

ICML 2005  
Tentative Program

\*Note: all papers will be presented both orally and in the evening poster sessions. Talks are allocated 20 minutes for presentation and 5 minutes for questions.

Monday August, 8

8:40	Opening			
9:00	Invited Talk Michael Jordan			
10:00	Best Papers			
10:30	<i>Coffee Break</i>			
11:00-12:40	Session 1	Session 2	Session 3	Session 4
	Probabilistic Approaches 1	Reinforcement Learning 1	Decision Tree Learning	Dimensionality Reduction
	Harmonic mixtures: combining mixture models and graph-based methods for inductive and scalable semi-supervised learning <i>Xiaojin Zhu, John Lafferty</i>	A Causal Approach to Hierarchical Decomposition of Factored MDPs <i>Anders Jonsson, Andrew Barto</i>	Tempering for Bayesian C&T <i>Nicos Angelopoulos, James Cussens</i>	Supervised dimensionality reduction using mixture models <i>Sajama, Alon Orlitsky</i>
	Discriminative versus Generative Parameter and Structure Learning of Bayesian Network Classifiers <i>Franz Pernkopf, Jeff Bilmes</i>	Identifying Useful Subgoals in Reinforcement Learning by Local Graph Partitioning <i>Ozgur Simsek, Alicia Wolfe, Andrew Barto</i>	Why Skewing Works: Learning Difficult Boolean Functions with Greedy Tree Learners <i>Bernard Rosell, Lisa Hellerstein, Soumya Ray, David Page</i>	Statistical and Computational Analysis of Locality Preserving Projection <i>Xiaofei He, Deng Cai, Wanli Min</i>
	Learning Class-Discriminative Dynamic Bayesian Networks <i>John Burge, Terran Lane</i>	Learning Predictive State Representations in Dynamical Systems Without Reset <i>Britton Wolfe, Michael R. James, Satinder Singh</i>	Closed-form dual perturb and combine for tree-based models <i>Pierre Geurts, Louis Wehenkel</i>	Large Margin Non-Linear Embedding <i>Alexander Zien, Joaquín Quinonero-Candela</i>
	Efficient discriminative learning of Bayesian network classifier via Boosted Augmented Naive Bayes <i>Yushi Jing, Vladimir Pavlovic, James M. Rehg</i>	Learning Predictive Representations from a History <i>Eric Wiewiora</i>	Generalized Skewing for Functions with Continuous and Nominal Attributes <i>Soumya Ray, David Page</i>	Independent Subspace Analysis Using Geodesic Spanning Trees <i>Barnabas Poczos, Andras Lorincz</i>
12:40-14:14	<i>Lunch</i>			
14:15-15:55	Session 5	Session 6	Session 7	Session 8
	Probabilistic Approaches 2	Reinforcement Learning 2	Learning and Bioinformatics	Gaussian Processes
	Computational Aspects of Bayesian Partition Models <i>Mikko Koivisto, Kismat Sood</i>	Exploration and Apprenticeship Learning in Reinforcement Learning <i>Pieter Abbeel, Andrew Y. Ng</i>	New Kernels for Protein Structural Motif Discovery and Function Classification <i>Chanq Wang, Stephen Scott</i>	Learning Gaussian Processes from Multiple Tasks <i>Kai Yu, Volker Tresp, Anton Schwaighofer</i>
	Hierarchical Bayesian Models for Kernel Learning <i>Mark Girolami, Simon Rogers</i>	Bayesian Sparse Sampling for On-line Reward Optimization <i>Tao Wang, Daniel Lizotte, Michael Bowling, Dale Schuurmans</i>	Large Scale Genomic Sequence SVM Classifiers <i>Soren Sonnenburg, Gunnar Rätsch, Bernhard Schölkopf</i>	Preference Learning with Gaussian Processes <i>Wei Chu, Zoubin Ghahramani</i>
	Expectation Maximization Algorithms for Conditional Likelihoods <i>Jarkko Salojärvi, Kai Puolamäki, Samuel Kaski</i>	Bounded Real-Time Dynamic Programming: RTDP with monotone upper bounds and performance guarantees <i>H. Brendan McMahan, Maxim Likhachev, Geoffrey J. Gordon</i>	Predicting Protein Folds with Structural Repeats Using a Chain Graph Model <i>Yan Liu, Eric Xing, Jaime Carbonell</i>	Heteroscedastic Gaussian Process Regression <i>Quoc V. Le, Alex J. Smola, Stephane Canu</i>
	Compact approximations to Bayesian predictive distributions <i>Edward Snelson, Zoubin Ghahramani</i>	Finite Time Bounds for Sampling Based Fitted Value Iteration <i>Csaba Szepesvari, Remi Munos</i>	Multi-Class protein fold detection using adaptive codes <i>Eugene Ie, Jason Weston, William Stafford Noble, Christina Leslie</i>	Near-Optimal Sensor Placements in Gaussian Processes <i>Carlos Guestrin, Andreas Krause, Ajit Pauk Singh</i>
15:55-16:25	<i>Coffee Break</i>			

