

Improved day ahead heating demand forecasting by online correction methods

Felix Bünning^{1,2}, Andrew Bollinger¹, Philipp Heer¹, Roy S. Smith², John Lygeros²

¹*Empa, Urban Energy Systems Laboratory, Überlandstrasse 129, 8600 Dübendorf, Switzerland*

²*Automatic Control Laboratory, Department of Electrical Engineering and Information Technology, ETH Zürich, Switzerland*

Abstract

To reduce the heating and cooling energy demand of buildings and districts novel control strategies are constantly being developed that require information on the future demand of the controlled entity. Demand forecasting is commonly done with deterministic white box models or fitted grey-box models, however, recently more and more data and machine learning based approaches are being developed. All approaches have weaknesses: white-box models require major modelling effort, grey-box approaches are limited by their model or parameter complexity and machine learning is dependent on hyperparameters, some of which are randomly chosen, and therefore considered unreliable. Here we develop a forecasting approach based on Artificial Neural Networks (ANN) and introduce error correction methods based on online learning and the learned autocorrelation of the forecasting error. We compare the approach to other regression based and grey-box methods in a real case study of a small-scale district energy system with mixed use and unknown lower-level control. We show that the proposed method outperforms the other forecasting methods in terms of average error and coefficient of determination. We further demonstrate that in our case study the error correction methods significantly reduce variance in ANN performance created by randomly initialized parameters in the networks.

Keywords:

Artificial Neural Networks, demand forecasting, online learning, error autocorrelation, building energy demand, district energy demand

1. Introduction

Building and district heating and cooling energy amounts to 70% of the total energy demand in households in Switzerland (BFE, 2014) and to more than 44% of the total primary energy consumption worldwide (Ölz et al., 2007). To reduce this demand, novel control strategies are constantly being developed (see for example Široký et al. (2011); Yudong Ma et al. (2012); Oldewurtel et al. (2010); Sturzenegger et al. (2016); Bünning et al. (2017), or Hameed Shaikh et al. (2014) and the references therein). Many of them rely on information about the future evolution of the system, especially on the future energy demand.

This has motivated an increasing interest in demand forecasting for the building energy domain, see for example the survey of different methods and models under Mat Daut et al. (2017); Zhao and Magoulès (2012); Harish and Kumar (2016); Wang and Srinivasan (2017); Fouquier et al. (2013); Suganthi and Samuel

(2012). Roughly speaking one can distinguish the approaches in the literature into white-box and grey-box approaches (Reinhart and Davila, 2016; Coakley et al., 2014; Kramer et al., 2012) (also referred to as first principle modelling and reverse modelling respectively) and black-box methods (Amasyali and El-Gohary, 2018; Deb et al., 2017) (also referred to as regression-based methods, Machine Learning (ML) or Artificial Intelligence (AI) methods).

Conventionally, demand forecasting is done with white-box models or grey-box models. In white-box models, the system is described in terms of differential equations that are motivated by the physical nature of the system and parameters that are set based on material or geometric properties. The resulting models are transparent and show good accuracy. However, the modelling effort is high and therefore expensive.

In grey-box modelling, differential equations are also used, but the parameters are fitted with the help of measured demand data. This significantly reduces the

modelling effort. However, as the form of the equations are usually fixed, for example as interconnected resistor-capacitor networks (Sturzenegger et al., 2014; Dimitriou et al., 2015; Ramallo-González et al., 2013), certain aspects of the system are often not captured. For example heating and ventilation systems and their control, occupant behaviour and other phenomena that result in time-dependent demand patterns require additional effort.

Black-box methods also use measured data to fit a model. However, the structure of the model is not inspired by the physical nature of the modelled systems. This allows for higher flexibility in the learned system behaviour on the one hand, on the other it does in principle allow non-physical behaviour of the model.

Methods such as Artificial Neural Networks (ANN) have been proven to perform well in a variety of tasks, such as gaming (Silver et al., 2017), computer vision (Wojna et al., 2018) and robotics (Tobin et al., 2017). Several studies have also demonstrated that ANN perform well on the task of doing sub-hourly short-term demand forecasts in the building domain, (Bagnasco et al., 2015; Kamaev et al., 2012; Escrivá-Escrivá et al., 2011; Paudel et al., 2014; Chae et al., 2016). However, there are to the authors' knowledge only two studies (Geysen et al., 2018; Idowu et al., 2016) on the forecast of larger thermal energy systems, such as multiple buildings connected in a district heating system. Moreover, none of these studies analyse previous forecasting errors to correct the forecast during the online phase of the demand prediction; Geysen et al. (2018) makes use of tracking the best performing forecasting model, but does not use the history of forecasting errors directly. Oldewurtel et al. (2012); Sturzenegger et al. (2016) and Darivianakis et al. (2017), who do predictive control on building and district energy systems, use Kalman filters to correct model inputs but not the forecast itself. Furthermore, there is a lack of studies that compare data-driven methods such as ANN to conventional grey-box and white-box models on the same system, with the exception of Idowu et al. (2016) who investigate whole-day forecasts.

The sensitivity of ANN to their parameters, of which some are randomly chosen and the implications for control are emphasized by Recht (2018); Arulkumaran et al. (2017); Bengio (2012); Jamieson and Talwalkar (2016). This is an issue that also needs to be addressed in building energy and district energy control.

1.1. Contribution

We make the following contributions with this work: A prediction model based on ANN to make 24 hour

ahead forecasts of the heating demand for complex buildings and districts with individual uses, individual occupant patterns and unknown control systems, is developed. Two simple forecast correction methods, based on the error-autocorrelation and online learning, which make use of the history of previous forecasting errors during the online phase, are introduced. The methods increase prediction accuracy and significantly reduce the ANN dependence on random initialization parameters. These methods are verified in a case study based on data from the NEST building, which comprises independent modular control zones that allow it to mimic a district network. The prediction results are compared to other machine learning (or regression-based) approaches and to conventional R-C building models. It is demonstrated that online error corrections lead to ANN models that not only outperform conventional grey-box models and other machine learning approaches, but also lead to an increased robustness towards random initialization parameters.

