ORIGINAL RESEARCH **Open access**

An LSTM-Based Recommendation Framework for Predictive Energy Management in Smart Buildings

Aisha Bello[1*], Ibrahim Musa[2], Zainab Sule[1]

Abstract

Precise forecasting and load refinement in building energy management are vital for improving efficiency and reducing consumption. Conventional approaches are often inadequate when faced with the nonlinear and fluctuating nature of building energy usage trends. As sophisticated data collection via smart sensors has advanced, there is an increasing need for intelligent systems that can harness this data to generate actionable guidance. The present study addresses this gap by developing a recommendation system based on three machine and deep learning algorithms—Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and XGBoost—to forecast and optimize efficiency levels for space heating, ceiling cooling, and free cooling installations. Our solution capitalizes on the strengths of these algorithms, with LSTM achieving the best overall performance, to project energy usage at both hourly and daily scales, enabling precise adjustments and streamlined energy oversight. The methodology comprises extensive data preparation, including hierarchical gap completion and label encoding of categorical attributes, followed by converting raw inputs into efficiency ratings. The deep learning design, composed of sequential layers, captures long-range dependencies inherent in the data, while performance is maximized through grid-search-driven hyperparameter tuning. Outcomes indicate elevated predictive precision, with R-squared values indicating that the model accounts for up to 97.2% of the variance in hourly space heating, 95.2% in daily ceiling cooling, and 93 % in daily free cooling energy usage. In addition, graphs are decoded using OpenAI's GPT-4 model, deepening comprehension and enabling actionable insights. This added layer of interpretation renders predictive outputs more transparent, underpinning more enlightened energy management choices. The value of this work lies in its capacity to reshape energy management routines in built environments, providing a robust tool for refining heating and cooling systems and advancing holistic energy efficiency.

Keywords LSTM, Deep learning, Energy prediction, HVAC optimization, Smart buildings

*Correspondence:
Aisha Bello
aisha.bello@gmail.com
[1] Department of Smart Energy Engineering, Faculty of Engineering, University of Ibadan, Ibadan, Nigeria
[2] Department of Intelligent Energy Analytics, Faculty of Engineering, Ahmadu Bello University, Zaria, Nigeria

Introduction

Buildings account for a considerable share of global primary energy demand (40%) and greenhouse gas emissions (30%) [1]. Strengthening building energy performance, a cornerstone for reducing ecological impact, focuses on principal energy consumers such as lighting, cooling, and heating [2-4]. Capturing dynamic thermal responses under non-stationary conditions is indispensable [5]. ASHRAE GreenGuide underscores the need to understand energy sources to streamline building operations, singling out space conditioning, lighting, and plug loads as the predominant end uses [6]. Passive solutions (e.g., double-skin façades) and active solutions (e.g., variable refrigerant flow systems) both offer avenues for

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

advancing energy performance [7]. Simulation-based inquiries reveal prospective savings and operational enhancements in building systems [8, 9].

Consequently, to curtail energy usage, energy modeling should be embedded in the initial design stage, while ongoing building use should inform energy predictions across varied and uncertain scenarios [6]. Anticipating demands shaped by numerous building parameters with nonlinear temporal traits poses a challenge; nonetheless, the proliferation of intelligent sensors and big data technologies has elevated the application of data-driven models for thermal load and energy forecasting in buildings [10]. Such models have attracted interest owing to their adaptability, simplicity, and ease of deployment; a range of algorithms exists to assess and forecast building energy requirements [11].

Multiple avenues have been pursued to investigate machine learning and deep learning models for energy use prediction in a recommendation context. Traditional tools like linear regression and decision trees offer clarity but can falter when applied to datasets exhibiting pronounced temporal interdependence and nonlinearity in building energy records [12, 13]. Alternatives such as random forests and gradient boosting machines accommodate nonlinearities yet may require substantial feature engineering and lack the capacity to represent sequential information effectively. Convolutional neural networks (CNNs) show promise for spatial data tasks but may struggle to capture long-term temporal relationships [14]. In comparison, deep learning structures such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) have demonstrated prowess in modeling sequential data, particularly in energy consumption domains [15].

The present study develops an LSTM-based model and compares its performance with GRU and XGBoost, establishing a comparative framework to identify the most effective approach for optimizing building energy performance. LSTMs can learn long-range dependencies and move between short- and long-term signals, a capability essential for predicting energy usage shaped by ambient conditions and occupancy rhythms in buildings [16].

Our inquiry focused on devising a recommendation system to enhance building energy efficiency, with attention to three vital categories of energy expenditure: space heating (SH), cooling ceiling (CC), and free cooling (FC). Space Heating denotes the energy consumed to heat indoor spaces during chilly periods and raise temperatures to occupant-comfortable levels [17]. Cooling Ceiling refers to the energy used to chill interior zones, particularly in hotter seasons, via overhead cooling systems to maintain comfortable indoor climates [18]. Free Cooling, conversely, draws upon an external air or water supply to lower internal temperatures without mechanical chilling, exploiting favorable outdoor conditions to conserve energy [19].

Sustainability within building operations is becoming increasingly pivotal against the backdrop of global efforts to combat climate change and reduce environmental harm [20, 21]. Hence, our work aims to advance sustainable targets by deploying effective energy predictors, such as the LSTM model. In this way, beyond offering a chance to curtail energy waste, the proposed recommendation system for SH, CC, and FC is engineered to anticipate and dynamically guide energy use.

A central facet of this investigation is the generation of graphical outputs using OpenAI's Generative Pretrained 4.0 (GPT-4.0) model, a step that presents both considerable hurdles and opportunities. One significant hurdle is faithfully conveying intricate motifs in energy consumption data, which are shaped by diverse factors such as seasonal variations, occupancy cycles, and meteorological influences. Visual portrayals of space heating, ceiling cooling, and free-cooling setups often reveal complex interdependencies that are difficult to disentangle without advanced aids. The assimilation of the GPT-4 model enriches this undertaking by furnishing nuanced, context-specific interpretations and suggestions grounded in the illustrated data. This tactic demystifies convoluted charts and facilitates operational decision-making by spotlighting salient trends and anomalies. To the best of our knowledge, this is the first recommendation system to incorporate such interpretive faculties, yielding a more refined and precise understanding of energy use dynamics. The weight of this interpretive layer lies in its power to buttress more effective and enlightened energy management approaches, culminating in greater efficiency and lower operational expenditures within built contexts. The contributions of the study are as follows:

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

i. Establishment of a novel recommendation system harnessing the LSTM model to project and fine-tune SH, CC, and FC energy usage in buildings.

ii. Benchmarking LSTM, GRU and XGBoost models to determine the most effective architecture for forecasting building energy consumption.

iii. Progression of energy efficiency and sustainability by wielding advanced machine learning methodologies to dynamically tailor building energy operations in response to real-time environmental stimuli and occupancy profiles.

iv. Unveiling of a novel recommendation system that melds OpenAI's GPT-4 for the interpretation of graphical energy consumption analyses, supplying granular, context-specific intelligence to fortify decision-making and operational performance.

Related Work

Coordinating heating and cooling system operations through tailored suggestions has been shown to improve building energy performance [22]. Earlier research has used data-driven analytics and machine learning to formulate bespoke HVAC refinement directives rooted in past usage patterns, surrounding environmental signals, and occupant behavior [23], generally proposing alterations to temperature targets, operational timing, and other procedural tweaks to balance occupant comfort and resource efficiency.

Work employing reinforcement learning (RL) has advanced these platforms by enabling flexible, situation-sensitive regulatory schemes. As one illustration, recEnergyW draws building users into energy-frugal habits via Markov decision process formulations [24]. In contrast, parallel investigations have deployed deep reinforcement learning (DRL) to adjust indoor thermal setpoints on the fly to increase solar (PV) uptake, yielding demonstrable reductions and comfort dividends [25]. Complementary RL-rooted systems, such as HERS [26], incorporate occupant preferences for demand-curbing measures, although these designs typically omit PV arrays, storage banks, or microgrid scenarios. Broad-gauge RL-oriented HVAC tuning blueprints have likewise been set forth [27], yet translating these to real-world deployments and keeping computational loads manageable remain persistent hurdles.

Established recommendation paradigms, spanning knowledge-engineered expert shells and condition-action rule analytics, have been used in both office and residential buildings [28, 29], often drawing on live building telemetry to flag wasteful patterns and nudge behavioral shifts. For all their practical utility, such methods commonly require extensive manual tailoring, offer limited flexibility across heterogeneous building stocks, and lack integration with more powerful learning-centric engines.

Contemporary lines of inquiry have woven deep learning into recommender architectures, drawing on IoT infrastructure and edge processing to enable elastic, instantaneous insight generation [30]. Even so, a sizable portion still rests on fixed, predetermined suggestion templates and does not fully mine collaborative sifting or outcome-driven individualization. LSTM-anchored approaches have also been used to project building energy trajectories, bolstered at times by transfer learning to extend predictive reach across disparate climatic settings [28]. That said, these contributions have not addressed the task of autonomously making sense of prediction streams, and their applicability across structurally distinct buildings remains circumscribed.

While LSTM and kindred forecasting engines have been widely adopted for energy outlooks, the landscape lacks frameworks that can independently convert projection results into ready-to-use operational counsel without human mediation. To the best of our knowledge, the present study marks the first instance of integrating OpenAI’s GPT-4 into an LSTM-driven conduit for building energy orchestration. The inherent language comprehension strengths of GPT-4 make it feasible to transform high-dimensional forecast readouts into setting-aware operational plans, thereby spanning the chasm between bare numerical predictions and strategic action. This fusion remedies a deficiency in earlier contributions

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

where sense-making and advisory generation remained either operator-dependent or strictly rule-prescribed. The innovation lies in combining LSTM-based forecasting with GPT-4-facilitated explanations for live, elastic, and architecture-tailored energy optimization **(Table 1).**

Table 1. Comparison of existing recommendation systems.

Ref.	Year	Research Focus	Dataset	Method	Outcome	Limitations
[24]	2020	Energy-saving recommendation system	Not available (NA)	Deep reinforcement learning	Achieved energy reduction between 19% and 26%	Requires manual post-processing due to insufficient recommendation acceptance data
[25]	2021	Smart home energy management	Collected every 30 minutes daily over a 3-month period	Deep reinforcement learning (DRL)	Average energy savings of 8%, with peaks reaching 16% during summer	Difficulty in capturing user preferences and adapting to photovoltaic (PV) policy changes
[26]	2022	Smart home energy optimization	ARAS Activity Dataset	Deep reinforcement learning	Enhanced electricity cost efficiency and user comfort	Does not consider PV-enabled devices, storage systems, microgrids, or data-sharing mechanisms
[27]	2020	HVAC control and demand response optimization	TMY3 dataset	Reinforcement learning	Achieved up to 22% weekly energy savings	Faces challenges in real-world deployment and computational efficiency
[28]	2022	Energy-saving recommendation system	Data from ten pilot buildings	Model-based expert system	Delivered an overall energy saving of 12.2%	Requires manual parameter configuration and has limited metadata support
[29]	2018	Occupant-aware energy-saving recommendations	Historical location records	Gaussian distribution fitting	Increased energy savings by 25%	Strong dependence on historical location data
[22]	2020	Human engagement-based recommendation systems	Real-time sensor data	Sensor- and actuator-driven recommendation engine	Improved recommendation acceptance rate by 26%	Scalability issues and uncertain generalization beyond initial experiments
[30]	2023	Energy efficiency using IoT systems	Data from 5 office environments	LSTM model	Achieved 89.06% accuracy	Needs further improvement for goal-driven and collaborative filtering applications
[31]	2023	Building energy load forecasting	Hospital prototype	LSTM with transfer learning	R^2 score of 0.97	Limited robustness across different

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

			dataset (U.S. DOE)			building types and weather conditions

Dataset

The climatological variable records used to feed our recommender framework for cooling, free cooling, and heating loads were obtained from a purpose-built meteorological station installed on the rooftop of the European Central Bank (ECB) tower. This collection provides pertinent atmospheric readings—covering air temperature, moisture content, airflow velocity, and other parameters relevant to the structure's ambient envelope. The sensor elevation is deliberately set to register measurements that closely track the actual atmospheric envelope surrounding the building.

Pairing this feed with the Building Automation System (BAS) widens the conduit for data retrieval and exploitation. By virtue of the BAS's unified framework for signal routing and archival, supervisory stations can equip facility stewards and system handlers with current weather intelligence. This competency supports immediate decision-making cycles informed by real-time meteorological data.

Furthermore, the arrangement spotlights ancillary upsides of the BAS's authority over HVAC circuits. By internalizing longitudinal shifts and recurring motifs, occupants and other invested observers within the building can trace how distinct weather drivers modulate cooling, free cooling, and heating draws over time. Amassing this statistical reservoir proves instrumental for fine-grained dissection to surface correlations linking atmospheric inputs with energy uptake indices.

Space heating consumption and air temperature

The spatial mapping of space-heating draw against outdoor air thermal level is shown in **Figure 1**, as logged by the Building Weather Station (BWS). The plotted evidence discloses a robust negative coupling, plainly visible across the collected samples. When air thermal readings decline, heating draw rises sharply, topping out within the −10 to 0 °C thermal bracket. Once readings exceed 0 and continue to rise, a discernible tapering of heating draw sets in, subsiding stepwise to a marginal level as the exterior thermal mark nears 15 °C. Examining the concentration clusters below 15 °C points to an elevated heating imperative in colder regimes, indicating that exterior temperature governance is part of the building's energy calculus.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

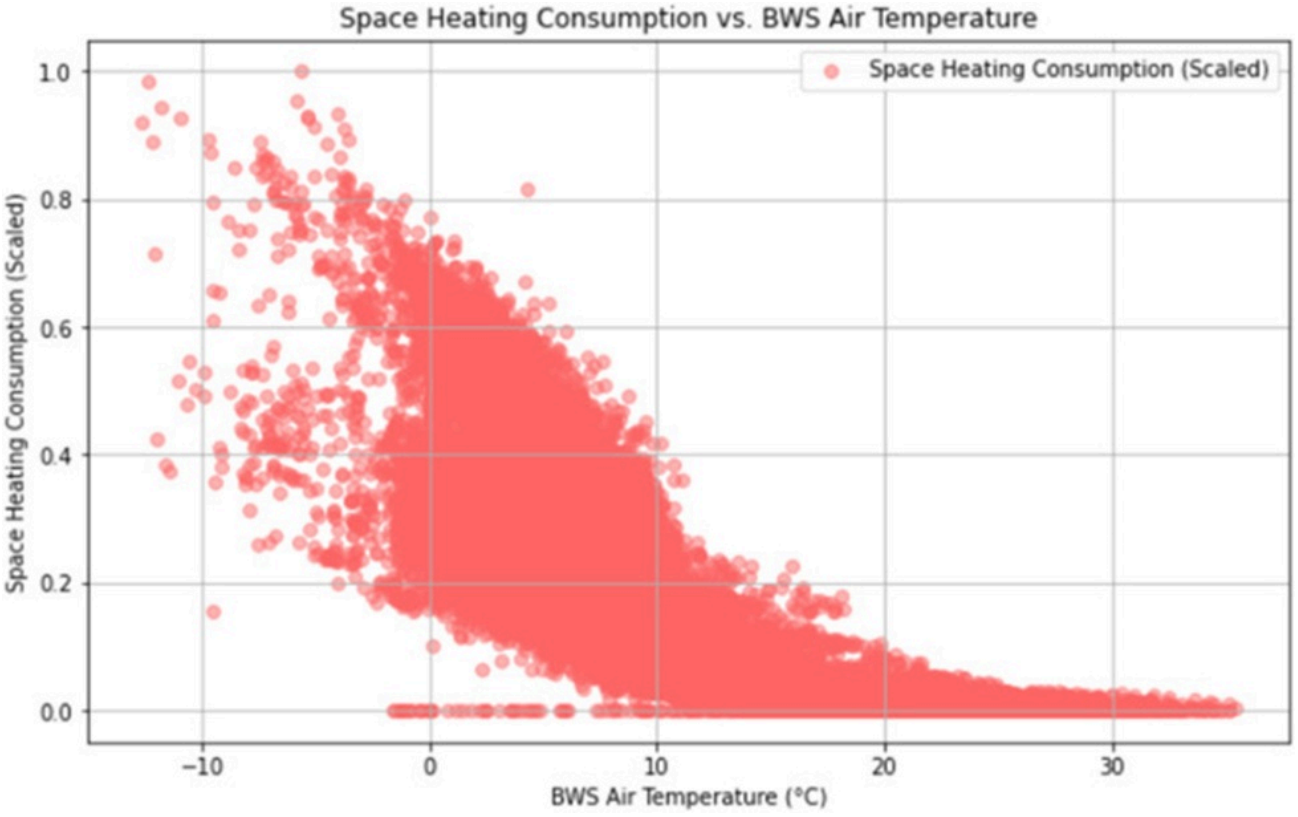

Figure 1. Relationship between space heating consumption and external air temperature.

Ceiling cooling and air temperature

As shown in **Figure 2**, the ceiling cooling draw aligns with the BWS air thermal readings. The scatter reveals an evident gradient: cooling draw hovers near zero until roughly 7 °C. Past this inflection, a steep surge in cooling draw materializes as air temperature ascends, with the thickest point clusters recorded in the 15 to 30 °C corridor. The contour of this relationship suggests that the cooling apparatus is chiefly called into action at higher temperatures to secure indoor conditions that are comfortable for occupants. The remarkably broad dispersal of readings, most notably above 7 °C, underscores that cooling work rates intensify substantially at higher thermal bands and underscores the need for effective interior heat extraction. This signals that the predominant share of cooling burden arises within a confined temperature envelope, indicating that structural behavior tracks the outdoor thermal environment.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

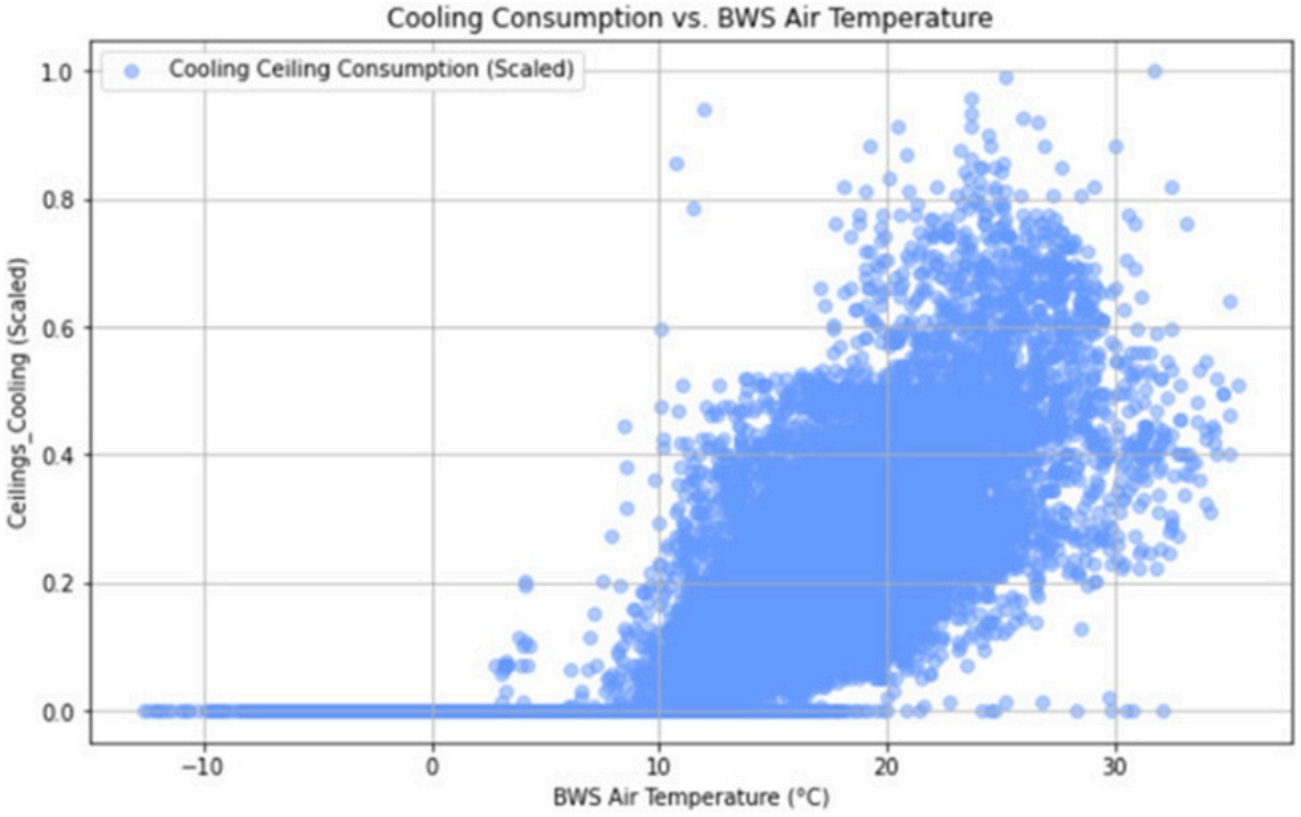

Figure 2. Relationship between cooling consumption and external air temperature.

