

2012 IEEE 12th International Conference on Data Mining Workshops

(ICDMW 2012)

**Brussels, Belgium
10 December 2012**

Pages 1-485



**IEEE Catalog Number: CFP1256B-PRT
ISBN: 978-1-4799-1707-5**

2012 IEEE 12th International Conference on Data Mining Workshops

ICDMW 2012

Table of Contents

Message from the Conference General Chairs.....	xvi
Message from the Workshop Co-Chairs.....	xviii
Preface to the Third Workshop on Biological Data Mining and Its Applications in Healthcare—BioDM 2012.....	xx
Preface to the First Workshop on Co-clustering and Applications—CoClus’12.....	xxiii
Preface to the Workshop on Cost Sensitive Data Mining—COSTS 2012.....	xxv
Preface to the Second International Workshop on Data Mining in Networks—DaMNet 2012.....	xxvii
Preface to the Third Workshop on Data Mining for Service—DMS 2012.....	xxx
Preface to the International Workshop on Discrimination and Privacy-Aware Data Mining—DPADM 2012.....	xxxii
Preface to the Third International Workshop on Knowledge Discovery Using Cloud and Distributed Computing Platforms—KDCloud 2012.....	xxxv
Preface to the Seventh Workshop on Optimization Based Techniques for Emerging Data Mining—OEDM 2012.....	xxxviii
Preface to the First Workshop on Privacy in Social Data—PinSoDa 2012.....	xli
Preface to the Workshop on Practical Theories for Exploratory Data Mining—PTDM 2012.....	xliv
Preface to the Fourth Workshop on Reliability Issues in Knowledge Discovery—RIKD 2012.....	xlvii
Preface to the Workshop on Sentiment Elicitation from Natural Text for Information Retrieval and Extraction—SENTIRE 2012.....	l
Preface to the Workshop on Social Media Analysis and Mining—SMAM 2012.....	lii
Preface to the Seventh Workshop on Spatial and Spatio-Temporal Data Mining—SSTDM 2012.....	liv
Preface to the First Workshop on Web Entity Modeling and Applications—WEMA 2012.....	lviii
Preface to the Workshop on the ICDM 2012 Contest on Consumer PRODUCT Mention Recognition and Linking—CPROD1.....	lx

Preface to the Demo Session	lxi
Preface to the Industry and Government Track	lxii
Preface to the PhD Forum on Data Mining	lxiii

BioDM 2012: Third Workshop on Biological Data Mining and Its Applications in Healthcare

Coupled Matrix Factorization with Sparse Factors to Identify Potential Biomarkers in Metabolomics	1
<i>Evrin Acar, Gözde Gürdeniz, Morten A. Rasmussen, Daniela Rago, Lars O. Dragsted, and Rasmus Bro</i>	
Diagnosis of Coronary Artery Disease Using Cost-Sensitive Algorithms	9
<i>Roohallah Alizadehsani, Mohammad Javad Hosseini, Zahra Alizadeh Sani, Asma Ghandeharioun, and Reihane Boghrati</i>	
Predicting Hospital Length of Stay (PHLOS): A Multi-tiered Data Mining Approach	17
<i>Ali Azari, Vandana P. Janeja, and Alex Mohseni</i>	
Figure Retrieval in Biomedical Literature	25
<i>K. Sai Deepak, Harikrishna G.N. Rai, and P. Radhakrishna</i>	
Cerebral Palsy EEG Signals Classification: Facial Expressions and Thoughts for Driving an Intelligent Wheelchair	33
<i>Brigida Monica Faria, Luis Paulo Reis, and Nuno Lau</i>	
Improved Feature Selection by Incorporating Gene Similarity Into the LASSO	41
<i>C.E. Gillies, X. Gao, N.V. Patel, M.R. Siadat, and G.D. Wilson</i>	
Evaluation of Feature Ranking Ensembles for High-Dimensional Biomedical Data: A Case Study	49
<i>Ludmila I. Kuncheva, Christopher J. Smith, Yasir Syed, Christopher O. Phillips, and Keir E. Lewis</i>	
Adapting Surgical Models to Individual Hospitals Using Transfer Learning	57
<i>Gyemin Lee, Ilan Rubinfeld, and Zeeshan Syed</i>	
Clustering Tandem Repeats via Trinucleotides	64
<i>Yupu Liang, Dina Sokol, and Sarah Zelikovitz</i>	
Using Perspective Wall to Visualize Medical Data in the Intensive Care Unit	72
<i>Hela Ltifi, Mounir Ben Ayed, Ghada Trabelsi, and Adel M. Alimi</i>	
Evidence Theory-Based Approach for Epileptic Seizure Detection Using EEG Signals	79
<i>Abduljalil Mohamed, Khaled Bashir Shaban, and Amr Mohamed</i>	
Discovering Aberrant Patterns of Human Connectome in Alzheimer’s Disease via Subgraph Mining	86
<i>Junming Shao, Qinli Yang, Afra Wohlschlaeger, and Christian Sorg</i>	
Discovering Aging-Genes by Topological Features in <i>Drosophila melanogaster</i> Protein-Protein Interaction Network	94
<i>Xin Song, Yuan-Chun Zhou, Kai Feng, Yan-Hui Li, and Jian-Hui Li</i>	
Mining Medical Data to Develop Clinical Decision Making Tools in Hemodialysis	99
<i>J. Ion Titapiccolo, M. Ferrario, S. Cerutti, M.G. Signorini, C. Barbieri, F. Mari, and E. Gatti</i>	

