

Online feature selection for rapid, low-overhead learning in networked systems

Xiaoxuan Wang[†] Forough Shahab Samani^{†‡} and Rolf Stadler^{†‡}

[†] Dept. of Computer Science, KTH Royal Institute of Technology, Sweden

[‡] RISE Research Institutes of Sweden

Email: {xiaoxuan, foro, stadler}@kth.se

Abstract— Data-driven functions for operation and management often require measurements collected through monitoring for model training and prediction. The number of data sources can be very large, which requires a significant communication and computing overhead to continuously extract and collect this data, as well as to train and update the machine-learning models. We present an online algorithm, called OSFS, that selects a small feature set from a large number of available data sources, which allows for rapid, low-overhead, and effective learning and prediction. OSFS is instantiated with a feature ranking algorithm and applies the concept of a stable feature set, which we introduce in the paper. We perform extensive, experimental evaluation of our method on data from an in-house testbed. We find that OSFS requires several hundreds measurements to reduce the number of data sources by two orders of magnitude, from which models are trained with acceptable prediction accuracy. While our method is heuristic and can be improved in many ways, the results clearly suggests that many learning tasks do not require a lengthy monitoring phase and expensive offline training.

Index Terms—Data-driven engineering, Machine learning (ML), Dimensionality reduction

I. INTRODUCTION

Data-driven network and systems engineering is based upon applying AI/ML methods to data collected from an infrastructure in order to build novel functionality and management capabilities. This is achieved through learning tasks that use this data for training. Examples are KPI prediction and forecasting through regression and anomaly detection through clustering techniques.

Data sources that feed the learning tasks include system logs, telemetry data, and real-time measurements collected through monitoring. The number of available data sources can be very high, even in small systems. For example, on our testbed at KTH, which includes 10 compute servers, we can extract several thousand metrics from the operating system and orchestration layers. Since these metrics are dynamic, we monitor them periodically, e.g., once per second. It requires a significant overhead to extract and collect this data, as well as a significant computational overhead to train and update the machine-learning models that underlie the learning tasks. Considering the fact that the monitoring and computational overhead increases at least linearly with the number of measurements and the dimensionality of the input, i.e., the number of (one-dimensional) data sources, it becomes vital to reduce the number of data sources to the extent possible.

The focus of this paper is on a novel online source-selection method that requires only a small number of measurements to significantly reduce the number of sources needed for training models that are effective for learning tasks. As a result, the communication overhead for monitoring and the computational overhead and time needed for the model training are significantly reduced.

Using the terminology of machine learning, we call a (one-dimensional, scalar) data source also a *feature*, and we refer to measurements taken from a set of data sources at a specific time as a *sample*.

Our approach consists of (1) ranking the available data sources using (unsupervised) feature selection algorithms and (2) identifying *stable feature* sets that include only the top k features. We call a feature set stable, if it remains sufficiently similar when additional samples are considered.

We evaluate our approach using traces from an in-house testbed that runs two services under different load conditions. The results show that our method can reduce the number of data sources needed for learning tasks by two orders of magnitude, while still achieving acceptable errors for a prediction task. The reduction in input dimensionality is consistent with results from our earlier work, which studies non-linear methods for dimensionality reduction in an offline setting [1]. We find that a stable feature set can often be identified with only a few hundred samples. As a consequence, the monitoring effort required for collecting the data for a machine-learning task can be reduced quickly. We consider this capability key for increasing the acceptance of data-driven engineering solutions.

With this paper, we make the following contributions:

- We present an online algorithm, which we call OSFS, that selects a small feature set from a large number of available data sources using a small number of measurements, which allows for rapid, low-overhead, and effective learning. The algorithm is initialized with a feature ranking algorithm and applies the concept of a stable feature set, which we introduce in the paper.
- We perform an extensive, experimental evaluation of our method on an in-house testbed.

The significance of our findings lies in the prospect that many data-driven functions in networked systems can be trained rapidly and with low overhead and thus do not require a lengthy monitoring phase and expensive offline training.

This paper contains results from a master thesis project conducted at KTH [2].

