GRDT: Enhancing Model-Based Learning for Its Application in Robot Navigation

Volker Klingspor

Universität Dortmund
Lehrstuhl Informatik VIII
D-44221 Dortmund
volker@ls8.informatik.uni-dortmund.de

Abstract

One central point of machine learning in general and inductive logic programming in special is the search space of the algorithms, defined by the control structure of the algorithms and additional knowledge. Since the sensible search space differs from domain to domain, a flexible way to describe this space is desired. To demonstrate problems occuring while using existing algorithms, we introduce learning tasks in a real world domain: concept learning for navigation of autonomous mobile robots. We point out differences between three existing algorithms used within this framework and their results. Since all of these algorithms have problems in solving the tasks, we developed GRDT *(grammar based rule discovery tool)*, an algorithm combining their ideas and techniques. In GRDT a two level description language is used for describing the hypothesis space. A grammar is used to define a set of second order rule schemata and these schemata then define the hypothesis space itself.

1 Introduction

In contrast to most of the robot systems actually used, future systems must have more user-friendly human-system interfaces, must be adaptable and portable to new, previously unknown environments, and must be able to handle uncertain and unknown events (Dillmann, 1993). Presently, commands are most often given on the level of real coordinates or at best by names of objects of which the positions the system knows. Flexible and user-friendly systems demand for commands in terms of *concepts*. Then, the user is not called to specify concrete objects like a specific door in the command. Instead, he can tell the robot to perform an action with any instance of a concept, e.g., he can tell the robot to go to the next a priori unknown door and enter the room. To execute this kind of commands, the robot must be able to classify the objects it perceives with its sensors, i.e., to assign them to concepts. Additionally, the system must be able to perform actions with these objects like moving through the doorway of the door found. K. Morik and A. Rieger (Morik and Rieger, 1993) have shown that this requires perceptual features and action features to be integrated, and perceptual features to be action oriented and action features to be perception oriented. They developed a representation for these *operational concepts* based on first order language.

To simplify the adaptation and the transmission of systems to new environments and new tasks, the rules to classify the perceptions should be learned by the system. In recent

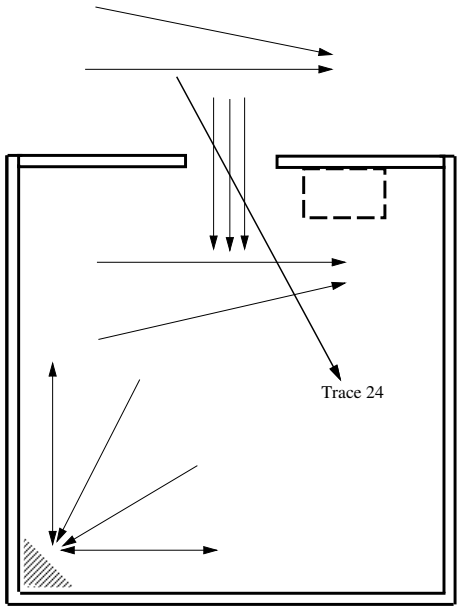

Figure 1: Room with traces

years, many algorithms for machine learning were developed, some of them able to learn in restricted first order logic. Nevertheless, *applying* these algorithms to real systems like autonomous mobile robots is a another challenging task.

In this paper, we will first shortly present the representation of the perceptions of a mobile robot and a scenario used for learning, both developed by K. Morik and A. Rieger. In Section 3, we will describe three different *inductive logic programming (ILP)* algorithms. We will also describe the results we got from applying these learning algorithms to our learning tasks. In Section 4, the new algorithm GRDT combining several ideas of these algorithms is presented, able to solve the learning problems we introduced. The paper ends with a short conclusion in Section 5.

2 Representation and Scenario

Suppose we have a robot moving through a known environment and measuring distances to objects while moving, using ultrasonic sensors. In our experiments, we used data of traces through a simple room (Figure 1)[1]. Figure 2 shows the measurements the robot gets from a left side sensor of the robot, called sensor 5, while entering the room diagonally through the doorway. This path is indicated as thicker line in Figure 1. The x-axis displays the time points of the measurements, the y-axis shows the measured distances. In the first part of the trace, the sensor gets no echo because the left hand side of the robot is empty. Then, the robot senses two times the door frames and three times the side of the cupboard. During the next three measurements, again the robot gets no echo, because the front side of the cupboard does not reflect the sonar beam back to the robot. The last measurements represent the wall.

[1] The data were prepared by the University of Karlsruhe.

