Abstract

In machine learning, the principle of selective action can be understood as decision-making under some specific constraints, usually related to performance. A simple example is: predict the next value in a time series only if the confidence level of the prediction is higher than 80%. This means that predictions are sparse, i.e. they can be considered valid only for a subset of the problem space. Regarding the existing methods in machine learning, the literature has imposed Kernel methods as the current state of the art. Not until recently, the principle of selective action has been incorporated in several kernel-based models, called sparse kernel machines. Kernel methods are a fairly young and rapidly evolving domain, with research in sparse kernel machines being a rather new field. The literature lacks a set of studies on the applicability of sparse kernel machines to the problem of time series prediction. The main purpose of this research is to partially fill this gap, with advances in and experimental studies of the Sparse Kernel Machine model for time series prediction. The experiments carried out focus mainly on financial data, notorious for the difficulty in prediction and on traffic data in the area of Bucharest and Ilfov County.

We propose two Sparse Kernel Machine models. One model focuses on improving the directional accuracy of the predictions while the second model focuses on the reduction of the Root Mean Squared Error. The second model also uses tools of High Performance Computing for feature selection. The studies show that the proposed methods can outperform the existing state of the art. Compared to the proposed models, the studies involve a significant body of existing methods for time series prediction. Data covers chaotic financial time series as well as fundamental information data. The experiments show that chaotic data contains a weak deterministic component, detectable only with the proposed Sparse Kernel Machine models, while fundamental information provides a statistically significant effect size for the directional prediction of monthly currency prices.

The second focus of the thesis is on the new and promising field of Kernel Based Reinforcement Learning (KBRL), where significant advances have been achieved only since as early as 2002. Existing methods in KBRL extend the problem of decision making from discrete Markov Decision Processes (MDPs) to continuous-state MDPs by the use of a kernel function, but still have some problems. First, the choice of the kernel function for a particular problem is still an open issue. Relevant literature suggests the kernel design problem should be solved by a human operator, but this is not always possible. Second, existing KBRL inference models lack a confidence measure attached to each inference of the KBRL agent. There may be applications, as suggested in this thesis, which require the agent to use an inference filter to avoid penalties in its operating environment.

In the last part of the thesis we focus on solving the problem of automatic kernel function parameter selection and providing a set of inference filters. First, we propose a novel inference metaheuristic for automatic parameter selection, using two inference strategies. Second, we propose a set of inference filters, collectively referred to as QoINF (Quality of Inference) filters. The proposed inference metaheuristic and the inference filters are tested in three scenarios: a trading experiment, a synthetic data set experiment and an image recognition experiment. The trading experiment shows a KBRL agent can achieve positive returns using the proposed inference metaheuristic and fundamental economic information. The second experiment uses synthetic data and outlines the advantage of the second inference strategy. In the third experiment we cast the problem of classification as a reinforcement learning problem and illustrate the possibility of achieving 80%-100% image recognition accuracies. Two parallel KBRL inference algorithms are also proposed and tested on the UPB NCIT cluster and on an NVIDIA GeForce GPU, in an image recognition application for intelligent buildings.