IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 9 March 2023, accepted 8 April 2023, date of publication 17 April 2023, date of current version 21 April 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3267810

RESEARCH ARTICLE

Online Machine Learning for 1-Day-Ahead Prediction of Indoor Photovoltaic Energy

FRANK ALEXANDER KRAEMER[1], (Member, IEEE), HAFIZ AREEB ASAD[1], KERSTIN BACH[2], (Member, IEEE), AND BERND-CHRISTIAN RENNER[3]

[1]Department of Information Security and Communication Technology, Norwegian University of Science and Technology (NTNU), 7034 Trondheim, Norway
[2]Department of Computer Science, Norwegian University of Science and Technology (NTNU), 7034 Trondheim, Norway
[3]Department of Information Security and Communication Technology, Hamburg University of Technology (TUHH), 21073 Hamburg, Germany

Corresponding author: Frank Alexander Kraemer (kraemer@ntnu.no)

This work was supported in part by the Norwegian Research Center for AI Innovation (NorwAI) under Grant 309834, and in part by the Research Council of Norway.

ABSTRACT We explore the potential for predicting indoor photovoltaic energy on a forecasting horizon of up to 24 hours. The objective is to enable energy management approaches that exploit harvesting opportunities more strategically, for which they require more accurate energy intake predictions. Our study is based on a data set covering over 3 years, for which we simulate online machine learning algorithms with different amounts of training data and input features. Our results show that relatively simple machine learning methods can outperform a persistent predictor considerably, and we observed a reduction of errors of up to 56%. When devices obtain a significant amount of sunlight, adding the weather forecast improves the prediction accuracy. We discuss prediction features, the amount of training data and analyze the sources of errors to understand the potential of indoor photovoltaic energy harvesting predictions.

INDEX TERMS Energy harvesting, photovoltaic systems, energy predictions, machine learning.

I. INTRODUCTION

Energy harvesting is a key enabler for large-scale deployments of wireless, embedded sensor devices, as it removes the need for manual interventions in the form of battery replacements. One downside of energy harvesting, however, is that energy is not always available, and that the devices therefore need to do energy management and be strategic about when to use how much energy. In the future, we expect energy management to get even more significant. Embedded devices are continually gaining computational capabilities, so that they are able to perform machine learning inference and training [1], [2]. This enables use cases like anomaly detection [3] or image recognition [4] directly on the device. In the context of energy harvesting, such computationally complex tasks require and enable more strategic behavior with respect to energy management: The training of machine learning models is often delay-tolerant and can be performed at a later time when there is an energy surplus. Other tasks may require fully charged energy buffers at a specific time, like the transmission of large amounts of data when a transmission window is available, for instance through low-orbit satellites [5]. Yet other tasks may be performed at different quality or accuracy levels, depending on the available energy [6]. This implies that instead of just using whatever energy is available, devices can act more strategically when spending their energy and decide which tasks and when achieve the best outcomes.

For power management to be more strategic, having more accurate predictions for the incoming energy is therefore important, since more accurate predictions enable to perform more advanced scheduling [7], [8], [9], [10], [11]. Instead of using larger energy buffers and accepting the unavoidable losses when buffering energy, delay-tolerant tasks may be executed in phases of energy surplus and bypass the energy buffer. Likewise, expected energy shortages can be mitigated by early reduction of consumption if predictions are sufficiently accurate [12].

Since photovoltaic cells are a promising option for energy harvesting in indoor settings [13], we assess to which degree and how their energy intake can be predicted. The prediction of solar energy has been studied extensively in outdoor settings (see, for instance [14]), but there is little work

The associate editor coordinating the review of this manuscript and approving it for publication was Valerio Freschi.

related to indoor settings. Additional factors like artificial lighting, room utilization, and obstructions play a significant role. Most existing approaches focus on short time horizons, and only predict the immediate next time period ahead, like 10 minutes or an hour, which is insufficient when planning ahead for the scenarios mentioned above. Strategic planning requires longer forecasting horizons [15], [16]. Therefore, we analyze how predictions can be produced at midnight for the entire upcoming day, so that they are available for medium-term energy planning over the next 24 hours. In addition, many existing approaches use a training-test split that implies to initially collect data for a long period. Instead, we study an online setting that starts with little training data.

In the following, we use the comprehensive data set from Sigrist et al. [17] to illustrate, discuss and investigate the predictability of indoor photovoltaic energy. We use simulations of online learning with various parameter settings such as the amount of training data and used features. We analyze the sources of errors, which also reveal general challenges with indoor predictions. The scope of the study is predicability and understanding its limits, not the implementation or computational complexity to which we, however, provide some comments in the end.

Our results show that relatively simple machine learning models significantly outperform a persistence-based predictor for all locations. Accuracy varies with the individual locations, for which we selected representative devices from the data sets. Some locations may be subject to unpredictable environment changes typical for indoor deployments (like curtains or blinds), which lead to inaccuracies, both for machine-learning-based and persistence-based methods. When devices obtain a significant portion from natural sunlight, adding the weather forecast to the prediction improves accuracy. We also show that only a few days of training data are sufficient, so that we can observe a reduction in errors of up to 56% compared to persistence-based methods.

In the following, we provide an overview of related work and present the data set in Sect. III. We describe the metrics, baseline and machine learning models for harvest prediction in Sect. IV, analyze the causes or errors in Sect. V and the amount of training data in Sect. VI-A, before we close in Sect. VII.

