

Article

Building Heat Demand Forecasting by Training a Common Machine Learning Model with Physics-Based Simulator

Lotta Kannari, Jussi Kiljander ^{*}, Kalevi Piira, Jouko Piippo and Pekka Koponen

VTT Technical Research Centre of Finland, 02044 VTT Espoo, Finland; lotta.kannari@vtt.fi (L.K.); kalevi.piira@vtt.fi (K.P.); jouko.piippo@vtt.fi (J.P.); pekka.koponen@vtt.fi (P.K.)

* Correspondence: Jussi.Kiljander@vtt.fi

Abstract: Accurate short-term forecasts of building energy consumption are necessary for profitable demand response. Short-term forecasting methods can be roughly classified into physics-based modelling and data-based modelling. Both of these approaches have their advantages and disadvantages and it would be therefore ideal to combine them. This paper proposes a novel approach that allows us to combine the best parts of physics-based modelling and machine learning while avoiding many of their drawbacks. A key idea in the approach is to provide a variety of building parameters as input for an Artificial Neural Network (ANN) and train the model with data from a large group of simulated buildings. The hypothesis is that this forces the ANN model to learn the underlying simulation model-based physics, and thus enables the ANN model to be used in place of the simulator. The advantages of this type of model is the combination of robustness and accuracy from a high-detail physics-based model with the inference speed, ease of deployment, and support for gradient based optimization provided by the ANN model. To evaluate the approach, an ANN model was developed and trained with simulated data from 900–11,700 buildings, including equal distribution of office buildings, apartment buildings, and detached houses. The performance of the ANN model was evaluated with a test set consisting of 60 buildings (20 buildings for each category). The normalized root mean square errors (NRMSE) were on average 0.050, 0.026, 0.052 for apartment buildings, office buildings, and detached houses, respectively. The results show that the model was able to approximate the simulator with good accuracy also outside of the training data distribution and generalize to new buildings in new geographical locations without any building specific heat demand data.



Citation: Kannari, L.; Kiljander, J.; Piira, K.; Piippo, J.; Koponen, P. Building Heat Demand Forecasting by Training a Common Machine Learning Model with Physics-Based Simulator. *Forecasting* **2021**, *3*, 290–302. <https://doi.org/10.3390/forecast3020019>

Academic Editors: Cong Feng and Ted Soubdhan

Received: 12 March 2021

Accepted: 16 April 2021

Published: 21 April 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: building energy modelling; machine learning; artificial neural networks; demand response; short-term forecasting; simulation

1. Introduction

The energy systems increasingly need demand response for balancing and reduction of environmental impacts such as greenhouse gas emissions. Buildings have excellent demand response potential [1–3]. Tapping this potential requires the ability to accurately forecast short-term energy demand, its flexibility, and the load control responses. Accurate modeling and forecasting are essential to utilize a model-based optimal control for demand response and peak demand shaving.

Reviews of building energy modelling for control and operation, by Xiwang and Ji [4], and by Harish and Kumar [5], made the following conclusions. First, detailed physical models have high accuracy, but are difficult to utilize in an on-line building operation because they have many parameters and require large computation time and power. Second, statistics and machine learning (ML) models are fast, and their accuracy good enough for model-based control, but these methods require a large amount of training data that covers the building operation range. Third, reducing the computational cost and memory demand for building energy modeling and optimal control, while maintaining the accuracy, is an urgent issue for on-line practical applications.

Based on the analysis above, it would be beneficial to develop methods that combine the best parts of physical and statistical modelling while avoiding their drawbacks. To elaborate, high-detail physical models based in tools such as Modelica [6], EnergyPlus [7], and IDA [8] require a lot of manual work, making it difficult to deploy and utilize them at large scale in a cost-efficient way. This type of physical models are also computationally heavy, and it is therefore difficult to utilize them in model-based optimization. Moreover, their deployment to production use is also a challenge due to the simulation tools used to implement these models. On the other hand, machine learning-based methods are scalable, accurate, and require less human effort in the modelling [9–13]. Their deployment is also easier and supported by existing tools and infrastructure [14–16]. Machine learning models also typically provide fast enough inference for model-based optimization. Moreover, gradient-based methods such as neural networks make it possible to utilize gradient information and thus converge much faster than gradient-free methods [17–19]. However, ML models also have their limitations. They typically require a lot of data to provide good results and even with large amount of training data, it is impossible to cover all situations that need to be forecast [20]. This makes it risky to utilize machine learning methods as they can produce large errors in exceptional situations.

There are several ways to combine physical-based modelling with machine learning. For example, Koponen et al. [21] developed and studied combinations of several model hybridization approaches such as residential hybrid, constraining model, and physically based input forecasts. Another typical way to combine physics-based modelling and ML is to utilize a physics-based simulator to generate training data for a ML model [22–28]. This is a natural way to combine physics-based modelling and ML, because a main limitation of ML is the lack of training data on all relevant situations. This is especially important in situations where the ML model is used for demand-side management. A common feature of existing studies that use a physics-based simulator for generating training data is that a separate ML model is trained for every building. These type of approaches can combine some of the benefits of physic- and ML-based modelling, but some drawbacks still remain. To elaborate, these types of approaches make it possible to utilize faster and more easily deployable ML models in model-based control, but still require manual work for building modelling. This is because the ML model is only trained for one type of building and it cannot generalize to new buildings. For the same reason, the ML model does not have to learn to approximate the underlying physics encoded into the simulation models, and therefore the model is still not able to perform well in situations not covered by the training data.

