

Editors' Introduction

Marcus Hutter, Frank Stephan, Vladimir Vovk, and Thomas Zeugmann

The conference “Algorithmic Learning Theory 2010” is dedicated to studies of learning from a mathematical and algorithmic perspective. Researchers consider various abstract models of the problem of learning and investigate how the learning goal in such a setting can be formulated and achieved. These models describe ways to define

- the goal of learning,
- how the learner retrieves information about its environment,
- how to form of the learner’s models of the world (in some cases).

Retrieving information in some models is passive where the learner just views a stream of data. In other models, the learner is more active, asking questions or learning from its actions. Besides explicit formulation of hypotheses in an abstract language with respect to some indexing system, there are also more implicit methods like making predictions according to the current hypothesis on some arguments which then are evaluated with respect to their correctness, and wrong predictions (coming from wrong hypotheses) incur some loss on the learner. In the following, a more detailed introduction is given to the five invited talks and then to the regular contributions.

Invited Talks. The 5 joint invited speakers of the conferences ALT 2010 and DS 2010 are eminent researchers in their fields and give either an introduction to their specific research area or talk about a topic of wide general interest.

Alexander Clark (Royal Holloway University of London) received his bachelor degree from Trinity College, Cambridge, and his Ph.D. from the University of Sussex. He has throughout his career put a large emphasis on applying theoretical insights to solve the corresponding practical problems. In particular, he studied the unsupervised learning of natural languages; his findings are also relevant to first language acquisition in humans. His work included finding the right definition of learnability in various linguistic contexts, designing learning algorithms and implementing them. These algorithms were tested on both synthetic and natural examples. He also studied the learnability of regular languages and context-free languages; a sample result, obtained in collaboration with Franck Thollard, is that the class of regular languages can be PAC-learned using a polynomial amount of data and processing time, provided that the distributions of the samples are restricted to be generated by one of a large family of related probabilistic deterministic finite state automata. In his invited talk *Towards General Algorithms for Grammatical Inference*, Alexander Clark deals with the learning of context-free languages and multiple context-free languages. He formulates a general framework for a large class of learning algorithms for such languages and, using this framework, he reviews Angluin’s classical LSTAR algorithm and compares it with various contemporary approaches.

Manfred K. Warmuth (University of California at Santa Cruz) is a leading expert in computational learning theory. In his groundbreaking article “Learnability and the Vapnik-Chervonenkis dimension”, he showed, jointly with Anselm Blumer, Andrzej Ehrenfeucht and David Haussler, that a class is PAC learnable if and only if it has finite Vapnik-Chervonenkis dimension, which established a close connection between these two fields. He further showed, jointly with Leonard Pitt, that there is no approximation algorithm to find the minimum consistent DFA within a polynomial bound. He made important contributions to the boosting algorithm, a very successful method to construct powerful learners. One of his main interests is on-line learning, where he originated several novel approaches and obtained fundamental theoretical results. His most recent work applies theoretical insights to the study of evolution. In his invited talk *The Blessings and the Curse of the Multiplicative Updates*, Manfred K. Warmuth considers learning settings in which parameters are updated in a multiplicative way; the advantage is that the importance of major patterns might grow exponentially, the disadvantage is that the importance of some minor pattern might go down too fast so that this pattern is wiped out although the information it contains might be needed later. The talk describes how modern machine learning algorithms try to preserve relevant information and compares this to the strategies nature has to preserve relevant genetic information during evolution. In his conclusion, the author states that there are still various strategies which one can take over from nature in order to use them in learning algorithms.

Ivan Bratko (University of Ljubljana) received his bachelor, masters and doctoral degrees from the University of Ljubljana. He is an eminent researcher in machine learning, knowledge-based systems, heuristic programming, qualitative modelling, intelligent robotics and computer chess. He is the author of the standard reference “Prolog Programming for Artificial Intelligence”, which has been translated into German, Italian, French, Slovenian, Japanese and Russian; he has furthermore authored about 200 research articles.

Rao Kotagiri (University of Melbourne) received the Bachelor of Engineering from Andhra University, the Master of Engineering from the Indian Institute of Science and the Ph.D. from Monash University. He is a leading expert in machine learning and data mining, robust agent systems, information retrieval, intrusion detection, logic programming and deductive databases, distributed systems, bioinformatics and medical imaging.