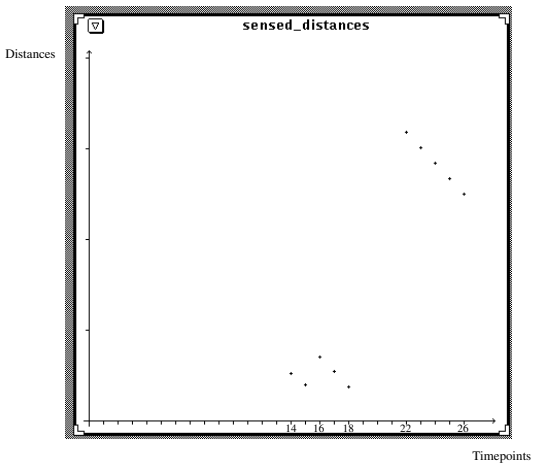

Figure 2: Distances measured by sensor 5 in trace 24

To be able to apply the sensed distances to more than one object, it is necessary to construct features from these measurements, i.e., we need a first step of abstraction. For this task, a sequence of measurements of a single sensor is divided into intervals of linear changes of measured distances. A symbol is attached to every interval dependent on the gradient of its function. The resulting sequence of these *basic features* for the example is:[2]

```
no_measurement (trace24, Or, sensor5, 1, 14, _).
decreasing     (trace24, Or, sensor5, 14, 15, -22).
incr_peak      (trace24, Or, sensor5, 15, 16, 47).
decreasing     (trace24, Or, sensor5, 16, 18, -30).
no_measurement (trace24, Or, sensor5, 18, 22, _).
decreasing     (trace24, Or, sensor5, 22, 26, -30).
```

The arguments of these predicates are the trace number, the orientation of the sensor relative to the world coordinates, the sensor number, the start and the end time point of the sequence and a value equivalent to the gradient. For example, the last line represents a sequence of decreasing measurements, measured by sensor 5 in trace 24 from the 22th to the 26th time point.

For characterizing classes of objects, chaining of basic features allows different numbers of measurements while perceiving one of the objects of a class, e.g., if the length of the objects or the speed of the robot differs from case to case. The chain of time points determines the need of a representation language higher than propositional logic, we have chosen Horn logic.

Since we want to learn from classified examples, the user must classify the traces, e.g.: "in trace 24, the robot moved diagonally through the doorway". From this classification, from information about the environment, and from the measurements, examples at different representation levels are generated. The first level describes situations, in which a single sensor sensed a particular constellation of edges. For example, the fact s_jump(trace24, sensor5, 14, 18, diagonal) represents that sensor 5 sensed a jump in trace 24 from time

[2]Learning the meaning of basic features, i.e., how to compute them, is another learning task, that will be solved by the University Dortmund within the project BLEARN-II.

point 14 to 18, i.e., two parallel walls got sensed in sequence, while the robot moved diagonally along this jump. In our example, these two walls are the door frame and the side of the cupboard. At the next representation level, situations are gathered, in which different sensors belonging to the same group of sensors out of several given groups sensed the same single sensor features. E.g., the fact sg_jump (trace24, left_side, 14, 18, diagonal) expresses that the left side sensors sensed a jump, while the robot moved diagonally along the jump. At the highest representation level, we combine features of different sensor groups, e.g., the left side and the right side sensors, to describe the perceptual part of operational concepts.

Our aim is to use machine learning to find rules describing the concepts of each level in terms of the next lower level. First, we need rules describing single sensor features by patterns of basic features. Such a rule could be:

decreasing (Tr, Or, Se, T_1, T_2, _) &
incr_peak (Tr, Or, Se, T_2, T_3, _) &
decreasing (Tr, Or, Se, T_3, T_4, _)
$\rightarrow$ s_jump (Tr, Se, T_1, T_4, diagonal).

The second kind of rules are sensor group feature descriptions. For these rules, we need additional knowledge about the class membership of the sensors in the particular traces. We try to learn rules combining single sensor features like in the following rule:

s_jump (Tr, Se_1, T_1, T_2, diagonal) &
s_jump (Tr, Se_2, T_3, T_4, diagonal) &
sclass (Tr, Se_1, left_side) & sclass (Tr, Se_2, left_side) &
d1succ (T_1, T_3) & d1succ (T_2, T_4)
$\rightarrow$ sg_jump (Tr, left_side, T_1, T_4, diagonal).

The first two premises express that two sensors perceived a jump while the robot moved diagonally along this jump. The next two premises guarantee that these two sensors belong to the sensor class left_side. The last two premises relate the start and end time points to be sure that the perceptions are chronologically close enough. These rules should classify the perceptions of the robot while it moves in new and unknown environments enabling the robot to interpret commands of the conceptual level and to execute actions according to these commands. Also, it can monitor whether the actions are performed correctly, because it can compare the measurements it got with the expected perceptions and it can change the further planned actions, if it sensed unexpected values (Klingspor, 1994b).

In this paper, two learning task are used to describe the behavior of the learning algorithms. In the first task, we try to learn descriptions for the single sensor feature s_jump from positive examples. We calculated 2129 basic features about 9 different predicates from 17472 measurements in 28 traces. In these traces, 206 times a sensor perceived a jump. For the second task, we used the 206 occurrencies of s_jump to learn descriptions for sensor group features. Because of the large number of sensor classes we defined, we have 225 positive examples of the sensor group feature sg_jump.

3 Different ILP algorithms

Efficient machine learning in full first order logic is impossible because first order logic is undecidable. Most first order logic learners operate on more or less restricted Horn logic.

The most frequent restriction used by these ILP-algorithms is to abstain from proper functions. There are additional syntactical restrictions like learning linked clauses only, i.e., clauses with all the arguments linked via the premises to the head arguments of the clause. The restriction used determines the complexity of the learning task and possibly the need of further semantical heuristics leading the search into areas of special interest and pruning parts that seem to be uninteresting.

