

ORIGINAL RESEARCH

Open access

Comparative Evaluation of Machine Learning Models for Building Power Prediction under Different Data Resolutions

Nora Schneider^{1*}, Tobias Frank¹

Abstract

Forecasting electricity usage is essential for boosting operational efficiency and advancing sustainability in buildings. The present work explores how the interval of data acquisition and the choice of modeling technique influence prediction precision for electricity demand in two contrasting building categories: an academic facility providing 15-minute data and an industrial facility providing hourly data. Multiple machine learning algorithms were tested—among them support vector machines (SVMs) with radial and sigmoid kernels, random forests (RFs), and deep neural networks (DNNs)—across various training/validation splits and feature combinations. Findings demonstrate that finer temporal resolution in data collection typically yields better model performance, reflected in lower RMSE, MAPE, and CV metrics, as well as higher R^2 values. Adding more lagged power consumption variables exerted a considerably stronger beneficial effect on forecast quality than incorporating weather-related variables. Notably, the SVM model equipped with the Radial kernel delivered the most consistent and superior results, proving especially effective at representing intricate nonlinear dependencies. By comparison, the DNN approach, despite acceptable performance on selected measures, exhibited notably higher MAPE values, raising concerns about potential overfitting. Collectively, these observations emphasize the need for deliberate decisions regarding data granularity, predictor variables, and algorithm selection to maximize forecasting reliability and thereby support smarter, more resource-efficient energy management in built environments.

Keywords Power consumption prediction, Machine learning models, Predictive analytics, Feature analysis, Random forest (RF), Support vector machine (SVM)

*Correspondence:

Nora Schneider
nora.schneider@rwth-aachen.de

¹ Department of Electrical Energy Systems, Faculty of Engineering, RWTH Aachen University, Aachen, Germany

Introduction

Escalating global energy requirements, alongside the critical need to combat climate change and global warming, have elevated energy conservation and greenhouse gas mitigation to top priorities in sustainable development agendas. In 2021, building operations accounted for 30% of total global energy consumption. It generated 27% of energy-linked CO₂ emissions, making the built environment one of the largest single sources of greenhouse gases [1]. High-quality short-term load forecasting (STLF) for buildings is indispensable for intelligent energy allocation and optimization—an especially pressing requirement in rapidly urbanizing territories such as China, where emission reductions are urgently needed [2]. In recent years, the proliferation of advanced energy-monitoring infrastructure has produced enormous datasets that hold great promise for refining building performance and reducing resource waste [3].

The ongoing move toward widespread electrification—largely motivated by decarbonization objectives—has further intensified demand for reliable electricity demand forecasts in the building sector [4]. Combining on-site generation, energy storage technologies, and the principles underlying Zero-Emission Buildings (ZEB) creates complex, variable operating conditions that require flexible and resilient forecasting tools [5]. Techniques rooted in Artificial Intelligence (AI), particularly those that exploit high-volume data streams from smart metering systems and IoT sensors, offer powerful avenues to improve energy performance in buildings [6].

Reviewing the historical trajectory of energy forecasting methodologies also sheds light on how contemporary hybrid approaches have matured. Early efforts in the 1970s depended on straightforward deterministic calculations; the 1980s saw the adoption of more formal statistical procedures; from the mid-1990s onward, machine learning (ML) and AI methods steadily gained prominence, each era contributing foundational elements to today's sophisticated solutions [5, 7]. Successfully deploying AI-driven forecasting depends critically on robust mechanisms for gathering, storing, and preprocessing large-scale data [8, 9]. Concurrently, government policies and regulatory frameworks that incentivize energy-efficient practices strongly shape both the uptake and the real-world impact of predictive modeling tools [10].

Building function and geographic location exert substantial influence on appropriate model selection because they dictate characteristic load shapes and variability [11, 12]. Demand profiles diverge markedly across categories (office, factory, university, dwelling) and climatic zones (hot versus cold, wet versus arid). For example, an office block in a cold region will exhibit heating and electricity patterns distinctly different from those of a residence in a warm climate. Model effectiveness, therefore, hinges on these contextual differences, making it imperative to adapt forecasting methods to the particular traits of each building class and its local setting for reliable energy oversight.

Existing research [13] indicates a relative scarcity of publications addressing load prediction specifically for industrial plants and academic campuses. The dominant approaches currently applied to STLF rely heavily on recurrent neural network architectures—particularly long short-term memory (LSTM) cells and gated recurrent units (GRU)—as well as convolutional neural networks (CNNs) and various autoencoder designs.

One study [14] benchmarks three statistical techniques—the Changepoint model, Support Vector Machine (SVM), and Basis function expansions—for general-purpose industrial load prediction, feeding the models with meteorological records, operational timetables, and equipment parameters. Similarly, references [15, 16] systematically compare a range of deep learning and conventional machine learning methods to identify the most promising approach for industrial power forecasting. Beyond pure performance ranking of recurrent models, the investigation reported in Ungureanu *et al.* [17] additionally quantifies the sensitivity of predictions to different input configurations. Campus-level load forecasting is presented in Seefoo Jarquin *et al.* [18], where an artificial neural network (ANN) using the similar-day strategy generates both point and probabilistic predictions for the next 24 hours, leveraging weather data, calendar variables, and recency indicators. Another campus-oriented effort [19] applies ANN architectures to produce day-ahead hourly load estimates, using predictors based on aggregate load, ambient temperature, and selected sub-metered loads. Extending the horizon, the methodology described in Chévez and Martini [20] uses an ANN to anticipate energy consumption 1 hour ahead and at every hour of a complete annual cycle, incorporating temperature, calendar day, type of day, and the immediately preceding hour's consumption as inputs.

Across the broad spectrum of machine learning strategies, Artificial Neural Networks (ANNs), Support Vector Regression (SVR), and Deep Neural Networks (DNNs) remain among the most frequently adopted tools for energy demand estimation [21]. A persistent limitation, however, is their difficulty in fully capturing the sequential and autocorrelated nature intrinsic to energy time series. More advanced deep architectures—most notably LSTM networks and convolutional neural networks (CNN)—generally excel at automatically distilling informative features from voluminous datasets and uncovering concealed temporal structures [22]. While these methods offer clear strengths in modeling long-range dependencies and intricate patterns, they often impose heavy computational demands and require extensive training data. These constraints can hinder deployment in data-scarce environments or under limited processing budgets.

