

# Table of Contents

Preface .....	xiii
Organization .....	xviii
Program Committees .....	xix
Funding.....	xxiii
Invited Talks .....	xxv
Tutorial and Workshop Summaries .....	xxxiii
<b>The Role of Machine Learning in Business Optimization .....</b>	<b>1</b>
<i>Chid Apte</i>	
<b>FAB-MAP: Appearance-Based Place Recognition and Mapping using a Learned Visual Vocabulary Model .....</b>	<b>3</b>
<i>Mark Cummins, Paul Newman</i>	
<b>Discriminative Latent Variable Models for Object Detection .....</b>	<b>11</b>
<i>Pedro Felzenszwalb, Ross Girshick, David McAllester, Deva Ramanan</i>	
<b>Web-Scale Bayesian Click-Through rate Prediction for Sponsored Search Advertising in Microsoft's Bing Search Engine.....</b>	<b>13</b>
<i>Thore Graepel, Joaquin Quiñero Candela, Thomas Borchert, Ralf Herbrich</i>	
<b>Music Plus One and Machine Learning.....</b>	<b>21</b>
<i>Christopher Raphael</i>	
<b>Climbing the Tower of Babel: Unsupervised Multilingual Learning .....</b>	<b>29</b>
<i>Benjamin Snyder, Regina Barzilay</i>	
<b>Detecting Large-Scale System Problems by Mining Console Logs.....</b>	<b>37</b>
<i>Wei Xu, Ling Huang, Armando Fox, David Patterson, Michael I. Jordan</i>	
<b>Particle Filtered MCMC-MLE with Connections to Contrastive Divergence.....</b>	<b>47</b>
<i>Arthur U. Asuncion, Qiang Liu, Alexander T. Ihler, Padhraic Smyth</i>	
<b>Surrogating the surrogate: accelerating Gaussian-process-based global optimization with a mixture cross-entropy algorithm .....</b>	<b>55</b>
<i>Remi Bardenet, Balazs Kegl</i>	
<b>Forgetting Counts : Constant Memory Inference for a Dependent Hierarchical Pitman-Yor Process.....</b>	<b>63</b>
<i>Nicholas Bartlett, David Pfau, Frank Wood</i>	
<b>Robust Formulations for Handling Uncertainty in Kernel Matrices .....</b>	<b>71</b>
<i>Sahely Bhadra, Sourangshu Bhattacharya, Chiranjib Bhattacharyya, Aharon Ben-tal</i>	
<b>Active Learning for Networked Data.....</b>	<b>79</b>
<i>Mustafa Bilgic, Lilyana Mihalkova, Lise Getoor</i>	
<b>Distance dependent Chinese restaurant processes .....</b>	<b>87</b>
<i>David M. Blei, Peter Frazier</i>	
<b>Causal filter selection in microarray data .....</b>	<b>95</b>
<i>Gianluca Bontempi, Patrick E. Meyer</i>	
<b>Label Ranking under Ambiguous Supervision for Learning Semantic Correspondences .....</b>	<b>103</b>
<i>Antoine Bordes, Nicolas Usunier, Jason Weston</i>	
<b>A Theoretical Analysis of Feature Pooling in Visual Recognition .....</b>	<b>111</b>
<i>Y-Lan Boureau, Jean Ponce, Yann LeCun</i>	
<b>Multi-agent Learning Experiments on Repeated Matrix Games .....</b>	<b>119</b>
<i>Bruno Bouzy, Marc Metivier</i>	

<b>Learning Tree Conditional Random Fields .....</b>	<b>127</b>
<i>Joseph K. Bradley, Carlos Guestrin</i>	
<b>Finding Planted Partitions in Nearly Linear Time using Arrested Spectral Clustering.....</b>	<b>135</b>
<i>Nader H. Bshouty, Philip M. Long</i>	
<b>Fast boosting using adversarial bandits .....</b>	<b>143</b>
<i>Róbert Busa-Fekete, Balázs Kégl</i>	
