

Table of Contents

To Use this Index: Scroll down or use the bookmarks in the left-hand frame to move to a new location in this index. Click on a **blue paper title** to view that paper. To return to this index after viewing a paper, click the PREVIOUS MENU bookmark in the left frame.

Preface	xv
ICML 2006 Organization	xvii
Program Committee	xviii
Local Arrangement Volunteers	xix
Board of the International Machine Learning Society 2006	xix
ICML 2006 Funding	xx
Invited Talks	xxi
Workshop and Tutorial Summaries	xxvii

ICML 2006 Papers

Using Inaccurate Models in Reinforcement Learning	1
<i>Pieter Abbeel, Morgan Quigley, Andrew Y. Ng</i>	
Algorithms for Portfolio Management Based on the Newton Method	9
<i>Amit Agarwal, Elad Hazan, Satyen Kale, Robert E. Schapire</i>	
Higher Order Learning with Graphs	17
<i>Sameer Agarwal, Kristin Branson, Serge Belongie</i>	
Ranking on Graph Data	25
<i>Shivani Agarwal</i>	
Robust Probabilistic Projections	33
<i>Cédric Archambeau, Nicolas Delannay, Michel Verleysen</i>	
A DC-Programming Algorithm for Kernel Selection	41
<i>Andreas Argyriou, Raphael Hauser, Charles A. Micchelli, Massimiliano Pontil</i>	
Relational Temporal Difference Learning	49
<i>Nima Asgharbeygi, David Stracuzzi, Pat Langley</i>	
A New Approach to Data Driven Clustering	57
<i>Arik Azran, Zoubin Ghahramani</i>	
Agnostic Active Learning	65
<i>Maria-Florina Balcan, Alina Beygelzimer, John Langford</i>	
On a Theory of Learning with Similarity Functions	73
<i>Maria-Florina Balcan, Avrim Blum</i>	
On Bayesian Bounds	81
<i>Arindam Banerjee</i>	

Convex Optimization Techniques for Fitting Sparse Gaussian Graphical Models	89
<i>Onureena Banerjee, Laurent El Ghaoui, Alexandre d'Aspremont, Georges Natsoulis</i>	
Cover Trees for Nearest Neighbor	97
<i>Alina Beygelzimer, Sham Kakade, John Langford</i>	
Graph Model Selection using Maximum Likelihood	105
<i>Ivona Bezáková, Adam Kalai, Rahul Santhanam</i>	
Dynamic Topic Models	113
<i>David M. Blei, John D. Lafferty</i>	
Predictive Search Distributions	121
<i>Edwin V. Bonilla, Christopher K.I. Williams, Felix V. Agakov, John Cavazos, John Thomson, Michael F.P. O'Boyle</i>	
Learning Predictive State Representations using Non-Blind Policies	129
<i>Michael Bowling, Peter McCracken, Michael James, James Neufeld, Dana Wilkinson</i>	
Efficient Co-Regularized Least Squares Regression	137
<i>Ulf Brefeld, Thomas Gärtner, Tobias Scheffer, Stefan Wrobel</i>	
Semi-Supervised Learning for Structured Output Variables	145
<i>Ulf Brefeld, Tobias Scheffer</i>	
Fast Nonparametric Clustering with Gaussian Blurring Mean-Shift	153
<i>Miguel Á. Carreira-Perpiñán</i>	
An Empirical Comparison of Supervised Learning Algorithms	161
<i>Rich Caruana, Alexandru Niculescu-Mizil</i>	
Robust Euclidean Embedding	169
<i>Lawrence Cayton, Sanjoy Dasgupta</i>	
Hierarchical Classification: Combining Bayes with SVM	177
<i>Nicolò Cesa-Bianchi, Claudio Gentile, Luca Zaniboni</i>	
A Continuation Method for Semi-Supervised SVMs	185
<i>Olivier Chapelle, Mingmin Chi, Alexander Zien</i>	
A Regularization Framework for Multiple-Instance Learning	193
<i>Pak-Ming Cheung, James T. Kwok</i>	
Trading Convexity for Scalability	201
<i>Ronan Collobert, Fabian Sinz, Jason Weston, Léon Bottou</i>	
Learning Algorithms for Online Principal-Agent Problems (and Selling Goods Online)	209
<i>Vincent Conitzer, Nikesh Garera</i>	
Dealing with Non-Stationary Environments using Context Detection	217
<i>Bruno C. da Silva, Eduardo W. Basso, Ana L. C. Bazzan, Paulo M. Engel</i>	

