Anticipating Teachers' Performance

J. BARRACOSA

Instituto Superior Técnico – Technical University of Lisbon, Portugal

AND

C. ANTUNES

Instituto Superior Técnico – Technical University of Lisbon, Portugal

Data mining faces multiple defies, with improving its results through the use of domain knowledge at the top. The advent of domain driven data mining brings new techniques to use that knowledge. Educational data mining is a favored domain to explore such tools, both due to the advances in the area and the existence of domain knowledge. In this work, we propose a new methodology for anticipating teachers' performance based on the analysis of pedagogical surveys. Our approach combines classification and sequential pattern mining to identify a set of meta-patterns that can be used to enrich source data, and in this manner we better describe teachers and consequently improve classification models accuracy. The use of domain knowledge is used in two steps: in the discovery of frequent behaviors (through the use of constraints) and in the enrichment of original data. A case study on mining pedagogical surveys is presented, corroborating our argument, and showing a significant improvement in classification accuracy.

Key Words: Educational data mining, Domain knowledge, Sequential pattern mining, Pedagogical surveys

1. INTRODUCTION

Domain driven data mining (D3M) has reached a significant attention in the last five years, being considered one of the most promising paths to follow for improving mining results (Cao2010). In this context, mining is performed maintaining domain knowledge as background. Educational data mining is a privileged field to apply D3M, since domain knowledge has been gained for years in the studies performed on each particular aspect of the educational process. Among the relevant work in this field, students' modeling has deserved a considerable attention: see for example the work (Antunes2008), where a method for mining students' behaviors is proposed.

Evaluating teachers' performance is a more difficult task, being to some extent more subjective, but is made for years through pedagogical surveys. The most usual kind of *pedagogical survey* is just a set of closed-ended questions, with multiple choices that follow some order (as in the Likert scale (Likert1932)). In order to assess teacher's performance, several students individually fill a pedagogical survey, and the mean of the answers corresponds to the teacher evaluation. Despite the spread of such tools, to our knowledge there is no work on automatically predict teachers' performance based on surveys.

Authors' addresses: J. Barracosa and C. Antunes, Dep. Computer Science and Engineering, Instituto Superior Técnico, Technical University of Lisbon, Portugal. E-mail: joana.barracosa@ist.utl.pt, claudia.antunes@ist.utl.pt

Fig. 1. Methodology for mining surveys

In this work we propose a new methodology to anticipate teachers' performance based on past survey results. Our methodology explores the discovery of frequent patterns to characterize each teacher as described in (Barracosa2011), enriching his set of features, allowing for training a more informed classifier. Our argument is that pattern mining can be used as a pre-processing tool for classification, improving its accuracy.

The rest of this paper is organized as follows: section 2 proposes the new methodology, detailed descriptions of its main steps are given in sections 3, 4 and 5. A case study is presented in section 6, and the paper concludes with a critical analysis of achieved results and some guidelines for future work.

2. MINING PEDAGOGICAL SURVEYS

Returning our attention to the educational context, pedagogical surveys may be a powerful and efficient tool to gather information about teachers, when correctly used. Our goal is to predict teachers' performance, based on the analysis of past pedagogical surveys collected among students on evaluating teachers' performance. Any methodology to deal with surveys issues two major problems: how to process pedagogical surveys to characterize teachers, and which domain knowledge can be useful to the mining process.

Having previous studies (Antunes2008) demonstrated that sequential pattern mining can be successfully applied for mining students' frequent behaviors, we propose to apply the same methodology for teachers. In this context we can represent each teacher as a sequence of survey results along curricular years. On the other side, each survey result is an itemset containing the different questions, with questions being ordinal attributes.

In our methodology, we will start by pre-processing the data for solving the first problem, and then we will apply sequential pattern mining to find teachers' frequent results – *patterns*. From the set of the identified patterns is then possible to define a set of *meta-patterns*, useful for introducing some knowledge to enrich the training step. The last step is just to incorporate acquired knowledge and train a classifier for anticipating the performance of each teacher in the following semester (Fig. 1).

3. SURVEYS PRE-PROCESSING

In this work we are just interested in surveys with closed-ended questions with multiple choices that follow some order. Since pedagogical surveys are filled by several students individually, the first step is to get the global evaluation collected for each teacher on each time period. From the application of basic statistics (*mean* or *mode*, for example) to the set of individual questionnaires evaluating a single teacher for a single time period, we determine the set of *items* that characterize teacher's performance in that period. In particular, each item corresponds to the proposition verified for each question. For example, a question about security on presenting contents that can be filled with *Negative*, *Neutral* or *Positive*, would raise three different items.

Since surveys in general are long sets of questions with a reasonable odd number of choices, it is expected to have a very large number of distinct items. Indeed, in the worst case there could exist q^n items, with q the number of questions and n the number of choices for each question. For 20 questions with just 5 options it would exist 3.200.000 items. In order to control this number, and the usual pattern explosion, we propose the reduction of the number of possible values to three fixed values, following an *equidistant strategy*. Equidistant means that something is equally distant from some point: so, we create a new scale just by compacting n possible contiguous values, with n the ratio between the number of choices in the initial and final scale.