I will now present three ILP-algorithms. The first one learns linked Horn clause programs using a semantical heuristic useful in many domains. The other two algorithms use different kinds of an explicit definition of the hypothesis space, so that it can be reduced. In contrast to the first two algorithms, the third one performs a complete search in the restricted hypothesis space.

3.1 FOIL

FOIL (Quinlan, 1990) is presently one of the most cited ILP-algorithms. It attracts attention because of applicability and good results in various domains without the need of additional knowledge, e.g., in form of syntactical structures as described in the next two sections.

3.1.1 General description of FOIL

FOIL's hypothesis space is linked Horn clause programs. To handle this huge space, parts of it will be pruned semantically during learning by a heuristic. This heuristic leading search through the hypothesis space to find a short path to a good hypothesis, is the *information gain*, well-known from the propositional learner ID3 (Quinlan, 1986).

FOIL's algorithm can be described as follows:

- while positive examples are still uncovered by the previously learned rules
  - while negative examples are still covered by the recent hypothesis
    - build all specializations of the recent hypothesis by extending the clause by a linked literal
    - calculate the information gain for all these specializations
    - choose the specialization with the best information gain for the next iteration

The hypothesis space is searched depth-first, the choice of the next hypothesis only depends on the heuristic. Clearly, using a heuristic, it is possible that interesting parts of the hypothesis space are pruned erroneously. FOIL's heuristic evaluates a hypothesis only dependent on the newly added literal. This makes the system shortsighted, since it rejects the addition of literals of which the positive effect is seen only after more literals have been added. So FOIL cannot learn all rules of its hypothesis space: it is an incomplete algorithm.

3.1.2 Learning results in the doorway domain

We applied FOIL to our two learning tasks. In the first case, learning patterns of basic features to describe single sensor features, FOIL learned a single rule:

no_measurement	(Tr, _,	Se,	_T_k,	T_1,	_) &
stable	(Tr, _,	Se,	T_1,	T_2,	_) &
incr_peak	(Tr, _,	Se,	T_2,	T_3,	_) &
something_happened	(Tr, _,	Se,	T_3,	_T_n,	_) &
straight_to	(Tr, _,	_Se_i,	_T_i,	_T_j,	_)
→ s_jump	(Tr,	Se,	T_1,	T_3,	M).[3]

This rule might separate the positive examples from the negative examples[4], but it is an unsatisfying rule for our application for two reasons. First, with this rule, the classification of s_jump not only depends on the measurements during the perception of the jump, but also on the measurements before and after the jump (premises 1 and 4). Second, the pattern of the specified sensor is combined with perceptions of an arbitrary other sensor possibly directed to another orientation. This constellation is only given in the environment we learn from, it cannot be expected in other environments.

FOIL was also applied to the second learning task: learning descriptions for sensor group features. The result shows that exclusively using of information gain is insufficient to lead the search. FOIL needs much time for learning, we cancelled the experiment after 26 hours with only 3 learned rules. Additionally, the learned rules cannot be used for our task, because they relate a given situation with arbitrary sensors, traces and so on. The longer the rules become, the more premises are added relating many uninteresting time points of arbitrary sensor features.[5]

Since FOIL, one of the best learning algorithms with heuristic search, is unable to learn the intended concept descriptions, we need algorithms restricting the hypothesis space by syntactical structures to lead the learning into areas of expected rules.

3.2 GRENDEL

3.2.1 General description of GRENDEL

GRENDEL (Cohen, 1991; Cohen, 1993) is a learning algorithm based on FOIL. Representation language, search strategy, and hypothesis evaluation are identical. The important difference is the way GRENDEL specializes hypotheses. This specialization can be determined by the user with a *antecedent description grammar* explicitly defining the hypothesis space. This grammar is an extended context free grammar: instead of usual terminal and nonterminal symbols, terminal and nonterminal *literals* are used. In each specialization step, only those specializations will be generated and tested that can be derived by the grammar in one step. Then, all these new hypotheses will be evaluated. GRENDEL uses a heuristic similar to the information gain used by FOIL. In difference to FOIL's heuristic this one does not only depend on the part of the hypothesis generated up to now, i.e., on the terminal literals of the hypothesis. Instead, GRENDEL builds all further grammatical derivations of the hypothesis and calculates the information gain also based on the further covered and uncovered examples. So GRENDEL is more foresighted than FOIL.

[3] Variables beginning with an underscore do not occur twice, they are anonymous.

[4] FOIL uses the closed world assumption, so examples not given are considered negative ones.

[5] A more detailed description of the learned rules and the used hypothesis space restrictions can be found in (Klingspor, 1994a).

3.2.2 Learning results in the doorway domain

We applied GRENDEL to our two learning tasks. We had to define grammars for both tasks. For the first task, defining the grammar was easy, but the resulting grammar was large, because for each predicate, a grammar rule deriving this predicate from a nonterminal is needed. GRENDEL learned all patterns describing sensor features in short time. The learned rules corresponds exactly to the expected rules described in Section 2.

