UNIVERSITÀ
DEGLI STUDI
FIRENZE

DOCTORAL PROGRAMME IN INDUSTRIAL ENGINEERING

Energy and Innovative Industrial and Environmental Technologies

XXXV

Developing of machine learning-based energy monitoring methodologies for the building energy demand of healthcare facilities: an Italian case study

ING/IND-09

Doctoral Candidate

Marco Zini

Supervisor

Prof. Carlo Carcasci

Coordinator

Prof. Giovanni Ferrara

Years 2019/2022

Abstract

The field of energy efficiency in buildings is gaining increasing interest from the scientific and industrial communities. Due to the recent instability in the energy market and the increased consciousness of environmental and climate related concerns, companies are increasingly interested in finding new ways to develop strategies to monitor and optimise the consumption of complex systems such as commercial or residential buildings. On the one hand, renewable energy communities are being studied to reverse the producer-consumer paradigm aiming to pursue a rational synergy between energy production and its uses. On the other hand, the optimization of buildings energy consumption could lead to significant benefits for the overall energy demand and a marked reduction in energy wastes.

In this context, numerical energy simulation techniques play a crucial role in achieving a clear understanding of the energy performance of buildings. The difficulty in applying innovative solutions and the need to take economically viable decisions, make the dynamic energy simulation models an excellent tool for obtaining insights and predictions regarding the effectiveness of new system solutions or advanced management strategies. On the other hand, optimal management of complex systems such as commercial buildings can be a challenging task, requiring constant monitoring and verification of the proper system operations. There is therefore the need to develop numerical methods that can be directly implemented in centralized energy management systems, being capable to offer support to operators for constant monitoring of building energy demand. Automated diagnostic methods able to detect deviations or mild and progressive decrease in energy performance results crucial for buildings equipped with sophisticated energy systems.

Numerical tools for the thermodynamic simulation of buildings require a large amount of data on the building envelope and systems, which may not always be available in real-world contexts. Therefore, it is necessary to reverse the traditional view of energy modelling, moving beyond the realization of a building model starting from information on its configuration, and instead starting from the final effect: the energy requirement. In this context, machine learning-based techniques offer an excellent opportunity to create a mathematical model that can predict the energy demand of a building based on various influential parameters: the so-called Energy Drivers. However, the biggest drawback of

these methodologies lies in the need to have users with a decent level of experience and knowledge both in the energy field and in the field of Machine Learning, that is not always available in an industrial or commercial reality (Especially for those activities whose core business is completely different).

The aim of this work is to create a systematic methodology to proceed from the raw data to a ready-to-implement monitoring method, minimizing the knowledge required to the user and, contemporary, ensure reliable performance of the monitoring method itself, starting from methods with low computational effort that require high user experience, to method that requires significantly higher computational resources, but aiming to minimize the required background knowledge to be applied. The proposed monitoring methods are based on machine learning models able to predict the energy demand. This work explores several machine learning algorithms (Multiple Linear Regression and Artificial Neural Network), evaluating different methods to perform the Feature Selection and Hyperparameters Tuning. The models prediction are exploited to perform statistical residual analyses through the well-known technique of cumulative sum of differences.

This work was developed in collaboration with *Casa di Cura Villa Donatello S.p.A.* (located in Sesto Fiorentino, Florence, Italy), which made available its healthcare facility and all the data about energy demand and systems operations in order to reach the goal of increasing the energy efficiency of the whole building. This study analyses the results obtained by means of six procedures designed to offer different combinations of prior knowledge and computational cost necessary to go from raw building energy consumptions data to an implementable monitoring method, and is therefore also intended as a "map" that allows the end user to orient themselves on the methodology to follow based on their specific needs.

The six procedures have been applied on whole building electricity demand data, being able to define the procedure which allow to obtain the best prediction performance minimizing the knowledge required by the final user.The best procedure is then applied also to more detailed data about the electricity consumptions of Heating, Ventilation and Air Conditioning system installed in the facility. As a result, the obtained models will be exploited to perform the monitoring of the whole building and detailed systems energy consumptions, obtaining insight of the effective usefulness of the implementation of this methods in a real activity context.

Chapters organization

The present thesis work is divided into four main chapters. In chapter 1, an introduction is provided on the energy consumption of hospital facilities, the main tools for conducting

energy audit are analyzed, and the topic of building energy monitoring is explored.

Chapter 2 provides a comprehensive background on the machine learning techniques investigated in this work, finally presenting the proposed systematic monitoring method.

Chapter 3 thoroughly analyzes the case study under examination, providing details regarding the construction, systems, climatic conditions, and finally analyzing the energy consumption at the whole-building level and for individual components of the HVAC system.

