Computer Networks 53 (2009) 191–201

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Support vector regression for link load prediction

Paola Bermolen [a], Dario Rossi [a,*]

[a] TELECOM ParisTech, INFRES Department, 46 rue Barrault, Paris 75013, France

ARTICLE INFO

Article history:
Available online 10 October 2008

Keywords:
Network forecast
Traffic measurement
Supervised learning
Support vector machines

ABSTRACT

From weather to networks, forecasting techniques constitute an interesting challenge: rather than giving a faithful description of the current reality, as a looking glass would do, researchers seek crystal-ball models to speculate on the future. This work explores the use of Support Vector Regression (SVR) for the purpose of link load forecast. SVR works well in many learning situations, because they generalize to unseen data, and are amenable to continuous and adaptive online learning – an extremely desirable property in network environments. Motivated by the encouraging results recently gathered by means of SVR on other networking applications, our aim is to enlighten whether SVR is also successful for the prediction of network links load at short time scales. We consider the problem of link load forecast based only on its past measurements, which is referred to as "embedded process" regression in the SVR lingo, and adopt a hands-on approach to evaluate SVR performance. In more detail, we perform a sensitivity analysis of the parameters involved, assess the computational complexity for training and validation, dig into the correlation structure of the prediction errors and evaluate techniques to extend the forecasting horizon. Our finding is that accuracy results are close enough to be tempting, but not enough to be convincing. Yet, as SVR exhibit a number of advantages, such as good robustness and flexibility properties, furthermore at a price of a limited complexity, we then speculate on what directions can be undertaken to ameliorate its performance in this context.

1. Introduction

It is fairly well accepted that, as a result of network services and Internet applications evolution, network traffic is becoming increasingly complex. On the one hand, transport networks are challenged by the current convergence trend of voice/video/data services on an all-IP network, and by the fact that user-mobility will likely translate into service-mobility as well. On the other hand, the explosion of Internet telephony, television and gaming applications implies that we may be forced to re-think what we mean by "data" traffic. Moreover, the widespread usage of application layer overlays directly translates into a much higher variability of the data traffic injected into the network.

In this paper, we question whether such variability can be efficiently forecasted, and if so, with what level of accuracy. The supervised prediction technique we selected is Support Vector Machines (SVM), a set of classification and regression techniques, introduced in the early nineties [1], that are grounded in the framework of statistical learning theory. Basically, Support Vector Regression (SVR) uses training data to build a forecast model which works well in many learning situations because it generalizes to unseen data and is amenable to continuous and adaptive online learning, an extremely desirable property in network environments. Initially bound to the optical character recognition context, the use of SVM rapidly spread to other fields, including time series prediction [2] and, more recently, networking [3–6]. Motivated by such encouraging results, we focus on link load forecast based only on past measurements, following an approach known as "embedded process" [2]. This problem is of great interest in

* Corresponding author. Tel.: +33 145817563.
E-mail address: dario.rossi@telecom-paristech.fr (D. Rossi).

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.09.018

networking for both capacity planning and self-management application (e.g. bandwidth provisioning, admission control, trigger of backpressure mechanisms, etc.).

Though the SVM approach fits well to longer timescales as well, which are more of a concern for capacity planning, in this paper we focus on the estimation of load variation at short time scales: adopting a hands-on approach to the SVM regression, we evaluate the effectiveness of SVR for link load forecast by exploring a rather extensive parameter and design space. Our aim is twofold: first, we want to evaluate the SVM accuracy and robustness and, second, we want to provide useful insights on the tuning of the SVM parameters, an aspect not always clear in previous work. We compare the performance with those achievable using Moving Average and Auto-Regressive models: our results show that, despite a good accordance with the actual data, the SVR gain achievable over simple prediction methods is not enough to justify its deployment for link load prediction at short time scales. Yet, we have to tribute SVR of a number of extremely positive aspects: for instance, SVR models are rather robust to parameter variation, and their computational complexity is far from being prohibitive, which makes them suitable for online prediction. Moreover, we experimentally verify that errors calculated over consecutive samples are independent and identically distributed, which allows the evaluation of confidence intervals. Finally, we also investigate methods to extend the forecast horizon using forecasted values as input for a new prediction: interestingly, this approach of recursive SVR may significantly extend the achievable forecast horizon, entailing only a very limited accuracy degradation.