II. RELATED WORK

Energy harvest prediction has been discussed for various settings and harvesting techniques. A fundamental prediction technique are persistence-based models, which take the observed intake from previous intervals, typically referred to as *time slots*, directly as estimates for future ones. For solar energy this is often a suitable approach since it captures the diurnal cycle of the Sun. These approaches are especially relevant for embedded systems as they only require previous observations, that can be accumulated locally, and few operations. Hence, they have a small memory footprint and low computational effort. The exponentially weighted moving average (EWMA) predictor [10] computes the moving average per time slot from the past days, which is directly used as prediction. The weather-conditioned moving average (WCMA) [9] combines the value of the previous time slot and the current time slot's mean value for a given number of past days, where an additional scaling of an observed (weather) trend is applied. Pro-Energy [18] is a harvesting predictor which identifies energy profiles of ''typical'' days. It selects the energy profile as prediction that is most similar to previously observed time slots.

For an outdoor setting, Ahmed et al. [19] use historic observations to build a probabilistic model of the energy harvest at a specific location. Buchli et al. [20] use an astronomical model to estimate the available energy and adjust device operation to the actually received energy. The historical or astronomical models only apply for outdoor settings and do not take indoor effects like artificial lights into account. Machine learning has also been used for the forecast of outdoor solar energy, usually within the domain of renewable energy production. Voyant et al. [14] provide a comprehensive review over these methods. For machine learning approaches for embedded devices, Kraemer et al. [21] have shown the positive impact of weather information into medium-term energy predictions.

In contrast to the outdoor setting, the prediction of indoor photovoltaic energy received much less attention. Berner [22] applies various machine learning models on the data set, but only studies short-term predictions, i.e., one time slot ahead. The models use training data that implies collecting data over a longer period before any predictions are made. Instead, we focus on the application of online learning, where models are frequently retrained as more observations become available.

Stricker and Thiele [23] apply random forests on the same indoor data set. They train models initially offline, deploy them on the devices, and then apply pruning and weighting of trees through online learning. This learning process increases prediction accuracy only marginally, but reduces the number of trees to use for computation. In contrast to our work, they only use previously harvested energy as input feature, and focus on a shorter prediction intervals (10 minutes) and forecasting horizon (1 hour).

An entirely different family of approaches tries to predict energy harvesting opportunities *before* devices are deployed. Kim et al. [24] predict energy intake based on ray tracing of light based on building plans. Schneider [25] presents an analytical model for indoor solar energy based on illuminance measurements. Both approaches require detailed knowledge of the room geometry. In such a setting, Ma et al. [26] show the significance of taking the spectrum of light into account, especially in conditions with mixed light sources.

These last works show the abundance of information and factors to take into account when forecasting indoor energy intake using manually constructed models. In contrast, we study the application of a data-driven approach where the incoming energy is directly predicted based on a set of prediction variables. This indirectly covers effects like room

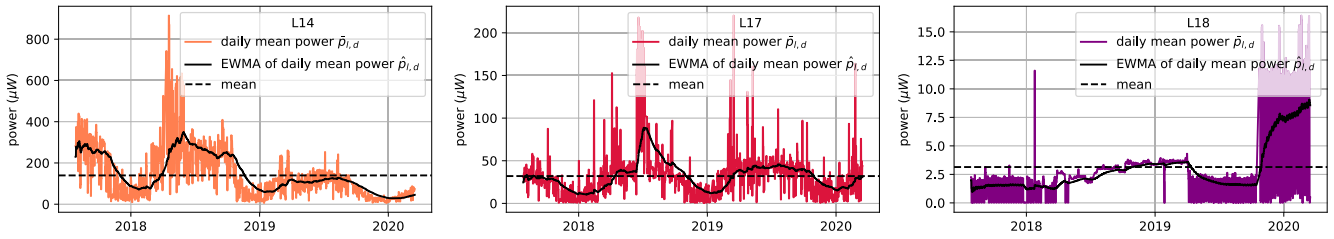

FIGURE 1. Daily average power, EWMA of daily average power and total mean power.

geometry, mixtures of different light sources, translation from illuminance to energy intake and solar panel efficiency. This hence enables a more autonomous approach required when many devices need to be deployed or when conditions change after deployment.

III. INDOOR HARVESTING DATA

We use the data set by Sigrist et al. [17], which contains power traces of device locations l, covering a period of 961 days from 2017 to 2020.[1] With $p_{l,d,i}$, we denote the average power at location l on day d during time slot i, $i \in \{0..23\}$. The mean power within a day is $\bar{p}_{l,d} = \frac{1}{24}\sum_{i\in 0..23} p_{l,d,i}$. We also calculate the EWMA $\tilde{p}_{l,d}$ of the daily means

$$\tilde{p}_{l,d} = \begin{cases} \bar{p}_{l,d} & d = 1 \\ \alpha\, \bar{p}_{l,d} + (1-\alpha)\, \tilde{p}_{l,d-1} & d > 1. \end{cases} \quad (1)$$

We use $\alpha = 0.095$, which means that mainly the last 20 days contribute to the average. This value is a good trade-off between following the variations but smoothing sufficiently to see the general trend of the data. We chose the three locations L14, L17, L18 that are exposed differently to natural and artificial light. Fig. 1 shows the daily average $\bar{p}_{l,d}$, the moving average over the daily means $\tilde{p}_{l,d}$ and the total mean power over the entire period. Note the different scaling of the vertical axes. The locations receive different amounts of artificial light, indirect and direct sunlight, summarized in Table 1. The power harvested at the different locations varies in the order of two magnitudes. The mean power of L18, which receives only indirect, artificial lighting, is 3.1 μW. In contrast, the mean power of L14 is 140 μW.

When the influence of natural light is significant, like for L14 and L17, harvested power is subject to seasonality, and it is lower during winter than summer periods. In addition, we notice concept drifts (not further elaborated in [17]), indicating changes in the environment: For L14, the summer of 2020 shows significantly less energy intake than the summer of 2018. L18 shows a notable change in the autumn of 2019 (Fig. 1).

To study the data on a finer time scale, Fig. 2 shows two weeks for each of the locations. The power available at L14 varies considerably from years 2018 to 2020. The artificial lighting for L18 must have changed, as the power profile considerably differs between the two weeks. For L18 we see a dependency on the weekday; power from Monday to Friday follows the working hours which affect the artificial lighting.

