Theoretical Computer Science 473 (2013) 1–3

Contents lists available at SciVerse ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Preface

# Guest Editors' foreword

This special issue of *Theoretical Computer Science* is dedicated to the 21st International Conference on Algorithmic Learning Theory (ALT 2010) held in Canberra, Australia, October 6–8, 2010. It contains nine articles that were among the best in the conference.[1] The authors of these papers have been invited by the Special Issue Editors to submit completed versions of their work for this special issue. Once received, these papers underwent the usual refereeing process of *Theoretical Computer Science*. In the following, we briefly introduce each of the papers.

Algorithmic learning theory focuses on theoretical aspects of machine learning. It is dedicated to studies of learning from a mathematical and algorithmic perspective. Depending on the learning task considered, considerable interaction between various mathematical theories including statistics, probability theory, combinatorics, linguistics, and the theory of computation is required. These studies comprise the investigation of various formal models of machine learning and statistical learning, and the design and analysis of learning algorithms. This also leads to a fruitful interaction with the practical fields of machine learning, linguistics, psychology, and philosophy of science.

The first paper in this special issue belongs to the area of PAC learning. In this model the learner observes the data concerning a target concept generated according to an unknown probability distribution, but it does not have to figure out aspects of the concept to be learned which are unlikely to be observed. That is, when learning a concept $L$, the learner observes randomly drawn data according to some unknown probability distribution $D$ and the learner has to find with high probability a hypothesis $H$ such that $H$ is similar to $L$ with respect to the distribution $D$, i.e., $D(\{x \mid H(x) \neq L(x)\})$ is below a bound given to the algorithm as a parameter.

In their paper *Tighter PAC–Bayes bounds through distribution-dependent priors*, Lever, Laviolette, and Shawe-Taylor prove sharp risk bounds for stochastic exponential weight algorithms. The idea is here to base the analysis on a prior defined in terms of the data generating distribution. The authors derive a number of PAC–Bayes bounds for Gibbs classifiers using prior and posterior distributions which are defined, respectively, in terms of regularized empirical and true risks for a problem. The results rely on a key bound on the Kullback–Leibler divergence between distributions of this form. Furthermore, this bound introduces a new complexity measure.

The topic of Pestov's paper is already explained by its title, *PAC learnability under non-atomic measures: A problem by Vidyasagar*. In the PAC learning model the standard lower bounds for the sample complexity in terms of the Vapnik–Chervonenkis dimension and similar quantities are based on very adversarial probability distributions. These distributions are defined in such a way that all the probability is allocated to a small number of "difficult" points. So, it is only natural to consider the problem of how the lower bounds extend to cases where the distributions are a bit more reasonable. In 1997 Vidyasagar posed the problem of characterizing learnability under non-atomic distributions, where a distribution $D$ is said to be *non-atomic* if every set $A$ with $D(A) > 0$ has a subset $B$ with $0 < D(B) < D(A)$. Pestov resolves this problem by introducing the notion of Vapnik–Chervonenkis dimension modulo countable sets. This allows him to obtain a complete characterization. Note that instead of the usual measure-theoretic assumptions about the concept class, the author assumes Martin's Axiom, a set-theoretic axiom weaker than the continuum hypothesis. The results are also extended to the case of learning real-valued functions.

The next paper belongs to the area of learning probabilistic automata and of query learning. In the query learning model the learner aims to identify a concept which a teacher is teaching. So the learner is allowed to ask queries which the teacher has to answer truthfully, but not more helpfully than required. In most settings of query learning, the queries are of a fixed form. One type of such queries are statistical queries where an underlying distribution is assumed and the teacher returns a polynomial-time program which has — with respect to the underlying distribution — an error probability below a parameter given in the query. The paper *Learning probabilistic automata: A study in state distinguishability* by Balle, Castro, and

[1] The conference proceedings, including preliminary versions of these papers, appeared as "Algorithmic Learning Theory, 21st International Conference, ALT 2010, Canberra, Australia, October 6–8, 2010. Proceedings," *Lecture Notes in Artificial Intelligence*, vol. 6331, Springer, 2010.

