A Polynomial Approach to the
Constructive Induction of Structural Knowledge

Jörg-Uwe Kietz

German National Research Centre for Computer
Science (GMD)
Institute for Applied Information Technology
P.O. Box 1240
D-5205 St.Augustin
Germany

Katharina Morik

University Dortmund
Dept. of Computing Science
Artificial Intelligence (LS VIII)
P.O. Box 500 500
D-4600 Dortmund 50
Germany

Abstract

The representation formalism as well as the representation language is of great importance for the success of machine learning. The representation formalism should be expressive, efficient, useful, and applicable. First-order logic needs to be restricted in order to be efficient for inductive and deductive reasoning. In the field of knowledge representation term subsumption formalisms have been developed which are efficient and expressive. In this paper, a learning algorithm, KLUSTER, is described which represents concept definitions in this formalism. KLUSTER enhances the representation language if this is necessary for the discrimination of concepts. Hence, KLUSTER is a constructive induction program. KLUSTER builds the most specific generalization and a most general discrimination in polynomial time. It embeds these concept learning problems into the overall task of learning a hierarchy of concepts.

1   Introduction

Concept learning can be described as inductively forming hypotheses expressed using a hypothesis language such that they deductively cover observations expressed using an observation language. The choice of an appropriate formalism for the hypotheses is crucial for the success of learning. On one hand, the representation formalism should be powerful enough to express at least relations between concepts. On the other hand, it should be efficient with respect to deductive and inductive inference. Moreover, it should be easily understandable so that experts can inspect the results of learning and it should be in the framework of standard representations so that researchers and practicioners from other fields of computer science can easily apply the learning system.

Attribute-value representations have been in the focus of interest for several years as they are easily understandable and applicable. Algorithms for inductive and deductive reasoning in polynomial time have been investigated (e.g., learning monomials (Kearns, 1990)). The expressive power of such representations, however, is very restricted. Therefore, first-order logic moved into the foreground. The advantages of first-order logic are its

expressive power, its understandability, and its applicability in the framework of logic programming. The disadvantage is its complexity. Without restrictions, first-order logic is not efficient - neither for deductive nor for inductive inference. Deduction even in Horn logic is efficiently computable only if the clauses are programmed in a programming language with a fixed evaluation strategy (e.g., Prolog). Induction (e.g., deciding whether there is a hypothesis consistent with the examples) is polynomially computable only for a minimal subset of predicate logic (Kietz, 1992). So, first-order logic has been restricted in several ways for its use in machine learning (e.g., a restricted higher-order logic (Emde, et al., 1983; Wrobel, 1987; Kietz & Wrobel, 1991, Morik, et al., in press) or datalog (Ceri, et al., 1990) as used by FOIL (Quinlan, 1990) or ij-determinate Horn clauses (Muggleton & Feng, 1990)).

An alternative restriction of first-order logic has been developed in the field of knowledge representation: term subsumption formalisms or terminological logics (Brachman & Schmolze, 1985). This representation formalism has a well-defined formal semantics. It is a greatest subset of first-order logic with deduction being still efficiently computable (Donini, et al., 1991). The representation of observations and concepts is easily understandable. Several concepts can be represented by their relations to each other. The formalism is easily applicable and about to become a standard in knowledge representation. However, no learning algorithms which use a term subsumption formalism have been developed until recently. KLUSTER is the first system which learns within this framework (Morik & Kietz, 1989).

In this paper, we first describe the term subsumption formalism (section 2). Then we present the learning algorithm (section 3). Its evaluation with respect to related work and in terms of a theoretical assessment is section 4.

2   The Term Subsumption Formalism Used by KLUSTER

Starting with KL-ONE (Brachman, 1977; Brachman & Schmolze, 1985) an increasing effort has been spent in the development of knowledge representation systems in the framework of term subsumption formalisms (also called terminological logic or description logic), e.g. NIKL (Moser, 1983), KL-TWO (Vilain, 1985), KRYPTON (Brachman, et al., 1985), CLASSIC (Borgida, et al., 1989), BACK (Luck, et al., 1987; Peltason, et al., 1989). Recently, these system have been sucessfully applied to a number of real-world applications, cf. (Peltason, et al., 1991).

The representation formalism corresponds to a rather classical view of concept descriptions, where first a set of superconcepts is referenced and then distinguishing statements are made. For instance, a motor-cycle is defined as a vehicle with exactly two parts, that are wheels. A car is defined as a vehicle with at least three and at most four wheels. The roles of the superconcept $_{vehicle}$ are inherited by the sub-concepts which are distinguished by number restrictions on the $_{part-of}$ role with the concept $_{wheels}$ in its range. The concept representation, i.e. the hypothesis language, is called TBox. A TBox is a semi-lattice with defined meets. Concepts are classified within this structure according to their super-/sub-concept relation. The formalism distinguishes between primitive concepts (concept :<

conditions), where the conditions are necessary, but not sufficient, and defined concepts (concept := conditions), where the conditions are necessary as well as sufficient.

The observations are represented in the so-called ABox. The ABox represents assertions about individual terms. These are classified with respect to their concept membership, i.e. by their link with the TBox.

The main inferences supported by term subsumption formalisms are the classification of concepts and instances into a concept hierarchy. The classification process is formalized by the subsumption relation between concepts. This subsumption goes beyond θ-subsumption in that it respects the overall concept structure. Hence, it is similar to generalized subsumption (Buntine, 1988). The subsumption provides for a partial ordering (generality) which corresponds to logic implication within the term subsumption formalism. Term subsumption formalisms offer an expressiveness in the middle of attribute-value representations and first-order logic. They enhance the quantification of first order logic in that they allow the specification of the minimal and the maximal number of instances for existentially quantified variables. The formal properties of various implementations of term subsumption formalisms are investigated and work on revisions in concept structures has been put forward (Nebel, 1990).

KLUSTER uses a formalism built from a standard set of concept and role forming operators proposed in the literature (e.g. Nebel, 1990; Donini et al., 1991) for representing hypotheses. The syntax follows the representation of the BACK system (Peltason, et al., 1989).

<TBox>	::=	<term-proposition> *
<term-proposition>	::=	<term-restriction>
	|	<term-introduction>
<term-introduction>	::=	<concept-introduction>
	|	<role-introduction>
<concept-introduction>	::=	<concept-name> **:<** <concept>
	|	<concept-name> **:=** <concept>
<role-introduction>	::=	<role-name> **:<** <role>
	|	<role-name> **:=** <role>
<term-restriction>	::=	**disjoint(**<concept-name>+**)**
<concept>	::=	<concept-ref>
	|	**anything**
	|	**nothing**
	|	**all(**<role-ref>**,**<concept-ref>**)**
	|	**atleast(**<integer>**,**<role-ref>**)**
	|	**atmost(**<integer>**,**<role-ref>**)**
<concept-ref>	::=	<concept-name>
	|	<concept-name> **and** <concept-ref>
<role>	::=	<role-ref>
	|	<role> **and** <role>

| **domain**(<concept-ref>)

| **range**(<concept-ref>)

<role-ref>   ::= <role-name>

| **inverse**(<role-name>)

The only difference between this syntax and the one of other TBox formalisms is the restriction in building complex expressions. Only a concept name or a conjunction of concept names is allowed in **all**, **domain** and **range** restrictions. This eases the readability of the concept definitions and helps to avoid problems with terminological cycles (Nebel, 1990). It has no effect on the complexity of the concept learning task. Only role names or the inverse of named roles are allowed in **all**, **atleast** and **atmost** restrictions. Not allowing complex role expressions or defined roles guarantees that the basic algorithm can compute a most specific generalization in polynomial time. If, however, defined roles are needed in order to distinguish between two disjoint concepts, they are introduced via constructive induction. This introduction of defined roles is bounded by parameters such that only polynomially many roles are constructed. So, constructive induction is our way out of the contradiction between the two requirements: expressiveness and efficiency (see section 3.5).

The assertional formalism (ABox) is used as the observation language by KLUSTER. Within the ABox it is expressible that an object belongs to a concept and that two objects are related by a role.

<ABox>   ::= <assertion>+

<assertion>   ::= <object-description>

| <relation-description>

<object-description>   ::= <concept-name>(<object>)

<relation-description>   ::= <role-name>(<object>,<object>)

KLUSTER's formalism has a standard model-theoretic semantics as follows.