In contrast to the good result of the first task, the result of learning sensor group feature descriptions was unsatisfactory. First, it was very complicated to create a correct grammar. Only *lazy macros*[6] enabled us to specify that some predicate names must be the same in different steps of specialization. Second, GRENDEL needs much time for each hypothesis test, because it calculates the information gain dependent on further derivations of the grammar and the number of these derivations is very high. We cancelled the experiment after a long time without having learned any rule. But the generated hypotheses look promising.

GRENDEL's grammars have two weaknesses. First, the predicate names itself occur in the grammar. This results in large grammars, which must be changed whenever the set of predicates of the knowledge base changes. The second problem concerns rules where the same predicate multiple occurs in the premise. For a single predicate it is easy to generate a grammar. The following rule, e.g., generates a sequence of chained predicates $\mathsf{p}\ (\mathsf{t}_i,\ \mathsf{t}_{i+1})$:[7]

$$nt\ (t_1) \rightsquigarrow \mathsf{p}\ (\mathsf{t}_1,\mathsf{t}_2),\ nt\ (t_2).$$

But if more than one predicate exist, for which these sequences should be built, the grammar becomes complex, because for each sequence an own grammar branch must be generated:

$$\begin{aligned}
&nt\ (t_1) \rightsquigarrow nt_1\ (t_1).\\
&nt\ (t_1) \rightsquigarrow nt_2\ (t_1).\\
&\vdots\\
&nt\ (t_1) \rightsquigarrow nt_n\ (t_1).\\
&nt_1\ (t_1) \rightsquigarrow \mathsf{p}_1\ (\mathsf{t}_1,\mathsf{t}_2),\ nt_1\ (t_2).\\
&nt_2\ (t_1) \rightsquigarrow \mathsf{p}_2\ (\mathsf{t}_1,\mathsf{t}_2),\ nt_2\ (t_2).\\
&\vdots\\
&nt_n\ (t_1) \rightsquigarrow \mathsf{p}_n\ (\mathsf{t}_1,\mathsf{t}_2),\ nt_n\ (t_2).
\end{aligned}$$

The grammar becomes even more complex, if each recursion is more complex. This problem can be weaken by using lazy macros, too. But the resulting grammars remains complex. This problem could be solved, if it is possible to specify, that an actually occurring predicate should also occur in a further recursion at a specific place.

---

[6] GRENDEL expands a lazy macro $\mathsf{A} \rightsquigarrow \mathsf{B}$ `where` P by proving P and adding a grammar rule $\mathsf{A}\sigma \rightsquigarrow \mathsf{B}\sigma$, where σ is the used substitution, for every possible proof (Cohen, 1993). Since only a finite set of context free rules are added, the language remains context free.

[7] In all examples, variables are written upper case, objects and predicates lower case. Objects and predicates are written sans serif, nonterminals *slanted*, and learnable variables used by RDT are written **bold face**. To distinguish them from implication, arrows in grammar rules are written $\rightsquigarrow$.

BF (Tr, _, Se, T_1, T_2, _)
→ SF (Tr, Se, T_1, T_2, **Movement**)

BF (Tr, _, Se, T_1, T_2, _) &
BF (Tr, _, Se, T_2, T_3, _)
→ SF (Tr, Se, T_1, T_3, **Movement**)

⋮

BF (Tr, _, Se, T_1, T_2, _) &
BF (Tr, _, Se, T_2, T_3, _) &
⋮
BF (Tr, _, Se, T_5, T_6, _) &
BF (Tr, _, Se, T_6, T_7, _)
→ SF (Tr, Se, T_1, T_7, **Movement**)

Figure 3: RDT's rule schemata to learn single sensor feature descriptions

3.3 RDT

The last learning algorithm to be presented, RDT (Kietz and Wrobel, 1992), is based on ideas of W. Emde (Emde, 1987). RDT differs significantly from the two previously described algorithms. The representation formalism of MOBAL (Morik et al., 1993), into which RDT is integrated, is more powerful than FOIL's and GRENDEL's, because negated literals are allowed to occur in the clauses. Additionally, RDT is able to distinguish positive, negative, unknown, and contradictory examples.

3.3.1 General description of RDT

To restrict the hypothesis space, RDT uses second order rule schemata (also called rule models) describing sets of learnable rules. Their predicate variables are successively instantiated during learning. After each instantiation, RDT evaluates the new possibly partially instantiated hypothesis. This evaluation is quite different from the way the other algorithms evaluate the hypotheses. It counts the number of covered positive and negative examples, the number of uncovered positive examples, the total number of examples, and how many additional facts can be derived by this hypothesis. The user can build an arbitrarily complex expression using these items, and hence decide in which case a hypothesis will be accepted.

In contrast to the algorithms previously described, RDT searches through the whole hypothesis space. Only if no further specialization of a hypothesis leads to an acceptable rule, this part of the hypothesis space will be pruned, since completeness is preserved in this case. This case occurs, for example, if the number of covered positive examples of the tested hypothesis is less than the number required to accept the hypothesis.

3.3.2 Learning results in the doorway domain

Since in RDT the structure of the expected rules can be specified and this structure is well known in our domain, RDT learns good rules in both applications. But we need a

large number of rule schemata to solve the learning task. For the first task, we need one schema for each expected pattern length (Figure 3). A second problem concerns the time RDT needed. It orders the rule schemata by θ-subsumption (Plotkin, 1970). This allows pruning, if a hypothesis is already too special to be accepted. But the used schemata are not comparable by θ-subsumption, none of two schemata is more general than the other, because of the chain of time points and there binding to the conclusion. Therefore, exactly the same partial patterns were tested for each schema. Clearly, this behavior slows RDT down. It needed about twenty hours to learn exactly the same rules as GRENDEL did.

