Short-Term Ambient Temperature Forecasting for Smart Heaters

Danilo Carastan-Santos, Anderson Andrei da Silva, Denis Trystram, Alfredo Goldman, Angan Mitra, Yanik Ngoko, Clement Mommessin

► To cite this version:

Danilo Carastan-Santos, Anderson Andrei da Silva, Denis Trystram, Alfredo Goldman, Angan Mitra, et al.. Short-Term Ambient Temperature Forecasting for Smart Heaters. 2021. hal-03364728

HAL Id: hal-03364728

<https://hal.inria.fr/hal-03364728>

Preprint submitted on 4 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Short-Term Ambient Temperature Forecasting for Smart Heaters

Danilo Carastan-Santos*

Informatics Institute

Federal University of Rio Grande do Sul

Porto Alegre, Brazil

danilo.csantos@inf.ufrgs.br

Anderson Andrei Da Silva*, Denis Trystram*

Univ. Grenoble Alpes,

CNRS, Inria, Grenoble INP, LIG

Grenoble, France

{anderson-andrei.da-silva, denis.trystram}@inria.fr

Alfredo Goldman*

Institute of Mathematics and Statistics,

University of São Paulo, Brazil

São Paulo, Brazil

gold@ime.usp.br

Angan Mitra*, Yanik Ngoko*

Qarnot Computing

Montrouge, France

{angan.mitra, yanik.ngoko}@qarnot-computing.com

Clément Mommessin*

School of Computing

University of Leeds

Leeds, U.K.

c.mommessin@leeds.ac.uk

Abstract—Maintaining Cloud data centers is a worrying challenge in terms of energy efficiency. This challenge leads to solutions such as deploying Edge nodes that operate inside buildings without massive cooling systems. Edge nodes can act as smart heaters by recycling their consumed energy to heat these buildings. We propose a novel technique to perform temperature forecasting for Edge Computing smart heater environments. Our approach uses time series algorithms to exploit historical air temperature data with smart heaters’ power consumption and heat-sink temperatures to create models to predict short-term ambient temperatures. We implemented our approach on top of Facebook’s Prophet time series forecasting framework, and we used the real-time logs from Qarnot Computing as a use-case of a smart heater Edge platform. Our best trained model yields ambient temperature forecasts with less than 2.66% Mean Absolute Percentage Error showing the feasibility of near real-time forecasting.

Index Terms—Temperature forecasting, Edge computing, Smart heater

I. INTRODUCTION

Heating is an essential requirement for countries with cold climate and is mainly achieved by hot water or electrical heaters in residential or commercial buildings. A survey of Gudmundsson *et al.* [1] reports that the energy consumption due to heating demands constitutes 40% of the net energy produced in European countries. Another study from Awada *et al.* [2] states that the energy consumption of data centers worldwide is estimated at $2.6 \cdot 10^{10}$ Watts – which is about 1.4% of worldwide electrical energy consumption – though with an alarming growth rate of 12% per year. For instance, the power and cooling demands to operate a cloud data center constitute close to 25% of the amortized total yearly costs [3].

Solutions focused to these problems have been proposed by companies such as Qarnot Computing¹, which offers heating as a service by executing computationally intensive

applications in Cloud/Edge infrastructures composed by computing machines placed into general environments such as houses, apartments or offices [4]. This task is achieved by deploying smart heaters – which is a computing machine specialized to dissipate its produced heat in the same fashion as a traditional heater – into several heat demanding environments such as studios, residential and commercial buildings. A key hypothesis is that the heat produced as result of the smart heater executing an application can fulfill the heating demands of these environments. Applications can come either from the Internet (in a Cloud Computing fashion), or from the environments themselves, with data generated from the Internet-of-Things (IoT) devices present at these environments, thus characterizing as an Edge Computing infrastructure.

Forecasting the heat interaction between the smart heater and the environment is a key requirement for an efficient resource management of this smart heaters infrastructure. With accurate predictions, an adequate load balancing of applications into the smart heaters can be performed, to optimally fulfill the heating demands of the environments. A major hindrance to perform these predictions is that each environment is different in terms of heating insulation, climate properties, and other rather unpredictable factors such as the number of persons in the environment, the opening/closing of windows, etc.