Chapter 4 is divided into three main sections. Section 4.1 is dedicated to the application of the proposed method to the energy consumption of the entire building. Each step of the method will be analyzed, establishing the best procedure in terms of prediction performance and computational cost.

This will then be applied to energy consumption data related to specific HVAC components, and the results will be presented in Section 4.2.

The models obtained and described in the previous sections will be used to complete the monitoring method, and the final results will be discussed and compared in Section 4.3, following a scientific article style.

Publications

Conference papers

- F. Balduzzi, M. Zini, G. Ferrara, and A. Bianchini, '*Development of a CFD Methodology to Reproduce the Effects of Macro Turbulence on Wind Turbines and its Application to the Particular Case of a VAWT*', presented at the ASME Turbo Expo 2019: Turbomachinery Technical Conference and Exposition, Nov. 2019, https://doi.org/10.1115/GT2019-90889.

- A. Bianchini, C. Carcasci, G. Manfrida, and M. Zini, '*Reconstruction and Analysis of the Energy Demand of a Healthcare Facility in Italy*', E3S Web Conf., vol. 197, p. 02009, 2020, https://doi.org/10.1051/e3sconf/202019702009.

- M. Zini and C. Carcasci, '*Developing of an Offline Monitoring Method for the Electric Energy Demand of a Healthcare Facility in Italy*', presented at the 16th Conference of Sustainable Development of Energy, Water and Environment Systems, 2021.

- M. Zini and C. Carcasci, '*Dynamic Building Energy Modelling of Healthcare Facilities: An Italian Case Study*', presented at the 17th conference of Sustainable Development of Energy, Water and Environment Systems, Paphos, Cyprus, Oct. 2022.

- M. Zini, C. Carcasci, and R. Sodini, '*Modelling and Optimization of a Hospital Gas Turbine-Based Cogeneration System*', presented at the ASME Turbo Expo 2022: Turbomachinery Technical Conference and Exposition, Oct. 2022, https://doi.org/10.1115/GT2022-80828.

- M. Zini and C. Carcasci, '*Analysis and monitoring of HVAC systems energy demand in an italian healthcare facility*', 18th Conference of Sustainable Development of Energy, Water and Environment Systems, 2023, under review.

International Journals

- F. Balduzzi, M. Zini, G. Ferrara, and A. Bianchini, '*Development of a Computational Fluid Dynamics Methodology to Reproduce the Effects of Macroturbulence on Wind Turbines and Its Application to the Particular Case of a VAWT*', J. Eng. Gas Turbines Power, vol. 141, no. 11, Nov. 2019, https://doi.org/10.1115/1.4044231.

- F. Balduzzi et al., '*Understanding the Aerodynamic Behavior and Energy Conversion Capability of Small Darrieus Vertical Axis Wind Turbines in Turbulent Flows*', Energies, vol. 13, no. 11, Art. no. 11, Jan. 2020, https://doi.org/10.3390/en13112936.

- M. Zini and C. Carcasci, '*Developing of an Offline Monitoring Method for the Energy Demand of a Healthcare Facility in Italy*', [Journal of Sustainable Development of Energy, Water and Environment Systems], vol. [10], no. [4], p. [1]-[22], Dec. 2022, https://doi.org/10.13044/j.sdewes.d10.0421.

- M. Zini and C. Carcasci, '*Machine learning-based monitoring method for the electricity consumption of a healthcare facility in Italy*', Energy, p. 125576, Sep. 2022, https://doi.org/10.1016/j.energy.2022.125576.

- M. Zini, R. Sodini, and C. Carcasci, '*Modelling and Optimization of a Hospital Gas Turbine Based Cogeneration System*', Journal of Engineering for Gas Turbines and Power, Aug. 2022, https://doi.org/10.1115/1.4055418.

- M. Zini and C. Carcasci, '*Machine learning-based energy monitoring of buildings: development of a systematic method applied to an Italian case study*', Energy and Buildings, under review.

- M. Zini, S. Bellavia, C. Carcasci, S. Rebegoldi, '*Tuning-free Stochastic Optimization For Neural Network Training in Building Energy Prediction*', in preparation.

