

Contents – Part I

Anomaly Detection

Concentration Free Outlier Detection	3
<i>Fabrizio Angiulli</i>	
Efficient Top Rank Optimization with Gradient Boosting for Supervised Anomaly Detection	20
<i>Jordan Frery, Amaury Habrard, Marc Sebban, Olivier Caelen, and Liyun He-Guelton</i>	
Robust, Deep and Inductive Anomaly Detection	36
<i>Raghavendra Chalapathy, Aditya Krishna Menon, and Sanjay Chawla</i>	
Sentiment Informed Cyberbullying Detection in Social Media.	52
<i>Harsh Dani, Jundong Li, and Huan Liu</i>	
ZOORANK: Ranking Suspicious Entities in Time-Evolving Tensors	68
<i>Hemank Lamba, Bryan Hooi, Kijung Shin, Christos Faloutsos, and Jürgen Pfeffer</i>	

Computer Vision

Alternative Semantic Representations for Zero-Shot Human Action Recognition	87
<i>Qian Wang and Ke Chen</i>	
Early Active Learning with Pairwise Constraint for Person Re-identification	103
<i>Wenhe Liu, Xiaojun Chang, Ling Chen, and Yi Yang</i>	
Guiding InfoGAN with Semi-supervision.	119
<i>Adrian Spurr, Emre Aksan, and Otmar Hilliges</i>	
Scatteract: Automated Extraction of Data from Scatter Plots	135
<i>Mathieu Cliche, David Rosenberg, Dhruv Madeka, and Connie Yee</i>	
Unsupervised Diverse Colorization via Generative Adversarial Networks	151
<i>Yun Cao, Zhiming Zhou, Weinan Zhang, and Yong Yu</i>	

Ensembles and Meta Learning

Dynamic Ensemble Selection with Probabilistic Classifier Chains	169
<i>Anil Narassiguin, Haytham Elghazel, and Alex Aussem</i>	
Ensemble-Compression: A New Method for Parallel Training of Deep Neural Networks.	187
<i>Shizhao Sun, Wei Chen, Jiang Bian, Xiaoguang Liu, and Tie-Yan Liu</i>	
Fast and Accurate Density Estimation with Extremely Randomized Cutset Networks.	203
<i>Nicola Di Mauro, Antonio Vergari, Teresa M. A. Basile, and Floriana Esposito</i>	

Feature Selection and Extraction

Deep Discrete Hashing with Self-supervised Pairwise Labels	223
<i>Jingkuan Song, Tao He, Hangbo Fan, and Lianli Gao</i>	
Including Multi-feature Interactions and Redundancy for Feature Ranking in Mixed Datasets	239
<i>Arvind Kumar Shekar, Tom Bocklisch, Patricia Iglesias Sánchez, Christoph Nikolas Straehle, and Emmanuel Müller</i>	
Non-redundant Spectral Dimensionality Reduction	256
<i>Yochai Blau and Tomer Michaeli</i>	
Rethinking Unsupervised Feature Selection: From Pseudo Labels to Pseudo Must-Links	272
<i>Xiaokai Wei, Sihong Xie, Bokai Cao, and Philip S. Yu</i>	
SetExpan: Corpus-Based Set Expansion via Context Feature Selection and Rank Ensemble.	288
<i>Jiaming Shen, Zeqiu Wu, Dongming Lei, Jingbo Shang, Xiang Ren, and Jiawei Han</i>	

Kernel Methods

Bayesian Nonlinear Support Vector Machines for Big Data	307
<i>Florian Wenzel, Théo Galy-Fajou, Matthäus Deutsch, and Marius Kloft</i>	
Entropic Trace Estimates for Log Determinants.	323
<i>Jack Fitzsimons, Diego Granziol, Kurt Cutajar, Michael Osborne, Maurizio Filippone, and Stephen Roberts</i>	