In this paper, we present a method for forecasting the ambient temperature of environments heated by smart heaters, such as the ones composing the Qarnot Computing infrastructure. More specifically, this paper highlights the two following contributions:

- 1) We exploit short-term historical data from both the smart heaters and the environment’s temperature, and a time-series forecasting framework to perform predictions about the ambient temperature evolution, when the ambient air is being heated by a smart heater. Our experiments showed that a model combining the

*Authors’ names are sorted alphabetical order

¹<https://qarnot.com/>

smart heater’s power data with ambient temperature data performs accurate forecasts, with a Mean Absolute Percentage Error (MAPE) of $2.66\% \pm 2.52\%$ and Root Mean Squared Error (RMSE) $0.92^{\circ}\text{C} \pm 1.00^{\circ}\text{C}$, when compared with the ambient data measured *a posteriori*.

- 2) We performed an analysis about the smart heater’s feature data – notably the heater’s power consumption and heat-sink temperature – in terms of their efficiency as ambient temperature predictors. Our results showed a similar accuracy with a MAPE of $2.75\% \pm 2.55\%$ and RMSE of $0.92^{\circ}\text{C} \pm 1.03^{\circ}\text{C}$, when compared with the ambient data measured *a posteriori*, for a forecasting model that uses all of the aforementioned smart heater’s features.

The remainder of this paper is organized as follows: We discuss related work in Section II. In Section III, we explain the modeling of temperature forecast based on time series, followed by automating the model hyper-tuning process. Finally, we discuss in Section IV the model performance on actual sensor logs, ending with a conclusion in Section V that points towards the limitation and scope of future improvement. We open-source the methodology used in our work and the code base for analysis and forecasting is made publicly available at a Gitlab repository².

II. RELATED WORK

The energy demand has lead to an increased adoption of Heating, Ventilation and Air Conditioning (HVAC) systems to counter the thermal stress and provide ambient comfort [5]. Such systems rely on accurate estimate of room temperature which is prone to activity related fluctuations. Although estimating ambient temperature with phone sensors has been tried [6], the results show high correlation with the phone internal thermal state rather than the environment. Instead, building characteristics such as spatial-orientation, wall-thermal conductivity, specific heat capacity show better predictive capabilities [7]. Although traditionally room heating has been through hot water supply or electrical coil heating, there has recently been a trend for smart heaters. A smart heater is defined to be a thermal radiating unit that can auto adjust to heating requirements of a space. A Qarnot smart heater derives heat by executing programs on computational processing units like CPU and graphics cards. The programs are usually high performance compute loads or cloud jobs. We propose a novel method of a near real-time forecast by incorporating the intrinsic properties of heating systems such as the heat-sink temperature and power consumption.

There are many contexts where time series forecasting have been applied such as to predict wind speed [8], or enrolments [9]. And there are studies that compare methodologies taken as conventional with other considered modern [10]. Focused on temperature time series forecasting, it is also found in the literature several works based on different techniques

such as Neural Networks [11]–[14], Support Vector Machines [15], Fuzzy Time Series [16]–[18], Auto Regressive (AR) processes [19], ARIMA and SARIMA [20].

When the focus turns to houses indoor temperature forecasting, there are works based on Deep Learning techniques, such as Deep Neural Networks [21] or Recurrent Neural Network (RNN) and Convolutional Neural Network (CNN) [22]. Zamora-Martínez *et al.* [23] created a house, with a range of several technologies to improve the energy efficiency, and they used Artificial Neural Networks for temperature forecasting [23]. Eredics [14] pointed out that such controlled indoor environments need to take into account its surroundings, which can be done using Neural Networks. Gustin *et al.* showed that it is possible to use the heatwaves periods to predict the temperature of buildings [19]. Taylor and Letham [24] proposed Prophet, a flexible tool to decompose time-series into components like trend, seasonality and irregularity. Such tool has been used and compared with other techniques in several contexts, as air pollution forecasts [25], daily or monthly stream-flow forecasting [26], [27], and water precipitation [28].

The proposed work in this paper leverages Prophet to design predictive thermal models for smart heaters, enhanced with intrinsic device readings. Our models are validated against real-time temperature logs of Qarnot smart heaters and achieves an accuracy of 2.66% Mean Absolute Percentage Error and 0.92°C of Root Mean Squared Error.