Contents

List of Figures

List of Tables

Nomenclature

Acronyms

ACH	Air Changes per Hour
$Adam$	Adaptive Moment Estimation
AHU	Air Handling Unit
ANN	Artificial Neural Network
BCL	Building Components Library
BEM	Building Energy Modelling
BTO	Building Technologies Office
$CAPEX$	Capital Expenditure
CAT	Computerized Axial Termography
CF	Cash Flow
CHP	Combined Heat and Power
COP	Coefficient of Performance
DCT	Dry Cooling Towers
ECT	Evaporative Cooling Towers
FE	Feature Engineering
FNN	Feedforward Neural Network
FS	Feature Selection
HLN	Number of neurons in hidden layer
$HVAC$	Heating, Ventilation and Air Conditioning

LR	Linear Regression
lr	Learning Rate
MAE	Mean Absolute Error
ML	Machine Learning
MLP	Multi Layer Perceptron
MLR	Multiple Linear Regression
MRI	Magnetic Resonance Imaging
MSE	Mean Squared Error
NPV	Net Present Value
$OPEX$	Operational Expenditure
$P2G$	Power to Gas
PBP	Payback Period
PI	Profit Index
PMV	Predicted Mean Vote
PV	Photovoltaic
R^2	Coefficient of Determination
$RACHP$	Refrigeration, Air Conditioning and Heat Pump
$ReLU$	Rectifier Linear Unit
$RMSE$	Root Mean Squared Error
SGD	Stochastic Gradient Descent
SVD	Singular Value Decomposition
UPS	Uninterruptible Power Supply
$VCCC$	Ventilation, Climate Control, Contaminant Control

Latin Symbols

CDD	Cooling Degree Days, HDD

E	Electrical energy, kWh
e	Residual
h	enthalpy, kJ/kg
HDD	Heating Degree Days, HDD
HDh	Heating Degree hours, HDh
HR	Humidity Ratio, -
I	Illuminance, Lx
p	Pressure, bar
$Prcp$	Precipitation, mm
Q	Thermal Power, kW
q	Thermal energy, kWh
RH	Relative Humidity, $\%$
S	Cumulated sum of differences
SR	Solar Radiation, W/m^2
T	Temperature, $°C$
$tanh$	Hyperbolic tangent
v	Water volume, m^3
W	Electrical power, kW
WD	Wind direction, $°$
WS	Wind speed, m/s

Greek Symbols

α	Azimuthal angle
ρ	Spearman's rank correlation coefficient
σ	Standard deviation
τ_i	Irrelevant features threshold

τ_r	Redundant features threshold
θ	Tilt angle

Subscripts and Superscripts

0	Conventional
c	Current
$cool$	Cooling
db	Dry bulb
e	External, Outdoor
el	Electrical
max	Maximum
min	Minimum
p	Peak
$ptod$	Production
sim	Simulation, Simulated
th	Thermal
amb	Ambient
$diff$	Diffusive
dp	Dew point
gl	Global
wb	Wet bulb

Chapter 1

Introduction

The investigation of energy usage in Europe has been a crucial area of interest among both scholars and policymakers in contemporary times. According to Balaras et al. [1], the total final energy consumption in the *EU-27* member states in *2019* was 937.9 million tons of oil equivalent (Mtoe). One of the key drivers of energy consumption in Europe is the regional heavy reliance on fossil fuels. In *2019*, the primary energy consumption from fossil fuels (including oil, gas, and coal) accounted for 79.1% of the total primary energy consumption in the *EU-27* member states [2]. However, there has been a growing focus on reducing Europe dependence on fossil fuels and increasing the use of renewable energy sources. The share of renewable energy in gross final consumption of energy increased from 17.5% in *2015* to 18.9% in *2019*.

Although the current energy demand is not negligible, some projections anticipate a marked increase in demand peaks and, in general, in total energy demand by *2050*. A report by *Terna* [3] report the Italian electricity demand trend over the years referring to the amount of the electricity consumed by the end users, plus national grid losses.

Figure 1.1 shows the evolution of electricity demand in *2040* for the italian *National Integrated Plan for Energy and Climate (PNIEC)* scenario [4] for the years *2025*, *2030*, and *2040*. The indicated variability is relative to the different values resulting from the application of the climatic years in the *1982-2016* interval. Note that the *2040* data does not include electricity consumption for hydrogen production through power-to-gas plants (estimated at 12 TWh).

The graph in Figure 1.2 shows the evolution of the electric power demand peaks. For the forecasting horizon years, the data was obtained through elaborations by *Terna*, using the technology diffusion data reported in the National Integrated Plan for Energy and Climate. The peak is expected to increase in *2030* and *2040*, mainly due to the spread of technologies for heating and cooling environments, electric vehicles, and the increased electrification of consumption in general.