In contrast, discrete data offer several advantages over continuous time-series data. Discrete datasets generally carry a lower risk of overfitting because they contain fewer observations and exhibit simpler underlying structures, which promotes greater model robustness and better generalization to unseen data [23]. Preparing discrete data is far less complicated than the detailed preprocessing steps often required for time-series formats, enabling more efficient, streamlined data management processes [24]. Applying models to discrete data is usually more straightforward, supporting rapid implementation, evaluation, and deployment even with modest computational capabilities [25]. Furthermore, models developed using discrete data often exhibit improved interpretability, delivering more transparent understandings of the key drivers of energy use and yielding practical recommendations for stakeholders. That said, multiple considerations, including the problem domain, data availability, and the specific goals of the forecasting effort, should guide the decision to favor discrete data over time-series data.

Random forest (RF), support vector machine (SVM), and deep neural network (DNN) algorithms are particularly well-suited to discrete data formats, delivering clear performance gains over conventional techniques thanks to their ability to detect complex interactions and structures in the input data [26, 27]. Specifically, SVM—particularly when configured with Radial and Sigmoid kernels—was included because of its proven strength in addressing nonlinear behaviors within modestly sized datasets, which matches the discrete character of the present data. RF was selected for its resilience with intricate, high-dimensional inputs and its built-in mechanisms—via ensemble averaging—to reduce overfitting risks. DNN was added to leverage its strengths in learning complex nonlinear mappings and deriving high-level representations from larger feature sets.

Although the bulk of the short-term load forecasting (STLF) literature centers on residential structures, this work focuses on an industrial facility and an academic facility located in separate geographic areas with differing sampling resolutions. This investigation conducts an in-depth examination of how various predictor variables, machine learning techniques, and proportions for training, validation, and testing partitions affect demand forecasts across these two distinct building categories. In contrast to many prior works that incorporate weather variables, the operational profiles of the studied buildings render climatic influences negligible for predictive purposes. Overall, the research evaluates the performance of diverse ML approaches for anticipating building electricity demand 24 hours ahead. To establish a reference point, a persistent model (PM)—which projects future usage based on the value recorded exactly 24 hours earlier—serves as the benchmark for gauging the relative effectiveness of the ML alternatives.

Drawing from the broader backdrop of escalating global energy pressures and the pivotal importance of reliable electricity demand forecasting, this investigation pursues several targeted research questions designed to enhance forecast quality for varied building categories and climatic settings:

- RQ1. Which predictors (e.g., lagged power usage, temperature, irradiation) exert the greatest influence on electricity demand forecasts in academic buildings versus industrial buildings?
- RQ2. To what degree do prevailing weather patterns (oceanic versus continental Mediterranean) shape the forecasting precision of power demand models across different building categories?
- RQ3. How do data granularity (15-minute intervals in the Academic building versus hourly in the Industrial building) and the allocation ratios among training, validation, and test subsets affect the efficacy of the applied machine learning models?
- RQ4. In what ways does the SVM kernel type (Radial versus Sigmoid) influence the model's capacity to represent nonlinear dynamics in electricity consumption patterns from the two building types?
- RQ5. How do prediction accuracies differ across machine learning approaches (RF, SVM with Radial and Sigmoid kernels, DNN) when deployed on buildings serving contrasting purposes (academic versus industrial)?

Drawing from these inquiries, the study advances the following hypotheses:

Hypothesis 1

(H1). In the Academic building, humidity levels and occupancy patterns are expected to emerge as the dominant predictors of power consumption. In contrast, temperature and equipment operating schedules will play a more decisive role in the Industrial building.

Hypothesis 2

(H2). Weather conditions are expected to exert a stronger effect on forecast precision in the Industrial building than in the Academic building, owing to the pronounced temperature swings characteristic of continental Mediterranean environments.

Hypothesis 3

(H3). Models applied to the Academic building should achieve superior prediction accuracy compared to those for the Industrial building, attributable to the finer temporal resolution of its data collection (15-minute intervals) versus hourly sampling.

Hypothesis 4

(H4). Employing the Radial kernel in SVM is expected to yield higher forecasting accuracy than the Sigmoid kernel for power consumption across both building types, due to its enhanced ability to capture intricate nonlinear structures.

Hypothesis 5

(H5). The DNN approach is expected to surpass both RF and SVM in predictive performance for electricity consumption in both Academic and Industrial buildings, owing to its proficiency in modeling complex nonlinear associations.

These research questions and associated hypotheses will guide the evaluation and comparative assessment of the chosen ML techniques, providing valuable guidance on their suitability and strengths for power demand forecasting across diverse building contexts.

Materials and Methods

Accurate prediction of building electricity usage requires careful feature engineering and thorough model development. The workflow commences with data acquisition and cleaning, proceeds to identify the most relevant influencing factors, and culminates in model fitting and performance validation to ensure reliable outcomes. This methodical framework (**Figure 1**) is indispensable for constructing high-precision models customized to the unique attributes of individual building types.




Figure 1. Diagram of the methodology.

Data collection

Collecting electricity consumption records from diverse building categories forms a critical foundation for reliable forecasting. The Academic and Industrial buildings rely on separate monitoring systems, each configured with sampling rates adapted to their operational requirements. For the Academic building, measurements were recorded every 15 minutes across 3 full academic years, spanning September 2016 to July 2019. This produced a dataset comprising 92,160 individual readings, equivalent to 960 complete days (Table 1). By comparison, the Industrial building's data were gathered hourly over a single year, from March 2022 to February 2023, yielding 8759 readings or 365 full days. In both cases, the electricity consumption figures were obtained from certified metering devices installed and operated by the regional utility provider, in accordance with IEC 62053-11, the international standard for metering precision and dependability.

Table 1. Raw data. Statistics of the academic building.

Building type	Data records (length)	Sampling interval	Total days	Minimum (kW)	Maximum (kW)	Average (kW)	Standard deviation
Academic	92,160	15 min	960	0	10,410	20.24	40.87

Site-specific weather information for each location was obtained from the Solcast online platform, with a temporal resolution aligned with the corresponding consumption data series for seamless integration and subsequent analysis. This synchronized strategy supports robust feature engineering and effective model development, both of which are vital for achieving high-quality power demand predictions.

