[1] Qasim Mustafa Zainel,

[2] Parviz Rashidi-Khazaee,

[1] Leila Sharifi

Residential Building Heating Load Prediction using Deep Learning TabNet Model

Journal of Electrical Systems

Abstract: - Efficient heating load (HL) predictions in residential buildings are vital in energy optimization and cost savings. Deep learning models with a high ability to solve complex tasks could be good tools for this purpose. Therefore, in this study, we proposed a new deep learning model for heating load prediction based on an attentive interpretable tabular learning model (TabNet). The in-hand dataset contains only 768 records which causes deep learning models' weak performance in comparison with classical machine learning tools. To solve the problem and utilize the ability of deep learning models, a new hybrid model that combines TabNet and a tabular data augmentation module based on GAN and CGAN methods has been proposed. The data augmentation module increased the size of the training dataset 5 times. The performance results indicated that the BiLSTM outperforms other well-known deep learning models without data augmentation, including ResNet, Fully Connected Neural Networks (FCNN), TabNet, And LSTM. By utilizing the data augmentation module the TabNet-GAN model outperformed other deep learning models and brought comparable results with other classical machine learning HL prediction models. Therefore, the TabNet-GAN model could be used for residential building HL prediction and help engineers select the best plan/design from the energy usage perspective.

Keywords: Heating Load Prediction, Fully Connected Neural Networks, Tabular Data Augmentation Tool, Attentive Interpretable Tabular Learning Model, LSTM, BiLSTM, ResNet.

INTRODUCTION

Deep learning methods bring the opportunity to solve complex engineering tasks like residential building heating load estimation at the early stage of design. Reducing energy consumption in residential buildings has become a global priority, as these structures account for 39% of worldwide energy production [1]. Recent progress in machine learning (ML), and deep learning (DL) have introduced powerful methods for predicting energy consumption, including heating requirements, based on building structural and design characteristics [2, 3]. Deep learning has proven highly effective in managing complex, high-dimensional data, offering enhanced predictive accuracy compared to conventional machine learning techniques [4]. These approaches can model intricate relationships in tabular datasets, often characterized by mixed data types and imbalanced distributions. Accurate predictions from such tools enable optimal energy utilization, reduced waste, and improved residential comfort.

Heating load prediction has been extensively studied, leveraging various machine learning and statistical approaches. Traditional methods, such as regression models have been employed to predict energy consumption based on physical and environmental parameters [5]. Kalogirou et al. utilized artificial neural networks (ANNs) to predict energy demands in buildings, demonstrating superior performance compared to regression techniques [6]. More recently, ensemble methods like Random Forests and eXtreme Gradient Boosting Machines have been applied with promising results in capturing non-linear relationships among features [5, 7, 8]. Chou and Bui successfully proposed a hybrid model combining a Support Vector Regressor and ANN (SVR+ANN) for heating load prediction. Al-Rakhami et al. used the optimized state-of-the-art eXtreme Gradient Boosting Method (XGB) using grid search hyperparameter tuning method and successfully used it for HL prediction [7]. Zhou et al. optimized the Multi-Layer

[1] Computer Engineering Department, Urmia University, Urmia, Iran

[2] Information Technology and Computer Engineering Department, Urmia University of Technology, Urmia, Iran.

*Corresponding Author: p.rashidi@uut.ac.ir

Perceptron Network with Particle Swarm Optimiztion (MLP-PSO) [9], Afzal et al. optimized MLP with PSO and Gray Wolf Optimiztion Techniques [10], Salami et al. Optimized XGB with the Bayesian Method, Alawi compared XGB model performance with other models [11], and Sadaghat et al proposed the XACM method [12] for HL prediction and estimation. Qasim et al. successfully leveraged Large Language Models and proposed a linguistic model for HL estimation [13]. With great progress in ML tools development, however, these models often require extensive feature engineering and cannot dynamically adapt to new data.

Deep learning models, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), have been investigated for their potential in energy prediction tasks. Ahmad et al. utilized RNNs to forecast hourly energy consumption, highlighting the advantages of sequence modeling in capturing temporal dependencies [14]. Despite their success, these models often struggle with interpretability and computational complexity, which are crucial for practical deployment. Fan et al. developed long short-term memory (LSTM) for building yearly energy usage based on time series data [15]. Gao et al. highlight the potential of combining LSTM networks with self-attention mechanisms to develop interpretable and accurate models for building energy consumption prediction [16]. These models are designed to process high-dimensional data effectively while providing insights into decision-making, addressing the common issue of interpretability in deep learning applications. Olu-Ajayi et al. analyzed the performance of various machine learning (ML) and deep learning (DL) models, such as Long Short-Term Memory (LSTM) networks and Bidirectional LSTM (BLSTM) networks, for forecasting energy consumption in buildings [4]. Their findings reveal that Bidirectional LSTM (BLSTM) models surpass LSTM models in predicting energy consumption across diverse datasets, highlighting the advantages of bidirectional weight updating in improving model performance.

None of the above deep learning methods proposed for energy prediction was utilized for selecting the best design/plan of the building at the early design stage. The main reason behind its weakness and unavailability of deep learning methods in building energy prediction based on building design structure is the lack of available hand data [5]. The in-hand dataset only contains 768 records that are insufficient for deep learning methods training.

To fill the gap and develop successful deep learning models for selecting the best plan/design from an energy usage perspective, this study introduces a new hybrid HL prediction model. The model integrates data augmentation techniques with deep learning approaches to enhance accuracy and improve predictive performance.