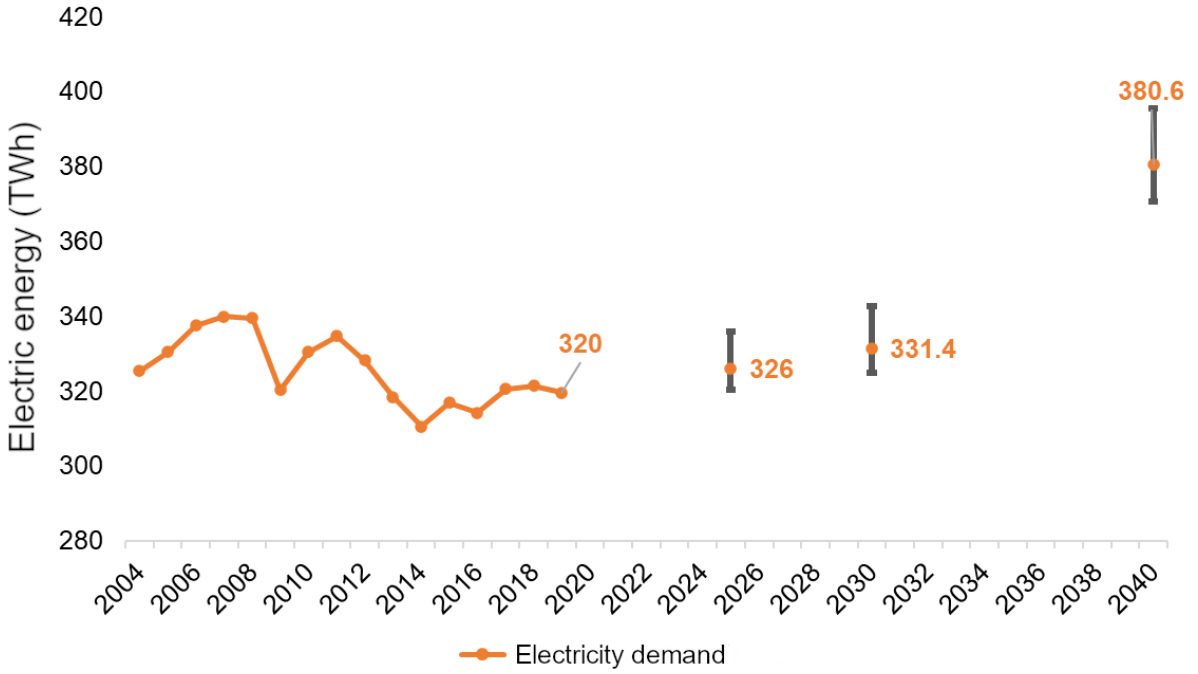

Figure 1.1: Italian electricity demand trend between 2004 and 2040

The annual gas demand in the *NT Italy scenario* [3] shows a decreasing trend in line with what is predicted by the *PNIEC*. By *2025*, the gas demand remains above 70 billion cubic meters (72.2 bcm) and the decrease in consumption is expected mainly after *2025*, reaching 62.3 billion cubic meters by *2030* and dropping to 60.6 billion cubic meters by *2040*. The decreasing trend in gas demand is indeed conditioned by both the energy efficiency measures foreseen in final uses and by the progressive penetration of renewable sources in electricity generation. The described gas demand is made up of both natural gas and biomethane, which, according to the *PNIEC*, is assumed to be about 1 billion cubic meters by *2030*, rising to about 7 billion cubic meters by *2040*.

In the *NT Italy scenario*, in addition to natural gas and biomethane, the quantities of hydrogen predicted by the *PNIEC* are also taken into account, which are equal to about 0.1 billions of cubic meters of methane equivalent by *2030* and will grow to about 3.9 bcm by *2040* (3.2 Mtoe). The predicted quantities of hydrogen can be obtained in various ways and can contribute both as pure hydrogen and as synthetic methane to the decarbonization of high energy-intensive industrial sectors and long-haul commercial transport. In particular, these values do not yet include the developments predicted by the recent "Guidelines for the National Hydrogen Strategy", which anticipates the development of hydrogen by about a decade compared to the *PNIEC* scenario assumed as a reference for the *NT Italy* scenario. Taking into account these quantities as well, the demand for gas and green gas is equal to 65 billion cubic meters by *2040*, as reported in Figure 1.3.

