

Evaluating Short-Term Forecasting of Multiple Time Series in IoT Environments

Christos Tzagkarakis¹, Pavlos Charalampidis¹, Stylianos Roubakis¹,
Alexandros Fragkiadakis¹ and Sotiris Ioannidis^{2,1}

¹*Foundation for Research and Technology-Hellas, Institute of Computer Science, Heraklion, Greece*

²*Technical University of Crete, Chania, Greece*

E-mails: {tzagarak, pcharala, roub, alfrag}@ics.forth.gr, sotiris@ece.tuc.gr

Abstract—Modern Internet of Things (IoT) environments are monitored via a large number of IoT enabled sensing devices, with the data acquisition and processing infrastructure setting restrictions in terms of computational power and energy resources. To alleviate this issue, sensors are often configured to operate at relatively low sampling frequencies, yielding a reduced set of observations. Nevertheless, this can hamper dramatically subsequent decision-making, such as forecasting. To address this problem, in this work we evaluate short-term forecasting in highly underdetermined cases, i.e., the number of sensor streams is much higher than the number of observations. Several statistical, machine learning and neural network-based models are thoroughly examined with respect to the resulting forecasting accuracy on five different real-world datasets. The focus is given on a unified experimental protocol especially designed for short-term prediction of multiple time series at the IoT edge. The proposed framework can be considered as an important step towards establishing a solid forecasting strategy in resource constrained IoT applications.

Index Terms—Internet of Things, multiple time series, short-term forecasting, rolling window tuning, machine learning, neural networks

I. INTRODUCTION

With the advent of the Internet of Things (IoT) era, the concept of Internet networking has been shifted into a ubiquitous interconnection between users, data and smart devices or sensors in a seamless way [1], [2]. A wide range of applications has emerged at the forefront of IoT technology spanning from smart cities [3], [4], smart homes [5], [6] and wearables [7], [8] to energy management [9], predictive maintenance [10], automotive driving [11], etc. The modern IoT systems are typically decomposed into three levels, i.e., highly heterogeneous data are captured from the surrounding environment via several IoT devices/sensors at the edge level, which then can be transmitted through the fog layer up to the cloud for further storage and processing. However, the rapidly growing use and realization of IoT technology comes at the cost of resolving major technical and business impediments as reflected in dynamicity, scalability and heterogeneity [12], [13].

Specifically, a dynamically adaptive behaviour is followed at the IoT infrastructure, at the IoT applications and at the IoT devices, and thus it is imperative to promote a (semi)-automatic behaviour within all IoT layers. This gives rise to the pursuit of high scalability properties from the network layers

as well as from the IoT infrastructure. In addition, enhanced heterogeneous behaviour as a result of the extensive use and interconnection of a large volume of diverse IoT devices should be addressed through the concept of efficient semantic interoperability within IoT applications and platforms. These IoT inter-layer characteristics could promote a direct data analysis at the IoT edge as a very important factor within IoT systems.

Typically, a large number of sensing devices is used to monitor IoT environments. The data acquisition and processing modules impose restrictions in terms of computational power and energy resources at the edge level. These constraints can be handled by configuring the sensors to operate at relatively low sampling frequencies, yielding a reduced set of observations, which can be further analyzed by adopting lightweight algorithmic procedures. This motivates us to investigate the short-term forecasting of multiple sensor streams at the IoT edge. Specifically, we assume that a set of sensing devices is used to collect time series data from the IoT environment, where the number of time series is much higher than the number of observations per sensor. Instead of predicting the future values of each sensor stream separately, each sensor's collected time series can be transmitted from the edge to an aggregation mechanism such as an edge gateway. As such, an overall (global) forecasting model can be computed based on the aggregated multiple time series, considering in this way the inherent relationships between the acquired time series.

