energies

MDPI

Article

EECO: An AI-Based Algorithm for Energy-Efficient Comfort Optimisation

Giacomo Segala [1,2,3]*, Roberto Doriguzzi-Corin [2], Claudio Peroni [1], Matteo Gerola [1] and Domenico Siracusa [2]

[1] Energenius s.r.l., 38068 Rovereto, Italy; claudio.peroni@mapsgroup.it (C.P.); matteo.gerola@mapsgroup.it (M.G.)
[2] Fondazione Bruno Kessler (FBK), 38123 Trento, Italy; rdoriguzzi@fbk.eu (R.D.-C.); dsiracusa@fbk.eu (D.S.)
[3] Department of Information Engineering and Computer Science (DISI), University of Trento, 38123 Trento, Italy
* Correspondence: giacomo.segala@mapsgroup.it

Abstract: Environmental comfort takes a central role in the well-being and health of people. In modern industrial, commercial, and residential buildings, passive energy sources (such as solar irradiance and heat exchangers) and heating, ventilation, and air conditioning (HVAC) systems are usually employed to achieve the required comfort. While passive strategies can effectively enhance the livability of indoor spaces with limited or no energy cost, active strategies based on HVAC machines are often preferred to have direct control over the environment. Commonly, the working parameters of such machines are manually tuned to a fixed set point during working hours or throughout the whole day, leading to inefficiencies in terms of comfort and energy consumption. Albeit effective, previous works that tackle the comfort–energy tradeoff are tailored to the specific environment under study (in terms of geometry, characteristics of the building, etc.) and thus cannot be applied on a large industrial scale. We address the problem from a different angle and propose an adaptive and practical solution for comfort optimisation. It does not require the intervention of expert personnel or any customisations around the environment while it implicitly analyses the influence of different agents (e.g., passive phenomena) on the monitored parameters. A convolutional neural network (CNN) predicts the long-term impact on thermal comfort and energy consumption of a range of possible actuation strategies for the HVAC system. The decision on the best HVAC settings is taken by choosing the combination of ON/OFF and set point (SP), which optimises thermal comfort and, at the same time, minimises energy consumption. We validate our solution in a real-world scenario and through software simulations, providing a performance comparison against the fixed set point strategy and a greedy approach. The evaluation results show that our solution achieves the desired thermal comfort while reducing the energy footprint by up to approximately 16% in a real environment.

Keywords: thermal comfort; energy efficiency; automated HVAC configuration; deep learning

check for updates

Citation: Segala, G.; Doriguzzi-Corin, R.; Peroni, C.; Gerola, M.; Siracusa, D. EECO: An AI-Based Algorithm for Energy-Efficient Comfort Optimisation. *Energies* **2023**, *16*, 7334. https://doi.org/10.3390/en16217334

Academic Editor: Tullio De Rubeis

Received: 14 September 2023
Revised: 17 October 2023
Accepted: 27 October 2023
Published: 29 October 2023

1. Introduction

Governments, regulatory agencies, and public bodies have been promoting policies and measures for healthy and energy-efficient buildings, issuing directives such as the Directive (2018/844) [1] developed by the European Parliament. Despite these efforts, a significant amount of energy is often wasted in industrial, commercial, and residential buildings, causing uncomfortable conditions for occupants. Typical examples of such inefficiencies are retail stores, in which energy managers struggle to find a good balance in the tradeoff between optimisation of comfort and minimisation of energy consumption by HVAC systems. Indeed, managers need customers to feel comfortable while shopping at any time of the year and in all environmental conditions. Furthermore, they seek to gain important indoor air quality (IAQ) certifications [2,3] in order to distinguish themselves against their competitors.

Currently, HVAC systems are often managed directly on site in buildings, manually operating on thermostats to regulate thermal comfort with not enough attention to energy efficiency. For instance, in commercial buildings, HVAC systems are often left active at the end of the day, thus continuing to work and consuming unnecessary energy during the night when shops are closed. Another common issue in the vast majority of buildings is that whatever HVAC configuration in terms of ON/OFF and set point (i.e., desired target temperature) is configured in the early morning is usually left unchanged throughout the day. However, according to outdoor weather and its impact on the indoor environments, it might be convenient applying different HVAC settings during the day (e.g., switching the devices off). In such scenarios, automated systems capable of continuously optimising HVAC devices over time have drawn the interest of managers. Indeed, these solutions aim to address the aforementioned challenges, efficiently meeting the comfort requirements throughout the day while reducing the energy footprint. This takes a central role in scenarios where a large amount of geographically distributed and physically heterogeneous sites are managed as each of them might require a different control strategy.

Multiple works from the literature deal with the topic of comfort maximisation jointly with energy consumption minimisation in buildings. Despite the progress in this field resulting in innovative solutions, e.g., based on advanced passive strategies [4,5] or model predictive control (MPC) [6–9], their major drawback is the limited scalability. This refers to the capability to replicate and automate a specific approach across different environments regardless of their physical characteristics. These solutions often require a comprehensive analysis of each building and the machinery installed therein to define tailored physical or mathematical models (e.g., [4,5,8,10–12]), which typically need manual updates over time. In parallel, other solutions use complex building-related information (e.g., [6,7,13,14]) and might require customisations within the monitored environment. While designing smart and adaptive solutions using data collected from Internet of Things (IoT) sensors is essential to optimise HVAC systems [15], it is crucial to ensure that their deployment and replicability involve automated operations. This makes them attractive from a business perspective, especially for managers who need to control tens or hundreds of buildings. Finally, different works provide theoretical analysis techniques for comfort and energy optimisation but lack real-world validation [16–18]. In this regard, Ngarambe et al. [19] underline that experimental studies demonstrating the benefit of artificial intelligence (AI) control strategies (e.g., MPC approaches) in real environments take a central role.

In this work, we propose a novel solution called energy-efficient comfort optimisation (EECO) based on deep learning (DL) to regulate HVAC systems in an automated manner. It does not require any intervention of expert personnel or prior information of buildings (e.g., installed HVAC devices, layout, and materials) as it works on real data collected from the environment. In this regard, our aim is to analyse how the different agents, including passive phenomena, impact the parameters within the environment through the collected data. This allows us to provide an adaptive solution for the monitored environment that implicitly considers the influence of the different sources. From an applicability perspective, the proposed solution holds the potential for being applied in any building equipped with a control system capable of gathering environmental and energy consumption data and interfacing with local HVAC devices. Basically, after an initial configuration of the main parameters (e.g., the comfort interval throughout the day and some parameters of the comfort model), the proposed solution can effectively work just after its deployment, and it keeps up to date independently over time, resulting in an automated and practical solution. The objective of HVAC optimisation is to guarantee the comfort requirements, at least during opening hours, and then balance both thermal comfort and energy consumption concerns. Indoor comfort is modelled by means of predicted mean vote (PMV) [20,21], a thermal comfort index referenced by different indoor comfort standards all over the world, including European Standard EN 16798. A shallow 1D convolutional neural network (CNN) is used as DL architecture to predict the short-term evolution of future indoor environmental parameters (i.e., temperature, humidity, and carbon dioxide (CO_2)) and the energy consumption of the HVAC system. The idea behind the DL model is to predict the environmental and energy impact of a set of possible device configurations (ON/OFF and set point) for m upcoming time periods. Basically, a tree of possible actuation strategies that keeps track of the environment evolution in the next future based on past (real or predicted) conditions is generated. Each branch of the resulting tree is then evaluated to select the strategy that maintains the best expected comfort at minimal energy cost.

Our work can contribute to reducing the carbon footprint of buildings caused by HVAC systems, improving the comfort conditions to occupants and saving on operating costs required to control thermal comfort. The designed approach has been tested during the summer and winter periods in a real environment of a small production plant belonging to a large retail company in northern Italy. Furthermore, an additional analysis based on software simulations is proposed.

The main contributions of this work are the following:

- A practical solution, with no prior information of the local environment (e.g., installed HVAC devices and building features) or need for customisation or intervention of expert personnel, capable of selecting an efficient HVAC configuration in terms of ON/OFF and set point that aims to guarantee the given thermal comfort while minimising energy consumption.
- An adaptive and continuous update of the actuations through short-term decisions based on long-term predictions of the environment.
- A comparison analysis in terms of tradeoff between thermal comfort and energy consumption with the manual approach, which sets a static set point temperature throughout the day, and a greedy PMV-based solution, which configures the HVAC devices according to the current environmental conditions.

The remainder of this paper is organised as follows. Section 2 describes the relevant literature. Section 3 provides background regarding the neural network architecture and predicted mean vote (PMV) index. Section 4 presents the proposed methodology. Section 5 illustrates the experimental setup, while Section 6 describes and presents the results. Section 7 discusses the limitations of the proposed solution. Finally, conclusions are provided in Section 8.

2. State of the Art

In the recent scientific literature, a number of research works have been proposed to achieve thermal comfort, trying to solve the problem of the tradeoff between comfort maximisation and energy minimisation from different perspectives. In this regard, different works [16–18] tackle the problem through Pareto analysis. This approach provides a set of possible tradeoffs between comfort and energy consumption, each of which might be a feasible solution for the deployment. However, the mentioned research works provide static analysis with a limited number of software simulations and do not consider any prediction in the future for proactive decision making. Additionally, they model the objective functions through ad hoc mathematical models for the specific environment under evaluation, thus limiting their applicability across multiple sites. Finally, while calculating the Pareto front can be useful, a proper strategy is necessary to select a single configuration that guarantees good comfort at a low cost, and this is missing in these works.

Other research initiatives have tackled the problem from another perspective: they physically model the buildings through simulation software to provide either simulated environments for analysis or generate a large amount of data to train AI models [7,10–12]. For instance, Gao et al. [12] propose a DL solution based on reinforcement learning validated by means of a simulated building thermal environment and an HVAC system; a large amount of hourly simulated data are used to train their AI models. Another solution based on reinforcement learning is presented by Valladares et al. [11]. In their study, a reinforcement learning model is first trained with 10 years of simulated data, following a similar approach to Gao et al. [12], before being deployed in real environments to evaluate the performance.

By means of training data collected over a large time interval, they achieve a balance among indoor comfort, air quality, and the energy consumption of the air conditioning and ventilation systems. Unlike the research works proposed by Valladares et al. [11] and Gao et al. [12], a different solution based on model predictive control (MPC) is proposed by Ascione et al. [10]. However, even in this case, it relies on simulation-based physical models to optimise the hourly set point temperature for the next 24 h. Furthermore, Jing et al. [7] propose a simple PMV-based approach to keep the environment within the comfort level and overcome the typical temperature-based mechanism. Despite improvements in terms of daily energy savings, the proposed solution only focuses on thermal comfort, with no attention for a tradeoff between PMV index and energy consumption in the HVAC control strategy. Additionally, the proposed solution is validated and analysed using simulation models, without any validation in real environments. Finally, other works rely on advanced passive strategies. For instance, Liu et al. [4] analyse the applicability and effectiveness of these technologies in residential buildings through physical models, resulting in significant energy savings. Additionally, de Araujo Passos et al. [5], in their study, define a mathematical model to optimise a novel HVAC system by relying on passive energy sources (e.g., solar irradiance and heat exchangers) as much as possible. Significant energy-saving results have been achieved, demonstrating that over half of the energy demand can be met passively.

All the research works described above are based on building modelling. In addition to a significant manual effort to model various aspects of the environment (e.g., layout, materials, location, and installed HVAC machinery), this approach provides clear limitations. Firstly, detailed modelling of individual buildings impacts scalability, limiting their replicability across multiple sites with limited effort. Secondly, the usage of simulated data might hinder a faithful replication of real-world environments, which can be affected by unexpected events (e.g., windows or doors being opened or rapid increases in occupancy). In this regard, the validation of AI-control solutions in real environments is fundamental to demonstrate their benefits in the intelligent control of HVAC systems [19].

