Relational Learning from Drug Adverse Events Reports

Ergun M. Biçici
North Carolina State University,
Department of Computer Science,
Campus Box 8206 Raleigh, NC 27695-8206

Abstract

We applied relational learning to discover rules from adverse events reports. We used the FOIL relational learning system to find a set of rules for withdrawn drugs. We compared our results with FDA's reasons for withdrawal.

1 Introduction

A drug adverse event is any unintended response in a patient's body during or after the use of a drug. Reports of adverse events can be used to reach a greater understanding of the causes and the background of reactions with particular drugs, which is useful for both the FDA [1] and the pharmaceutical industry. These reports contain valuable information, which can be harnessed by machine learning.

In this paper, we show how relational learning, the type of machine learning that can discover rules by using multiple relations in the language of first-order logic, can help in discovering knowledge from drug adverse events reports. We use adverse event reports pertaining to drugs that are withdrawn from the market, for which the causal link is established by the FDA.

In section 2, we review the structure and content of the reports and discuss the difficulties involved in learning from them. In section 3 we present the relational learning system FOIL, which learns function-free Horn clauses from a data set that contains relations with multiple arguments. In section 4, we present our experimental setup and in the next section we discuss and present our results.

2 Drug Adverse Event Reports

An adverse events report may contain information regarding the dose and the frequency of intake of the drug, the demographics of the patient (e.g. age, weight, gender), the set of adverse reactions observed (e.g. eye infection, difficulty in breathing), and the concomitant medication (e.g. aspirin, antibiotics). For this study, we chose to use adverse event reports pertaining to the drugs withdrawn from the market. There were two main reasons for this choice. Adverse events reports of withdrawn drugs are more likely to contain information that is of value in understanding the causal links. Another reason was that the causal link established by the FDA for withdrawal provides us a way to evaluate the effectiveness of our approach. These reports were taken from Health Canada's Canadian Adverse Drug Reaction Information System (CADRIS) [2].

From this set of reports, we were able to produce a data set for each drug that contains tuples of:

- $Drug\text{-}AdverseEvent\text{-}ConcomitantMedication$
- $Drug\text{-}AdverseEvent\text{-}ConcomitantMedication\text{-}SeriousEffect\text{-}Gender$
- $Drug\text{-}AdverseEvent\text{-}ConcomitantMedication\text{-}SeriousEffect\text{-}Age\text{-}Gender\text{-}Weight$

which enter the learning as arguments to the relation "$associatedWith$" depending on the experiment.

In addition, we used ontologies about drugs (WHO Drug Dictionary [4]) and adverse events (MedDRA [3]), and added the "isA" relations to the relational learning technique as a form of background knowledge. An *ontology* is a thesaurus [8] that answers the question of "what there is" [6] in a domain. Ontologies usually structured according to the subsumption relation ("isA"), but they may contain other relations between the terms as well. We used Babylon Knowledge Explorer (BKE) [5] to access the ontologies.

Identification and selection of terms in the corresponding ontologies was a challenge due to problems caused by: (1) variations in the spelling and misspellings, (2) phrases, that is terms that are recognized in groups, (3) the use of abbreviations in the text etc. Examples include the use of "novo-cimetine" instead of "novocimetine" that is in WHO DD, "IV fluids" instead of "I.V. solutions", "meperidine hcl" instead of "meperidine hydrochloride". In some cases, we used the main ingredient as in the case of "coffee" instead of "caffeine". We resolved these manually by selecting the appropriate term.

3 Relational Learning

Reports of drug adverse events contain information that belong to various domains such as, drugs, adverse events, concomitant medication, and patient demographics. Multiple sources of information are related to each other within the context of each report. Therefore reports lend themselves to a relational representation. We framed the task of learning drug adverse reactions as learning the concept of causality of adverse events under the reported conditions.

Our learning task was *to find a classification of the adverse events each drug is associated with, with or without the presence of other medication and with or without the presence of other information including the gender, age, weight, and the seriousness of the adverse event.* In addition, we measured the effects of incorporating additional sources of knowledge on the accuracy and coverage of the rules generated by the relational learning algorithm FOIL.

3.1 First Order Inductive Learning

We used the learning algorithm FOIL [7] to learn first-order relations from the data. FOIL is an efficient and a widely used top-down inductive learning system that can learn function-free Horn clauses. It uses the information-gain metric to select the best literal to add to a clause.