Related work. Recent works have dealt with the multiple time series forecasting task by focusing on IoT-related data. A block Hankel tensor-based autoregressive integrated moving average (BHT-ARIMA) method is described in [14] that exploits the intrinsic correlations among multiple time series. The authors in [15] introduce a long short-term memory multiseasonal network (LSTM-MSNet), i.e., a three-layered forecasting framework using LSTMs that accounts for the multiple seasonal periods present in time series. Short-term electricity load forecasting is examined in [16], where random forest (RF) models are used to forecast time series with multiple seasonal variations. In [17] a comprehensive time series forecasting benchmarking archive is proposed that contains twenty-five publicly available time series datasets from various domains, with different characteristics in terms of series lengths, time resolution, and inclusion of missing values.

Contribution. The contribution of the current paper is twofold. Firstly, we establish a solid multiple time series forecasting protocol for single-step prediction, especially targeted in IoT use cases. Secondly, a set of five different real-world IoT datasets is thoroughly examined by using off-the-shelf statistical, machine learning and neural network methods as well as scale-free and percentage error-based accuracy metrics.

The rest of the paper is organized as follows: Section II overviews the problem formulation. The proposed multiple time series prediction protocol is described in Section III, while Section IV demonstrates the experimental evaluation in light of the forecasting accuracy on five real-world datasets. Finally, Section V summarizes the main results and gives directions for further extensions.

II. PROBLEM FORMULATION

Let us assume that a collection of N equal length time series can be written in matrix form

$$\mathbf{Y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_l \\ \vdots \\ \mathbf{y}_L \end{bmatrix} = \begin{bmatrix} y_1^1 & y_1^2 & \cdots & y_1^N \\ \vdots & \vdots & & \vdots \\ y_l^1 & y_l^2 & \cdots & y_l^N \\ \vdots & \vdots & & \vdots \\ y_L^1 & y_L^2 & \cdots & y_L^N \end{bmatrix} \in \mathbb{R}^{L \times N}, \quad (1)$$

where each row $\mathbf{y}_l = [y_l^1 \dots y_l^N] \in \mathbb{R}^{1 \times N}$, $l \in \{1, \dots, L\}$, represents the l -th observation of all time series and each column $\mathbf{y}^n = [y_1^n \dots y_L^n]^T \in \mathbb{R}^{N \times 1}$, $n \in \{1, \dots, N\}$, denotes the time series captured by the n -th sensor. Our main goal is to compute a forecasting model \mathcal{M} , which uses historical observations of the multiple time series \mathbf{Y} to predict the future values

$$\mathbf{Y}^{(f)} = \begin{bmatrix} \mathbf{y}_{L+1} \\ \vdots \\ \mathbf{y}_{L+h} \end{bmatrix} = \begin{bmatrix} y_{L+1}^1 & y_{L+1}^2 & \cdots & y_{L+1}^N \\ \vdots & \vdots & & \vdots \\ y_{L+h}^1 & y_{L+h}^2 & \cdots & y_{L+h}^N \end{bmatrix} \in \mathbb{R}^{h \times N} \quad (2)$$

for a given forecasting horizon h . In this work, we assume a single-step forecasting horizon $h = 1$, and thus the matrix (2) collapses into a vector

$$\mathbf{y}^{(f)} = [y_{L+1}^1 y_{L+1}^2 \cdots y_{L+1}^N] \in \mathbb{R}^{1 \times N}. \quad (3)$$

A prediction model \mathcal{M} can then be defined as

$$\hat{\mathbf{y}}^{(f)} = \mathcal{M}(\mathbf{Y}, \boldsymbol{\theta}), \quad (4)$$

where

$$\hat{\mathbf{y}}^{(f)} = [\hat{y}_{L+1}^1 \hat{y}_{L+1}^2 \cdots \hat{y}_{L+1}^N] \in \mathbb{R}^{1 \times N} \quad (5)$$

denotes the predicted samples within the single-step forecasting horizon, and $\boldsymbol{\theta}$ are the model parameters. It is obvious that the smaller the difference between the actual $\mathbf{y}^{(f)}$ and the predicted $\hat{\mathbf{y}}^{(f)}$ samples, the more accurate the future samples prediction is.

