

Introduction: special issue of selected papers from ACML 2015

Geoffrey Holmes¹ · Tie-Yan Liu² · Hang Li³ ·
Irwin King⁴ · Masashi Sugiyama⁵ · Zhi-Hua Zhou⁶

Published online: 5 April 2017
© The Author(s) 2017

We are delighted to present this special issue of Machine Learning Journal with selected papers from the Seventh Asian Conference on Machine Learning (ACML 2015) held in Hong Kong, from 20 to 22 November 2015. ACML aims at providing a leading international forum for researchers in machine learning and related fields to share their new ideas and achievements. While located in Asia, the conference has a wide visibility to the international community. ACML was the first machine learning conference with two cycles of submissions with a strict double-blind review process, and this tradition continues.

ACML 2015 received 96 submissions. Each paper was assigned with one meta-reviewer and at least three reviewers. In the end, 28 papers were accepted into the main program, accounting for an acceptance rate of 29%. This special issue contains extended versions of

✉ Tie-Yan Liu
tyliu@microsoft.com
Geoffrey Holmes
geoffrey.holmes@waikato.ac.nz
Hang Li
hangli65@hotmail.com
Irwin King
king@cse.cuhk.edu.hk
Masashi Sugiyama
sugi@k.u-tokyo.ac.jp
Zhi-Hua Zhou
zhouzh@lamda.nju.edu.cn

- 1 University of Waikato, Hamilton, New Zealand
- 2 Microsoft Research Asia, Beijing, China
- 3 Noah's Ark Lab, Huawei Technologies, Hong Kong, China
- 4 The Chinese University of Hong Kong, Shatin, Hong Kong
- 5 RIKEN and The University of Tokyo, Tokyo, Japan
- 6 Nanjing University, Nanjing, China

selected papers. The selection was made by the team of guest editors consisting of Program Chairs, General Chairs, and Steering Committee Chairs of ACML 2015, on the basis of the scientific quality and potential impact of these papers. These extended papers have been reviewed again according to the peer-review process set out by the journal criteria. In the end, six papers were selected for this special issue.

- The paper, Class-prior Estimation for Learning from Positive and Unlabeled Data, by Marthinus C. du Plessis, Gang Niu, and Masashi Sugiyama, studies the problem of estimating the class prior in an unlabeled dataset. They show that with additional samples coming only from the positive class, the class prior of the unlabeled dataset can be estimated correctly. The key idea is to use properly penalized divergences for model fitting to cancel the error caused by the absence of negative samples. The consistency, stability, and estimation error are theoretically analyzed, and the usefulness of the proposed method is experimentally demonstrated.
- The paper, Geometry-Aware Principal Component Analysis for Symmetric Positive Definite Matrices, by Inbal Horev, Florian Yger, and Masashi Sugiyama, develops a new Riemannian geometry based formulation of principal component analysis (PCA) for symmetric positive definite (SPD) matrices that preserves more data variance by appropriately extending PCA to matrix data, and extends the standard definition from the Euclidean to the Riemannian geometries. The usefulness of the proposed approach as preprocessing for electroencephalogram (EEG) signals and for texture image classification is experimentally demonstrated.
- The paper, Preference Relation-based Markov Random Fields for Recommender Systems, by Shaowu Li, Gang Li, Truyen Tran, and Yuan Jiang, tackles the problem of top-N recommendation in a novel way. Specifically, it removes the assumptions of availability of explicit feedbacks (such as ratings) and the equivalence between optimizing the ratings and optimizing the item ranking. Instead, it directly exploits preference relations, a more practical user feedback. Furthermore, the proposed approach enjoys the representational power of Markov Random Fields thus side information such as item and user attributes can be easily incorporated.
- The paper, Surrogate regret bounds for generalized classification performance metrics, by Wojciech Kotlowski and Krzysztof Dembczynski, studies the optimization of generalized performance metrics for binary classification by means of surrogate losses. They focus on a class of metrics, which are linear-fractional functions of the false positive and false negative rates, and consider a two-step learning procedure: first, a real-valued function is learned by minimizing a surrogate on the training sample, and then given the learned function, a threshold is tuned on a separate validation sample, by direct optimization of the target performance metric. They show that the regret of the resulting classifier measured by the target metric is upper bounded by the regret of the learned function measured by the surrogate loss. The theoretical findings are further analyzed in a computational study on both synthetic and real data sets.
- The paper, Maximum Margin Partial Label Learning, by Fei Yu and Min-Ling Zhang, proposes a new maximum margin formulation for partial label learning, which directly optimizes the margin between the ground-truth label and all other labels. Specifically, the predictive model is learned via an alternating optimization procedure which coordinates the task of ground-truth label identification and margin maximization iteratively. Experiments on artificial as well as real-world datasets show that the proposed approach is highly competitive to other well-established partial label learning approaches.

- The paper, Proximal Average Approximated Incremental Gradient Descent for Composite Penalty Regularized Empirical Risk Minimization, by Yiu-ming Cheung and Jian Lou, proposes a new proximal average (PA) based algorithm called IncrePA by incorporating proximal average approximation into an incremental gradient framework. The proposed method is a better PA-based method that features lower per-iteration cost, a faster convergence rate for convex composite penalties, and guaranteed convergence for even nonconvex composite penalties. Experiments on both synthetic and real datasets demonstrate the efficacy of the proposed method in solving both convex and nonconvex ERM problems with composite penalties.

This special issue would not have been possible without the contribution of many people. We wish to thank all authors for their contributions to this special issue. We also would like to express our sincere gratitude to all the referees for the time and effort in ensuring the quality of the submissions for this issue. We also wish to thank Peter Flach, editor-in-chief for MLJ, and Dragos Margineantu, editor of special issues for MLJ, for their guidance and support, as well as Melissa Fearon from the Springer team for her assistance, throughout the organization and production of this special issue.