<b>Modeling Transfer Learning in Human Categorization with the Hierarchical Dirichlet Process.....</b>	<b>151</b>
<i>Kevin R. Canini, Mikhail M. Shashkov, Thomas L. Griffiths</i>	
<b>Transfer Learning for Collective Link Prediction in Multiple Heterogenous Domains .....</b>	<b>159</b>
<i>Bin Cao, Nathan Nan Liu, Qiang Yang</i>	
<b>The Elastic Embedding Algorithm for Dimensionality Reduction.....</b>	<b>167</b>
<i>Miguel Á. Carreira-Perpiñán</i>	
<b>Random Spanning Trees and the Prediction of Weighted Graphs .....</b>	<b>175</b>
<i>Nicolo Cesa-Bianchi, Claudio Gentile, Fabio Vitale, Giovanni Zappella</i>	
<b>Efficient Learning with Partially Observed Attributes .....</b>	<b>183</b>
<i>Nicolo Cesa-Bianchi, Shai Shalev-Shwartz, Ohad Shamir</i>	
<b>Convergence, Targeted Optimality, and Safety in Multiagent Learning .....</b>	<b>191</b>
<i>Doran Chakraborty, Peter Stone</i>	
<b>Structured Output Learning with Indirect Supervision .....</b>	<b>199</b>
<i>Ming-Wei Chang, Vivek Srikumar, Dan Goldwasser, Dan Roth</i>	
<b>Dynamical Products of Experts for Modeling Financial Time Series.....</b>	<b>207</b>
<i>Yutian Chen, Max Welling</i>	
<b>Label Ranking Methods based on the Plackett-Luce Model.....</b>	<b>215</b>
<i>Weiwei Cheng, Krzysztof Dembczynski, Eyke Hüllermeier</i>	
<b>Graded Multilabel Classification: The Ordinal Case .....</b>	<b>223</b>
<i>Weiwei Cheng, Krzysztof Dembczynski, Eyke Hüllermeier</i>	
<b>Comparing Clusterings in Space .....</b>	<b>231</b>
<i>Michael H. Coen, M. Hidayath Ansari, Nathanael Fillmore</i>	
<b>Two-Stage Learning Kernel Algorithms .....</b>	<b>239</b>
<i>Corinna Cortes, Mehryar Mohri, Afshin Rostamizadeh</i>	
<b>Generalization Bounds for Learning Kernels .....</b>	<b>247</b>
<i>Corinna Cortes, Mehryar Mohri, Afshin Rostamizadeh</i>	
<b>Fast Neighborhood Subgraph Pairwise Distance Kernel.....</b>	<b>255</b>
<i>Fabrizio Costa, Kurt De Grave</i>	
<b>Mining Clustering Dimensions.....</b>	<b>263</b>
<i>Sajib Dasgupta, Vincent Ng</i>	
<b>Bottom-Up Learning of Markov Network Structure .....</b>	<b>271</b>
<i>Jesse Davis, Pedro Domingos</i>	
<b>Bayes Optimal Multilabel Classification via Probabilistic Classifier Chains .....</b>	<b>279</b>
<i>Krzysztof Dembczynski, Weiwei Cheng, Eyke Hüllermeier</i>	
<b>A Conditional Random Field for Multiple-Instance Learning.....</b>	<b>287</b>
<i>Thomas Deselaers, Vittorio Ferrari</i>	
<b>Asymptotic Analysis of Generative Semi-Supervised Learning .....</b>	<b>295</b>
<i>Joshua V. Dillon, Krishnakumar Balasubramanian, Guy Lebanon</i>	

<b>Heterogeneous Continuous Dynamic Bayesian Networks with Flexible Structure and Inter-Time Segment Information Sharing .....</b>	<b>303</b>
<i>Frank Dondelinger, Sophie Lebre, Dirk Husmeier</i>	
<b>Temporal Difference Bayesian Model Averaging: A Bayesian Perspective on Adapting Lambda.....</b>	<b>311</b>
<i>Carlton Downey, Scott Sanner</i>	
<b>High-Performance Semi-Supervised Learning using Discriminatively Constrained Generative Models .....</b>	<b>319</b>