A. Forecasting Models

The forecasting model \mathcal{M} can be estimated through various approaches grouped into three main categories, namely the statistical, the machine learning and the neural network methods. BHT-ARIMA [14] is the statistical model used in the current work, while support vector regression (SVR) [18] and RF model [16] correspond to the adopted machine learning approaches. The used neural network-based methods correspond to long short-term memory (LSTM), bidirectional LSTM (Bi-LSTM) and convolutional neural networks (CNN) [19], [20] as well as to echo state network (ESN) [21]. Each forecasting model's input data should be transformed into the appropriate shape for training, hyper-parameter tuning and testing purposes. Next, more details on the data shaping process and the overall forecasting protocol are provided.

III. FORECASTING PROTOCOL

A specific sequence of steps must be followed to estimate the future samples $\mathbf{y}^{(f)}$ through the forecasting model \mathcal{M} . As a first step, it is important to perform data scaling, since in real-world IoT applications, the collected time series contain observations with different value ranges. The min-max scaler

$$\mathbf{y}_{sc}^n = \frac{\mathbf{y}^n - \mathbf{y}_{min}^n}{\mathbf{y}_{max}^n - \mathbf{y}_{min}^n}, \quad n = 1, \dots, N \quad (6)$$

is used for data normalization, with \mathbf{y}_{sc}^n representing the scaled observations of the n -th time series, while \mathbf{y}_{min}^n and \mathbf{y}_{max}^n are the minimum and maximum values of the n -th time series, respectively. The scaled version of the multiple time series (1) can be written as $\mathbf{Y}_{sc} \in \mathbb{R}^{L \times N}$.

The second step involves data training/validation of each forecasting model mentioned in Section II-A. To properly train and validate each forecasting model based on the available data, we need to follow a rolling window-based cross-validation procedure towards identifying the best hyper-parameter configuration per model. In the following subsections, we describe the different types of cross-validation methods per forecasting model category. Before proceeding, let us define the notations: $^{[i]} \mathbf{A}$ denotes the i -th row of matrix \mathbf{A} , $^{[i:j]} \mathbf{A}$ is the submatrix of \mathbf{A} from row i to row j and $\hat{\mathbf{Y}}_{sc}$ is the predicted counterpart of \mathbf{Y}_{sc} .

A. Rolling Window: Matrix Case

In the case of SVR and RF machine learning models, it is necessary to perform hyper-parameter tuning by using the pair of training data submatrices $(^{[1:L_{tr}-1]} \mathbf{Y}_{sc}, ^{[L_{tr}]} \mathbf{Y}_{sc})$, where $L_{tr} < L$ defines the amount of training data in each cross-validation iteration. In particular, each machine learning model is trained on this pair of data submatrices, where the trained (for a specific hyper-parameter configuration) model's forecasting performance is then evaluated on the first validation fold $^{[L_{tr}+1]} \mathbf{Y}_{sc}$, obtaining the validation error of the form

$$e_v^1 = ^{[L_{tr}+1]} \mathbf{Y}_{sc} - ^{[L_{tr}+1]} \hat{\mathbf{Y}}_{sc}. \quad (7)$$

Next, the second training data pair $([2:L_{tr}]\mathbf{Y}_{sc}, [L_{tr}+1]\mathbf{Y}_{sc})$ is composed by sliding the training window one step forward, yielding the second validation error

$$e_v^2 = [L_{tr}+2]\mathbf{Y}_{sc} - [L_{tr}+2]\hat{\mathbf{Y}}_{sc} \quad (8)$$

for the same hyper-parameter configuration. We keep moving the window forward one step at a time until the computation of the last validation error

$$e_v^{L_v} = [L_{tr}+L_v]\mathbf{Y}_{sc} - [L_{tr}+L_v]\hat{\mathbf{Y}}_{sc} \quad (9)$$

based on the trained model fed with the data $([L_{tr}-1:L_{tr}+L_v-2]\mathbf{Y}_{sc}, [L_{tr}+L_v-1]\mathbf{Y}_{sc})$, where $L_v = L - L_{tr}$ denotes the amount of validation data. The best hyper-parameter configuration corresponds to the minimum mean validation error