An extended version of this paper can be found at arxiv [3]. It contains the description of the experimental infrastructure and the load pattern used for creating the traces that are used for evaluation in this paper, a related work section, and additional experimental results.

II. PROBLEM FORMULATION AND APPROACH

We consider a monitoring infrastructure that collects readings from a set F of n distributed data sources (or features). Each feature has a one-dimensional, numerical value that changes over time. We collect readings at discrete times t and store them in sample vectors $X_t \in \mathbb{R}^n, t = 1, 2, 3, \dots$. Our plan is to identify a subset $F_k \subset F$ with $k \ll n$ features using the samples X_1, X_2, \dots, X_{t_k} .

Second, we consider a learning task, like KPI prediction or anomaly detection, whose model is trained using the samples $X_t \in \mathbb{R}^k$ with the features from F_k .

In order to keep low the monitoring overhead for collecting the samples and the computational overhead for training the model associated with the learning task, the numbers for k and t_k should be small. Note that k indicates the number of data sources that need to be monitored to train the model and that t_k refers to the number of measurements that are needed to compute F_k . Assuming periodic measurements, t_k further indicates the time it takes until the feature set F_k is available.

Our objective thus is to select k and t_k as small as possible, while enabling the models trained using F_k to be equally (or similarly) effective for prediction as those trained using the complete feature set F and a large number of samples.

The task of selecting a subset of features from a larger set is called *feature selection* in machine learning and data mining and is a well-studied topic area. We are specifically interested in unsupervised feature selection methods, whereby the values of the target are not known during the feature selection process, i.e., the process to compute F_k . This allows us to keep the feature selection process independent from the the learning task and will enable different learning tasks in a system to share the same feature subset.

The problem we address in this paper is to find an online algorithm that reads a sequence of n -dimensional sample vectors X_1, X_2, \dots one by one, computes k and the feature set F_k , and terminates after step t_k . The values for k and t_k should be small, while F_k must be effective in training models for learning tasks.

In our approach, we choose an unsupervised feature selection method that ranks the n features in every step t of the online algorithm and checks how the top i features ($i = 1, \dots, n$) change with increasing t . We introduce a similarity metric that captures this change (see Section IV). If the similarity between the top k features in consecutive steps is high and not increasing anymore, the algorithm terminates and the values for F_k , k and t_k are returned (see Section V).

Note that we assume here that the feature set F is fixed. In a real system that runs over some time, changes to the physical

configuration or the virtualization layer occur, which result in changes to the set of available measurement points, i.e. the feature set. In such a case, the online algorithm must be restarted. Note also that we do not investigate in this work how many samples are needed to train the model of the learning task. This will be done in future work.

III. CREATING RANKED FEATURE LISTS

In this section, we describe two algorithms (ARR, LS) that produce a ranked feature list from a list of samples. They are based on unsupervised feature selection methods from the literature. In addition, we include a third ranking algorithm (TB), which is supervised and will serve as a baseline. We will evaluate the suitability of these algorithms for our online feature selection method OSFS in Section V.

Table 1 shows the notation we use in the paper. The available data for computing the feature set F_k is presented as a design matrix $X \in \mathbb{R}^{m \times n}$, whose n columns represent the feature vectors and m rows represent the samples in the data set. Since we assume that the samples arrive in sequence one-by-one, m is increasing over time and can be interpreted as time index.

TABLE 1
Table of notation

X	data set
n	number of features
m	number of samples
$X_{i,:}$ or $X_i (i = 1, \dots, m)$	i -th row or i -th sample of X
$X_{:,j} (j = 1, \dots, n)$	j -th column or j -th feature vector of X
$X_{i,j}$	element of the i -th row and the j -th column of X
k	number of selected features
t_k	number of samples used for feature selection
F	set of all available features
F_k	subset with k selected features

