

Using Taxonomic Background Knowledge in Propositionalization and Rule Learning ^{*}

Monika Žáková, Filip Železný

Czech Technical University
Technická 2, 16627 Prague 6, Czech Republic
zakovm1@fel.cvut.cz, zelezny@fel.cvut.cz

Abstract. Knowledge representations using semantic web technologies often provide information which translates to explicit term and predicate taxonomies in relational learning. Here we show how to speed up the process of propositionalization of relational data by orders of magnitude, by exploiting such ontologies through a novel refinement operator used in the construction of conjunctive relational features. Moreover, we accelerate the subsequent search conducted by a propositional learning algorithm by providing it with information on feature generality taxonomy, determined from the initial term and predicate taxonomies but also accounting for traditional θ -subsumption between features. This information enables the propositional rule learner to prevent the exploration of useless conjunctions containing a feature together with any of its subsumees and to specialize a rule by replacing a feature by its subsumee. We investigate our approach with a propositionalization algorithm, a deterministic top-down propositional rule learner, and a recently proposed propositional rule learner based on stochastic local search. Experimental results on genomic and engineering data [2] indicate striking runtime improvements of the propositionalization process and the subsequent propositional learning.

1 Introduction

With the development of semantic web technologies and knowledge management using ontologies, increasing amounts of expert knowledge in important knowledge-intensive domains such as bioinformatics is becoming available in the form of ontologies and semantic annotations. One of the well known examples is the Gene Ontology¹. However, semantic representation is becoming popular even in industrial use [2] for sharing and efficient searching of information in production enterprises. Knowledge representation formalisms used to capture ontologies and semantic annotations are based on description logics, which have convenient properties with regard to complexity and decidability of reasoning [3].

^{*} An extended version of a paper appearing in the ECML/PKDD'07 proceedings.

¹ <http://www.geneontology.org>

Inductive logic programming (ILP) aims at learning a theory in a subset of first-order logic from given examples, taking background knowledge into account. It has been considerably successful in various knowledge discovery problems such as in bioinformatics [4]. Standard ILP techniques cannot efficiently exploit explicit taxonomies on concepts and relations, which are typically available in semantic knowledge representations [2]. While in principle, any taxonomy can be encoded in background knowledge, there is good reason to view ontologies as meta-information, which can be directly exploited to guide the refinement operator used to search through the space of first-order rules.

We illustrate this statement through an example. The Gene Function Ontology declares a concept `binding` and its subconcept `protein_binding`. Such concepts are reflected by terms in ILP. It is possible to declare in background knowledge e.g.

```
subclass(binding, protein_binding).
geneFunction(G, F1) :- geneFunction(G, F2), subclassTC(F1, F2).
```

(where `subclassTC/2` is defined as the transitive closure of `subclass/2`). Unfortunately, in such an approach, for the following two exemplary clauses (hypotheses)

```
C = activeGene(G) :- geneFunction(G, binding).
D = activeGene(G) :- geneFunction(G, protein_binding).
```

it does not hold $C\theta \subseteq D$, so clause D is not obtained by applying a specialization refinement operator onto clause C . Similar reasoning applies to taxonomies on relations (predicates), which are also often present in ontologies (they are a standard part of the Resource Description Framework [5]). Informally, declaring taxonomies only in background knowledge has the consequence that some clause pairs with comparable generality are treated as incomparable by the clause search algorithm. This has significant implications to efficiency of search, as we demonstrate in this work.