Peter L. Bartlett (University of California at Berkeley) is an outstanding researcher in the areas of machine learning and statistical learning theory. Jointly with Martin Anthony, he co-authored the excellent textbook “Learning in Neural Networks: Theoretical Foundations”. For his work in statistical learning theory, in 2001 he was awarded the Malcolm McIntosh Prize for the Physical Scientist of the Year in Australia. His research interests include, besides neural networks and statistical learning theory, also privacy and security aspects of learning, on-line learning algorithms and kernel methods. In his invited talk *Optimal Online Prediction in Adversarial Environments*, Peter L. Bartlett deals with prediction in statistical settings and investigates strategies to minimise the regret in a

prediction game against an adversarial environment. He explains that, although not every environment is adversarial, it is often the mathematically most elegant way to model the prediction situation.

Statistical Learning. Statistical learning theory studies methods of assigning labels to data vectors under the statistical assumption that the labelled data vectors are generated independently by an arbitrary unknown probability distribution. Sometimes, the labels depend only on a small fraction of the variables present while most components of the data vectors are irrelevant; the area of feature selection studies methods of selecting the relevant variables (vector components) so that in the future the learning algorithm can concentrate on the data vectors compressed in this way and so easier to label. In active learning, the learner does not get the labels together with the training data; instead the learner has to request the labels from a teacher. Naturally this can be done only with a small fraction of the data presented. Boosting is a method which improves the properties of a learner by combining various primitive learners into a better one.

Pierre Alquier's paper *An Algorithm for Iterative Selection of Blocks and Features* is about selecting variables from very long vectors of variables where in the data, almost all variables are 0 and neighbouring variables are with high probability equal. This topic has been studied previously, but its theoretical treatment so far has been insufficient. To obtain better results in the area, the author proposes an alternative approach, based on the Iterative Feature Selection method. This method is based on an iterative algorithm which takes the general form of the vector to be learnt into account, but does not know the positions where the blocks start and end. The algorithm improves the statistical performance of its current guess (estimator) at each step. The obtained results are justified both theoretically and through simulations on practically relevant data.

Liu Yang, Steve Hanneke and Jaime Carbonell study in their paper *Bayesian Active Learning Using Arbitrary Binary Valued Queries* how to learn a concept to precision ϵ using as few binary queries as possible. The authors provide an upper and lower bound on how many queries may be required to learn successfully. The model is generalised from the usual one in the sense that arbitrary binary valued queries are taken into consideration and not only membership queries. The analysis is Bayesian in the sense that the bound depends on a prior distribution on the concept space.

In their paper *Approximation Stability and Boosting*, Wei Gao and Zhi-Hua Zhou revisit the notion of stability for boosting algorithms. It is known that algorithms like AdaBoost have almost-everywhere uniform stability when the base algorithm has L_1 stability. The latter is however too restrictive: the authors show that AdaBoost using such a learner becomes a constant learner unless the underlying algorithm is a real-valued learner. Therefore the authors dedicate themselves to the question on what can be said when the base learner is not real-valued. For this analysis, they introduce a property called "approximation stability". They show that AdaBoost has this property and prove that this

property is sufficient for generalisation and for the learnability of asymptotic empirical risk minimisation in the general learning setting.

Grammatical Inference and Graph Learning. This section is dedicated to learning specific structures such as formal languages, trees or graphs. These types of structures are important in mathematics and computer science and also play an important role in learning theory. Most prominent graphs occur in the internet and social networks; for example, search machines collect a lot of information on the graph structure of the internet where nodes are given by webpages and edges by the hyperlinks.

In their paper *A Spectral Approach for Probabilistic Grammatical Inference of Trees*, Raphaël Bally, François Denis and Amaury Habrard consider distributions over the set of trees which are computed by weighted automata. This is a quite natural class of distributions which has an algebraic characterisation. By concentrating on the finite dimensional subspace containing all the residuals of such a distribution, the authors find an approach which allows them to define a global solution for their inference problem, so that they can avoid to construct the automaton to be built iteratively step by step.

Balázs Csáji, Raphaël Jungers and Vincent Blondel dedicate their paper *PageRank Optimization in Polynomial Time by Stochastic Shortest Path Reformulation* to the question on how a member-node of a network can increase its importance and visibility by small modifications of the network. The underlying idea is that the nodes in the network are evaluated using the well-known PageRank algorithm which, roughly speaking, assigns to every node the expected time which one spends on the node during a random walk. The task is now the following: Given a set of possible new edges, one has to select a fixed number of them and add them to the network in order to increase the PageRank. Csáji, Jungers and Blondel show that the general problem on how to select these edges is polynomial time solvable; they do this by reformulating the algorithm as a stochastic shortest path problem and they then show that this new problem is well-suited for reinforcement learning methods.

