Editors' Introduction

Peter Auer, Alexander Clark, Thomas Zeugmann, and Sandra Zilles

The aim of the series of conferences on Algorithmic Learning Theory (ALT) is to look at learning from an algorithmic and mathematical perspective. Over time several models of learning have been developed which study different aspects of learning. In the following we describe in brief the invited talks and the contributed papers for ALT 2014 held in Bled, Slovenia..

Invited Talks. Following the tradition of the co-located conferences ALT and DS all invited lectures are shared by the two conferences. The invited speakers are eminent researchers in their fields and present either their specific research area or lecture about a topic of broader interest.

This year's joint invited speaker for ALT 2014 and DS 2014 is Zoubin Ghahramani, who is Professor of Information Engineering at the University of Cambridge, UK, where he leads a group of about 30 researchers. He studied computer science and cognitive science at the University of Pennsylvania, obtained his PhD from MIT in 1995 under the supervision of Michael Jordan, and was a postdoctoral fellow at the University of Toronto with Geoffrey Hinton. His academic career includes concurrent appointments as one of the founding members of the Gatsby Computational Neuroscience Unit in London, and as a faculty member of CMU's Machine Learning Department for over 10 years. His current research focuses on nonparametric Bayesian modeling and statistical machine learning. He has also worked on applications to bioinformatics, econometrics, and a variety of large-scale data modeling problems. He has published over 200 papers, receiving 25,000 citations (an h-index of 68). His work has been funded by grants and donations from EPSRC, DARPA, Microsoft, Google, Infosys, Facebook, Amazon, FX Concepts and a number of other industrial partners. In 2013, he received a $750,000 Google Award for research on building the Automatic Statistician. In his invited talk *Building an Automated Statistician* (joint work with James Lloyd, David Duvenaud, Roger Grosse, and Josh Tenenbaum) Zoubin Ghahramani addresses the problem of abundant data and the increasing need for methods to automate data analysis and statistics. The Automated Statistician project aims to automate the exploratory analysis and modeling of data. The approach uses Bayesian marginal likelihood computations to search over a large space of related probabilistic models. Once a good model has been found, the Automated Statistician generates a natural language summary of the analysis, producing a 10-15 page report with plots and tables describing the analysis. Zoubin Ghahramani discusses challenges such as: how to trade off predictive performance and interpretability, how to translate complex statistical concepts into natural language text that is understandable by a numerate non-statistician, and how to integrate model checking.

The invited speaker for ALT 2014 is Luc Devroye, who is a James McGill Professor in the School of Computer Science of McGill University in Montreal. He studied at Katholieke Universiteit Leuven and subsequently at Osaka University and in 1976 received his PhD from University of Texas at Austin under the supervision of Terry Wagner. Luc Devroye specializes in the probabilistic analysis of algorithms, random number generation and enjoys typography. Since joining the McGill faculty in 1977 he has won numerous awards, including an E.W.R. Steacie Memorial Fellowship (1987), a Humboldt Research Award (2004), the Killam Prize (2005) and the Statistical Society of Canada gold medal (2008). He received an honorary doctorate from the Université catholique de Louvain in 2002, and an honorary doctorate from Universiteit Antwerpen in 2012. The invited paper *Cellular Tree Classifiers* (joint work with Gerard Biau) deals with classification by decision trees, where the decision trees are constructed recursively by using only two local rules: (1) given the input data to a node, it must decide whether it will become a leaf or not, and (2) a non-leaf node needs to decide how to split the data for sending them downstream. The important point is that each node can make these decisions based only on its local data, such that the decision tree construction can be carried out in parallel. Somewhat surprisingly there are such local rules that guarantee convergence of the decision tree error to the Bayes optimal error. Luc Devroye discusses the design and properties of such classifiers.