CoClus 2012: Workshop on Co-clustering and Applications

TeamFinder: A Co-clustering Based Framework for Finding an Effective Team of Experts in Social Networks	107
<i>Farnoush Farhadi, Elham Hoseini, Sattar Hashemi, and Ali Hamzeh</i>	
A Triclustering Approach for Time Evolving Graphs	115
<i>Romain Guigourès, Marc Boullé, and Fabrice Rossi</i>	
Concept-Based Biclustering for Internet Advertisement	123
<i>Dmitry I. Ignatov, Sergei O. Kuznetsov, and Jonas Poelmans</i>	
Biclustering of High-throughput Gene Expression Data with Bicluster Miner	131
<i>András Király, János Abonyi, Asta Laiho, and Attila Gyenesei</i>	
Mining Local Staircase Patterns in Noisy Data	139
<i>Thanh Le Van, Ana Carolina Fierroy, Tias Guns, Matthijs van Leeuwen, Siegfried Nijssen, Luc De Raedt, and Kathleen Marchal</i>	
An Approximation of the Integrated Classification Likelihood for the Latent Block Model	147
<i>Aurore Lomet, Gérard Govaert, and Yves Grandvalet</i>	

COSTS 2012: Workshop on Cost Sensitive Data Mining

Learning Cost-Sensitive Rules for Non-forced Classification	154
<i>Arjun Bakshi and Raj Bhatnagar</i>	
A Multi-armed Bandit Approach to Cost-Sensitive Decision Tree Learning	162
<i>Susan Lomax, Sunil Vadera, and Mohammed Saraee</i>	
When Additional Views are Not Free: Active View Completion for Multi-view Semi-Supervised Learning	169
<i>Brian Quanz and Jun Huan</i>	
Towards Utility Maximization in Regression	179
<i>Rita P. Ribeiro</i>	
A Weighted SOM for Classifying Data with Instance-Varying Importance	187
<i>Peter Sarlin</i>	
Learning in the Class Imbalance Problem When Costs are Unknown for Errors and Rejects	194
<i>Xiaowan Zhang and Bao-Gang Hu</i>	

DaMNet 2012: Second Workshop on Data Mining in Networks

Sensor Network Localization for Moving Sensors	202
<i>Arvind Agarwal, Hal Daumé, Jeff M. Phillips, and Suresh Venkatasubramanian</i>	
Motif Mining in Weighted Networks	210
<i>Sarvenaz Choobdar, Pedro Ribeiro, and Fernando Silva</i>	
Comparison of the Efficiency of MapReduce and Bulk Synchronous Parallel Approaches to Large Network Processing	218
<i>Tomasz Kajdanowicz, Wojciech Indyk, Przemyslaw Kazienko, and Jakub Kukul</i>	