Other approaches that do not rely on physical models of buildings are proposed in the literature. Chen et al. [8] propose an MPC solution by modelling the building through mathematical models. However, complex building-specific information is used (e.g., conduction/convection coefficient, wall thickness, air mass flow rate, etc.). It is worth noting that, in this work, feedback from occupants takes a central role to adapt the thermal comfort based on personal perception, resulting in improved comfort outcomes. In this regard, other studies based on MPC delve into how personal preferences affect the optimisation of energy consumption and the well-being of occupants [9]. To address the limitations of physical-based models, as per our goal, Manjarres et al. [13] introduce a framework aimed at minimising energy consumption while ensuring indoor temperatures remain within predefined ranges. The proposed framework outlines an optimal schedule for HVAC ON/OFF and mechanical ventilation (MV) operation for the next 24 h. However, it requires the installation of specific sensors (e.g., in the outlet conduct of the air handling unit within the HVAC device). Additionally, it primarily considers indoor temperature rather than thermal comfort (e.g., PMV index) and does not account for updates to the operating schedule throughout the day in response to potential environmental changes. Similarly, Yang et al. [6] propose an MPC approach designed to overcome the constraints associated with physical models by integrating AI. Additionally, they introduce an update mechanism over time to capture any possible environmental change. However, their solution requires customisations within the environments in terms of advanced sensors (e.g., combined temperature–humidity–pressure–lux (THPL) sensors) to be installed in specific locations as well as detailed information regarding chilled water of HVAC devices. This bounds their approach to the specific environment being evaluated. Another approach that effectively keeps up with with environmental changes but includes complex building-related information is proposed by Martell et al. [14]. Indeed, the authors propose a multi-objective control architecture to estimate optimal set points where the computed

Pareto front is updated hourly, thus selecting optimal temperature set points for each hour of the day. Despite the update mechanism, even in this case, complex parameters closely tied to the evaluated environment are considered. For instance, the authors use heat gains resulting from different natural phenomena (e.g., convection, ventilation, and infiltration) to model the indoor temperature behaviour, which might be different across different sites.

In summary, existing solutions for comfort optimisation present various limitations that might impact their applicability to real-world scenarios. Indeed, they provide theoretical analysis with no HVAC strategy selection and real-world validation [16–18], rely on tailored physical (or mathematical) models [4,5,7,10–12], or use complex information of the local environment [6,8,9,13,14]. Furthermore, no clear update mechanisms of HVAC settings over time are taken into account, except in rare cases [6,10,14]. Our solution aims to overcome the above limitations. On one hand, it does not require preliminary analysis to define physical or mathematical models of the environment or gather building-specific information. Instead, it adapts to the monitored environment by learning from the collected data. On the other hand, we rely on long-term predictions to make short-term decisions and continuously select the actuation strategy that optimises comfort and minimises energy consumption over time.

3. Background

The solution proposed in this paper is grounded on a one-dimensional convolutional neural network (CNN) [22] and on the predicted mean vote (PMV), a thermal comfort index introduced by Fanger et al. [23].

3.1. Dataset

In Table 1, we report an overview of the variables used by the proposed solution. They can be grouped into

- HVAC parameters (i.e., ON/OFF, SP, fan speed, operating mode), which are collected for each device installed in the environment through Modbus protocol.
- Outdoor environment parameters (i.e., temperature, humidity), which are collected from the OpenWeatherMap platform [24] through APIs.
- Indoor environment parameters (i.e., temperature, humidity, CO_2, energy consumption). Indoor temperature is sensed by the thermostats, while humidity and CO_2 are collected through an IoT sensor installed in the environment. Finally, energy consumption is collected from a smart energy meter. All these variables are collected through Modbus protocol.
- Supporting variables (i.e., day of the week, hour of the day).

Table 1. Overview of the input variables.

Variable	Type	Description
ON/OFF	Categorical	ON/OFF of HVAC devices.
Set Point	Number	Set point temperature [°C] of HVAC devices.
Fan Speed	Categorical	Fan speed of HVAC devices (1 = low, 2 = high, 3 = very high).
Operating Mode	Categorical	Operating mode of HVAC devices, i.e., cooling or heating.
Outdoor Temperature	Number	Outdoor temperature [°C] collected from the OpenWeatherMap platform [24].

Table 1. *Cont.*

Variable	Type	Description
Outdoor Humidity	Number	Outdoor humidity [%] collected from the OpenWeatherMap platform [24].
Indoor Temperature	Number	Indoor temperature [°C] sensed by the thermostats, which feed their readings to the installed HVAC devices. The mean value is used within the algorithm.
Indoor Humidity	Number	Indoor humidity [%] collected through an IoT sensor installed in the environment.
Indoor CO_2	Number	Indoor CO_2 [ppm] collected through an IoT sensor installed in the environment.
Energy Consumption	Number	Energy consumption [kWh] due to HVAC devices collected from smart energy meters.
Day of the week	Number	Information on the day of the week.
Hour of the day	Number	Information on the hour of the day.

All the above variables are collected every 15 min, with the exception of the outdoor temperature and humidity, which are affected by a one-hour time granularity. In this case, a linear interpolation has been applied to fill the missing quarters of hours. Furthermore, before feeding the neural network, the values are normalised by defining a maximum and minimum value for each of them. No missing values to be handled have been encountered in the collected data.

3.2. Neural Network

A one-dimensional (in short 1D) CNN is a neural network model composed of one or more 1D convolutional layers. Like in our previous work [25], we use 1D convolutions to extract fine-grained information from one-dimensional data (such as indoor temperature, humidity, or energy consumption) along the temporal dimension. Compared to 2D CNNs and other DL models such as multi-layer perceptrons (MLPs) or recurrent neural networks (RNNs), 1D CNNs are less computationally complex and can perform well even when using shallow architectures [22].

3.2.1. Architecture

In this work, we propose a 1D CNN to predict future values of temperature, humidity, CO_2 level, and energy consumption with the objective to timely tune the HVAC system so that the given comfort level is followed with minimal energy consumption. The architecture consists of four layers, as sketched in Figure 1:

- *Input Layer*. The first layer takes as input an $n \times f$ array, where n is the duration of the observation time window (expressed in quarters of an hour) and f is the number of features. That is, an input sample consists of the values of f variables, including temperature, humidity, and CO_2, energy consumption, and timestamp HVAC operating parameters, collected during a time window of n quarters of an hour.

Each sample is normalised along the temporal axis by using the nominal minimum and maximum values of each variable.

- *1D Convolutional Layer*. Each sample is operated by a convolutional layer with k filters, each filter of size $h \times f$, with h being the height and f the width of the input sample (i.e., the number of features). Each of these k kernels slides over the input matrix with a step of 1 to extract the temporal properties of the f features. This layer outputs a matrix of size $(n - h + 1) \times k$, in which i-th column is a feature vector extracted by the i-th filter. The rectified linear activation function (*ReLU(x) = max{0,x}*) is used to break linearity in the model, as per convention for CNNs.
- *Max Pooling Layer*. The max pooling operation downsamples the temporal properties extracted with the convolution by keeping only the largest values. This operation involves a filter that slides along each feature map with a pre-defined step (also called *stride*) and applies a maximum operator to a number of elements equal to the *pool size* parameter. As we set *stride* equal to *pool size*, the size of the resulting matrix is equal to $((n - h + 1) / pool\_size) \cdot k$.
- *Flatten Layer*. It reshapes the output of the max pooling operation into a one-dimensional feature vector.
- *Output Layer*. The vector is processed by a final fully connected layer. The output of this layer is a vector of four elements, with the predictions of indoor temperature, humidity, CO_2, and energy consumption for the quarter of an hour that follows the input time window.

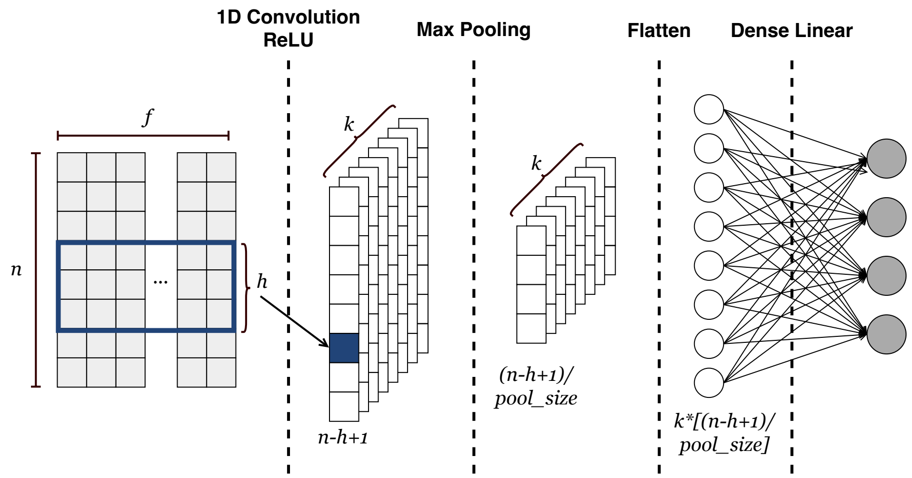

Figure 1. The designed neural network architecture.

3.2.2. Model Training

Table 2 reports the values of the main hyper-parameters of the AI model. Considering the obtained performance, we use largely the same configuration reported in our previous work [25], differing only in a couple of parameters: we define different loss weights for the output variables with the aim to balance their contribution equally within the overall loss, and we reduce the maximum number of training epochs to 1000.

Table 2. Model parameters.

Parameter	Value
Time granularity of data	15 min
Quarters of an hour in a sample	8
Number of kernel filters, Convolutional layer	64

Table 2. *Cont.*

Parameter	Value
Kernel size, Convolutional layer	3
Pool size, Max Pooling layer	2
Learning rate	0.001
Batch size	32
Optimizer	Adam [26]
Loss function	Mean squared error
Validation split	0.3
Loss weights	CO_2: 0.1, indoor temperature: 100, indoor humidity: 1, energy: 1000
Maximum number of epochs	1000
Patience	25

Every day after midnight, we train the AI model by using a window of data (i.e., mobile window) collected over the last 30 days and handled in samples covering a window of eight quarters of an hour. As reported under Sections 1 and 3.1, data are made available thanks to a data collection system, which is necessary to regularly sense the environment through sensors and smart energy meters. The objective of the update mechanism leveraging mobile window [25] is twofold:

1. Enabling the model to learn the impact of actuations (ON/OFF, set point) on the indoor environment and energy consumption over the recent period.
2. Keeping up with seasonal variations in environmental conditions. Indeed, the same HVAC settings might lead to different effects on the environment under different conditions, e.g., outdoor weather.

3.3. Predicted Mean Vote

The predicted mean vote (PMV) is a thermal comfort index introduced by Fanger [23] and used by many research works to model the thermal comfort of the occupants. In particular, as observed by Tartarini et al. [27], PMV computes the mean thermal sensation vote of a large number of people according to a sensation scale ranging from −3 to +3, respectively, from cold to hot passing through a value equal to 0, which means a neutral condition. In this research work, the open-source Python library *pythermalcomfort* [28] is used to compute the PMV index. This library comprises a range of functions for modelling indoor environmental comfort and its associated parameters. In addition, we refer to an online tool [27] to dynamically find the boundaries for the different thermal comfort categories defined by the EN 16798 standard. The PMV index is computed as a function of environmental and personal variables [20], in particular:

- Air temperature [°C], the indoor temperature in the environment.
- Mean radiant temperature [°C], defined as the temperature due to radiant heat exchange between a human body and a given environment [29]. For the sake of simplicity, we assume the mean radiant temperature equal to the air temperature.
- Relative humidity [%], Indoor relative humidity in the environment.
- Metabolic rate [W/m^2]. It is associated with the activity performed by the occupants in the environment. We set it equal to 1.6, which corresponds to light activity in the environment coherently with our real test case.
- Relative air velocity [m/s]. It includes the air speed within the considered environment as well as the air speed due to body movement. It is computed by using function *v_relative(v, met)* from library *pythermalcomfort*, with *v* equal to 0.15 m/s (heating) or 0.25 m/s (cooling) according to standard limits defined by ISO 7730 [30].
- Clothing insulation [clo]. It is the thermal insulation provided by clothing worn by people in the environment. As estimating a single value for each person requires advanced sensors as well as possible customisations within the environment, we modeled such a parameter with a unique value. To cope with this task, we used

the function *clo_dynamic(clo, met)* from library *pythermalcomfort*. In this case, *clo* is computed by means of function *clo_tout(tout, units="SI")*, which computes the daily clothing insulation based on outdoor temperature at 6.00 a.m. In this way, we provide a dynamic estimate of the clothing insulation throughout the year.

Standard EN 16798 ([31]) defines specific categories for indoor comfort based on the PMV, which are reported in Table 3:

Table 3. Comfort categories and the related PMV range.

