NON-CUMULATIVE LEARNING IN METAXA.3

Werner Emde *

Technische Universitaet Berlin, Project KIT-LERNER,
Sekr. FR 5-8, Franklinstr. 28/29, D-1000 Berlin 10 (Fed. Rep. Germany)
emde%dbOtuil1.bitnetGwiscvm.wise.edu

ABSTRACT

Knowledge revision in incremental learning systems will usually be restricted by some external criteria to achieve a conservative behavior of the system. Unfortunately, conservatism has some well known drawbacks. Therefore, it can become necessary to drop these restrictions and to change over to a non-cumulative learning mode. In this paper the incremental learning system METAXA.3 is described which is able to perform a special kind of non-cumulative knowledge revision enabling it to learn without requiring unrestricted resources or the absence of noisy data. The generalization approach is sketched and knowledge revision in METAXA.3 is described.

I INTRODUCTION AND MOTIVATION

The appropriate generalization approach to be used in a learning system is determined by the learning task. The presence of noisy data, for example, requires the use of induction methods which are insensitive to noise [Quinlan 83]. The learning task also places demands on the knowledge revision abilities of the learning system.

Suppose a learning task description allows us to assume the availability of unrestricted resources for knowledge maintenance and knowledge revision, the existence of an 'oracle' (cf. [Shapiro 81]) or the possibility to make critical experiments, and the perfect reliability of incoming data. Then, knowledge revision is "no" problem. Unfortunately, the following scenario is possible as well: The learning task is such that the knowledge base of the learning system becomes contaminated with noisy data, 'oracles' are not available, resources are restricted, and incremental learning is required. In addition, the learning task may be defined in a domain with counterexamples for all theories claiming a minimum of elegance.

In the latter case, difficulties will arise if the "no problem" solutions for generalization and knowledge revision of the first case are applied to acquire a domain theory: It will hardly be possible to find any theory at all, neither a simple one which can be improved later, nor a complex one, because counterexamples will cause the refutation of these theories. This problem is well known and several researchers have proposed and/or used a heuristic called 'conservatism' to overcome this problem in their learning programs (cf. [Salzberg 85], [Emde/Habel/Rollinger 83], [Rose/Langley 86]).

Conservatism includes both: First, to exclude uncommon counterexamples in the induction process and second, to choose the smallest changes to rectify a theory. It is advisibe to complete this strategy with another one known from philosophy and psychology: The system should look for confirmation rather than disconfirmation as it would follow from a logical point of view. Confirmatory strategies together with conservatism constitute the 'confirmation bias' (cf. [Tweny/Doherty/Myatt 81]) of a learning system.

The use of a confirmation bias helps to develop a fully operational performance element in the early stages of induction. On the other hand confirmation bias also has many drawbacks, including convergence to local maxima. Therefore, we assume that a learning system should have the possibility to choose between at least two learning modes: the (quasi-)cumulative and the non-cumulative learning mode. A learning system is said to be in the non-cumulative learning mode if the extent of modifications made by knowledge revision procedures is not restricted by conservatism. At the least, the learning system should leave the (quasi-)cumulative learning mode if it has reached a local maxima or a domain theory which consists of a description of exceptions rather than a description of regularities.

In this paper the incremental learning system METAXA.3 is described which performs a specific kind of non-cumulative learning. In contrast to other approaches, non-cumulative learning in METAXA.3 relies neither on the simplicity of the learning task ('oracels' are available, backtracking is possible, ...) (cf. [Shapiro 81], [Rose/Langley 86]), nor it is simply a 'learning by scratch' [Michalski 85] where the knowledge of the system is deleted and redeveloped all over again. This approach is influenced by the work of Thomas Kuhn and Paul K. Feyerabend and tries to take advantage of their analysis of scientific discovery processes (cf. [Emde 86], [Tweny/Doherty/Myatt 81]).

* The work of the author is supported by the German Ministry for Research and Technology (BMFT) under contract ITW8501B1.

II THE LEARNING TASK

The learning task of METAXA.3 can informally be described as follows: Data about properties of objects of a world and about relations among these objects are continuously presented to the system. The data may be 'noisy*. The system is supposed to be able to answer questions about the facts of the world after each input. A question should be answered with 'yes* if a corresponding input was accepted or if a corresponding fact can be deduced with inference rules induced from regularities found in the factual knowledge. Furthermore, the system is supposed to use less resources than complete search in inductive and deductive processes would require.