1.2. Structure

The remaining article is structured as follows: In Section 2 our ANN based methodology and the forecast corrections that are based on the error-autocorrelation and online learning are described. In Section 3 the case study is defined. The energy system under consideration is introduced, and the structure of the applied ANN and the grey-box and black-box methods that are used for benchmarking are defined. In Section 4 the results regarding prediction accuracy and variance are discussed, and the ANN methods are compared to the benchmarking methods. In Section 5 the article is concluded.

2. Methodology

2.1. Forecasting task

In this study the following forecasting task is considered: A heating demand forecast for a building or a district is made at midnight for the next 24 hours, sampled every 15 minutes. The training and validation data are assumed to be sampled at the same rate. Furthermore, perfect knowledge of the ambient temperature for the forecasting period is assumed. This is done as it avoids uncertainty in the inputs when comparing the different methods.

2.2. Forecast correction methods

The forecast is made with a feed-forward ANN, which will be described in detail in Section 3. To improve the forecasting accuracy, we introduce two meth-



Figure 2: NEST building at Empa in Switzerland, Copyright: Zooey Braun - Stuttgart

It should be noted, that in line 10 a second forecast is made at the end of the day, after the ANN has been re-trained on the inputs and demand realization of the day. The error of this second forecast is then added to the error database. This is done because next day's error estimation needs to be performed for the retrained ANN. If this is not done, the error for the next day is overestimated. The results for this combined method are shown in Section 4.

Algorithm 2

```

1: procedure COMBINED CORRECTION
2:   for  $T$  in runtime do
3:     for  $t$  in  $[1, 96]$  do
4:        $\text{forecast}[T,t] := \text{ANN.predict}(\text{inputs}[T,t])$ 
5:        $\text{error}[T,t] := \text{AvError}[T-1] * R_{ee}(t,E)$ 
6:        $\text{correctedFcast}[T,t] := \text{forecast}[T,t]$ 
         +  $\text{error}[T,t]$ 
7:       wait until day passes
8:        $\text{ANN} := \text{ANN.train}(\text{inputs}[T], \text{realization}[T])$ 
9:       for  $t$  in  $[1, 96]$  do
10:         $\text{forecast2}[T,t] := \text{ANN.predict}(\text{inputs}[T,t])$ 
11:         $\text{measuredError}[T] := \text{realization}[T]$ 
         -  $\text{forecast2}[T]$ 
12:         $E := E.\text{add}(\text{measuredError}[T])$ 

```

3. Case study

3.1. System description

The NEST building at Empa in Switzerland (Richner et al., 2018), shown in Figure 2, is used as a case study. The building contains multiple individual units with different uses that can be added and removed from the building, in addition to office and meeting rooms that are permanently installed. At different times during the period covered by our data, two, three or four of these

individual units were in operation. One of the units was added during the training set (August 2017) and one during the test set (February 2018) of the case study. The configuration resembles a district heating system because different units have different use patterns, some of which resemble those of residential buildings while others those of office spaces. Moreover, all of the individual units are equipped with local controllers, whose details are unknown to the forecasting system. The demand forecast is therefore made for a collection of individual closed-loop energy systems. The units are connected via substations to a central heating system with a supply temperature of 35 °C and a return temperature of 25 °C. The heating system is served by a central heat pump with a maximum thermal capacity of 100 kW. A simplified schematic of the system can be seen in Figure 3. The heating demand that our algorithm is designed to forecast is the energy flow difference between supply line (marked with S) and the return line (marked with R). For more information on the complete system see Lydon et al. (2017).

The building is set-up as a test bed for district heating and cooling and building energy experiments. It integrates approximately 540 sensors, whose data is stored in an SQL database with a temporal granularity of one minute, and approximately 200 controllable actuators. The measurement data used in this study contain 411 days (01.04.2017 to 16.05.2018) and include the measurements for ambient temperature and total heat consumption. The measured heating demand is shown in Figure 4. Figure 5 shows the ambient air temperature measured at the roof top of the building. Though, strictly speaking, weather forecasts should be used in demand forecasting, in this study we use the actual ambient air temperature directly as an input. This is done to avoid the additional uncertainty in the inputs and allow direct comparing the different forecasting methods.

3.2. Artificial Neural Networks

We use feed-forward ANN with the Python package Keras (Chollet, 2018), which is a higher-level API to TensorFlow (Abadi et al., 2016). We refrain from fundamental introduction to the concept here and refer the reader to other sources (for example Shanmuganathan (2016); Basheer and Hajmeer (2000)). To fit the ANN to the training data, we use the optimizer Adam (Yun et al., 2018) with standard parameters. As a loss function, the mean squared error and as activation functions, Rectified Linear Units (ReLU) are used.




Figure 3: NEST heating system schematic

3.2.1. Structure

As we want to forecast the heating demand in a day-ahead fashion, we forecast 96 data points at once. (One value every 15 minutes for one day.) For the structure of the ANN this gives two possibilities: 1) The network is built with 96 outputs and the forecast for the day is performed in one single step. 2) The network is built with one output and the forecast is performed with 96 different input vectors to forecast 96 single output values.