Free cooling and air temperature

Figure 3 plots the free cooling draw against the air thermal level captured at the structure through BWS. The portrayal sketches a rather contrasting picture, in which free cooling is primarily used when the outdoor temperature is between 0 and 15 °C. This bracket carries weight, as becomes apparent at 15 °C, beyond which the free-cooling activation drops to a very modest level. This stretch indicates that the arrangement achieves its highest point load when operating within its free-cooling activation band, which exploits cooler outdoor air to reduce reliance on compressor-based chilling. Crucially, the relative point density peaks exceptionally high within the approximate span of 0–10 °C, which seems to coincide with the crest of free cooling engagements. When the outside thermal state falls beneath 0 °C or rises above 15 °C, free cooling engagement retreats, signaling that its effectiveness is maximized at middling external temperatures; it demonstrates that free cooling tactics deserve embedding within power deployment blueprints, above all when the outdoors sits within the sweet-spot interval, as shown here to be the case the bulk of the time.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

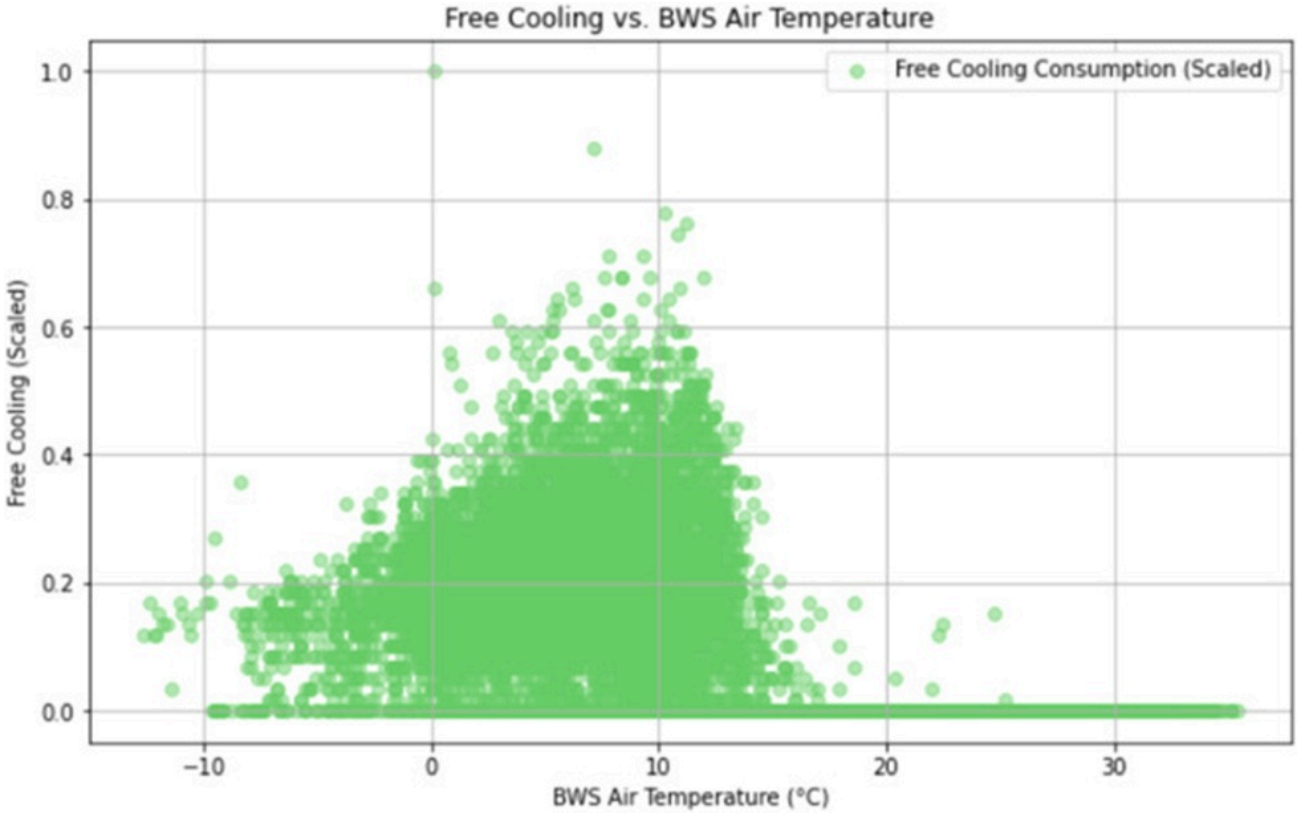

Figure 3. Relationship between free cooling and external air temperature.

Interaction between three energy consumption

Figure 4 presents standardized consumption of three energy streams—space heating, ceiling cooling, and free cooling—across the winter, spring, summer, and fall seasons from January 2020 to October 2022. Over the full monitoring stretch, a pronounced seasonal cadence stands out. Space heating draws crests in the winter months, mirroring the need for greater warming to safeguard indoor habitability under chillier boundary conditions. By contrast, ceiling cooling draw escalates sharply during the summer phase, keeping pace with higher outdoor thermal loads and the resulting demand for heat rejection. Free cooling exhibits seasonal swings. Still, it asserts itself most vigorously in the shoulder seasons, such as spring and fall, when middling outdoor temperatures permit harnessing external air for cooling. Tellingly, the stretches where elevated heating and cooling demands coincide point toward unrealized optimization leeway, especially around peak clipping and the bolstering of whole-building energy yield.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

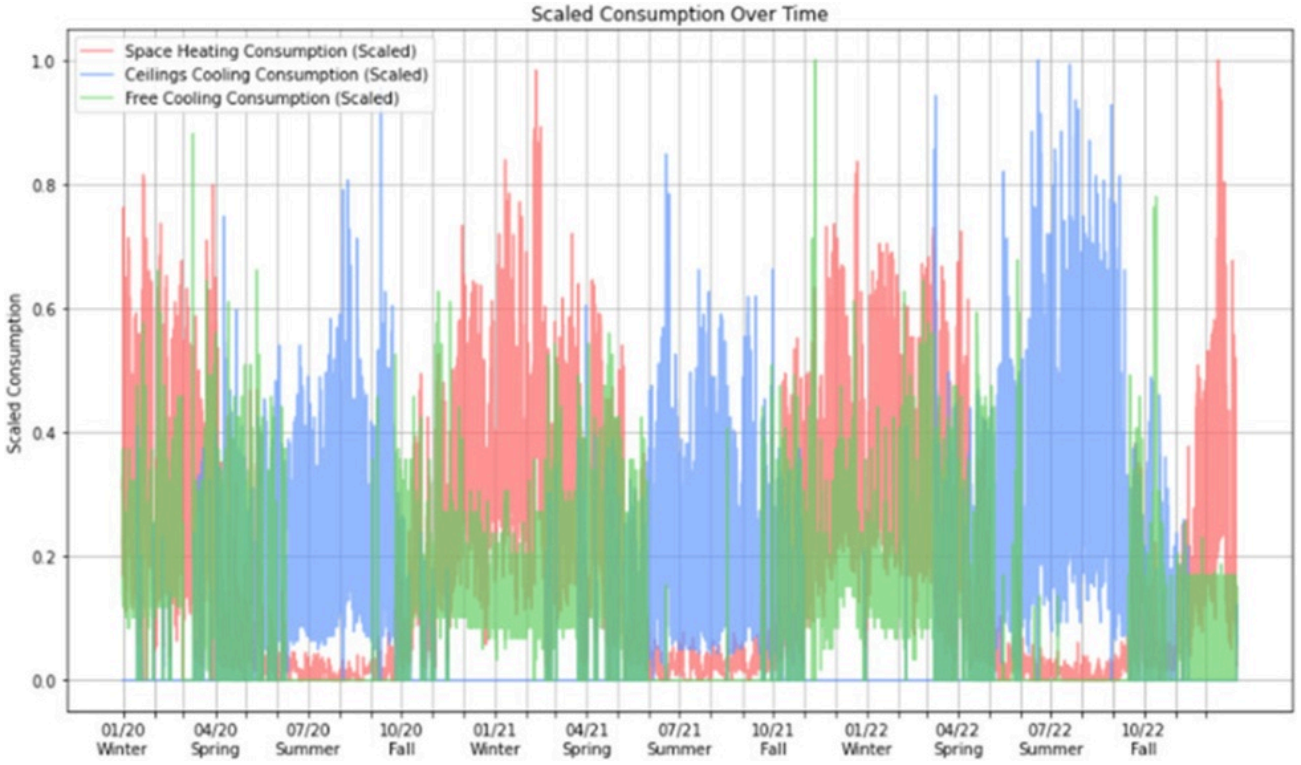

Figure 4. Scaled consumption of space heating, ceiling cooling, and free cooling over time across different seasons.

Proposed Methodology

The present effort is directed at constructing a Long Short-Term Memory (LSTM) neural predictor to estimate efficiency bands for ceiling cooling, free cooling, and heating draw, with the overarching aim of powering an energy advisory platform. The LSTM construct is measured against two well-known benchmarks: the Gated Recurrent Unit (GRU) and XGBoost, enabling an apples-to-apples comparison of their strengths across both daily and hourly forecasting horizons. These side-by-side evaluations help assess each learner's readiness to streamline energy stewardship in built spaces. **Figure 5** supplies a bird’s-eye view of the suggested pipeline.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

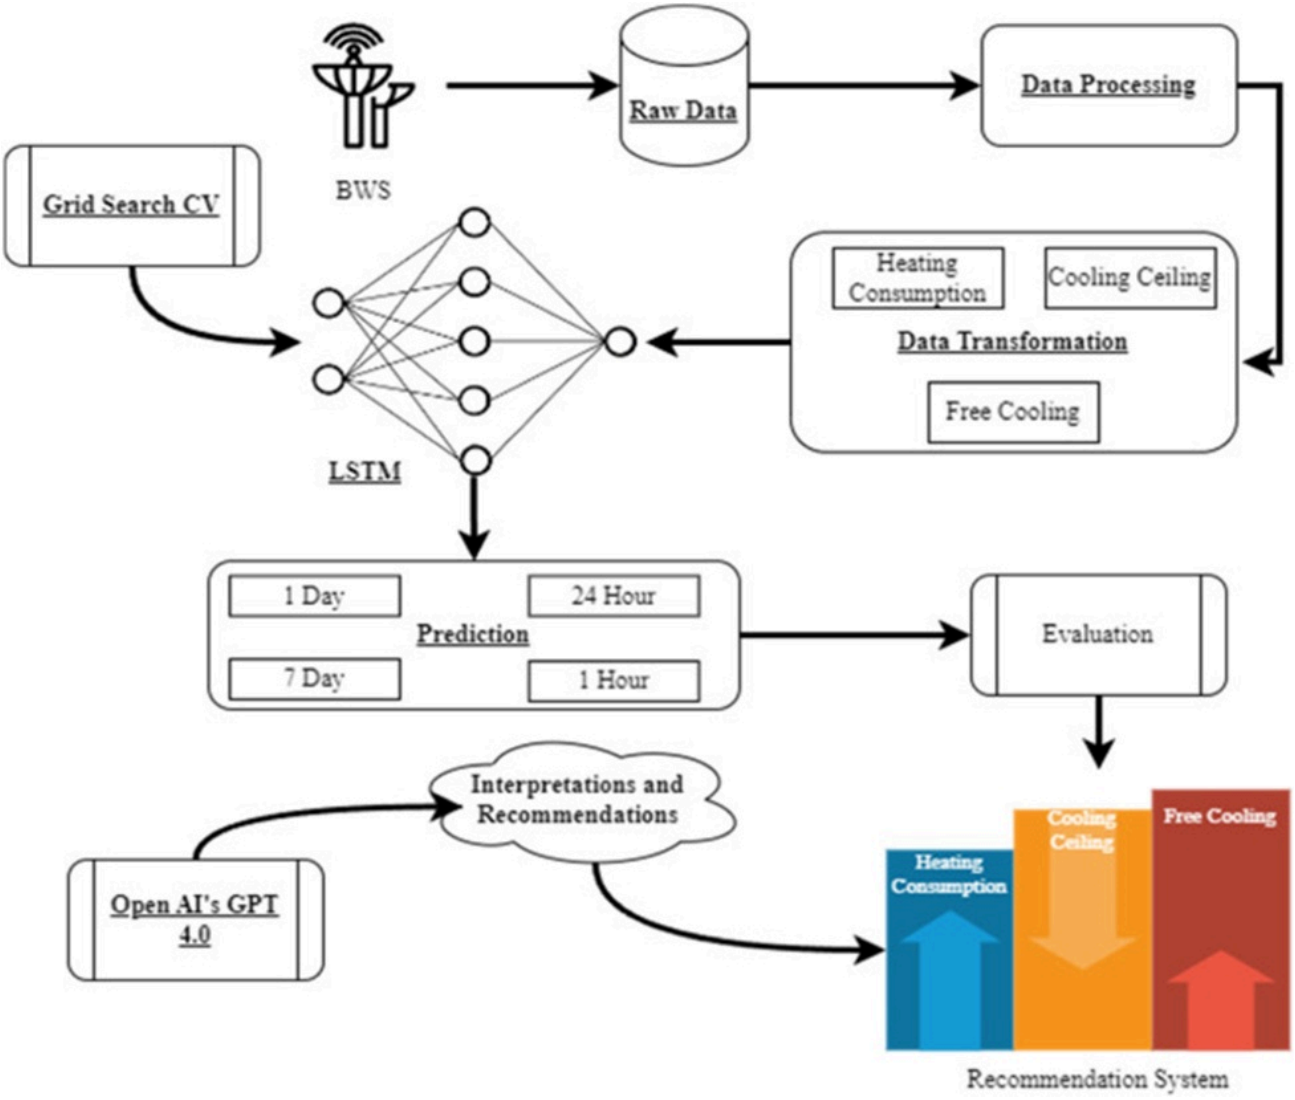

Figure 5. Proposed methodology.

Data preparation and transformation

Our study draws upon archival meteorological and energy metering logs sourced from the ECB tower's Building Weather Station (BWS). This corpus spans impactful drivers such as air temperature, moisture saturation, airflow speed, and the energy allocated to ceiling cooling, free cooling, and heating.

The cleaning sequence commences by isolating the numeric fields beyond 'Year,' 'Month,' and 'Hour' as the non-numeric attributes. Absent entries are then populated through a tiered group-mean strategy, with 'Year,' 'Month,' and 'Hour' dictating the pecking order. The grouping successively fills residual gaps, meaning estimates are initially derived via 'Year' and 'Month,' and thereafter via 'Year' and 'Season.' As a last resort, any variable still containing blank slots is replaced with the annual average. Certain attributes are mapped from class designations to numeric codes via label encoding applied to 'Season' and 'Weekday.'

To sharpen model appraisal and legibility, source figures are first re-expressed as efficiency percentages bounded between 0% and 100%. This rescaling normalizes the magnitude of energy-hungry operations, thereby helping the model gauge the relative intensity of energy uptake against typical seasonal and operational benchmarks.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

The assembled dataset is split into training, validation, and test sets to train and validate the model. The chronologically ordered attributes are arranged into sequences to honor the temporal interdependencies woven into the observed energy-use signatures.

Model architecture

A sequential LSTM design was chosen for its proven track record of handling time-series streams that exhibit protracted temporal dependencies. The setup chains multiple LSTM layers, then terminates with a dense projection layer. Each LSTM cell uses the ‘tanh’ activation function, which supports stable gradient propagation and captures the nonlinear temporal interplay across successive energy consumption readings. The final dense module employs ‘ReLU’ activation to yield three output nodes representing the forecasted efficiency scores for ceiling cooling (CC), free cooling (FC), and space heating (SH) draw. Within this LSTM architecture, concerns about excessive capacity and memorization of spurious patterns were mitigated by adding a recurrent layer, thereby strengthening the model's ability to learn temporal structure without destabilizing the training trajectory. Alongside this, cross-validation was used to assess model consistency across alternative data partitions, reinforcing the model's reliability and transferability within the target building dataset. This overall scheme yielded a favorable equilibrium between expressive depth and restraint against overfitting, thereby circumventing the need for dropout-based suppression or early-termination rules.

In parallel to the LSTM, a comparator model rooted in the Gated Recurrent Unit (GRU) was formulated with an equivalent topological structure to enable direct performance comparison. The GRU, while reminiscent of LSTM, adopts a more parsimonious gating regime and was scrutinized to map the trade-off space between computational lightness and predictive fidelity. Both the LSTM and GRU variants were subjected to a structured sweep over hyperparameter landscapes.

For baseline reference, an XGBoost implementation was likewise fielded. Though not natively equipped to model sequential order, it was included due to its competitive record on tabular regression challenges. The XGBoost instance was configured for multi-response output and stands as a reference point to illuminate the relative merits of deep temporally aware architectures. The structural outlines of the LSTM and GRU proposals are summarized in **Table 2.**

Table 2. Proposed model architecture.

