

Nader H. Bshouty Gilles Stoltz  
Nicolas Vayatis Thomas Zeugmann (Eds.)

# Algorithmic Learning Theory

23rd International Conference, ALT 2012  
Lyon, France, October 29-31, 2012  
Proceedings



Springer

## Series Editors

Randy Goebel, University of Alberta, Edmonton, Canada  
Jörg Siekmann, University of Saarland, Saarbrücken, Germany  
Wolfgang Wahlster, DFKI and University of Saarland, Saarbrücken, Germany

## Volume Editors

Nader H. Bshouty  
Technion, Haifa, Israel  
E-mail: bshouty@cs.technion.ac.il

Gilles Stoltz  
Ecole Normale Supérieure, CNRS, INRIA, Paris, France  
E-mail: gilles.stoltz@ens.fr

Nicolas Vayatis  
Ecole Normale Supérieure de Cachan, France  
E-mail: vayatis@cmla.ens-cachan.fr

Thomas Zeugmann  
Hokkaido University, Sapporo, Japan  
E-mail: thomas@ist.hokudai.ac.jp

ISSN 0302-9743 e-ISSN 1611-3349  
ISBN 978-3-642-34105-2 e-ISBN 978-3-642-34106-9  
DOI 10.1007/978-3-642-34106-9  
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2012949059

CR Subject Classification (1998): I.2, F.4.1, F.1, F.2, I.2.3, I.2.6

LNCS Sublibrary: SL 7 – Artificial Intelligence

© Springer-Verlag Berlin Heidelberg 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

*Typesetting:* Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

# Preface

This volume contains the papers presented at the 23rd International Conference on Algorithmic Learning Theory (ALT 2012), which was held in Lyon, France, October 29–31, 2012. The conference was co-located and held in parallel with the 15th International Conference on Discovery Science (DS 2012). The technical program of ALT 2012 contained 23 papers selected from 47 submissions, and five invited talks. The invited talks were presented in joint sessions of both conferences.

ALT 2012 was dedicated to the theoretical foundations of machine learning and took place in the historical building of the Université Lumière Lyon 2 (berges du Rhône). ALT provides a forum for high-quality talks with a strong theoretical background and scientific interchange in areas such as inductive inference, universal prediction, teaching models, grammatical inference, complexity of learning, online learning, semi-supervised and unsupervised learning, clustering, statistical learning, regression, bandit problems, Vapnik–Chervonenkis dimension, probably approximately correct learning, information-based methods, and applications of algorithmic learning theory.

The present volume contains the texts of the 23 papers presented at ALT 2012, divided into groups of papers on inductive inference, teaching and PAC-learning, statistical learning theory and classification, relations between models and data, bandit problems, online learning of individual sequences, and on other models of online learning. The volume also contains the texts or abstracts of the invited talks:

- Luc De Raedt (Katholieke Universiteit Leuven, Belgium), “Declarative Modeling for Machine Learning and Data Mining” (joint invited speaker for ALT 2012 and DS 2012)
- Shai Shalev-Shwartz (The Hebrew University of Jerusalem, Israel), “Learnability Beyond Uniform Convergence” (invited speaker for ALT 2012)
- Pascal Massart (Université Paris-Sud, France), “Some Rates of Convergence for the Selected Lasso Estimator” (tutorial speaker for ALT 2012)
- Toon Calders (Eindhoven University of Technology, The Netherlands), “Recent Developments in Pattern Mining” (invited speaker for DS 2012)
- Gilbert Ritschard (Université de Genève, Switzerland), “Exploring Sequential Data” (tutorial speaker for DS 2012).

Since 1999, ALT has been awarding the *E. M. Gold Award* for the most outstanding student contribution. This year, the award was given to Ziyuan Gao for his paper “Confident and Consistent Partial Learning of Recursive Functions” co-authored by Frank Stephan.

ALT 2012 was the 23rd in the ALT conference series, established in Japan in 1990. The ALT series is supervised by its Steering Committee: Naoki Abe (IBM Thomas J. Watson Research Center, Yorktown, USA), Shai Ben-David

(University of Waterloo, Canada), Nader Bshouty (Technion - Israel Institute of Technology, Israel) Marcus Hutter (Australian National University, Canberra, Australia), Jyrki Kivinen (University of Helsinki, Finland), Philip M. Long (Google, Mountain View, USA), Akira Maruoka (Ishinomaki Senshu University, Japan), Takeshi Shinohara (Kyushu Institute of Technology, Iizuka, Japan), Frank Stephan (National University of Singapore, Republic of Singapore), Gilles Stoltz (Ecole Normale Supérieure, France), Einoshin Suzuki (Kyushu University, Fukuoka, Japan), Csaba Szepesvári (University of Alberta, Canada), Eiji Takimoto (Kyushu University, Fukuoka, Japan), György Turán (University of Illinois at Chicago, USA and University of Szeged, Hungary), Osamu Watanabe (Tokyo Institute of Technology, Japan), Thomas Zeugmann (Chair, Hokkaido University, Japan), and Sandra Zilles (Publicity Chair, University of Regina, Saskatchewan, Canada).

We would like to thank the many people and institutions who contributed to the success of the conference. In particular, we want to thank our authors for contributing to the conference and for coming to Lyon in October 2012. Without their efforts and their willingness to choose ALT 2012 as a forum to report on their research, this conference would not have been possible.