Category	PMV
I	$-0.2 < \text{PMV} < 0.2$
II	$-0.5 < \text{PMV} < 0.5$
III	$-0.7 < \text{PMV} < 0.7$
IV	$-1 < \text{PMV} < 1$

It is worth noting that negative ranges of PMV index are typical for heating mode, while positive values are obtained when cooling mode is active.

As underlined by Yau et al. [20], the PMV index is an objective measure that can be computed in any indoor environment regardless of installed HVAC systems and conditions of the outdoor environment. Therefore, considering its widespread use in the literature as well as its reference in different standards (e.g., ISO 7730, EN 16798), we have chosen to use such methodology.

4. Methodology

In this work, we tackle the problem of energy-efficient comfort optimisation in indoor environments. That is, we study and develop a methodology for the automated control of HVAC systems so that the defined comfort requirements within the considered environment (e.g., by managers) during the day are respected with minimal energy consumption. As introduced in Section 3.3, the thermal comfort index (PMV), as defined by Fanger [23], depends on a set of parameters (such as air temperature and humidity of the environment), which, in a real-world environment, can be influenced by the outdoor conditions. In this regard, adapting the HVAC optimisation according to outdoor weather takes a central role from a research perspective [32]. At first glance, a trivial greedy PMV-based mechanism that activates the HVAC system when the thermal comfort level is outside the desired range, similar to the approach proposed by Jing et al. [7], might be viewed as a viable solution. However, such an approach, which makes decisions only considering the current state, might not work as desired. In particular, let us first define four comfort states (represented in Figure 2) based on the comfort interval defined throughout the day:

- **No Comfort (NC)**: the shop is closed (e.g., at night or on Sundays).
- **No Comfort then Comfort (NC-C)**: usually early morning before the opening.
- **Comfort (C)**: the shop is open (e.g., during a working day).
- **Comfort then No Comfort (C-NC)**: generally late afternoon before closing.

Indeed, a greedy approach might not be able to achieve the target comfort at the beginning of the working time (NC-C state), i.e., when the comfort level is far from the target value because of a long inactivity period (e.g., night closure, holiday, etc.). For the same reason, it might activate the HVAC system when the store closure is approaching (C-NC state), leading to inefficient energy utilisation.

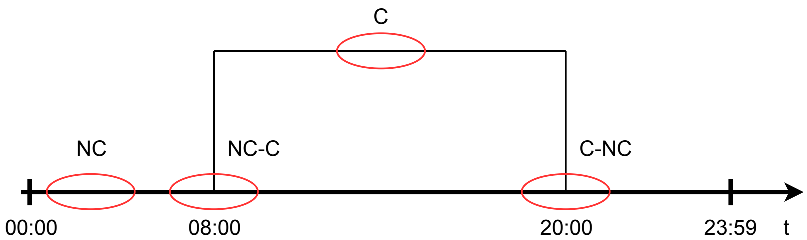

Figure 2. Representation of the four control states throughout the day. In this example, the comfort interval is defined between 8 a.m. and 8 p.m..

Based on these premises, we propose an AI-based solution called EECO, in which a CNN is used to predict the future comfort level and energy consumption of the HVAC. Specifically, given a range of possible HVAC configurations (meaning, ON/OFF, and SP), the CNN predicts the effects of each choice on future comfort and energy consumption. At every quarter of an hour, the system computes the predictions for the next m quarters of an hour, generating an m-level tree of candidate sequences of HVAC configurations.

The ultimate goal of EECO is to select the branch of the tree (hence, a sequence of future HVAC configurations), which, based on the CNN predictions, will minimise an objective function defined as the weighted summation of thermal comfort index PMV and energy.

In the remainder of this section, we describe the whole process of comfort optimisation, including input/output of the CNN, the structure of the decision tree, and the logic behind the choice on the HVAC settings. This process is described in Algorithms 1 and 2 and illustrated in Figures 3 and 4.

Algorithm 1 Tree building

Input: Root node (n_{000}), Historical data (X), Tree depth (m), Target comfort ($\bar{C}$), Operating mode (o)
Output: Tree (t)

procedure BUILDTREE(n_{000},X,m,$\bar{C}$,o)
  $t \leftarrow L_0(n_{000})$ ▷ Init tree node n_{000} at level 0
  for $i = 1, ..., m$ **do** ▷ Loop over tree levels
    $L_i \leftarrow \emptyset$ ▷ Init level i
    $X_i \leftarrow X[-n+i, i-1]$ ▷ Extract *n-1* rows from X
    for $j = 0, ..., |L_{i-1}| - 1$ **do** ▷ Loop over parents p_{ij}
      $P_{ij} \leftarrow \emptyset$ ▷ Init list of children of parent node p_{ij}
      $A_{ij} \leftarrow [(\text{OFF}, p_{ij}(\text{SP}))]$ ▷ Init list of actuations
      if $\bar{C} \neq NC$ **then**
        $A_{ij} \leftarrow A_{ij} \cup \text{GETACT}(p_{ij}(\text{T}), p_{ij}(\text{H}), \bar{C}, o)$
      end if
      for $k = 0, ..., K_j - 1$ **do** ▷ with $K_j = |A_{ij}|$
        $(\text{ON/OFF}_{ijk}, \text{SP}_{ijk}) = A_{ij}[k]$
        $V_{ij} = [p_{ij}(T), p_{ij}(H), p_{ij}(CO2), p_{ij}(E))]$
        $X_{ijk} \leftarrow X_i \cup [\text{ON/OFF}_{ijk}, \text{SP}_{ijk}, \cdots, V_{ij}]$
        $n_{ijk} \leftarrow \text{GETNODE}(X_{ijk})$
        P_{ij}.insert(n_{ijk})
      end for
      L_i.insert(P_{ij}) ▷ Add nodes of list P_{ij} to level i
    end for
    t.insert(L_i) ▷ Add level i to the tree
  end for
  return t
end procedure

Algorithm 2 Get actuations

Input: Current temperature (T), Current humidity (H), Target comfort ($\bar{C}$), Operating mode (o)
Output: List of actuations (A)

```
procedure GETACT(T, H, C̄, o)
    A ← ∅                                          ▷ Init list of actuations
    T_min, T_max ← GETRANGETEMPERATURE(T, H, C̄)
    if T_min <= T <= T_max then
        SP ← T
        if o = HEATING then
            while SP <= T_max and SP <= T + 1 do
                A.insert((ON,SP))
                SP ← SP + 1
            end while
        else if o = COOLING then
            while SP >= T_min and SP >= T − 1 do
                A.insert((ON,SP))
                SP ← SP − 1
            end while
        end if
    else
        if o = HEATING then
            SP_min ← ceil(T)
            A.insert((ON, SP_min))
        else if o = COOLING then
            SP_max ← floor(T)
            A.insert((ON, SP_max))
        end if
    end if
    return A
end procedure
```

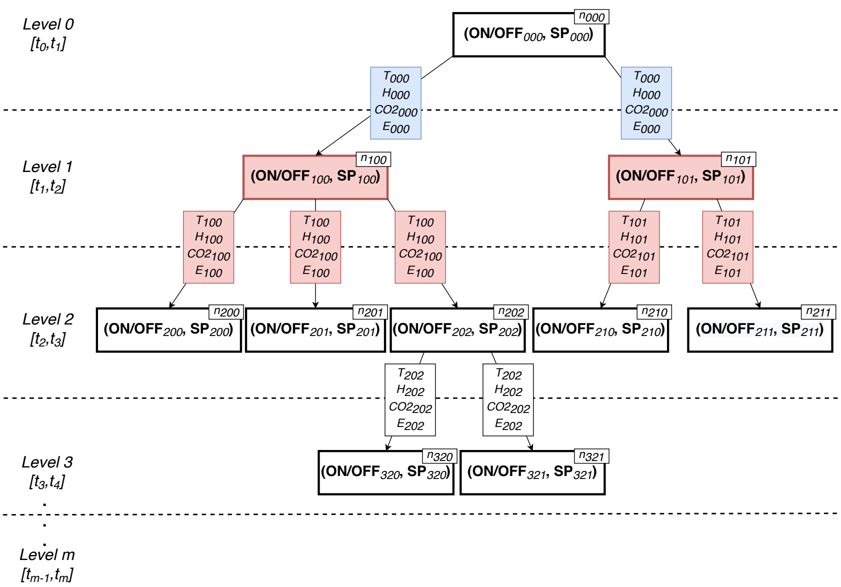

Figure 3. The decision tree. Node's attributes are HVAC configurations (ON/OFF$_{ijk}$, SP$_{ijk}$), which are labelled with 3−digit numbers: the level of the tree (i), the index of the parent node (j), and the index of the node (k). T_{ijk}, H_{ijk}, $CO2_{ijk}$, and E_{ijk} refer to predicted values of temperature, humidity, CO_2, and energy consumption for node n_{ijk} at Level i in the time slot $[t_i, t_{i+1}]$.

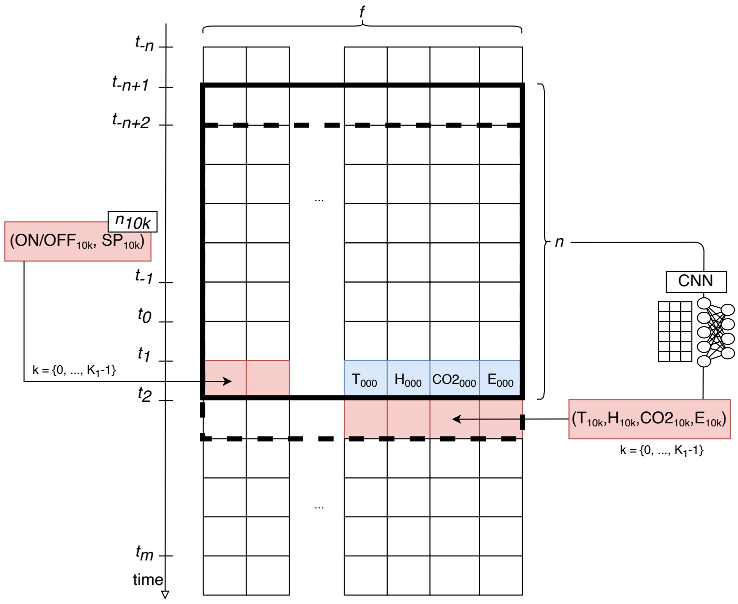

Figure 4. The first step of the building decision process at time slot $[t_0, t_1]$, during which a tree of possible HVAC configurations is built iteratively from time slot $[t_1, t_2]$ to time slot $[t_{m-1}, t_m]$. In this example, $k = \{0, 1\}$ for the two nodes at Level 1 of the tree (Figure 3).

The decision tree is built every quarter of an hour (or time slot) using the output from the previous time slot as a root node. The process that builds the tree is formulated in Algorithm 1 BUILDTREE. BUILDTREE takes as input the current root node n_{000}, historical data of HVAC settings, weather conditions, energy consumption, the target comfort value $\bar{C}$ (a PMV value), and the operating mode o (either heating or cooling). The root node's attributes include the current HVAC settings, i.e., the operational settings in time slot $[t_0, t_1] = [t_0, t_0 + 15 \text{ min}]$. In general, a node n_{ijk} of the tree is characterised by a 3-digit label and range of attributes. The first digit of the label indicates the level of the tree to which the node belongs, the second digit is the index of the parent node, and the third digit is the index of the node. The attributes are the current HVAC settings at time t_i, i.e., the pair of values $(\text{ON/OFF}_{ijk}, \text{SP}_{ijk})$, fan speed, and operating mode. Node's attributes also include average energy consumption E_{ijk}, indoor temperature T_{ijk}, indoor humidity H_{ijk}, and indoor CO_{2ijk}.

Figure 3 illustrates a portion of the tree built during time slot $[t_0, t_1]$, starting from Level 0, which consists of root node n_{000}. Level 1 of the tree is populated with a set of children nodes n_{10k}, $k = \{0, \cdots, K_1 - 1\}$, each one defined with pair $(\text{ON/OFF}_{10k}, \text{SP}_{10k})$, i.e., a set of possible HVAC configurations that could be applied during time slot $[t_1, t_2]$ (Level 1 in Figure 3). Like all the other tree levels, Level 1 includes the OFF actuation (line 8 of Algorithm 1) and a set of actuations that are computed with Algorithm 2 (called at line 10 of Algorithm 1) using the indoor temperature of the parent node ($p_{00}(T) = T_{000}$ for Level 1), the indoor humidity of the parent node ($p_{00}(H) = H_{000}$ for Level 1), the target comfort level $\bar{C}$, and the HVAC's operating mode o (either HEATING or COOLING). Algorithm 2 defines the temperature range to be within the desired target comfort $\bar{C}$ and, based on that information and HVAC's operating mode o, selects the strategy to enter the comfort range or move within that through a couple of actuations.