Energy consumption in buildings is a significant contributor to overall energy consumption in Europe. According to the European Commission (2022) [5], buildings account for

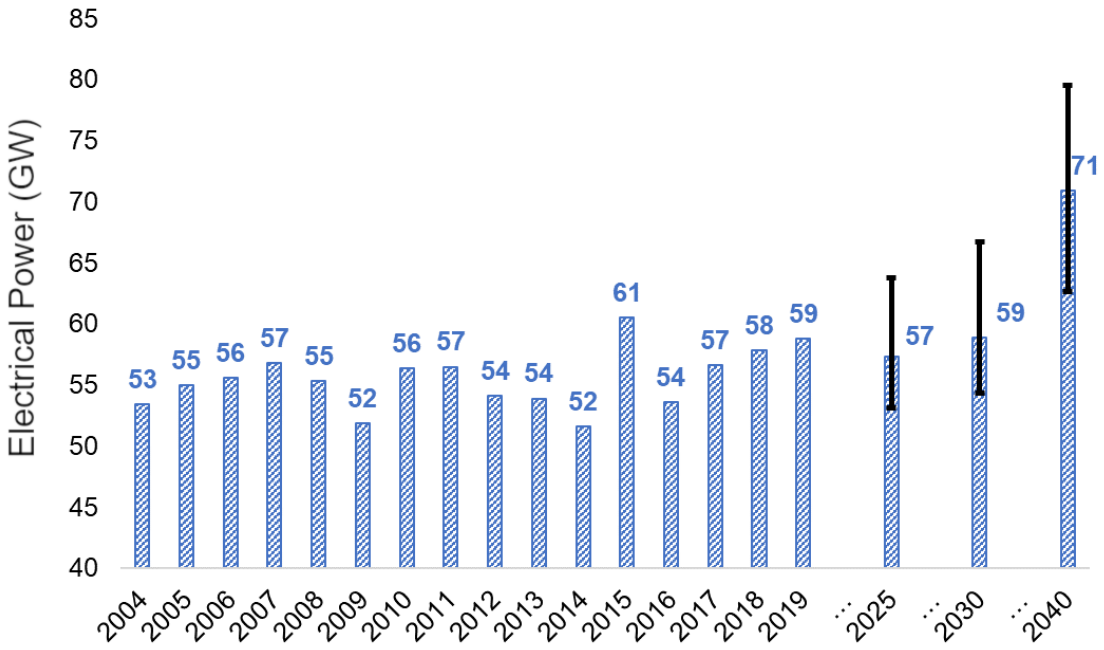

Figure 1.2: Italian peak electric power demand trend between 2004 and 2040

40% of the total energy consumption and 36% of the total CO_2 emissions in the *EU-27* countries. The residential sector is the largest consumer of energy in the building sector, accounting for 63% of the total energy consumption in buildings in the *EU*, while the commercial sector, which includes offices and other non-residential buildings, accounted for the remaining 37% of the total buildings energy consumption. Moreover, the majority of the buildings energy uses is related to the space heating (45%), followed by domestic hot water generation (18%) and lighting purposes (17%).

The European Union has set a target to increase the energy efficiency of buildings by at least 20% by 2020, as part of its efforts to meet the goals of the Energy Efficiency Directive (*2012/27/EU*) and the Paris Agreement on climate change [6]. To achieve this target, the *EU* has implemented a number of policy measures, including the *Energy Performance of Buildings Directive* (**2018/844/EU**), which requires member states to establish and implement minimum energy performance requirements for buildings, and the *Renewable Energy Directive* (**2018/2001/EU**), which sets targets for the use of renewable energy in the building sector. In addition to policy measures, new technologies and innovative design approaches are also playing a key role in reducing energy consumption in buildings.

Even if the main part of the buildings energy consumption is due to the residential buildings, also energy consumption in non-residential buildings is a significant contributor to overall energy consumption and greenhouse gas emissions. Heating, Ventilation and Air Conditioning systems *HVAC* are the main drivers of energy demand in non-residential buildings like office buildings, schools, and hospitals, accounting for approximately 40% of total energy consumption [7]. Consequently, the use of advanced building control systems

