<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *DASC-2015*.

Citation for the original published paper:

de Fréin, R. (2015)

Take off a Load: Load-Adjusted Video Quality Prediction and Measurement.

In: *13th IEEE International Conference on Dependable, Autonomic and Secure Computing: 2015 IEEE International Conference on Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing*, (pp. 1-9). Liverpool: IEEE Computer Society

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-175434>

Take off a Load: Load-Adjusted Video Quality Prediction and Measurement

Ruairí de Fréin^{† ††}

[†]KTH - Royal Institute of Technology, Stockholm,
Sweden

^{††}Waterford Institute of Technology,
Ireland

web: <https://robustandscalable.wordpress.com>

in: 2015 IEEE International Conference on Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing, to appear. See also $\text{BIB}_{\text{E}}\text{X}$ entry below.

$\text{BIB}_{\text{E}}\text{X}$:

```
@article{rdefrein15DASCTake,  
author={Ruairí de Fréin† ††},  
journal={2015 IEEE International Conference on Computer and Information Technology;  
Ubiquitous Computing and Communications;  
Dependable, Autonomic and Secure Computing;  
Pervasive Intelligence and Computing, to appear},  
title={Take off a Load: Load-Adjusted Video Quality Prediction and Measurement},  
year={2015},  
pages={9},  
month={Oct}, }
```

© 2015 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Take off a Load: Load-Adjusted Video Quality Prediction and Measurement

Ruairí de Fréin

KTH Royal Institute of Technology, Sweden
rdefrein@gmail.com

Abstract—An algorithm for predicting the quality of video received by a client from a shared server is presented. A statistical model for this client-server system, in the presence of other clients, is proposed. Our contribution is that we explicitly *account for* the interfering clients, namely the load. Once the load on the system is understood, accurate client-server predictions are possible with an accuracy of 12.4% load adjusted normalized mean absolute error. We continue by showing that performance measurement is a challenging sub-problem in this scenario. Using the correct measure of prediction performance is crucial. Performance measurement is miss-leading, leading to potential over-confidence in the results, if the effect of the load is ignored. We show that previous predictors have over (and under) estimated the quality of their prediction performance by up to 50% in some cases, due to the use of an inappropriate measure. These predictors are not performing as well as stated for about 60% of the service levels predicted. In summary we achieve predictions which are $\approx 50\%$ more accurate than previous work using just $\approx 2\%$ of the data to achieve this performance gain—a significant reduction in computational complexity results.

I. INTRODUCTION

Understanding and predicting performance metrics for telecom clouds services is a challenging, open problem [1], [2]. The authors of [1] take a Statistical Learning (SL) approach—they apply variants of the Lasso [3], Random Forests [4] and Ridge Regression [5]—to predict the client-side metrics for a video streaming service. Their approach is significant as they collect statistics from the Linux kernel of a server machine to achieve this; their initial prediction performance results are promising; and finally, they have made their traces publicly available. To evaluate the performance of their video quality metric prediction algorithms they designed three sub-components: a test-bed, a video quality metric prediction model and learning algorithm, and finally, a measurement approach for evaluating the quality of the predictions in a fair way. The focus of this paper is on improving the performance of the second and third component above given that the test-bed, described in [1], and resulting traces, have been used in several papers and have gained acceptance.

Related work: The position of Yanggratoke et al., in [1] is that 1) by collecting thousands of kernel variables their prediction approach is service independent (they can omit service specific instrumentation etc.). We illustrate that this service-independent-prediction assumption does not always hold, that comparison of the predictor’s performance across services is

not justified, and that unless the service is modelled correctly, artifacts are introduced into the predictor and the measurement system. 2) The prediction of client-side metrics such as RTP packet rates (with 10%–15% error) across different scenarios and loads are possible. We demonstrate that if the approach in [1] is adopted the results are in fact dominated by the choice of scenario and load when prediction is performed. In short, this approach may inadvertently *game* the performance of the predictor positively based on the selection of a favourable scenario; or, on the other hand lead to the unfair dismissal of an approach due to an unfavourable scenario.

Domain knowledge is often inferred and then used in Signal Processing [6] and Computational Finance [7]; these techniques are referred to as *Blind* inference, learning and prediction [8]. Blind inference has not yet been widely embraced by the Network Management community. This is due to the number of different active network services, and also, whether or not it is feasible to model highly dynamic network services automatically. We desire learning algorithms that are powerful enough to learn from data without any domain knowledge or human intervention, namely *Blind* [9], [10] or *autonomic* approaches. With out loss of generality, to evaluate our approach we place an adaptive sinusoidal user request load on a video server; in practice the load is an arbitrary trace. To improve prediction performance it is reasonable to assume that we need accurate knowledge of this trace. Somewhat suprisingly we show that a Blind inference procedure is not necessary in our scenario to estimate the trace. We obtain a good estimate of it from the TCP socket count of the server. Our non-blind approach uses exactly the same information as the previous work [1], prediction accuracy improvement is obtained *for-free*, and the approach is applicable irrespective of the load on the system. We agree with the authors in [1] that collecting statistics from the Linux kernel and client-side, and learning the mapping between them, is a promising first approach. Intrusion detection systems have a long history of using such parameters—sequences of system call executed by running processes—to discriminate between normal and abnormal operating characteristics of UNIX programs [11]. However, a general model, which accounts for network, service and client delay is needed. The authors focused on the simple instantaneous case in [1] as the lab configuration considered has sufficient resources for these delays to have little affect.

The current trend of running software systems on general purpose platforms without real-time guarantees with the expectation that one can safeguard revenues, is dichotomous. The choice of video service level prediction by [1], as an exemplar instance of this problem, is timely given that Cisco

Dr de Fréin is affiliated with TSSG, Waterford Institute of Technology, Ireland. This work was supported by an ELEVATE Irish Research Council International Career Development Fellowship co-funded by Marie Curie Actions award “EOLAS”: ELEVATEPD/2014/62.