Recently, effort has been exerted to utilize the information available in ontologies out of the scope of the traditional ILP settings. There are 3 main approaches to this problem: introduce learning mechanisms into description logics [6], [7], hybrid languages integrating Horn logic and description logics [8] and learning in a more expressive formalism [9], [10]. Learning in description logics is useful mainly for the refinement of existing description hierarchies; however, here we are constrained to limitations of description logic and therefore e.g. unable to express formulas with a free variable. Therefore the works investigating learning in description logics are valuable especially in their results on dealing with the open-world assumption in learning and in transformations of description logics into some other subset of the first-order logic. Reasoning and learning in hybrid languages attempts to loosely couple descriptions of concept hierarchies in description logics with rules expressed in function-free Horn logic. Learning in hybrid languages is split into two phases, each utilizing well-known algorithms particular to each of the respective formalisms. Reasoning in hybrid languages is more complex than reasoning in both its constituent formalisms. E.g. even if

reasoning in the chosen subsets of both DL and HL formalisms separately is decidable, reasoning in the corresponding hybrid formalism may be undecidable. A growth in computational complexity is obviously also a deficiency of approaches using representation formalisms with expressivity exceeding that of first-order logic.

This paper concentrates on the problem of exploiting taxonomies in the framework of propositionalization of relational data by constructing relational features and subsequent learning from the propositionalized representation. This approach to relational learning has been paid significant attention in the last few years [11]. In particular, we demonstrate that explicit taxonomies can be exploited with significant runtime benefits at two different stages of this approach.

First, we exploit the term and predicate taxonomies in the process of relational feature construction conducted by the propositionalization algorithm, by adopting the formalism of sorted logic for feature representation and by adapting a refinement operator for first-order features to be co-guided by the taxonomies. The generated feature set plays the role of a boolean attribute set describing the examples from which the propositional algorithm learns.

Second, note that construction of features is implemented through systematic search using the mentioned refinement operator. In the search space, some elements comply to the notion of a feature [1] and they are used as such. Naturally, some of the resulting features are specializations of other resulting features. In contrast to state-of-the-art logic-based propositionalization systems [11], we explicitly store the information that a feature has been obtained by specializing another feature. The propositional algorithm then receives, apart from the propositionalized data matrix, also the resulting feature taxonomy information. This efficiently enables it to prevent the exploration of a conjunction containing a feature together with any of its subsumees and to specialize a rule by replacing a feature by its subsumee.

The first step above assumes that relation and term taxonomies are available beforehand, e.g. from ontology information. In the second step, the feature taxonomy is co-determined by these relation and term taxonomies, but also by standard θ -subsumption among first-order formulas. Consequently, the feature taxonomy is created and can be exploited whether or not relation and term taxonomies were available.

2 Method

In this section we describe our methodological improvements in algorithms for propositionalization and subsequent propositional learning.

2.1 Features

Propositionalization can be understood as a transformation method, where a relational learning problem is compiled to an attribute-value problem, which one can solve using propositional learners [12, 11]. During propositionalization, *features* are constructed from the background knowledge, addressing structural properties of individuals. Each feature is defined as a clause in the form

$$f_i(X) \leftarrow Lit_{i,1}, \dots, Lit_{i,n}$$

where the literals in the body are derived from the background knowledge and the argument in clause’s head refers to an individual as an example identifier. The features then correspond to attributes which form the basis for columns in single-table (propositional) representations of the data. If the clause defining a feature is called for a particular individual and this call succeeds, the feature is set to “true” in the corresponding column of the given example; otherwise it is set to “false”. Recently several propositionalization systems have been proposed. Examples include: RSD [1] and SINUS [13] among others.

Propositionalization systems differ in their constraints applied to select which clauses of the form above form an admissible feature. Our approach stems from the constraints used in the RSD system. While the exact details and their motivations can be found in [1], the general conditions making a conjunction of literals a correct feature body can be briefly stated as follows.

1. Each literal in the conjunction corresponds to a predicate from a declared set of predicates. A predicate declaration also specifies for each argument whether it is an *input* or an *output* argument.
2. Every variable appearing at an input argument in a literal must also appear at an output argument in a preceding literal in the conjunction or in the feature head. Every variable appearing at an output argument in a literal is distinct from all variables in the literal and all preceding literals.
3. The conjunction is not a literal-wise union of two or more conjunctions which themselves are correct feature bodies.
4. The conjunction has at most l literals, for a prescribed $l > 0$.