The ALT 2014 tutorial speaker is Eyke Hüllermeier, who is professor and head of the Intelligent Systems Group at the Department of Computer Science of the University of Paderborn. He received his PhD in Computer Science from the University of Paderborn in 1997 and he also holds an MSc degree in business informatics. He was a researcher in artificial intelligence, knowledge-based systems, and statistics at the University of Paderborn and the University of Dortmund and a Marie Curie fellow at the Institut de Recherche en Informatique de Toulouse. He has held already a full professorship in the Department of Mathematics and Computer Science at Marburg University before rejoining the University of Paderborn. In his tutorial *A Survey of Preference-based Online Learning with Bandit Algorithms* (joint work with Róbert Busa-Fekete) Eyke Hüllermeier reports on learning with bandit feedback that is weaker than the usual real-value reward. When learning with bandit feedback the learning algorithm receives feedback only from the decisions it makes, but no information from other alternatives. Thus the learning algorithm needs to simultaneously explore and exploit a given set of alternatives in the course of a sequential decision process. In many applications the feedback is not a numerical reward signal but some weaker information, in particular relative preferences in the form of qualitative comparisons between pairs of alternatives. This observation has motivated the study of variants of the multi-armed bandit problem, in which more general representations are used both for the type of feedback to learn from and the target of prediction. The aim of the tutorial is to provide a survey of the state-of-the-art in this area which is referred to as preference-based multi-armed bandits. To this end, Eyke Hüllermeier provides an overview of problems that have been considered in the literature as well as methods for tackling them. His systematization is mainly based on the assumptions made by these methods about the data-generating process and, related to this, the properties of the preference-based feedback.

The DS 2014 tutorial speaker is Anuška Ferligoj, who is professor of Multivariate Statistical Methods at the University of Ljubljana. She is a Slovenian mathematician who earned international recognition by her research work on network analysis. Her interests include multivariate analysis (constrained and multicriteria clustering), social networks (measurement quality and blockmodeling), and survey methodology (reliability and validity of measurement). She is a fellow of the European Academy of Sociology. She has also been an editor of the journal Advances in Methodology and Statistics (Metodoloski zvezki) since 2004 and is a member of the editorial boards of the Journal of Mathematical Sociology, Journal of Classification, Social Networks, Statistic in Transition, Methodology, Structure and Dynamics: eJournal of Anthropology and Related Sciences. She was a Fulbright scholar in 1990 and visiting professor at the University of Pittsburgh. She was awarded the title of Ambassador of Science of the Republic of Slovenia in 1997. Social network analysis has attracted considerable interest from the social and behavioral science community in recent decades. Much of this interest can be attributed to the focus of social network analysis on relationship among units, and on the patterns of these relationships. Social network analysis is a rapidly expanding and changing field with broad range of approaches, methods, models and substantive applications. In her tutorial *Social Network Analysis* Anuška Ferligoj gives a general introduction to social network analysis and an overview of tasks and corresponding methods, accompanied by pointers to software for social network analysis.

Inductive Inference. There are a number of papers in the field of inductive inference, the most classical branch of algorithmic learning theory. First, *A Map of Update Constraints in Inductive Inference* by Timo Kötzing and Raphaela Palenta provides a systematic overview of various constraints on learners in inductive inference problems. They focus on the question of which constraints and combinations of constraints reduce the learning power, meaning the class of languages that are learnable with respect to certain criteria.

On a related theme, the paper *On the Role of Update Constraints and Text-Types in Iterative Learning* by Sanjay Jain, Timo Kötzing, Junqi Ma, and Frank Stephan looks more specifically at the case where the learner has no memory beyond the current hypothesis. In this situation the paper is able to completely characterize the relations between the various constraints.

The paper *Parallel Learning of Automatic Classes of Languages* by Sanjay Jain and Efim Kinber continues the line of research on learning automatic classes of languages initiated by Jain, Luo and Stephan in 2012, in this case by considering the problem of learning multiple distinct languages at the same time.

Laurent Bienvenu, Benoît Monin and Alexander Shen present a negative result in their paper *Algorithmic Identification of Probabilities is Hard.* They show that it is impossible to identify in the limit the exact parameter—in the sense of the Turing code for a computable real number—of a Bernoulli distribution, though it is of course easy to approximate it.

Exact Learning from Queries. In cases where the instance space is discrete, it is reasonable to aim at exact learning algorithms where the learner is required to produce a hypothesis that is exactly correct.

The paper winning the E.M. Gold Award, *Learning Boolean Halfspaces with Small Weights from Membership Queries* by the student authors Hasan Abasi and Ali Z. Abdi and co-authored by Nader H. Bshouty, presents a significantly improved algorithm for learning Boolean Halfspaces in $\{0,1\}^n$ with integer weights $\{0,\ldots,t\}$ from membership queries only. It is shown that this algorithm needs only $n^{O(t)}$ membership queries, which improves over previous algorithms with $n^{O(t^5)}$ queries and closes the gap to the known lower bound n^t.