Layer Type	Number of Units	Activation Function	Function
LSTM / GRU (first layer)	150	tanh / ReLU	Learns long-range temporal dependencies from sequential data
LSTM / GRU (second layer)	150	tanh / ReLU	Further processes and refines the extracted time-series features
Dense (output layer)	3	ReLU	Produces efficiency level predictions for CC, FC, and SH systems

Training and hyperparameter tuning

All three learners (LSTM, GRU, and XGBoost) underwent a rigorous training and hyperparameter-optimization regimen. For the sequence-aware models, temporally ordered cross-validation (3-fold) was applied through GridSearchCV to

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

promote generalizability across distinct chronological slices. Each of the LSTM and GRU setups was conditioned under four separate forecasting configurations:

1. Daily LSTM/GRU (1-day) – Forecasts the following day's energy efficiency from daily roll-up figures.

2. Daily LSTM/GRU (7-day) – Projects consumption over the upcoming week to facilitate weekly energy orchestration.

3. Hourly LSTM/GRU (1 h) – Delivers high-granularity estimates for prompt operational fine-tuning.

4. Hourly LSTM/GRU (24 h) – Supports next-day programming via hour-by-hour outlooks spanning the forthcoming 24-hour window.

The search grid for the LSTM and GRU pipelines encompassed unit count, activation function class, optimizer algorithm, and learning rate. The top-performing configuration comprised 150 units, ‘relu’ activation, the ‘adam’ optimizer, a learning rate fixed at 0.001, and training over 20 epochs with a batch size of 16. The champion hyperparameter set, which yielded the strongest results, is presented in **Table 3.**

Table 3. Optimal hyperparameters for LSTM/GRU (via GridSearchCV).

Parameter	Best Value
Units	150
Activation Function	relu
Optimizer	adam
Learning Rate	0.001
Epochs	20
Batch Size	16
Cross-Validation	3-fold

For the XGBoost baseline, RandomizedSearchCV was used to navigate the expansive parameter space within the available computational resources **(Table 4)**. The model was assessed using 3-fold cross-validation, and the leading hyperparameters were selected based on the smallest negative mean squared error. The exploration domain took in core knobs such as the number of base learners, learning rate, maximum tree depth, and penalty coefficients.

Table 4. Hyperparameter ranges for XGBoost (via RandomizedSearchCV).

Parameter	Range
n_estimators	[100, 200, 500, 1000]
learning_rate	[0.01, 0.05, 0.1, 0.2]
max_depth	[3, 5, 7, 9]
subsample	[0.6, 0.8, 1.0]
colsample_bytree	[0.6, 0.8, 1.0]

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

gamma	[0, 0.1, 0.2]
reg_alpha	[0, 0.1, 1.0]
reg_lambda	[1.0, 5.0, 10.0]
Cross-Validation	3-fold

Across all architectures, training aimed at driving down the Mean Squared Error (MSE). Ancillary evaluation metrics, namely Mean Absolute Error (MAE) and R^2 (coefficient of determination), were brought in to gauge both predictive sharpness and overall model fit.

Integration of GPT-4 into the recommendation system

GPT-4 was brought aboard as a downstream intelligibility enhancer to illuminate the control heuristics derived from an LSTM-grounded temperature-prediction pipeline. It played no part in the training, forecasting, or scoring phases. Rather, its role was to unpack condensed statistical overviews of HVAC activity across three subsystems:

• Space heating (SH)

• Cooling ceilings (CC)

• Free cooling (FC)

Technical workflow

Folding GPT-4 into the advisory framework followed a staged technical pipeline conceived to deepen the interpretability of LSTM-derived indoor temperature projections. In the prediction stage, the LSTM engine outputs indoor thermal estimates, which serve as the primary substrate for all downstream reasoning. This phase gave way to a statistical post-processing stage, during which engagement cut-ins for each HVAC loop were derived adaptively using percentile-based logic, and intensity envelopes were demarcated using the Interquartile Range (IQR) to characterize nominal operating bands.

Once these quantities were extracted, the data structuring stage sorted the assembled outputs by seasonal block and by occupancy regime (weekday versus weekend), enabling the systematic examination of recurring motifs and deviations in system behavior. In the concluding interpretation stage, this curated dataset was fed into GPT-4, which carried out three principal tasks: rendering natural-language accounts of system dynamics, surfacing wasteful overlaps, such as coincident heating and cooling activation bands, and proposing concrete operational refinements to bolster energy efficiency and thermal comfort. The illustrative prompt template for GPT-4 is shown in **Table 5.**

Table 5. Example prompt structure.

Prompt Section	Objective	Example Variables
Role definition	Specifies GPT-4's function in interpreting post-processed HVAC performance results	"You are assisting in interpreting energy performance…"
Context	Clarifies that all values are derived from an LSTM-based forecasting model	{season}, {threshold_percentiles_heating}
Data	Provides system-specific details such as thresholds, IQR	{heating_iqr_lb_weekend},

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

presentation	ranges, and intensity levels for different days	{cooling_threshold_weekday}
Task specification	Instructs the model to generate concise interpretations, technical insights, and actionable recommendations	“Please analyze this information and provide…”

Example GPT-4 Output (Winter, excerpt)

On weekdays, SH engages once the temperature dips below 18 °C, with intensity reaching 92.79%. This preserves thermal comfort but inflates energy demand. Easing the trigger toward approximately 17 °C may lift efficiency. CC stays off, fittingly for winter conditions. FC cuts in below 16 °C at up to 36.44% intensity; reining in the control envelope may enhance stability.

Graph generation and interpretation

In the outlined approach, graphical examination serves as a linchpin for transmuting LSTM model projections into concrete energy management interventions. Normalized consumption traces for ceiling cooling, free cooling, and heating systems are mapped against BWS air temperature readings across varied seasonal and occupancy scenarios, with dynamically computed decision boundaries and usage intensity contours overlaid. These depictions expose system responses under heterogeneous conditions, paving the way for targeted operational course corrections. GPT-4 is then tasked with making sense of these plots—spotlighting ideal operational corridors, flagging anomalous patterns such as heating and cooling threshold clashes, and formulating context-sensitive counsel. This efficient pipeline distills complex model outputs into transparent, actionable intelligence that supports sound judgment and elevated energy performance.

Role of GPT-4 beyond visual summaries

GPT-4 goes well beyond producing visual synopses by reshaping multidimensional temporal output arrays into organized, situation-grounded advisories. While a domain practitioner might fix attention on the most conspicuous trends apparent in graphical displays, GPT-4 can concurrently weigh and cross-compare multiple variables, dredging up inefficiencies or subtle operational signals that might otherwise escape notice. For example, it can associate elevated ambient temperature readings with control boundary placements to highlight regimes where energy expenditure is excessive even though thermal comfort constraints are formally met.

Example: human expert: “Temperature is around 22 °C; heating is high”

GPT-4: “Even though the comfort envelope is maintained, the indoor temperature persistently sits at the upper edge, driving up heating outlay. A marginally reduced setpoint strategy could pare energy consumption while keeping comfort intact.”

Results and Discussion

The present inquiry assesses the capability of a recommendation platform designed to project efficiency scores for space heating, ceiling cooling, and free-cooling circuits. Our strategy capitalizes on LSTM architectures to estimate energy draw at differing temporal resolutions. In practice, we partition our projections into two principal classes based on aggregation granularity: hourly and daily. Under the hourly umbrella, two configurations are probed: 1 h ahead (conditioned on the preceding 24 h window) and 24 h ahead (fed by a 48 h sequence). Within the daily category, both 1-day and 7-day outlooks are scrutinized, yielding perspectives suited to immediate and week-scale energy programming.

Hourly

1-hour prediction

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

For the 1 h energy draw forecasting task, the merits of three algorithmic candidates—XGBoost (XGB), Gated Recurrent Unit (GRU), and Long Short-Term Memory (LSTM)—were weighed over three distinct circuits: space heating, ceiling cooling, and free cooling **(Table 6)**. XGB consistently outperformed its deep learning counterparts in accuracy, achieving R^2 values of 0.973 for space heating, 0.962 for ceiling cooling, and 0.947 for free cooling, with correspondingly modest RMSE and MAE values. These results indicate that XGB proved notably proficient at modeling near-instantaneous energy-use rhythms, while imposing far lighter burdens on training clock cycles and memory than GRU and LSTM.

Table 6. Evaluation of models in 1 h energy prediction.

Model	Training Time	Memory	Space Heating Consumption (R^2 / RMSE / MAE)	Cooling Ceilings (R^2 / RMSE / MAE)	Free Cooling (R^2 / RMSE / MAE)
XGB	6.25 s	41.0 MiB	0.973/0.029/0.014	0.962/0.032/0.014	0.947/0.027/0.015
GRU	292.9 s	141.3 MiB	0.962/0.186/0.107	0.949/0.052/0.025	0.926/0.019/0.010
LSTM	369.6 s	302.3 MiB	0.969/0.168/0.097	0.948/0.053/0.025	0.924/0.019/0.010

The LSTM and GRU frameworks likewise yielded commendable predictive results, with R^2 values exceeding 0.92 across all circuits. LSTM achieved marginally higher accuracy for space heating than GRU, whereas the two ran neck-and-neck on the cooling circuits. Despite their longer training schedules and greater memory requirements, the recurrent architectures faithfully encoded temporal structure, most notably in the free-cooling case, where the RMSE and MAE metrics reached their minimums. These observations support choosing LSTM and GRU in contexts that privilege sequential learning. At the same time, XGB retains its standing as a stalwart baseline for latency-sensitive applications operating under constrained hardware budgets.

• Training and Validation Performance Comparison (1 h)

Figure 6 charts the training and validation loss trajectories for one-hour-ahead energy draw estimation produced by the LSTM, GRU, and XGBoost learners. The recurrent nets, LSTM and GRU, trace a familiar convergence contour, with both training and validation error diminishing in a steady cadence across epochs. GRU converges modestly ahead of LSTM, though the two ultimately settle at closely matched validation loss plateaus, suggesting disciplined learning free of notable overfitting. The XGBoost entry, depicted via the MSE curve against boosting stages, shows rapid tightening within the first 50 rounds, after which the validation metric stabilizes, suggesting that incremental rounds deliver diminishing returns. On balance, the picture that emerges is that all three formulations are well-equipped to latch onto short-horizon energy fluctuations, with the deep learning alternatives supplying more uniform performance trajectories across the training span.

• Residual Distribution for 1 h Prediction Models

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

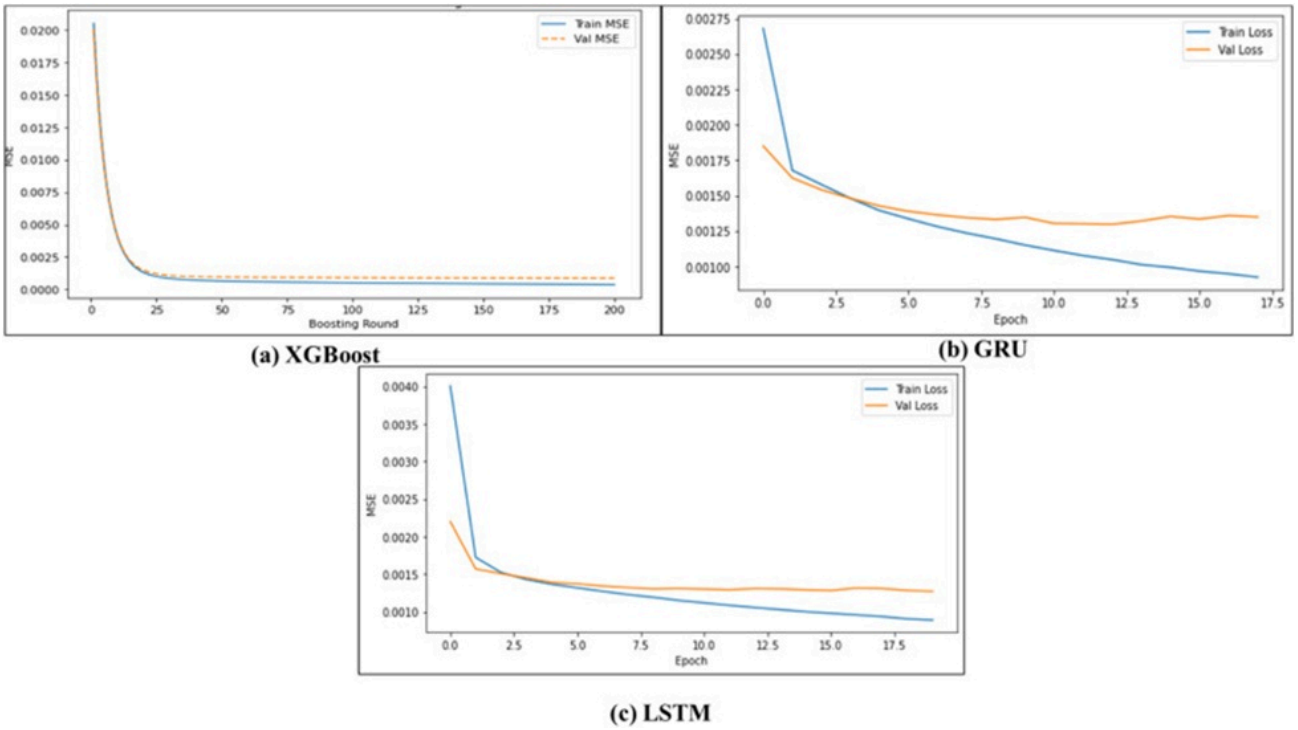

Figure 6. Training and validation performance of models in 1 h.

Figure 7 maps the residual spreads for three categories of energy draw—Space Heating Consumption, Cooling Ceilings, and Free Cooling—across the three model variants. Each cell plots the disparity between measured and estimated values (residuals), with a dashed vertical red line anchoring the zero-error mark. All distributions mass tightly around the zero origin, reflecting minimal systematic skew and broadly dependable estimation. Space Heating Consumption nevertheless exhibits mildly thicker tails, hinting at amplified scatter or sporadic large misses relative to the other two circuits—cooling Ceiling and Free Cooling exhibit tighter residual envelopes, denoting steadier model output.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

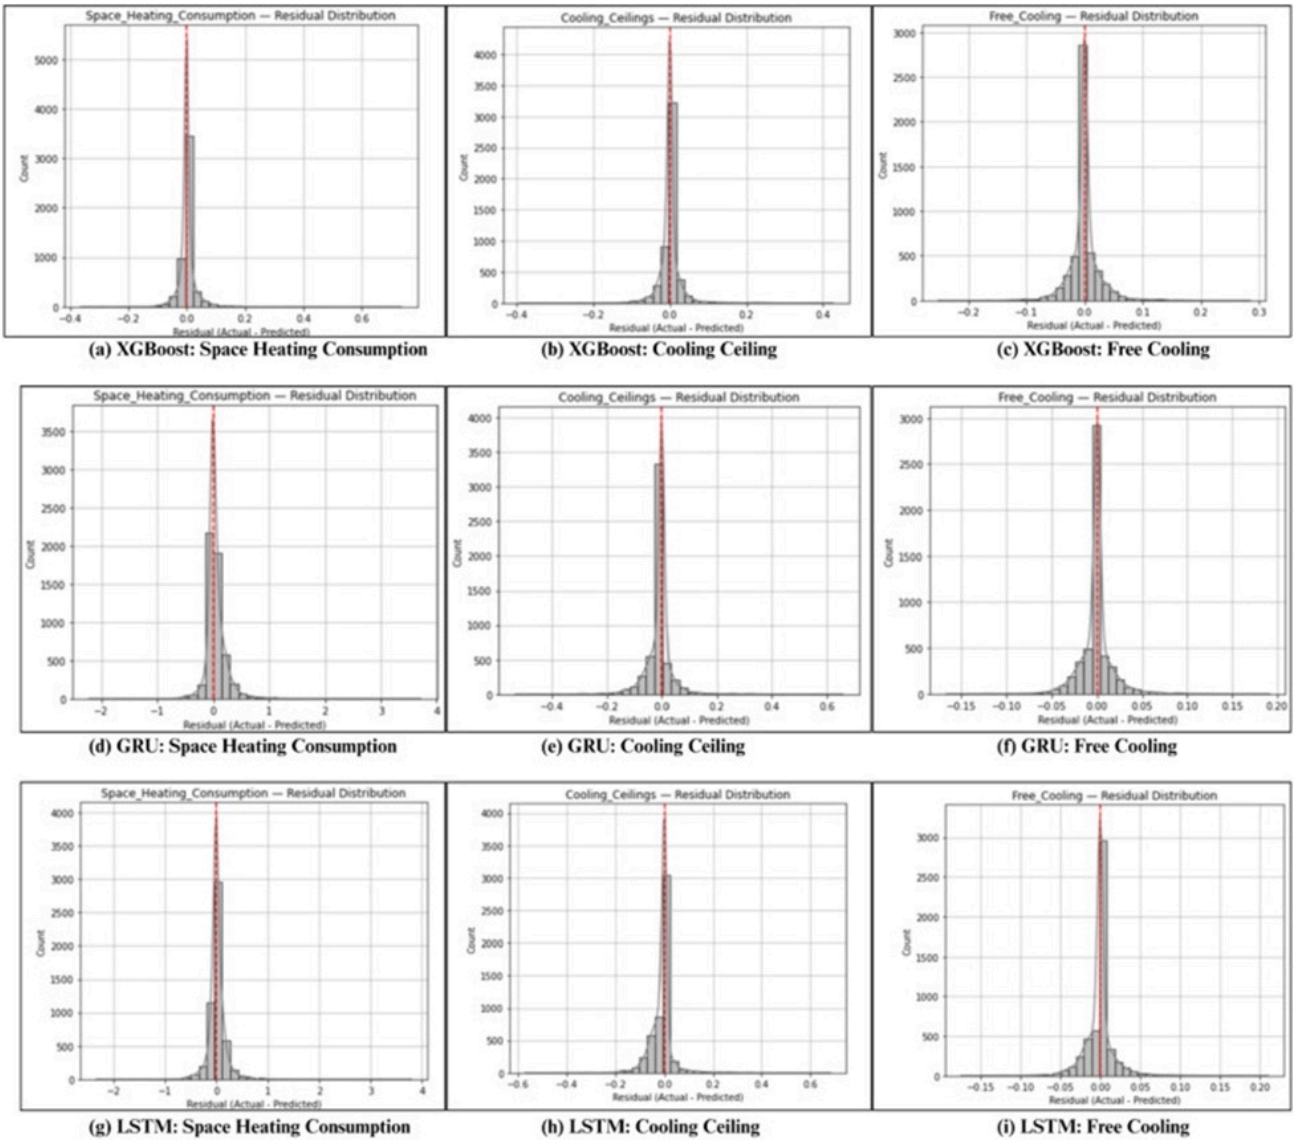

Figure 7. Residual distributions of predicted vs. actual values (1 h).