III. A METHOD FOR FORECASTING TEMPERATURE FOR SMART HEATER ENVIRONMENTS

In this section we present our method of exploiting the data that is present in the Qarnot Computing logs to perform predictions of the ambient air for future periods. Figure 1 gives a schematic of ambient temperature forecasting from the real-life logs, where we first pre-process the data followed by a hyperparameter tuning to yield the final forecasting model. We start by the 1. *Data Preprocessing*, which consists of a descriptive analysis and the removal of outliers. First we investigate the time interval of the measurements to verify the length and if there were gaps among the measurements. Then, we filtered out the outliers in the data, such as negative power measurements. Since the chosen forecasting model has several hyperparameters with many possible values, we selected the most important ones to perform our second step, the 2. *Hyperparameter Tuning* process with cross validations. For that, we also split our available data in many intervals of time with a length of two days. The first day was used for the training phase and the second day for the test phase. At the end, we selected the achieved best set of hyperparameters for our case. Finally, in our third step, the 3. *Sliding Window Temperature Forecasting*, we used the best set of hyperparameters selected in the previous step to perform two forecasts regarding the ambient air temperature. The first one considers as extra regressor the power measurements, and the second one considers as extra regressors the heat-sink temperature in addition to the power measurements. For that we use

²https://gitlab.com/andersonandreiforecasting_smart_heaters_temperature/

Fig. 1. Methodology workflow

Prophet [24], an open source tool developed by Facebook for forecasting time series data.

The remainder of this section presents the details of the Prophet model, the hyperparameters tuning process, the methodology for our short term forecasting, and the metrics utilized in the evaluation phase.

A. Prophet Forecasting Model

We implemented our method on top of the Prophet [24] time-series forecasting framework. For the sake of completeness, we present the Prophet's time-series forecasting model (Equation 1) for decomposing a time series into non-periodic changes (trend), season recurrences (seasonality) and effect of holidays:

$$y(t) = g(t) + s(t) + h(t) + \epsilon_t \quad (1)$$

where the observed signal $y(t)$ follows the additive property in combining $g(t)$, $s(t)$, $h(t)$ representing the trend, seasonality, holiday behaviour respectively, and ϵ_t represents a Gaussian white noise with 0 mean. In a real context, the ambient temperature does not keep on increasing forever, rather a heated space starts cooling when it is not supplied with more thermal energy. The set of points where the trend function has local minima or maxima or is discontinuous, are denoted as change-points ($S = \{s_1, s_2 \dots s_n\}$) and each change-point s_i comes with a rate of adjustment (d_i). We define a binary vector $a(t) = \{a_1(t), a_2(t), \dots a_n(t)\}$ with $a_i(t) = 1$ if $t \geq s_i$ or 0 otherwise. The vector $\delta = \{d_1, d_2, \dots d_n\}$ is the collection of adjustment rates or equivalently $\delta \in R^S$. $(k + a(t)^\top \delta)t$ represents the total growth starting from $t = 0$ to t where k is the initial growth rate. We want $g(t)$ to be continuous and hence the correction factor is given by $(m + a(t)^\top \gamma)$ where $\gamma = \{\gamma_1, \gamma_2 \dots \gamma_n\}$, $\gamma_j = -s_j \delta_j$ and m is a real-numbered offset. Equation 2 shows the linear growth trend used for our methodology.

$$g(t) = (k + a(t)^\top \delta)t + (m + a(t)^\top \gamma) \quad (2)$$

TABLE I
HYPERPARAMETERS SUMMARY FOR THE GRID OF EXPERIMENTS

Parameter	Set of Values
Changepoint Prior Scale	{0.01, 0.1, 1}
Seasonality Prior Scale	{0.01, 0.1, 1, 10}
Fourier Order	{1, 3, 5, 10}
Mode	{'additive', 'multiplicative'}
Size of the data	2 days

Next we model seasonal components $s(t)$ like weekly or yearly repetitive patterns with Fourier series given by Equation 3 where the argument of n^{th} cosine and sinusoidal functions are given by $f_n \times t$ where $f_n = \frac{2\pi n}{P}$, $P = 365.25$ for yearly and 7 for weekly occurrences. The variables a_n and b_n reflect the magnitude of the n^{th} harmonic and N represents the degree of approximation.

