

Contents

Foreword.....	vii
----------------------	------------

Invited Presentations

Peter Grunwald	1
<i>The Catch-Up Phenomenon in Bayesian Inference</i>	
Robin Hanson	3
<i>Combinatorial Prediction Markets</i>	
Dan Klein	5
<i>Unsupervised Learning for Natural Language Processing</i>	
Gabor Lugosi.....	7
<i>Concentration Inequalities</i>	

Unsupervised, Semi-Supervised and Active Learning

Kamalika Chaudhuri and Satish Rao	9
<i>Learning Mixtures of Product Distributions Using Correlations and Independence</i>	
Kamalika Chaudhuri and Satish Rao	21
<i>Beyond Gaussians: Spectral Methods for Learning Mixtures of Heavy-Tailed Distributions</i>	
Shai Ben-David, Tyler Lu and David Pal	33
<i>Does Unlabeled Data Provably Help? Worst-case Analysis of the Sample Complexity of Semi-Supervised Learning</i>	
Maria-Florina Balcan, Steve Hanneke and Jennifer Wortman	45
<i>The True Sample Complexity of Active Learning</i>	

On-Line Learning

Elad Hazan and Satyen Kale	57
<i>Extracting Certainty from Uncertainty: Regret Bounded by Variation in Costs</i>	

Kosuke Ishibashi, Kohei Hatano and Masayuki Takeda..... 69
Online Learning of Maximum p -Norm Margin Classifiers with Bias

Subhash Khot and Ashok Kumar Ponnuswami 81
Minimizing Wide Range Regret with Time Selection Functions

Other Directions

Nir Ailon and Mehryar Mohri..... 87
An Efficient Reduction of Ranking to Classification

Michael Kearns and Jennifer Wortman 99
Learning from Collective Behavior

Bharath Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Gert Lanckriet and Bernhard Schölkopf..... 111
Injective Hilbert Space Embeddings of Probability Measures

Complexity and Boolean Functions

Sung-Soon Choi, Kyomin Jung and Jeong Han Kim 123
Almost Tight Upper Bound for Finding Fourier Coefficients of Bounded Pseudo-Boolean Functions

Robert Holte, Steffen Lange, Sandra Zilles and Martin Zinkevich..... 135
Teaching Dimensions based on Cooperative Learning

Vitaly Feldman 147
On the Power of Membership Queries in Agnostic Learning

Complexity and Boolean Functions

Thorsten Doliwa, Michael Kallweit and Hans Ulrich Simon..... 157
Dimension and Margin Bounds for Reflection-invariant Kernels

Dana Angluin, James Aspnes, Jiang Chen, David Eisenstat and Lev Reyzin..... 169
Learning Acyclic Probabilistic Circuits Using Test Paths

Linda Sellie 181
Learning Random Monotone DNF Under the Uniform Distribution

Eric Blais, Ryan O'Donnell and Karl Wimmer 193
Polynomial Regression under Arbitrary Product Distributions

Generalization and Statistics

Alon Zakai and Yaacov Ritov	205
<i>How Local Should a Learning Method Be?</i>	
Yiming Ying and Colin Campbell	217
<i>Learning Coordinate Gradients with Multi-Task Kernels</i>	
Vladimir Koltchinskii and Ming Yuan	229
<i>Sparse Recovery in Large Ensembles of Kernel Machines</i>	

On-Line Learning and Bandits

Amy Greenwald, Zheng Li and Warren Schudy	239
<i>More Efficient Internal-Regret-Minimizing Algorithms</i>	
Giovanni Cavallanti, Nicolo' Cesa-Bianchi and Claudio Gentile	251
<i>Linear Algorithms for Online Multitask Classification</i>	
Jacob Abernethy, Elad Hazan and Alexander Rakhlin	263
<i>Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization</i>	

Other Directions

Wouter M. Koolen and Steven De Rooij	275
<i>Combining Expert Advice Efficiently</i>	
Maria-Florina Balcan, Avrim Blum and Nathan Srebro	287
<i>Improved Guarantees for Learning via Similarity Functions</i>	
J. Hyam Rubinstein and Benjamin I. P. Rubinstein	299
<i>Geometric & Topological Representations of Maximum Classes with Applications to Sample Compression</i>	
Shai Shalev-Shwartz and Yoram Singer	311
<i>On the Equivalence of Weak Learnability and Linear Separability: New Relaxations and Efficient Boosting Algorithms</i>	

Bandits and Reinforcement Learning

- Andrey Bernstein and Nahum Shimkin**..... 323
*Adaptive Aggregation for Reinforcement Learning with Efficient Exploration:
Deterministic Domains*
- Peter Bartlett, Varsha Dani, Thomas Hayes, Sham Kakade, Alexander Rakhlin and
Ambuj Tewari** 335
High-Probability Regret Bounds for Bandit Online Linear Optimization
- Aleksandrs Slivkins and Eli Upfal**..... 343
Adapting to a Changing Environment: the Brownian Restless Bandits
- Varsha Dani, Thomas P. Hayes and Sham M. Kakade**..... 355
Stochastic Linear Optimization under Bandit Feedback

Unsupervised and Semi-Supervised Learning

- Ohad Shamir and Naftali Tishby** 367
Model Selection and Stability in k -means Clustering
- Shai Ben-David and Ulrike von Luxburg** 379
Relating Clustering Stability to Properties of Cluster Boundaries
- Kamalika Chaudhuri and Andrew McGregor**..... 391
Finding Metric Structure in Information Theoretic Clustering
- Karthik Sridharan and Sham M. Kakade**..... 403
An Information Theoretic Framework for Multi-view Learning