24-hour prediction

The comparative audit of the candidate frameworks for 24 h forward space heating draw prediction reveals that LSTM reliably surpasses both XGBoost (XGB) and GRU across most hours, with the advantage becoming more pronounced in later hours, when extended temporal linkages carry greater weight **(Table 7).** LSTM captures the highest mean R^2 value of 0.9384 along with the smallest mean RMSE (0.0441) and MAE (0.0295), signaling top-tier forecast fidelity and error suppression over the entire prediction span. GRU yields a mean R^2 of 0.934 and somewhat elevated error indices, whereas XGB—despite its brisk turnaround and lean resource footprint (Train time: 279.89 s, memory usage: 204.8 MiB)—trails with a lower mean R^2 of 0.9067 and larger mean error readings. These patterns highlight the strength of recurrent architectures, particularly LSTM, in distilling the temporal cadence of energy draw across protracted forecast windows.

Table 7. Evaluation of proposed model in 24 h energy prediction (space heating).

	LSTM			GRU			XGB		
	Space Heating			Space Heating			Space Heating		

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

Hour	RMSE	MAE	(R^2)	RMSE	MAE	(R^2)	RMSE	MAE	(R^2)
1	0.0392	0.0263	0.9523	0.0354	0.025	0.961	0.1557	0.0783	0.9739
2	0.0407	0.0264	0.9484	0.0419	0.0277	0.9452	0.2102	0.1122	0.9523
3	0.0405	0.0271	0.9488	0.0415	0.028	0.9461	0.2169	0.1268	0.949
4	0.0422	0.0283	0.9445	0.042	0.0289	0.945	0.2357	0.1406	0.94
5	0.0417	0.0275	0.9461	0.042	0.0287	0.9452	0.2527	0.1511	0.9313
6	0.0428	0.0284	0.9431	0.0425	0.0295	0.9439	0.2698	0.1617	0.9217
7	0.0424	0.0286	0.9436	0.042	0.0286	0.9446	0.276	0.1659	0.917
8	0.0421	0.0283	0.9438	0.0421	0.0286	0.9437	0.2839	0.1741	0.9113
9	0.0414	0.0278	0.9453	0.0429	0.0292	0.9413	0.2888	0.1775	0.9078
10	0.0439	0.029	0.9383	0.0451	0.0306	0.9348	0.2961	0.1813	0.9025
11	0.0447	0.0303	0.9362	0.0446	0.0306	0.9364	0.3029	0.1867	0.8982
12	0.0433	0.0293	0.9409	0.0453	0.031	0.9352	0.3078	0.1913	0.8964
13	0.0426	0.029	0.9427	0.045	0.0303	0.9361	0.3153	0.1946	0.8912
14	0.0437	0.0294	0.9399	0.0456	0.0308	0.9346	0.3162	0.1971	0.8909
15	0.0433	0.029	0.9411	0.0452	0.0308	0.9359	0.3191	0.1976	0.8892
16	0.0438	0.0293	0.9398	0.0463	0.0316	0.9329	0.3194	0.2013	0.8892
17	0.0439	0.0293	0.9395	0.0466	0.0316	0.9318	0.3154	0.1991	0.8919
18	0.0434	0.0293	0.941	0.0471	0.0323	0.9304	0.3182	0.2019	0.8899
19	0.0443	0.0299	0.9379	0.049	0.0336	0.9237	0.3192	0.2044	0.888
20	0.046	0.031	0.9326	0.0492	0.0339	0.9232	0.3263	0.2067	0.8827
21	0.0474	0.0317	0.9292	0.0507	0.0345	0.9189	0.3251	0.2054	0.8845
22	0.0492	0.0331	0.9241	0.0519	0.0351	0.9155	0.3221	0.205	0.8874
23	0.0516	0.0343	0.9159	0.0537	0.036	0.9089	0.3225	0.2056	0.8863
24	0.0544	0.0356	0.9071	0.0561	0.037	0.9011	0.3211	0.2046	0.8876
Mean	**0.0441**	**0.0295**	**0.9384**	**0.0456**	**0.031**	**0.934**	**0.289**	**0.178**	**0.9067**

Turning to the 24 h ceiling cooling energy draw forecast, the LSTM variant achieves the highest aggregate fidelity among the three rivals, with a mean R^2 of 0.9153 and the smallest average RMSE (0.0487) and MAE (0.0254). GRU trails at a short distance with a mean R^2 of 0.8969, outperforming XGB, which records a lower mean R^2 of 0.8806 and heftier error readings. Although XGB initially holds its own in the early hours, its precision decays more steeply in the latter part of the day. LSTM maintains unwavering performance across the entire 24-hour period, more adeptly encoding the temporal signatures of ceiling cooling activity. Even after accounting for the additional computational expense (training time: 364.6 s, memory usage: 136.0 MiB), LSTM represents a more difficult solution for day-ahead cooling projection, a trait especially prized in applications that demand reliability and exactness across extended intervals, as enumerated in **Table 8.**

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

Table 8. Evaluation of proposed model in 24 h energy prediction (cooling ceiling).

	LSTM			GRU			XGB		
	Cooling Ceiling			Cooling Ceiling			Cooling Ceiling		
Hour	RMSE	MAE	(R^2)	RMSE	MAE	(R^2)	RMSE	MAE	(R^2)
1	0.0451	0.0224	0.9267	0.0405	0.0201	0.9407	0.047	0.0206	0.9588
2	0.0464	0.0237	0.9245	0.0455	0.0225	0.9274	0.0602	0.0284	0.9342
3	0.0467	0.0236	0.9246	0.0493	0.0239	0.9161	0.0654	0.0316	0.9235
4	0.0469	0.0237	0.924	0.0501	0.0245	0.9134	0.0689	0.034	0.9152
5	0.0467	0.0246	0.9243	0.0491	0.0251	0.9162	0.0721	0.0374	0.9066
6	0.0493	0.0258	0.9144	0.0503	0.0253	0.9108	0.0745	0.0392	0.8988
7	0.0478	0.0248	0.9187	0.0506	0.0259	0.909	0.0769	0.0409	0.8914
8	0.047	0.0245	0.9222	0.0508	0.0264	0.9091	0.0786	0.0425	0.8873
9	0.0482	0.025	0.918	0.0525	0.0269	0.903	0.0805	0.0442	0.8817
10	0.0473	0.0253	0.921	0.052	0.0277	0.9043	0.0812	0.0453	0.8793
11	0.0477	0.0253	0.9194	0.0531	0.0274	0.9001	0.084	0.0461	0.8707
12	0.0477	0.025	0.9189	0.0543	0.0276	0.895	0.0852	0.047	0.8661
13	0.0489	0.0257	0.9144	0.0553	0.0283	0.8907	0.0852	0.0475	0.8657
14	0.0482	0.0251	0.9167	0.0564	0.0294	0.8861	0.0867	0.0483	0.8609
15	0.0495	0.0262	0.9131	0.0564	0.0292	0.8875	0.087	0.0485	0.8613
16	0.0491	0.025	0.9146	0.0572	0.0299	0.8843	0.0896	0.0494	0.853
17	0.0478	0.025	0.9172	0.0547	0.0288	0.8916	0.0884	0.0488	0.8539
18	0.0487	0.0257	0.9146	0.0571	0.0297	0.8825	0.0861	0.0476	0.8616
19	0.0487	0.0258	0.9135	0.0562	0.0294	0.8848	0.0852	0.0479	0.8629
20	0.0498	0.0261	0.9096	0.0571	0.0295	0.8812	0.0843	0.0479	0.866
21	0.0501	0.0271	0.9085	0.0572	0.0301	0.8808	0.0858	0.0487	0.8611
22	0.052	0.0277	0.9028	0.0586	0.0315	0.8763	0.087	0.0493	0.8591
23	0.0526	0.028	0.8988	0.0591	0.0318	0.8723	0.0861	0.0485	0.8595
24	0.0565	0.0297	0.8859	0.0623	0.0336	0.8617	0.0883	0.0498	0.8561
Mean	**0.0487**	**0.0254**	**0.9153**	**0.0536**	**0.0277**	**0.8969**	**0.0798**	**0.0433**	**0.8806**

Turning to the 24-h free-cooling energy draw forecast, the LSTM formulation yields the smallest average RMSE (0.0339) and MAE (0.019), indicating superior error containment relative to both GRU and XGBoost (XGB) **(Table 9)**. However, LSTM records a diminished mean R^2 of 0.7676, betraying reduced explanatory strength, especially through the day's concluding hours. GRU occupies the middle tier with a mean R^2 of 0.8232, whereas XGB secures the peak mean R^2 of

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

0.846, though at the cost of inflated predictive error. Strikingly, LSTM excels during the opening hour, achieving an R^2 of 0.9426 in Hour 1, yet its potency fades gradually as the forecast horizon extends. Even while boasting the smallest memory appetite (71.8 MiB) and moderate compute time (Training time: 308.5 s), the evidence suggests that although LSTM keeps raw errors in check, its capacity to account for free cooling variability wanes over time, perhaps owing to the more capricious and less rhythmic operational character of this particular circuit.

• Training and Validation Performance Comparison (24 h)

Table 9. Evaluation of proposed model in 24 h energy prediction (free cooling).

	LSTM Free Cooling			GRU Free Cooling			XGB Free Cooling		
Hour	**RMSE**	**MAE**	**(R^2)**	**RMSE**	**MAE**	**(R^2)**	**RMSE**	**MAE**	**(R^2)**
1	0.0169	0.0089	0.9426	0.0368	0.0201	0.9046	0.0421	0.0227	0.8755
2	0.0245	0.0117	0.8788	0.0415	0.0204	0.8789	0.0434	0.0217	0.8676
3	0.0274	0.0138	0.8502	0.0439	0.0222	0.866	0.0436	0.0223	0.8673
4	0.0282	0.0147	0.8422	0.0442	0.0228	0.8656	0.0434	0.022	0.8701
5	0.0301	0.016	0.8218	0.0447	0.0232	0.8627	0.044	0.0221	0.8673
6	0.0314	0.017	0.806	0.0462	0.0243	0.8537	0.0446	0.0231	0.864
7	0.0329	0.0181	0.7864	0.0488	0.0257	0.8365	0.0455	0.0238	0.8577
8	0.0332	0.0185	0.7823	0.0495	0.0269	0.8309	0.0442	0.0233	0.8655
9	0.034	0.0192	0.7716	0.0489	0.0262	0.8351	0.0456	0.0242	0.8565
10	0.0343	0.0196	0.7683	0.0491	0.0262	0.8347	0.0453	0.0237	0.8591
11	0.0343	0.0198	0.7669	0.0495	0.0266	0.831	0.0463	0.0237	0.8521
12	0.0352	0.0202	0.7607	0.0511	0.0273	0.825	0.0473	0.0246	0.8502
13	0.0363	0.0205	0.7447	0.052	0.0278	0.8183	0.047	0.0251	0.8511
14	0.0379	0.0215	0.7256	0.0529	0.0285	0.8136	0.0478	0.0251	0.8478
15	0.0368	0.0213	0.7332	0.0508	0.0284	0.823	0.0461	0.025	0.8537
16	0.0374	0.0214	0.7253	0.0523	0.0283	0.8129	0.0467	0.0246	0.8507
17	0.0385	0.0218	0.7114	0.0538	0.0288	0.8034	0.0481	0.0252	0.8427
18	0.0386	0.0218	0.7065	0.0547	0.0289	0.7946	0.0486	0.0254	0.8384
19	0.0372	0.0216	0.7215	0.0533	0.0289	0.8011	0.0476	0.0256	0.8412
20	0.0373	0.0212	0.723	0.0548	0.0292	0.7918	0.0497	0.0257	0.8292
21	0.0375	0.0214	0.7197	0.0557	0.0297	0.7846	0.0522	0.0276	0.8109
22	0.0375	0.0218	0.7184	0.056	0.0311	0.7806	0.0521	0.0281	0.8105
23	0.0373	0.0218	0.7161	0.0572	0.0318	0.7676	0.0524	0.0288	0.8054
24	0.0387	0.0223	0.6998	0.0612	0.0337	0.7395	0.0575	0.0313	0.77

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

Mean	**0.0339**	**0.019**	**0.7676**	**0.0504**	**0.0269**	**0.8232**	**0.0471**	**0.0248**	**0.846**

The training and validation loss plots for the 24 h forward prediction assignment convey that all three learners—XGBoost, GRU, and LSTM—manifested orderly convergence behavior devoid of overfitting hallmarks **(Figure 8)**. XGBoost, tuned over 160 boosting cycles, displayed a gradual downward glide in both training and validation mean squared error (MSE), with the two contours drawing closely into alignment. GRU and LSTM likewise yielded consistent reductions in MSE across the training epochs, with both achieving broadly comparable validation outcomes. The narrow divergence between training and validation losses throughout indicates robust generalization.

• Residual Distribution for 24 H Prediction Models

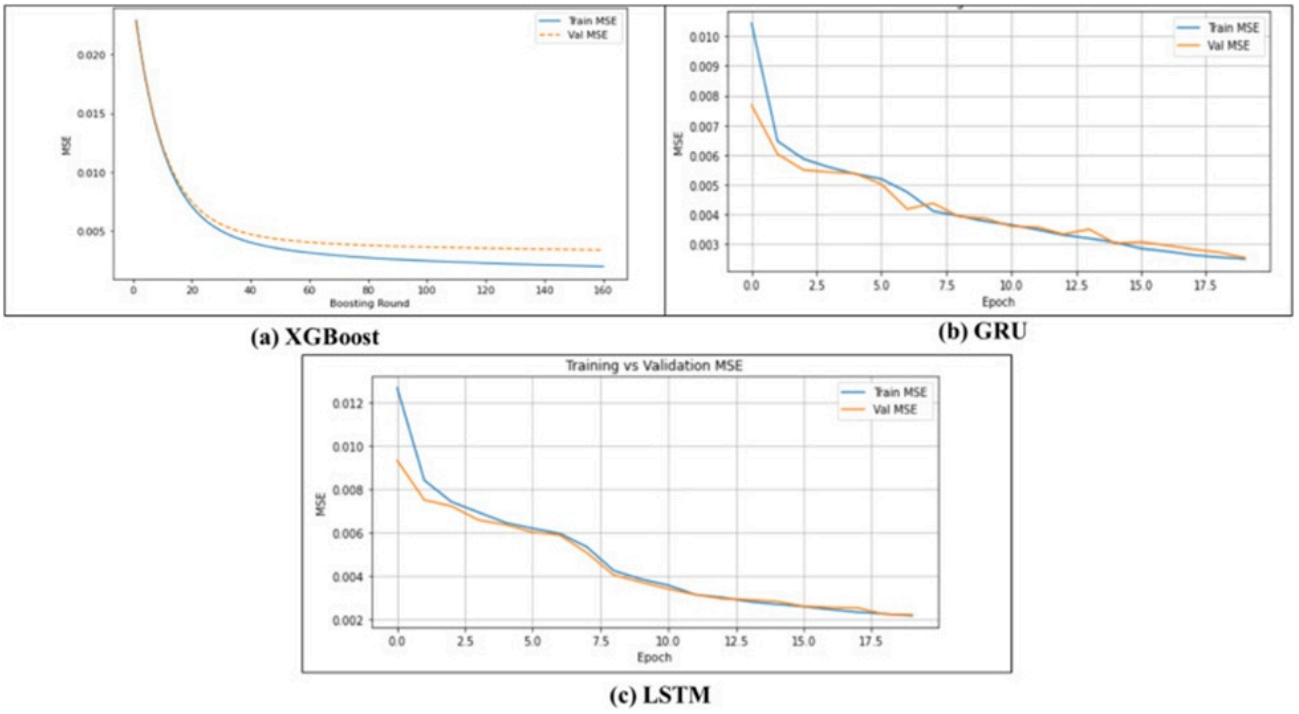

Figure 8. Training and validation performance of models in 24 h.

Figure 9 depicts the residual profiles for 24 h forward estimates of Space Heating Consumption, Cooling Ceilings, and Free Cooling produced by three separate modeling arrangements. Each panel charts the dispersion of forecast miscalculations (observed minus projected), with the dashed red vertical bar anchoring the zero-residual mark. Every distribution exhibits a pronounced peak and symmetrical form centered at zero, underscoring that the models retain considerable precision with negligible systematic distortion across longer projection intervals. The Cooling Ceilings and Free Cooling configurations yield noticeably tighter residual bounds than Space Heating, where occasional elongated tails surface, suggesting greater inherent fluctuation or anomalous spikes in heating load.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

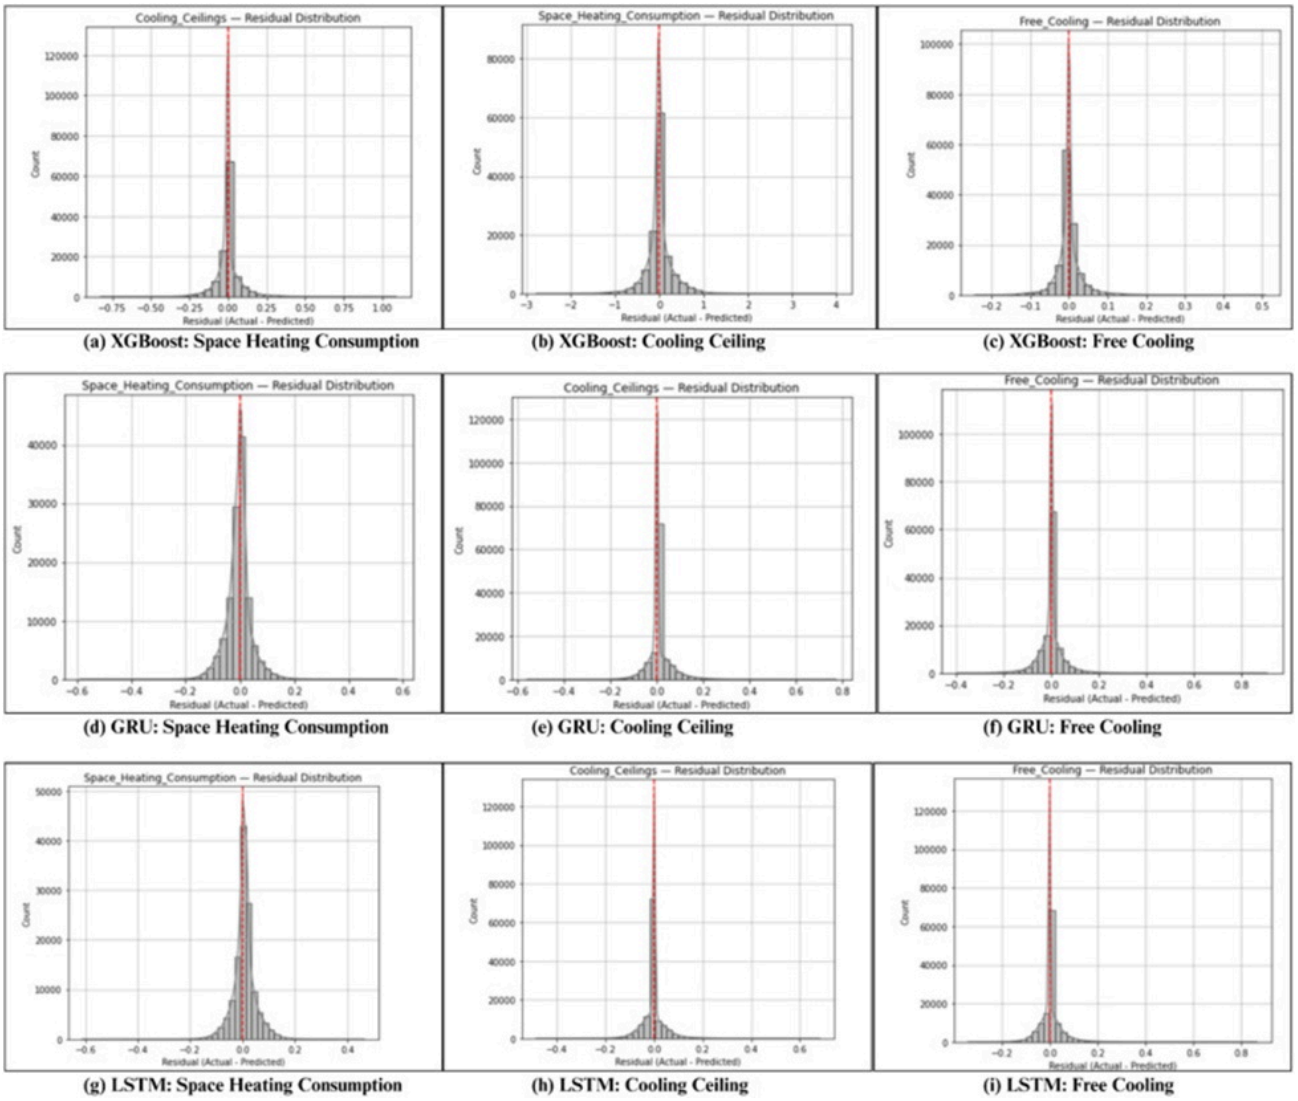

Figure 9. Residual distributions of predicted vs. actual values (24 H).

