energies

Article

An Automated Machine Learning Approach towards Energy Saving Estimates in Public Buildings

Felix Biessmann [1,2,†], Bhaskar Kamble [3,†] and Rita Streblow [2,4,*]

1 FB VI (Computer Science Department), Berlin University of Applied Science and Technology, Luxemburger Str. 10, 13353 Berlin, Germany; felix.biessmann@bht-berlin.de
2 Einstein Center Digital Future, Wilhelmstraße 67, 10117 Berlin, Germany
3 Senercon GmbH, Hochkirchstr. 11, 10829 Berlin, Germany; bhaskar.kamble@senercon.de
4 Faculty III, TU Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany
* Correspondence: r.streblow@tu-berlin.de
† These authors contributed equally to this work.

Abstract: Reducing the energy consumption of buildings in the public sector is an important component in our efforts towards reaching our sustainability goals. In this context, a decisive prerequisite for administrations and policy makers is a tool for estimating the effectiveness of measures to reduce energy consumption. Estimating the impact of planned investments in building technology at scale, however, remains challenging, mainly for two reasons. For one, accurate physical modeling requires detailed building data, which can be difficult to obtain. Second, adapting established building models to novel measures aiming at energy consumption reduction is difficult. Hence, modeling building consumption patterns after retrofitting is a non-trivial task, and more research is needed to improve modeling techniques as well as to assess their effectiveness across a wide range of application scenarios. Modeling tools need to be generic enough to enable modeling of a variety of building types, they should ideally require as few input features as possible and they should allow for a high degree of automation in the selection and calibration of building modeling tools. Here, we propose a novel machine learning approach that does not require detailed building data and can automatically adapt to retrofitting measures. We evaluate our method on a data set of 113 public buildings in 4 building categories in Berlin, Germany. The data set contains energy consumption data in the initial state and after implementation of a weather-predictive heating control system. Despite being fully automated and requiring only minimal information about the building, our model can reliably predict the energy consumption of large public buildings better than established methods. All code and data are publicly released.

Keywords: machine learning; automated machine learning; energy demand prediction; energy efficiency measure; public buildings

Citation: Biessmann, F.; Kamble, B.; Streblow, R. An Automated Machine Learning Approach towards Energy Saving Estimates in Public Buildings. *Energies* **2023**, *16*, 6799. https://doi.org/10.3390/en16196799

Academic Editor: Denis N. Sidorov

Received: 14 August 2023
Revised: 9 September 2023
Accepted: 21 September 2023
Published: 25 September 2023

1. Introduction

Buildings are responsible for about a third of global energy consumption and a quarter of CO_2 emissions [1]. Heating is a major factor in energy consumption in buildings. Especially in public buildings, the energy savings that could be achieved with improvements involved in the heating of buildings are substantial.

Various options for retrofitting or improving the efficiency of heating systems exist [2]. However, the impact of these measures is often not clear. In part, this is due to the complexity of modeling the energy consumption in buildings. In the past decades, significant improvements were achieved in both building structure as well as building technology. On the technology side, the control of heating systems has been profiting from a number of advancements in the fields of physical models, model predictive control, time series forecasting and machine learning. Each of these approaches has advantages and disadvantages, but what most of them have in common is that they can be difficult to calibrate to different building categories and usage patterns therein. One reason for that is the missing availability of high-quality data and, consequently, the lack of studies dedicated to the modeling of energy demand in public buildings on an urban level across a wide range of building types [3]. The research challenge addressed in this study is to develop data-driven methods (as opposed to detailed physics-based modeling) that can predict energy demands for a given building with a minimal amount of information about the specific building properties.

The need for increased energy efficiency in buildings has led to many strategies for retrofitting buildings with energy consumption optimization (henceforth referred to as *ECO*). These methods cover a wide range of invasiveness and financial cost. Determining which measures should be taken to increase the building efficiency requires robust and accurate models of energy consumption in buildings. Given the complexity of the modeling task, the heterogeneity of ECO measures and the lack of data and studies in the field, more research is needed to better support decision makers with modeling options that generalize well and are simple to calibrate in practice.

In this study, we investigate the potential of ML methods for modeling both energy consumption as well as energy saving potential in large public buildings. Our main contributions are threefold. First, we adopt an automated machine learning (AutoML) approach to predict energy consumption in large public buildings with a minimal set of building variables across a heterogeneous spectrum of functionalities. The method was developed with a special focus on automation in order to enable practitioners to quickly adapt it to other buildings and usage patterns with a minimal amount of training data. Second, we demonstrate how the proposed approach can be used to robustly predict the impact of energy consumption optimization measures across a wide range of building categories. Third, we conduct an extensive empirical comparison of established methods and our proposed approach on 113 large public buildings in 4 different building categories. These contributions and the structure of the paper are highlighted in Figure 1.