$$s(t) = \sum_{n=1}^N \left(a_n \cos\left(\frac{2\pi n t}{P}\right) + b_n \sin\left(\frac{2\pi n t}{P}\right) \right) \quad (3)$$

Weekends or holidays are special days where one can expect either a low or high degree of activity depending on an office space having no employees on a Sunday or employees staying back at home respectively. An indicator function is used to label time t as 1 if t falls on a holiday i or 0 otherwise. Let D_i represent the day of the years corresponding to a holiday i . For example Christmas happens on 25th December every year, thus yielding $D = \{\dots, 25/12/2019, 25/12/2020, 25/12/2021, \dots\}$. At time t , for L holidays, we construct a binary matrix $Z(t) = [1(t \in D_1), \dots, 1(t \in D_L)]$. The effect of holidays are assumed independent while modeling temperature on such days and magnitude of change κ_i is drawn from $\kappa \sim Normal(0, v^2)$. Formally the effect of the holidays is given by Equation 4.

$$h(t) = Z(t)\kappa \quad (4)$$

B. Hyperparameters Tuning

To achieve a model with the least generalization error, we investigate the following hyperparameters from the Prophet model: a) *Changepoint prior scale* that controls the trend functions modeling the non-periodic changes in the data, i.e $g(t)$ in Equation 2; b) *Seasonality prior scale* which models the periodic changes in the data, i.e $s(t)$ in Equation 3; c) *Fourier order* that indicates resolution of frequency decomposition and works in tandem with the seasonality prior scale; d) *Mode* which indicates the effect of the regressors in the model.

For evaluating good hyperparameter combinations, we designed a grid experiment as shown in Table I which resulted in $2 \times 3 \times 4 \times 4 = 96$ combinations. For each combination, we perform a time series forecasting cross-validation test. For this task a forecast model needs to be trained on sequential partitions of data, labeled as train and test. We select the model with the least generalization error on the cross-validation test, according to the RMSE and MAPE metrics. Experimental

Fig. 2. Sliding window procedure for forecasting short-term air temperatures.

results regarding the aforementioned grid evaluation of the hyperparameters are presented in Section IV-B.

C. Metrics Evaluation

The forecasting model is trained using Stan's implementation of the Limited-Memory Broyden-Fletcher-Goldfarb-Shanno Algorithm (L-BFGS) [29] to find a maximum *a posteriori* estimate of ambient temperature. The model accuracy is evaluated by calculating the RMSE (Root Mean Square Error) $= \sqrt{\frac{1}{n} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2}$ and MAPE (Mean Absolute Percentage Error) $= \frac{1}{n} \sum_{i=1}^n \left| \frac{Y_i - \hat{Y}_i}{Y_i} \right|$ considering n predictions where Y_i is the vector of observed values of the variable being predicted, and $\hat{Y}_i$ being the predicted values. The RMSE metric is useful in highlighting major deviations from the true sensor value and a low RMSE score ensures less signal to noise ratio for temperature prediction, while MAPE provides intuitive performance measure by averaging the absolute deviation over the set of data points.

D. Short Term Forecasting

We hypothesize that ambient temperature is an interaction between the smart heater's electrical power consumption and the thermal radiation of its heat-sink. In particular we address the question of forecasting the air temperature over one hour horizon. We illustrate the forecasting performance by comparing the two following model variants: *model 1*, which depends on the power consumption of the smart heater and *model 2*, which relies on dual channel inputs of power consumption and heat-sink temperature. We follow a sliding window protocol length of 24 hours to predict the next hour and hence generate a forecast for the desired time-interval.

Figure 2 illustrates an example of the step-wise temperature forecasting with an one hour sliding window protocol. The first forecast, at *step 1*, takes the historical data of 24 hours from 23rd November 09:00 am to 24th November 09:00 am as input to predict the values for the 25th hour. For the next steps, we shift the window of the training data by one hour. The combined forecast of these 4 steps is shown at the bottom facet of Figure 2 at *steps combination*.

IV. EXPERIMENTAL RESULTS

In this section we present the results for each step of our methodology presented on the previous one. We analysed several data-sets constituted by logs from Qarnot's smart heaters, from those we selected the most consistent one in terms of gaps and outliers. It consists of 510567 samples, with the following columns of interest namely *time-stamp*, *ambient air temperature*, *heat-sink temperature* and *power consumed*. As introduced on Section I we achieved very good and similar results for our models 1 and 2, then we will not present both details. Since we do not have enough evidences to distinguish their accuracies, we selected model 2 to be detailed due to the number of features that it treats, such as the power and the heat-sink. In the following of this section, we discuss the results of data pre-processing along with hyperparameter evaluation and forecasting efficiency.

