Empirical validation of a data-driven heating demand simulation with error correction methods

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

¹Urban Energy Systems Laboratory, Swiss Federal Laboratories for
Materials Science and Technology, Empa, Dübendorf, Switzerland

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

Abstract

Advanced control concepts for building energy systems, such as Model Predictive Control, often require models that forecast the energy demand of a building. Such models are commonly based on first principles, however the cost and effort required to develop such models may be prohibitive for real-life applications. As an alternative, we introduce and validate a data-driven simulation approach based on Artificial Neural Networks to forecast the heating demand of buildings. The forecast is enhanced with the help of two correction methods, based on online learning and forecast error auto-correlation. Validation results based on data from four office buildings suggest that our method shows better forecasting performance than a fitted 5R3C building model.

Introduction

Predictive control of buildings allows for energy and money savings as well as reduction of CO₂ emissions (Halvgaard et al., 2012). This can be achieved by predicting future states of a system and optimally determining according set-points, making use of renewable energy when available or taking advantage of time-dependent electricity prices for example. The concept of Model Predictive Control (MPC) has successfully been demonstrated in the building energy domain (Sturzenegger et al., 2016; Oldewurtel et al., 2012). Several authors, however, (see for example (Jain et al., 2018; Smarra et al., 2018)) argue that a major bottleneck with this approach is the need to develop and maintain first principles models. Indeed the time and effort necessary to develop, customise and maintain first principles models can often be prohibitive for real life deployment to building energy management.

With increasing availability of high-resolution energy monitoring data in the building domain, the possibility of data-driven modelling approaches (or Machine Learning based approaches) arises. With the help of methods such as Artificial Neural Networks (ANN) previously measured data of the ambient temperature and the heating demand can be used to predict

the future heating demand of the building.

In (Bünning et al., 2019) a heating demand forecasting approach based on ANN with forecast correction methods based on online learning and forecast error auto-correlation was introduced. The aim is to generate a 24h forecast of the heating demand in 15-minute intervals. The method showed high prediction accuracy on a modern multi-use building and outperformed other regression-based methods such as random forests or support vector machines, as well as fitted resistor-capacitor models with varying complexity. Moreover, the dependence of ANN prediction performance on initialization parameters, which are commonly randomly set, as criticized by Recht (2018); Henderson et al. (2018) and others, is significantly lowered with the shown methods.

Here we explore this approach further and show that the concept generalises well to different types of data and buildings. In particular we apply the method to four different buildings, each with different building characteristics and data availability. By comparing the prediction accuracy to a fitted 5R3C building model, which gives a reasonable trade-off between detail and modelling/fitting effort, we demonstrate that the approach outperforms the R-C model in all presented cases. Furthermore, the robustness to initializing parameters of the ANN is also preserved for all test cases.

Methodology

This section is divided into four subsections. First, the forecasting task is introduced. Second, the forecasting method based on ANN with forecasting correction based on online learning and error auto-correlation is explained. Third, the used R-C building model is introduced. Finally, an overview of the validation study, key-performance indicators and buildings is provided.

Forecasting task

In this study the following forecasting task is assumed: A heating demand forecast is made at midnight for the next 24 hours, sampled every 15 minutes. The training and validation data are assumed to be

Figure 1: Artificial neuron.

sampled at the same frequency. 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.

Artificial Neural Network with forecast correction

We use ANN for the heating demand forecast. An ANN is a sequence of artificial neurons, also called nodes. A schematic of a node with three inputs is shown in Figure 1. The output of a node is the activation function applied to the weighted sum of all inputs. The nodes weights are the parameters that are fitted in the training process of the ANN. For a more detailed explanation of ANN, see (Basheer and Hajmeer, 2000) for example.

In (Bünning et al., 2019) two forecast correction methods were developed. The first forecast correction method is based on error auto-correlation and is not specific to ANN. The use of the error autocorrelation is motivated by the assumption that errors in building energy forecasts are correlated over time, because the source of the error persists for more time than one sample: for example if a window is opened or a failure in a fluid pump occurs, the forecasting error will last for more than one interval, because the source of the error is not always eliminated during this interval (window closed or fluid pump fixed). Thus, if an error is measured, it can be used to make an estimate of the error in the following forecasting steps.