The third condition prevents the formation of redundant features as it is assumed that the learner employed subsequently on the propositionalized representation (such as a conjunctive rule or decision tree learner) is itself able to conjugate two or more constructed features. To simplify the explanations to follow, we will also assume that all literals $Lit_{i,j}$ are non-negated atoms.

2.2 Sorted Logic

The first difference of the current approach to propositionalization from RSD is in dealing with terms. In RSD, a predicate declaration assigns a type symbol to each argument, from a finite set of type symbols. On top of the constraints enumerated above, a further condition then dictates that in a correct feature body, a single variable may not appear simultaneously in two arguments with different types.

The present approach replaces the notion of type with that of *sort* borrowed from the formalism of sorted logic, which is suitable for encoding term taxonomies. Sorted logic contains in addition to predicate and function symbols also a disjoint set of sort symbols. A sort symbol denotes a subset of the domain called a sort [14].

A *sorted variable* is a pair, $x : \tau$, where x is a variable name and τ is a sort symbol. Semantically, a sorted variable ranges over only the subset of the domain

denoted by its sort symbol. The semantics of universally-quantified sorted formulas can be defined in terms of their equivalence to ordinary formulas: $\forall x : \tau \phi$ is logically equivalent to $\forall x : \neg\tau(x) \vee \phi'$ where ϕ' is the result of substituting x for all free occurrences of $x : \tau \in \phi$.

A *sort theory* Σ is a finite set of formulas containing function formulas and subsort formulas. A *function formula* has the form

$$\forall x_1, \dots, x_n \tau_1(x_1) \wedge \dots \wedge \tau_n(x_n) \rightarrow \tau(f(x_1, \dots, x_n)) \quad (1)$$

where, in this paper, we constrain ourselves to $n = 0$, thus reducing function formulas to the form $\tau(f)$ reflecting that constant f is of sort τ . A *subsort formula* has the form

$$\forall x \tau_1(x) \rightarrow \tau_2(x) \quad (2)$$

reflecting that τ_1 is a direct subsort of τ_2 . It is required that the directed graph corresponding to the subsort theory is acyclic and has a single root denoted *univ*.

For a sort theory Σ , a Σ -sorted substitution is a mapping from variables to terms such that for every variable $x : \tau$, it holds that $\Sigma \models \forall x \tau(t)$ where t is $(x : \tau)\theta$. Informally, this is a substitution that does not violate the sort theory.

In the present propositionalization approach, terms in features are constants or sorted variables. Background knowledge consists of an ordinary first-order theory and a sort theory Σ . A *declaration* for a predicate of symbol π and arity n has the form

$$\pi(m_1\tau_1, \dots, m_n\tau_n)$$

where $m_i \in \{+, -\}$ denotes whether i -th argument is an input (+) or an output (-), as defined in Sec. 2.1. Besides the constraints stated in Sec. 2.1, correct features must respect the sort relationships. Formally, a literal *Lit* may appear in a feature only if there is a declaration $\pi(m_1\tau_1, \dots, m_n\tau_n)$ and a Σ -sorted substitution θ such that $\pi(\tau_1, \dots, \tau_n)\theta = Lit$.

Example. Assume a learning task where examples are genes and background knowledge consists of the gene function ontology and further of gene-gene interaction data. Classification may e.g. distinguish genes expressed in various biological situations. The following predicate declarations may then be stated:

```
geneFunction(+gene, -function).
interacts(+gene, -gene).
```

The background knowledge consists of a first-order theory, such as

```
geneFunction(il2ra, 'interleukin-2 receptor activity').
interacts(cd4, il2ra).
```

and a sort theory, which, for technical convenience, is represented using auxiliary predicates `is_a/2` (for function formulas) and `subsort/2` (for subsort formulas), for example:

```

is_a(il2ra, gene).
subsort('interleukin-2 receptor activity', 'receptor activity').
subsort('receptor activity', function).

```

Then the following exemplary features

```

f0(G:gene) :- interacts(G:gene, H:gene)
f1(G:gene) :- geneFunction(G:gene, F:'receptor activity').