Daily

1-Day

The one-day forward forecast results highlight the LSTM's commanding lead across all three circuits: space heating, ceiling cooling, and free cooling. LSTM yields the highest R^2 values, especially for space heating (0.969) and cooling ceilings (0.947), indicating strong variance-accounting performance. Its RMSE and MAE figures also register at the floor level for space heating (0.031 and 0.018, respectively) and cooling ceilings (0.039 and 0.018), underscoring its ability to produce highly pinpoint forecasts. On the free-cooling front, LSTM maintains its position with an R^2 of 0.905 and modest error, though it trails GRU by a fraction in explained variance. This illustrates LSTM's ability to untangle the intricate temporal threads woven through building energy sequences over a full diurnal cycle.

The GRU alternative likewise acquits itself well, trailing LSTM by a narrow margin, most evidently for space heating (R^2 = 0.967) and cooling ceilings (R^2 = 0.945). It slightly outperforms LSTM on the free-cooling R^2 metric (0.920 versus 0.905), albeit with somewhat larger error terms. XGBoost, while markedly more economical in training overhead and memory footprint, delivers considerably weaker outcomes across all yardsticks, with R^2 figures slipping below 0.88 for space heating and dipping even lower for the cooling circuits. These patterns suggest that although XGBoost offers a speed

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

advantage, it may lack the expressive breadth needed to encode the temporal sophistication characteristic of building energy draw. In contrast, LSTM, despite steeper training costs, provides the most accurate and uniform forecasts across the three circuits **(Table 10).**

• Training and Validation Performance Comparison (1 day)

Table 10. Evaluation of proposed model in 1-day energy prediction.

Model	Free Cooling Performance (R^2 / RMSE / MAE)	Cooling Ceilings Performance (R^2 / RMSE / MAE)	Space Heating Performance (R^2 / RMSE / MAE)	Memory Usage	Training Duration
XGBoost	0.795 / 0.056 / 0.031	0.846 / 0.066 / 0.036	0.879 / 0.061 / 0.039	3.6 MiB	19.2 s
GRU	0.920 / 0.021 / 0.009	0.945 / 0.055 / 0.022	0.967 / 0.171 / 0.106	200.0 MiB	271.1 s
LSTM	0.905 / 0.038 / 0.017	0.947 / 0.039 / 0.018	0.969 / 0.031 / 0.018	74.1 MiB	294.1 s

The one-day forward forecast evidence indicates that all three learners—XGBoost, GRU, and LSTM—attained satisfactory convergence with minimal overfitting signals **(Figure 10)**. XGBoost traced a gradual and uniform drop in both training and validation mean squared error (MSE) over 550 boosting cycles, with the two trajectories remaining tightly coupled, a hallmark of strong generalization. GRU showed a smooth descent in both training and validation loss, with a modest and stable separation, reflecting controlled variance throughout the learning phase. LSTM followed a similar course, achieving swift loss compression inside the initial epochs before flattening into steady behavior across later iterations. The slim divide between training and validation losses across the board confirms that overfitting is absent and that the models are well-calibrated for short-horizon energy draw forecasting.

• Residual Distribution for 1 Day Prediction Models

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

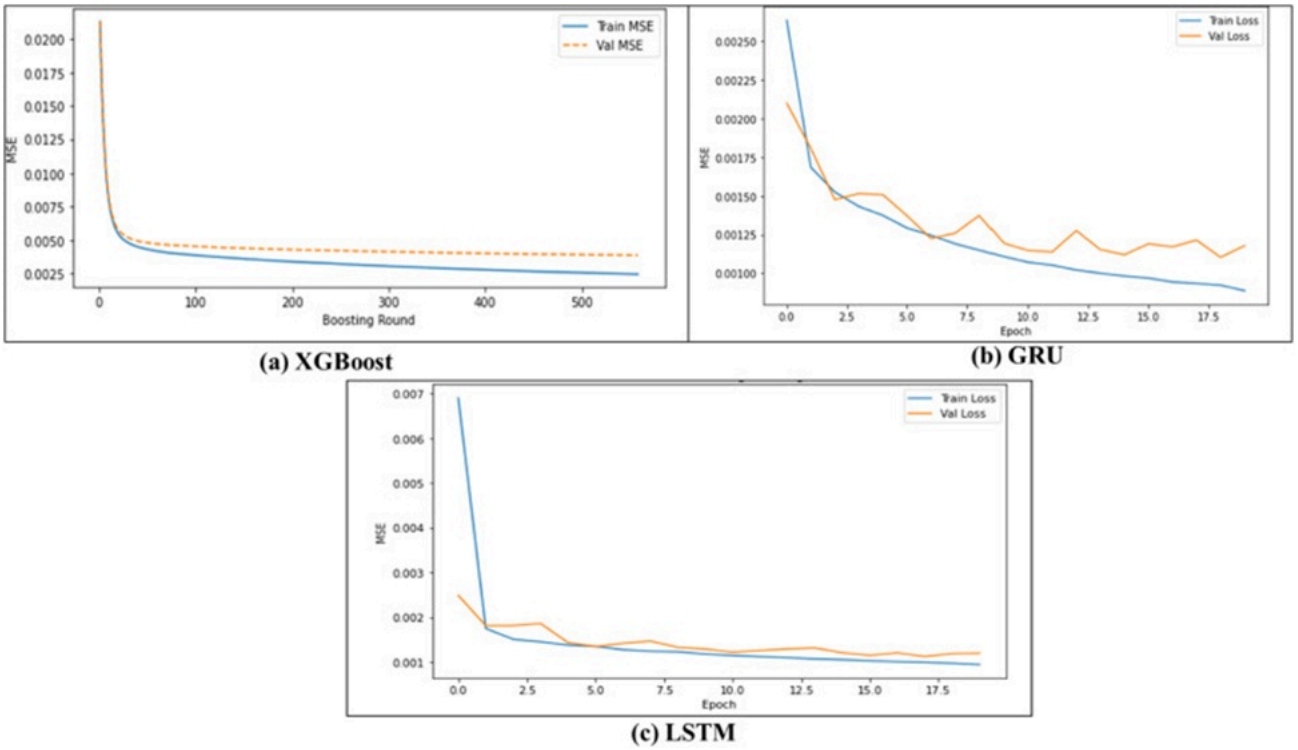

Figure 10. Training and validation performance of models in 1 day.

Figure 11 presents the residual profiles for one-day forward predictions of Space Heating Consumption, Cooling Ceilings, and Free Cooling, arrayed across three different modeling configurations. Each histogram captures the delta between measured and projected values, with a red dashed line denoting zero residual. The distributions are markedly peaked, centered at zero, and broadly symmetric, indicative of minimal predictive distortion. Free Cooling, nonetheless, shows somewhat extended right tails in selected panels, suggesting intermittent underestimation occurrences. Cooling Ceilings and Space Heating, conversely, maintain narrower error bands, although certain panels betray subtle skewness deviations.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

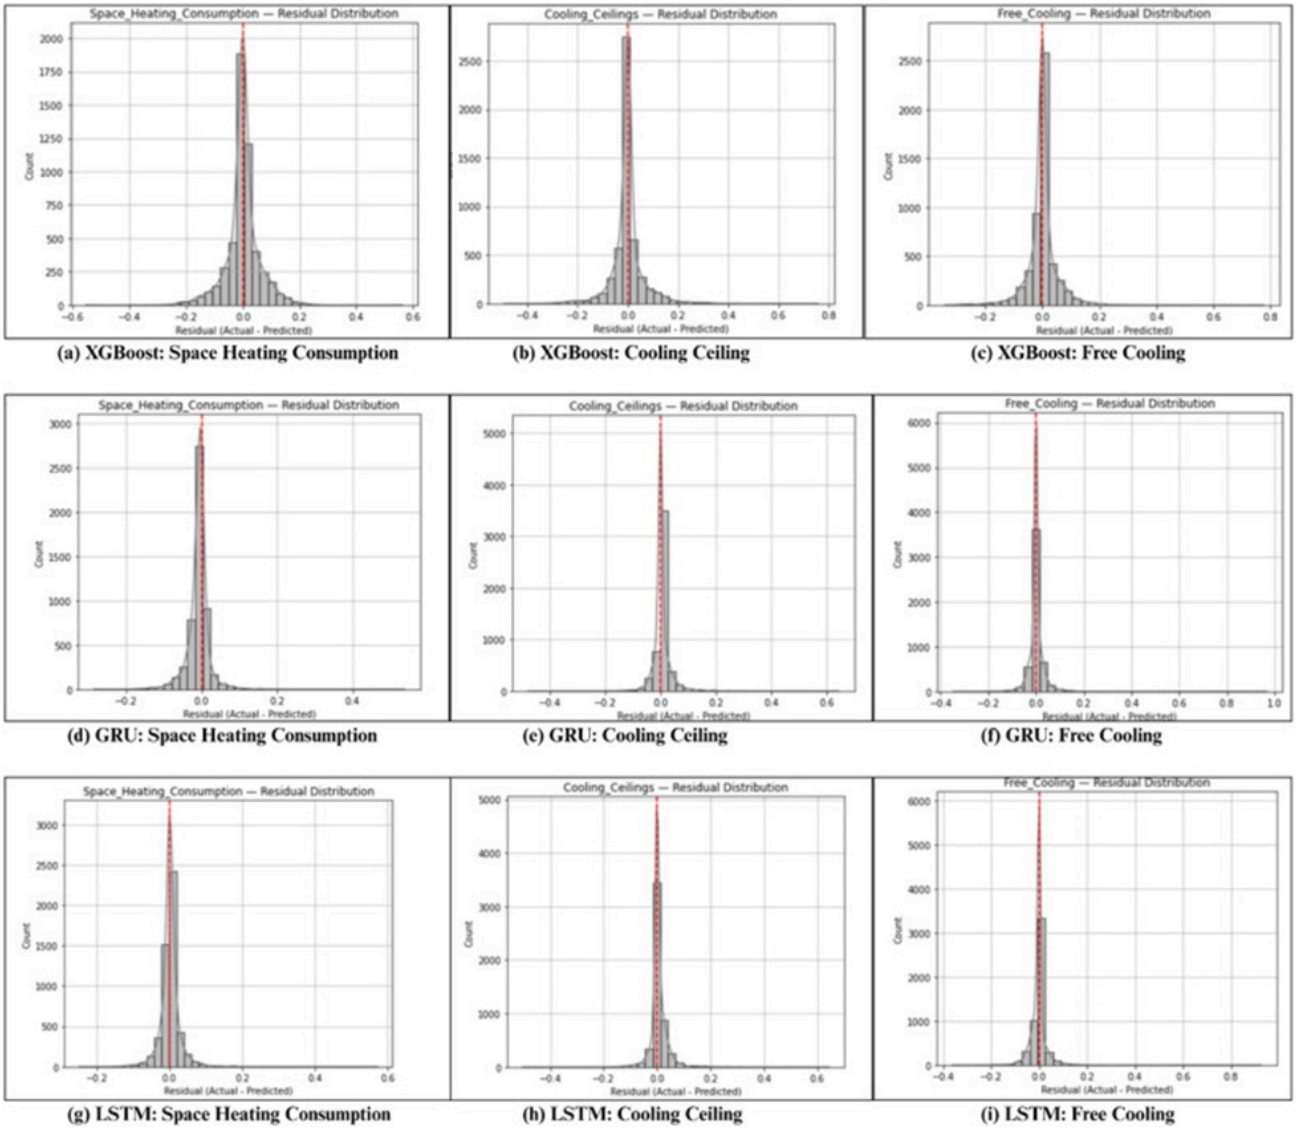

Figure 11. Residual distributions of predicted vs. actual values (1 Day)

7-Day

For the 7-day space heating forecasting task, the LSTM architecture achieves strong performance, with a mean R^2 of 0.9303 and trimmed RMSE (0.0462) and MAE (0.0287), outperforming GRU and XGBoost (XGB) **(Table 11).** While XGB achieves the highest R^2 values on Days 1 and 2 (0.9739 and 0.9523), its fidelity tapers steadily as the week progresses, leading to larger average errors (MAE: 0.1431, RMSE: 0.2307). GRU delivers adequate performance, with a mean R^2 of 0.9108, yet its error indices remain higher than those of LSTM. Importantly, LSTM maintains uniform forecast quality across the full seven-day stretch, suggesting superior generalization and a firmer grasp of long-term temporal couplings. Though LSTM and GRU require longer training times and greater memory demands than XGB, their higher predictive accuracy makes them more suitable for multi-day energy outlooks under volatile building conditions.

Table 11. Evaluation of proposed model in 7 day energy prediction (space heating).

	LSTM Space Heating			**GRU Space Heating**			**XGB Space Heating**		

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

Day	RMSE	MAE	(R^2)	RMSE	MAE	(R^2)	RMSE	MAE	(R^2)
1	0.0318	0.0197	0.9685	0.031	0.0191	0.97	0.1552	0.0819	0.9739
2	0.0393	0.0241	0.9512	0.0436	0.0267	0.94	0.2087	0.1246	0.9523
3	0.0442	0.0272	0.938	0.0513	0.031	0.9166	0.2378	0.1451	0.9379
4	0.0471	0.0289	0.93	0.0562	0.0338	0.9	0.2449	0.1547	0.9342
5	0.0504	0.0313	0.9195	0.058	0.0356	0.8934	0.2509	0.1605	0.9308
6	0.0535	0.0336	0.9084	0.0597	0.0369	0.8859	0.253	0.1646	0.9289
7	0.0569	0.0359	0.8966	0.0638	0.0396	0.8701	0.2644	0.1704	0.9226
Mean	**0.0462**	**0.0287**	**0.9303**	**0.0519**	**0.0318**	**0.9108**	**0.2307**	**0.1431**	**0.9401**

On the 7-day ceiling cooling energy draw forecast, the LSTM variant delivers the strongest overall result, posting the lowest mean RMSE (0.0515), a competitive MAE (0.0258), and the highest average R^2 (0.9036), indicating dependable predictive accuracy and stability **(Table 12)**. GRU follows at close quarters, posting a comparable RMSE (0.0519) and a marginally slimmer MAE (0.0250), alongside a mean R^2 of 0.9021. XGBoost (XGB), despite opening strongly on Day 1 with an R^2 of 0.9522, gradually declines in performance over the week, culminating in a lower mean R^2 of 0.8936 and substantially higher error terms (MAE: 0.0413, RMSE: 0.075). This suggests that recurrent architectures, with LSTM at the fore, are better suited to encoding the temporal regularities embedded in ceiling cooling dynamics over a longer forecast span. Even though training time and memory usage are higher for LSTM (108.3 s, 58.1 MiB) and GRU (92.4 s, 71.5 MiB) than for XGB (9.94 s, 93.1 MiB), the resulting boost in forecast fidelity vindicates their use in predictive control settings.

Table 12. Evaluation of proposed model in 7 day energy prediction (cooling ceiling).

	LSTM Cooling Ceiling			GRU Cooling Ceiling			XGB Cooling Ceiling		
Day	**RMSE**	**MAE**	**(R^2)**	**RMSE**	**MAE**	**(R^2)**	**RMSE**	**MAE**	**(R^2)**
1	0.0389	0.0181	0.9468	0.0376	0.0168	0.9503	0.0512	0.0241	0.9522
2	0.0479	0.0226	0.9197	0.0476	0.0216	0.9208	0.0685	0.0354	0.9149
3	0.0495	0.0246	0.9121	0.0513	0.0238	0.9057	0.0742	0.0405	0.8978
4	0.0504	0.0253	0.9082	0.0515	0.0247	0.9043	0.0755	0.0424	0.8934
5	0.0541	0.0273	0.895	0.0549	0.0273	0.8919	0.0803	0.0458	0.8802
6	0.0587	0.0303	0.8781	0.0598	0.0295	0.8737	0.0874	0.0495	0.8602
7	0.0612	0.0327	0.8655	0.0606	0.0312	0.8684	0.0878	0.0516	0.8568
Mean	**0.0515**	**0.0258**	**0.9036**	**0.0519**	**0.025**	**0.9021**	**0.075**	**0.0413**	**0.8936**

In the 7-day free-cooling energy draw prediction, the LSTM architecture consistently outperforms GRU and XGBoost across combined accuracy and error benchmarks. It yields the highest average R^2 of 0.8364, indicating a broader explanation of variance across the entire forecast window. In addition, LSTM achieves the lowest mean RMSE (0.049) and a competitive MAE (0.0245), highlighting its ability to deliver precise estimates with a narrower spread. GRU trails modestly with an R^2 of 0.8102 and a slightly inflated RMSE (0.0525). In contrast, XGBoost suffers a steep decline in

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

accuracy beyond Day 1, posting the lowest average R^2 (0.7901) and the highest error indices (MAE: 0.0176, RMSE: 0.0324). These patterns confirm that recurrent neural network topologies, above all LSTM, are better suited to capturing the temporal fluctuations and nonlinear interdependencies that steer free cooling demand, particularly over extended durations.

• Training and Validation Performance Comparison (7 days)

Table 13. Evaluation of proposed model in 7 day energy prediction (free cooling).