The first algorithm is Adapted Relevance Redundancy Feature Selection (ARR). It is based on the Relevance Redundancy Feature Selection (RRFS) method [4], which we adapted to compute a ranked feature list. ARR uses two criteria in assessing the rank of a feature: (a) relevance, which relates to the distance of a feature vector to the mean of all feature vector vectors, and (b) redundancy, which relates to the cosine similarity between a feature vector and the vectors of all other features. High relevance and low similarity result in a high score. The pseudo-code of ARR is given in Algorithm 1. First, the relevance of a feature is computed as the mean absolute difference of its feature vector from the mean (line 3). Then, the relevance of a feature is computed as the sum of the cosine similarity of its feature vector and each feature vector of the data set (lines 5-6). The score for ranking a feature is the relevance value divided by the redundancy value (line 7). The computational complexity of ARR is $O(n^2m)$. Recall that the cosine similarity between two features vectors $X_{:,a}$ and $X_{:,b}$ is calculated as:

$$\text{cosim}(X_{:,a}, X_{:,b}) = \left| \frac{\sum_{i=1}^m (X_{i,a} X_{i,b})}{(\sqrt{\sum_{i=1}^m X_{i,a}^2})(\sqrt{\sum_{i=1}^m X_{i,b}^2})} \right| \quad (1)$$

Algorithm 1: Adapted Relevance Redundancy Feature Selection (ARR)

Input: Data matrix $X \in \mathbb{R}^{m \times n}$
Output: Ranked feature list F'

- 1 $F' = []$;
- 2 **for** $i = 1; i \leq n; i++$ **do**
- 3 $relevance_i = \sum_{j=1}^m |X_{j,i} - \bar{X}_{:,i}|$;
- 4 $sim_sum_i = 0$;
- 5 **for** $j = 1; j \leq n; j++$ **do**
- 6 $sim_sum_i += \cosim(X_{:,i}, X_{:,j})$;
- 7 $score_i = \frac{relevance_i}{sim_sum_i}$;
- 8 Construct F' as list of all features sorted by $score_i$ in descending order;
- 9 **return** F' ;

The second algorithm is Laplacian Score (LS) [5]. It follows the so-called filter method which examines intrinsic properties of the data to evaluate the features. LS ranks those features high that preserve locality with respect to a neighborhood graph. Algorithm 2 shows the pseudocode of LS. Input parameters for this algorithm are the design matrix and the number of local neighbors K . First, using the m sample vectors as the nodes of the graph, a neighborhood graph is constructed with the K nearest neighbors of each node as the links of the graph (line 1). Then, the graph connectivity and the distance between node pairs is used to compute the weight matrix (line 2). The graph Laplacian matrix is computed in line 4. The Laplacian score of all n features is obtained in the for loop (lines 5-7). A low score for a feature signifies high locality preservation (see [5] for justification), and the features are ranked according to increasing score (line 8). The computational complexity of the LS algorithm is $O(nm^2)$.

In the evaluations reported in Sections V and VI of this paper, we choose K in relation to the value of m : $K = 2$ for $0 < m \leq 16$; $K = 5$ for $16 < m \leq 128$; $K = 10$ for $m > 128$.

The third algorithm for feature ranking is tree-based feature selection (TB). TB is a supervised algorithm in which the feature importance is computed using random forest prediction [6] [7]. We consider TB a baseline to better assess the performance of ARR and LS. In this work, scikit-learn library is used to compute the TB feature ranking [8] [9]. We set the number of trees to 100 and use default values for the remaining parameters. The computational complexity of TB is $O(Tnm \log_2 m)$ where T is the number of trees, n is the number of features, and m is the number of samples [10].

IV. COMPUTING STABLE FEATURE SETS

In Section III we discussed the algorithms ARR, LS, and TB, which produce ranked feature lists after reading m samples. In an online setting, where samples become available one-by-one at discrete times $t = 1, 2, \dots$, the value of m can be interpreted as time. Let's assume we run a feature

Algorithm 2: Laplacian Score (LS)

Input: Data matrix $X \in \mathbb{R}^{m \times n}$, $K \in \mathbb{N}$ nearest neighbors
Output: Ranked feature list F'