Figure 2 demonstrates the first correction procedure. Offline, based on the measured errors of all past forecasts (which are stored in a data-base), an estimation of the error auto-correlation is made. As a first step in the online phase, a forecast is made with an ANN. We will call this the uncorrected forecast. With the help of the auto-correlation and the forecasting error of the previous day (which is the uncorrected forecast of the previous day minus the measured load of that day), an estimation for the forecasting error of the uncorrected forecast can be made. The uncorrected forecast is then corrected by adding the estimation of the forecasting error, giving rise to the corrected forecast.

Figure 2: Forecast correction scheme.

Figure 3: 5R3C building model.

The second forecast correction method is specific to ANN and makes use of online learning. Instead of only training on the basis of the training set, the ANN is retrained based on the measurement of the latest heating demand realization and the corresponding network inputs at the end of every day. This biases the network towards the recent behaviour of the building and appears to improve the forecasting accuracy. The two forecast correction methods can of course be combined.

R-C building model

In order to benchmark the performance of the ANN and the correction mechanisms, we use a 5R3C building model. This model was chosen because on the one hand it should offer enough complexity to capture the governing thermal behaviour of the buildings and on the other it should perform well with a modelling effort similar to the neural networks. More detailed models, for example an EnergyPlus model combined with an occupancy model adapted to each building, may outperform the 5R3C and ANN models in terms of forecasting accuracy, but would require considerably more effort to develop and maintain, as discussed above.

Figure 3 shows a schematic of the 5R3C model. The model features a capacitor for outer walls, a thermal zone and inner walls. Furthermore, there are five resistors - two for the outer walls, one for the inner walls, one for heat conductions towards the ground and one for building infiltration. All capacitor and resistor values are optimization variables in the fitting

process. The total heating demand of the building

$$\dot{Q}_{total} = \dot{Q}_1 + \dot{Q}_2, \quad (1)$$

is distributed between $\dot{Q}_1$ (directly to the zone) and $\dot{Q}_2$ (to the core) to be able to model both direct heating systems such as radiators and indirect ones such as concrete core activation. The share, c_f , between $\dot{Q}_1$ and $\dot{Q}_2$ is also a decision variable in the fitting process:

$$\dot{Q}_1 = c_f \times \dot{Q}_{total}. \quad (2)$$

To emulate a simple control system, the temperature of C_{zone} , T_{zone} , is controlled with a P controller,

$$\dot{Q}_{total} = k_p \times (T_{set} - T_{zone}), \quad (3)$$

where T_{set} denotes the temperature set point for the zone. The gain, k_p , is also a decision variable in the fitting process.

Validation study

General description

To validate the neural networks and the correction methods and to benchmark them against the R-C model, four different buildings were selected as case studies.

For each building heating demand measurements and ambient temperature measurements in 15 minute intervals were available over a period of one to three years. In each case 70% of the available dataset was used for model training (for both the ANN and the R-C models) and 30% of the dataset was used for model validation.

As a key performance indicator the coefficient of determination,

$$R^2 = 1 - \frac{\sum_{i \in N} (y_i - f(x_i))^2}{\sum_{i \in N} (y_i - \bar{y}_N)^2}, \quad (4)$$

is used. It is *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 *one* if the forecast is exactly the same as the validation data y_i . R^2 becomes negative if the forecast is worse than taking the average.

The different buildings are introduced in the following and are depicted in Figure 4.

Building (a) is the NEST building at Empa in Dübendorf, Switzerland. Opened in 2016, it is a multi-use building with residential units, office units as well as meeting rooms and a fitness center. Its heating and cooling system is build to imitate a district energy system. The building is therefore referred to as a 'vertical district' (as opposed to a conventional district, which is built horizontally). 13.5 months of measured demand data was used for this building.

(a) NEST at Empa, © Zooey Braun.

(b) ETL at ETH Zürich (ETH Zürich, 2019).

