DTU Library

Localizing weather forecasts for enhanced heat load forecast accuracy in urban district heating systems

Bergsteinsson, Hjörleifur G.; Sørensen, Mikkel L.; Møller, Jan Kloppenborg; Madsen, Henrik

Published in:
International Journal of Sustainable Energy

Link to article, DOI:
10.1080/14786451.2023.2271261

Publication date:
2024

Document Version
Publisher's PDF, also known as Version of record

International Journal of Sustainable Energy

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/gsol20

Localizing weather forecasts for enhanced heat load forecast accuracy in urban district heating systems

Hjörleifur G. Bergsteinsson, Mikkel L. Sørensen, Jan Kloppenborg Møller & Henrik Madsen

To cite this article: Hjörleifur G. Bergsteinsson, Mikkel L. Sørensen, Jan Kloppenborg Møller & Henrik Madsen (2024) Localizing weather forecasts for enhanced heat load forecast accuracy in urban district heating systems, International Journal of Sustainable Energy, 43:1, 2271261, DOI: 10.1080/14786451.2023.2271261

To link to this article: https://doi.org/10.1080/14786451.2023.2271261

INTERNATIONAL JOURNAL OF SUSTAINABLE ENERGY
2024, VOL. 43, NO. 1, 2271261
https://doi.org/10.1080/14786451.2023.2271261

Localizing weather forecasts for enhanced heat load forecast accuracy in urban district heating systems

Hjörleifur G. Bergsteinsson, Mikkel L. Sørensen, Jan Kloppenborg Møller and Henrik Madsen

Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kongens Lyngby, Denmark

ABSTRACT
Weather forecasts are essential for district heating (DH) utility operations as they prepare the utility for future consumption, thus ensuring optimal operation by supplying sufficient heat while keeping costs low. Weather forecasts are usually converted into heat demand forecasts, which are used for production planning and control of the temperatures in the network. Hence, increasing the accuracy of weather forecasts will lead to improvements in the system's operational performance. However, numerical weather predictions (NWPs) are computed over the earth as grid values, and NWPs are designed for rural areas, not urban areas. Therefore, we propose to localise the weather forecasts to the urban environment by calibrating them using Model Output Statistics. We show that localising weather forecasts (removing the bias) leads to enhanced accuracy in the heat demand forecasts. In our case study, localised weather forecasts lead to an error reduction between 1.5% and 2.5% when compared to forecasts using uncalibrated NWPs.

ARTICLE HISTORY
Received 21 September 2022
Accepted 11 September 2023

KEYWORDS
Forecast calibration; recursive estimation; heat demand; localised weather forecast; urban heat island

1. Introduction

The role of district heating is changing dramatically, with more sector coupling and increasing amounts of energy being generated by renewable energy sources. Traditionally, the role of district heating has been to provide heat to consumers when needed, usually by burning fossil fuels. Sometimes, district heating production plants also produce electricity. These are called cogeneration plants. However, with awareness of the global climate crisis due to increasing CO_2 concentrations in the air, more climate-friendly operation of energy production facilities is needed. Hence, district heating needs to shift away from fossil fuels towards renewables, and due to its flexibility as a means of energy storage, district heating has become a crucial element of the efficient operation of the overall integrated energy system. To maximise the potential of district heating, more advanced methods are needed for its operation. Advanced methods for production optimisation of plants have been proposed for optimal scheduling of heating units and for bidding into the electricity markets to increase the share of renewables and lower overall system costs (Blanco et al. 2018). Advanced methods for delivering heat to consumers have also been proposed. For instance, Madsen et al. (1994) propose advanced optimisation methods to control supply temperature and flow in district heating networks in order to reduce heat production costs and heat losses in the network.

CONTACT Hjörleifur G. Bergsteinsson hgbe@dtu.dk

Heat demand forecasts are needed for efficient production planning and for temperature control of the network. For some applications, like for optimal use of storage systems, the forecasts are most conveniently provided as scenarios. Hence, accurate future scenarios of heat consumption are desirable to lower the overall cost of operating a district heating system. Heat consumption is directly related to the local climate, and therefore heat demand forecast models usually use forecasts of weather variables as inputs. Thus, the accuracy of heat demand forecasts can be increased by improving the weather forecasts. Additionally, district heating is typically applied in urban areas, where the climate is generally different from rural areas due to human activities, buildings, and infrastructures. This is a well-known effect termed Urban Heat Islands (UHI). Therefore, the operation of district heating needs to account for this local climate to achieve efficient operation. Heating consumption is highly correlated with climate variables, especially the ambient air temperature, as has been shown and discussed by multiple previous studies, see e.g. Dotzauer (2002) or Dahl et al. (2018). Most high-quality models for heat demand thus include forecasts of temperature when forecasting beyond the 1-hour mark (Frederiksen and Werner 2013). Predictions of the weather are therefore desired by district heating companies to ensure that heat production is sufficient to meet future heating consumption and to ensure efficient production. Numerical Weather Prediction (NWP) is used to forecast the weather by simulating physics-based partial differential equations of the atmospheric processes (Bauer, Thorpe, and Brunet 2015). Most often the NWPs do not represent the UHI effect and hence the weather inside cities, and consequently a systematic bias between the NWP and local weather stations is often observed (Crochet 2004). It would therefore be beneficial for district heating applications to localise the weather forecast, thus removing the bias and thereby reducing the prediction uncertainty of the weather forecasts.