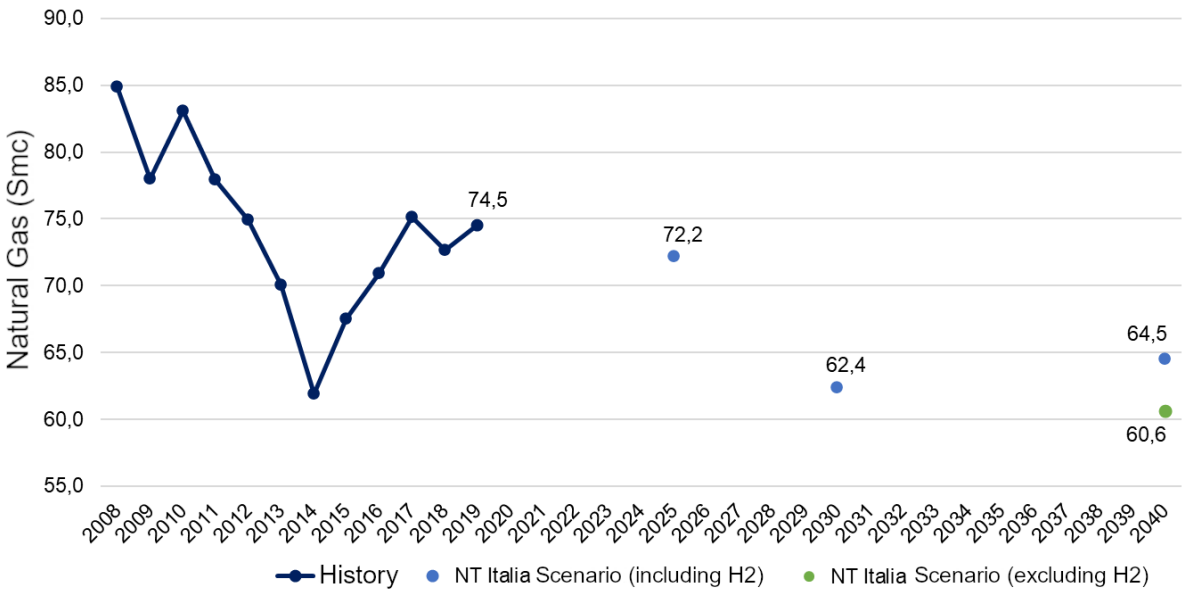

Figure 1.3: Italian natural gas demand trend between 2008 and 2040

can also lead to significant energy savings in non-residential buildings.

The economic crisis caused a sharp drop in electricity consumption, the most affected sector in Italy being undoubtedly the industry, which in *2008* consumed 151.4 TWh while in *2009*, due to the crisis, this value decreased to 130.5 TWh (13.8% less). The only sector that has consistently increased its demand for electricity in recent years is the tertiary sector, which, according to *Assoelettrica* data, in *2010* was worth 96.3 TWh (31% of national demand), while in *2012* it was responsible for consuming 101.1 TWh of electricity (32.9% of national demand). Also according to *Assoelettrica* estimates, the tertiary sector will see a 30 TWh increase in demand for electricity over the next 20 years. This growth will mainly come from the use of electronic and computer equipment and the increasing use of air conditioning in work environments during the summer months.

Figure 1.4 shows the trend of electricity consumption in the tertiary sector between 1990 and 2020. The sectors of the tertiary sector with the highest consumption are hotels, restaurants, communications, health and other social services. A significant percentage of the energy demand in the tertiary sector is due to the obsolescence of facilities and structures. It is often possible to improve the devices performance in these work environments with managerial interventions.

1.1 Healthcare Facilities

According to a report from the Ministry of Health, in *2010* there were 146 Local Health Companies distributed throughout the country. In *2010*, hospital care made use of 1′165 care institutions, 54% of which were public and the remaining 46% were accredited