[12] predicts that network traffic volumes in the order of tens of exabytes are not that far off, and 90% will be video related [13]. A SL approach –similar to [1]– is preferable to developing and fitting complex analytical models for the different layers of soft/hardware in these complex systems. The authors of [14] make the case that modern multi-core (parallel online) learning algorithms are limited by the bandwidth bottleneck. It is hard to justify expending bandwidth resources on predicting why a service is not meeting service level agreements if this bandwidth could be used to meet the service delivery shortfall. If a SL approach with low complexity, which increases linearly in the feature set size, is unable to perform predictions with low enough latency (using one of the computational architectures in [15]), it is unlikely that a significantly more complex, hierarchical analytical model will exhibit sufficiently good performance; our philosophy is to explore the simplest approach in depth first before we discard it. As a first result we demonstrate that we can outperform the results in [1] by using just 2% of the data, which contradicts the assertion that we need large amounts of data for successful SL. This performance gain is achieved by incorporating a small amount of –already present– knowledge into the SL algorithms.

The application of SL for prediction in cloud and network environments is in its infancy. A method for identifying and ranking servers with problematic behavior is proposed in [16]. The authors use Random Forest classifiers to select candidate servers for modernization. A predictive model is then used to determine the impact of modernization actions. A Support Vector Regression predictor is used in [17] to perform lightweight TCP throughput prediction. Prediction is based on prior file transfer history and measurements of path properties. A method for modeling application servers in order to detect performance degradation due to aging is presented by [18]. The authors use classification algorithms to perform proactive detection of performance degradation. Finally, the authors attempt to reduce the size of the data-stream that is forwarded to an operators’ operations support system by removing uncorrelated noise events in [19]. A heuristic cross-correlation function determines the degree of inter-relationship between the events in the data-stream. To the best of our knowledge there is no work which explicitly deals with the effects of adaptive loads on these systems.

Contribution 1: We claim that a different prediction model should be used when there is a different load on the system under observation. We propose a simple hierarchical model, namely Load-Adjusted RR (LA-RR), and demonstrate performance gains of up to 30% are achievable using our model over traditional Ridge Regression (T-RR). **Contribution 2:** We propose a performance measure for analyzing the performance of the new model, namely Load-Adjusted Normalized Mean Absolute Error (LA-NMAE). Empirical results support the claim that the Traditional Normalized Mean Absolute Error (T-NMAE) is an inappropriate performance measure. We quantify how big of an issue this is. **Contribution 3:** We complete our study by proposing a new hierarchical prediction algorithm, LA-RR. We compare the performance of the measure LA-NMAE when evaluating the RR algorithm [20] used in [1], namely T-RR, with our new load-adjusted hierarchical solver, LA-RR. We also compare the performance of LA-RR to the previously proposed RR technique in [1], T-RR, using the performance measure in [1], T-NMAE. We do not exhaustively

TABLE I. ACRONYMS: LOAD-ADJUSTED & TRADITIONAL STATISTICAL MODEL/ALGORITHM & PERFORMANCE MEASURE.

	Statistical Model/Algorithm	Performance Measure
	Ridge Regression	NMAE
Load-adjusted (new)	LA-RR	LA-NMAE
Traditional (old)	T-RR	T-NMAE

evaluate each of the methods in [1] because: 1) the solver cannot correct the formulation of the problem; 2) T-RR gives the best performance on the periodic load traces used in [1]; 3) a thorough empirical comparison of the Lasso [3], for example, with RR, involves the selection of different regression parameters for each algorithm. We have focused on comparing RR for both the model in [1], T-RR, and our hierarchical model, LA-RR, using the same regularization parameter for both, as the purpose of this paper is to motivate the candidacy of our hierarchical load-adjusted statistical model.

Organization: This paper makes both a theoretical and practical contribution. It starts by introducing the theoretical tools we need to perform improved predictions in Section II, III and IV. We introduce a statistical model for the client-server system in Section II. In Section III we support this model empirically using a statistical test which compares the probability of our model being valid given the data (in [1]), with the probability of the state-of-the-art model being valid given the same data. This test finds that the state-of-the-art model is implausible, given the data, and that our new model is more plausible. The second part of our theoretical contribution, in Section IV, demonstrates that prediction performance measurement is challenging. In Section IV we illustrate that prediction performance measurement is highly dependent on the load on the system. We introduce a new measure to account for this type of error. We continue by showing, using the data in [1], that using the inappropriate performance measure may unjustifiably inflate or deflate the quality of our predictions. In the final two sections we introduce some practical tools for making service level predictions. We introduce a practical hierarchical RR prediction technique in Section V which follows from the analysis in Section II, III and IV. We perform a thorough simulation study that empirically evaluates and compares this technique with T-RR in Section VI. In this empirical study we use exactly the same information as [1] and our improved model yields significant performance gains.

II. SYSTEM MODEL: LA-RR

A client is connected to a server via a network and s/he requests Video on Demand, which runs on the server in [1]. Assume for the purpose of exposition that the system is operating under a light to medium load (we relax this assumption later); we, like the authors of [1], do not model the network state. What happens when a client requests video? The response of the server, with respect to kernel metric n , the n -th feature, to one request for video at time i is expressed as:

$$\mathbf{x}_i[n] = \hat{\mathbf{u}}_i[n] + \epsilon_i[n], \quad \text{where } i \in \mathbb{Z}, \mathbf{x}_i[n], \hat{\mathbf{u}}_i[n] \in \mathbb{R}. \quad (1)$$

A feature refers to a metric on the operating system level, for example, the number of active TCP connections. The feature set $\mathbf{x}_i[n]$ is constructed using the System Activity Report^{1a}

^{1a}<http://linux.die.net/man/1/sar>; ^{1b}<http://www.videolan.org/vlc>;

^{1c}<http://www.ntp.org/>

Fig. 1. Service level metric and system load trace for a periodic-type load.

(SAR) which computes system metrics over a given time interval. The term $\mathbf{x}_i[n]$ denotes the n -th feature at time index i . On the client-side we observe an application level metric, the RTP packet rate, y_i at time i . VLC^{1b} media player provides Video-on-Demand on the test-bed in [1].

