

What Kinds of Relational Features are Useful for Statistical Learning?

Amrita Saha¹, Ashwin Srinivasan², and Ganesh Ramakrishnan¹

¹ School of Computer Science and Engineering, Indian Institute of Technology, Bombay, India.

² Department of Computer Science
South Asian University, New Delhi, India.

Abstract. A workmanlike, but nevertheless very effective combination of statistical and relational learning uses a statistical learner to construct models with features identified (quite often, separately) by a relational learner. This form of model-building has a long history in Inductive Logic Programming (ILP), with roots in the early 1990s with the LINUS system. Additional work has also been done in the field under the categories of propositionalisation and relational subgroup discovery, where a distinction has been made between *elementary* and *non-elementary* features, and statistical models have been constructed using one or the other kind of feature. More recently, constructing relational features has become an essential step in many model-building programs in the emerging area of Statistical Relational Learning (SRL). To date, not much work—theoretical or empirical—has been done on what kinds of relational features are sufficient to build good statistical models. On the face of it, the features that are needed are those that capture diverse and complex relational structure. This suggests that the feature-constructor should examine as rich a space as possible, in terms of relational descriptions. One example is the space of all possible features in first-order logic, given constraints of the problem being addressed. Practically, it may be intractable for a relational learner to search such a space effectively for features that may be potentially useful for a statistical learner. Additionally, the statistical learner may also be able to capture some kinds of complex structure by combining simpler features together. Based on these observations, we investigate empirically whether it is acceptable for a relational learner to examine a more restricted space of features than that actually necessary for the full statistical model. Specifically, we consider five sets of features, partially ordered by the subset relation, bounded on top by the set F_d , the set of features corresponding to definite clauses subject to domain-specific restrictions; and bounded at the bottom by F_e , the set of “elementary” features with substantial additional constraints. Our results suggest that: (a) For relational datasets used in the ILP literature, features from F_d may not be required; and (b) Models obtained with a standard statistical learner with features from subsets of features are comparable to the best obtained to date.

1 Introduction

The emerging area of statistical relational learning (SRL) is characterised by a number of distinct strands of research. Especially prominent is research concerned with the construction of models from data that use representations based on either first-order logic programs augmented with probabilities or probabilistic graphical models. The concerns here are the usual ones to do with expressive power, estimation, and inference: what kinds of probabilistic models can be constructed with one representation or the other; how can we estimate the structure and parameters in the model; how do we answer queries exactly, given data that is observed, and perhaps missing; and so on. The combination of relational learning with statistical modeling, however, has a longer history within Inductive Logic Programming (ILP), with origins at least as early as 1990, with the LINUS system [8]. Since then, there have been regular reports in the literature on the use of ILP systems, as a tool for constructing relational features for use in statistical modeling [15].

An argument can be made that construction of relational features must necessarily require some form of first-order learning, of which ILP is an instance (for example, see [7]). Arguments in-principle aside, the literature also suggests that augmenting any existing features with ILP-constructed relational ones can substantially improve the predictive power of a statistical model. There are thus good practical reasons to persist with this variant of statistical and logical learning. On the other hand there has been some work done on comparing the different kinds of propositionalisation techniques used to transform the search space from the space of first-order hypothesis to the space of propositional features which can be handled by more scalable propositional/statistical learners. [6] claims that of the two main kinds of propositionalization methods, namely logic oriented and database-oriented, both have their specific advantages. While logic-oriented methods can handle complex relational structures in the form of background knowledge and provide more expressive relational models, database-oriented models are much more scalable. According to their empirical findings, a combination of features from these two groups are necessary for learning good models.