4. PATTERN MINING OVER SURVEYS

Sequential pattern mining aims to find all the frequent sequences that exist in a dataset. In this context, a *sequence* is an ordered set of *itemsets*, with an itemset being just a set of propositions (*items*) that are verified simultaneously. Sequences are perfect for modeling behaviors, since their itemsets can represent co-occurring behaviors along time.

With teachers represented as temporal sequences is then possible to apply sequential pattern mining to identify frequent performances. At first, we will apply sequential pattern mining algorithm to our data, in order to discover frequent patterns. The patterns discovered are just frequent results revealed by teachers in the past.

While the surveys considered by our methodology only include *ordinal attributes* (attributes with values among which there is a partial order), those patterns will describe trends (how teachers evolve along time), and the correlation among attributes.

Indeed, each sequence derived from surveys can be viewed as a conjunction of a sequence of propositions for the same variable. For example, if we consider only two questions in a survey for three semesters, we may have the following sequence

<(*Q1=good*, *Q2=bad*), (*Q1=good*, *Q2=regular*), (*Q1=good*, *Q2=good*)>

which can be seen as the conjunction of the two individual sequences

<(*Q1=good*),(*Q1=good*),(*Q1=good*)> and <(*Q2=bad*),(*Q2=regular*),(*Q2=good*)>

This pattern reveals that the value of *Q1* does not affect the value of *Q2*, and that *Q1* tends to be constant.

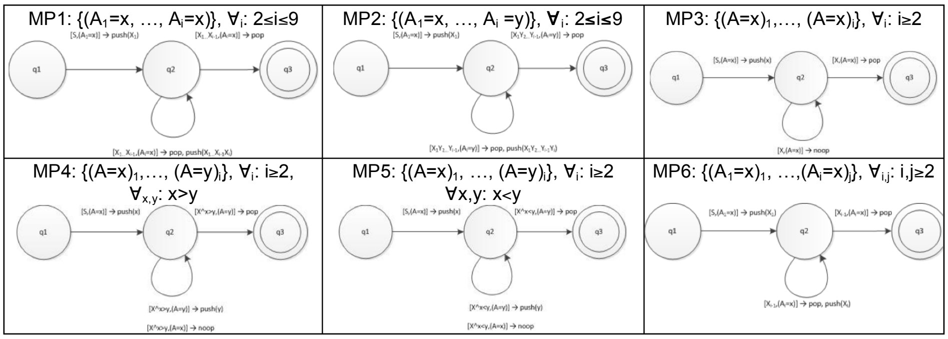

Fig. 2. Meta-patterns for mining surveys

Given that our ultimate goal is to predict teachers' evolution, this means, their following performances, the most interesting patterns are the ones that can contribute to distinguish among teachers that improve their performance against the ones that do not.

Having this is mind we are interested on describing those patterns in a more compact way, for example as a context-free language (represented by automata as used in (Antunes2008)). From the patterns discovered with sequential pattern mining, we identified a set of meta-patterns. In this context, a *meta-pattern* (*MP* for short) is just a generalization of a set of patterns, and can be seen as the trend that those patterns follow.

Pushdown automata in Fig. 2 establish the set of meta-patterns identified. In each meta-pattern, relations among the values of specific questions are established, both for a single semester and over time. *MP1* establishes that some attributes (questions) have the same value for a particular semester, while *MP2* establishes that the presence of a particular value for one attribute implies a different value for other attribute in the same semester. *MP3*, *MP4* and *MP5* establish that a particular attribute have the same, higher or lower values, along consecutive semesters. At last, MP6 describe the impact of the value of an attribute in one semester, in the value of other attributes in the following semesters.

5. CLASSIFICATION AND ANTICIPATION

Classification is usually performed through the training of a classification model, for example using algorithms for mining decision trees. In the case of pedagogical surveys, a training dataset can be created using the *n* propositions (*question=value*, with the last one corresponding to the global result) resulting from each survey for the *m* instants of evaluation (semesters for example), totalizing *n*x*m* attributes, with the class being the global result for the last semester. In this manner is possible to train a model to anticipate the next performance for each teacher, given his past behavior.

However, since traditional classifiers only deal with tabular data, all teachers have to be defined by the same number of evaluation instants, and the relation among the same variable along time and the relation among different co-occurring variables are both lost. In order to overpass this lost, meta-patterns can be used efficiently to enrich the dataset. A new dataset can be obtained from the previous one, enlarged by more *k* Boolean attributes, one for each meta-pattern. Each instance attribute *n+i* is then filled with the *true* value whenever MP_i accepts it. This way, the new instances take the next format:

q34Good,q35Neutral,q36Good,… ,T,F,F,F,F,T,F,F,F,T,F,

Since excellence is rare, as well as mediocrity, datasets resulting from preprocessing pedagogical surveys tend to be unbalanced (a considerable difference in the number of instances for each class). In order to deal with that, balancing techniques, such resampling and SMOTE (Chawla2002), should be used to train the classification model.