Data preprocessing

Proper preprocessing transforms raw measurements into a format suitable for machine learning, thereby improving model precision and computational performance. Key activities include data cleansing, outlier detection, scaling, feature

engineering, and dataset partitioning to construct a dependable foundation for electricity consumption forecasting. For the Industrial building, prior preprocessing steps had already been completed, removing the necessity for additional cleaning or outlier treatment (**Table 2**).

Table 2. Clean data. Statistics of both academic and industrial buildings.

Building type	Number of readings	Sampling interval	Number of days	Minimum	Maximum	Average	Standard deviation
Academic	69,216	15 minutes	721	0.15	64.66	23.65	14.74
Industrial	8,759	1 hour	365	10	119	41.82	23.21

Data cleaning

Initial processing focused on the Academic building dataset. Duplicate entries were detected and eliminated. Overall, 1168 records containing NaN values were found, accounting for 1.27% of the total observations. Isolated or paired missing values were imputed using the nearest preceding or following valid entry. Sequences with three or more consecutive NaN values prompted the exclusion of the entire corresponding day, resulting in the removal of 28 such days.

Outlier identification

Subsequently, outlier detection was performed on the Academic building records. Any readings falling below 100 W or above 100,000 W were discarded. This process eliminated 15,745 measurements under 100 W, which stemmed from a known malfunction in the data logging system during that period. Additionally, four readings surpassing 100 kW were identified and excluded. Five further days lacking a complete 15-minute resolution were also removed.

Outlier detection employed a standard deviation-based approach. Data points were grouped according to contextual factors—holidays versus regular weekdays, working hours versus non-working periods, and chiller operational status. Suspect values within these categories were treated using the same imputation rule applied to missing data. Among working-hour observations, 1234 out of 21,666 samples were flagged as outliers (5.7%). During non-working hours, 540 of 51,582 samples (1%) qualified as outliers.

Normalization

Power consumption values for both the Academic and Industrial buildings were scaled to the [0, 1] interval through min-max normalization, as expressed in **Eq.1**.

$$P_{normalized} = \frac{P - P_{min}}{P_{max} - P_{min}} \quad (1)$$

Feature selection

To maximize model performance, only the most pertinent variables were retained as input to the algorithm. This selection was guided by visualizing daily 24-hour consumption profiles, enabling an informed, graphical identification of the most representative subset of available features.

Data division

To assess how different partitioning ratios influence the learning performance of the models, a thorough examination was carried out involving two specific partitioning strategies for each dataset:

- Case 1. Approximately 72% of the data was assigned to the training set, 18% to the validation set, and 10% to the test set.
- Case 2. Approximately 80% of the data was assigned to the training set, 10% to the validation set, and 10% to the test set.

To preserve temporal coherence and enable full-day visualization during model evaluation, the datasets were divided into contiguous 24-hour segments. This approach also ensured that weekdays, weekends, and holidays remained proportionally balanced within each subset. As a result, the exact percentages differed slightly between the Academic and Industrial datasets, yet the overall distribution patterns remained comparable, supporting a consistent and fair comparative assessment (Figure 2).




Figure 2. Dataset split for training, validation, and test.

Prediction models

The models implemented in this study are as follows [26, 27]:

- RF proves especially valuable owing to its capacity to process numerous input variables and detect sophisticated patterns. It builds an ensemble of decision trees, each trained on randomly sampled portions of the data, improving overall accuracy while reducing the risk of overfitting. Moreover, RF offers straightforward interpretability, allowing identification and ranking of the most important predictors contributing to the forecasts.
- SVM excels in forecasting tasks by effectively managing both linear and nonlinear dependencies within the data. It identifies the optimal separating hyperplane that maximizes the margin between different classes or patterns, delivering reliable performance even in high-dimensional spaces.

The selection of an appropriate kernel function is crucial in SVM, as it determines how the input space is mapped and separated, thereby greatly enhancing the model's predictive capability and generalization.

– The Radial Basis Function (RBF) kernel—commonly referred to as the Gaussian or simply Radial kernel—stands out for its effectiveness in capturing nonlinear relationships. Nevertheless, it carries a risk of overfitting when applied to limited or noisy datasets.

– The Sigmoid kernel provides an alternative that accommodates a range of relationship types, although it tends to yield inferior results compared to RBF in scenarios dominated by strong nonlinearity.

- DNN comprises a multi-layered neural network architecture with interconnected nodes across several hidden layers. Such networks are highly adept at learning elaborate, nonlinear feature representations, rendering them well-suited for problems characterized by complex interdependencies. A key strength of DNNs lies in their capacity to automatically extract features directly from raw inputs, thereby boosting their ability to model intricate structures.

The Rectified Linear Unit (ReLU) activation function was selected for its proven efficiency and effectiveness: it accelerates convergence during training, reduces the risk of overfitting, and helps avoid the vanishing gradient problem, thereby supporting the discovery of sophisticated patterns.

Evaluation metrics of the models

Model performance was quantified using six widely accepted error and goodness-of-fit measures: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), Coefficient of Determination (R^2), and Coefficient of Variation (CV). These metrics are calculated according to [Eqs.2–7](#).

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

$$\text{RMSE} = \sqrt{\text{MSE}} \quad (3)$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i| \quad (4)$$

$$\text{MAPE} = \frac{1}{n} \sum_{i=1}^n \left| \frac{\hat{y}_i - y_i}{y_i} \right| \cdot 100 \quad (5)$$

$$R^2 = 1 - \frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (6)$$

$$\text{CV} = \frac{\text{RMSE}}{\bar{y}} \cdot 100 \quad (7)$$

Where y represents the actual observed value, $\hat{y}$ the predicted value, $\bar{y}$ the mean of the observed values, and n the total number of samples (with $n = 24$ for the Industrial building and $n = 96$ for the Academic building).

Case Study

Buildings description

The two buildings examined in this research are:

Academic building

Situated in Arrasate, Spain, this facility experiences an oceanic climate characterized by mild-to-cool temperatures, limited seasonal variation, and high levels of precipitation. The structure comprises five floors and has a gross floor area of 7640 m², accommodating offices, lecture halls, and research laboratories. Beyond standard educational activities, it serves as an active research center, featuring energy-intensive laboratories that require substantial cooling capacity, including a dedicated chiller unit within the HVAC infrastructure. Designed as a smart building, it incorporates automated energy management systems and heat recovery technologies to promote efficient resource use and reliable performance.