ALT 2012 and DS 2012 were organized by the Université Lumière Lyon 2, France. We are very grateful to the General Chair Djamel Abdelkader Zighed and the General Local Arrangements Chair Stéphane Lallich. We would like to thank them and their team for the tremendous amount of work they have dedicated to making ALT 2012 and DS 2012 a success. We are grateful for the continuous collaboration with the series Discovery Science. In particular, we would like to thank the Conference Chair Jean-Gabriel Ganascia and the Program Committee Chairs Philippe Lenca and Jean-Marc Petit of Discovery Science 2012.

We are also grateful that we could use the excellent conference management system EasyChair for putting together the program for ALT 2011; EasyChair was developed mainly by Andrei Voronkov and is hosted at the University of Manchester. The system is cost-free.

We are grateful to the members of the Program Committee for ALT 2012 and the subreferees for their hard work in selecting a good program for ALT 2012. Reviewing papers and checking the correctness of results is demanding in time and skills and we very much appreciate this contribution to the conference. Last but not least we thank Springer for their support in preparing and publishing this volume in the *Lecture Notes in Artificial Intelligence* series.

August 2012

Nader H. Bshouty  
Gilles Stoltz  
Nicolas Vayatis  
Thomas Zeugmann

# Organization

## General Chair (ALT 2012 and DS 2012)

Djamel Zighed                      Université Lumière Lyon 2, France

## General Local Chair (ALT 2012 and DS 2012)

Stéphane Lallich                      Université Lumière Lyon 2, France

## ALT 2012 Conference Chair

Nicolas Vayatis                      Ecole Normale Supérieure de Cachan, France

## Program Committee

Jake Abernethy	University of Pennsylvania, USA
András Antos	MTA SZTAKI, Hungary
Shai Ben-David	University of Waterloo, Canada
Avrim Blum	Carnegie Mellon University, USA
Nader H. Bshouty (Chair)	Technion–Israel Institute of Technology, Israel
Koby Crammer	Technion–Israel Institute of Technology, Israel
Sanjoy Dasgupta	UC San Diego, USA
Vitaly Feldman	IBM Research, USA
Claudio Gentile	Università dell’Insubria, Varese, Italy
Timo Kötzing	Max-Planck-Institut für Informatik, Germany
Wouter M. Koolen	Royal Holloway, University of London, UK
Phil M. Long	Google, Mountain View, USA
Claire Monteleoni	George Washington University, USA
Lev Reyzin	Georgia Institute of Technology, USA
Steven de Rooij	CWI Amsterdam, The Netherlands
Daniil Ryabko	INRIA, France
Rocco Servedio	Columbia University, USA
Hans Ulrich Simon	Ruhr-Universität Bochum, Germany
Frank Stephan	National University of Singapore
Gilles Stoltz (Chair)	Ecole Normale Supérieure, Paris, France
Csaba Szepesvári	University of Alberta, Edmonton, Canada
Eiji Takimoto	Kyushu University, Japan
Ambuj Tewari	University of Texas at Austin, USA
Vladimir Vovk	Royal Holloway, University of London, UK
Bob Williamson	Australian National University, Australia
Sandra Zilles	University of Regina, Canada

## ALT 2012 Local Arrangements Chair

Fabrice Muhlenbach                      Université Jean Monnet de Saint-Etienne, France

## Subreferees

Mohammad Gheshlaghi Azar  
Peter Antal  
Frank J. Balbach  
Sébastien Bubeck  
John Case  
Nicolò Cesa-Bianchi  
Karthikeyan Chandrasekaran  
Archie Chapman  
François Coste  
Sanmay Das  
John Duchi  
Rafael M. Frongillo  
András György  
Zaid Harchaoui  
Kohei Hatano  
David P. Helmbold  
Daniel Hsu  
Benjamin Jourdain  
Satyen Kale

Yuri Kalnishkan  
Nikos Karampatziakis  
Azadeh Khaleghi  
Roni Khardon  
Aryeh Kontorovich  
Nathaniel Korda  
Luigi Malagò  
Scott McQuade  
Samuel E. Moelius III  
Edward Moroshko  
Francesco Orabona  
Ronald Ortner  
Vikas Sindhwani  
István Szita  
Matus Telgarsky  
Nina Vaitz  
Vladimir Vovk  
Andre Wibisono  
Thomas Zeugmann

## Sponsoring Institutions



# Table of Contents

Editors' Introduction .....	1
<i>Nader H. Bshouty, Gilles Stoltz, Nicolas Vayatis, and Thomas Zeugmann</i>	

## Invited Papers

Declarative Modeling for Machine Learning and Data Mining .....	12
<i>Luc De Raedt</i>	
Learnability beyond Uniform Convergence .....	13
<i>Shai Shalev-Shwartz</i>	
Some Rates of Convergence for the Selected Lasso Estimator .....	17
<i>Pascal Massart and Caroline Meynet</i>	
Recent Developments in Pattern Mining .....	34
<i>Toon Calders</i>	
Exploring Sequential Data .....	35
<i>Gilbert Ritschard</i>	