Even within this well-trodden corner of statistical relational learning (no more, perhaps, than a “poor man’s SRL”), there are some issues that remain unaddressed. To date, not much work—theoretical or empirical—has been done on what kinds of relational features are sufficient to build good statistical models. On the face of it, the features that are needed are those that capture diverse and complex relational structure. This suggests that the feature-constructor should examine as rich a space as possible, in terms of relational descriptions. One example is the space of all possible features in first-order logic, given constraints of the problem being addressed. Practically, it may be intractable for a relational learner to search such a space effectively for features that may be potentially useful for a statistical learner. Additionally, the statistical learner may also be able to capture some kinds of complex structure by combining simpler features together. For example, a statistical learner like a support vector machine or lo-

gistic regression that may also be able to approximate the effect of a conjunction of these features using their weighted sum. Reports in the ILP literature suggest at least 5 kinds of relational feature classes: (1) F_d : (the set of) features from definite clauses with no restrictions other than those of the domain [17, 15]; (2) F_i : features from “independent” clauses that place restrictions on the sharing of existential variables [1]; (3) F_r : features denoting a class of relational subgroup that place additional restrictions on the use of existential variables in independent clauses [9]; (4) F_s : features from “simple” clauses in the sense described in [10]; and (5) F_e : features developed from the class of “elementary” clauses described in [3]. In this paper, we show certain subset relationships hold between these sets. These are shown diagrammatically in Fig. 1. When exploring whether smaller feature-spaces are adequate, we use these relationships to investigate empirically whether exploring larger sets of features adds any significant predictivity to a statistical learner.




Fig. 1. Relationships between the sets of features considered in this paper.

The rest of the paper is organised as follows. In Section 2 we describe the mapping between clauses and relational features. We also describe in greater detail the five feature classes F_d – F_e . In Section 3 we derive the relationships shown in Fig. 1 and enumerate some consequences that follow directly from them. Section 4 describes an empirical investigation, using standard statistical learners of the smallest size feature class that appears to be useful for constructing predictive models for several ILP benchmark datasets. Section 5 concludes the paper.

2 Feature Classes

In this paper, we will take first-order relational features simply to be first-order clauses that either maximize some objective either collectively or individually (effectively, a restatement of the distinction in [4] between strong and weakly relevant features, in optimisation terms). Specifically, we will assume that is a one-to-one correspondence between first-order relational features and definite clauses (alphabetic variants of a clause are treated as being the same clause). Formally, we adopt the same notation as in [15] for a function that maps a first-order definite clause to a feature. The set of examples provided to an ILP system

can be defined as a binary relation which is a subset of the Cartesian product $\mathcal{X} \times \mathcal{Y}$ where \mathcal{X} denotes the set of individuals (*i.e.* structured objects consisting of first-order predicates) and \mathcal{Y} denotes the finite set of classes. The definite clauses obtained as an output of the ILP system can be represented in the form $h_j(x, c) : \text{Class}(x, c) \leftarrow \text{Cp}_j(x)$ where $\text{Cp}_j : \mathcal{X} \rightarrow \{0, 1\}$ is a nonempty conjunction of predicates on the variable $x \in \mathcal{X}$ and c is the class variable specified in the head predicate. For convenience, we will say $\text{Head}(h_j(x, c)) = \text{Class}(x, c)$ and $\text{Body}(h_j(x, c)) = \text{Cp}_j(x)$. Given a clause $h_j(x, c)$, a first-order feature can be defined as $f_j(x) = 1$ iff $\text{Body}(h_j(x, c)) = 1$ and 0 otherwise. Constraints on $\text{Body}(h_j(x, c))$ allow us to define several kinds of features. Following [10] (with a small difference) we distinguish between “source” predicates and “sink” predicates in the language of clauses allowed in the domain. The former are those that contain at least one output argument and any number of input arguments (in the sense used by the mode declarations in [11]) and the latter are those that contain no output arguments. In a clause, thus all new existentially quantified variables should be introduced in the source literals and not in the sink literals.³ Additionally, following [1], for a clause $h_j(x, c)$, the independent components in $\text{Body}(h_j(x, c))$ are partitions of the literals in $\text{Body}(h_j(x, c))$ into sets, such that each partition consists of a “connected” set of literals (once again, using input and output variables in the sense of [11]), and that literals across partitions only share the head variable x . We are then able to define the following kinds of relational feature classes:

The class F_d : This consists of first-order features obtained from definite clauses $h_j(x, c)$ in the functional manner described above. That is, no constraints are placed on $\text{Body}(h_j(x, c))$. Features used in [17], for example, belong to this class.

The class F_i : This is a restricted version of the class F_d , consisting of features from clauses $h_j(x, c)$ such that $\text{Body}(h_j(x, c))$ consists of exactly 1 independent component.