Dana Angluin, James Aspnes and Lev Reyzin explore in their paper *Inferring Social Networks from Outbreaks* a learning setting which stems from the study of diseases. In a network, a disease might have travelled along the edges of the network. Hence whenever there is an outbreak of the disease, the disease is tracked down at the locations of its appearance and the observed locations can be considered to be connected through the network. The learning task is to build a model of the network which explains how during an outbreak the illness propagated in the network. Formally, an outbreak is a set S_i of nodes, called constraints, and the goal of the learner is to find a subset E of edges, called connections, such that each constraint S_i is connected by those members of E which are between nodes from S_i ; then it is assumed that diseases can travel along the so selected edges. Here the choice of E should be optimised with respect to its minimum loglikelihood cost. In the off-line learning problem, the learner receives all the constraints S_i at the start of the algorithm; in the on-line learning problem, the learner reads the constraints one by one and each time

has to add some edges in order to meet the constraint immediately. Angluin, Aspnes and Reyzin obtain a lower bound of $\Omega(\log(n))$ for the off-line version of the problem and an upper $O(n \log(n))$ bound for the on-line version. Better bounds are obtained for various special cases.

Probably Approximately Correct Learning. The basic idea of PAC learning is that the learner observes the data according to a distribution, but it does not need to figure out aspects of the concept to be learnt which are unlikely to be observed. In other words, 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 , that is, $D(\{x : H(x) \neq L(x)\})$ is below a bound given to the algorithm as a parameter.

Guy Lever, François Laviolette and John Shawe-Taylor derive in their paper *Distribution-Dependent PAC-Bayes Priors* a number of PAC-Bayes bounds for Gibbs classifiers using prior and posterior distributions which are defined, respectively, in terms of regularised empirical and true risks for a problem. The results rely on a key bound on the Kolmogorov-Loveland divergence between distributions of this form; this bound introduces a new complexity measure.

The purpose of Vladimir Pestov's work is explained already in its title, *PAC Learnability of a Concept Class under Non-Atomic Measures: A Problem by Vidyasagar*. The characterisation of PAC learnability under the class of all non-atomic measures is achieved by introducing an appropriate combinatorial parameter modifying the Vapnik-Chervonenkis dimension. This is a natural problem (in fact it was asked by Vidyasagar 13 years ago) and the solution is non-trivial, involving techniques from set theory and measure theory.

In *A PAC-Bayes Bound for Tailored Density Estimation*, Matthew Higgs and John Shawe-Taylor consider the problem of density estimation with an unusual twist: they want their solution to be tailored to the larger inference process of which this problem is part. Formalization of this idea involves the theory of reproducing kernel Hilbert spaces. Error bounds are stated in the framework of PAC-Bayes theory, which is standard in the case of classification but rarely used in density estimation.

In their paper *Compressed Learning with Regular Concept*, Jiawei Lv, Jianwen Zhang, Fei Wang, Zheng Wang and Changshui Zhang study the PAC learnability of half spaces where, for any given distribution, only those half spaces are considered where the measure of the bounding hyperplane is 0. The learning is called compressed as some of the components of the data presented are replaced by a random value; such compression is done for two reasons: (a) privacy as parts of the data might otherwise reveal information which should not be compromised; (b) efficiency reasons in order to reduce the overfitting in the learning process. The algorithm used is the voted-perceptron algorithm invented by Freund and Schapire.

Query Learning and Algorithmic Teaching. The basic scenario is that a learner wants to learn a concept which a teacher is teaching; in that scenario the

learner has to be successful whenever the teacher satisfies the minimum requirements, that is, gives correct answers although those need not to be more helpful than required. In most settings of query learning, the queries are of a fixed form, for example the learner can ask equivalence queries to which the teacher provides a counter example in the case that the hypothesis does not match the concept to be learnt and also membership queries where the teacher answers “yes” or “no”, depending on whether the queried item is an element of the concept to be learnt or not. More recent research also looks at 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.

Borja Balle, Jorge Castro and Ricard Gavaldà investigate in their paper *A Lower Bound for Learning Distributions Generated by Probabilistic Automata* the limitations on the learnability of certain distributions. These distributions are generated by probabilistic deterministic finite automata (PDFA). The authors show that the learnability of such distributions using statistical queries depends on a parameter μ which is quite frequently studied in the literature, and they show that this parameter cannot be omitted without losing polynomial time learnability for various important classes; in other words, the number of queries needed depends on this parameter μ . For their results, they use in addition to statistical queries also a new variant of these called L_∞ -queries.