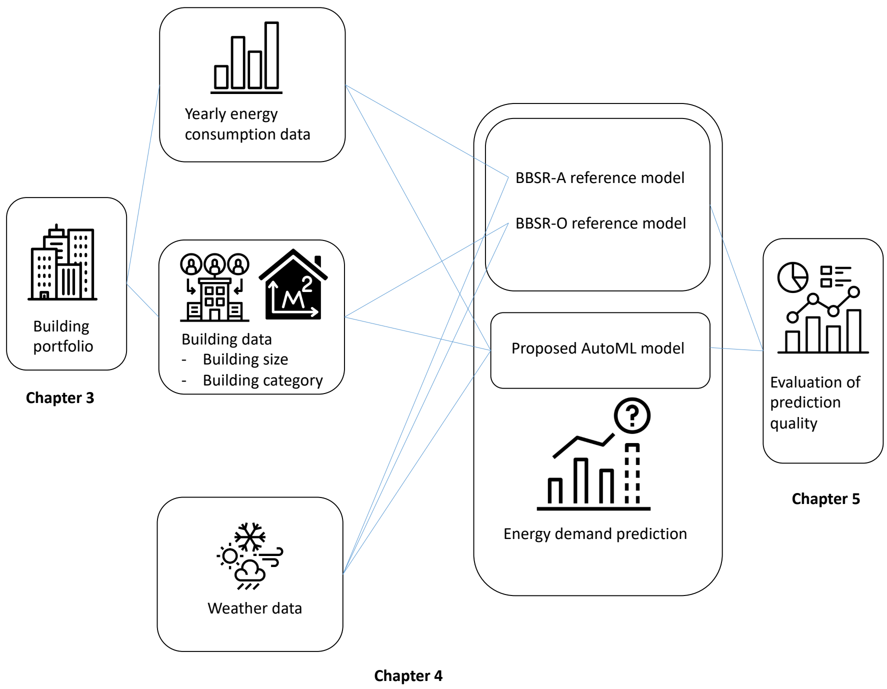

Figure 1. Outline of the experimental setting in this study. In Section 3, we illustrate the data set of 113 public buildings in Berlin; Section 4 details the methodology for the reference models as well as the proposed AutoML approach; Section 5 presents an evaluation and comparison of the predictive performance of the two reference models and the proposed AutoML model. Note that all methods obtain the same amount of data, but the reference models aggregate the data in a different manner than the AutoML approach.

2. Related Work

Modeling of energy consumption in buildings can be accomplished with a variety of approaches. A promising solution is to use physical models that capture all relevant details of buildings, such as their building structure, size and material [4–8]. With appropriate physics models and enough information about each individual building, it is possible to model the thermal inertia or the impact of sunlight on heating energy demand.

A central problem with these models, however, is that it can be difficult to include human usage patterns. This is why machine learning (ML) approaches have become popular also in the field of building technology to model energy consumption [3,9]. In an extensive review of the state of the art in modern machine learning (ML) methods for building technology, only one public building type (schools) is considered [3]. Considering the importance of high-quality data availability for ML methods, this highlights the need for more studies that consider more and more fine-grained functionalities of public buildings. In this study, we investigate the potential of ML models in a large data set of 113 large public buildings over several years with and without energy optimization measures in the form of a weather-predictive heating control algorithm.

A key advantage of ML methods in the building technology sector is that while classical physical models require detailed information about building properties, ML methods can model building energy consumption implicitly via extracting the relevant building properties from historic observations. Previous studies have used a variety of building features to model energy consumption, such as 3D geometry data, GIS-based building footprint, building height and number of stories for each building [4]. Here, we aim at an approach that requires as little information as possible.

When using ML methods, selecting the right ML model class can be difficult. Many studies investigated different ML models, for instance, neural networks [9–12], SVM [13] and random forests [14–16]. However, finding the right model type for each individual setting requires automation in the model selection procedure, especially in the building sector. Here, we follow the approach by [17] and extend their work to fully automated model selection and hyperparameter tuning with an automated ML (AutoML) approach.

3. Data Set

The data set comprised 113 public buildings in total and covered a wide range of building functions and properties. All data were provided by the Berliner Immobilienmanagement GmbH (BIM) (https://www.bim-berlin.de/ (accessed on 13 August 2023)), a company owned by the state of Berlin managing public buildings in the city of Berlin. In the data shared along with the manuscript, records were anonymized by removing addresses of buildings. Our data set has annual consumption values for 47 fire stations, 12 cultural institutions (e.g., theaters), 33 schools of higher education and 21 police stations. The building yearly consumption for heating (including hot water) for all categories and the building areas as well as the distribution of buildings in the data set are listed in Figure 2. Heating systems in these buildings were either based on gas or district heating. The energy consumption was monitored in the intial state of the building and after the implementation of a weather-predictive heating control algorithm as an ECO measure.