To tackle these limitations of the existing work, this paper proposes a novel approach for combining physics-based simulator and ML based modelling. The key idea and main novelty in the proposed approach is that instead of training a separate ML model for each building we use simulated data from a large pool of buildings to train a single Artificial Neural Network (ANN). In this way, the ML models must learn to approximate the physics of the simulation model in order to generalize to new buildings and weather conditions. This approach makes it possible to forecast buildings energy demand just based on building's characteristics and thus makes it possible to use ML models in situations with no energy consumption data of a specific building. Moreover, the main benefit of the proposed approach is that once the ANN model has been trained, we can completely replace the original simulator. This is important since it makes it possible to have the benefits of ANNs (i.e., ease of deployment, fast inference and gradient-based optimization in model-predictive control) while having the benefits of physics-based models at the same time.

To evaluate the proposed approach, we developed a Feed Forward Neural Network (FFNN) to forecast buildings heat demand based on building characteristics, weather, and temporal data. We used a dynamic hourly based energy model, based on EN ISO 13790:2008 and EN 15241:2007 standards, to generate training, validation, and test data sets. The generated data sets consist of three types of buildings with different parameter

combinations for building construction year, dimensions, ventilation heat recovery, and geographical location. The model was evaluated with a test set consisting of a full year of data from 60 buildings (20 from each building type).

The rest of the paper is structured as follows. Section 2 introduces the novel approach proposed in the paper. It describes the original physical heating and cooling simulator, generated datasets, and the neural network developed for the task. Section 3 presents the validation and results of the study. In Section 4, there is discussion about the results and ideas for future work. Section 5 concludes the paper.

2. Methodology and Approach

2.1. Overview

The main aim of the approach was to develop an ANN model that could be trained to represent the knowledge of an existing building simulator, the presuppositions usually made for the modelled buildings and the post processing of the simulator output. It should be clarified, that the physics-based simulator is not directly utilized in the final model, but instead used for producing training data for the ANN model. A key idea in the approach is also that the experiment is designed in a way that the ANN model has to learn to approximate the simulator in order to perform well. This includes, for example, using different buildings and geographical locations (with different weather conditions) in the training and test data sets so that the model must be able to generalize to achieve good accuracy. An overview of the methodology is illustrated in Figure 1.

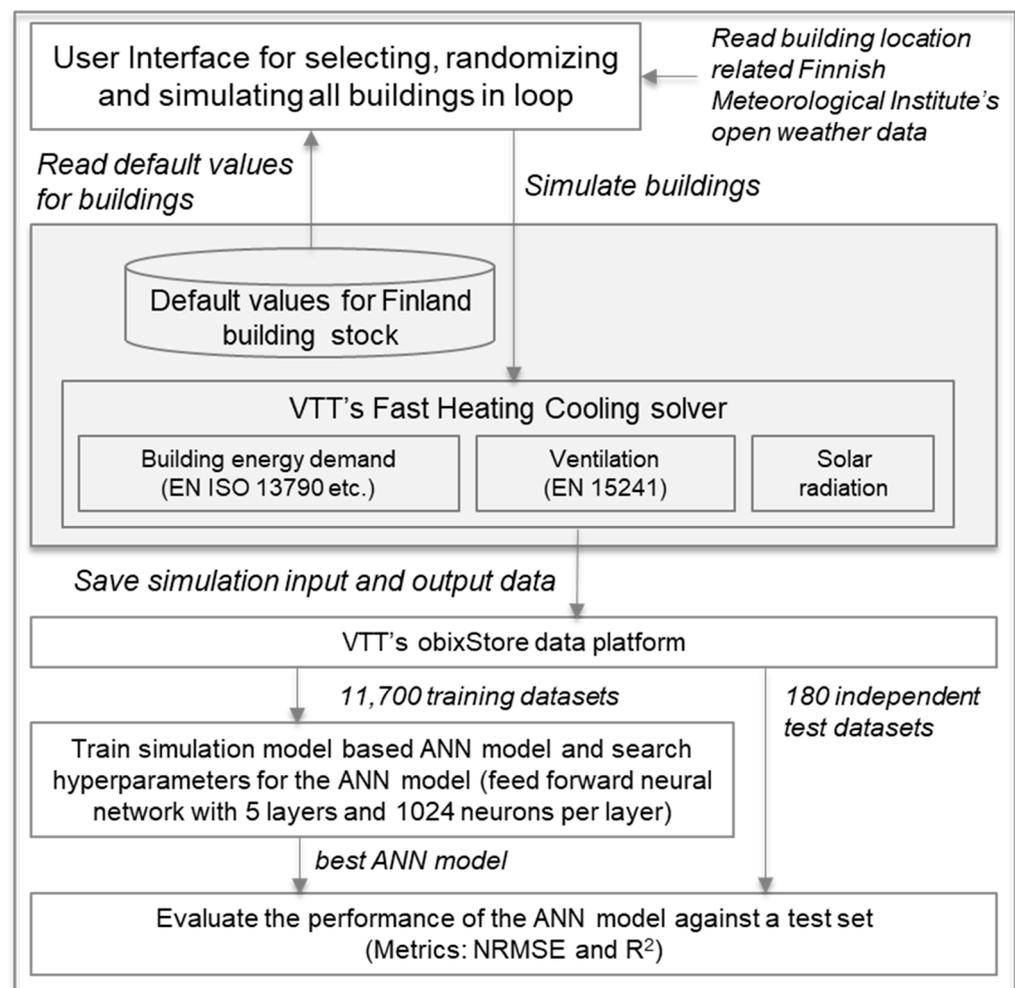


Figure 1. An overview of the approach and methodology.

The main stages of the developing process were as follows. First, typical building parameters for the selected building types (office buildings, apartment buildings, and one family houses) were fetched from VTT's Finland building stock default value database. The selected building model input parameters (e.g., floor areas, number of floors) were randomized to create as much variety to the training data as possible. Second, location related weather data was collected from Finnish Meteorological Institute's (FMI) open weather data platform via RESTful Application Programming Interface (API). Third, the randomized buildings were simulated by calling VTT's Fast Heating Cooling solver in a loop via a RESTful API. Both simulation inputs and results were saved as a three different training sets (includes validation set for hyperparameter tuning) (11,700 simulations) and a test set (180 independent simulations). Fourth, based on the training datasets, different ANN models with different number of layers and neurons per layer were trained for each training set and the model that performed best with a validation set was selected for final evaluation. The model architecture tuning, along with other hyperparameters were searched with grid search using 10% of the training data as a validation set. It should be noted, that the main point of the paper was not to find the most optimal ANN architecture for the task, but to study and evaluate whether an ANN model learns to approximate the physics-based model. Fifth, the best ANN model was evaluated using an independent test set. These steps are described more detailed in Sections 2.2–2.6.