Problem Statement: The objective of this paper is to predict unseen values of y_i using the features $\mathbf{x}_i[n], \forall n$. We assume that a global clock can be read on both the client and server to match up the $\{\mathbf{x}_i[n], y_i\}$ pairs.

The signal $\hat{\mathbf{u}}_i[n]$ in (Eqn. 1), a square-wave (off-on-off) signal, corresponds to an increase in the CPU workload, for example, an extra X units per additional user for the duration of the video requested by the user. We assume that $\hat{\mathbf{u}}_i[n]$ is scaled in order to account for the sensitivity of a given feature to the effect of adding a new user to the system. This scaling is specific to each service, feature and machine. The signal, $\epsilon_i[n]$, captures deviations from the ideal performance of the server with respect to the n -th feature. We assume that this deviations signal is normally distributed with 0 mean and variance σ^2 . If there is more than one deviation signal, they are uncorrelated. For example, for two simultaneous video requests (of the same duration) the response of the n -th feature is

$$\mathbf{x}_i[n] = 2\hat{\mathbf{u}}_i[n] + \epsilon_i[n, 1] + \epsilon_i[n, 2]. \quad (2)$$

The deviation from the ideal performance due to the second user is $\epsilon_i[n, 2]$. A video server is only really useful if (up to K) clients can start and stop watching video at arbitrary times, simultaneously. That is, the server must be able to deal with time-varying loads. Let $K(i)$ be the number of user requests being serviced at time i . It follows that the response of the n -th feature to this load is

$$\mathbf{x}_i[n] = K(i)X + \sum_{k=1}^{K(i)} \epsilon_i[n, k]. \quad (3)$$

We drop the square-wave and use the more flexible and general notation $K(i)X$, the number of active users at time i times the resources one user uses, X . We call this signal $\mathbf{l}_i[n] = K(i)X$, the load signal. Traces are available from an independent study in which the traces have a strong sinusoidal-like component [1]. The service level, RTP Packet Count, and load, TCP socket count, are plotted in Fig. 1. To fix ideas, we propose that a simple model for the load in these traces has the form

$$\mathbf{l}_i[n] = K(i)X \approx a[n] \cos(\omega_n i + \phi_n) + c. \quad (4)$$

The observed n -th feature is the linear combination:

$$\mathbf{x}_i[n] = a[n] \cos(\omega_n i + \phi_n) + c + \hat{\mathbf{x}}_i[n]. \quad (5)$$

The real-valued scalars, $a[n] \in \mathbb{R}$, $\omega_n \in \mathbb{R}$ and $\phi_n \in \mathbb{R}$, are the amplitude, radial frequency, and phase of the load –they

describe the user video request pattern. The constant c ensures that $\mathbf{l}_i[n]$ is a positive signal; the demand for resources should not be negative. We make the simplifying assumption that these parameters are constant. In a real-world system this is unlikely to be true, but it serves as a good first approximation. We also simplify our notation by introducing the notation $\hat{\mathbf{x}}_i[n] = \sum_{k=1}^{K(i)} \epsilon_i[n, k]$, for the aggregate deviation.

This model is general: The amplitude scales the load to give it a response in the correct range for the n -th feature; if the n -th feature is not a function of the load, $a[n] = 0$, and the n -th feature is $\mathbf{x}_i[n] = \hat{\mathbf{x}}_i[n]$ in system (Eqn. 5). The phase ϕ_n may capture the network and machine delay between when the request is made and the response given (cf. Fig. 1). This model is further generalized by considering loads which are parametric signals and/or stochastic processes. The service level metric is a linear function of the set of features (where the effects of the network are ignored as it is assumed to be sufficiently well resourced). The service level metric is:

$$y_i = \sum_n \mathbf{w}[n] \left\{ \frac{\mathbf{l}_i[n]}{(a[n] \cos(\omega_n i + \phi_n + \varphi_n) + c)} + \sum_{k=1}^{K(i)} \epsilon_i[n, k] \right\} \quad (6)$$

The additional phase terms $\varphi_n, \forall n$ capture the delay in the effect of the load due to client requests on the server machine, and network and server delays. In this paper we assume $\phi_n = \varphi = 0$ as the bandwidth is assumed to be large enough (due to the light load assumption). The clocks of the server and client are synchronized in [1] using NTP^{1c} and samples are collected every second. Note that we can substitute in an arbitrary expression for the load in (Eqn. 6) and the analysis in the rest of the paper holds.

System Characterization: We have introduced a deviations signal for each feature, $\hat{\mathbf{x}}_i[n]$, to explain the deviation of each feature from its ideal performance for a given load. A deviations signal is also required for the service level metric, $\hat{y}_i$. We want to explain the signal that captures the deviation of the service level metric from its ideal performance, as a function of the effect of the user requests on each of the video server's features (the deviation of each of the features from their ideal performance). The following model states that the observed deviation in the service metric is a weighted sum of the external causes, e.g. deviations in the features, plus internal causes, η_i , which captures non-idealities on the client's side.

$$\hat{y}_i = \mathbf{w}^T \hat{\mathbf{x}}_i + \eta_i. \quad (7)$$

However this model is significantly different from the model in (Eqn. 6). We make this explicit by indicating what we want (don't want) to model on both sides of the system:

$$\underbrace{y_i}_{\text{observed}} = \underbrace{\mathbf{w}^T \mathbf{l}_i}_{\text{don't want}} + \underbrace{\hat{y}_i}_{\text{want}} = \underbrace{\mathbf{w}^T \mathbf{l}_i}_{\text{don't want}} + \underbrace{\mathbf{w}^T \hat{\mathbf{x}}_i + \eta_i}_{\text{observed features}}. \quad (8)$$

The problem is that the signals that we observe, the pairs $\{\mathbf{x}_i, y_i\}$, are mixtures of what we want, $\hat{\mathbf{x}}$, and a high energy load component, $\mathbf{l}$, which we do not want. The reason why we distinguish between these two problems is that the load may potentially drown-out the deviation signals, $\hat{y}_i, \hat{\mathbf{x}}_i[n], \forall n$. In general a good approximation of the load is known, and therefore, there is little point in approximating it if it is already known, or worse, letting it bias the learning algorithm. The ability to estimate and predict deviations from ideal performance is the problem that is crucial to solve. The

approximation of the load comes from the TCPSCK field of the UNIX SAR command. The TCP socket count gives a good indication of the load on the kernel.