(c) Bauhalle at Empa, © Heinrich Helfenstein.

(d) Verwaltungsgebäude at Empa (Empa, 2019).

Figure 4: Buildings used in the validation study.

Building (b) is the ETL building located at the centrum campus of ETH Zürich. It is an eleven-story office building that comprises several research institutes, experimental laboratories and a machine shop. Six stories lie below ground. 36 months of measured demand data was used for this building.

Building (c) is the Bauhalle at Empa. The building has offices and testing facilities for material science experiments. The offices are spread over three stories while some of the laboratories have the ceiling height of these three stories. The building also has a basement. 24 months of measured demand data was used for this building.

Building (d) is the Verwaltungsgebäude at Empa. This is a three-storey office building with some meeting rooms. 24 months of measured demand data was used for this building.

Specific models for the case studies

Artificial Neural Networks with two hidden layers, each containing eight nodes were chosen as the network architecture. They are trained with ten epochs on the training set, meaning that the network's weights are updated ten times based on each sample in the training data. The choice of parameters for the architecture and training process is based on a sensitivity analysis regarding these parameters.

As inputs to the network the following features were chosen:

1. ambient temperature
2. hour of the day
3. weekday/weekend
4. heating load of the previous day
5. heating load of the previous week

The feature *hour of the day* was one-hot encoded, meaning that instead of one input with a continuous

value, 24 inputs with binary values (either 0 or 1) were used. Our experiments suggest that this improves prediction accuracy compared to a continuous input.¹ The feature *weekday/weekend* is also binary. All other features are continuous. The Python package Keras (Chollet, 2018) was used to implement the network and the solver adam (Kingma and Lei Ba, 2015) was used for the optimization of the network parameters. The training of one network takes between five and twelve minutes.

The 5R3C model was implemented in Modelica (Mattsson et al., 1998) and simulated in Dymola (Brück et al., 2002). To fit the parameters, the model was coupled to Python and optimized with a CMA-ES (Covariance Matrix Adaption Evolution Strategy) optimizer to minimize the sum of squared errors between the predicted load and the reference load. The optimization budget was fixed at three hours for each model. As the model does not allow for the other inputs, only the ambient temperature is used. The room temperature set point is $T_{set} = 20^{\circ}\text{C}$ and T_{ground} is set to constant 12°C in all cases.²

Results and discussion

Variance

The prediction performance of ANN is dependent on their parameters, some of which are random, for example the initialization parameters of the node weights. If ANN are used in a control context, a reduction of the variance in prediction performance is desirable, as the controller needs to be designed for the expected accuracy of the forecast.

To investigate the variance in prediction accuracy, 100 instances of ANN were trained and tested for each building. Each individual network is initialized with random node weights. As the training process is non-convex this leads to a different prediction accuracy of each network.

Figure 5 (a)-(d) shows the variance of the coefficient of determination for 100 different instances of neural networks for each building (with boxes describing the interquartile range, whiskers of 1.5 times the interquartile range, circles indicating outliers and the orange line indicating the median). The first column depicts uncorrected networks, the second column networks with error auto-correlation correction, the third column networks with online learning and the last column a combination of both correction measures.

¹This can be explained with the monotonic increase of the ReLu activation functions that are used in the network. A continuous implementation of the time as an input would imply that the influence of the time on the demand is either monotonically increasing or decreasing.

²This might have an effect on the accuracy of the forecast. However, a more detailed implementation would require a human-in-the-loop model for the temperature set-point and a ground model for the ground temperature, both of which require extensive modelling effort.

Subplot (a) corresponds to the NEST building (building (a)) that was also used in (Bünning et al., 2019). It can be seen that the median of the coefficient of determination increases whereas the variance decreases from no correction to the combination of both correction methods. Online learning gives the bigger improvement when compared to the correction based on the error auto-correlation. However, the combination of both gives the best results: even though the variance does not improve much, the median of the coefficient of determination further increases. The same trends are observed for buildings (b) and (c).