The paper by Hasan Abasi, Nader H. Bshouty and Hanna Mazzawi *On Exact Learning Monotone DNF from Membership Queries* presents learning results on learnability by membership queries of monotone DNF (disjunctive normal forms) with a bounded number of terms and a bounded number of variables per term.

Dana Angluin and Dana Fisman look at exact learning using membership queries and equivalence queries in their paper *Learning Regular Omega Languages.* Here the class concerned is that of regular languages over infinite words; the authors consider three different representations which vary in their succinctness. This problem has applications in verification and synthesis of reactive systems.

Reinforcement Learning. Reinforcement learning continues to be a centrally important area of learning theory, and this conference contains a number of contributions in this field. Ronald Ortner, Odalric-Ambrym Maillard and Daniil Ryabko present a paper *Selecting Near-Optimal Approximate State Representations in Reinforcement Learning*, which looks at the problem where the learner does not have direct information about the states in the underlying Markov Decision Process (MDP); in contrast to Partially Observable MDPs, here the information is via various models that map the histories to states.

L.A. Prashanth considers risk constrained reinforcement learning in his paper *Policy Gradients for CVaR-Constrained MDPs*, focusing on the stochastic shortest path problem. For a risk constrained problem not only the expected sum of costs per step $\mathbb{E}[\sum_m g(s_m, a_m)]$ is to be minimized, but also the sum of an additional cost measure $C = \sum_m c(s_m, a_m)$ needs to be bounded from above. Usually the Value at Risk, $\mathrm{VaR}_\alpha = \inf\{\xi | \mathbb{P}(C \leq \xi) \geq \alpha\}$, is constrained, but such constrained problems are hard to optimize. Instead, the paper proposes to constrain the Conditional Value at Risk, $\mathrm{CVaR}_\alpha = \mathbb{E}[C | C \geq \mathrm{VaR}_\alpha]$, which allows to apply standard optimization techniques. Two algorithms are presented that converge to a locally risk-optimal policy using stochastic approximation, mini batches, policy gradients, and importance sampling.

In contrast to the usual MDP setting for reinforcement learning, the two following papers consider more general reinforcement learning. *Bayesian Reinforcement Learning with Exploration* by Tor Lattimore and Marcus Hutter improves some of their earlier work on general reinforcement learning. Here the true environment does not need to be Markovian, but it is known to be drawn at random from a finite class of possible environments. An algorithm is presented that alternates between periods of playing the Bayes optimal policy and periods of forced experimentation. Upper bounds on the sample complexity are established, and it is shown that for some classes of environments this bound cannot be improved by more than a logarithmic factor.

Marcus Hutter's paper *Extreme State Aggregation beyond MDPs* considers how an arbitrary (non-Markov) decision process with a finite number of actions can be approximated by a finite-state MDP. For a given feature function $\phi : H \to S$ mapping histories h of the general process to some finite state space S, the transition probabilities of the MDP can be defined appropriately. It is shown that the MDP approximates the general process well if the optimal policy for the general process is consistent with the feature function, $\pi^*(h_1) = \pi^*(h_2)$ for $\phi(h_1) = \phi(h_2)$, or if the optimal Q-value function is consistent with the feature function, $|Q^*(h_1, a) - Q^*(h_2, a)| < \varepsilon$ for $\phi(h_1) = \phi(h_2)$ and all a. It is also shown that such a feature function always exists.

Online Learning and Learning with Bandit Information. The paper *On Learning the Optimal Waiting Time* by Tor Lattimore, András György, and Csaba Szepesvári, addresses the problem of how long to wait for an event with independent and identically distributed (i.i.d.) arrival times from an unknown distribution. If the event occurs during the waiting time, then the cost is the time until arrival. If the event occurs after the waiting time, then the cost is the waiting time plus a fixed and known amount. Algorithms for the full information setting and for bandit information are presented that sequentially choose waiting times over several rounds in order to minimize the regret in respect to an optimal waiting time. For bandit information the arrival time is only revealed if it is smaller than the waiting time, and in the full information setting it is revealed always. The performance of the algorithms nearly matches the minimax lower bound on the regret.

In many application areas, e.g. recommendation systems, the learning algorithm should return a ranking: a permutation of some finite set of elements. This problem is studied in the paper by Nir Ailon, Kohei Hatano, and Eiji Takimoto titled *Bandit Online Optimization Over the Permutahedron* when the cost of a rankings is calculated as $\sum_{i=1}^{n} \pi(i)s(i)$, where $\pi(i)$ is the rank of item i and $s(i)$ is its cost. In the bandit setting in each iteration an unknown cost vector $\mathbf{s}_t$ is chosen, and the goal of the algorithm is to minimize the regret in respect to the best fixed ranking of the items.

