

Lecture Notes in Artificial Intelligence
Subseries of Lecture Notes in Computer Science
Edited by J. Siekmann

Editorial

Artificial Intelligence has become a major discipline under the roof of Computer Science. This is also reflected by a growing number of titles devoted to this fast developing field to be published in our Lecture Notes in Computer Science. To make these volumes immediately visible we have decided to distinguish them by a special cover as Lecture Notes in Artificial Intelligence, constituting a subseries of the Lecture Notes in Computer Science. This subseries is edited by an Editorial Board of experts from all areas of AI, chaired by Jörg Siekmann, who are looking forward to consider further AI monographs and proceedings of high scientific quality for publication.

We hope that the constitution of this subseries will be well accepted by the audience of the Lecture Notes in Computer Science, and we feel confident that the subseries will be recognized as an outstanding opportunity for publication by authors and editors of the AI community.

Editors and publisher

Lecture Notes in Artificial Intelligence

Subseries of Lecture Notes in Computer Science

Edited by J. Siekmann

347

K. Morik (Ed.)

Knowledge Representation and Organization in Machine Learning



Springer-Verlag

Berlin Heidelberg New York London Paris Tokyo

Editor

Katharina Morik

Gesellschaft für Mathematik und Datenverarbeitung

Schloß Birlinghoven, 5205 Sankt Augustin, FRG

CR Subject Classification (1987): I.2.6

ISBN 3-540-50768-X Springer-Verlag Berlin Heidelberg New York

ISBN 0-387-50768-X Springer-Verlag New York Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in other ways, and storage in data banks. Duplication of this publication or parts thereof is only permitted under the provisions of the German Copyright Law of September 9, 1965, in its version of June 24, 1985, and a copyright fee must always be paid. Violations fall under the prosecution act of the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1989

Printed in Germany

Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr.
2145/3140-543210

Preface

Machine learning has become a rapidly growing field of Artificial Intelligence (AI). Since the First International Workshop on Machine Learning held at Carnegie-Mellon University in Pittsburgh (USA) in 1980 ¹, the number of scientists working in the field of machine learning has increased steadily, a fact indicated, e.g., by a new journal devoted to machine learning which has been appearing since 1986 ². This situation allows for specializing within the field. There are two types of specialization: on subfields of machine learning such as *learning from examples* or *learning from observation*, or, orthogonal to them, on themes of interest such as, e.g., *learning and logic* or *knowledge representation in machine learning*. This book follows the thematic orientation. It contains research papers, each of which brings to light the relation between knowledge representation, knowledge acquisition, and machine learning from a different direction. The book is based on the Workshop on Knowledge Representation and Organization in Machine Learning held in 1987 in Schloss Eringerfeld (Germany). It should be useful for researchers in the fields of knowledge representation, its application in knowledge-based systems, knowledge acquisition, and machine learning; for teachers and students of artificial intelligence as well as for developers of expert systems who are interested in new methods of knowledge acquisition and maintenance.

Before we give an overview of the content of the book and describe the topics more carefully, we will first give a short introduction to machine learning and clarify some technical terms.

Introduction to Machine Learning

The first question, of course, is: **What is Machine Learning?** Since the very earliest days of computer science, the need for machine learning as a prerequisite of intelligent machine behavior has been known. Turing, for instance, describes machine learning as using indexed problems and their solutions to solve new problems (Turing 59). The machine generates solutions to problems by trial and error, a human teacher then assesses the solutions, and the machine remembers the good solutions and uses them for further problems. Learning in this case is not only remembering a good solution but also accessing the solution when it is appropriate. From this point of view, the important ingredients of learning are

¹ This year's international workshop is the fifth one and was held at the University of Michigan in Ann Arbor (USA). In addition to the series of international workshops, a series of European Working Sessions on Machine Learning was started in 1985 at the Universite de Paris-Sud in Orsay (France). The International Meeting on Advances in Machine Learning held in 1986 in Les Arcs (France) was another event demonstrating both the increased interest in machine learning and the advances achieved in the field of machine learning.