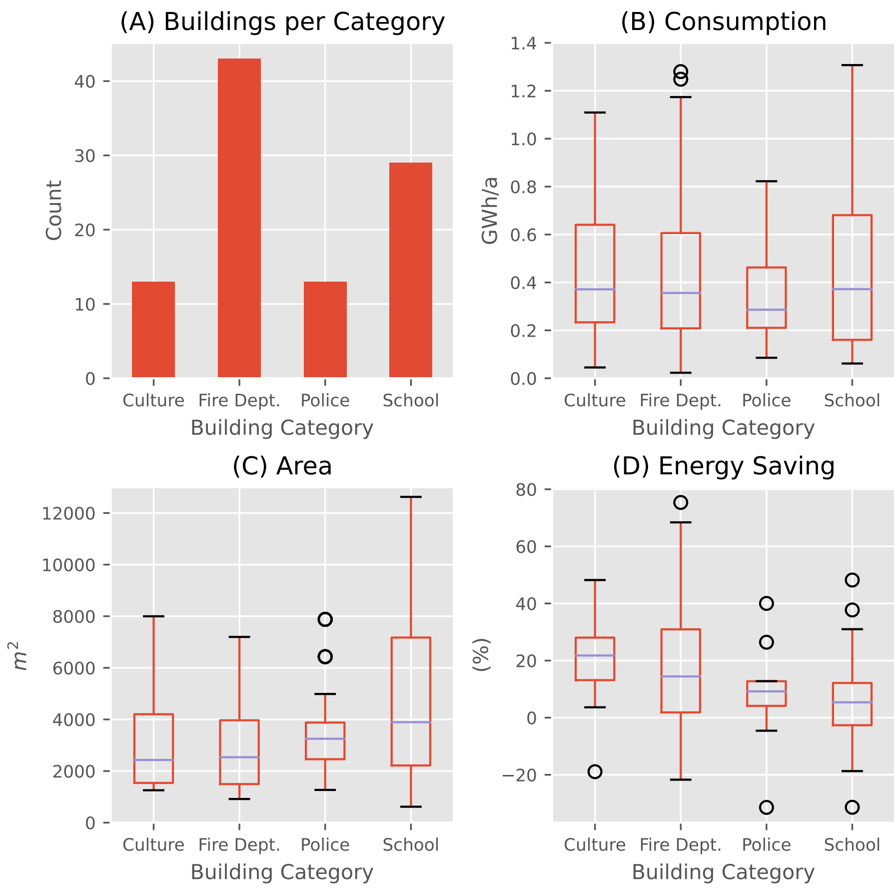

Figure 2. The data set used in this study contains yearly energy consumption for heating in 113 public buildings in Berlin, Germany. Shown are (**A**) counts of buildings in each building category, (**B**) yearly consumption of buildings per category, (**C**) building area in each category and (**D**) energy savings obtained with a retrofitting measure. Boxes indicate 25th/75th percentiles, whiskers 5th/95th percentiles.

4. Methods

The goal of our study is to investigate whether yearly energy consumption of public buildings can be predicted with a limited feature set across a wide range of building characteristics and functionalities. This goal is motivated by the limited availability of building data. A second goal is to develop methods to estimate energy savings when buildings are retrofitted with energy efficiency measures for the operation of the building. This second goal is motivated by the difficulty of existing modeling approaches to account for a change in the consumption characteristics after implementation of a new control strategy. Given these constraints, we selected two different approaches that appeared most promising from a scientific perspective as well as from a practitioner's viewpoint. Both methods are intended to use a minimal set of features, which, in the context of this study, was limited to building size, building usage and climatic constraints. The first approach is based on the regulatory standards put forward by the German Federal Institute for Research on Building, Urban Affairs and Spatial Development, in German, Bundesinstitut für Bau-, Stadt- und Raumforschung (BBSR), henceforth referred to as the BBSR approach. In order to adapt this approach to the settings investigated in this study and to allow for a fair comparison between methods, we developed two extensions of this approach. Next to this rather simple model class, we also developed a novel modeling approach based on recent achievements in automated machine learning (AutoML). The advantage of this approach is that it does not require extensive modeling efforts, such as feature engineering and model selection, and instead learns which ML model or ensemble of ML models is optimal for a given prediction scenario. In the following, we will denote the measured and predicted building consumption for heating as

$$\text{Measured yearly energy consumption } \; y \in \mathrm{R}^1 \tag{1}$$

$$\text{Predicted energy demand } \; \hat{y} = v(x) \tag{2}$$

where v refers to one of the prediction models outlined below and $x \in \mathrm{R}^d$ refers to a d-dimensional feature vector specific to a building. For simple models (see Sections 4.1–4.3), the building-specific features x are merely the building type and building area; for ML models (Section 4.4), the feature vector x comprises also building type and building area as building features and, additionally, climate features. The range of measured yearly energy consumption values y, the target variable, is illustrated in the boxplot in Figure 2, middle panel.