- 1 Construct graph G of K nearest neighbors from nodes $X_{i,:}$, $i = 1, \dots, m$;
- 2 Compute weight matrix $S \in \mathbb{R}^{m \times m}$ from G

$$S_{ij} = \begin{cases} e^{-\|X_{i,:} - X_{j,:}\|^2} & \text{if nodes } i \text{ and } j \text{ are connected,} \\ 0 & \text{otherwise} \end{cases}$$

- 3 $D = \text{diag}(SI)$ where $I = [1, \dots, 1]^T$, $D \in \mathbb{R}^{m \times m}$;
- 4 $L = D - S$;
- 5 **for** $i = 1; i \leq n; i++$ **do**
- 6 $V_i = X_{:,i} - \frac{X_{:,i}^T D I}{I^T D I} I$ where $V_i \in \mathbb{R}^{m \times 1}$;
- 7 $lscore_i = \frac{V_i^T L V_i}{V_i^T D V_i}$;
- 8 Construct F' as list of all features sorted by $lscore_i$ in ascending order;
- 9 **return** F' ;

ranking algorithm at time t_1 and compute the set with the top k features, F_{k,t_1} . This set will generally be different from the set F_{k,t_2} produced by the same algorithm at a later time t_2 . Using the standard assumption from statistical learning that samples are drawn from static distributions [11], we can assume that the sets $F_{k,t}$ converge to a set F^*_k with growing time t . Our objective for the remainder of the section is to find a heuristic criterion to determine at which time the feature sets $F_{k,t}$, $t = 1, 2, 3, \dots$ have sufficiently converged or, as we also say, have become “stable”.

We first introduce a metric $sim()$ that captures the similarity between two sets A and B, which both contain $k > 0$ elements:

$$sim(A, B) := |A \cap B|/k \quad (2)$$

The value of sim is between 0 and 1. 0 means no similarity between A and B, 1 means maximum similarity, i.e. A and B are identical.

We study the evolution of $sim(F_{k,t}, F^*_k)$ over time using the data traces from our testbed. In this case, F^*_k denotes a feature set that has been computed using an entire trace of samples (i.e. 20 000 - 50 000 samples, depending on the trace).

Figure 1 gives in the top row the value of $sim(F_{k,t}, F^*_k)$ in function of time, computed by the ranking algorithms ARR and LS on the KV flash-crowd trace. The curves show the values for $k = 4, 16, 64, 256$. Each point on the curves is the mean of 10 values from different start times, whereby one start time is $t=1$, and the other nine are chosen uniformly at random between $t=2$ and $t=10\,000$. For each start time, $F_{k,t}$ is computed and compared with F^*_k . (This applies to the top and bottom row. The curves in the middle row are based on a single start point with $t = 1$ to better illustrate values encountered in an online scenario.)

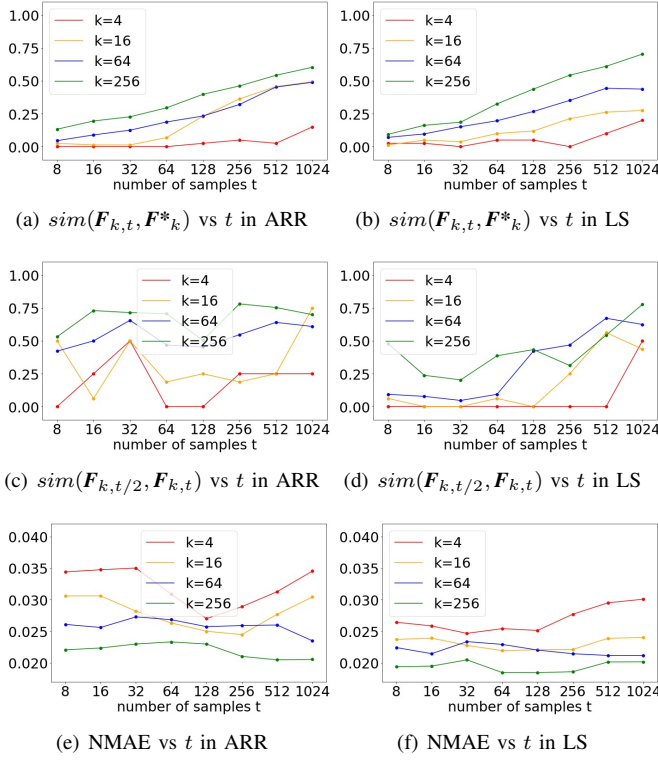


Fig. 1. The similarity of consecutive feature sets over time using ARR and LS with samples from the KV flash-crowd data set and start point $t = 1$. The bottom row shows the prediction accuracy of a random-forest regressor using the feature set $\mathbf{F}_{k,t}$ for training.