Several studies have also investigated the possibility of increasing the forecasting capabilities of NWPs by including the UHI effects in the model. See e.g. Ronda et al. (2017) for their suggestion of incorporating UHI in NWPs for the city of Amsterdam.

More broadly however there exist multiple methods of localising or calibrating forecasts. One of the more prominent methods is the Model Output Statistics (MOS) technique as proposed by Glahn and Lowry (1972). MOS is a well-known method and is often applied by forecast providers to calibrate or localise their forecasts, e.g. MOS has been used to calibrate NWPs for use in wind power forecasting (Giebel and Denhard 2017). Furthermore, an extension of the MOS technique for probabilistic forecasting called Ensemble Model Output Statistics (EMOS) (Gneiting et al. 2005) is commonly used for a number of purposes, e.g. in solar power forecasting Gneiting, Lerch, and Schulz (2023) or hydropower forecasting Graham et al. (2022). However, to the knowledge of the authors, this method has not been widely applied in the context of heat demand forecasting.

We propose to use Model Output Statistics on NWPs of ambient air temperature to calibrate them to the local conditions, thus incorporating the UHI effect in the forecasts used for heat demand forecasting.

This study is structured such that Section 2 describes the methods used for localising the NWP as well as the modelling framework. Section 3 describes the data used for testing the proposed method of localisation. Section 4 presents the results of the localisation as well as the effects of the localisation on the heat demand forecasts. Lastly, Section 5 summarises and concludes the findings.

2. Method

Heat demand forecast models have been studied extensively for the past decades and both grey-box and black-box models have been proposed. Black-box models are methods where the modeller feeds data through the model, trains or calibrates the model parameters using these data, and finally produces forecasts without considering any physical knowledge of the system. As an example, Idowu et al. (2016) propose four different methods to forecast heat demand. Grey-box models are based on a combination of physical knowledge and information embedded in data.

The methods use statistical methods to estimate the parameters and reduce the model dimensions compared to a purely physical model of the system. Aalborg Nielsen and Madsen (2000) have discussed a physical-based model of the heat consumption in a district heating network and they suggested a simplified version of the full (complex) model. They demonstrate that a rather simple model, based on the physical nature but calibrated using data on heat consumption, is adequate for forecasting heat demand.

The forecasting models that will be used to both localise the numerical weather predictions and produce heat demand forecasts are generated using recursive and adaptive techniques, as implemented in the R package, `onlineforecast` (Bacher et al. 2021). The package provides a simple but efficient framework for constructing and optimising multi-step forecast models. The methodology is based on modelling the output variable as a linear combination of the input. Different transformations are possible for the input variables, either by a function or non-parametrically through e.g. splines. The coefficients of the linear model are then estimated, either by simple least squares or by recursive least squares. In the case of recursive least squares, the coefficients are updated adaptively when new data arrives through the use of an exponential forgetting factor. For further details on the forecasting framework provided by `onlineforecast`, see Bacher et al. (2021).

Numerical weather predictions, such as those used for the proposed models, are generally created from a system of nonlinear differential equations based on the laws of physics and predict the weather by describing the physical processes in the atmosphere. These differential equations cannot be solved for each point on the earth, and therefore a spatial grid of the earth is used as points for the equations to be solved at. However, these systems model urban areas with the same formulations as rural areas, hence discarding the Urban Heat Island effect (Baklanov et al. 2009).

Therefore, the raw forecast of the ambient air temperature does not consider the local climate in the area where the heat demand forecast is needed. However, if there is a local climate station that measures the temperature inside the area, then it will be possible to adjust the temperature forecast using such local measurements. Multiple methods have been used to localise numerical weather predictions, however, one of the more widely used is the Model Output Statistics method by Glahn and Lowry (1972). Model Output Statistics is a post-processing method to calibrate or localise numerical weather predictions using a regression model. In this regression model, the dependent variable could be the observed weather variable at a local weather station, and the possible explanatory variables are e.g. the forecast weather variable (i.e. the same variable as the dependent) and additional forecast weather variables.

In our case, we suggest,

$$y = \beta_0 + \beta_1 T_a^{\mathrm{nwp}} + \epsilon, \tag{1}$$

where the dependent variable, y is the observed ambient air temperature, T_a^{nwp} is is the forecasted ambient air temperature while the β are coefficients of the regression model. This is indeed a very simple model, since it assumes that the bias introduced by the UHI effect only scales linearly with the temperature forecast and is independent of other potential variables. This method of localisation does however have the added benefit of updating the weather forecasts between forecast times, i.e. since the NWPs are usually created only a few times a day, the localisation model, which receives measurements more frequently, helps in updating the increasingly uncertain weather prediction.