4.1. Linear BBSR Model

The German Federal Institute for Research on Building, Urban Affairs and Spatial Development, in German, Bundesinstitut für Bau-, Stadt- und Raumforschung (BBSR), has developed an adaptation of a linear model for estimating the energy consumption in public buildings [18]. The model uses an estimate of building-specific energy consumption constants (ECCs), reflecting usage patterns and building characteristics. Table 1 lists the ECCs in kWh/(m^{2}year) for heating $\mathrm{ECC_H}$ and hot water $\mathrm{ECC_W}$. These ECCs are scaled by the respective building area A and some non-linear factor f dependent on the area to obtain the building category-specific consumption prediction v_{BBSR}:

$$v_{\mathrm{BBSR}}(x) = fA\,\mathrm{ECC_H} + fA\,\mathrm{ECC_W} \tag{3}$$

where f is a factor depending on the net floor area A (in m^2) of the building and defined in [18] as

$$f = \begin{cases} 1.46 & \text{if } A \le 500\ \mathrm{m}^2 \\ 4.53 \cdot A^{-0.215} + 0.27 & \text{if } 500 < A \le 50{,}000\ \mathrm{m}^2 \\ 0.71 & \text{if } A > 50{,}000\ \mathrm{m}^2 \end{cases} \tag{4}$$

Table 1. Overview of energy consumption constants in kWh/(m^{2}year) for heating ($\mathrm{ECC_H}$) and hot water ($\mathrm{ECC_W}$) for each building category used in Equation (3) to estimate the building energy demand v_{BBSR}. The right column shows the factors used to adapt the original BBSR model to buildings retrofitted with energy consumption optimization (ECO) technology.

Category	**$\mathrm{ECC_H}$**	**$\mathrm{ECC_W}$**	**ECO Factor**
Police Stations	52.4	7.4	0.86
Fire Stations	50.8	7.1	0.86
Schools	49.3	22.4	0.96
Cultural	55.9	7.5	0.87

Note that this model does not require any data for calibration or training. This can be seen as an advantage to other (ML-based) approaches, as it only requires the building area and category to provide an estimate of the energy consumption. However, it does not capture building-specific changes in usage patterns or building properties. In order to adapt this approach to the settings investigated in this study, prediction of energy savings, and to allow for a fair comparison between this approach and other ML-based methods, we developed two extensions of v_{BBSR}.

4.2. Adapted BBSR Model (BBSR-A)

As the previous model v_{BBSR} does not take into account energy savings obtained by improved operation strategies with energy consumption optimization (ECO) technology, we adapted v_{BBSR} slightly. For each building category, we measured the energy demand $y_{\text{before}}, y_{\text{after}}$ before and after the installation of the respective ECO technology. We calculated the ECO factor e as the ratio of energy consumption before and after the installation

$$\text{ECO factor } \; e := \frac{y_{\text{after}}}{y_{\text{before}}} \tag{5}$$

where the measured consumption y was climate corrected with the factors listed in Table 2. The ECO factor e was then used to scale the v_{BBSR} estimate

$$v_{\text{BBSR-A}}(x) = e \, v_{\text{BBSR}}(x). \tag{6}$$

Estimates of the ECO factors e for each building category are listed in the right column of Table 1. This adaptation enabled the BBSR model to take energy efficiency measures into account in its estimate. Note that the BBSR-A approach is based on buildings and their respective energy demands in the data set of the report published by the BBSR [18]. These buildings could exhibit slightly different patterns of consumption compared to our data set of public buildings. To ensure a fair comparison, we developed another variant of the linear BBSR model in which we used the same modeling assumptions, but we calibrated the building category-specific measures to our data set of public buildings.

Table 2. Climate features (TM: average temperature, SO: sunshine hours, NM: cloud cover, FM: wind, RFM: humidity) and climate correction factors (KF) for the years considered in this study.

Year	KF	TM	SO	NM	FM	RFM
2021	1.08	5.72	3.31	6.39	4.29	82.02
2020	1.20	7.37	4.45	5.67	4.38	76.78
2019	1.19	7.01	3.87	5.59	4.53	78.50
2018	1.17	5.15	3.53	5.72	4.39	80.24
2017	1.11	6.19	2.98	6.09	4.50	83.31
2016	1.10	5.19	3.01	6.06	4.23	81.88
2015	1.13	7.56	4.24	5.53	4.54	80.07

4.3. Optimized BBSR Model

In order to allow for a fair comparison with ML models, which leverage a dedicated training data set drawn from the same distribution as is used to make predictions, we developed another variation of the linear v_{BBSR} model. As most buildings in our data set had a similar area, we omitted the non-linearity captured by the factor f in Equation (4). For the sake of simplicity, we also did not differentiate between hot water and heating consumption. Instead, we only computed the mean energy consumption per square meter $\text{ECC}_{\text{total}}$ in each building category and obtained an estimate of the building energy consumption $v_{\text{BBSR-O}}$ as

$$v_{\text{BBSR-O}}(x) = A \, \text{ECC}_{\text{total}}. \tag{7}$$

We computed the building category-specific constants $\text{ECC}_{\text{total}}$ for each category before and after the implementation of the energy efficiency measure. This allows the $v_{\text{BBSR-O}}$ estimated to account for the retrofitting measures, such as $v_{\text{BBSR-A}}$, but, in contrast to the above two modeling approaches, $v_{\text{BBSR-O}}$ is calibrated specifically to the data set investigated in this study. In other words, the two BBSR predictors are identical in terms of their modeling assumptions, but the ECC values in the BBSR-A model are calibrated on the data set used for the BBSR report [18], whereas the ECC values for the BBSR-O predictor are calibrated on the specific data set of 113 public buildings used in this study.