Let D, the domain, be any set and E a function, mapping objects to elements of D, concepts to subsets of D and roles to subsets of D×D:

$$
E: \begin{cases} <object> \to D \\ <concept> \to 2^D \\ <role> \to 2^{D \times D} \end{cases}
$$

E is an extension function of a TBox T, if and only if for all $C_i \in$ <concept>, $R_i \in$ <role>, CA $\in$ <concept-ref>, CN $\in$ <concept-name>, RA $\in$ <role-ref>:

$$
\begin{aligned}
E(\textbf{anything}) &= D \\
E(\textbf{nothing}) &= \emptyset \\
E(C_1 \textbf{ and } C_2) &= E(C_1) \cap E(C_2) \\
E(\textbf{all}(RA,CA)) &= \{x \in D \mid \forall y: <x,y> \in E(RA) \to y \in E(CA)\} \\
E(\textbf{atleast}(n,RA)) &= \{x \in D \mid |\{ y \in D \mid <x,y> \in E(RA) \}| = n \} \\
E(\textbf{atmost}(n,RA)) &= \{x \in D \mid |\{ y \in D \mid <x,y> \in E(RA) \}| = n \} \\
E(R_1 \textbf{ and } R_2) &= E(R_1) \cap E(R_2)
\end{aligned}
$$

$$E(\textbf{domain}(CA)) \;=\; E(CA) \times D$$

$$E(\textbf{range}(CA)) \;=\; D \times E(CA)$$

$$E(\textbf{inverse}(RN)) \;=\; \{<x,y> \in D \times D \mid <y,x> \in E(RN)\}$$

A Pair <D,E>, where D is a domain and E is an extension function, is a model of a TBox T and an ABox A, if and only if:

$$E(CN) \;\subseteq\; E(C), \text{ for all } CN \textbf{:<} C \quad T$$

$$E(CN) \;=\; E(C), \text{ for all } CN \textbf{:=} C \quad T$$

$$E(RN) \;\subseteq\; E(R), \text{ for all } RN \textbf{:<} R \quad T$$

$$E(RN) \;=\; E(R), \text{ for all } RN \textbf{:=} R \quad T$$

$$E(CN_i) \cap E(CN_j) \;=\; \emptyset, \text{ for all } CN_i, CN_j \text{ with } i\;?\;j, \text{ and } 1 = i,j = n \text{ in a}$$

$$\textbf{disjoint}(CN_1, CN_2, \ldots, CN_n) \in T.$$

$$E(t_1) = E(t_2) \;\leftrightarrow\; t_1 = t_2, \text{ for all objects } t_1, t_2 \text{ in A}$$

$$E(t) \;\in\; E(C), \text{ for all } C(t) \in A$$

$$<E(t_1),E(t_2)> \;\in\; E(R), \text{ for all } R(t_1,t_2) \in A$$

The syntax and the model-theoretic semantics together define a logic. Before we define the inferences performed by the terminological reasoners, let us give an example of a well-formed concept definition and its equivalent in first-order logic with equality:

motorcycle **:=** vehicle **and all**(base_part, wheel) **and atleast**(2, base_part) **and atmost**(2, base_part)

$\forall x$ (motocycle(x) $\rightarrow$ vehicle(x) $\land$ $\forall y$ (base_part(x,y) $\rightarrow$ wheel(y)) $\land$

$\quad \exists y1\ y2$ (base_part(x,y1) $\land$ base_part(x,y2) $\land$ y1 ? y2 $\land$

$\quad\quad \neg\exists y3$ (base_part(x,y3) $\land$ y1 ? y3 $\land$ y2 ? y3)))

Now, let us precisely define what we mean by subsumption, equivalence, disjointness and incoherence of terms within a TBox T, by entailment of assertions from a TBox T and an ABox A and inconsistency of a TBox T and an ABox A.

- Within a TBox T a term t *is subsumed by* a term t', written $t =_T t'$, iff for every model <D,E> of T it holds that $E(t) \subseteq E(t')$.

- Within a TBox T two terms t and t' are *equivalent*, written $t \; \tilde{}\;_T t'$, iff for every model <D,E> of T it holds that $E(t) = E(t')$.

- Within a TBox T two terms t and t' are *disjoint*, iff for every model <D,E> of T it holds that $E(t) \cap E(t') = \emptyset$.

- Within a TBox T a term t is *incoherent*, iff for every model <D,E> of T it holds that $E(t) = \emptyset$.

- An assertion f is entailed by a TBox T and an ABox A, written $A \models_T f$, iff for every model <D,E> of T and A it holds that $E(t) \in E(C)$ if $f = C(t)$, or $<E(t_1),E(t_2)> \in E(R)$, if $f = R(t_1,t_2)$.

- A TBox T and an ABox A are *inconsistent*, iff there exists no model <D,E> of T and A.

Note, that subsumption as defined above is a semantic relation like implication or generalized subsumption (Buntine, 1988), which takes into account background knowledge. It is not a pure syntactic relation like θ-subsumption (Plotkin, 1970).

In our TBox formalism we can compute the disjointness and incoherence using subsumption or equivalence alone:

- $t \sim_T t'$, iff $t =_T t'$ and $t' =_T t$

- t is *incoherent*, iff $t =_T$ **nothing**

- t and t' are *disjoint*, iff $(t$ **and** $t') =_T$ **nothing**

- $t =_T t'$, iff $t \sim_T (t$ **and** $t')$

It is known (Donini et al.,1991), that subsumption between two concepts with respect to a TBox T in the formalism above can be decided in polynomial time, if T does not contain any role introductions and all disjoint restrictions contain only names of primitive concepts (concept names introduced by <concept-name> **:<** <concept>). It is also known, that the formalism cannot be extended without losing the polynomial time decidability or completeness[1]. Thus, the learning result of KLUSTER cannot be classified completely polynomially, if constructive induction has introduced new roles.

3   KLUSTER

In this section, we present the system KLUSTER which is an inductive learning system for constructing a concept structure in the term subsumption formalism presented in the last section. A deductive reasoning system (e.g. BACK, CLASSIC) for this term subsumption formalism is assumed to be given.

The overall learning task of KLUSTER is:

Given a set of assertions in the ABox (the examples), and an empty TBox. If a partially filled TBox (the background knowledge) is given, the assertions are assumed to be saturated by entailment. Clearly, ABox and background knowledge must be consistent.

Goal: A TBox, i.e. a hierarchy of concept definitions, organizing the factual knowledge such that the concept definitions of the TBox are true in the minimal model of the ABox. The TBox can be used for inferring by entailment further descriptions about objects newly entered into the ABox.

We will use a domain of side-effects of drugs for illustrating our approach. The following set of assertions is given as input to KLUSTER:

contains(aspirin,asa)	contains(anxiolit,oxazepun)	contains(placo,nhc)
contains(alka-seltzer,asa)	contains(anxiolit,finalin)	contains(placo,sugar)
contains(alka-seltzer,nhc)	contains(adolorin,phenazetin)	affects(asa,headache)
contains(adumbran,coffein)	contains(adolorin,prophymazon)	affects(oxazepun,stress)
contains(adumbran,oxazepun)	contains(adolorin,nhc)	affects(finalin,stress)

affects(prophymazon,headach)	anodyne(aspirin)	active(phenazetin)
affects(phenazetin,headache)	anodyne(alka-seltzer)	active(oxazepun)
placebo(placo)	anodyne(adolorin)	add_on(nhc)
combidrug(anxiolit)	sedative(anxiolit)	add_on(coffein)
combidrug(adolorin)	sedative(adumbran)	add_on(sugar)
monodrug(aspirin)	active(asa)	excitement(stress)
monodrug(alka-seltzer)	active(finalin)	pain(headache)
monodrug(adumbran)	active(prophymazon)	

These are the given observations. No background knowledge is provided. Note, that the relation $_{contains}$ is a n to m relation. The first step of KLUSTER is to compute a basic taxonomy which is a hierarchy of primitive concepts and roles based on set inclusion between the known extensions of concepts and roles. The computed basic taxonomy is used for structuring the overall task of KLUSTER into a set of concept learning problems. The concepts which KLUSTER tries to define are taken top-down and breadth-first from the basic taxonomy. This search strategy is implemented by an agenda of concept learning problems. Each agenda entry is a cluster of concepts (called MDC, mutually disjoint concepts) which have the same superconcept and which are mutually disjoint. This enables KLUSTER to define concepts not in isolation, but in the context in which they occur.

A concept learning problem of KLUSTER is to build discriminating definitions for the concepts of a MDC. A definition is discriminating if the number of misclassified examples is lower or equal than a given threshold ($F_{MDC} = \varepsilon$). To test if such a discriminating definition exists, KLUSTER first builds *most specific generalizations* (MSGs) for all examples of a concept. If the available concepts and roles are not sufficient for a discriminating characterization, the representation language is expanded. This means, more complex expressions are only built if simpler ones are not sufficient. The introduction of new concepts and roles is bounded by two parameters (rlength and refinement, see section 3.5). As the concept learning goal is to find discriminating concept definitions for the concepts of a MDC, the best (most predictive) definition is the *most general discrimination* (MGD). Therefore, KLUSTER generalizes all discriminating MSGs to MGDs. This two step approach of learning concepts is preferred to learning MGDs directly, as the MSGs have some useful properties that MGDs do not have:

- The MSG is unique in our formalism and simple to built (see section 3.3.1).

- If the MSG is not discriminating, then no concept expression covering all positive examples is discriminating.

