

# Next-Generation Air Conditioning: Leveraging Machine Learning For Advanced System Modeling and Optimization

<sup>1</sup>Prince N Nwankwo, <sup>2</sup>A.N Isizoh

<sup>1</sup>Department of Computer Engineering, Federal Polytechnic, Oke, Anambra State, Nigeria

<sup>2</sup>Department of Electronic and Computer Engineering, Nnamdi Azikiwe University, Awka, Nigeria

Date of Submission: 10-12-2024

Date of Acceptance: 20-12-2024

## ABSTRACT:

The introduction of next-generation HVAC (Heating, Ventilation, and Air Conditioning) systems has spurred a revolution in HVAC technology, especially with regard to the incorporation of machine learning algorithms for improved system optimization. The ideas of innovation (next generation), technology (machine learning), and practical application (system modeling and optimization) are all skillfully combined in this study. It investigates how to use cutting-edge machine learning approaches to enhance the operational effectiveness, energy efficiency, and adaptability of air conditioning systems. Using time series analysis and machine learning approaches, the study focuses on developing an advanced model and algorithm to estimate the energy usage of air conditioners. The goal is to develop a system that can forecast energy consumption in response to time and environment dependent factors like temperature, humidity, and daytime. Precise estimation of energy usage can enhance the effectiveness of air conditioning systems, minimize energy wastage, and boost financial viability, particularly in residential and commercial structures. The outcomes of this study underline the feasibility of machine learning as a vital instrument in defining the future of HVAC technology and the broader sustainability of building management systems. These results offer a way forward for creating intelligent, cutting-edge air conditioning systems that tackle environmental and financial issues.

**Keywords:** Machine learning, air conditioning, energy efficiency, HVAC, adaptive control.

## 1. INTRODUCTION:

Every day, HVAC (Heating, Ventilation, and Air Conditioning) systems contribute significantly to our comfort. The introduction of technology, particularly artificial intelligence (AI), has ushered in a digital era for the HVAC industry. AI algorithms and machine learning are being used in HVAC systems to make them smarter, more

efficient, and energy-optimized. Given the rising worldwide energy consumption, there is a greater need than ever for eco-friendly and efficient air conditioning systems (K.A. Akpado, P. N. Nwankwo, et al., 2018). Actually, a significant portion of the energy needed by buildings for living or working is accounted for by air conditioning, which increases CO<sub>2</sub> emissions in the environment and puts more strain on power networks (U.S. Department of Energy 2020; Zhou et al. 2020).

Thus, the goal of developing next-generation air conditioning systems is to significantly reduce energy usage while preserving or enhancing occupant comfort (Lee & Park, 2021; Wang et al., 2022). Conventional AC control approaches, including PID (Proportional-Integral-Derivative) controllers or rule-based approaches, frequently find it difficult to effectively handle dynamic and complicated interior situations (Johnson et al., 2019; Sun & Zhang, 2020).

These methods frequently are not flexible enough to adjust in real time, which might result in inefficiencies during periods of high demand or quickly shifting circumstances (Chaudhary et al., 2018). In order to improve HVAC (Heating, Ventilation, and Air Conditioning) systems' capacity to model, anticipate, and optimize performance, there has been a spike in interest in incorporating artificial intelligence (AI), particularly machine learning (ML) (Kim et al., 2021; Zhou & Li, 2020).

The potential of machine learning to transform the modeling and optimization of next-generation air conditioning systems is examined in this research. We seek to show how these technologies can convert AC systems from passive, energy-intensive devices into intelligent, adaptive systems that strike a balance between comfort and sustainability through a thorough investigation of machine learning models and algorithms. Additionally, we provide case examples that demonstrate how ML-driven AC systems are used in the real world and emphasize the considerable

energy savings, increased performance, and increased adaptability that these systems provide.

### A. Problem Statement

Power grids are under a great deal of strain due to the rising worldwide energy demand brought on by increased urbanization and climate change, with air conditioning systems using a large amount of electricity. Because they rely on static models that don't take into account changing occupancy patterns or real-time environmental changes, traditional HVAC (heating, ventilation, and air conditioning) systems are frequently inefficient. Higher operating expenses and excessive energy use are both caused by this inefficiency. Adaptive systems that can automatically control energy use and preserve ideal indoor air quality and comfort levels are desperately needed. Because machine learning (ML) can handle large volumes of data and build predictive models that can improve system performance, it has become a promising option.

Recent developments in machine learning (ML) techniques, including hybrid approaches, neural networks, and reinforcement learning, present the possibility of creating highly optimal, self-learning air conditioning systems that dynamically adapt to user behaviour and the surrounding environment. This study intends to solve the inefficiencies of standard air conditioning systems by presenting a next-generation HVAC solution, employing machine learning to increase system modeling, maximize operational efficiency, and reduce energy consumption.

### B. Aim of the Study

To explore the development of next-generation air conditioning systems that leverage advanced machine learning techniques for improved system modeling, and optimization.

### C. The objectives of the Study are

1. To develop an advanced machine learning algorithm to improve air conditioning systems' performance and adaptability in response to changing inputs.
2. To put into practice ML-driven tactics that maximize overall operational efficiency and cost-effectiveness by lowering energy consumption while preserving indoor comfort.
3. To integrate data from environmental sensors (temperature, humidity, occupancy) to enable continuous system adaptation, allowing the air conditioning system to learn and self-optimize over time.

4. To use simulations to test and validate the suggested ML-enhanced air conditioning system and evaluate its scalability, accuracy, and efficiency in different scenarios.

5. To evaluate the energy savings, responsiveness, and performance of the machine learning-driven air conditioning model in comparison to traditional HVAC systems.

## II CONCEPT OF THE PROJECT

In order to improve the functionality and energy efficiency of air conditioning systems, the study "Next-Generation Air Conditioning: Leveraging Machine Learning for Advanced System Modeling and Optimization" looks into the integration of machine learning techniques. The goal of the technique is to minimize energy waste and optimize operational settings by developing an advanced model that forecast energy consumption and system behaviour. The research demonstrates how machine learning can dynamically alter cooling parameters in real-time, leading to more sustainable and cost-effective climate management solutions. The study emphasizes how effectively AI-driven air conditioning systems can adjust to shifting user preferences and environmental changes.




Figure 1: Neural network training for a HVAC control system

### A. Artificial Intelligence (AI) in HVAC Systems

Artificial Intelligence involves the development of intelligent algorithms that enable machines to learn from data, make decisions, and execute tasks without being explicitly programmed. In the realm of HVAC (Heating, Ventilation, and Air Conditioning) systems, AI analyzes data from a range of sensors and inputs to optimize system performance, control airflow and temperature, and minimize energy use.

Some basic needs for the integration of Artificial Intelligence (AI) in HVAC Systems are:

### 1. Enhanced Comfort and Personalization

AI-driven HVAC systems can learn user preferences and automatically adjust settings to enhance comfort. By examining patterns in user behaviour and preferences, AI can generate customized comfort profiles, providing tailored temperature and airflow control across different building zones, thereby improving overall occupant comfort and satisfaction.

### 2. Energy Efficiency and Cost Savings

By continuously monitoring and adapting to real-time elements including occupancy, thermal loads, and exterior temperature, artificial intelligence (AI) improves the energy efficiency of HVAC systems. These systems dynamically adjust settings to cut down on energy use, which lowers utility costs and has a less negative effect on the environment (Prince. N Nwankwo, K. A. Akpado, et al., 2021).