4.4. Automated Machine Learning Model

While the above approaches had the advantage of providing robust estimates of energy demand with simple models, some variation in the data will be difficult to capture. Most importantly, there could be factors that are not related to the building category alone, but are specific to a single building.

In order to account for building-specific consumption patterns with a minimal amount of building data, we developed an AutoML-based approach that leverages all measurements of yearly consumptions optimally. More concretely, the approach uses data from buildings with and without energy efficiency measures and is designed to capture building-specific consumption patterns without explicitly providing any building features. The proposed AutoML model predicts the yearly demand $v_{ML}(x)$ as an ensemble of ML models, including linear models, K-Nearest Neighbor regression, decision trees, gradient boosted trees, random forests as well as neural networks, trained with an AutoML package. In this study, we used the AutoML package AutoGluon Tabular [19]. The feature vector x here refers to a concatenation of features, listed in Table 3, that captures both building characteristics and climate variation.

Table 3. Features used for AutoML approach. Building features are illustrated in Figure 2, and climate features are listed in Table 2.

Building Features	Climate Features (Yearly Aggregates)
Area in m^2	Outside temperature (2 m above ground)
Building category	Sum of sunshine hours
Energy efficiency measure (True or False)	Cloud cover
Consumption in past years	Wind
	Humidity

Climate features were obtained from the Deutscher Wetterdienst (DWD) [20]. All features are concatenated into a feature vector x and provided as input to an AutoML model with a regression loss function that optimizes the mean-squared error of the predicted yearly energy demand $\hat{y}$ and the measured yearly energy consumption y.

4.5. Modeling Seasonal Climate Changes

Comparing and aggregating yearly consumption values requires a correction for climate variation. In warmer years, the energy consumption for heating is lower. If such climate variations correlate with the installation of energy consumption optimization technology, the resulting energy savings will be confounded by the climate variation.

4.5.1. Correcting for Seasonal Climate Changes

A well-established [20] technique to correct for such climate variation is to multiply the measured yearly consumption by a climate correction factor. This factor grows linearly with yearly average temperature. We used the climate correction factors of the Deutscher Wetterdienst to correct for climate variations in the measured and predicted yearly consumptions. For the years considered in this study, we list the climate correction factors in Table 2.

4.5.2. Sampling Realistic Climate Characteristics for Forecasts

Using climate features as inputs to the AutoML models to predict energy demand and energy savings in future years requires sampling climate features. In order to obtain realistic estimates of the climate variation for predictions with the AutoML model, we performed principal component analysis (PCA) on the climate features in Table 2 and considered only the first principal component, which accounted for 80% of the variance in the data. We sampled a linear space grid of values between -1 and 1 standard deviation of the data along the first principal component and projected these synthesized data points to the original climate feature space to obtain realistic samples of climate variation for forecasts of energy consumption and energy savings with the AutoML model.

5. Results

In the following, we compare the results of the two modeling approaches. Note that all comparisons were conducted on climate-corrected consumption values, meaning each yearly consumption was multiplied with its corresponding climate correction factor as shown in Table 2. This ensured that errors in consumption predictions or energy saving predictions were less confounded with climate variation.

5.1. Modeling Energy Consumption

Figure 3 shows a comparison of the true energy consumption of each building on the x-axis and the energy consumption predicted by different models on the y-axis. Both values are corrected for climate variations. Symbols on the black line indicate accurate energy consumption prediction, symbols below the line indicate that the model underestimated the consumption and symbols above the line indicate that the model overestimated the energy consumption. Note that there are multiple data points for each building, corresponding to different years. For our test data, we only considered years and buildings after the installation of the weather-predictive heating control algorithm as an energy efficiency measure in a leave-one-building-out cross-validation setting. This cross-validation setting was chosen because it authentically reflects the real world setting in which we want to make a prediction for one building for which we have historical data and some other buildings, some of which were already retrofitted. Our results demonstrate that AutoML-based methods achieve lower errors when predicting energy demand after the implementation of energy efficiency measures. This is also reflected in the quantitative comparison in Table 4, where, for each building category but also for the mean across all categories, the proposed AutoML approach achieves significantly lower mean absolute errors when predicting energy consumption patterns after energy efficiency measures.