$$\bar{e}_v = \frac{1}{L_v} \sum_{k=1}^{L_v} e_v^k. \quad (10)$$

For BHT-ARIMA, we follow a similar hyper-parameter tuning concept by providing $[1:L_{tr}]\mathbf{Y}_{sc}$ as training data input to the model during the first cross-validation iteration, and computing the first validation error (7). The training window is then shifted forward by a single step to obtain the second training data input $[2:L_{tr}+1]\mathbf{Y}_{sc}$, and thus providing the second validation error (8). The rolling window process is repeated until the last training segment $[L_{tr}-1:L_{tr}+L_v-1]\mathbf{Y}_{sc}$, with the last validation error computed as in (9). The optimal hyper-parameter combination corresponds to the minimum mean validation error (10).

B. Rolling Window: Multi-dimensional Matrix Case

In the case of LSTM, Bi-LSTM and CNN neural network models, a slightly different rolling window-based cross-validation methodology is followed. Specifically, an additional hyper-parameter $S < L_{tr}$ is introduced, which determines the rolling window length. Next, a pair of data $([1:S]\mathbf{Y}_{sc}, [S+1]\mathbf{Y}_{sc})$ is formed. The rolling window of length S is then shifted forward by a single step composing the second pair $([2:S+1]\mathbf{Y}_{sc}, [S+2]\mathbf{Y}_{sc})$ of data. This rolling window scheme is repeated until the assembly of the last pair $([L_{tr}-S:L_{tr}-1]\mathbf{Y}_{sc}, [L_{tr}]\mathbf{Y}_{sc})$. To properly feed the LSTM, Bi-LSTM and CNN model with the training input data, we need to concatenate these pairs as follows

$$\left(\left[[1:S]\mathbf{Y}_{sc}, [2:S+1]\mathbf{Y}_{sc}, \dots, [L_{tr}-S:L_{tr}-1]\mathbf{Y}_{sc} \right], \right. \\ \left. \left[[S+1]\mathbf{Y}_{sc}, [S+2]\mathbf{Y}_{sc}, \dots, [L_{tr}]\mathbf{Y}_{sc} \right] \right). \quad (11)$$

The trained neural network model (for a particular hyper-parameter combination) is fed with $[L_{tr}-S+1:L_{tr}]\mathbf{Y}_{sc}$ to compute the first validation error (7). The second validation error (8) is estimated based on the data $[L_{tr}-S+2:L_{tr}+1]\mathbf{Y}_{sc}$, while the last validation error (9) is obtained given the data $[L_{tr}-S+L_v:L_{tr}+L_v-1]\mathbf{Y}_{sc}$, and the optimal hyper-parameter configuration corresponds to the minimum mean validation error (10).

C. Rolling Window: Matrix List Case

Here, we overview the rolling window-based cross-validation methodology in the case of ESN model. In particular, let us introduce an extra hyper-parameter $S < L_{tr}$ as in the previous subsection. The following training pairs

$$\left(\left([1:S]\mathbf{Y}_{sc}, [2:S+1]\mathbf{Y}_{sc} \right), \left([2:S+1]\mathbf{Y}_{sc}, [3:S+2]\mathbf{Y}_{sc} \right), \dots, \right. \\ \left. \left([L_{tr}-S:L_{tr}-1]\mathbf{Y}_{sc}, [L_{tr}-S+1:L_{tr}]\mathbf{Y}_{sc} \right) \right) \quad (12)$$

are then composed and used as input to train the ESN model (given a specified hyper-parameter combination). After the ESN training process, the pair of data $([L_{tr}-S:L_{tr}-1]\mathbf{Y}_{sc}, [L_{tr}-S+1:L_{tr}]\mathbf{Y}_{sc})$ is fed as input to the trained ESN model in order to compute the first validation error (7). The second validation error (8) is estimated based on the pair of data $([L_{tr}-S+1:L_{tr}]\mathbf{Y}_{sc}, [L_{tr}-S+2:L_{tr}+1]\mathbf{Y}_{sc})$. This rolling window process is repeated until the last training data input $([L_{tr}+L_v-S-1:L_{tr}+L_v-2]\mathbf{Y}_{sc}, [L_{tr}+L_v-S:L_{tr}+L_v-1]\mathbf{Y}_{sc})$ provides the last validation error (9). As in the two rolling window schemes above, the minimum mean validation error (10) provides the optimal hyper-parameter configuration.