Usually, MOS is applied in a multiple-linear regression fashion, where multiple weather forecasts are used (Gneiting 2014). So this model is simpler than what is usually applied. However, it was nonetheless chosen for this study to show that even simple methods of incorporating the UHI effect can affect the subsequent heat demand forecasts.

The forecasts produced by the localisation model are used as inputs for the heat demand forecast model to quantify the accuracy improvement by localising the NWP. This method provides the possibility to locally adjust the NWP and therefore update it with information on the local climate, thereby providing better information for predicting the heat demand. The method also makes it possible to estimate the uncertainty of the NWP for this location.

The proposed heat demand forecasting method in this paper uses physical knowledge on how heating consumption and weather are linked as suggested by Aalborg Nielsen and Madsen (2000). A final model was identified using the forward selection approach as shown in Bacher et al. (2013), where inputs are added one by one to the model, and for each input, it is investigated if the forecasting performance increases significantly. The Root-Mean-Squared-Error (RMSE) is used to compare the forecasting performance for each k-step horizon. Different transformations of the inputs are also investigated. Models for heat demand are usually based on weather forecast inputs (e.g. air temperature and wind) and social components (e.g. Fourier Harmonics).

3. Data

Data is required to apply the methods proposed in this study and investigate if heat demand forecasts can be improved by localising the weather forecasts. Therefore, the case study from the IDASC project for the area of Tingbjerg in Copenhagen is used Bergsteinsson et al. (2021). This case study contains heat demand measurements, temperature measurements from a local climate station, and weather forecasts.

Two years of hourly measurements from 2019 and 2020 are used in this study. Data from 2019 will be used in-sample to estimate model parameters and forgetting factors, while data from 2020 is used out-of-sample to test and validate performance.

Ordinarily, cross-validation would be used to avoid overfitting, however since estimation in `onlineforecast` is done by recursive least squares only past data is used when calculating the regression coefficient Bacher et al. (2021).

From the point of view of the utility provider (HOFOR), the district heating system in Tingbjerg can be thought of as an island in their system, as the Tingbjergs distribution network has only one heat source, a heat exchanger that receives heat from the transmission network. The top plot in Figure 1 shows the historical measurement of the heat demand from the heat exchanger at Tingbjerg.

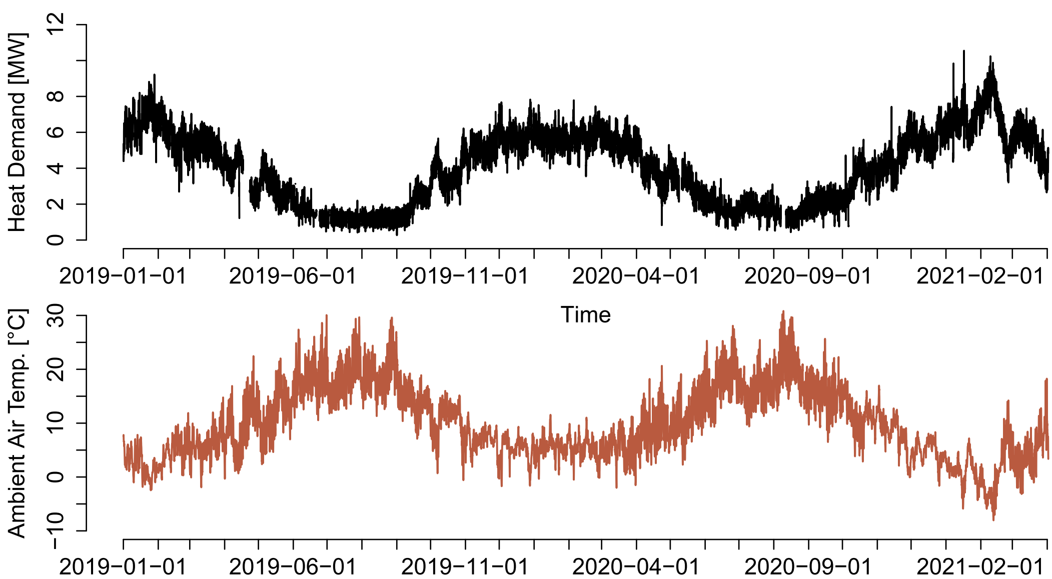

Figure 1. Figure shows measurements of heat demand (top) and ambient air temperature (bottom).

The climate station that will be used to provide local weather measurements is located at the Landbohøjskolen in Copenhagen (Station ID: 06186). The data is provided by the Danish Meteorological Institute (DMI) through their Open Data initiatives (Danish Meteorological Institute (DMI) 2021). We will only be using the ambient air temperature (the *temp_mean_past1h* variable) in this work. The lower plot in Figure 1 shows the measured ambient air temperature at the local climate station.

It is clear from these plots that the heat demand and the ambient air temperature are highly correlated. Hence, accurate forecasts of the ambient air temperature are of great importance when forecasting heat demand.