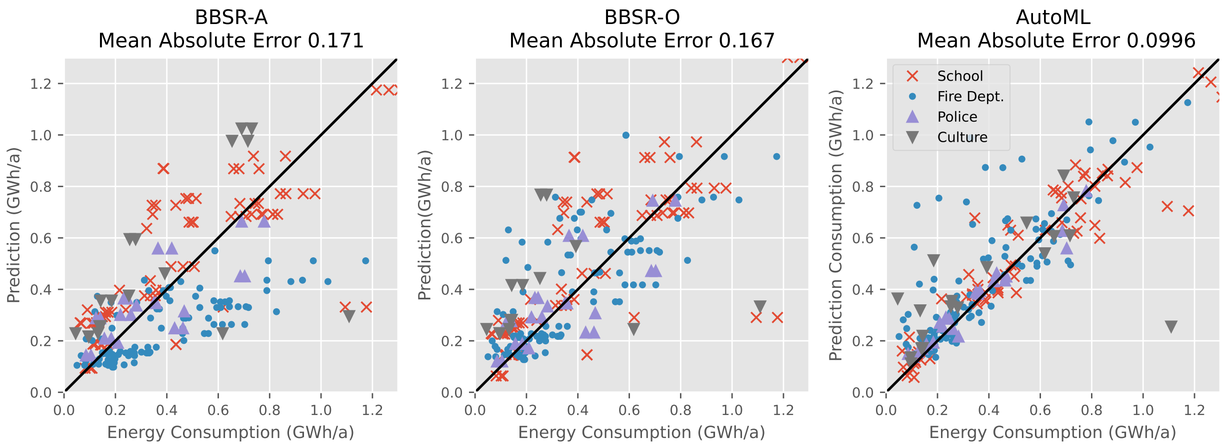

Figure 3. Comparison of measured yearly energy consumption (x-axis) and the consumption predicted (y-axis) with the respective model described in Section 4.2 (**left**), Section 4.3 (**middle**) and Section 4.4 (**right**). Symbols below the line indicate that the respective model underestimated the consumption, symbols above the line indicate that the model overestimated the energy consumption. AutoML-based predictions (**right**) achieve lower errors compared to the linear methods $v_{\text{BBSR-A}}$ (**left**) and $v_{\text{BBSR-O}}$ (**middle**).

Table 4. Comparison of mean absolute errors (in kWh) for building energy consumption after retrofitting; lowest errors are indicated with bold-faced font. The AutoML approach achieves significantly lower absolute errors in individual building categories as well as in the average across all building categories.

	MAE BBSR-A	MAE BBSR-O	MAE AutoML
School	156,697	150,102	**71,109**
Fire Dept.	156,521	122,545	**106,553**
Police	106,684	103,978	**38,904**
Culture	375,970	541,350	**236,521**
Mean	198,968	229,494	**113,272**

5.2. Modeling Energy Savings

Predicting energy demand accurately after the implementation of energy efficiency measures is a fundamental prerequisite for estimating energy consumption savings. Building on the results in Section 5.1, we compared the accuracy of different methods when predicting the energy demand savings after energy efficiency measures. To that end, we computed the *true*/*predicted* energy savings per building by dividing the *true*/*predicted* consumption after retrofitting the buildings by the *measured* consumption before retrofitting y_{before} and transformed this ratio (which is equivalent to the ECO factor in Equation (5)) to percent

$$\text{true energy saving} = 100 \cdot \left(1 - \frac{y_{\text{after}}}{y_{\text{before}}}\right) \tag{8}$$

$$\text{predicted energy saving} = 100 \cdot \left(1 - \frac{\hat{y}_{\text{after}}}{y_{\text{before}}}\right) \tag{9}$$

In Table 5, we compare the overall prediction accuracy of the energy savings predicted by different methods. Investigating the mean absolute errors (in % energy savings), we find that simple linear methods ($v_{\text{BBSR-A}}$, $v_{\text{BBSR-O}}$) result in substantially larger errors than the proposed AutoML approach. The AutoML approach achieves a higher saving prediction accuracy in each individual building category as well as across all categories. We also examine the prediction difference in more detail in Figure 4. Comparing the energy saving predictions for the best linear method BBSR-O and the AutoML approach, we find that in most categories, the linear BBSR-O method results in a substantially larger error than the AutoML method. Only for some buildings in the categories *Fire Dept.* and *School*, the BBSR-O method achieved smaller errors, which, however, does not impact the overall comparison of the methods.

Table 5. Comparison of mean absolute errors of saving predictions (in %) for each building category; lowest errors are indicated with bold-faced font. In all categories as well as overall, the AutoML approach achieves lower energy saving prediction errors.

	BBSR-O	BBSR-A	AutoML
School	45	52	**14**
Fire Dept.	24	28	**22**
Police	30	32	**12**
Culture	108	70	**29**
Mean	41	40	**19**

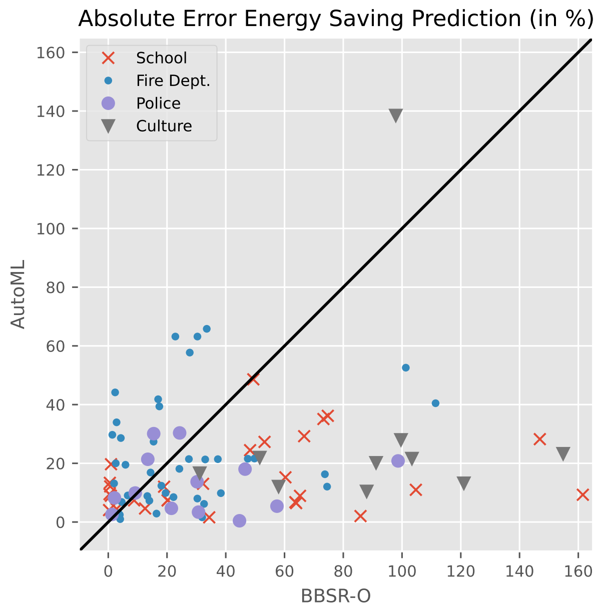

Figure 4. Comparison of absolute errors (in % of energy saving predictions for the BBSR-O model (x-axis) and for the AutoML model (y-axis)). For most buildings, the linear BBSR-O model overestimates the savings.