Fair Kernel Learning	339
<i>Adrián Pérez-Suay, Valero Laparra, Gonzalo Mateo-García, Jordi Muñoz-Marí, Luis Gómez-Chova, and Gustau Camps-Valls</i>	
GaKCo: A Fast Gapped k -mer String Kernel Using Counting	356
<i>Ritambhara Singh, Arshdeep Sekhon, Kamran Kowsari, Jack Lanchantin, Beilun Wang, and Yanjun Qi</i>	
Graph Enhanced Memory Networks for Sentiment Analysis	374
<i>Zhao Xu, Romain Vial, and Kristian Kersting</i>	
Kernel Sequential Monte Carlo	390
<i>Ingmar Schuster, Heiko Strathmann, Brooks Paige, and Dino Sejdinovic</i>	
Learning Łukasiewicz Logic Fragments by Quadratic Programming	410
<i>Francesco Giannini, Michelangelo Diligenti, Marco Gori, and Marco Maggini</i>	
Nyström Sketches	427
<i>Daniel J. Perry, Braxton Osting, and Ross T. Whitaker</i>	

Learning and Optimization

Crossprop: Learning Representations by Stochastic Meta-Gradient Descent in Neural Networks	445
<i>Vivek Veeriah, Shangdong Zhang, and Richard S. Sutton</i>	
Distributed Stochastic Optimization of Regularized Risk via Saddle-Point Problem	460
<i>Shin Matsushima, Hyokun Yun, Xinhua Zhang, and S. V. N. Vishwanathan</i>	
Speeding up Hyper-parameter Optimization by Extrapolation of Learning Curves Using Previous Builds	477
<i>Akshay Chandrashekar and Ian R. Lane</i>	
Thompson Sampling for Optimizing Stochastic Local Search	493
<i>Tong Yu, Branislav Kveton, and Ole J. Mengshoel</i>	

Matrix and Tensor Factorization

Comparative Study of Inference Methods for Bayesian Nonnegative Matrix Factorisation	513
<i>Thomas Brouwer, Jes Frellsen, and Pietro Lió</i>	

Content-Based Social Recommendation with Poisson Matrix Factorization . . .	530
<i>Eliezer de Souza da Silva, Helge Langseth, and Heri Ramampiaro</i>	
C-SALT: Mining Class-Specific ALTerations in Boolean Matrix Factorization	547
<i>Sibylle Hess and Katharina Morik</i>	
Feature Extraction for Incomplete Data via Low-rank Tucker Decomposition	564
<i>Qiquan Shi, Yiu-ming Cheung, and Qibin Zhao</i>	
Structurally Regularized Non-negative Tensor Factorization for Spatio-Temporal Pattern Discoveries.	582
<i>Koh Takeuchi, Yoshinobu Kawahara, and Tomoharu Iwata</i>	
Networks and Graphs	
Attributed Graph Clustering with Unimodal Normalized Cut	601
<i>Wei Ye, Linfei Zhou, Xin Sun, Claudia Plant, and Christian Böhm</i>	
K-Clique-Graphs for Dense Subgraph Discovery.	617
<i>Giannis Nikolentzos, Polykarpos Meladianos, Yannis Stavrakas, and Michalis Vazirgiannis</i>	
Learning and Scaling Directed Networks via Graph Embedding	634
<i>Mikhail Drobyshevskiy, Anton Korshunov, and Denis Turdakov</i>	
Local Lanczos Spectral Approximation for Community Detection	651
<i>Pan Shi, Kun He, David Bindel, and John E. Hopcroft</i>	
Regularizing Knowledge Graph Embeddings via Equivalence and Inversion Axioms	668
<i>Pasquale Minervini, Luca Costabello, Emir Muñoz, Vít Nováček, and Pierre-Yves Vandenbussche</i>	
Survival Factorization on Diffusion Networks.	684
<i>Nicola Barbieri, Giuseppe Manco, and Ettore Ritacco</i>	
The Network-Untangling Problem: From Interactions to Activity Timelines.	701
<i>Polina Rozenshtein, Nikolaj Tatti, and Aristides Gionis</i>	
TransT: Type-Based Multiple Embedding Representations for Knowledge Graph Completion.	717
<i>Shiheng Ma, Jianhui Ding, Weijia Jia, Kun Wang, and Minyi Guo</i>	