Industrial building

Located in Burgos, Spain, the site is subject to a continental Mediterranean climate marked by severe winters, hot summers, and limited, uneven rainfall. The facility comprises a spacious production hall and administrative offices, totaling 11,800 m². Given on-site industrial operations, electricity demand is closely tied to manufacturing activities and remains consistently high. The building must accommodate the region's pronounced seasonal temperature swings while maintaining stable, productive working conditions year-round.

Data analysis and correlations

Following the preprocessing steps ([Table 2](#)), a detailed examination of the datasets was conducted to identify the most suitable input features.

The consumption profile of the Academic building ([Figure 3a](#)) reveals a noticeable rise in electricity usage linked to the operation of a chiller within the HVAC system, persisting over multiple consecutive days. It is also easy to distinguish between holiday periods (highlighted in purple) and regular working days (highlighted in green), as well as the standard operating schedule, which runs from 7 a.m. to 6 p.m.




Figure 3. Daily consumption patterns of both (a) Academic and (b) Industrial buildings.

Visual inspection of the Industrial building's daily load patterns (**Figure 3b**) allows separating days ending at 7 p.m. (marked in orange) from those extending to 11 p.m. (marked in green), with all workdays commencing at 7 a.m. Additionally, periods exhibiting consistently elevated demand—defined as power exceeding 78 kW between 7 a.m. and 7 p.m.—are classified as base-load conditions. Non-operational days stand out clearly due to near-zero consumption and are shown in purple.

All of these identifiable characteristics, given their direct influence on electricity usage, were incorporated as input variables for the predictive models.

Pearson Correlation: The relationships between lagged consumption values, meteorological variables, and current power demand were evaluated. Pearson correlation coefficients were computed over two different time horizons: one month and one full year. The environmental variables considered included outdoor air temperature, global horizontal irradiation, cloud opacity, and outdoor relative humidity.

In both the Industrial and Academic buildings, the strongest correlation with current consumption is observed for the power value recorded 24 hours earlier (**Table 3**), followed closely by that recorded 48 hours earlier.

Table 3. Correlation between factors concerning the consumption for both industrial and academic buildings.

	Academic Sector		Industrial Sector	
	Working Hours	Off-Hours	Working Hours	Off-Hours
Temperature	0.205508	0.012566	−0.196301	−0.177106
Irradiation	0.176214	0	0.033515	−0.033712
	Workday	Holiday	Workday	Holiday
Pprev (24 h)	0.816096	0.926576	0.887726	0.154234
Pprev (48 h)	0.710268	0.915227	0.854311	0.272240

Overall, the correlations between the examined environmental parameters and power consumption remain weak, as illustrated by the results for air temperature and irradiation in [Table 3](#). The Academic building benefits from an advanced automated energy management system that promotes highly efficient, sustainable operations. Combined with substantial internal heat gains, sophisticated control mechanisms, localized temperature regulation, and optimized energy strategies, this setup renders external climatic conditions largely irrelevant to its electricity demand. In contrast, the Industrial building exhibits a modest correlation between outdoor air temperature and power consumption during operational hours. However, the strength of this relationship fluctuates by month—for instance, reaching 0.89 in June and 0.93 in October. This factor was therefore retained as a supplementary predictor in the modeling process.

Input data

Certain feature definitions are shared between the two buildings but adapted to reflect their distinct operational profiles. The complete set of predictors considered is described below and summarized in [Table 4](#).

- $P_{prev,24h}$. Power consumption recorded at the equivalent time point 24 hours earlier, scaled to the range [0, 1].
- $P_{prev,48h}$. Power consumption recorded at the equivalent time point 48 hours earlier, scaled to the range [0, 1].
- Holiday. Binary indicator: 1 if the forecast timestamp falls on a holiday, 0 otherwise.
- Base power. Binary indicator: 1 if base-load conditions are present at the forecast timestamp, 0 otherwise.
- Academic: Applies to days where consumption remains continuously above 15 kW ([Figure 3a](#)).
- Industrial: Applies to days where consumption exceeds 78 kW ([Figure 3b](#)).
- Working hours. Binary indicator: 1 if the forecast timestamp occurs during working hours, 0 otherwise.
- Academic: Defined as 7 a.m. to 6:30 p.m. based on [Figure 3a](#).
- Industrial: Defined as 7 a.m. to 7 p.m. (orange periods) or 7 a.m. to 11 p.m. (green periods) based on [Figure 3b](#).
- Air temperature. Outdoor air temperature at the forecast timestamp, scaled to the range [0, 1].

Table 4. Input features are combinations of academic and industrial buildings.

Label type	Feature description	Academic building configurations			Industrial building configurations		
		Basics	+48h	+AirTemp	Basics	+48h	Complete
P _{prev,24 h}	Lagged power (24 h earlier) – Fractional (normalized)	•	•	•	•	•	•
Holiday	Holiday indicator – Boolean	•	•	•	•	•	•
Base Power	Base load indicator – Boolean	•	•	•	•	•	•
Work hours	Working period indicator – Boolean	•	•	•	•	•	•
Air temp.	Outdoor air temperature – Normalized			•			•
P _{prev,48 h}	Lagged power (48 h earlier) – Fractional (normalized)		•			•	•

Analysis of results

Table 5 presents a comprehensive summary of all performance metrics obtained across every combination of partitioning scheme, feature set, model type, and building category. Figure 4 offers a graphical representation of these outcomes, enabling clearer visual comparison and interpretation.




Figure 4. Evaluation metrics results of both (a) Industrial and (b) Academic building.

Table 5. Performance comparison of ML models for power prediction in Academic and Industrial buildings.