To study the influence of sunlight, we plot the power based on the position of the Sun. The plots in Fig. 3 show the maximum power observed within combinations of the solar angles azimuth and elevation. The azimuth describes the position of the Sun during the year and elevation during the day. We see that there is a significant dependency on the solar angles at positions L14 and L17. We also see the direct sunlight in the morning hours on L14, and the sunlight in the afternoons at L17. The variations after sunset show dependencies on artificial lighting. In contrast, L18 is rather independent from the Sun's position.

[1] All data and code used in our experiments, including weather observations and forecasts,is available from https://github.com/falkr/indoor-photovoltaic-energy-prediction

TABLE 1. Overview of the locations with a qualitative description of lighting conditions provided by [17] and [27].

Location	Light Exposure (from [17])
L14 "Window"	Laboratory, wall mounted at 2.1 m height, significant natural light with increased level and potential direct sunlight in morning hours and during summer.
L17 "Curtain"	Employee office, wall mounted at 2.4 m height, significant natural light with increased level in the afternoon, no direct sun exposure. According to [28], there is a curtain that sometimes partly obstructs the window.
L18 "Hallway"	Hallway, wall mounted at 2.2 m height, no natural light, only little, indirect artificial light due to high wall mount mounting position.

IV. ENERGY HARVESTING PREDICTION

Aiming at medium-term power management, we are interested in predictions that cover an entire day with a resolution of 1 hour. Lower resolutions will likely not improve energy planning, and tend to be more inaccurate [21].

A. PREDICTION PERFORMANCE METRICS

We use metrics that assign scores for a predictor for an entire day, that means, the entire prediction horizon of 24 hours. We define the scaled total absolute percentage error (STAPE) and the scaled mean percentage error (SMAPE):

$$\text{STAPE}(l,d) = \frac{100}{24 \cdot \tilde{p}_{l,d}} \left| \sum_{i\in 0..23} p_{l,d,i} - \sum_{i\in 0..23} \hat{p}_{l,d,i} \right| \quad (2)$$

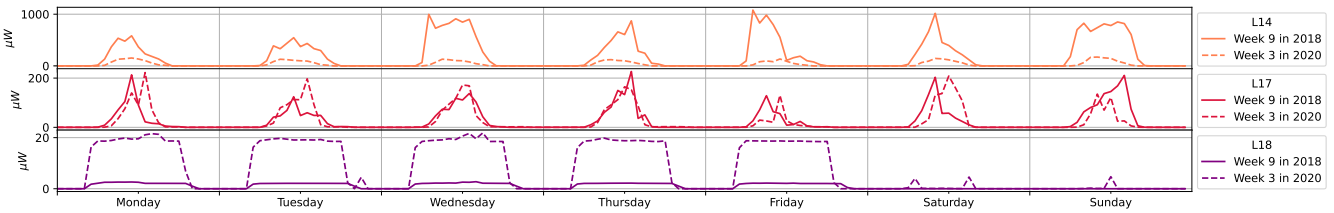

FIGURE 2. Detailed power over two distinct weeks for each location.

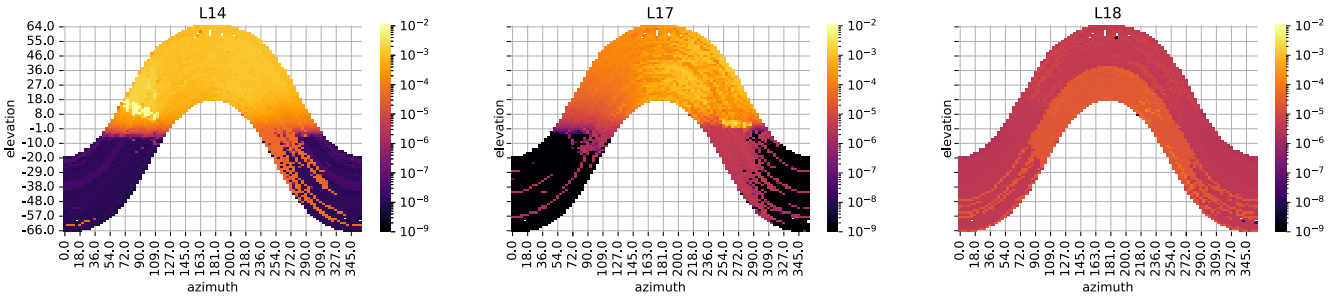

FIGURE 3. Maximum observed power over solar angles azimuth and elevation.

$$\text{SMAPE(l, d)} = \frac{100}{24 \cdot \tilde{p}_{l,d}} \sum_{i \in 0..23} \left| p_{l,d,i} - \hat{p}_{l,d,i} \right| . \tag{3}$$

Both are scaled as percentages to the moving average of the daily mean values introduced in (1). This avoids outliers of percentage errors and makes the metrics scale-independent [28]. The two metrics differ in how they handle summation within days: STAPE first sums over predictions and actual values and compares those sums. It is hence a measure of how many percent, relative to the average intake during a period, the total energy for a day is off. SMAPE first determines the absolute error for each time slot and then summarizes the results. This means SMAPE also takes timeliness into account, i.e., how accurate the predictions for each time slot are. A good SMAPE implies a good STAPE, but not vice-versa. Since both metrics are absolute, we can evaluate the quality of a predictor by averaging the scores for all days of a simulation run.

B. BASELINE PREDICTORS

A persistent predictor uses the value of the time slot from the previous day, that means,

$$\hat{p}^{\text{LAG24}}_{l,d,i} = p_{l,d-1,i} . \tag{4}$$

We also evaluated the EWMA predictor [10] explained in Sect. II, but it only slightly outperformed the LAG predictor for one location (L17) for the average STAPE, while it scored considerably worse in terms of average SMAPE for all locations and all parameter settings. We did not consider WCMA since it did not show to improve predictions beyond the next time slot [29], while our prediction should address the entire day ahead. In Pro-Energy, predictions have a strong dependency on similarity within the earlier time slots and is not suitable in our setting as the time slots around midnight are usually similar and close to zero, and hence not a good indicator for the energy during the day ahead. We therefore use here $\hat{p}^{\text{LAG24}}_{l,d,i}$ as simple prediction baseline and refer to it as the LAG predictor.