Neural Networks and Deep Learning

A Network Architecture for Multi-Multi-Instance Learning.	737
<i>Alessandro Tibo, Paolo Frasconi, and Manfred Jaeger</i>	
CON-S2V: A Generic Framework for Incorporating Extra-Sentential Context into Sen2Vec	753
<i>Tanay Kumar Saha, Shafiq Joty, and Mohammad Al Hasan</i>	
Deep Over-sampling Framework for Classifying Imbalanced Data.	770
<i>Shin Ando and Chun Yuan Huang</i>	
FCNN: Fourier Convolutional Neural Networks	786
<i>Harry Pratt, Bryan Williams, Frans Coenen, and Yalin Zheng</i>	
Joint User Modeling Across Aligned Heterogeneous Sites Using Neural Networks	799
<i>Xuezhi Cao and Yong Yu</i>	
Sequence Generation with Target Attention	816
<i>Yingce Xia, Fei Tian, Tao Qin, Nenghai Yu, and Tie-Yan Liu</i>	
Wikipedia Vandal Early Detection: From User Behavior to User Embedding	832
<i>Shuhan Yuan, Panpan Zheng, Xintao Wu, and Yang Xiang</i>	
Author Index	847

Contents – Part II

Pattern and Sequence Mining

BEATLEX: Summarizing and Forecasting Time Series with Patterns	3
<i>Bryan Hooi, Shenghua Liu, Asim Smailagic, and Christos Faloutsos</i>	
Behavioral Constraint Template-Based Sequence Classification	20
<i>Johannes De Smedt, Galina Deeva, and Jochen De Weerd</i>	
Efficient Sequence Regression by Learning Linear Models in All-Subsequence Space	37
<i>Severin Gsponer, Barry Smyth, and Georgiana Ifrim</i>	
Subjectively Interesting Connecting Trees	53
<i>Florian Adriaens, Jefrey Lijffijt, and Tijl De Bie</i>	

Privacy and Security

Malware Detection by Analysing Encrypted Network Traffic with Neural Networks	73
<i>Paul Prasse, Lukáš Machlica, Tomáš Pevný, Jiří Havelka, and Tobias Scheffer</i>	
PEM: A Practical Differentially Private System for Large-Scale Cross-Institutional Data Mining.	89
<i>Yi Li, Yitao Duan, and Wei Xu</i>	

Probabilistic Models and Methods

Bayesian Heatmaps: Probabilistic Classification with Multiple Unreliable Information Sources	109
<i>Edwin Simpson, Steven Reece, and Stephen J. Roberts</i>	
Bayesian Inference for Least Squares Temporal Difference Regularization . . .	126
<i>Nikolaos Tziortziotis and Christos Dimitrakakis</i>	
Discovery of Causal Models that Contain Latent Variables Through Bayesian Scoring of Independence Constraints	142
<i>Fattaneh Jabbari, Joseph Ramsey, Peter Spirtes, and Gregory Cooper</i>	

Labeled DBN Learning with Community Structure Knowledge.	158
<i>E. Auclair, N. Peyrard, and R. Sabbadin</i>	
Multi-view Generative Adversarial Networks	175
<i>Mickaël Chen and Ludovic Denoyer</i>	
Online Sparse Collapsed Hybrid Variational-Gibbs Algorithm for Hierarchical Dirichlet Process Topic Models	189
<i>Sophie Burkhardt and Stefan Kramer</i>	
PAC-Bayesian Analysis for a Two-Step Hierarchical Multiview Learning Approach	205
<i>Anil Goyal, Emilie Morvant, Pascal Germain, and Massih-Reza Amini</i>	
Partial Device Fingerprints	222
<i>Michael Ciere, Carlos Gañán, and Michel van Eeten</i>	
Robust Multi-view Topic Modeling by Incorporating Detecting Anomalies. . .	238
<i>Guoxi Zhang, Tomoharu Iwata, and Hisashi Kashima</i>	