We used the following background relations to represent reports:

- $associatedWith(Drug, AdverseEvent, ConcomitantMedication)$: This predicate indicates that there exists a report in which drug $Drug$ is associated with the adverse reaction named $AdverseEvent$ when used with $ConcomitantMedication$.
- $isaD(Drug1, Drug2)$: Each of these Boolean predicates indicate that $Drug1$ is related to $Drug2$ by an "isa" relation, which we inferred using ontologies.
- $isaA(AdverseEvent1, AdverseEvent2)$: Similarly, these Boolean predicates indicate that $AdverseEvent1$ is related to $AdverseEvent2$ by an "isa" relation inferred from ontologies.
- $isaO(ConcomitantMedication1, ConcomitantMedication2)$: These Boolean predicates indicate that $ConcomitantMedication1$ is related to $ConcomitantMedication2$ by an "isa" relation inferred from ontologies.

4 Experiments

We can think of each adverse events report as evidence for possible association between the entities described in it. For example, given that we have an adverse event report for BAYCOL- a drug used to treat high cholesterol-, the adverse events specified in the report may share a causal relation with BAYCOL and other terms that are mentioned in the report.

The addition of "isA" relations from the WHO Drug Dictionary and MedDRA to the data set generated from the reports enabled FOIL to relate terms that are at different levels of generalization. We had three separate relations defined for drugs (isaD), adverse events (isaA), and concomitant or other medication (isaO). In addition, we defined the target relation for learning as $associatedWith$, which can take up to 7 attributes ($Drug$, $AdverseEvent$, $ConcomitantMedication$, $SeriousEffect$, Age, $Gender$, $Weight$) depending on the experiment.

The assumption of a closed world is used in FOIL to generate negative examples that do not belong to a relation when negative examples are not given explicitly. However, this assumption does not hold when continuous variables are used. Selection of appropriate sized negative example set is important for learning. We performed experiments to tailor this idea according to the data set we have.

We conducted the following set of experiments. Initially, for each separate withdrawn drug, we learned the $associatedWith$ relation. Later, we attempted to find rules that are close to the causal links that lead to the withdrawal. For example, Pondimin was withdrawn from the market due to its link to heart valve damage. Thus, from our experiments, we expect to find rules of the form:

$$associatedWith(Pondimin, A, C) \leftarrow isaA(A, HeartDamage), isaO(C, Diuretics),$$ [1]

which states that Pondimin is associated with an adverse event that is of type heart damage when it is taken with a medication that is of type diuretics. For the experiments, we used reports for the following set of drugs:

CISAPRIDE (85 reports), HISMANAL (73 reports), PONDIMIN (24 reports), RAXAR (3 reports), REDUX (16 reports), REZULIN (18 reports), and SELDANE (147 reports).

Then we explored the change in accuracy, coverage, and computational cost when we incorporate different sources of information into learning such as age and gender. Lastly, we used the rules learned from the set of all withdrawn drugs to predict the withdrawal chance of other drugs. For each experiment, 5-fold cross-validation is performed.

[1] The $\leftarrow$ sign represents the implication operator in which the head appears on the left hand side and the body on the right hand side. The "," sign "and"s the relations specified in the body.

$associatedWith(Pondimin, B) \leftarrow$
$isaA(B, C),\ isaA(C, D),\ isaA(D, RenalandUrinaryDisorders).$
$associatedWith(Hismanal, Hypertension, C, D) \leftarrow$
$C <> D.$
$associatedWith(Rezulin, B, Serious, Female) \leftarrow$
$associatedWith(Rezulin, B, NotSerious, Male),\ B <>$
$HAEMOGLOBINDECREASED.$

Table 1: Some rules induced by FOIL.

$associatedWith(Hismanal, ECG\ abnormal, Lidocaine, Y, N)$
$associatedWith(Hismanal, Dyspnoea, Oxygen, Y, N)$
$associatedWith(Hismanal, QT\ prolonged, Licodaine, Y, N)$
$associatedWith(Hismanal, Cardiac\ arrest, Licodaine, Y, N)$
$associatedWith(Hismanal, Arrhythmia\ ventricular,$
$Licodaine, Y, N)$

Table 2: Some rules induced by FOIL when learning from *Hismanal* data with DAOSG attribute set.

5 Results

In Table 1, we presented some of the rules learned by FOIL, which showed the power of relational representation in exploring rules that can make use of various sources of information. The first rule found that *Pondimin* is associated with an adverse event which is a type of renal and urinary disorder. In exploring this relation, FOIL is able to use the "isaA" relation multiple times to reveal the higher level term.