Dana Angluin, David Eisenstat, Aryeh Kontorovich and Lev Reyzin study *Lower Bounds on Learning Random Structures with Statistical Queries*. The researchers consider randomly composed DNF formulas, randomly selected decision trees of logarithmic depth and randomly constructed deterministic finite automata. They show that each of these three concept cannot be weakly learned with a polynomial number of statistical queries, where the underlying distribution on the examples is arbitrary.

The paper *Recursive Teaching Dimension, Learning Complexity and Maximum Classes* by Thorsten Doliwa, Hans Ulrich Simon and Sandra Zilles deals with the recursive teaching dimension, which is the smallest number n such that for each concept C in the class to be learnt there are n examples x_1, x_2, \dots, x_n such that C is the only concept D in the class satisfying $C(y) = D(y)$ for all $y \in \{x_1, x_2, \dots, x_n\}$. The authors show, among other results, that for maximum classes the recursive teaching dimension equals the Vapnik-Chervonenkis dimension. In addition the authors show that the sequences defining the recursive teaching dimension also arise from various famous algorithms.

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. Typical results in this areas are relative loss bounds: the goal is to design prediction algorithms that are guaranteed to suffer a loss that is not much worse than the loss suffered by the best experts. To achieve this goal, the learner

keeps some statistics on the reliability of each expert, which is taken into account when making decisions.

Student Gábor Bartók received the *E. Mark Gold Award* for his paper *Toward a Classification of Finite Partial-Monitoring Games* which is joint work with Dávid Pál and Csaba Szepesvári. A finite partial-monitoring game is a two player game; the two players are called Learner and Nature in order to express that a learner explores and studies its natural environment which reacts to the learner's actions. In this game, Learner repeatedly chooses one of finitely many actions and Nature reacts to the learner by choosing one of finitely many possible outcomes. Depending on the action and outcome, the learner receives a feedback signal and suffers a loss; the goal of the learner is to choose the actions such that the overall loss is minimised. The authors make significant progress in classifying the games with two outcomes.

The paper *Switching Investments* by Wouter K. Koolen and Steven de Rooij is devoted to mathematical finance. As usual in on-line learning, the authors do not make any statistical assumptions about the financial market, and design investment algorithms competitive with a wide class of investment strategies that “buy low and sell high”. One of their algorithms, in addition, possesses linear time and space complexity.

Alexey Chernov and Fedor Zhdanov explore in their paper *Prediction with Expert Advice under Discounted Loss* a relatively new type of performance guarantees in on-line learning. In the standard approach, the learner's goal is to be competitive with the best experts according to the learner's and experts' cumulative losses. Chernov and Zhdanov, following earlier work by Freund and Hsu, establish similar results for cumulative discounted losses, where more recent losses are taken with greater weights. The framework of discounted losses provides an elegant alternative to Herbster and Warmuth's framework of “tracking the best expert”.

Jacob Abernethy, Peter Bartlett, Niv Buchbinder and Isabelle Stanton address in their paper *A Regularization Approach to Metrical Task Systems* the construction of randomised on-line algorithms for metrical task systems, where the learner always follows one expert and where it incurs a cost for switching from one expert to another. In the general case, the costs are an arbitrary metric among states. The authors restrict themselves to the class of “work-based” algorithms and obtain for this special case various algorithms.

Inductive Inference. The basic scenario of inductive inference is that 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: when reading the elements of L in arbitrary order from an infinite list, the learner 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 with each other, and the topic remains active more than 40 years after Gold's papers started it.

John Case and Timo Kötzing investigate in their paper *Solutions to Open Questions for Non-U-Shaped Learning with Memory Limitations* the question

when U-shaped learning behaviour can be avoided without losing learning power. Here a learner is said to be *U-shaped* on a text for a language L in the class to be learnt if at some time it conjectures L , later conjectures a language different from L and at the end returns to conjecturing L . For basic learning criteria it is known whether U-shaped learning behaviour can be avoided: in the case of explanatory learning the answer is “yes” and in the case of behaviourally correct learning the answer is “no”. But for various other learning criteria, in particular those with limitations of the long term memory, this question remained open. Case and Kötzing solve several of these open questions. One sample result is that there are classes which are learnable with three memory states such that every learner using only finitely many memory states for these classes has U-shaped learning behaviour on some text for some language to be learnt.

Samuel E. Moelius III and Sandra Zilles study in their paper *Learning Without Coding* notions of iterative learning which hinder or reduce the abilities of the learner to code. 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. As there is some temptation for the learner to code observed data into the hypothesis, the authors look for learning models which minimise such coding by the learners. The authors investigate to which extent one can overcome such behaviour by requiring that the learner uses a one-to-one hypothesis space. Furthermore, they generalise 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 learnt iteratively; conversely there are also classes learnable using a one-to-one hypothesis space which are not learnable under this new model.