- The MSG is useful for a possible extension of KLUSTER to incremental learning as $msg(\{o_1,o_2,\ldots,o_n\}) = msg(\ldots msg(o_1,o_2)\ldots,o_n)$.

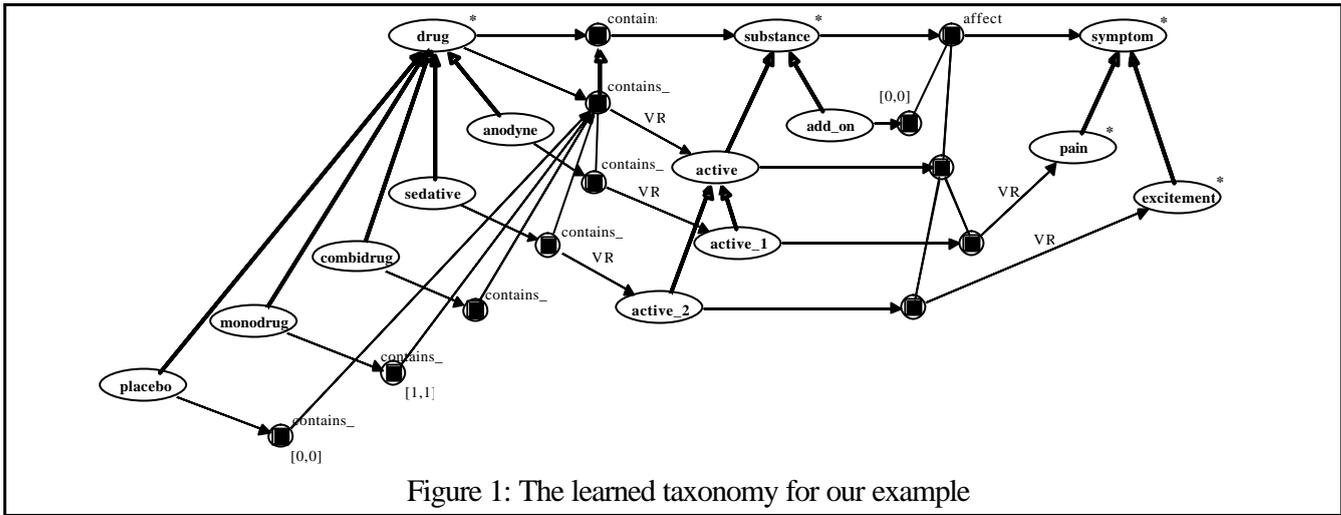In our example the following concept definitions (MGD's) are learned (see Figure 1):

Figure 1: The learned taxonomy for our example

active	:=	substance **and atleast**(1, affects)
add_on	:=	substance **and atmost**(0, affects)
placebo	:=	drug **and atmost**(0, contains_active)
monodrug	:=	drug **and atleast**(1, contains_active)
		and atmost(1, contains_active)
combidrug	:=	drug **and atleast**(2, contains_active)
anodyne	:=	drug **and all**(contains_active, active_1)
		and atleast(1, contains_active)
sedative	:=	drug **and all**(contains_active, active_2)
		and atleast(1, contains_active)

The above definitions use the following defined concepts and roles which are introduced by KLUSTER's constructive induction:

contains_active	:=	contains **and range**(active)
active_1	:=	active **and all**(affects, pain)
active_2	:=	active **and all**(affects, excitement)

The overall method of KLUSTER is summarized in Table 1.

In section 3.1. we show how KLUSTER aggregates objects into primitive concepts and how the basic taxonomy of these primitive concepts is built. In section 3.2 we describe the computation of MDCs and the agenda mechanism. MSGs and the evaluation functions are defined in section 3.3. Section 3.4 presents the generalisation from characterizations (MSGs) to definitions of concepts (MGD). The constructive induction of new concepts and relations for defining a concept is described in section 3.5.

```
learn_TBox(ε, maxrefinement, maxrlength):

  begin

    compute_basic_taxonomy                                          ; building the basic taxonomy

    initialize_agenda

    repeat

      select_best_active_MDC(mdc, refinement, rlength)

      for all c   mdc

        compute_and_store_MSG(c)                                    ;building MSGs

      if FMDC(mdc) = ε                                              ;evaluating MDC

      then   set_definable_MDC(mdc)

      else   if refinement > maxrefinement   rlength = maxrlength

             then   set_undefinable_MDC(mdc)

             else   build_refinements(mdc, refinement, rlength)     ;constructive induction of concepts, roles

    until all m dc   agenda: definable_MDC(mdc) Δ undefinable_MDC(mdc)

    for all definable_MDC(mdc)

      for all c   mdc

        compute_and_store_MSG(c)                                    ;building MSGs with enhanced language

        generalize_MSG_to_MGD(c)                                    ;building MGDs

    delete_all_refinements_not_used_in_MGDs

  end
```

Table 1: An outline of the learning algorithm

3.1   Building the Basic Taxonomy

As the first step of learning, KLUSTER aggregates objects of the ABox into primitive concepts of the TBox. Objects which occur in the ABox as argument of a one-place predicate are collected as the known extension of a primitive concept in the TBox named by the predicate symbol. Tuples of objects, which occur in a two-place predicate of the ABox are interpreted as the known extension of a primitive role in the TBox named by the predicate symbol. The domains and ranges of the primitive roles are also determined. The domain of a role is the set of objects occurring at the first place of the role. The range of a role is the set of objects occurring at the second place of the role s.

Let us describe this more formally. Let ext be an extension function as defined in section 2,

$$\text{ext:} \begin{cases} <object> \rightarrow <object> \\ <concept> \rightarrow 2^{<object>} \\ <role> \rightarrow 2^{<object> \times <object>} \end{cases}$$

where $\text{ext}_{<object>}$ is the identity function between the objects in the ABox, i.e. the objects of the ABox are the domain of the interpretation. Then the pair $<<object>,\text{ext}>$ is a minimal model of the given TBox and ABox, if

TBox and ABox are consistent and the ABox is complete with respect to the given TBox. This is always the case if the TBox is empty, i.e. if no background knowledge is given.

The system then builds root concepts as the union of all extensionally overlapping domains and ranges of roles and primitive concepts. The root concepts are similar to the sorts or types which other learning systems (e.g. FOIL (Quinlan, 1990), GOLEM (Muggleton and Feng, 1990) take as input.

Then, the primitive concepts are arranged into a hierarchy based on set inclusion of the extensions. This means, the subsumption relationships valid in the minimal model <<object>,ext> are induced. As subsumption is a partial ordering, a minimal representation consist of the direct subsumptions. KLUSTER uses standard algorithms to compute the direct subsumption from subsumption.

Disjointness of primitive concepts is also determined based on the extensions, i.e. all disjoint relationships valid in the minimal model <<object>,ext> are induced. As disjoint restrictions are inherited along subsumption, the system computes the minimal set of disjoint restrictions, necessary to infer the inherited ones. The disjointness of anodyne and pain for example can be inferred from the disjointness of drug and symptom, as drug subsumes anodyne and symptom subsumes pain.

In our example, root concepts (the predecessors of **anything**) and primitive concepts are[2]:

drug	:< **anything**,	ext(drug) = {adolorin,adumbran,alka_seltzer,anxiolit,aspirin,placo}
placebo	:< drug,	ext(placebo) = {placo},
monodrug	:< drug,	ext(monodrug) = {adumbran,alka_seltzer,aspirin},
combidrug	:< drug,	ext(drug) = {adolorin,anxiolit},
anodyne	:< drug,	ext(anodyne) = {adolorin,alka_seltzer,aspirin},
sedative	:< drug,	ext(sedative) = {adumbran,anxiolit},
substance	:< **anything**,	ext(substance)= {asa,coffein,finalin,nhc,oxazepun,phenazetin,prophymazon,sugar}
active	:< substance,	ext(active) = {asa,finalin,oxazepun,phenazetin,prophymazon},
add_on	:< substance,	ext(add_on) = {coffein,nhc,sugar},
symptom	:< **anything**,	ext(symptom) = {bellyache,headache,stress},
pain	:< symptom,	ext(pain) = {bellyache,headache},
excitement	:< symptom,	ext(excitement) = {stress}.