Effect of Data Repair on Mining Network Streams	226
<i>Ji Meng Loh and Tamraparni Dasu</i>	
EigenSP: A More Accurate Shortest Path Distance Estimation on Large-Scale Networks	234
<i>Koji Maruhashi, Junichi Shigezumi, Nobuhiro Yugami, and Christos Faloutsos</i>	
Uncovering the Spatio-temporal Structure of Social Networks Using Cell Phone Records	242
<i>Luis G. Moyano, Oscar R. Moll Thomaes, and Enrique Frias-Martinez</i>	
Canonical Correlation Analysis for Detecting Changes in Network Structure	250
<i>Aidan O’Sullivan, Niall M. Adams, and Iead Rezek</i>	
Maximizing Information Spread through Influence Structures in Social Networks	258
<i>Saurav Pandit, Yang Yang, and Nitesh V. Chawla</i>	
Sampling Online Social Networks Using Coupling from the Past	266
<i>Kenton White, Guichong Li, and Nathalie Japkowicz</i>	
 DMS 2012: Third Workshop on Data Mining for Service	
Data Mining in the Age of Curation	273
<i>Akinori Abe</i>	
Rigid or Flexible? A New Navigation Approach for Better Consumer Service Based on Knowledge Enhancement	280
<i>Hongyun Bao, Qiudan Li, Daniel Zeng, and Heng Gao</i>	
Temporary Staffing Services: A Data Mining Perspective	287
<i>Jeroen D’Haen and Dirk Van Den Poel</i>	
Data-Oriented Construction and Maintenance of Clinical Pathway Using Similarity-Based Data Mining Methods	293
<i>Haruko Iwata, Shusaku Tsumoto, and Shoji Hirano</i>	
Viewers’ Side Analysis of Social Interests	301
<i>Takeshi Mitamura and Kenichi Yoshida</i>	
Streamlining Service Levels for IT Infrastructure Support	309
<i>Girish Keshav Palshikar, Mohammed Mudassar, Harrick M. Vin, and Maitreya Natu</i>	
An Analytics Approach for Proactively Combating Voluntary Attrition of Employees	317
<i>Moninder Singh, Kush R. Varshney, Jun Wang, Aleksandra Mojsilovic, Alisia R. Gill, Patricia I. Faur, and Raphael Ezry</i>	
Exploration of Dependencies among Sections in a Supermarket Using a Tree-Structured Undirected Graphical Model	324
<i>Keiji Takai</i>	
Evaluating Fraud Detection Algorithms Using an Auction Data Generator	332
<i>Sidney Tsang, Gillian Dobbie, and Yun Sing Koh</i>	
Extracting Information from Sequences of Financial Ratios with Markov for Discrimination: An Application to Bankruptcy Prediction	340
<i>Andrey Volkov and Dirk Van den Poel</i>	

DPADM 2012: Workshop on Discrimination and Privacy-Aware Data Mining

Exploring Discrimination: A User-centric Evaluation of Discrimination-Aware Data Mining	344
<i>Bettina Berendt and Sören Preibusch</i>	
A Study on the Impact of Data Anonymization on Anti-discrimination	352
<i>Sara Hajian and Josep Domingo-Ferrer</i>	
Injecting Discrimination and Privacy Awareness Into Pattern Discovery	360
<i>Sara Hajian, Anna Monreale, Dino Pedreschi, Josep Domingo-Ferrer, and Fosca Giannotti</i>	
Classifying Socially Sensitive Data Without Discrimination: An Analysis of a Crime Suspect Dataset	370
<i>Faisal Kamiran, Asim Karim, Sicco Verwer, and Heike Goudriaan</i>	
Considerations on Fairness-Aware Data Mining	378
<i>Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma</i>	
Discriminatory Decision Policy Aware Classification	386
<i>Koray Mancuhan and Chris Clifton</i>	
Discovering Gender Discrimination in Project Funding	394
<i>Andrea Romei, Salvatore Ruggieri, and Franco Turini</i>	

KDCloud 2012: Third Workshop on Knowledge Discovery Using Cloud and Distributed Computing Platforms