A. Data Preprocessing

In the first step, we reject data points with abnormal values of temperature and power based on valid pre-defined intervals of $[0, 85]^{\circ}\text{C}$ and $[0, 1000]$ Watts respectively. These interval limits are derived from the typical operating conditions en-listed in the data-sheet of Qarnot's smart heaters. This resulted in an 11 % decrease in data-set size. A major challenge in non-synthetic data-set is the occurrence of missing values for substantial amounts of time. Just to illustrate such behavior, out of the 510567 samples, the two largest continuous periods found contain 90721 and 253320 samples, representing about 18% and 50% of the data respectively.

B. Hyperparameters Evaluation

Figure 3 summarizes the hyperparameters performance on forecasting with a variation of the changepoint prior scale from the set $\{0.01, 0.1, 1\}$ in the x axis, with each subplot representing the results of the seasonality prior scale from the set $\{0.01, 0.1, 1, 10\}$. The geometrical elements (dots, triangles, squares, and crosses) represent the mean RMSE of a certain Fourier order from the set $\{1, 3, 5, 10\}$, and the vertical lines represent confidence intervals as standard deviation of RMSE values on the y axis. The effect of altering *mode* from $\{\text{'additive'}, \text{'multiplicative'}\}$ showed no effect on both RMSE mean or standard deviation. Hence, we only discuss the results for the additive mode hereafter, using $\frac{96}{2} = 48$ hyperparameter combinations. We shall denote the changepoint, seasonality and Fourier order as a triplet of 3 values for a configuration. Some good performing configurations found are $(1, 0.01, 3)$ with RMSE of $0.98^{\circ}\text{C} \pm 1.05^{\circ}\text{C}$, or $(0.1, 0.01, 3)$ with error

Fig. 3. Hyperparameters Performance Evaluation with Mean RMSE and Standard Deviation

margin of $1.00^{\circ}\text{C} \pm 1.04^{\circ}\text{C}$. The utility of the cross validation is evident by observing bad configurations such as (1, 10, 1) with RMSE of $2.05^{\circ}\text{C} \pm 3.47^{\circ}\text{C}$, 45 % performing less than the best found configuration or (1, 1, 1) with error of $2.06^{\circ}\text{C} \pm 3.57^{\circ}\text{C}$, a 1 % increase in standard deviation error compared to the former.

In order to analyze the effect of each hyperparameter, we re-grouped the results presented in Figure 3 and present them in Table II. The changepoint prior scale variation captures the non-periodic abrupt changes in the time series. The longer the interval of consideration, more likely are the chances of abnormal patterns. Thus, we lower the scope of error by considering only a day ahead of training while forecasting for the next day. From historical data, we observe that the major of pre-seen ambient temperature pattern stays in an interval larger than two days and hence our methodology is reactive towards capturing a new pattern and adjusting the forecast accordingly. Nevertheless, it is possible to see that as far as we increase the value of this hyperparameter the errors increase as well, such as from $1.33^{\circ}\text{C} \pm 2.02^{\circ}\text{C}$ to $1.54^{\circ}\text{C} \pm 2.30^{\circ}\text{C}$.

Low values of seasonality prior scale exhibit lesser generalization error. Higher values reflect greater temperature fluctuations thereby increasing the error margin, such as from $1.22^{\circ}\text{C} \pm 1.65^{\circ}\text{C}$ to $1.57^{\circ}\text{C} \pm 2.41^{\circ}\text{C}$. Notably as per our approach, alternate day training-forecasting reduces the possibility to register a high number of periodic changes within a span of 48 hours.

The Fourier order value of 1 resulted in average RMSE of $1.91^{\circ}\text{C} \pm 3.16^{\circ}\text{C}$ where as setting the value to 3 leads to $1.27^{\circ}\text{C} \pm 1.76^{\circ}\text{C}$. The values of 5 and 10 yields mean RMSE of $1.20^{\circ}\text{C} \pm 1.56^{\circ}\text{C}$ and $1.43^{\circ}\text{C} \pm 1.86^{\circ}\text{C}$ respectively.