### 3. Predictive Maintenance and Fault Detection

AI makes predictive maintenance possible for HVAC systems by examining performance data to find wear indicators or possible problems early on. This proactive approach allows for timely maintenance, minimizing downtime, and ensuring system reliability. Furthermore, AI aids in defect diagnosis and detection, assisting in the discovery of issues' underlying causes and enabling prompt, precise fixes.

### 4. Smart Load Management and Demand Response

AI algorithms enhance energy efficiency by intelligently managing HVAC loads. During times of peak energy demand, AI-powered systems can automatically adjust temperature settings or implement load-shedding strategies to reduce energy consumption, all while maintaining occupant comfort. This smart load management supports grid stability and helps lower overall energy costs.

### 5. Future Trends in AI Integration

The integration of Artificial Intelligence (AI) into HVAC systems is an evolving process with promising future developments. As technology progresses, AI will drive several transformative trends, further improving HVAC

performance, energy efficiency, and user experience.

### B. The Structure of Machine Learning (ML)



Figure 2(a). Machine learning techniques

Machine learning uses two types of techniques: supervised learning, which trains a model on known input and output data so that it can predict future outputs, and unsupervised learning, which finds hidden patterns or intrinsic structures in input data.

#### 1. Supervised Learning

In the face of uncertainty, supervised machine learning creates a model that bases predictions on data. A supervised learning technique trains a model to produce plausible predictions for the response to incoming data given a known set of input data and known responses to the data (output). It is a prominent strategy in machine learning that involves training a model on a dataset with known labels.

Labeled data is used in supervised learning, where each training example has an output label associated with it. By spotting patterns in the data, the model has the ability to translate inputs into outputs. Making accurate predictions on new, unknown data is the goal of generalizing from the training set (Ng, 2021).

In Supervised learning, model performance is evaluated using metrics such as accuracy, precision, recall, F1-score for classification tasks, and mean squared error (MSE) or mean absolute error (MAE) for regression tasks (Chicco & Jurman, 2021). Supervised learning is widely used in various domains, including finance for credit scoring, healthcare for disease prediction, and marketing for customer segmentation (Kotsiantis, 2019).

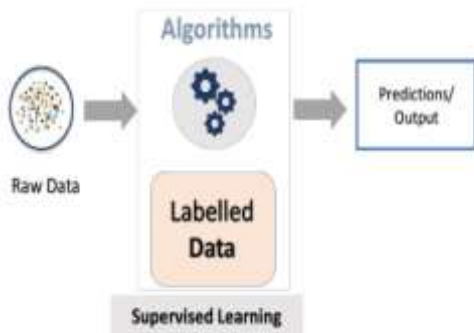


Figure 2(b): The structure of Supervised Machine Learning

## 2. Unsupervised Learning

In unsupervised learning, patterns are discovered by the algorithm from unlabeled data (Russell & Norvig, 2016). It differs from supervised learning, which bases prediction-making on labeled datasets (Goodfellow et al., 2016). Without human assistance, the machine finds innate patterns and relationships in the data through unsupervised learning (Murphy, 2012). Algorithms used in this area try to group data items that share common features (Bishop, 2006).

Clustering and association are two key techniques of unsupervised learning (Aggarwal & Reddy, 2014). Clustering algorithms, like K-means and hierarchical clustering, group similar data points based on certain features or metrics (Jain, 2010). On the other hand, association rule learning discovers relationships between variables in large datasets, such as in market basket analysis (Agrawal et al., 1993).

Scalability is a crucial feature of unsupervised learning since it allows for the application to huge datasets without requiring costly labeling efforts (Hastie, Tibshirani, & Friedman, 2009). However, because there are no labels or ground truth to confirm the results, evaluating the outputs of unsupervised learning can be difficult (Zhou & Mao, 2020). In unsupervised learning, dimensionality reduction methods like principal component analysis (PCA) and t-SNE are frequently employed to deconstruct high-dimensional data (van der Maaten, 2008).




Figure 2(c): The structure of Unsupervised Machine Learning

One of the practical applications of unsupervised learning is anomaly detection, where the algorithm identifies data points that deviate from expected patterns (Chandola, Banerjee, & Kumar, 2009). It is also widely utilized in picture identification, natural language processing, and bioinformatics (Manning, Raghavan, & Schütze, 2008).

Applications for cluster analysis include DNA sequence analysis, market research, and object recognition. To improve the locations for cell phone towers, a mobile phone company, for instance, can use machine learning to determine the amount of clusters of people dependent on their towers. Because a phone can only talk with one tower at a time, the team uses clustering algorithms to figure out where to put the cell towers so that client clusters may receive the strongest signal possible.

## 3. Reinforcement Learning

Reinforcement Learning is another part of Machine Learning that is gaining a lot of prestige in how it helps the machine learn from its progress.




Figure 3: The Structure of Reinforcement Machine Learning

Reinforcement learning (RL) is a branch of machine learning where an agent learns to make decisions by interacting with its environment (Sutton, R. S., Barto, A. G., 2018). With incentives or penalties based on the tasks it does, the agent's



goal is to maximize the cumulative reward over time (Kaelbling, L. P. et al., 1996). Unlike supervised learning, which depends on labeled data, reinforcement learning learns from the outcomes of the agent's activities (Szepesvári, C., 2010).

One of the most popular algorithms in reinforcement learning is Q-learning, which helps agents determine the optimal course of action based on expected future rewards (Watkins, 1992). Deep reinforcement learning allows agents to learn directly from high-dimensional sensory inputs such as images by combining deep neural networks with reinforcement learning (Mnih, V., et al., 2015). Reinforcement learning finds applications in robotics, gaming, financial portfolio management, and autonomous vehicles.

### III. METHODOLOGY

Designing a next-generation air conditioning system using machine learning

```
import numpy as np
import pandas as pd
# Simulating some data for air conditioning system
np.random.seed(42)
n_samples = 1000
outdoor_temp = np.random.normal(30, 5, n_samples) # Outdoor temperature in Celsius
indoor_temp = np.random.normal(22, 2, n_samples) # Indoor temperature in Celsius
humidity = np.random.uniform(30, 90, n_samples) # Humidity in percentage
power_consumption = (outdoor_temp - indoor_temp) * np.random.uniform(0.1, 0.2, n_samples) + np.random.normal(5, 1, n_samples)
# Create a DataFrame
data = pd.DataFrame({
    'Outdoor_Temperature': outdoor_temp,
    'Indoor_Temperature': indoor_temp,
    'Humidity': humidity,
    'Power_Consumption': power_consumption
})
print(data.head())
```

The above programme is the Python code for the dataset of the HVAC. From the code, Power consumption is modeled as a function of the temperature difference between outdoor and indoor temperatures. The difference is multiplied by a random efficiency factor between 0.1 and 0.2, simulating how much power is consumed based on the temperature gap. An additional random noise (normal distribution with mean 5 and standard

involves multiple steps, including data collection, preprocessing, model building, system optimization, and visualization. Below is a step-by-step breakdown of how to approach the project.

#### 1. Problem Definition

The goal is to design an air conditioning system that leverages machine learning to optimize energy efficiency and maintain a comfortable indoor temperature. The model will predict energy consumption based on various factors (such as outdoor temperature, humidity, user preferences) and then optimize the system's operations to minimize energy use.

#### 2. Data Collection

We need a dataset with features such as: Outdoor temperature, Indoor temperature, Humidity, Power consumption, Cooling system settings (fan speed, compressor state, etc.).

deviation 1) is added to simulate other factors affecting power consumption.