2.2. Assumptions

It should be emphasized, that the methodology proposed in the paper assumes that a building simulator can accurately estimate building's energy consumption in different conditions. It is especially important that the simulator can predict the response of a building better than a machine learning model in situations not covered by the training data. Examples of such situations include, for instance, new buildings, extreme weather conditions, and most importantly demand response (i.e., when the building's heating is actively controller with respect to market signals). This assumption is backed up by existing work [4,5], but nevertheless should be carefully considered in future research.

2.3. Building Energy Simulation Models

The dynamic hourly based building energy demand simulations are performed with VTT's Fast Heating-Cooling demand solver (Fast HC solver). It is based on EN ISO 13790:2008 (Energy performance of buildings: Calculation of energy use for space heating and cooling) and EN 15241:2007 (Ventilation for buildings: Calculation methods for energy losses due to ventilation and infiltration in buildings) standards and dynamic models for estimating solar radiation. The model considers heat losses through building envelope, ventilation, and air leakages, as well as heat gains from appliances, occupancy, and solar radiation.

The model requires a lot of detailed information of a simulated building. If all input parameters are not known, it can use representative values from VTT's default value database [29]. The database includes typical values for different ages of Finnish office buildings, apartment buildings and one family houses, such as building structures thermal transmittances (U-values), ventilation details and user profiles. Requirements from U-values for each construction year come from the Finnish building regulations. User profiles define running times for ventilation, and schedules for heat gains from people and appliances, these are specific for each building type.

2.4. Dataset Generation

Training and testing data are generated with the simulator described in Section 2.2. Different building parameters are chosen randomly from a realistic range of each building type (apartment buildings, office buildings, one family houses). The data resolution for energy demand and weather measurements is 60-min.

Weather measurements from FMI of three different locations are utilized in the simulations. These locations represent various temperature zones of Finland. Helsinki is located in southern Finland by the seaside, Jyväskylä in the middle, and Sodankylä in the northern part of the country. Weather measurements are utilized instead of forecast because the goal is to study whether the physics-based simulator can be approximated by an ANN model. FMI does not also provide historical data on forecasts, just on measurements.

Three separate training sets and one test set are generated with the following configurations:

- Training set 1:
 - Three year simulation for three locations (Helsinki, Jyväskylä, Sodankylä), 100 buildings per type per location.
 - Weather data includes years 2016–2018.
 - Contains a total of 900 buildings.
 - Training set is divided so that 90% is used for training and 10% for validation.
- Training set 2:
 - One month simulation for Helsinki only, 300 buildings per building type.
 - Simulated months are sampled randomly from 2015 weather data.
 - Contains a total of 10,800 buildings.
 - Training set is divided so that 90% is used for training and 10% for validation.
- Training set 3:
 - Training set 1 and 2 combined (added together as such).
- Test set:
 - One year simulation for each of the three locations, 20 building per building type.
 - Year 2019 measurements are used as the weather data.
 - Contains a total of 180 buildings.

2.5. Feed Forward Neural Network

In order to find a good ANN model for the task, a FFNN with different number of layers, and neurons per layer were trained and evaluated with the three training (and validation) datasets. Based on the average RMSE of the validation sets, the best model was an FFNN with five layers and 1024 neurons per layer. Features were normalized with MinMaxScaler, Rectified Linear Unit (RELU) was used as the activation function for hidden layers, and Adam [30] as the optimization method. Table 1 summarizes the main attributes of the FFNN.

Table 1. Key attributes of the FFNN.

Attribute	Value
Number of layers	5
Number of neurons in total	5121
Activation function for hidden layers	RELU
Optimization method	Adam [22]
Input scaling	MinMaxScaler

The model has identical input features to the original simulator. These 12 features include: day of the year, hour of the day, outside temperature, direct solar radiation, diffuse solar radiation, construction year, building type, floor count, cross floor area, building shape, heat recovery efficiency for ventilation, and heat capacity. Sin and cos values for both time features was used. One-hot encoding was utilized for building type (apartment, office, single family), building shape (rectangle, U-shape, L-shape, closed block, between two buildings), and heat capacity (light, medium, heavy). As typically building properties

are related to the existing regulations, construction year is one-hot encoded to match known time spans of different regulations.

In order to perform well in the task, the FFNN model needs to learn both the physical building model and the associated default value database. Heating plays a much bigger role than cooling in Finland due to the relatively short cooling season. For this reason, space heating was chosen as an output for the model. Single family homes have a significantly smaller heating need, and with the current model, they had to be trained separately from the other building types.

2.6. Metrics

Normalized root-mean-square error (NRMSE) and coefficient of determination (R^2) are calculated for each building in the test set separately. In the results, the average of these values are presented based on the building type and location. Formulas used for calculating the metrics are presented in Equations (1)–(3).

$$R^2 = \frac{\sum(\hat{y}_i - \bar{y})^2}{\sum(y_i - \bar{y})^2} \quad (1)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (2)$$

$$\text{NRMSE} = \frac{\text{RMSE}}{y_{\max} - y_{\min}} \quad (3)$$

where y_i is the observation, \bar{y} is its mean, and \hat{y}_i is predicted value, y_{\max} and y_{\min} are the maximum and minimum values of y , respectively.