Option 1 has one major advantage: As the heating demand (forecast target) is highly autocorrelated for small time delays, the accuracy of the prediction is very precise for the first few forecast data points at the beginning of the day. This is because the network makes use of the measured heating demands at the end of the previous day. In option 2 this is not possible as, for example, we cannot use the demand of $t - 4$ as an input for the network because for forecasts $t = 5, \dots, 96$ this input has not yet been measured at the time the forecast is made. However, in case of option 1 the network needs a much higher complexity in terms of number of layers and nodes, while the training set is not increased when compared to option 2. Preliminary forecast runs have shown that this results in reduced prediction performance. For option 2, the disadvantage of not be-

ing able to use the last measured heating demands of the previous day as inputs is compensated by the error-autocorrelation based correction method that was introduced in Section 2.2.1. We therefore use option 2 in this study.

3.2.2. Feature selection

Although NEST also has measurements for wind speed and solar radiation, the ambient temperature was selected as the only weather-related feature², because preliminary experiments suggested that the other measurements do not substantially improve forecasting accuracy.

With the assumption that the demand of the system shows daily and weekly patterns, as the units are used for living and as office spaces, two more features were added: the total heat consumption with a time delay of one day (again following the assumption “tomorrow will be like today”) and the total heat consumption with a time delay of one week (motivated by the assumption that weekly routines can be expected). This allows the prediction of the energy demand today based on the demand at the corresponding time yesterday and last week. Finally, the time stamp of all features was implemented as two inputs: hour of the day and weekday/weekend³. Hour of the day was one-hot encoded (Aggarwal, 2018), meaning that instead of one continuous input 24 binary inputs are used. Preliminary experiments suggested that this improves the forecast accuracy.

The resulting network has 28 inputs, of which three are continuous and 25 binary, and one continuous output.

3.2.3. Parameter selection

There is no proven rule on how to best determine the parameters of ANN. These values are commonly set based on heuristics and an iterative approach that involves a substantial amount of trial and error. Generally, a trade-off between overfitting and underfitting of the training data has to be found.

A preliminary sensitivity analysis was conducted to decide on the number of hidden layers, number of nodes per layer, number of training epochs (number of times

²The difference between “feature” and “input” is the level of detail in the description: E.g. “daytime” can be a feature and the corresponding inputs could be “hour of the day, encoded continuously” and “minute of the hour, encoded continuously”.

³Weekday/weekend was chosen instead of working day/non-working day as the difference is marginal and the implementation effort substantially lower.




Figure 4: Measured total heating load of the NEST building




Figure 5: Measured ambient temperature at the roof of NEST building

Table 1: ANN parameters

Parameter	Value
Number of hidden layers	2
Number of nodes per layer	8
Number of Epochs	10
Input scaling method	Linear [-1,1]

the network’s weights are updated based on each individual sample in the training set), and on the input scaling method. The results for the suggested optimal parameter values are given in Table 1. They were used in the case study.

3.3. Forecasting methods for benchmarking

To assess the performance of the ANN with the introduced correction methods for forecasting the thermal demand of NEST, the prediction accuracy is compared to other state of the art prediction methods: These are grey-box models in the form of resistor-capacitor building models, and black-box models in the form of other regression-based or machine-learning based approaches.

3.3.1. Grey-box method (R-C models)

For the grey-box approach, conventional resistor-capacitor models are used. Four different topologies were considered as shown in Figure 6, inspired by the models used in De Coninck et al. (2015). In this approach, the thermal capacity of wall materials, floor materials and air volumes is represented by capacitors, whereas the thermal resistance of walls and floors is represented by resistors. The parameters of all resistors and capacitors are estimated to fit a heating demand measurement curve. As the R-C models are purely thermal models, the ambient temperature is the only used input.

The models were implemented in the modelling language Modelica (Mattsson et al., 1998) and simulated in Dymola (Brück et al., 2002). To estimate the parameters, a CMA-ES (Hansen et al., 2003) optimizer was used in Python.

In models 1R1C and 2R2C, the parameters were estimated such that the mean squared error between $\dot{Q}$ and the target heating demand is minimized. In model 1R1C, $\dot{Q}$ was set equal to the heat flux through R_{wall} and in model 2R2C $\dot{Q}$ was set equal to the heat flux through R_{wall2} , which keeps the temperature of C_{zone} at a constant 20 °C.

As an additional capacitor for the building core C_{core} was added for models 4R3C and 5R3C, the mapping of




Figure 6: R-C building models

the target heating demand is more complex. The optimization was set up to minimize the mean squared error between the target heating and $\dot{Q}_{total}$,

$$\dot{Q}_{total} = \dot{Q}_1 + \dot{Q}_2, \quad (3)$$

which is the sum of the heat distributed via air conditioning to C_{zone} ($\dot{Q}_1$) and the heat distributed via floor heating to C_{core} ($\dot{Q}_2$). Moreover, as the direction of heat loss is not unique in these models (losses can go from C_{zone} to C_{core} and/or to T_{ground}), a P-controller was set-up to keep the temperature of C_{zone} at a constant $T_{set} = 20^\circ\text{C}$:

$$\dot{Q}_{total} = k_p \times (T_{set} - T_{zone}). \quad (4)$$

Additionally a fraction coefficient c_f was introduced to distribute the heat between air conditioning and floor heating, such that

$$\dot{Q}_1 = c_f \times \dot{Q}_{total} \quad (5)$$

and

$$\dot{Q}_2 = (1 - c_f) \times \dot{Q}_{total}. \quad (6)$$

The values of k_p and c_f are also determined in the fitting process.

3.3.2. Black-box method (regression-based methods)

For the black-box methods that are used for benchmarking, we use Python's Scikit-learn package (Pedregosa et al., 2011). The package features a variety of different regression or machine learning methods. In the case study, the following methods are used:

1. Least squares Linear Regression
2. Support Vector Machine
3. Huber Regressor
4. Orthogonal Matching Pursuit
5. SGD Regressor
6. Decision Tree Regression
7. Random Forest

The data sets used for these models were identical to that used for the ANN method, including the scaling of inputs and outputs and the one-hot encoding of categorical features. All methods were used with default parameters.

