

Foreword

Algorithmic learning theory is mathematics about computer programs which learn from experience. This involves considerable interaction between various mathematical disciplines including theory of computation, statistics, and combinatorics. There is also considerable interaction with the practical, empirical fields of machine and statistical learning in which a principal aim is to predict, from past data about phenomena, useful features of future data from the same phenomena.

The papers in this volume cover a broad range of topics of current research in the field of algorithmic learning theory. We have divided the 29 technical, contributed papers in this volume into eight categories (corresponding to eight sessions) reflecting this broad range. The categories featured are Inductive Inference, Approximate Optimization Algorithms, Online Sequence Prediction, Statistical Analysis of Unlabeled Data, PAC Learning & Boosting, Statistical Supervised Learning, Logic Based Learning, and Query & Reinforcement Learning.

Below we give a brief overview of the field, placing each of these topics in the general context of the field. Formal models of automated learning reflect various facets of the wide range of activities that can be viewed as *learning*.

A first dichotomy is between viewing learning as an indefinite process and viewing it as a finite activity with a defined termination. Inductive Inference models focus on indefinite learning processes, requiring only eventual success of the learner to converge to a satisfactory conclusion.

When one wishes to predict future data, success can be enhanced by making some restrictive but true assumptions about the nature (or regularities) of the data stream. In the learning theory community, this problem is addressed in two different ways. The first is by assuming that the data to be predicted is generated by an operator that belongs to a restricted set of operators that is known to the learner a priori. The PAC model and some of the work under the Inductive Inference framework follow this path. Alternatively, one could manage without any such prior assumptions by relaxing the success requirements of the learner: rather than opting for some absolute degree of accuracy, the learner is only required to perform as well as any learner in some fixed family of learners. Thus, if the data is erratic or otherwise hard to predict, the learner can ignore poor accuracy as long as no member of the fixed reference family of learners can do no better. This is the approach taken by some Online Sequence Prediction models in the indefinite learning setting and, also, by most of the models of Statistical Learning in the finite horizon framework.

Boosting is a general technique that applies a given type of learner iteratively to improve its performance. Boosting approaches have been shown to be effective for a wide range of learning algorithms and have been implemented by a variety of methods.

A second dichotomy is between *Supervised* and *Un-Supervised* learning. The latter we refer to as *learning from Unlabeled Data*. In the first scenario, the data has the form of an example-label pairs. The learner is trained on a set of such pairs and then, upon seeing some fresh examples, has to predict their labels. In the latter model, the data points lack any such labeling, and the learner has to find some persistent regularities in the data, on the basis of the examples it has seen. Such regularities often take the form of partitioning the data into clusters of similar points, but in some cases take other forms, such as locating the boundaries of the support of the data generating distribution.

Many learning algorithms can be viewed as searching for an object that fits the given training data best. Such *optimization* tasks are often computationally infeasible. To overcome such computational hurdles, it is useful to apply algorithms that search for approximations to the optimal objects. The study of such algorithms, in the context of learning tasks, is the subject of our Approximate Optimization Algorithms session.

There is a large body of research that examines different *representations* of data and of learners' conclusions. This research direction is the focus of our Logic Based Learning session.

A final important dichotomy separates models of interactive learning from those that model passive learners. In the first type of learning scenarios the actions of the learner affect the training data available to it. In the Query Learning model this interaction takes the form of queries of certain (pre-indicated) type(s) that the learner can pose. Then the data upon which the learner bases its conclusions are the responses to these queries. The other model that addresses interactive learning is Reinforcement Learning, a model that assumes that the learner takes actions and receives *rewards* that are a function of these actions. These rewards in turn are used by the learner to determine its future actions.

August 2004

Shai Ben-David
John Case
Akira Marouka

Preface

This volume contains the papers presented at the 15th Annual International Conference on Algorithmic Learning Theory (ALT 2004), which was held in Padova (Italy) October 2–5, 2004. The main objective of the conference was to provide an interdisciplinary forum for discussing the theoretical foundations of machine learning as well as their relevance to practical applications. The conference was co-located with the 7th International Conference on Discovery Science (DS 2004) and the 11th Conference on String Processing and Information Retrieval (SPIRE 2004) under the general title “The Padova Dialogues 2004”.

The volume includes 29 technical contributions that were selected by the program committee from 91 submissions. It also contains the invited lecture for ALT and DS 2004 presented by Ayumi Shinohara (Kyushu University, Fukuoka, Japan) on “String Pattern Discovery”. Furthermore, this volume contains the ALT 2004 invited talks presented by Nicolò Cesa-Bianchi (Università degli Studi di Milano, Italy) on “Applications of Regularized Least Squares in Classification Problems”, and by Luc De Raedt (Universität Freiburg, Germany) on “Probabilistic Inductive Logic Programming”.

Additionally, it contains the invited lecture presented by Esko Ukkonen (University of Helsinki, Finland) on “Hidden Markov Modelling Techniques for Haplotype Analysis” (joint invited talk with DS 2004). Moreover, this volume includes the abstract of the joint invited lecture with DS 2004 presented by Pedro Domingos (University of Washington, Seattle, USA) on “Learning, Logic, and Probability: A Unified View”.

Finally, this volume contains the papers of the research tutorials on *Statistical Mechanical Methods in Learning* by Toshiyuki Tanaka (Tokyo Metropolitan University, Japan) on “Statistical Learning in Digital Wireless Communications”, by Yoshiyuki Kabashima and Shinsuke Uda (Tokyo Institute of Technology, Japan), on “A BP-Based Algorithm for Performing Bayesian Inference in Large Perceptron-like Networks”, and by Manfred Opper and Ole Winther on “Approximate Inference in Probabilistic Models”.