IV. EXPERIMENTAL EVALUATION

In this section, we evaluate the short-term multiple time series forecasting models on real-world time series datasets based on the protocol described in Section III.

A. Datasets Description

We use five publicly available multiple time series datasets that correspond to various practical IoT applications. Specifically, the *Energy Consumption Fraunhofer* dataset¹ contains the measured energy consumption, with an hourly time resolution, of 499 customers (each customer is assigned to one of the 68 customer profiles such as private households, shops, bakeries) in Spain. The energy consumption is measured in kilowatt hour (kWh). Each time series contains the consumption values within the measurement period of Jan. 1 to Dec. 31, 2019. In the current experimental framework, the mean daily energy consumption of each customer's household is used, and thus the final dataset consists of 314 time series (customers) and 365 (daily) measurements per time series.

The *Electricity Load Diagrams* dataset² corresponds to electricity measurements across 370 households in Portugal from 2011 to 2014, with a 15 minutes time resolution. Here, we use the total amount of energy consumption per customer in a daily basis during the period Jan. 1, 2012 to Dec. 31, 2014, providing an overall dataset of size 324 time series by 1096 measurements.

The *Guangzhou Traffic* dataset³ is an urban traffic speed dataset collected in Guangzhou, China, which consists of 214

¹<https://fordatis.fraunhofer.de/handle/fordatis/215>

²<https://archive.ics.uci.edu/ml/datasets/ElectricityLoadDiagrams20112014>

³<https://zenodo.org/record/1205229/#.YguzWN9BxaR>

TABLE I
DATASETS INFORMATION

Dataset	Number of time series	Number of samples per time series	Time resolution
Energy Consumption Fraunhofer (D1)	314	365	daily
Electricity Load Diagrams (D2)	320	1096	daily
Guangzhou Traffic (D3)	206	1464	hourly
San Francisco Traffic (D4)	862	104	weekly
London Smart Meters (D5)	504	9983	half hourly

anonymous road segments within the period Aug. 1 to Sep. 30, 2016 at 10 minutes interval. Since the original dataset contains missing values, the road segments containing non-zero values are only kept. During the experimental evaluation, we use the mean hourly traffic speed per road segment, and thus the final dataset's size is 206 time series (road segments) by 1464 measurements.

The *San Francisco Traffic* dataset⁴ is a collection of hourly time series, representing the traffic occupancy rate of different car lanes of San Francisco bay area freeways between 2015 and 2016. The size of the used dataset is 862 time series by 104 measurements.

The *London Smart Meters* dataset⁵ consists of 5560 half hourly time series that represent the energy consumption readings of London households in kWh from Nov. 2011 to Feb. 2014. Here, we select 504 households with energy consumption during Oct. 23, 2012 until May 18, 2013 leading to 9983 measurements per time series. The overall datasets information is presented in Table I.

B. Performance Metrics

We employ percentage error-based and scale-free error metrics to assess the accuracy of the forecasting models. Three error metrics are used in the context of multiple time series short-term forecasting within a prediction horizon $h = 1$, namely the symmetric mean absolute percentage error (sMAPE), the recently proposed mean arctangent absolute percentage error (MAAPE) [22] and the mean absolute scaled error (MASE) given by

$$\text{sMAPE}(\mathbf{y}^{(f)}, \hat{\mathbf{y}}^{(f)}) = \frac{1}{N} \sum_{n=1}^N 2 \frac{|y_{L+1}^n - \hat{y}_{L+1}^n|}{|y_{L+1}^n| + |\hat{y}_{L+1}^n|} \quad (13)$$