Ideally, the location of the climate station would be inside the area where the heat is being consumed, unfortunately however, there are no climate stations in Tingbjerg. Therefore, we found the climate station that is closest to Tingbjerg and also inside Copenhagen. Figure 2 shows where Tingbjerg is located and the placement of the climate station. The distance between them is approximately 5.6 kilometres.

The NWPs are forecasts from the ECMWF HRES model and they are updated twice a day; at noon and midnight.[2] The forecasts used were created by selecting the longitude and latitude point for Tingbjerg (55°43’12”N, 12°28’48”E, the blue point in Figure 2) and are then bilinearly interpolated from the four nearest neighbour forecasts from the $0.1^\circ \times 0.1^\circ$ grid made by the ECMW HRES model.

Examining the difference between the forecasted temperature and the locally measured ambient air temperature in Figure 3, the bias in the weather forecast due to the UHI effect can be seen clearly. On average for the training period, the weather forecast underestimates the local ambient air temperature by approximately 0.65°C. It is natural that such a difference will affect the performance of a heat demand model.

4. Results

The NWPs are localised by estimating parameters in Equation (1) during 2019 with an initial burn-in period of 3 months to estimate the forgetting factor. The parameters and forgetting factor

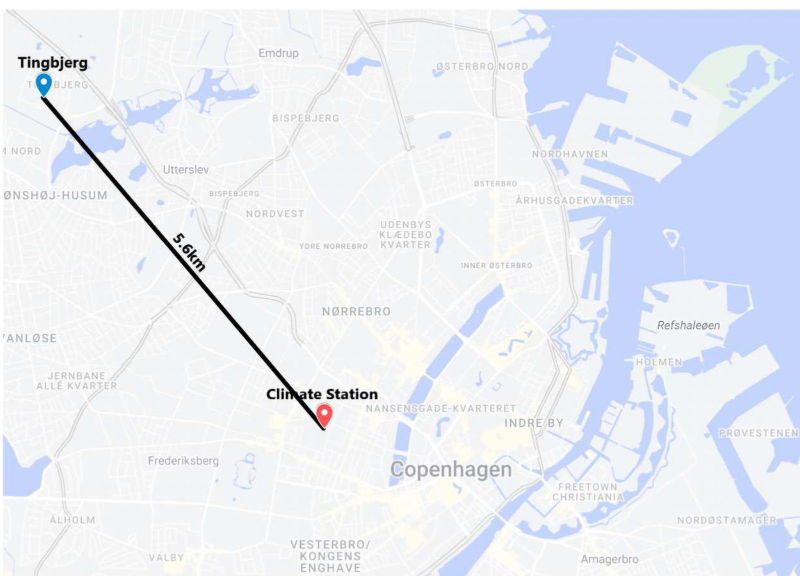

Figure 2. Figure shows the location of the climate station and Tingbjerg, where the district heating is located. The distance between the two places is approximately 5.6 km.[1]

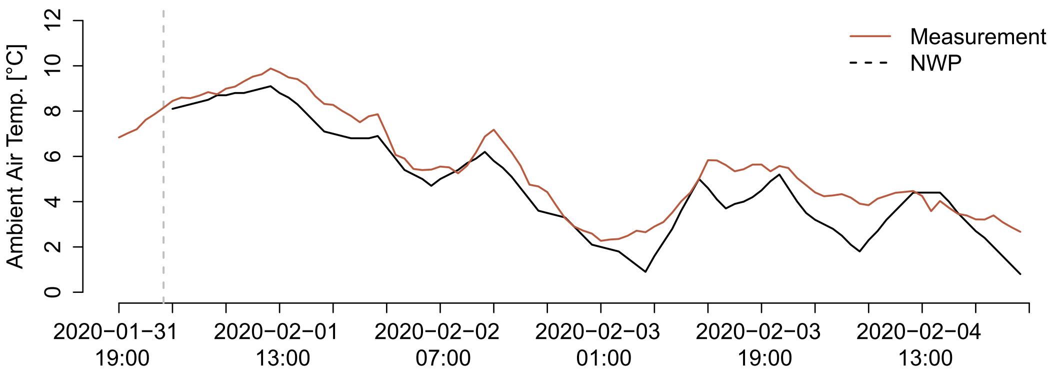

Figure 3. Figure shows local ambient air temperature measurements (red) and corresponding forecasted temperature (black). The vertical grey dashed line indicates when the NWP was produced, i.e. the prediction time.

are then used out-of-sample to forecast the weather variables during 2020. The parameters of the model are estimated and the forecasts of Equation (1) are obtained using the R-package `online-forecast`. This R-package for forecasting uses either least squares (constant parameters over time) or recursive least squares with exponential forgetting factor (adaptive parameters, i.e. varying over time). The performance of the raw NWPs and the NWPs that have been locally adjusted to the climate of the ambient air temperature can be seen in Figure 4. Three different localised NWPs were created:

- Parameters estimated using least squares (LS) estimation, i.e. constant parameters (denoted as *lm*)

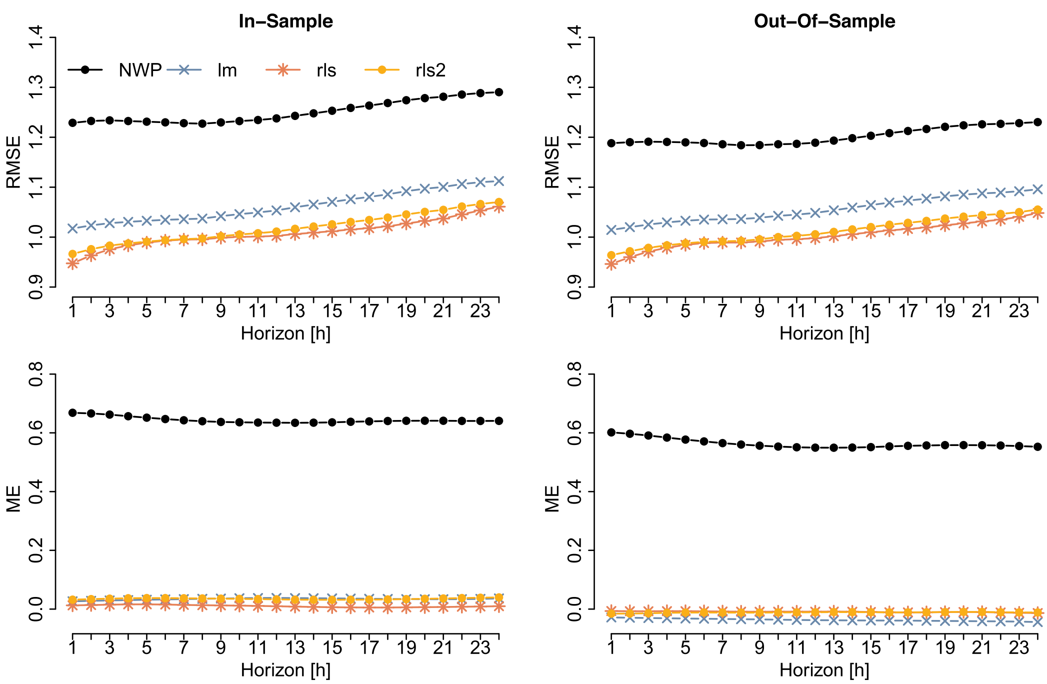

Figure 4. Figure shows the accuracy of the NWPs in Tingbjerg. The left plots show the in-sample and the right plots show the out-sample. The Root Mean Square Error (RMSE) is shown in the top plots and Mean Error (ME) in the bottom plots.

- Parameters estimated using recursive and adaptive estimation, with forgetting factor ($\lambda = 0.994$) (denoted as *rls*)
- Parameters estimated using recursive and adaptive estimation, with forgetting factor ($\lambda = 0.998$), i.e. making the parameters change slower over time than case 2 (denoted as *rls2*).

The forgetting factor for the second method, *rls*, was found by minimising the RMSE of the prediction on horizons, $k = \{1, \ldots, 24\}$ hours and the forgetting factor in *rls2* was selected by hand. The results from the in-sample period are shown in the left plots, while the right plots show the out-of-sample results. Two error scores are used; the RMSE in the top plots and Mean Error (ME) in the bottom plots. Here the RMSE shows how accurately the localisation model is able to forecast the local ambient air temperature, and the ME shows how biased the forecasts are.

The plots in Figure 4 demonstrate how the suggested method for locally adjusting forecasts produces forecasts that are better suited for use in heat demand forecasting than the original NWP. The plots of the mean error confirm what was seen in Figure 3, i.e. that the NWP is biased and underestimates the temperature in the city due to the UHI effect. However, all methods that were tested successfully removed the bias in the weather forecast, as seen by the mean error being almost zero for all horizons both in and out-of-sample. The plots in Figure 4 also show the importance of adaptive parameter estimation, since the *rls* method with optimal selection of the forgetting factor shows the best performance.

The next step is to examine of the localised weather forecast affects heat demand forecasts. Here, the optimal model found for forecasting the heat demand was,

$$\begin{aligned}\hat{y}_{\mathrm{t+k\,|\,t}} &= \beta_{0,\mathrm{k}} + \beta_{1,\mathrm{k}} y_{\mathrm{t}} + \mu_{\mathrm{k}}(t,\ n_{\mathrm{har}},\ \alpha_{\mathrm{diu}}) + \beta_{3,\mathrm{k}} H(q) T^{\mathrm{obs,nwp}}_{a,t+k|t} \\ &\quad + \beta_{4,\mathrm{k}} H(q) W^{\mathrm{nwp}}_{s,t+k|t} + \beta_{5,\mathrm{k}} H(q) G^{\mathrm{nwp}}_{t+k|t},\end{aligned} \tag{2}$$