ALT has been awarding the *E. Mark Gold Award* for the most outstanding paper by a student author since 1999. This year the award was given to Hubie Chen for his paper “Learnability of Relatively Quantified Generalized Formulas”, co-authored by Andrei Bulatov and Victor Dalmau.

This conference was the 15th in a series of annual conferences established in 1990. Continuation of the ALT series is supervised by its steering committee consisting of: Thomas Zeugmann (Hokkaido Univ., Sapporo, Japan), Chair, Arun Sharma (Queensland Univ. of Technology, Australia), Co-chair, Naoki Abe (IBM T.J. Watson Research Center, USA), Klaus Peter Jantke (DFKI, Germany), Roni Khardon (Tufts Univ., USA), Phil Long (National Univ. of Singapore), Hiroshi Motoda (Osaka University, Japan), Akira Maruoka (Tohoku Univ., Japan), Luc De Raedt (Albert-Ludwigs-Univ., Germany), Takeshi Shi-

nohara (Kyushu Institute of Technology, Japan), and Osamu Watanabe (Tokyo Institute of Technology, Japan).

We would like to thank all individuals and institutions who contributed to the success of the conference: the authors for submitting papers, the invited speakers for accepting our invitation and lending us their insight into recent developments in their research areas, as well as the sponsors for their generous financial and logistical support.

We would also like to thank Thomas Zeugmann for assisting us via his experience in the publication of previous ALT proceedings, for providing the ALT 2004 logo, and for managing the ALT 2004 Web site. We are very grateful to Frank Balbach who developed the ALT 2004 electronic submission page.

Furthermore, we would like to express our gratitude to all program committee members for their hard work in reviewing the submitted papers and participating in online discussions. We are also grateful to the external referees whose reviews made a considerable contribution to this process.

We are also grateful to the DS 2004 chairs Einoshin Suzuki (PC Chair, Yokohama National University, Japan) and Setsuo Arikawa (Conference Chair, Kyushu University, Japan) for their effort in coordinating with ALT 2004, and to Massimo Melucci (University of Padova, Italy) for his excellent work as the local arrangements chair. Last but not least, Springer provided excellent support in preparing this volume.

August 2004

Shai Ben-David
John Case
Akira Maruoka

Organization

Conference Chair

Akira Marouka Tohoku University, Sendai, Japan

Program Committee

Shai Ben-David	Univ. of Waterloo, Canada, Co-chair
John Case	Univ. of Delaware, USA, Co-chair
Nader Bshouty	Technion, Israel
Michael Collins	MIT, USA
Sanjoy Dasgupta	UC San Diego, USA
Peter Flach	Univ. of Bristol, UK
Steffen Lange	DFKI, Saarbrücken, Germany)
Jon Langford	TTI Chicago, USA
Gabor Lugosi	UPF, Barcelona, Spain
Rüdiger Reischuk	Univ. at Lübeck, Germany
Rocco Servedio	Columbia Univ., USA
Arun Sharma	Queensland Univ. of Technology, Brisbane, Australia
Ayumi Shinohara	Kyushu Univ., Japan
Eiji Takimoto	Tohoku Univ., Japan
Sandra Zilles	Univ. Kaiserslautern, Germany

Local Arrangements

Massimo Melucci University of Padova, Italy

Subreferees

Kazuyuki Amano	Jochen Nessel
Pedro Felzenswalb	Ryan O'Donnell
Kouichi Hirata	Daniel Reidenbach
Klaus P. Jantke	Yoshifumi Sakai
Adam Klivans	Gilles Stoltz
Stephen Kwek	Rolf Wiehagen
Martin Memmel	

Sponsoring Institutions

Institute for Theoretical Computer Science, University at Lübeck
Division of Computer Science, Hokkaido University

Remembering Carl Smith, 1950–2004

Sadly, Carl Smith passed away 10:30PM, July 21, 2004. He had had a 1.5 year battle with an aggressive brain tumor. He fought this battle with calm optimism, dignity, and grace. He is survived by his wife, Patricia, his son, Austin, and his sister, Karen Martin.

Carl was very active in the algorithmic or computational learning communities, especially in the inductive inference subarea which applies recursive function theory techniques.

I first met Carl when I interviewed for my faculty position at SUNY/Buffalo in the Spring of 1973. He was then a graduate student there and told me he was interested in recursive function theory. After I joined there, he naturally became my Ph.D. student, and that's when we both began working on inductive inference. We spent a lot of time together, pleasantly blurring the distinction between the relationships of friendship and advisor-student.

After Buffalo, Carl had faculty positions at Purdue and, then, the University of Maryland.

Carl had a very productive career. He was a master collaborator working with many teams around the world. Of course he also produced a number of papers about inductive inference by teams — as well as papers about anomalies, queries, memory limitation, procrastination, and measuring mind changes by counting down from notations for ordinals. I had the reaction to some of his papers of wishing I'd thought of the idea. This especially struck me with his 1989 *TCS* paper (with Angluin and Gasarch) in which it is elegantly shown that the learning of some classes of tasks can be done only sequentially after or in parallel with other classes.

Carl played a significant leadership role in theoretical computer science. In 1981, with the help of Paul Young, Carl organized the Workshop on Recursion Theoretic Aspects of Computer Science. This became the well known, continuing series of Computational Complexity conferences. Carl provided an improvement in general theoretical computer science funding level during his year as Theory Program Director at NSF. He was involved, in many cases from the beginning, in the COLT, AII, ALT, EuroCOLT, and DS conferences, as a presenter of papers, as a member of many of their program committees and, in some cases, steering committees. He spearheaded the development of COLT's Mark Fulk Award for best student papers and managed the finances.

Carl was very likable. He had a knack for finding funding to make good things happen. He was a good friend and colleague. He is missed.