Parallel Concept Drift Detection with Online Map-Reduce	402
<i>Artur Andrzejak and João Bártolo Gomes</i>	
Using Storm to Perform Dynamic Egocentric Network Motif Analysis	408
<i>Martin Harrigan, Pádraig Cunningham, and Lorcan Coyle</i>	
Genetic Algorithm Based Feature Selection Algorithm for Effective Intrusion Detection in Cloud Networks	416
<i>Anand Kannan, Gerald Q. Maguire Jr., Ayush Sharma, and Peter Schoo</i>	
Scalable Clustering Using PACT Programming Model	424
<i>Sharanjit Kaur, Dhriti Khanna, Tripti Gupta, and Vasudha Bhatnagar</i>	
Convex-Concave Hull for Classification with Support Vector Machine	431
<i>Asdrúbal López-Chau, Xiaou Li, and Wen Yu</i>	
Large Scale KNN-Graph Approximation	439
<i>Mohamed Riadh Trad, Alexis Joly, and Nozha Boujemaa</i>	
MapReduce-based Closed Frequent Itemset Mining with Efficient Redundancy Filtering	449
<i>Su-Qi Wang, Yu-Bin Yang, Yang Gao, Guang-Peng Chen, and Yao Zhang</i>	

OEDM 2012: Seventh Workshop on Optimization Based Techniques for Emerging Data Mining Problems

OCCAMS—An Optimal Combinatorial Covering Algorithm for Multi-document Summarization	454
<i>Sashka T. Davis, John M. Conroy, and Judith D. Schlesinger</i>	
Rare Events Forecasting Using a Residual-Feedback GMDH Neural Network	464
<i>Simon Fong, Zhou Nannan, Raymond K. Wong, and Xin-She Yang</i>	
Anomalous Neighborhood Selection	474
<i>Satoshi Hara and Takashi Washio</i>	
Employing the Principal Hessian Direction for Building Hinging Hyperplane Models	481
<i>Anca Maria Ivanescu, Thivaharan Albin, Dirk Abel, and Thomas Seidl</i>	
Overlapping Clustering with Sparseness Constraints	486
<i>Haibing Lu, Yuan Hong, W. Nick Street, Fei Wang, and Hanghang Tong</i>	
Nonlinear L-1 Support Vector Machines for Learning Using Privileged Information	495
<i>Lingfeng Niu and Jianmin Wu</i>	
Regular Multiple Criteria Linear Programming for Semi-supervised Classification	500
<i>Zhiquan Qi, Yingjie Tian, and Yong Shi</i>	
Nonlinear Unsupervised Feature Learning: How Local Similarities Lead to Global Coding	506
<i>Amirreza Shaban, Hamid R. Rabiee, Marzieh S. Tahaei, and Erfan Salavati</i>	
Learning from Multiple Annotators: When Data is Hard and Annotators are Unreliable	514
<i>Chirine Wolley and Mohamed Quafafou</i>	
Robust Kernel Nonnegative Matrix Factorization	522
<i>Zhichen Xia, Chris Ding, and Edmond Chow</i>	
The Performance of Alternative Exchange Rate Regimes and their Countries' Condition: Matching Analysis and Selection Model Building	530
<i>Haizhen Yang, Yunpeng Song, and Kun Guo</i>	
The Transfer Learning Based on Relationships between Attributes	535
<i>Jinwei Zhao, Boqin Feng, Guirong Yan, and Longlei Dong</i>	

PinSoDa 2012: Workshop on Privacy in Social Data

Exploiting Dynamic Privacy in Socially Regularized Recommenders	539
<i>Ramona Bunea, Shahab Mokarizadeh, Nima Dokoohaki, and Mihhail Matskin</i>	
A Practical System for Privacy-Preserving Collaborative Filtering	547
<i>Richard Chow, Manas A. Pathak, and Cong Wang</i>	
Interactive Grouping of Friends in OSN: Towards Online Context Management	555
<i>Bo Gao, Bettina Berendt, Dave Clarke, Ralf De Wolf, Thomas Peetz, Jo Pierson, and Rula Sayaf</i>	
Measuring Local Topological Anonymity in Social Networks	563
<i>Gábor György Gulyás and Sándor Imre</i>	