<i>Gregory Druck, Andrew McCallum</i>	
<b>On the Consistency of Ranking Algorithms .....</b>	<b>327</b>
<i>John C. Duchi, Lester W. Mackey, Michael I. Jordan</i>	
<b>Inverse Optimal Control with Linearly-Solvable MDPs.....</b>	<b>335</b>
<i>Krishnamurthy Dvijotham, Emanuel Todorov</i>	
<b>Continuous-Time Belief Propagation .....</b>	<b>343</b>
<i>Tal El-Hay, Ido Cohn, Nir Friedman, Raz Kupferman</i>	
<b>Nonparametric Information Theoretic Clustering Algorithm .....</b>	<b>351</b>
<i>Lev Faivishevsky, Jacob Goldberger</i>	
<b>Feature Selection as a One-Player Game .....</b>	<b>359</b>
<i>Romarc Gaudel, Michele Sebag</i>	
<b>Multiscale Wavelets on Trees, Graphs and High Dimensional Data: Theory and Applications to Semi Supervised Learning .....</b>	<b>367</b>
<i>Matan Gavish, Boaz Nadler, Ronald R. Coifman</i>	
<b>A Language-based Approach to Measuring Scholarly Impact.....</b>	<b>375</b>
<i>Sean M. Gerrish, David M. Blei</i>	
<b>Boosting Classifiers with Tightened L0-Relaxation Penalties.....</b>	<b>383</b>
<i>Noam Goldberg, Jonathan Eckstein</i>	
<b>Budgeted Nonparametric Learning from Data Streams .....</b>	<b>391</b>
<i>Ryan Gomes, Andreas Krause</i>	
<b>Learning Fast Approximations of Sparse Coding.....</b>	<b>399</b>
<i>Karol Gregor, Yann LeCun</i>	
<b>Boosted Backpropagation Learning for Training Deep Modular Networks .....</b>	<b>407</b>
<i>Alexander Grubb, J. Andrew Bagnell</i>	
<b>Interactive Submodular Set Cover.....</b>	<b>415</b>
<i>Andrew Guillory, Jeff Bilmes</i>	
<b>Large Scale Max-Margin Multi-Label Classification with Priors .....</b>	<b>423</b>
<i>Bharath Hariharan, Lih Zelnik-Manor, S. V. N. Vishwanathan, Manik Varma</i>	
<b>Active Learning for Multi-Task Adaptive Filtering .....</b>	<b>431</b>
<i>Abhay Harpale, Yiming Yang</i>	
<b>Bayesian Nonparametric Matrix Factorization for Recorded Music .....</b>	<b>439</b>
<i>Matthew D. Hoffman, David M. Blei, Perry R. Cook</i>	
<b>Multi-Task Learning of Gaussian Graphical Models.....</b>	<b>447</b>
<i>Jean Honorio, Dimitris Samaras</i>	
<b>Learning Hierarchical Riffle Independent Groupings from Rankings.....</b>	<b>455</b>
<i>Jonathan Huang, Carlos Guestrin</i>	
<b>On learning with kernels for unordered pairs .....</b>	<b>463</b>
<i>Martial Hue, Jean-Philippe Vert</i>	

<b>A Simple Algorithm for Nuclear Norm Regularized Problems.....</b>	<b>471</b>
<i>Martin Jaggi, Marek Sulovský</i>	
<b>Telling cause from effect based on high-dimensional observations.....</b>	<b>479</b>
<i>Dominik Janzing, Patrik O. Hoyer, Bernhard Schölkopf</i>	
<b>Proximal Methods for Sparse Hierarchical Dictionary Learning .....</b>	<b>487</b>
<i>Rodolphe Jenatton, Julien Mairal, Guillaume Obozinski, Francis Bach</i>	
<b>3D Convolutional Neural Networks for Human Action Recognition .....</b>	<b>495</b>
<i>Shuiwang Ji, Wei Xu, Ming Yang, Kai Yu</i>	
<b>Accelerated dual decomposition for MAP inference .....</b>	<b>503</b>
<i>Vladimir Jojic, Stephen Gould, Daphne Koller</i>	