Row 0 of Table 2 shows the results of the LAG predictor for the different locations. The STAPE is between 31.52% and 58.54%, and the SMAPE is even higher, since it also requires the predictor to be accurate for each hour.

C. ONLINE MACHINE LEARNING MODELS

We have seen in the previous sections that the energy intake in the given indoor setting is highly specific for the individual locations, and in practice only known after the deployment. We hence study machine learning models for the energy prediction in an *online* learning setting, that means, a setting where data is used as it becomes available over time and new machine learning models are trained frequently. Training of the models as new data arrives takes care of the non-stationarity, and we split between training and test data according to what real deployments would have available. We produce predictions at midnight for the entire next day ahead, which can be exploited for energy planning with hourly intervals as input.

To evaluate a prediction method, we create a series of prediction models $M_{l,d}$. We train new models $M_{l,d}$ at the beginning of each day d for position l to produce predictions $\hat{p}^{M}_{l,d,i}$ for the 24 time slots ahead.[2] The training and test data

[2]We assume here that the training is executed at midnight, at the same instance data from the previous day $d-1$ becomes available and the predictions for the next day are due. This is an idealization to make the discussion simpler but is no significant limitation in practice.

TABLE 2. Mean STAPE and SMAPE of the different predictors.

	Predictor	Features	$t_{\max}$	L14 "Window" STAPE	L14 "Window" SMAPE	L17 "Curtain" STAPE	L17 "Curtain" SMAPE	L18 "Hallway" STAPE	L18 "Hallway" SMAPE	Sect.
0	LAG	power_lag24		38.11	49.67	58.54	76.6	31.52	34.96	IV-B
1	RFR	workday, timecount, dayofweek, power_lag24	10	38.42	52.2	56.27	77.28	13.92	19.76	IV-E
2			30	37.55	51.28	53.8	75.0	13.87	18.68	
3		workday, timecount, dayofweek, power_lag24, elevation, azimuth	10	36.92	49.07	56.63	74.69	13.63	18.3	
4			30	35.66	48.25	55.82	74.82	13.85	18.61	
5		workday, timecount, dayofweek, power_lag24, elevation, azimuth, radiation_observation	10	(15.55)	(23.98)	(33.41)	(48.61)	(13.59)	(18.18)	IV-F1
6			30	(13.57)	(22.13)	(34.52)	(48.57)	(13.9)	(18.62)	
7		workday, timecount, dayofweek, power_lag24, elevation, azimuth, radiation_forecast	10	25.91	38.01	**45.12**	**64.38**	**13.65**	**18.23**	IV-F2
8			30	**25.58**	**37.79**	48.58	66.53	14.0	18.7	
9		Percentage improvement of **best** RFR over LAG		-32.9 %	-23.9 %	-22.9 %	-16.0 %	-56.7 %	-47.9 %	

available for a model at the beginning of day d is hence

$$\text{train}_{l,d} = \{p_{l,\delta,i} \forall \delta, i \mid \delta < d \ \wedge \ \delta \geq d - \text{train}_{\max}\}, \quad (5)$$

$$\text{test}_{l,d} = \{p_{l,\delta,i} \forall \delta, i \mid \delta = d\}. \quad (6)$$

The value of $\text{train}_{\max}$ allows us to constrain how many days of training data we keep.

D. CHOICE OF MACHINE LEARNING MODEL

We have experimented with a variety of machine learning models, including support vector machines (SVMs), XGBoost, neural networks and random forests. For neural networks, we explored shallow and deep models, and searched through a wide range of layer sizes and hyper parameters. However, despite of the high search effort, neural networks and the other models did in general not score better than the random forest regressors. For brevity, we don't quantify these results here but refer to similar experiences for outdoor predictions [21], where random forests outperform other models. Another benefit of random forests is their low number of parameters. We therefore use random forests regressors for our study, and select the implementation from Scikit-learn [30].

In the following, we describe our experiments and the feature sets. We discuss the results and explore sources of errors in the next section.

E. BASIC FEATURES

Table 2 shows the results for the different locations. The second column shows which features are provided to the predictor:

- *power_lag24* is the power harvested in the time slot 24 hours ago.

Since the harvested energy shows correlations with time in Sect. III, we add the following features:

- *workday* encodes if the day is a working day, as 1 or 0,
- *dayofweek* encodes the day of the week, with Monday encoded as 0 and Sunday as 6,
- *timecount* encodes the time slot, 0 to 23.

Column $t_{\max}$ shows the maximum number of training days that were used to build the predictor model. Before we consider their significance systematically in Sect. VI-A, we use $t_{\max}$ with 10 and 30 days as the maximum number of days worth of training data. In this run, we set the number of estimators (trees) to 40, and leave the maximum depth unconstrained, which results in a maximum observed tree depth of 31. A more detailed analysis in Sect. VI will show that much lower numbers for these parameters and hence less complex models perform well, too. Depending on the days available for each location in the data set and $t_{\max}$, each simulation run evaluates between $n = 873$ and 926 days.

With the features *workday*, *timecount*, *dayofweek*, *power_lag24* (rows 1 and 2) the RFR predictor scores similar to the LAG predictor for locations L14 and L17, but is significantly better for L18. Motivated by the dependency of the harvested energy to the solar angles, we further add solar angles *elevation* and *azimuth* as features. Rows 3 and 4 show that the RFR predictor performs now slightly better than the LAG predictor for each location.