As expected, the similarity values tend to increase over time, which means that the sets $\mathbf{F}_{k,t}$ and $\mathbf{F}_{k,t}^*$ share more and more features as time progresses. We observe further that the similarity increases with increasing k . Since in an online setting $\mathbf{F}_{k,t}^*$ is not available for a very long time, we study the evolution of $sim(\mathbf{F}_{k,t}, \mathbf{F}_{k,2t})$ instead. The results are shown in the second row of Figure 1. We observe the same qualitative behavior of the curves as in the top row: the prediction error tends to decrease with increasing values of k and t .

The bottom row of the figure gives the prediction error of a random-forest regressor that has been trained using the feature set $\mathbf{F}_{k,t}$; it uses 100 trees. The prediction error is expressed as Normalized Mean Absolute Error (NMAE), which is computed as follows:

$$NMAE = \frac{1}{\bar{y}} \left(\frac{1}{q} \sum_{i=1}^q |y_i - \hat{y}_i| \right) \quad (3)$$

\bar{y} is the mean value of targets in test set, q is the number of samples and \hat{y}_i is the predicted value.

From the evaluations presented in Figure 2 and additional ones we performed on other traces collected from our testbed, we decide on the following heuristic notion of a *stable feature set*:

$$\mathbf{F}_{k,t} \text{ is stable, if } sim(\mathbf{F}_{k,t}, \mathbf{F}_{k,2t}) > 0.5$$

which means that $\mathbf{F}_{k,t}$ must share at least half of the features with $\mathbf{F}_{k,2t}$, which is computed with double the number of samples.




Fig. 2. Time for computing $\mathbf{F}_{k,t}$ for all possible k on a server using ARR, LS or TB. Data trace is KV flash-crowd.

Finally, some numbers on the computing time required to obtain the features sets. Figure 2 gives the time to execute the feature ranking algorithms ARR, LS, and TB on a compute server of our testbed. t is the number of samples the algorithms take as input. The time corresponds to computing the feature sets $\mathbf{F}_{k,t}$ for all possible $k = 1 \dots n$. To give a specific example, the execution time for processing 256 samples is 32 ms for ARR, 311 ms for LS, and 8532 ms for TB. The measurements suggest that computing $\mathbf{F}_{k,t}$ with ARR or LS in real time is feasible, also on smaller machines, e.g. on edge nodes.

V. ONLINE FEATURE SELECTION WITH LOW OVERHEAD

In this section, we introduce the Online Stable Feature Set algorithm (OSFS), which reads a stream of samples and returns the number of features k , the number of samples t_k needed to determine k , and the feature set \mathbf{F}_k . OSFS is instantiated with a feature ranking algorithm of choice. We evaluate the effectiveness of OSFS for several feature ranking algorithms using data traces from our testbed.

Recall that the purpose of OSFS is to select a subset of available data sources (i.e., features), in order to reduce the monitoring costs and the training overhead for learning. To keep costs and overhead low, we want k and t_k to be small while still allowing for effective learning and prediction.

The standard use case for OSFS plays out as follows. We start monitoring the values of n features at time $t = 1$ and collect a sequence of samples with index $t = 1, 2, \dots$. We use the collected samples to find small values for k ($k \ll n$) and t_k so that the first t_k samples allow us to compute a stable feature set \mathbf{F}_{t_k} . At time $t_k + 1$, we reduce the number of data sources from n to k and continue motoring only sources from the set \mathbf{F}_k . Once we have collected l samples, we take them as input to train the model for a learning task. The model is then applied for prediction using the samples collected after $t=l$. (In this work, we do not automatically determine the value for l but set it to 1024 based on experience.)