In the second learning task, we have more complicated rule schemata. Nevertheless, learning is fast, because the breadth of the hypothesis space, i.e., the number of instantiable predicates, is small. Unfortunately, it is again necessary to enter rule schemata of different length, thus increasing the effort of entering and changing schemata and decreasing the clarity of the whole set of schemata.

4 GRDT: Learning with rule schemata defined by grammars

Testing and comparing the different learning methods yields the following results:

- The information gain heuristic does not guide learning to the desired hypotheses. Additional knowledge about the syntactical structure of the preferred rules is necessary.
- Rule schemata and grammars can be used to describe these structures, but rule schemata are unwieldy because of the static length of each schema. Grammars of GRENDEL are unwieldy because of the predicate names to be fixed in the grammar and the problem of demanding for the same predicate names in different iterations.
- The extended information gain heuristic of GRENDEL may be very costly.
- The choice of the confirmation criterion of RDT is much more flexible than the pruning algorithms used by FOIL or GRENDEL.

Why not combining ideas of GRENDEL and RDT to get a more powerful learning algorithm for our domain? The result is an algorithm using grammars to describe rule schemata instead of defining rules immediately. These rule schemata will be created successively during learning and will be used by an RDT-like algorithm instantiating the predicate variables and testing the generated partial hypotheses.

4.1 Grammars defining sets of rule schemata

We will now introduce the use of grammars defining sets of rule schemata with a sequence of examples. In our first example, rules should be learned where an object must have two specific properties to derive the goal concept. For this task the following rule schema can be used:

$$\mathbf{P}\,(\mathsf{X})\ \&\ \mathbf{Q}\,(\mathsf{X}) \rightarrow \mathbf{Concl}\,(\mathsf{X})$$

RDT tries to find predicate names for the predicate variables **Q** and **P** satisfying the confirmation criterion. This rule schema can be described with a single context free grammar rule:

$\mathbf{Concl}\,(\mathsf{X}) \rightsquigarrow \mathbf{P}\,(\mathsf{X}),\ \mathbf{Q}\,(\mathsf{X}).$

In the case that the object X is not classified to a *fixed* number of concepts to derive the conclusion, rule schemata of different length are necessary:

$$\begin{array}{ll}
\mathbf{P}_1\,(\mathsf{X}) & \rightarrow \mathbf{Concl}\,(\mathsf{X}) \\
\mathbf{P}_1\,(\mathsf{X}) \,\&\, \mathbf{P}_2\,(\mathsf{X}) & \rightarrow \mathbf{Concl}\,(\mathsf{X}) \\
\mathbf{P}_1\,(\mathsf{X}) \,\&\, \mathbf{P}_2\,(\mathsf{X}) \,\&\, \mathbf{P}_3\,(\mathsf{X}) & \rightarrow \mathbf{Concl}\,(\mathsf{X}) \\
\vdots &
\end{array}$$

The maximal necessary length of the rule schemata must be known a priori to learn all suggested rules. The following grammar defines an infinite set of rule schemata of the previous structure:

$\mathbf{Concl}\,(\mathsf{X}) \rightsquigarrow \mathit{it}\,(\mathsf{X}).$
$\mathit{it}\,(\mathsf{X}) \rightsquigarrow \mathbf{P}\,(\mathsf{X}).$
$\mathit{it}\,(\mathsf{X}) \rightsquigarrow \mathbf{P}\,(\mathsf{X}),\ \mathit{it}\,(\mathsf{X}).$[8]

Both previously described learning tasks can also be easily described with GRENDEL's grammars. The only difference is that GRENDEL requires fixing the predicate names to be used in the grammar. Instead of this, GRDT determines them in runtime during learning.

In the next learning task rules describing transitivity should be learned. Therefore, an arbitrary predicate occurs multiple times as premise of the rule:

$$\begin{array}{ll}
\mathbf{P}\,(\mathsf{X}_1,\mathsf{X}_2) & \rightarrow \mathbf{Concl}\,(\mathsf{X}_1,\mathsf{X}_2). \\
\mathbf{P}\,(\mathsf{X}_1,\mathsf{X}_2) \,\&\, \mathbf{P}\,(\mathsf{X}_2,\mathsf{X}_3) & \rightarrow \mathbf{Concl}\,(\mathsf{X}_1,\mathsf{X}_3). \\
\mathbf{P}\,(\mathsf{X}_1,\mathsf{X}_2) \,\&\, \mathbf{P}\,(\mathsf{X}_2,\mathsf{X}_3) \,\&\, \mathbf{P}\,(\mathsf{X}_3,\mathsf{X}_4) & \rightarrow \mathbf{Concl}\,(\mathsf{X}_1,\mathsf{X}_4). \\
\vdots &
\end{array}$$