It should be noted that the paper does not use the typical lead time, forecasting horizon, and update rate parameters. This is because the proposed ANN does not utilize energy consumption of the building as input, and since we utilize measured weather conditions during the evaluations (instead of weather forecasts) lead time, horizon, and update rate do not affect the accuracy. These parameters are of course important when the impact of the weather forecast accuracy to the forecast accuracy is analyzed. However, since the main point of the paper is to study whether the physics-based simulator can be approximated, it is better to use the actual measurements so that unnecessary noise is not included into the evaluation.

3. Results

Tables 2 and 3 show the performance of the ANN model separately for the different building types using the standard metrics introduced in Section 2.5. Office buildings have a bit better performance than the other building types. There is only small differences in the performance between the different training sets.

As illustrated in Tables 4 and 5, the forecasting results are also quite similar for all three geographical locations. The accuracy in Jyväskylä, located in the middle of Finland, is slightly better than in the other two locations. Even absence of colder locations in the training data (training set 2) does not affect the results significantly (see values marked with * (No buildings from this location were present in the training data). in the following tables).

Table 2. NRMSE in the test set for different building types.

Training Set	Apartment Buildings	Offices	Single Family Houses
(1) training set 1	0.064	0.031	0.059
(2) training set 2	0.063	0.035	0.060
(3) training set 3	0.058	0.026	0.052

Table 3. R^2 in the test set for different building types.

Training Set	Apartment Buildings	Offices	Single Family Houses
(1) training set 1	0.89	0.95	0.92
(2) training set 2	0.91	0.93	0.93
(3) training set 3	0.92	0.96	0.94

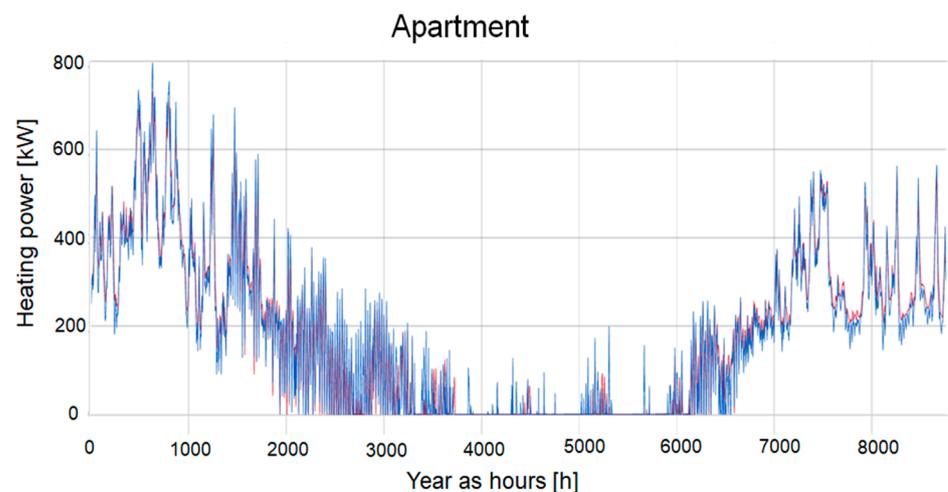
Table 4. NRMSE in the test set for different geographical locations.

Training Set	Helsinki	Jyväskylä	Sodankylä
(1) training set 1	0.053	0.048	0.054
(2) training set 2	0.051	0.050 *	0.057 *
(3) training set 3	0.048	0.042	0.047

Table 5. R^2 in the test set for different geographical locations.

Training Set	Helsinki	Jyväskylä	Sodankylä
(1) training set 1	0.91	0.93	0.92
(2) training set 2	0.92	0.93 *	0.93 *
(3) training set 3	0.94	0.95	0.95

Figures 2 and 3 visualize the forecast against measurement data for different building types (example building were selected randomly). Real energy demand values are presented in red and the forecasted values in blue. Figure 2 includes whole test year. Figure 3 illustrates the heat demand during a two week period that was sampled randomly from the heating period. As can be seen from the figures, the model seems to learn well also the periodic schedules defined in the original default value database. This can be seen especially from the figures depicting office buildings (e.g., the middle chart in Figure 3), where the pattern from ventilation schedule is most visible (the ventilation is controlled more effective during office hours and less effective during evenings and weekends).