4. Results and discussion

For the case study, the measurement data from NEST is divided into a training set and a testing set. The training set consists of 70% of the whole data set, which corresponds to 287 days (27552 data points). The testing set consists of 124 days (11904 data points).

The data was post-processed with the help of the pandas package (McKinney, 2011) in Python 3. Missing data points were first linearly interpolated or extrapolated from neighbouring data points. This excludes data points that have already been handled by the NEST database (e.g. the constant values around September 2017 in Figure 5). The data set was then re-sampled to a 15-minute interval using the mean value of all relevant data points.

To benchmark the prediction performance, the coefficient of determination,

$$R^2 = 1 - \frac{\sum_{i \in N} (y_i - f(x_i))^2}{\sum_{i \in N} (y_i - \bar{y}_N)^2}, \quad (7)$$

is used. It becomes *zero* if the forecast $f(x_i)$ is as good as taking the average $\bar{y}_N$ of the data in the considered set N as a forecast, and becomes *one* if the forecast exactly matches the validation data y_i . In principle, R^2 may also become negative if the forecast is worse than taking the average, though this was not observed with any of the methods tested here. In the following, if more than one network is used, R^2 describes the average of all R^2 achieved, unless otherwise stated.

We train and forecast with 100 individual ANN as the initial input weights of each node in the network are initialized with a random value. This randomness leads to different prediction performance for each network as the training process involves the solution of a non-convex optimisation problem that generally converges to a local minimum. By simulating 100 networks, a more confident statement about R^2 can be made and the quartiles can be studied.

4.1. Prediction accuracy, quartiles and improvements through forecast corrections

Figure 7 shows the box plot of the coefficient of determination in the test set of 100 different networks for the uncorrected ANN, ANN with error-autocorrelation correction, ANN with online learning and ANN with error-autocorrelation + online learning. It can be seen that the quartiles regarding prediction accuracy amongst different network instances are significantly reduced when correction methods are used. While the interquartile range is 0.038 in the uncorrected case, it reduces to 0.023 in the case with error-autocorrelation correction,




Figure 7: Coefficient of determination for ANN and forecast correction methods for 100 individual networks (with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median)

0.009 in the case with online learning and 0.008 in the case where both correction methods are used. With both corrections in place all R^2 lie between 0.872 and 0.898. The average achieved R^2 are 0.818, 0.860, 0.878 and 0.885 subsequently.

These results suggest that the corrections introduced in Section 2 can to a large extent alleviate a main disadvantage of ANN, namely the dependence of the prediction performance on randomly initialized node input weights. Moreover, the average prediction performance is significantly improved.

4.2. Prediction error analysis

To further evaluate the impact of the correction terms introduced in Section 2 we analysed the statistics of the forecast error with and without the corrections.⁴

Figure 8 shows the autocorrelation of the error for one instance of an uncorrected ANN forecast (with a relatively high R^2 of 0.841) and the remaining error after correction for one instance of a corrected ANN with error-autocorrelation and online learning (with R^2 of 0.887). It can be seen that the signal for the corrected ANN is less correlated than the uncorrected signal. However, the residual is still not “white”, as the autocorrelation is still substantial for many values of the lag. We conjecture that this residual autocorrelation is

⁴The absolute error is used here instead of the relative error, as the latter is undefined for many of the data points where demand is close to zero.




Figure 8: Residual autocorrelation for uncorrected ANN and ANN with error-autocorrelation correction + online learning




Figure 9: Forecasting error histogram for uncorrected ANN and ANN with error-autocorrelation correction + online learning (Standard deviations are indicated by the vertical lines)

an intrinsic artefact of the day-ahead forecast: Disturbances that occur during the forecasting period will introduce a systematic forecasting error for the rest of the period, as there is no chance to correct the forecast until the next forecasting period begins.⁵ Moreover, one can see that there is no peak at $t = 96$, i.e. after one day. This suggests that the network captured all regular daily time-dependencies of the demand and has no systematic error regarding this aspect.

Figure 9 compares the forecasting error histogram of one instance of ANN without correction to the one using both error correction methods.⁶ The plot is cut at a frequency of 10, because by far most errors are 0 for both cases. The interquartile range reduces from 0.054 in the uncorrected case to 0.033 in the corrected case. Furthermore, the median improves from -0.025 to 0, which shows that the bias in the forecast is effectively eliminated by the forecast corrections.

4.3. Forecasting trajectory examples

Figure 10 shows examples of the forecast heating demand trajectory at different temporal resolutions for an ANN that includes the error-autocorrelation and online learning corrections. The figure also shows the empirical confidence bounds for the forecast computed with

⁵To use the example of an opened window again: If the window is opened at day T-1, we can use the resulting forecast error in day T-1 in our forecast corrections for day T. However, if the window is opened at day T, for example at $t=5$, this will introduce a systematic error in the forecast for $t=[6...96]$. This error will show up in the autocorrelation plot and is impossible to avoid in this forecasting setting.

⁶The number of bins is set according to the Freedman-Diaconis rule.

a posteriori analysis of the absolute error distribution of the testing set. The dark-grey background depicts a confidence bound of 68% (one standard deviation) and the light-grey background depicts a confidence bound of 95% (two standard deviations). Such confidence intervals could be useful if one uses the forecast for predictive building energy management based on robust or stochastic model predictive control.