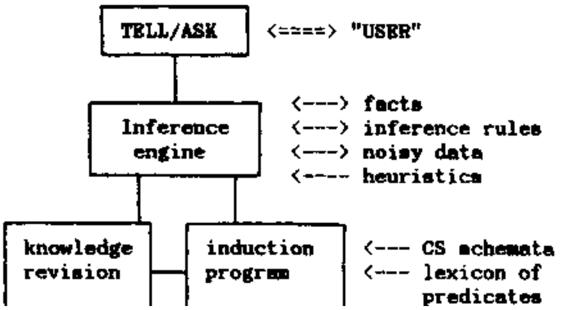

Fig. 1: General architecture of METAXA.3

This learning task is illustrated with the general architecture of the system shown in figure 1. A user-interface is provided by two functions "TELL" and "ASK". If an input, like

(FI) heavier_than(block3,blockl),

is consistent to the systems's knowledge it is stored in the factual knowledge base maintained by the inference engine with reasoning maintenance capabilities. After each input the induction process is triggered, which may cause the addition of one or more new inference rules, e.g. like RI (X,Wx,Y and Wy are variables):

(RI) weight(X,W_x) , weight(Y,W_y) ,
$W_x > W_y$ -> heavier_than(X,Y),

to the rule knowledge base of the inference engine. Inference rules which have been added to the rule base will be used in further inductive or deductive processes. If an input is inconsistent to the knowledge of the system, it will be rejected and stored as noisy data.

III THE GENERALIZATION APPROACH

The generalization approach of METAXA.3 is based on higher concepts like transitivity and conversity (cf. [Emde/Habel/Rollinger 83]). The higher concept 'num(erical)-comparative*, e.g., is defined to acquire rules like rule RI above. The definition of a higher concept like 'num_comparative' includes

- the rule-schema:
  numcomparative(P,Q):
  Q(X,Nx) , Q(Y,Ny) , Nx > Ny -> P(X,Y)
  (with Q and P as predicate variables),
- a schema for (positive) characteristic situations:
  Q(X,Nx) & Q(Y,Ny) & Nx > N_y & P(X,Y),
- the set of already induced "meta-facts" describing the inferential knowledge of the system according to the rule-schema, e.g.:
  num_comparative(heavier_than,weight,al1)
  num comparative(longer_than,length,all)**,
- and meta-rules describing the semantic relations between higher concepts, e.g.:
  num_comparative(P,Q,S) -> transitive(P,S)
  symmetrical(P,S) -> not(num_comparative(P,Q,S))

The generalization process of METAXA.3 can roughly be described as follows: The process is triggered by each input to the inference engine. First, hypotheses about possible inference rules are generated. Next, the hypotheses are ordered and finally tested using the corresponding characteristic situation (CS) schemata. If a) more than 'n' (usually 6) instances for the positive CS schema of a hypothesis can been found, b) no (or only a few) instances for negative CS schema, and c) only positive evidence for the corresponding meta-fact can be deduced with meta-rules ***, the hypothesis is confirmed. Then, the rule is generated using the rule-schema and entered into the rule base.

IV NON-CUMULATIVE KNOWLEDGE REVISION

Usually, a learning system will learn in the (quasi-)cumulative learning mode. Some knowledge revision strategies which may be used in this mode have been described in [Emde/Habel/Rollinger 83]. The non-cumulative learning mode should be entered either if a "crisis" (as known from Kuhn) takes place or if the quasi-cumulative learning mode has lead to the discovery of rules or concepts which might be used to restructure the system's knowledge.

The METAXA.3 system changes over to this learning mode if the experimental threshold of the maximum number of noisy data entries has been exceeded (a crisis) or if the system has induced a new rule with particular higher concepts. In both cases the system will then look for regularities in the factual knowledge which might be used to form a new theory. If a regularity (inconsistent to the old theory and factual knowledge) has been found METAXA.3 will try to develop a new theory and revise its knowledge according to this theory. In the following an example of the non-cumulative knowledge revision in METAXA.3 is given by a description of an actual run (see [Emde 86] for more details).

** The last argument of a meta-fact in METAXA.3 is used to specify the support set of the corresponding rule. A support set description not equal to 'all* will define a restriction to the applicability of the rule with regard to the arguments of the premises.

*** When METAXA.3 is looking for counterexamples of a rule the search depth is more restricted than in the search for positive instances. Thus, METAXA.3 is using a confirmatory strategy.

METAXA.3 induced a set of inference rules from a factual description of a world of floating and non-floating objects which can be used to answer questions about whether or not a particular object is floatable- The system's first "theory" sight be summarized as follows: Small objects are light; big objects are heavy; the weight of objects determines whether an object is heavy; light objects are able to float.

With the above theory the input about a non-floating needle 'needle1* has been rejected because it was described as small object. On the other hand the noisy input about a floating 'cablel' has been accepted because it was described as small.