One of the nodes at Level 1 is the output of the process executed during time slot $[t_0, t_1]$ and contains the HVAC configuration for time slot $[t_1, t_2]$. Moreover, such a node will be the root node when the process is executed in time slot $[t_1, t_2]$. Which is the right node? The selection of the most appropriate node is completed by populating the tree up to Level m using the predictions of the CNN to simulate the behaviour of the system in different conditions over the time (until time slot $[t_{m-1}, t_m]$). The solution is the node at Level 1 that belongs to the branch of the tree whose sequence of actuations guarantees the best comfort at the minimum energy consumption in the long term. The logic behind this decision is explained in the following steps:

- Given a level L_i, and a parent node p_{ij}, with $j \in \{0, \cdots, |L_{i-1}|\}$ and $|L_{i-1}|$, the cardinality of Level L_{i-1}, A_{ij} is the list of possible HVAC actuations for time slot $[t_i, t_{i+1}]$ applied to the children nodes of parent p_{ij} (lines 8, 10). In Figure 3, $A_{21} = \{(\mathrm{ON/OFF}_{210}, \mathrm{SP}_{210}), (\mathrm{ON/OFF}_{211}, \mathrm{SP}_{211})\}$. For each HVAC configuration $A_{ij}[k]$ $k \in \{0, \cdots, K_j - 1\}$, the system predicts the effects of such configuration on comfort and energy consumption starting from the parent's conditions V_{ij} of indoor temperature $p_{ij}(T)$, indoor humidity $p_{ij}(H)$, indoor CO_2 $p_{ij}(CO2)$, and energy consumption $p_{ij}(E)$ (line 14).
- As sketched in Figure 4, the prediction for node k is obtained by feeding the CNN with an array of $n-1$ rows of historical HVAC settings, environmental values, and other features (see Table 1) observed from t_{i-1} to t_{-n+i}. The nth line contains the node's attributes A_{ij} and other features related to $[t_i, t_{i+1}]$. This operation is summarised at line 16 with function GETNODE. Node n_{ijk}, generated using actuation $A_{ij}[k]$, is added to the list of children nodes P_{ij} of parent p_{ij} (line 17).
- The list of children nodes P_{ij} is added to Level L_i, which is then added to the tree when all the parents of the previous Level L_{i-1} have been processed.
- The above steps are repeated until the maximum tree depth m is reached.

The result of the process is a set B of simulated sequences of HVAC configurations from time t_1 to time t_m, which can also be viewed as a set of paths across the decision tree (or branches) from the root node to the leaves. The final step consists of choosing the best path, i.e., the path that minimises both PMV and energy values, as formally expressed in Equation (1):

$$f_\alpha(C_b, E_b) = \alpha \cdot C_b + (1-\alpha) \cdot \frac{E_b}{E_{max} \cdot m} \qquad \forall b \in B \tag{1}$$

The objective function $f_\alpha(C_b, E_b)$ is the weighted sum of predicted comfort and energy for branch b, where C_b is a sum of the predicted values of thermal comfort on each node of the branch, while E_b is the sum of the predicted values of energy consumption. More precisely, C_b and E_b are computed as follows:

$$C_b = \sum_{i=0}^{m} (|C_{b,i}| - |\bar{C}|) \cdot \beta^i \quad E_b = \sum_{i=0}^{m} E_{b,i} \quad \forall b \in B \tag{2}$$

where β is a positive number smaller than 1, so that β^i (β at the power of i) decreases as the tree level i increases to provide less importance to the nodes far from the root (i.e., far in the future).

The energy is normalised with the estimation of the maximum energy E_{max} consumed by the HVAC system in a quarter of an hour and multiplied by the number of time slots in a branch (m). α controls the relative weight of comfort and energy values. In our analysis, we focus on a scenario where comfort holds priority. In this regard, we set $\alpha = 0.9$.

For a given value of α, the solution is represented by the branch $\bar{b} \in B$ such that

$$\bar{b} = \underset{b \in B}{\operatorname{argmin}} \, f_\alpha(C_b, E_b) \tag{3}$$

Hence, the output of the whole process is the HVAC configuration $(\mathrm{ON/OFF}_{10k_{\bar{b}}}, \mathrm{SP}_{10k_{\bar{b}}})$ for the next time slot $[t_1, t_2]$, i.e., the attributes of node $k_{\bar{b}}$ at Level 1 of branch $\bar{b}$. The above process is executed every 15 min.

5. Experimental Setup

The proposed solution has been tested in a real warehouse of about 250 m^2 owned by an international retail company in northern Italy where two HVAC devices are installed.

The latter are interfaced through the Modbus protocol [33] and regulated with same settings in terms of ON/OFF and set point.

The building under consideration, as reported in Section 1, is located in northern Italy and generally experiences a Mediterranean climate, with outdoor temperatures ranging from approximately 0 °C to 15 °C in the winter and between 18 °C and 35 °C in the summer. Both real-world experiments and software simulations have been performed to validate our algorithm, considering both winter and summer days to provide a performance overview both in heating and cooling mode.

We have compared *EECO*, which dynamically sets the HVAC configuration to achieve a tradeoff between thermal comfort and energy consumption, with other two approaches:

- The *Fixed Set Point* approach, which configures the same set point value throughout the whole day.
- The *PMV-based* approach, a greedy strategy that controls the HVAC devices by just analysing the current value of the PMV index. In this regard, the greedy approach aims to achieve the best comfort conditions within the desired range (i.e., the lower bound), switching the HVAC devices off once such objective is addressed.

All the operations in the real environment are performed on an Intel NUC [34] with Intel Core i3 as CPU and 8 GB of memory. For all the experiments, we set a comfort time interval of twelve hours between 8 a.m. and 8 p.m.in accordance with our partner's requirements. During this interval, a pre-defined comfort level needs to be guaranteed (cf. Table 3). Due to lack of activities within the environment on the weekends, we performed our experiments only from Monday to Friday. For evaluation purposes, we set the tree depth m equal to 10, starting from the current quarter of hour. This enables us to align the operational schedule of HVAC devices according to EECO with that of the fixed set point strategy (from 6 a.m. to 8 p.m.), which is defined by the partner company.

6. Results

In this section, we present the benefits of *EECO* from different perspectives, both through experiments in a real environment and by software simulation.

6.1. Indoor Environment Forecast

One of the strengths of *eeco* is the capability to predict the evolution of the indoor environment and to make proper short-term decisions. To this aim, let us consider a summer early morning scenario when the HVAC devices must be turned on in advance to reach the desired comfort level at a given time (e.g., 8 a.m.). Also, let us select the *category III* (cf. Table 3) as the target comfort, i.e., $0.5 <= \text{PMV} < 0.7$ in cooling mode and $\alpha = 0.9$ in Equation (1). With this experiment, we show the ability of our predictive methods in activating the HVAC devices in advance so that the target comfort is achieved before 8 a.m..

Figure 5 shows the actuation strategy in terms of ON/OFF and predicted/actual values of PMV index and energy consumption recorded during two overlapping time intervals of two hours, one starting at 5:45 a.m. and the following one starting at 6:00 a.m. From the plots in the figure, we can observe that the algorithm turns on the HVAC devices at 6 a.m. and plans to keep them active over the following two hours to achieve the comfort range on time (i.e., PMV $<$ 0.7 before 8 a.m.). Indeed, as reported in Table 4, such a configuration is expected to achieve a PMV value just below 0.7 at 8 a.m. Actually, analysing the dashed line in the PMV plots, which corresponds to the actual comfort values recorded every quarter of an hour, it is possible to notice that the environment meets the comfort requirements from 7 a.m., resulting in a PMV value equal to 0.65 at 8 a.m. thanks to an accurate prediction of the environment evolution (Table 4). It is worth noting that this result would not have been possible with a basic greedy strategy, mentioned in Section 4, which would have activated the HVAC devices starting from 8 a.m., thus leaving the indoor comfort outside the required range at the beginning of the working day.

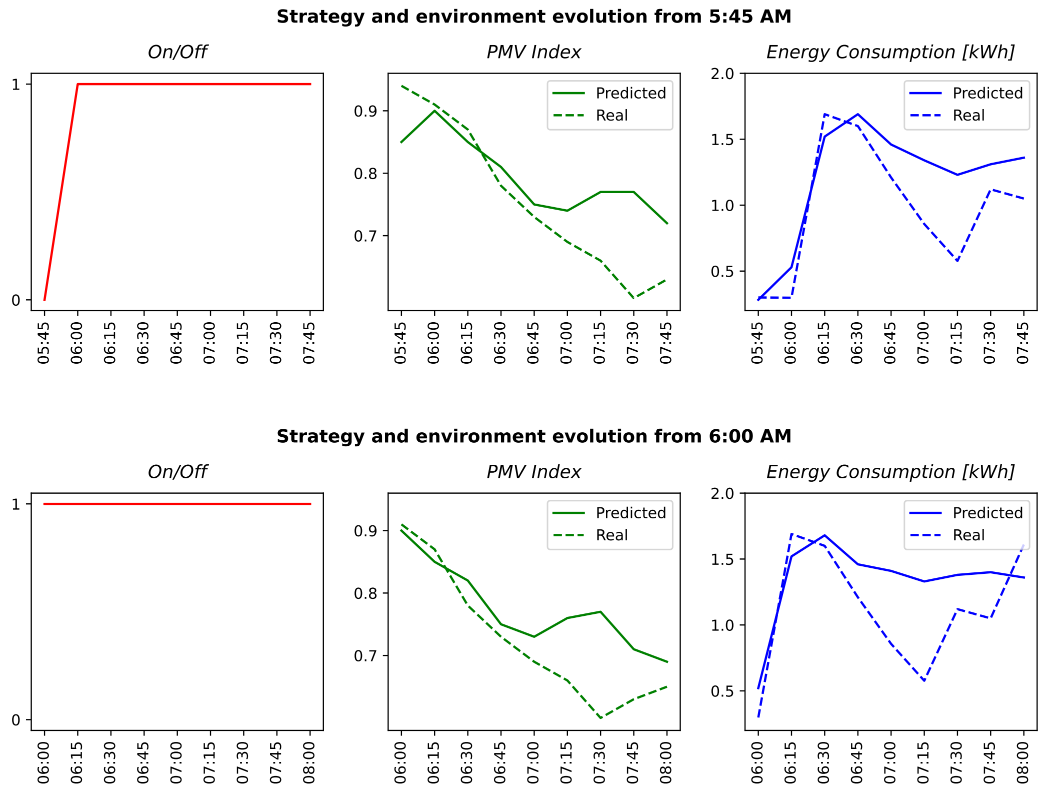

Figure 5. Example of actuation strategy in terms of ON/OFF (**left**) and related predicted environment evolution in terms of PMV index (**middle**) and energy consumption (**right**) in the early morning at 5:45 a.m. (**top**) and 6:00 a.m. (**bottom**) to achieve the comfort requirements at 8:00 a.m.

Table 4. Accuracy of the predicted environment in terms of energy consumption and PMV as well as difference between the real PMV and the predicted value at 5:45 a.m. and 6:00 a.m., respectively.

Time	Energy-RMSE	PMV-RMSE	Real PMV-8 a.m.	Predicted PMV-8 a.m.
5:45 a.m.	0.32	0.08	0.65	0.7
6:00 a.m.	0.37	0.07	0.65	0.68

6.2. Indoor Comfort and Energy Consumption Optimisation

In this experiment, our objective is to evaluate the sensitivity of our solution to term α of Equation (1), which controls the relative weight between comfort and energy in the HVAC strategy selection. To this aim, let us choose for this test $\alpha = \{0.3, 0.5, 0.7, 0.9\}$, with a target comfort level set to *category I* in cooling mode, i.e., $0 <= \text{PMV} < 0.2$ (cf. Table 3). In this regard, because of the consistently low outdoor temperatures observed during the period of these experiments, the indoor environment would have exceeded the boundaries set by comfort categories classified as II or III (which are more appropriate for this specific scenario), regardless of the value of α, while keeping the HVAC devices off (as reported in Section 4). This would have affected the performance analysis of our solution on different values of α both in terms of energy consumption and thermal comfort. As a result, in such conditions, category I (i.e., which is typically defined in environments with vulnerable people) enables us to analyse the behaviour of our solution while varying the α parameter.