A Regression Tree, Random Forest [4], Lasso [3], RR or any other valid solver for problems of the form $y_i = \mathbf{w}\mathbf{x}_i$ or $\hat{y}_i = \mathbf{w}\hat{\mathbf{x}}_i$ will learn weights that solve the problem put to it. If the load is present in a model, when it should not be, e.g. if we pass $\{y_i, \mathbf{x}_i\}$ instead of $\{\hat{y}_i, \hat{\mathbf{x}}_i\}$ to the solver, we cannot expect the solver to correct the problem that is being asked. Asking the wrong question will generally yield the right answer for the wrong question. How can we learn a mapping between the kernel and service metrics which is independent of the load such that we can ask the right questions? Even more crucially, how can we measure the success of an approach that asks the correct questions? Are off-the-shelf measurement functions adversely affected by artifacts in the learning algorithm that arise due to the inappropriate presence of the load? In the next two sections we show that they are.

III. USING THE DATA TO SUPPORT THE MODEL

We investigate the extent to which the mean of the samples of the service level metric y depend on the underlying load on the system when these samples were drawn. Fig. 2 summarizes the statistics that characterize the values of y for different values of the load on the system, e.g. $l = 19, \dots, k, \dots$. The set of points used to construct each box-plot is the set of values of y corresponding to a given value of the load signal, $l = k$, and each set, and the set of the associated indices, are denoted

$$\mathcal{H}(y)|_{l=k}, \text{ and } \mathcal{I}(y)|_{l=k} \text{ respectively.} \quad (9)$$

Firstly, the mean of each of the sets, $\mathcal{H}(y)|_{l=k}$, which we denote $\mu(\mathcal{H}(y)|_{l=k})$, is different for each value of the load. For $19 \leq k \leq 30$ it is reasonable to assume that the model described above (in Eqn. 8) holds. However above $k = 30$, the values obtained by y generally decrease. In summary,

$$\begin{aligned} \mu(\mathcal{H}(y)|_{l=k}) &\propto k, & \text{for } 22 \leq k \leq 30 \\ \mu(\mathcal{H}(y)|_{l=k}) &\propto -k, & \text{for } 31 \leq k \leq 87 \end{aligned} \quad (10)$$

In words, the mean of the set of points of y for a given value of the load, $\mu(\mathcal{H}(y)|_{l=k})$, is proportional to the value of the load for loads less than 30 active requests, and proportional to the negative of the load when the load is greater than 30 active requests. Alternatively fitting a quadratic of the form

$$\mu(\mathcal{H}(y)|_{l=k}) \propto -a(k - 30)^2 + b \quad (11)$$

using the scalars $\{a, b\}$ gives a more concise description of the behaviour of the mean of y . Increasing the order of the polynomial (in Eqn. 10) increases the quality of the fit. Fig. 2 is significant because (Eqn. 7) assumes that both $\hat{y}_i$ and $\hat{\mathbf{x}}_i[n]$ are zero-mean signals. What is clear from Fig. 2 is that the mean of the y_i depends on the value of the load that was present on the system when it was observed. We do not attempt to fit parameters to the model (Eqn. 8) or assume that the mean of y_i and the features $\mathbf{x}_i[n]$ are the same irrespective of what the load was on the system.

Hypothesis testing: We use a hypothesis test as a first demonstration that the model (Eqn. 7) is of interest. The null hypothesis, namely ‘the load has no effect’, as used in [1], is that the service level metric has a mean which is approximately

equal to the $\mu(y) = 119.44$ irrespective of which samples are used to approximate it. We assume that the population standard deviation $\sigma(y)$ is unknown; we approximate it with the sample standard deviation. The value 119.44 is the mean of the $\approx 50k$ observed values of y . In words if we select any N_s samples of the signal y it should give a good estimate of $\mu(y)$. The alternative hypothesis, ‘the load has an effect’, is that we believe that the load has an effect on the values of y . The mean of the signal conditional on the load $l = k$ is $\mu(\mathcal{H}(y)|_{l=k})$. We also need the sample standard deviation of the service metric y conditional on the load, which is $\sigma(\mathcal{H}(y)|_{l=k})$. In summary, our hypotheses are

$$\begin{aligned} H_o : & \mu(y) \equiv \mu(\mathcal{H}(y)|_{l=k}) = 119.44 \forall k \\ H_a : & \mu(y) \neq \mu(\mathcal{H}(y)|_{l=k}), \forall k, k \neq k^* \end{aligned} \quad (12)$$

Does the load have an effect on the mean? Let us consider whether or not to accept or reject the null hypothesis. If the null hypothesis is true, what is the probability that we would have measured $\mu(\mathcal{H}(y)|_{l=k})$ as our estimate of the mean, $\mu(y)$. If the probability of the null hypothesis is really small we can reject it. We compute the z -statistic for each k

$$Z = \sqrt{N_s} \left(\frac{|\mu(y) - \mu(\mathcal{H}(y)|_{l=k})|}{\sigma(\mathcal{H}(y)|_{l=k})} \right) \quad (13)$$

and tabulate the associated probabilities in Table II. We compute the probability that the null hypothesis is true for values of the load that arise more than 100 times in the traces (and our choice of the z -statistic is justified). The values of the load for when this is true are indicated. The probability that the null hypothesis is true is zero in every case except for when the load is $k \in \{41, 42, 43, 44\}$, that is when $\mu(\mathcal{H}(y)|_{l=k}) \approx \mu(y)$. Fig. 3 illustrates the difference between the means. When $k = k^* = 42$, then $\mu(y) \approx \mu(\mathcal{H}(y)|_{l=k^*})$ with $\approx 14\%$ chance. In every other case applying RR to the data assuming that $\mu(y) \approx \mu(\mathcal{H}(y)|_{l=k^*})$, and thus identically distributed, is valid with less than 2% chance in one case, and $\approx 0\%$ in all other cases². This analysis motivates the following conclusions: 1) Different conditional means and variances for each load value imply that we need to learn a different regression model for each value of the load. 2) The linear model described in the previous section is insufficient because the RTP Packet Rate can increase with the number of active users until $k = 30$. Above this, the system begins to become saturated and the RTP Packet Rate begins to decrease as the number of active users increases. A piece-wise linear model, or some higher order polynomial model is required. Fig. 3 illustrates how good an estimate of the mean of the entire set of samples, $\mathcal{H}(y)|_{l=k}$ is.