TABLE II
EFFECT OF HYPERPARAMETERS MEASURED BY RMSE AND STANDARD DEVIATION

Changepoint Values	Mean RMSE	Standard Deviation
0.01	1.33	2.02
0.10	1.49	2.25
1.00	1.54	2.30
Seasonality Values	Mean RMSE	Standard Deviation
0.01	1.22	1.65
0.10	1.47	2.19
1.00	1.56	2.43
10.00	1.57	2.41
Fourier Order	Mean RMSE	Standard Deviation
1	1.91	3.16
3	1.27	1.76
5	1.20	1.56
10	1.43	1.86

The optimal value of 3 has been stated by the authors of Prophet [24] to capture weekly periodic changes and matches the expectation of our office space setting, where one can expect repetition of similar activities every 7 days.

Finally, comparing the results of the groups of Table II, we can see that the Fourier order is the hyperparameter which affects more the models RMSE accuracy varying from the worst to the best results in 0.71°C , followed by the seasonality prior scale with 0.35°C and the changepoint prior scale with 0.21°C .

C. Time Series Temperature Forecasting

The placement of Qarnot's smart heaters ubiquitously in buildings acts on the thermal state of a space. Dynamic thermal patterns are evident from Figure 4 where one can see higher power and heat-sink temperature may not necessarily increase the ambient temperature. One way to understand this phenomenon is thermal equilibrium between temperature outside the space considered and the spatial environment around it. For instance, let us consider the period around the beginning of December, when we have high values regarding the power and the heat-sink temperature but the ambient temperature is low. However, we generally see that the air temperature captures most part of the power and heat-sink variations, for example around middle of December, when the air temperature exhibits a periodic pattern of highs and low which correspond to similar behaviors on the power and the heat-sink. The explanation is that the smart heater was turned off and as consequence the air temperature has decreased. Regarding the abrupt change in the power and heat-sink, at such moments the smart heater was turned on again and the air temperature started to increase as well. The figure also illustrates that it is indeed possible to predict the ambient temperature of smart heaters and schedule computational-loads accordingly. A prominent example is the last two weeks of December where office utility was low and hence more jobs were run, resulting in heating and cooling cycles with zero emission thermal dissipation through walls or connecting surfaces.

The third facet of Figure 4 shows the comparative forecast of two models: *power* versus *power and heat-sink* combination

Fig. 4. Comparison of temperature forecasts for the two models. The first and second facets represent the power and heat-sink inputs to the forecasting models. The results of such models are illustrated in the third facet in addition to the original air temperature data.

in forecasting ambient air temperature for 4 months (x axis) and the pre-processed device-intrinsic values of power (Watts), heat-sink ($^{\circ}\text{C}$) and ambient air temperature ($^{\circ}\text{C}$) on the y axis. For both, power and heat-sink, there is only the original values and no forecasts. The optimal uni-channel *model 1* recorded a MAPE of $2.66\% \pm 2.52\%$ and RMSE of $0.92^{\circ}\text{C} \pm 1.00^{\circ}\text{C}$ versus *model 2*'s MAPE of $2.75\% \pm 2.55\%$ and RMSE of $0.92^{\circ}\text{C} \pm 1.03^{\circ}\text{C}$.

In addition, we analyse the distribution of the absolute error $|Y_t - \hat{Y}_t|$ between the forecast $\hat{Y}_t$ and the original value Y_t as shown in Figure 5 for *model 2*. We see that 93.53% of the net predictions are lower or equal to 0.5°C with maximum error difference at 3.64°C . This is most likely attributed to unpredictable environmental usages like opening of a door or window.

V. CONCLUSIONS AND FUTURE WORK

In this work, we shed light on understanding how one can perform ambient temperature predictions of environments heated by a Cloud/Edge infrastructure of smart heaters, a network of computing machines designed to dissipate their heat in the same manner as a standard heater. We developed an ambient temperature forecasting method based on short-term time series prediction, which considers the historical data of the ambient temperature and the smart heater's properties, notably its power consumption and heat-sink temperature.