Figure 6 outlines the behaviour of the PMV index and total energy consumption at different values of α. We can notice that, setting $\alpha = 0.9$, the PMV value stays close to zero for most of the day, which corresponds to optimal comfort conditions for the occupants. On the other hand, this strategy impacts the building's energy footprint as HVAC devices are constantly activated to achieve maximum comfort. Decreasing the value of α to 0.7 impacts the comfort, especially in the morning. Nevertheless, equivalent performance to higher α values is achieved in the late afternoon due to a decreasing trend in outdoor temperature. Compared to $\alpha = 0.9$, with these settings, the energy consumption is reduced by about 5 kWh. Finally, with lower values of α (i.e., $0.3, 0.5$), we penalise the indoor comfort, resulting in higher PMV values ranging between 0.1 and 0.2. We can notice that the reduction in energy consumption is negligible compared to the tests performed with higher values of α. As reported in Table 5, low values of α not only impact the daily comfort conditions but result in higher PMV values the next early morning as they leave the environment in worse comfort conditions at the end of the day. This means that HVAC devices will have to bring the environment within the desired comfort level, with a consequent impact on the energy footprint. As can be noticed in the last column of the table, higher values of α ensure the desired comfort level from the early hours of the next day.

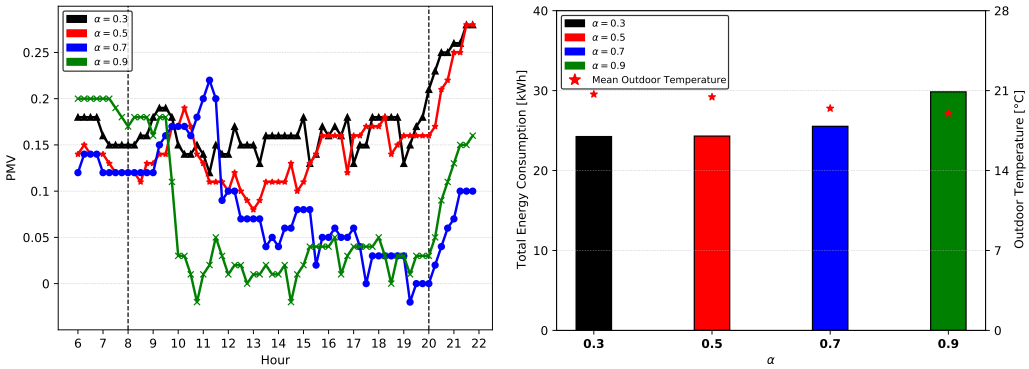

Figure 6. Behaviour of (**left**) hourly PMV index and (**right**) total energy consumption as α parameter changes.

Table 5. Average PMV, total energy consumption [kWh], and PMV at 6 a.m. the next day at different α values.

α	PMV	Total Energy Consumption	PMV 6 a.m.
0.3	0.15	24 kWh	0.27
0.5	0.13	24 kWh	0.22
0.7	0.08	25 kWh	0.09
0.9	0.04	29 kWh	0.12

6.3. *Performance Comparison: Real Environment*

The results of the previous analysis highlighted that our solution can effectively achieve the comfort objective at different values of α. In this experiment, we focus on guaranteeing a high comfort level with minimum attention on the energy footprint as well, comparing our algorithm with $\alpha = 0.9$ in Equation (1) and a manual approach, where the set point is fixed to the same value throughout the day. We configure our algorithm with *category III* as the target comfort, i.e., with $0.5 <= \text{PMV} < 0.7$, and we consider both days in heating and cooling mode.

Finally, we activate the HVAC devices at 6 a.m. in the manual approach to guarantee the same operating interval between the two strategies (see discussion of the experiment in Section 6.1).

The evaluation of the two strategies has been performed over 10 days, five days each, in cooling mode, while we have considered 20 days, 10 days each, in heating mode.

6.3.1. Cooling Mode

In the manual approach, we set SP = 27 °C, i.e., the value typically set within the considered environment during cooling periods when the indoor temperature is usually higher. It roughly corresponds to *category III* of the target comfort according to Tartarini et al. [27] and the configured parameters of PMV index described in Section 3.3, which define our evaluated scenario.

Figure 7 reports results in terms of average (a) PMV during opening hours, (b) total energy consumption over the operational interval of HVAC devices against the average daily PMV index, (c) indoor temperature, and (d) variation in PMV during opening hours (i.e., the difference between PMV at 6 a.m. and PMV at the different quarters of hours).

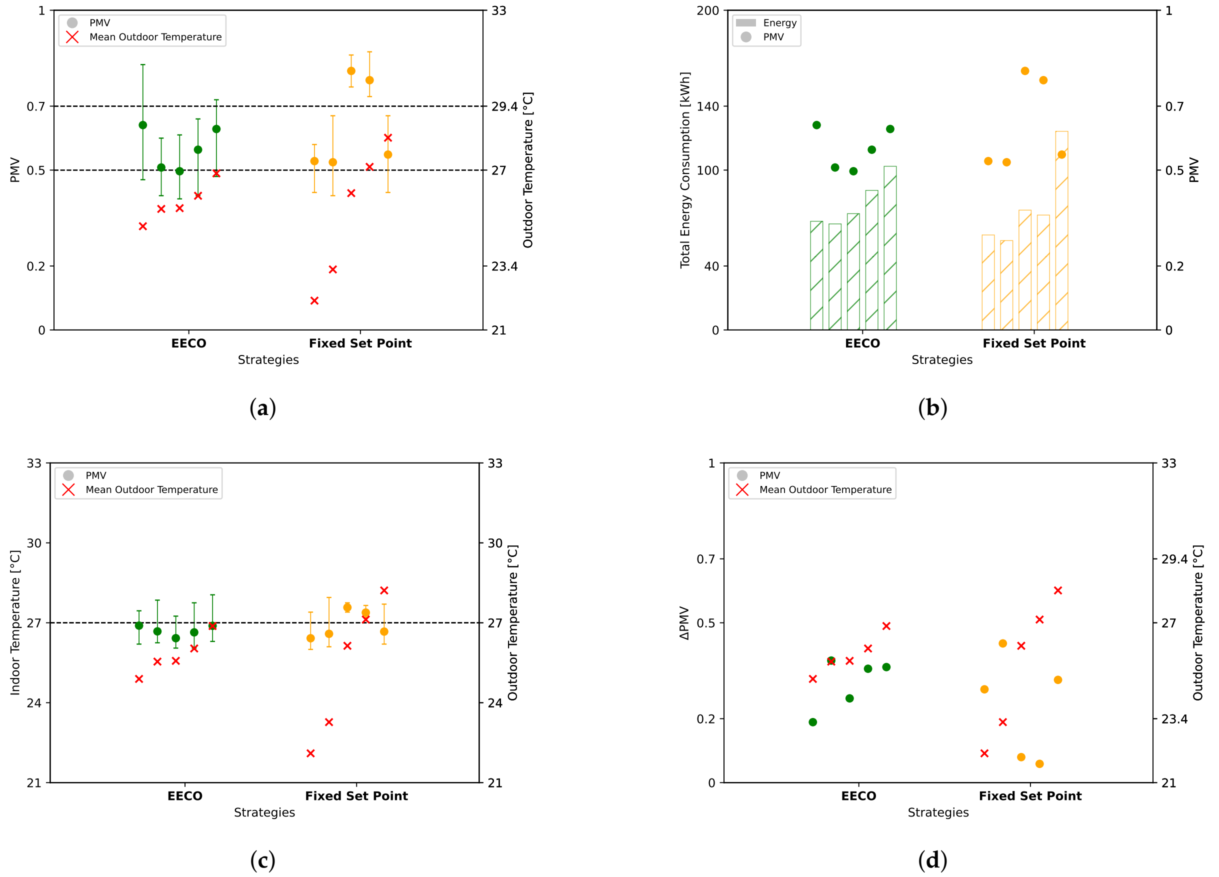

Figure 7. Average daily (**a**) PMV index, (**b**) total daily energy consumption against the average daily PMV index, (**c**) indoor temperature, and (**d**) variation in PMV index compared to value at 6 a.m. (PMV_{6AM}-PMV) for both EECO and the Fixed Set Point approach in cooling mode.

In the case of PMV and indoor temperature, Figure 7 reports the confidence interval bounded by the maximum and minimum values. From a comfort perspective, our solution constantly keeps the PMV value close to the lower bound (i.e., 0.5), which corresponds to the best possible indoor comfort conditions within the desired range (Figure 7a). The process of comfort optimisation generally requires more activity by HVAC devices (Figure 7d), hence affecting the total energy footprint (Figure 7b). However, such energy impact is generally limited, with no relevant peaks on the energy footprint. A slight increase in energy consumption can be noticed during the last two days (day four and five) due to a slightly higher outdoor temperature.

On the other hand, the *Fixed Set Point* approach produces unstable results in terms of thermal comfort (Figure 7a). Indeed, with static settings, the HVAC devices are activated only based on the indoor temperature. However, the configured set point value (i.e., 27 °C) might not be correct to constantly achieve the target comfort requirements, which, in the evaluated scenario, represent a constraint of the problem; as reported in Section 3.3, PMV index depends on a set of parameters. For instance, clothing insulation, whose value is computed daily using the outdoor temperature at 6 a.m., affects the weight of the indoor temperature on the computation of thermal comfort. Obviously, this parameter is not considered when using a static set point, possibly leading to discomfort conditions. For instance, day four of the *EECO* experiment and day three of the *Fixed Set Point* experiment are expected to have similar clothing insulation values. However, during the former, the set point is automatically set to SP = 26 °C for a large part of the day. Meanwhile, with the *Fixed Set Point*, the set point value is close to the indoor temperature of 27 °C, leaving the HVAC devices in economic mode. The HVAC devices typically have a dead band equal to 0.5 °C around the control set point, which keeps them active to maintain the target temperature with no need for strong activity. In terms of energy consumption, if we compare the two days in Figure 7b, we can observe that our solution consumed slightly more energy than the *Fixed Set Point* configuration; as underlined previously, the *Fixed Set Point* configuration never reached the required comfort range during that day due to the static set point value that never changes during the day, while *EECO* was able to drive the HVAC devices in a way that the comfort requirements were respected throughout the whole day.

In Figure 7b, we can notice a peak in energy consumption on day five of the *Fixed Set Point* experiment. Compared to days one and two of the same experiment, day five reaches a similar comfort level with a much higher energy consumption. The reason for that is the high outdoor temperature, which requires higher and longer activity from the HVAC devices to keep the indoor temperature close to the set point of SP = 27 °C. On the other hand, on days one and two, the low outdoor temperature helped the cooling operations, pushing the indoor temperature under the set point value with a lower impact on the building's energy footprint.

In conclusion, as reported in Table 6, in cooling mode, our solution guarantees a slightly better comfort conditions approach with a minimal impact on the building's energy footprint compared to the *Fixed Set Point* approach.

Table 6. Overall performance of *EECO* and the *Fixed Set Point* approach in terms of PMV and energy consumption in cooling mode.

	PMV	**Energy Consumption**
EECO	0.56	78 kWh
Fixed Set Point	0.63	77 kWh
Difference	**11%**	**−1%**

6.3.2. Heating Mode

In the manual approach, we set SP = 21 °C, i.e., the value typically set within the considered environment during heating periods when the indoor temperature is usually lower and roughly corresponding to *category III* of the target comfort. Even in this case, we refer to Tartarini et al. [27] and the parameters of PMV index described in Section 3.3 we set for our scenario.

As in the previous subsection, Figure 8 reports results in terms of average (a) PMV during opening hours, (b) total energy consumption over the operational interval of HVAC devices against the average daily PMV index, (c) indoor temperature, and (d) variation in PMV during opening hours (i.e., the difference between PMV at 6 a.m. and PMV at the different quarters of hour).