0304-3975/$ – see front matter 
doi:10.1016/j.tcs.2012.10.007

Gavaldà distinguishes between algorithms that learn probabilistic deterministic finite automata from (a) independent and identically distributed samples from the target distribution, (b) statistical queries, and (c) statistical queries of a new type called $L_\infty$ queries. In particular, Balle et al. show that many existing algorithms can be reformulated in terms of type (b) and (c) algorithms, and that type (c) algorithms can be simulated using type (b) algorithms. Thus, many existing algorithms can be reduced to statistical query algorithms. Furthermore, a direct lower bound on the number of required type (c) queries is shown and a new type (c) algorithm is proposed that is efficient in the PAC sense. Moreover, the difficulty of learning distributions generated by probabilistic deterministic finite automata using statistical queries depends on a parameter $\mu$ which is quite frequently studied in the literature. The authors show that this parameter cannot be omitted without losing polynomial-time learnability for various important classes.

Koolen and de Rooij study in their paper *Switching investments* a learning problem from mathematical finance. The relevant learning model is on-line learning. The basic idea of on-line algorithms is that a learner combines expert advice in the process of decision making. In each round, the experts are asked which action to take, and then the learner makes its own decision based on this advice. Experts can be free agents or just decision or prediction strategies. As usual in on-line learning, the authors do not make any statistical assumptions about the financial market, except that it functions in discrete time. They consider a game between two players, *Investor* and *Nature*. At each point of time Investor is allowed to invest its capital (or a fraction of it) in the share, and Nature decides the price of the share. Koolen and de Rooij develop two algorithms, "mix" and "fix", which define trading strategies for Investor in a discrete-time setting. For each strategy an explicit lower bound on the payoff achieved by that strategy is provided. The bound is expressed in terms of the optimal payoff that could be achieved if the path of the share price was known in advance. The "mix" algorithm runs in time $O(T^2)$ and space $O(T)$, where $T$ is the total number of time periods considered. In contrast, the "fix" algorithm runs in time $O(T)$ and constant space. On the one hand, the "mix" algorithm offers more flexibility and may perform better. On the other hand, the "fix" algorithm provides better bounds.

The paper *Toward a classification of finite partial-monitoring games* by Antos, Bartók, Pál, and Szepesvári also studies on-line learning. A finite partial-monitoring game is a two-player game; the two players are called *Learner* and *Opponent*. In each round, Learner chooses one of finitely many actions and simultaneously Opponent chooses one of finitely many possible outcomes. Depending on the action and outcome, Learner receives a limited feedback. This could be the outcome, the loss of its chosen action, or some other function of the action/outcome pair. Furthermore, Learner suffers a loss. The goal of Learner is to choose the actions in such a way that the cumulative loss is small. However, since Opponent could choose the outcome sequence such that Learner suffers an arbitrarily high loss, it is too much to ask for an absolute guarantee for the cumulative loss. Therefore, the cumulative loss of Learner is compared with the cumulative loss of the best of the strategies which choose the same action in every round (constant strategies). The difference between the cumulative loss of Learner and the cumulative loss of the best constant strategy is called the regret. The authors make significant progress in classifying the games with two outcomes. In particular, they show that in this case there are just four possible regret rates, aptly called "trivial," "easy," "hard," and "hopeless." A geometric criterion for classifying games based on the loss and feedback matrices is developed and it is shown that the criterion is computationally efficient. Some of the regret rate bounds were known, and the missing ones are provided in this paper.

The following two papers study problems belonging to the area of algorithmic learning theory known as inductive inference of formal languages. The basic scenario is easily explained. A class $\mathcal{C}$ of recursively enumerable languages is called learnable from positive data if there is a recursive learner which can identify every language $L \in \mathcal{C}$ in the following sense. The learner, when receiving the elements of $L$ in arbitrary order from an infinite list (called text), outputs in parallel finitely many hypotheses such that the last of these generates $L$. Many variants of this notion of learning have been introduced and compared to one another.

Case and Kötzing investigate in their paper *Memory-limited non-U-shaped learning with solved open problems* the question of when U-shaped learning behavior can be avoided without losing learning power. Here by *U-shaped* behavior the following learning behavior is meant. When successively fed a text for a language $L$ in the target class $\mathcal{C}$, the learner at some time conjectures $L$, later conjectures a language different from $L$, and then returns at the end to a hypothesis that is correct for $L$. Such learning behavior has been observed for human language learning, and thus it is only natural to ask whether or not it is really necessary to achieve the best possible learning behavior. Addressing this problem in inductive inference, it was shown previously that U-shaped learning is *not* necessary for explanatory learning. On the other hand, it was also proved that for behaviorally correct learning, U-shapes are necessary. But for various other learning criteria, in particular those with limitations of the long term memory, this question of whether or not U-shaped learning is unavoidable remained open. Case and Kötzing solve many of these open problems.