## Regular Contributions

### Inductive Inference

Enlarging Learnable Classes .....	36
<i>Sanjay Jain, Timo Kötzing, and Frank Stephan</i>	
Confident and Consistent Partial Learning of Recursive Functions .....	51
<i>Ziyuan Gao and Frank Stephan</i>	
Automatic Learning from Positive Data and Negative Counterexamples .....	66
<i>Sanjay Jain and Efim Kimber</i>	
Regular Inference as Vertex Coloring .....	81
<i>Christophe Costa Florêncio and Sicco Verwer</i>	

## Teaching and PAC-Learning

- Sauer’s Bound for a Notion of Teaching Complexity ..... 96  
*Rahim Samei, Pavel Semukhin, Boting Yang, and Sandra Zilles*
- On the Learnability of Shuffle Ideals ..... 111  
*Dana Angluin, James Aspnes, and Aryeh Kontorovich*

## Statistical Learning Theory and Classification

- New Analysis and Algorithm for Learning with Drifting  
Distributions ..... 124  
*Mehryar Mohri and Andres Muñoz Medina*
- On the Hardness of Domain Adaptation and the Utility of Unlabeled  
Target Samples ..... 139  
*Shai Ben-David and Ruth Urner*
- Efficient Protocols for Distributed Classification and Optimization ..... 154  
*Hal Daumé III, Jeff M. Phillips, Avishek Saha, and  
Suresh Venkatasubramanian*

## Relations between Models and Data

- The Safe Bayesian: Learning the Learning Rate via the Mixability  
Gap ..... 169  
*Peter Grünwald*
- Data Stability in Clustering: A Closer Look ..... 184  
*Lev Reyzin*

## Bandit Problems

- Thompson Sampling: An Asymptotically Optimal Finite-Time  
Analysis ..... 199  
*Emilie Kaufmann, Nathaniel Korda, and Rémi Munos*
- Regret Bounds for Restless Markov Bandits ..... 214  
*Ronald Ortner, Daniil Ryabko, Peter Auer, and Rémi Munos*
- Minimax Number of Strata for Online Stratified Sampling Given Noisy  
Samples ..... 229  
*Alexandra Carpentier and Rémi Munos*



## Online Prediction of Individual Sequences

Weighted Last-Step Min-Max Algorithm with Improved Sub-logarithmic Regret . . . . .	245
<i>Edward Moroshko and Koby Crammer</i>	
Online Prediction under Submodular Constraints . . . . .	260
<i>Daiki Suehiro, Kohei Hatano, Shuji Kijima, Eiji Takimoto, and Kiyohito Nagano</i>	
Lower Bounds on Individual Sequence Regret . . . . .	275
<i>Eyal Gofer and Yishay Mansour</i>	
A Closer Look at Adaptive Regret . . . . .	290
<i>Dmitry Adamskiy, Wouter M. Koolen, Alexey Chernov, and Vladimir Vovk</i>	
Partial Monitoring with Side Information . . . . .	305
<i>Gábor Bartók and Csaba Szepesvári</i>	

## Other Models of Online Learning

PAC Bounds for Discounted MDPs . . . . .	320
<i>Tor Lattimore and Marcus Hutter</i>	
Buy Low, Sell High . . . . .	335
<i>Wouter M. Koolen and Vladimir Vovk</i>	
Kernelization of Matrix Updates, When and How? . . . . .	350
<i>Manfred K. Warmuth, Wojciech Kottowski, and Shuisheng Zhou</i>	
Predictive Complexity and Generalized Entropy Rate of Stationary Ergodic Processes . . . . .	365
<i>Mrinalkanti Ghosh and Satyadev Nandakumar</i>	
<b>Author Index</b> . . . . .	381

# Editors' Introduction

Nader H. Bshouty, Gilles Stoltz, Nicolas Vayatis, and Thomas Zeugmann

The ALT-conference series is dedicated to studies on learning from an algorithmic and mathematical perspective. In the following, the five invited lectures and the regular contributions are introduced in some more detail.

**Invited Talks.** It is now a tradition of the co-located conferences ALT and DS to have a joint invited speaker—namely this year, Luc De Raedt. Since 2006 he is a full research professor at the Department of Computer Science of the Katholieke Universiteit Leuven (Belgium). His research interests are in artificial intelligence, machine learning and data mining, as well as their applications. He is currently working on probabilistic logic learning (sometimes called statistical relational learning), which combines probabilistic reasoning methods with logical representations and machine learning, the integration of constraint programming with data mining and machine learning principles, the development of programming languages for machine learning, and analyzing graph and network data. In his talk *Declarative Modeling for Machine Learning and Data Mining* he notes that despite the popularity of machine learning and data mining today, it remains challenging to develop applications and software that incorporates machine learning or data mining techniques. This is because machine learning and data mining have focused on developing high-performance algorithms for solving particular tasks rather than on developing general principles and techniques. He thus proposes to alleviate these problems by applying the constraint programming methodology to machine learning and data mining and to specify machine learning and data mining problems as constraint satisfaction and optimization problems. The aim is that the user be provided with a way to declaratively specify what the machine learning or data mining problem is rather than having to outline how that solution needs to be computed.

Four other invited talks are also given by eminent researchers in their fields, who present either an introduction to their specific research area or give a lecture of wide general interest.