IV. PREDICTION QUALITY MEASUREMENT: LA-NMAE

Using the correct measure of prediction performance is crucial if we are to distinguish between the performance of

²Remark: Our application of the z -statistic has a number of drawbacks. There is correlation between the samples which affects the conditional standard deviation of each sample. The sets $\mathcal{H}(y)|_{l=k}$ are in some sense anti-correlated and so the samples that we choose are not independent. The sample size used to generate each p -value is in general different, which affects the resolution of our estimates (under/over estimation of the sample standard deviation). Despite these shortcomings, the p -value gives a very strong recommendation that we reject the null hypothesis; for all loads but one. The assumption that the trace values are identically distributed holds with probability 0.

TABLE II. WHAT IS THE PROBABILITY THAT THE MEAN VALUE OF A SET OF SAMPLES DRAWN UNDER LOAD l HAVE THE MEAN $\mu(y)$? ONLY WHEN $l \in \{41, 42, 43, 44\}$ IS THE LIKELIHOOD NON-ZERO. IF THE NULL HYPOTHESIS IS INDICATIVE OF THE ASSUMPTIONS MADE BY THE LEARNING ALGORITHM, THEY DO NOT HOLD.

Fig. 2. The distribution of the service level metric y is illustrated for a range of loads on the system $19 \leq l \leq 88$. The largest load observed was 110, but above $l > 88$ there were too few samples to generate box-plots. X-tick labels are removed (for plotting purposes) as certain load values are never observed. The load values are indicated above each box-plot. We plot the mean (full line) and 1 standard-deviation (dashed lines).

Fig. 3. Difference between the mean of the entire set of service level samples, $\mu(y)$, and the service level samples conditional on the mean, $\mu(\mathcal{H}(y)|_{l=k})$. For each load value a stem is drawn from zero to the amplitude of the difference. In most cases, the stem length is greater than ± 20 units.

competing prediction algorithms with confidence. Previous works have considered the T-NMAE between the signal to be predicted, y_i and the prediction estimate, $\hat{y}_i$

$$S = \frac{100}{\mu(y)N_s} \left(\sum_i |y_i - \hat{y}_i| \right) = \sum_i \frac{100}{\mu(y)N_s} |y_i - \hat{y}_i| \quad (14)$$

They scale the score by 100 to obtain a percentage. Re-ordering the summation and the constant is useful as it makes the argument below more intuitive.

In this section we show that 1) the load dependence of the mean described above renders the T-NMAE measure suspect to overestimation of the prediction accuracy in many cases; 2) the dependence of the T-NMAE on the population mean $\mu(y)$ potentially dominates the sensitivity of the T-NMAE measurement to just one (unimportant) statistic of the signal being predicted. Given that the first step of many estimation procedures is to center the signal, it is troubling for the mean to have such a dominant effect on prediction performance.

1) Load dependence: The T-NMAE produces an aggregate score for the prediction error between y_i and $\hat{y}_i$ for a set of signal values y_i , $i = 1, \dots, N_s$. The underpinning assumption is that the signal values to be estimated are picked from the same distribution. Each prediction produces an error $\epsilon_i = y_i - \hat{y}_i$. The T-NMAE assumes that each prediction error is equally important, because the signal values y_i are taken from the same distribution. Therefore each error is scaled by $\frac{100}{N_s}$. Finally, in order to give this number context, the T-NMAE scales the weighted error by the typical value that the signal, y_i , achieves, e.g. $\mu(y)$, and sums up the values. Herein lies the

problem however. When the load affects the mean in the manner we described above, $\mu(\mathcal{H}(y)|_{l=k})$, the value of the “typical value” of y_i changes too, as a function of the load. For a given error in our prediction, the global mean, $\mu(y)$ and the load dependent mean scale the error differently. So how does the load, in particular the mean dependent load, $\mu(\mathcal{H}(y)|_{l=k})$, affect the contribution of the error in one sample to the T-NMAE? To answer this question, we consider the maximum possible prediction error in the trace above. For a given value of the load, e.g. $l = k$, the maximum value is $\max \mathcal{H}(y)|_{l=k}$ (assuming y is non-negative). The NMAE should depend on the load, because failing this, some errors are scaled unfairly by a mean value which is not representative of the distribution from which they were drawn. Therefore the load-adjusted NMAE (LA-NMAE) is

$$S_k = \sum_{i \in \mathcal{I}(y)|_{l=k}} \frac{100}{N_s \mu(\mathcal{H}(y)|_{l=k})} |\epsilon_i|. \quad (15)$$

In this new measure, the LA-NMAE, the errors ϵ_i due to the predictions $\hat{y}_i$ correspond to the samples y_i which were generated under the load condition $l = k$. To evaluate how important it is to select the correct mean in the LA-NMAE, we vary the error ϵ_i , as a function of each value of the load, from 0 to $\max \mathcal{H}(y)|_{l=k}$. We compute the NMAE using the two definitions above, for each value of the error $0 \leq \epsilon_i \leq \max \mathcal{H}(y)|_{l=k}$. We plot the pairs of resulting NMAEs for each value of the load l . Fig. 4 illustrates how both NMAEs penalize errors in the case where the load is $49 \leq l \leq 56$. These results should be read as follows. If the absolute error in the prediction, *irrespective of which prediction algorithm was used to generate the prediction*, for a given value of the load l , was e , the associated contribution of this value to the total T-NMAE is $p\%$ if the global mean was used, e.g. $\mu(y)$. The percentage p can be found by finding the y -value of e on the black dashed line. Whereas if the load-adjusted mean was used $\mu(\mathcal{H}(y)|_{l=k})$, the percentage, p , may be obtained by finding the y value corresponding to e using the blue line.