The class F_r : This is a restricted version of the class F_i , consisting of features obtained from clauses $h_j(x, c)$ such that $\text{Body}(h_j(x, c))$ consists of exactly 1 independent component, and with an additional constraint that all new existential variables introduced by a source literal appear in source or sink literals in $\text{Body}(h_j(x, c))$. The features in [9] are from this class, and the first-order features described in [3] are a special case of features in F_r (the special case arising from restrictions on sources and sinks to structural and property predicates as described earlier).

The class F_s : This consists of features obtained from clauses $h_j(x, c)$ such that h_j is a simple clause in the sense defined by [10]. That is, $\text{Body}(h_j(x, c))$ contains exactly 1 sink literal.

The class F_e : This is a restricted version of the of the class F_r , consisting of features obtained from clauses $h_j(x, c)$ such that $\text{Body}(h_j(x, c))$ contains

³ “Structural” predicates used in [3] are thus binary source literals that introduce a single new variable, and “property” predicates are sink literals.

exactly 1 sink literal. “Elementary” features described in [3] are a special case of features in F_e .

3 Relationships Between Feature Classes

Some of the relationships⁴ between feature classes are evident from the definitions. That is: $F_r \subseteq F_i \subseteq F_d$. The following additional statements hold

- $F_s \subseteq F_i$
- $F_s \not\subseteq F_r$ and $F_r \not\subseteq F_s$
- $F_e = F_r \cap F_s$

Thus, the feature classes exhibit the following hierarchical structure:

$$F_e = (F_r \cap F_s) \subseteq F_r \subseteq F_i \subseteq F_d$$

$$F_e = (F_r \cap F_s) \subseteq F_s \subseteq F_i \subseteq F_d$$

Given these subset relationships between feature classes, it is also of some importance to consider whether there a way of reconstructing, using logical operations, every feature in a superset class by combining features from a subset (or smaller) class (clearly, the reverse is always possible: a feature in a subset class can always be constructed from a feature in the superset class). The interest here is of course that if such logical relationships hold, then features in the larger class may be approximated by statistical learners using weighted combinations of features from the smaller class. The following logical relationships hold between the feature classes:

- Every feature in F_i can be constructed from features in F_s
- Every feature in F_d can be constructed from features in F_i

Thus, every feature in F_d can be constructed from F_s . In addition:

- Not every feature in F_i can be constructed from features in F_r
- Not every feature in F_r can be constructed from features in F_e
- Not every feature in F_s can be constructed from features in F_e

We now evaluate empirically the utility of features from the different classes, proceeding from the smallest (F_e) to the largest (F_d).

4 Empirical Evaluation

4.1 Aims

Our aim is to obtain empirical evidence of the smallest feature class that is found to yield good statistical models. Specifically, we mean that we wish to find using a set of well-studied benchmark data sets and popular statistical learners, whether there is a feature class such that adding features from larger classes yields no significant increases in predictive accuracy.

⁴ Proofs of these results have not been elaborated here. They are accessible at this location:<http://www.cse.iitb.ac.in/~amrita/ilp/appendix.pdf>

4.2 Materials

Domains We use biochemical datasets that have been used in a range of papers in the ILP literature. These are tabulated below. These datasets have been used widely: see, for example [7, 14, 18]).

1. Alz (Amine)
2. Alz (Acetyl)
3. Alz (Memory)
4. Carcin
5. DSSTox
6. Mut188

4.3 Algorithms and Machines

The statistical learners used in the paper are these:

- A support vector machine(SVM). Specifically, we examine SVMs with both the L1-norm ,also known in literature as L1-SVM (the LibLINEAR implementation) and the L2-norm ,known as L2-SVM,(the LibSVM implementation) on the weight vector used as the regularizer. The L1-norm is known to induce sparsity on the feature space, forcing a form of feature selection which is important when the number of features is large.
- Logistic regression. Specifically, we consider a standard version of this technique and a faster and more efficient variant called SMLR (Sparse Multinomial Logistic Regression) [5]
- Rule ensemble learning using maximum likelihood estimation (MLRules) [2]. This employs a greedy approach of maximising likelihood to construct an ensemble of rules from the features. Like logistic regression, the weights of the rules are derived from the conditional probability distribution learnt. So it can be thought of as a generalization of Logistic Regression, where the space explored to construct rule ensembles can be considered as an approximation of of conjunctions of rules/features.