When interpreting these empirical results, it is important to keep in mind that not all energy savings (or, more generally, all changes in the building energy consumption, not all buildings exhibited a reduction in energy consumption) can be directly attributed to the new heating control strategy as an energy efficiency measure. When the weather-predictive heating control algorithm was implemented, it was usually accompanied by an overall check of the heating system. This often leads to improvements or changes in the energy consumption that might not necessarily be attributed to the specific energy efficiency measure but rather to heating maintenance work. In our experiments, we excluded buildings for which the heating system was not functioning properly according to the building management company.

5.3. Modeling Energy Consumption for Individual Buildings

With the methods developed in this study, we can make predictions for energy demand of individual buildings after retrofitting buildings with energy efficiency measures. In Figure 5, we show such a prediction for a building in the category *Police Station* as a function of the mean heating period outside temperature, defined here as days with temperatures lower than 18 °C. We show the linear predictions with the BBSR-O model as well as the more flexible predictions with the AutoML approach for weather features simulated with the method detailed in Section 4.5.2. Both models are able to predict the savings obtained with the retrofitting measure, however, the AutoML approach extrapolates differently compared to the linear BBSR-O approach: the AutoML approach is more conservative than the linear modeling approach for temperatures outside of the range of the training data. This could explain the better empirical performance of the AutoML approach compared to the classical linear method.

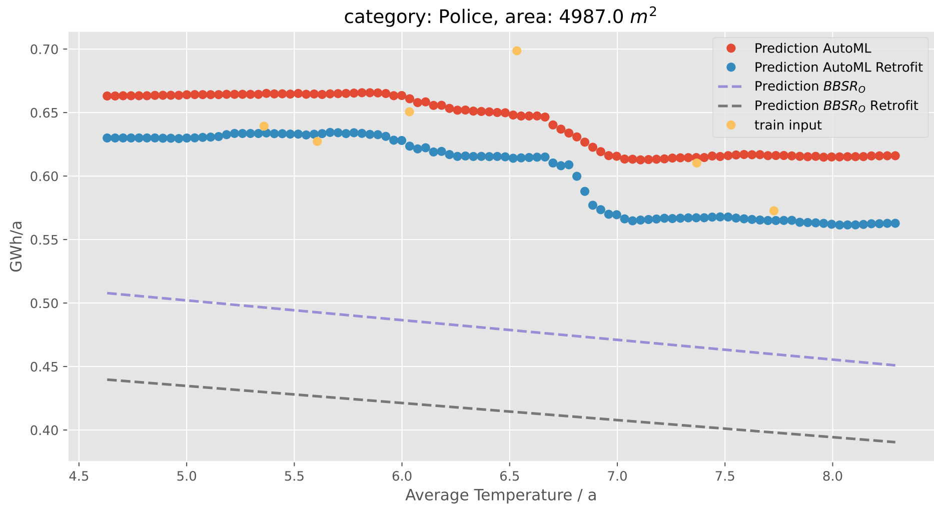

Figure 5. Climate-corrected consumption predictions for the best linear model $v_{\text{BBSR-O}}$ (dashed lines) and the ML-based approach (dotted lines) with and without the energy efficiency measure. The ML-based approach captures the building-specific consumption patterns before and after the retrofitting better than $v_{\text{BBSR-O}}$. Training data for this particular building, measured before the energy efficiency measure, are indicated as yellow dots.

6. Conclusions

In this study, we developed methods to predict the energy demand of large public buildings across a wide spectrum of building categories. We compared different methods in a comprehensive empirical evaluation on a data set of yearly energy consumption in over 100 buildings. The methods compared were selected and adapted to meet the needs of policy makers and administrations working towards fulfilling our sustainability goals. The methods investigated can be applied to arbitrary energy efficiency measures and require almost no building data. This means that these methods can be used for scalable and automated data-driven decisions on energy saving measures for public buildings. We demonstrate that AutoML-based approaches achieve substantially higher accuracies when modeling energy demand patterns after installation of energy efficiency measures in each individual building category investigated, as well as across all categories. The methodology used in our approach is readily applicable to other application scenarios beyond heating energy: as the method learns from measurements available through smart meters without requiring detailed knowledge about the building structure, it is straightforward to transfer the methodology to, for instance, energy demand for lighting.

Author Contributions: Data Preparation—B.K., Modelling—F.B., Analyses—F.B., Writing—original draft, F.B. and B.K.; Writing—review & editing, F.B., B.K. and R.S. All authors have read and agreed to the published version of the manuscript.