Beware of What You Share: Inferring Home Location in Social Networks	571
<i>Tatiana Pontes, Gabriel Magno, Marisa Vasconcelos, Aditi Gupta, Jussara Almeida, Ponnurangam Kumaraguru, and Virgilio Almeida</i>	

PTDM 2012: Workshop on Practical Theories for Exploratory Data Mining

Thorough Analysis of Log Data with Dependency Rules: Practical Solutions and Theoretical Challenges	579
<i>Wilhelmiina Hämmäläinen</i>	
Generalized Expansion Dimension	587
<i>Michael E. Houle, Hisashi Kashima, and Michael Nett</i>	
Enhancing the Analysis of Large Multimedia Applications Execution Traces with FrameMiner	595
<i>C. Kamdem Kengne, L.C. Fopa, N. Ibrahim, A. Termier, M.C. Rousset, and T. Washio</i>	
Logical Itemset Mining	603
<i>Shailesh Kumar, Chandrashekar V., and C.V. Jawahar</i>	
Generating Diverse Realistic Data Sets for Episode Mining	611
<i>Albrecht Zimmermann</i>	

RIKD 2012: Fourth Workshop on Reliability Issues in Knowledge Discovery

The PerfSim Algorithm for Concept Drift Detection in Imbalanced Data	619
<i>Daniel K. Antwi, Herna L. Viktor, and Nathalie Japkowicz</i>	
Reliable Knowledge Discovery with a Minimal Causal Model Inducer	629
<i>Honghua Dai, Sarah Keble-Johnston, and Min Gan</i>	
A Weighted Support Vector Data Description Based on Rough Neighborhood Approximation	635
<i>Yanxing Hu, James N.K. Liu, Yuan Wang, and Lucas Lai</i>	
Outlier Detection in Logistic Regression: A Quest for Reliable Knowledge from Predictive Modeling and Classification	643
<i>Abdul Nurunnabi and Geoff West</i>	
Model Selection with Combining Valid and Optimal Prediction Intervals	653
<i>Darko Pevec and Igor Kononenko</i>	
Bootstrap Confidence Intervals in DirectLiNGAM	659
<i>Kittitat Thamvitayakul, Shohei Shimizu, Tsuyoshi Ueno, Takashi Washio, and Tatsuya Tashiro</i>	

SENTIRE 2012: Second Workshop on Sentiment Elicitation from Natural Text for Information Retrieval and Extraction

Adaptation and Use of Subjectivity Lexicons for Domain Dependent Sentiment Classification	669
<i>Rahim Dehkharghani, Berrin Yanikoglu, Dilek Tapucu, and Yucel Saygin</i>	
Learning Domain-Specific Polarity Lexicons	674
<i>Gulsen Demiroz, Berrin Yanikoglu, Dilek Tapucu, and Yucel Saygin</i>	
Fine-grained Product Features Extraction and Categorization in Reviews Opinion Mining	680
<i>Sheng Huang, Xinlan Liu, Xueping Peng, and Zhendong Niu</i>	