All statistical learners here, compute the decision function as a weighted linear combination of the features (which can be thought of as an approximation to a logical conjunction). Features in each feature class are constructed using the ILP system Aleph [16]. It is possible to enforce the constraints associated with each feature class in a straightforward manner as part of background knowledge provided to this system.

All experiments were conducted on a machine equipped with a 8-core Intel i7 2.67GHz processors and 8 gigabytes of random access memory.

4.4 Method

The methodology adopted for providing the statistical learners features from feature classes F_e to F_d is as follows:

1. Select a total ordering on the classes $F_e \prec \dots \prec F_d$ that is consistent with the partial ordering imposed by subset relationships between these sets.
2. For each dataset and each statistical learner:
 - (a) For sets S_j from smallest (F_e) to the largest (F_d), in the total ordering \prec :
 - (b) Let $F_0 = \emptyset$
 - i. Repeat R times:
 - A. Obtain a set of features F from S_j
 - B. $F_j = F_{j-1} \cup F$
 - C. Obtain an estimate of the predictive accuracy A_j of the statistical learner with features F_j
 - ii. Obtain the mean predictive accuracy $\overline{A_j}$ across the R repeats
3. Determine the smallest set S_k after which there are no significant changes in mean predictive accuracy $\overline{A_k}$

The following details are relevant:

1. There are only 2 total orderings possible, given the subset relationships between the feature classes: $F_e \prec_1 F_r \prec_1 F_s \prec_1 F_i \prec_1 F_d$ and $F_e \prec_2 F_s \prec_2 F_r \prec_2 F_i \prec_2 F_d$. According to our incremental algorithm, here, F_s in \prec_1 and F_r in \prec_2 actually refer to sets of features from $F_r \cup F_s$, because of their order of appearance. The results we report here are with $\prec = \prec_1$. Our conclusions do not change with either \prec_1 or \prec_2 .
2. All predictive accuracies are estimated using 10-fold cross-validation.
3. The results are averaged over $R = 4$ repetitions. Larger values of R would result in smaller standard errors of the mean estimate.
4. Statistical learners have parameters that require optimisation. It has been shown elsewhere [18] that using default values of parameters can result in sub-optimal models, which can clearly confound the conclusions that can be drawn here. For each set of training data in a cross-validation run, we set aside some small part of the training data as a “validation” set, and use this to tune the parameters of the statistical learner. The “best” parameter value that results is then used to construct a model on the training data; and then evaluated on the test dataset of that cross-validation run. This does not necessarily yield the best model possible, but accuracies are usually higher than those obtained with the default settings for the parameters.
5. The feature-construction procedure employed by the ILP engine requires an upper-bound k on the number of features produced for each class label in the example data. For experiments reported here, $k = 500$. There is a randomised element to the feature-construction procedure in Aleph, which can be seen as selecting (non-uniformly) upto k features, from the set of all possible features allowed. The selection of features is controlled by parameters that correspond to precision and recall in the data mining literature. For experiments here, these values have been deliberately left low (and hence, easy to obtain for the ILP engine). The intent is that the statistical learner should be able to combine these to obtain higher values.

6. We have elected to assess the utility of features from each feature class by examining the improvements in mean accuracy by augmenting features already present from feature classes earlier in the ordering \prec . The alternative of simply comparing accuracies with a new set of features drawn from the from each feature class X in the ordering would have confounded matters, since this set could contain features from a class $Y \subseteq X$. It would not then be known whether the increases in accuracy, if any, are due to features from a class X , or those from class Y .
7. It is known that for large R , mean accuracies are distributed normally. For small R , it is known that the sampling distribution of the mean is a t -distribution with $R - 1$ degrees of freedom. This is used in statistical comparisons when needed (as will be seen, it is often evident whether differences are significant, and no undue statistical testing is needed).