² The journal "Machine Learning" appears quarterly and is published by Kluwer Academic Publishers.

- examples (here: problem-solution pairs),
- assessments of the examples by a teacher, and
- indexing or classifying the examples so that they can be accessed in future situations.

Turing did not claim that people learn in this manner or that machine learning simulates human learning, nor do most AI researchers nowadays. Rather, learning is viewed as a basic intelligent behavior observable among machines, animals, and, at the highest level, among human beings. A theory of learning should describe learning wherever it occurs and allow for pointing out the differences between human and machine learning or human and animal learning ³. The definition of "learning" therefore should cover animal learning, human learning, and machine learning, each being a specialization of learning. Simon proposed such a definition:

"Learning is any change in a system that allows it to perform better the second time on repetition of the same task or another task drawn from the same population" (Simon 83).

This definition has been criticized for restricting itself to task performance. Michalski (86) argued that the notion of a goal has to be introduced explicitly into the definition. Under certain circumstances, decreasing performance can be the goal of the learning subject (e.g., in a labor camp, the workers learn to work less while appearing to do more). Michalski also argues that performance can be improved by means other than learning. A sharper knife cuts better, but sharpening the knife is not learning. Scott (83) also argues against Simon's definition. The usefulness of learned information need not be known at the time learning occurs. For instance, you may happen to notice the location of a library in a city without knowing that someone will ask you later in the street how to get to the library. However, now that you know where the library is you can answer the question. Thus, the knowledge turned out to be useful. Scott comes up with a definition very close to the one by Michalski:

"It is a process in which a system builds a retrievable representation of its past interactions with its environment" (Scott 83).

Whereas Scott restricts the definition to learning machines, Michalski puts it more generally:

"Learning is constructing or modifying representations of what is being experienced" (Michalski 86).

Ever since knowledge-based systems became an important paradigm in AI, the problem of acquiring the knowledge for such systems has been urgent (Hayes-Roth et al. 83). **Knowledge-**

³ This is a typical AI view of intelligent behavior such as communication or language, problem solving or inferences, adaptation or learning. Whereas the notion of language has become independent from the human subject of that intelligent behavior, the notion of learning is just beginning to become a term describing a system of behavior in its own right.

based systems are those which perform their task on the basis of knowledge. Knowledge is explicable and represents a domain theory and problem solving methods in a form closely corresponding to human descriptions. The performance of the system increases with more (correct) knowledge. Knowledge is supposed to be more easily added or changed than a program can be changed to implement a new feature of a system. In this paradigm, the notion of **knowledge representation** is prominent. A representation is a mapping from a model of a part of the world to expressions of a representation language. The representation language is constructed by using a representation formalism which allows interpreting the language by the inverse mapping from expressions of the language to assertions of the model.

The paradigm of knowledge-based systems established a **need for machine learning**. Machine learning is now viewed as a means of

- automatically building up parts of a knowledge base for a knowledge-based system and
- enhancing a given knowledge base.

Corresponding to these two demands, there are two main approaches in the field of machine learning. One is often called **similarity-based learning**. Famous learning algorithms of this type are AQn (Michalski 83) and ID3 (Quinlan 83). From given instances of concepts, they learn the concepts. The concepts are then used in order to recognize a new object as an instance of a particular concept. They are similarity-based in that they detect the attributes which all instances of a concept have in common (what is similar among all instances?) and attributes which an object may not have in order to be an instance of that concept (what is dissimilar among instances and not-instances?). This approach has proved its applicability especially for acquiring knowledge in **heuristic classification** systems (Clancey 85). These are knowledge-based systems whose knowledge base consists of a classification of problems and already known solutions to them. A problem is solved by recognizing it as belonging to a certain class and applying the known solution to it. As is easily seen, from instances of classified problems, similarity-based learning algorithms can build up a classification of problems which can then be used by the knowledge-based system for solving new problems by heuristic classification.