## Monday August, 8

16:25-18:10

### Session 9

#### Graphed Based Data

Learning from Labeled and Unlabeled Data on a Directed Graph  
*Dengyong Zhou, Jiayuan Huang, Bernhard Schölkopf*

Semi-supervised Graph Clustering: A Kernel Approach  
*Brian Kulis, Sugato Basu, Inderjit Dhillon, Raymond Mooney*

Online Learning over Graphs  
*Mark Herbster, Massimiliano Pontil, Lisa Wainer*

Optimal Assignment Kernels For Attributed Molecular Graphs  
*Holger Fröhlich, Jörg Wegner, Florian Sieker, Andreas Zell*

### Session 10

#### Ensemble Methods 1

Ensembles of Biased Classifiers  
*Rinat Khossainov, Andreas Hess, Nicholas Kushmerick*

Experimental Comparison between Bagging and Monte Carlo Ensemble Classification  
*Roberto Esposito, Lorenza Saitta*

A Practical Generalization of Fourier-based Learning  
*Adam Drake, Dan Ventura*

Using Additive Expert Ensembles to Cope with Concept Drift  
*Jeremy Kolter, Marcus Maloof*

### Session 11

#### Applications 1

A Graphical Model for Chord Progressions Embedded in a Psychoacoustic Space  
*Jean-Francois Paiement, Douglas Eck, Samy Bengio, David Barber*

Predicting Probability Distributions for Surf Height Using an Ensemble of Mixture Density Networks  
*Michael Carney, Padraig Cunningham, Jim Dowling, Ciaran Lee*

A Brain Computer Interface with Online Feedback based on Magnetoencephalography  
*Thomas Navin Lal, Michael Schroeder, N. Jeremy Hill, Hubert Preissl, Thilo Hinterberger, Juergen Mellinger, Martin Bogdan, Wolfgang Rosenstiel, Niels Birbaumer, Bernhard Schoelkopf*

Learning Strategies for Story Comprehension: A Reinforcement Learning Approach  
*Eugene Grois, David C. Wilkins*

### Session 12

#### Learning and Vision 1

Non-Negative Tensor Factorization with Applications to Statistics and Computer Vision  
*Amnon Shashua, Tamir Hazan*

Q-Learning of Sequential Attention for Visual Object Recognition from Informative Local Descriptors  
*Lucas Paletta, Gerald Fritz, Christin Seifert*

Object Correspondence as a Machine Learning Problem  
*Bernhard Schölkopf, Florian Steinke, Volker Blanz*

Interactive Learning of Mappings from Visual Percepts to Actions  
*Sébastien Jodogne, Justus H. Piater*