4.5 Results and Discussion

The comparative performance of the models is shown in Fig. 2. The principal result from this tabulation is this: broadly, there is little value in including features from F_d (the largest feature class considered here). Examining the table in greater detail, we are able to obtain the number of “wins” for each feature class (that is, the number of times the highest predictive accuracy results from using features in that class). A cross-tabulation of this against the learners is shown in Fig. 3:

The data in Figs. 2 and 3 suggest that it is sufficient to consider features from F_i (that is, features from clauses containing one independent component). Now, while it is possible to reconstruct every feature in F_d exactly as a simple logical conjunction, of some features in F_i , it is also possible to reconstruct several (but not all) features in F_d by simple logical conjunction of select features in F_r , F_s and F_e . This appears to be exploited by all of the statistical learners here, since it is by no means necessary for any of them that features from F_d are consistently required to produce the best results. This is reinforced further, if rather than considering outright wins, we consider a model as being good enough if there no (statistically) significant difference to the best model. Then, the performance tallies can be summarized as in Fig. 4.

These tabulations do suggest that of the classes $F_e \dots F_d$, the class F_i may be the most useful. But how good are the models obtained with F_i , when compared against the ILP models reported in the literature? Fig. 5 shows a comparison of the best statistical models against the predictive accuracies of the ILP models obtained after parameter selection and optimisation [18].

Finally, although not relevant to the aims of the experiment here, we note some exceptional behaviour on the “Carcin” dataset, in which there is a fairly consistent trend of decreasing predictive accuracies as we progress from F_e to F_d . An examination of re-substitution (training-set) accuracies for this data shows the opposite trend, suggesting that these data may be especially prone to overfitting.

Data	Learner	Using Features From				
		F_e	F_r	F_s	F_i	F_d
Alz(Amine)	L2-SVM	66.14±0.84	65.71±0.69	67.23±0.86	81.62±0.92	81.45±0.61
	L1-SVM	65.50±0.15	65.84±0.14	64.91±1.65	78.34±4.18	80.31±1.57
	LR	65.42±0.0	65.42±0.0	64.38±0.0	79.32±3.93	81.28±0.31
	SMLR	65.42±0.0	66.05±0.22	67.12±0.38	81.20±2.56	81.63±2.77
	MLRules	64.39±0.0	64.39±0.0	68.67±0.0	82.32±1.18	82.22±0.35
Alz(Acetyl)	L2-SVM	69.42±0.33	68.52±0.70	68.59±2.11	74.40±0.80	72.97±0.47
	L1-SVM	69.85±0.05	68.55±0.0	66.20±0.0	74.06±0.91	72.78±1.44
	LR	69.82±0.0	69.82±0.0	71.93±0.0	73.87±0.44	74.16±0.24
	SMLR	70.36±0.11	70.32±0.14	72.17±0.12	73.34±1.38	73.20±0.65
	MLRules	69.59±0.0	69.59±0.0	70.95±0.0	71.67±0.47	72.44±0.52
Alz(Memory)	L2-SVM	59.80±0.22	61.09±0.45	63.46±1.07	68.39±1.15	68.07±1.75
	L1-SVM	56.99±1.29	62.89±1.07	64.33±0.25	71.29±1.04	67.96±0.41
	LR	59.07±0.08	59.07±0.08	65.87±0.0	70.31±1.54	69.30±0.55
	SMLR	59.69±0.0	58.44±1.32	66.34±0.73	70.10±1.26	71.83±1.67
	MLRules	58.83±0.19	58.83±0.19	64.76±0.08	70.13±1.14	67.91±0.64
Alz(Toxic)	L2-SVM	70.72±1.0	71.97±1.99	77.17±0.82	83.01±0.73	82.04±2.40
	L1-SVM	70.06±0.0	71.73±0.0	80.94±3.15	82.59±0.70	82.83±2.21
	LR	72.84±0.0	72.09±1.51	77.91±0.0	81.65±0.79	83.32±0.87
	SMLR	74.50±2.35	74.07±0.79	76.76±0.16	82.59±0.39	84.21±0.69
	MLRules	73.18±0.0	73.18±0.0	77.44±0.0	84.50±0.44	84.19±0.56
Carcin	L2-SVM	60.59±1.58	62.15±1.75	59.49±1.75	62.03±1.28	59.85±1.67
	L1-SVM	61.21±0.0	59.71±0.0	58.66±3.10	60.26±0.71	59.21±1.29
	LR	59.00±0.0	57.71±0.16	54.01±2.79	51.80±1.37	52.52±3.49
	SMLR	60.88±0.0	60.17±0.38	59.51±2.91	58.80±3.42	57.94±2.84
	MLRules	59.39±0.25	59.20±0.15	58.83±1.96	57.94±2.31	56.64±0.94
DSSTox	L2-SVM	69.36±1.15	69.86±1.19	72.21±1.22	73.12±0.94	70.90±3.15
	L1-SVM	69.25±0.25	69.00±0.25	70.77±1.60	71.91±1.77	69.21±1.34
	LR	70.49±0.0	70.49±0.0	72.18±1.71	71.01±0.87	70.32±0.47
	SMLR	72.24±0.0	72.29±0.09	70.71±2.62	70.95±2.94	71.02±1.38
	MLRules	71.55±0.0	71.55±0.0	72.85±1.43	72.73±1.52	72.16±1.09
Mut(188)	L2-SVM	–	65.80±1.25	85.84±1.36	86.49±0.28	84.63±1.58
	L1-SVM	–	63.70±0.30	84.92±0.65	85.27±0.31	85.52±0.48
	LR	–	67.12±0.0	74.85±1.74	77.48±3.38	77.61±2.88
	SMLR	–	73.44±0.0	88.06±1.57	85.83±1.81	84.10±2.02
	MLRules	–	71.86±0.0	85.04±0.76	86.50±1.40	85.70±1.55