Mahito Sugiyama, Eiju Hirowatari, Hideki Tsuiki and Akihiro Yamamoto give in their paper *Learning Figures with the Hausdorff Metric by Fractals* a theoretical foundation in the framework of inductive inference for learning with discretisation of analog data. They study the learnability of geometric figures, that is, fractals. The two main learnability notions employed are identification in the limit as well as closer and closer approximations of the object to be learnt where the approximations are measured with a Hausdorff metric.

Sanjay Jain and Efim Kimber analyse in their paper *Inductive Inference of Languages from Samplings* the scenario where the learner is not given all the data about the set to be learnt but only some part of it. In prior work they studied the scenario where for every language L to be learnt and every subset X of L , the learner has to converge on a text for X to a target language L' satisfying $X \subseteq L' \subseteq L$. In this paper, rather than considering all subsets of the target language, they consider A -sampling of the target language, where A is some fixed set and A -sampling of a language L is the set formed by taking the i th least-elements of L , for $i \in A$. Results show that the choice of A has a large influence on the learnability of the class. Furthermore, the authors consider

when such a learner can be constructed independently of or uniformly in A , for a collection of such sets A .

Reinforcement Learning. In reinforcement learning, a decision maker (agent) interacts with an environment (world) by an alternating sequence of actions and observations, including (occasional) rewards that should be maximised 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.

It has been argued that the AIXI theory represents the first general and formal “optimal” but incomputable “solution” to this problem. Laurent Orseau in his paper *Optimality Issues of Universal Greedy Agents with Static Priors* challenges the optimality of AIXI. Unlike passive Solomonoff induction 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. One suggested notion is to require that the probability of a suboptimal actions tends to zero, where an action is called suboptimal if it differs from the optimal action of the informed agent on the same history. Environments with this property are called asymptotically learnable. Orseau shows that there exist histories and environments that AIXI cannot learn asymptotically, hence establishing that this optimality notion is too strong.

At the other end of the spectrum are efficient but limited reinforcement learning algorithms. In particular, efficient learning and planning algorithms exist for completely observable finite state Markov decision processes (MDPs). Real-world problems can often be approximately modeled or reduced to finite MDPs. A natural idea is to formally define the quality of a reduction and to automatically learn good reductions by optimizing the quality criterion. Peter Sunehag and Marcus Hutter in their paper *Consistency of Feature Markov Processes* investigate a recently introduced such criterion. They show asymptotic consistency in the sequence prediction case, and extend their result to prediction with side information and to the active case.

Multi-armed bandit problems can be regarded as (reinforcement) learning problems with a single state. Despite their apparent simplicity, they constitute prototypical active learning problems that already require trading off exploration and exploitation. Taishi Uchiya, Atsuyoshi Nakamura and Mineichi Kudo in their paper *Algorithms for Adversarial Bandit Problems with Multiple Plays* consider the non-stochastic / online / adversarial setting and the case where multiple arms are played simultaneously, which is relevant, e.g., for multiple advertisement placement. They analyze and present bounds for extensions of the Exp3 and CompBand algorithms in terms of time and space efficiency and regret for the best fixed action set.

On-line Learning and Kernel Methods. The papers in this part of the proceedings are in the intersection of two areas: on-line learning and kernel methods. Kernel methods were popularised in machine learning by the work of Vladimir Vapnik on support vector machines, but later became a powerful

tool used in many other areas of learning theory. The basic idea is the so-called “kernel trick”: the instances (typically low-dimensional) are mapped to a high-dimensional (often infinite-dimensional) *feature space*, where the prediction is done and its analysis is performed. 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). Even conventional linear methods become a powerful tool when applied to high-dimensional feature spaces, and when mapped back to the original instance space, they may become highly non-linear and yet computationally efficient.

The paper *Online Multiple Kernel Learning: Algorithms and Mistake Bounds* by Rong Jin, Steven Hoi and Tianbao Yang constructs a number of on-line kernel algorithms for classification and provides relative loss bounds for them. Its goal is to merge classifiers based on several different kernels. The performance of the resulting algorithms should be comparable with that of the algorithm based on the best kernel; this is a non-trivial problem since which kernel is best becomes known only after we see the data. The authors construct both deterministic and randomised versions of such algorithms, the latter achieving computational efficiency by applying ideas from the popular area of “multi-armed bandits” (see above).

Fedor Zhdanov and Yuri Kalnishkan analyze in their work *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 penalised). This new identity makes it possible to derive, in an elegant way, upper bounds for the cumulative quadratic loss of online kernel ridge regression.