Representing and Resolving Negation for Sentiment Analysis	687
<i>Emanuele Lapponi, Jonathon Read, and Lilja Øvrelid</i>	
Full Spectrum Opinion Mining: Integrating Domain, Syntactic and Lexical Knowledge	693
<i>Daniel J. Olsher</i>	
A Regularized Recommendation Algorithm with Probabilistic Sentiment-Ratings	701
<i>Filipa Peleja, Pedro Dias, and João Magalhães</i>	
Enriching SenticNet Polarity Scores through Semi-Supervised Fuzzy Clustering	709
<i>Soujanya Poria, Alexander Gelbukh, Erik Cambria, Dipankar Das, and Sivaji Bandyopadhyay</i>	
Domain Adaptation Using Domain Similarity- and Domain Complexity-Based Instance Selection for Cross-Domain Sentiment Analysis	717
<i>Robert Remus</i>	
How Much Supervision? Corpus-Based Lexeme Sentiment Estimation	724
<i>Aleksander Wawer and Dominika Rogozińska</i>	
Sentiment Polarity Classification Using Statistical Data Compression Models	731
<i>Dominique Ziegelmayer and Rainer Schrader</i>	
 SMAM 2012: Workshop on Social Media Analysis and Mining	
Online Social Behavior in Twitter: A Literature Review	739
<i>Olav Aarts, Peter-Paul van Maanen, Tanneke Ouboter, and Jan Maarten Schraagen</i>	
Twitter Keyword Volume, Current Spending, and Weekday Spending Norms Predict Consumer Spending	747
<i>Justin Stewart, Homer Strong, Jeffery Parker, and Mark A. Bedau</i>	
Geosocial Graph-Based Community Detection	754
<i>Yves van Gennip, Huiyi Hu, Blake Hunter, and Mason A. Porter</i>	
 SSTDM 2012: Seventh Workshop on Spatial and Spatio-Temporal Data Mining	
Performance-Optimizing Classification of Time-Series Based on Nearest Neighbor Density Approximation	759
<i>Shin Ando</i>	
Mining Spatio-temporal Patterns in the Presence of Concept Hierarchies	765
<i>Le Van Quoc Anh and Michael Gertz</i>	
Approximate Search on Massive Spatiotemporal Datasets	773
<i>Ivan Brugere, Karsten Steinhaeuser, Shyam Boriah, and Vipin Kumar</i>	
Hierarchical Classifier-Regression Ensemble for Multi-phase Non-linear Dynamic System Response Prediction: Application to Climate Analysis	781
<i>Doel L. Gonzalez, Zhengzhang Chen, Isaac K. Tetteh, Tatdow Pansombut, Fredrick Semazzi, Vipin Kumar, Anatoli Melechko, and Nagiza F. Samatova</i>	
Toward Geographic Information Harvesting: Extraction of Spatial Relational Facts from Web Documents	789
<i>Corrado Loglisci, Dino Ienco, Mathieu Roche, Maguelonne Teisseire, and Donato Malerba</i>	

Geo-referenced Time-Series Summarization Using <i>k</i> -Full Trees: A Summary of Results	797
<i>Dev Oliver, Shashi Shekhar, James M. Kang, Renee Laubscher, Veronica Carlan, and Michael R. Evans</i>	
Spatio-temporal Co-occurrence Pattern Mining in Data Sets with Evolving Regions	805
<i>Karthik Ganesan Pillai, Rafal A. Angryk, Juan M. Banda, Michael A. Schuh, and Tim Wylie</i>	
Comparative Study of Association Rule Mining and MiSTIC in Extracting Spatio-temporal Disease Occurrences Patterns	813
<i>Vipul Raheja and K.S. Rajan</i>	
Spatial Interestingness Measures for Co-location Pattern Mining	821
<i>Christian Sengstock, Michael Gertz, and Tran Van Canh</i>	
Location Extraction from Social Networks with Commodity Software and Online Data	827
<i>George Valkanas and Dimitrios Gunopoulos</i>	
 WEMA 2012: Workshop on Web Entity Modeling and Applications	
Endless and Scalable Knowledge Table Extraction from Semi-structured Websites	835
<i>Yingqin Gu, Lei Ji, Ziheng Jiang, and Jun He</i>	
Constructing and Exploring Composite Items Using Max-valid Bundles	843
<i>Gowtham Srinivas P., Sreyantha Chary M., Satheesh Kumar D., and Santhi Thilagam P.</i>	
Learning to Extract Entity Uniqueness from Web for Helping User Decision Making	850
<i>Wenhan Wang, Ning Liu, and Yiran Xie</i>	
 CPROD1 2012: Workshop on the ICDM 2012 Contest on Consumer PRODUCT Mention Recognition and Linking	
Rule Based Product Name Recognition and Disambiguation	858
<i>Balázs Gódcény</i>	
An Overview of the CPROD1 Contest on Consumer Product Recognition within User Generated Postings and Normalization against a Large Product Catalog	861
<i>Gabor Melli and Christian Romming</i>	
An Ensemble-Based Named Entity Recognition Solution for Detecting Consumer Products	865
<i>Lukasz Romaszko</i>	
An NER-based Product Identification and Lucene-based Product Linking Approach to CPROD1 Challenge: Description of Submission System to CPROD1 Challenge	869
<i>Zhiqiang Toh, Wenting Wang, Man Lan, and Xiaoli Li</i>	
Identification and Disambiguation of Product Mentions with Information Retrieval and Problem Specific Methods	872
<i>Olexandr Topchylo</i>	
Accurate Product Name Recognition from User Generated Content	874
<i>Sen Wu, Zhanpeng Fang, and Jie Tang</i>	