Recommendation

A Regularization Method with Inference of Trust and Distrust in Recommender Systems	253
<i>Dimitrios Rafailidis and Fabio Crestani</i>	
A Unified Contextual Bandit Framework for Long- and Short-Term Recommendations	269
<i>M. Tavakol and U. Brefeld</i>	
Perceiving the Next Choice with Comprehensive Transaction Embeddings for Online Recommendation	285
<i>Shoujin Wang, Liang Hu, and Longbing Cao</i>	

Regression

Adaptive Skip-Train Structured Regression for Temporal Networks.	305
<i>Martin Pavlovski, Fang Zhou, Ivan Stojkovic, Ljupco Kocarev, and Zoran Obradovic</i>	
ALADIN: A New Approach for Drug–Target Interaction Prediction	322
<i>Krisztian Buza and Ladislav Peska</i>	
Co-Regularised Support Vector Regression.	338
<i>Katrin Ullrich, Michael Kamp, Thomas Gärtner, Martin Vogt, and Stefan Wrobel</i>	

Online Regression with Controlled Label Noise Rate.	355
<i>Edward Moroshko and Koby Crammer</i>	

Reinforcement Learning

Generalized Inverse Reinforcement Learning with Linearly Solvable MDP. . .	373
<i>Masahiro Kohjima, Tatsushi Matsubayashi, and Hiroshi Sawada</i>	
Max K -Armed Bandit: On the ExtremeHunter Algorithm and Beyond.	389
<i>Mastane Achab, Stephan Cl��men��on, Aur��lien Garivier, Anne Sabourin, and Claire Vernade</i>	
Variational Thompson Sampling for Relational Recurrent Bandits.	405
<i>Sylvain Lamprier, Thibault Gisselbrecht, and Patrick Gallinari</i>	

Subgroup Discovery

Explaining Deviating Subsets Through Explanation Networks.	425
<i>Antti Ukkonen, Vladimir Dzyuba, and Matthijs van Leeuwen</i>	
Flash Points: Discovering Exceptional Pairwise Behaviors in Vote or Rating Data	442
<i>Adnene Belfodil, Sylvie Cazalens, Philippe Lamarre, and Marc Plantevit</i>	

Time Series and Streams

A Multiscale Bezier-Representation for Time Series that Supports Elastic Matching.	461
<i>F. H��ppner and T. Sobek</i>	
Arbitrated Ensemble for Time Series Forecasting	478
<i>V��tor Cerqueira, Lu��s Torgo, F��bio Pinto, and Carlos Soares</i>	
Cost Sensitive Time-Series Classification	495
<i>Shoumik Roychoudhury, Mohamed Ghalwash, and Zoran Obradovic</i>	
Cost-Sensitive Perceptron Decision Trees for Imbalanced Drifting Data Streams	512
<i>Bartosz Krawczyk and Przemys��aw Skryjowski</i>	

Efficient Temporal Kernels Between Feature Sets for Time Series Classification	528
<i>Romain Tavenard, Simon Malinowski, Laetitia Chapel, Adeline Bailly, Heider Sanchez, and Benjamin Bustos</i>	
Forecasting and Granger Modelling with Non-linear Dynamical Dependencies	544
<i>Magda Gregorová, Alexandros Kalousis, and Stéphane Marchand-Maillet</i>	
Learning TSK Fuzzy Rules from Data Streams	559
<i>Ammar Shaker, Waleri Heldt, and Eyke Hüllermeier</i>	
Non-parametric Online AUC Maximization	575
<i>Balázs Szörényi, Snir Cohen, and Shie Mannor</i>	
On-Line Dynamic Time Warping for Streaming Time Series	591
<i>Izaskun Oregi, Aritz Pérez, Javier Del Ser, and José A. Lozano</i>	
PowerCast: Mining and Forecasting Power Grid Sequences	606
<i>Hyun Ah Song, Bryan Hooi, Marko Jereminov, Amritanshu Pandey, Larry Pileggi, and Christos Faloutsos</i>	
UAPD: Predicting Urban Anomalies from Spatial-Temporal Data	622
<i>Xian Wu, Yuxiao Dong, Chao Huang, Jian Xu, Dong Wang, and Nitesh V. Chawla</i>	
Transfer and Multi-task Learning	
LKT-FM: A Novel Rating Pattern Transfer Model for Improving Non-overlapping Cross-Domain Collaborative Filtering	641
<i>Yizhou Zang and Xiaohua Hu</i>	
Distributed Multi-task Learning for Sensor Network	657
<i>Jiyi Li, Tomohiro Arai, Yukino Baba, Hisashi Kashima, and Shotaro Miwa</i>	
Learning Task Clusters via Sparsity Grouped Multitask Learning	673
<i>Meghana Kshirsagar, Eunho Yang, and Aurélie C. Lozano</i>	
Lifelong Learning with Gaussian Processes	690
<i>Christopher Clingerman and Eric Eaton</i>	
Personalized Tag Recommendation for Images Using Deep Transfer Learning	705
<i>Hanh T. H. Nguyen, Martin Wistuba, and Lars Schmidt-Thieme</i>	

