

2011 IEEE 11th International Conference on Data Mining (ICDM 2011)

**Vancouver, British Columbia, Canada
11 – 14 December 2011**

Pages 1-625



**IEEE Catalog Number: CFP11278-PRT
ISBN: 978-1-4577-2075-8**

2011 11th IEEE International Conference on Data Mining

ICDM 2011

Table of Contents

Message from the Conference General Chairs.....	xv
Message from the Program Co-Chairs.....	xvii
Organizing Committee.....	xix
Steering Committee.....	xxi
Program Committee.....	xxii
List of Reviewers.....	xxx
Sponsors.....	xxxiv

Regular Papers

Algorithms for Mining the Evolution of Conserved Relational States in Dynamic Networks	1
<i>Rezwan Ahmed and George Karypis</i>	
Infrastructure Pattern Discovery in Configuration Management Databases via Large Sparse Graph Mining	11
<i>Pranay Anchuri, Mohammed J. Zaki, Omer Barkol, Ruth Bergman, Yifat Felder, Shahar Golan, and Arik Sityon</i>	
Role-Behavior Analysis from Trajectory Data by Cross-Domain Learning	21
<i>Shin Ando and Einoshin Suzuki</i>	
Semi-supervised Feature Importance Evaluation with Ensemble Learning	31
<i>Hasna Barkia, Haytham Elghazel, and Alex Aussem</i>	
COMET: A Recipe for Learning and Using Large Ensembles on Massive Data	41
<i>Justin D. Basilico, M. Arthur Munson, Tamara G. Kolda, Kevin R. Dixon, and W. Philip Kegelmeyer</i>	
Overlapping Correlation Clustering	51
<i>Francesco Bonchi, Aristides Gionis, and Antti Ukkonen</i>	
Learning with Minimum Supervision: A General Framework for Transductive Transfer Learning	61
<i>Mohammad Taha Bahadori, Yan Liu, and Dan Zhang</i>	

Confidence in Predictions from Random Tree Ensembles	71
<i>Siddhartha Bhattacharyya</i>	
Mining Heavy Subgraphs in Time-Evolving Networks	81
<i>Petko Bogdanov, Misael Mongiovì, and Ambuj K. Singh</i>	
Multi-Class L2,1-Norm Support Vector Machine	91
<i>Xiao Cai, Feiping Nie, Heng Huang, and Chris Ding</i>	
SolarMap: Multifaceted Visual Analytics for Topic Exploration	101
<i>Nan Cao, David Gotz, Jimeng Sun, Yu-Ru Lin, and Huamin Qu</i>	
Efficiently Mining Unordered Trees	111
<i>Mostafa Haghiri Chehreghani</i>	
CEMiner—An Efficient Algorithm for Mining Closed Patterns from Time Interval-Based Data	121
<i>Yi-Cheng Chen, Wen-Chih Peng, and Suh-Yin Lee</i>	
LinkBoost: A Novel Cost-Sensitive Boosting Framework for Community-Level Network Link Prediction	131
<i>Prakash Mandayam Comar, Pang-Ning Tan, and Anil K. Jain</i>	
Learning Dirichlet Processes from Partially Observed Groups	141
<i>Avinava Dubey, Indrajit Bhattacharya, Mrinal Das, Tanveer Faruque, and Chiranjib Bhattacharyya</i>	
Exploiting False Discoveries—Statistical Validation of Patterns and Quality Measures in Subgroup Discovery	151
<i>Wouter Duivesteijn and Arno Knobbe</i>	
An Efficient Greedy Method for Unsupervised Feature Selection	161
<i>Ahmed K. Farahat, Ali Ghodsi, and Mohamed S. Kamel</i>	
Structured Feature Selection and Task Relationship Inference for Multi-task Learning	171
<i>Hongliang Fei and Jun Huan</i>	
A Taxi Driving Fraud Detection System	181
<i>Yong Ge, Hui Xiong, Chuanren Liu, and Zhi-Hua Zhou</i>	
Isograph: Neighbourhood Graph Construction Based on Geodesic Distance for Semi-supervised Learning	191
<i>Marjan Ghazvininejad, Mostafa Mahdiah, Hamid R. Rabiee, Parisa Khanipour Roshan, and Mohammad Hossein Rohban</i>	
D-cores: Measuring Collaboration of Directed Graphs Based on Degeneracy	201
<i>Christos Giatsidis, Dimitrios M. Thilikos, and Michalis Vazirgiannis</i>	
SIMPACT: An Efficient Algorithm for Influence Maximization under the Linear Threshold Model	211
<i>Amit Goyal, Wei Lu, and Laks V.S. Lakshmanan</i>	