```

are correct (in the second case, 'receptor activity' is a subsort of 'function'), and true for gene `il2ra`, whereas the following expression

```

f2(G:gene) :- geneFunction(G:gene, F:function), interacts(G:gene,
H:gene).

```

is not a correct feature, despite its compliance with the sort theory. This is due to condition 3 in Section 2.1; clearly, the body of `f2` is a union of two correct feature bodies. In other words, $f2(X)$ is logically equivalent to $f0(X) \wedge f1(X)$ for any X . Note however, that `f2` can be *specialized by refinement* into a correct feature

```

f3(G:gene) :- geneFunction(G:gene, F:function), interacts(G:gene,
H:gene), geneFunction(H:gene, F:function).

```

expressing that gene `G` interacts with another gene with whom it shares some function. Clearly, `f3` cannot be expressed as a conjunction of any two or more correct features. (*End of example*).

Next we turn attention to the refinement operator through which features are constructed.

2.3 Refinement

We have adapted the *sorted downward refinement* from [14], which accounts for term taxonomies, to further account for the earlier defined feature constraints and predicate declarations used in propositionalization, and for a further kind of taxonomy – the *relation taxonomy* – often available in ontology data. This taxonomy is encoded through meta-predicates in the form

```

subrelation(pred1/n, pred2/n).

```

providing the explicit meta-information that goal $pred_1(Arg_1, \dots, Arg_n)$ succeeds whenever goal $pred_2(Arg_1, \dots, Arg_n)$ succeeds, i.e. $pred_1$ is more general. The directed graph corresponding to the entire set of the `subrelation/2` statements (where direction is such that edges start in the more general goal) is assumed to be a forest. The set of its roots is exactly the set of predicates declared through the predicate declarations defined in the previous section. It is assumed that the non-root predicates inherit the argument declarations from their respective roots.

As feature heads are fixed in our propositionalization framework, we are concerned with refinement of their bodies, i.e. conjunctions. We will use the notion of an *elementary Σ -substitution*. Its general definition can be found in [14], however, adapted to our framework, the definition simplifies.

An *elementary Σ -substitution* for a sorted conjunction C is $\{x : \tau_1\} \rightarrow \{x : \tau_2\}$ where $\{x : \tau_1\}$ occurs in C and Σ contains the subsort formula $\forall\psi\tau_2(\psi) \rightarrow \tau_1(\psi)$ for some variable ψ . If $\{x : \tau_2\}$ already occurs in C , then x is deterministically renamed² to a variable not occurring in C . Unlike in [14], we can disregard the case of substituting a sorted variable by a function (as we work with function-free features) and, similarly to RSD [1], we neither allow to unify two distinct variables (an equality theory can be defined instead in background knowledge).

Let C be a conjunction, possibly empty, of non-negated atoms where any term is either a constant or a sorted variable, Σ be a sort theory, and Δ a set of predicate declarations. We define the *downward Δ, Σ -refinement*, written $\rho_{\Delta, \Sigma}(C)$, as the smallest set such that:

1. For each θ that is an elementary Σ -substitution for C , $\rho_{\Delta, \Sigma}(C)$ contains $C\theta$.
2. Let $\pi(m_1\tau_1, \dots, m_n\tau_n)$ be a declaration in Δ such that for each i for which $m_i = +$, C contains a variable (denote it x_i) of sort τ'_i which equals or is a subsort of τ_i . Let further $\{x_i | m_i = -\}$ be a set of distinct variables not appearing in C . Then $\rho_{\Delta, \Sigma}(C)$ contains $C \wedge \pi(x_1 : v_1, \dots, x_n : v_n)$, where $v_i = \tau'_i$ if $m_i = +$ and $v_i = \tau_i$ otherwise.
3. Let C contain a literal $pred_1(x_1\tau_1, \dots, x_n\tau_n)$ and let $pred_2$ be a direct subrelation of $pred_1$. Then $\rho_{\Delta, \Sigma}(C)$ contains C' , which is acquired by replacing $pred_1(x_1\tau_1, \dots, x_n\tau_n)$ with $pred_2(x_1\tau_1, \dots, x_n\tau_n)$ in C .