Ranking Based Multitask Learning of Scoring Functions	721
<i>Ivan Stojkovic, Mohamed Ghalwash, and Zoran Obradovic</i>	
Theoretical Analysis of Domain Adaptation with Optimal Transport	737
<i>Ievgen Redko, Amaury Habrard, and Marc Sebban</i>	
TSP: Learning Task-Specific Pivots for Unsupervised Domain Adaptation . . .	754
<i>Xia Cui, Frans Coenen, and Danushka Bollegala</i>	
Unsupervised and Semisupervised Learning	
k^2 -means for Fast and Accurate Large Scale Clustering	775
<i>Eirikur Agustsson, Radu Timofte, and Luc Van Gool</i>	
A Simple Exponential Family Framework for Zero-Shot Learning.	792
<i>Vinay Kumar Verma and Piyush Rai</i>	
DeepCluster: A General Clustering Framework Based on Deep Learning	809
<i>Kai Tian, Shuigeng Zhou, and Jihong Guan</i>	
Multi-view Spectral Clustering on Conflicting Views.	826
<i>Xiao He, Limin Li, Damian Roqueiro, and Karsten Borgwardt</i>	
Pivot-Based Distributed K-Nearest Neighbor Mining	843
<i>Caitlin Kuhlman, Yizhou Yan, Lei Cao, and Elke Rundensteiner</i>	
Author Index	861

Contents – Part III

Applied Data Science Track

A Novel Framework for Online Sales Burst Prediction	3
<i>Rui Chen and Jiajun Liu</i>	
Analyzing Granger Causality in Climate Data with Time Series Classification Methods.	15
<i>Christina Papagiannopoulou, Stijn Decubber, Diego G. Miralles, Matthias Demuzere, Niko E. C. Verhoest, and Willem Waegeman</i>	
Automatic Detection and Recognition of Individuals in Patterned Species . . .	27
<i>Gullal Singh Cheema and Saket Anand</i>	
Boosting Based Multiple Kernel Learning and Transfer Regression for Electricity Load Forecasting	39
<i>Di Wu, Boyu Wang, Doina Precup, and Benoit Boulet</i>	
CREST - Risk Prediction for Clostridium Difficile Infection Using Multimodal Data Mining	52
<i>Cansu Sen, Thomas Hartvigsen, Elke Rundensteiner, and Kajal Claypool</i>	
DC-Prophet: Predicting Catastrophic Machine Failures in <u>Data</u> Centers.	64
<i>You-Luen Lee, Da-Cheng Juan, Xuan-An Tseng, Yu-Ting Chen, and Shih-Chieh Chang</i>	
Disjoint-Support Factors and Seasonality Estimation in E-Commerce.	77
<i>Abhay Jha</i>	
Event Detection and Summarization Using Phrase Network	89
<i>Sara Melvin, Wenchao Yu, Peng Ju, Sean Young, and Wei Wang</i>	
Generalising Random Forest Parameter Optimisation to Include Stability and Cost	102
<i>C. H. Bryan Liu, Benjamin Paul Chamberlain, Duncan A. Little, and Ângelo Cardoso</i>	