Algorithm 3 shows the pseudo-code of OSFS. The algorithm takes as input the sequence of arriving samples X_t and is initialized with a feature ranking algorithm, for instance *ARR*, *LS*, or *TB*. The ranking algorithm is used in the function *subset()*, which takes as input k and the first t samples and

Algorithm 3: Online Stable Feature Set (OSFS)

Input: Sample sequence X_1, X_2, X_3, \dots ; feature ranking algorithm $Rank$

Output: Feature subset F_k, k, t_k
 $subset(k, t, Rank)$ returns top k features computed with $Rank$ using samples with index $1, \dots, t$.

```
1  $\eta = 0.5$  (threshold for stable feature set);
2  $read = false$ ;
3 for  $k$  in  $[4, 16, 64, 256]$  do
4   if not read then
5      $\lfloor$  read and store  $X_t, t = 1, \dots, 16$ 
6      $F_{k1} = subset(k, 8, Rank)$ ;
7      $F_{k2} = subset(k, 16, Rank)$ ;
8      $sim_{k12} = sim(F_{k1}, F_{k2})$ ;
9     for  $t = 17, \dots, 1024$  do
10      if not read then
11         $\lfloor$  read and store  $X_t$ 
12      if  $t$  in  $[32, 64, 128, 256, 512, 1024]$  then
13         $F_{kt} = subset(k, t, Rank)$ ;
14         $sim_{kt} = sim(F_{k2}, F_{kt})$ ;
15        if  $sim_{kt} < sim_{k12}$  and  $sim_{k12} > \eta$  then
16           $\lfloor$  return  $F_{k1}, k, t/4$ ;
17        else if  $sim_{kt} > \eta$  and  $t == 1024$  then
18           $\lfloor$  return  $F_{k2}, k, t/2$ ;
19        else
20           $F_{k1} = F_{k2}$ ;
21           $F_{k2} = F_{kt}$ ;
22           $sim_{k12} = sim_{kt}$ ;
23       $\lfloor$   $read = true$ 
24 return  $F_{k2}, 256, 1024$ 
```

returns a set with the top k features as ranked by the algorithm. OSFS has two main loops: an outer loop (lines 3-23) that iterates over a subspace of k and an inner loop (lines 9-22) that iterates over a subspace of t . The index values of k and t increase exponentially to enable exploration of a large space with a small number of evaluations. The algorithm performs a grid search in the space of tuples (k, t) with the termination condition A: $sim_{kt} < sim_{k12}$ and $sim_{k12} > \eta$ (line 15) or B: $sim_{kt} > \eta$ and $t == 1024$ (line 17). (η is the threshold value for a stable feature set.) In case the conditions A or B are never met, the algorithm terminates after the search on the grid sector $[4 - 256] * [8 - 1024]$ has been completed. The key termination condition A expresses the case where (1) the similarity of two consecutive feature sets is above the threshold η and (2) the similarity declines when the subsequent feature set is considered.

OSFS reads at least 32 samples in order to ensure statistical viability and terminates after at most 1024 samples. Since the outer loop is indexed by increasing k , OSFS favors a smaller k at the possible expense of a larger t . This means that we prefer

reducing the monitoring overhead over reducing the time to compute F_k .

We evaluate OSFS for the feature ranking algorithms ARR, LS, and TB using the four data traces described in [3].

Table 2 summarizes the main results. The columns k and t_k give the figures for the feature selection algorithms we consider in this paper. The values for k and t_k show the means and standard deviations from 10 start times for OSFS, whereby one start time is $t=1$ and the other 9 are chosen uniformly at random between $t=2$ and $t=10\,000$.