In the case of building (d), the trend is not as strong, but still present. The median increases with each individual correction method, but the median R^2 of the online learning correction is lower than the one for error auto-correlation correction. The median of the combined method is again the highest one. The maximum achieved coefficient of performance is lower for the online learning and the combination of both correction methods compared to the uncorrected network and the correction method based on error auto-correlation. This is due to overcorrection of already well predicting networks.

The difference in the result between building (d) and the other buildings could be explained by the quality of the available data set. The measured heating demand of building (d) has a quantization of 0.5 kW at a maximum load of 26.5 kW, while for all other buildings the measured demand has a quantization of 0.5 kW at a maximum load of at least 86 kW. Quantisation error effects are therefore likely to be more pronounced for building (d).

These results generally support the findings of Bünning et al. (2019) and indicate that the correction methods can be beneficial for a range of different buildings. An interesting question is why the combination of both correction methods tends to improve over the individual methods. Both correction methods address a similar problem: the mismatch between forecast and realization due to recently changed behaviour of the building. It seems, however, that the two methods correct for slightly different, complementary phenomena. The correction based on error auto-correlation can react very efficiently to errors that have occurred shortly before the end of the last day's forecast and correct the beginning of the next day's forecast, when the auto-correlation, hence the confidence in the error estimation, is high. However, strong corrections are only possible for the first few intervals of the day, as the auto-correlation of the forecasting error decreases during the course of the day. The correction based on online learning, on the other hand corrects to a lesser extend on each individual interval, as the ANN still needs to be able to generalize. However, this correction can be applied over the course of the full day's forecast and not just

Figure 5: Variance of the coefficient of determination for all buildings. The label of the subplots (a)-(d) reflects the number of the building used in the case study.

Table 1: R^2 of different modelling approaches for all buildings.

Building	ANN uncorr.	ANN auto	ANN online	ANN auto+online	5R3C model
Building (a)	0.818	0.860	0.878	0.885	0.761
Building (b)	0.886	0.907	0.933	0.936	0.890
Building (c)	0.674	0.752	0.793	0.809	0.676
Building (d)	0.856	0.858	0.860	0.862	0.747

Figure 6: Forecast trajectory example from building (b) test set with one instance of an ANN and both correction methods.

at the beginning of the day. The combination of both therefore gives the best result.

Accuracy

Table 1 shows the mean coefficient of determination (as opposed to the median in Figure 5) of the different ANN approaches for all buildings as well as the coefficient of determination achieved by the 5R3C building model. It can be seen that the average R^2 increases for all four buildings when the individual correction methods are applied. A combination of both approaches (auto + online) gives a further improvement in all cases. The ANN with the applied correction methods outperform the 5R3C model in all investigated cases.

These results are of course by no means a proof that the ANN with corrections outperforms R-C type models in general. The R-C model could for example be made more complex or the parameter fitting could be improved leading to higher coefficients of determination. However, the results do give an indication that the corrected ANN perform reasonably well compared to conventional methods and generalize to different types of buildings.

We note in passing that it is not possible to compare the coefficient of determination between buildings, because it depends on the distribution of heating loads in the data set, which is different for each building.

Physical behaviour of the forecast

Figure 6 shows the scaled forecast and real heating load of building (b) from an excerpt of the test set. The figure demonstrates that the real load is captured reasonably well by the forecast, although the peaks in the last third of the excerpt are not predicted well. This lends further support to our earlier observations on the accuracy of the forecasts based on the coefficient of determination. Moreover, even though ANN are known to sometimes have weak extrapolation capabilities, in our case the forecast does not show any unrealistic behaviour in the sense of unreasonably high or low forecasting values. The results for all other buildings³ also point to the same conclusion. While this does not guarantee that non-physical behaviour is impossible with other forecasting inputs or for other building cases, the positive results for all four buildings indicate a robustness that could be sufficient for non-safety-critical control tasks.

To further increase the confidence in the forecast, it could be limited in its derivative and maximum/minimum values.

³The full forecasting plots of both training and test set for one instance of an ANN with both correction methods applied to all four buildings are shown in Appendix A and Appendix B.