All concepts are primitive, i.e. they still need to be defined. The minimal set of disjoint restrictions is the following:

disjoint(drug, substance),	**disjoint**(drug, symptom),	**disjoint**(substance,symptom),
disjoint(placebo, monodrug),	**disjoint**(placebo, combidrug),	**disjoint**(monodrug, combidrug),
disjoint(placebo, anodyne),	**disjoint**(placebo, sedative),	**disjoint**(anodyne, sedative),
disjoint(active, add_on),	**disjoint**(pain, stress)	

The roles of the basic taxonomy are:

contains   :< **domain**(drug) **and range**(substance), with

ext(contains)={ adolorin,nhc),(adolorin,phenazetin),(adolorin,prophymazon),(adumbran,coffein),

(adumbran,oxazepun),(alka_seltzer,asa),(alka_seltzer,nhc),(anxiolit,finalin),

Figure 2: The basic taxonomy for our example
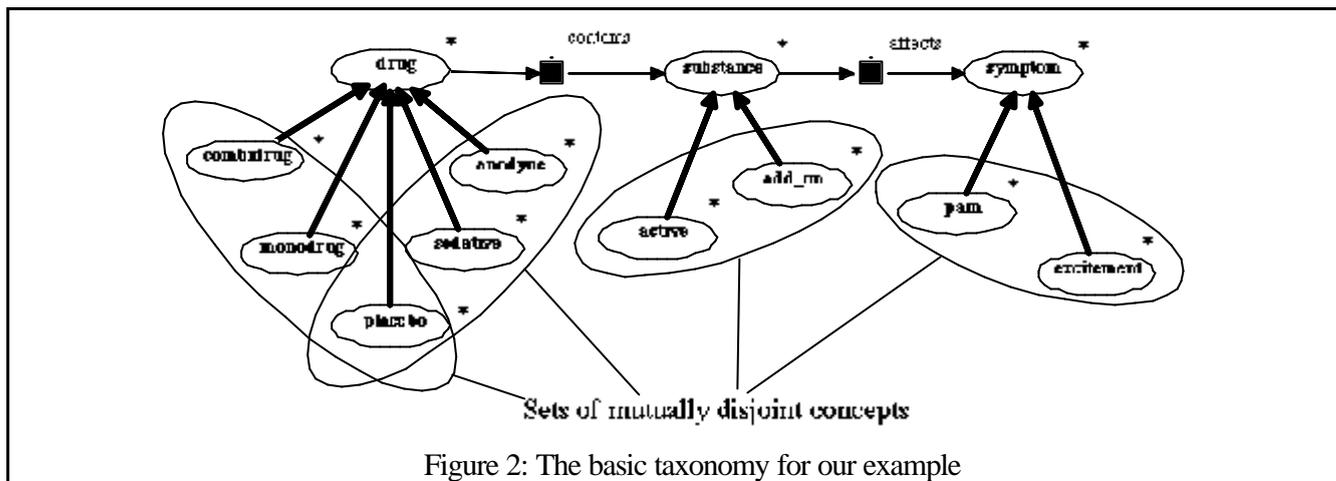
(anxiolit,oxazepun),(aspirin,asa),(placo,nhc),(placo,sugar)}

affects   :< **domain**(active) **and range**(symptom), with

ext(affects) = {(asa,headache),(finalin,stress),(oxazepun,stress),(phenazetin,bellyache),(prophymazon,headache)}

Figure 2 shows the basic taxonomy which is the result of the first step of KLUSTER for our example.

3.2   The Concept Learning Problems of KLUSTER

Having computed the basic taxonomy, KLUSTER sets up concept learning problems. The concept learning goal is to define primitive concepts preserving the discrimination from their sister concepts. Sister concepts are the mutually disjoint subconcepts of a common superconcept. They are called *mutually disjoint concepts*, MDC. There can be more than one MDC for a superconcept. This is the case, if a concept can be specialized with respect to diverse aspects. For instance, in our example domain, drugs are classified with respect to the combinations of substances into monodrugs, which consist of only one effective substance, combidrugs, which consist of more than one effective substance, and placebos, which consist of no effective substance. These three primitive concepts together form a MDC. Drugs are also classified according to the effect they have on the human body into painkillers (anodyne), stress removers (sedative), and placebos (with no effect at all). These primitive concepts together form another MDC. Both classifications are appropriate, i.e. a cross classification of drugs is desired. Both MDCs are to be defined. To define concepts of a MDC such that no instance is covered by more than one concept of a MDC is the concept learning problem of KLUSTER.

In order to set up concept learning problems, MDCs are first built for each root concept. Pairwise disjointness is already computed. Mutual disjointness of several primitive concepts is computed by first establishing the complementary list of non-disjoint pairs. Then, a list of all primitive concepts, which occur in any of the disjoint pairs is split according to the non-disjoint pairs. The list of non-disjoint pairs is checked exhaustively. The lists resulting from splitting are the MDCs. Computing these maximal sets of mutually disjoint concepts is computational expensive in the worst case. The computational costs for m MDC is $m \log_2 m$. However, for n concepts there are in the worst case $\binom{n}{n/2}$ different MDCs.[3]

MDCs are then ordered on an agenda. The agenda determines a top-down, breadth-first order of concepts to be defined. In terms of the graphic representation of the concept structure, MDCs are set up as concept learning problems from left to right, one level at a time. In our example, there are two MDCs for $_{drugs}$, one for $_{substances}$, and one for $_{symptoms}$ at the beginning:

MDC_1: {placebo, monodrug, combidrug}

MDC_2: {placebo, anodyne, sedative}

MDC_3: {active, add_on}

MDC_4: {pain, excitement}

As root concepts are not defined, the concept structure will never consist completely of defined concepts. Most often, not all of the concept learning tasks are accomplished by KLUSTER. Some concepts remain primitive. They assist the definition of other concepts without being defined themselves. The reason for this is that we want to prohibit circular definitions. Hence, there are always concepts which are used in **all** restrictions of concept definition but are not defined themselves. In the graphic representation of our example, these are the right-most concepts. Only if the right-most concepts are definable using number restrictions, they can be defined avoiding the pitfall of circular definitions.

3.3   Characterizing Concepts

The definition of concepts is performed in three steps by KLUSTER. First, concepts are characterized by the *most specific generalization*, MSG. Then, the characterization is evaluated. Finally, the characterization is further generalized to become the *most general discrimination*, MGD, which is the definition of the concept. A characterization selects all relevant roles where a definition selects the most discriminating ones among the relevant roles. We first describe the induction resulting in a most specific generalization. Then we describe the evaluation of all concept characterisations of a MDC. Finally (section 3.4), we describe how the acceptable MSGs are further generalized to a most general discrimination.

3.3.1   Building the Most Specific Generalizations (MSG)

The MSG of a concept $_c$ is a set of most specific concept expression $_{ce}$, such that $_c :_< ce$ is true in the minimal model <<object>,ext>, or more formally:

$$\text{MSG}(c) = \{ \ ce \in <concept> \ | \ ext(c) \subseteq ext(ce)$$
$$\neg \ \exists \ ce_2 \in <concept>: ext(c) \subseteq ext(ce_2) \subset ext(ce) \ \}$$

In our term subsumption formalism there are several semantically equivalent concept expressions, which are not syntactically identical. Fortunately, there exists a unique normalized concept expression for every equivalence class in our term subsumption formalism. To determine this unique normalized expression, let us look at the possible concept expressions in our formalism. A concept expression is a conjunction (**and**) of concept names, **all**-, **atleast**-, and **atmost**-restrictions. Clearly, the **and** operator is commutative, associative, and idempotent. This

means, the order of the restrictions is irrelevant. Now, let us look at the particular restrictions possible in concept expressions and their normalization:

- Concept names: these are the superconcepts of the concept. Normalization selects the direct predecessors of the concept in the subsumption partial ordering.

- **all**-restrictions: these are the expressions of the form: **all**(r, c_1 **and** ...**and** c_n), where r is a role name or the inverse of a named role, and the c_i are concept names. Two **all**-restrictions with the same role r in a concept expression: **all**(r, c_1 **and** ...**and** c_n) **and all**(r, c_{n+1} **and** ...**and** c_{n+m}), are normalized to the equivalent expression: **all**(r, c_1 **and** ...**and** c_n **and** c_{n+1} **and** ...**and** c_{n+m}). The c_1 **and** ...**and** c_n in **all** restrictions are also normalized based on the subsumption partial ordering. If the range of a role r is C, an **all**(r, C) restriction is equivalent to **anything** and can be dropped.

- **atleast**-restrictions: these are the expressions of the form: **atleast**(l,r), where r is a role name or the inverse of a named role. Two **atleast**-restrictions with the same role r : **atleast**(l_1,r) **and atleast**(l_2,r), are normalized to the equivalent expression: **atleast**(maximum(l_1,l_2),r). An **atleast**(0, r) restriction is equivalent to **anything** and can be dropped.

- **atmost**-restrictions: these are the expressions of the form: **atmost**(m,r), where r is a role name or the inverse of a named role. Two **atmost**-restrictions with the same role r **atmost**(m_1,r) **and atmost**(m_2,r), are normalized to the equivalent expression: **atmost**(minimum(m_1,m_2),r). An **atmost**(0,r) restriction can be dropped, if the domain of r and a superconcept of the concept under normalization are necessarily disjoint.

It is clear, that any concept expression normalized as above contains for every role name r at most one **all**-restriction, at most one **atleast**-restriction and at most one **atmost**-restriction for r and for **inverse**(r). This means, that the size of any concept expression in our formalism is polynomially bound in the number of concept names and role names and the coding size of the greatest integer used in the TBox.