18:30

*City Hall Reception*

## Tuesday August, 9

9:00

10:00

10:30

11:00-12:40

Invited Talk Johannes Gehrke

Best Papers

*Coffee Break*

	Session 13	Session 14	Session 15	Session 16
	Kernel Methods and SVMs (1)	Graphical Models for Text, Language and Wen	Reinforcement Learning 3	Feature Selection and Dimensionality Reduction
	Explanation-Augmented SVM: an Approach to Incorporating Domain Knowledge into SVM Learning <i>Qiang Sun, Gerald DeJong</i>	2D Conditional Random Fields for Web Information Extraction <i>Jun Zhu, Zaiqing Nie, Ji-Rong Wen, Bo Zhang, Wei-Ying Ma</i>	Reinforcement learning with Gaussian processes <i>Yaakov Engel, Shie Mannor, Ron Meir</i>	Action Respecting Embedding <i>Michael Bowling, Ali Ghodsi, Dana Wilkinson</i>
	New Approaches to Support Vector Ordinal Regression <i>Wei Chu, S. Sathiya Keerthi</i>	Exploiting Syntactic, Semantic and Lexical Regularities in Language Modeling via Directed Markov Random Fields <i>Shaojun Wang, Shaomin Wang, Russell Greiner, Dale Schuurmans, Li Chena</i>	Proto-Value Functions: Developmental Reinforcement Learning <i>Sridhar Mahadevan</i>	Multimodal Oriented Discriminant Analysis <i>Fernando De la Torre, Takeo Kanade</i>
	Predictive low-rank decomposition for kernel methods <i>Francis R. Bach, Michael I. Jordan</i>	Integer Linear Programming Inference for Conditional Random Fields <i>Dan Roth, Wen-tau Yih</i>	TD(lambda) Networks: Temporal-Difference Networks with Eligibility Traces <i>Brian Tanner, Richard Sutton</i>	Generalized LARS as an Effective Feature Selection Tool for Text Classification With SVMs <i>S. Sathiya Keerthi</i>
	The cross entropy method for classification <i>Shie Mannor, Dori Peleg, Reuven Rubinfeld</i>	Active Learning for Hidden Markov Models: Objective Functions and Algorithms <i>Brigham Anderson, Andrew Moore</i>	Learning to Compete, Compromise, and Cooperate in Repeated General-Sum Games <i>Jacob W. Crandall, Michael A. Goodrich</i>	Online Feature Selection for Pixel Classification <i>Karen Glocer, Damian Eads, James Theiler</i>

12:40-14:14

14:15-15:55

*Lunch*

	Session 17	Session 18	Session 19	Session 20
	Learning from Structured Data	Theory	Probability Estimation and Ranking	Scientific Discovery, Meta-Learning and Instance Based Learning
	Weighted Decomposition Kernels <i>Sauro Menchetti, Fabrizio Costa, Paolo Frasconi</i>	PAC-Bayes Risk Bounds for Sample-Compressed Gibbs Classifiers <i>François Laviolette, Mario Marchand</i>	Naive Bayes Models for Probability Estimation <i>Daniel Lowd, Pedro Domingos</i>	New D-Separation Identification Results for Learning Continuous Latent Variable Models <i>Ricardo Silva, Richard Scheines</i>
	Learning Structured Prediction Models: A Large Margin Approach <i>Ben Taskar, Vassil Chatalbashev, Daphne Koller, Carlos Guestrin</i>	Error Bounds for Correlation Clustering <i>Thorsten Joachims, John Hopcroft</i>	Predicting Good Probabilities With Supervised Learning <i>Alexandru Niculescu-Mizil, Rich Caruana</i>	Reducing Overfitting in Process Model Induction <i>Will Bridewell, Narges Asani, Pat Langley, Ljupco Todorovski</i>
	Learning as Search Optimization: Approximate Large Margin Methods for Structured Prediction <i>Hal Daume III, Daniel Marcu</i>	A Comparison of Tight Generalization Error Bounds <i>Matti Käariäinen, John Langford</i>	Augmenting Naive Bayes for Ranking <i>Harry Zhang, Liangxiao Jiang, Jiang Su</i>	Predicting Relative Performance of Classifiers from Samples <i>Rui Leite, Pavel Brazdil</i>
	Propagating Distributions on a Hypergraph by Dual Information Regularization <i>Koji Tsuda</i>	Error Limiting Reductions Between Classification Tasks <i>Alina Beygelzimer, Varsha Dani, Tom Hayes, John Langford, Bianca Zadrozny</i>	Learning to Rank using Gradient Descent <i>Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, Greg Hullender</i>	Fast Condensed Nearest Neighbor Rule <i>Fabrizio Angiulli</i>