The performance of an air cooling system is simulated by the python code, which generates random data as indicated in Table 1 and focuses on the correlation between indoor and outdoor temperatures, humidity, and power consumption. For later study, the simulated data can be found stored in a pandas DataFrame.

Table 1: Data to simulate an air-conditioning system's performance

	Outdoor_Temperature	Indoor_Temperature	Humidity	Power_Consumption
0	32.483571	24.798711	54.426389	7.069270
1	29.308678	23.849267	33.960591	6.457650
2	33.238443	22.119261	50.929232	9.309322
3	37.615149	20.706126	36.659886	5.443363
4	28.829233	23.396447	78.494113	7.117993

### 3. Data Preprocessing

The next stage in the modeling and optimization of an HVAC system is to preprocess the data by normalizing the features and splitting

the data into training and test sets as shown in the code below. The code demonstrates how to prepare data for training a machine learning model using scikit-learn.

```
from sklearn.model_selection import train_test_split
from sklearn.preprocessing import StandardScaler

# Splitting the data into input features (X) and output target (y)
X = data[['Outdoor_Temperature', 'Indoor_Temperature', 'Humidity']]
y = data['Power_Consumption']

# Normalize the input features
scaler = StandardScaler()
X_scaled = scaler.fit_transform(X)

# Split the dataset into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X_scaled, y, test_size=0.2, random_state=42)
```

\* The input characteristics (X) are standardized using StandardScaler. It turns each feature into a common scale by subtracting the mean and dividing by the standard deviation. This is especially useful for machine learning models that depend on the assumption of regularly distributed data or features that are sensitive to varying magnitudes.

\* fit\_transform (X) computes the mean and standard deviation from X and then applies the transformation, fitting the scaler to the data and transforming it in the process.

By standardizing the input features and dividing the data into training and testing sets, the code gets a dataset ready for training. Standardization guarantees that features are on the same scale, while data splitting permits appropriate model evaluation on data that has not yet been seen (test set).

### 4. Model Building

For this system, we use a **regression model** to predict the power consumption based on environmental conditions. Below we use Random Forest Regressor to build the model.

```
from sklearn.ensemble import RandomForestRegressor
from sklearn.metrics import mean_squared_error

# Initialize the Random Forest Regressor
rf_model = RandomForestRegressor(n_estimators=100, random_state=42)

# Train the model
rf_model.fit(X_train, y_train)

# Predicting the power consumption
y_pred = rf_model.predict(X_test)

# Evaluate the model
mse = mean_squared_error(y_test, y_pred)
print(f'Mean Squared Error: {mse}')
```

#### Explanation of the code (Model Building)

##### A. Importing Required Libraries:

□ **RandomForestRegressor:** A machine learning model that carries out regression tasks using a group of decision trees.

□ **mean\_squared\_error:** A metric that calculates the average squared difference between actual and predicted values. It is used to measure how well the model performs.

### B. Initializing the Random Forest Regressor:

□ **n\_estimators = 100**: The number of decision trees in the forest. In this case, the model uses 100 decision trees.

□ **random\_state = 42**: Ensures reproducibility of the model by controlling the randomness involved in building the trees.

### C. Training the Model

`fit(X_train, y_train)`: This trains the `RandomForestRegressor` using the training data (`X_train`) and the corresponding target values (`y_train`). The model learns patterns from the training data to predict power consumption.

### D. Making Predictions

`Predict(X_test)`: Once the model is trained, it uses the testing set (`X_test`) to make predictions. These predictions are stored in `y_pred`. The values in `y_pred` are the model's estimates for power consumption.

### E. Evaluating the Model

`mean_squared_error(y_test, y_pred)`: This function computes the Mean Squared Error (MSE), which is a measure of how close the model's predictions (`y_pred`) are to the actual target values (`y_test`).

□ **Formula**:  $MSE =$

$$\frac{1}{n} \sum_{i=1}^n (y_{test} - y_{pred})^2 \dots\dots Equ (1)$$

```
from sklearn.ensemble import RandomForestRegressor
from sklearn.preprocessing import StandardScaler
from sklearn.model_selection import train_test_split

# Example of preparing data (X_train should be your actual data)
X = np.random.rand(1000, 3) # Simulate 3 feature columns: Outdoor_Temp, Indoor_Temp, Humidity
y = np.random.rand(1000)    # Simulate target: Power Consumption

# Splitting data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)

# Scaling the data
scaler = StandardScaler()
X_train_scaled = scaler.fit_transform(X_train)
X_test_scaled = scaler.transform(X_test)

# Train the RandomForest model
rf_model = RandomForestRegressor()
rf_model.fit(X_train_scaled, y_train)
```

### Line-by-line Explanation of the code:

#### A. Imports:

□ **RandomForestRegressor**: This is an ensemble machine learning model that builds multiple decision trees during training and outputs the

□ **Interpretation**: The smaller the MSE, the better the model's predictions. However, MSE can be harder to interpret directly because it's in the squared units of the target variable.

□ **Result**: The calculated MSE is 1.355467332768074, which means that, on average, the squared difference between the true and predicted power consumption values is about 1.36.

### F. Summary on Model Building

□ The Random Forest Regressor was trained on the training data to predict power consumption based on features like outdoor and indoor temperature and humidity.

□ The model's performance was evaluated using Mean Squared Error (MSE), which gives an idea of how well the model predicts the power consumption on the test set. An MSE of 1.355 means the average squared error between the predicted and actual power consumption values is relatively low.

### 5. Pre-train Model and Scaler

Train the `RandomForestRegressor` model. The python code train the model and fit the scaler before running the optimization.

average prediction of individual trees for regression tasks.

□ **StandardScaler**: This is used to scale features so that they have zero mean and unit variance, which

helps some models perform better, especially when features are measured in different units.

□ **train\_test\_split**: The dataset is divided into training and testing sets using this technique. The testing set is used to assess the model after it has been fitted using the training set.

### B. Simulating data

**X**: This is a simulated dataset with 1000 rows and 3 columns representing features such as "Outdoor Temperature," "Indoor Temperature," and "Humidity."

**Y**: This is a simulated target variable representing "Power Consumption" for each instance.

### C. Splitting the data

□ The data is split into training and testing sets.  
□ **test\_size = 0.2**: 20% of the data will be used as the testing set, while 80% will be used for training.  
□ **random\_state = 42**: Ensures the split is reproducible, meaning the same data split will occur each time the code is run.

### D. Scaling the data

□ The features in both the training and test sets are scaled to have zero mean and unit variance.  
□ **fit\_transform(X\_train)**: The fit\_transform method computes the mean and variance from the training data and applies scaling to it.

□ **transform(X\_test)**: The transform method uses the same mean and variance from the training data to scale the test data.

### E. Training the RandomForest model:

□ **RandomForestRegressor()**: Initializes a random forest regressor model.

□ **fit(X\_train\_scaled, y\_train)**: Trains the model using the scaled training data (X\_train\_scaled) and the corresponding target values (y\_train).

### Pre-train Model and Scaler Summary:

The Pre-train Model and Scaler code shown above demonstrates how to create a machine learning pipeline where the data is split into training and test sets, scaled using StandardScaler, and then used to train a RandomForestRegressor model for predicting a continuous target variable (e.g., power consumption).