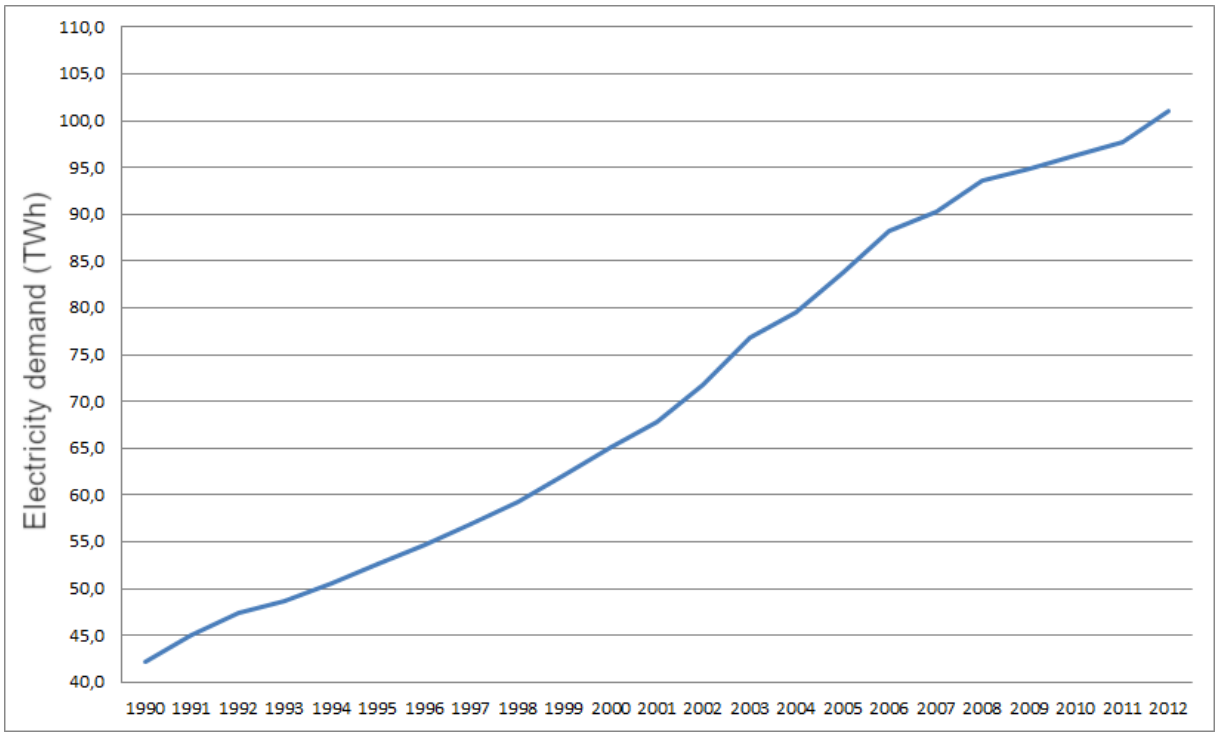

Figure 1.4: Tertiary sectorn electricity demand trend

private. 68% of public facilities were constituted by hospitals directly managed by Local Health Companies, 10% by Hospital Companies, and the remaining 22% by other types of public hospitals. The National Health Service had a more than 215′000 beds for ordinary hospitalization, 21% of which were in accredited private facilities, 21′761 beds for day hospital, almost entirely public (91%), and 8′230 beds for day surgery, predominantly public (80%).

Table 1.1 [8] shows the distribution of facilities by type of assistance provided (hospital, outpatient specialist, residential territorial, semi-residential territorial, other territorial assistance, and rehabilitation under Art. 26 of Law 833/78) and by nature (public and accredited private) in 2020.

Table 1.1: Distribution of facilities by type of assistance provided and by nature

Assistance	**Public (%)**	**Accredited priveate (%)**	**Total (#)**
Hospitals	51.4%	48.6%	1'004
Outpatient specialist	40.1%	59.9%	8'803
Residential territorial	16.8%	83.2%	7'858
Semi-residential territorial	28.3%	71.7%	3'189
Other territorial	86.9%	13.1%	5'557
Specialized ambulatories	40.1%	59.9%	8'803
Rehabilitation	21.8%	78.2%	1'151

The facilities counted are: $1'004$ for hospital care, $8'803$ for outpatient specialist care, $7,858$ for residential territorial care, $3,189$ for semi-residential territorial care, $5,557$ for other territorial care, and $1,151$ for rehabilitation care (under Art. 26 of Law 833/78). Regarding the nature of the facilities, hospital care is predominantly provided by public facilities (51.4%), and other territorial care is predominantly provided by public facilities (86.9%). Residential territorial care and semi-residential territorial care are predominantly provided by accredited private facilities (83.2% and 71.7%, respectively), and rehabilitation care under Art. 26 of Law 833/78 is predominantly provided by accredited private facilities (78.2%).

Among all the tertiary sector buildings, the hospitals and healthcare facilities energy demand is characterized by very specific behaviour due to the necessary service continuity of this kind of activity [9]. Among U.S. commercial buildings, hospitals and healthcare facilities present one of the higher energy requirements per square meter [10]. In European Union, healthcare facilities and hospitals represent 7% of the non-residential building and are responsible for about 10% of the total energy consumption in this sector [11].

In hospitals, energy is used in different sectors in the form of thermal energy (for both heating and cooling), and electrical energy. The complexity of the energy system of each structure can vary from hospital to hospital, and is closely related to:

- the health activities that take place in it
- its geographical location
- the type of structure and facility adopted

The use of different energy sources also leads to the use of different technological options, resulting in specific operating costs and different emission levels.

The thermal energy consumption in hospitals is often at low temperatures ($< 100\ ^{\circ}C$), required for heating the environments, compensating for heat losses from ventilation needs, and for producing domestic hot water for sanitary purposes. Another sector that exploits thermal energy, but at higher enthalpies, is the one for disinfection and sterilization. The need to ensure high levels of service even in case of sudden failures or emergencies requires having redundant energy production sources. To increase reliability, it is often present a connection between the steam system and the heating circuit, which however represents an energy loss point.