To expand dataset size and improve re-training performance, data augmentation techniques are employed to generate new but contextually relevant data [17]. The available dataset are tabular and therefore the tabular data augmentation could be used for this purpose. Table-GAN employs the vanilla GAN (VGAN) approach to generate synthetic data [18]. The Conditional Tabular GAN (CTGAN) [19] and Synthetic Minority Over-sampling Technique (SMOTE) [20] have been developed and bring surprising results in imbalanced datasets and datasets with categorical features. Additionally, GANBLR modifies traditional GAN architectures by incorporating Bayesian networks into both the generator and discriminator [21]. The Tabular Variational Autoencoder (TVAE), an adaptation of the VAE for tabular data, has demonstrated significant performance improvements in classification tasks [19]. TimeGAN tackles data scarcity issues and enhances the accuracy of heating load prediction models [22]. Conditional Variational Autoencoders (CVAE) have also been successfully used to generate synthetic data, offering valuable inputs for constructing 24-hour energy forecast models [15, 23].

Because the fact that in-hand dataset is tabular, therefore, it is preferred to choose deep learning models that process tabular data directly. The newly developed TabNet [24], has emerged as a powerful solution tailored for tabular data, leveraging attention mechanisms to balance performance and interpretability. Unlike traditional deep learning models, TabNet directly processes tabular data without extensive preprocessing, preserving feature semantics. Its performance will be compared with other well-known deep learning methods including Bi-LSTM [4], LSTM [4, 15], FCNN [25] and Residual Network (ResNet) [26] and also be compared with state-of-the-art eXtreme Gradient boosting (XGB) [7] and random forest [5] models.

This paper investigates how integrating TabNet with advanced data augmentation methods, specifically, GAN and CGAN, can further enhance its predictive capabilities. GAN and CGAN are utilized to generate synthetic data, effectively mitigating overfitting and enhancing model generalization. This approach tackles the challenges associated with limited datasets in the domain.

The remainder of this work is structured as follows: Section 2 reviews related works. Section 3 details the methodology, including the proposed model architecture, dataset, background information, and the design of the new synthetic data generation tools. Section 4 presents the implementation results, comparing the model's performance with similar studies. Finally, Section 5 concludes the paper.

Related Work

This section examines the advanced machine learning and deep learning techniques previously employed for energy prediction.

The random forest (RF) algorithm is an ensemble learning technique commonly applied to classification and regression tasks [5]. It builds multiple decision trees during training and aggregates their predictions to enhance accuracy and robustness. RF is known for handling high-dimensional data, managing missing values, and providing feature importance rankings, making it a versatile and powerful machine learning tool. Tsanas and Xifara employ an RF model for the heating load of residential buildings based on eight design parameters, including glazing area, relative compactness, and roof area. Using a dataset of 768 simulated building configurations [5]. The RF outperformed linear regression, achieving mean absolute errors of 0.51 for HL, demonstrating their superior predictive accuracy. The study showcases the technical efficacy of random forests in handling complex relationships in building energy data. It offers a reliable framework for precise energy load prediction in residential building design.

The eXGBoost, an advanced ensemble learning algorithm that combines multiple weak learners to form a robust predictive model, is known for its regularization techniques, which help prevent overfitting and enhance generalization performance [7]. The eXGBoost model outperformed several state-of-the-art models across all evaluation metrics, including RFs, Ensemble Models, Genetic Programming, Neural Networks, and Gaussian Processes. This work and [5] used the same dataset that was considered in our work.

ResNet is a deep learning architecture, that addresses the challenges of training very deep neural networks, particularly the vanishing and exploding gradient problems, which make it difficult for deep networks to learn effectively [27]. Choi et al. integrated ResNet and LSTM to capture both spatial and temporal dependencies in load data. ResNet efficiently extracts hierarchical features using residual learning, while LSTM captures long-term temporal relationships, improving the model's accuracy in predicting future loads [26]. The integrated ResNet-LSTM model outperforms traditional forecasting methods, demonstrating superior accuracy in short-term load predictions. This enhancement is credited to the model's ability to effectively capture complex patterns within the load data [26].

LSTM (Long Short-Term Memory) is a specialized type of recurrent neural network (RNN) designed to handle sequential data effectively. It overcomes the vanishing gradient problem by employing memory cells and gates (forget, input, and output) to regulate the flow of information [28]. LSTM effectively captures long-term dependencies, making it well-suited for applications such as time series forecasting, natural language processing, and speech recognition. Bi-LSTM (Bidirectional LSTM) [4] extends LSTM by incorporating two layers that process the input sequence in both forward and backward directions, allowing the model to capture both past and future context for each time step [29]. Kim and Cho propose a hybrid model that integrates Convolutional Neural Networks (CNN) with Long Short-Term Memory (LSTM) networks to predict residential energy consumption [30]. The CNN component identifies spatial features by analyzing relationships among factors affecting energy use, while the LSTM component captures temporal dependencies present in time-series data. This integrated approach effectively models complex patterns in energy consumption, leading to highly accurate predictions. Empirical results show that the CNN-LSTM model surpasses traditional forecasting methods, delivering near-perfect prediction accuracy with the lowest root mean square error on datasets of individual household power consumption [30].

A Fully Connected Neural Network (FCNN) is a type of neural network in which each neuron in one layer is connected to every neuron in the subsequent layer. It comprises an input layer for raw data, one or more hidden layers that process the data using weighted sums, biases, and activation functions, and an output layer that generates the final prediction. The FCNN architecture has been successfully applied to energy prediction tasks [30].