where β are the coefficients of the model, $T^{\mathrm{obs,nwp}}_{a,t+k|t}$ is a combined sequence of measured and forecast ambient air temperature (°C) including current measurements and NWP of the ambientair temperature, i.e.$T^{\mathrm{obs,nwp}}_{a,t+k|t} = \{\ldots, T^{\mathrm{obs}}_{a,t-1}, T^{\mathrm{obs}}_{a,t}, T^{\mathrm{nwp}}_{a,t+1|t}, T^{\mathrm{nwp}}_{a,t+2|t}, \ldots, T^{\mathrm{nwp}}_{a,t+K|t}\}$, $W^{\mathrm{nwp}}_{s,t+k|t}$ is the NWP of winds peed (m/s), $G^{\mathrm{nwp}}_{t+k|t}$ is the NWP of global radiation (W/m2), $\mu_{\mathrm{k}}(t, n_{\mathrm{har}}, \alpha_{\mathrm{diu}})$ describes the diurnal curve using Fourier harmonic series where t is the time of the day (hour), n_{har} is the number of harmonics, chosen to be $n_{\mathrm{har}} = 3$, and finally α_{diu} is a vector consisting of the coefficients for the included harmonics. H(q) is a transfer function that acts as a low-pass filter with a stationary gain equal to one,

$$H(q) = \frac{1-a}{1-q^{-1}}, \tag{3}$$

where a is the time constant describing how the buildings are affected by changes in the corresponding climate variable, e.g. ambient air temperature.

The model is fitted to the data using the recursive least squares method, with the optimal forgetting factor being $\lambda = 0.9929$. Figure 5 shows the RMSE for each horizon obtained when fitting the model to the data with the raw NWP.

It is clear from Figure 5 that the model performs quite well. The errors jump after the first two hours and then increase more slowly for the higher horizons.

The model coefficients and the optimal low pass filter parameters of this fit are reported in Table 1. Since the fit is made using recursive least squares, and the coefficients thus change in each time step only the mean value is given. The standard deviation listed here are of the entire parameter trace, and hence cannot be used for significance tests.

This is of course not a perfect model, and improvements could still be made. However, the purpose of this model is only to provide a baseline for comparing the effect of localisation.

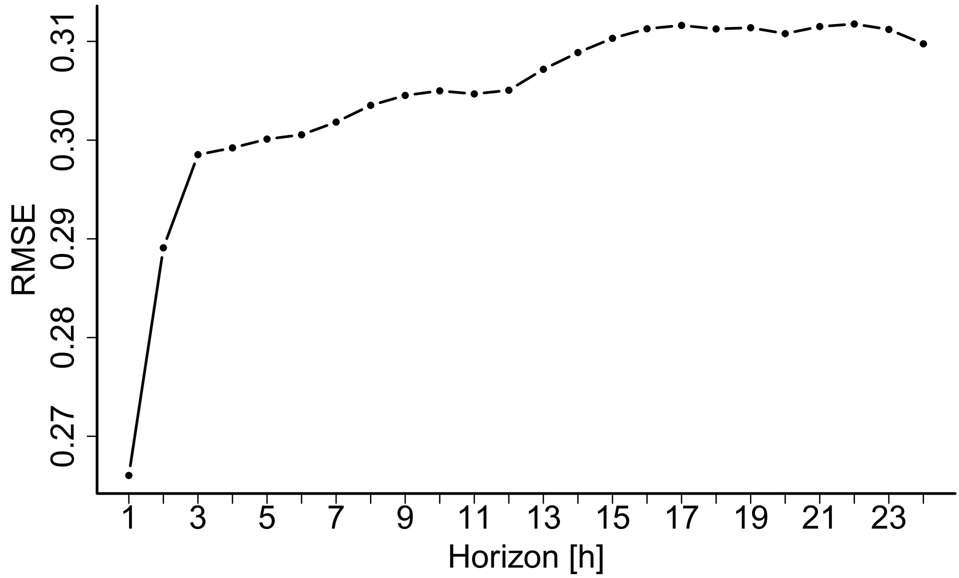

Figure 5. RMSE for each horizon of the model with the raw NWP.

The four NWPs of the ambient air temperature are then used in Equation (2) to forecast the heat demand in Tingbjerg. To evaluate the accuracy improvements of the heat demand forecast, the Relative Root Mean Square Error (RRMSE) is examined. This compares the RMSE of the heat demand forecast with localised weather forecasts to the heat demand forecast using the raw NWP as input $\text{RMSE}_{\text{BASE}}$,

$$\text{RRMSE} = \left(\frac{\text{RMSE}}{\text{RMSE}_{\text{BASE}}} - 1\right) * 100\%. \tag{4}$$

Here negative values indicate that the localised forecast has improved the accuracy of the heat demand forecast.

The results in Figure 6 show that all three localised forecasts improved the accuracy of the heat demand forecasts when compared to the raw NWP model. The *rls* method of localisation, which adapts faster, has the best performance overall, especially for horizons up to 12 hours ahead, which are the most important horizons for temperature optimisation, where the objective is to lower the supply temperature. For the last horizons, the *rls* performs similarly to the other localised forecasts, although it still demonstrates around 1.5% higher accuracy than using the raw NWPs as input.