Locally Adaptive Classification Piloted by Uncertainty	225
<i>Juan Dai, Shuicheng Yan, Xiaoou Tang, James T. Kwok</i>	
The Relationship Between Precision-Recall and ROC Curves	233
<i>Jesse Davis, Mark Goadrich</i>	
Discriminative Cluster Analysis	241
<i>Fernando De la Torre, Takeo Kanade</i>	
Collaborative Prediction using Ensembles of Maximum Margin Matrix Factorizations	249
<i>Dennis DeCoste</i>	
Learning the Structure of Factored Markov Decision Processes in Reinforcement Learning Problems	257
<i>Thomas Degris, Olivier Sigaud, Pierre-Henri Wuillemin</i>	
Efficient Learning of Naive Bayes Classifiers under Class-Conditional Classification Noise	265
<i>Francois Denis, Christophe Nicolas Magnan, Liva Ralaivola</i>	
Learning User Preferences for Sets of Objects	273
<i>Marie desJardins, Eric Eaton, Kiri L. Wagstaff</i>	
R_1-PCA: Rotational Invariant L1-norm Principal Component Analysis for Robust Subspace Factorization	281
<i>Chris Ding, Ding Zhou, Xiaofeng He, Hongyuan Zha</i>	
Clustering Documents with an Exponential-Family Approximation of the Dirichlet Compound Multinomial Distribution	289
<i>Charles Elkan</i>	
A Graphical Model for Predicting Protein Molecular Function	297
<i>Barbara Engelhardt, Michael Jordan, Steven Brenner</i>	
Qualitative Reinforcement Learning	305
<i>Arkady Epshteyn, Gerald DeJong</i>	
Online Multiclass Learning by Interclass Hypothesis Sharing	313
<i>Michael Fink, Shai Shalev-Shwartz, Yoram Singer, Shimon Ullman</i>	
Regression with the Optimised Combination Technique	321
<i>Jochen Garcke</i>	
A Note on Mixtures of Experts for Multiclass Responses: Approximation Rate and Consistent Bayesian Inference	329
<i>Yang Ge, Wenxin Jiang</i>	
The Rate Adapting Poisson Model for Information Retrieval and Object Recognition	337
<i>Peter V. Gehler, Alex D. Holub, Max Welling</i>	
Kernelizing the Output of Tree-Based Methods	345
<i>Pierre Geurts, Louis Wehenkel, Florence d'Alché-Buc</i>	

Nightmare at Test Time: Robust Learning by Feature Deletion	353
<i>Amir Globerson, Sam Roweis</i>	
A Choice Model with Infinitely Many Latent Features	361
<i>Dilan Görür, Frank Jäkel, Carl Edward Rasmussen</i>	
Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks	369
<i>Alex Graves, Santiago Fernández, Faustino Gomez, Juergen Schmidhuber</i>	
Practical Solutions to the Problem of Diagonal Dominance in Kernel Document Clustering	377
<i>Derek Greene, Pádraig Cunningham</i>	
Fast Transpose Methods for Kernel Learning on Sparse Data	385
<i>Patrick Haffner</i>	
An Analysis of Graph Cut Size for Transductive Learning	393
<i>Steve Hanneke</i>	
Learning a Kernel Function for Classification with Small Training Samples	401
<i>Tomer Hertz, Aharon Bar Hillel, Daphna Weinshall</i>	
Looping Suffix Tree-Based Inference of Partially Observable Hidden State	409
<i>Michael Holmes, Charles Lee Isbell, Jr.</i>	
Batch Mode Active Learning and Its Application to Medical Image Classification	417
<i>Steven C. H. Hoi, Rong Jin, Jianke Zhu, Michael R. Lyu</i>	
Ranking Individuals by Group Comparisons	425
<i>Tzu-Kuo Huang, Chih-Jen Lin, Ruby C. Weng</i>	
Hidden Process Models	433
<i>Rebecca A. Hutchinson, Tom M. Mitchell, Indrayana Rustandi</i>	
Estimating Relatedness via Data Compression	441
<i>Brendan Juba</i>	
Automatic Basis Function Construction for Approximate Dynamic Programming and Reinforcement Learning	449
<i>Philipp W. Keller, Shie Mannor, Doina Precup</i>	
Personalized Handwriting Recognition via Biased Regularization	457
<i>Wolf Kienzle, Kumar Chellapilla</i>	
Optimal Kernel Selection in Kernel Fisher Discriminant Analysis	465
<i>Seung-Jean Kim, Alessandro Magnani, Stephen Boyd</i>	
Pareto Optimal Linear Classification	473
<i>Seung-Jean Kim, Alessandro Magnani, Sikandar Samar, Stephen Boyd, Johan Lim</i>	
Fast Particle Smoothing: If I Had a Million Particles	481
<i>Mike Klaas, Mark Briers, Nando de Freitas, Arnaud Doucet, Simon Maskell, Dustin Lang</i>	