## 6. System Optimization

After building the model, the next step is to optimize the system's operation to minimize power consumption. One approach is to use Reinforcement Learning or an Optimization Algorithm to control the system settings. For simplicity, we use a grid search to find the optimal indoor temperature and humidity set points that minimize power consumption. The programme below shows the system optimization code.

```
import numpy as np
import pandas as pd
# Assuming rf_model and scaler are defined and trained earlier in your code
# Define the correctly capitalized feature names used when training the scaler and the model
feature_names = ['Outdoor_Temperature', 'Indoor_Temperature', 'Humidity']
# Simulating different indoor temperature set points and humidity levels
indoor_temp_range = np.arange(18, 28, 0.5)
humidity_range = np.arange(40, 60, 1)
optimal_temp = None
optimal_humidity = None
min_power = float('inf')
for temp in indoor_temp_range:
    for hum in humidity_range:
        try:
            # Create a DataFrame for the new input with correctly capitalized feature names
            X_new = pd.DataFrame([[30, temp, hum]], columns=feature_names)

            # Transform the input using the scaler
            X_new_scaled = scaler.transform(X_new)

            # Predict power consumption with fixed settings
            power = rf_model.predict(X_new_scaled)[0]
            if power < min_power:
                min_power = power
                optimal_temp = temp
                optimal_humidity = hum
        except Exception as e:
            print(f"Error during prediction: {e}")
print(f"Optimal Indoor Temperature: {optimal_temp}")
print(f"Optimal Humidity: {optimal_humidity}")
print(f"Predicted Power Consumption: {min_power}")
```



The code is designed to optimize indoor temperature and humidity settings to minimize power consumption, using a machine learning model and pre-fitted scaler.

**Here is a brief breakdown:**

**A. Imports**

□ numpy and pandas are imported to handle arrays and data frames.

**B. Variables and Definitions**

□ rf\_model (Random Forest model) and scaler (for feature scaling) are assumed to be already trained.  
□ Feature\_names stores the names of the features used to train the model and scaler.

**C. Simulating Inputs**

□ Indoor\_temp\_range and humidity\_range define the range of indoor temperatures to be tested.  
□ optimal\_temp, optimal\_humidity, and min\_power initialize variables to store the optimal values.

**D. Optimization Loop**

□ A nested loop iterates over all combinations of indoor temperatures and humidity levels.

**For each combination:**

□ A new DataFrame X\_new is created with current values for outdoor temperature (fixed at 30°C), indoor temperature, and humidity.  
□ The input is scaled using the scaler.  
□ Power consumption is predicted using rf\_model with the scaled input.  
□ If the predicted power is lower than the current minimum (min\_power), the corresponding temperature and humidity are stored as optimal.

```
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
from sklearn.model_selection import train_test_split
from sklearn.linear_model import LinearRegression
from sklearn.metrics import mean_squared_error
# Simulate some data
np.random.seed(0)
data_size = 1000
temperature = np.random.uniform(15, 35, data_size)
humidity = np.random.uniform(30, 90, data_size)
time_of_day = np.random.uniform(0, 24, data_size)
energy_consumption = 0.5 * temperature + 0.3 * humidity + 0.2 * time_of_day + np.random.normal(0, 2, data_size)
# Create a DataFrame
df = pd.DataFrame({
    'Temperature': temperature,
    'Humidity': humidity,
    'Time_of_Day': time_of_day,
    'Energy_Consumption': energy_consumption
})
```

**E. Error Handling**

If any error occurs during prediction, it is caught and printed.

**F. Output:**

The optimal indoor temperature, humidity, and the minimum predicted power consumption are printed as shown below:

```
Optimal Indoor Temperature: 27.0
Optimal Humidity: 55
Predicted Power Consumption: 5.210413200238203
```

**7. Visualize the data and model performance using plots**

The Python code to visualize the model performance and feature importance is shown below. The code uses Python's pandas, scikit-learn, and matplotlib libraries.

**A. Evaluating the Model**

□ The model's performance is evaluated using Mean Squared Error (MSE), which measures the average squared difference between actual and predicted values. The MSE is printed to the console as shown on the top left hand side of figure 4(a).

Overall, the Python code below shows how to apply linear regression simply to predict a target variable using a variety of features. It also includes visualizations to evaluate the effectiveness of the model and the significance of the features.

```
# Split the data into training and testing sets
X = df[['Temperature', 'Humidity', 'Time_of_Day']]
y = df['Energy_Consumption']
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=0)
# Train a Linear Regression model
model = LinearRegression()
model.fit(X_train, y_train)
# Predict on the test set
y_pred = model.predict(X_test)
# Evaluate the model
mse = mean_squared_error(y_test, y_pred)
print(f'Mean Squared Error: {mse}')
# Plot the predictions vs true values
plt.scatter(y_test, y_pred)
plt.xlabel('True Energy Consumption')
plt.ylabel('Predicted Energy Consumption')
plt.title('True vs Predicted Energy Consumption')
plt.plot([min(y_test), max(y_test)], [min(y_test), max(y_test)], 'r--')
plt.show()
# Plot feature importance (simple linear regression coefficients)
features = X.columns
importance = model.coef_
plt.bar(features, importance)
plt.xlabel('Features')
plt.ylabel('Importance')
plt.title('Feature Importance')
plt.show()
```

## B. Plotting Predictions

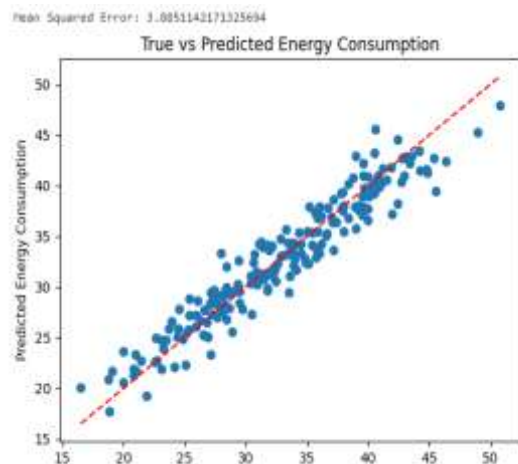


Figure 4(a): Predictions Plot and Mean Squared Error (MSE) of the system

- A scatter plot compares true energy consumption against predicted values.
- A reference line (r--) is drawn to indicate perfect predictions (where predicted values equal true values).

## C. Feature Importance Plot

- The coefficients from the trained model are used to create a bar plot showing the importance of each feature. This gives an idea of how much each feature contributes to the prediction.

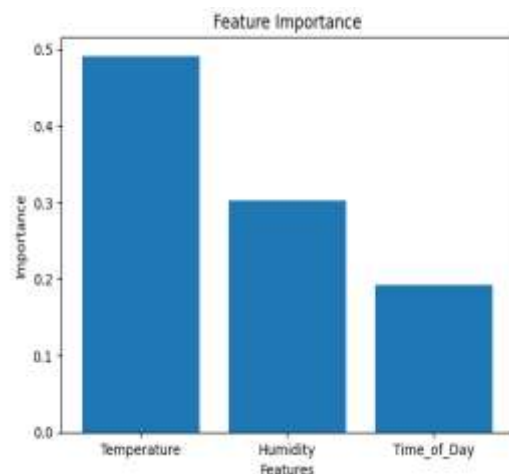


Figure 4(b): Feature Importance Plot

## IV. FURTHER ANALYSIS

The goal is to build a model that can forecast energy consumption depending on variables such as temperature, humidity, and time of day. To fully comprehend the data and the workings of the model, it is imperative to comprehend each graphic.

To visualize different data points and model performance for a deeper comprehension of the air conditioning system, more plots and graphs can be helpful. These are some noteworthy charts and graphs from the developed system.

## 1. Distribution of Features

The distribution of features plot (often shown using histograms or Kernel Density Estimates) displays the spread of the feature values, helping us understand the underlying data patterns for individual variables such as temperature, humidity, or time of day.