**Figure 2.** *Cont.*

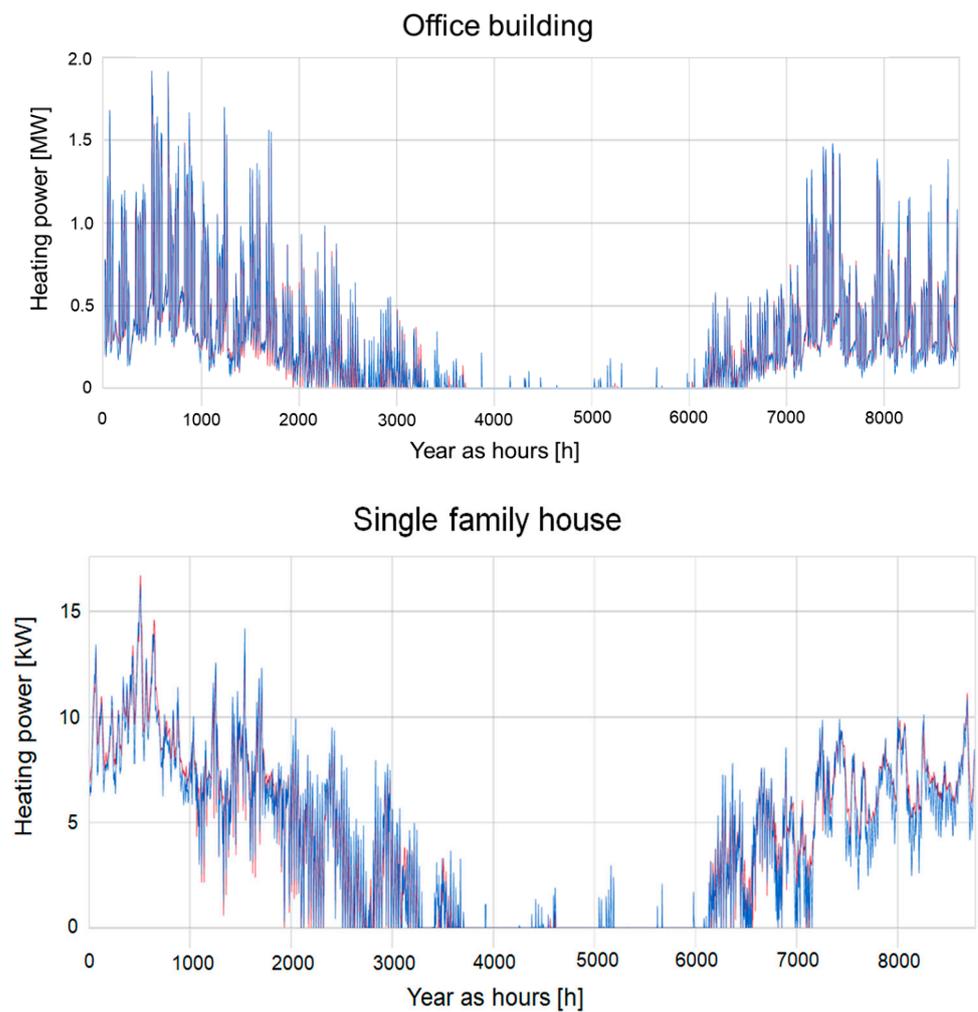


Figure 2. One year examples of measured (red) and forecast (blue) heating load for different building types (apartment, office, single family home).

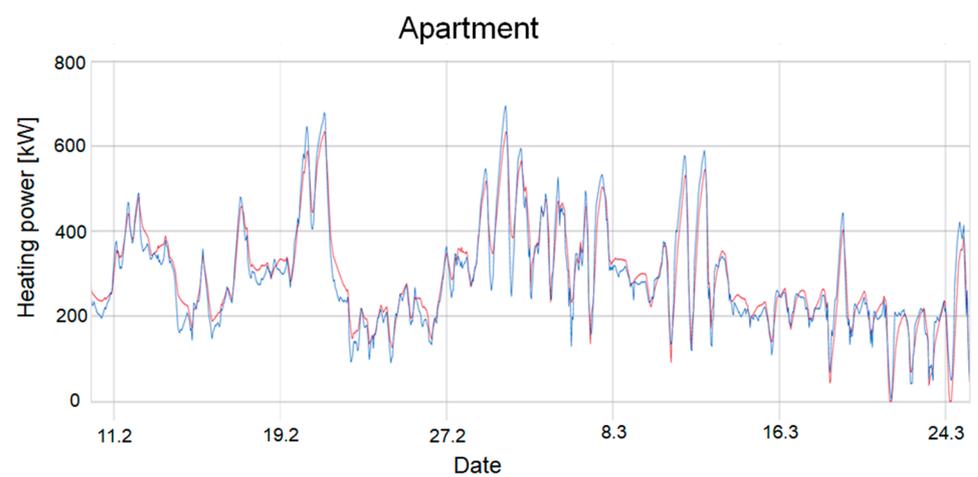


Figure 3. *Cont.*

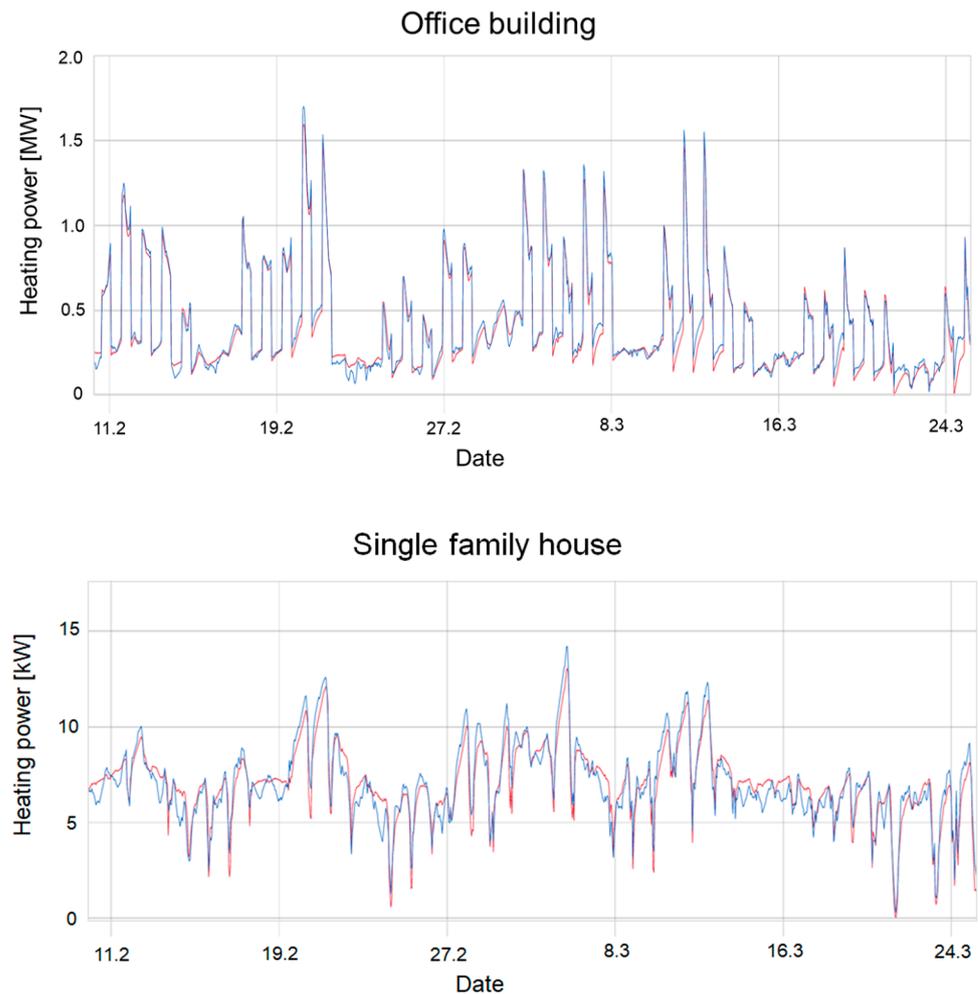


Figure 3. Six week examples of measured (red) and forecast (blue) heating load for different building types (apartment, office, single family home).