Autonomous Shaping: Knowledge Transfer in Reinforcement Learning	489
<i>George Konidaris, Andrew Barto</i>	
Data Association for Topic Intensity Tracking	497
<i>Andreas Krause, Jure Leskovec, Carlos Guestrin</i>	
Learning Low-Rank Kernel Matrices	505
<i>Brian Kulis, Mátyás Sustik, Inderjit Dhillon</i>	
Local Distance Preservation in the GP-LVM through Back Constraints	513
<i>Neil D. Lawrence, Joaquin Quiñero-Candela</i>	
Simpler Knowledge-based Support Vector Machines	521
<i>Quoc V. Le, Alex J. Smola, Thomas Gärtner</i>	
Using Query-Specific Variance Estimates to Combine Bayesian Classifiers	529
<i>Chi-Hoon Lee, Russ Greiner, Shaojun Wang</i>	
A Probabilistic Model for Text Kernels	537
<i>Alain Lehmann, John Shawe-Taylor</i>	
Efficient MAP Approximation for Dense Energy Functions	545
<i>Marius Lordeanu, Martial Hebert</i>	
Nonstationary Kernel Combination	553
<i>Darrin P. Lewis, Tony Jebara, William Stafford Noble</i>	
Region-Based Value Iteration for Partially Observable Markov Decision Processes	561
<i>Hui Li, Xuejun Liao, Lawrence Carin</i>	
Multiclass Boosting with Repartitioning	569
<i>Ling Li</i>	
Pachinko Allocation: DAG-Structured Mixture Models of Topic Correlations	577
<i>Wei Li, Andrew McCallum</i>	
Spectral Clustering for Multi-type Relational Data	585
<i>Bo Long, Mark Zhang, Xiaoyun Wu, Philip S. Yu</i>	
Combined Central and Subspace Clustering for Computer Vision Applications	593
<i>Le Lu, René Vidal</i>	
Fast Direct Policy Evaluation using Multiscale Analysis of Markov Diffusion Processes	601
<i>Mauro Maggioni, Sridhar Mahadevan</i>	
Pruning in Ordered Bagging Ensembles	609
<i>Gonzalo Martínez-Muñoz, Alberto Suárez</i>	
Learning High-Order MRF Priors of Color Images	617
<i>Julian J. McAuley, Tibério S. Caetano, Alex J. Smola, Matthias O. Franz</i>	