From a comfort perspective, as depicted in Figure 8a, it can be noted that our algorithm provides non-optimal results during the coldest days (days one, two, and three in both approaches) unlike the *Fixed Set Point* approach. This behaviour is due to some shutdowns of the HVAC devices configured by our approach throughout the day with the aim to effectively achieve a tradeoff between the thermal comfort and the energy consumption. However, due to the low outdoor temperature, just a shutdown for a quarter of hour prevents our solution from achieving good PMV values in the long term like *Fixed Set Point*, which constantly keeps the HVAC devices active at 21 °C as a set point value. Even if the target temperature is not reached, the given comfort category is effectively achieved when the indoor temperature is between 20 °C and 21 °C. Despite such a scenario, the minimum comfort requirements are satisfied with *EECO* at some point during the day (i.e., in the afternoon), resulting in about 20 kWh on average of saved energy (Figure 8b).

When the average outdoor temperature is around 5–6 °C (from day four to seven in both approaches), *EECO* generally provides better results in terms of thermal comfort compared to the *Fixed Set Point* approach. Indeed, as depicted in Figure 8a, larger improvements in PMV are performed by our approach in order to improve the comfort conditions within the desired range. As a result, HVAC devices are forced to have a higher activity, resulting in consuming a little bit more energy (Figure 8b) compared to the *Fixed Set Point* approach. On the other hand, the latter struggles in such conditions from a comfort perspective (e.g., day five provides out-of-comfort conditions). Indeed, the *Fixed Set Point* strategy is driven only by the indoor temperature, which is always really close to the set point value (i.e., 21 °C) throughout the whole day. This forces the HVAC devices to constantly work in economic mode (as reported in the previous Section 6.3.1), limiting their energy consumption but potentially compromising optimal comfort. In such conditions, achieving the desired comfort levels typically requires setting a higher target temperature (e.g., 22 °C). In this regard, *EECO* dynamically tunes the set point value to enhance comfort, resulting in only a marginal increase in energy consumption compared to the *Fixed Set Point* approach.

During warmer winter days (from day eight to ten in both approaches), the higher outdoor temperature values help the environment to achieve good comfort conditions with both approaches. In such a scenario, our algorithm is capable of reducing the energy footprint by almost 20% compared to the *Fixed Set Point* approach. In this regard, *EECO* not only configures some strategic shutdowns of the HVAC devices to take advantage of the high outdoor temperature to optimise heating operations but switches the HVAC devices off when the lower bound of the expected comfort level (i.e., in our case, 0.5) is achieved. In contrast, such behaviour is not included in the *Fixed Set Point* approach. Indeed, the latter is driven solely by the indoor temperature with no attention for thermal comfort. As a result, it always keeps the HVAC devices active, even after achieving the desired thermal comfort level with the configured set point value, as observed on day eight. Such behaviour results in a significant amount of wasted energy over the long term, which can be effectively limited through an intelligent approach like *EECO*.

To sum up, as presented in Table 7, in heating mode, *EECO* provides slightly worse overall performance in terms of thermal comfort compared to the *Fixed Set Point* strategy, mainly due to cold days, as discussed previously. While a portion of the energy savings equal to 16% can be attributed to the slightly higher PMV value, it is worth noting that, considering the overall amount of energy saved, a better tradeoff between thermal comfort and energy consumption is achieved.

In heating mode, we provide an additional analysis to compare the two approaches on Mondays, which is a particular case in terms of thermal comfort and energy consumption. This is because the initial environmental conditions in the early morning are often considerably far from the desired range due to inactivity during the weekend, thereby requiring the HVAC devices to operate more intensively throughout the day to ensure the desired comfort level. Although both the approaches are not able to address this task, providing out-of-comfort conditions, the *Fixed Set Point* achieves an average PMV value closer to the upper bound of the range than our solution, as reported in Table 8. On the other hand, *EECO* provides slightly worse environmental conditions but reduces the energy footprint of the building by more than 40 kWh on average, as underlined by the results reported in Table 9.

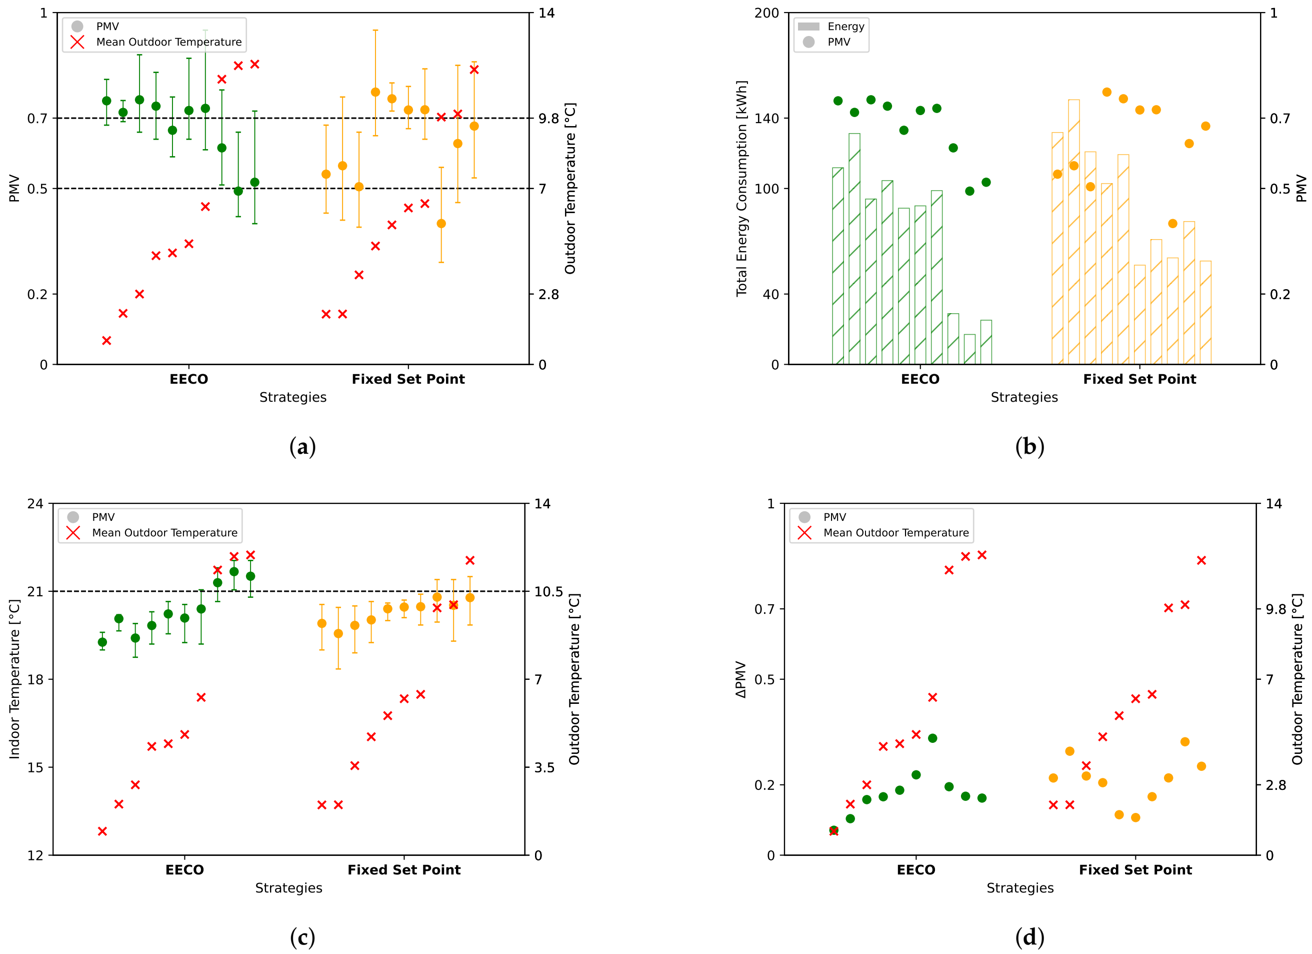

Figure 8. Average daily (**a**) PMV index, (**b**) total daily energy consumption, (**c**) indoor temperature, and (**d**) variation in PMV index compared to value at 6 a.m. (PMV_{6AM}-PMV) for both EECO and the Fixed Set Point approach on working days in heating mode.

Table 7. Overall performance of *EECO* and the *Fixed Set Point* approach in terms of PMV and energy consumption in heating mode.

	PMV	Energy Consumption
EECO	0.67	79 kWh
Fixed Set Point	0.63	95 kWh
Difference	**−6%**	**16%**

Table 8. Results of the *Fixed Set Point* approach during Mondays.

Days	PMV	PMV 6 a.m.	Indoor Temperature	Outdoor Temperature	Energy Consumption
Day 1	0.75	1.04	19.1 °C	2.8 °C	204 kWh
Day 2	0.77	1.07	19.2 °C	4.5 °C	125 kWh
Day 3	0.74	1.18	19.6 °C	9.1 °C	122 kWh
Day 4	0.77	1.14	20.6 °C	10.3 °C	82 kWh
Average	**0.76**	**1.11**	**19.6 °C**	**6.7 °C**	**133 kWh**

Table 9. Results of *EECO* during Mondays.

Days	PMV	PMV 6 a.m.	Indoor Temperature	Outdoor Temperature	Energy Consumption
Day 1	0.81	1.03	19.1 °C	4.3 °C	84 kWh
Day 2	0.89	1.28	19.7 °C	4.8 °C	154 kWh
Day 3	0.82	1.09	19.9 °C	6.6 °C	61 kWh
Day 4	0.72	1.12	20.9 °C	12.5 °C	51 kWh
Average	**0.81**	**1.13**	**19.9 °C**	**7.1 °C**	**87 kWh**

6.4. Performance Comparison: Simulated Environment

The validation in real-world scenarios makes it difficult to compare our solution with multiple approaches due to different reasons. Firstly, we are forced to complete experiments only from Monday to Friday as the building is closed in the weekends, thus requiring many weeks to collect a good amount of results for each approach. Secondly, even with a great availability of experimental data, it is difficult to compare more strategies over multiple days due to different outdoor weather conditions.

Considering the above limitations, we propose a comparison analysis among our solution, the *Fixed Set Point* approach (in this case, we test it with two different set point values) and the greedy *PMV-based* approach through software simulations. We simulate the behaviour of each strategy throughout each selected day in order to simulate the response of the environment upon the specific HVAC configuration at each quarter of an hour. We leverage an AI model, which we call *Global Model*, based on 1D CNN trained by using all the data collected from the warehouse during the evaluated operating mode (from June to October in cooling mode and from November to February in heating mode). The idea is to include as much information as possible to make the *Global Model* learn the response of the environment under all the different situations (e.g., due to outdoor weather), thus maximising the accuracy in the simulation.

In terms of results, we first demonstrate the reliability of the simulated environment by emulating the HVAC settings (i.e., ON/OFF and SP) occurring in the real environment during the day by means of the designed *Global Model*. Then, we simulate the three different approaches i.e., *EECO*, *PMV-based*, and *Fixed Set Point*. Consistent with the previous analysis in the real environment, we consider *category III* as the target comfort, i.e., $0.5 <= \text{PMV} < 0.7$. For a thorough comparison, we examine days in both cooling and heating modes, selecting eight random days for each month under consideration. Furthermore, we split the results between normal working days and Mondays, as in Section 6.3.2. Finally, as observed in the manual approach outlined in Section 6.3, both for the *PMV-based* and *Fixed Set Point* approaches, we activate the HVAC devices at 6 a.m., thus ensuring the same operating interval of our solution.

6.4.1. Simulator Validation

In this subsection, we aim to validate the simulated environment. In this regard, we analyse the accuracy of the *Global Model* mentioned in Section 6.4. To cope with this task, we analyse the behaviour of the predicted environment (indoor temperature, humidity, CO_2, and energy consumption) at each quarter of an hour upon the actuations (ON/OFF, SP) that occurred in the real environment. Basically, each time, we rely on the previous predictions to forecast the future evolution of the environment for the next quarter of an hour. This approach enables us to evaluate the reliability of the simulator by comparing the simulated behaviour of the warehouse throughout each chosen day with the real one.

In Tables 10 and 11, we show the overall monthly results of the simulated environment during the summer (cooling mode) and the winter months (heating mode), respectively. To provide a detailed overview of the accuracy of our simulated environment, for each variable, we report the root mean squared error (RMSE) between the predicted and real values as well as the percentage error (i.e., the percentage difference between the mean of the predicted values and the real ones). Unlike the percentage error used by Mancini et al. [35] in a similar scenario, the RMSE enables us to highlight possible deviations in the behaviour of each variable.