Cross Domain Random Walk for Query Intent Pattern Mining from Search Engine	
Log	221
<i>Siyu Gu, Jun Yan, Lei Ji, Shuicheng Yan, Junshi Huang, Ning Liu, Ying Chen, and Zheng Chen</i>	
Flexible Fault Tolerant Subspace Clustering for Data with Missing Values	231
<i>Stephan Günnemann, Emmanuel Müller, Sebastian Raubach, and Thomas Seidl</i>	
Heuristic Updatable Weighted Random Subspaces for Non-stationary Environments	241
<i>T. Ryan Hoens, Nitesh V. Chawla, and Robi Polikar</i>	
Learning Tags from Unsegmented Videos of Multiple Human Actions	251
<i>Timothy M. Hospedales, Shaogang Gong, and Tao Xiang</i>	
Generating Breakpoint-based Timeline Overview for News Topic Retrospection	260
<i>Po Hu, Minlie Huang, Peng Xu, Weichang Li, Adam K. Usadi, and Xiaoyan Zhu</i>	
A Robust Clustering Algorithm Based on Aggregated Heat Kernel Mapping	270
<i>Hao Huang, Shinjae Yoo, Hong Qin, and Dantong Yu</i>	
Patent Maintenance Recommendation with Patent Information Network Model	280
<i>Xin Jin, Scott Spangler, Ying Chen, Keke Cai, Rui Ma, Li Zhang, Xian Wu, and Jiawei Han</i>	
S-preconditioner for Multi-fold Data Reduction with Guaranteed User-Controlled Accuracy	290
<i>Ye Jin, Sriram Lakshminarasimhan, Neil Shah, Zhenhuan Gong, C.S. Chang, Jackie Chen, Stephane Ethier, Hemanth Kolla, Seung-Hoe Ku, Scott Klasky, Robert Latham, Robert Ross, Karen Schuchardt, and Nagiza F. Samatova</i>	
Beyond ‘Caveman Communities’: Hubs and Spokes for Graph Compression and Mining	300
<i>U. Kang and Christos Faloutsos</i>	
Improving Product Classification Using Images	310
<i>Anitha Kannan, Partha Pratim Talukdar, Nikhil Rasiwasia, and Qifa Ke</i>	
Learning Markov Logic Networks via Functional Gradient Boosting	320
<i>Tushar Khot, Sriraam Natarajan, Kristian Kersting, and Jude Shavlik</i>	
Signature Pattern Covering via Local Greedy Algorithm and Pattern Shrink	330
<i>Hyungsul Kim, Sungjin Im, Tarek Abdelzaher, Jiawei Han, David Sheridan, and Shobha Vasudevan</i>	
TWITObI: A Recommendation System for Twitter Using Probabilistic Modeling	340
<i>Younghoon Kim and Kyuseok Shim</i>	
Maximum Entropy Modelling for Assessing Results on Real-Valued Data	350
<i>Kleanthis-Nikolaos Kontonasis, Jilles Vreeken, and Tijl De Bie</i>	
Local Models for Expectation-Driven Subgroup Discovery	360
<i>Florian Lemmerich and Frank Puppe</i>	
Context-Aware Multi-instance Learning Based on Hierarchical Sparse Representation	370
<i>Bing Li, Weihua Xiong, and Weiming Hu</i>	