The kernel inside similarity-based learning algorithms is **inductive inference**. From a finite set of examples a general concept is induced. Of course, this inference is not a safe one; it is not truth preserving but only falsity preserving (Michalski 86). Therefore, the learned concepts must be revisable, the domain of a concept must be narrowed in order to exclude counter-examples, or it must be widened in order to include new positive examples. Induction is a justified generalization from an instance to a class. There are always several generalizations possible ⁴. Normally, the most specific generalization is looked for. The most specific generalizations

⁴ For a careful analysis of generalization see (Kodratoff, Ganascia 86), (Kodratoff 88).

possible can be further restricted by concept quality measures (Michalski 83) and can be improved by using **background knowledge**. Background knowledge is knowledge about the domain which is used by the learning algorithm. A learning program is knowledge-intensive if it uses knowledge in addition to the examples for its learning task. **Closed-loop learning** is a particular way of using knowledge for learning. The learning program uses the knowledge it has already acquired, i.e. it feeds back the learning results into the background knowledge. Of course, knowledge-intensive approaches require a more ambitious knowledge representation than those which have different languages for the representation of the examples and for the representation of the learning results and do not use any more knowledge for learning. We have said that induction is a *justified* generalization. If the justification is purely statistical, it is more likely to verify arbitrary concepts. A theoretical justification, on the other hand, presupposes what was to be achieved: a domain model.

The other main approach in machine learning is **explanation-based learning**. Here, the justification for the new, learned concept is theoretically based. An explanation of why a certain example is a member of a certain concept is generalized to represent the concept (Mitchell et al. 86; DeJong, Mooney 86). Most often, the explanation is a proof, thus requiring the domain theory to be represented in (restricted) predicate logic. The main advantage is that one example together with the explanation is sufficient for acquiring an operational concept description. The example focuses the learning activity. As opposed to inductive learning, here most of the knowledge is already present. The learning step is to put the knowledge to use or, in other words, to find an *operational* representation for a concept. Explanation-based learning has been applied for enhancing knowledge bases in several scenarios. One scenario is represented by the LEAP system (Mitchell et al. 85). Whenever a rule in the knowledge base is missing and thus preventing the problem solver from coming to a solution, the user gives a solution to the problem. The learning system verifies the solution and then generalizes the verification (explanation). This produces a new rule for solving new problems of the same kind. Another scenario is presented by Wilkins in this book.

Overview of the Book

Above, we presented machine learning as a method to acquire knowledge and enhance knowledge bases. In this book, we consider machine learning embedded in the knowledge acquisition and maintenance process. Therefore, we first look at this context. We focus on knowledge representations which ease knowledge acquisition. Coming closer to machine learning in knowledge acquisition, we look at knowledge representation for machine learning. Here are the various aspects:

- The representation of the events from which a system learns. This representation determines what can be learned. How can this representation be enhanced?
- The representation of the learning results. What is necessary in order to enable the learning of complex concepts?
- Knowledge especially needed for a particular learning algorithm. How should this be represented? Moreover, how can it be acquired?
- The states of the learning program. How can they be represented?

The relation between knowledge representation and machine learning is two-fold: a suitable knowledge representation eases machine learning, and machine learning builds up representations of knowledge.

The knowledge acquisition problem demands explainable, maintainable, and learnable representations and therefore has led to new developments in expert systems. **William Swartout** and **Stephen Smoliar** show how the requirement for a maintainable knowledge base can be used as a guideline for an expert system's representation ⁵. Their Explainable Expert System (EES) offers a KL-ONE-like representation formalism for building a representation language by which a domain model is expressed. This representation can then be compiled into a more performance-oriented representation. **Walter Van de Velde** also presents a two-level expert system: a representation for causal relations of a domain can be transformed into another one which is better suited for quick problem solving. As opposed to EES, his system transforms each solution found by reasoning on the causal model into performance-oriented rules, thus incrementally building up a rule base on the basis of a domain model.