Table 10. Simulator results during summer months.

	Energy [kWh]		Indoor Temperature [°C]		Indoor Humidity [%]		CO_2 [ppm]	
Months	**RMSE**	**% Error**	**RMSE**	**% Error**	**RMSE**	**% Error**	**RMSE**	**% Error**
July	0.42	10.9	0.43	1.1	2.0	3.1	27.7	3.2
August	0.37	11.9	0.40	1.0	2.0	2.6	24.2	2.6
September	0.36	36	0.34	0.9	1.9	2.7	47.7	4.5
Average	**0.38 kWh**	**19.6%**	**0.39 °C**	**1.0%**	**2.0%**	**2.8%**	**33.2 ppm**	**3.4%**

Table 11. Simulator results during winter months.

	Energy [kWh]		Indoor Temperature [°C]		Indoor Humidity [%]		CO_2 [ppm]	
Months	**RMSE**	**% Error**	**RMSE**	**% Error**	**RMSE**	**% Error**	**RMSE**	**% Error**
December	0.34	22.1	0.32	1.1	1.4	2.5	35.3	3.6
January	0.33	10.2	0.29	1.1	1.3	2.3	25.0	2.4
February	0.31	11.9	0.49	1.9	1.3	2.8	38.8	3.0
Average	**0.33 kWh**	**14.7%**	**0.37 °C**	**1.3%**	**1.3%**	**2.5%**	**33.0 ppm**	**3.0%**

To the best of our knowledge, few research works in the literature demonstrate the goodness of their simulated environments in a systematic way. Nevertheless, we try to compare the accuracy of our simulator with the results of other works in the same research domain. Unlike the validation results presented by Mancini et al. [35], our simulated environment provides better performance in terms of percentage error for all the output variables (except for CO_2, which was not considered by the authors). However, even in our case, the energy consumption due to HVAC devices results in the most challenging variable to be simulated. In this regard, an average percentage error between 14% and 20% results in a slight underestimation or overestimation of the total predicted energy consumption compared to the real behaviour. However, the same amount of error is expected among all the evaluated approaches, guaranteeing a fair comparison between them. Nevertheless, it is worth noting that the RMSE value is limited for all the variables, thus demonstrating the capability of the simulator to simulate the environment throughout the day in an accurate way.

6.4.2. Cooling Mode

In cooling mode, we have completed experiments from July to September. As depicted in Figure 9, *EECO* provides slightly worse comfort conditions while reducing the energy footprint by at least 6 kWh on average compared to all the other evaluated approaches. In this scenario, the *Fixed Set Point* approach provides better thermal comfort conditions than requested but impacts the energy footprint of the building more than *EECO* and the *PMV-based* approach, as reported in Table 12. Such behaviour is particularly clear with a lower set point value (e.g., 26 °C), which forces the HVAC devices to work more, especially during very hot months such as July and August. Conversely, the greedy *PMV-based* approach achieves PMV values close to the lower bound of the comfort range but consumes a little bit more energy than *EECO*. Finally, comparable performance both in terms of thermal comfort and energy consumption is achieved during Mondays, as reported in Table 13.

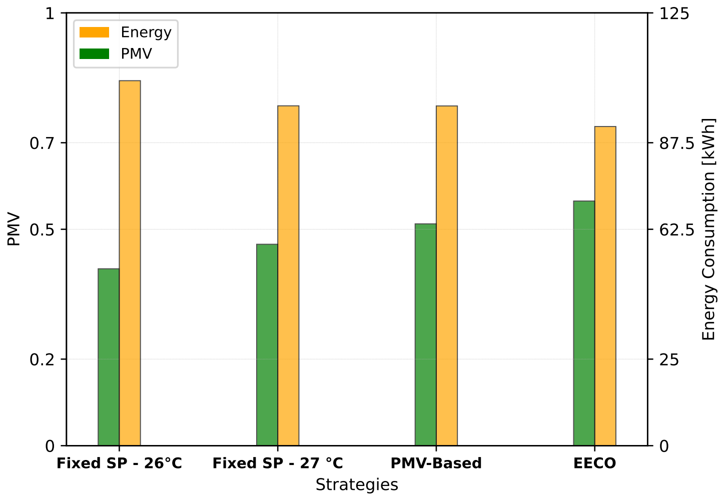

Figure 9. Overall average behaviour of PMV index and energy consumption for the evaluated approaches on working days.

Table 12. Average monthly behaviour of PMV index and energy consumption for the evaluated approaches on working days.

	Fixed SP-26 °C		**Fixed SP-27 °C**		**PMV-based**		**EECO**	
Months	**Energy**	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**
July	131 kWh	0.51	123 kWh	0.57	126 kWh	0.57	**117 kWh**	**0.65**
August	112 kWh	0.44	113 kWh	0.50	115 kWh	0.54	**105 kWh**	**0.63**
September	67 kWh	0.25	60 kWh	0.33	56 kWh	0.43	**56 kWh**	**0.42**
Average	*104 kWh*	*0.40*	*99 kWh*	*0.47*	*99 kWh*	*0.51*	***93 kWh***	***0.57***

Table 13. Average behaviour of PMV index and energy consumption for the evaluated approaches on Mondays.

Fixed SP-26 °C		**Fixed SP-27 °C**		**PMV-based**		**EECO**	
Energy	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**
114 kWh	*0.46*	*108 kWh*	*0.52*	*109 kWh*	*0.53*	***108 kWh***	***0.54***

To sum up the results from simulation data in cooling mode, Table 14 shows the improvement of *EECO* compared to the other approaches both from a comfort and energy perspective. In cooling mode, our solution guarantees 6% to 13% energy savings while providing slightly worse comfort conditions in terms of distance from the lower bound of the comfort range.

Table 14. Overall performance of *EECO* compared to the *Fixed Set Point* and *PMV-based* approaches in terms of absolute PMV difference from the lower bound of the comfort range and percentage difference in energy saving.

	PMV Distance	**Energy Consumption**
Fixed SP-26 °C	0.03	11%
Fixed SP-27 °C	−0.04	6%
PMV-based	−0.06	6%

6.4.3. Heating Mode

In heating mode, we have compared the different approaches in the period between December and February. In such a scenario, as depicted in Figure 10, our solution is capable of guaranteeing similar comfort conditions compared to the *Fixed Set Point* and *PMV-based* approaches but reducing the energy footprint by about 15 kWh on average, thus resulting in significant energy savings. This is due to some shutdowns of the HVAC devices configured by our approach at some quarters of hours during the day: this enables the environment to guarantee the comfort category but limits the impact on the energy consumption. On the other hand, such intelligence is not provided by the *Fixed Set Point* approach, which is prone to energy inefficiencies if the set point value is not manually set properly, as reported in Table 15. For instance, increasing the set point value from 21 °C to 22 °C does not improve the expected thermal comfort while impacting the building's energy footprint in a significant way. The same behaviour, as reported in Table 16, can be noticed on Mondays as well.

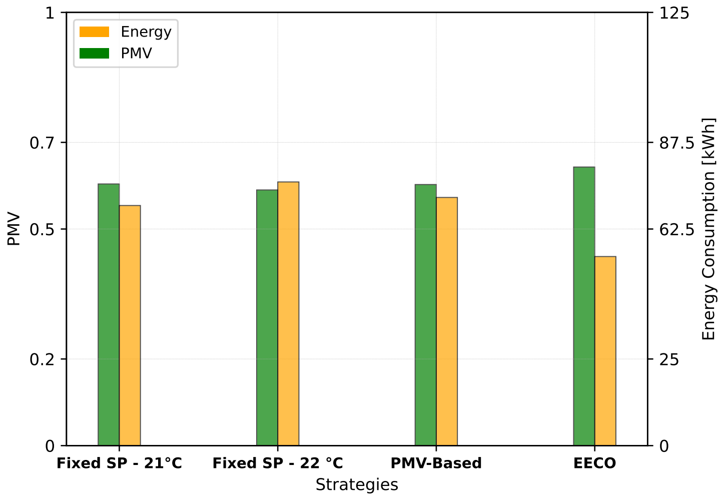

Figure 10. Overall average behaviour of PMV index and energy consumption for the evaluated approaches on working days.

As highlighted in cooling mode, better results in terms of PMV value are obtained during the last evaluated month (i.e., February). This is due to December and January being the coldest winter months. It is worth noting that heating operations within the environment prove to be challenging, as underlined by the PMV values obtained during the experiments.

Table 15. Average monthly behaviour of PMV index and energy consumption for the evaluated approaches on working days.

	Fixed SP-21 °C		**Fixed SP-22 °C**		**PMV-based**		**EECO**	
Months	**Energy**	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**
December	71 kWh	0.64	79 KWh	0.62	76 kWh	0.62	**56 kWh**	**0.68**
January	65 kWh	0.62	71 kWh	0.61	69 kWh	0.61	**48 kWh**	**0.67**
February	72 kWh	0.54	79 kWh	0.53	69 kWh	0.57	**60 kWh**	**0.58**
Average	*69 kWh*	*0.60*	*76 kWh*	*0.59*	*72 kWh*	*0.60*	***55 kWh***	***0.64***

Table 16. Average behaviour of PMV index and energy consumption for the evaluated approaches on Mondays.

Fixed SP-21 °C		**Fixed SP-22 °C**		**PMV-based**		**EECO**	
Energy	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**	**Energy**	**PMV**
85 kWh	*0.79*	*92 kWh*	*0.78*	*90 kWh*	*0.78*	***75 kWh***	***0.83***

To summarise, as reported in Table 17, in heating mode, *EECO* guarantees roughly the same comfort level provided by basic approaches, i.e., *PMV-based* and *Fixed Set Point*. However, unlike these strategies, our solution achieves an amount of energy saving greater than 20%. This confirms that these solutions have a limited overview of the problem as they just follow a single objective function with no attention for the energy footprint. In this regard, an intelligent approach such as *EECO* can guarantee a proper tradeoff between the thermal comfort and energy consumption. It worth noting that such simulated results confirm the results obtained in the real environment and described in Section 6.3.2.

Table 17. Overall performance of *EECO* compared to the *Fixed Set Point* and *PMV-based* approaches in terms of absolute PMV difference from the lower bound of the comfort range and percentage difference in energy saving.

	PMV Distance	Energy Consumption
Fixed SP-21 °C	−0.04	20%
Fixed SP-22 °C	−0.05	28%
PMV-based	−0.04	24%

7. Discussion

7.1. Comfort Model

Despite the benefits described in the previous sections, the proposed solution is affected by some limitations. For instance, some parameters of the PMV index (i.e., metabolic rate, clothing insulation, and air velocity) have been configured in a static way. However, some research works recommend dynamically adjusting comfort model parameters (e.g., air velocity [7] or clothing insulation [36]) in response to local environmental conditions. Additionally, in the literature, some studies [8,9] underline that feedback from building occupants takes a central role in meeting the requirements of a large number of people and accurately modelling the comfort.

In this study, our primary emphasis has centered around refining the algorithm responsible for the management of HVAC devices. While we delve into an extensive examination of EECO performance using static PMV parameter values, we acknowledge that the usage of dynamic values is beyond the current scope and does not impact the designed methodology. Nevertheless, we consider these aspects as opportunities for further investigation and exploration in the future.

7.2. Methodology

The proposed algorithm optimises a single comfort objective as the input. However, different stakeholders might be present within the environment (e.g., local personnel or customers), potentially leading to conflicting comfort requirements. In such cases, a decision is needed because the algorithm lacks the capability to address multiple comfort objectives simultaneously. In this regard, this solution particularly fits scenarios where occupants have similar comfort needs or a preference for one stakeholder over the other is exhibited. As a result, the solution's applicability might be better suited to scenarios with a less diverse environment of stakeholders.

Further improvements could also be integrated into the proposed solution; e.g., additional input variables could be considered. In this regard, different research studies in the literature include solar irradiation in their solutions [9,12]. Such information, integrated into the proposed algorithm, might allow for a more accurate selection of the HVAC configuration. This could indeed fine-tune the use of natural effects (passive methods) instead of activating HVAC devices at certain times during the day.