We implemented our method using the Prophet [24] time-series forecasting framework, and we considered the Cloud/Edge smart heaters infrastructure of Qarnot Computing as a case example. Our experimental results show that a relatively simple trained time-series prediction model is

Fig. 5. Histogram distribution (log scale on y axis) of the absolute difference between the original temperature data and the temperature forecasts performed by model 2. Samples at the left of the vertical dotted line account for 93.53% of the total number of samples (forecasts), and the samples at the right of the same line account for the remaining 6.47%.

capable of accurately predicting the ambient temperature of an environment heated by a smart heater. Our simplest evaluated model – which uses the ambient temperature and the heater's power consumption historical data – performed ambient temperature predictions with a Mean Absolute Percentage Error of $2.66\% \pm 2.52\%$ and Root Mean Squared Error $0.92^{\circ}\text{C} \pm 1.00^{\circ}\text{C}$, when compared with the ambient data measured *a posteriori*.

We also performed an experimental campaign to evaluate which potential smart heater's features would better perform

for ambient temperature forecasting. With the Qarnot Computing data, we had access to two good features, the smart heater's power consumption, and the heat-sink temperature. Our experiments showed that, in comparison with a model that considers only the ambient temperature and the heater's power consumption historical data, adding the heat-sink temperature achieved similar results.

This work is a step towards learning how well standard processors can recycle their dissipated heat to fulfill the everlasting heat requirements of buildings, thus using energy more efficiently and contributing with ways to deal with the growth of power demands for both heating and maintaining data centers. An ambient temperature prediction method is a crucial component for better resource management of a smart heater computing infrastructure. With accurate predictions, one can better balance the smart heaters' computing workload to provide optimal heat.

For future work, we plan to advance in two main axes. The first axis is to evaluate more sophisticated time-series prediction models such as Long Short-Term Memory (LSTM) and better ways to deal with missing data sections in the historical data. The second axis is to characterize the smart heater power consumption of Cloud/Edge computing workloads to perform a better load balancing, taking into account the temperature predictions.

ACKNOWLEDGMENTS

This work was supported by the ANR Greco project 16-CE25-0016-01 and by the research program on Edge Intelligence of the Multi-disciplinary Institute on Artificial Intelligence MIAI at Grenoble Alpes (ANR-19-P3IA-0003).