The requirement of *learnable representations* exemplified by a vision system is discussed by **Michael Mohnhaupt** and **Bernd Neumann**. In vision systems, very complex examples have to be abstracted in order to find concepts. Mohnhaupt and Neumann describe a series of transformations from the original, close-to-pixel representation to a prototypical representation of car movements.

Even information retrieval systems set up a context for the *transformation of representations*. **Roy Rada** and **Hafedh Mili** reorganize different databases with different indexing schemata into one database. Their procedure is based on occurrences of the same term in both databases and uses analogy for mapping two indices.

The knowledge acquisition context is not restricted to systems. There is always a knowledge engineer, an expert, or an AI programmer involved. **David Littman** observed expert system designers at work and proposes some categories for cognitive entities and operations. A better understanding of knowledge engineers' cognitive processes when they represent knowledge

⁵ Swartout's use of "explanation" is different from that in the context of machine learning: the system can produce explanations which make sense to a user. This requires the knowledge base to be well structured.

could also provide a guideline for adequate representation formalisms, and these, in turn, would ease knowledge acquisition.

With the context now established the next articles of the book consider knowledge acquisition with a closer look at learning. The introduction to this view is presented by **Katharina Morik**. A unifying view of knowledge acquisition and machine learning is developed thus allowing their integration. The reversibility of knowledge is stressed, and the conclusions for an integrated system are drawn. As an example of such a system BLIP is introduced.

The DISCIPLE system of **Yves Kodratoff** and **Gheorghe Tecuci** integrates knowledge acquisition by questioning the user, machine learning using various techniques (similarity-based learning, explanation-based learning, and learning by analogy), and a performance element.

The next section of the book focusses on *knowledge representation for machine learning*. **Werner Emde** examines requirements for knowledge representation formalisms posed by different learning tasks and demonstrates an inference machine which is specially designed to support incremental machine learning and knowledge revision. This deals with the last question of our list of topics above.

Sabine Thieme discusses the knowledge which a machine learning algorithm needs, thus contributing to the third topic of our list above. In her case, the learning algorithm is model-driven, i.e. the induction is restricted by models indicating what to look for. She presents not only the representation of these models but also a manual acquisition method for them.

The next papers all deal with particular learning algorithms. Similarity-based learning is looked at with respect to the representation of the examples and with respect to the representation of the learned concepts. **Maarten van Someren** discusses the dependency of the *representation of examples* from which the system learns and the learning results (the first topic from our list above). His careful description of problems for simple rule learning leads to an approach that uses knowledge about the domain in order to construct a more complex representation of the examples.

Michel Manago and **Jim Blythe** enhance Mitchell's version space method (Mitchell 82). As *representations of the concepts being learned*, disjunctive concept descriptions are allowed. This requires heuristics to avoid a combinatorial explosion of possible specific generalizations. Thus, the paper deals with the second aspect from our list above by showing the implications of enhancing the representation of the concepts to be learned.

Learning by analogy is investigated by **Christel Vrain** and **Yves Kodratoff** with respect to problem solving. If the structure of a new problem is similar to that of an already solved problem, how can the known solution be transferred to the new problem? This gives the notion

of *similarity* in similarity-based learning a new meaning ⁶. It can be used not only to generalize but also to directly transform a solution in order to transfer it to a new problem. The authors present similarity-based learning by analogy which makes good use of the dissimilarities of the analogous problems.

David Wilkins and Michael Pazzani describe two kinds of explanation-based learning. Stefan Wrobel presents a model-driven concept formation algorithm.

David Wilkins stresses the value of *learning for knowledge-base repair*. His system observes an expert at work and creates "explanations" for the expert's actions. If the system cannot derive a goal with respect to the problem state and the domain theory that corresponds to the expert's action - this being an explanation - then the knowledge base is suspected of being incomplete. Learning is then undertaken to complete the knowledge base by creating a rule that would help to explain the expert's action.