### Air Conditioner Example:

- Temperature might range from 15°C to 35°C.
- Humidity could be between 30% and 90%.
- Energy consumption might have a broad range depending on how much energy the system consumes based on these factors.

The python code to plot the distribution of features is shown below, and resulting plots are shown in figure 5.

```
# Plot distribution of each feature
fig, axs = plt.subplots(3, 1, figsize=(10, 15))
axs[0].hist(df['Temperature'], bins=30, color='blue', edgecolor='black')
axs[0].set_title('Distribution of Temperature')
axs[0].set_xlabel('Temperature (°C)')
axs[0].set_ylabel('Frequency')

axs[1].hist(df['Humidity'], bins=30, color='green', edgecolor='black')
axs[1].set_title('Distribution of Humidity')
axs[1].set_xlabel('Humidity (%)')
axs[1].set_ylabel('Frequency')

axs[2].hist(df['Time_of_Day'], bins=24, color='red', edgecolor='black')
axs[2].set_title('Distribution of Time of Day')
axs[2].set_xlabel('Time of Day (Hours)')
axs[2].set_ylabel('Frequency')

plt.tight_layout()
plt.show()
```

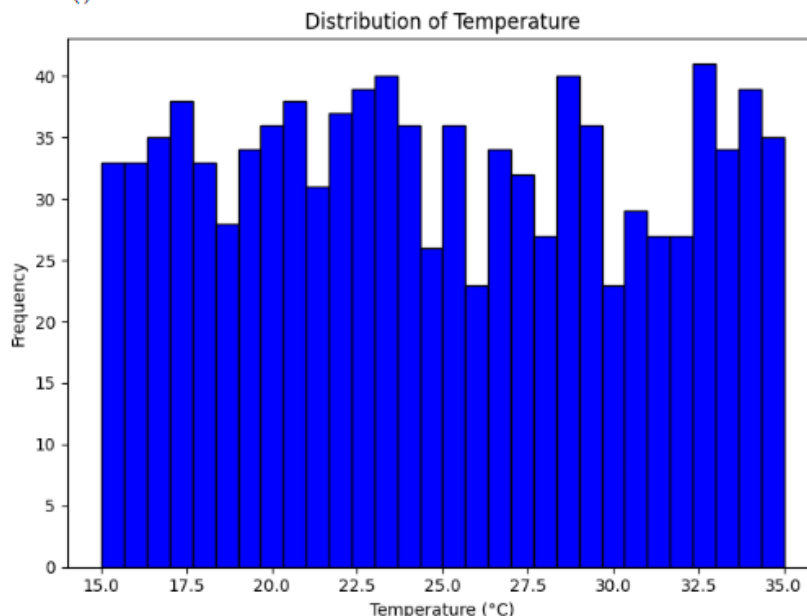


Figure 5(a): Distribution of Temperature




Figure 5(b): Distribution of Humidity




Figure 5(c): Distribution of Time of Day

## 2. Correlation Matrix

The correlation between characteristics and the target variable is displayed in a correlation matrix. Each pair of variables (such as temperature,

humidity, time of day, and energy use) is displayed along with their correlation coefficients. The python code to plot the correlation matrix is shown below:

```
import seaborn as sns
# Compute correlation matrix
corr = df.corr()

# Plot heatmap
plt.figure(figsize=(10, 8))
sns.heatmap(corr, annot=True, cmap='coolwarm', fmt='.2f')
plt.title('Correlation Matrix')
plt.show()
```




Figure 6: The Correlation Matrix

This matrix visually shows how each feature is related to every other feature, with a focus on how temperature, humidity, and time of day are correlated with energy consumption.

A high correlation between temperature and energy consumption means that temperature is a strong predictor, while low correlations indicate less influence.

#### Air Conditioner Example:

□ The correlation matrix helps us see how temperature and humidity influence energy consumption.

□ For example, if temperature and energy consumption have a strong positive connection (around +1), then the energy consumption of the air conditioner increases as the temperature rises.

□ An inverse association is shown by negative correlations (near -1). For instance, at some times of the day, there may be a weak negative association between the time of day and energy use (e.g., lower consumption during nighttime).

### 3. Actual vs. Predicted Values Plot

This plot compares the actual energy consumption values (from test data) with the predicted values (from the model). The goal is to see how closely the model's predictions align with the true values.

#### Air Conditioner Example:

For each test data point, the predicted energy consumption based on temperature, humidity, and time of day is plotted against the actual recorded energy consumption.

If the model is good, the points should fall close to the 45-degree line (where predicted equals actual). Points above the line represent overpredictions (the model overestimated consumption), and points below the line represent underpredictions (the model underestimated consumption).

#### Why It's Important:

The plot provides a visual way to see how well the model is performing across a range of actual values. It reveals if the model is systematically over or underpredicting energy consumption, which can be a sign of bias.

#### Example Insight:

If many points are far from the 45-degree line, it suggests the model may not generalize well or might struggle with specific conditions (e.g., high temperature or humidity).

The python code that plots the Actual vs. Predicted Values is shown below, while figure 7 shows the graph of Actual vs. Predicted Values of the air conditioner.



```
plt.figure(figsize=(10, 6))
plt.scatter(y_test, y_pred, alpha=0.5)
plt.xlabel('True Energy Consumption')
plt.ylabel('Predicted Energy Consumption')
plt.title('Actual vs. Predicted Energy Consumption')
plt.plot([min(y_test), max(y_test)], [min(y_test), max(y_test)], 'r--')
plt.show()
```

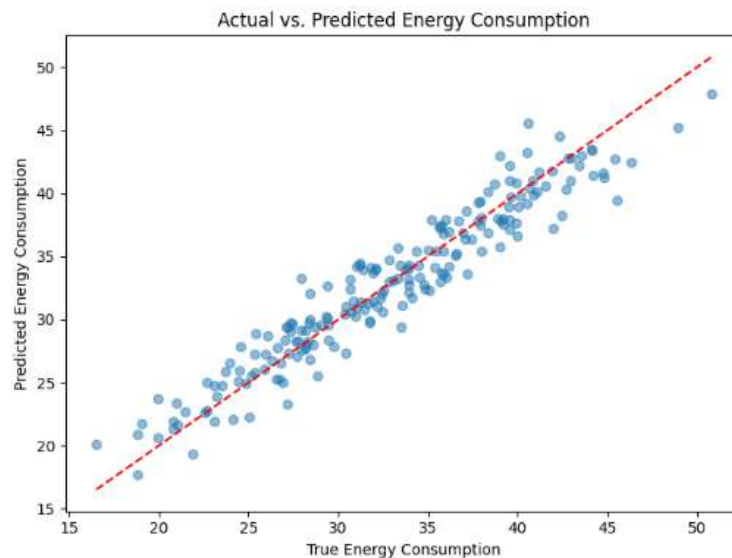


Figure 7: Actual vs. Predicted Values

This graph shows how accurately the model is predicting energy consumption values. Ideally, the points should cluster around the diagonal red line.

#### 4. Residuals Plot

Visualize residuals to check for patterns that might indicate problems with the model. A residual plot visualizes the errors (residuals) between the actual and predicted values. Residuals are calculated as:  
Residual = Actual - Predicted.  
Ideally, residuals should be randomly distributed around zero.

##### Air Conditioner Example:

□ The plot shows how much the model's predictions deviate from actual energy consumption.

□ If the residuals are randomly scattered, the model likely performs well. However, patterns in the residuals (e.g., clustering or trends) suggest the model is missing certain aspects of the data (e.g., non-linear relationships).