Shai Shalev-Shwartz is the ALT invited speaker; since 2009 he is a senior lecturer at the School of Computer Science and Engineering of The Hebrew university (Jerusalem, Israel). His research interests include machine learning and learning theory at broad, with an emphasis on online algorithms, large-scale learning, information retrieval, and optimization. In his talk *Learnability Beyond Uniform Convergence* he discusses the problem of characterizing learnability, which in his view is the most basic question of statistical learning theory. He indicates that a fundamental result is that learnability is equivalent to uniform convergence of the empirical risk to the population risk, and that if a problem is learnable, it is learnable via empirical risk minimization. However, the equivalence of uniform convergence and learnability was formally established only in

the supervised classification and regression setting. He then shows that in (even slightly) more complex prediction problems learnability does not imply uniform convergence. He thus presents several alternative attempts to characterize learnability. The results obtained are based on joint researches with Ohad Shamir, Nati Srebro, Karthik Sridharan, and with Amit Daniely, Sivan Sabato, and Shai Ben-David, respectively.

Pascal Massart is the ALT tutorial speaker; since 1990 he is a full professor at the Department of Mathematics of the Université Paris-Sud (Orsay, France). He dedicated most of his work in the past 20 years to elaborate a non-asymptotic theory for model selection and made contributions also to related fields, like the theory of empirical processes, concentration-of-the-measure inequalities, and non-parametric statistics. He also established connections between model selection theory and statistical learning theory. His tutorial is based on the paper *Some Rates of Convergence for the Selected Lasso Estimator* co-authored with Caroline Meynet. He illustrates on the example of the Lasso estimator how the theory of model selection in statistics can shed some light and improve some results in learning. More precisely he considers the estimation of a function in some ordered finite or infinite dictionary, that is, in some (non necessarily orthonormal) family of elements in a Hilbert space. He focuses on a variant of the Lasso, the selected Lasso estimator, which he introduced in an earlier paper with Caroline Meynet. This estimator is an adaptation of the Lasso suited to infinite dictionaries. He uses the oracle inequality established therein to derive rates of convergence of this estimator on a wide range of function classes (Besov-type spaces). The results highlight that the selected Lasso estimator is adaptive to the smoothness of the function to be estimated, contrary to the classical Lasso or to other algorithms considered in the literature.

Toon Calders, the invited speaker for DS, received his PhD in Mathematics in 2003 from the University of Antwerp. Since 2006 he is assistant professor in the Information Systems Group at the Department of Mathematics and Computer Science of the Eindhoven University of Technology. His lecture *Recent Developments in Pattern Mining* gives an overview of the many techniques developed to solve pattern mining problems. Many methods have been proposed to enumerate all frequent itemsets. The basic difficulty is the pattern explosion problem, i.e., millions of patterns may be generated. Though this problem is widely recognized, it still lacks a satisfactory solution. Toon Calders surveys promising methods based upon the minimal description length principle, information theory, and statistical models, and discusses the advantages and disadvantages of these new methods. The final part of his lecture addresses more complex patterns such as sequences and graphs, and concludes with important open problems in this challenging area.

Gilbert Ritschard, the DS tutorial speaker, graduated in econometrics and got his PhD in Econometrics and Statistics at the University of Geneva in 1979. He also taught as invited professor in Toronto, Montreal, Lyon, Lausanne and Fribourg and participated as a statistical expert in several large statistical modeling projects of International Organizations. He is a full professor of statistics at the

Department of Economics of the University of Geneva, where he is responsible for the program of statistics and quantitative methods for the social sciences and runs his researches within the Institute for Demographic and Life Course Studies and acts as vice-dean of the Faculty of Economics and Social Sciences since 2007. Several funded applied researches were headed or co-headed by him. He also published papers in economics as well as on more applied topics in the field of social sciences, especially in demography, sociology and social science history. With his team he developed the world wide used TraMineR toolbox for exploring and analyzing sequence data in R. His present research interests are in categorical and numerical longitudinal data analysis and their application to life course analysis. His tutorial *Exploring Sequential Data* gives an introduction to sequence analysis as it is practiced for life course analysis. Examples comprise successive buys of customers, working states of devices, visited web pages, or professional careers, and addressed topics are the rendering of state and event sequences, longitudinal characteristics of sequences, measuring pairwise dissimilarities and dissimilarity-based analysis of sequence data such as clustering, representative sequences, and regression trees. All the methods employed are available in TraMineR R-package.

We now turn our attention to the regular contributions contained in this volume.

***Inductive Inference.*** One of the classical areas of algorithmic learning is inductive inference of recursive functions. In this setting the learner is usually fed augmenting finite sequences  $f(0), f(1), f(2), \dots$  of the target function  $f$ . For each finite sequence the learner has to compute a hypothesis, i.e., a natural number. These numbers are interpreted with respect to a given enumeration of partial recursive functions comprising the target function. Then the number  $j$  output by the learner is interpreted as a program computing the  $j$ th function enumerated. The sequence of all hypotheses output by the learner has then to converge (to stabilize) on a program that, under the given correctness criterion, correctly computes the target function. This learning scenario is commonly called *explanatory inference* or *learning in the limit*. Since only finitely many values of the function have been seen by the learner up to the unknown point of convergence, some form of learning must have taken place. Usually, the goal is then to construct a learner that can infer all functions from a given target class  $U$ . Many variations of this model are possible. For finite learning one requires the point of convergence to be decidable. Another variation is to allow the learner to converge semantically, i.e., instead of stabilizing to a correct program, the learner is allowed to output infinitely many *different programs* which, however, beyond some point, all must correctly compute the target function. This model is commonly referred to as *behaviorally correct learning*. In the context of the first paper introduced below also the notions of *confidence* and *reliability* are of particular interest. A confident learner is required to converge on every function, even it is not in the target class (but may stabilize to a special symbol “?”). In contrast, a reliable learner must signal its inability to learn a target