Limitations

The results of the case studies indicate that the methods are sensitive to the quality of the measurement data. A quantitative statement cannot be made yet. This will be considered in following studies.

Furthermore and as mentioned before, this validation study gives an indication that the methods developed in (Bünning et al., 2019) generalize to different buildings. However, it does not give a theoretical proof of any kind. The confidence in the methods is strengthened by the consistency of the results.

Conclusion

Generating and maintaining building models based on first principles is often considered prohibitively expensive and time consuming for applications such as real-time optimal control. The increasing availability of data in the building domain allows the use of data-driven modeling approaches. However, such approaches have the disadvantage that they - in theory - allow non-physical behaviour and that they depend on random parameters such as the initialization of node weights in ANN.

In this case study we have validated two forecast correction methods for ANN for the forecasting of sub-hourly heating demands in buildings. The results for four different buildings have shown that the resulting ANN perform well when compared to a 5R3C building model in these cases. Furthermore, the correction methods avoid non-physical behaviours and significantly reduce the variance in ANN performance, which can improve confidence in using ANN in the frame of control tasks.

In future work we want to quantify the quality of input data vs. accuracy of forecast trade-off. Moreover, we want to couple the introduced approaches to randomised optimisation methods (treating forecasts from different networks as samples) and robust optimisation methods (using the intervals).

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

- Basheer, I., Hajmeer, M., 2000. Artificial neural networks: fundamentals, computing, design, and application. *Journal of Microbiological Methods* 43,

- 3–31. doi:10.1016/S0167-7012(00)00201-3.
- 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.
- Bünning, F., Bollinger, A., Heer, P., Smith, R.S., Lygeros, J., 2019. Improved day ahead heating demand forecasting by online correction methods. doi:10.31224/osf.io/vxw6g.
- Chollet, F., 2018. Keras: The Python Deep Learning library. Astrophysics Source Code Library .
- Empa, 2019. Empa - Communication - Ein Preis für die «Pfeilbogenbrücke».
- ETH Zürich, 2019. Gebäude – Services und Ressourcen | ETH Zürich.
- Halvgaard, R., Poulsen, N.K., Madsen, H., Jorgensen, J.B., 2012. Economic Model Predictive Control for building climate control in a Smart Grid, in: 2012 IEEE PES Innovative Smart Grid Technologies (ISGT), IEEE. pp. 1–6. doi:10.1109/ISGT.2012.6175631.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D., 2018. Deep Reinforcement Learning that Matters, in: The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), pp. 3207–3214.
- Jain, A., Nghiem, T., Morari, M., Mangharam, R., 2018. Learning and Control Using Gaussian Processes, in: 2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS), IEEE. pp. 140–149. doi:10.1109/ICCPS.2018.00022.
- Kingma, D.P., Lei Ba, J., 2015. ADAM: A METHOD FOR STOCHASTIC OPTIMIZATION, in: 3rd Int. Conf. Learn. Representations. arXiv:1412.6980v9.
- Mattsson, S.E., Elmqvist, H., Otter, M., 1998. Physical system modeling with Modelica. Control Engineering Practice 6, 501–510. doi:10.1016/S0967-0661(98)00047-1.
- 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. doi:10.1016/J.ENBUILD.2011.09.022.
- Recht, B., 2018. Updates on Policy Gradients – arg min blog.
- Smarra, F., Jain, A., de Rubeis, T., Ambrosini, D., D’Innocenzo, A., Mangharam, R., 2018. Data-driven model predictive control using random forests for building energy optimization and climate control. Applied Energy 226, 1252–1272. doi:10.1016/j.apenergy.2018.02.126.
- 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. doi:10.1109/TCST.2015.2415411.

Appendix A

Demand forecasts with one instance of an ANN and both correction methods for buildings (a)-(d) **on the training set**.

Figure 7: Scaled forecasting trajectories on the training set.

Appendix B

Demand forecasts with one instance of an ANN and both correction methods for buildings (a)-(d) **on the testing set**.

Figure 8: Scaled forecasting trajectories on the test set.