THE 18TH EUROPEAN CONFERENCE ON MACHINE LEARNING
AND
THE 11TH EUROPEAN CONFERENCE ON PRINCIPLES AND PRACTICE
OF KNOWLEDGE DISCOVERY IN DATABASES

PROCEEDINGS OF THE PLANNING TO LEARN WORKSHOP

PlanLearn-07

September 17, 2007

Warsaw, Poland

Editors:
Pavel Brazdil, University of Porto
Abraham Bernstein, University of Zurich

Typesetting:
Rui Leite, University of Porto

Preface

The task of constructing composite systems, that is systems composed of more than one part, can be seen as interdisciplinary area which builds on expertise in different domains. The aim of this workshop is to explore the possibilities of constructing such systems with the aid of Machine Learning and exploiting the know-how of Data Mining. One way of producing composite systems is by inducing the constituents and then by putting the individual parts together.

For instance, a text extraction system may be composed of various subsystems, some oriented towards tagging, morphosyntactic analysis or word sense disambiguation. This may be followed by selection of informative attributes and finally generation of the system for the extraction of the relevant information. Machine Learning techniques may be employed in various stages of this process.

The problem of constructing complex systems can thus be seen as a problem of planning to resolve multiple (possibly interacting) tasks. So, one important issue that needs to be addressed is how these multiple learning processes can be coordinated. Each task is resolved using certain ordering of operations. Meta-learning can be useful in this process. It can help us to retrieve previous solutions conceived in the past and re-use them in new settings.

Of particular interest are methods and proposals that address the following issues:

- Planning to construct composite systems,
- Exploitation of ontologies of tasks and methods,
- Representation of learning goals and states in learning,
- Control and coordination of learning processes,
- Recovering / adapting sequences of DM operations,
- Meta-learning and exploitation of meta-knowledge,
- Layered learning,
- Multi-task learning,
- Transfer learning,
- Multi-predicate learning (and other relevant ILP methods),
- Combining induction and abduction,
- Multi-strategy learning,
- Learning to learn.

The aim of the workshop is to explore the possibilities of this new area, offer a forum for exchanging ideas and experience concerning the state-of-the art, permit to bring in knowledge gathered in different but related and relevant areas and outline new directions for research.

Warsaw, September 2007

Pavel Brazdil
Abraham Bernstein

Workshop Organization

Workshop Chairs

Pavel Brazdil (University of Porto)
Abraham Bernstein (University of Zurich)

ECML/PKDD Workshop Chair

Marzena Kryszkiewicz (Warsaw University of Technology)

Workshop Program Committee

Abraham Bernstein, Switzerland
Pavel Brazdil, Portugal
Christophe Giraud-Carrier, USA
Peter Flach, Great Britain
Larry Hunter, USA
Rui Leite, Portugal
Katharina Morik, Germany
Oliver Ray, Great Britain
Ashwin Ram, USA
Luc de Raedt, Belgium
Michele Sebag, France
Carlos Soares, Portugal
Maarten van Someren, The Netherlands
Ricardo Vilalta, USA

Table of Contents

Special Contributions

Main Contributions

Work in Progress

20 Years of Planning to Learn

Lawrence Hunter

Center for Computational Pharmacology
University of Colorado School of Medicine
Box 6511, Mail stop 8303, Aurora, CO 80045-9802, USA
Larry.Hunter@uchsc.edu

Full article in PDF

Towards Intelligent Assistance for a Data Mining Process

Abraham Bernstein

University of Zurich
bernstein@ifi.uzh.ch

Full article in PDF

Designing Complex Systems: Role of Learning and Domain Specific Meta-knowledge

Pavel Brazdil

[1] LIAAD - INESC Porto L.A., University of Porto, Rua de Ceuta, 118-6Âş, 4050-190, Porto, Portugal; pbrazdil@liacc.up.pt
[2] Faculty of Economics, University of Porto, Porto, Portugal