Marcus Hutter's paper *Offline to Online Conversion* introduces the problem of turning a sequence of distributions q_n on strings in X^n, $n = 1, \ldots, n$, into a stochastic online predictor for the next symbol $\tilde{q}(x_n|x_1, \ldots, x_{n-1})$, such that the induced probabilities $\tilde{q}(x_1, \ldots, x_n)$ are close to $q_n(x_1, \ldots, x_n)$ for all sequences $x_1, x_2, \ldots$ The paper considers four strategies for doing such a conversion, showing that naïve approaches might not be satisfactory but that a good predictor can always be constructed, at the cost of possible computational inefficiency. One examples of such a conversion gives a simple combinatorial derivation of the Good-Turing estimator.

Statistical Learning Theory. Andreas Maurer's paper *A Chain Rule for the Expected Suprema of Gaussian Processes* investigates the problem of assessing generalization of a learner who is adapting a feature space while also learning the target function. The approach taken is to consider extensions of bounds on Gaussian averages to the case where there is a class of functions that create features and a class of mappings from those features to outputs. In the applications considered in the paper this corresponds to a two layer kernel machine, multitask learning, and through an iteration of the application of the bound to multilayer networks and deep learners.

A standard assumption in statistical learning theory is that the data are generated independently and identically distributed from some fixed distribution; in practice, this assumption is often violated and a more realistic assumption is that the data are generated by a process which is only sufficiently fast mixing, and maybe even non-stationary. Vitaly Kuznetsov and Mehryar Mohri in their paper *Generalization Bounds for Time Series Prediction with Non-stationary Processes* consider this case and are able to prove new generalization bounds that depend on the mixing coefficients and the shift of the distribution.

Rahim Samei, Boting Yang, and Sandra Zilles in their paper *Generalizing Labeled and Unlabeled Sample Compression to Multi-label Concept Classes* consider generalizations of the binary VC-dimension to multi-label classification, such that maximum classes of dimension d allow a tight compression scheme of size d. Sufficient conditions for notions of dimensions with this property are derived, and it is shown that some multi-label generalizations of the VC-dimension allow tight compression schemes, while other generalizations do not.

Privacy, Clustering, MDL, and Kolmogorov Complexity. Christos Dimitrakakis, Blaine Nelson, Aikaterini Mitrokotsa, and Benjamin I.P. Rubinstein present the paper *Robust and Private Bayesian Inference*. This paper looks at the problem of privacy in machine learning, where an agent, a statistician for example, might want to reveal information derived from a data set, but without revealing information about the particular data points in the set, which might contain confidential information. The authors show that it is possible to do Bayesian inference in this setting, satisfying differential privacy, provided that the likelihoods and conjugate priors satisfy some properties.

Behnam Neyshabur, Yury Makarychev, and Nathan Srebro in their paper *Clustering, Hamming Embedding, Generalized LSH and the Max Norm* look at asymmetric locality sensitive hashing (LSH) which is useful in many types of machine learning applications. Locality sensitive hashing, which is closely related to the problem of clustering, is a method of probabilistically reducing the dimension of high dimensional data sets; assigning each data point a hash such that similar data points will be mapped to the same hash. The paper shows that by shifting to co-clustering and asymmetric LSH the problem admits a tractable relaxation.

Jan Leike and Marcus Hutter look at martingale theory in *Indefinitely Oscillating Martingales*; as a consequence of their analysis they show a negative result in the theory of Minimum Description Length (MDL) learning, namely that the MDL estimator is in general inductively inconsistent: it will not necessarily converge. The MDL estimator gives the regularized code length, $\mathrm{MDL}(u) = \min_Q\{Q(u) + K(Q)\}$, where Q is a coding function, $K(Q)$ its complexity, and $Q(u)$ the code length for the string u. It is shown that the family of coding functions Q can be constructed such that $\lim_{n\to\infty} \mathrm{MDL}(u_{1:n})$ does not converge for most infinite words u.

As is well-known, the Kolmogorov complexity is not computable. Peter Bloem, Francisco Mota, Steven de Rooij, Luís Antunes, and Pieter Adriaans in their paper *A Safe Approximation for Kolmogorov Complexity* study the problem of approximating this quantity using a restriction to a particular class of models, and a probabilistic bound on the approximation error.