$$\text{MAAPE}(\mathbf{y}^{(f)}, \hat{\mathbf{y}}^{(f)}) = \frac{1}{N} \sum_{n=1}^N \arctan \left(\left| \frac{y_{L+1}^n - \hat{y}_{L+1}^n}{y_{L+1}^n} \right| \right) \quad (14)$$

$$\text{MASE}(\mathbf{y}^{(f)}, \hat{\mathbf{y}}^{(f)}) = \frac{1}{N} \sum_{n=1}^N \frac{|y_{L=1}^n - \hat{y}_{L+1}^n|}{\frac{1}{L-1} \sum_{l=2}^L |y_l^n - \hat{y}_{L+1}^n|}, \quad (15)$$

where $\mathbf{y}^{(f)}$ and $\hat{\mathbf{y}}^{(f)}$ is the original and predicted samples vector, as defined in (3) and (5), respectively. The sMAPE is used as a metric during the cross-validation process, while

sMAPE, MAAPE and MASE are used to evaluate the forecasting generalization (out-of-sample) performance of each prediction model.

C. Hyper-parameter Tuning

Based on the forecasting protocol mentioned in Section III, each dataset is split into a training (80%) and a validation (20%) partition. For each prediction model, a different rolling window-based cross-validation approach is followed to identify the best hyper-parameter combination, as discussed in Section III-A- III-C. Hyper-parameter grid search is performed for each model per dataset. Due to lack of space, we provide a detailed description of the hyper-parameter grids per model in <https://github.com/pcharala/multiple-timeseries-forecasting>.

D. Results

In this section, we compare the forecasting performance of each prediction model \mathcal{M} with the performance obtained by predicting the future single-step $\mathbf{y}^{(f)}$ based on the scaled multiple time series \mathbf{Y}_{sc} . For this purpose, we perform simulations on five real-world datasets described in Section IV-A. We investigate the forecasting accuracy of each prediction model for a varying length L of \mathbf{Y}_{sc} , i.e., the accuracy for a total length $L = 40$ and $L = 90$ is examined for all datasets. The single-step forecasting results are averaged over fifteen Monte Carlo iterations, where during each Monte Carlo run, a different part of the time series dataset is examined. The average error metric per dataset is reported to show the overall mean errors for each prediction model.

Table II shows the forecasting results measured in terms of sMAPE, MAAPE and MASE. The datasets names D1 (Energy Consumption Fraunhofer), D2 (Electricity Load Diagrams), D3 (Guangzhou Traffic), D4 (San Francisco Traffic), D5 (London Smart Meters) and the respective experimental window configurations are depicted in the first two rows of the table. The best sMAPE, MAAPE and MASE value is depicted in blue, green and gray colour, respectively. As it can be seen, the statistical (BHT-ARIMA) and machine learning (SVR, RF) models achieve in general a better forecasting performance against the neural network models (LSTM, Bi-LSTM, CNN, ESN). This behaviour has been observed in other works (e.g. in [14]) and can be attributed to the fact that neural networks need more data to achieve a higher performance. Our analysis code and data can be found on GitHub: <https://github.com/pcharala/multiple-timeseries-forecasting>.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we address the problem of short-term forecasting of multiple time series in IoT environments, when the number of time series is much higher than the number of observations. Seven different prediction models are thoroughly examined on five real-world IoT datasets, applying a unified experimental protocol especially designed for short-term prediction of multiple time series at the IoT edge. It is experimentally shown that the statistical and machine learning models are generally outperforming the neural network-based

