FISHER, Douglas H., PAZZANI, Michael J., LANGLEY, Pat (Eds.): **Concept Formation: Knowledge and Experience in Unsupervised Learning.** San Mateo, California: Morgan Kaufmann Publishers Inc. 1991. 472+XIII p with figures, tables, an index and detailed references to each of the contribution. ISBN 1-55860-201-1.

Concept formation systems discover a classification scheme over a stream of observations (events, objects, instances) (1) by an *unsupervised* and (2) by an *incremental* learning, i.e., (1) observations are not classified a priori by an external teacher as in the supervised case which would mean that the sole learning task consists in summarizing the commonality among members of the same categories and differences among competing ones, rather the system must invoke internalized heuristics to organize itself, its observations into categories; and (2) observations are not available en masse from the outset of learning, rather learning occurs with each new observation. After their processing the observations are no longer available, and no revision can be performed in case of inconsistencies. Clustering methods based on computing the similarity between all pairs of observations, e.g., are typically nonincremental. Because an observation encounters only the classification scheme left by its predecessors, incremental learning systems suffer from ordering effects, in that they may discover different categories depending on the order in which they process observations.

Unsupervised and incremental constraints are met in many real-world situations: humans observe a never-ending sequence of instances, and they can use their learned knowledge at any point in time. Adaptability under these same conditions is increasingly important for artificial intelligence which are characterized by the ability to explore alternative actions, and, in many cases, to acquire new information during this exploration. This latter capability is called machine learning, and it enables an "intelligent" system to improve its own performance. Thus, concept formation is the subject of psychological research as well as of machine learning. However, psychological experiments tend to incremental but supervised approaches whereas AI systems prefer nonincremental but unsupervised ones.

To come together and to promote interdisciplinary interaction in this field, a Symposium on Computational Approaches to Concept Formation was held at Stanford University in 1990. The desire for cross-disciplinary interaction has also guided the organization of the book, which largely grew out of the presentations at the symposium (Preface). The 15 contributions included in this volume are subdivided into three sections: (I) Inductive Approaches to Concept Formation, (II) Knowledge and Experience in Concept Formation, and (III) the Utility of Concept Formation in Intelligent Behavior. Each section has an introductory chapter. D. FISHER & M. PAZZANI (*Computational Models of Concept Learning*) open section (I) with an overview on inductive concept learning models in machine learning and psychology and explain basic principles and concepts. D. FISHER & M. PAZZANI introduce the next two sections with *Theory-Guided Concept Formation and Concept Formation in Context*, respectively.

That method-oriented subdivision reveals a methodical weakness often found in AI publications: Not the problem to be solved is the starting point of investigations, but a special approach. There is no search for the best method for a given problem, rather it is looked for applications of a given method. Doing so, the reader is confronted with a technical apparatus, and, at the end, with the (foreseeable) confession that the method described proved to be inadequate. Some contributions in the volume show this typical pattern. Because of the method-oriented view, the problems inherent in the applications are often overlooked, the difficulties are underestimated, and one desists automatically from a theoretical foundation.

Normally one should expect that the product of a concept formation system is a concept, and that in order to develop such a system the question 'what is a concept' must be answered beforehand. However, as to be learnt from the above, not concepts, rather a classification scheme as a knowledge baseis the system's learning outcome. The input of the system is an observation, its output is the corresponding class name possibly combined with the description of the class that allow a useful supplementing of the normally incomplete input information. Thus e.g., one might want to predict whether an observed object will be dangerous or not. In addition, the observation may cause a modification of the internal knowledge base. The following questions arise from these four problems: (1) How to describe the observations? (2) How to represent the knowledge base? (3) Which search strategy, and (4) which learning strategy should be used? The concept formation systems described in this volume differ from each other in their solutions of these problems, and some aspects are singled out at the expense of the others.

(1) "*Observations*" used are either fictive items, like binary digits, having nothing to do with real world situations, or they are in one or another form a list of property-value pairs describing real objects. E.g., the list

color: red
form: circle   material: plastic

may describe a toy-button. COBWEB, UNIMEM, CYRUS are all systems restricted to attribute-value languages. Some few systems as LABYRINTH operate on composite objects in which also relations can appear like "there are three objects, X, Y, Z; X is on top of Y, and both X and Y are to the left of Z" (K. THOMPSON & P. LANGLEY: *Concept Formation in Structured Domains*).