TabNet provides a novel approach that combines high performance with interpretability. Leveraging sequential attention mechanisms, it dynamically selects relevant features during the learning process [24]. Arik and Pfister demonstrated its effectiveness across multiple domains, including healthcare and finance, where data interpretability is critical. Its adaptability to imbalanced datasets and native processing of mixed-type tabular data make it an ideal candidate for heating load prediction tasks. Based on our knowledge, there are no published papers that used TabNet in energy prediction tasks. Therefore, the main novelty of this work is utilizing TabNet for heating load prediction for the first time.

METHODOLOGY

Dataset

The dataset used in this study contains the structural characteristics of 768 buildings [5], with each building described by eight features: overall height, surface area, relative compactness, roof area, wall area, orientation, glazing area, and glazing area distribution. Additionally, the dataset includes the heating load (HL) consumed in kilowatts (kW) for each building. Since the materials used in these buildings were identical, they were excluded from the analysis. This dataset is widely used in research for predicting HL [5, 7]. It was selected for this study primarily because it includes structural design features, which are crucial for evaluating and selecting optimal building designs. Although other datasets may contain more records, they lack detailed structural information and are thus unsuitable for research that selects the best design based on physical characteristics.

Data Augmentation Module

To enhance the performance of TabNet-based heating load (HL) predictions while reducing model re-training time and cost, a new data augmentation module was used in this study [13]The data augmentation module combines GAN and CGAN to generate synthetic data, marking the first time these techniques and their combination have been applied to energy consumption prediction models.

By utilizing the TabGAN module [13], The training data size expanded from 691 records (90% of the dataset) to 9,469 records through augmentation, which were used for model training. The model's performance will be evaluated on 77 unseen testing records (10% of the dataset) using 10-fold cross-validation.

Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GANs) are widely known for generating synthetic image data but have also been adapted to create synthetic tabular data. As illustrated in Fig. 1, a GAN comprises two deep neural networks, a generator and a discriminator, that are trained simultaneously [31].The generator creates data that mimics real data, aiming to generate outputs that are indistinguishable from actual data. The discriminator assesses these outputs to identify differences. If discrepancies are found, both networks are updated to enhance the generator's capability to produce more realistic and accurate data.

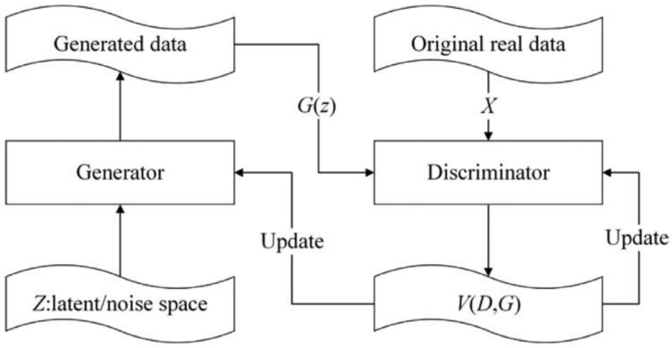

Figure 1: Architecture of GAN [31]

The value function V (D, G) of the GAN can be defined: Z is the noise space, G(z) represents a mapping from the noise space to the generated data space, and X is the original data space. The V (D, G) is defined as (1) [31]:

$$V(D,G) = E_{x \sim p_{\text{data}}(x)}[\log D(x)] + E_{z \sim p_z(z)}[\log(1 - D(G(z)))] \quad (1)$$

Here, z represents noise sampled from the noise space Z, where $p_z(z)$ defines a prior distribution over the input noise variables. G is a differentiable function implemented as a multilayer perceptron that maps z to the generated data space. x denotes a sample from the original data space X, $p_{data}(x)$ represents the distribution of the original data, and D(x) indicates the probability that xxx originates from the real data rather than being generated by G.

The objective function is optimized during training, with the discriminator maximizing it and the generator minimizing it. In essence, the generator and discriminator are trained by solving the following optimization problem:

$$\min_G \max_D V(D,G) = E_{X \sim p_{\text{data}}(x)}[\log D(x)] + E_{z \sim p_z(z)}[\log(1 - D(G(z)))] \quad (2)$$

The discriminator and generator engage in a two-step min-max optimization game to achieve a well-behaved GAN. Initially, the discriminator (D) is optimized while keeping the generator (G) fixed, aiming to maximize the accuracy of distinguishing real data from generated data. In the subsequent step, the generator (G) is adjusted to minimize the discriminator's accuracy while keeping D fixed. This iterative process continues until convergence. When the generator is fixed, the value function V(D,G) in continuous space can be formally described as (3):

$$\begin{aligned} V(D) &= \int_x p_{\text{data}}(x) \log(D(x)) dx + \int_z p_z(z) \log(1 - D(G(z))) dz \\ &= \int_x \left[p_{\text{data}}(x) \log(D(x)) + p_g(x) \log(1 - D(x)) \right] dx \end{aligned} \quad (3)$$

Where the generative distribution, $p_g(x)$, is picked up from the initial data set, x.