These rules can be described by GRENDEL's grammars only for a fixed set of predicate names, for each one an own branch must be built. GRDT's grammars allows the user to use predicate variables as arguments of nonterminal literals. Then the predicate variables of the different iterations which are normally different by definition are unified. The following example shows a grammar defining the previously displayed rule schemata:

$\mathbf{Concl}\,(\mathsf{X}_s,\mathsf{X}_e) \rightsquigarrow \mathit{it}\,(\mathbf{P},\mathsf{X}_s,\mathsf{X}_e).$
$\mathit{it}\,(\mathbf{P},\mathsf{X}_p,\mathsf{X}_e) \rightsquigarrow \mathbf{P}\,(\mathsf{X}_p,\mathsf{X}_e).$
$\mathit{it}\,(\mathbf{P},\mathsf{X}_p,\mathsf{X}_e) \rightsquigarrow \mathbf{P}\,(\mathsf{X}_p,\mathsf{X}_2),\ \mathit{it}\,(\mathbf{P},\mathsf{X}_2,\mathsf{X}_e).$

4.2 The algorithm of GRDT

This section briefly describes the new algorithm GRDT. The algorithm starts with the goal predicate as conclusion and generates the first hypothesis depending on the grammar. It then searches the hypothesis space in depth-first order from general to special hypotheses. Every iteration consists of three steps:

1. Specialization of the hypothesis
2. Test of the hypothesis
3. Evaluation of the test

[8] Predicate variables of two grammars rules are regarded to be different.

Specialization The specialization step consists of four alternatives:

- if the hypothesis contains an uninstantiated predicate variable,
  try to find an admissible instantiation of this variable; (backtrackable)
- else if the hypothesis contains a nonterminal symbol,
  try to expand the nonterminal; (backtrackable)
- else if the hypothesis contains a constant to be learned,
  try to find an admissible value for this constant; (backtrackable)
- else
  no further specialization is possible, start backtracking.

Hypothesis test The following items will be counted like RDT does:

- number of covered positive examples;
- number of faultily covered negative examples;
- number of uncovered positive examples;
- number of previously unknown facts that can be derived from the hypothesis;
- total number of examples.

If not all the predicate variables, constants to learn, and nonterminal symbols of the hypothesis are instantiated, the hypothesis is called partial.

Evaluation of the test Depending on the hypothesis test and the parameter set, the hypothesis will be evaluated. The following cases may appear:

- if the hypothesis is too special (i.e., the pruning criterion is satisfied)
  stop specializing and start backtracking;
- else if the hypothesis is partial,
  continue specialization;
- else if the hypothesis is accepted (i.e., the confirmation criterion is satisfied)
  store the rule and start backtracking. Further specialization is not sensible because all more special rules are subsumed by the learned one;
- else
  the hypothesis is too general, but cannot be expanded; start backtracking.

4.3 Learning results

We tested GRDT with our two learning tasks. The used grammars are displayed in Figures 4 and 5. First we will describe the way the grammar to learn single sensor features (Fig. 4) will be expanded and the predicate variables instantiated. The algorithm starts with the first grammar rule:

$$\mathbf{Concl}\,(\mathsf{Tr},\,\mathsf{Se},\,\mathsf{T}_1,\,\mathsf{T}_2,\,\mathbf{M}) \rightsquigarrow \mathit{seq}\,(\mathsf{Tr},\,\mathsf{Se},\,\mathsf{T}_1,\,\mathsf{T}_2,\,\mathbf{M}).$$

Since *seq* is a nonterminal, the first specialization is the expansion of *seq*, using the second rule. This results in the rule schema:

$$\mathsf{BF}\,(\mathsf{Tr},\,_,\,\mathsf{Se},\,\mathsf{T}_1,\,\mathsf{T}_2,\,_) \rightarrow \mathbf{Concl}\,(\mathsf{Tr},\,\mathsf{Se},\,\mathsf{T}_1,\,\mathsf{T}_2,\,\mathbf{M}).$$

Concl (Tr, Se, T_1, T_2, **M**)	$\rightsquigarrow$ *seq* (Tr, Se, T_1, T_2, **M**).
seq (Tr, Se, T_i, T_2, **M**)	$\rightsquigarrow$ **BF** (Tr, _, Se, T_i, T_2, _).
seq (Tr, Se, T_i, T_2, **M**)	$\rightsquigarrow$ **BF** (Tr, _, Se, T_i, T_j, _), *seq* (Tr, Se, T_j, T_2, **M**).