The remainder of the paper is organized as follows. After discussing related work in Section 2, we briefly overview the Support Vector Regression theory in Section 3. In Section 4 we specify the methodology we follow in applying SVR models to link load forecast, as well as describing the other forecasting techniques that will be used for comparison purposes. A complete and extensive sensitivity analysis of SVR performance is reported in Section 5, whereas further details on the temporal evolution of the error, computational complexity considerations and result of recursive SVR are reported in Section 6. Finally, concluding remarks and future work are addressed in Section 7.

2. Related work

Most of the work related to network load forecast is based on the analysis of time series properties. In this context, a number of very different models [7–9] have been proposed, ranging from very simple to very complex ones. However, the majority of these approaches relies on specific assumptions and underlying models for the network traffic (e.g., they are tailored to capture Long Range Dependence (LRD) [10] at short and long timescales, etc.). A first drawback is that such models will no longer be applicable if the assumption no longer holds (e.g., considering other timescales). A second drawback is that such models usually rely on the precise estimation of some traffic parameters, whose computation can be a very intensive and delicate task (e.g., Hurst parameter of the arrival time series). Rather, as in [11,12], we prefer to focus on techniques that, avoid making any assumptions on the phenomenon under observation, allow for intrinsically more robust and flexible prediction. A simple local Gaussian predictor is provided in [11] as a core tool to guide the bandwidth provisioning in the hose model: interestingly, the model is able (but not *forced*) to embed assumptions on the LRD properties of the traffic, by an appropriate tuning of the parameters. TCP throughput prediction is the object of [12], where the authors compare formula-based versus history-based prediction schemes, showing that even simple moving-average models are able to yield satisfactory results (provided that one copes with major error sources).

The forecasting technique that we plan to evaluate in this paper falls in the class of Support Vector Machines (SVM) [13]: despite its relatively short existence, the literature of SVM is already full blown. At the same time, while the use of SVM for *classification* is relatively more popular in networking research, especially in the context of anomaly and intrusion detection [5,6], the use of SVM for *regression* is largely left unexplored. To the best of our knowledge, the only work that explores the use of SVR techniques in the networking field is [3,4]. TCP throughput prediction on a given path is the object of [3], where the forecast is based on a combination of path properties (such as queueing delays and available bandwidth) and on the performance of prior file transfers as well. Authors show that when the path properties are *precisely* known (e.g., when they are provided by an "oracle"), SVR is able to predict TCP throughput within 10% of the actual value in 90% of the cases – which represents nearly a 3-fold improvement in accuracy over prior history-based methods. Also, in more realistic scenarios and using less accurate measurements of path properties (e.g., gathered by means of active probes), the predictions can be made within 10% of the actual value nearly 50% of the time – which still represents a 60% improvement, with a furthermore much lower impact on end-to-end paths.

The authors of [4] focus instead on the prediction of the latency toward an unknown IP address, based on the latency knowledge toward other previously contacted IP addresses. Using as input features vectors of IP address bits (transformed into a 32 dimension input space, where each bit of the address corresponds to a different dimension), authors show that the estimation performance is within the 30% of the true value for approximately three-quarters of the latency prediction on a large Internet data set. More in detail, SVM regression on a large randomly collected data set of 30,000 (IP,latency) couples, yield a mean prediction error of 30 ms (25 ms) using only 6% (20%) of the samples for training.

In the context of SVM regression [13], the problem of forecasting future values of a series based only on previous observation of the same phenomenon is known as an "embedding process" [2]. However, its application has usually targeted domains other than the networking context, and the series that SVR has been fed with up to now are very much different from those representing the packet arrival process at a router queue: thus, our aim is to test whether SVR can prove to be a useful tool also for link load forecast.