<b>Efficient Selection of Multiple Bandit Arms: Theory and Practice .....</b>	<b>511</b>
<i>Shivaram Kalyanakrishnan, Peter Stone</i>	
<b>A scalable trust-region algorithm with application to mixed-norm regression .....</b>	<b>519</b>
<i>Dongmin Kim, Suvrit Sra, Inderjit Dhillon</i>	
<b>Local Minima Embedding .....</b>	<b>527</b>
<i>Minyoung Kim, Fernando De la Torre</i>	
<b>Gaussian Processes Multiple Instance Learning.....</b>	<b>535</b>
<i>Minyoung Kim, Fernando De la Torre</i>	
<b>Tree-Guided Group Lasso for Multi-Task Regression with Structured Sparsity.....</b>	<b>543</b>
<i>Seyoung Kim, Eric P. Xing</i>	
<b>Learning Markov Logic Networks Using Structural Motifs .....</b>	<b>551</b>
<i>Stanley Kok, Pedro Domingos</i>	
<b>On Sparse Nonparametric Conditional Covariance Selection.....</b>	<b>559</b>
<i>Mladen Kolar, Ankur P. Parikh, Eric P. Xing</i>	
<b>Submodular Dictionary Selection for Sparse Representation.....</b>	<b>567</b>
<i>Andreas Krause, Volkan Cevher</i>	
<b>Implicit Online Learning.....</b>	<b>575</b>
<i>Brian Kulis, Peter L. Bartlett</i>	
<b>Probabilistic Backward and Forward Reasoning in Stochastic Relational Worlds.....</b>	<b>583</b>
<i>Tobias Lang, Marc Toussaint</i>	
<b>Supervised Aggregation of Classifiers using Artificial Prediction Markets.....</b>	<b>591</b>
<i>Nathan Lay, Adrian Barbu</i>	
<b>Bayesian Multi-Task Reinforcement Learning .....</b>	<b>599</b>
<i>Alessandro Lazaric, Mohammad Ghavamzadeh</i>	
<b>Analysis of a Classification-based Policy Iteration Algorithm .....</b>	<b>607</b>
<i>Alessandro Lazaric, Mohammad Ghavamzadeh, Rémi Munos</i>	
<b>Finite-Sample Analysis of LSTD .....</b>	<b>615</b>
<i>Alessandro Lazaric, Mohammad Ghavamzadeh, Rémi Munos</i>	
<b>A fast natural Newton method.....</b>	<b>623</b>
<i>Nicolas Le Roux, Andrew Fitzgibbon</i>	
<b>Making Large-Scale Nyström Approximation Possible .....</b>	<b>631</b>
<i>Mu Li, James T. Kwok, Bao-Liang Lu</i>	
<b>Learning Programs: A Hierarchical Bayesian Approach .....</b>	<b>639</b>
<i>Percy Liang, Michael I. Jordan, Dan Klein</i>	

<b>On the Interaction between Norm and Dimensionality: Multiple Regimes in Learning</b> .....	647
<i>Percy Liang, Nati Srebro</i>	
<b>Power Iteration Clustering</b> .....	655
<i>Frank Lin, William W. Cohen</i>	
<b>Robust Subspace Segmentation by Low-Rank Representation</b> .....	663
<i>Guangcan Liu, Zhouchen Lin, Yong Yu</i>	
<b>Robust Graph Mode Seeking by Graph Shift</b> .....	671
<i>Hairong Liu, Shuicheng Yan</i>	
<b>Large Graph Construction for Scalable Semi-Supervised Learning</b> .....	679
<i>Wei Liu, Junfeng He, Shih-Fu Chang</i>	
<b>Learning Temporal Causal Graphs for Relational Time-Series Analysis</b> .....	687
<i>Yan Liu, Alexandru Niculescu-Mizil, Aurélie Lozano , Yong Lu</i>	
<b>Efficient Reinforcement Learning with Multiple Reward Functions for Randomized Controlled Trial Analysis</b> .....	695
<i>Daniel J. Lizotte, Michael Bowling, Susan A. Murphy</i>	