15:55-16:25

*Coffee Break*

## Tuesday August, 9

16:25-17.15

### Session 21

#### Multi-Instance Learning

Multi-Instance Tree Learning

*Hendrik Blockeel, David Page, Ashwin Srinivasan*

### Session 22

#### ROC Analysis

Optimizing Abstaining Classifiers using ROC  
Analysis

*Tadeusz Pietraszekry*

### Session 23

#### Logistic Regressions

Incomplete-Data Classification using Logistic  
Regression

*David Williams, Xuejun Liao, Ya Xue, Lawrence  
Carin*

### Session 24

#### Learning and Vision 2

Linear Asymmetric Classifier for Cascade Detectors

*Jianxin Wu, Matthew D. Mullin, James M. Rehg*

Supervised versus Multiple Instance Learning: An  
Empirical Comparison

*Soumya Ray, Mark Craven*

ROC Confidence Bands : An Empirical Evaluation

*Sofus A. Macskassy, Foster Provost, Saharon  
Rosset*

Logistic Regression with an Auxiliary Data Source

*Xuejun Liao, Ya Xue, Lawrence Carin*

Variational Bayesian Image Modelling

*Li Cheng, Feng Jiao, Dale Schuurmans, Shaojun  
Wang*

18:00

***Boat Trip Departure***

## Wednesday August, 10

	Session 25	Session 26	Session 27	Session 28
9:00-10:40	Kernel Methods and SVMs 2	Reinforcement and Agent Learning	Clustering	Text Classification and Information Extraction
	Implicit Surface Modelling as an Eigenvalue Problem <i>Christian Walder, Olivier Chapelle, Bernhard Schölkopf</i>	Hedged learning: Regret minimization with learning experts <i>Yu-Han Chang, Leslie Kaelbling</i>	Robust One-Class Clustering Using Hybrid Global and Local Search <i>Gunjan Gupta, Joydeep Ghosh</i>	Learning Hierarchical Multi-Category Text Classification Models <i>Juho Rousu, Craig Saunders, Sandor Szedmak, John Shawe-Taylor</i>
	Building Sparse Large Margin Classifiers <i>Mingrui Wu, Bernhard Schoelkopf, Goekhan Bakir</i>	Dynamic Preferences in Multi-Criteria Reinforcement Learning <i>Sriraam Natarajan, Prasad Tadepalli</i>	Multi-Way Distributional Clustering via Pairwise Interactions <i>Ron Bekkerman, Ran El-Yaniv, Andrew McCallum</i>	A Model for Handling Approximate, Noisy or Incomplete Labeling in Text Classification <i>Ganesh Ramakrishnan, Krishna Prasad Chitrapura, Raghu Krishnapuram, Pushpak Bhattacharyya</i>
	Adapting Two-Class Classification Methods to Many Class Problems <i>Simon I. Hill, Arnaud Doucet</i>	Learning Approximate Preconditions for Methods in Hierarchical Plans <i>Okhtay Ilghami, Hector Munoz-Avila, Dana S. Nau, David W. Aha</i>	Comparing Clusterings - An Axiomatic View <i>Marina Meila</i>	Evaluating Machine Learning for Information Extraction <i>Neil Ireson, Fabio Ciravegna, Mary Elaine Califf, Dayne Freitag, Nicholas Kushmerick, Alberto Lavelli</i>
	An Efficient Method for Simplifying Support Vector Machines <i>DucDung Nguyen, TuBao Ho</i>	Recycling Data for Multi-Agent Learning <i>Santi Ontanon, Enric Plaza</i>	Bayesian Hierarchical Clustering <i>Katherine Heller, Zubin Ghahramani</i>	Learn to Weight Terms in Information Retrieval Using Category Information <i>Ronq Jin, Joyce Y. Chai, Luo Si</i>
10:40-11:10		<i>Coffee Break</i>		
11:10-12:25	Session 29	Session 30	Session 31	Session 32
	Kernel Methods and SVMs 3	Dirichlet Models	Learning in Bioinformatics and Robotics	Distances and Measures
	Healing the Relevance Vector Machine through Augmentation <i>Carl Edward Rasmussen, Joaquin Quinero-Candela</i>	Dirichlet Enhanced Relational Learning <i>Zhao Xu, Volker Tresp, Kai Yu, Shipeng Yu, Hans-Peter Kriegel</i>	Learning Discontinuities with Products-of-Sigmoids for Switching between Local Models <i>Marc Toussaint, Sethu Vijayakumar</i>	A New Mallows Distance Based Metric For Comparing Clusterings <i>Ding Zhou, Jia Li, Hongyuan Zha</i>
	Supervised Clustering with Support Vector Machines <i>Thomas Finley, Thorsten Joachims</i>	Modeling Word Burstiness Using the Dirichlet Distribution <i>Rasmus Madsen, David Kauchak, Charles Elkan</i>	Active Learning for Sampling in Time-Series Experiments With Application to Gene Expression Analysis <i>Rohit Singh, Nathan Palmer, David Gifford, Bonnie Berger, Ziv Bar-Joseph</i>	Estimating and computing density based distance metrics <i>Sajama, Alon Orlitsky</i>
	A Support Vector Method for Multivariate Performance Measures <i>Thorsten Joachims</i>	Hierarchical Dirichlet Model for Document Classification <i>Sriharsha Veeramachaneni, Diego Sona, Paolo Avesani</i>	Unsupervised Evidence Integration <i>Philip M. Long, Vinay Varadan, Sarah Gilman, Mark Treshock, Rocco A. Servedioy</i>	A Martingale Framework for Concept Change Detection in Time-Varying Data Streams <i>Shen-Shyang Ho</i>
12:40-14:14		<i>Lunch</i>		
14:15-15:15		Invited Talk Gerhard Widmer		
15:15-15:45		<i>Coffee Break</i>		