F. ADDING WEATHER FEATURES

Fig. 3 reveals a significant dependency of the harvested energy with the Sun's position for some locations. However, the solar angles alone could only slightly improve the prediction as shown above, as it also depends on the intensity of the sunlight, which in turn depends on the weather. Similar to the prediction of outdoor solar power [21], we hence add a weather forecast to the prediction variables.

1) RESULTS WITH RADIATION OBSERVATIONS

To first explore the dependency of the radiation from the Sun, we add the *observed* radiation as feature. These observations are available from the weather station Zürich Fluntern, in about 1 km distance of the building where the energy harvesting data set was collected, through the Federal Office of Meteorology and Climatology (MeteoSwiss) [31]. Note that this would not be feasible in practice, as the values are *observations* and hence not available at model prediction time. The results are shown in rows 5 and 6 of Table 2. For L14, the median STAPE improves considerably and is now as low as 13.57%. Also L17 improves, but less significantly than for L14. The scores for L18 stay almost unaffected.

2) RESULTS WITH WEATHER FORECASTS

Since the radiation observations are not available at prediction time, we need to use radiation *forecasts* instead. We obtained the forecasts based on a COSMO-1 model from MeteoSwiss [31]. Forecasts were computed retroactively by MeteoSwiss for the coordinates of the university building, eight times a day. This means forecasts were produced for every three hours, starting at midnight, so that every observation time slot is covered by several forecasts with different lead times. For our experiments, we selected the forecasts that would be available at or before midnight, to keep causality with our online setting; we hence only use data that would be indeed available at prediction time. The forecasts contain the prediction of the downward shortwave radiation flux at the surface, here shortly described as *radiation_forecast*, which correspond to the value described by the radiation observations from above. (We have also experimented with using the cloud coverage at different levels as input, but leave them out here for brevity.)

The results using the weather forecasts are shown in rows 7 and 8, for different amount of training days. For L14 and L17 we still see a significant improvement of the predictions, although not as good as with the radiation observations. For all locations, RFR now outperforms the baseline considerably. Row 9 shows the percentage improvement of the best available RFR compared to the LAG predictor baseline. For L14, metrics are improved by 32% and 23%, for L17 by 22% and 16%. For L18, metrics are improved by 56% and 47%. This means for instance that the RFR for L18 estimates the total energy harvested within a day with an average error of 13,65% compared to the moving average from (1). Fig. 4 shows example predictions, together with the daily values for STAPE and SMAPE.

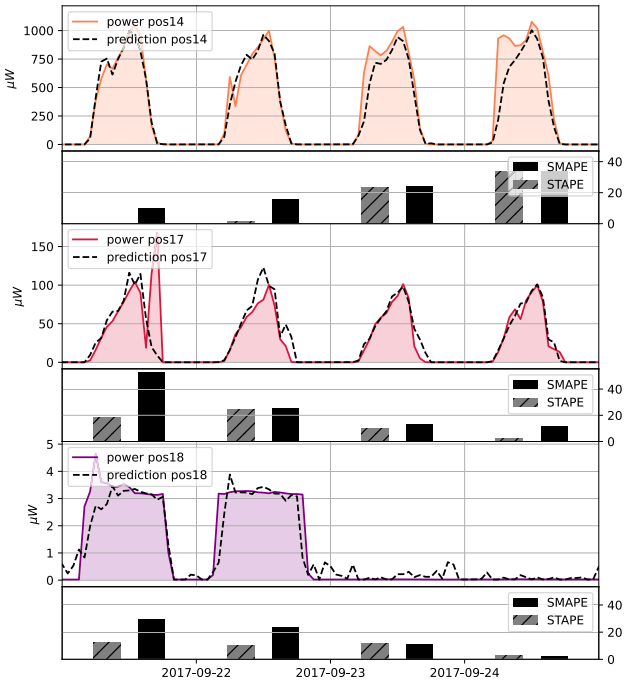

FIGURE 4. Sample predictions for four consecutive days for the different locations, using the predictors from row 8 of Table 2. The bar plots show SMAPE and STAPE for each day.

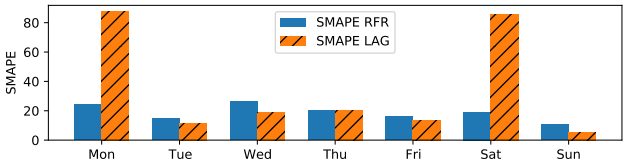

FIGURE 5. Mean SMAPE for LAG and RFR predictor for position L18 for each weekday.

V. ANALYSIS OF PREDICTION ERRORS

In Table 2 we see that the three locations score differently for the various prediction variables. In the following, we discuss these differences, which mainly depend on different sources for errors affecting the locations.

A. EFFECT OF WEEKENDS

For L18, artificial lighting is dominant. The energy intake follows the schedule of the lights in the hallway, which are in turn correlated with the working hours during a day. For this location, we found a significant dependency of the error metrics on the weekday, as shown in Fig. 5. This explains the large error of the LAG predictor, since it takes data from a Friday (a working day) as input for Saturday (where lights are off) and hence overestimates energy for Saturday. Vice versa, it takes the time slots from Sunday and uses them for forecasting the ones for Monday, hence underestimating the energy intake.

In contrast, the RFR predictor scores much more similar over the weekdays. Compared with the LAG predictor, RFR learns the effect of the weekend. One could of course build a LAG-predictor that is aware of working days, but this would require this insight, creating the predictor, and enabling it for locations with that characteristic; all manual steps that are to be avoided once a large number of devices should be deployed. Also, if working patterns change, this manually created predictor would be obsolete, while the online RFR predictor would adapt.