		Academic				Industrial						
		Case 1		Case 2		Case 1				Case 2		
		Basics	Basics+ 48h	Basics	Basics+ 48h	Basics	Basics+ 48h	Basics+ AirTemp	Compl.	Basics	Basics+ 48h	Basics+ AirTemp
MSE	Persistent	0.0091	0.0091	0.0091	0.0091	0.0058	0.0058	0.0058	0.0058	0.0063	0.0063	0.0063
	DNN	0.0050	0.0050	0.0052	0.0049	0.0055	0.0053	0.0061	0.0056	0.0060	0.0054	0.0054
	SVM-Sig.	0.0059	0.0056	0.0069	0.0064	0.0057	0.0054	0.0064	0.0053	0.0061	0.0057	0.0057
	SV-Rad.	0.0045	0.0045	0.0050	0.0045	0.0051	0.0050	0.0055	0.0053	0.0063	0.0055	0.0055
	RF	0.0050	0.0050	0.0050	0.0050	0.0052	0.0049	0.0053	0.0051	0.0057	0.0054	0.0054
RMSE	Persistent	0.0954	0.0954	0.0954	0.0954	0.0763	0.0763	0.0763	0.0763	0.0795	0.0795	0.0795
	DNN	0.0704	0.0704	0.0721	0.0701	0.0739	0.0726	0.0784	0.0745	0.0774	0.0733	0.0733
	SVM-Sig.	0.0768	0.0748	0.0831	0.0802	0.0754	0.0732	0.0801	0.0728	0.0780	0.0755	0.0755
	SVM-Rad.	0.0670	0.0670	0.0708	0.0669	0.0714	0.0705	0.0741	0.0727	0.0796	0.0739	0.0739
	RF	0.0707	0.0706	0.0709	0.0708	0.0719	0.0702	0.0726	0.0716	0.0756	0.0737	0.0737
MAE	Persistent	0.0480	0.0480	0.0480	0.0480	0.0531	0.0531	0.0531	0.0531	0.0551	0.0551	0.0551
	DNN	0.0424	0.0402	0.0427	0.0404	0.0551	0.0549	0.0584	0.0557	0.0580	0.0536	0.0536
	SVM-Sig.	0.0468	0.0497	0.0443	0.0429	0.0557	0.0550	0.0601	0.0531	0.0575	0.0543	0.0543
	SVM-Rad.	0.0398	0.0398	0.0410	0.0397	0.0536	0.0526	0.0553	0.0542	0.0640	0.0543	0.0543
	RF	0.0405	0.0405	0.0406	0.0406	0.0531	0.0515	0.0543	0.0528	0.0556	0.0539	0.0539
MAPE	Persistent	18.76	18.76	18.76	18.76	25.49	25.49	25.49	25.49	26.24	26.24	26.24
	DNN	119.73	115.84	118.50	114.61	150.45	153.03	151.20	150.70	142.91	144.44	144.44
	SVM-Sig.	20.16	23.73	16.92	16.09	30.51	32.64	33.64	28.76	31.18	28.21	28.21
	SVM-Rad.	15.34	15.34	16.75	15.28	30.61	29.92	31.36	31.59	44.41	28.97	28.97
	RF	15.34	15.28	15.39	15.36	29.18	28.06	31.16	29.97	29.89	28.31	28.31
R ²	Persistent	0.8202	0.8202	0.8202	0.8202	0.8774	0.8774	0.8774	0.8774	0.8588	0.8588	0.8588
	DNN	0.9021	0.9022	0.8975	0.9030	0.8850	0.8890	0.8706	0.8830	0.8662	0.8800	0.8800
	SVM-Sig.	0.8834	0.8895	0.8637	0.8730	0.8804	0.8871	0.8568	0.8883	0.8640	0.8726	0.8726
	SVM-Rad.	0.9113	0.9113	0.9011	0.9115	0.8925	0.8954	0.8842	0.8887	0.8584	0.8780	0.8780

	RF	0.9015	0.9016	0.9008	0.9009	0.8910	0.8962	0.8890	0.8921	0.8724	0.8788	0.8
CV	Persistent	63.69	63.69	63.69	63.69	74.70	74.70	74.70	74.70	73.35	73.35	73
	DNN	60.89	60.12	59.91	61.34	68.22	67.96	66.92	68.22	65.23	69.44	70
	SVM-Sig.	59.06	55.93	62.70	62.44	69.47	68.05	66.86	69.19	68.39	68.41	66
	SVM-Rad.	59.59	59.59	58.86	59.59	67.71	67.90	68.47	68.25	58.66	67.57	67
	RF	60.24	60.29	60.25	60.26	68.72	68.87	68.43	68.85	67.26	67.35	67

For the Academic building, the results shown in [Figure 4a](#) and [Table 5](#) indicate that the Basics+48h feature set typically yields modest improvements over the Basics configuration in most scenarios. Nonetheless, the gaps in the top-performing RMSE and R² values across configurations remain small, with limited differentiation among the tested conditions. Model-specific performance differences, however, are more pronounced.

In the case of the Industrial building ([Figure 4b](#)), the Basics+48h configuration delivers superior RMSE values across nearly all models and partitioning cases, except for SVM-Sigmoid in Case 1. Comparable improvements appear in R² and MAE. Notably, the Basics+AirTemp setup consistently ranks as the weakest performer, occasionally falling below even the persistent baseline model. The Complete and Basic configurations vie closely for second position across most combinations. In terms of MAPE, the DNN model exhibits a clear deviation from the remaining approaches.

The degree of alignment between actual and predicted values was assessed using linear regression, as in [Eq.8](#), where β captures the linear relationship, and x_0 represents the intercept.

$$\hat{y} = \beta \cdot y + x_0 \tag{8}$$

Predicted values approach the true values most closely when $\beta = 1$ and $x_0 = 0$; smaller deviations correspond to R² values nearer to 1.

[Figure 5](#) illustrates the relationship between predicted and observed consumption for the Academic building. Three distinct consumption regimes—evident in [Figure 3a](#)—can be recognized: the lower band (0–15 kW) corresponds to non-working periods with the chiller inactive; the intermediate band (15–30 kW) reflects non-working periods with the chiller active; and the upper band (above 30 kW) aligns with working hours.




Figure 5. Correlation between P_{true} and $P_{predicted}$ consumption in the Academic building across data division proportions. (a) Case 1 and (b) Case 2.

When examining the strongest configurations (SVM with Radial kernel, Random Forest, and Deep Neural Networks), the performance gap narrows considerably between Case 1 and Case 2, as well as between the basic feature set and the version augmented with 48-hour historical values. This narrow margin points to substantial stability across these modeling approaches. Even so, the single most accurate outcome is delivered by SVM-Radial under Case 2 when paired with the Basics+48h predictors.