⁴<https://zenodo.org/record/4656135/#.Ygu-o99BxaQ>

⁵<https://zenodo.org/record/4656091/#.YgvTLd9BxaS>

TABLE II
FORECASTING RESULTS

DATASET	# Sim. data	D1		D2		D3		D4		D5	
		40	90	40	90	40	90	40	90	40	90
BHT-ARIMA	sMAPE	0.3369	0.3554	0.0969	0.0819	0.0797	0.1102	0.0742	0.1022	0.5209	0.4683
	MAAPE	0.2888	0.3035	0.0986	0.0756	0.0765	0.1050	0.0724	0.1025	0.4298	0.3988
	MASE	0.1806	0.1905	0.0503	0.0247	0.3169	0.3569	0.1877	0.2576	0.7366	0.3829
SVR	sMAPE	0.3683	0.3816	0.0931	0.0765	0.0655	0.0930	0.0795	0.0946	0.4705	0.5047
	MAAPE	0.3275	0.3424	0.0933	0.0732	0.0666	0.0935	0.0766	0.0952	0.4383	0.4898
	MASE	0.1867	0.1945	0.0421	0.0250	0.2697	0.3190	0.1978	0.2267	0.3878	0.4271
RF	sMAPE	0.3200	0.3417	0.1026	0.0813	0.0601	0.0921	0.0814	0.1076	0.4649	0.4353
	MAAPE	0.2907	0.3086	0.1033	0.0782	0.0615	0.0923	0.0788	0.1068	0.4461	0.4281
	MASE	0.1892	0.1927	0.0459	0.0248	0.2324	0.2978	0.2037	0.2617	0.3935	0.3673
LSTM	sMAPE	0.5551	0.6310	0.1668	0.2151	0.4762	0.5625	0.1771	0.2652	0.9009	0.8736
	MAAPE	0.3672	0.4038	0.1373	0.1597	0.3281	0.3675	0.1488	0.2204	0.4959	0.4842
	MASE	0.3111	0.3557	0.0849	0.0885	1.8338	2.0560	0.4068	0.5719	0.6320	0.6127
Bi-LSTM	sMAPE	0.5557	0.6299	0.1667	0.2138	0.4761	0.5515	0.1770	0.3246	0.9017	0.8734
	MAAPE	0.3682	0.4048	0.1373	0.1590	0.3280	0.3623	0.1487	0.2577	0.4963	0.4843
	MASE	0.3112	0.3558	0.0849	0.0883	1.8338	2.0253	0.4067	0.7376	0.6321	0.6124
CNN	sMAPE	0.5555	0.6314	0.1668	0.2150	0.4738	0.5619	0.1771	0.2639	0.9010	0.8732
	MAAPE	0.3680	0.4046	0.1374	0.1596	0.3266	0.3672	0.1488	0.2191	0.4958	0.4845
	MASE	0.3113	0.3562	0.0849	0.0883	1.8235	2.0542	0.4069	0.5678	0.6322	0.6124
ESN	sMAPE	0.3510	0.4262	0.1021	0.0749	0.0588	0.0825	0.0825	0.1216	0.5154	0.6125
	MAAPE	0.3084	0.3657	0.1030	0.0719	0.0601	0.0830	0.0801	0.1206	0.4562	0.5251
	MASE	0.1897	0.2240	0.0474	0.0256	0.2322	0.2954	0.2054	0.2714	0.4239	0.5068

methods. As a future work, we intend to conduct extensive simulations using a larger number of IoT datasets and lightweight multiple time-series prediction models as well as investigating the training/validation performance of additional rolling window-based cross-validation techniques.

ACKNOWLEDGMENT

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 957337 (project MARVEL) and the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH-CREATE-INNOVATE (project code: T1EDK-00070).

