

Classifier Generation by Combining Domain Knowledge and Machine Learning

Violeta Mirchevska

Department of Intelligent Systems, Jožef Stefan Institute, Jamova 39, Ljubljana, Slovenia

E-mail: violeta.mircevaska@ijs.si

Thesis Summary

Keywords: prior domain knowledge, inductive machine learning, small-dataset learning, ambient intelligence

Received: March 25, 2014

This article presents a summary of the doctoral dissertation of the author, which addresses the task of classifier generation by combining domain knowledge and machine learning.

Povzetek: Prispevek predstavlja povzetek doktorske disertacije avtorice, ki obravnava naloge kreiranja klasifikatorjev s kombiniranjem domenskega znanja in strojnega učenja.

1 Introduction

The field of machine learning (ML) is concerned with the development of algorithms that enable computer programs to learn and automatically improve with experience [1]. ML algorithms have been successfully applied to a wide variety of domains, such as credit-card fraud detection, book recommendations and creating helicopter control logic. They may automatically extract comprehensive concept models solely from concept examples, finding even patterns that are too subtle to be detected by humans. However, their success greatly depends on the quality and the completeness of the available concept examples.

Despite the exponential growth of digital data, there are still domains for which data is scarce. We assume there are at least two reasons for scarce data: (1) sufficient general-purpose data may be costly or otherwise difficult to obtain, possibly due to great domain variation, and (2) general-purpose data may be inappropriate for some deployments, for example, because they are user-specific. One such domain is fall detection. The available training data in the fall-detection domain partially captures the domain's properties because it is difficult to record fall examples due to ethical issues and injury danger. In addition, for reliable performance, fall-detection classifiers need to be tuned to user-specific data. A classifier which suits a user with diminished motor skills may not be appropriate for a user which regularly exercises in the living room.

When learning from training examples which partially capture the domain properties, the learner may create a classifier from patterns, which although representative of the available examples, are not characteristic for the learned concept [2]. Such classifier would perform poorly in real life because it does not capture the essence of the learned concept. This issue may be partially tackled by introducing domain

knowledge (DK) as an additional information source in the learning process. Expert DK complements ML as it may contain patterns which are not captured in the available concept examples. An expert may verify a classifier's patterns and/or supplement them with patterns from DK. Therefore, classifier generation by a combination of DK and ML is beneficial in domains with insufficient general-purpose data.

Classifier adaptation to user needs may be performed online after system deployment when real-life user-specific data becomes available. In order to pose minimal burden on the user, we consider obtaining user-specific data through occasionally given user feedback which contains information about false negatives (i.e., the system did not detect the class of interest when there was one) or false positives (i.e., the system detected the class of interest when there was none). Such user feedback may be considered as a reward signal given to the system. Learning from rewards is often applied in sequential decision making domains, where the reward function is considered as the most parsimonious description of a task. Online classifier adaptation may be represented as a sequential decision-making task. Therefore, rewards extracted from user feedback may be used for online classifier adaptation in the cases when user-specific data may be obtained after deployment.

The dissertation [2] proposes a novel method, named CDKML (Combining Domain Knowledge and Machine Learning), for classifier generation in the case of scarce data. It combines DK and ML when learning from insufficient general-purpose data, and leverages user feedback for online classifier adaptation to user needs.

The article is organized as follows. Section 2 gives an overview of the CDKML method. Section 3 summarizes the evaluation results. The dissertation's

scientific contributions are outlined in Section 4 together with plans for future work.

2 The CDKML method

The CDKML method [3] is a three-phase approach to learning consisting of initialization, refinement and online adaptation. In the initialization phase, an expert specifies a set of patterns important for distinguishing the concept of interest. The patterns may be extracted from DK or be obtained using interactive data mining. In the refinement phase, an optimization algorithm is used for finding the most suitable general-purpose pattern-parameter values by maximizing the classifier's accuracy on available general-purpose data. In the online adaptation phase, user feedback is used to fine-tune the pattern-parameter values to user needs. The online adaptation problem is formulated as a Markov decision process.

3 Evaluation

CDKML was evaluated in three behavior modeling domains: behavioral cloning, posture recognition and fall detection. Its performance was compared to the performance of five classifiers built using ML algorithms in Weka [4]: SMO, RandomForest, NaiveBayes, JRip and J48.

CDKML's refined classifier showed the best performance in the fall-detection domain where it considerably outperformed all five ML algorithms, the posture-recognition domain followed, while it did not show improvement in comparison to standard ML in the behavioral-cloning domain. We attribute the improvement in performance primarily to the contribution of the expert in CDKML's initialization phase where the expert extracted the classifier's patterns using DK and interactive data mining. The improvement was the most evident in the fall-detection domain where DK provided clear instructions: "If a person is lying or sitting on the ground for a longer period of time then a fall happened". Formulating the patterns for the posture-recognition classifier was, however, not simple. In this case, interactive data mining played an important role, helping the expert to incorporate DK into the classifier. In the behavioral-cloning domain, we did not have available DK.

The evaluation of CDKML's online adaptation phase showed that the proposed online adaptation approach is capable of adjusting the refined classifier to correctly recognize events not present in the available general-purpose examples, making tradeoffs between contradictory user feedback based on the cost of each misclassification.

4 Conclusions

The dissertation addresses the problem of classifier generation from scarce data. It proposes a new, three-phase method, named CDKML, for extraction of reliable classifiers in domains where the training examples partially represent the domain properties, but human

experts can contribute with their DK. The main contributions of the dissertation are:

- A novel method, named CDKML, for classifier generation and online adaptation which leverages both ML and DK. The novelty is in the way of integration of three phases: initialization, refinement and online adaptation;
- A novel classifier adaptation based on user feedback using Markov decision processes. This, third phase of the CDKML method, is novel on its own.

CDKML achieved higher accuracy than classical ML algorithms when learning from scarce data by leveraging the available DK and user feedback.

As future work, we plan to examine two CDKML improvements. First, exploitation of DK captured in ontologies needs to be considered. The Web offers huge amounts of unstructured, textual data, and approaches to extracting domain patterns and ontology development from that kind of data are emerging [5]. It would be interesting to research possibilities for automating CDKML's initialization by utilizing DK available on the Web. Second, CDKML's online classifier adaptation relies only on user feedback. However, the more real-life examples of the learned concept become available, the better the capability of ML to induce a reliable concept classifier. CDKML's online adaptation may be accompanied with ML classifier re-induction. A combination of the two classifiers in which the ML classifier's influence on the final classification increases as more data becomes available seems reasonable.

Acknowledgement

The research leading to the dissertation was partially financed by the European Union, European Social Fund.

References

- [1] T. M. Mitchell (1997). *Machine Learning*. McGraw-Hill, Inc., New York, NY, USA.
- [2] V. Mirchevska (2013) *Behavior Modeling by Combining Machine Learning and Domain Knowledge*, PhD Thesis, IPS Jožef Stefan, Ljubljana, Slovenia.
- [3] V. Mirchevska, M. Luštrek, M. Gams (2013) Combining domain knowledge and machine learning for robust fall detection. *Expert Systems*, preprint published online.
- [4] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, I. H. Witten (2009) The Weka data mining software: An update. *SIGKDD Explorations Newsletter* 11, pp. 10-18.
- [5] B. Dalvi, W. W. Cohen, J. Callan (2012) Collectively representing semi-structured data from the Web. *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-Scale Knowledge Extraction*, Association of Computational Linguistics, Stroudsburg, PA, USA, pp. 7-12.