4. Discussion

The main purpose of the study was to evaluate whether a ML model can be trained to approximate a physics-base simulator so that we can combine the best parts of both worlds. The results show that this is possible, as the proposed FFNN model achieved accurate forecast with all building types. The best accuracy with NRMSE of 0.026 was achieved with office buildings. The model NRMSE with apartment buildings and single family houses were 0.050 and 0.052, respectively.

The model performed well with all three training sets and in all geographical locations. This is especially interesting in the case of training set 2 for two reasons. First, the model had only one month of training data from each building. Second, all the buildings in training set 2 were in the Helsinki area, which is located 900 km south from Sodankylä. Helsinki and Sodankylä have naturally quite different outdoor temperature conditions during the winter. Additionally, the angle of solar radiation is different. However, solar radiation angles have typically greater impact during summer than during the heating season. Nevertheless, the results indicate that the ANN model is able to learn to approximate the physics-based simulator, as it is able to extrapolate the heat demand also to new type of weather conditions.

The main benefits of the proposed approach is that it makes it possible to replace the original simulator and the associated pre- and post-processing computations with a neural network. This makes it much easier to deploy and use the model in model-based optimization, which is important in optimal energy management and demand

response. When compared to other ML-based methods proposed in the literature, the main advantages of the proposed approach are (1) that it can work without any energy measurement data (important, e.g., with new buildings), and (2) that it is more robust with respect to situations not seen in the training data due to the fact that the model has to learn to approximate the underlying physics to perform well. A good example of this robustness is that the model also achieved good performance with buildings from Sodankylä even though it was trained just with buildings from Helsinki (training set 2).

When compared to other approaches [22–28] where a physics-based simulator is used to train a machine learning model, the main difference is that in all of the previous approaches, a separate ML model is trained for each building, whereas the approach proposed in this paper is generic for all buildings. This means that the ANN model can completely replace the physics-based simulator. The results also show that the model is able to generalize to inputs not seen in the training data, which is a positive sign that the model has learned to approximate the physics encoded into the simulator. Clear evidence of the generalization capability is that the model worked well with buildings it had not seen before. Moreover, the model was able to forecast energy consumption of colder locations (Jyväskylä and Sodankylä), although training set 2 included data only from Helsinki.

It should be emphasized that the approach proposed in the paper is not restricted to ANNs. In theory, any generic function approximator could be used instead of an ANN to approximate the physics-based simulator. An ANN model was selected for this study mainly for following reasons: (1) ANN models are highly parallel and thus provide fast inference, (2) ANN models provide gradient based optimization through back propagation, which is important when the model is used in planning and control [17–19], and (3) ANN platforms such as TensorFlow also provide good support for model deployment in edge environments [14,15].

There are also limitations in the current study, and future research is needed. For instance, the current model is strictly tied to Finland due to the default value database used in the generation of simulation inputs that are typical for Finnish buildings. However, this limitation can be easily tackled by training the ANN model with a default value database that covers other types of buildings as well. Additionally, the current approach was only validated against a simulated dataset. Therefore, scientific studies are needed to validate the approach also with real building data. In this context, it would be interesting to study whether the ANN model could be fine-tuned with building specific data. This is important since there is always some differences between a real building and a simulated one. Additionally, the behavior of people can influence the heat demand and cause differences between the simulation model and the real building. In practice, this fine-tuning could be done in many ways. For instance, one approach would be to re-train the whole ANN model carefully (i.e., smaller learning rate to avoid overfitting) with building specific heat demand data. Another interesting, and more data-efficient, approach would be to freeze the original ANN model and use it to fine-tune the building-specific input parameters by backpropagating gradients through the model.

It should also be noted, that the study was implemented with weather measurement data instead of forecast. Errors in weather forecast naturally cause errors in the forecast made by the ANN model. However, the main purpose of the paper is to show that the physics-based model can be approximated with an ANN (with loss of some accuracy) in order to gain the benefits of ANN models. In this case, it is not appropriate to include additional noise to the evaluation. Nevertheless, analyzing the impact of errors in the weather input is a relevant topic for future work.

It should be also noted that the proposed approach assumes that a high detail physics-based simulator can accurately simulate the building's energy demand as presented in the reviews by Xiwang and Ji [4], and by Harish and Kumar [5]. However, detailed physics-based simulation models are not automatically more accurate in forecasting when compared to compact and dedicated forecasting and on-line optimization models. The main reasons for this are the difficulty to calibrate the model parameters, possibilities

for overfitting, and structural inaccuracies in the models that cannot be compensated by parameter calibration. [31,32]. These are valid concerns that have to be taken into account in future research, also with the approach proposed in the paper. However, we believe that the approach can actually help to overcome these challenges. This is because the ANN model can be fine-tuned with building specific data with efficient gradient based training, which can help to overcome problems related to structural inaccuracies in the original physics-based model.

Other interesting directions for future research include (1) evaluation of the approach with different type of ANN architectures and new inputs, and (2) extending the model so that it can predict also thermal comfort inside the building as well as cooling demand of the building. In particular, having the capability to predict the indoor temperature is required if the model is used in energy efficiency and demand response and optimizations. Proper modelling of indoor temperatures might require better handling of building dynamics, e.g., by utilizing Recurrent Neural Networks (RNN) such as LSTM [33] for weather data analytics.