Demo Papers

OpinioNetIt: A Structured and Faceted Knowledge-base of Opinions	878
<i>Rawia Awadallah, Maya Ramanath, and Gerhard Weikum</i>	
Topick: Accurate Topic Distillation for User Streams	882
<i>Anton Dimitrov, Alexandra Olteanu, Luke Mcdowell, and Karl Aberer</i>	
A Subspace Clustering Extension for the KNIME Data Mining Framework	886
<i>Stephan Günnemann, Hardy Kremer, Richard Musiol, Roman Haag, and Thomas Seidl</i>	
Mining Temporal Profiles of Mobile Applications for Usage Prediction	890
<i>Zhung-Xun Liao, Po-Ruey Lei, Tsu-Jou Shen, Shou-Chung Li, and Wen-Chih Peng</i>	
Cubix: A Visual Analytics Tool for Conceptual and Semantic Data	894
<i>Cassio Melo, Alexander Mikheev, Bénédicte Le-Grand, and Marie-Aude Aufaure</i>	
Supporting the Discovery of Relevant Topological Patterns in Attributed Graphs	898
<i>Julien Salotti, Marc Plantevit, Céline Robardet, and Jean-François Boulicaut</i>	
AD-MAD: Automated Development and Optimization of Online Advertising Campaigns	902
<i>Stamatina Thomaidou, Konstantinos Leymonis, Kyriakos Liakopoulos, and Michalis Vazirgiannis</i>	
Rapid Damage eXplorer (RDX): A Probabilistic Framework for Learning Changes from Bitemporal Images	906
<i>Ranga Raju Vatsavai</i>	
Mining Trajectories for Spatio-temporal Analytics	910
<i>Songhua Xing, Xuan Liu, Qing He, and Arun Hampapur</i>	

Industry and Government Track

Distributed Big Advertiser Data Mining	914
<i>Ashish Bindra, Srinivasulu Pokuri, Krishna Uppala, and Ankur Teredesai</i>	
Data Mining for Official Statistics: Challenges and Opportunities	915
<i>Bart Buelens, Piet Daas, and Jan van den Brakel</i>	
Tectonic Shifts in Television Advertising	916
<i>Brendan Kitts, Dyng Au, and Brian Burdick</i>	
A Data Mining Framework for Monitoring Nuclear Facilities	917
<i>Ranga Raju Vatsavai</i>	

PhD Forum

Effective Text Classification by a Supervised Feature Selection Approach	918
<i>Tanmay Basu and C.A. Murthy</i>	
Active Learning Based Rule Extraction for Regression	926
<i>Enric Junqué de Fortuny and David Martens</i>	
Multislice Modularity Optimization in Community Detection and Image Segmentation	934
<i>Huiyi Hu, Yves van Gennip, Blake Hunter, Andrea L. Bertozzi, and Mason A. Porter</i>	
A Sorted Neighborhood Approach to Multidimensional Privacy Preserving Blocking	937
<i>Alexandros Karakasidis and Vassilios S. Verykios</i>	

Modeling of Collective Synchronous Behavior on Social Media	945
<i>Victor C. Liang and Vincent T.Y. Ng</i>	
Selecting Accurate and Comprehensible Regression Algorithms through Meta Learning	953
<i>Gert Loterman and Christophe Mues</i>	
Towards a Particle Swarm Optimization-Based Regression Rule Miner	961
<i>Bart Minnaert and David Martens</i>	
Prediction of HLA Genes from SNP Data and HLA Haplotype Frequencies	964
<i>Vanja Paunić, Michael Steinbach, Vipin Kumar, and Martin Maiers</i>	

Author Index