The system then was supplied with facts about the materials of several objects. This lead to the generalization of a rule (R2) which uses the specific weight of materials to infer the volume of objects. The rule is also represented declaratively as met a-fact (MF1) using the name of the corresponding higher concept.

(R2) weight(X,Wx) , volume(X,V_x) ,
Z is Wx/Vx , material(X,Mxy) ,
material(Y,Mxy) , weight(Y,Wy) ,
V_y is Wy/Z -> volume(Y,Vy)
(MF1) const ratio(material,weight,volume,all)

METAXA.3 interprets this generalization as an interesting new rule which might be used to revise the system's knowledge because the following heuristic (cp. figure 1) will fire:

(HI) const_ratio(0,P,Q,S) -> eval(treat(
try_shift(thresh_ratio(P,Q,_,_,all))))****

This heuristic might be interpreted as: "If constant ratios have been found then try to develop a new theory with a 'thresh-ratio' meta-fact as its core hypothesis, i.e., look for an extreme value of all ratios (as an implicit new concept, like e.g., 'specific weight') to revise the old theory". After some search which is caused by unbound variables in HI MBTAXA.3 finds a regularity described by the meta-fact MF2 and can be read as: "If the ratio of weight to volume of an object is smaller than this ratio for ice-objects then this object will float. This rule is inconsistent to the system's first theory and would not be introduced in the quasi-cumulative learning mode.

(MF2) thresh-ratio(weight,volume,
is_ice_object,floating,all)

MF2 and the mete-facts derivable from it by using meta-rules are then added to this meta-fact forming the core of the new theory. Bach meta-fact of the old theory is classified according to whether or not it is consistent with the core of the new theory (once again by applying meta-rules). In the next step, a part of the factual knowledge is re-classified: Facts probably classified as "not noisy" by deleted rules (e.g., the input about the floating 'cablel') and facts that might be classified as 'noisy' by now rules are classified once more. Furthermore, the 'noisy data' entries are classified once more to rehabilitate data which have erronously been classified as 'noisy'. Then, MBTAXA.3 returns to the non-cumulative learning mode to work out the new theory.

**** Conclusions with 'eval' as main operator are evaluated by the inference engine rather stored as fact. The program 'treat* adds a task to the agenda.

V DISCUSSION AND ACKNOWLEDGEMENTS

In this paper the non-cumulative knowledge revision on MBTAXA.3 has been described. In contrast to the 'learning by scratch' approach which is another (simple) kind of non-cumulative knowledge revision, MBTAXA.3 takes advantage of parts of the old knowledge in two different ways: First, the old knowledge is applied to develop the core of the next theory and second, the old knowledge consistent to the new theory is incorporated into the new theory. Knowledge revision in MBTAXA.3 takes place on the meta-level without requiring complete bookkeeping as it is necessary, e.g. in the (quasi-)cumulative discovery program STAHLp [Rose/Langley 86]. In contrast to UNIMBM [Lebowitz 86] and other programs METAXA.3 applies rules generalized at one stage in subsequent learning stages. Many questions remain unanswered, such as how different theories can be compared and which heuristics can be used to estimate the usefulness of a theory-shift proposed by a heuristic in advance.

I would like to thank all members of the project group KIT (in particular K. v. Luck, K. Morik, B. Nebel, Ch. Peltason, S. Thieme and St. Wrobel) and the referees for helpful discussions and comments.

REFERENCES

Bmde, W. "Big Flood in the Blocks World (or Non-Cumulative Learning)." Proc. ECAI-86. Brighton, U.K., 1986, pp. 569-575

Bmde, W./ Habel, Ch./ Bollinger, C.-R. "The Discovery of the Equator (or Concept Driven Learning)." In Proc. IJCAI-83. Karlsruhe, F.R.G, 1983, pp. 455-458

Lebowitz, M. "UNIMBM, A General Learning System: An Overview." Proc. BCAI-86. Brighton, U.K., 1986, pp. 32-42

Michalski, R.S. "Knowledge Repair Mechanisms." In: Proc. 3rd International Machine Learning Workshop, Skytop, PA, 1985, pp. 116-117

Quinlan, J.R. "Learning from Noisy Data."; In Proc. 2nd International Maschine Learning Workshop, Monticello, Illinois, 1983, pp. 58-64

Rose, D. / Langley, P. "Chemical Discovery as Belief Revision." Machine Learning 1:4 (1986) 423-452

Salzberg, S. "Heuristics for Inductive Learning." In Proc. IJCAI-85. Los Angeles, CA, 1985, pp. 603-609

Shapiro, B.H. "Inductive Inference of Theories from Facts."; Research Report 192, Yale University, 1981

Tweny, R.D./ Doherty, M.B./ Myatt, C.R. (eds.) "On Scientific Thinking." Columbia University Press, 1981