Figure 4: A GRDT grammar to learn single sensor feature descriptions

Concl (Tr, SCl, T_1, T_2, **M**)	$\rightsquigarrow$ *sf_comb* (Tr, **SCl**, T_1, T_2, **M**).
sf_comb (Tr, **SCl**, T_1, T_2, **M**)	$\rightsquigarrow$ **SF** (Tr, Se, T_1, T_2, **M**), sclass (Tr, Se, **SCl**), *sf_comb1* (SF, Tr, **SCl**, T_1, T_2, **M**, Se).
sf_comb (Tr, **SCl**, T_1, T_2, **M**)	$\rightsquigarrow$ **SF** (Tr, Se, T_1, T_i, **M**), sclass (Tr, Se, **SCl**), *sf_comb2* (SF, **Diff**, Tr, **SCl**, **M**, T_1, T_2).
sf_comb1 (SF, Tr, **SCl**, T_1, T_2, **M**, Se_{old})	$\rightsquigarrow$ [ ].
sf_comb1 (SF, Tr, **SCl**, T_1, T_2, **M**, Se_{old})	$\rightsquigarrow$ **SF** (Tr, Se, T_1, T_2, **M**), sclass (Tr, Se, **SCl**), $Se_{old} < Se$, *sf_comb1* (SF, Tr, **SCl**, T_1, T_2, **M**, Se).
sf_comb2 (SF, **Diff**, Tr, **SCl**, **M**, T_{old}, T_{last})	$\rightsquigarrow$ **SF** (Tr, Se, T_1, T_{last}, **M**), sclass (Tr, Se, **SCl**), **Diff** (T_{old}, T_1).
sf_comb2 (SF, **Diff**, Tr, **SCl**, **M**, T_{old}, T_{last})	$\rightsquigarrow$ **SF** (Tr, Se, T_1, T_2, **M**), sclass (Tr, Se, **SCl**), **Diff** (T_{old}, T_1), *sf_comb2* (SF, **Diff**, Tr, **SCl**, **M**, T_1, T_{last}).

Figure 5: A GRDT grammar to learn sensor group feature descriptions

This rule schema contains the predicate variable **BF**, for which an acceptable instantiation must be found. In our domain, this may be one of the basic features, e.g., increasing.

After this instantiation, the rule schema cannot be further specialized, because no further nonterminal, no predicate variable and no constant to be learned exists. So the algorithms starts backtracking using the third rule. Applying this rule results in the rule schema:

BF (Tr, _, Se, T_1, T_2, _) & *seq* (Tr, Se, T_2, T_3, **M**) $\rightarrow$ **Concl** (Tr, Se, T_1, T_e, **M**).

Once again, GRDT tries to instantiate the predicate variable **BF**. If an acceptable predicate could be found (increasing in our example), the hypothesis will be further specialized, because the nonterminal *seq* can be expanded. The following two rule schemata will be generated in the next steps:

increasing (Tr, _, Se, T_1, T_i, _) &
BF (Tr, _, Se, T_2, T_3, _)
→ **Concl** (Tr, Se, T_1, T_3, **M**).

increasing (Tr, _, Se, T_1, T_2, _) &
BF (Tr, _, Se, T_2, T_3, _) &
seq (Tr, Se, T_3, T_e, **M**)
→ **Concl** (Tr, Se, T_1, T_e, **M**).

This specialization step iterates until a hypothesis is found that is acceptable or too special.

For testing GRDT we used a confirmation criterion accepting a hypothesis if it covers at least one positive example. RDT, GRDT and GRENDEL[9] learned exactly the same rules, because RDT and GRDT always searches the complete hypothesis space and in our domain GRENDEL is complete, too. Only the time they needed, displayed in Table 1, differs. FOIL cannot be compared with the other algorithms. Because of its huge effort of memory, we had to restrict the learning to only one of the 28 traces to get a result at all.

The grammar used for the second learning task (Fig. 5) must be explained, too. It consists of two branches, *sf_comb1* and *sf_comb2*, alternatively generated by the nonterminal *sf_comb*. **SF** and **DIFF** are predicate variables, which should be equal in each iteration step. To reach this goal, they are passed to the next iteration as an argument of a nonterminal literal. sclass is a fixed domain predicate which have not to be learned. For the argument variables **M** and **SCL** constants must be learned. This grammar allows to learn two kinds of rules. First, cases are covered where the different sensors of a class sense the same feature at the same time, e.g.:

s_line (Tr, Se_1, T_1, T_2, parallel) & sclass (Tr, Se_1, front_right_corner) &
s_line (Tr, Se_2, T_1, T_2, parallel) & sclass (Tr, Se_2, front_right_corner) &
$Se_1 < Se_2$
→ sg_line (Tr, front_right_corner, T_1, T_2, parallel).

The second kind of rules covers cases, where the different sensors sense the features with a fixed delay, e.g.

s_jump (Tr, Se_1, T_1, T_2, diagonal) & sclass (Tr, Se_1, left_side) &
s_jump (Tr, Se_2, T_3, T_4, diagonal) & sclass (Tr, Se_2, left_side) &
s_jump (Tr, Se_2, T_5, T_6, diagonal) & sclass (Tr, Se_3, left_side) &
d2succ (T_1, T_3) & d2succ (T_3, T_5) &
d2succ (T_2, T_4) & d2succ (T_4, T_6)
→ sg_jump (Tr, left_side, T_1, T_6, diagonal).

Again, we compared the learning results of the algorithms and again we used a confirmation criterion accepting hypotheses, if they cover one positive example for RDT. GRENDEL is applied without use of the noise handlings parts, also accepting hypothesis covering one example. FOIL learned rules senseless for our domain and GRENDEL does not terminate within two days. GRDT learned the same rules as RDT in one third of the time needed by RDT. The time used by the algorithms for solving this task is also displayed in Table 1.

[9] GRENDEL was applied without use of its abilities to handle noise.

