Intelligent Data Analysis 8 (2004) 211–212
IOS Press

Guest-editorial

Incremental learning and concept drift: Editor's introduction

Miroslav Kubat[a], João Gama[b] and Paul Utgoff[c]
[a]*University of Miami, Coral Gables, FL, USA*
[b]*University of Porto, Porto, Portugal*
[c]*University of Massachusetts, Amherst, MA, USA*

A complex problem in data analysis is the time-varying nature of many realistic domains. In many real-world learning problems, training data become available in batches over time, or even flow steadily, as in user-modeling tasks, dynamic control systems, web-mining, and times series analysis. In these applications, learning algorithms should be able to adjust the decision model dynamically whenever new data become available. This is the scenario that motivates this special issue of *Intelligent Data Analysis* to machine learning systems capable of dealing with concept drift.

To narrow the domain of interest, we focus on those learning scenarios in which the system must induce the concept from timestamped training data. A brute-force algorithm relearns the concept from scratch each time a new example becomes available. This poses several problems. Learning from the scratch wastes computational resources. Moreover, in non-stationary environments, the system should take into account the fact that only the most recent examples are relevant to the actual target concept. A less expensive approach would employ an *incremental* learning technique that adapts the previously induced concept model by incorporating the experience obtained from newly available examples.

An incremental learning system can be used with some success for domains in which the underlying instance distribution evolves, especially if there is an abundance of examples that are representative of the most recent version of the target concept. Yet, in domains where the change is substantial and there is a paucity of recent examples, the system needs to be able to discount or even **forget** older examples, and **adjust** what has been induced from them. The task is more difficult than it appears. When learning in time-varying domains, the system needs to modify the internal concept representation not only as more examples become available, but also in response to suspected changes in the definition of the target concept. It is of paramount importance that the system be able to distinguish between the situation in which new examples only help to **fine-tune** the existing concept model, and the situation in which the new examples are indicative of a shift in the target concept. To complicate matters even further, the system should not be misled by noise.

Over the past decade, many researchers have become interested in this task, and the results of their work have appeared in diverse journals and conferences. By organizing this special issue, we wanted to concentrate several alternative approaches in the same volume in order to give the interested reader a better idea about the state-of-the-art of the relevant algorithms, applications, and evaluation methods. We believe that the five articles that appear here satisfy this goal.

Three of the papers present, discuss, and evaluate incremental algorithms that dynamically react to changes in the distribution of the learned concepts. In *Incremental Learning and Concept Drift in INTHELEX*, Esposito et al. describe a system for the induction of first order logic theories. INTHELEX detects a stable context and adapts the induced theory to a restricted number of the most recent examples. In *Adaptive Ripple Down Rules Method Based on Minimum Description Length Principle*, the authors present an incremental algorithm for knowledge acquisition using an approach they dub ripple down rules (RDR). The approach ensures that the integrity of the accumulated knowledge always remains intact after every addition of a new piece of knowledge. In *Towards a Machine Learning Approach Based on Incremental Concept Formation*, Maddouri uses formal concept analysis to facilitate incremental learning by the introduction of operators for data addition, deletion, and updating. In *Learning Drifting Concepts: Example Selection Versus Example Weighting*, Klinkenberg presents an excellent overview of methods used to detect concept drift, and he evaluates the different techniques using support vector machines. In *Sequential Learning in Neural Networks: A Review and a Discussion of Pseudorehearsal Based Methods*, Robins studies the effect of *catastrophic forgetting* in sequential learning of neural networks and he discusses possible solutions based on pseudorehearsal based methods.

We express our thanks to the *Intelligent Data Analysis* journal as well as to all the authors who submitted their work. Importantly, the special issue would never have been possible without the commitment of the anonymous reviewers who carefully reviewed the submissions.