B. ACCURACY OF WEATHER FORECASTS

For L14, the SMAPE when using the observed radiation from the remote weather station is with 22.13% for $t_{\max} = 30$ quite low (Table 2, row 6). With the use of the actually available radiation forecast, this error increases to 37.79% (row 8), which implies that the prediction error is mainly due to the inaccuracy of the weather forecast, i.e., the difference between forecasted and observed radiation. We can confirm this by considering the correlation of weather forecast errors and prediction errors. To distinguish between over- and underestimation, we calculate the scaled percentage

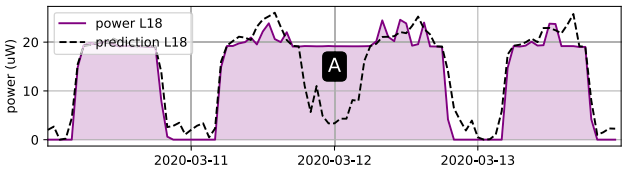

FIGURE 6. Power trace at location L18, located in the hallway. During one night (A), the light was left on. The prediction does not capture this.

error as $\text{SPE}(l, d, i) = 100 \cdot (p_{l,d,i} - \hat{p}_{l,d,i})/\tilde{p}_{l,d}$. We then calculate the weather forecast error as radiation_observation−radiation_forecast. For L14, the correlation of SPE and weather forecast error is 0.44, for L17 it is 0.35. For L18, where natural sunlight has little influence, the correlation is only 0.03. This further indicates that much of the difference between the prediction score with radiation observation and prediction in L14 can be explained by inaccuracies of the weather forecast.

C. UNPREDICTABLE ENVIRONMENT CHANGES

One source of unpredictability is the artificial lighting. Usually, artificial lights are correlated with working hours and predictable. In some cases, however, we observe that that the energy harvested during night is larger than usual, like in Fig. 6. Here we see the harvested power of L18 in the dark hallway follows the working hours, but during one night we see that the power does not reduce, probably because someone left the light on.

For L17, even the predictor with access to the radiation observations (rows 5 and 6 in Table 2) shows a significantly higher error than for location L14. A potential explanation for this can be found in Draskovic and Thiele [27], who mention that there is a curtain that "sometimes partly obstructs the window" at location L17. The curtain is probably controlled manually by people in the room, and as such impossible to predict. This is an example of an inherent source of unpredictability that significantly affects the harvesting prediction.

VI. COMPLEXITY OF THE PREDICTION MODEL

A. AMOUNT OF TRAINING DATA

In Table 2 we only show results for $t_{\max} = 10$ and 30 days for brevity. To study more values for $t_{\max}$, Fig. 7 shows the mean SMAPE for different values of $t_{\max}$, using the feature set that includes the radiation forecast. The dashed lines show the mean SMAPE of the LAG predictor, which uses one day of training data. For L14, RFR is better than LAG already with a single training day, and improves with further training days. It reaches the best results with around 20 days of training data. For L17, RFR follows the same trend as L14, but does not benefit from more than 7 days of training data. For L18, RFR is initially worse than the LAG predictor but increases once it receives 7 days of training data, indicating that learning about the weekend (ref. Sect. V-A) is most critical. For both L14 and L17 we observe that keeping more training data does not necessarily help. Seasonality and non-stationarity in the data is one explanation for this. This strategy is also called *managed forgetting* [32].

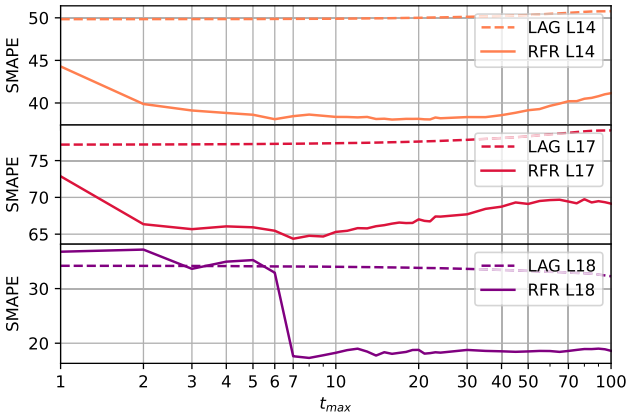

FIGURE 7. Mean SMAPE over the amount of training days.

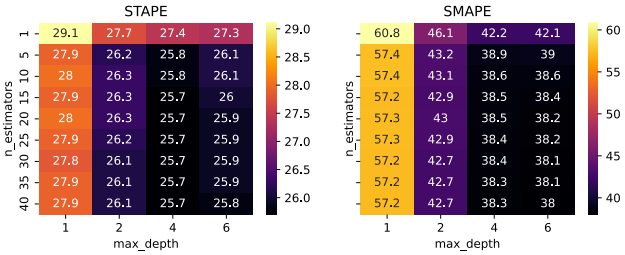

FIGURE 8. STAPE and SMAPE for L14 for different depths and number of estimators.

B. COMPLEXITY OF RANDOM FORESTS

Although we are not primarily concerned with the computational complexity of the predictors, we are interested in the most relevant parameter settings for the random forests. Fig. 8 shows the results of the prediction for location L14 with $t_{\max} = 10$ training days and different values for the the number of estimators (trees) and their maximum depth. Already a depth of 2 delivers results close to the best values, and for both STAPE and SMAPE a depth of 4 and 5 estimators seem like a good tradeoff between achieved accuracy and model complexity. This means, that relatively simple machine learning models in terms of model complexity and required training data are sufficient to achieve the scores presented above.

VII. CONCLUSION

We have not addressed computational effort in this study, as we are initially interested in the accuracy of machine-learning-based methods and their application in principle. Even though machine learning approaches become increasingly feasible on embedded devices [2], they are certainly more resource-intensive than the LAG predictor, which only has to store a value for each of the 24 time slots.

The feasibility of using RFR with weather forecasts will therefore depend on the overall energy budget of the device, and how significant a more efficient planning algorithm through more accurate intake predictions turns out to be. For future scenarios where wireless embedded sensors also execute other, more computationally demanding machine learning tasks, such tradeoffs may be significantly different than for the typical operations of today's devices, and hence justify approaches like the one presented.