function (which may be again outside the class) by performing infinitely many mind changes. Thus, if a reliable learner converges then it learns. In this context it remains to specify what is meant by “outside” the target class. Above all total functions (including the non-computable ones) are considered. If one allows only the total recursive functions then the resulting models are called *weakly confident* and *weakly reliable*, respectively.

The problems studied by Sanjay Jain, Timo Kötzing, and Frank Stephan in their paper *Enlarging Learnable Classes* are easily described as follows. Suppose we have already a learner for a target class  $U_1$  and another one for a class  $U_2$ . Then it is only natural to ask under which circumstances one can obtain a more powerful learner that simultaneously infers all functions from  $U_1 \cup U_2$ . A classical result shows that this is not always possible, even for behaviorally correct learning. If one can obtain such a more powerful learner then it is also interesting to ask whether or not it can be done effectively. That is, given programs for learners  $M_1$  and  $M_2$  inferring  $U_1$  and  $U_2$ , respectively, one studies the problem whether or not one can compute from these programs a learner  $M$  for the union  $U_1 \cup U_2$ . Still, it is imaginable that one cannot compute such a learner but show it to exist (this is the non-effective case). The interesting new modification of this problem introduced by the authors is to ask which classes  $U_1$  have the property that  $U_1 \cup U_2$  is learnable for *all* classes  $U_2$ . As shown by Jain *et al.*, if  $U_1$  has a weakly confident and reliable learner then the union  $U_1 \cup U_2$  is always explanatory learnable and the learner is effective. Moreover, they show the effective case and the non-effective case separate and a sufficient criterion is shown for the effective case. A closely related problem is to ask the same questions when the second class is restricted to be any singleton class. In this case it suffices that the learner for  $U_1$  is weakly confident to obtain the effective case. In contrast, for finite learning there is no non-empty class  $U_1$  satisfying the non-constructive case for classes and the constructive case for singleton classes. Furthermore, the authors investigate the problem how much information is needed to enlarge a learnable class by infinitely many functions while maintaining its learnability. In this context, two questions that remained open in a paper by Mark Fulk and John Case in 1999 are completely answered.

The next paper in this section is the Gold Award winning paper *Confident and Consistent Partial Learning of Recursive Functions* by Ziyuan Gao and Frank Stephan for the best paper co-authored by a student, who is Ziyuan Gao. As the discussion above shows there is no single learner that can infer the whole class of recursive functions. Therefore, it is interesting to consider further variations. Osherson, Stob, and Weinstein (1986) considered partial learning, where the learner is required to output a correct program for the target function infinitely often and any other hypothesis only finitely often. Gao and Stephan refine this model by combining it with the confidence demand discussed above and with consistent learning. A consistent learner has to correctly reflect the information already obtained, and this demand is posed to all but finitely many of the hypotheses output. The resulting models are called confident partial learning and consistent partial learning, respectively. The paper contains many interesting

results and masters several complicated proof techniques. In particular, it is shown that confident partial learning is more powerful than explanatory learning. On the other hand, the authors show that there are behaviorally correct learnable classes which are *not* confidently partially learnable. So, the learning model is also not trivial in the sense that it can infer every recursive function. Moreover, confident partial learning has another interesting property, i.e., it is closed under finite unions. The authors then study confident partial learning with respect to oracles, and obtain some deep results. That is, in addition to the successively fed graph of the the target function, the learner has access to an oracle. The second part of the paper combines partial learning with consistency. Since a consistent learner is preferable, these results deserve attention. On the positive site it is shown that every behaviorally correct learnable class is also essentially class consistently partially learnable. On the other hand, the set of all recursive predicates is not essentially class consistently partially learnable. Finally, it is shown that PA-oracles are sufficient in order to partially learn every recursive function essentially class consistently.

The paper *Automatic Learning from Positive Data and Negative Counterexamples* by Sanjay Jain and Efim Kinber deals with the inductive inference of languages. So the target is a formal language and the information given to the learner may be eventually all strings in the language (positive examples only), all strings over the underlying alphabet which are then marked with respect to their containment in the target language, or, as in the present paper, positive examples and negative counterexamples (but not all). This source of information is justified by two facts. First, learning from positive examples only is often too weak, and receiving potentially all strings does not reflect, e.g., natural language acquisition. Again, one has to study the problem of inferring all languages from a given target class of languages by one learner. In their paper Jain and Kinber consider classes of target languages that are required to be automatic ones. That is, the authors consider classes of regular languages of the form  $(L_i)_{i \in I}$  such that  $\{(i, x) \mid x \in L_i\}$  and  $I$  itself are regular sets. So automatic classes of languages are a particular type of an automatic structure. In this context it should be noted that the family of automatic classes which are inferable from positive examples only is rather restricted. Thus, it is natural to ask under which conditions *all* automatic classes are automatically learnable. Here automatically learnable means that the learner itself must be describable by an automatic structure. The rest of the model is *mutatis mutandis* the same as in the inductive inference of recursive functions. However, since the learner is required to be automatic, it can obviously not memorize all data. So, the constraint to learn iteratively and/or with a bounded long term memory is very natural in this context. Here iterative means that the learner can just store the last example seen. The authors distinguish between least counterexamples (a shortest possible one), bounded counterexamples (bounded in the size of the longest positive example seen so far) and arbitrary counterexamples. The first main result is that all automatic classes are automatically learnable (iteratively and with bounded long term memory, respectively) from positive examples and arbitrary counterexamples. Furthermore, there are