Inspection of the Industrial building scatter plots (Figure 6) likewise reveals three broad clusters, albeit with softer boundaries: a low-demand region hovering near 25 kW that corresponds to off-hours baseline usage; a moderate band running from 50 kW to 75 kW that aligns with typical working periods of lighter load; and a high-demand region above 75 kW that captures peak operational intervals.




Figure 6. Correlation between true and predicted power consumption in an industrial building across data division proportions (a) Case 1 and (b) Case 2.

Taken together, Case 1 produces slope values (β) closer to the ideal of 1.0 and smaller intercepts (x_0), which is in line with its generally superior R^2 scores. From this, it follows that Case 1 tends to outperform Case 2 overall. Moreover, appending the 48-hour lagged consumption (Basics+48h) consistently yields the strongest forecasts, except for SVM-Sigmoid. Among the tested algorithms, Random Forest achieves the highest accuracy, although it has only a slight edge over SVM-Radial.

Table 6. Absolute error summary for both Academic and Industrial buildings.

		Academic				Industrial						
		Case 1		Case 2		Case 1				Case 2		
		Basics	Basics+48h	Basics	Basics+48h	Basics	Basics+48h	Basics+AirTemp	Compl.	Basics	Basics+48h	Basics+AirTemp
Mean	SVM-Rad.	0.07	0.07	-0.15	0.07	-0.43	-0.57	0.49	-0.64	-1.62	-0.09	-0.02
	SVM-Sig.	-0.57	-0.41	-1.09	0.08	0.60	-0.09	0.95	0.62	0.56	0.81	0.76

Median	RF	-0.26	-0.26	-0.27	-0.27	-0.15	-0.33	-0.27	-0.47	0.08	0.06	0.09
	DNN	0.43	0.34	-0.06	-0.86	1.35	-0.94	-0.73	-0.77	-1.43	0.10	1.50
	PM	-0.11	-0.11	-0.11	-0.11	0.21	0.21	0.21	0.21	0.46	0.46	0.46
	SVM-Rad.	-0.32	-0.32	-0.47	-0.32	-1.33	-1.33	-0.25	-1.38	-2.87	-0.66	-0.73
	SVM-Sig.	-1.10	-0.58	-1.33	-0.10	-0.06	-0.32	0.19	0.14	0.09	0.46	0.24
	RF	-0.37	-0.37	-0.36	-0.36	-0.82	-1.06	-1.01	-1.05	-0.49	-0.84	-0.54
	DNN	0.17	0.12	-0.19	-0.73	0.36	-1.46	-1.28	-1.25	-2.21	-0.59	0.67
	PM	0.04	0.04	0.04	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	SVM-Rad.	4.33	4.33	4.56	4.32	7.78	7.66	8.22	7.90	8.53	8.06	8.47
Std	SVM-Sig.	7.53	4.97	5.05	4.96	8.20	7.98	8.32	7.92	8.50	8.20	8.70
	RF	4.55	4.55	4.57	4.56	7.84	7.65	7.91	7.79	8.24	8.03	8.33
	DNN	4.62	4.53	4.53	4.52	7.87	7.63	8.39	7.97	8.39	7.99	8.42
	PM	6.16	6.16	6.16	6.16	8.32	8.32	8.32	8.32	8.66	8.66	8.66

Table 6 compiles key descriptive statistics (mean, median, standard deviation) of the absolute prediction errors for each building. The closer these three statistics are to zero, the more reliable the forecasts become: a near-zero mean signals low average deviation, a median near zero reflects roughly symmetric over- and under-predictions, and a low standard deviation indicates that errors remain tightly grouped with minimal dispersion.

The degree to which error distributions concentrate around zero directly reflects model precision. In the Academic building results (**Figure 7a**), Random Forest and DNN produce the narrowest, most peaked distributions; Random Forest further stands out for showing the least sensitivity to changes in partitioning scheme or feature combination. For the Industrial building (**Figure 7b**), Random Forest, SVM-Radial, and DNN all achieve similarly tight error profiles. Yet, the spread among different configurations is noticeably wider in this case.




Figure 7. Absolute error distribution of both (a) academic and (b) industrial buildings.

Figures 8 and 9 illustrate single-day consumption trajectories, both with and without base-load conditions, for the Academic and Industrial facilities, respectively. Every panel employs Case 1 together with the Basics+48h feature set. Although these example days do not uncover novel patterns on their own, they provide concrete, side-by-side views of how closely each algorithm tracks the building's actual load shape under representative conditions. Direct comparison of the predicted curves with measured values reveals the models' ability (or limitations) to reproduce real operational behavior across different usage regimes faithfully.




Figure 8. Prediction of a day for the Academic building in both (a) with and (b) without base load.




Figure 9. Prediction of a day for an industrial building in both (a) with and (b) without base load.

Discussion

Building upon the quantitative evaluation detailed in the preceding section, the following discussion interprets the key outcomes and their broader significance in direct relation to the five research questions and associated hypotheses outlined earlier. Each hypothesis is systematically evaluated for its validity, using empirical evidence to determine the degree of alignment with initial expectations. This critical review sheds light on both the capabilities and constraints of the employed predictive approaches, while also highlighting their practical value for advancing short-term electricity demand forecasting in diverse building contexts.

- Examination of the consumption profiles and their statistical relationships with external variables revealed that predictors included in the Basics set—those tied closely to occupancy patterns and equipment operation—emerged as the dominant drivers of electricity use in both the Academic and Industrial facilities. Environmental variables, in contrast, displayed generally weak associations with demand. The sole notable exception was a pronounced month-dependent correlation between outdoor air temperature and power consumption during operating hours in the Industrial building, which supported its selective inclusion as a predictor in that specific setting.
- As previously highlighted, the Industrial building exhibited a strong, seasonally varying link between air temperature and electricity demand during work periods—a relationship absent in the Academic building. Yet, when performance was assessed under the Basics+AirTemp configuration (using metrics such as RMSE, MAE, and R^2), improvement over the simpler Basics setup occurred in only one instance: Case 2 with SVM-Radial. In numerous other scenarios, results were actually inferior to those of the Persistent baseline model.