Another important consequence of the normalization is that there are always at most finitely many concept expressions not within subsumption order. By definition, the MSG contains only expressions which are not in a subsumption relation. So, the MSG contains at most finitely many concept expressions. Suppose, that there would be two different concept expressions within the MSG. This implies, that also the conjunction of both is a generalization of all examples. Clearly the conjunction of two concepts is more specialized than the concepts it is built from. This implies, that both are not MSGs, but the conjunction is. As the conjunction can be built, this will be the MSG. This proves that the MSG of a concept is unique (under equivalence) in our term subsumption formalism.

From this it becomes clear how the unique, normalized MSG of a concept c is constructed:

The superconcepts of c are already computed in the basic taxonomy.

For each role name r,
    If the domain of r is not disjoint to a superconcept of c add
        all(r, c_1 **and** ...**and** c_n), for all smallest c_i,

which fulfil $\{ y \mid (x,y) \in r \wedge x \in ext(c) \} \subseteq ext(c_i)$

atleast(l,r), where $l = minimum(\mid \{ y \mid (x,y) \in ext(r) \} \mid$, for all $x \in ext(c)$)

atmost(m,r), where $m = maximum(\mid \{ y \mid (x,y) \in ext(r) \} \mid$, for all $x \in ext(c)$)

to the MSG of c.

If the range of r is not disjoint to a superconcept of c add

all(**inverse**(r), c_1 **and** …**and** c_n), for all smallest c_i,

which fulfil $\{ y \mid (y,x) \in r \wedge x \in ext(c) \} \subseteq ext(c_i)$

atleast$(l,$**inverse**$(r))$, where $l = minimum(\mid \{ y \mid (y,x) \in ext(r) \} \mid$, for all $x \in ext(c)$)

atmost$(m,$**inverse**$(r))$, where $m = maximum(\mid \{ y \mid (y,x) \in ext(r) \} \mid$, for all $x \in ext(c)$)   to

the MSG of c.

In the example, the characterizations for the concepts are:

MSG(placebo) = drug **and all**(contains, add_on) **and atleast**(2,contains) **and atmost**(2,contains)

MSG(monodrug) = drug **and atleast**(1,contains) **and atmost**(2,contains)

MSG(combidrug) = drug **and atleast**(2,contains) **and atmost**(3,contains)

MSG(sedative) = drug **and atleast**(2,contains) **and atmost**(2,contains)

MSG(anodyne) = drug **and and atleast**(1,contains) **and atmost**(3,contains)

MSG(active) = substance **and atleast**(1, affects) **and atmost**(1, affects) **and**

　　　　　　atleast(1,**inverse**(contains)) **and atmost**(2,**inverse**(contains))

MSG(add_on) = substance **and atmost**(0, affects) **and**

　　　　　　atleast(1,**inverse**(contains)) **and atmost**(2,**inverse**(contains))

MSG(pain) = symptom **and atleast**(1,**inverse**(affects)) **and atmost**(2,**inverse**(affects))

MSG(excitement) = symptom **and atleast**(2,**inverse**(affects)) **and atmost**(2,**inverse**(affects))

3.3.2　Evaluating MSGs in Context

The evaluation of most specific generalizations is performed in the context of the MDC. The purpose of the evaluation is to accept or reject a concept characterization. As opposed to decision tree induction or conceptual clustering, the evaluation is not concerned with the selection of the best characterization among other alternatives because there is always exactly one MSG for a concept. If a MSG does not get a perfect evaluation, we know that using the given representational entities there can be no acceptable generalization. This is the criterion for introducing new concepts or relations (see 3.5).

It is evaluated,

- how well the overall MDC is characterized, i.e. the MDC failure,

- how well a MSG separates a concept from the other concepts of the same MDC, i.e. the MSG failure,

- how much the restrictions of a particular role within the MSGs of the concepts of the MDC contributes to the separation within a MDC, i.e. the role failure,

• how well the restrictions of a particular role within the MSG describe a concept, i.e. the restrictions failure.

The overall evaluation of characterizations of a MDC, the MDC failure, is simply the sum of all MSG failures divided by the number of concepts of the MDC. The MSG failure counts how many instances of a MSG are also instances of other concepts of the MDC and normalizes this number by dividing it by the number of objects of the MDC. The formula for the MSG failure is:

$$F_{MSG}(c, mdc) = \frac{\left| ext(MSG(c)) \cap \bigcup_{mc \in mdc, c \neq mc}^{mc} ext(mc) \right|}{\left| \bigcup_{mc \neq mdc}^{mc} ext(mc) \right|}$$

The role failure measures the contribution of a role R to discriminating the concepts of a MDC. It sums up all restrictions failures for a role and normalizes this by dividing the sum by the number of objects of the MDC. Hence, the basis for the role failure is the restrictions failure. The restrictions failure counts how many objects of another concept of the MDC are covered by the restrictions in the MSG using only a particular role. This count is then divided by the number of objects of the MDC. We use this for formalizing the restriction failure :

$$F_R(r, c, mdc) = \frac{\left| ext(rr(r,c)) \cap \bigcup_{mc \in mdc, c \neq mc}^{mc} ext(mc) \right|}{\left| \bigcup_{mc \neq mdc}^{mc} ext(mc) \right|}$$

where rr(r,c) := **all**(r,vc) **and atleast**(l,r) **and atmost**(m,r) within the MSG(c).

In the example, for MDC_1 the extensions of $placebo$, $monodrug$, and $combidrug$ according to their characterization are:

ext (rr(contains, placebo)) = ext(MSG(placebo)) = { placo }

ext (rr(contains, monodrug)) = ext(MSG(monodrug)) = { alka-seltzer, aspirin, adumbran, <u>placo</u> }

ext (rr(contains, combidrug)) = ext(MSG(combidrug)) = {adolorin, anxiolit, <u>alka-seltzer, adumbran, placo</u>}

The underlined objects are intersections with other concepts of MDC_1. They are misclassified by the characterization. The restrictions failures are:

F_R (contains, placebo, MDC_1) $= F_{MSG}$(placebo,MDC_1) = 0

F_R (contains, monodrug, MDC_1) $= F_{MSG}$(monodrug,MDC_1)   = 1/6

F_R (contains, combidrug, MDC_1) $= F_{MSG}$(combidrug,MDC_1)   = 3/6

The role failure for $contains$ is the sum of the restrictions failures divided by the number of concepts of MDC_1. This is equal to the overall MDC failure as $contains$ is the only role involved:

F_{RMDC} (contains, MDC_1) = F_{MDC} (MDC_1) = 4/18

For MDC_2 the extensions of $placebo$, $sedative$, and $anodyne$ according to their characterization are:

ext (rr(contains, placebo)) = ext(MSG(placebo)) = { placo }

ext (rr(contains, sedative)) = ext(MSG(sedative)) = { adumbran, anxiolit, <u>alka-seltzer</u>, <u>placo</u> }

ext (rr(contains, anodyne)) = ext(MSG(anodyne)) = {adolorin, alka-seltzer<u>,</u> aspirin, <u>adumbran, anxiolit, placo</u>}

The restrictions failures are:

F_R (contains, placebo, MDC_1) = F_{MSG}(placebo,MDC_1) = 0

F_R (contains, sedative, MDC_1) = F_{MSG}(sedative,MDC_1) = 2/6

F_R (contains, anodyne, MDC_1) = F_{MSG}(anodyne,MDC_1) = 3/6

The role failure for $_{contains}$ as well as the overall MDC failure of $_{MDC\_2}$ is:

F_{RMDC} (contains, MDC_1) = F_{MDC} (MDC_1) = 5/18

These failures are rather high. It shows that the characterizations are not specific enough. But, the most specific generalizations are already built. Using the concepts and roles given, there are no more specific characterizations. Thus, either new concepts and/or roles can be built, which contribute to a better discrimination, or the MDCs must be marked as undefinable and taken away from the agenda of concept learning tasks. When describing the introduction of new concepts and relations (3.5), we shall come back to these examples.

$_{MDC\_1}$ uses only one role for its concepts. Thus, the role failure is the same as the MDC failure. $_{MDC\_2}$ is more interesting as there are two roles involved in the characterizations of each concept:

ext(RR (affects, active))   = {phenacetin, asa, prophymacon, oxacepun,finalin}

ext(RR (inverse (contains), active))   = {phenacetin, asa, prophymacon, oxacepun,finalin, <u>sugar, coffein, nhc</u>}

ext (RR (affects, add_on))   = {sugar, coffein, nhc}

ext (RR (inverse (contains), add_on))   = {<u>phenacetin, asa, prophymacon, oxacepun, finalin</u>, sugar, coffein, nhc}

Using the role $_{affects}$ gives no misclassification for $_{active}$ nor for $_{add\_on}$. The restrictions failure is 0 in both cases. Characterizing the concepts by the inverse of the role $_{contains}$, however, makes for a restrictions failure of 3/8 for $_{active}$ and 5/8 for $_{add\_on}$. The overall role failure is 0 for $_{affects}$ and 1/2 for the inverse of $_{contains}$. As the MSG failure measures the failure from the conjunction of the restrictions for each concept, it is 0, too, for both $_{active}$ and $_{add\_on}$. Therefore, also the MDC failure is 0. From the comparison of the role failure with the MDC failure it becomes clear that the concepts can be defined without using the relation $_{inverse(contains)}$.