The top graph of Figure 10 shows the prediction and the actual heating demand for the complete testing set. The general trend is captured well. The middle graph shows a more detailed view of five days with a medium heating demand in March. At this scale the capabilities and limitations of the forecast correction methods can be observed. At the beginning of the 18th of March and 19th of March, the forecast signal is corrected towards the true demand, because an error was observed at the end of the previous days. The time of correction is indicated by the vertical grey lines. On the other hand, in the middle of 17th of March the forecast deviates from the actual demand for several hours. However, this error can not be corrected, as the correction is applied only at the beginning of each 24 hour forecasting interval. In the bottom graph, a system shut-down or failure occurs in the afternoon of the 26th of February. The forecast is not corrected until the start of the next day. At the beginning of the 28th of February and 1st of March a correction of the forecast towards the real demand at the beginning of the day can be observed again.






Figure 10: Forecasting examples with varying temporal resolution. The time of correction is indicated by the vertical grey lines.

Table 2: Coefficient of determination for all forecasting methods

Method	R^2
ANN online + autocorr.	0.885
Huber Regressor	0.722
Orthogonal Matching Pursuit	0.765
Least squares Linear Regression	0.770
SGD Regressor	0.772
Support Vector Machine	0.797
Decision Tree Regression	0.844
Random Forest	0.850
Model 1R1C	0.691
Model 2R2C	0.702
Model 4R3C	0.739
Model 5R3C	0.761

4.4. Comparison of different forecasting methods

Table 2 compares the coefficient of determination in the testing set for all forecasting methods. It can be seen that the ANN in all three corrected cases outperform the other forecasting methods. However, the regression based methods were used with standard parameters. With some tuning effort this could possibly be improved.

Decision Tree Regression and Random Forest reach a high R^2 without any additional corrections. It is possible that with appropriate correction methods (similar to the ones introduced in this study) the performance of our corrected ANN can be reached or exceeded. With respect to the correction methods introduced here, we note that online learning is not a viable option for regression trees and random forests, because of the way these are constructed. Corrections based on error-autocorrelation could, however, be used to correct forecasts made by any forecasting method.

The resistor-capacitor based models perform worse than the majority of the regression based methods in this case study. This could be due to the fact that these models can only make use of the ambient temperature as an input, making it impossible to capture time-dependent phenomena, such as occupancy for example. This shortcoming could be improved by an occupancy model. Moreover, these models involve only a small number of parameters, in the form of resistor and capacitor values. Increasing the model complexity further could improve forecasting, as suggested by the fact that the R^2 of the R-C models in our results continue to increase as more parameters are added. The addition of more R-C com-

ponents of course implies increased modelling effort. Finally, R-C models are normally used in this context for modelling building thermal dynamics (Sturzenegger et al., 2014) and not for demand forecasting. This study suggests that they are not equally suitable for the latter.

4.5. Generalization to other buildings

To strengthen the confidence in the correction methods, the case study was extended to three other buildings: An eleven-storey office building at ETH Zürich with measurement data of 36 months, a three-storey laboratory building at Empa Dübendorf with data of 24 months and a three-storey office building at Empa with data of 24 months. The data quality of the latter differs from the other buildings, as the quantization error in the heating demand measurement is at least three times higher than that of the other case studies.

For all cases, the data was divided into 70% for the training set and 30% for the testing set. The ANN used are identical to those in the NEST case studies. Again, 100 instances of ANN were trained and tested for each case in order to compensate for the variance in prediction performance.

Table 3 shows the average R^2 for each building and each level of correction. It can be seen that the error correction methods subsequently increase the coefficient of determination in each case. The relative improvement in case of the office building at Empa is small compared to the other buildings. The fact that the measurement data quality of this building is inferior to the other ones suggests that the methods are sensitive to data quality. A quantitative statement regarding this issue is the topic of current research. Although these results are by no means a theoretical proof, they give confidence that the methods generalize to different types of buildings.

5. Conclusion

Demand forecasting in the domain of district and building energy systems is conventionally done with white-box models. These models require a time consuming and expensive modelling effort. Data-driven approaches such as Artificial Neural Networks make use of increasing data availability in the building energy domain and require little engineering effort. However, their prediction accuracy has high variance and depends on network parameters that are commonly randomly initialized.

We have introduced two simple forecast correction methods that make use of the information in the forecasting error to tackle this problem: One based on

Table 3: R^2 of different modelling approaches for all buildings

Building	ANN uncorrected	ANN auto	ANN online	ANN auto+online
Office ETH	0.886	0.907	0.933	0.936
Laboratory Empa	0.674	0.752	0.793	0.809
Office Empa	0.856	0.858	0.860	0.862

the error-autocorrelation and the other based on online learning.

With the help of a case study of a district energy system with unknown low-level control, we have shown that the methods significantly reduce the quartiles in the prediction performance and also improve prediction accuracy for a sub-hourly day-ahead heat demand forecasting task. Furthermore, we demonstrated that ANN with both correction methods outperform R-C models and regression-based forecasting methods in the case study.

We are currently investigating the use of the forecasting approach in a variety of Model Predictive Control and Data Predictive Control settings for building and district energy systems. We intend to use the forecasts as an input for control schemes to provide thermal reserves in low temperature district heating and cooling networks or secondary frequency control in electrical grids based on demand-side management in buildings.

6. Acknowledgements

We would like to thank Kristina Orehounig and Viktor Dorer for their valuable help and support. We are also grateful to Ahmed Aboudonia, Annika Eichler, Benjamin Flamm, Reto Fricker, Marc Hohmann, Mohammad Khosravi, Ralf Knechtel, Sascha Stoller, Bratislav Svetozarevic, Christoph Waibel and Joseph Warrington for fruitful discussions.

This research project is financially supported by the Swiss Innovation Agency Innosuisse and is part of the Swiss Competence Center for Energy Research SCCER FEEB&D.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X., Brain, G., 2016. TensorFlow: A System for Large-Scale Machine Learning, in: Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI '16), pp. 265–283. URL: <https://tensorflow.org>.