The discriminator's optimal value, $D_G^*(x)$, is found in equation (4) when the generator is fixed:

$$D_G^*(x) = \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_g(x)} \quad (4)$$

When the original data is used as input, the discriminator evaluates the conditional probability of the input data by maximizing the log-likelihood. Consequently, the min-max optimization game described in (2) is reformulated as:

$$C(G) = \max_D V(D,G) = E_{x\sim p_{\text{data}}(x)}\left[\log \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_g(x)}\right] + E_{x\sim p_g}\left[\log \frac{p_g(x)}{p_{\text{data}}(x) + p_g(x)}\right] \quad (5)$$

The global optimal solution of V(D, G) and the minimum value of the virtual training criterion C(G) will be reached if and only if $p_{data} = p_g$.

Conditional Tabular GAN (CGAN)

A Conditional Tabular GAN (CTGAN) is a GAN-based technique designed to model tabular data distributions using sample rows from the data. Xu et al. introduced mode-specific normalization to address challenges arising from multimodal and non-Gaussian distributions in tabular data [19] . They also developed a conditional generator to handle imbalanced discrete columns and trained the model using advanced techniques and fully connected networks.

In a Conditional GAN (CGAN), the discriminator and generator are conditioned on additional information, y, such as class labels or auxiliary data from various modalities [32]. This conditioning is implemented by adding y as an additional input layer to both the generator and the discriminator. The generator constructs a joint hidden representation combining the prior input noise, $p_z(z)$, and the conditional input y. The adversarial training framework enables the generator to leverage this hidden representation flexibly, enhancing its ability to generate realistic and relevant synthetic data.

Proposed TabNet-GAN model structure

TabNet is a deep learning architecture tailored for tabular data, merging the interpretability of traditional machine learning models with the predictive capabilities of neural networks [24]. TabNet employs a sequential attention mechanism, unlike standard neural networks, to dynamically select the most relevant features, enabling the model to focus on different subsets of features at each decision step. This architecture provides both high accuracy and interpretability, highlighting which features contribute most to the prediction. TabNet processes data efficiently by combining a sparsity-inducing regularization mechanism and gradient-based learning, reducing overfitting and improving generalization. Its ability to handle mixed data types (numerical and categorical) without extensive preprocessing, such as normalization or one-hot encoding, makes it versatile and practical for real-world applications in finance, healthcare, and energy prediction.

To improve TabNet prediction performance, a new model based on the TabGan Module and TabNet networks was proposed. Fig 2 shows the proposed TabNet-Gan model structure.

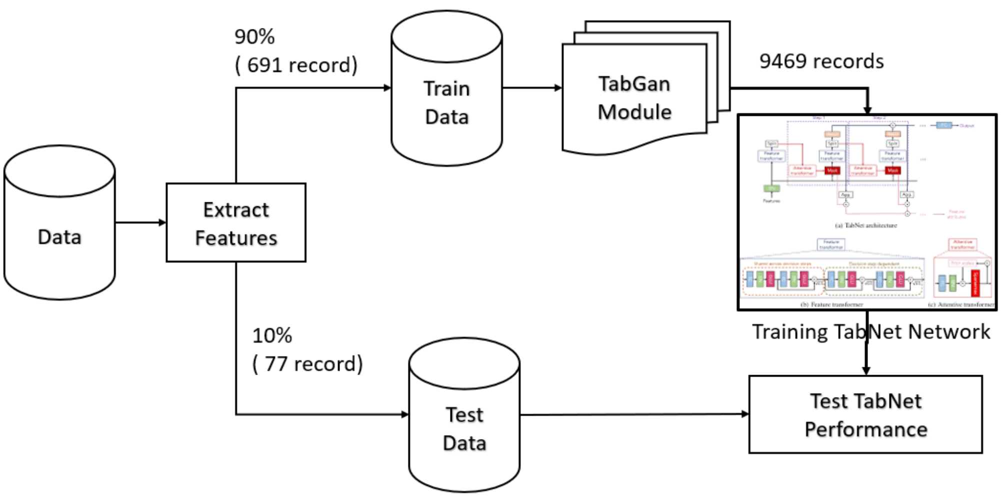

Figure 2: The structure of the proposed TabNet-GAN model structure

Model Training and Testing

This study employed the 10-fold cross-validation (CV) technique for training and testing the model. In 10-fold cross-validation, the dataset is divided into ten folds, with nine folds used for training and the remaining one for testing in each iteration. This process is repeated ten times, ensuring the model is evaluated on different unseen folds in each iteration, thereby assessing its generalization ability. The average performance across these ten iterations is reported as the model's performance metric, providing a more robust evaluation than relying on a single test fold. Additionally, the performance of the proposed TabNet model will be evaluated both with and without data augmentation to analyze its impact.

Testing Model Performance Metrics

Various statistical metrics were calculated to assess the performance and accuracy of the prediction models, evaluating how closely the predicted values match the actual values. Common metrics for evaluating continuous target variables prediction include Mean Absolute Error (MAE), Mean Square Error (MSE), and Root Mean Square Error (RMSE) [5, 7]. Smaller values for these metrics indicate higher model accuracy and lower prediction errors. The calculations for these metrics are based on Equations (6) to (8).

$$MAE = \frac{1}{N}\sum_{i=1}^{N}|o_i - y_i| \quad (6)$$

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(o_i - y_i)^2 \quad (7)$$

$$RMSE = \sqrt{MSE} \quad (8)$$

Here, N represents the total number of samples in the testing dataset, o_i is the predicted value of HL generated by the proposed model, y_i is the actual value of HL or CL, and $\overline{y_i}$ is the mean of the actual values. The overall model performance is determined by averaging the results of 10 runs using a 10-fold cross-validation process to ensure a fair comparison.