##### Why It's Important:

□ Randomly distributed residuals around zero indicate that the model is unbiased and captures the data well as shown in figure 15.  
□ Systematic patterns in the residuals can point to issues such as overfitting, underfitting, or failure to capture important trends.

##### Example Insight:

If residuals are consistently higher or lower for certain ranges of energy consumption, the model may be over or underpredicting in specific situations (e.g., during very hot or cold periods).

```
residuals = y_test - y_pred

plt.figure(figsize=(10, 6))
plt.scatter(y_pred, residuals, alpha=0.5)
plt.axhline(y=0, color='r', linestyle='--')
plt.xlabel('Predicted Energy Consumption')
plt.ylabel('Residuals')
plt.title('Residuals vs. Predicted Values')
plt.show()
```




Figure 8: Residuals Plot

## 5. Time Series Analysis (Energy Consumption Plot)

The time series plot shows how energy consumption changes over time, providing insights into trends and patterns, such as daily cycles, seasonal variations, or long-term increases or decreases in energy usage. This plot shows the fluctuation of energy consumption over time, helping us detect trends, peaks, or cyclical patterns.

### Air Conditioner Example:

□ Energy consumption might be higher during the day when temperatures are higher and people are more active, and lower at night when the cooling demand is lower.

□ Over a longer period, you might observe patterns, such as energy consumption spikes during hot summer months and decreases during cooler seasons.

### Why It's Important:

Time series analysis helps in understanding trends (overall direction), seasonality (repeating cycles like daily or weekly), and random noise.

Identifying these components is important for both understanding the data and making accurate future predictions using time series forecasting models.




Figure 9: Energy Consumption Plot

#### Example Insight:

The energy consumption plot could reveal that energy usage peaks around midday when temperatures are highest, or it might show a weekly pattern where consumption is lower on weekends. This plot shows the fluctuation of energy

consumption over time, helping us detect trends, peaks, or cyclical patterns.

The python code below combines all the plots; from the distribution of features to energy consumption plot, to create a comprehensive set of visualizations, arranged in a grid for easy comparison.

```
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
import seaborn as sns
from sklearn.model_selection import train_test_split
from sklearn.linear_model import LinearRegression
from sklearn.metrics import mean_squared_error

# Simulate some data
np.random.seed(0)
data_size = 1000
temperature = np.random.uniform(15, 35, data_size)
humidity = np.random.uniform(30, 90, data_size)
time_of_day = np.random.uniform(0, 24, data_size)
energy_consumption = 0.5 * temperature + 0.3 * humidity + 0.2 * time_of_day + np.random.normal(0, 2, data_size)

# Create a DataFrame
df = pd.DataFrame({
    'Temperature': temperature,
    'Humidity': humidity,
    'Time_of_Day': time_of_day,
    'Energy_Consumption': energy_consumption
})

# Split the data into training and testing sets
X = df[['Temperature', 'Humidity', 'Time_of_Day']]
y = df['Energy_Consumption']
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=0)

# Train a Linear Regression model
model = LinearRegression()
model.fit(X_train, y_train)
```

```
# Predict on the test set
y_pred = model.predict(X_test)

# Evaluation metrics
mse = mean_squared_error(y_test, y_pred)

# Create plots
fig, axs = plt.subplots(3, 2, figsize=(15, 18))

# Distribution of Features
axs[0, 0].hist(df['Temperature'], bins=30, color='blue', edgecolor='black')
axs[0, 0].set_title('Distribution of Temperature')
axs[0, 0].set_xlabel('Temperature (°C)')
axs[0, 0].set_ylabel('Frequency')

axs[0, 1].hist(df['Humidity'], bins=30, color='green', edgecolor='black')
axs[0, 1].set_title('Distribution of Humidity')
axs[0, 1].set_xlabel('Humidity (%)')
axs[0, 1].set_ylabel('Frequency')

axs[1, 0].hist(df['Time_of_Day'], bins=24, color='red', edgecolor='black')
axs[1, 0].set_title('Distribution of Time of Day')
axs[1, 0].set_xlabel('Time of Day (Hours)')
axs[1, 0].set_ylabel('Frequency')

# Correlation Matrix
corr = df.corr()
sns.heatmap(corr, annot=True, cmap='coolwarm', fmt='.2f', ax=axs[1, 1])
axs[1, 1].set_title('Correlation Matrix')

# Actual vs. Predicted Values
axs[2, 0].scatter(y_test, y_pred, alpha=0.5)
axs[2, 0].set_xlabel('True Energy Consumption')
axs[2, 0].set_ylabel('Predicted Energy Consumption')
axs[2, 0].set_title('Actual vs. Predicted Energy Consumption')
axs[2, 0].plot([min(y_test), max(y_test)], [min(y_test), max(y_test)], 'r--')

# Residuals Plot
residuals = y_test - y_pred
axs[2, 1].scatter(y_pred, residuals, alpha=0.5)
axs[2, 1].axhline(y=0, color='r', linestyle='--')
axs[2, 1].set_xlabel('Predicted Energy Consumption')
axs[2, 1].set_ylabel('Residuals')
axs[2, 1].set_title('Residuals vs. Predicted Values')

plt.tight_layout()
plt.show()

# Time Series Analysis (if the DataFrame is indexed by time or similar)
# For demonstration purposes, we're creating a time index
df.index = pd.date_range(start='2024-01-01', periods=data_size, freq='H')

plt.figure(figsize=(10, 6))
plt.plot(df.index, df['Energy_Consumption'], label='Energy Consumption')
plt.xlabel('Time')
plt.ylabel('Energy Consumption')
plt.title('Energy Consumption Over Time')
plt.show()
```

This script gives us a comprehensive view of the data and model performance.

#### Summary of the plots:

- **Distribution of Features:** Plots the distribution of temperature, humidity, and time of day.
- **Correlation Matrix:** Shows the correlation between different features and the target variable.
- **Actual vs. Predicted Values:** Compares the model's predicted energy consumption with the actual values.
- **Residuals Plot:** Plots the residuals to help identify any patterns or issues with the model.
- **Time Series Analysis:** Plots energy consumption over time if you have a time-based index.

### V. ADVANCED GRAPHS AND PLOTS OF THE HVAC

Here we explained other graphs and plots to give a comprehensive understanding of our model and data. These graphs and plots are:

#### 1. Pair Plot

A pair plot visualizes relationships between all pairs of features, as well as the distribution of individual features.

□ **Usage:** This is useful for exploring correlations and patterns between multiple variables. You can see whether two features have a linear, exponential, or no relationship at all.

□ **How to Read:** Diagonal plots show the distribution (histograms or kernel density estimates, KDE) of individual features. The off-diagonal plots show the scatter plots of feature pairs.

□ **Example:** As shown in figure 10, as Temperature increases, Energy Consumption increases as well. Such insights help in feature engineering and model selection.