The paper *Learning without coding* by Jain, Moelius, and Zilles also studies inductive inference of formal languages from positive data. The authors focus their research on iterative learning. Here an iterative learner is a learner which starts with a default hypothesis and maps each datum plus the old hypothesis to the new hypothesis; the hypothesis itself is the only memory the learner has of the previously observed data. So there may be some temptation for the learner to code observed data into the hypothesis, e.g., if the hypothesis space has a redundancy. The question is then to what extent such "coding tricks" are necessary. To answer this question the authors look for learning models which minimize such coding by the learners. First, they investigate to which extent one can overcome such behavior by requiring that the learner has to use a one-to-one hypothesis space. Additionally, Jain et al. generalize learnability by considering learners which are coded as enumeration operators and which do not need hypothesis spaces. One sample result of the authors is that such learners can infer various classes which cannot be learned iteratively; conversely there are also classes learnable using a one-to-one hypothesis space which are not learnable under this new model.

Next, we turn our attention to reinforcement learning. In reinforcement learning, a decision maker (agent) interacts with an environment (world) through an alternating sequence of actions and observations, including (occasional) rewards that should be maximized in the long run. The environment is stochastic and unknown and has to be learned. This setting encompasses most other learning scenarios, including active and passive learning.

In this context, it has been argued that the AIXI theory represents the first general and formal ''optimal'' but incomputable ''solution'' to this problem. The paper *Asymptotic non-learnability of universal agents with computable horizon functions* by Orseau challenges the optimality of AIXI. Note that the AIXI model extends the passive Solomonoff induction to the reinforcement learning framework. But unlike the case for the former, it is quite non-trivial to come up with notions of optimality that are simultaneously strong enough to be interesting and weak enough to be satisfiable by any agent at all. In particular, Orseau shows that active Bayesian agents suffer from the exploration/exploitation dilemma. More precisely, he proves that any Bayesian agent that takes its environment class to be the class of all computable environments (deterministic or stochastic) must eventually cease exploration. Using this insight, it is then shown that there exist environments for which the active Bayesian agent never converges to an optimal policy.

Finally, this special issue presents a paper that is at the intersection of two areas: on-line learning and kernel methods. Kernel methods constitute a powerful tool that is used in many areas of machine learning. The idea is to map the typically low-dimensional instances to a high-dimensional (often infinite-dimensional) *feature space*, where then the analysis can be carried out and the prediction can be performed. This is the so-called ''kernel trick.'' The method was popularized by the work of Vladimir Vapnik on support vector machines and has found numerous applications ever since. Popular methods used in the feature space are separating positive and negative examples with a large-margin hyperplane (in the case of classification) and fitting a linear function to the data (in the case of regression).

Zhdanov and Kalnishkan analyze in their paper *An identity for kernel ridge regression* properties of the popular method of kernel ridge regression as an on-line prediction algorithm. The main result of the paper is the equality between the quadratic loss (suitably reduced) of the kernel ridge regression algorithm applied in the on-line mode and the quadratic loss of the best regressor (suitably penalized). This new identity makes it possible to derive, in an elegant way, upper bounds for the cumulative quadratic loss of online kernel ridge regression.

We would like to express our immense gratitude to all authors for contributing their papers to this special issue and for all their efforts to improve and to polish them. Furthermore, we would like to thank all the referees for their fine reports and their efficient work. We are very grateful to all the members of the ALT 2010 Program Committee for selecting the papers. Finally, we are particularly grateful to Giorgio Ausiello for the opportunity to compile this special issue.

Canberra, Singapore, London, and Sapporo, September 4, 2012.

*Special Issue Editors*
Marcus Hutter
Frank Stephan
Vladimir Vovk
Thomas Zeugmann *
*E-mail address:* thomas@ist.hokudai.ac.jp (T. Zeugmann).

* Corresponding editor.