5. Conclusions

This paper proposed a novel approach for combining a physics-based simulator with machine learning. The key concept in the approach is to use the physics-based simulator to generate a large dataset on different types of buildings, which is used to train a single ML model. The idea is that this would force the machine learning model to better learn to approximate the physics-based simulator and not just the particularities of a single building. This way, the approach would be more robust [34] (i.e., also able to perform well in situations not covered in the training data) when compared to typical machine learning approaches. The advantages compared to physics-based models in turn are ease of deployment, inference speed, better support for model-predictive control, and the possibility to train and fine-tune the model with building specific data. To our knowledge, this is the first approach to learn the whole building simulator with a machine learning model. Previous work [22–28] has focused on learning building specific models, which is a similar problem compared to learning the behavior of a single building. The advantage of the proposed approach is that it does not require a building specific model, as the model is able to generalize to new buildings. Additionally, in our setting, the model is also forced to better approximate the physics as it needs to be able to generalize to different type of buildings and weather conditions.

In order to evaluate the approach, we implemented an ANN model that takes building parameters, in addition to the typical weather and temporal data, as input. Existing physics-based model and a database of building attributes were used for generating the training, validation, and test datasets. The datasets included equal distribution of office buildings, apartment buildings, and single family houses. The NRMSE were on average 0.026, 0.050, 0.052 for office buildings, apartment buildings, and single family houses, respectively. The results show that the ANN model is able to accurately forecast the heat demand of different type of buildings without any energy consumption data from a given building. In order to do this the machine learning model had to learn to approximate the physics-based simulator. A good example of the extrapolation capabilities of the model was that it was able to accurately forecast the energy consumption of buildings located in Sodankylä, even though only buildings in Helsinki were present in the training set. Sodankylä is located approximately 900 km north from Helsinki, and has much colder temperatures during the winter.

Although the results are promising, more research is still needed. For instance, an interesting future research direction is to investigate approaches for fine-tuning the model with building specific data. In this context, it will be interesting to compare the approach to state-of-the-art ML models and compare the accuracy and data efficiency. Moreover, extending the approach for modelling indoor temperature and cooling demand are needed.

These extensions will likely require experimentation with different type of ANN model architectures such as LSTMs to better handle the building dynamics.

Author Contributions: Conceptualization, J.P., K.P., and J.K.; methodology, L.K., J.P., K.P., and J.K.; software, L.K. and J.P.; validation, L.K. and J.P.; writing—original draft preparation, L.K., J.K., P.K., and K.P.; supervision, J.K., K.P., and P.K.; funding acquisition, P.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research is partly done in project Analytics funded by the Academy of Finland, and has also received funding from the European Union’s Horizon 2020 research and innovation program under grant agreements 957670.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The dataset used for the research is available for research purposes. Please contact Kalevi Piira (kalevi.piira@vtt.fi) for requesting access.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. Rotger-Griful, S.; Jacobsen, R.H.; Nguyen, D.; Sørensen, G. Demand response potential of ventilation systems in residential buildings. *Energy Build.* **2016**, *121*, 1–10. [CrossRef]
2. Li, W.; Yang, L.; Ji, Y.; Xu, P. Estimating demand response potential under coupled thermal inertia of building and air-conditioning system. *Energy Build.* **2019**, *182*, 19–29. [CrossRef]
3. Goy, S.; Finn, D. Estimating Demand Response Potential in Building Clusters. *Energy Procedia* **2015**, *78*, 3391–3396. [CrossRef]
4. Li, X.; Wen, J. Review of building energy modeling for control and operation. *Renew. Sustain. Energy Rev.* **2014**, *37*, 517–537. Available online: <https://www.sciencedirect.com/science/article/pii/S1364032114003815?via%3Dihub> (accessed on 20 April 2021). [CrossRef]
5. Harish, V.; Kumar, A. A review on modeling and simulation of building energy systems. *Renew. Sustain. Energy Rev.* **2016**, *56*, 1272–1292. Available online: <https://www.sciencedirect.com/science/article/pii/S1364032115014239?via%3Dihub> (accessed on 12 March 2021). [CrossRef]
6. Wetter, M.; Zuo, W.; Nouidui, T. Recent Developments of the Modelica Buildings Library for Building Heating, Ventilation and Air-Conditioning Systems. In Proceedings of the 8th Modelica Conference, Dresden, Germany, 20–22 March 2011.
7. Pang, X.; Wetter, M.; Bhattacharya, P.; Haves, P. A framework for simulation-based real-time whole building performance assessment. *Build. Environ.* **2012**, *54*, 100–108. [CrossRef]
8. Hilliäho, K.; Lahdensivu, J.; Vinha, J. Glazed space thermal simulation with IDA-ICE 4.61 software—Suitability analysis with case study. *Energy Build.* **2015**, *89*, 132–141. [CrossRef]
9. Amasyali, K.; El-Gohary, N.M. A review of data-driven building energy consumption prediction studies. *Renew. Sustain. Energy Rev.* **2018**, *81*, 1192–1205. [CrossRef]
10. Salmi, T.; Kiljander, J.; Pakkala, D. Stacked Boosters Network Architecture for Short-Term Load Forecasting in Buildings. *Energies* **2020**, *13*, 2370. [CrossRef]
11. Fang, T.; Lahdelma, R. Evaluation of a multiple linear regression model and SARIMA model in forecasting heat demand for district heating system. *Appl. Energy* **2016**, *179*, 544–552. [CrossRef]
12. Mocanu, E.; Nguyen, P.H.; Gibescu, M.; Kling, W.L. Deep learning for estimating building energy consumption. *Sustain. Energy Grids Netw.* **2016**, *6*, 91–99. [CrossRef]
13. Marino, D.L.; Amarasinghe, K.; Manic, M. Building Energy Load Forecasting Using Deep Neural Networks. In Proceedings of the IECON 2016–42nd Annual Conference of the IEEE Industrial Electronics Society, Florence, Italy, 23–26 October 2016; IEEE: Piscataway Township, NJ, USA, 2016. [CrossRef]
14. Pääkkönen, P.; Pakkala, D.; Kiljander, J.; Sarala, R. Architecture for Enabling Edge Inference via Model Transfer from Cloud Domain in a Kubernetes Environment. *Future Internet* **2021**, *13*, 5. [CrossRef]
15. Lee, S.H.; Lee, T.; Kim, S.; Park, S. Energy Consumption Prediction System Based on Deep Learning with Edge Computing. In Proceedings of the 2019 2nd International Conference on Electronics Technology, Chengdu, China, 10–13 May 2019; IEEE: Piscataway Township, NJ, USA, 2019. [CrossRef]
16. Zhou, L.; Wen, H.; Teodorescu, R.; Du, D.H.C. Distributing Deep Neural Networks with Containerized Partitions at the Edge. In Proceedings of the 2nd USENIX Workshop on Hot Topics in Edge Computing, Renton, WA, USA, 9 July 2019; HotEdge: Renton, WA, USA, 2019.