These observations lead to the conclusion that external weather conditions exert only a marginal influence on electricity consumption in either building. In the Academic facility, elevated internal heat gains, combined with sophisticated heat recovery mechanisms and automated controls, largely insulate energy use from meteorological or orientation-related effects; instead, demand is primarily governed by occupancy schedules and the academic calendar. This interpretation is reinforced by the consistent and substantial gains in forecast accuracy achieved when extending the Basics set with 48-hour lagged consumption (Basics+48h) across most tested conditions. Consequently, under circumstances characterized by dominant internal loads, predictive models may reasonably exclude climatic inputs without compromising reliability. Overall, Hypothesis H1 receives partial confirmation, whereas Hypothesis H2 is disproven, particularly when accounting for the high proportion of internal loads relative to total consumption observed in both cases.

- Across the full range of experiments, the Academic building consistently yielded more favorable RMSE values than the Industrial building. It also produced lower MAPE and CV values, along with higher R^2 scores, indicating markedly superior

model performance in the finer-grained (15-minute) dataset compared with the coarser (hourly) dataset. This superiority was further evident in the much tighter clustering of absolute errors around zero within the Academic results.

These patterns confirm that increasing the temporal resolution of data collection enhances the precision of power demand forecasts, thereby supporting Hypothesis H3. With respect to dataset partitioning, Case 1 demonstrated greater generalization capability, attributable to its more evenly balanced validation subset relative to Case 2.

- In nearly every configuration across both buildings—except for Case 2 in the Industrial setting—the SVM model equipped with the Radial kernel outperformed its Sigmoid counterpart. The absolute error distributions underscored this advantage particularly clearly in the Academic building, where SVM-Radial produced a noticeably tighter distribution of errors. Although the margin was narrower in the Industrial building, a similar trend toward reduced dispersion around zero remained apparent.

Collectively, these findings uphold Hypothesis H4, affirming that the Radial kernel confers superior predictive capability to SVM by more effectively capturing the intricate, nonlinear dynamics present in the consumption data of both building types.

- For the Academic building, evaluation metrics indicated broadly comparable behavior between DNN and RF. Yet, RF consistently edged ahead, recording the lowest RMSE values and exhibiting the most compact, sharply peaked absolute error distribution. In the Industrial building, DNN lagged behind both SVM-Radial and RF, with RF again showing a slight but consistent advantage.

In light of these results, SVM-Radial emerges as the most appropriate choice for the Academic building, delivering the strongest overall performance. In contrast, RF proves better suited to the Industrial context. Thus, Hypothesis H5 is rejected: contrary to expectations, the DNN architecture did not outperform the other methods in either scenario. Moreover, DNN showed markedly higher MAPE values than RF and SVM-Radial. Given its reasonable performance across other metrics, issues related to data scaling, preprocessing, or architectural design appear unlikely; instead, the elevated MAPE most plausibly stems from overfitting. This would imply that while DNN often generates accurate predictions, it occasionally produces large deviations in specific instances, resulting in disproportionately high relative errors that inflate the MAPE statistic.

Future research directions

As research into building energy consumption forecasting continues to evolve, several promising directions merit deeper investigation. These avenues aim not only to strengthen the real-world viability of the developed models but also to tackle the multifaceted challenges present in operational energy environments. The principal areas proposed for subsequent work include:

Deployment-related practical issues

Future efforts will examine the real-world obstacles encountered when implementing predictive models within actual energy management systems. This encompasses analyzing hardware costs associated with various machine learning techniques, evaluating the computational demands of different algorithms, and formulating effective approaches to embed forecasting tools into existing building energy control frameworks. Overcoming these barriers is vital to translating high theoretical performance into dependable, field-ready solutions for energy prediction.

Reducing the impact of input data imperfections

The current study excluded entire days featuring three or more successive missing values or outlier flags. Upcoming investigations will prioritize techniques to mitigate the adverse effects of erroneous or incomplete input data, thereby increasing the robustness and trustworthiness of model outputs. Improving data reliability and resilience will be essential to ensuring stable performance across a wide range of real operating scenarios.

In-depth examination of load behavior

Although the Industrial building exhibits relatively stable consumption driven by a fixed production timetable, the Academic building displays considerably greater variability driven by fluctuating occupancy levels and diverse internal operations. Additional research is required to characterize better how these irregular usage dynamics shape energy demand in educational facilities. Gaining clearer insight into this variability will be key to designing specialized forecasting models that accurately represent diverse behavioral patterns.

Development and application of hybrid approaches

Plans include exploring hybrid modeling frameworks that merge physics-based simulation with live measurement data to achieve superior prediction precision and enhanced operational control. Such integrated strategies hold particular promise for constructing digital twins—virtual replicas of physical systems that support advanced energy oversight and adaptive responses. By integrating continuous real-time inputs with simulated scenarios, digital twins can dynamically respond to changing circumstances, thereby enabling greater energy savings and more intelligent resource distribution. This line of inquiry will also facilitate exploration of predictive modeling's broader utility in complex, multifaceted energy contexts, ultimately amplifying the real-world relevance and reach of the present findings.

Conclusion

First, the analysis demonstrates that external weather conditions exert only a limited influence on electricity consumption in the studied buildings. Nevertheless, augmenting the basic feature set with consumption values from 48 hours earlier consistently yields substantial gains in forecasting precision across the majority of tested conditions. The outcomes also substantiate that increasing the temporal granularity of data acquisition markedly improves model accuracy. Furthermore, adopting a more equitable, carefully balanced dataset partitioning substantially enhances overall model effectiveness.

Second, the Support Vector Machine with the Radial kernel consistently delivered higher predictive performance than its Sigmoid counterpart across both building types, owing to its greater ability to represent intricate nonlinear dependencies in the data. More specifically, SVM-Radial achieved the top results in the Academic building, whereas Random Forest proved marginally more effective in the Industrial building.

Third, the Deep Neural Network model exhibited distinctly higher MAPE values than the other approaches, even though it performed adequately on alternative metrics. This pattern points to a susceptibility to overfitting: the model tends to produce accurate predictions in most cases but generates large deviations in some, thereby disproportionately increasing the mean absolute percentage error through outsized relative discrepancies.

Taken together, these results emphasize the critical need to carefully match data sampling intervals, predictor variables, and modeling techniques to the specific functional and operational traits of each building to maximize the precision and reliability of electricity demand forecasts.

Finally, the baseline Persistent model was surpassed by the machine learning approaches in the vast majority of configurations and scenarios examined. This indicates that even modest refinement of these techniques already produces highly valuable forecasting capabilities worthy of serious consideration. Such performance underscores the substantial promise these methods hold for improving energy management strategies across a range of building categories, affirming their practicality and effectiveness for deployment in everyday settings.