The next rule found that *Hismanal* was associated with hypertension when the attributes for seriousness of the report and the gender are not equivalent. FOIL finds this relation because we use the same type for representing gender and seriousness i.e. "Y" for both serious and the gender male. Thus, inequalities are also discovered while learning.

In Figure 1, histograms present the corresponding values for the accuracy and coverage of the rules learned by FOIL for different learning tasks. At the top of the histograms, we listed the drug names. The results are separated in columns that correspond to the selection of attributes for the *associatedWith* relation. DA was used for drugs and adverse events, DAO was used when we have added information about other medication, DASG was used when we added information about the seriousness of the adverse event and the gender to learning. DAOSG was interpreted similarly as mentioned. DASAGWd corresponded to the case when we added discretized values for age and gender. The corresponding results for the experiments are given in Table 3.

We were able to get interesting rules when we performed learning on the *Hismanal* data set with the attribute set of DAOSG. Some of the rules induced by FOIL are listed in Table 2. We compared the results

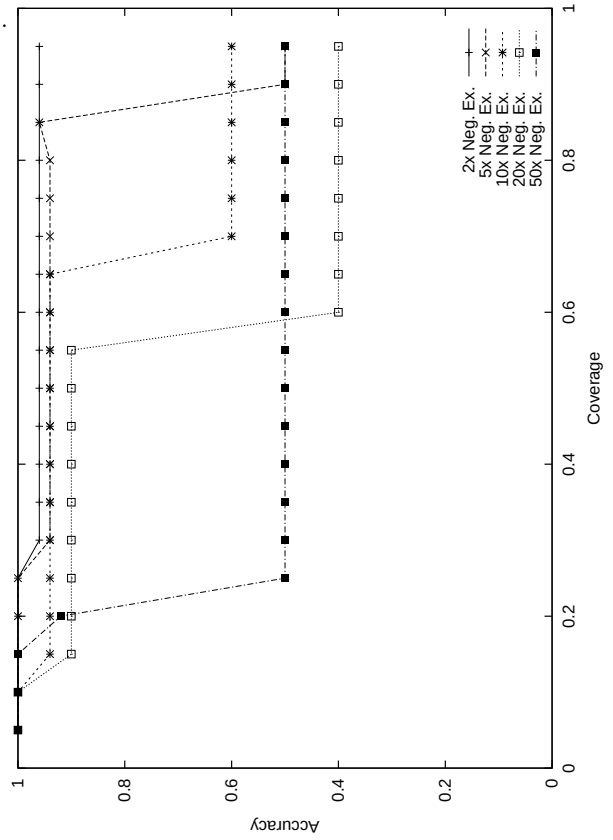

Figure 2: Pondimin

with the causal links specified by the FDA. This set is interesting because *Hismanal* was withdrawn due to cardiac arrest and arrhythmias and each of the adverse events specified in the rule set is a type of cardiac disorders.

We learn a separate set of clauses for each withdrawn drug. The accuracy-coverage plot for leaning adverse events of *Pondimin* when we vary the threshold on the minimum accuracy of clauses is given in Figure 2.

Although we can find very interesting rules from the data (at this level), their accuracy level is low compared to the other levels. This is due to the large space in which the learning is performed and the scarcity of examples to learn rules. Providing a set of negative examples also helps in learning. But we do not have any negative example available to us. We explored the effect of randomly generating the negative examples from the data in Figure 2. FOIL cannot generate negative examples for data with continuous values. This is the reason we chose to discretize the continuous values in some of our experiments.