The data set of Sigrist et al. [17] reveals the large variance of the energy harvested by indoor photovoltaic cells, both in amount and the patterns in which they depend on natural and artificial lighting and other factors of their environment. We have shown that a machine learning model in form of a random forest regressor (RFR) can outperform a persistence-based predictor for all locations when taking the weather forecast into account, as shown in row 9 of Table 2. We have also presented an analysis to explain relevant sources of errors.

The potential benefit of machine-learning-based approaches is their better accuracy while working autonomously. The proposed prediction models require no manual adjustment or oversight and the same setup can capture quite diverse locations, from a dark hallway to a sunny place near a window. In addition, they are operational with little training data, as discussed in the previous section. For a setting with many embedded, wireless devices these facts are crucial and promising, as sensor devices can indeed learn autonomously. Overall, our study reveals that in future scenarios where prediction of energy intake is a critical component, model complexity or amount of training data are not the main problems; it is rather the availability of accurate weather forecasts and that some locations are subject to unpredictable obstructions.

REFERENCES

[1] R. David, J. Duke, A. Jain, V. J. Reddi, N. Jeffries, J. Li, N. Kreeger, I. Nappier, M. Natraj, T. Wang, and P. Warden, "TensorFlow lite micro: Embedded machine learning for TinyML systems," in *Proc. Mach. Learn. Syst.*, vol. 3, A. Smola, A. Dimakis, and I. Stoica, Eds. Mar. 2021, pp. 800–811. [Online]. Available: https://proceedings.mlsys.org/paper/2021/file/d2ddea18f00665ce8623e36bd4e3c7c5-Paper.pdf

[2] F. Kuppers, J. Albers, and A. Haselhoff, "Random forest on an embedded device for real-time machine state classification," in *Proc. 27th Eur. Signal Process. Conf. (EUSIPCO)*, 2019, pp. 1–5.

[3] H. Ren, D. Anicic, and T. A. Runkler, "TinyOL: TinyML with online-learning on microcontrollers," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, 2021, pp. 1–8.

[4] (2022). *Edge Impulse*. Accessed: Sep. 2022. [Online]. Available: https://docs.edgeimpulse.com/docs/edge-impulse-studio/learning-blocks/object-detection/fomo-object-detection-for-constrained-devices

[5] D. Palma and R. Birkeland, "Enabling the Internet of Arctic Things with freely-drifting small-satellite swarms," *IEEE Access*, vol. 6, pp. 71435–71443, 2018.

[6] A. Montanari, M. Sharma, D. Jenkus, M. Alloulah, L. Qendro, and F. Kawsar, "ePerceptive: Energy reactive embedded intelligence for batteryless sensors," in *Proc. 18th Conf. Embedded Networked Sensor Syst.*, 2020, pp. 382–394.

[7] C. Renner, S. Unterschutz, V. Turau, and K. Römer, "Perpetual data collection with energy-harvesting sensor networks," *ACM Trans. Sensor Netw.*, vol. 11, no. 1, pp. 1–45, Nov. 2014.

[8] A. Cammarano, C. Petrioli, and D. Spenza, "Online energy harvesting prediction in environmentally powered wireless sensor networks," *IEEE Sensors J.*, vol. 16, no. 17, pp. 6793–6804, Sep. 2016.

[9] J. R. Piorno, C. Bergonzini, D. Atienza, and T. S. Rosing, "Prediction and management in energy harvested wireless sensor nodes," in *Proc. 1st Int. Conf. Wireless Commun., Veh. Technol., Inf. Theory Aerosp. Electron. Syst. Technol.*, San Diego, CA, USA, Nov. 2009, pp. 6–10.

[10] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava, "Power management in energy harvesting sensor networks," *ACM Trans. Embedded Comput. Syst.*, vol. 6, no. 4, p. 32, Sep. 2007.

[11] A. Sharma and A. Kakkar, "A review on solar forecasting and power management approaches for energy-harvesting wireless sensor networks," *Int. J. Commun. Syst.*, vol. 33, no. 8, p. e4366, 2020.

[12] N. Stricker and L. Thiele, "Analysing and improving robustness of predictive energy harvesting systems," in *Proc. 8th Int. Workshop Energy Harvesting Energy-Neutral Sens. Syst.*, 2020, pp. 64–70.

[13] I. Mathews, S. N. Kantareddy, T. Buonassisi, and I. M. Peters, "Technology and market perspective for indoor photovoltaic cells," *Joule*, vol. 3, no. 6, pp. 1415–1426, 2019.

[14] C. Voyant, G. Notton, S. Kalogirou, M.-L. Nivet, C. Paoli, F. Motte, and A. Fouilloy, "Machine learning methods for solar radiation forecasting: A review," *Renew. Energy*, vol. 105, pp. 569–582, May 2017.

[15] V. Cionca, A. McGibney, and S. Rea, "MAllEC: Fast and optimal scheduling of energy consumption for energy harvesting devices," *IEEE Internet Things J.*, vol. 5, no. 6, pp. 5132–5140, Jan. 2019.

[16] L. Hanschke and C. Renner, "EmRep: Energy management relying on state-of-charge extrema prediction," *IET Comput. Digit. Techn.*, vol. 16, no. 4, pp. 91–105, 2022.

[17] L. Sigrist, A. Gomez, and L. Thiele, "Dataset: Tracing indoor solar harvesting," in *Proc. 2nd Workshop Data Acquisition Anal.*, 2019, pp. 47–50.

[18] A. Cammarano, C. Petrioli, and D. Spenza, "Pro-Energy: A novel energy prediction model for solar and wind energy-harvesting wireless sensor networks," in *Proc. IEEE 9th Int. Conf. Mobile Ad-Hoc Sensor Syst. (MASS)*, Mar. 2012, pp. 75–83.

