

Table of Contents

Editors' Introduction	1
<i>Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann</i>	

Invited Papers

Models for Autonomously Motivated Exploration in Reinforcement Learning (Extended Abstract)	14
<i>Peter Auer, Shiao Hong Lim, and Chris Watkins</i>	
On the Expressive Power of Deep Architectures	18
<i>Yoshua Bengio and Olivier Delalleau</i>	
Optimal Estimation	37
<i>Jorma Rissanen</i>	
Learning from Label Preferences	38
<i>Eyke Hüllermeier and Johannes Fürnkranz</i>	
Information Distance and Its Extensions	39
<i>Ming Li</i>	

Inductive Inference

Iterative Learning from Positive Data and Counters	40
<i>Timo Kötzing</i>	
Robust Learning of Automatic Classes of Languages	55
<i>Sanjay Jain, Eric Martin, and Frank Stephan</i>	
Learning and Classifying	70
<i>Sanjay Jain, Eric Martin, and Frank Stephan</i>	
Learning Relational Patterns	84
<i>Michael Geilke and Sandra Zilles</i>	

Regression

Adaptive and Optimal Online Linear Regression on ℓ^1 -Balls	99
<i>Sébastien Gerchinovitz and Jia Yuan Yu</i>	

Re-adapting the Regularization of Weights for Non-stationary Regression	114
<i>Nina Vaitis and Koby Crammer</i>	

Competing against the Best Nearest Neighbor Filter in Regression	129
<i>Arnak S. Dalalyan and Joseph Salmon</i>	

Bandit Problems

Lipschitz Bandits without the Lipschitz Constant	144
<i>Sébastien Bubeck, Gilles Stoltz, and Jia Yuan Yu</i>	

Deviations of Stochastic Bandit Regret	159
<i>Antoine Salomon and Jean-Yves Audibert</i>	

On Upper-Confidence Bound Policies for Switching Bandit Problems ...	174
<i>Aurélien Garivier and Eric Moulines</i>	

Upper-Confidence-Bound Algorithms for Active Learning in Multi-armed Bandits	189
<i>Alexandra Carpentier, Alessandro Lazaric, Mohammad Ghavamzadeh, Rémi Munos, and Peter Auer</i>	

Online Learning

The Perceptron with Dynamic Margin	204
<i>Constantinos Panagiotakopoulos and Petroula Tsampouka</i>	

Combining Initial Segments of Lists	219
<i>Manfred K. Warmuth, Wouter M. Koolen, and David P. Helmbold</i>	

Regret Minimization Algorithms for Pricing Lookback Options	234
<i>Eyal Gofer and Yishay Mansour</i>	

Making Online Decisions with Bounded Memory	249
<i>Chi-Jen Lu and Wei-Fu Lu</i>	

Universal Prediction of Selected Bits	262
<i>Tor Lattimore, Marcus Hutter, and Vaibhav Gavane</i>	

Semantic Communication for Simple Goals Is Equivalent to On-line Learning	277
<i>Brendan Juba and Santosh Vempala</i>	

Kernel and Margin Based Methods

Accelerated Training of Max-Margin Markov Networks with Kernels	292
<i>Xinhua Zhang, Ankan Saha, and S.V.N. Vishwanathan</i>	

Domain Adaptation in Regression	308
<i>Corinna Cortes and Mehryar Mohri</i>	

Approximate Reduction from AUC Maximization to 1-Norm Soft Margin Optimization	324
<i>Daiki Suehiro, Kohei Hatano, and Eiji Takimoto</i>	

Intelligent Agents

Axioms for Rational Reinforcement Learning	338
<i>Peter Sunehag and Marcus Hutter</i>	

Universal Knowledge-Seeking Agents	353
<i>Laurent Orseau</i>	

Asymptotically Optimal Agents	368
<i>Tor Lattimore and Marcus Hutter</i>	

Time Consistent Discounting	383
<i>Tor Lattimore and Marcus Hutter</i>	

Other Learning Models

Distributional Learning of Simple Context-Free Tree Grammars	398
<i>Anna Kasprzik and Ryo Yoshinaka</i>	

On Noise-Tolerant Learning of Sparse Parities and Related Problems ...	413
<i>Elena Grigorescu, Lev Reyzin, and Santosh Vempala</i>	

Supervised Learning and Co-training	425
<i>Malte Darnstädt, Hans Ulrich Simon, and Balázs Szörényi</i>	

Learning a Classifier when the Labeling Is Known	440
<i>Shalev Ben-David and Shai Ben-David</i>	

Erratum

Erratum: Learning without Coding	452
<i>Samuel E. Moelius III and Sandra Zilles</i>	

Author Index	453
--------------------	-----