Full article in PDF

Meta-Learning Rule Learning Heuristics

Frederik Janssen and Johannes Fürnkranz

TU Darmstadt, Knowledge Engineering Group
Hochschulstraße 10, D-64289 Darmstadt, Germany
[janssen,juffi]@ke.informatik.tu-darmstadt.de

Abstract. The goal of this paper is to investigate to what extent a rule learning heuristi c can be learned from experience. Our basic approach is to learn a large number of rules and record their performance on the test set. Subsequently, we train regression algorithms on predicting the test set performa nce from training set characteristics. We investigate several variations of this basic scenario, including the question whether it is better to predict the performance of the candidate rule itself or of the resulting final rule. Our experiments on a number of independent evaluation sets show that the learned heuristics outperform standard rule learning heuristics. We also analyze their behavior in coverage space.

Full article in PDF

Evolutionary Learning with Cross-Class Knowledge Reuse for Handwritten Character Recognition

Wojciech Jaśkowski, Krzysztof Krawiec, and Bartosz Wieloch

Institute of Computing Science, Poznan University of Technolo gy
Piotrowo 2, 60965 Poznań, Poland
{wjaskowski|kkrawiec|bwiel och}@cs.put.poznan.pl

Abstract. We propose a learning algorithm that reuses knowledge acquired in past learning sessions to improve its performance on a new learning task. The method concerns visual learning and uses genetic programming to represent hypotheses, each of them being a procedure that processes visual primitives derived from the training images. The process of recognition is generative, i.e., a procedure is supposed to restore the shape of the processed object by drawing its reproduction on a separate canvas. This basic method is extended with a knowledge reuse mechanism that allows learners to import genetic material from hypotheses that evolved for the other decision classes (object classes). We compare both methods on a task of handwritten character recognition, and conclude that knowledge reuse leads to significant improvement of classification accuracy and reduces the risk of overfitting.

Full article in PDF

An Iterative Process for Building Learning Curves and Predicting Relative Performance of Classifiers

Rui Leite and Pavel Brazdil

rleite@liacc.up.pt, pbrazdil@liacc.u.pt
LIAAD-INESC Porto L.A./Faculty of Economics,
University of Porto, Rua de Ceuta, 118-6,
4050-190 Porto, Portugal

Abstract. This paper concerns the problem of predicting the relative performance of classification algorithms. Our approach requires that experiments are conducted on small samples. The information gathered is used to identify the nearest learning curve for which the sampling procedure was fully carried out. This allows the generation of a prediction regarding the relative performance of the algorithms. The method automatically establishes how many samples are needed and their sizes. This is done iteratively by taking into account the results of all previous experiments - both on other datasets and on the new dataset obtained so far. Experimental evaluation has shown that the method achieves better performance than previous approaches.

Full article in PDF

Towards Automating Goal-driven Learning

Maarten W. van Someren

Department of Human-Computer Studies
University of Amsterdam
Kruislaan 419
1098 VA Amsterdam
The Netherlands
maarten@science.uva.nl

Abstract. A number of computational tools exist to perform parts of the task of constructing knowledge-based systems. These tools come from research and development in Machine Learning and Knowledge Acquisition. More fully automating this process requires a language for specifying the goal of the construction process and knowledge for when and how to invoke a tool. This paper presents a language, LML, for formulating what we call learning goals, in the sense of descriptions of target systems. LML can also be used to express functional models of learning and acquisition tools. These models can be used to select tools to achieve a goal and to prove that the resulting knowledge system satisfies the goal. To demonstrate the use of the language, we present a system that effectively uses these functional models to select combinations of tools for automated construction of knowledge-based systems.

Full article in PDF

Learning to Evaluate Conditional Partial Plans

Sławomir Nowaczyk and Jacek Malec

Department of Computer Science, Lund University
{slawek|jacek}@cs.lth.se

Full article in PDF

Author Index