out *all* extensions of a partial model. The operator \vdash_{δ} is simple logical deduction, that is, for $T \vdash_{\delta} T'$, T' is the set of logical consequences of T , and hence also deductive consequences, by the completeness theorem for first-order logic.

Because conjectures pick out a subset of the possible models of T , the inference operator has the reflexive property

$$T \vdash_{\alpha} T. \quad (2.6)$$

Conjectures are thus appropriate for default reasoning or defeasible reasoning in general, where the initial facts, though sparse, are assumed to be accurate. There are, of course, other types of non-monotonic reasoning that are not naturally expressed as conjectures: for example, events are often treated formally as arbitrary transformations on models, and the revision of belief on the basis of new information requires changing a theory to admit models it did not originally have.

There is no guarantee that partial models exist, or if they do, that their extensions fully cover the set of models \mathbf{M} . \mathbf{M}' is designed to take up the slack in these situations, so that all models of \mathbf{M} are 'accounted for'. This, and the fact that conjectures are a pruning operation on sets of models, yield the following consistency property for the inference operator: if the initial theory T is consistent, then any set of inferred sentences is also consistent; that is, it is impossible to have

$$T \vdash_{\alpha} p \wedge \neg p. \quad (2.7)$$

The notion of the *coverage* of partial models is an important one, and is in some sense a completeness criterion for this method. If there are *no* partial models for a given theory, then for every conjecture α the operator \vdash_{α} becomes logical deduction, and no non-monotonic inference takes place. If the partial models of a theory fully cover the intended models (that is, every intended model is an extension of some partial model), then a conjecture on the partial models takes into account all of the interpretations of the theory. For example, the two partial models (2.2) cover all the models of $T = \text{Bird}(\text{Tweety}) \vee \text{Bird}(\text{Opus})$, and so the conjecture (2.3) gives us the maximum restriction on the models of T . An important feature of conjectures is that they degrade gracefully when some models are not covered, either because of theoretical or computational limitations. If for some reason only m_1 is used as a partial model of T , then the conjecture (2.3) produces the weaker result

$$T \vdash_{\alpha} \text{Bird}(\text{Tweety}) \supset \text{Fly}(\text{Tweety}). \quad (2.8)$$

One of the strengths of the method is that there are many different ways to construct partial models of the world. The type of partiality we choose to represent will influence the nature of the non-monotonic operator \vdash_α . For example, we might take partial models to be a subset of each relation's (positive) extension, as we did in (2.2); data bases are often viewed in this way (Gallaire *et al.*, 1978). A partial model of this sort covers a set of models that agree on the common subset, but can otherwise disagree. It invites the conjecture that the subset is the complete extension: there are no other true positive facts about the world (sometimes referred to as the *closed-world assumption*; see Section 4).

An important type of partiality, and one we will exploit for most of the remainder of this paper, is the ability to leave unspecified the equality (or inequality) of terms in a theory. One way to do this is by introducing syntactic elements into the partial models, as we do with Hintikka sets in Section 3. Partial models then become sets of atoms and their negations, including equality predications. For example, the set $\{Bird(Tweety), Bird(Opus)\}$ has extensions in which Tweety and Opus are the same individual, and in which they are different. Assumptions about the uniqueness of named individuals can be framed in terms of conjectures on this partial model.

2.2. Circumscription

Predicate circumscription is a proof-theoretic technique in which an FOT T is augmented by a circumscription formula. We can summarize its current formulation (from Etherington *et al.*, 1984) as follows: let P be a predicate, and P' a finite sequences of predicates of a finite theory T . Then $Circ(T, P, P')$ is a particular second-order formula expressing the circumscription of P , letting the predicates P' vary.

The semantics of circumscription come from the notion of P -minimal models. A model M is P -minimal if there is no other model N , agreeing with M everywhere except for the predicates P and P' , such that the extension of P in N is a proper subset of that in M . Circumscription is sound with respect to minimal models, in the sense that $Circ(T, P, P')$ is true in all P -minimal models of T ; however it is known to be incomplete (these results are summarized in Minker and Perlis, 1985).

Partial model conjectures have close ties with reasoning about minimal models. In fact, we can express the intended semantics of circumscription as a conjecture in the following way. We take partial models to be the P -minimal models, where the extension of a partial model M is the set of all models N which agree with M , except possibly on P and P' . The conjecture α is to pick only the minimal model itself.

Reasoning about minimal models was first employed in AI by McCarthy (1980) in an attempt to deal with what he called the *qualification problem*. In brief, this is the problem of stating formally

what objects and conditions *do not obtain* in a given situation. Using minimal models is a means of applying Occam's razor: only those objects are assumed to exist that are actually required by the statements of a theory.

It is not clear, however, that reasoning in minimal models is the best means of performing defeasible reasoning in general. For example, it can lead to a complicated statement of defaults by means of an abnormality predicate. Compare the compact formulation of Example 4.9 with the corresponding circumscriptive rendering on pp. 300–302 of McCarthy (1984). But the evidence here is not yet in, and awaits a fuller exploration of the application of circumscription.

With regard to assumptions about equality, certain inherent limitations are already known (see Etherington *et al.*, 1984). Because minimal models are defined with respect to a fixed denotation function for the terms of a theory, it is impossible to perform non-monotonic inferences about the equality of terms by reasoning in such models. However, there have been attempts to account for equality by importing names and their denotations as objects of the domain (Lifschitz, 1984; McCarthy, 1984).

By contrast, we can choose partial models in such a way that non-monotonic inferences about equality are possible. As we show in the next section, partial model conjectures enjoy a natural treatment of assumptions about equality, including domain closure and the unique names hypothesis.

Finally, non-monotonic inference using partial model conjectures has been defined to always yield a consistent extension for a theory. For circumscription this is not the case, unless every model of the theory is an extension of a minimal model. In those instances where this is not the case, it has been shown that the circumscription formula can be inconsistent with an originally consistent theory (Etherington *et al.*, 1984).

3. HINTIKKA SETS AS PARTIAL MODELS

We now introduce a particular type of partial model, based on the method of analytic tableaux. [We do not give more than a cursory presentation of this method. See Smullyan (1968) for a general introduction; the method used here is based on work by Hintikka (1955).] Consider a theory T , which may be infinite. A tableau for T is a tree whose nodes are sentences, constructed in the following manner. The root of the tree is an arbitrary element of T . The tree is grown in a systematic manner from its leaves by adding new nodes, either elements of T or sentences derived from previous nodes by a small set of rules. Some of these rules, those dealing with disjunction, cause splits in the tree. The end result is a (perhaps infinite, but finitely branching) tree