RESULTS

This section presents the HL prediction performance of different deep learning models. To compare fairly and reduce the biased data splitting mechanism on model performance. The models run five times independently with different random_state parameters of the 10-Fold Cross Validation method. In addition, the shuffle parameter of the KFold method is set to true to generate different training and testing datasets each time. Finally, the average performance of five different runs is considered the model's final performance metric and evaluated with other models. Tables 1, 2,3,4, and 5 showed LSTM, BiLSTM, FCNN, ResNet, and TabNet models with and without data augmentation.

Table 1: Hating Load Prediction Result with LSTM and LSTM-GAN Model

	LSTM				LSTM-GAN			
Random State	**MAE**	**MSE**	**RMSE**	**R2**	**MAE**	**MSE**	**RMSE**	**R2**
42	2.16	10.01	2.82	0.90	1.32	2.96	1.69	0.97
86	1.00	1.67	1.25	0.98	1.03	1.87	1.32	0.98
92	0.80	1.03	1.00	0.99	1.29	2.91	1.66	0.97
102	0.94	1.35	1.15	0.99	1.33	3.49	1.75	0.97
76	0.79	0.98	0.98	0.99	1.21	2.80	1.63	0.97
Average	**1.14**	**3.01**	**1.44**	**0.97**	**1.24**	**2.81**	**1.61**	**0.97**

Table 2: Hating Load Prediction Result with BiLSTM and BiLSTM-GAN Model

	BiLSTM				BiLSTM-GAN			
Random State	**MAE**	**MSE**	**RMSE**	**R2**	**MAE**	**MSE**	**RMSE**	**R2**
42	2.10	8.96	2.81	0.91	1.02	1.86	1.30	0.98
86	1.00	1.61	1.25	0.98	0.87	1.36	1.14	0.99
92	0.77	1.12	0.99	0.99	0.90	1.37	1.13	0.99
102	0.61	0.58	0.76	0.99	1.00	1.75	1.31	0.98
76	0.59	0.55	0.74	0.99	0.92	1.54	1.21	0.99
Average	**1.01**	**2.56**	**1.31**	**0.97**	**0.94**	**1.58**	**1.22**	**0.98**

Table 3: Hating Load Prediction Result with FCNN and FCNN-GAN Model

	FCNN				FCNN-GAN			
Random State	**MAE**	**MSE**	**RMSE**	**R2**	**MAE**	**MSE**	**RMSE**	**R2**
42	2.39	9.99	3.12	0.90	0.96	1.57	1.22	0.98
86	1.76	5.78	2.38	0.94	0.86	1.30	1.13	0.99
92	1.49	4.30	2.02	0.96	0.80	1.11	1.03	0.99
102	0.65	0.75	0.84	0.99	0.94	1.56	1.23	0.98
76	0.59	0.69	0.76	0.99	0.86	1.31	1.12	0.99
Average	**1.38**	**4.30**	**1.82**	**0.96**	**0.88**	**1.37**	**1.15**	**0.99**

Table 4: Hating Load Prediction Result with REsNet and ResNet-GAN Model

	ResNet				ResNet-GAN			
Random State	**MAE**	**MSE**	**RMSE**	**R2**	**MAE**	**MSE**	**RMSE**	**R2**
42	2.47	10.48	3.20	0.90	1.17	2.31	1.47	0.98
86	1.65	4.44	2.09	0.96	0.88	1.26	1.09	0.99
92	1.36	3.09	1.69	0.97	0.95	1.41	1.17	0.99
102	1.19	2.40	1.54	0.98	1.21	2.54	1.47	0.98
76	1.15	2.35	1.48	0.98	0.90	1.27	1.11	0.99
Average	**1.57**	**4.55**	**2.00**	**0.95**	**1.02**	**1.76**	**1.26**	**0.98**

Table 5: Hating Load Prediction Result with TabNet and TabNet-GAN Model

	TabNet				TabNet-GAN			
Random State	**MAE**	**MSE**	**RMSE**	**R2**	**MAE**	**MSE**	**RMSE**	**R2**
42	5.04	67.29	6.15	0.35	0.47	0.40	0.62	1.00
86	5.26	56.96	6.37	0.43	0.40	0.28	0.53	1.00
92	5.40	57.65	6.65	0.40	0.42	0.31	0.55	1.00
102	4.55	52.00	5.62	0.49	0.42	0.32	0.55	1.00
76	4.83	79.46	6.34	0.13	0.40	0.29	0.54	1.00
Average	**5.02**	**62.67**	**6.23**	**0.36**	**0.42**	**0.32**	**0.56**	**1.00**

Table 6. and Fig 3. summarize all model's performance with and without the data augmentation module. It has been shown that without data augmentation the BiLSTM outperformed the other four models. In addition, the average prediction result presented in Table 5 showed that by using the data augmentation module, all models' performance improved. The TabNet-GAN model obtained the largest prediction performance improvement. The results indicated that the proposed TabNet-GAN model outperformed in all five different executions, Table 5 and Fig 4. Therefore, the proposed TabNet-GAN model has stable results and could be used to solve real-world problems like HL estimation and the prediction of residential buildings.