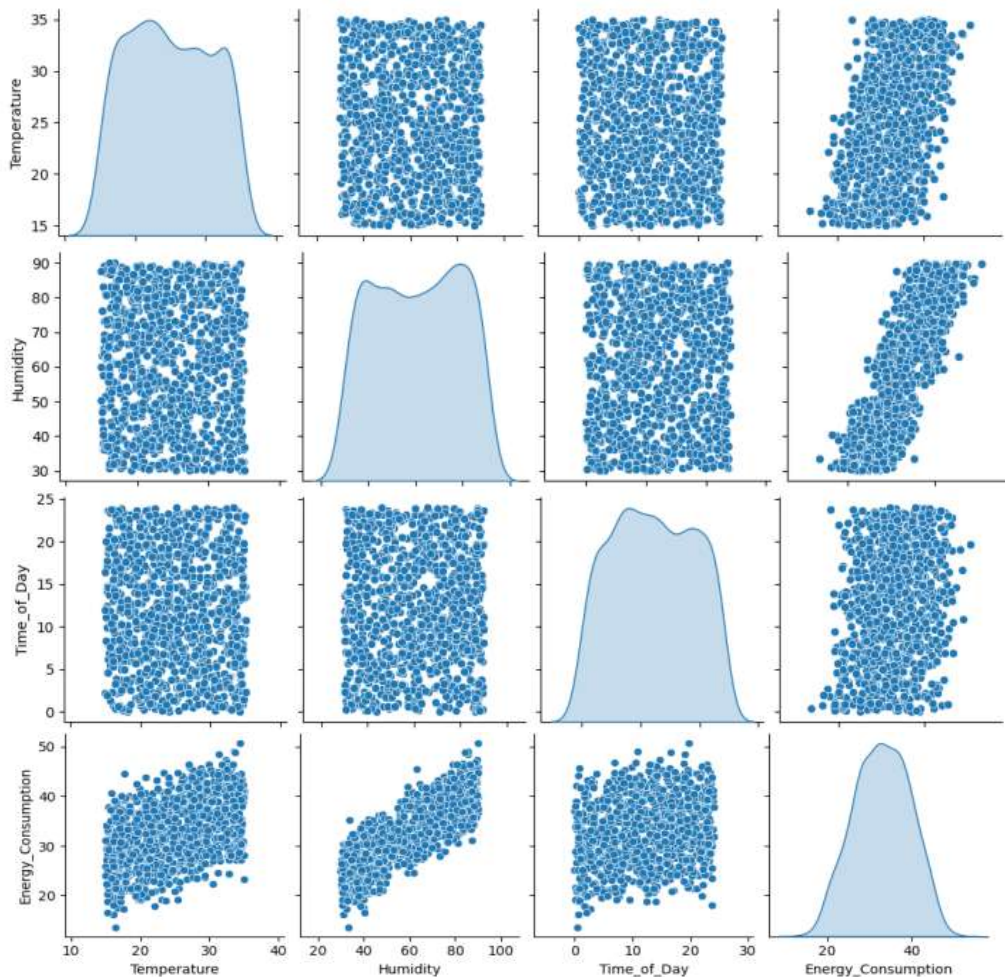


Figure 10: Pair plot of the HVAC



## 2. Feature Importance (Random Forest):

Feature importance plots display how much each feature contributes to predictions in a Random Forest model.

**Usage:** This is used to understand which features have the most impact on the target variable. Random Forest uses multiple decision trees, and the importance score reflects how often a feature is used to split nodes in trees.

**How to Read:** Features are ranked by importance (calculated as the reduction in error caused by each feature). Features with higher bars are more influential in predicting the target. In our model, humidity is influential as shown in figure 11.

**Example:** If Humidity is the most important feature, it means that variations in humidity are responsible for the largest changes in energy consumption.




Figure 11: Random forest plot

## 3. SHAP Summary Plot

SHAP (SHapley Additive exPlanations) values quantify how much each feature contributes to individual predictions.

**Usage:** SHAP values explain how each feature impacts a model's prediction for a specific data point. It's an advanced, model-agnostic technique to interpret complex models like Random Forest, Gradient Boosting, etc.

**How to Read:** The plot summarizes the overall feature importance (left-to-right: most important

features on the top). Each dot represents a SHAP value for a feature, where colors (blue to red) represent the feature's value (low to high).

**Example:** For figure 12, one can see the effect of both low and high values of the feature on the model's prediction. If the dots are spread widely along the X-axis, it indicates that the feature can have a significant positive or negative effect on the prediction.




Figure 12: SHAP Summary Plot

#### 4. Learning Curves

A learning curve shows how a model's performance evolves as it is trained on larger datasets.

**Usage:** It is useful for diagnosing underfitting or overfitting in a model. It tells us whether the model needs more data or if it's already learning well.

**How to Read:** The X-axis shows the number of training examples, and the Y-axis shows the performance metric (accuracy, error, etc.). Two

lines are typically plotted: one for the training set and one for the validation set. If both lines converge and flatten out, the model is likely well-trained.

**Example:** If the training score is high but the validation score is low, the model is likely overfitting. If both scores are low, the model may be underfitting and too simple to capture the relationships in the data.

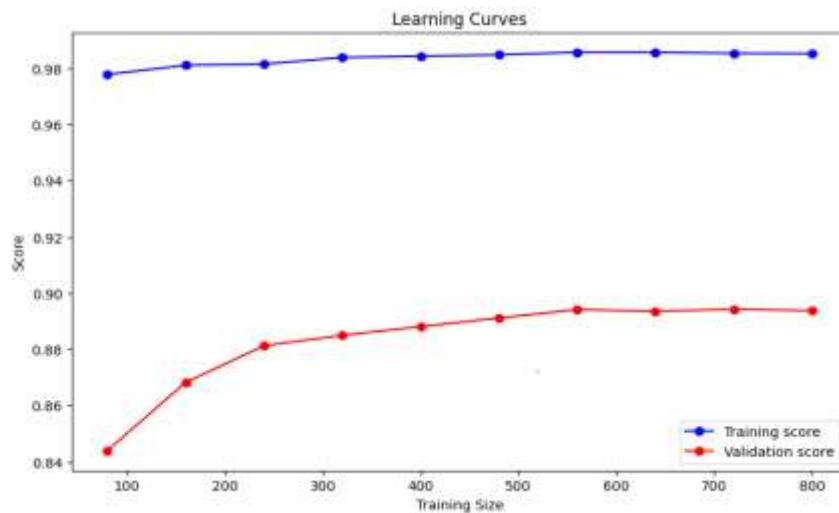


Figure 13: Learning curves

#### 5. Error Distribution (Residuals)

This plot shows the distribution of residuals, which are the differences between the actual and predicted values of the target variable.

**Usage:** Residual plots are used to diagnose problems with a model's predictions. Ideally, the residuals should be centered around zero and normally distributed, which indicates that the model is making unbiased predictions.

**How to Read:** The X-axis shows the residuals, and the Y-axis shows the frequency of those residuals. A bell-shaped curve centered around zero indicates good model performance.

**Example:** If the residuals are skewed or not centered around zero, it suggests that the model is systematically under or overpredicting in certain cases, meaning there might be bias in the model.




Figure 14: The Error Distribution (Residuals) Plot

## 6. Time Series Decomposition

A time series decomposition separates a time series into three components: trend, seasonality, and residuals.

**Usage:** It helps analyze time series data to understand the underlying patterns like long-term trends, periodic fluctuations, and irregular variations.

**How to Read:**

**Trend:** Shows the long-term direction of the data.

**Seasonal:** Captures regular, repeating patterns (daily, monthly, yearly, etc.).

**Residual:** Represents random noise or unexplained fluctuations in the data.

**Example:** For energy consumption data, you might see a daily or weekly seasonal component, a long-term trend of increasing or decreasing consumption, and some residual random noise.




Figure 15: Time Series Decomposition Plot

This combination of visualizations gives us deep insights into the model's performance, feature impacts, and the overall patterns in your data.

## VI: MACHINE LEARNING MODEL AND TRADITIONAL METHOD

When comparing the results obtained from the machine learning-based model for predicting air conditioner energy consumption to traditional methods, several key distinctions can be made in terms of accuracy, efficiency, flexibility, and adaptability.