Some crucial observations are listed as follows: 1) The full

Fig. 4. For a given prediction error and system load condition, what is the contribution of the error to the prediction accuracy score? The traditional NMAE (T-NMAE), dashed line, under-estimates the contribution of a prediction error to the score for many load conditions compared to the Load-Adjusted-NMAE (LA-NMAE). We illustrate this for loads $49 \leq l \leq 56$.

TABLE III. PATHOLOGICAL PREDICTION PROBLEM SOLVER

1. Initialization: Set the constant $a = 0$, store the serv. metrics: $y_{orig} = y$.
2. Pick N_s values from any distribution (e.g. a normal dist.), assign $\hat{y}_i \leftarrow \mathcal{N}(0, 1)$.
3. Assign $a = a + 1$;
4. Assign $y_i \leftarrow y_i + a$ and $\hat{y}_i \leftarrow \hat{y}_i + a$
5. Compute the T-NMAE $S = \sum_i |y_i - \hat{y}_i| \frac{100}{N_s \mu(y)}$
6. If $S < 1$ break (the prediction is $\hat{y}_i$); else return to step 2.

line is higher for all prediction errors, ϵ_i , for approximately 74.7% of the values the load can take and 61.295% of all of the signal values; 2) The scaled-error values computed using the global mean $\mu(y)$ are significantly smaller for approximately 61% of the values used to compute the T-NMAE than they should be; 3) The prediction performance given using the global mean scaled T-NMAE is quoted to be better (using the measure S) than it actually is (using the measure S_k). It is straightforward to compute by how much S over-inflates the accuracy of the prediction algorithm compared to S_k . For a given load value $l = k$ let the weight of proportionality be α

$$\alpha S_k = S, \quad \text{which implies } \alpha = \frac{\mu(\mathcal{H}(y)|_{l=k})}{\mu(y)}. \quad (16)$$

In words, if we divide a given T-NMAE, S , by the inflation weight α we get the correct LA-NMAE. For example in Fig. 4, $\mu(\mathcal{H}(y)|_{l=k}) < \mu(y)$ when the load is above $k^* = 42$. Taking the case when $k = 70$, $\mu(y) \approx 119$ and $\mu(\mathcal{H}(y)|_{l=k}) \approx 80$. These values imply that $\alpha = \frac{80}{119} \approx 0.66$, which means that an T-NMAE of $S = 11\%$ equals a LA-NMAE of $S_k = 17\%$, an T-NMAE of $S = 20\%$ equals a LA-NMAE of $S_k = 30\%$ and so on. In short, the correct LA-NMAE is 50% worse in many cases above, when the load-adjusted mean is used.

This is a very practical result. If the absolute prediction error ϵ is 1 unit for a particular sample the difference between the NMAEs, S and S_k , is .5%; for $\epsilon = 5$, the difference is 2.5%; for $\epsilon = 20$, the difference is 10% and so on. A break down of these differences for given values of the load can be obtained in Table IV. The numbers of samples that fall into each category are listed along with the load values. For example, 822 samples in the traces are drawn under a load $k = 70$. The error in the reported error using S , when the errors range from 1 to 40 units, ranges from .48% to 19.11%. Note that the error can easily be greater than 40, and in this case the error in the reported percentage is larger.

2) Dominance of $\mu(y)$: It is not reasonable to claim that the T-NMAE allows for the comparison of prediction accuracy of different service level metrics across different scenarios

Fig. 5. During the learning phase, kernel and service metrics enter on the upper LHS arrow into a switch that determines –based on the load value $x_i[n^*]$ – which learning model to learn. During the prediction phase, the same 2-level approach is taken. The load, $x_{i+1}[n^*]$, and features, $x_{i+1}[n]$, enter on the lower LHS arrow, and the appropriate prediction model is chosen based on $x_{i+1}[n^*]$, to produce the prediction $\hat{y}_{i+1}$ on the RHS.

and loads. We demonstrate this by considering the following pathological problem-solver pair, which illustrates the counter-intuitive behaviour of the T-NMAE S .

Problem 1: Predict N_s values of y_i , using any prediction algorithm, such that the T-NMAE, S , of the errors $y_i - \hat{y}_i$ are less than 1%. Consider the valid approach in Table III. At first glance, Problem 1 looks like a reasonable statement of the video service level prediction we are interested in solving. Note however that the value of the mean $\mu(y)$ is crucial. As a increases, the T-NMAE goes to zero $S \mapsto 0$, in general, irrespective of what $\hat{y}_i$ is, or how it was generated. This is because $\mu(y) = \mu(y_{orig}) + a$. In the more general setting of comparing the performance of predictions of service level metrics it is clear that the mean of the service level metric observed is crucial as it sets the sensitivity of the performance to deviations in performance. In summary, the comparison of prediction performance across services is not meaningful unless the services have the same mean. If they do not have the same mean, what value is the appropriate value for the mean so that the comparison is fair? We cannot use 0 as this gives a T-NMAE of ∞ . We risk inflating the performance of our predictor by picking an arbitrary value. Therefore it is not reasonable to claim that the T-NMAE allows for the comparison of prediction accuracy of different service level metrics. We draw the following conclusions: 1) prediction performance across the samples using the T-NMAE is unreliable; performance should be measured relative to the load on the system. 2) The prediction performance across services is unreliable using the T-NMAE due to the difference in the mean values of the traces. The LA-NMAE does fix the first problem with prediction performance measurement, the load; and we have raised awareness of the problems associated with comparing prediction performance across services, loads and scenarios. As an aside, we note that Signal-to-Noise-Ratio-like (SNR) measures and the Root Mean Square Error (RMSE) suffer from a similar dependence on the load. We do not give a full treatment for these measures in this paper as measures derived from the NMAE are sufficient to provide a like-for-like comparison with the state-of-the-art.