## Wednesday August, 10

15:45-17:00	Session 33	Session 34	Session 35	Session 36
	Kernel Methods and SVMs	Learning and Robotics	Manifolds and Dimensionality Reduction	Inductive Logic Programming Track
	A General Regression Technique for Learning Transductions <i>Corinna Cortes, Mehryar Mohri, Jason Weston</i>	Coarticulation: An Approach for Generating Concurrent Plans in Markov Decision Processes <i>Khashayar Rohanimanesh, Sridhar Mahadevan</i>	Analysis and Extension of Spectral Methods for Nonlinear Dimensionality Reduction <i>Fei Sha, Lawrence K. Saul</i>	(ICML paper) Learning First-Order Probabilistic Models with Combining Rules <i>Sriram Natarajan, Prasad Tadepalli, Eric Altendorf, Thomas G. Dietterich, Alan Fern, Angelo Restifcary</i>
	Beyond the Point Cloud: from Transductive to Semi-supervised Learning <i>Vikas Sindhwani, Partha Niyogi, Mikhail Belkin</i>	Recognition and Reproduction of Gestures using a Probabilistic Framework combining PCA, ICA and HMM <i>Sylvain Calinon, Aude Billard</i>	Clustering Through Ranking On Manifolds <i>Markus Breitenbach, Gregory Z. Grudic</i>	(ILP Paper) Logical Bayesian Networks and Their Relation to Other Probabilistic Logical Models <i>Daan Fierens, Hendrik Blockeel, Maurice Bruynooghe, Jan Ramon</i>
	Fast Maximum Margin Matrix Factorization for Collaborative Prediction <i>Jason D. M. Rennie, Nati Srebro - University of Toronto</i>	High Speed Obstacle Avoidance using Monocular Vision and Reinforcement learning <i>Jeff Michels, Ashutosh Saxena, Andrew Y. Ng</i>	Intrinsic Dimensionality Estimation of Submanifolds in $\mathbb{R}^d$ <i>Matthias Hein, Jean-Yves Audibert</i>	(ICML paper) Learning the Structure of Markov Logic Networks <i>Stanley Kok, Pedro Domingos</i>
17:00-17:15	<i>Coffee Break</i>			
17:15-18:05	Session 37	Session 38	Session 39	Session 40
	Ensemble Methods and Output Codes	Reinforcement Learning - Theory	Scalibility	Statistical Relational Learning and ILP
	A Smoothed Boosting Algorithm Using Probabilistic Output Codes <i>Rong Jin, Jian Zhang</i>	A Theoretical Analysis of Model-Based Interval Estimation <i>Alexander L. Strehl, Michael L. Littman</i>	Fast Inference and Learning in Large-State-Space HMMs <i>Sajid M. Siddiqi, Andrew W. Moore</i>	(ILP Paper) Probabilistic First-Order Theory Revision from Examples <i>Aline Paes, Kate Revoredo, Gerson Zaverucha, Vitor Santos Costa</i>
	Unifying the Error-Correcting and Output-Code AdaBoost within the Margin Framework <i>Yijun Sun, Sinisa Todorovic, Jian Li, Dapeng Wu</i>	Relating Reinforcement Learning Performance to Classification Performance <i>John Langford, Bianca Zadrozny</i>	Core Vector Regression for Very Large Regression Problems <i>Ivor W. Tsang, James T. Kwok, Kimo T. Lai</i>	(ICML Paper) Combining Model-Based and Instance-Based Learning for First Order Regression <i>Kurt Driessens, Saso Dzeroski</i>
18:05-19:00	<b>Business Meeting</b>			