Table 1. Left: Mean coefficient values and the standard deviation of their trace.

Coefficient	Mean value	Standard Deviation
β_0	4.900	1.800
β_1	0.130	0.170
β_3	-0.120	0.079
β_4	0.019	0.026
β_5	-0.014	0.012

Parameter	Value
a_T	0.760
a_W	0.760
a_G	0.999

Note: Right: Optimal low-pass filter parameters for the fit.

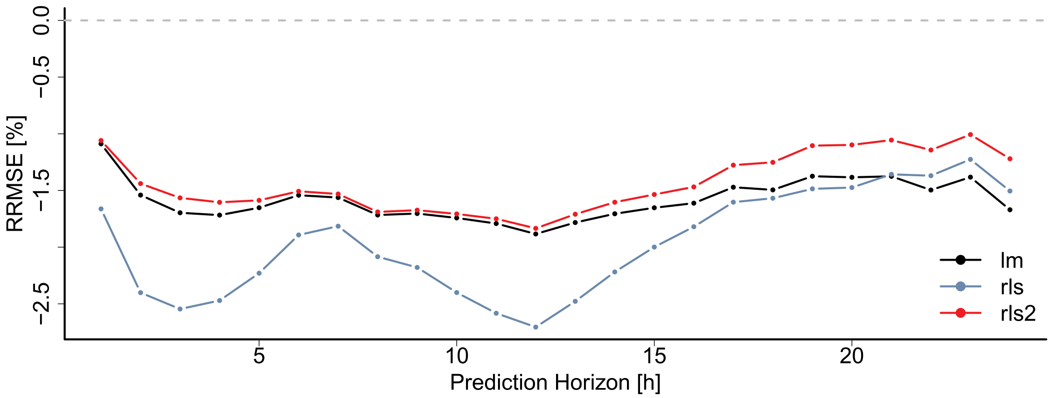

Figure 6. Accuracy improvements in heat demand forecast by localising the meteorological forecast of the ambient air temperature to a local climate station. The RRMSE for the three localised methods compared against the forecasting model with the raw NWP as inputs.

5. Conclusion

In this study, we demonstrated that the numerical weather prediction from ECMWF delivers a biased forecast of the temperature for the local climate station in Tingbjerg, and on average it predicts a lower temperature than is measured locally. This bias is likely due to the urban heat island effect. Therefore, we proposed adjusting the raw numerical weather prediction to measurements of the temperature from a local climate station in Copenhagen. This reduces the bias and gives a more accurate representation of the local climate.

To make this adjustment, we proposed using the model output statistics technique to calibrate weather forecasts, as this method has seen a lot of use in other fields of forecasting. Three different methods were proposed to locally adjust the numerical weather prediction, one where the parameters are constant and two methods where the parameters vary over time with different forgetting factors. Each of these methods effectively removes the bias in the weather forecast. Of the three methods the two adaptive methods result in forecasts most closely resembling the local weather measurements.

Furthermore, it was investigated how locally adjusting weather forecasts affects the accuracy of heat load forecasts. For this purpose, a simple adaptive heat demand model was constructed which uses the ambient temperature forecasts as an explanatory variable. With this model, the unadjusted weather forecast can be compared to the three locally adjusted forecasts in terms of the accuracy of the resulting heat load forecast in Tingbjerg (Copenhagen). The results showed that using local climate data to adjust temperature forecasts outperforms the forecasts based solely on the raw numerical weather prediction for all forecasting horizons. This suggests that having a better understanding and forecasting of the local climate will improve the heat demand forecast, which in turn will contribute to cost savings achieved due to better operation of the district heating system.

We conclude that, in order to optimise the operation of the district heating systems, it would be beneficial to install and use a local climate station to localise the weather forecast to enhance the heat demand forecasts.

Notes

1. Figure created using Google Maps.
2. European Centre for Medium-Range Weather Forecasts (ECMWF), HRES – High-Resolution Forecast.

Acknowledgments

The authors would like to thank HOFOR for making their data available and for their support.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work is funded by the Innovation Fund Denmark through the project HEAT 4.0 (8090-00046B), TOP-UP (9045-00017B), FED (8090-00069B), and SEM4Cities (0143-00032B). The project 'IEA HTP Annex – IoT Annex – Digitalization and IoT for Heat Pumps' (EUDP – 64020-2071) and the IDASC (18012745) project funded by the Region Hovedstaden and the IEA TS4 annex on Digitalisation of District Heating: Optimised Operation (and Maintenance) of District heating and Cooling Schemes via Digital Processes Management funded by the EUDP Denmark, have also supported this work.

References

Aalborg Nielsen, H., and H. Madsen. 2000. "Predicting the Heat Consumption in District Heating Systems Using Meteorological Forecasts."

