

1. Record Nr.	UNINA9910484000803321
Titolo	Machine Learning and Knowledge Discovery in Databases : European Conference, ECML PKDD 2015, Porto, Portugal, September 7-11, 2015, Proceedings, Part I // edited by Annalisa Appice, Pedro Pereira Rodrigues, Vítor Santos Costa, Carlos Soares, João Gama, Alípio Jorge
Pubbl/distr/stampa	Cham : , : Springer International Publishing : , : Imprint : Springer, , 2015
ISBN	3-319-23528-1
Edizione	[1st ed. 2015.]
Descrizione fisica	1 online resource (LVIII, 709 p. 160 illus.)
Collana	Lecture Notes in Artificial Intelligence ; ; 9284
Disciplina	006.31
Soggetti	Data mining Artificial intelligence Pattern recognition Information storage and retrieval Database management Application software Data Mining and Knowledge Discovery Artificial Intelligence Pattern Recognition Information Storage and Retrieval Database Management Information Systems Applications (incl. Internet)
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Bibliographic Level Mode of Issuance: Monograph
Nota di contenuto	Intro -- Preface -- Organization -- Abstracts of Invited Talks -- Towards Declarative, Domain-Oriented Data Analysis -- Sum-Product Networks: Deep Models with Tractable Inference -- Mining Online Networks and Communities -- Learning to Acquire Knowledge in a SmartGrid Environment -- Untangling the Web's Invisible Net -- Towards a Digital Time Machine Fueled by BigData and Social Mining -- Abstracts of Journal Track Articles -- Contents - Part I -- Contents - Part II -- Contents - Part III -- Research Track Classification,

Regression and Supervised Learning -- Data Split Strategies for Evolving Predictive Models -- 1 Introduction -- 2 Data Splits for Model Fitting, Selection, and Assessment -- 3 Issues with Evolving Models -- 4 Data Splits for Evolving Models -- 4.1 Parallel Dump Workflow -- 4.2 Serial Waterfall Workflow -- 4.3 Hybrid Workflow -- 5 Bias Due to Test Set Reuse -- 6 Illustration on Synthetic Data -- 7 Case Study: Paraphrase Detection -- 8 Related Work -- 9 Conclusions -- A Appendix: Bias Due to Test Set Reuse -- References -- Discriminative Interpolation for Classification of Functional Data -- 1 Introduction -- 2 Function Representations and Wavelets -- 3 Related Work -- 4 Classification by Discriminative Interpolation -- 4.1 Training Formulation -- 4.2 Testing Formulation -- 5 Experiments -- 6 Conclusion -- References -- Fast Label Embeddings via Randomized Linear Algebra -- 1 Introduction -- 1.1 Contributions -- 2 Algorithm Derivation -- 2.1 Notation -- 2.2 Background -- 2.3 Rank-Constrained Estimation and Embedding -- 2.4 Rembrandt -- 3 Related Work -- 4 Experiments -- 4.1 ALOI -- 4.2 ODP -- 4.3 LSHTC -- 5 Discussion -- References -- Maximum Entropy Linear Manifold for Learning Discriminative Low-Dimensional Representation -- 1 Introduction -- 2 General Idea -- 3 Theory -- 4 Closed form Solution for Objective and its Gradient. 5 Experiments -- 6 Conclusions -- References -- Novel Decompositions of Proper Scoring Rules for Classification: Score Adjustment as Precursor to Calibration -- 1 Introduction -- 2 Proper Scoring Rules -- 2.1 Scoring Rules -- 2.2 Divergence, Entropy and Properness -- 2.3 Expected Loss and Empirical Loss -- 3 Decompositions with Ideal Scores and Calibrated Scores -- 3.1 Ideal Scores  $Q$  and the Decomposition  $L=EL+IL$  -- 3.2 Calibrated Scores  $C$  and the Decomposition  $L=CL+RL$  -- 4 Adjusted Scores  $A$  and the Decomposition  $L=AL+PL$  -- 4.1 Adjustment -- 4.2 The Right Adjustment Procedure Guarantees Decreased Loss -- 5 Decomposition Theorems and Terminology -- 5.1 Decompositions with  $S,C,Q,Y$  -- 5.2 Decompositions with  $S,A,C,Q,Y$  and Terminology -- 6 Algorithms and Experiments -- 7 Related Work -- 8 Conclusions -- References -- Parameter Learning of Bayesian Network Classifiers Under Computational Constraints -- 1 Introduction -- 2 Related Work -- 3 Background and Notation -- 4 Algorithms for Online Learning of Reduced-Precision Parameters -- 4.1 Learning Maximum Likelihood Parameters -- 4.2 Learning Maximum Margin Parameters -- 5 Experiments -- 5.1 Datasets -- 5.2 Results -- 6 Discussions -- References -- Predicting Unseen Labels Using Label Hierarchies in Large-Scale Multi-label Learning -- 1 Introduction -- 2 Multi-label Classification -- 3 Model Description -- 3.1 Joint Space Embeddings -- 3.2 Learning with Hierarchical Structures Over Labels -- 3.3 Efficient Gradients Computation -- 3.4 Label Ranking to Binary Predictions -- 4 Experimental Setup -- 5 Experimental Results -- 5.1 Learning All Labels Together -- 5.2 Learning to Predict Unseen Labels -- 6 Pretrained Label Embeddings as Good Initial Guess -- 6.1 Understanding Label Embeddings -- 6.2 Results -- 7 Conclusions -- Regression with Linear Factored Functions -- 1 Introduction -- 1.1 Kernel Regression. 1.2 Factored Basis Functions -- 2 Regression -- 3 Linear Factored Functions -- 3.1 Function Class -- 3.2 Constraints -- 3.3 Regularization -- 3.4 Optimization -- 4 Empirical Evaluation -- 4.1 Demonstration -- 4.2 Evaluation -- 5 Discussion -- Appendix A LFF Definition and Properties -- Appendix B Inner Loop Derivation -- Appendix C Proofs of the Propositions -- References -- Ridge Regression, Hubness, and Zero-Shot Learning -- 1 Introduction -- 1.1 Background -- 1.2 Research Objective and Contributions -- 2 Zero-