F_{RMDC} (inverse(contains), MDC_2) = 1/2, but

F_{MDC} (MDC_2) = 0

This information will be used by the shift from characterizations to definitions.

3.4   The Shift from Characterizations to Definitions or Building the MGD

Definitions of concepts are intended to cover more than the observed objects but no objects that are classified into a disjoint concept. They are supposed to be as short as possible, and they should all use the same roles, if possible. Finally, they should not be cyclic.

The generalization of MSGs to MGDs is performed by dropping and generalizing restrictions as long as the discrimination is preserved. In our example, the MGDs for $_{active}$ and $_{add\_on}$ are

MGD(active) = substance **and atleast**(1, affects)

MGD(add_on) = substance **and atmost**(0,affects)

All the restrictions involving the inverse relation $_{contains}$ are dropped and for $_{active}$ the **atmost**-restriction of $_{affects}$ is also dropped.

When no further restriction can be dropped, KLUSTER tries to generalize the restrictions. **all**-restrictions are generalized by generalizing the concept reference, i.e. replacing a concept by its superconcepts or simply dropping a conjunct. **atleast**-restrictions are generalized by decreasing the number, **atmost**-restrictions by increasing the number. KLUSTER generalizes as long as no misclassification is introduced.

There can be several MGDs. In principle, from n relevant restrictions m restrictions are sufficient for discrimination, i.e. in the worst case there are $\binom{n}{n/2}$ different minimal concept definitions. But, KLUSTER enters the first found MGD into the concept structure, instead of looking for the best one. Therefore, no combinatorial explosion can occur. The algorithm drops a restriction and checks whether the remaining definition leads to a misclassification. If no misclassification occurs the restriction is dropped, otherwise it is kept. Then the next restriction is tested in a similar manner. This guarantees that a most general still discriminating generalization is achieved.

3.5   Forming New Concepts and Relations

As is illustrated by the example of $_{placebos, monodrugs}$ and $_{combidrugs}$ as well as $_{placebos, sedatives}$ and $_{anodynes}$, sometimes a good MSG cannot be built using the given concepts and roles. However, if a role or a concept in **all**-restrictions can be specialized, the new, specialized roles or concepts can be used for a more special characterization. The MDC which was not definable before is marked as $_{waiting}$ on the agenda and the new roles or concepts are put on the agenda.

Specialization is performed using two rules:

- If two concepts (C11 and C12 in fig. 3) of a MDC have the same concepts in the **all**-restriction of a role (C2 in fig. 3), but the range of the role is in fact disjoint for the two concepts, then introduce new subconcepts of the concept in the **all**-restriction and describe the **all**-restrictions in terms of these new concepts .

- If the concept (C2 in fig. 3) in the **all**-restriction of a role have disjoint subconcepts, introduce new relations that are restricted to these subconcepts and try them for characterization.
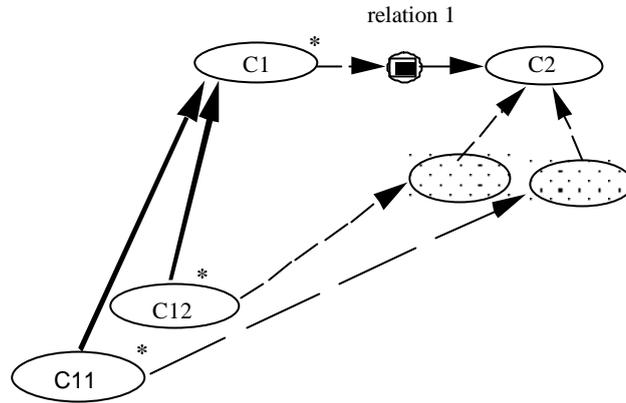
relation 1

Figure 3: Illustration of the refinement rules.

In the example of monodrug, combidrug, and placebos, the second rule applies. The concept substance has two disjoint subconcepts, active and add_on. The relation contains is specialized into contains_active which relates drugs and active, and the relation contains_add_on which relates drugs and add_on. Then, the MDC_1 is put back on the agenda as active, with the counter refinement increased by one. Based one the parameter max_refinement, at most $(| \text{role} | * | \text{concept} |)^{\text{max\_refinement}}$ different new roles are introduced by the second refinement rule. When MDC_1 is next selected from the agenda, these new roles are also tried for characterization. This leads to the following MSGs:

MSG(placebo) = drug **and all**(contains, add_on) **and atleast**(2,contains) **and atmost**(2,contains)

atmost(0,contains_active) **and**

atleast(2,contains_add_on) **and atmost**(2,contains_add_on)

MSG (monodrug) = drug **and atleast**(1,contains) **and atmost**(2,contains) **and**

atleast(1,contains_active) **and atmost**(1,contains_active) **and**

atmost(1,contains_add_on)

MSG (combidrug) = drug **and atleast**(2,contains) **and atmost**(3,contains) **and**

atleast(2,contains_active) **and atmost**(2,contains_active) **and**

atmost(1,contains_add_on)

The evaluation results in a role failure of 0 for contains_active and of 1/3 for contains_add_on, the failure for contains remaining the same as above. Therefore, the MGDs as presented in section 3 are built using uniformly contains_active. As the definitions are entered, MDC_1 is marked as definable.

In the example of monodrug, anodyne, and sedative, now the the first rule applies. The concept active used in the **all**-restriction of the role contains_active can be specialized into two disjoint subconcepts, active_1 and active_2. The following extensions are assigned:

ext(active_1) = {finalin,oxazepun},

ext(active_2) = {asa,phenazetin,prophymazon}.

These new concepts form MDC_5. This new MDC_5 is put on the agenda as active, and the counter rlength (for role chain length) is set to the one for MDC_2 plus 1. Then MDC_2 is marked on the agenda as waiting for a definition of MDC_5. Based on the parameter max_rlength at most $(| \text{role} | * | \text{concept} |)^{\text{max\_rlength}}$ different new

MDCs can be introduced by the first refinement rule. When $_{MDC\_5}$ is selected from the agenda, the built MSGs are without failure, as the **all**-restriction of $_{affects}$ is sufficient for discrimination. Therefore $_{MDC\_5}$ is marked $_{definable}$ and $_{MDC\_2}$ is set to $_{active}$. The new concepts are then sufficient to discriminate the concepts of $_{MDC\_2}$ based on an **all**-restriction of the role $_{contains\_active}$ (see the MGDs in section 3.0).

4   Evaluating KLUSTER

In the following sections, we evaluate our approach. First, we describe the theoretical properties of KLUSTER. Then, we compare KLUSTER with other work on conceptual clustering, learning relational concept definitions, and constructive induction.

4.1   Theoretical Evaluation

In the following, we want to evaluate KLUSTER theoretically. We first characterize the learning result of KLUSTER, then indicate the certainty of finding the MSG for a set of facts and the time complexity of the algorithm.

KLUSTER's learning result consists of root concepts (which correspond to user given sorts of other learning systems), several hierarchies and their interrelations, and newly constructed concepts and roles. Number restrictions are also learned. The learning result is represented within a term subsumption formalism, which has a well-defined semantics. It provides classification with inheritance. The **all**-restrictions and the number-restrictions of the formalism are computable in polynomial time. In this respect, the term subsumption formalism goes beyond the common restrictions of first-order logics. In particular, the formalism is not restricted to ij-determinate clauses (Muggleton & Feng, 1990). It has been shown that the term subsumption formalism is one of the greatest subsets of first-order logic with decidability in polynomial time (Donini et al., 1991). Therefore, it is a promising alternative to other restrictions of first-order logics. The learning result is easily understandable because the concept structure corresponds to a classical view of concept definitions. Hybrid representation systems with a TBox in the term subsumption formalism and facts in the ABox are becoming widely used. The KLUSTER algorithm can be incorporated into such hybrid systems in order to make them easier to use.

The restrictions of the formalism concern truly disjunctive concepts and the transitivity of relations. This includes recursive concepts which require a termination condition. So, for example, $_{member}$ cannot be learned by KLUSTER. Recursive concepts such as $_{ancestor}$ can be learned. However, many term subsumption system do not allow recursive concepts (terminological cycles). These systems cannot fully use KLUSTER's learning result. Another restriction of the term subsumption formalism is that it cannot express transitivity where more than two variables need be bounded within the same expression. For instance, it cannot be stated directly that a drug which contains a substance which increases blood pressure, also increases blood pressure. In order to express this information, a new sub-concept of drugs must be defined by its relation to those substances which raise blood pressure.