Machine learning models outperform traditional methods in predicting air conditioner energy consumption across several key dimensions. Here is a breakdown of the comparison:

□ **Accuracy:** The machine learning models produce more accurate predictions by capturing complex, non-linear relationships between temperature, humidity, and energy consumption.

□ **Adaptability:** These models are flexible and can be easily retrained on new data, making them more scalable and suitable for different environments.

□ **Handling Complexity:** Advanced models handle interactions between features and capture intricate patterns like seasonality, which traditional models often miss.

□ **Future Potential:** Machine learning models can be enhanced further by integrating additional real-time data, such as occupancy or real-time air quality sensors, to improve predictive accuracy.

Ultimately, while traditional methods may still be useful in simple, well-understood environments, the machine learning approach provides far greater power, flexibility, and accuracy for modern, data-driven energy management systems. This makes it the preferred solution in dynamic and complex real-world settings.

## VII: CONCLUSION AND FURTHER WORK

This study underscores the transformative potential of machine learning in redefining the landscape of next-generation air conditioning systems. By integrating advanced algorithms with system modeling, we have demonstrated how predictive analytics can significantly enhance energy efficiency, operational reliability, and adaptability to dynamic environmental conditions. For both residential and commercial applications, the suggested framework, which makes use of time series analysis and machine learning techniques, offers precise energy usage projections based on variables like temperature, humidity, and time of day. The results open the door to more environmentally friendly building management techniques by demonstrating the viability of machine learning as a key enabler of intelligent HVAC systems. These developments significantly meet the global demand for ecologically responsible technologies in addition to cutting down on energy waste and enhancing financial results. This work could be expanded in the future by investigating the incorporation of diverse datasets from different climates. This study lays the groundwork for a new age of intelligent, effective, and sustainable air conditioning solutions by bridging the gap between technological innovation and real-world application.

### REFERENCES:

- [1]. Aggarwal, C. C., & Reddy, C. K. (2014). *Data Clustering: Algorithms and Applications*. CRC Press.
- [2]. Aghajanzadeh, A., Baniassadi, A., & Haghighat, F. (2019). Sensor data for predicting HVAC energy consumption using machine learning models. *Energy Performance Journal*, 19(2), 83-95.
- [3]. Agrawal, R., Imielinski, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data* (pp. 207-216).
- [4]. Bertsekas, D. P., & Tsitsiklis, J. N. (1996). *Neuro-dynamic programming*. Athena Scientific.
- [5]. Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
- [6]. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 15.
- [7]. Chaudhary, V., Singh, R., & Kumar, A. (2018). Energy efficiency in HVAC systems: A review of control mechanisms. *Journal of Sustainable Building Technology*, 22(3), 101-114.
- [8]. Chicco, D., & Jurman, G. (2021). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 22(1), 6.
- [9]. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [10]. Guo, Y., Chen, L., & Zhu, Z. (2021). Machine learning-based fault detection and predictive maintenance in air conditioning systems. *HVAC Research*, 27(1), 34-45.
- [11]. Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.). Springer.
- [12]. International Energy Agency. (2019). *The future of cooling: Opportunities for energy-efficient air conditioning*. Paris: IEA Publications.
- [13]. Jain, A. K. (2010). Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*, 31(8), 651-666.
- [14]. Jain, P., & Singh, K. (2020). Predictive maintenance in HVAC systems using machine learning. *Journal of Sustainable Engineering*, 12(3), 75-81.
- [15]. Johnson, M., Davis, T., & Liu, X. (2019). Challenges and limitations of traditional HVAC control systems. *IEEE Transactions on Control Systems*, 36(5), 57-64.
- [16]. K.A. Akpado, P. N. Nwankwo, et al. (2018). *International Journal of Engineering and Applied Sciences (IJEAS)*. ISSN: 2394-3661, Volume-5, Issue-8, August 2018.
- [17]. Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). "Reinforcement learning: A survey." *Journal of Artificial Intelligence Research*, 4, 237-285.
- [18]. Kim, H., Choi, Y., & Park, D. (2021). Integrating AI into HVAC systems: Trends and case studies. *Energy and Buildings*, 219, 110796.
- [19]. Kotsiantis, S. B. (2019). Supervised machine learning: A review of classification techniques. *Frontiers in Artificial Intelligence and Applications*, 272, 25-50.

- [20]. Lee, S., & Park, J. (2021). Energy-saving strategies for next-generation air conditioning systems. *Building Energy Research*, 28(4), 299-310.
- [21]. Liang, Y., & Chen, Y. (2022). Case studies of machine learning-driven HVAC systems in commercial buildings. *Energy Efficiency & Buildings*, 17(2), 229-240.
- [22]. Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press.
- [23]. Messung, 2023. How Artificial Intelligence is Revolutionizing HVAC Systems (on July 17, 2023).
- [24]. Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). "Human-level control through deep reinforcement learning." *Nature*, 518(7540), 529-533.
- [25]. Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press.
- [26]. Nguyen, T., & Das, A. (2019). Machine learning applications in real-time control of HVAC systems. *Applied Energy*, 238, 1524-1531.
- [27]. Pereira, F., & Silva, R. (2022). Deep learning applications in HVAC system optimization. *Journal of Energy Informatics*, 5(1), 10-21.
- [28]. Prince. N Nwankwo, K. A. Akpado, et al. (2021). *International Journal for Research in Applied Science & Engineering Technology (IJRASET)*. Volume 9 Issue X Oct 2021. <https://doi.org/10.22214/ijraset.2021.38665>
- [29]. Russell, S., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson Education.
- [30]. Schulman, J., Wolski, F., Dhariwal, P., et al. (2017). "Proximal policy optimization algorithms." *arXiv preprint, arXiv:1707.06347*.
- [31]. Silver, D., Lever, G., Heess, N., et al. (2014). "Deterministic policy gradient algorithms." *arXiv preprint, arXiv:1404.0039*.
- [32]. Smith, J. (2021). Global energy consumption trends and their impact on sustainability. *Energy Journal*, 45(2), 120-133.
- [33]. [Sun, J., & Zhang, W. (2020). Dynamic control strategies for air conditioning in complex environments. *Energy Efficiency Journal*, 16(7), 204-217.
- [34]. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.
- [35]. Szepesvári, C. (2010). "Algorithms for Reinforcement Learning." *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 4(1), 1-103.
- [36]. Szepesvári, C. (2010). "Algorithms for Reinforcement Learning." *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 4(1), 1-103.
- [37]. U.S. Department of Energy. (2020). *Energy efficiency in buildings: HVAC systems and energy use*. Retrieved from <https://www.energy.gov>
- [38]. Watkins, C. J. C. H., & Dayan, P. (1992). "Q-learning." *Machine Learning*, 8(3-4), 279-292.
- [39]. Wang, K., Zhang, Y., & Liu, X. (2022). Optimizing HVAC systems for energy efficiency and occupant comfort. *Renewable Energy Reports*, 14(3), 44-56.
- [40]. Zhou, H., Wang, Y., & Liu, Z. (2020). The environmental impact of air conditioning in urban environments. *Journal of Energy Management*, 33(1), 88-97.
- [41]. Zhou, T., & Li, J. (2020). AI-driven optimization in HVAC systems. *Journal of Smart Systems*, 31(2), 55-67.
- [42]. Zhou, Z.-H., & Mao, C. X. (2020). On unsupervised learning: When, how, and why should we care? *National Science Review*, 7(1), 10-13.
- [43]. van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9, 2579-2605.