REFERENCES

- [1] O. Gudmundsson, J. Thorsen, and L. Zhang, "Cost analysis of district heating compared to its competing technologies," *WIT Transactions on Ecology and the Environment*, vol. 176, pp. 3–13, 2013.
- [2] U. Awada, K. Li, and Y. Shen, "Energy consumption in cloud computing data centers," *International Journal of Cloud Computing and Services Science (IJ-CLOSER)*, vol. 3, 06 2014.
- [3] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks," 2008.
- [4] Y. Ngoko, N. Saintherant, C. Cerin, and D. Trystram, "Invited paper: How future buildings could redefine distributed computing," in *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, May 2018, pp. 1232–1240.
- [5] L. Klein, J.-y. Kwak, G. Kavulya, F. Jazizadeh, B. Becerik-Gerber, P. Varakantham, and M. Tambe, "Coordinating occupant behavior for building energy and comfort management using multi-agent systems," *Automation in construction*, vol. 22, pp. 525–536, 2012.
- [6] P. Wijukkana, A. Julsereewong, and T. Thepmanee, "Temperature variation modeling for improving building hvac control using built-in temperature sensors in smartphones," in *2017 14th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology (ECTI-CON)*. IEEE, 2017, pp. 790–793.
- [7] T. Oreszczyn, S. H. Hong, I. Ridley, P. Wilkinson, W. F. S. Group *et al.*, "Determinants of winter indoor temperatures in low income households in england," *Energy and Buildings*, vol. 38, no. 3, pp. 245–252, 2006.
- [8] B. U. Islam, "Comparison of conventional and modern load forecasting techniques based on artificial intelligence and expert systems," *International Journal of Computer Science Issues (IJCSI)*, vol. 8, no. 5, p. 504, 2011.
- [9] Kurniawan Tanuwijaya and S.-M. Chen, "A new method to forecast enrollments using fuzzy time series and clustering techniques," in *2009 International Conference on Machine Learning and Cybernetics*, vol. 5, 2009, pp. 3026–3029.
- [10] B. Brown, R. Katz, and A. Murphy, "Time series models to simulate and forecast wind speed and wind power," *Journal of Climate and Applied Meteorology*, vol. 23, pp. 1184–1195, 07 1984.
- [11] A. Khotanzad, M. H. Davis, A. Abaye, and D. J. Maratukulam, "An artificial neural network hourly temperature forecast with applications in load forecasting," *IEEE Transactions on Power Systems*, vol. 11, no. 2, pp. 870–876, 1996.
- [12] K. Zhang, A. Guliani, S. Memik, G. Memik, K. Yoshii, R. Sankaran, and P. Beckman, "Machine learning-based temperature prediction for runtime thermal management across system components," *IEEE Transactions on Parallel and Distributed Systems*, vol. PP, pp. 1–1, 07 2017.
- [13] L. Nengbao, V. Babushkin, and A. Afshari, "Short-term forecasting of temperature driven electricity load using time series and neural network model," *Journal of Clean Energy Technologies*, vol. 2, 01 2014.
- [14] P. Eredics, "Short-term external air temperature prediction for an intelligent greenhouse by mining climatic time series," *2009 IEEE International Symposium on Intelligent Signal Processing*, pp. 317–322, 2009.
- [15] Y. Radhika and M. Shashi, "Atmospheric temperature prediction using support vector machines," *International journal of computer theory and engineering*, vol. 1, no. 1, p. 55, 2009.
- [16] N.-Y. Wang and S.-M. Chen, "Temperature prediction and taifex forecasting based on automatic clustering techniques and two-factors high-order fuzzy time series," *Expert Systems with Applications*, vol. 36, pp. 2143–2154, 03 2009.
- [17] K.-B. Song, S.-K. Ha, J.-W. Park, D.-J. Kweon, and K.-H. Kim, "Hybrid load forecasting method with analysis of temperature sensitivities," *IEEE Transactions on Power Systems*, vol. 21, no. 2, pp. 869–876, 2006.
- [18] S.-M. Chen and J.-R. Hwang, "Temperature prediction using fuzzy time series," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 30, no. 2, pp. 263–275, 2000.
- [19] M. Gustin, R. Mcleod, and K. Lomas, "Forecasting indoor temperatures during heatwaves using time series models," *Building and Environment*, vol. 143, 07 2018.
- [20] P. Chen, A. Niu, D. Liu, W. Jiang, and B. Ma, "Time series forecasting of temperatures using sarima: an example from nanjing," *IOP Conference Series: Materials Science and Engineering*, vol. 394, p. 052024, 08 2018.
- [21] P. Romeu Guallart, F. Zamora-Martínez, P. Botella-Rocamora, and J. Pardo Albiach, "Time-series forecasting of indoor temperature using pre-trained deep neural networks," vol. 8131, 09 2013, pp. 451–458.
- [22] M. Cai, M. Pipattanasomporn, and S. Rahman, "Day-ahead building-level load forecasts using deep learning vs. traditional time-series techniques," *Applied Energy*, vol. 236, pp. 1078–1088, 02 2019.
- [23] F. Zamora-Martínez, P. Romeu Guallart, P. Botella-Rocamora, and J. Pardo Albiach, "On-line learning of indoor temperature forecasting models towards energy efficiency," *Energy and Buildings*, pp. –, 11 2014.
- [24] S. Taylor and B. Letham, "Forecasting at scale," *The American Statistician*, vol. 72, 09 2017.
- [25] K. Samal, K. Babu, S. Das, and A. Acharya, "Time series based air pollution forecasting using sarima and prophet model," in *proceedings of the 2019 international conference on information technology and computer communications*, 08 2019, pp. 80–85.
- [26] G. Papacharalampous and H. Tyralis, "Evaluation of random forests and prophet for daily streamflow forecasting," *Advances in Geosciences*, vol. 45, pp. 201–208, 08 2018.
- [27] H. Tyralis and G. Papacharalampous, "Large-scale assessment of prophet for multi-step ahead forecasting of monthly streamflow," *Advances in Geosciences*, vol. 45, pp. 147–153, 08 2018.
- [28] G. Papacharalampous, H. Tyralis, and D. Koutsoyiannis, "Predictability of monthly temperature and precipitation using automatic time series forecasting methods," *Acta Geophysica*, vol. 66, pp. 807–831, 08 2018.
- [29] D. C. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," pp. 503–528, 1989.