Fig. 2. Mean predictive accuracies of statistical models including features from F_e to F_d . “–” indicates no features in this class were possible given the domain constraints. Here, refers to the F_s in the incremental order \prec_1 , hence it actually denotes the set of features from $F_r \cup F_s$. Same notation has been followed in Fig. 3. and Fig. 4.

Feature Class	Number of Wins					Total Wins
	L1-SVM	L2-SVM	LR	SMLR	MLRules	
F_e	1	0	1	1	1	4/35
F_r	0	1	1	1	0	3/35
F_s	0	0	1	1	1	3/35
F_i	3	6	1	1	4	15/35
F_d	3	0	4	3	1	11/35

Fig. 3. Number of outright wins for a feature class. This is number of occasions out of the total number of possible occasions (*i.e.* 35) on which a statistical learners achieves the highest mean predictive accuracy using features from that class.

Feature Class	Number of Good Enough Models					Total No. of Good Models
	L1-SVM	L2-SVM	LR	SMLR	MLRules	
F_e	2	1	2	2	2	9/35
F_r	2	2	1	4	4	13/35
F_s	3	2	1	4	4	14/35
F_i	7	7	6	7	7	34/35
F_d	4	5	6	7	7	25/35

Fig. 4. Number of good enough models (out of all possible models *i.e.* 35), using a feature class. A model is taken to be good enough if its predictive accuracy is not statistically different to the model with the highest predictive accuracy.

We also experimented with $\prec=\prec_2$ and found that the results were quite similar with the ones with $\prec=\prec_1$ tabulated here. In interest of space, those results have been omitted.

5 Concluding Remarks

In this paper we have explored the relationship between several feature spaces that have been reported in ILP literature and examined empirically whether it is possible to construct good statistical models using features from smaller spaces. The intuition underlying this is that statistical models may often be able to approximate the effect of more elaborate features by weighted combinations of simpler features. Our results suggests that the class F_i , consisting of features constructed from clauses containing exactly one independent component seems to be particularly useful. This makes some sense: a linear combination of multiple features from F_i can approximate the reconstruction of a full first-order feature, since no variable sharing is required between such features. In fact, this leads us to hypothesize that statistical learners like [12] that perform conjunctive feature learning will not perform any better than learners using weighted combinations

Data	Statistical Model	ILP Model With Parameter Selection & Optimization
Alz (Amine)	82.32±1.18	80.20
Alz (Acetyl)	74.16±0.24	77.40
Alz (Memory)	71.83±1.67	67.40
Alz (Toxic)	84.50±0.44	87.20
Carcin	62.15±1.75	59.10
DSSTox	73.12±0.94	73.10
Mut(188)	88.06±1.57	88.30