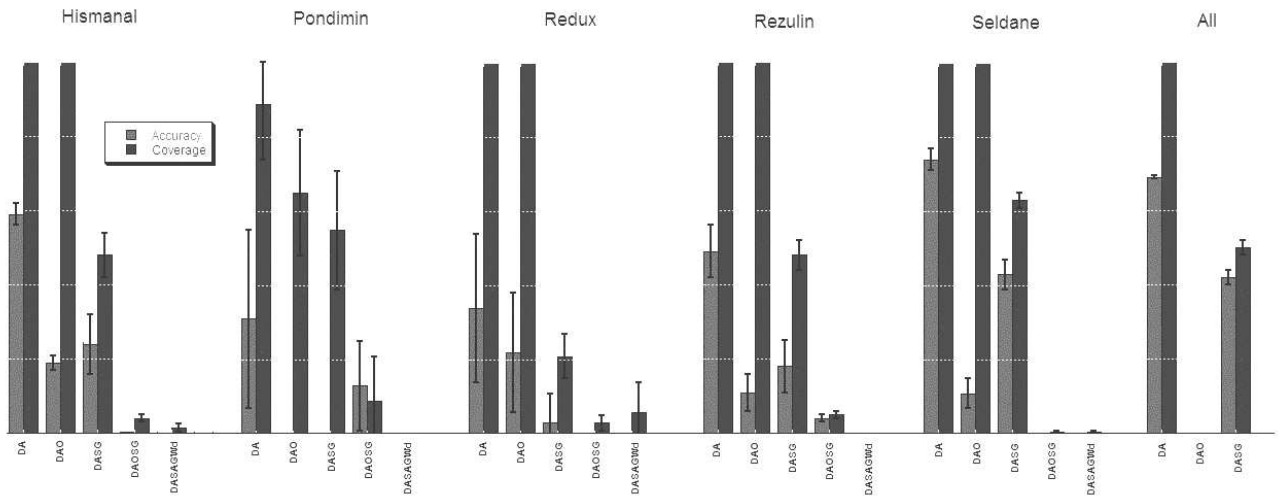

Figure 1: Accuracy-coverage of the rules learned by FOIL corresponding to the withdrawn drugs with increasing sources of information included from the adverse events reports.

Drugs	DA	DAO	DASG	DAOSG	DASAGWd
Hismanal	0.59(∓0.03)	0.19(∓0.02)	0.24(∓0.08)	0.004(∓0.005)	0
	1	1	0.48(∓0.06)	0.04(∓0.01)	0.014(∓0.01)
Pondimin	0.31(∓0.24)	0	0.0(∓0.03)	0.13(∓0.12)	0
	0.89(∓0.15)	0.65(∓0.17)	0.55(∓0.16)	0.09(∓0.12)	0
Redux	0.34(∓0.2)	0.22(∓0.16)	0.03(∓0.08)	0	0
	1	1	0.21(∓0.06)	0.03(∓0.02)	0.06(∓0.08)
Rezulin	0.49(∓0.07)	0.11(∓0.05)	0.18(∓0.07)	0.04(∓0.01)	0
	1	1	0.46(∓0.04)	0.05(∓0.01)	0
Seldane	0.74(∓0.03)	0.11(∓0.04)	0.43(∓0.04)	0	0
	1	1	0.63(∓0.02)	0.002(∓0.004)	0.004(∓0.001)
AllWithdrawn	0.69(∓0.005)	0.69(∓0.005)	0.42(∓0.02)		
	1	1	0.5(∓0.02)		

Table 3: Accuracy/coverage of rules learned by FOIL corresponding to drugs and sources of information involved.

6 Discussion

We studied how relational learning, the type of machine learning that can discover rules by using multiple relations in the language of first-order logic, can discover the relational structure inherent in drug adverse events reports. Although relational learning currently takes a long time due to its computational cost, in terms of accuracy and coverage, we were able to find good predictions that showed us the structure of the data.

Acknowledgements

This work is supported in part by GlaxoSmithKline. The author would like to thank Gary H. Merrill and Alan J. Menius of GlaxoSmithKline Data Exploration Sciences for helpful discussions and for their guidance and support during the term of this research.

References

[1] Food and Drug Administration. URL: http://www.fda.gov.

[2] Health Canada's Canadian Adverse Drug Reaction Information System (CADRIS). URL: http://www.hc-sc.gc.ca/hpfb-dgpsa/tpd-dpt/.

[3] MedDRA: Medical Dictionary for Regulatory Activities. URL: http://www.meddramsso.com.

[4] The World Health Organization Drug Dictionary (WHO DD). http://www.who-umc.org/faqs/faqdd.html.

[5] Gary H. Merrill. The Babylon Project: Toward an extensible text-mining platform. *IEEE IT Professional*, 5(2):23–30, 2003.

[6] W. V. O. Quine. *Philosophy of Logic*. Harvard University Press, second edition, 1986.

[7] J. R. Quinlan and R. M. Cameron-Jones. FOIL: A midterm report. In P. B Brazdil, editor, *Machine Learning (ECML-93) European Conference on Machine Learning Proceedings; Vienna, Austria*, pages 3–20, Berlin, Germany, 1993. Springer-Verlag.

[8] Dana S. Scott. Capturing concepts with data structures. In *Proc.DS-2, IFIP TC-2 Conference on Knowledge and Data, Portugal*, November 1986.