automatic classes that *cannot* be learned from positive examples and bounded counterexamples. The authors show many more results for which we refer the reader to the paper.

Christophe Costa Florêncio and Sicco Verwer in *Regular Inference as Vertex Coloring* also study a problem that belongs to the inductive inference of formal languages, i.e., learning the class of all regular languages from complete data in the limit. The hypothesis space chosen is the set of all deterministic finite automata (abbr. DFA). In this context it is known that it suffices to output in each learning step a minimal DFA that is consistent with all the data seen so far. This is, however, easier said than done, since the problem is known to be *NP*-complete. Thus the idea is to reduce the learning problem to satisfiability and to exploit the enormous progress made for satisfiability solvers. The approach undertaken previously is to perform this in two steps, i.e., first the learning problem is translated into a graph coloring problem, and second the graph coloring problem obtained is translated into a satisfiability problem. Here the first step included some inequality constraints (requiring the constraint vertices to have a different color) as well as some equality constraints. So, these constraints had to be translated into the resulting satisfiability problem. The main contribution of the present paper is an improvement for the first step that allows for a direct translation of the inference problem into a graph coloring problem. In this way, one can also directly use sophisticated solvers for graph coloring.

**Teaching and PAC-Learning.** Each learning model specifies the learner, the learning domain, the source of information, the hypothesis space, what background knowledge is available and how it can be used, and finally, the criterion of success. While the learner is always an algorithm, it may also be restricted in one way or another, e.g., by requiring it to be space and/or time efficient.

A significant line of work over the past decade studies combinatorial measures of the complexity of teaching. In this framework a helpful teacher chooses an informative sequence of labeled examples and provides them to the learner, with the goal of uniquely specifying the target concept from some a priori concept class of possible target functions. Several different combinatorial parameters related to this framework have been defined and studied, including the worst-case teaching dimension, the average teaching dimension, and the "recursive teaching dimension".

Rahim Samei, Pavel Semukhin, Boting Yang and Sandra Zilles in *Sauer's Bound for a Notion of Teaching Complexity* show that Sauer's Lemma can be adjusted to the recursive teaching dimension of the concept. This paper establishes an upper bound on the size of a concept class with given recursive teaching dimension. The upper bound coincides with Sauer's well-known bound on classes with a fixed VC-dimension. They further introduce and study classes whose size meets the upper bound and other properties of this measure.

It is well known that the language accepted by an unknown deterministic finite automata can be efficiently PAC-learnable if membership queries are allowed. It is also well known that cryptographic lower bounds preclude the efficient PAC learnability of arbitrary DFAs when membership queries are not allowed and

learning must be based only on random examples. It is natural to ask about whether specific restricted types of regular languages are PAC learnable. A shuffle ideal generated by a string  $u$  is simply the collection of all strings containing  $u$  as a (discontiguous) subsequence.

The paper *On the Learnability of Shuffle Ideals* by Dana Angluin, James Aspnes, and Aryeh Kontorovich shows that shuffle ideal languages are efficiently learnable from statistical queries under the uniform distribution, but not efficiently PAC-learnable, unless  $RP = NP$ .

***Statistical Learning Theory and Classification.*** The simplest setting in which statistical learning theory takes place is the following. A training set  $(X_t, Y_t)$  of samples is given, where the outcomes  $Y_t$  depend on the instances  $X_t$ ; the learner then has to construct some rule to predict new outcomes  $Y$  for new instances  $X$ . Often the training set is formed by independent and identically distributed samples and the new instance–outcome pairs are drawn independently from the same distribution. The simplest task is (binary) classification, which corresponds to the case where the outcomes  $Y$  are  $\{0, 1\}$ -valued. More abstract formulations of the learning task can be provided, based on a hypothesis space (gathering functions  $h$  that map instances to outcomes) and on a loss function (associating with each pair of predicted outcome  $h(X)$  and observed outcome  $Y$  a measure of their divergence). We call generalization bounds the bounds on the expected loss of a prediction function  $\hat{h}$  constructed on the training set and evaluated on a new independent random pair  $(X, Y)$ . These bounds are often in terms of the hypothesis set (and in particular, of its so-called Vapnik-Chervonenkis dimension or of its Rademacher complexity).

Mehryar Mohri and Andres Muñoz Medina present a *New Analysis and Algorithm for Learning with Drifting Distributions*, that is, they consider the case where the distribution of the instance–outcomes pairs evolves with  $t$ . Their analysis relies on the notion of discrepancy, which is a loss-based measure of divergence. They prove performance bounds based on the Rademacher complexity of the hypothesis set and the discrepancy of distributions; these bounds improve upon previous ones based on the  $\mathbb{L}_1$ -distances between distributions.