The Joint Inference of Topic Diffusion and Evolution in Social Communities	378
<i>Cindy Xide Lin, Qiaozhu Mei, Jiawei Han, Yunliang Jiang, and Marina Danilevsky</i>	
Towards Optimal Discriminating Order for Multiclass Classification	388
<i>Dong Liu, Shuicheng Yan, Yadong Mu, Xian-Sheng Hua, Shih-Fu Chang, and Hong-Jiang Zhang</i>	
A Hypergraph-based Method for Discovering Semantically Associated Itemsets	398
<i>Haishan Liu, Paea Le Pendu, Ruoming Jin, and Dejing Dou</i>	
Personalized Travel Package Recommendation	407
<i>Qi Liu, Yong Ge, Zhongmou Li, Enhong Chen, and Hui Xiong</i>	
Tag Clustering and Refinement on Semantic Unity Graph	417
<i>Yang Liu, Fei Wu, Yin Zhang, Jian Shao, and Yueting Zhuang</i>	
Minimizing Seed Set for Viral Marketing	427
<i>Cheng Long and Raymond Chi-Wing Wong</i>	
Privacy Risk in Graph Stream Publishing for Social Network Data	437
<i>Nigel Medforth and Ke Wang</i>	
Boolean Tensor Factorizations	447
<i>Pauli Miettinen</i>	
Sparse Domain Adaptation in Projection Spaces Based on Good Similarity Functions	457
<i>Emilie Morvant, Amaury Habrard, and Stéphane Ayache</i>	
Incremental Elliptical Boundary Estimation for Anomaly Detection in Wireless Sensor Networks	467
<i>Masud Moshtaghi, Christopher Leckie, Shanika Karunasekera, James C. Bezdek, Sutharshan Rajasegarar, and Marimuthu Palaniswami</i>	
Learning Classification with Auxiliary Probabilistic Information	477
<i>Quang Nguyen, Hamed Valizadegan, and Milos Hauskrecht</i>	
Word Cloud Model for Text Categorization	487
<i>Tam T. Nguyen, Kuiyu Chang, and Siu Cheung Hui</i>	
SLIM: Sparse Linear Methods for Top-N Recommender Systems	497
<i>Xia Ning and George Karypis</i>	
Novel Recommendation Based on Personal Popularity Tendency	507
<i>Jinoh Oh, Sun Park, Hwanjo Yu, Min Song, and Seung-Taek Park</i>	
An Analysis of Performance Measures for Binary Classifiers	517
<i>Charles Parker</i>	
Detection of Cross-Channel Anomalies from Multiple Data Channels	527
<i>Duc Son Pham, Budhaditya Saha, Dinh Q. Phung, and Svetha Venkatesh</i>	
Threshold Conditions for Arbitrary Cascade Models on Arbitrary Networks	537
<i>B. Aditya Prakash, Deepayan Chakrabarti, Michalis Faloutsos, Nicholas Valler, and Christos Faloutsos</i>	

Time Series Epenthesis: Clustering Time Series Streams Requires Ignoring Some Data	547
<i>Thanawin Rakthanmanon, Eamonn J. Keogh, Stefano Lonardi, and Scott Evans</i>	
Mining Historical Documents for Near-Duplicate Figures	557
<i>Thanawin Rakthanmanon, Qiang Zhu, and Eamonn J. Keogh</i>	
Analysis of Textual Variation by Latent Tree Structures	567
<i>Teemu Roos and Yuan Zou</i>	
Healing Sample Selection Bias by Source Classifier Selection	577
<i>Chun-Wei Seah, Ivor Wai-Hung Tsang, and Yew-Soon Ong</i>	
An In-depth Study of Stochastic Kronecker Graphs	587
<i>C. Seshadhri, Ali Pinar, and Tamara G. Kolda</i>	
Learning Spectral Embedding for Semi-supervised Clustering	597
<i>Fanhua Shang, Yuanyuan Liu, and Fei Wang</i>	
Detection of Arbitrarily Oriented Synchronized Clusters in High-Dimensional Data	607
<i>Junming Shao, Claudia Plant, Qinli Yang, and Christian Böhm</i>	
A Generalized Fast Subset Sums Framework for Bayesian Event Detection	617
<i>Kan Shao, Yandong Liu, and Daniel B. Neill</i>	
Learning to Rank for Query-Focused Multi-document Summarization	626
<i>Chao Shen and Tao Li</i>	
Simple Multiple Noisy Label Utilization Strategies	635
<i>Victor S. Sheng</i>	
Partitionable Kernels for Mapping Kernels	645
<i>Kilho Shin</i>	
Mining Dominant Patterns in the Sky	655
<i>Arnaud Soulet, Chedy Raïssi, Marc Plantevit, and Bruno Crémilleux</i>	
Ranking Web-Based Partial Orders by Significance Using a Markov Reference Model	665
<i>Michel Speiser, Gianluca Antonini, and Abderrahim Labbi</i>	
Interesting Multi-relational Patterns	675
<i>Eirini Spyropoulou and Tijl De Bie</i>	
On Generating All Optimal Monotone Classifications	685
<i>Luite Stegeman and Ad Feelders</i>	
Recursive Multi-step Time Series Forecasting by Perturbing Data	695
<i>Souhaib Ben Taieb and Gianluca Bontempi</i>	
Finding Robust Itemsets under Subsampling	705
<i>Nikolaj Tatti and Fabian Moerchen</i>	
Density Estimation Based on Mass	715
<i>Kai Ming Ting, Takashi Washio, Jonathan R. Wells, and Fei Tony Liu</i>	
Diverse Dimension Decomposition of an Itemset Space	725
<i>Mikalai Tsytsarau, Francesco Bonchi, Aristides Gionis, and Themis Palpanas</i>	