<b>Restricted Boltzmann Machines are Hard to Approximately Evaluate or Simulate</b> .....	703
<i>Philip M. Long, Rocco A. Servedio</i>	
<b>Mixed Membership Matrix Factorization</b> .....	711
<i>Lester Mackey, David Weiss, Michael I. Jordan</i>	
<b>Toward Off-Policy Learning Control with Function Approximation</b> .....	719
<i>Hamid Reza Maei, Csaba Szepesvari, Shalabh Bhatnagar, Richard S. Sutton</i>	
<b>Constructing States for Reinforcement Learning</b> .....	727
<i>M. M. Hassan Mahmud</i>	
<b>Deep learning via Hessian-free optimization</b> .....	735
<i>James Martens</i>	
<b>Learning the Linear Dynamical System with ASOS</b> .....	743
<i>James Martens</i>	
<b>From Transformation-Based Dimensionality Reduction to Feature Selection</b> .....	751
<i>Mahdokht Masaali, Glenn Fung, Jennifer G. Dy</i>	
<b>Risk minimization, probability elicitation, and cost-sensitive SVMs</b> .....	759
<i>Hamed Masnadi-Shirazi, Nuno Vasconcelos</i>	
<b>Exploiting Data-Independence for Fast Belief-Propagation</b> .....	767
<i>Julian J. McAuley, Tibério S. Caetano</i>	
<b>Metric Learning to Rank</b> .....	775
<i>Brian McFee, Gert Lanckriet</i>	
<b>Learning Efficiently with Approximate Inference via Dual Losses</b> .....	783
<i>Ofer Meshi, David Sontag, Tommi Jaakkola, Amir Globerson</i>	
<b>Deep Supervised t-Distributed Embedding</b> .....	791
<i>Renqiang Min, Laurens van der Maaten, Zineng Yuan, Anthony Bonner, Zhaolei Zhang</i>	
<b>Nonparametric Return Distribution Approximation for Reinforcement Learning</b> .....	799
<i>Tetsuro Morimura, Masashi Sugiyama, Hisashi Kashima , Hirotaka Hachiya, Toshiyuki Tanaka</i>	
<b>Rectified Linear Units Improve Restricted Boltzmann Machines</b> .....	807
<i>Vinod Nair, Geoffrey E. Hinton</i>	
<b>Implicit Regularization in Variational Bayesian Matrix Factorization</b> .....	815
<i>Shinichi Nakajima, Masashi Sugiyama</i>	

<b>Estimation of (near) low-rank matrices with noise and high-dimensional scaling .....</b>	<b>823</b>
<i>Sahand Negahban, Martin J. Wainwright</i>	
<b>Multiple Non-Redundant Spectral Clustering Views.....</b>	<b>831</b>
<i>Donglin Niu, Jennifer G. Dy, Michael I. Jordan</i>	
<b>Multiagent Inductive Learning: an Argumentation-based Approach.....</b>	<b>839</b>
<i>Santiago Ontañón, Enric Plaza</i>	
<b>A Stick-Breaking Construction of the Beta Process.....</b>	<b>847</b>
<i>John Paisley, Aimee Zaas, Christopher W. Woods, Geoffrey S. Ginsburg, Lawrence Carin</i>	
<b>The Margin Perceptron with Unlearning .....</b>	<b>855</b>
<i>Constantinos Panagiotakopoulos, Petroula Tsampouka</i>	
<b>Boosting for Regression Transfer .....</b>	<b>863</b>
<i>David Pardoe, Peter Stone</i>	
<b>Feature Selection Using Regularization in Approximate Linear Programs for Markov Decision Processes.....</b>	<b>871</b>
<i>Marek Petrik, Gavin Taylor, Ron Parr, Shlomo Zilberstein*</i>	
<b>Budgeted Distribution Learning of Belief Net Parameters.....</b>	<b>879</b>
<i>Liuyang Li, Barnabás Póczos, Csaba Szepesvári, Russ Greiner</i>	
