

Guest editor's introduction: special issue of the ECML PKDD 2013 journal track

Hendrik Blockeel · Kristian Kersting ·
Siegfried Nijssen · Filip Železný

Received: 23 June 2013 / Accepted: 26 June 2013 / Published online: 4 July 2013
© The Author(s) 2013

The 2013 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECMLPKDD) was held in Prague, September 2013. Following the long-standing tradition of the series, it brought together researchers in Machine Learning and Data Mining for an exciting program of cutting-edge research results, aiming in particular at a cross-fertilization between these two areas. The 2013 edition featured a record-breaking 138 full presentations of novel research results, selected through a careful peer review process from 629 submissions.

For the first time, the conference used a mixed submission model. Work could be submitted as a journal article to one of two journals participating in this model (*Machine Learning*

H. Blockeel (✉) · S. Nijssen
Dept. of Computer Science, KU Leuven, Celestijnenlaan 200A, 3001 Leuven, Belgium
e-mail: hendrik.blockeel@cs.kuleuven.be

K. Kersting
Institute for Geodesy and Geoinformation, University of Bonn, Meckenheimer Allee 172, 53115 Bonn,
Germany
e-mail: kersting@igg.uni-bonn.de

K. Kersting
Dept. of Knowledge Discovery, Fraunhofer IAIS, Schloss Birlinghoven, 53754 Sankt Augustin,
Germany
e-mail: kristian.kersting@iais.fraunhofer.de

S. Nijssen
Leiden Institute for Advanced Computer Science, Universiteit Leiden, Niels Bohrweg 1, 2333 CA,
Leiden, The Netherlands
e-mail: siegfried.nijssen@cs.kuleuven.be

F. Železný
Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, 16627 Prague,
Czech Republic
e-mail: zelezny@fel.cvut.cz

and *Data Mining and Knowledge Discovery*), or it could be submitted for publication in the conference proceedings. A total of 86 original manuscripts were submitted to Machine Learning. Eight of these were accepted for inclusion in this issue, and for presentation at the conference. We briefly summarize them in alphabetical order of authors.

Nicola Barbieri, Giuseppe Manco, Ettore Ritacco, Marco Carnuccio and Antonio Bevacqua, in *Probabilistic Topic Models for Sequence Data*, propose and explore extensions of Latent Dirichlet Allocation that do not treat texts as bags of words, as is usually done, but as sequences of words. They demonstrate that by taking sequence information into account in this manner, better results can be obtained in terms of recommendation or ranking of documents.

In *Block coordinate descent algorithms for large-scale sparse multiclass classification*, Mathieu Blondel, Kazuhiro Seki and Kuniaki Uehara propose novel ways to learn compact and fast-to-evaluate multiclass models via mixed-norm regularization. Specifically, based on a multiclass extension of the squared hinge loss, they develop block coordinate descent algorithms tailored towards large-scale and high-dimensional problems such as document classification. Experimentally, the resulting approach is shown to perform favorably to other solvers.

The flip-the-state transition operator for restricted Boltzmann machines, by Kai Brügge, Asja Fischer, and Christian Igel, advances the state of the art in learning restricted Boltzmann machines. The authors observe that the Markov chain Monte Carlo methods that are typically used in this context, and which are based on Gibbs sampling, may perform suboptimally due to a low mixing rate. They propose a novel transition operator, called *flip-the-state*, which can improve mixing without introducing computational overhead. A theoretical analysis confirms that the operator induces a converging Markov chain, and an experimental evaluation confirms that the better mixing properties of the operator lead to better learning results, compared to standard Gibbs sampling.

ROC Curves in Cost Space, by José Hernández-Orallo, Peter Flach and Cèsar Ferri, is a study of the connections between two different spaces used for evaluating and visualizing classifier performance: ROC space, which visualizes false and true positive rates in terms of operating conditions, and cost space, which visualizes misclassification costs in terms of these conditions. The authors clarify the relationship between these two by introducing so-called rate-driven curves in cost space, which are argued to be the equivalent of ROC curves in cost space. The article introduces concepts that allow for a richer use of the cost space.

In *A Comparative Evaluation of Stochastic-based Inference Methods for Gaussian Process Models*, Maurizio Filippone, Mingjun Zhong, and Mark Girolami discuss the difficulties involved in using Markov chain Monte Carlo methods for inference in Gaussian process models, and experimentally compare multiple strategies that try to tackle the mentioned problems. In addition to making their own contribution to inference in Gaussian process models, they provide an insightful overview of the current state of the art in this area, and identify directions for further research.

Nico Piatkowski, Sangkyun Lee and Katharina Morik, in *Spatio-Temporal Random Fields: Compressible Representation and Distributed Estimation*, investigate how to make discrete probabilistic graphical models practical for predicting sensor states in a spatio-temporal setting. A set of new ideas allows keeping the advantages of such models while achieving scalability. The prediction quality of the suggested methods is comparable to those of standard Markov random fields and spatio-temporal k -nearest neighbor algorithms, while using much less computational resources.

Pairwise Meta-rules for Better Meta-learning-based Algorithm Ranking, by Quan Sun and Bernhard Pfahringer, is a study in meta-learning. The authors propose a novel method

for generating meta-features, which is based on comparing the performance of base-learners on the dataset one-against-one. They identify the separate optimization of base learners for each specific dataset as a key point in this approach. In addition to this, they introduce a new meta-learner called Approximate Ranking Tree Forests. They show that these methods improve the performance of meta-learning for algorithm ranking.

Finally, Zhanglong Ji and Charles Elkan, in *Differential Privacy Based on Importance Weighting*, propose a novel method for publishing data while protecting privacy. In essence, the method adapts a dataset for which there are no confidentiality issues, by adding weights to it in such a way that the dataset becomes analogous to a given confidential dataset, in terms of statistical queries, while guaranteeing differential privacy.

The double-track publication model was introduced in an attempt to bring the thorough and efficient reviewing process of journals to the conference context, while safeguarding the possibility to have innovative work presented at the conference at an early stage. It enables authors to immediately publish their newest results in a journal, without giving up the opportunity of presenting them at a conference. We believe that this model can result in a faster, more efficient and higher-quality review process, of which all benefit: journals, conferences, authors, reviewers, and ultimately the reader.

This special issue would not have been possible without the help of many people. We thank the members of the ECML PKDD 2013 “guest editorial board”, as well as the additional reviewers, for the hundreds of reviews that they have written. Their reviewing work was exceptional in terms of both quality and timeliness, and many of the articles in this issue have improved significantly through their efforts. We thank the editor-in-chief and Springer’s staff for being open to this new publication model, which, at times, put their submission procedures under a substantial amount of stress. Finally, we thank the authors for choosing Machine Learning and ECML PKDD 2013 to publish their work.