Shot Learning as a Regression Problem -- 3 Hubness Phenomenon and the Variance of Data -- 4 Hubness in Regression-Based Zero-Shot Learning -- 4.1 Shrinkage of Projected Objects -- 4.2 Influence of Shrinkage on Nearest Neighbor Search -- 4.3 Additional Argument for Placing Target Objects Closer to the Origin -- 4.4 Summary of the Proposed Approach -- 5 Related Work -- 6 Experiments -- 6.1 Experimental Setups -- 6.2 Task Descriptions and Datasets -- 6.3 Experimental Results -- 7 Conclusion -- References -- Solving Prediction Games with Parallel Batch Gradient Descent -- 1 Introduction -- 2 Problem Setting and Data Transformation Model -- 3 Analysis of Equilibrium Points -- 3.1 Existence of Equilibrium Points -- 3.2 Uniqueness of Equilibrium Points -- 4 Finding the Unique Equilibrium Point Efficiently -- 4.1 Inexact Line Search -- 4.2 Arrow-Hurwicz-Uzawa Method -- 4.3 Parallelized Methods -- 5 Experimental Results -- 5.1 Reference Methods -- 5.2 Performance of the Parameterized Transformation Model -- 5.3 Optimization Algorithms -- 5.4 Parallelized Models -- 6 Conclusion -- References -- Structured Regularizer for Neural Higher-Order Sequence Models -- 1 Introduction -- 2 Related Work -- 3 Higher-Order Conditional Random Fields -- 3.1 Parameter Learning -- 3.2 Forward Algorithm for 2nd-Order CRFs -- 4 Structured Regularizer -- 5 Experiments. 5.1 TIMIT Data Set -- 5.2 Experimental Setup -- 5.3 Labeling Results Using Only MLP Networks -- 5.4 Labeling Results Using LC-CRFs with Linear or Neural Higher-Order Factors -- 6 Conclusion -- References -- Versatile Decision Trees for Learning Over Multiple Contexts -- 1 Introduction -- 2 Dataset Shift -- 3 Versatile Decision Trees -- 3.1 Constructing Splits Using Percentiles -- 3.2 Adapting for Output Shifts -- 3.3 Versatile Model for Decision Trees -- 4 Experimental Results -- 4.1 Generating Synthetic Shifts -- 4.2 Results of the Synthetic Shifts -- 4.3 Results on Non-synthetic Shifts -- 5 Conclusion -- References -- When is Undersampling Effective in Unbalanced Classification Tasks? -- 1 Introduction -- 2 The Warping Effect of Undersampling on the Posterior Probability -- 3 The Interaction Between Warping and Variance of the Estimator -- 4 Experimental Validation -- 4.1 Synthetic Datasets -- 4.2 Real Datasets -- 5 Conclusion -- References -- Clustering and Unsupervised Learning -- A Kernel-Learning Approach to Semi-supervised Clustering with Relative Distance Comparisons -- 1 Introduction -- 2 Related Work -- 3 Kernel Learning with Relative Distances -- 3.1 Basic Definitions -- 3.2 Relative Distance Constraints -- 3.3 Extension to a Kernel Space -- 3.4 Log Determinant Divergence for Kernel Learning -- 3.5 Problem Definition -- 4 Semi-supervised Kernel Learning -- 4.1 Bregman Projections for Constrained Optimization -- 4.2 Semi-supervised Kernel Learning with Relative Comparisons -- Selecting the Bandwidth Parameter. -- Semi-Supervised Kernel Learning with Relative Comparisons. -- Clustering Method. -- 5 Experimental Results -- 5.1 Datasets -- 5.2 Relative Constraints vs. Pairwise Constraints -- 5.3 Multi-resolution Analysis -- 5.4 Generalization Performance -- 5.5 Effect of Equality Constraints -- 6 Conclusion -- References. Bayesian Active Clustering with Pairwise Constraints -- 1 Introduction -- 2 Problem Statement -- 3 Bayesian Active Clustering -- 3.1 The Bayesian Clustering Model -- Marginalization of Cluster Labels. -- 3.2 Active Query Selection -- Selection Criteria. -- Computing the Selection Objectives. -- 3.3 The Sequential MCMC Sampling of  $W$  -- 3.4 Find the MAP Solution -- 4 Experiments -- 4.1 Dataset and Setup -- 4.2 Effectiveness of the Proposed Clustering Model -- 4.3 Effectiveness of the Overall Active Clustering Model -- 4.4 Analysis of the Acyclic Graph Restriction -- 5 Related Work -- 6 Conclusion -- References --