<b>Variable Selection in Model-Based Clustering: To Do or To Facilitate.....</b>	<b>887</b>
<i>Leonard K. M. Poon, Nevin L. Zhang, Tao Chen, Yi Wang</i>	
<b>Approximate Predictive Representations of Partially Observable Systems .....</b>	<b>895</b>
<i>Monica Dinculescu, Doina Precup</i>	
<b>Spherical Topic Models .....</b>	<b>903</b>
<i>Joseph Reisinger, Austin Waters, Bryan Silverthorn, Raymond J. Mooney</i>	
<b>SVM Classifier Estimation from Group Probabilities.....</b>	<b>911</b>
<i>Stefan Rueping</i>	
<b>Clustering processes .....</b>	<b>919</b>
<i>Daniil Ryabko</i>	
<b>Gaussian Process Change Point Models.....</b>	<b>927</b>
<i>Yunus Saatçi, Ryan Turner, Carl Edward Rasmussen</i>	
<b>Online Prediction with Privacy.....</b>	<b>935</b>
<i>Jun Sakuma, Hiromi Arai</i>	
<b>Learning Deep Boltzmann Machines using Adaptive MCMC .....</b>	<b>943</b>
<i>Ruslan Salakhutdinov</i>	
<b>Active Risk Estimation.....</b>	<b>951</b>
<i>Christoph Sawade, Niels Landwehr, Steffen Bickel, Tobias Scheffer</i>	
<b>Should one compute the Temporal Difference fix point or minimize the Bellman Residual ? The unified oblique projection view.....</b>	<b>959</b>
<i>Bruno Scherrer</i>	
<b>Gaussian Covariance and Scalable Variational Inference .....</b>	<b>967</b>
<i>Matthias W. Seeger</i>	
<b>Application of Machine Learning To Epileptic Seizure Detection .....</b>	<b>975</b>
<i>Ali Shoeb, John Guttag</i>	
<b>Learning optimally diverse rankings over large document collections .....</b>	<b>983</b>
<i>Aleksandrs Slivkins, Filip Radlinski, Sreenivas Gollapudi</i>	

<b>Hilbert Space Embeddings of Hidden Markov Models</b> .....	991
<i>Le Song, Sajid M. Siddiqi, Geoffrey Gordon, Alex Smola</i>	
<b>COFFIN : A Computational Framework for Linear SVMs</b> .....	999
<i>Soeren Sonnenburg, Vojtech Franc</i>	
<b>Internal Rewards Mitigate Agent Boundedness</b> .....	1007
<i>Jonathan Sorg, Satinder Singh, Richard Lewis</i>	
<b>Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design</b> .....	1015
<i>Niranjn Srinivas, Andreas Krause, Sham Kakade, Matthias Seeger</i>	
<b>Unsupervised Risk Stratification in Clinical Datasets: Identifying Patients at Risk of Rare Outcomes</b> .....	1023
<i>Zeeshan Syed, Ilan Rubinfeld</i>	
<b>Model-based reinforcement learning with nearly tight exploration complexity bounds</b> .....	1031
<i>István Szita, Csaba Szepesvári</i>	
<b>Total Variation and Cheeger Cuts</b> .....	1039
<i>Arthur Szlam, Xavier Bresson</i>	
<b>Learning Sparse SVM for Feature Selection on Very High Dimensional Datasets</b> .....	1047
<i>Mingkui Tan, Li Wang, Ivor W. Tsang</i>	
<b>Deep networks for robust visual recognition</b> .....	1055
<i>Yichuan Tang, Chris Eliasmith</i>	
<b>A DC Programming Approach for Sparse Eigenvalue Problem</b> .....	1063
<i>Mamadou Thiao, Tao Pham Dinh, Hoai An Le Thi</i>	
<b>Least-Squares Policy Iteration: Bias-Variance Trade-off in Control Problems</b> .....	1071
<i>Christophe Thiery, Bruno Scherrer</i>	
<b>An Analysis of the Convergence of Graph Laplacians</b> .....	1079