(2) An appropriate *knowledge base structure* can significantly improve the efficiency of learning. The knowledge base can be viewed as an optimized representation of all observations so that the language used to describe the observations determines decisively its structure. Observations are clustered into classes and their features are stored exactly one time. This leads automatically to an inherity model of features, i.e., to a hierarchical structure in one or another form. H.B. RICHMAN (*Discrimination Net Models of Concept Formation*) recommends with EPAM a classical n-ary discrimination net model. Most systems use as their nodes

probabilistic summaries which associate a probability or weight with properties. Such approaches are more flexible in the face of inconsistencies (or typicalities) that are liable to arise during incremental processing. COBWEB summarizes its observations in a hierarchy, in which each node is a probabilistic concept indexed by is-a links from its parents. A concept is described as a set of attributes and their possible values, along with the conditional probability that a value will occur in an instance of the concept. The system also stores the overall probability of each concept. Whereas CLUSTER/2 can represent an attribute 'color' with alternative values 'blue or red', a COBWEB concept C can represent the observed conditional probabilities, P(color = blue | C) = 0,6 and P(color=red | C) =0,4. The tree structure also appears when observations are described in a complex probabilistic manner like in LABYRINTH. However, it is rarely clear when a representational unit is general or specifc. J.D. MARTIN & D. BILLMAN (*Representational Specificity and Concept Learning*) examine, therefore, three different measures of specificity - (i) the number of observed instances that have contributed to the node's description, (ii) the number of attributes specified in a node, and (iii) the number of observed instances that match the node. Although related, the three definitions of specifity do differ.

(3) Different organizations of the knowledge base facilitate different types of *search strategies*. Most of the systems like EPAM, CYRUS and UNIMEM follow a top-down classification strategy; some systems also employ a bottom-up or bidirectional search. The search criterion defines what is found in the knowledge base. Some approaches incorporate an observation along a path of best-matching nodes which leads to strict categories. Overlapping categories may be formed if an observation is classified along all nodes that "adequately" match. Discrimination network systems use monothetic (polythetic) classifiers, in that classification at each step in the hierarchy is guided by a single attribute value (by many attribute values). Formal measures like similarity, search costs etc. enter into the decision rule. J.R. ANDERSEN & M. MATESSA (An *Incremental Bayesian Algorithm for Categorization*) use a probability measure to suggest candidate categories for classification.

(4) *Learning* is triggered in supervised systems by misclassification of an observation indicated by an external teacher; in unsupervised systems the learning starts whenever the classification of a new observation failed. The knowledge base may be constructed with a leaf node at the outset, and a generalizing process, which forms then step by step the abstract superstructure (specific-to-general learner); or it may be constructed just in the reversed order (general-to-specific learner). There are three kinds of knowledge base reorganization: Creation of a new class (node), merging existing classes into larger ones, and splitting classes into groups of finer granularity. Which of them will be the best choice in a given situation is decided (a) by minimizing / maximizing a quantity, or (b) by means of additional knowledge.

(a) COBWEB's evaluation function, e.g., favors clusters that maximize the potential for inferring information. Whether to add a new category or to use an existing one is decided in other systems by estimating which would most improve the probability of making correct predictions. There are also information-theoretic measures in use, or the known observations are partitioned into subsets, each corresponding to the value of a 'best' divisive attribute. The methods take into account the possibility of noise or variation of data that blurs the boundaries between categories, or that can confound feature intercorrelations.

(b) Inductive approaches in concept formation assume that observations are independent; the only link among them is implicit in the common language used for describing surface features such as size and shape. However, an important insight is that observations may interact, and these interactions influence clustering and characterization just as intrinsic properties do. Inductive approaches furthermore assume that surface similarity indicates functional similarity and vice versa. No doubt that surface features may provide a good heuristic guide to more fundamental similarities, however, form and function often interact in subtle ways, each suggesting alternative categorizations. Thus, animals may be grouped morphologically into classes for mammals, reptiles, fishes, birds; but they can be grouped into herbivores, omnivores, carnivores as well according to their feeding habits. Normally there is no information in the input data about the classification point of view. Therefore, prior to clustering, an inference preprocessing step is required to fill in the gaps from internal background knowledge. E. J. WISNIEWSKI & D. L. MEDIN (*Harpoons and Long Sticks: The Interaction of Theory and Similarity in Rule Induction*) confirmed experimentally that theories

- cause the people's attention to the features in the training items,
- cause people to interpret the same data differently, and
- allow people to view different features as similar at a higher level of abstraction.

Papers dealing with knowledge-based concept formation are therefore the most interesting contributions in this volume. An explanation-based system uses its background knowledge to explain or prove why a training example is a member of a given category. It then generalizes the explanation so that it will apply to future examples. According to R. J. MOONEY (*Explanation-Based Learning as Concept Formation*) the standard explanation-based methods can therefore be seen as forming a new specialized concept from a single unclassified instance, and, like the knowledge acquired by concept formation systems, that specialized concept can be used to predict missing information. B.H. ROSS & T.L. SPALDING (*Some Influences of Instance Comparisons on Concept Formation*) propose a reminding-based approach. When people are categorizing new instances, they may be reminded of some earlier instances and use it to categorize the new instance. Such reminding-based categorization may lead to a generalization essentially incorporating the commonalities of the to instances. J. YOO & D. FISHER (*Concept Forma-*

tion over Problem-Solving Experience) desribe EXOR which carries out concept formation over explanations and incorporates explanation-based und case-based research. EXOR abstracts redundant explanation substructures and organizes them hierarchical for reuse.

