Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE @ DIRECT®
Theoretical Computer Science 313 (2004) 173 – 174
Theoretical Computer Science
www.elsevier.com/locate/tcs

Foreword

This special issue contains six of the 21 regular and 3 invited papers presented at the Twelfth Annual conference on Algorithmic Learning Theory (ALT ’01), which was held in Washington DC, USA, during November 25–28, 2001. The conference proceedings, including preliminary versions of these papers, appeared as Lecture Notes in Artificial Intelligence, Vol. 2225 from Springer. Following the conference, authors were invited to submit full versions of their papers to the special issue. All the papers were subjected to the standard refereeing procedure of Theoretical Computer Science.

Algorithmic learning theory aims at theoretical study of learning problems and algorithms, and it does so by providing formal models of learning and analyzing complexity of learning problems and performance of learning algorithms in the models set forth. It provides foundations for such varied branches of artificial intelligence as machine learning, inductive logic programming, data mining, natural language processing, bioinformatics and others.

The first paper, an invited paper by Dana Angluin, presents a comprehensive survey of the state of the art in the so-called query learning model, a field she has initiated. Major emphasis is put on the number of queries needed to learn a class of concepts. Work in this area has identified several characterizations of learnability using combinatorial properties of concept classes. Thus the complexity of learning a class can be determined without directly considering algorithmic issues but purely in terms of such characterizations. The paper puts these results in a single framework clarifying and showing the relations among them to give a clear picture of the state of the art in this area.

The next paper by Kalnishkan et al. concerns the on-line learning model, or more specifically the on-line prediction games. In this framework an iterative game is played where in each step the learner makes a prediction, observes the true value and suffers a loss based on these two values and a fixed loss function. The prediction complexity of a sequence gives a lower bound on the loss of any prediction algorithm for the sequence and can thus be seen as another way to characterize the inherent complexity of strings. Kalnishkan et al. derive results on the average complexity when the sequence is i.i.d. Bernoulli and relates this to the information complexity of the sequence. As a result it is shown that the Kolmogorov complexity does not coincide with the prediction complexity for the binary game.

The next two papers are in the sub-area generally known as inductive inference, which follows seminal work by Gold who introduced the model of learning in the limit. Here finite complexity rather than polynomial complexity defines the notion of

doi:10.1016/j.tcs.2003.11.003

feasibility. Jain and Stephan study the problem of learning to separate pairs of disjoint sets which do not necessarily cover the instance space. Several restrictions on the learner are studied within this framework, for example: conservative learners who only abandon hypotheses which were contradicted by data, and set-driven learners whose hypotheses depend exclusively on the range of the input. The effect of these restrictions and their interaction is extensively studied. The two notions mentioned here are not comparable if the learner converges on all data sequences.

Zilles studies a notion of meta-learning where a single learning algorithm can learn several concept classes by being given an index of the class as a parameter. Two scenarios are discussed where the learner either uses a single representation for hypotheses in all classes or can change the representation with the class. The paper studies the effect of restricting the learner, for example to be conservative as described above, on the learnable classes. Various results are given separating different models. Interestingly, the concept classes used in the separation (i.e. learnable in one setting but not the other) have finite cardinality.

The final two papers study learning problems for formal languages. The problem of learning regular languages has been extensively studied with several representation schemes. Dennis et al. study learnability of regular languages using a non-deterministic representation based on residuals—completion languages for prefixes of words in the language. While the representation is shown not to be polynomially learnable, parameters of the representation are studied empirically and these suggest a new learning algorithm with desirable properties. Experiments show that the algorithm compares favorably with other systems.

The class of Büchi automata defines languages over infinite strings. When modeling learnability of such languages one is faced with the question of examples of infinite size. The paper by de la Higuera and Janodet introduces a model of learning such languages from finite prefixes of examples. While the complete class is not learnable a sub-class is identified and shown learnable in the limit with polynomial update on each example.

We would like to thank the members of the program committee of ALT'01 and the reviewers for their immense efforts.

N. Abe
IBM Thomas J. Watson Research Center
Yorktown Heights, NY 10598, USA
E-mail address: nabe@us.ibm.com

R. Khardon
Department of Computer Science, Tufts University
Medford, MA 02155, USA
E-mail address: roni@cs.tufts.edu