Aggarwal, C.C., 2018. Text Sequence Modeling and Deep Learning, in: Machine Learning for Text. Springer International Publishing, Cham, pp. 305–360. URL: http://link.springer.com/10.1007/978-3-319-73531-3_10, doi:10.1007/978-3-319-73531-3_10.

Amasyali, K., El-Gohary, N.M., 2018. A review of data-driven building energy consumption prediction studies. Renewable and Sustainable Energy Reviews 81, 1192–1205. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364032117306093>, doi:10.1016/j.rser.2017.04.095.

Arulkumaran, K., Deisenroth, M.P., Brundage, M., Bharath, A.A., 2017. Deep reinforcement learning: A brief survey. IEEE Signal Processing Magazine 34, 26–38. URL: <http://www.aaai.org/http://arxiv.org/abs/1709.06560>, doi:10.1109/MSP.2017.2743240, arXiv:1708.05866.

Bagnasco, A., Fresi, F., Saviozzi, M., Silvestro, F., Vinci, A., 2015. Electrical consumption forecasting in hospital facilities: An application case. Energy and Buildings 103, 261–270. URL: <http://dx.doi.org/10.1016/j.enbuild.2015.05.056>, doi:10.1016/j.enbuild.2015.05.056.

Basheer, I., Hajmeer, M., 2000. Artificial neural networks: fundamentals, computing, design, and application. Journal of Microbiological Methods 43, 3–31. URL: <https://www.sciencedirect.com/science/article/pii/S0167701200002013>, doi:10.1016/S0167-7012(00)00201-3.

Bengio, Y., 2012. Practical Recommendations for Gradient-Based Training of Deep Architectures, in: Montavon, G., Orr, G., Müller, K.R. (Eds.), Neural Networks: Tricks of the Trade. 2 ed.. Springer Berlin Heidelberg, pp. 437–478. doi:10.1016/j.ijantimicag.2014.12.030, arXiv:1206.5533.

BFE, 2014. Energieverbrauch in der Schweiz und weltweit. Technical Report 4. Bundesamt für Energie BFE. URL: http://www.bfe.admin.ch/php/modules/publikationen/stream.php?extlang=de&name=de_802079548.pdf, doi:10.1016/j.cortex.2006.01.002.

Brück, D., Elmqvist, H., Erik Mattsson, S., Olsson, H., 2002. Dymola for Multi-Engineering Modeling and Simulation, in: 2nd International Modelica Conference, pp. 1–8. URL: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.217.1782&rep=rep1&type=pdf>.

Bünning, F., Sangi, R., Müller, D., 2017. A Modelica library for the agent-based control of building energy systems. Applied Energy 193, 52–59. URL: <https://www.sciencedirect.com/science/article/pii/S0306261917300612>, doi:10.1016/J.APENERGY.2017.01.053.

Chae, Y.T., Horesh, R., Hwang, Y., Lee, Y.M., 2016. Artificial neural network model for forecasting sub-hourly electricity usage in commercial buildings. Energy and Buildings 111, 184–194. URL: <http://dx.doi.org/10.1016/j.enbuild.2015.11.045>, doi:10.1016/j.enbuild.2015.11.045.

Chollet, F., 2018. Keras: The Python Deep Learning library. Astrophysics Source Code Library URL: <http://adsabs.harvard.edu/abs/2018ascl.soft06022C>.

Coakley, D., Raftery, P., Keane, M., 2014. A review of methods to match building energy simulation models to measured data. Renewable and Sustainable Energy