Funding: Felix Bießmann acknowledges funding from the Einstein Center Digital Future, Berlin as well as the Project "Prognose Wettergeführte Heizungssteuerung" (PWH, Project Number 1294-B5-O). Bhaskar Kamble acknowledges funding from the Project "Prognose Wettergeführte Heizungssteuerung" (PWH, Project Number 1294-B5-O). Rita Streblow acknowledges funding from the Einstein Center Digital Future.

Data Availability Statement: All code and data is publicly available at https://github.com/calgo-lab/PWH (accessed on 13 August 2023).

Acknowledgments: We thank Joachim Seifert for valuable feedback on an earlier version of this work.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. U.S. Energy Information Administration. *Energy Statistics*; U.S. Energy Information Administration: Washington, DC, USA, 2023.
2. Lee, J.; Shepley, M.M.; Choi, J. Exploring the effects of a building retrofit to improve energy performance and sustainability: A case study of Korean public buildings. *J. Build. Eng.* **2019**, *25*, 100822. [CrossRef]
3. Fathi, S. Machine learning applications in urban building energy performance forecasting: A systematic review. *Renew. Sustain. Energy Rev.* **2020**, *133*, 110287. [CrossRef]
4. Chen, Y.; Hong, T. Impacts of building geometry modeling methods on the simulation results of urban building energy models. *Appl. Energy* **2018**, *215*, 717–735. [CrossRef]
5. Fonseca, J.A.; Schlueter, A. Integrated model for characterization of spatiotemporal building energy consumption patterns in neighborhoods and city districts. *Appl. Energy* **2015**, *142*, 247–265. [CrossRef]
6. Reinhart, C.F.; Davila, C.C. Urban building energy modeling—A review of a nascent field. *Build. Environ.* **2016**, *97*, 196–202. [CrossRef]
7. Chalal, M.L.; Benachir, M.; White, M.; Shrahily, R. Energy planning and forecasting approaches for supporting physical improvement strategies in the building sector: A review. *Renew. Sustain. Energy Rev.* **2016**, *64*, 761–776. [CrossRef]
8. Zhao, H.x.; Magoulès, F. A review on the prediction of building energy consumption. *Renew. Sustain. Energy Rev.* **2012**, *16*, 3586–3592. [CrossRef]
9. Seyedzadeh, S.; Rahimian, F.P.; Glesk, I.; Roper, M. Machine learning for estimation of building energy consumption and performance: A review. *Vis. Eng.* **2018**, *6*, 5. [CrossRef]
10. Elbeltagi, E.; Wefki, H. Predicting energy consumption for residential buildings using ANN through parametric modeling. *Energy Rep.* **2021**, *7*, 2534–2545. [CrossRef]
11. Truong, N.S.; Ngo, N.T.; Pham, A.D. Forecasting time-series energy data in buildings using an additive artificial intelligence model for improving energy efficiency. *Comput. Intell. Neurosci.* **2021**, *2021*, 6028573. [CrossRef] [PubMed]
12. Olu-Ajayi, R.; Alaka, H.; Sulaimon, I.; Sunmola, F.; Ajayi, S. Building energy consumption prediction for residential buildings using deep learning and other machine learning techniques. *J. Build. Eng.* **2022**, *45*, 103406. [CrossRef]
13. Ma, Z.; Ye, C.; Li, H.; Ma, W. Applying support vector machines to predict building energy consumption in China. *Energy Procedia* **2018**, *152*, 780–786. [CrossRef]
14. Zhang, X.; Zhang, J.; Zhang, J.; Zhang, Y. Research on the combined prediction model of residential building energy consumption based on random forest and BP neural network. *Geofluids* **2021**, *2021*, 7271383. [CrossRef]
15. Wang, Z.; Wang, Y.; Zeng, R.; Srinivasan, R.S.; Ahrentzen, S. Random Forest based hourly building energy prediction. *Energy Build.* **2018**, *171*, 11–25. [CrossRef]
16. Liu, Y.; Chen, H.; Zhang, L.; Feng, Z. Enhancing building energy efficiency using a random forest model: A hybrid prediction approach. *Energy Rep.* **2021**, *7*, 5003–5012. [CrossRef]
17. Rätz, M.; Javadi, A.P.; Baranski, M.; Finkbeiner, K.; Müller, D. Automated data-driven modeling of building energy systems via machine learning algorithms. *Energy Build.* **2019**, *202*, 109384. [CrossRef]
18. Bundesministerium des Innern, für Bau und Heimat. Bekanntmachung der Regeln für Energieverbrauchswerte und der Vergleichswerte im Nichtwohngebäudebestand. Available online: https://geg-info.de/geg/210503_bmwi_bmi_regeln_energieverbrauchskennwerte_nichtwohnbestand.pdf (accessed on 1 April 2021).
19. Erickson, N.; Mueller, J.; Shirkov, A.; Zhang, H.; Larroy, P.; Li, M.; Smola, A. AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data. *arXiv* **2020**, arXiv:2003.06505.
20. Deutscher Wetterdienst. *Klimafaktoren (KF) für Energieverbrauchsausweise*; Deutscher Wetterdienst: Offenbach am Main, Germany, 2020.