The columns NMAE1 and NMAE2 indicate the effectiveness of the feature sets produced by OSFS. They measure the error of a random-forest predictor using the produced feature sets. The values for NMAE1 relate to the case where the predictor models are trained using the 1024 samples after the respective start times, and the error is computed using all samples of the trace following those used for model training. In contrast, the values for NMAE2 relate to the case where the predictor models are trained in an offline fashion; i.e. 70% of the samples of the trace is used for training and the rest for evaluation. “No FS” refers to the case where all features are considered for training. Therefore, the column NMAE2 shows the effectiveness of the feature sets produced by OSFS compared to a baseline that considers all features. The column NMAE1 gives insights about how OSFS performs in conjunction with different feature selection algorithms for a specific learning task in a realistic online setting.

The general conclusions we draw from Table 2 are as follows. Regarding the number of features k and samples t_k :

- OSFS instantiated with ARR, LS, or TB achieves a massive reduction of features of 1-2 orders of magnitude. Most prominent is ARR, which produces feature sets with 10-20 features, down from 1255-1723 features.
- The number of samples needed to compute the feature sets averages around 100-400. For our tesbed, where metrics are monitored once per second, this means that the computed feature sets are available after 2-6 minutes of monitoring.

Regarding the accuracy of a specific predictor using the computed feature set:

- Online OSFS with ARR, LS, or TB incurs a 50%-100% larger error than an offline-trained predictor with access to the entire feature set and the entire data trace. This is primarily due to the online computation of the feature sets. Further, comparing the columns NMAE1 and NMAE2 provides the cost of online prediction vs. offline prediction.
- Consistent with our earlier results (e.g., [12], [13]), we find that the type of service and the load pattern significantly affect the prediction error.

When focusing on unsupervised feature selection, the comparison of ARR versus LS shows that ARR and LS exhibit similar performance regarding online feature selection and prediction (NMAE1). Since ARR has a significantly lower computing overhead (Table 2, Figure 2), we conclude that

TABLE 2

Evaluation of OSFS for feature selection algorithms ARR, LS, and TB for a the task of service metrics prediction on four data sets. The prediction method is random forest. No FS refers to the base line where all features are used by the predictor. NMAE1 is prediction accuracy where the prediction model is built with the first 1024 samples. NMAE2 is prediction accuracy where the prediction model is built through offline analysis with knowledge of the complete data set.

Dataset	Method	k	t_k	NMAE1	NMAE2
KV flash-crowd	ARR	25.6 ± 25.6	251.2 ± 224.0	0.0494 ± 0.0010	0.0280 ± 0.0056
	LS	10 ± 6	64 ± 72.3	0.0691 ± 0.0592	0.0232 ± 0.0014
	TB	97.6 ± 106.6	332.8 ± 153.6	0.0221 ± 0.0012	0.0191 ± 0.0008
	No FS	1723			0.0184
KV periodic	ARR	10 ± 6	249.6 ± 194.4	0.0573 ± 0.0114	0.0325 ± 0.0071
	LS	12.4 ± 5.5	195.2 ± 177.3	0.0711 ± 0.0473	0.0268 ± 0.0016
	TB	156.4 ± 122.0	409.6 ± 159.9	0.0336 ± 0.0052	0.0232 ± 0.0028
	No FS	1751			0.0214
VoD flash-crowd	ARR	12.4 ± 5.5	131.2 ± 142.7	0.1772 ± 0.0245	0.1629 ± 0.0139
	LS	17.2 ± 16.5	226.4 ± 198.2	0.1631 ± 0.0354	0.0969 ± 0.0405
	TB	54.4 ± 72.5	142.4 ± 198.2	0.1454 ± 0.0179	0.0999 ± 0.0306
	No FS	1255			0.0771
VoD periodic	ARR	10 ± 6	363.2 ± 196.5	0.2172 ± 0.0473	0.2085 ± 0.0203
	LS	25.6 ± 25.6	108 ± 85.5	0.2145 ± 0.0527	0.1162 ± 0.0295
	TB	47.2 ± 74.6	216 ± 242.2	0.2214 ± 0.0430	0.1473 ± 0.0429
	No FS	1296			0.119

ARR is the preferred algorithm in a resource-constrained environment (e.g. edge node or sensor node).