V. LOAD-ADJUSTED LEARNING AND PREDICTION

We have contributed practical results for modeling video metrics under different load conditions. We have also demonstrated how to measure the performance of the predictor under different load conditions. We now illustrate the flow of control of a learning and prediction algorithm pair that use knowledge

TABLE V. LA-RR: COMPARISON OF LA-NMAE, T-NMAE, DIFFERENCE IN NMAE MEASUREMENT.

[illegible]

TABLE VI. T-RR: COMPARISON OF LA-NMAE, T-NMAE, DIFFERENCE IN NMAE MEASUREMENT (NAN: NOT A NUMBER).

		load values l																																																																																																																																																																																																																																																																																																																																																																																																															
$\frac{v_1}{v_2}$	$\frac{v_1}{v_2}$ (%)	$\frac{v_1}{v_2}$																																																																																																																																																																																																																																																																																																																																																																																																															
		> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100	> 100

TABLE VII. LA-RR PERFORMANCE GAIN

Measure	μ_{gain}	σ_{gain}	min gain	max gain
LA-NMAE Gain	3.76%	5.57	0.03%	32.20%
T-NMAE Gain	2.77%	3.52	0.048%	20.30%

tics which describe the performance of the LA-RR and T-RR algorithms in Table VII. We use both the T-NMAE and LA-NMAE to evaluate the performance. The mean performance gain achieved by the LA-RR is 3.76% if the LA-NMAE is used and 2.77% if the T-NMAE is used to compare performance. These mean values are computed by taking the mean of the mean LA-NMAE for each load value, and the mean of the mean T-NMAE for each load value. For completeness, we list the minimum and maximum performance gain achieved by the LA-RR routine over all load values. These performance gains are positive and range from 0.03% to 32.30% for the LA-NMAE measure and 0.048% to 20.30% for the T-NMAE measure. We remove the outliers when we compute the maximum performance gain, e.g. if the performance gain is greater than 100% this has generally occurred because the T-RR routine has failed to find good prediction weights. Even with a regularization parameter which is as large as 5000, the T-RR algorithm is ill-conditioned. The LA-RR is generally successful with a much lower regularization parameter –the problem is better behaved because the load has been removed. The minimum performance gain corresponds to the case where the load is $\mathbf{l} = k^*$, in this case, the load conditional mean is approximately equal to the mean value of the trace. We also list the standard deviations of the performance gains for both NMAE measures. In summary, for all values of the load, k , both the T-NMAE and LA-NMAE scores support the claim that LA-RR outperforms T-RR. In some cases, the performance gain is 30%, which is significant.

Comparison of prediction performance metrics: We have established that on average the LA-RR outperforms the T-RR routine for all values of the load irrespective of whether the LA-NMAE measurement function is used, or the T-NMAE measurement function is used. We examine the performance of the LA-NMAE and T-NMAE measurement functions in more detail. Tables VI and V list the LA-NMAE, S_k , the T-NMAE, S , and the difference between the T-NMAE and LA-NMAE, $S - S_k$. We draw the following conclusions:

1: In both Table VI and V, S_k the LA-NMAE, is lower than S , the T-NMAE, when the load is less than $l = 39$. When the load is low, the LA-NMAE states that the prediction performance is better than the T-NMAE. Using the measure, S , penalizes the

Fig. 6. T-NMAE is correlated with the $\mu(\mathcal{H}(y)_1)$, LA-NMAE is not.

prediction algorithm unfairly for when the load on the system is small. Unfortunately there are more load values greater than $l = 39$ than less than $l = 39$.

2: In both Table VI and V, S_k the LA-NMAE is higher than S , when the load is $l > 39$. When the load is high, the T-NMAE states that the prediction performance is better than it actually is, giving rise to inflated but unjustified confidence in the performance of the predictor. In row three of both tables, $S - S_k$ lists the difference in the prediction performance measurements for the T-NMAE and LA-NMAE. A negative percentage for a given load value, means that on average the error in the measurement of the prediction performance by the T-NMAE acts to reduce the error in the prediction performance. For example, when $l = 71$ active users, the T-NMAE gives the performance as being 13.41% accurate, whereas the LA-NMAE gives the accuracy as 21.7%. This discrepancy, -8.33% leads to over-confidence, it is founded on a measure which has reduced the sensitivity of the measurement function to errors by standardizing the score with a mean value which has no relation to the distribution from which the predictions are made. Similarly, when the discrepancy is positive, it follows that the T-NMAE gives a pessimistic measure of the prediction performance.

3: In general more result inflation than result deflation occurs. For 64% of the load values, which have greater than 231 samples associated with them, the percentage of the T-NMAE, S , measurement is made smaller which inflates our confidence of the success of the prediction algorithm. Generally, the inflation is in the range 4–7%.

4: A number of high level trends warrant reporting. We plot box-plots of the spread of LA-NMAE and T-NMAE

measurements in our Monte Carlo trials in Fig. 6. As the load on the system increases its behaviour becomes more volatile; prediction becomes harder because the behaviour of the system is more unpredictable. What is significant is that the LA-NMAE does not seem to penalize as a function of the parabolic nature of the relationship between the RTP Packet Count metric and the underlying load on the system (cf. Fig. 2). We posit that this indicates that the LA-NMAE measures the error in the predictions, *independently* of the artifacts introduced by considering the load. Note that the T-NMAE measure, S , increases initially until $k = 30$ and then decreases in Fig. 6. The T-NMAE's measurements are clearly correlated with the parabolic relationship between the RTP Packet Count and the underlying load on the system. It is unusual that the T-NMAE's measurement of prediction error should decrease with the load, that is, that we can predict the behaviour of the system better if the load is increased. If we accept this measure as being a reasonable measure of performance, a *misguided* conclusion we could draw from this plot is that we should be able to make better predictions of how resource behave if we have fewer resources per user. In comparison the LA-NMAE, S_k , is an approximately linearly increasing function of the load; it does not exhibit change-points like the T-NMAE measure S . When there are more active users using a fixed set of resources the performance of the resource becomes less predicible.