Abbreviations

AI: Artificial Intelligence

ANN: Artificial Neural Network

CNN: Convolutional Neural Network

CV: Coefficient of Variation

DL: Deep Learning

DNN: Deep Neuronal Network

GRU: Gated Recurrent Unit

IoT: Internet of Things

LSTM: Long Short-Term Memory

MAE: Mean Absolute Error

MAPE: Mean Absolute Percentage Error

ML: Machine Learning

MSE: Mean Squared Error

PM: Persistent Model

R2: Coefficient of Determination

RBF: Radial Basis Function

ReLU: Rectified Linear Unit

RF: Random Forest

RMSE: Root Mean Squared Error

RNN: Recurrent Neural Networks

STLF: Short-term Load Forecasting

SVM: Support Vector Machine

ZEB: Zero-Emission Building

Acknowledgements

None

Conflict of interest

None

Financial support

This work was developed within the ELKARTEK program, supported by the Basque Government under the Projects RESINET (KK2023/00042) and BERDEA (KK-2023/00037).

Ethics statement

None

Received: 08 Oct 2022

Revised: 24 Nov 2022

Accepted: 06 Jan 2023

Published online: 20 March 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

IEA. Energy Efficiency; Technical Report; International Energy Agency: Paris, France, 2023.

Zhao T, Xu J, Zhang C, Wang P. A monitoring data based bottom-up modeling method and its application for energy consumption prediction of campus building. *J Build Eng*. 2021;35:101962.

Liu X, Liu X, Luo X, Fu H, Wang M, Li L. Impact of different policy instruments on diffusing energy consumption monitoring technology in public buildings: evidence from Xi'an, China. *J Clean Prod*. 2020;251:119693.

Wu Y, Wang W, Sun Y, Cui Y, Duan D, Deng S. An equivalent temperature drop method for evaluating the operating performances of ASHP units jointly affected by ambient air temperature and relative humidity. *Energy Build*. 2020;224:110211.

Nematchoua MK, Nishimwe AMR, Reiter S. Towards nearly zero-energy residential neighbourhoods in the European Union: a case study. *Renew Sustain Energy Rev*. 2021;135:110198.

Zhang Q, Yang LT, Chen Z, Li P. A survey on deep learning for big data. *Inf Fusion*. 2018;42:146-57.

Amasyali K, El-Gohary NM. A review of data-driven building energy consumption prediction studies. *Renew Sustain Energy Rev*. 2018;81:1192-205.

Somu N, R GRM, Ramamritham K. A hybrid model for building energy consumption forecasting using long short-term memory networks. *Appl Energy*. 2020;261:114131.

Bakar NNA, Hassan MY, Abdullah H, Rahman HA, Abdullah MP, Hussin F et al. Energy efficiency index as an indicator for measuring building energy performance: a review. *Renew Sustain Energy Rev*. 2015;44:1-11.

Bourdeau M, Zhai XQ, Nefzaoui E, Guo X, Chatellier P. Modeling and forecasting building energy consumption: A review of data-driven techniques. *Sustain Cities Soc*. 2019;48:101533.

Wei Y, Zhang X, Shi Y, Xia L, Pan S, Wu J et al. A review of data-driven approaches for prediction and classification of building energy consumption. *Renew Sustain Energy Rev*. 2018;82:1027-47.

Ramokone A, Popoola O, Awelewa A, Temitope A. A review on behavioural propensity for building load and energy profile development – model inadequacy and improved approach. *Sustain Energy Technol Assess*. 2021;45:101235.

Eren Y, Küçükdemiral I. A comprehensive review on deep learning approaches for short-term load forecasting. *Renew Sustain Energy Rev*. 2024;189:114031.

Kapp S, Choi JK, Hong T. Predicting industrial building energy consumption with statistical and machine-learning models informed by physical system parameters. *Renew Sustain Energy Rev*. 2023;172:113045.

Oliveira N, Sousa N, Praça I. Deep learning for short-term instant energy consumption forecasting in the manufacturing sector. In: Omatu S, Mehmood R, Sitek P, Cicerone S, Rodríguez S, eds. *Distributed computing and artificial intelligence, 19th International Conference*. Cham (Switzerland): Springer; 2023. p. 165-75.

El Aziz Ahmed EA, Ibrahim RA, Abdelsalam AK. A comparative analysis for machine learning-based short-term load forecasting techniques. In: *Proceedings of the 2023 IEEE 6th International Electrical and Energy Conference (CIEEC)*. Hefei (China); 2023. p. 1166-71.

Ungureanu S, Topa V, Czikier AC. Deep learning for short-term load forecasting—industrial consumer case study. *Appl Sci*. 2021;11:10126.

Seefoo Jarquin CS, Gandelli A, Grimaccia F, Mussetta M. Short-term probabilistic load forecasting in university buildings by means of artificial neural networks. *Forecasting*. 2023;5:390-404.

Baek SJ, Yoon SG. Short-term load forecasting for campus building with small-scale loads by types using artificial neural network. In: *Proceedings of the 2019 IEEE Power and Energy Society Innovative Smart Grid Technologies Conference (ISGT)*. Washington, DC (USA); 2019. p. 1-5.

Chávez P, Martini I. Applying neural networks for short and long-term hourly electricity consumption forecasting in universities: a simultaneous approach for energy management. *J Build Eng*. 2024;97:110612.

Ghoddusi H, Creamer GG, Rafizadeh N. Machine learning in energy economics and finance: A review. *Energy Econ*. 2019;81:709-27.

Al-Saudi K, Degeler V, Medema M. Energy consumption patterns and load forecasting with profiled CNN-LSTM networks. *Processes*. 2021;9:1870.

Johnson JM, Khoshgoftaar TM. Survey on deep learning with class imbalance. *J Big Data*. 2019;6:27.

Al-Saudi K, Degeler V, Medema M. Energy consumption patterns and load forecasting with profiled CNN-LSTM networks. *Processes*. 2021;9:1870.

Adadi A. A survey on data-efficient algorithms in big data era. *J Big Data*. 2021;8:24.

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.

Zheng G, Feng Z, Jiang M, Tan L, Wang Z. Predicting the energy consumption of commercial buildings based on deep forest model and its interpretability. *Buildings*. 2023;13:2162.