- Reviews 37, 123–141. URL: <http://dx.doi.org/10.1016/j.rser.2014.05.007>, doi:10.1016/j.rser.2014.05.007, arXiv:j.rser.2014.05.007.
- Darivianakis, G., Georghiou, A., Smith, R.S., Lygeros, J., 2017. The Power of Diversity: Data-Driven Robust Predictive Control for Energy-Efficient Buildings and Districts. *IEEE Transactions on Control Systems Technology* 99, 1–14. URL: <http://ieeexplore.ieee.org/document/8100652/>, doi:10.1109/TCST.2017.2765625.
- De Coninck, R., Magnusson, F., Åkesson, J., Helsen, L., 2015. Tool-box for development and validation of grey-box building models for forecasting and control. *Journal of Building Performance Simulation* 9, 1–16. URL: <http://www.tandfonline.com/doi/full/10.1080/19401493.2015.1046933>, doi:10.1080/19401493.2015.1046933.
- Deb, C., Zhang, F., Yang, J., Lee, S.E., Shah, K.W., 2017. A review on time series forecasting techniques for building energy consumption. *Renewable and Sustainable Energy Reviews* 74, 902–924. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364032117303155>, doi:10.1016/j.rser.2017.02.085.
- Dimitriou, V., Firth, S.K., Hassan, T.M., Kane, T., Coleman, M., 2015. Data-driven simple thermal models: The importance of the parameter estimates. *Energy Procedia* 78, 2614–2619. URL: <https://www.sciencedirect.com/science/article/pii/S1876610215020548>, doi:10.1016/j.egypro.2015.11.322, arXiv:j.rser.2014.05.007.
- Escrivá-Escrivá, G., Álvarez-Bel, C., Roldán-Blay, C., Alcázar-Ortega, M., 2011. New artificial neural network prediction method for electrical consumption forecasting based on building end-uses. *Energy and Buildings* 43, 3112–3119. URL: https://ac.els-cdn.com/S0378778811003422/1-s2.0-S0378778811003422-main.pdf?_tid=29d377cd-f9fa-405a-bef8-e504638b79e3&acdnat=1540288745f92407d2f0a7d8804ed29f6951c0068c9, doi:10.1016/j.enbuild.2011.08.008.
- Foucquier, A., Robert, S., Suard, F., Stéphan, L., Jay, A., 2013. State of the art in building modelling and energy performances prediction: A review. *Renewable and Sustainable Energy Reviews* 23, 272–288. URL: <http://dx.doi.org/10.1016/j.rser.2013.03.004>, doi:10.1016/j.rser.2013.03.004.
- Geysen, D., De Somer, O., Johansson, C., Brage, J., Vanhoudt, D., 2018. Operational thermal load forecasting in district heating networks using machine learning and expert advice. *Energy and Buildings* 162, 144–153. URL: <http://dx.doi.org/10.1016/j.enbuild.2017.12.042>, doi:10.1016/j.enbuild.2017.12.042, arXiv:1710.06134.
- Hameed Shaikh, P., Bin Mohd Nor, N., Nallagownden, P., Elamvazuthi, I., Ibrahim, T., 2014. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews* 34, 409–429. URL: <http://dx.doi.org/10.1016/j.rser.2014.03.027>, doi:10.1016/j.rser.2014.03.027.
- Hansen, N., Müller, S.D., Koumoutsakos, P., 2003. Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evolutionary Computation* 11, 1–18. URL: <http://www.mitpressjournals.org/doi/10.1162/106365603321828970>, doi:10.1162/106365603321828970.
- Harish, V.S., Kumar, A., 2016. A review on modeling and simulation of building energy systems. *Renewable and Sustainable Energy Reviews* 56, 1272–1292. URL: <http://dx.doi.org/10.1016/j.rser.2015.12.040>, doi:10.1016/j.rser.2015.12.040.
- Huang, J., Smola, A.J., Gretton, A., Borgwardt, K.M., Schölkopf, B., 2006. Correcting Sample Selection Bias by Unlabeled Data, in: *NIPS'06 Proceedings of the 19th International Conference on Neural Information Processing Systems*, pp. 601–608. URL: <http://papers.nips.cc/paper/3075-correcting-sample-selection-bias-by-unlabeled-data.pdf>, doi:10.1.1.114.6412.
- Idowu, S., Saguna, S., Åhlund, C., Schelén, O., 2016. Applied machine learning: Forecasting heat load in district heating system. *Energy and Buildings* 133, 478–488. URL: <https://www.sciencedirect.com/science/article/pii/S0378778816310155>, doi:10.1016/j.enbuild.2016.09.068.
- Jamieson, K., Talwalkar, A., 2016. Non-stochastic Best Arm Identification and Hyperparameter Optimization, in: *19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 240–248. URL: <https://arxiv.org/pdf/1502.07943.pdf>, <http://arxiv.org/abs/1502.07943>, doi:10.1002/0471266388, arXiv:1502.07943.
- Kamaev, V.A., Shcherbakov, M.V., Panchenko, D.P., Shcherbakova, N.L., Brebels, A., 2012. Using connectionist systems for electric energy consumption forecasting in shopping centers. *Automation and Remote Control* 73, 1075–1084. URL: <https://link.springer.com/content/pdf/10.1134/S0005117912060124.pdf>, <http://link.springer.com/10.1134/S0005117912060124>, doi:10.1134/S0005117912060124.
- Kramer, R., Van Schijndel, J., Schellen, H., 2012. Simplified thermal and hygric building models: A literature review. *Frontiers of Architectural Research* 1, 318–325. URL: www.elsevier.com/locate/foar, Available online at www.sciencedirect.com, doi:10.1016/j.foar.2012.09.001.
- Lydon, G.P., Hofer, J., Svetozarevic, B., Nagy, Z., Schlueter, A., 2017. Coupling energy systems with lightweight structures for a net plus energy building. *Applied Energy* 189, 310–326. URL: <https://www.sciencedirect.com/science/article/pii/S030626191631741X>, doi:10.1016/j.apenergy.2016.11.110.
- Mat Daut, M.A., Hassan, M.Y., Abdullah, H., Rahman, H.A., Abdullah, M.P., Hussin, F., 2017. Building electrical energy consumption forecasting analysis using conventional and artificial intelligence methods: A review. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364032116310619>, doi:10.1016/j.rser.2016.12.015.
- Mattsson, S.E., Elmqvist, H., Otter, M., 1998. Physical system modeling with Modelica. *Control Engineering Practice* 6, 501–510. URL: <https://www.sciencedirect.com/science/article/pii/S0967066198000471>, doi:10.1016/S0967-0661(98)00047-1.
- McKinney, W., 2011. pandas: a Foundational Python Library for Data Analysis and Statistics, in: *Python for High Performance and Scientific Computing*, pp. 1–9. URL: <http://pandas.sf.net>.
- Oldewurtel, F., Parisio, A., Jones, C., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Wirth, K., 2010. Energy Efficient Building Climate Control using Stochastic Model Predictive Control and Weather Predictions, in: *Proceedings of the 2010 American Control Conference*, Ieee Service Center, 445 Hoes Lane, Po Box 1331, Piscataway, Nj 08855-1331 Usa. pp. 5100–5105. URL: <https://infoscience.epfl.ch/record/169733>.
- Oldewurtel, F., Parisio, A., Jones, C.N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Morari, M., 2012. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings* 45, 15–27. URL: <https://www.sciencedirect.com/science/article/pii/S0378778811004105>, doi:10.1016/J.ENBUILD.2011.09.022.
- Ölz, S., Sims, R., Kirchner, N., 2007. CONTRIBUTION