The explanation-based and case-based paradigms provide some guidance on how inference, categorization, and learning interact, though considerable research remains to be done before the field realizes a robust coupling of these processes within a single model. However, our ability to learn from instances and use knowledge requires a hybrid concept learning which involves both the application of such prior knowledge and the learning from similarities among instances (i.e., inductive learning). OCCAM illustrates several ways in which knowledge-driven and inductive mechanisms can interact. A hybrid model is also suggested from the experience with such pretentious applications like discovery and exploration, problem solving and planning, engineering applications, natural language processing, and efficient and intelligent information retrieval. Some examples for application are presented by Y. REICH & S. S. FENVES (*The Formation and Use of Abstract Concepts in Design*), W. IBA & J. H. GENNARI (*Learning to Recognize Movements*), P.D. SCOTT & S. MARKOVITCH (*Representation Generation in an Exploratory Learning System*) and T.SIMON, A.NEWELL & D. KLAHR (A *Computational Account of Children's Learning About Number Conservation*).

The volume gives a detailed overview on the activities in concept formation. The latter is a complex domain, and the authors make efforts for a clear representation of their ideas. However, there are obviously some methodical deficiencies. Concept formation systems as well as their psychological models are language-bound in the sense that they start with "representational units", i.e. with the description of objects anticipating in such a way the knowledge base structure for the most part. But how to get such a description is left open, and it seems that for all systems this task has to be performed by the user. The representation of real objects in a technical language may be hard work normally done in pattern recognition, and what is called there a 'pattern' is in fact a concept. It seems that when a real object is described then the most difficult task is already past. Concept formation systems restrict themselves to group hierarchically language objects according to heuristical principles, i.e. the known observations are partitioned into subsets. Doing so, they apparently produce at best meta-concepts. Many experimental results indicate that the human memory also forms such hierarchical structures for a more efficient use of the capacity and a faster access (known as priming effect). That may be an application of concept formation, too. However, the meaning of 'concept' is left unclear. Sometimes the nodes in the hierarchical constructs are regarded as concepts. But a node gets its whole significance from its integration with other nodes, and, as a consequence, we would have to revise our idea of a concept as an autonomous entity; what is called 'concept' seems to be only a short name of an ability: The system behaves for an external observer in such a way as it would have concepts to its disposal. Concepts in a conventional sense are then apparently the description of that ability.

An explicit consideration of suitable validation and performance tasks can have significant implications on the design of both psychological and computational models of unsupervised learning, but (as D. FISHER & M. PAZZANI p. 34f critically note) the importance of this observation is sometimes overlooked and the methods are often left implicit or not considered at all. In the supervised approach, the presence of a teacher with its possibility for a direct intervention mitigates the validation problem. At best, the system will be as subtle as its teacher. Unsupervised systems, however, get their ability from general, hypothetical principles. Thus, the prediction of unknown attributes or missing information about new observations, or augmentation of existing information is considered as an important efficiency proof in concept formation. But it seems to be questionable whether such formal measures can ensure that a system accomplishes in fact that task, for which it was developed.

Validation and performance are therefore important research desiderata. Other research areas concern more complete representation languages for objects and concepts, and the development of more global reorganization strategies for hierarchical methods. Complications caused by noise in the environment and overlapping categories are traditional research topics. Most important is the development of robust control and flexible representations that can mitigate ordering effects (D. FISHER & M. PAZZANI, p. 35). Some promising research is under way, but considerable work remains. The volume contains a variety of stimulations about an interesting domain.

Peter Jaenecke

Dr. Peter Jaenecke, An der Mauer 32
D-75334 Straubenhardt, Germany

WILLE, Rudolf; ZICKWOLFF, Monika (Eds.): **Begriffliche Wissensverarbeitung** (Conceptual knowledge processing): Grundfragen und Aufgaben (Basic questions and tasks). Mannheim-Leipzig-Vienna-Zürich: BI Wissenschaftsverlag 1994. 324 p. ISBN 3-411-17241-X

The volume on hand contains the papers read at the conference held at the Technische Hochschule Darmstadt (Darmstadt Technical University) from Febr.23-26, 1994, on the state of the art of Conceptual Knowledge Processing. Scholars from a variety of disciplines spoke out in favor of "a fitting way, from a human point of view, of handling media and tools for the processing and transmission of data and knowledge". Accordingly they advocated "such methods of and instruments for conceptual knowledge processing as assist man in his rational thinking, judging and acting, and as promote critical discourse".

In the paper by R. Wille, Darmstadt, attention is paid to the philosophical foundations of knowledge processing, which generally have been somewhat neglected so far.

The different views on the concept 'knowledge' and, in line therewith, the various definitions thereof encoutered in the literatrue, which definitions frequently reflect highly incongruous aspects, have, in my opinion, led to a certain chaos in the conceptual realm and to uncertainty in the terminological