Have It Both Ways—From A/B Testing to A&B Testing with Exceptional Model Mining	114
<i>Wouter Duivesteijn, Tara Farzami, Thijs Putman, Evertjan Peer, Hilde J. P. Weerts, Jasper N. Adegeest, Gerson Foks, and Mykola Pechenizkiy</i>	
Koopman Spectral Kernels for Comparing Complex Dynamics: Application to Multiagent Sport Plays	127
<i>Keisuke Fujii, Yuki Inaba, and Yoshinobu Kawahara</i>	
Modeling the Temporal Nature of Human Behavior for Demographics Prediction	140
<i>Bjarke Felbo, Pål Sundsøy, Alex ‘Sandy’ Pentland, Sune Lehmann, and Yves-Alexandre de Montjoye</i>	
MRNet-Product2Vec: A Multi-task Recurrent Neural Network for Product Embeddings.	153
<i>Arijit Biswas, Mukul Bhutani, and Subhajit Sanyal</i>	
Optimal Client Recommendation for Market Makers in Illiquid Financial Products	166
<i>Dieter Hendricks and Stephen J. Roberts</i>	
Predicting Self-reported Customer Satisfaction of Interactions with a Corporate Call Center	179
<i>Joseph Bockhorst, Shi Yu, Luisa Polania, and Glenn Fung</i>	
Probabilistic Inference of Twitter Users’ Age Based on What They Follow	191
<i>Benjamin Paul Chamberlain, Clive Humby, and Marc Peter Deisenroth</i>	
Quantifying Heterogeneous Causal Treatment Effects in World Bank Development Finance Projects	204
<i>Jianing Zhao, Daniel M. Runfola, and Peter Kemper</i>	
RSSI-Based Supervised Learning for Uncooperative Direction-Finding	216
<i>Tathagata Mukherjee, Michael Duckett, Piyush Kumar, Jared Devin Paquet, Daniel Rodriguez, Mallory Haulcomb, Kevin George, and Eduardo Pasilliao</i>	
Sequential Keystroke Behavioral Biometrics for Mobile User Identification via Multi-view Deep Learning	228
<i>Lichao Sun, Yuqi Wang, Bokai Cao, Philip S. Yu, Witawas Srisa-an, and Alex D. Leow</i>	

Session-Based Fraud Detection in Online E-Commerce Transactions Using Recurrent Neural Networks	241
<i>Shuhao Wang, Cancheng Liu, Xiang Gao, Hongtao Qu, and Wei Xu</i>	
SINAS: Suspect Investigation Using Offenders' Activity Space	253
<i>Mohammad A. Tayebi, Uwe Glässer, Patricia L. Brantingham, and Hamed Yaghoubi Shahir</i>	
Stance Classification of Tweets Using Skip Char Ngrams	266
<i>Yaakov HaCohen-kerner, Ziv Ido, and Ronen Ya'akovov</i>	
Structural Semantic Models for Automatic Analysis of Urban Areas	279
<i>Gianni Barlacchi, Alberto Rossi, Bruno Lepri, and Alessandro Moschitti</i>	
Taking It for a Test Drive: A Hybrid Spatio-Temporal Model for Wildlife Poaching Prediction Evaluated Through a Controlled Field Test	292
<i>Shahrazad Gholami, Benjamin Ford, Fei Fang, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, and Joshua Mabonga</i>	
Unsupervised Signature Extraction from Forensic Logs	305
<i>Stefan Thaler, Vlado Menkovski, and Milan Petkovic</i>	
Urban Water Flow and Water Level Prediction Based on Deep Learning	317
<i>Haytham Assem, Salem Ghariba, Gabor Makrai, Paul Johnston, Laurence Gill, and Francesco Pilla</i>	
Using Machine Learning for Labour Market Intelligence	330
<i>Roberto Boselli, Mirko Cesarini, Fabio Mercorio, and Mario Mezzanzanica</i>	