Fig. 5. Comparison of mean predictive accuracies of statistical models against the ILP models constructed with parameter selection and optimisation (see [18]).

of features from F_i , and will incur a greater computational cost. Further, this also leads us to believe that weighted linear combinations of first order clauses in relational learning models such as MLN [13] and Relational Markov Networks [19] could be efficiently and effectively approximated by weighted linear combination of clauses from simpler classes such as F_i and is part of our ongoing work.

References

1. Vitor Santos Costa, Ashwin Srinivasan, Rui Camacho, Hendrik Blockeel, Bart De-moen, Gerda Janssens, Wim Van Laer, James Cussens, and Alan Frisch. Query transformations for improving the efficiency of ilp systems. *Journal of Machine Learning Research*, 4:491, 2002.
2. Krzysztof Dembczynski, Wojciech Kotlowski, and Roman Slowinski. Maximum likelihood rule ensembles. In *ICML*, pages 224–231, 2008.
3. Peter A. Flach and Nicolas Lachiche. First-order bayesian classification with 1bc. 2000.
4. George H. John, Ron Kohavi, and Karl Pfleger. Irrelevant features and the subset selection problem. In *MACHINE LEARNING: PROCEEDINGS OF THE ELEVENTH INTERNATIONAL*, pages 121–129. Morgan Kaufmann, 1994.
5. Balaji Krishnapuram, Lawrence Carin, Mario A. T. Figueiredo, and Alexander J. Hartemink. Sparse multinomial logistic regression: Fast algorithms and generalization bounds. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(6):957–968, June 2005.
6. Mark-A. Krogel, Simon Rawles, Filip Zelezny, Peter A. Flach, Nada Lavrac, and Stefan Wrobel. Comparative evaluation of approaches to propositionalization, 2003.
7. Niels Landwehr, Andrea Passerini, Luc De Raedt, and Paolo Frasconi. kfoil: Learning simple relational kernels. In *AAAI*, pages 389–394. AAAI Press, 2006.
8. Nada Lavrac and Saso Dzeroski. *Inductive Logic Programming: Techniques and Applications*. Routledge, New York, NY, 10001, 1993.
9. Nada Lavrac, Filip Zelezny, and Peter A/ Flach. Rsd: Relational subgroup discovery through first-order feature construction. pages 149–165, July 2002.

10. Eric McCreath and Arun Sharma. Lime: A system for learning relations. pages 336–374, 1998.
11. Stephen Muggleton. Inverse entailment and prolog, 1995.
12. Jagarlapudi Saketha Nath Pratik Jawanpuria and Ganesh Ramakrishnan. Efficient rule ensemble learning using hierarchical kernels. pages 161–168, 2011.
13. Matthew Richardson and Pedro Domingos. Markov logic networks. *Mach. Learn.*, 62(1-2):107–136, February 2006.
14. José Carlos Almeida Santos. Toplog: Ilp using a logic program declarative bias. In *ICLP'08*, pages 818–819, 2008.
15. Lucia Specia, Ashwin Srinivasan, Sachindra Joshi, Ganesh Ramakrishnan, and Maria das Graças Volpe Nunes. An investigation into feature construction to assist word sense disambiguation. *Machine Learning*, 76(1):109–136, 2009.
16. Ashwin Srinivasan. *The aleph manual*. 1999.
17. Ashwin Srinivasan, Stephen Muggleton, Michael J. E. Sternberg, and Ross D. King. Theories for mutagenicity: A study in first-order and feature-based induction. *Artif. Intell.*, 85(1-2):277–299, 1996.
18. Ashwin Srinivasan and Ganesh Ramakrishnan. Parameter screening and optimisation for ilp using designed experiments. *Journal of Machine Learning Research*, 12:627–662, 2011.
19. B. Taskar, P. Abbeel, M.-F. Wong, and D. Koller. Relational markov networks. In L. Getoor and B. Taskar, editors, *Introduction to Statistical Relational Learning*. MIT Press, 2007.