17. Kiljander, J.; Sarala, R.; Rehu, J.; Pakkala, D.; Pääkkönen, P.; Takalo-Mattila, J.; Känsälä, K. Intelligent Consumer Flexibility Management with Neural Network-Based Planning and Control. *IEEE Access* **2021**, *9*, 40755–40767. [CrossRef]
18. Afram, A.; Janabi-Sharifi, F.; Fung, A.S.; Raahemifar, K. Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. *Energy Build.* **2017**, *141*, 96–113. [CrossRef]
19. Jain, A.; Smarra, F.; Reticcioli, E.; D’Innocenzo, A.; Morari, M. NeurOpt: Neural Network Based Optimization for Building Energy Management and Climate Control. In Proceedings of the 2nd Conference on Learning for Dynamics and Control, Berkeley, CA, USA, 30–31 May 2019; Bayen, A.M., Jadbabaie, A., Pappas, G., Parrilo, P.A., Recht, B., Tomlin, C., Zeilinger, M., Eds.; PMLR The Cloud, 2020; Volume 120, pp. 445–454. Available online: <http://proceedings.mlr.press/> (accessed on 12 March 2021).
20. Bengio, Y.; Deleu, T.; Rahaman, N.; Ke, N.R.; Lachapelle, S.; Bilaniuk, O.; Goyal, A.; Pal, C. A Meta-Transfer Objective for Learning to Disentangle Causal Mechanisms. *arXiv* **2019**, arXiv:1901.10912.
21. Koponen, P.; Niska, H.; Mutanen, A. Mitigating the Weaknesses of Machine Learning in Short-Term Forecasting of Aggregated Power System Active Loads. In Proceedings of the 17th IEEE International Conference on Industrial Informatics (INDIN 2019), Espoo, Finland, 23–25 July 2019; IEEE: Piscataway Township, NJ, USA, 2019; pp. 303–310.
22. De Wilde, P.; Martinez-Ortiz, C.; Pearson, D.; Beynon, I.; Beck, M.; Barlow, N. Building simulation approaches for the training of automated data analysis tools in building energy management. *Adv. Eng. Inform.* **2013**, *27*, 457–465. [CrossRef]
23. Li, Q.; Meng, Q.; Cai, J.; Yoshino, H.; Mochida, A. Applying support vector machine to predict hourly cooling load in the building. *Appl. Energy* **2009**, *86*, 2249–2256. [CrossRef]
24. Li, X.; Lü, J.H.; Ding, L.; Xu, G.; Li, J. Building Cooling Load Forecasting Model Based on LS-SVM. In Proceedings of the 2009 Asia-Pacific Conference on Information Processing, APCIP 2009, Shenzhen, China, 18–19 July 2009; IEEE: Piscataway Township, NJ, USA, 2009. [CrossRef]
25. Zhao, H.X.; Magoulès, F. Parallel Support Vector Machines Applied to the Prediction of Multiple Buildings Energy Consumption. *J. Algorithms Comput. Technol.* **2010**, *4*, 231–249. [CrossRef]
26. Ben-Nakhi, A.E.; Mahmoud, M.A. Cooling load prediction for buildings using general regression neural networks. *Energy Convers. Manag.* **2004**, *45*, 2127–2141. [CrossRef]
27. Turhan, C.; Kazanasmaz, T.; Uygun, I.E.; Ekmen, K.E.; Akkurt, G.G. Comparative study of a building energy performance software (KEP-IYTE-ESS) and ANN-based building heat load estimation. *Energy Build.* **2014**, *85*, 115–125. [CrossRef]
28. Yun, K.; Luck, R.; Mago, P.J.; Cho, H. Building hourly thermal load prediction using an indexed ARX model. *Energy Build.* **2012**, *54*, 225–233. [CrossRef]
29. Häkkinen, T.; Ala-Juusela, M.; Shemeikka, J. Usability of energy performance assessment tools for different use purposes with the focus on refurbishment projects. *Energy Build.* **2016**, *127*, 217–228. [CrossRef]
30. Kingma, D.P.; Ba, J.L. Adam: A Method for Stochastic Optimization. In Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015-Conference Track Proceedings, San Diego, CA, USA, 7–9 May 2015; DBLP: Wadern, Germany, 2015.
31. Sastry, S.; Bodson, M.; Bartram, J.F. Adaptive Control: Stability, Convergence, and Robustness. *J. Acoust. Soc. Am.* **1990**, *88*, 588–589. [CrossRef]
32. Akaike, H. A new look at the statistical model identification. *IEEE Trans. Autom. Control.* **1974**, *19*, 716–723. [CrossRef]
33. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef]
34. Guinot, V.; Gourbesville, P. Calibration of physically based models: Back to basics? *J. Hydroinformatics* **2003**, *5*, 233–244. [CrossRef]