This definition of downward refinement is adapted to our feature construction framework and differs from the refinement operator defined in [14] in several respects, even if no **subrelation** information is assumed. First, we use a simplified notion of the *elementary Σ -substitution* as explained above. Second, our feature bodies are conjunctions of non-negated atoms, and thus the range of our $\rho_{\Delta, \Sigma}(C)$ is reduced in comparison to [14], which refines clauses and therefore also produces negated atoms. Thirdly, the range of $\rho_{\Delta, \Sigma}(C)$ is further reduced because a new atom is added to C only if its input variables can be matched to existing output variables in C , while respecting sort relationships. Lastly, while [14] adds a new atom to C always with all arguments being a variable assigned the root-sort *univ*, we instead ‘initiate’ each variable with the sort defined for it in the corresponding predicate declaration, which usually are proper subsorts of *univ*. This fact obviously reduces the total volume of search space traversed by the refinement operator.

Confined to 12 pages, we skip the proof that, under very general assumptions on Δ , the defined refinement operator is (i) finite, (ii) complete, in that all correct features (as defined in Sections 2.1 and 2.2) up to variable renaming are enumerated by its recursive closure, whenever the initial C in the recursive application of $\rho_{\Delta, \Sigma}(C)$ is the empty conjunction, and also (iii) non-redundant, in that $\rho_{\Delta, \Sigma}(C_1) \cap \rho_{\Delta, \Sigma}(C_2) = \{\}$ if $C_1 \neq C_2$. However, the operator is not necessarily correct, in that all its products would be correct feature bodies. In

² That is, we do not allow several elementary substitutions differing only in the chosen renaming.

$Propositionalize(\Delta, B, \Sigma, E, l)$: **Given**, a set Δ of predicate declarations, a first-order theory (background knowledge) B , a sort theory Σ , a set of unary ground facts (examples) $E = \{e_1, \dots, e_m\}$ and a natural number l ; **returns** a set $\{f_1, \dots, f_n\}$ of constructed features, each with at most l atoms in the body, an elementary subsumption matrix \mathbf{E} , an exclusion matrix \mathbf{X} , and an attribute-value matrix \mathbf{A} where $\mathbf{A}_{i,j} = 1$ whenever f_i is true for e_j and $\mathbf{A}_{i,j} = 0$ otherwise.

1. $n = 0$; $Agenda$ = a single element list $[(C, 0)]$, where C = empty conjunction;
2. If $Agenda = []$: go to 10
3. $(Curr, Parent) := \mathbf{Head}(Agenda)$; $Tail := \mathbf{Tail}(Agenda)$
4. If **Nonempty** $(Curr)$ and **Undecomposable** $(Curr)$:
5. $n := n + 1$; $f_n = \mathbf{AddFeatureHead}(Curr)$;
6. $\mathbf{E}_{n,Parent} = 1$; $Parent = n$;
7. $\mathbf{A}_{n,1..l} = \mathbf{Coverage}(Curr, E, B, \Sigma, \mathbf{A}_{Parent,1..l})$
8. $Rfs := \rho_{\Delta, \Sigma}(Curr)$; $Rfs := \{(Cnj, Parent) | Cnj \in Rfs, |Cnj| \leq l\}$
9. $Agenda := \mathbf{Append}(Rfs, Tail)$; go to 2
10. $\mathbf{X} = \mathbf{Closure}(\mathbf{E})$
11. **Return** $f_1, \dots, f_n, \mathbf{E}, \mathbf{X}, \mathbf{A}$

Fig. 1. A skeleton of the algorithm for propositionalization through relational feature construction using the sorted refinement operator $\rho_{\Delta, \Sigma}$.

particular, it may produce a conjunction violating condition 3 in Sec. 2.1, such as the decomposable expression f_2 shown in Sec. 2.1. However, as we have illustrated, a correct feature body may be obtained by refining a decomposable conjunction. For this reason, we deliberately keep the refinement incorrect in this sense and conduct a further non-decomposability check in the feature construction algorithm (Sec. 2.5).