In his system OCCAM **Michael Pazzani** combines similarity-based learning and explanation-based learning. In his setting of story-understanding an explanation is a justification for a prediction of what will happen. Events have certain outcomes and to understand events also means to be able to predict these outcomes. Events are represented as schemata, which are acquired by similarity-based learning from input events and are used by explanation-based learning to explain new events. Thus, the *representation needed for explanation-based learning*, that is, the schemata as compact representations of goal-directed sequences of actions, is *built up by similarity-based learning*. In this respect, Pazzani contributes to the first aspect of knowledge representation for machine learning.

Stefan Wrobel applies model-driven rule-learning. He deals with both knowledge-base repair - as does David Wilkins - and the introduction of new concepts into the representation language of the domain theory - as does Michael Pazzani. The new concepts are used by further learning, thus enabling the system to learn rules that were previously excluded from the hypothesis space because of representational limits. In this respect, Wrobel's contribution refers to the first topic of our list above. In respect of knowledge repair it is an attempt to cope with the risk of inductive learning. Inductively acquired rules can be too general. They need to be restricted as soon as a contradiction indicates the overgeneralization of the rule. Explicitly representing the applicability of a rule allows for narrowing the rule's domain and thus for revising and working out the results of inductive rule learning (Emde, Habel, Rollinger 83).

All the papers of this book investigating particular learning algorithms are also contributions to the knowledge acquisition and repair problem in that the problem of knowledge acquisition is the problem of representing a domain and task, and the aim of machine learning is to build up a

⁶ Similarity-based learning is commonly abbreviated SBL. Other common abbreviations are EBL for explanation-based learning and EBG for explanation-based generalization.

representation for a domain and/or task. Machine learning and knowledge acquisition meet at the topic of knowledge representation. The representation problem has to be confronted at various levels:

- the representation formalism has to be adequate for knowledge acquisition and machine learning as is discussed in this volume particularly by Emde and by Swartout and Smoliar;
- the representation language has to be adequate for machine learning and can be built up by machine learning as is particularly investigated by Manago, van Someren, Pazzani and Wrobel in this volume;
- the representation of domain knowledge is a complex process even if the representation language and formalism are given and adequate as is argued by Morik, Kodratoff and Tecuci, and Littman in this book.

The book by no means handles all the aspects of knowledge representation for machine learning and machine learning for knowledge representation. Nor does it fully cover the aspects it investigates, for this is impossible. However, a start is made, and the reader may feel inspired to think some more about the impact of representations on machine learning.

Acknowledgements

Without the support of the European Community's COST-13 project on "Machine Learning and Knowledge Acquisition", the workshop on "Knowledge Representation and Organization in Machine Learning " would not have been possible. The project, its vital leader Yves Kodratoff, and the scientific community it brought together supported the workshop in two respects: encouragement and financial support. Thanks also go to the project "Lerner", partially sponsored by the German Ministry of Research and Technology (BMFT) under contract ITW8501B1, in which scientists from Nixdorf Computer AG, Technical University of Berlin, and Stollmann GmbH have the opportunity to work on machine learning and knowledge acquisition. Thanks to all the colleagues!