It is certain that KLUSTER finds a MSG for any concept. This is due to the concept representation in which exactly one MSG can be constructed for any set of terms. As was shown in section 3.3. this MSG is constructed by KLUSTER. If there exists a concept definition which is consistent with the ABox, then KLUSTER will determine it in polynomial time. If KLUSTER does not find a concept definition, then no hypothesis exists which is consistent with the minimal model of the ABox.

It may happen, however, that the failure of a MSG is greater than the threshold of MSG failure. This means that the MSG covers all positive instances but also instances of a disjoint concept. In this case, KLUSTER does not shift from the MSG to the MGD and does not enter a concept definition into the concept structure. The concept is indicated to be explored but not defined. Three cases can be distinguished. First, the concept cannot be expressed by a conjunction but is a truly disjunctive concept. KLUSTER cannot learn disjunctive concepts. Second, the concept cannot be defined using the given representation language. In this case, the specialization rules may introduce new concepts or roles which then allow for defining the concept. However, in contrast with Shapiro's refinement operator (Shapiro, 1983), KLUSTER's specialization is not complete. Therefore, the third reason for a concept being undefinable is that its definition lies outside of the hypothesis space enlarged by the specialization for introducing new terms.

KLUSTER does not need very many instances for learning. KLUSTER delivers already a MSG for just one example. In this case, the MSG corresponds to the classification of instances as performed by term subsumption formalisms. The time for finding a MSG grows polynomially in the number of instances (and roles and concepts in the all-restriction). Therefore, KLUSTER is able to run on large example sets. The most time-consuming part is the calculation of the MDCs. This information, however, needs not be given by the user as is the case for many other learning systems but is acquired by KLUSTER.

KLUSTER does not require the user to build the background knowledge carefully in order to enable successful learning. Instead, KLUSTER acquires the information which is represented as background knowledge by other learning systems (e.g. DISCIPLE (Kodratoff & Tecuci, 1989)). As the most specific generalization is exactly determined with respect to the given examples, incomplete descriptions of objects (e.g., a combidrug which contains only one active instance) prevent KLUSTER from learning the user intended concept definition (e.g., combidrugs having more than one active substance). A user who is not content with KLUSTER's learning result may input additional facts. In this way, KLUSTER can be used as an aid in inspecting data.

Computing the basic taxonomy by KLUSTER is of polynomial complexity over the number of facts. The MDCs are computable in the average case, but in the worst case there are exponentially many different MDC. It is $m \log_2 m$ to compute m MDCs. But, for n concepts there are at most $\binom{n}{n/2}$ different MDCs. Building the MSG is polynomial over the number of instances, the number of roles, and the number of concepts for the **all**-restriction of a role. It is polynomial because only named concepts and roles are used for **all**-restrictions. This is an incompleteness with respect to the expressability of term subsumption formalisms which allow more complex expressions. As is often the case, incompleteness makes the task solvable in polynomial time. If no named concept or role can be found for restricting a role's range, then constructive induction can define such a concept or role by

specialization. The specialization step is bounded by two parameters: the depth of specialization (i.e. a specialized concept or role can be further specialized and so on, but there is a specialization which will not be further specialized) and the number of trials to define a MDC. These bounds prevent the specialization step from combinatorial explosion. Further work on the trade-off between the formalism's expressability and the complexity of the concept learning task and on relating this to complexity results of others (e.g., (Haussler, 1989)) is planned. A preliminary study is (Kietz, 1992).

4.2   Related Work

The learning result of KLUSTER is a concept structure which is capable of expressing cross-classifications, hierarchies for several root concepts, and the cardinality of roles. A concept structure of this type is not learned by any other learning system. Therefore, it is hard to compare KLUSTER with other systems. In the following, we compare KLUSTER with conceptual clustering algorithms because the overall task of the system is to learn a hierarchy of concepts. With respect to KLUSTER's concept learning problem, it is compared with other learning algorithms which acquire structural concept definitions. As KLUSTER introduces new terms into the hypothesis language, it is also compared with other constructive induction algorithms.

4.2.1   Conceptual Clustering

The learning goal of conceptual clustering methods as well as that of KLUSTER is a hierarchy of concepts. The attribute based conceptual clustering methods, e.g. COBWEB (Fisher, 1987), UNIMEM (Lebowitz, 1987), WITT (Hanson & Bauer, 1989) require that all instances are described along the same attributes. This approach is not suitable to describe really different - but nevertheless related - things like, e.g., drugs, substances and symptoms. The complete attribute vectors are also a kind of segmentation into completely described observations. Even the relational conceptual clustering system KBG (Bisson, 1990) needs a segmentation of the input into observations, and it clusters only the observations and not the objects involved in them. KLUSTER does not require such a segmentation of the input. KLUSTER does *learning from examples* instead of clustering observations. It intensionally defines sets of objects involved in examples. LABYRINTH (Thompson & Langley, 1989), another approach for relational clustering, also requires a segmentation into observations. Its main task is to cluster these observations, but it also tries to cluster the objects occurring in the observations. LABYRINTH suffers from combinatorial explosion when trying to find an optimal mapping between the different objects involved in an observation. KLUSTER does not encounter this explosion because it uses the **all**-restriction instead of an optimal mapping of the involved objects.

4.2.2   Learning Structural Descriptions

KLUSTER is comparable with logical concept learning approaches as, e.g., RLGG (Plotkin, 1970), GOLEM (Muggleton & Feng, 1990), FOIL (Quinlan, 1990) in that it learns relational concept definitions.

KLUSTER requires both unary and binary relations as input. The quantity and quality of given examples is irrelevant. KLUSTER reflects the quality of the examples by the output of MSGs that cover the examples. Kietz

(1992) shows that that learning MSGs (RLGGs) in Horn Logic is in general intractable. GOLEM's restriction to depth-bounded determinate Horn clauses is one possibility to come to polynomial learnability. KLUSTER's MSGs with the **all**-restriction are another possibility of polynomial learnability. The difference is that GOLEM requires all objects in the examples to be reachable by deterministic relations. This is why GOLEM is not applicable to our drug-example as contains is a non-determinate relation. A drug contains many substances, so contains is really a relation and not a function. If the substances of a drug are encoded as a list (so that contains becomes ij-determinate) then accessing one of the contained substances, ne cessary for defining anodyne and sedative, requires the non-determinate member relation. In contrast, KLUSTER allows non-deterministic relations (e.g. contains in the examples). The non-determinate relations in the examples are abstracted into one expression (the **all**-restriction) describing the similarities of all related objects.

The heuristic learning approach FOIL is also capable of using non-determinate relations. Running FOIL on our side-effect of drug data gave the following results[4]:

anodyne(A) :- contains(A,B) & affects(B,C) & pain(C)

sedative(A) :- contains (A,B) & affects(B,C) & excitement(C)

active(A) :- affects(A,B)

add_on(A) :- contains(B,A), placebo(B) with warning that this does not cover all tuples

placebo (A):- not(monodrug(A)) & not(combidrug(A))

combidrug (A):- not(monodrug(A)) & anodyne(A) with warning

monodrug (A) :- not(combidrug(A)) & anodyne(A) with warning

The definitions of monodrug, combidrug, and placebo cannot be found by FOIL. The rules found are not covering all positive instances. It is easily seen that the cross classification leads to some confusion: FOIL tries to use anodyne for the definition of combidrugs and monodrugs. This, however, does not lead to forming a MSG. The learning result of KLUSTER is of a different representation formalism and requires different input than FOIL[5]. The main difference between KLUSTER and FOIL, however, concerns the search in the hypothesis space. Whereas KLUSTER can construct a consistent MSG if one exists, FOIL's search heuristics cannot guarantee to find a hypothesis which is consistent with the data. The reason is that an encoding of the SAT problem (Garey & Johnson, 1979) is a possible learning problem of FOIL (cf. Haussler, 1989) but not of KLUSTER. As we know that SAT is intractable, any equivalent learning problem is intractable as well.

4.2.3   Constructive Induction

Approaches to constructive induction can be structured with respect to the reasons for the introduction of a new term. KLUSTER's reason for introducing a new term is the need to refer to a particular set of objects. This need is constituted by the definition of another concept. KLUSTER also introduces new relations as was shown in our example of section 3. The newly introduced terms are specializations of already given or learned terms of the hypothesis language.

The CIGOL system which implements induction as inverse resolution learns literals that can play the role of a missing premise, given the other premises and the conclusion from a resolution step (Muggleton & Buntine, 1988). CIGOL introduces new terms into the hypothesis language. The decision whether the newly introduced term should be kept or removed is left to the user. So, no criterion for the selection of a new term is formalized. Moreover, the search space for a new literal is $2^n - 1$ given a substitution with n elements. In the literature on inverse resolution there exists no formalized method to focus the search within this space. Finally, CIGOL does not define the newly introduced predicate. KLUSTER's introduction of a new concept can be viewed as learning a missing premise of a classification rule. However, the implemented method is more efficient than inverse resolution because the search space is limited and the search within it is focused. In KLUSTER, at most n concepts can be newly introduced, given n relevant roles. A new term is only introduced into the hypothesis language if KLUSTER is capable of defining it.