Table 6: Hating Load Prediction Models Five Different Run Average Performance

	Without Data Augmentation				**With Data Augmentation**			
Model	**MAE**	**MSE**	**RMSE**	**R2**	**MAE**	**MSE**	**RMSE**	**R2**
LSTM	1.14	3.01	1.44	0.97	1.24	2.81	1.61	0.97
BiLSTM	**1.01**	**2.56**	**1.31**	**0.97**	0.94	1.58	1.22	0.98
FCNN	1.38	4.30	1.82	0.96	0.88	1.37	1.15	0.99
ResNet	1.57	4.55	2.00	0.95	1.02	1.76	1.26	0.98
TabNet	5.02	62.67	6.23	0.36	**0.42**	**0.32**	**0.56**	**1.00**

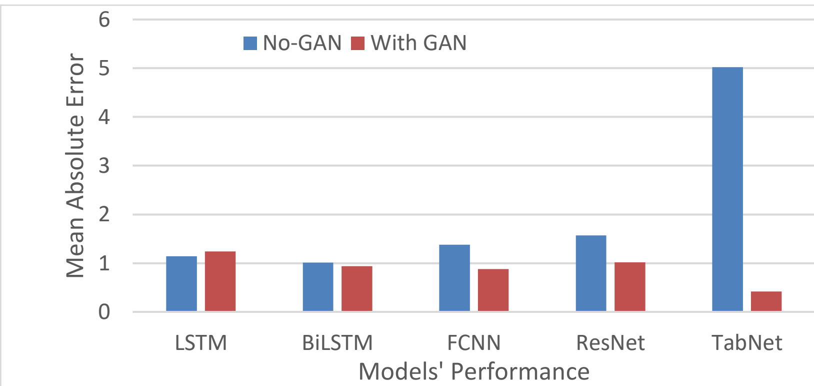

Figure 3: The HL estimation MAE Error of different models with and without GAN Module.

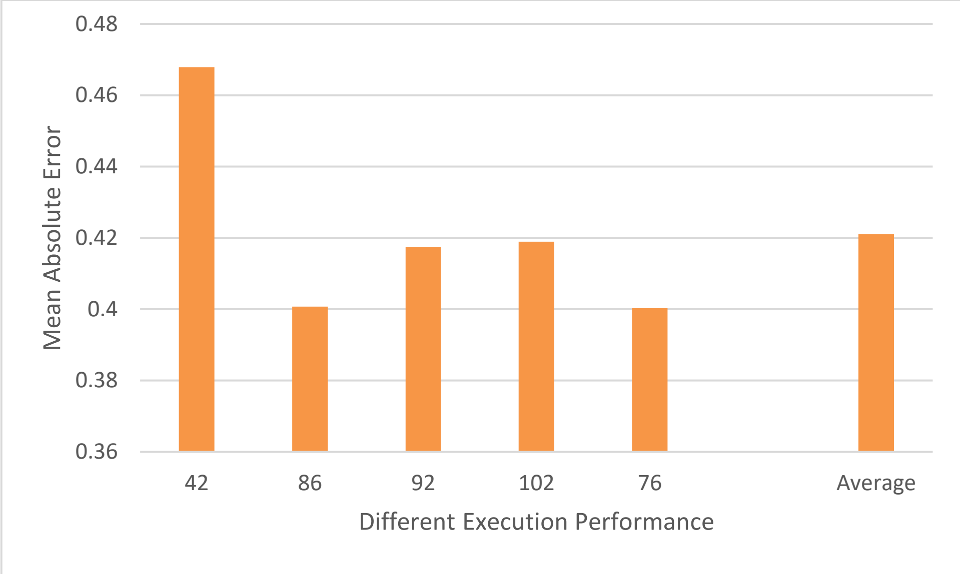

Figure 4: The TabNet-GAN model Performance in Different States

Table 7 and Fig 5. compares the proposed TabNet-GAN, HL prediction performance with some published work that used in-hand dataset [5]. The results show that despite very little data for use in training deep learning models, these models have achieved acceptable accuracy and outperformed many classic machine learning models.

Table 7: Comparison of different HL prediction models' performance

	HL Prediction	
HL Prediction Models	MAE	RMSE
RF (2012)[5]	0.51	
SVR+ANN (2014)[33]	0.236	0.35
RF (2017)[8]	0.351	0.22
GS-XGB (2019)[7]	0.175	0.265
RF (2019)[34]	0.557	1.589
MLP-PSO (2020)[9]	1.863	2.569
PSOGWO-MLP (2023)[10]	0.787	1.412
Bayesian-XGB (2023)[35]	0.247	0.380
XGB (2024)[11]	0.356	0.492
XACM (2024)[12]	~0.65	0.904
DAPM-LLM [13]	0.32	0.54
Proposed TabNet-GAN	**0.42**	**0.56**

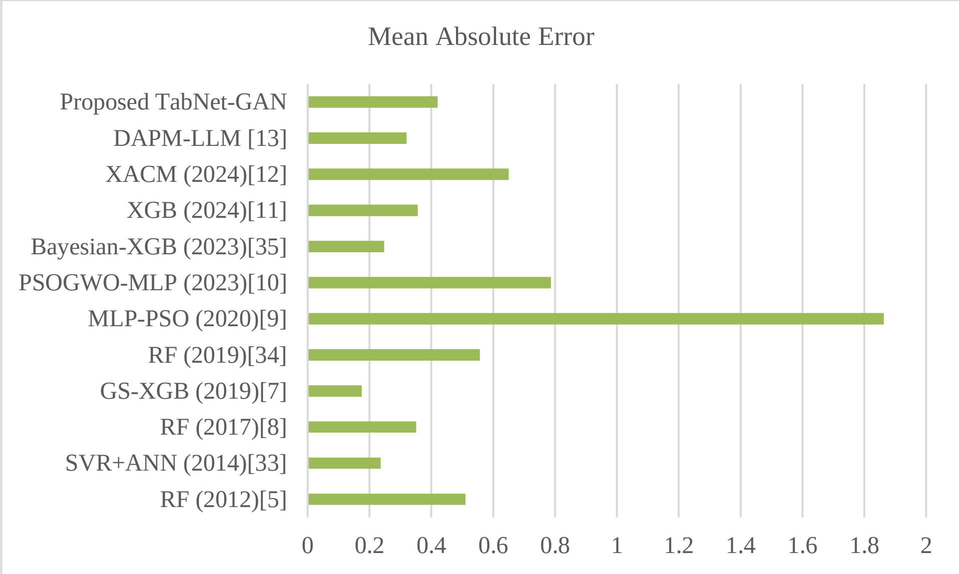

Figure 5: Different Models HL prediction performance.