Conditional Anomaly Detection with Soft Harmonic Functions	735
<i>Michal Valko, Branislav Kveton, Hamed Valizadegan, Gregory F. Cooper, and Milos Hauskrecht</i>	
Random Forest Based Feature Induction	744
<i>Celine Vens and Fabrizio Costa</i>	
Class Imbalance, Redux	754
<i>Byron C. Wallace, Kevin Small, Carla E. Brodley, and Thomas A. Trikalinos</i>	
Combining Feature Context and Spatial Context for Image Pattern Discovery	764
<i>Hongxing Wang, Junsong Yuan, and Yap-Peng Tan</i>	
Nonnegative Matrix Tri-factorization Based High-Order Co-clustering and Its Fast Implementation	774
<i>Hua Wang, Feiping Nie, Heng Huang, and Chris Ding</i>	
Detecting Community Kernels in Large Social Networks	784
<i>Liaoruo Wang, Tiancheng Lou, Jie Tang, and John E. Hopcroft</i>	
ADANA: Active Name Disambiguation	794
<i>Xuezhi Wang, Jie Tang, Hong Cheng, and Philip S. Yu</i>	
Document Clustering via Matrix Representation	804
<i>Xufei Wang, Jiliang Tang, and Huan Liu</i>	
Using Bayesian Network Learning Algorithm to Discover Causal Relations in Multivariate Time Series	814
<i>Zhenxing Wang and Laiwan Chan</i>	
Efficient Mining of a Concise and Lossless Representation of High Utility Itemsets	824
<i>Cheng Wei Wu, Philippe Fournier-Viger, Philip S. Yu, and Vincent S. Tseng</i>	
Understanding Propagation Error and Its Effect on Collective Classification	834
<i>Rongjing Xiang and Jennifer Neville</i>	
Direct Robust Matrix Factorization for Anomaly Detection	844
<i>Liang Xiong, Xi Chen, and Jeff Schneider</i>	
A New Markov Model for Clustering Categorical Sequences	854
<i>Tengke Xiong, Shengrui Wang, Qingshan Jiang, and Joshua Zhexue Huang</i>	
BibClus: A Clustering Algorithm of Bibliographic Networks by Message Passing on Center Linkage Structure	864
<i>Xiaoran Xu and Zhi-Hong Deng</i>	
Multi-instance Metric Learning	874
<i>Ye Xu, Wei Ping, and Andrew T. Campbell</i>	
Multi-task Learning with Task Relations	884
<i>Zhao Xu and Kristian Kersting</i>	
Secure Clustering in Private Networks	894
<i>Bin Yang, Issei Sato, and Hiroshi Nakagawa</i>	