Berlin, November 1988

Katharina Morik

References

- Clancey, W.J.(85): "Heuristic Classification"; *Artificial Intelligence* 27, p.289-350, 1985
- DeJong, G., Mooney, R. (86): "Explanation-Based Learning: An Alternative View"; *Machine Learning*, Vol.1, Number 2, p.145-176, 1986
- Emde, W.,Habel, Ch.,Rollinger, C.R.(83): "The Discovery of the Equator (or Concept Driven Learning)"; In: *Proc. IJCAI-83, Karlsruhe*, 1983
- Hayes-Roth,F.,Waterman, D.A., Lenat, D. (eds.): "Building Expert Systems"; Addison Wesley, London, 1983
- Kodratoff, Y.,Ganascia, J.G.(86): "Improving the Generalization Step of Learning"; In: Michalski/Carbonell/Mitchell (eds.): *Machine Learning, An Artificial Intelligence Approach*, Volume II; Morgan Kaufmann, Los Altos, 1986
- Kodratoff, Y. (88): "Introduction to Machine Learning"; Pitman, London, 1988
- Michalski, R.S. (86): "Understanding the Nature of Learning: Issues and Research Directions"; In: Michalski/Carbonell/Mitchell (eds.): *Machine Learning, An Artificial Intelligence Approach*, Volume II; Morgan Kaufmann, Los Altos, 1986
- Michalski, R.S., Stepp, R.E. (83): "Learning from Observation: Conceptual Clustering" In: Michalski/Carbonell/Mitchell (eds.): *Machine Learning, An Artificial Intelligence Approach*, Volume I; Tioga, Palo Alto, 1983, and Springer, Berlin, 1984
- Mitchell, T.M., Keller, R.M., Kedar-Cabelli, S.T. (86): "Explanation-Based Generalization: A Unifying View"; *Machine Learning*, Vol.1, Number 1, p. 47-80, 1986
- Mitchell,T.M.; Mahadevan,S.; Steinberg,L.I. (85): "LEAP: A Learning Apprentice for VLSI Design"; In: *Proc. IJCAI-85, Los Angeles, CA*, 1985.
- Quinlan, R.J. (83): "Learning Efficient Classification Procedures and their Application to Chess End Games"; In: Michalski/Mitchell/Carbonell (eds.): *Machine Learning, An Artificial Intelligence Approach*, Volume I; Palo Alto, p. 463 - 482, 1983, and Springer, Berlin, 1984
- Scott, P.D.(83): "Learning: The Construction of A Posteriori Knowledge Structures"; *Proc. AAAI-83*, Washington, 1983
- Simon, H.A.(83): "Why Should Machines Learn?"; In: Michalski/Carbonell/Mitchell (eds.): *Machine Learning, An Artificial Intelligence Approach*, Volume I; Tioga, Palo Alto, 1983, and Springer, Berlin, 1984

CONTENTS

Preface.....	v
<i>William R. Swartout, Stephen W. Smoliar</i>	
Explanation: A Source of Guidance for Knowledge Representation	1
<i>Walter Van de Velde</i>	
(Re)Presentation Issues in Second Generation Expert Systems	17
<i>Michael Mohnhaupt, Bernd Neumann</i>	
Some Aspects of Learning and Reorganization in an Analogical Representation	50
<i>Roy Rada, Hafedh Mili</i>	
A Knowledge-Intensive Learning System for Document Retrieval	65
<i>David Littman</i>	
Constructing Expert Systems as Building Mental Models or Toward a Cognitive Ontology for Expert Systems.....	88
<i>Katharina Morik</i>	
Sloppy Modeling.....	107
<i>Yves Kodratoff, Gheorghe Tecuci</i>	
The Central Role of Explanations in DISCIPLE	135
<i>Werner Emde</i>	
An Inference Engine for Representing Multiple Theories.....	148
<i>Sabine Thieme</i>	
The Acquisition of Model-Knowledge for a Model-Driven Machine Learning Approach.....	177
<i>Maarten W. van Someren</i>	
Using Attribute Dependencies for Rule Learning	192
<i>Michel Manago, Jim Blythe</i>	
Learning Disjunctive Concepts.....	211
<i>Christel Vrain, Yves Kodratoff</i>	
The Use of Analogy in Incremental SBL	231
<i>David C. Wilkins</i>	
Knowledge Base Refinement Using Apprenticeship Learning Techniques.....	247
<i>Michael Pazzani</i>	
Creating High Level Knowledge Structures from Simple Elements	258
<i>Stefan Wrobel</i>	
Demand-Driven Concept Formation.....	289