	LSTM Free Cooling			GRU Free Cooling			XGB Free Cooling		
Day	**RMSE**	**MAE**	**(R^2)**	**RMSE**	**MAE**	**(R^2)**	**RMSE**	**MAE**	**(R^2)**
1	0.0369	0.0198	0.9087	0.0356	0.0186	0.9148	0.0191	0.0105	0.9295
2	0.0451	0.0221	0.8639	0.0456	0.0213	0.8611	0.0276	0.0137	0.853
3	0.0458	0.0224	0.8579	0.0494	0.0232	0.8349	0.0307	0.0162	0.8162
4	0.049	0.0239	0.8415	0.0549	0.0261	0.8008	0.035	0.0187	0.7678
5	0.0514	0.0256	0.8241	0.0579	0.028	0.777	0.036	0.02	0.7515
6	0.0557	0.0281	0.7914	0.0607	0.0297	0.7519	0.038	0.0213	0.7202
7	0.059	0.0295	0.7677	0.0635	0.0312	0.7307	0.0401	0.0229	0.6923
Mean	**0.049**	**0.0245**	**0.8364**	**0.0525**	**0.0255**	**0.8102**	**0.0324**	**0.0176**	**0.7901**

The training and validation plots for the 7-day forecasting assignment convey that all three learners—XGBoost, GRU, and LSTM—succeeded in compressing mean squared error (MSE) throughout their respective learning cycles, signaling effective model calibration and generalization **(Figure 12)**. The GRU and LSTM architectures tracked well-aligned loss curves between training and validation, indicative of steady learning dynamics largely free of overfitting markers. Correspondingly, XGBoost maintained low, closely matched MSE values across both stages, underscoring its resilience. Among the trio, LSTM achieved the lowest validation MSE, highlighting its ability to disentangle the nuanced temporal structures woven through building energy consumption records.

• Residual Distribution for 1 Day Prediction Models

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

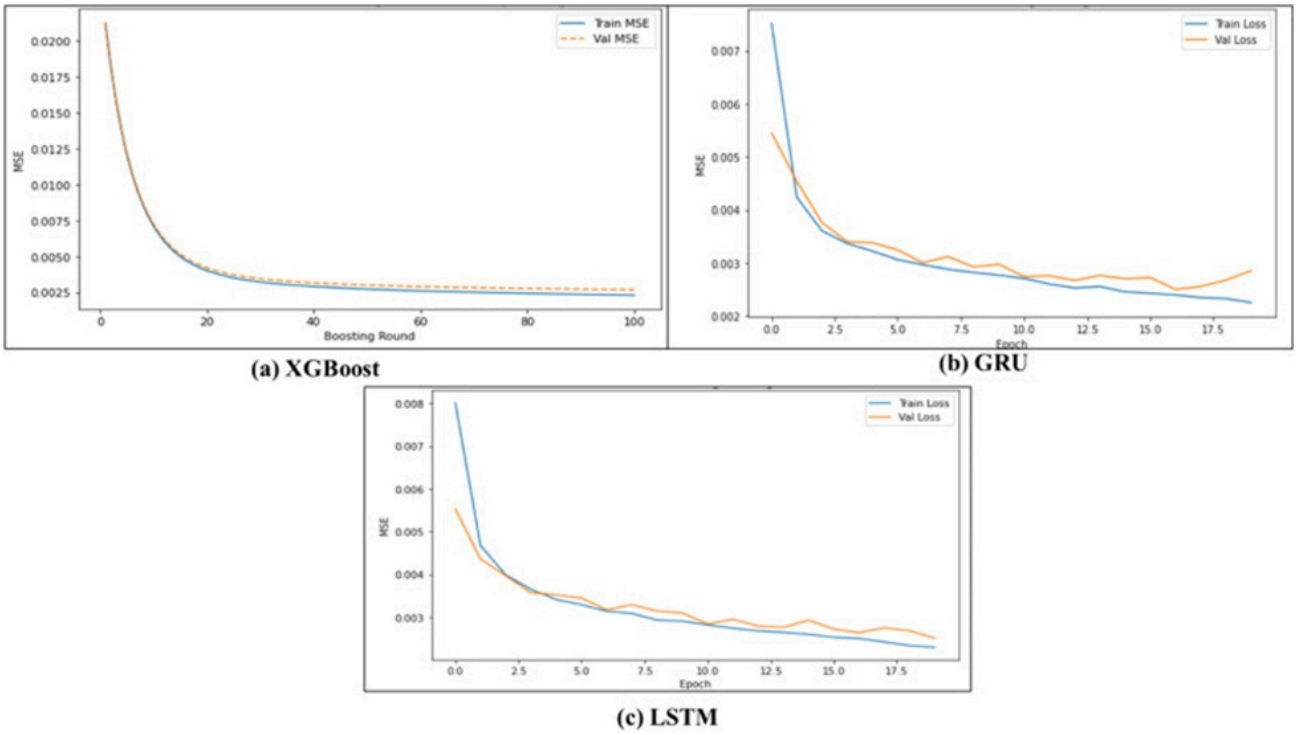

Figure 12. Training and validation performance of models in 7 days.

The residual distribution panels for the 7-day forecasting models—XGBoost, GRU, and LSTM—covering space heating, cooling ceiling, and free cooling circuits, reveal the extent to which each model reconciles the gaps between actual and projected energy draw **(Figure 12)**. All three configurations produce residuals densely packed around the zero mark, attesting to robust predictive sharpness. Notably, the LSTM variant yields the most compact clustering of residuals about the zero line, reflecting its heightened capacity to absorb nonlinear and temporally sequenced patterns in the data while keeping systematic bias to a minimum. The GRU counterpart performs similarly but with a somewhat broader residual envelope, while XGBoost, though still effective, generates slightly heavier tails, indicative of sporadic larger estimation lapses.

Recommendation system

What follows is a granular dissection of the advisory engine. The proposed recommendation platform targets boosting the energy economy and paring operational outlay by regulating space heating and cooling across multiple seasonal and tenancy scenarios. The engine draws on six maneuvers adapted to distinct circumstances, such as proportionately tempering heating output on weekends, singling out abnormal readings through interquartile range (IQR) inspection, and pinning outdoor activation marks to elevated percentile ranks (95th for weekends and 100th for weekdays). The IQR, which brackets the zone between the 25th and 75th percentiles, serves as a tool for surfacing atypical consumption patterns, thereby opening the door to rectifications that avert wasteful operations. Percentiles slice a data corpus into 100 equally sized portions, and by tapping upper percentile values as activation cues, the arrangement is well positioned to handle severe conditions, preserving both energy frugality and occupant well-being. For cooling, the underlying rationale is inverted, with trigger points extracted from the 2.5th and 0th percentiles. In concert, these measures ensure the platform operates at peak efficiency, honing heating and cooling interventions to track demand with precision while reining in energy consumption.

The interpretability layer, driven by OpenAI's Generative Pretrained 4.0 model, adds depth to the advisory pipeline by providing a thorough analysis of graphical renderings of energy usage logs. This module interprets visual outputs to tease apart convoluted relationships and dynamics, yielding actionable pointers that inform modification protocols. For example,

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

the GPT-4 model helps interpret the consequences of outlier flagging via IQR and the weight carried by elevated percentiles in determining activation marks. Such interpretive capacity ensures that the platform's guidance rests on a sophisticated grasp of the data, clearing the way for more exacting refinements to heating and cooling operations.

Hourly space heating consumption

The designated percentile boundaries for space heating draw are tailored to sharpen the platform's sensitivity to thermal fluctuations across varied seasons and weekday categories, as tabulated in **Table 14**. Over weekend intervals, the 97.5th percentile is used, indicating that the platform will summon heating when temperatures dip below a level that exceeds 97.5% of all logged readings for the given season. This tactic ensures heating activation remains measured, sparing energy during typically lighter occupancy periods. In contrast, on weekdays, the 100th percentile is used, implying that heating is only called upon at the very outermost thermal readings, reflecting denser human presence and activity levels.

Table 14. Percentile thresholds in hourly space heating consumption.

Season	Weekday Percentile	Weekend Percentile
Winter	100	97.5
Spring	100	97.5
Summer	100	97.5
Autumn	100	97.5

The winter weekend chart **(Figure 13)** shows that the framework engages space heating once the Building Weather Station (BWS) air temperature falls below 14 °C, with normalized draw oscillating within the 0%–50% band, thereby shielding indoor comfort during colder stretches. The green dashed trace at 55% marks the restrained curtailment of heating intensity on weekends, consistent with reduced occupancy requirements. On weekdays, the heating trigger elevates to 18 °C in response to heavier tenancy and bustle, leading to an escalation of heating draw from 0% to almost 100% as thermal readings fall. The weekday heating cap is 93%, showcasing the platform's knack for efficiently handling steeper heating loads.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

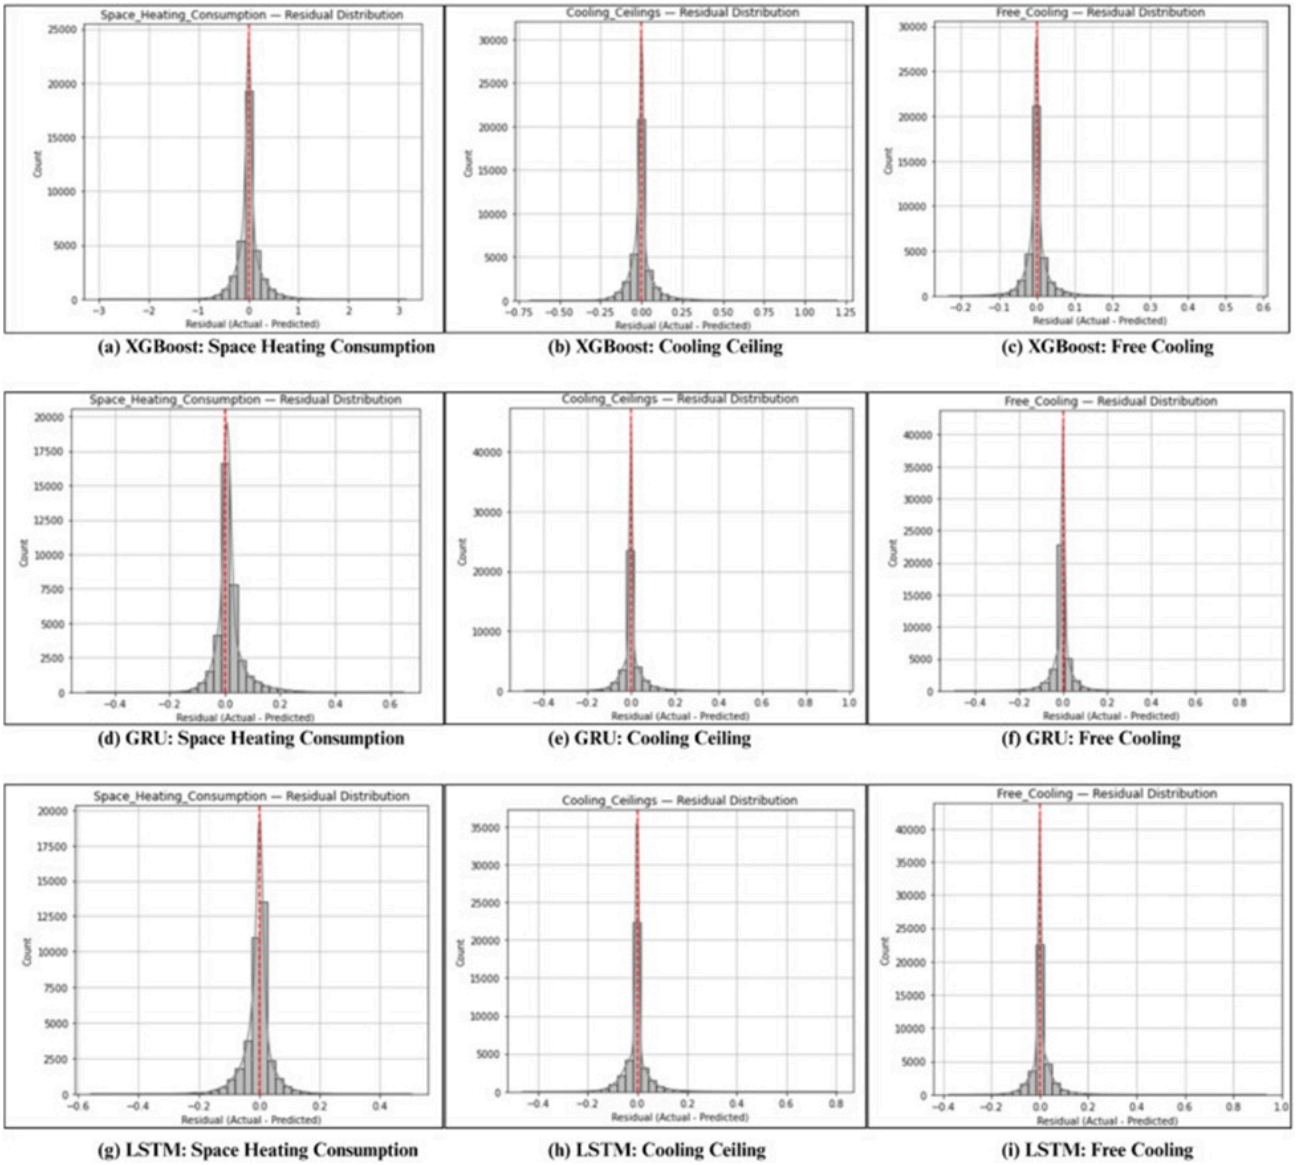

Figure 13. Residual distributions of predicted vs. actual values (7 Days).

The spring charts **(Figure 14)** map normalized space heating draw against BWS air temperature for the spring season, separated into weekend and weekday slices. In both scenarios, a strong inverse association between space-heating draw and air temperature is evident. Over weekends, as shown in the left panel, the trace captures the platform's intensity signature, with heating draw swelling at depressed temperatures and subsiding in a measured fashion, fading to near zero near the 20 °C threshold (blue dashed line). The evidence points to a well-functioning regulatory mechanism, as draw rarely breaches the 30.00% upper limit. In parallel, on weekdays (right panel), the platform's intensity declines as thermal readings rise, dropping to close to 20 °C. A stricter posture is assumed for the weekday case by invoking the 100th percentile (stripping out negative readings). As a result, heating remains enabled even at temperatures up to 28 °C. Still, the intensity remains trivial until the temperature falls below 15 °C, at which juncture it climbs past 20%. The weekday draw readings cling closely to the 61.00% upper limit, signaling disciplined heating stewardship.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

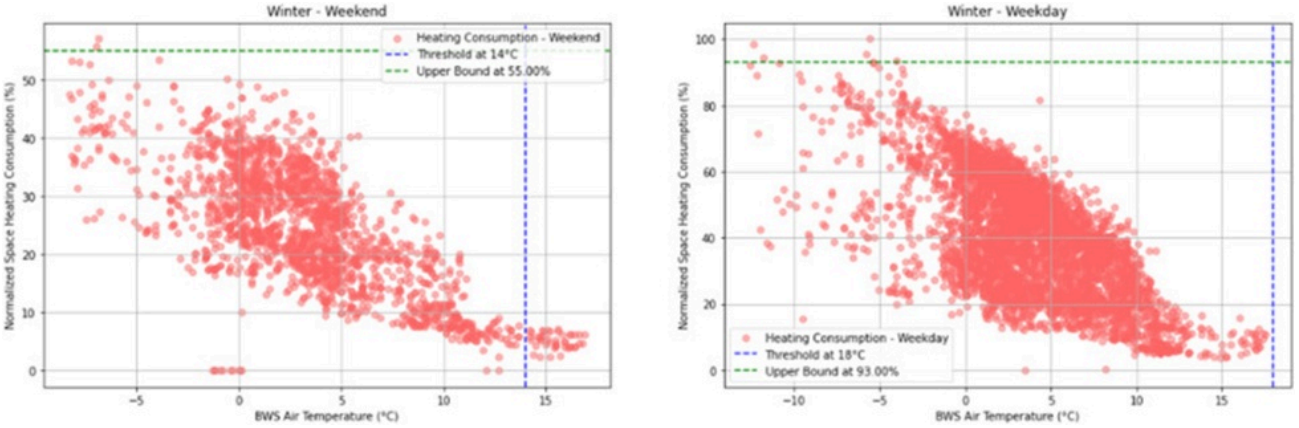

Figure 14. Normalized space heating consumption vs external temperature (winter).

Hourly cooling ceiling consumption

The percentile selections underpinning our inquiry—namely, 2.5% for weekends and 0% for weekdays across the full seasonal spectrum (Winter, Spring, Summer, and Autumn)—play a decisive role in establishing the engagement thresholds for hourly cooling ceiling draw, as enumerated in **Table 15**. These percentile picks were adopted to ensure that the cooling circuit functions optimally amid fluctuating tenancy cycles, exercising tighter restraint on weekdays when building bustle typically peaks, and adopting a slightly more permissive stance on weekends.

Table 15. Percentile thresholds in hourly cooling ceiling consumption.

Season	Weekday Percentile	Weekend Percentile
Winter	0	2.5
Spring	0	2.5
Summer	0	2.5
Autumn	0	2.5

The appraisal of cooling ceiling draw throughout summer brings into focus the capability of the recommended scheme in juggling cooling loads across a broad swath of thermal conditions **(Figure 15)**. On weekends, the platform energizes cooling solely when readings climb beyond 14 °C, dovetailing with energy-conservation imperatives to suppress usage at lower temperatures and streamline energy outlay. Even though the majority of draw readings remain within permissible corridors, a handful exceed the 56% upper limit, indicating spells of heightened load that warrant further optimization. That said, activating air conditioning at an outdoor temperature of 10 °C may be unnecessary, leading to a gratuitous energy drain. Such behavior could be rooted in the system's need to counterbalance after excessive heating, which later triggers cooling, or it might be necessitated by the need to hold specific technical enclosures at lower temperatures. On weekdays, the platform maintains indoor comfort when outdoor temperatures exceed 10 °C, with the draw rising during business hours. We counsel upholding comfort benchmarks on weekdays to serve building occupants, while tuning the system over weekends when the premises are largely vacant, thereby enabling energy savings without sacrificing comfort.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

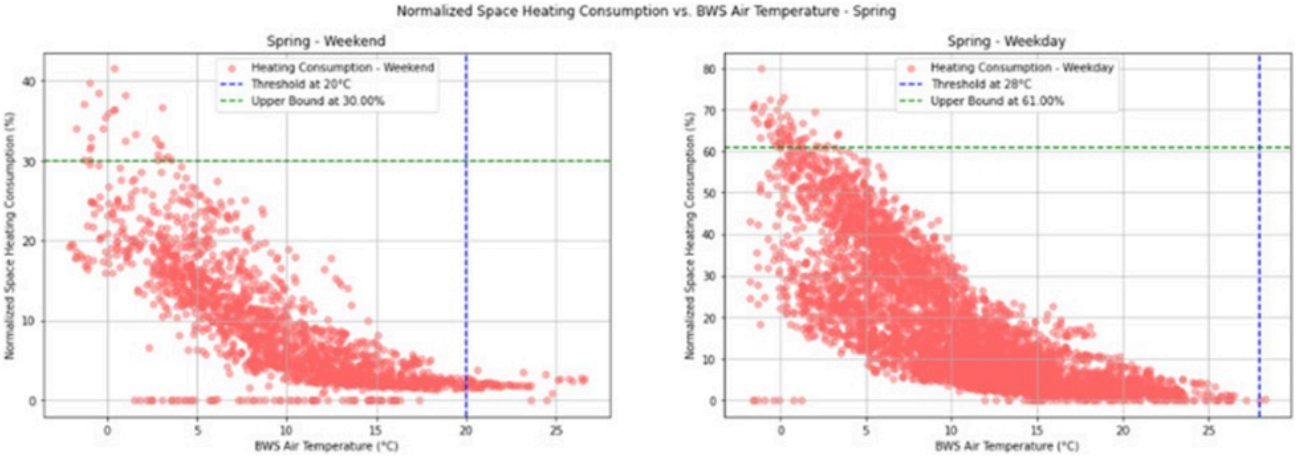

Figure 15. Normalized space heating consumption vs external temperature (spring).