	GRENDEL	RDT	GRDT	# Examples	# Facts
s_line	35 min	8 h 42 min	1 h 42 min	718	2847
s_jump	24 min	7 h 10 min	1 h 37 min	206	2335
sg_line	> 2 days	3 h 9 min	1 h 13 min	567	3488
sg_jump	> 2 days	30 min	9 min	225	2534

Table 1: Time used for learning

	FOIL	GRENDEL	RDT	GRDT
repr. formalism	function free Horn logic	see FOIL	function free Horn logic (neg. literals allowed)	see RDT
hypoth. space restrictions	linked clauses	antecedent descr. grammars	rule schemata	grammars defining rule schemata
search control	top-down depth first controlled by heuristic	top-down depth first contr. by grammar and heuristic	top-down breadth first contr. by rule schemata, compl.	top-down depth first contr. by grammar, complete
learning task 1	insufficient result	good result, fast	good result, slow	good result, fast, simple grammar
learning task 2	insufficient result	no result, complex grammar	good result, fast, large set of rule schemata	good result, fast, grammar simpler than GRENDEL's

Table 2: Comparison of the different learning algorithms

5 Conclusion

In this paper, we have presented a learning scenario in the domain of navigation of autonomous mobile robots. We have compared the characteristics of three learning algorithms, FOIL, GRENDEL and RDT (see also Table 2). Each of these algorithms lacks some properties useful for our tasks. So we combined ideas of the algorithms and built a new one, GRDT. We tested GRDT in our doorway domain and indicated that it is able to learn efficiently in large domains, if the user can determine the syntactical structure of the expected rules by a grammar. Thus, the algorithm can be applied to real domains like learning descriptions for operational concepts.

The advantages of GRDT in contrast to GRENDEL are that the syntactical structure can be defined more easily and the confirmation criterion be chosen more flexibly. Nevertheless, there are problems left.

The specialization step, presented in Section 4.2, does not guarantee that the new hypothesis is really more special than the one before. Having grammars defining derivations without any specialization of the derived rules, the algorithm not necessarily terminates. If it terminates, it fulfills a complete search through the hypothesis space defined by the grammar. This behavior of GRDT is derived from the completeness of RDT.

Another problem concerns the way how to exclude multiple search of the same part of the hypothesis space, if different derivations of the grammar yield the same hypotheses.

Acknowledgements

This work is partially funded by the European Community under the project B-Learn II (P7274) and the Ministry for Sciences and Research of the German federal state Nordrhein-Westfalen. I want to thank my advisor Katharina Morik for her ideas, suggestions and critics, Joachim Hertzberg for reading drafts of the paper and all the colleagues at the University of Dortmund, of the machine learning groups at GMD (St. Augustin) and the University of Stuttgart (especially Birgit Tausend) for interesting discussions.

References

Cohen, W. W. (1991). Grammatically biased learning: Learning horn theories using an explicit antecedent description language. Technical report, AT & T Bell Laboratories. (submitted to AIJ).

Cohen, W. W. (1993). Rapid prototyping of ILP systems using explicit bias. Submitted to the 1993 IJCAI workshop on ILP.

Dillmann, R. (1993). Entwicklungstendenzen und Anwendungen symbolischer Lernverfahren in der Robotik. In Herzog, O., Christaller, T., and Schütt, D., editors, *Grundlagen und Anwendungen der Künstlichen Intelligenz – 17. Fachtagung für KI*, Informatik aktuell, pages 28–43, Berlin. Springer-Verlag.

Emde, W. (1987). Non-cumulative learning in METAXA.3. In *Proceedings of IJCAI-87*, pages 208–210, Los Altos, CA. Morgan Kaufman. An extended version appeared as KIT-Report 56, Techn. Univ. Berlin.

Kietz, J.-U. and Wrobel, S. (1992). Controlling the complexity of learning in logic through syntactic and task-oriented models. In Muggleton, S., editor, *Inductive Logic Programming*, chapter 16, pages 335–360. Academic Press, London. Also available as Arbeitspapiere der GMD No. 503, 1991.

Klingspor, V. (1994a). GRDT: Enhancing model-based learning for its application in robot navigation. LS-8 Report No. 5, University of Dortmund, Lehrstuhl Informatik VIII, D-44221 Dortmund.

Klingspor, V. (1994b). Repräsentation operationaler Begriffe. In Bergmann, R., Paulokat, J., Schoeller, A.-M., and Wache, H., editors, *Beiträge zum 8. Workshop Planen und Konfigurieren*, pages 88–97, Kaiserslautern.

Morik, K. and Rieger, A. (1993). Learning action-oriented perceptual features for robot navigation. In Giordana, A., editor, *Workshop notes: Learning Robots of the ECML-93*. Also available as Research Report 3, Univ. Dortmund, Informatik VIII, D-44221 Dortmund.

Morik, K., Wrobel, S., Kietz, J.-U., and Emde, W. (1993). *Knowledge Acquisition and Machine Learning – Theory, Methods, and Applications*. Academic Press, London.

Plotkin, G. D. (1970). A note on inductive generalisation. *Machine Intelligence*, 5:153–163.

Quinlan, R. (1986). Induction of decision trees. *Machine Learning*, 1(1):81–106.

Quinlan, R. (1990). Learning logical definitions from relations. *Machine Learning*, 5:239–266.