Nectar Track

Activity-Driven Influence Maximization in Social Networks	345
<i>Rohit Kumar, Muhammad Aamir Saleem, Toon Calders, Xike Xie, and Torben Bach Pedersen</i>	
An AI Planning System for Data Cleaning	349
<i>Roberto Boselli, Mirko Cesarini, Fabio Mercorio, and Mario Mezzanzanica</i>	
Comparing Hypotheses About Sequential Data: A Bayesian Approach and Its Applications.	354
<i>Florian Lemmerich, Philipp Singer, Martin Becker, Lisette Espin-Noboa, Dimitar Dimitrov, Denis Helic, Andreas Hotho, and Markus Strohmaier</i>	

Data-Driven Approaches for Smart Parking	358
<i>Fabian Bock, Sergio Di Martino, and Monika Sester</i>	
Image Representation, Annotation and Retrieval with Predictive Clustering Trees	363
<i>Ivica Dimitrovski, Dragi Kocev, Suzana Loskovska, and Sašo Džeroski</i>	
Music Generation Using Bayesian Networks	368
<i>Tetsuro Kitahara</i>	
Phenotype Inference from Text and Genomic Data	373
<i>Maria Brbić, Matija Piškorec, Vedrana Vidulin, Anita Kriško, Tomislav Šmuc, and Fran Supek</i>	
Process-Based Modeling and Design of Dynamical Systems	378
<i>Jovan Tanevski, Nikola Simidjievski, Ljupčo Todorovski, and Sašo Džeroski</i>	
QuickScorer: Efficient Traversal of Large Ensembles of Decision Trees	383
<i>Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Nicola Tonellotto, and Rossano Venturini</i>	
Recent Advances in Kernel-Based Graph Classification	388
<i>Nils M. Kriege and Christopher Morris</i>	
Demo Track	
ASK-the-Expert: Active Learning Based Knowledge Discovery Using the Expert	395
<i>Kamalika Das, Ilya Avrekh, Bryan Matthews, Manali Sharma, and Nikunj Oza</i>	
Delve: A Data Set Retrieval and Document Analysis System	400
<i>Uchenna Akujuobi and Xiangliang Zhang</i>	
Framework for Exploring and Understanding Multivariate Correlations	404
<i>Louis Kirsch, Niklas Riekenbrauck, Daniel Thevessen, Marcus Pappik, Axel Stebner, Julius Kunze, Alexander Meissner, Arvind Kumar Shekar, and Emmanuel Müller</i>	
Lit@EVE: Explainable Recommendation Based on Wikipedia Concept Vectors	409
<i>M. Atif Qureshi and Derek Greene</i>	

Monitoring Physical Activity and Mental Stress Using Wrist-Worn Device and a Smartphone.	414
<i>Božidara Cvetković, Martin Gjoreski, Jure Šorn, Pavel Maslov, and Mitja Luštrek</i>	
Tetrahedron: Barycentric Measure Visualizer	419
<i>Dariusz Brzezinski, Jerzy Stefanowski, Robert Susmaga, and Izabela Szczęch</i>	
TF Boosted Trees: A Scalable TensorFlow Based Framework for Gradient Boosting	423
<i>Natalia Ponomareva, Soroush Radpour, Gilbert Hendry, Salem Haykal, Thomas Colthurst, Petr Mitrichev, and Alexander Grushetsky</i>	
TrajViz: A Tool for Visualizing Patterns and Anomalies in Trajectory.	428
<i>Yifeng Gao, Qingzhe Li, Xiaosheng Li, Jessica Lin, and Huzefa Rangwala</i>	
TrAnET: Tracking and Analyzing the Evolution of Topics in Information Networks	432
<i>Livio Bioglio, Ruggero G. Pensa, and Valentina Rho</i>	
WHODID: Web-Based Interface for Human-Assisted Factory Operations in Fault Detection, Identification and Diagnosis	437
<i>Pierre Blanchart and Cédric Gouy-Pailler</i>	
Author Index	443

<http://www.springer.com/978-3-319-71248-2>

Machine Learning and Knowledge Discovery in
Databases

European Conference, ECML PKDD 2017, Skopje,
Macedonia, September 18–22, 2017, Proceedings, Part
I

Ceci, M.; Hollmén, J.; Todorovski, L.; Vens, C.; Džeroski,
S. (Eds.)

2017, LXIII, 852 p. 245 illus., Softcover

ISBN: 978-3-319-71248-2