LPTA: A Probabilistic Model for Latent Periodic Topic Analysis	904
<i>Zhijun Yin, Liangliang Cao, Jiawei Han, Chengxiang Zhai, and Thomas Huang</i>	
Causal Associative Classification	914
<i>Kui Yu, Xindong Wu, Wei Ding, Hao Wang, and Hongliang Yao</i>	
Multi-task Learning for Bayesian Matrix Factorization	924
<i>Chao Yuan</i>	
Enabling Fast Lazy Learning for Data Streams	932
<i>Peng Zhang, Byron J. Gao, Xingquan Zhu, and Li Guo</i>	
Clusterability Analysis and Incremental Sampling for Nyström Extension Based Spectral Clustering	942
<i>Xianchao Zhang and Quanzeng You</i>	
Fast and Robust Graph-based Transductive Learning via Minimum Tree Cut	952
<i>Yan-Ming Zhang, Kaizhu Huang, and Cheng-Lin Liu</i>	
Positive and Unlabeled Learning for Graph Classification	962
<i>Yuchen Zhao, Xiangnan Kong, and Philip S. Yu</i>	
Finding Novel Diagnostic Gene Patterns Based on Interesting Non-redundant Contrast Sequence Rules	972
<i>Yuhai Zhao, Guoren Wang, Yuan Li, and Zhanghui Wang</i>	
Semi-supervised Hierarchical Clustering	982
<i>Li Zheng and Tao Li</i>	
Handling Conditional Discrimination	992
<i>Indre Žliobaite, Faisal Kamiran, and Toon Calders</i>	
 Short Papers	
On the Hardness of Graph Anonymization	1002
<i>Charu C. Aggarwal, Yao Li, and Philip S. Yu</i>	
SPO: Structure Preserving Oversampling for Imbalanced Time Series Classification	1008
<i>Hong Cao, Xiao-Li Li, Yew-Kwong Woon, and See-Kiong Ng</i>	
Manifold Learning and Missing Data Recovery through Unsupervised Regression	1014
<i>Miguel Á. Carreira-Perpiñán and Zhengdong Lu</i>	
Characterizing Inverse Time Dependency in Multi-class Learning	1020
<i>Danqi Chen, Weizhu Chen, and Qiang Yang</i>	
Supervised Lazy Random Walk for Topic-Focused Multi-document Summarization	1026
<i>Pan Du, Jiafeng Guo, and Xueqi Cheng</i>	
Unsupervised Anomaly Intrusion Detection via Localized Bayesian Feature Selection	1032
<i>Wentao Fan, Nizar Bouguila, and Djemel Ziou</i>	
Identifying Differentially Expressed Genes via Weighted Rank Aggregation	1038
<i>Qiong Fang, Jianlin Feng, and Wilfred Ng</i>	

Efficient Mining of Closed Sequential Patterns on Stream Sliding Window	1044
<i>Chuancong Gao, Jianyong Wang, and Qingyan Yang</i>	
A Spectral Framework for Detecting Inconsistency across Multi-source Object Relationships	1050
<i>Jing Gao, Wei Fan, Deepak Turaga, Srinivasan Parthasarathy, and Jiawei Han</i>	
Tracking and Connecting Topics via Incremental Hierarchical Dirichlet Processes	1056
<i>Zekai J. Gao, Yangqiu Song, Shixia Liu, Haixun Wang, Hao Wei, Yang Chen, and Weiwei Cui</i>	
Learning Protein Folding Energy Functions	1062
<i>Wei Guan, Arkadas Ozakin, Alexander Gray, Jose Borreguero, Shashi Pandit, Anna Jagielska, Liliana Wroblewska, and Jeffrey Skolnick</i>	
How Does Research Evolve? Pattern Mining for Research Meme Cycles	1068
<i>Dan He, Xingquan Zhu, and D. Stott Parker</i>	
Twin Gaussian Processes for Binary Classification	1074
<i>Jianjun He, Hong Gu, and Shaorui Jiang</i>	
Constraint Selection-Based Semi-supervised Feature Selection	1080
<i>Mohammed Hindawi, Kais Allab, and Khalid Benabdeslem</i>	
Discovering the Intrinsic Cardinality and Dimensionality of Time Series Using MDL	1086
<i>Bing Hu, Thanawin Rakthanmanon, Yuan Hao, Scott Evans, Stefano Lonardi, and Eamonn Keogh</i>	
Discovery of Versatile Temporal Subspace Patterns in 3-D Datasets	1092
<i>Zhen Hu and Raj Bhatnagar</i>	
A Fixed Parameter Tractable Integer Program for Finding the Maximum Order Preserving Submatrix	1098
<i>Jens Humrich, Thomas Gärtner, and Gemma C. Garriga</i>	
ASAP: A Self-Adaptive Prediction System for Instant Cloud Resource Demand Provisioning	1104
<i>Yexi Jiang, Chang-shing Perng, Tao Li, and Rong Chang</i>	
Learning from Negative Examples in Set-Expansion	1110
<i>Prateek Jindal and Dan Roth</i>	
Entropy-Based Graph Clustering: Application to Biological and Social Networks	1116
<i>Edward Casey Kenley and Young-Rae Cho</i>	
Semi-supervised Discriminant Hashing	1122
<i>Saehoon Kim and Seungjin Choi</i>	
Using Frequent Closed Itemsets for Data Dimensionality Reduction	1128
<i>Petr Krajca, Jan Outrata, and Vilem Vychodil</i>	
Modeling High-Level Behavior Patterns for Precise Similarity Analysis of Software	1134
<i>Taeho Kwon and Zhendong Su</i>	