6. CASE STUDY

The data used to demonstrate the applicability of our proposal, was collected from pedagogical surveys taken place at a Portuguese university, along 8 years, corresponding to sequences with 15 events. Each tuple (teacher, semester) is represented as an itemset as described. Each survey has nine questions about teachers' behaviors, with five options each, totalizing 9^5 possible items. Classification was done with the data mining open source software Weka 3.6.4, with decision trees and the algorithm C4.5.

As we can see in Fig. 3 (left), this dataset is highly unbalanced. To help us balancing the data, we used the techniques referred earlier: *Resample*, *Spread SubSample* and *SMOTE*. All these techniques were applied before the creation of the classifiers and produce different datasets distributions. In order to achieve a balanced dataset it is necessary to

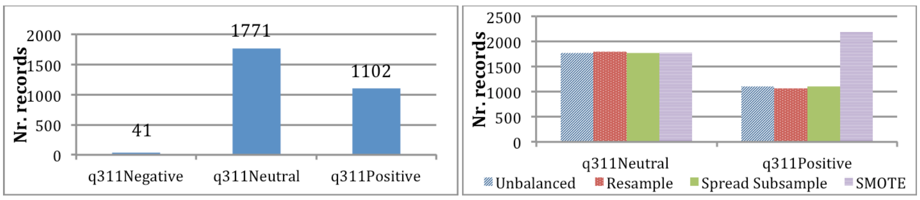

Fig. 3. Distribution of the dataset (left) and balanced dataset without the value Negative (right)

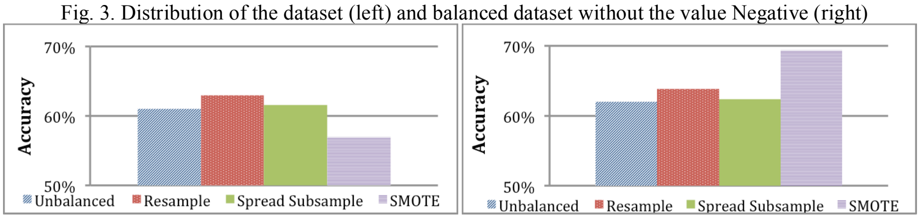

Fig. 4. Model accuracy in the original dataset (left) and in presence of MPs (right)

remove the instances classified as *Negative*. In Fig. 3-right we can see the results of different balancing techniques without negative instances, SMOTE was the technique that achieved the best balancing results.

In order to evaluate the importance of meta-patterns in classification, we measure the model accuracy for a dataset enriched with MPs as described above, and for the original one. In both cases, negative instances were discarded, obtaining a more balanced dataset.

As we can see through the results (Fig. 4), the presence of MPs increases the model accuracy for all techniques, but in particular for SMOTE, where it increased 12%, reaching 69% of global accuracy.

7. CONCLUSIONS AND FUTURE WORK

In this work we proposed to find teachers' frequent behaviors and use them to anticipate their performance in the following semester. It is shown that, by combining sequential pattern analysis with classification, is possible to improve the anticipation accuracy, demonstrating that the presence of relevant domain knowledge can contribute in a significant way to generally improve classification results.

The proposed methodology depends on the ability to generalize the discovered patterns identifying a set of meta-patterns. However, whenever data came from surveys as described above, in particular for data characterized by ordinal attributes, the meta-patterns proposed in this work are applicable. Discovered frequent patterns can be used to characterize each teacher in order to anticipate his future performance.

The identification of meta-patterns was done manually, and continues to be a hard and prohibitive task, in general situations. The next step in this work is the development of an automatic tool for identifying meta-patterns from any set of discovered patterns. First approaches will be based on the *Onto4AR* framework (Antunes2009).

8. ACKNOWLEDGEMENTS

This work is partially supported by Fundação para a Ciência e Tecnologia under research project *educare* (PTDC/EIA-EIA/110058/2009).

9. REFERENCES

ANTUNES, C. 2008. Acquiring Background Knowledge for Intelligent Tutoring Systems. In *Proc Int'l Conf on Educational Data Mining (EDM 2008)*, 18-27.

ANTUNES, C. 2009. Mining Patterns in the Presence of Domain Knowledge. In *Proc Int'l Conf on Enterprise Information Systems*, 188-193.

BARRACOSA, J. AND ANTUNES, C. 2011. Mining Teaching Behaviors from Pedagogical Surveys. In *Proc Int'l Conf on Educational Data Mining (EDM 2011)*.

CAO, L. 2010. Domain-Driven Data Mining: Challenges and Prospects. *IEEE Transactions on Knowledge and Data Engineering*, 22(6), 755-769.

CHAWLA, N.V., BOWYER, K.W., HALL, L.O. AND KEGELMEYER, W.P. 2002. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16, 321--357.

LIKERT, R. 1932. A Technique for the Measurement of Attitudes. *Archives of Psychology*, 22(140), 1-55.