The Uniqueness of a Good Optimum for K-means	625
<i>Marina Meila</i>	
Kernel Information Embeddings	633
<i>Roland Memisevic</i>	
Generalized Spectral Bounds for Sparse LDA	641
<i>Baback Moghaddam, Yair Weiss, Shai Avidan</i>	
Learning to Impersonate	649
<i>Moni Naor, Guy N. Rothblum</i>	
Online Decoding of Markov Models under Latency Constraints	657
<i>Mukund Narasimhan, Paul Viola, Michael Shilman</i>	
Learning Hierarchical Task Networks by Observation	665
<i>Negin Nejati, Pat Langley, Tolga Konik</i>	
Reinforcement Learning for Optimized Trade Execution	673
<i>Yuriy Nevmyvaka, Yi Feng, Michael Kearns</i>	
Concept Boundary Detection for Speeding up SVMs	681
<i>Navneet Panda, Edward Y. Chang, Gang Wu</i>	
The Support Vector Decomposition Machine	689
<i>Francisco Pereira, Geoffrey Gordon</i>	
An Analytic Solution to Discrete Bayesian Reinforcement Learning	697
<i>Pascal Poupart, Nikos Vlassis, Jesse Hoey, Kevin Regan</i>	
MISSL: Multiple-Instance Semi-Supervised Learning	705
<i>Rouhollah Rahmani, Sally A. Goldman</i>	
Constructing Informative Priors using Transfer Learning	713
<i>Rajat Raina, Andrew Y. Ng, Daphne Koller</i>	
CN=CPCN	721
<i>Liva Ralaivola, François Denis, Christophe N. Magnan</i>	
Maximum Margin Planning	729
<i>Nathan D. Ratliff, J. Andrew Bagnell, Martin A. Zinkevich</i>	
Quadratic Programming Relaxations for Metric Labeling and Markov Random Field MAP Estimation	737
<i>Pradeep Ravikumar, John Lafferty</i>	
Categorization in Multiple Category Systems	745
<i>Jean-Michel Renders, Eric Gaussier, Cyril Goutte, Francois Pacull, Gabriela Csurka</i>	
How Boosting the Margin Can Also Boost Classifier Complexity	753
<i>Lev Reyzin, Robert E. Schapire</i>	

Combining Discriminative Features to Infer Complex Trajectories	761
<i>David A. Ross, Simon Osindero, Richard S. Zemel</i>	
Sequential Update of ADtrees	769
<i>Josep Roure, Andrew W. Moore</i>	
Predictive Linear-Gaussian Models of Controlled Stochastic Dynamical Systems	777
<i>Matthew Rudary, Satinder Singh</i>	
A Statistical Approach to Rule Learning	785
<i>Ulrich Rückert, Stefan Kramer</i>	
Efficient Inference on Sequence Segmentation Models	793
<i>Sunita Sarawagi</i>	
Cost-Sensitive Learning with Conditional Markov Networks	801
<i>Prithviraj Sen, Lise Getoor</i>	
Feature Value Acquisition in Testing: A Sequential Batch Test Algorithm	809
<i>Victor S. Sheng, Charles X. Ling</i>	
Permutation Invariant SVMs	817
<i>Pannagadatta K. Shivaswamy, Tony Jebara</i>	
Bayesian Learning of Measurement and Structural Models	825
<i>Ricardo Silva, Richard Scheines</i>	
An Intrinsic Reward Mechanism for Efficient Exploration	833
<i>Özgür Simsek, Andrew G. Barto</i>	
Deterministic Annealing for Semi-supervised Kernel Machines	841
<i>Vikas Sindhwani, Sathya Keerthi, Olivier Chapelle</i>	
Feature Subset Selection Bias for Classification Learning	849
<i>Surendra K. Singhi, Huan Liu</i>	
Classifying EEG for Brain-Computer Interfaces: Learning Optimal Filters for Dynamical System Features	857
<i>Le Song, Julien Epps</i>	
An Investigation of Computational and Informational Limits in Gaussian Mixture Clustering	865
<i>Nathan Srebro, Gregory Shakhnarovich, Sam Roweis</i>	
Bayesian Pattern Ranking for Move Prediction in the Game of Go	873
<i>David Stern, Ralf Herbrich, Thore Graepel</i>	
PAC Model-free Reinforcement Learning	881
<i>Alexander L. Strehl, Lihong Li, Eric Wiewiora, John Langford, Michael L. Littman</i>	
Experience-Efficient Learning in Associative Bandit Problems	889
<i>Alexander L. Strehl, Chris Mesterharm, Michael L. Littman, Haym Hirsh</i>	