Looking at the spring period, the weekend panel **(Figure 16)** indicates that cooling draw remains near-negligible until the temperature exceeds 9 °C, beyond which a conspicuous surge occurs, particularly across the 15 to 25 °C band. The platform keeps drawing successfully, corralled beneath the 25% upper ceiling, demonstrating capable regulation. Meanwhile, the weekday panel unveils that cooling draw commences its upward march as temperatures clear the 6 °C mark, cresting markedly across 10 to 30 °C. Despite the wider thermal span, the weekday observations hew closely to the 41% upper ceiling, reflecting steady supervision. These patterns suggest that the cooling apparatus is proficient at squeezing energy use at lower temperatures and scaling up smoothly as readings climb, containing draw within the stipulated upper ceilings.

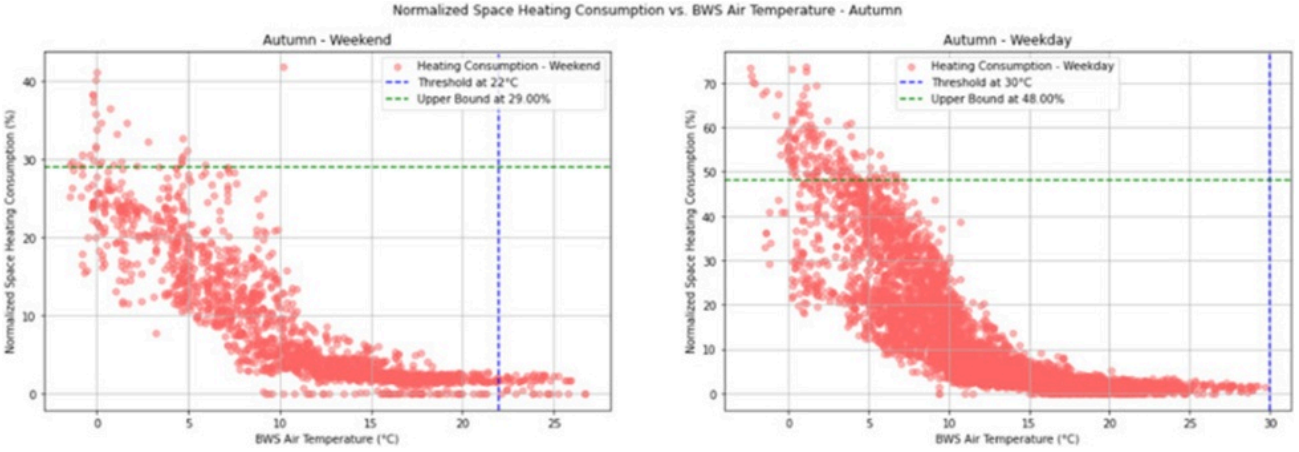

Figure 16. Normalized space heating consumption vs external temperature (autumn).

Regarding autumn, on weekend days, the visualization **(Figure 17)** reveals a well-defined behavior in which normalized ceiling cooling draw persists at comparatively restrained and steady magnitudes so long as the BWS air reading remains below roughly 13 °C. As thermal levels climb above this boundary, a clear escalation in cooling draw becomes apparent, reflected in a wider dispersion of plotted observations. The green dashed line indicates the ceiling at 25% draw, whereas the purple dashed line denotes the thermal cutoff at 12 °C. Weekdays, in a parallel vein, trace a comparable outline but feature steeper rises in cooling draw as temperatures mount. Notably, the scatter of observations is wider on working days than on weekends, suggesting greater swings in cooling load. The ceiling, pegged at 36% draw, lands above the weekend mark, potentially mirroring elevated utilization or load across business hours. The temperature cutoff holds steady at 8 °C.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

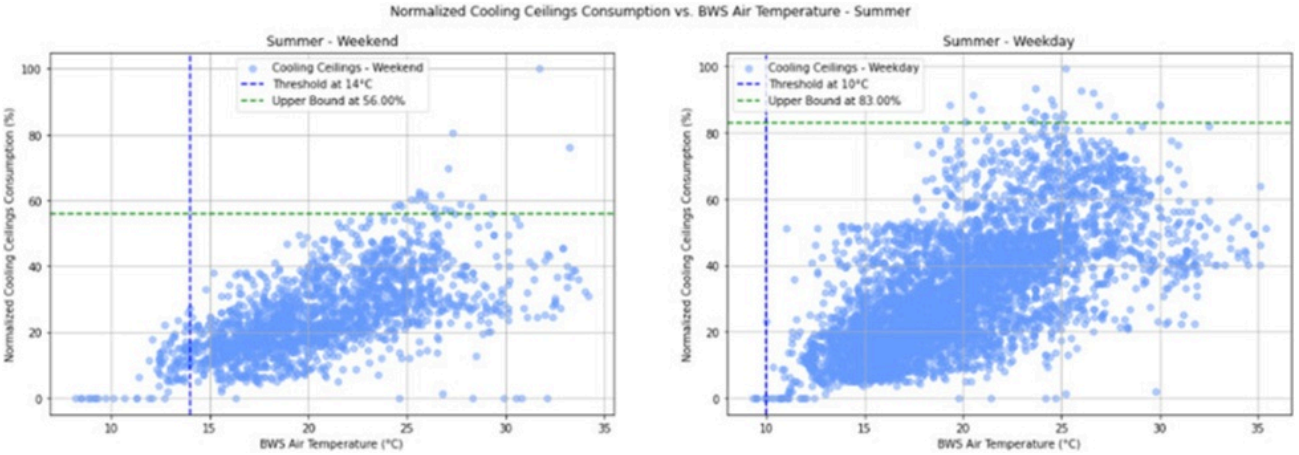

Figure 17. Normalized cooling ceiling vs external temperature (summer).

Hourly free cooling consumption

The percentile selections are fixed at 100% for both weekend and weekday periods across the full seasonal cycle (Winter, Spring, Summer, and Autumn), implying that no distinct cutoff markers or upper constraints are assigned to the hourly free-cooling draw, as recorded in **Table 16.** This posture suggests that the cooling circuit is meant to operate continuously at full capacity, regardless of whether it is a weekend or a weekday, and regardless of seasonal shifts in outdoor conditions. The primary objective is to harvest the greatest possible benefit from free cooling, exploiting its dependency on outdoor thermal states and its extremely low energy consumption profile. To accomplish this, the Interquartile Range (IQR) is expanded to its maximum breadth, affording latitude in operational demarcations. Correspondingly, percentiles are set to 100, ensuring that free cooling runs under all external thermal scenarios.

Table 16. Percentile thresholds in hourly free cooling consumption.

Season	Weekday Percentile	Weekend Percentile
Winter	100	100
Spring	100	100
Summer	100	100
Autumn	100	100

The investigation of free cooling draw during the winter months spotlights the advisory framework's adept handling of energy dispatch in alignment with exterior thermal readings, distinguishing weekend and weekday modes **(Figure 18)**. Across weekends, free cooling kicks in once temperatures drop below 10 °C, holding draw for the most part within the 0% to 40% band, with infrequent outliers yet broadly respecting the 31% upper boundary. This manner of operation signals anticipatory energy thrift during more tranquil stretches, confining activity within a narrow intensity span of 3%–31%. Conversely, on weekdays, free cooling is activated below 16 °C, with draw rates between 2.5% and 36%, and a denser grouping of observations around the 20%–30% range. Despite a wider scatter, the platform effectively keeps energy outlay within the 36% ceiling, ensuring streamlined management during busier operational periods.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

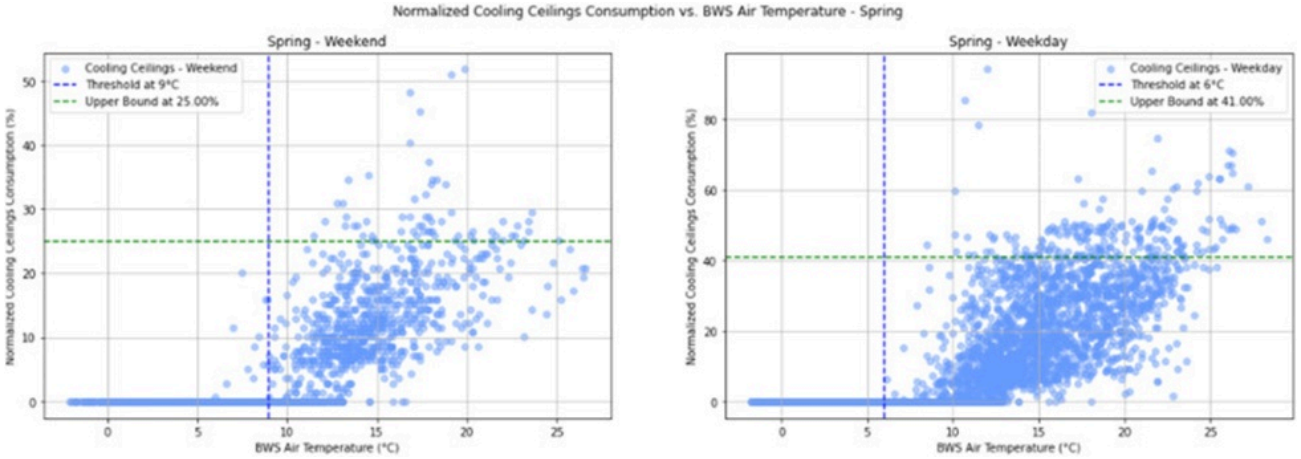

Figure 18. Normalized cooling ceiling vs external temperature (spring).

Throughout the spring period **(Figure 19)**, the weekend free-cooling draw seems to top out at around 17 °C BWS air temperature, reaching a maximum of 49%. Past this turning point, the draw falls away sharply, suggesting a favorable thermal corridor for effective free cooling. Correspondingly, the weekday panel shows free-cooling draw cresting near 19 °C, with more widely scattered observations and a peak value extending to 100%. This indicates greater engagement and wider fluctuations in free cooling on weekdays compared to weekends.

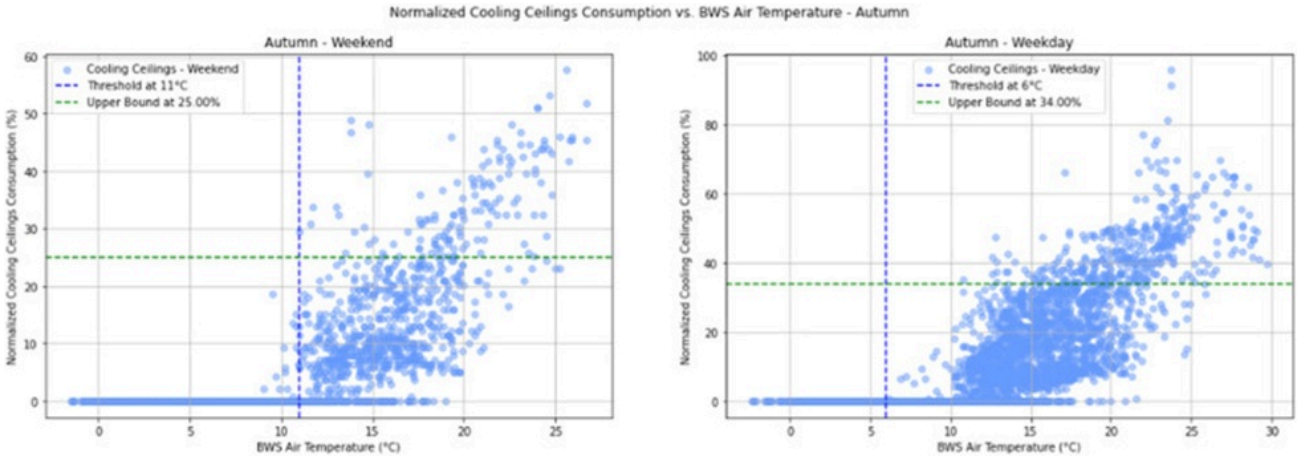

Figure 19. Normalized cooling ceiling vs external temperature (autumn).

Across autumn weekends, **Figure 20** shows that free-cooling draw summits near 15 °C, reaching a ceiling of 54 %. This contour squares with the autumn dataset, exhibiting a drop-off in draw as thermal readings climb past the inflection zone. On weekdays, the cutoff temperature for peak free-cooling draw is 18 °C, echoing autumn weekday patterns, though with a moderated maximum of 80%. The evidence once again spotlights the wider volatility and more intensive use of free cooling on weekdays versus weekends, with spring showing modestly displaced optimal thermal bands compared with autumn **(Figures 21 and 22).**

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

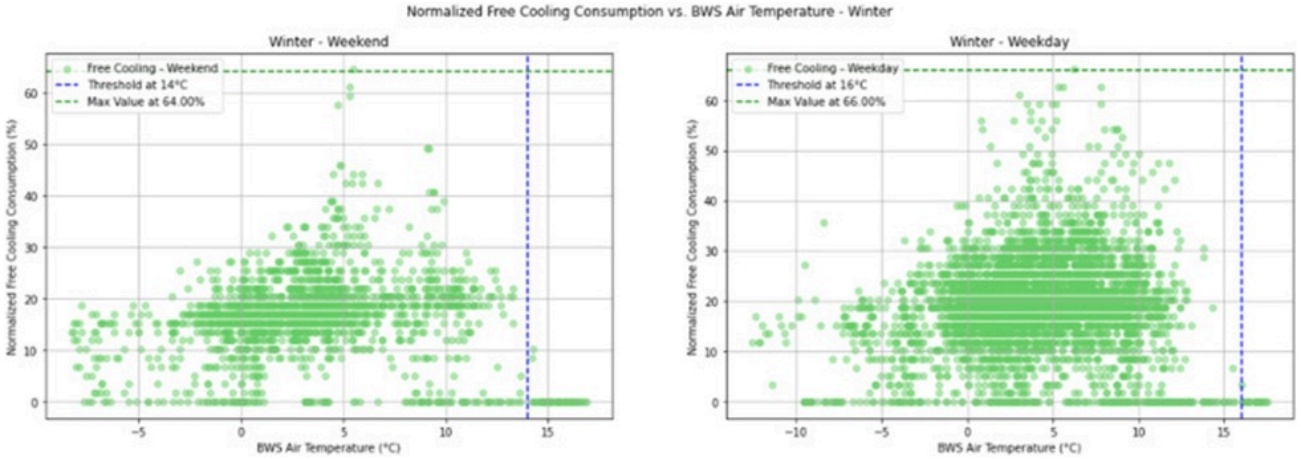

Figure 20. Normalized free cooling vs external temperature (winter).

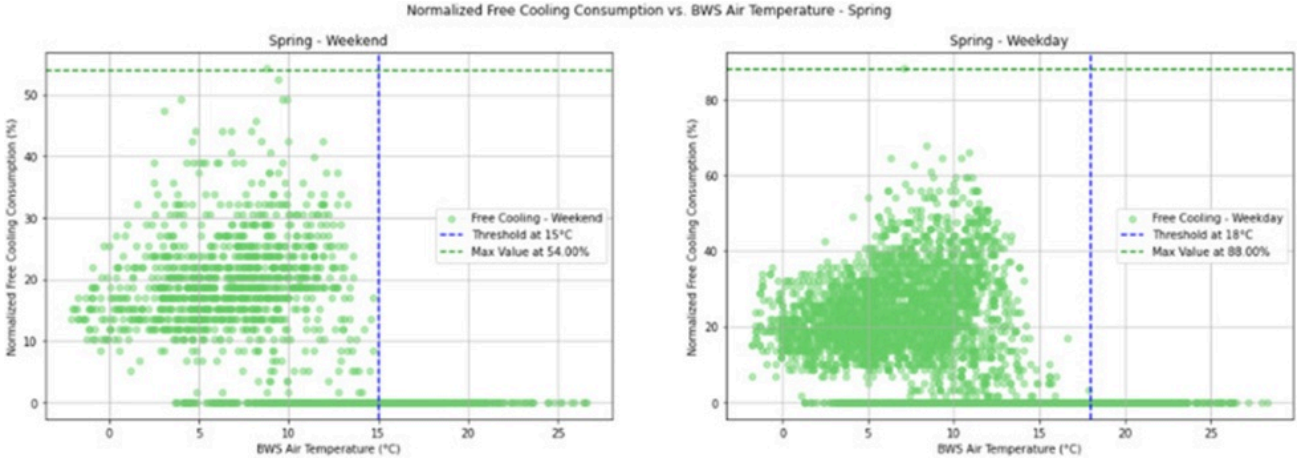

Figure 21. Normalized free cooling vs external temperature (spring).

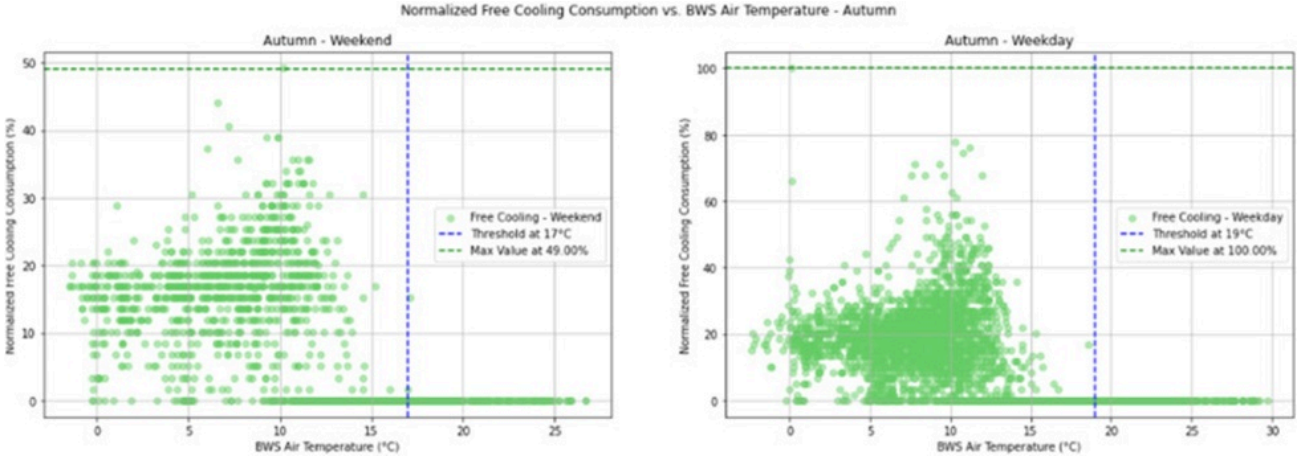

Figure 22. Normalized free cooling vs external temperature (autumn).