The presented results show that combining the GAN-based data augmentation method with the TabNet model has significantly improved the model's performance. They also show that even with very little training data, the performance of deep learning models can be improved by simulating training data and using it to solve complex engineering problems.

The main reason for outperforming TabNet over other models is its special characteristics. First, the TabNet is particularly developed and well-suited for tabular data. Second, it uses attention mechanisms to focus on the most relevant features during training, aligning well with augmented data that captures complex feature relationships. Third, TabNet inherently provides feature importance, enabling better utilization of enriched data. Finally, Unlike LSTM or ResNet, TabNet directly learns from raw tabular data, avoiding heavy preprocessing and leveraging the augmented data effectively. TabNet's attention-based architecture enables it to exploit the nuanced relationships in the augmented dataset, effectively capturing both linear and nonlinear interactions. This makes it more adaptable than FCNNs or sequential models like LSTMs. In addition, GAN-based data augmentation provides diverse and high-quality synthetic samples, addressing potential issues of small or imbalanced datasets. This improves the generalizability of TabNet by exposing it to a more representative training dataset, which traditional models like FCNN or LSTM cannot leverage directly.

The main limitation of this study is the use of a small dataset. The other available dataset doesn’t contain building design characteristics and, therefore, could not be used to develop an HL estimation model based on its characteristics.

CONCLUSION

Efficient heating load (HL) prediction is crucial for optimizing energy usage and reducing costs in residential buildings. This study introduced a novel hybrid model that integrates the TabNet deep learning framework with a GAN-based data augmentation module to tackle the issue of limited training data. The data augmentation module increased the training dataset size five times, significantly enhancing the performance of deep learning models, especially TabNet.

Without data augmentation, the BiLSTM model demonstrated superior performance to other deep learning models, including ResNet, FCNN, TabNet, and LSTM. However, the proposed TabNet-GAN hybrid model outperformed all other deep learning approaches and achieved results comparable to classical machine learning models, often preferred for small datasets. TabNet's ability to leverage enriched datasets generated by GANs and its interpretable architecture make it a robust and practical choice for HL prediction tasks.

These results suggest that the TabNet-GAN model can be a reliable and interpretable tool for predicting heating loads in residential buildings. Providing accurate predictions empowers engineers and architects to evaluate energy consumption patterns and optimize building designs from an energy efficiency perspective. Future research could explore integrating additional generative models or hybrid architectures to enhance prediction accuracy and generalizability further.

REFERENCES

[1] N. Somu, G. R. MR, and K. Ramamritham, "A hybrid model for building energy consumption forecasting using long short term memory networks," Applied Energy, vol. 261, p. 114131, 2020.

[2] Y. Sun, F. Haghighat, and B. C. Fung, "A review of the state-of-the-art in data-driven approaches for building energy prediction," Energy and Buildings, vol. 221, p. 110022, 2020.

[3] M. Khalil, A. S. McGough, Z. Pourmirza, M. Pazhoohesh, and S. Walker, "Machine Learning, Deep Learning and Statistical Analysis for forecasting building energy consumption—A systematic review," Engineering Applications of Artificial Intelligence, vol. 115, p. 105287, 2022.

[4] R. Olu-Ajayi, H. Alaka, I. Sulaimon, F. Sunmola, and S. Ajayi, "Building energy consumption prediction for residential buildings using deep learning and other machine learning techniques," Journal of Building Engineering, vol. 45, p. 103406, 2022.

[5] A. Tsanas and A. Xifara, "Accurate quantitative estimation of the energy performance of residential buildings using statistical machine learning tools," Energy and buildings, vol. 49, pp. 560-567, 2012.

[6] S. A. Kalogirou, "Artificial neural networks in energy applications in buildings," International Journal of Low-Carbon Technologies, vol. 1, no. 3, pp. 201-216, 2006.

[7] M. Al-Rakhami, A. Gumaei, A. Alsanad, A. Alamri, and M. M. Hassan, "An ensemble learning approach for accurate energy load prediction in residential buildings," IEEE Access, vol. 7, pp. 48328-48338, 2019.

[8] G. R. Duarte, L. G. da Fonseca, P. Goliatt, and A. C. de Castro Lemonge, "Comparison of machine learning techniques for predicting energy loads in buildings," Ambiente Construído, vol. 17, no. 3, pp. 103-115, 2017.

[9] G. Zhou, H. Moayedi, M. Bahiraei, and Z. Lyu, "Employing artificial bee colony and particle swarm techniques for optimizing a neural network in prediction of heating and cooling loads of residential buildings," Journal of Cleaner Production, vol. 254, p. 120082, 2020.