When feature selection is performed for a single predictor and target information is available, then TB should be considered for online feature selection. TB provides better prediction accuracy than ARR or LS, although at much higher computational costs (Table 2, Figure 2) and sometimes larger k .

From analysis and experimentation we know that increasing t_k , and to a lesser extent k , will improve the prediction accuracy, and OSFS can be extended for that purpose in a straightforward way. The particular technological use case determines the range of acceptable errors. For the services on our testbed, we believe we achieved monitoring and rapid training with low overhead at the cost of acceptable prediction errors.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we introduced an online algorithm, OSFS, that selects a small feature set from a large number of available data sources, which allows for rapid, low-overhead learning and prediction with acceptable errors. OSFS is instantiated with a feature ranking algorithm and performs a grid search that terminates when a stable feature set is identified and succeeding feature sets exhibit reduced similarity.

Regarding future work, there are many ways our method can be improved and extended. In addition to the three ranking algorithms considered in this work, many other candidates that can be evaluated, together with OSFS, with respect to their ability to produce small values for k and t_k and predictors with acceptable errors. Of particular interest are online ranking algorithms that require only a single processing pass over the samples and whose memory requirement does not increase with t . Also, the tradeoffs between the three metrics t_k , t_k , and prediction error warrant further study. Finally, a formal analysis of the performance of OSFS with a specific

ranking algorithm will be difficult due to dependencies on the particular data trace, but it should be attempted in order to fundamentally understand the conditions for efficient and effective online learning in networked systems.

VII. ACKNOWLEDGEMENTS

The authors are grateful to Andreas Johnsson, Hannes Larsson, and Jalil Taghia with Ericsson Research for fruitful discussion around this work, as well as to Kim Hammar and Rodolfo Villaça for comments on an earlier version of this paper. This research has been partially supported by the Swedish Governmental Agency for Innovation Systems, VINNOVA, through project AutoDC.

REFERENCES

- [1] F. S. Samani, H. Zhang, and R. Stadler, "Efficient learning on high-dimensional operational data," in *2019 15th International Conference on Network and Service Management (CNSM)*. IEEE, 2019, pp. 1–9.
- [2] X. Wang, "Dimensionality reduction for performance prediction in networked systems," master thesis, KTH Royal Institute of Technology, 2020.
- [3] X. Wang, F. Shahab Samani, and R. Stadler, "Online feature selection for rapid, low-overhead learning in networked systems," *arXiv preprint*, 2020.
- [4] A. J. Ferreira and M. A. Figueiredo, "An unsupervised approach to feature discretization and selection," *Pattern Recognition*, vol. 45, no. 9, pp. 3048–3060, 2012.
- [5] X. He, D. Cai, and P. Niyogi, "Laplacian score for feature selection," in *Advances in neural information processing systems*, 2006, pp. 507–514.
- [6] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
- [7] H. Deng and G. Runger, "Feature selection via regularized trees," in *The 2012 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2012, pp. 1–8.
- [8] scikit-learn developers, "Tree-based feature selection," 2007–2019. [Online]. Available: https://scikit-learn.org/stable/modules/feature_selection.html
- [9] —, "Feature importances with forests of trees," 2007–2019. [Online]. Available: https://scikit-learn.org/stable/auto_examples/ensemble/plot_forest_importances.html#sphx-glr-auto-examples-ensemble-plot-forest-importances-py
- [10] —, "Decision trees," 2007–2019. [Online]. Available: <http://mldata.org/repository/data/viewslug/realm-cnsm2015-vod-traces/>

- [11] V. N. Vapnik, *Statistical learning theory*. John Wiley, 1998.
- [12] R. Yanggratoke, J. Ahmed, J. Ardelius, C. Flinta, A. Johnsson, D. Gillblad, and R. Stadler, "Predicting Real-time Service-level Metrics from Device Statistics," in *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*. IEEE, 2015, pp. 414–422.
- [13] R. Stadler, R. Pasquini, and V. Fodor, "Learning from network device statistics," *Journal of Network and Systems Management*, vol. 25, no. 4, pp. 672–698, 2017.