ConDist: A Context-Driven Categorical Distance Measure -- 1  
Introduction -- 2 Related Work -- 3 The Distance Measure ConDist --  
3.1 Definition of ConDist -- 3.2 Attribute Distance  $d_X$  -- 3.3 Attribute  
Weighting Function  $w_X$  -- 3.4 Correlation, Context and Impact -- 3.5  
Heterogeneous Data Sets -- 4 Experiments -- 4.1 Evaluation  
Methodology -- 4.2 Experiment 1 -- Context Attribute Selection -- 4.3  
Experiment 2 -- Comparison in the Context of Classification -- 4.4  
Experiment 3 -- Comparison in the Context of Clustering -- 5  
Discussion -- 5.1 Experiment 1 -- Context Attribute Selection -- 5.2  
Experiment 2 -- Comparison in the Context of Classification -- 5.3  
Experiment 3 -- Comparison in the Context of Clustering -- 6  
Summary -- References -- Discovering Opinion Spammer Groups by  
Network Footprints -- 1 Introduction -- 2 Measuring Network  
Footprints -- 2.1 Neighbor Diversity of Nodes -- 2.2 Self-Similarity in  
Real-World Graphs -- 2.3 NFS Measure -- 3 Detecting Spammer  
Groups -- 4 Evaluation -- 4.1 Performance of NFS on Synthetic Data --  
4.2 Performance of GroupStrainer on Synthetic Data -- 4.3 Results on  
Real-World Data -- 5 Related Work -- 6 Conclusion -- References --  
Gamma Process Poisson Factorization for Joint Modeling of Network  
and Documents.  
1 Introduction.

---

## Sommario/riassunto

The three volume set LNAI 9284, 9285, and 9286 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2015, held in Porto, Portugal, in September 2015. The 131 papers presented in these proceedings were carefully reviewed and selected from a total of 483 submissions. These include 89 research papers, 11 industrial papers, 14 nectar papers, and 17 demo papers. They were organized in topical sections named: classification, regression and supervised learning; clustering and unsupervised learning; data preprocessing; data streams and online learning; deep learning; distance and metric learning; large scale learning and big data; matrix and tensor analysis; pattern and sequence mining; preference learning and label ranking; probabilistic, statistical, and graphical approaches; rich data; and social and graphs. Part III is structured in industrial track, nectar track, and demo track.

---