Another twist on the simplest problem described above is formed by domain adaptation, which corresponds to the case where the test and training data generating distributions differ. Shai Ben-David and Ruth Urner's contribution is *On the Hardness of Covariate Shift Learning (and the Utility of Unlabeled Target Samples)*; the covariate shift setting refers to the assumption that outcomes  $Y$  are solely determined by instances  $X$ , and that the function linking the two elements is the same in both domains. Algorithms had been proposed in this setting but often with very few generalization guarantees. The authors show that, without strong prior knowledge about the training task, such guarantees are actually unachievable, unless the training set and the set of new instance–outcomes pairs are prohibitively large. However, the (necessarily large) set of new elements can be formed rather by mostly unlabeled instances, which are often much cheaper to obtain than instance–outcome pairs.



Hal Daumé III, Jeff M. Phillips, Avishek Saha, and Suresh Venkatasubramanian study *Efficient Protocols for Distributed Classification and Optimization*. Their contribution takes places within a general model for distributed learning that bounds the communication required for learning classifiers with  $\varepsilon$  error on linearly separable data adversarially distributed across nodes; this model was introduced by the authors in an earlier article and they elaborate on it here. Their main result is a two-party multiplicative-weight-update based protocol that uses  $O(d^2 \log(1/\varepsilon))$  words of communication to  $\varepsilon$ -optimally classify distributed data in arbitrary dimension  $d$ . This result extends to classification over  $k$  nodes with  $O(kd^2 \log(1/\varepsilon))$  words of communication. The proposed protocol is simple to implement and empirical results show its improved efficiency.

***Relations between Models and Data.*** Data is the raw material for learning and is often handled through a model or a collection of models. But sometimes the available theoretical models can be partially wrong; or even worse, no such theoretical models exist and they need to be constructed from the data.

Standard Bayesian inference can behave suboptimally if the model is wrong. Peter Grünwald presents in his article *The Safe Bayesian: Learning the Learning Rate via the Mixability Gap* a modification of Bayesian inference which continues to achieve good rates with wrong models. The method adapts the Bayesian learning rate to the data, picking the rate minimizing the cumulative loss of sequential prediction by posterior randomization.

Clustering (the partition of data into meaningful categories) is one of the most widely used techniques in statistical data analysis. A recent trend of research in this field is concerned with so-called perturbation resilience assumptions. Lev Reyzin defines in *Data Stability in Clustering: A Closer Look* a new notion of stability that is implied by perturbation resilience and discusses the implications of assuming resilience or stability in the data; the strength of this resilience or stability is measured by a constant  $\alpha$ . He shows that for even fairly small constants  $\alpha$ , the data begins to have very strong structural properties, which makes the clustering task fairly trivial. When  $\alpha$  approaches  $\approx 5.7$ , the data begins to show what is called strict separation, where each point is closer to points in its own cluster than to points in other clusters.

***Bandit Problems.*** Bandit problems form a model of repeated interaction between a learner and a stochastic environment. In its simplest formulation the learner is given a finite number of arms, each associated with an unknown probability distribution with bounded support. Whenever he pulls an arm he gets some reward, drawn independently at random according to its associated distribution; his objective is to maximize the obtained cumulative reward. To do so, a trade-off between testing sufficiently often all the arms (exploration) and pulling more often the seemingly better arms (exploitation) needs to be performed. A popular strategy, the UCB strategy, constructs confidence bounds for the expected reward of each arm and pulls at each round the arm with best upper confidence bound.

In their paper *Thompson Sampling: An Optimal Finite Time Analysis*, Emilie Kaufmann, Nathaniel Korda, and Rémi Munos study another, older, strategy, called Thompson sampling; it relies on a Bayesian estimation of the expected reward. The authors show that in the case of Bernoulli distributions of the rewards, this strategy is asymptotically optimal.

Ronald Ortner, Daniil Ryabko, Peter Auer, and Rémi Munos consider in their contribution *Regret Bounds for Restless Markov Bandits* a more difficult scenario in which the rewards produced by each arm are not independent and identically distributed anymore; they are governed by Markov chains, which take transitions independently of whether the learner pulls this arm or not. They derive  $O(\sqrt{T})$  regret bounds, without formulating any assumption on the distributions of the Markov chains.

An application of these bandit problems is studied in *Minimax Number of Strata for Online Stratified Sampling given Noisy Samples* by Alexandra Carpentier and Rémi Munos: how to approximate the integral of a function  $f$  given a finite budget of  $n$  noisy evaluations of the function. This is done by resorting to an online stratified sampling performed by the algorithm Monte-Carlo UCB developed by the authors in an earlier article. In their contribution to this volume they show that this algorithm is minimax optimal both in terms of the number of samples  $n$  and in the number of strata  $K$ , up to logarithmic factors.

***Online Prediction of Individual Sequences.*** Another setting of repeated interactions between a learner and an environment is formed by the setting of online prediction of individual sequences. However, here, the environment may also use a strategy to pick his actions. At each round, the learner suffers a loss (or a gain) that only depends on the pair of actions taken by the two players. The quality of the strategy of the learner is measured through its regret, that is, the difference between his cumulative loss and the cumulative loss that the best constant choice of an action would have obtained on the same sequence of actions of the environment. Simple and efficient strategies exist to control the regret when the range of the losses is known and the number of actions is not too large. The first two papers described below relax these requirements. The three other papers deal with refinements of the basic situation presented above: the third one studies how sharp regret bounds can be, the fourth one focuses on a refined notion of regret, and the fifth one considers the case where only a partial monitoring of the actions taken by the environment is available.