<i>Daniel Ting, Ling Huang, Michael I. Jordan</i>	
<b>A Fast Augmented Lagrangian Algorithm for Learning Low-Rank Matrices</b> .....	1087
<i>Ryota Tomioka, Taiji Suzuki, Masashi Sugiyama, Hisashi Kashima</i>	
<b>One-sided Support Vector Regression for Multiclass Cost-sensitive Classification</b> .....	1095
<i>Han-Hsing Tu, Hsuan-Tien Lin</i>	
<b>Non-Local Contrastive Objectives</b> .....	1103
<i>David Vickrey, Cliff Chiung-Yu Lin, Daphne Koller</i>	
<b>The Translation-invariant Wishart-Dirichlet Process for Clustering Distance Data</b> .....	1111
<i>Julia E. Vogt, Sandhya Prabhakaran, Thomas J. Fuchs, Volker Roth</i>	
<b>Generalizing Apprenticeship Learning across Hypothesis Classes</b> .....	1119
<i>Thomas J. Walsh, Kaushik Subramanian, Michael L. Littman, Carlos Diuk</i>	
<b>Sequential Projection Learning for Hashing with Compact Codes</b> .....	1127
<i>Jun Wang, Sanjiv Kumar, Shih-Fu Chang</i>	
<b>A New Analysis of Co-Training</b> .....	1135
<i>Wei Wang, Zhi-Hua Zhou</i>	
<b>Multi-Class Pegasos on a Budget</b> .....	1143
<i>Zhuang Wang, Koby Crammer, Slobodan Vucetic</i>	
<b>The IBP Compound Dirichlet Process and its Application to Focused Topic Modeling</b> .....	1151
<i>Sinead Williamson, Chong Wang, Katherine A. Heller, David M. Blei</i>	
<b>Online Streaming Feature Selection</b> .....	1159
<i>Xindong Wu, Kui Yu, Hao Wang, Wei Ding</i>	

<b>Classes of Multiagent Q-learning Dynamics with <math>\epsilon</math>-greedy Exploration .....</b>	<b>1167</b>
<i>Michael Wunder, Michael Littman, Monica Babes</i>	
<b>Simple and Efficient Multiple Kernel Learning by Group Lasso .....</b>	<b>1175</b>
<i>Zenglin Xu, Rong Jin, Haiqin Yang, Irwin King, Michael R. Lyu</i>	
<b>Sparse Gaussian Process Regression via <math>l_1</math> Penalization .....</b>	<b>1183</b>
<i>Feng Yan, Yuan (Alan) Qi</i>	
<b>Online Learning for Group Lasso .....</b>	<b>1191</b>
<i>Haiqin Yang, Zenglin Xu, Irwin King, Michael R. Lyu</i>	
<b>Learning from Noisy Side Information by Generalized Maximum Entropy Model .....</b>	<b>1199</b>
<i>Tianbao Yang, Rong Jin, Anil K. Jain</i>	
<b>Convergence of Least Squares Temporal Difference Methods Under General Conditions .....</b>	<b>1207</b>
<i>Huizhen Yu</i>	
<b>Improved Local Coordinate Coding using Local Tangents .....</b>	<b>1215</b>
<i>Kai Yu, Tong Zhang</i>	
<b>Projection Penalties: Dimension Reduction without Loss.....</b>	<b>1223</b>
<i>Yi Zhang, Jeff Schneider</i>	
<b>OTL: A Framework of Online Transfer Learning .....</b>	<b>1231</b>
<i>Peilin Zhao, Steven C.H. Hoi</i>	
<b>Conditional Topic Random Fields .....</b>	<b>1239</b>
<i>Jun Zhu, Eric P. Xing</i>	
<b>Cognitive Models of Test-Item Effects in Human Category Learning .....</b>	<b>1247</b>
<i>Xiaojin Zhu, Bryan R. Gibson, Kwang-Sung Jun, Timothy T. Rogers, Joseph Harrison, Chuck Kalish</i>	
<b>Modeling Interaction via the Principle of Maximum Causal Entropy .....</b>	<b>1255</b>
<i>Brian D. Ziebart, J. Andrew Bagnell, Anind K. Dey</i>	