REFERENCES

- [1] L. D. Xu, W. He, and S. Li, “Internet of things in industries: A survey,” *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, pp. 2233–2243, Nov. 2014.
- [2] J. A. Stankovic, “Research directions for the internet of things,” *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 3–9, Feb. 2014.
- [3] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, “Internet of things for smart cities,” *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 22–32, Feb. 2014.
- [4] Y. Mehmood, F. Ahmad, I. Yaqoob, A. Adnane, M. Imran, and S. Guizani, “Internet-of-things-based smart cities: Recent advances and challenges,” *IEEE Communications Magazine*, vol. 55, no. 9, pp. 16–24, 2017.
- [5] B. L. R. Stojkoska and K. V. Trivodaliev, “A review of internet of things for smart home: Challenges and solutions,” *Journal of Cleaner Production*, vol. 140, pp. 1454–1464, 2017.
- [6] S. D. T. Kelly, N. K. Suryadevara, and S. C. Mukhopadhyay, “Towards the implementation of IoT for environmental condition monitoring in homes,” *IEEE Sensors Journal*, vol. 13, no. 10, pp. 3846–3853, Oct. 2013.
- [7] J. Wei, “How wearables intersect with the cloud and the internet of things : Considerations for the developers of wearables,” *IEEE Consumer Electronics Magazine*, vol. 3, no. 3, pp. 53–56, July 2014.
- [8] S. Amendola, R. Lodato, S. Manzari, C. Occhiuzzi, and G. Marrocco, “RFID technology for IoT-based personal healthcare in smart spaces,” *IEEE Internet of Things Journal*, vol. 1, no. 2, pp. 144–152, April 2014.

- [9] F. Shrouf and G. Miragliotta, “Energy management based on internet of things: practices and framework for adoption in production management,” *Journal of Cleaner Production*, vol. 100, pp. 235–246, 2015.
- [10] D. Kwon, M. R. Hodkiewicz, J. Fan, T. Shibutani, and M. G. Pecht, “IoT-based prognostics and systems health management for industrial applications,” *IEEE Access*, vol. 4, pp. 3659–3670, 2016.
- [11] W. He, G. Yan, and L. D. Xu, “Developing vehicular data cloud services in the IoT environment,” *IEEE Transactions on Industrial Informatics*, vol. 10, no. 2, pp. 1587–1595, May 2014.
- [12] M. A. Razzaque, M. Milojevic-Jevric, A. Palade, and S. Clarke, “Middleware for internet of things: A survey,” *IEEE Internet of Things Journal*, vol. 3, no. 1, pp. 70–95, Feb. 2016.
- [13] I. Lee and K. Lee, “The internet of things (IoT): Applications, investments, and challenges for enterprises,” *Business Horizons*, vol. 58, no. 4, pp. 431–440, 2015.
- [14] Q. Shi, J. Yin, J. Cai, A. Cichocki, T. Yokota, L. Chen, M. Yuan, and J. Zeng, “Block Hankel tensor ARIMA for multiple short time series forecasting,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, pp. 5758–5766, April 2020.
- [15] K. Bandara, C. Bergmeir, and H. Hewamalage, “LSTM-MSNet: Leveraging forecasts on sets of related time series with multiple seasonal patterns,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 4, pp. 1586–1599, 2021.
- [16] G. Dudek, “Short-term load forecasting using random forests,” *Advances in Intelligent Systems and Computing*, vol. 323, pp. 821–828, 2015.
- [17] R. Godahewa, C. Bergmeir, G. I. Webb, R. J. Hyndman, and P. Montero-Manso, “Monash time series forecasting archive,” *arXiv:2105.06643*, 2021.
- [18] Y. Bao, T. Xiong, and Z. Hu, “Multi-step-ahead time series prediction using multiple-output support vector regression,” *Neurocomputing*, vol. 129, pp. 482–493, 2014.
- [19] K. Benidis, S. S. Rangapuram, V. Flunkert, B. Wang, D. C. Maddix, A. C. Türkmen, J. Gasthaus, M. Bohlke-Schneider, D. Salinas, L. Stella, L. Callot, and T. Januschowski, “Neural forecasting: Introduction and literature overview,” *CoRR*, vol. abs/2004.10240, 2020.
- [20] S. Siami-Namini, N. Tavakoli, and A. S. Namin, “The performance of LSTM and BiLSTM in forecasting time series,” in *2019 IEEE International Conference on Big Data (Big Data)*, 2019, pp. 3285–3292.
- [21] C. Gallicchio, A. Micheli, and L. Pedrelli, “Design of deep echo state networks,” *Neural Networks*, vol. 108, pp. 33–47, 2018.
- [22] S. Kim and H. Kim, “A new metric of absolute percentage error for intermittent demand forecasts,” *International Journal of Forecasting*, vol. 32, no. 3, pp. 669–679, 2016.