2.4 Feature Taxonomy

During the recursive application of the refinement operator, a feature generality taxonomy becomes explicit. For purposes of enhancing the performance of the propositional learning algorithm applied subsequently on the propositionalized data, we pass the feature taxonomy information to the learner through two boolean matrices.³ Assume that features f_1, \dots, f_n have been generated with corresponding conjunctive bodies b_1, \dots, b_n . The *elementary subsumption matrix* \mathbf{E} of n rows and n columns is defined such that $\mathbf{E}_{i,j} = 1$ whenever $b_i \in \rho_{\Delta, \Sigma}(b_j)$ and $\mathbf{E}_{i,j} = 0$ otherwise. The *exclusion matrix* \mathbf{X} of n rows and n columns is defined such that $\mathbf{X}_{i,j} = 1$ whenever $i = j$ or $b_i \in \rho_{\Delta, \Sigma}(\rho_{\Delta, \Sigma}(\dots \rho_{\Delta, \Sigma}(b_j) \dots))$ and $\mathbf{X}_{i,j} = 0$ otherwise. Sec. 2.5 shows how the matrices are instantiated during feature construction and Sec. 2.6 shows how they are utilized by the propositional learner.

2.5 Feature Construction Algorithm

A skeleton of the propositionalization algorithm is shown in Fig. 1. The algorithm is a depth-first search generally similar to the feature constructor of RSD

³ While alternative data structures are of course possible for this sake, the elected binary matrix form requires modest space for encoding (our implementation uses one byte for each 8 matrix elements) and also is conveniently processed in the propositional algorithm implementation.

[1]. The main difference lies in using the novel sorted refinement operator $\rho_{\Delta, \Sigma}$ and also in creating the matrices \mathbf{E} and \mathbf{X} storing the generality taxonomy of constructed features. Some of the bold-faced functions have self-explaining names, others require comments. In particular, the **Undecomposable** procedure checks Condition 3 described in Sec. 2.1, through a method used in RSD and detailed in [1]. The **AddFeatureHead** forms a feature clause by formally attaching a head to the body, which consists of the constructed conjunction $Curr$. The **Coverage** procedure verifies the truth value of a conjunction for all examples in E returning a vector of Boolean values. The verification is done by a transformation of the sorted conjunction $Curr$ to an ordinary first-order conjunction as explained in Sec. 2.2 and then using a standard resolution procedure against a Prolog database consisting of B and Σ . Note that for efficiency, only nonempty undecomposable conjunctions, rather than all products of the refinement, are subjected to the coverage check. Again, for efficiency, the procedure obtains the coverage $\mathbf{A}_{Parent, 1 \dots l}$ of the closest ancestor (subsuming) conjunction whose coverage was tested: any example i such that $\mathbf{A}_{Parent, i}$ is false can be left out of testing as it makes the current conjunction necessarily false as well. The **Closure** procedure computes the transitive closure of the elementary subsumption relation captured in \mathbf{E} in the manner described in Sec. 2.4, and represents the closed relation analogically in matrix \mathbf{X} , in which it further sets $\mathbf{X}_{i, i} = 1$ for all $1 \leq i \leq n$.

2.6 Rule Learning

We have adapted two rule learning algorithms to account for the feature taxonomy information provided by the propositionalization algorithm.