Table 17 summarizes how the graphical reading of energy draw patterns informs the advisory guidance for each circuit and seasonal block.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

Table 17. Seasonal energy management recommendations based on graphical interpretation.

Season	System	Suggested Action	Weekend Interpretation	Weekday Interpretation
Winter	Space Heating (SH)	Apply automated and adaptive control strategies	System operates minimally when temperatures fall below 14 °C	Higher usage is observed even when outdoor temperatures rise up to 18 °C, driven by increased occupancy
	Cooling Ceilings (CC)	Kept inactive throughout winter	No operation above 0% intensity is advised	No operation above 0% intensity is advised
	Free Cooling (FC)	Increase utilization during moderate conditions	Runs at full capacity when activated, helping reduce heating demand	Runs at full capacity when activated, helping reduce heating demand
Summer	Space Heating (SH)	Improve energy efficiency strategies	Recommended intensity ranges between 0–4% when temperatures are below 29 °C	Recommended intensity ranges between 0–6.0% when temperatures are below 35 °C
	Cooling Ceilings (CC)	Adjust operation based on temperature conditions	Should remain off when temperatures exceed 14 °C; intensity ranges from 0–56%	Should remain off when temperatures exceed 10 °C; intensity ranges from 0–83%
	Free Cooling (FC)	Maximize deployment of Free Cooling systems	Used to offset heating and cooling demands during mild conditions	Used to offset heating and cooling demands during mild conditions
Spring	Space Heating (SH)	Modify system operation settings	Activated below 20 °C with intensity limited to 0–30%	Activated below 28 °C with intensity limited to 0–61%
	Cooling Ceilings (CC)	Tune operational thresholds	Inactive above 9 °C, with intensity ranging from 0–25%	Inactive above 6 °C, with intensity ranging from 0–41%
	Free Cooling (FC)	Optimize usage of Free Cooling	Intensively utilized across cooling demand range	Intensively utilized across cooling demand range
Autumn	Space Heating (SH)	Optimize heating intensity control	Intensity reduced to a maximum of 29% when temperatures fall below 22 °C	Intensity reduced to a maximum of 48% when temperatures fall below 30 °C
	Cooling Ceilings (CC)	Minimize cooling operations	Inactive when temperatures exceed 11 °C, with intensity between 0–25%	Inactive when temperatures exceed 6 °C, with intensity between 0–34%
	Free Cooling (FC)	Maximize Free Cooling utilization	Set to full intensity whenever required	Set to full intensity whenever required

The close evaluation of the proposed recommendation framework highlights its strong ability to balance heating and cooling loads across seasons and tenancy cycles. Throughout the winter phase, the platform demonstrates a robust

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

capacity to modulate heating activity in response to defined temperature markers and draw ceilings. This responsiveness ensures heating is used frugally on weekends, when occupancy levels customarily dip, while ramping up productively on weekdays to handle heavier loads. Such deliberate governance both trims energy use and boosts operational yield by tightly aligning heating dispatch with actual occupancy demand. The graphical interpretation enabled by GPT-4 has permitted the accurate isolation of these markers and recurring themes, supplying actionable intelligence that sharpens the platform's operating logic.

In parallel, throughout the summer, the advisory engine proves its worth in handling space-cooling loads by enforcing strict upper ceilings and temperature activation points. This method ensures that cooling is invoked only when necessary, slashing energy consumption while safeguarding agreeable indoor environments. By summoning cooling operations only after temperatures breach pre-set thresholds, the platform nimbly harmonizes energy conservation with comfort imperatives. The GPT-4 model's deciphering of seasonal cooling rhythms has shed light on optimal trigger thresholds and intensity bounds, revealing how targeted graphical intelligence translates into leaner operational costs and improved building management performance. This measured approach to deployment not only reduces running costs but also reinforces sustainable building practices, highlighting the platform's role in ensuring optimal energy deployment and raising the overall standard of building administration.

Moreover, the platform's orchestration of free cooling across both winter and summer seasons underscores its adaptive nature when confronted with shifting outdoor conditions. By confining free cooling within prescribed intensity bands and thermal boundaries shaped to weekend and weekday cadences, the platform achieves maximum energy efficiency without undermining comfort. The inclusion of GPT-4 in rendering free-cooling signatures intelligible has magnified its effectiveness at squeezing out energy gains while maintaining comfort, strengthening its place in sustainable energy practices and operational edge across distinct seasons. This stance underscores its potency in advancing ecologically mindful energy practices and sharpening operational throughput across varying seasonal conditions, delivering a material contribution to total energy savings and building output.

While ceiling cooling may not capture the entire cooling burden, it accounts for the dominant fraction of consumption and the most vital slice reserved for comfort conditioning. This highlights the pivotal role ceiling cooling plays in maintaining indoor environmental balance and boosting occupant satisfaction, underscoring the value of refining this component within the wider energy stewardship scheme. The GPT-4 model's analysis of ceiling cooling dynamics has also underscored the need for fine-grained corrections to achieve both peak comfort and operational efficiency.

Applications of the proposed recommendation system

The suggested advisory framework provides flexible solutions to streamline heating draw, ceiling cooling, and free-cooling strategies. Below, we delineate three specific application areas where this novel engine is poised to confer meaningful advantages across a range of sectors and functional arenas.

i. Building Management Companies

Property oversight firms stand to gain from this advisory engine, which they can use to manage their real estate assets more efficiently. Embedding the platform within the building's automation backbone enables the orchestration of heating, ceiling cooling, and free cooling in lockstep with prevailing atmospheric conditions and space-use densities [32]. This anticipatory methodology ensures occupants receive suitable conditioned air while simultaneously preventing energy waste during slack periods. Such adaptability further equips the platform to deliver steady performance throughout the full calendar year, slashing operating costs and carbon emissions. Beyond this, the rich diagnostic feeds generated by the engine give property supervisors the means to forecast energy requirements in advance and implement targeted upgrades that support environmentally sound facility operations and high tenant satisfaction.

ii. Facility Managers in Commercial Buildings

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

The advisory engine focused on heating draw, ceiling cooling, and free cooling offers notable upside for those responsible for facility oversight in commercial premises. It enables the unification of HVAC governance and the modulation of energy usage profiles in reaction to meteorological swings and space-use fractions. During periods of steep heating or cooling requirements, the engine's prognostic capacity helps forecast climatic trends, enabling parameter adjustments that reduce energy consumption while still meeting the specified thermal band [24]. Governance of this stripe not only reduces utility charges but also extends the working life of HVAC hardware by shielding it from undue stress under extreme thermal loads. What is more, the live-monitoring and logging dashboards enable facility overseers to scrutinize output metrics, spot emerging faults, and implement remedial interventions that elevate energy stewardship to a higher plane of both efficacy and operational rhythm.

iii. Energy Consultants and Sustainability Experts

Energy advisory practitioners and sustainability specialists can also call upon the advisory engine to provide tailored counsel to customers seeking to raise building energy grades and align with ecological benchmarks. Armed with deep energy appraisals and the engine's embedded predictive faculties, advisory professionals can craft bespoke regimens for heating, cooling, and free cooling that are attuned to each building's unique physical make-up and usage profile [33]. This method smooths the path to adopting energy-mindful customs that comply with existing legal codes and sector norms, thereby reducing operating costs and greenhouse gas emissions. In addition, the perpetual scrutiny and iterative fine-tuning afforded by the engine supply consultants, with the evidential basis to confirm and document energy gains and to keep advancing recalibrations, thereby lock in lasting beneficial shifts both for their clients' fiscal health and for the broader ecosystem.

Limitations

Notwithstanding the robust predictive performance of the proposed LSTM construct for the building under study, translating its performance to dissimilar building classes remains a challenge. Divergences in thermal behaviour, driven by differences in envelope design, HVAC layouts, and tenancy cycles, can substantially reshape energy-use signatures. Learners conditioned on a lone structure may prove unable to accommodate these idiosyncratic operational profiles once ported to new settings. On top of this, variety in regulatory logic constitutes a further brake, as differences in control schemas, setpoint philosophies, and operating doctrines can undercut the portability of learnt representations. A further pinch point is data scarcity: a large share of buildings lack the finely temporally resolved archival records indispensable for effective LSTM conditioning, thereby curbing the practicability of fielding such models in data-lean environments.

Conclusion

The present investigation provides a robust methodological framework for refining energy use in built spaces by constructing an LSTM neural network. The engine competently projects efficiency bands for space heating, ceiling cooling, and free-cooling loops at both hourly and daily granularities, enabling accurate, responsive energy governance. Principal outcomes show that the LSTM formulation achieves high predictive accuracy, with R-squared values indicating its capacity to account for up to 97% of the variation in energy draw across a range of contexts. These patterns demonstrate the engine's resilience and ability to forecast energy loads, thereby buttressing streamlined operational programming. To assess the relative merit of the LSTM architecture, it is systematically compared with two reference algorithms—Gated Recurrent Unit (GRU) and XGBoost—across several temporal resolutions. The results tell a consistent story: LSTM outperforms GRU and XGBoost in forecast precision, particularly for extended-horizon outlooks, underscoring its suitability for building energy management applications.

What is more, incorporating live meteorological telemetry from a dedicated measurement station perched at the ECB tower heightens the engine's practical performance and accuracy. By mining archival drawings and atmospheric readings, the platform yields actionable insights for property overseers, enabling well-informed decisions to refine HVAC performance. The decoding of graphical material via GPT-4 has added depth to the examination, offering nuanced

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

perspectives on energy-drawing motifs and system behavior. This capability enables delicate corrections and a richer appreciation of how diverse factors affect energy efficiency. The proposed advisory platform distinguishes itself by adapting to shifting seasonal and occupancy conditions, as evidenced by its deliberate management of heating and cooling loads. In winter conditions, the engine adapts heating activity to defined thermal boundaries, streamlining energy use across both weekend and weekday periods. Similarly, during summer, the engine prudently manages cooling functions, ensuring energy-sparing operation without eroding indoor comfort. The administration of free cooling further highlights the engine's pliable nature. By shepherding free cooling within stipulated intensity envelopes and thermal constraints tailored to weekend and weekday cadences, the engine extracts the maximum energy efficiency from operations without compromising comfort.

Acknowledgements

None

Conflict of interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Fernando Almeida reports article publishing charges was provided by NOVA University Lisbon NOVA Information Management School. Fernando Almeida reports a relationship with European Central Bank that includes: employment. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Financial support

This work was supported by national funds through FCT (Fundação para a Ciência e a Tecnologia), under the project - UIDB/04152 - Centro de Investigação em Gestão de Informação (MagIC)/NOVA IMS. This work was funded by the European Union throught the project 101084013 - DIGITAL4Business. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HADEA). Neither the European Union nor the granting authority can be held responsible for them.

Ethics statement

None

Received: 17 Jan 2022 Revised: 22 Mar 2022 Accepted: 08 Apr 2022
Published online: 15 November 2022

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

1. IEA. The latest trends in energy and emissions in 2018. Paris, France: International Energy Agency; 2019.

2. Chandel R, Chandel SS, Prasad D, Dwivedi RP. Prospects of sustainable photovoltaic powered thermoelectric cooling in zero energy buildings: a review. Int J Energy Res. 2022;46(14):19319-40.
https://doi.org/10.1002/er.8508.

3. Chandel R, Chandel SS, Prasad D, Dwivedi RP. Sustainable passive solar and photovoltaic integrated technology interventions for climate responsive net zero energy buildings in western Himalayan mountainous terrain of India. Next Sustain. 2024;3:100039.
https://doi.org/10.1016/j.nxsust.2024.100039.

4. Newell RG, Raimi D, Aldana G. Global energy outlook 2019: the next generation of energy. Resour Futures. 2019;:46.

5. Zhao HX, Magoulès F. A review on the prediction of building energy consumption. Renew Sustain Energy Rev. 2012;16(6):3586-92.
https://doi.org/10.1016/j.rser.2012.02.049.

6. Swift J, Lawrence TM. The ASHRAE GreenGuide: One means of establishing a link between sustainable design practitioners. In: Proceedings of Clima 2007 WellBeing Indoors. Helsinki, Finland: REHVA; 2007. p. 10-4.

7. Kim D, Yoon Y, Lee J, Mago PJ, Lee K, Cho H. Design and implementation of smart buildings: a review of current research trend. Energies. 2022;15(12).
https://doi.org/10.3390/en15124278.

8. Kim D, Cox SJ, Cho H, Yoon J. Comparative investigation on building energy performance of double skin façade (DSF) with interior or exterior slat blinds. J Build Eng. 2018;20:411-23.
https://doi.org/10.1016/j.jobe.2018.08.012.

9. Kim D, Cox SJ, Cho H, Im P. Evaluation of energy savings potential of variable refrigerant flow (VRF) from variable air volume (VAV) in the U.S. climate locations. Energy Rep. 2017;3:85-93.
https://doi.org/10.1016/j.egyr.2017.05.002.

10. Norouziasl S, Jafari A, Zhu Y. Modeling and simulation of energy-related human-building interaction: a systematic review. J Build Eng. 2021;44:102928.
https://doi.org/10.1016/j.jobe.2021.102928.

11. Chen Y, Guo M, Chen Z, Chen Z, Ji Y. Physical energy and data-driven models in building energy prediction: a review. Energy Rep. 2022;8:2656-71.
https://doi.org/10.1016/j.egyr.2022.01.162.

12. Nuthakki S, Kulkarni CS, Kathiriya S, Nuthakki Y. Artificial intelligence applications in natural gas industry: a literature review. Int J Eng Adv Technol. 2024;13(3):64-70.
https://doi.org/10.35940/ijeat.c4383.13030224.

13. Nuthakki S. Conversational ai and LLM ' S current and future impacts in improving and scaling health services. Int J Comput Eng Technol. 2023;14:149-55.

14. Nuthakk S, Buddiga SKP, Koganti S. Exploring deep learning models for image recognition: a comparative review. Signal Image Process Int J. 2024;15(3):1-11.
https://doi.org/10.5121/sipij.2024.15301.

15. Chen Y, Ye Y, Liu J, Zhang L, Li W, Mohtaram S. Machine learning approach to predict building thermal load considering feature variable dimensions: an office building case study. Buildings. 2023;13(2).

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

https://doi.org/10.3390/buildings13020312.

16. Dong F, Yu J, Quan W, Xiang Y, Li X, Sun F. Short-term building cooling load prediction model based on DwdAdam-ILSTM algorithm: a case study of a commercial building. Energy Build. 2022;272:112337. https://doi.org/10.1016/j.enbuild.2022.112337.

17. Aryal A, Becerik-Gerber B. Energy consequences of comfort-driven temperature setpoints in office buildings. Energy Build. 2018;177:33-46. https://doi.org/10.1016/j.enbuild.2018.08.013.

18. Kim J, Tzempelikos A, Braun JE. Review of modelling approaches for passive ceiling cooling systems. J Build Perform Simul. 2015;8(3):145-72. https://doi.org/10.1080/19401493.2014.899394.

19. Kamali S. Review of free cooling system using phase change material for building. Energy Build. 2014;80:131-6.

20. Tajjour S, Singh Chandel S. A comprehensive review on sustainable energy management systems for optimal operation of future-generation of solar microgrids. Sustain Energy Technol Assess. 2023;58:103377. https://doi.org/10.1016/j.seta.2023.103377.

21. Moradzadeh A, Mohammadi-Ivatloo B, Abapour M, Anvari-Moghaddam A, Roy SS. Heating and cooling loads forecasting for residential buildings based on hybrid machine learning applications: a comprehensive review and comparative analysis. IEEE Access. 2022;10:2196-215. https://doi.org/10.1109/ACCESS.2021.3136091.

22. Sardianos C, Varlamis I, Bouras G, Dimitrakopoulos G, Zalonis A, Koutsouri A, et al. Real-time personalised energy saving recommendations. In: Proceedings of the 2020 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData). 2020. p. 366-71. https://doi.org/10.1109/iThings-GreenCom-CPSCom-SmartData-Cybermatics50389.2020.00072.

23. Almeida R, Georgieva P, Martins N. Heating setpoint recommendation strategy for thermal comfort and energy consumption optimization. Energy Build. 2023;296:113406. https://doi.org/10.1016/j.enbuild.2023.113406.

24. Wei P, Xia S, Chen R, Qian J, Li C, Jiang X. A deep-reinforcement-learning-based recommender system for occupant-driven energy optimization in commercial buildings. IEEE Internet Things J. 2020;7(7):6402-13. https://doi.org/10.1109/JIOT.2020.2974848.

25. Lissa P, Deane C, Schukat M, Seri F, Keane M, Barrett E. Deep reinforcement learning for home energy management system control. Energy AI. 2021;3:100043. https://doi.org/10.1016/j.egyai.2020.100043.

26. Shuvo SS, Yilmaz Y. Home energy recommendation system (HERS): a deep reinforcement learning method based on residents' feedback and activity. IEEE Trans Smart Grid. 2022;13(4):2812-21. https://doi.org/10.1109/TSG.2022.3158814.

27. Azuatalam D, Lee WL, de Nijs F, Liebman A. Reinforcement learning for whole-building HVAC control and demand response. Energy AI. 2020;2:100020. https://doi.org/10.1016/j.egyai.2020.100020.

28. Pruvost H, Calleja-Rodríguez G, Enge-Rosenblatt O, Jiménez-Redondo N, Peralta-Escalante JJ. A recommendation system for energy saving and user engagement in existing buildings. Proc Inst Civ Eng Smart Infrastruct Constr. 2022;176(1):1-11. https://doi.org/10.1680/jsmic.21.00023.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press

29. Wei P, Xia S, Jiang X. Energy saving recommendations and user location modeling in commercial buildings. In: Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization (UMAP 2018). 2018. p. 3-11. https://doi.org/10.1145/3209219.3209244.

30. Varlamis I, Sardianos C, Tsakos K, Petrelli M, Kikis V, Tsakou G, et al. Using big data and federated learning for generating energy efficiency recommendations. Int J Data Sci Anal. 2023;16(3):353-69. https://doi.org/10.1007/s41060-022-00331-2.

31. Kim D, Lee Y, Chin K, Mago PJ, Cho H, Zhang J. Implementation of a long short-term memory transfer learning (LSTM-TL)-based data-driven model for building energy demand forecasting. Sustainability. 2023;15(3). https://doi.org/10.3390/su15032340.

32. Moradzadeh A, Mansour-Saatloo A, Mohammadi-Ivatloo B, Anvari-Moghaddam A. Performance evaluation of two machine learning techniques in heating and cooling loads forecasting of residential buildings. Appl Sci. 2020;10(11). https://doi.org/10.3390/app10113829.

33. Fang H, Tan H, Dai N, Liu Z, Kosonen R. Hourly building energy consumption prediction using a training sample selection method based on key feature search. Sustainability. 2023;15(9). https://doi.org/10.3390/su15097458.

© 2022 The Author(s). | Published by Institute of Smart Energy Systems Press