- OF RENEWABLES TO ENERGY SECURITY. Technical Report. International Energy Agency. URL: <https://www.iea.org/publications/freepublications/publication/so{ }contribution.pdf>.
- Paudel, S., Elmtiri, M., Kling, W.L., Corre, O.L., Lacarri re, B., 2014. Pseudo dynamic transitional modeling of building heating energy demand using artificial neural network. *Energy and Buildings* 70, 81–93. URL: <https://www.sciencedirect.com/science/article/pii/S0378778813007585?via{ }I3Dihub>, doi:10.1016/J.ENBUILD.2013.11.051.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay,  ., 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12, 2825–2830. URL: <http://www.jmlr.org/papers/v12/pedregosa11a.html>.
- Ramallo-Gonz lez, A.P., Eames, M.E., Coley, D.A., 2013. Lumped parameter models for building thermal modelling: An analytic approach to simplifying complex multi-layered constructions. *Energy and Buildings* 60, 174–184. URL: <https://www.sciencedirect.com/science/article/pii/S0378778813000315>, doi:10.1016/J.ENBUILD.2013.01.014.
- Recht, B., 2018. Updates on Policy Gradients – arg min blog. URL: <http://www.argmin.net/2018/03/13/pg-saga/>.
- Reinhart, C.F., Davila, C.C., 2016. Urban building energy modeling – A review of a nascent field. *Building and Environment* 97, 196–202. URL: <http://dx.doi.org/10.1016/j.buildenv.2015.12.001>, doi:10.1016/j.buildenv.2015.12.001.
- Richner, P., Heer, P., Largo, R., Marchesi, E., Zimmermann, M., Zimmermann, M., 2018. NEST – una plataforma para acelerar la innovaci n en edificios. *Informes de la Construcci n* 69, 222. URL: <http://informesdelaconstruccion.revistas.csic.es/index.php/informesdelaconstruccion/article/view/5879>, doi:10.3989/id.55380.
- Shanmuganathan, S., 2016. Artificial neural network modelling: An introduction, in: *Studies in Computational Intelligence*. volume 628, pp. 1–14. URL: http://link.springer.com/10.1007/978-3-319-28495-8_{ }1, doi:10.1007/978-3-319-28495-8_1.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., Van Den Driessche, G., Graepel, T., Hassabis, D., 2017. Mastering the game of Go without human knowledge. *Nature* 550, 354–359. URL: <http://www.nature.com/doifinder/10.1038/nature24270>, doi:10.1038/nature24270, arXiv:1610.00633.
-  irok y, J., Oldewurtel, F., Cigler, J., Pr vara, S., 2011. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy* 88, 3079–3087. URL: <https://www.sciencedirect.com/science/article/pii/S0306261911001668>, doi:10.1016/J.APENERGY.2011.03.009.
- Sturzenegger, D., Gyalistras, D., Morari, M., Smith, R.S., 2016. Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost–Benefit Analysis. *IEEE Transactions on Control Systems Technology* 24, 1–12. URL: <http://ieeexplore.ieee.org/document/7087366/>, doi:10.1109/TCST.2015.2415411.
- Sturzenegger, D., Gyalistras, D., Semeraro, V., Morari, M., Smith, R.S., 2014. BRCM Matlab Toolbox: Model generation for model predictive building control, in: *2014 American Control Conference, IEEE*. pp. 1063–1069. URL: <http://ieeexplore.ieee.org/document/6858967/>, doi:10.1109/ACC.2014.6858967.
- Suganthi, L., Samuel, A.A., 2012. Energy models for demand forecasting—A review. *Renewable and Sustainable Energy Reviews* 16, 1223–1240. URL: <https://www.sciencedirect.com/science/article/pii/S1364032111004242>, doi:10.1016/J.RSER.2011.08.014.
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P., 2017. Domain randomization for transferring deep neural networks from simulation to the real world, in: *IEEE International Conference on Intelligent Robots and Systems*, pp. 23–30. URL: <http://arxiv.org/abs/1703.06907>, doi:10.1109/IRROS.2017.8202133, arXiv:1703.06907.
- Wang, Z., Srinivasan, R.S., 2017. A review of artificial intelligence based building energy use prediction: Contrasting the capabilities of single and ensemble prediction models. *Renewable and Sustainable Energy Reviews* 75, 796–808. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364032116307420>, doi:10.1016/j.rser.2016.10.079.
- Wojna, Z., Gorban, A.N., Lee, D.S., Murphy, K., Yu, Q., Li, Y., Ibarz, J., 2018. Attention-Based Extraction of Structured Information from Street View Imagery, in: *Proceedings of the International Conference on Document Analysis and Recognition, ICDAR*, pp. 844–850. URL: <http://arxiv.org/abs/1704.03549>, doi:10.1109/ICDAR.2017.143, arXiv:1704.03549.
- Yudong Ma, Borrelli, F., Hencsey, B., Coffey, B., Bengesa, S., Haves, P., 2012. Model Predictive Control for the Operation of Building Cooling Systems. *IEEE Transactions on Control Systems Technology* 20, 796–803. URL: <http://ieeexplore.ieee.org/document/5739562/>, doi:10.1109/TCST.2011.2124461.
- Yun, D., Lee, H., Choi, S.H., 2018. A deep learning-based approach to non-intrusive objective speech intelligibility estimation. *IEICE Transactions on Information and Systems* E101D, 1207–1208. URL: <http://arxiv.org/abs/1412.6980>, doi:10.1587/transinf.2017EDL8225, arXiv:1412.6980.
- Zhao, H.X., Magoul s, F., 2012. A review on the prediction of building energy consumption. *Renewable and Sustainable Energy Reviews* 16, 3586–3592. URL: https://ac.els-cdn.com/S1364032112001438/1-s2.0-S1364032112001438-main.pdf?_{ }tid=d24ff951-1882-4611-b89e-5ee08d2c0510{ }&{ }acdnat=1540308093_{ }c9e893f0b56cb633368b99a0f42d87bc, doi:10.1016/j.rser.2012.02.049.