The first algorithm stems from the rule inducer of RSD [1]. It is based on a heuristic general-to-specific deterministic beam search for the induction of a single propositional conjunctive rule for a given target class, and a cover-set wrapper for the induction of the entire rule set for the class. The core part of the algorithm is a procedure for refinement (specialization) of a propositional conjunction. Given a set of features $F = \{f_1, \dots, f_n\}$, the standard algorithm refines a conjunction C of features into the set $\{C \wedge f_i | f_i \in F, f_i \notin C\}$. In our enhanced version, the algorithm is provided with the elementary subsumption matrix \mathbf{E} and the exclusion matrix \mathbf{X} . Using these matrices it can prevent the useless combination of a feature and its subsumee within the conjunction, and specialize a conjunction by replacing a feature with its elementary (direct) subsumee. Formally, the enhanced algorithm refines a conjunction $C = f_{k_1} \wedge \dots \wedge f_{k_p}$ into the set

$$\{C \wedge f_i | f_i \in F, \mathbf{X}_{i, j} = 0 \forall f_j \in C\} \cup \quad (3)$$

$$\{f_{k_1} \wedge \dots \wedge f_{k_{r-1}} \wedge f_i \wedge f_{k_{r+1}} \wedge \dots \wedge f_{k_p} | 1 \leq r \leq p, \mathbf{E}_{i, k_r} = 1\} \quad (4)$$

Furthermore, we have similarly enhanced the stochastic local DNF search algorithm first introduced in [15] and later transferred into the propositionalization framework by [16]. This algorithm conducts search in the space of DNF




Fig. 2. Sorted refinement vs. standard refinement on CAD data. Left: Nodes explored Right: Time taken. (Experiments exceeding 1000s were discarded)

formulas, i.e. it refines entire propositional rule sets. Refinement is done by local, non-deterministic DNF term changes, which are detailed in [15]. In our version, the \mathbf{X} matrix is used to prevent the combination of a feature and its subsumee within a DNF term.

3 Experimental Results

We designed experiments to assess the runtime impact of (i) the novel taxonomy-aware refinement operator in the propositionalization process, and (ii) the exploitation of the feature-taxonomy information in subsequent propositional learning.

We conducted tests in two domains. The first concerns genomics, where we used data and language declarations from [17]. The nature of this learning task has been illustrated in the Introduction and in Sec. 2.2. The second is concerned with learning from product design data. Here the examples are semantically annotated CAD documents. We used the same learning setting and ontology data as in [2].

Figures 2 and 3 illustrate on log scale the number of conjunctions searched (left) and the time spent on search (right) to enumerate all conjunctions true for at least 80% examples, for increasing maximum conjunction size l . Here, we distinguish the sorted refinement operator using a special sort theory Σ encoding the taxonomy information, against the standard refinement operator, which treats the taxonomy information only as part of background knowledge. While in both cases exactly the same set of conjunctions is produced, an order-of-magnitude runtime improvement is observed for the ‘taxonomy-aware’ operator.

Tables 1 and 2 show in turn the runtime spent of inducing a rule set by two algorithms (top-down and stochastic) through 10-fold cross validation in two scenarios: in the first, no feature taxonomy information is used by the algorithms, in the second, feature taxonomy is exploited as described in Sec. 2.6. A significant speedup is observed when feature taxonomy is used without compromising the predictive accuracy.




Fig. 3. Sorted refinement vs. standard refinement on Genomic data. Left: Nodes explored Right: Time taken. (Experiments exceeding 1000s were discarded)

Table 1. Propositional rule learning from CAD data

Algorithm	Time taken	Predictive accuracy
Top-down	0.223 \pm 0.0785	0.6599 \pm 0.2111
Top-down, feature taxonomy	0.061 \pm 0.0221	0.6599 \pm 0.2214
SLS	0.6268 \pm 1.4544	0.6154 \pm 0.1770
SLS, feature taxonomy	0.2758 \pm 0.83	0.6109 \pm 0.1864

4 Conclusions

In this work we have proposed principled methods to exploit term, predicate and feature taxonomies to increase the performance of propositionalization and subsequent propositional learning. The significance of our work is supported by three factors: (i) order-of-magnitude runtime improvements with no sacrifice in predictive accuracy, (ii) the practical value and common use [11] of the propositionalization strategy to relational machine learning, which was the target of our methodological enhancements, and (iii) the increasing volumes of semantic knowledge representations providing explicit taxonomies. In future work, we plan to extend the scope of meta-information exploitable by refinement operators beyond taxonomy information. For example, principled methods are needed to deal with meta-knowledge such as “relation R is a function” or “binary relation R is symmetrical,” etc.