[10] S. Afzal, B. M. Ziapour, A. Shokri, H. Shakibi, and B. Sobhani, "Building energy consumption prediction using multilayer perceptron neural network-assisted models; comparison of different optimization algorithms," Energy, vol. 282, p. 128446, 2023.

[11] O. A. Alawi, H. M. Kamar, and Z. M. Yaseen, "Optimizing building energy performance predictions: A comparative study of artificial intelligence models," Journal of Building Engineering, vol. 88, p. 109247, 2024.

[12] B. Sadaghat, S. Afzal, and A. J. Khiavi, "Residential building energy consumption estimation: A novel ensemble and hybrid machine learning approach," Expert Systems with Applications, vol. 251, p. 123934, 2024.

[13] P. R.-K. Qasim Mustafa Zainel, Leila Sharifi, "Leveraging Large Language Models and GAN-Augmented Data for Energy Load Prediction," Journal of Electrical Systems (JES) vol. 20, no. 11s, pp. 1926-1940, 2024.

[14] T. Ahmad, H. Chen, Y. Guo, and J. Wang, "A comprehensive overview on the data driven and large scale based approaches for forecasting of building energy demand: A review," Energy and Buildings, vol. 165, pp. 301-320, 2018.

[15] C. Fan, M. Chen, R. Tang, and J. Wang, "A novel deep generative modeling-based data augmentation strategy for improving short-term building energy predictions," in Building Simulation, 2022, vol. 15: Springer, pp. 197-211.

[16] Y. Gao and Y. Ruan, "Interpretable deep learning model for building energy consumption prediction based on attention mechanism," Energy and Buildings, vol. 252, p. 111379, 2021.

[17] H. Fang, H. Tan, R. Kosonen, X. Yuan, K. Jiang, and R. Ding, "Study of the Data Augmentation Approach for Building Energy Prediction beyond Historical Scenarios," Buildings, vol. 13, no. 2, p. 326, 2023.

[18] N. Park, M. Mohammadi, K. Gorde, S. Jajodia, H. Park, and Y. Kim, "Data synthesis based on generative sssadversarial networks," arXiv preprint arXiv:1806.03384, 2018.

[19] L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni, "Modeling tabular data using conditional gan," Advances in neural information processing systems, vol. 32, 2019.

[20] A. Fernández, S. Garcia, F. Herrera, and N. V. Chawla, "SMOTE for learning from imbalanced data: progress and challenges, marking the 15-year anniversary," Journal of artificial intelligence research, vol. 61, pp. 863-905, 2018.

[21] Y. Zhang, N. A. Zaidi, J. Zhou, and G. Li, "GANBLR: a tabular data generation model," in 2021 IEEE International Conference on Data Mining (ICDM), 2021: IEEE, pp. 181-190.

[22] Y. Zhang, Z. Zhou, J. Liu, and J. Yuan, "Data augmentation for improving heating load prediction of heating substation based on TimeGAN," Energy, vol. 260, p. 124919, 2022.

[23] C. Fan, Y. Lei, Y. Sun, M. S. Piscitelli, R. Chiosa, and A. Capozzoli, "Data-centric or algorithm-centric: Exploiting the performance of transfer learning for improving building energy predictions in data-scarce context," Energy, vol. 240, p. 122775, 2022.

[24] S. Ö. Arik and T. Pfister, "Tabnet: Attentive interpretable tabular learning," in Proceedings of the AAAI conference on artificial intelligence, 2021, vol. 35, no. 8, pp. 6679-6687.

[25] X. Song, Z. Wang, and H. Wang, "Short-term load prediction with LSTM and FCNN models Based on attention mechanisms," in Journal of Physics: Conference Series, 2024, vol. 2741, no. 1: IOP Publishing, p. 012026.

[26] H. Choi, S. Ryu, and H. Kim, "Short-term load forecasting based on ResNet and LSTM," in 2018 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), 2018: IEEE, pp. 1-6.

[27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770-778.

[28] J. Q. Wang, Y. Du, and J. Wang, "LSTM based long-term energy consumption prediction with periodicity," energy, vol. 197, p. 117197, 2020.

[29] S. Siami-Namini, N. Tavakoli, and A. S. Namin, "The performance of LSTM and BiLSTM in forecasting time series," in 2019 IEEE International conference on big data (Big Data), 2019: IEEE, pp. 3285-3292.

[30] T.-Y. Kim and S.-B. Cho, "Predicting residential energy consumption using CNN-LSTM neural networks," Energy, vol. 182, pp. 72-81, 2019.

[31] I. Goodfellow et al., "Generative adversarial nets," Advances in neural information processing systems, vol. 27, 2014.

[32] M. Mirza and S. Osindero, "Conditional generative adversarial nets," arXiv preprint arXiv:1411.1784, 2014.

[33] J.-S. Chou and D.-K. Bui, "Modeling heating and cooling loads by artificial intelligence for energy-efficient building design," Energy and Buildings, vol. 82, pp. 437-446, 2014.

[34] L. T. Le, H. Nguyen, J. Zhou, J. Dou, and H. Moayedi, "Estimating the heating load of buildings for smart city planning using a novel artificial intelligence technique PSO-XGBoost," Applied Sciences, vol. 9, no. 13, p. 2714, 2019.

[35] B. A. Salami, S. I. Abba, A. A. Adewumi, U. A. Dodo, G. K. Otukogbe, and L. O. Oyedele, "Building energy loads prediction using bayesian-based metaheuristic optimized-explainable tree-based model," Case Studies in Construction Materials, vol. 19, p. e02676, 2023.