Implementation Considerations: In terms of prediction performance, the LA-RR gives empirically more accurate predictions irrespective of which NMAE measure is used than T-RR. It is a well-principled approach; it acknowledges the effect changing the load has on the observations drawn from the system –it is a better model. In terms of complexity and the ability to perform real-time prediction, we achieve better results for all load values using only 2% of the data; what at first looks like a big data problem, is in fact manageable even with limited bandwidth and compute resources. Instead of using $\approx 50k$ samples training data, we use $\approx .2$ to 1k and achieve better results. Given that the asymptotic complexity of RR is $O(N^2 N_s)$, when $N_s > N$, where N_s is the number of training samples and $N = 231$ is the number of features, a comparison of the complexity of LA-RR, $O(230^2 \times 1000)$, with T-RR, $O(230^2 \times 50000)$, motives the selection of LA-RR. This reduction supports the claim that real-time prediction of video service level metrics are potentially computationally achievable. Finally, we observed that T-RR quickly runs into numerical difficulty due to the condition of its matrix inverse. This is due to the presence of the load. When the load signal is dominant the matrix inverse involved in computing T-RR becomes more singular and thus a large regularization parameter was necessary to ensure non-singularity. “NaN” is included in Table. VI when the solver is ill-conditioned. However, a large regression parameter increases the chances of having a large bias [3] in the prediction weights. The LA-RR matrix inverse is generally better conditioned.

VII. CONCLUSIONS

Many SL tools exist to organize data, measure prediction performance, make predictions about future video service level metrics –essentially to guess at the state of the system (server-client metrics) that gave rise to the data. The object of any such study should be the system, and not artifacts of the way we use SL tools to probe its secrets. We have demonstrated that

applying SL techniques and measurement functions, without first examining the underlying structure of the system, may introduce artifacts. Any subsequent effort to predict what the system will do next is quixotic and may lead to over-confidence. In response to this we introduced the idea of load-adjusted learning, prediction and measurement to remove these artifacts. Our empirical and analytical findings support the claim that adopting a load-adjusted processing philosophy gives more accurate prediction. They also support the assertion that if this approach is not adopted the load on the system can inadvertently affect the result. Given the recent interest in leveraging SL in Networking this results warrants reporting.

REFERENCES

- [1] R. Yanggratoke, J. Ahmed, J. Ardelius, C. Flinta, A. Johnsson, D. Gillblad, and R. Stadler, “Predicting real-time service-level metrics from device statistics,” *IFIP/IEEE Int. Sym. Int. Net. Man.*, pp. 1–8, 2015.
- [2] R. de Frin, C. Olariu, Y. Song, R. Brennan, P. McDonagh, A. Hava, C. Thorpe, J. Murphy, L. Murphy, and P. French, “Integration of QoS Metrics, Rules and Semantic Uplift for Advanced IPTV Monitoring,” *Journal of Network and Systems Management*, pp. 1–36, 2014.
- [3] R. Tibshirani, “Regression Shrinkage and Selection Via the Lasso,” *J. Roy. Stat. Soc., Series B*, vol. 58, pp. 267–88, 1994.
- [4] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [5] A. Hoerl and R. Kennard, “Ridge regression: Applications to nonorthogonal problems,” *Technometrics*, vol. 12, no. 1, pp. 69–82, 1970.
- [6] R. de Frin, “Effect of system load on video service metrics,” in *IEEE Irish Sig. Sys. Conf. (ISSC)*, 2015, pp. 1–6.
- [7] R. de Frin, K. Drakakis, and S. Rickard, “Portfolio diversification using subspace factorisations,” in *42nd An. IEEE Conf. Inf. Sc. Sys. (CISS)*, 2008, pp. 1075–80.
- [8] M. Zibulevsky and B. Pearlmutter, “Blind source separation by sparse decomposition in a signal dictionary,” *Neur. Comp.*, vol. 13, no. 4, pp. 863–82, 2001.
- [9] R. de Frin, “Ghostbusters: A parts-based NMF algorithm,” in *IET Irish Sig. Sys. Conf. (ISSC)*, 2013, pp. 1–8.
- [10] R. de Frin, “Formal concept analysis via atomic priming,” in *Formal Concept Analysis*, P. Cellier, F. Distel, and B. Ganter, Eds., vol. 7880 of *LNCS*, pp. 92–108. Springer Berlin Heidelberg, 2013.
- [11] S. Hofmeyr, S. Forrest, and A. Somayaji, “Intrusion detection using sequences of system calls,” *J. Comp. Sec.*, vol. 6, pp. 151–80, 1998.
- [12] Cisco, “Cisco visual networking index, Global IP Traffic Forecast, 2011–2016,” *white paper*, 2011.
- [13] A. Begen, T. Akgul, and M. Baugher, “Watching video over the web Part 2: Applications, standardization, and open issues,” *IEEE Internet Comp.*, vol. 15, no. 3, pp. 59–63, 2011.
- [14] M. Zinkevich, A. Smola, and J. Langford, “Slow learners are fast,” in *NIPS*, 2009, pp. 2331–9.
- [15] D. Hsu, N. Karampatziakis, J. Langford, and A. Smola, “Parallel online learning,” *CoRR*, vol. abs/1103.4204, 2011.
- [16] J. Bogojeska, D. Lanyi, I. Giurgiu, G. Stark, and D. Wiesmann, “Classifying server behavior and predicting impact of modernization actions,” in *Int. Conf. Net. and Serv. Man. (CNSM)*, 2013, pp. 59–66.
- [17] M. Mirza, J. Sommers, P. Barford, and X. Zhu, “A Machine Learning Approach to TCP Throughput Prediction,” *Networking, IEEE/ACM Trans.*, vol. 18, no. 4, pp. 1026–39, 2010.
- [18] A. Andrzejak and L. Silva, “Using machine learning for non-intrusive modeling and prediction of software aging,” in *IEEE Net. Op. Man. Symp. (NOMS)*, 2008, pp. 25–32.
- [19] F. Zaman, S. Robitzsch, Zhuo W., J. Keeney, S. van der Meer, and G. Muntean, “A heuristic correlation algorithm for data reduction through noise detection in stream-based communication management systems,” in *IEEE Net. Op. Man. Symp. (NOMS)*, 2014, pp. 1–8.
- [20] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: Applications in R*, Springer Texts in Stats. 2013.