[19] R. Ahmed, B. Buchli, S. Draskovic, L. Sigrist, P. Kumar, and L. Thiele, "Optimal power management with guaranteed minimum energy utilization for solar energy harvesting systems," *ACM Trans. Embedded Comput. Syst. (TECS)*, vol. 18, no. 4, p. 30, Jan. 2019.

[20] B. Buchli, F. Sutton, J. Beutel, and L. Thiele, "Dynamic power management for long-term energy neutral operation of solar energy harvesting systems," in *Proc. 12th ACM Conf. Embedded Netw. Sensor Syst.*, 2014, pp. 31–45.

[21] F. A. Kraemer, D. Palma, A. E. Braten, and D. Ammar, "Operationalizing solar energy predictions for sustainable, autonomous IoT device management," *IEEE Internet Things J.*, vol. 7, no. 12, pp. 11803–11814, Dec. 2020.

[22] C. Berner, "Prediction models for indoor solar energy harvesting," Semester thesis, Comput. Eng. Netw. Lab., ETH Zürich, Zürich, Switzerland, 2019.

[23] N. Stricker and L. Thiele, "Accurate onboard predictions for indoor energy harvesting using random forests," in *Proc. 11th Medit. Conf. Embedded Comput. (MECO)*, 2022, pp. 1–6.

[24] D. Kim, J. Ahn, J. Shin, and H. Cha, "Ray tracing-based light energy prediction for indoor batteryless sensors," *Proc. ACM Interact., Mobile, Wearable Ubiquitous Technol.*, vol. 5, no. 1, pp. 1–27, 2021.

[25] A. Schneider, "Analytical model for indoor solar energy harvesting," Bachelor Thesis, Comput. Eng. Netw. Lab., ETH Zürich, Zürich, Switzerland, Jul. 2021.

[26] X. Ma, S. Bader, and B. Oelmann, "Estimating harvestable energy in time-varying indoor light conditions," in *Proc. 8th Int. Workshop Energy Harvesting Energy-Neutral Sens. Syst.*, 2020, pp. 71–76.

[27] S. Draskovic and L. Thiele, "Optimal power management for energy harvesting systems with a backup power source," in *Proc. 10th Medit. Conf. Embedded Comput. (MECO)*, 2021, pp. 1–9.

[28] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *Int. J. Forecasting*, vol. 22, no. 4, pp. 679–688, Oct. 2006.

[29] M. I. Ali, B. M. Al-Hashimi, J. Recas, and D. Atienza, "Evaluation and design exploration of solar harvested-energy prediction algorithm," in *Proc. Design, Automat. Test Eur. Conf. Exhib.*, 2010, pp. 142–147.

[30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, no. 10, pp. 2825–2830, Jul. 2017.

[31] (Sep. 2022). *Federal Office of Meteorology and Climatology Meteoswiss*. [Online]. Available: https://www.meteoswiss.admin.ch
[32] C. Jilek, Y. Runge, C. Niederee, H. Maus, T. Tempel, A. Dengel, and C. Frings, "Managed forgetting to support information management and knowledge work," *KI-Künstliche Intelligenz*, vol. 33, no. 1, pp. 45–55, Nov. 2018.

FRANK ALEXANDER KRAEMER (Member, IEEE) received the M.Sc. degree in information technology and the Dipl.-Ing. degree in electrical engineering from the University of Stuttgart, Stuttgart, Germany, in 2003, and the Ph.D. degree in model-driven systems development from the Department of Telematics, Norwegian University of Science and Technology (NTNU), Trondheim, Norway, in 2008. As the Technology Manager, he co-founded a startup for the IoT software. He is currently an Associate Professor with the Department of Information Security and Communication Technology, NTNU. His current research interests include the Internet of Things architectures and application development, embedded and autonomous sensor systems, and the application of statistical methods and machine learning in constrained settings.

HAFIZ AREEB ASAD received the B.Sc. degree in computer science from the National University of Computer and Emerging Sciences, Islamabad, Pakistan, in 2017, and the M.Sc. degree in computer science from Uppsala University, Sweden, in 2020. He is currently pursuing the Ph.D. degree in information security and communications technology with the Norwegian University of Science and Technology, Trondheim, Norway. His current research interests include autonomous, cognitive, and the battery-less IoT. He was a recipient of the Swedish Institute (SI) Scholarship for Global Professionals.

KERSTIN BACH (Member, IEEE) received the M.Sc. degree in information management and technology and the Ph.D. degree from the University of Hildesheim, Germany, in 2007 and 2012, respectively. She was a Research Engineer with Verdande Technology, from 2013 to 2014. She is currently a Professor of artificial intelligence with the Department of Computer Science, Norwegian University of Science and Technology (NTNU). She is also the Deputy Head of the Data and Artificial Intelligence Group, NTNU, the Program Manager of the Norwegian Research Center for AI Innovation (NorwAI), and associated with the Norwegian Open AI Laboratory. Her research interests include artificial intelligence methods for developing intelligence decision support systems involving both domain experts and end-users to create explainable, interpretable, and trustworthy AI systems. In particular, she works on data-driven and knowledge-intensive case-based reasoning.

BERND-CHRISTIAN RENNER received the Ph.D. degree from the Hamburg University of Technology (TUHH), in 2013. He was a Postdoctoral Researcher with Beck University, from 2012 to 2016. He was an Assistant Professor with TUHH, from 2016 to 2020, and an Associate Professor with the University of Koblenz-Landau, from 2020 to 2022. He is currently a Full Professor and the Head of the Institute of Autonomous Cyber-Physical Systems (aCPS), TUHH. His research interests include the applications of networked embedded, cyber-physical, sensing systems, particularly aiming at energy harvesting and networking aspects, and mobile underwater swarm robotics, with a special focus on underwater acoustic communication and localization. In these domains, he covers theoretical, algorithmic, and practical aspects ranging from computer science over applied mathematics to electronics.