7.3. Validation

We underline the possibility to validate the designed solution from other perspectives. Indeed, our partner provided us strict constraints to test the proposed solution within a real production environment that includes human presence, preventing us from performing different experiments. In this regard, for instance, specific tests need to be conducted to evaluate the performance on days characterized by alternating comfort and no comfort requirements (e.g., similar to the approach proposed by Yang et al. [6]) as opposed to solely focusing on a single comfort interval throughout the day. Finally, it is fundamental to evaluate the effectiveness of the proposed solution across a number of buildings, encompassing different sizes, layout, materials, and potentially featuring multiple distinct environments, as analysed by Ascione et al. [10]. Additionally, the possibility to scale over multiple sites with minimal effort needs to be validated. In contrast to the existing literature, which often relies on complex physical or mathematical models tailored to the evaluated environment, our solution is expected to be rapidly deployed (with only a few parameters to configure) and adaptive to changes in terms of layout, HVAC machinery, outdoor weather, etc. as it continuously learns from the environment.

8. Conclusions

In this paper, we have presented an automated solution that leverages AI to continuously regulate HVAC devices with the aim to optimise comfort while minimising the energy footprint. It does not require any preliminary information of the local environment or any physical or mathematical modelling. Through the collected data, it implicitly evaluates the effect of different agents, including building features (e.g., wall thickness, orientation, and window presence) and passive phenomena (e.g., passive heating) on the monitored parameters, thus adapting to the observed environment.

We have tested our approach in a real warehouse of a small production plant belonging to an Italian retail company. Compared to a static approach where the HVAC set point is fixed at a specific temperature, the evaluation results in the real environment show that our solution can slightly improve the indoor comfort with minimal impact on the building's energy footprint in summer (i.e., cooling mode). On the other hand, during cold months (i.e., heating mode), it achieves higher energy savings (up to approximately 16%) while providing slightly worse comfort conditions.

Due to clear limitations in comparing multiple approaches in a real environment, we have provided an additional comparison analysis based on software simulations between our solution and two other approaches (i.e., the fixed set point and a greedy PMV-based approach). In this regard, the simulated results show significant improvements during the winter months compared to the summer period, confirming the results obtained in the real-world scenario. Indeed, the simulations show slightly reduced performance in terms of comfort requirements but underline substantial energy savings (exceeding 20%). Despite the promising results in our evaluated scenario, the application of our solution on a large scale is subject to overcoming some limitations mentioned in the previous section. Nevertheless, in contrast to non-intelligent approaches that follow a single objective function, the obtained results demonstrate the capability of our solution to guarantee a tradeoff between the comfort level and the energy consumption by dynamically selecting the configuration (ON/OFF and set point) of the HVAC devices.

Author Contributions: Conceptualization, G.S., C.P., M.G., and D.S.; methodology, G.S., R.D.-C., and D.S.; software, G.S.; validation, G.S., R.D.-C., M.G., and D.S.; formal analysis, G.S., R.D.-C., and D.S.; investigation, G.S., R.D.-C., and D.S.; resources, G.S.; data curation, G.S., R.D.-C., and D.S.; writing—original draft preparation, G.S.; writing—review and editing, G.S., R.D.-C., M.G., and D.S.; visualization, G.S. and M.G.; supervision, C.P. and D.S.; project administration, C.P. and D.S.; funding acquisition, C.P. All authors have read and agreed to the published version of the manuscript.

Funding: The work in this paper has been partially funded by "Programma operativo FESR 2014–2020" of Provincia autonoma di Trento, with co-financing of European ERDP funds, through project "GEM-Retail", by the European Union's Horizon 2020 Research and Innovation Programme under grant agreement no. 872548 (DIH4CPS: Fostering DIHs for Embedding Interoperability in Cyber-Physical Systems of European SMEs) and by the European Union's Horizon 2020 Research and

Innovation Programme under grant agreement no. 951771 (EUHubs4Data: European Federation Of Data Driven Innovation Hubs).

Data Availability Statement: Data will be made available on request due to privacy issues.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Parliament , T.E. Directive (EU) 2018/844 of the European Parliament and of the Council of 30 May 2018 a.m.ending Directive 2010/31/EU on the energy performance of buildings and Directive 2012/27/EU on energy efficiency (Text with EEA relevance), 2018, OJ L 156, 19.6.2018, p. 75–91. Available online: https://eur-lex.europa.eu/eli/dir/2018/844/oj (accessed on 14 March 2022).
2. BREEAM. Available online: https://www.breeam.com/ (accessed on 14 March 2022).
3. LEED. Available online: https://www.usgbc.org/leed (accessed on 14 March 2022).
4. Liu, S.; Kwok, Y.T.; Lau, K.K.L.; Ouyang, W.; Ng, E. Effectiveness of passive design strategies in responding to future climate change for residential buildings in hot and humid Hong Kong. *Energy Build.* **2020**, *228*, 110469. . [CrossRef]
5. De Araujo Passos, L.A.; van den Engel, P.; Baldi, S.; De Schutter, B. Dynamic optimization for minimal HVAC demand with latent heat storage, heat recovery, natural ventilation, and solar shadings. *Energy Convers. Manag.* **2023**, *276*, 116573. . [CrossRef]
6. Yang, S.; Wan, M.P.; Chen, W.; Ng, B.F.; Dubey, S. Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimization. *Appl. Energy* **2020**, *271*, 115147. . [CrossRef]
7. Wu, J.; Li, X.; Lin, Y.; Yan, Y.; Tu, J. A PMV-based HVAC control strategy for office rooms subjected to solar radiation. *Build. Environ.* **2020**, *177*, 106863. . [CrossRef]
8. Chen, X.; Wang, Q.; Srebric, J. Model predictive control for indoor thermal comfort and energy optimization using occupant feedback. *Energy Build.* **2015**, *102*, 357–369. . [CrossRef]
9. Zhang, H.; Tzempelikos, A.; Liu, X.; Lee, S.; Cappelletti, F.; Gasparella, A. The impact of personal preference-based thermal control on energy use and thermal comfort: Field implementation. *Energy Build.* **2023**, *284*, 112848. . [CrossRef]
10. Ascione, F.; Bianco, N.; De Stasio, C.; Mauro, G.M.; Vanoli, G.P. Simulation-based model predictive control by the multi-objective optimization of building energy performance and thermal comfort. *Energy Build.* **2016**, *111*, 131–144. . [CrossRef]
11. Valladares, W.; Galindo, M.; Gutiérrez, J.; Wu, W.C.; Liao, K.K.; Liao, J.C.; Lu, K.C.; Wang, C.C. Energy optimization associated with thermal comfort and indoor air control via a deep reinforcement learning algorithm. *Build. Environ.* **2019**, *155*, 105–117. . [CrossRef]
12. Gao, G.; Li, J.; Wen, Y. DeepComfort: Energy-Efficient Thermal Comfort Control in Buildings Via Reinforcement Learning. *IEEE Internet Things J.* **2020**, *7*, 8472–8484. [CrossRef]
13. Manjarres, D.; Mera, A.; Perea, E.; Lejarazu, A.; Gil-Lopez, S. An energy-efficient predictive control for HVAC systems applied to tertiary buildings based on regression techniques. *Energy Build.* **2017**, *152*, 409–417. . [CrossRef]
14. Martell, M.; Rodríguez, F.; Castilla, M.; Berenguel, M. Multiobjective control architecture to estimate optimal set points for user comfort and energy saving in buildings. *ISA Trans.* **2020**, *99*, 454–464. [CrossRef] [PubMed]
15. Halhoul Merabet, G.; Essaaidi, M.; Ben Haddou, M.; Qolomany, B.; Qadir, J.; Anan, M.; Al-Fuqaha, A.; Abid, M.R.; Benhaddou, D. Intelligent building control systems for thermal comfort and energy-efficiency: A systematic review of artificial intelligence-assisted techniques. *Renew. Sustain. Energy Rev.* **2021**, *144*, 110969. [CrossRef]
16. Yang, R.; Wang, L. Multi-objective optimization for decision-making of energy and comfort management in building automation and control. *Sustain. Cities Soc.* **2012**, *2*, 1–7. [CrossRef]
17. Wang, N.; Fang, F.; Feng, M. Multi-objective optimal analysis of comfort and energy management for intelligent buildings. In Proceedings of the 26th Chinese Control and Decision Conference (2014 CCDC), Changsha, China, 31 May–2 June 2014; pp. 2783–2788. [CrossRef]
18. Wahid, F.; Fayaz, M.; Aljarbouh, A.; Mir, M.; Aamir, M. Energy Consumption Optimization and User Comfort Maximization in Smart Buildings Using a Hybrid of the Firefly and Genetic Algorithms. *Energies* **2020**, *13*, 4363. . [CrossRef]
19. Ngarambe, J.; Yun, G.Y.; Santamouris, M. The use of artificial intelligence (AI) methods in the prediction of thermal comfort in buildings: Energy implications of AI-based thermal comfort controls. *Energy Build.* **2020**, *211*, 109807. . [CrossRef]
20. Yau, Y.; Chew, B. A review on predicted mean vote and adaptive thermal comfort models. *Build. Serv. Eng. Res. Technol.* **2014**, *35*, 23–35. [CrossRef]
21. Ul Haq Gilani, S.I.; Khan, M.H.; Pao, W. Thermal Comfort Analysis of PMV Model Prediction in Air Conditioned and Naturally Ventilated Buildings. *Energy Procedia* **2015**, *75*, 1373–1379. Clean, Efficient and Affordable Energy for a Sustainable Future: The 7th International Conference on Applied Energy (ICAE2015). [CrossRef]
22. Kiranyaz, S.; Avcı, O.; Abdeljaber, O.; Ince, T.; Gabbouj, M.; Inman, D.J. 1D Convolutional Neural Networks and Applications: A Survey. *arXiv* **2019**, arXiv:1905.03554.
23. Fanger, P. *Thermal Comfort: Analysis and Applications in Environmental Engineering*; Danish Technical Press: Copenhagen, Denmark, 1970.

24. OpenWeatherMap. Available online: https://openweathermap.org/ (accessed on 21 September 2022).
25. Segala, G.; Doriguzzi-Corin, R.; Peroni, C.; Gazzini, T.; Siracusa, D. A Practical and Adaptive Approach to Predicting Indoor CO_2. *Appl. Sci.* **2021**, *11*, 10771. [CrossRef]
26. Kingma, D.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2014**, arXiv:1412.6980.
27. Tartarini, F.; Cheung, C.T.; Hoyt, T. CBE Thermal Comfort Tool: Online tool for thermal comfort calculations and visualizations. *SoftwareX* **2020**, *12*, 100563. [CrossRef]
28. pythermalcomfort. Available online: https://pythermalcomfort.readthedocs.io/en/latest/readme.html (accessed on 21 September 2022).
29. Guo, H.; Aviv, D.; Loyola, M.; Teitelbaum, E.; Houchois, N.; Meggers, F. On the understanding of the mean radiant temperature within both the indoor and outdoor environment, a critical review. *Renew. Sustain. Energy Rev.* **2020**, *117*, 109207. [CrossRef]
30. ISO. ISO 7730. Available online: https://www.iso.org/standard/39155.html (accessed on 23 August 2023).
31. *CEN EN Standard 16798-1*. Indoor Environmental Input Parameters for Design and Assessment of Energy Performance of Buildings Addressing Indoor Air Quality; Thermal Environment, Lighting and Acoustics—Module M1-6. CEN: Brussels, Belgium , 2019.
32. Ala'raj, M.; Radi, M.; Abbod, M.F.; Majdalawieh, M.; Parodi, M. Data-driven based HVAC optimisation approaches: A Systematic Literature Review. *J. Build. Eng.* **2022**, *46*, 103678. .: 10.1016/j.jobe.2021.103678. [CrossRef]
33. Modbus. Available online: https://modbus.org/ (accessed on 21 September 2022).
34. Intel. Intel NUC. Available online: https://www.intel.com/content/www/us/en/products/sku/122498/intel-nuc-kit-nuc7i3dnhe/specifications.html (accessed on 8 August 2023).
35. Mancini, F.; Nardecchia, F.; Groppi, D.; Ruperto, F.; Romeo, C. Indoor Environmental Quality Analysis for Optimizing Energy Consumptions Varying Air Ventilation Rates. *Sustainability* **2020**, *12*, 482. [CrossRef]
36. Schiavon, S.; Lee, K.H. Dynamic predictive clothing insulation models based on outdoor air and indoor operative temperatures. *Build. Environ.* **2013**, *59*, 250–260. . [CrossRef]