Bacher, P., H. G. Bergsteinsson, L. Frölke, M. L. Sørensen, J. Lemos-Vinasco, J. Liisberg, J. K. Møller, H. A. Nielsen, and H. Madsen. 2021. "Onlineforecast: An R Package for Adaptive and Recursive Forecasting." Preprint arXiv:2109.12915.

Bacher, P., H. Madsen, H. A. Nielsen, and B. Perers. 2013. "Short-Term Heat Load Forecasting for Single Family Houses." *Energy and Buildings* 65:101–112. https://doi.org/10.1016/j.enbuild.2013.04.022.

Baklanov, A., G. Sue, M. Alexander, and M. Athanassiadou. 2009. *Meteorological and Air Quality Models for Urban Areas*. Vol. 184. Berlin: Springer.

Bauer, P., A. Thorpe, and G. Brunet. 2015. "The Quiet Revolution of Numerical Weather Prediction." *Nature* 525 (7567): 47–55. https://doi.org/10.1038/nature14956.

Bergsteinsson, H., S. Ben Amer, P. Nielsen, and H. Madsen. 2021. "Digitalization of District Heating." Technical University of Denmark. https://orbit.dtu.dk/en/publications/digitalization-of-district-heating.

Blanco, I., D. Guericke, A. N. Andersen, and H. Madsen. 2018. "Operational Planning and Bidding for District Heating Systems with Uncertain Renewable Energy Production." *Energies* 11 (12): 3310. https://doi.org/10.3390/en11123310.

Crochet, P. 2004. "Adaptive Kalman Filtering of 2-Metre Temperature and 10-Metre Wind-Speed Forecasts in Iceland." *Meteorological Applications* 11 (2): 173–187. https://doi.org/10.1017/S1350482704001252.

Dahl, M., A. Brun, O. S. Kirsebom, and G. B. Andresen. 2018. "Improving Short-Term Heat Load Forecasts with Calendar and Holiday Data." *Energies* 11 (7): 1678. https://doi.org/10.3390/en11071678.

Danish Meteorological Institute (DMI). 2021. "Meteorological Observation (metObs)." Accessed May 24, 2021. https://confluence.govcloud.dk/display/FDAPI/Danish+Meteorological+Institute+-+Open+Data.

Dotzauer, E. 2002. "Simple Model for Prediction of Loads in District – Heating Systems." *Applied Energy*73 (3-4): 277–284. https://doi.org/10.1016/S0306-2619(02)00078-8.

Frederiksen, S., and S. Werner. 2013. *District Heating and Cooling*. Lund: Studentlitteratur AB.

Giebel, G., and M. Denhard. 2017. "Short-Term Wind Power Forecasting." In *Electric Power Systems: Advanced Forecasting Techniques and Optimal Generation Scheduling*, 5–37. CRC Press.

Glahn, H. R., and D. A. Lowry. 1972. "The Use of Model Output Statistics (MOS) in Objective Weather Forecasting." *Journal of Applied Meteorology and Climatology* 11 (8): 1203–1211. https://doi.org/10.1175/1520-0450(1972)011¡1203:TUOMOS¿2.0.CO;2.

Gneiting, T. 2014. "Calibration of Medium-Range Weather Forecasts." European Centre for Medium-Range Weather Forecasts Reading, UK.

Gneiting, T., S. Lerch, and B. Schulz. 2023. "Probabilistic Solar Forecasting: Benchmarks, Post-Processing, Verification." *Solar Energy* 252: 72–80. https://doi.org/10.1016/j.solener.2022.12.054.

Gneiting, T., A. E. Raftery, A. H. Westveld, and T. Goldman. 2005. "Calibrated Probabilistic Forecasting Using Ensemble Model Output Statistics and Minimum CRPS Estimation." *Monthly Weather Review*133 (5): 1098–1118. https://doi.org/10.1175/MWR2904.1.

Graham, R. M., J. Browell, D. Bertram, and C. J. White. 2022. "The Application of Sub-Seasonal to Seasonal (S2S) Predictions for Hydropower Forecasting." *Meteorological Applications* 29 (1): e2047. https://doi.org/10.1002/met.v29.1.

Idowu, S., S. Saguna, C. Åhlund, and O. Schelén. 2016. “Applied Machine Learning: Forecasting Heat Load in District Heating System.” *Energy and Buildings* 133:478–488. https://doi.org/10.1016/j.enbuild.2016.09.068.

Madsen, H., K. Sejling, H. T. Søgaard, and O. P. Palsson. 1994. “On Flow and Supply Temperature Control in District Heating Systems.” *Heat Recovery Systems and CHP* 14 (6): 613–620. https://doi.org/10.1016/0890-4332(94)90031-0.

Ronda, R., G. Steeneveld, B. Heusinkveld, J. Attema, and A. Holtslag. 2017. “Urban Finescale Forecasting Reveals Weather Conditions with Unprecedented Detail.” *Bulletin of the American Meteorological Society* 98 (12): 2675–2688. https://doi.org/10.1175/BAMS-D-16-0297.1.