Online Learning

- Jacob Abernethy, Peter Bartlett, Alexander Rakhlin and Ambuj Tewari**..... 415
Optimal Strategies and Minimax Lower Bounds for Online Convex Games
- Robert D. Kleinberg, Alexandru Niculescu-Mizil and Yogeshwer Sharma**..... 425
Regret Bounds for Sleeping Experts and Bandits
- Jacob Abernethy, Manfred K. Warmuth and Joel Yellin**..... 437
When Random Play is Optimal Against an Adversary
- Andras Gyorgy, Gabor Lugosi and Gyorgy Ottucsak** 447
On-line Sequential Bin Packing

Generalization and Statistics

Shuheng Zhou, John Lafferty and Larry Wasserman 455
Time Varying Undirected Graphs

Constantine Caramanis and Shie Mannor 467
Learning in the Limit with Adversarial Disturbances

Liwei Wang..... 479
On the Margin Explanation of Boosting Algorithms

Aarti Singh and Robert Nowack and Clayton Scott..... 491
Adaptive Hausdorff Estimation of Density Level Sets

Satyaki Mahalanabis and Daniel Stefankovic..... 503
Density Estimation in Linear Time

Open Problems

Vitaly Feldman and Leslie G. Valiant..... 513
The Learning Power of Evolution

Parikshit Gopalan and Adam Kalai and Adam R. Klivans 515
A Query Algorithm for Agnostically Learning DNF?

Adam M. Smith and Manfred K. Warmuth..... 517
Learning Rotations

Author Index 519

Foreword

This volume contains papers presented at the 21st Annual Conference on Learning Theory (previously known as the Conference on Computational Learning Theory) held in Helsinki, Finland from July 9-12, 2008. The technical program contained 44 papers selected from 126 submissions, three open problems selected from among five contributed, and four invited lectures that were joint with UAI. The invited lectures were given by Peter Grünwald on “The Catch-Up Phenomenon in Bayesian Inference,” by Robin Hanson on “Combinatorial Prediction Markets,” by Dan Klein on “Unsupervised Learning for Natural Language Processing,” and by Gabor Lugosi on “Concentration Inequalities.” The abstracts of these lectures are included in this volume.

The Mark Fulk award is presented annually for the best paper co-authored by a student. This year the Mark Fulk award was supported in part by the *Machine Learning Journal*, which also supported two further awards. Thus three student papers were selected for prizes. The Mark Fulk Award was awarded to Maria-Florina Balcan, Steve Hanneke and Jennifer Wortman for their paper “The True Sample Complexity of Active Learning.” The two Machine Learning Journal Best Paper Awards were awarded to Jacob Abernathy for his paper “Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization” (co-authored by Elad Hazan and Alexander Rakhlin), and to Alexandru Niculescu-Mizil and Yogeshwer Sharma for their paper “Regret Bounds for Sleeping Experts and Bandits” (co-authored with Robert Kleinberg).

This year witnessed many COLT submissions and a very strong program of papers. The selected papers cover a wide range of topics including clustering, unsupervised and semi-supervised learning, active learning, boosting, online learning, bandit problems and reinforcement learning, complexity-theoretic aspects of learning, generalization and statistical learning, kernel methods, and other topics.

We would like to thank the many people who made COLT 2008 a success. We thank the members of the Program Committee for COLT 2008: Dana Angluin (Yale University), Jean-Yves Audibert (Ecole Nationale des Ponts), Peter Auer (University of Leoben), Peter Bartlett (UC Berkeley), Mikhail Belkin (Ohio State University), Shai Ben-David (University of Waterloo), Stéphane Boucheron (Universit Paris-Diderot), Nader Bshouty (Technion), Sanjoy Dasgupta (UC San Diego), Ran El-Yaniv (Technion), Vitaly Feldman (IBM Research), Sham M. Kakade (Toyota Technology Institute), Adam Kalai (Georgia Tech), Vladimir Koltchinskii (Georgia Tech), Sanjay Jain (National University of Singapore), John Langford (Yahoo! Research), Ping Li (Cornell University), Shie Mannor (McGill University), Mehryar Mohri (New York University), Massimiliano Pontil (University College, London), Rob Schapire (Princeton University), Shai Shalev-Shwartz (Hebrew University), Alex Smola (National ICT Australia), Nati Srebro (Toyota Technological Institute), Ingo Steinwart (Los Alamos National Laboratory), Nicolas Vayatis, (Ecole Normale Supérieure de Cachan), Volodya Vovk (Royal Holloway, University of London), and Bob Williamson (Australian National University). We are very grateful to all of them for their careful and thorough reviewing and for the detailed discussions that ensured a strong program for the conference. We thank the many sub-reviewers who assisted the Program Committee; unfortunately space constraints prevent us from including the long list of all their names, so we must ask them to accept our thanks anonymously.

We give special thanks to Jyrki Kivinen (University of Helsinki) who served as the Local Chair of COLT 2008. We thank Kati Kervinen for general administrative support of the conference, and Sanna Kettunen for his work in publicizing the conference. We thank Greger Lindén for creating and maintaining the

conference website, and Microsoft Research for providing the CMT software that was used in the Program Committee deliberations. We thank Nicolò Cesa-Bianchi for helping to organize the conference in his role as head of the COLT steering committee. We thank Ran Gilad-Bachrach for his work in updating and maintaining the www.learningtheory.org website. We also thank the ICML and UAI conference organizers for ensuring a smooth co-location of the three conferences, including overlap with UAI.

Finally, we would like to thank the Federation of Finnish Learned Societies, Google, Helsinki Institute for Information Technology, IBM, the *Machine Learning Journal*, the University of Helsinki, and Yahoo! for their support and sponsorship of the conference.

April 2008

Rocco Servedio and Tong Zhang
COLT 2008 Program Chairs