Full Bayesian Network Classifiers	897
<i>Jiang Su, Harry Zhang</i>	
Local Fisher Discriminant Analysis for Supervised Dimensionality Reduction	905
<i>Masashi Sugiyama</i>	
Iterative RELIEF for Feature Weighting	913
<i>Yijun Sun, Jian Li</i>	
Multiclass Reduced-Set Support Vector Machines	921
<i>Benyang Tang, Dominic Mazzoni</i>	
Fast and Space Efficient String Kernels using Suffix Arrays	929
<i>Choon Hui Teo, S.V.N. Vishwanathan</i>	
Bayesian Regression with Input Noise for High Dimensional Data	937
<i>Jo-Anne Ting, Aaron D'Souza, Stefan Schaal</i>	
Probabilistic Inference for Solving Discrete and Continuous State Markov Decision Processes	945
<i>Marc Toussaint, Amos Storkey</i>	
Clustering Graphs by Weighted Substructure Mining	953
<i>Koji Tsuda, Taku Kudo</i>	
Active Sampling for Detecting Irrelevant Features	961
<i>Sriharsha Veeramachaneni, Emanuele Olivetti, Paolo Avesani</i>	
Accelerated Training of Conditional Random Fields with Stochastic Gradient Methods	969
<i>S.V.N. Vishwanathan, Nicol N. Schraudolph, Mark W. Schmidt, Kevin P. Murphy</i>	
Topic Modeling: Beyond Bag-of-Words	977
<i>Hanna M. Wallach</i>	
Label Propagation through Linear Neighborhoods	985
<i>Fei Wang, Changshui Zhang</i>	
Two-Dimensional Solution Path for Support Vector Regression	993
<i>Gang Wang, Dit-Yan Yeung, Frederick H. Lochovsky</i>	
Totally Corrective Boosting Algorithms that Maximize the Margin	1001
<i>Manfred K. Warmuth, Jun Liao, Gunnar Rätsch</i>	
Inference with the Universum	1009
<i>Jason Weston, Ronan Collobert, Fabian Sinz, Léon Bottou, Vladimir Vapnik</i>	
Kernel Predictive Linear Gaussian Models for Nonlinear Stochastic Dynamical Systems	1017
<i>David Wingate, Satinder Singh</i>	
Predictive State Representations with Options	1025
<i>Britton Wolfe, Satinder Singh</i>	

Fast Time Series Classification using Numerosity Reduction	1033
<i>Xiaopeng Xi, Eamonn Keogh, Christian Shelton, Li Wei, Chotirat Ann Ratanamahatana</i>	
A Duality View of Spectral Methods for Dimensionality Reduction	1041
<i>Lin Xiao, Jun Sun, Stephen Boyd</i>	
Bayesian Multi-Population Haplotype Inference via a Hierarchical Dirichlet Process Mixture	1049
<i>Eric P. Xing, Kyung-Ah Sohn, Michael I. Jordan, Yee-Whye Teh</i>	
Discriminative Unsupervised Learning of Structured Predictors	1057
<i>Linli Xu, Dana Wilkinson, Finnegan Southey, Dale Schuurmans</i>	
Semi-Supervised Nonlinear Dimensionality Reduction	1065
<i>Xin Yang, Haoying Fu, Hongyuan Zha, Jesse Barlow</i>	
Null Space versus Orthogonal Linear Discriminant Analysis	1073
<i>Jieping Ye, Tao Xiong</i>	
Active Learning via Transductive Experimental Design	1081
<i>Kai Yu, Jinbo Bi, Volker Tresp</i>	
Collaborative Ordinal Regression	1089
<i>Shipeng Yu, Kai Yu, Volker Tresp, Hans-Peter Kriegel</i>	
Block-Quantized Kernel Matrix for Fast Spectral Embedding	1097
<i>Kai Zhang, James Kwok</i>	
Statistical Debugging: Simultaneous Identification of Multiple Bugs	1105
<i>Alice X. Zheng, Michael I. Jordan, Ben Liblit, Mayur Naik, Alex Aiken</i>	
Efficient Lazy Elimination for Averaged One-Dependence Estimators	1113
<i>Fei Zheng, Geoffrey I. Webb</i>	