Co-clustering for Binary and Categorical Data with Maximum Modularity	1140
<i>Lazhar Labiod and Mohamed Nadif</i>	
Calculating Feature Weights in Naive Bayes with Kullback-Leibler Measure	1146
<i>Chang-Hwan Lee, Fernando Gutierrez, and Dejing Dou</i>	
Scalable Diversified Ranking on Large Graphs	1152
<i>Rong-Hua Li and Jeffery Xu Yu</i>	
Web Horror Image Recognition Based on Context-Aware Multi-instance Learning	1158
<i>Bing Li, Weihua Xiong, and Weiming Hu</i>	
Mixture of Softmax sLDA	1164
<i>Xiaoxu Li, Junyu Zeng, Xiaojie Wang, and Yixin Zhong</i>	
Optimizing Performance Measures for Feature Selection	1170
<i>Qi Mao and Ivor Wai-Hung Tsang</i>	
Detecting Recurring and Novel Classes in Concept-Drifting Data Streams	1176
<i>Mohammad M. Masud, Tahseen M. Al-Khateeb, Latifur Khan, Charu Aggarwal, Jing Gao, Jiawei Han, and Bhavani Thuraisingham</i>	
Performances and Characteristics of DIGRank, Ranking in the Incomplete Networks	1182
<i>Xiang Niu, Lusong Li, Xiaobing Xiong, Daniel Tkach, He Li, and Ke Xu</i>	
Cross-Temporal Link Prediction	1188
<i>Satoshi Oyama, Kohei Hayashi, and Hisashi Kashima</i>	
Helix: Unsupervised Grammar Induction for Structured Activity Recognition	1194
<i>Huan-Kai Peng, Pang Wu, Jiang Zhu, and Joy Ying Zhang</i>	
Distance Preserving Graph Simplification	1200
<i>Ning Ruan, Ruoming Jin, and Yan Huang</i>	
Clustering with Attribute-Level Constraints	1206
<i>Jana Schmidt, Elisabeth Maria Brändle, and Stefan Kramer</i>	
A Fast and Flexible Clustering Algorithm Using Binary Discretization	1212
<i>Mahito Sugiyama and Akihiro Yamamoto</i>	
A New Multi-task Learning Method for Personalized Activity Recognition	1218
<i>Xu Sun, Hisashi Kashima, Ryota Tomioka, Naonori Ueda, and Ping Li</i>	
Identities Anonymization in Dynamic Social Networks	1224
<i>Chih-Hua Tai, Peng-Jui Tseng, Philip S. Yu, and Ming-Syan Chen</i>	
Discovering Emerging Topics in Social Streams via Link Anomaly Detection	1230
<i>Toshimitsu Takahashi, Ryota Tomioka, and Kenji Yamanishi</i>	
Finding Communities in Dynamic Social Networks	1236
<i>Chayant Tantipathananandh and Tanya Y. Berger-Wolf</i>	
Review Graph Based Online Store Review Spammer Detection	1242
<i>Guan Wang, Sihong Xie, Bing Liu, and Philip S. Yu</i>	

Classifying Categorical Data by Rule-Based Neighbors	1248
<i>Jiabing Wang, Pei Zhang, Guihua Wen, and Jia Wei</i>	
Tensor Fold-in Algorithms for Social Tagging Prediction	1254
<i>Miao Zhang, Chris Ding, and Zhifang Liao</i>	
Discovering Thematic Patterns in Videos via Cohesive Sub-graph Mining	1260
<i>Gangqiang Zhao and Junsong Yuan</i>	
Low Rank Metric Learning with Manifold Regularization	1266
<i>Guoqiang Zhong, Kaizhu Huang, and Cheng-Lin Liu</i>	
A Study of Laplacian Spectra of Graph for Subgraph Queries	1272
<i>Lei Zhu and Qinbao Song</i>	
Text Clustering via Constrained Nonnegative Matrix Factorization	1278
<i>Yan Zhu, Liping Jing, and Jian Yu</i>	

Author Index