Acknowledgements. This research was supported by the FP6-027473 Project SEVENPRO. The authors would like to thank Peter Flach and Simon Rawles for valuable advice concerning term taxonomies in ILP.

References

1. Železný, F. and Lavrač, N.: Propositionalization-based relational subgroup discovery with RSD. *Machine Learning* 62: 33-63, 2006.

Table 2. Propositional rule learning from Genomic data

Algorithm	Time taken	Predictive accuracy
Top-down search	0.991 ± 0.6548	0.787 ± 0.13
Top-down, feature taxonomy	0.335 ± 0.1927	0.755 ± 0.07
SLS	2.963 ± 2.587	0.787 ± 0.13
SLS, feature taxonomy	1.908 ± 1.689	0.755 ± 0.071

2. Žáková, M., Železný, F., Garcia-Sedano, J. A. et al.: Relational Data Mining Applied to Virtual Engineering of Product Designs. In *Proc. of the 16th Conference on ILP*, Springer, 2007.
3. Baader, F. and Calvanese, D. and McGuinness, D. et al. (Eds.): *The Description Logic Handbook: Theory, Implementation and Applications*, Cambridge University Press, 2003.
4. King, R. D. et. al.: Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427:247–252, 2004.
5. RDF Vocabulary Description Language 1.0: RDF Schema. W3C Recommendation 10 February 2004. Available at <http://www.w3.org/TR/rdf-schema/>.
6. Badea, L. and Neihuys-Cheng, S.-W.: A Refinement Operator for Description Logics. In Cusens, J. and Frich, A. (eds.): *Inductive Logic Programming*, Vol. 1866 of *Lecture Notes in Artificial Intelligence*, pp. 40-59, Springer-Verlag, 2000.
7. Kietz, J.-U.: Learnability of Description Logic Programs, *Inductive Logic Programming: 12th International Conference, ILP 2002, Sydney, Australia, LNCS Volume 2583/2003*, Springer, 2002.
8. Donini, F.M., Lenzerini, M. et. al.: AL-log: Integrating Datalog and Description Logics. *Journal of Intelligent Information Systems*, 10(3):227–252, 1998.
9. Lloyd, J.W.: *Logic for Learning: Learning Comprehensible Theories from Structured Data*, Series: *Cognitive Technologies*, Springer, 2003.
10. Kifer, M. and Lausen, G. and Wu, J.: Logical Foundations of Object-Oriented and Frame-Based Languages, *Journal of ACM*, vol. 42, 1995, pp. 741-843.
11. Mark-A. Krogel, Simon Rawles, Filip Železný, Peter A. Flach, Nada Lavrač, and Stefan Wrobel. Comparative evaluation of approaches to propositionalization. In *Proc. of the 13th ILP*, volume 2835 of *LNAI*, pages 197–214. Springer, 2003.
12. Nada Lavrač and Peter A. Flach. An extended transformation approach to inductive logic programming. *ACM Trans. on Comp. Logic*, 2(4):458–494, 2001.
13. N. Lavrač and S. Džeroski. *Inductive Logic Programming: Techniques and Applications*. Ellis Horwood, 1994.
14. Frisch, A.: Sorted downward refinement: Building background knowledge into a refinement operator for ILP. In *ILP, LNAI 1634:104-115*. Springer, 1999.
15. Ulrich Rückert and Stefan Kramer. Stochastic local search in k-term dnf learning. In *Proc. of the 20th ICML*, pages 648–655, 2003.
16. A. Paes, G. Zaverucha, F. Železný, D. Page, A. Srinivasan. ILP through Propositionalization and k-Term DNF Learning. In *Proc. of the 16th Conference on ILP*, Springer, 2007.
17. I. Trajkovski, F. Železný, J. Tolar, N. Lavrač. Relational Subgroup Discovery for Descriptive Analysis of Microarray Data. In *Procs of the 2nd Int Sympos on Computational Life Science*, Springer, 2006.