5   Conclusion

KLUSTER is the first learning algorithm which is capable of learning a concept structure in the framework of term subsumption formalisms. Concepts are defined by relations to other concepts which are uniformly represented within the same concept structure. Thus, a learned concept or role serves for defining another concept. There is no separation between background knowledge and learned knowledge. Concepts are represented in a structure which involves several roots. Cross-classification or forming subconcepts under diverse aspects is possible in KLUSTER. The interrelatedness of concepts is not only expressed by the concept representation but also by the way concepts are learned. Concepts are formed in the context of mutually disjoint concepts (MDCs). Refinements of concepts and roles are made in the course of defining a concept. In this way, the KLUSTER approach represents and exploits a rich concept structure.

KLUSTER learns most specific generalizations (MSGs) as well as most general discriminations (MGDs). With respect to a particular representation, it is guaranteed that KLUSTER finds the unique MSG in polynomial time. Finding the best MGD would be exponential. So, KLUSTER takes the first MGD found. The introduction and definition of new roles potentially makes classification exponential. Therefore, defined roles are excluded from the basic algorithm. Only some defined roles are introduced if they are really needed for the distinction between concepts whose extensions are disjoint. Learning new roles is polynomially bounded by two parameters. KLUSTER inductively learns in polynomial time. The use of KLUSTER's learning results (i.e. the deductive classification) cannot be performed completely in polynomial time because of the defined roles.

References

Bisson, G. (1990) KBG A Knowledge-Based Generalizer, *7th ICML-90*, Morgan Kauffmann, 9-15

Borgida, A., Brachman, R.J., McGuinness, D.L., & Resnick, L.A. (1989) Classic: A structural data model for objects, in *Proc. ACM SIGMOD-89*, Portland, Oregon, 58-67.

Brachman, R. J., & Schmolze, J.G. (1985) An Overview of the KL-ONE Knowledge Representation System, *Cognitive Science*, 9, 171-216.

Brachman, R.J. (1977) What's in a Concept: Structural Foundations for Semantic Networks, *Int. Journal of Man-Machine Studies*, 9, 127-152.

Brachman, R.J., Gilbert, V.P., & Levesque, H.J. (1985) An Essential Hybrid Reasoning System, in *IJCAI-85*, 532-539.

Buntine, W. (1988) Generalized Subsumption and Its Applications to Induction and Redundancy, *Artificial Intelligence*, 36, 149-176.

Ceri, S., Gottlob, G. & Tanca, L. (1990) *Logic Programming and Databases*, New York, Springer

Cohen, W.W., Borgida, A., & Hirsh, H. (1992) Computing Least Common Subsumers in Description Logic, in *Proc. of AAAI-92*, to appear.

Donini, F.M., Lenzerini, M., Nardi, C., & Nutt, W. (1991) Tractable Concept Languages, in *Proc. IJCAI-91*, 458-463

Emde, W., Habel, C. & Rollinger, C.R. (1983) The Discovery of the Equator or Concept-Driven Learning, in *Proc. IJCAI-83*, 455 - 458

Fisher, D.H. (1987) Knowledge Acquisition Via Incremental Conceptual Clustering, *Machine Learning*, 2, 139-172.

Garey, M. R. & Johnson, D.S. (1979) *Computers and Intractability - A Guide to the Theory of NP-Completeness,* New York, Freeman

Haussler, D. (1989) Learning conjunctive concepts in structural domains, *Machine Learning*, 4, 7-40.

Kearns, M.J. (1990). *The Computational Complexity of Machine Learning.* Cambridge, MA, London, MIT Press.

Kietz (1992) J.-U. Kietz. A Comparative Study of Structural Most Specific Generalizations Used in Machine Learning, Proc. of ECAI'92 Workshop W18.

Kietz, J.-U. & Wrobel, S. (1991) Controlling the Complexity of Learning in Logic Through Syntactic and Task-Oriented Models, Procs. *Inductive Logic Programming* Workshop, Porto, also in: Muggleton (ed): *Inductive Logic Programming*, Academic Press, 1992, 311-333

Kietz, J.-U. (1988) Incremental and Reversible Acquisition of Taxonomies, *Procs. European Knowledge Acquisition Workshop*, Birlinghoven: GMD-Studien No. 143.

Kodratoff, Y. & Tecuci, G. (1989) The Central Role of Explanations in DISCIPLE, in: K. Morik (ed): *Knowledge Representation and Organization in Machine Learning*, New York: Springer, 135-147

Lebowitz, M. (1987) Experiments with Incremental Concept Formation: UNIMEM, *Machine Learning*, 2, 103-138.

Luck, K.v., Nebel, B., Peltason, C., & Schmiedel, A. (1987) The Anatomy of the BACK System, KIT-Report No. 41, Technical University Berlin.

Michalski, R.S. (1983) A Theory and Methodology of Inductive Learning, in *Machine Learning —An Artificial Intelligence Approach*, Vol. I, Los Altos, CA:Morgan Kaufman, 83-134.

Michalski, R.S. (1990) Learning Flexible Concepts: Fundamental Ideas and a Method Based on Two-Tiered Representation, in Y. Kodratoff, R.S. Michalski (eds): *Machine Learning - An Artificial Intelligence Approach* Vol. III, San Mateo:Morgan Kaufmann, 63-111.

Morik, K. & Kietz, J.U. (1989) A Bootstrapping Approach to Conceptual Clustering, in *Proc. Sixth Intern. Workshop on Machine Learning.*

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

Moser, M.G. (1983) An Overview of NIKL, the New Implementation of KL-One, in *Research in Knowledge Representation and Natural Language Understanding*, B., B., Beranek and Newman Inc., Cambridge, Mass.

Muggleton, S., Buntine, W. (1988) Machine Invention of First-Order Predicates by Inverting Resolution, *Procs. IWML-88*, Ann Arbor, Michigan, Morgan Kaufmann.

Muggleton, S. (1990) Inductive Logic Programming, *Procs. of the First Conference on Algorithmic Learning Theory*, Tokyo, Ohmsha Publishers.

Muggleton, S., & Feng, C. (1990) Efficient Induction of Logic Programs, *Procs. of the First Conference on Algorithmic Learning Theory,* Tokio.

Nebel, B. (1990) *Reasoning and Revision in Hybrid Representation Systems*, New York: Springer.

Peltason, C., Luck, K., & Kindermann C. K. (1991) Terminological Logic Users Workshop, KIT-Report 95, Technical University Berlin.

Peltason, C., Schmiedel, A., Kindermann, C., & Quantz, J. (1989) The BACK System revisited, KIT-REPORT 75, Techn. Univ. Berlin, Berlin, W. Germany.

Plotkin, G.D. (1970) A note on inductive generalization,*Machine Intelligence*, 5, 153-163.

Quinlan, R. (1990) Learning Logical Definitions from Relations, in: *Machine Learning*, 5, 239-266.

Shapiro, E. (1983) *Algorithmic Program Debugging*. Cambridge, MA, MIT Press.

Stepp, R.E., & Michalski, R.S. (1986) Conceptual Clustering: Inventing Goal-Oriented Classifications of Structured Objects, in: R. Michalski, J. Carbonell, T. Mitchell (eds):*Machine learning - An AI Approach*, Vol.II, San Mateo, Morgan Kaufman, 471-498.

Thompson,K.,& Langley, P. (1989) Incremental concept formation with composite objects, *Procs. of the 6th International Workshop on Machine Learning*, 373-374.

Vilain, M. (1985) The Restricted Language Architecture of a Hybrid Reasoning System, *IJCAI-85*, 547-551.

Wrobel, S. (1987) Higher-Order Concepts in a Tractable Knowledge Representation, in: K. Morik (ed): Procs. German Workshop on Artificial Intelligence, Berlin, Springer, 129-138.

---

[1]For a discussion of the computational complexity of entailment see Nebel (1990, section 4.5).

[2] It is the user who names root concepts, the system generates an artificial name such as rootconcept_1. Primitive concepts are named based on the names in the ABox.

[3] Computing MDCs from pair wise disjointness of concepts corresponds to the NP-complete problem CLIQUE (Garey & Johnson, 1979). But, for KLUSTER the inheritance in the basic taxonomy restricts the number of concepts (nodes) in one CLIQUE. In the example, at most five of all twelve concepts are to be considered as a CLIQUE: the concepts subsumed by drug.

[4] FOIL has two modes, one with negated and one without negated literals in rule premises. We ran FOIL in both modes and show the best rules of both runs..

[5]When trying out KLUSTER on the senator votes domain, KLUSTER detected that the democratic senators all voted for south Africa sanctions whereas there was no topic on which the republic senators all gave the same vote.