The article *Weighted Last-Step Min-Max Algorithm with Improved Sub-Logarithmic Regret* by Edward Moroshko and Koby Crammer takes place within the framework of online linear regression with the square loss. It proposes a development of Forster's last-step min-max algorithm for the case where the range of the choices of the environment is unknown.

Daiki Suehiro, Kohei Hatano, Shuji Kijima, Eiji Takimoto, and Kiyohito Nagano deal with a case where the set of actions of the learner is large but bears some structure. More precisely, their contribution *Online Prediction under Submodular Constraints* focuses on the case of an action set formed by the vertices of a polyhedron described by a submodular function. Examples of the

general problem handled there include the cases of  $k$ -sets, (truncated) permutahedra, spanning trees, and  $k$ -forests.

Another line of research is to study how sharp the regret bounds can be. Eyal Gofer and Yishay Mansour focus in *Lower Bounds on Individual Sequence Regret* on lower bounds on the regret of algorithms only based on the cumulative losses of the actions, which include popular strategies. They characterize those with a nonnegative regret; they also show that any such algorithm obtaining in addition a refined  $O(\sqrt{Q})$  upper bound in terms of quadratic variations of the losses must also suffer an  $\Omega(\sqrt{Q})$  lower bound for any loss sequence with quadratic variation  $Q$ .

Dmitry Adamskiy, Wouter M. Koolen, Alexey Chernov, and Vladimir Vovk take *A Closer Look at Adaptive Regret*, which is a refined notion a regret. It corresponds to measuring regret only on subintervals of time, that is, to assessing how well the algorithm approximates the best experts locally. They investigate two existing intuitive methods to derive algorithms with low adaptive regret, one based on specialist experts and the other based on restarts; they show that both methods lead to the same algorithm, namely Fixed Share, which was known for its tracking regret. They then perform a thorough analysis of the adaptive regret of Fixed Share and prove the optimality of this strategy in this context.

The setting of *Partial Monitoring with Side Information* by Gábor Bartók and Csaba Szepesvári is the following. At every round the learner only receives a partial feedback about the choice of the action taken by the environment. The interaction protocol relies on a fixed function  $f$ , unknown to the learner: The action taken by the environment is a vector  $x_t$ , which is revealed to the learner, and is then used to draw at random the final action  $J_t$  of the environment according to the distribution  $f(x_t)$ . Simultaneously and based on  $x_t$ , the learner chooses his action  $I_t$ . The only feedback he gets is drawn at random according to a distribution that depends only on  $I_t$  and  $J_t$ , but he does not get to see  $J_t$ . The authors define a notion of regret in this setting and show an algorithm to minimize it.

***Other Models of Online Learning.*** This section gathers the contributions relative to online learning but that correspond neither to bandit problems nor to the prediction of individual sequences.

The goal of reinforcement learning is to construct algorithms that learn to act optimally, or nearly so, in unknown environments. Tor Lattimore and Marcus Hutter focus on finite-state discounted Markov decision processes (MDPs). More precisely, in *PAC Bounds for Discounted MDPs* they exhibit matching (up to logarithmic factors) upper and lower bounds on the sample-complexity of learning near-optimal behavior. These upper bounds are obtained for a modified version of algorithm UCRL.

Wouter M. Koolen and Vladimir Vovk present in *Buy Low, Sell High* a simplified setting of online trading where an investor trades in a single security. His objective is to get richer when the price of the security exhibits a large upcrossing without risking bankruptcy. They investigate payoff guarantees that are expressed in terms of the extremity of the upcrossings and obtain an exact and

elegant characterization of the guarantees that can be achieved. Moreover, they derive a simple canonical strategy for each attainable guarantee.

Kernels methods consist of mapping the data points from their original space to a feature space, where the analysis and prediction are performed more efficiently; the obtained results are then mapped back into the original space. In their paper *Kernelization of Matrix Updates, When and How?* Manfred Warmuth, Wojciech Kotłowski, and Shuisheng Zhou define what it means for a learning algorithm to be kernelizable in the case where the instances are vectors, asymmetric matrices, and symmetric matrices, respectively. They characterize kernelizability in terms of an invariance of the algorithm to certain orthogonal transformations. They provide a number of examples in the online setting of how to apply their methods.

The paper *Predictive Complexity and Generalized Entropy Rate of Stationary Ergodic Processes* by Mrinalkanti Ghosh and Satyadev Nandakumar takes place in the framework of online prediction of binary outcomes. They use generalized entropy to study the loss rate of predictors when these outcomes are drawn according to stationary ergodic distributions. They use a game-theoretic viewpoint and first show that a notion of generalized entropy of a regular game is well-defined for stationary ergodic distributions. They then study predictive complexity, a generalization of Kolmogorov complexity. More precisely, they prove that when the predictive complexity of a restricted regular game exists, the average predictive complexity converges to the generalized entropy of the game almost everywhere with respect to the stationary ergodic distribution.