The cooled water produced is used for cooling the inpatient departments, the operating block, and the special departments, and is also used for cooling special medical equipments that produces a lot of heat or has specific indoor air conditions requirements. The production of cooled water for cooling environments is mainly related to the summer period, while the production dedicated to cooling equipment and maintaining certain

Table 1.2: Classification of Hospital energy usage

Energy	**Application**
Thermal	Sterilization, Medical equipment, Laundry Heating, Air Conditioning, Ventilation, Domestic hot water production, Kitchens
Electric	Ventilation and Air Conditioning, Electromedical equipment, Diagnostic, Monitoring, ICT, Lighting, Security,
Electric (Service continuity granted)	Surgery, Intensive Care, Lighting, Sterilization, Security Monitoring, Emergency department

thermohygrometric conditions in critical rooms is required in during the whole year, regardless to the season.

Hospital buildings are often subject to structural/organizational interventions and may undergo changes over time in the volume, use of various areas, or thermo-hygrometric conditions required by regulatory adjustments. Often the plant designed to meet the needs of the original structure may no longer be optimal, as it no longer corresponds to the new needs. Typically, compression refrigeration units are used for producing cooled water, which due to their high electrical demand, significantly contribute to the hospital or healthcare facility peak loads.

In addition to air conditioning, electrical energy provides the operation of medical equipment and other machinery like the circulation of medical gases. Other electrical consumption arises from lighting, elevators, department kitchens, sterilization and disinfection equipment.

Table 1.2 shows the energy uses in a hospital categorized by type of energy required. The increase in the complexity of equipment used and the adjustment of regulations in terms of thermo-hygrometric conditions and air changes per hour (ACH) also leads to further increases in energy consumption and greater difficulties in plant management.

To understand the nature of energy consumption in a hospital, it is useful to differentiate its environments based on their needs. This division categorizes the premises into several typologies which are reported below:

1. **High intensity healthcare activities**: refer to zones with high air exchange rates or high usage factors. The first subgroup includes operating rooms (temperatures

20-22 $°C$, at least 15 air changes per hour from outside), pathological anatomy (temperature 18 $°C$, 15 air changes per hour), and the delivery room (temperatures 22-24 $°C$, 6 air changes per hour). The second subgroup has lower (though still high if compared to other hospital areas) air exchange requirements, but their high usage factor justifies their significant energy consumption. Examples include intensive care units, which operate twenty-four hours a day, seven days a week. These areas have the highest energy consumption, which results from the required special environmental conditions and hygiene conditions. These areas are where a significant number of treatment plants with high air flow rates are typically installed.

2. **Diagnostic laboratories**: This group can be divided into two subgroups: *Imaging Diagnostics* and *Laboratory Diagnostics*. The first is structured for environments equipped with medium to high-tech biomedical equipment, and includes reporting rooms and medical clinics. The strong incidence of biomedical equipment is one of the major causes of both electrical and continuous cooling consumption, due to the latter's significant thermal dissipation and the need to maintain the environment in which they are installed at a maximum temperature of 22-24 $°C$ for the duration of the service's operating activity. These outpatient hospital areas are normally used ten hours per day and five days per week, while in emergency services they are operative continuously. The second type of diagnostics laboratories is also characterized by the presence of automatic analysers, which require electrical and thermal energy and, moreover a continuous cooling. These areas can be structured as open spaces or can be composed by several individual rooms.

3. **Inpatient departments**: These are the areas characterized by environments intended for the continuous stay of patients and healthcare personnel. During the heating period, temperatures must be maintained at 22 $°C$, continuously. The required air exchange is at least 2 vol/h, achievable through mechanical ventilation. Summer conditioning is required to ensure the comfort of those inside such environments, although in less recent hospital realities, there is not always a control of thermo-hygrometric conditions during the summer season or a ventilation system. Inpatient departments have dedicated infrastructure: medical gas, nurse call, data transmission network, visit lighting, safety lighting, FM and TV power supply. They do not have a particular supply of biomedical equipment. A significant consumption of electrical energy is due to lighting.

4. **Outpatient clinics and Ambulatories**: These are often homogeneous portions of buildings or entire buildings intended for daytime use. In these environments, a heating temperature of 20-22 $°C$ must be guaranteed for twelve hours a day

hours, during the workdays. They are intended for medical offices, examination rooms, rehabilitation therapy rooms (including areas such as gyms, swimming pools), characterized by the presence of mainly ambulatory patients and healthcare personnel. There is some therapeutic and diagnostic equipment, although not of significant impact in terms of electrical consumption and heat dissipation.

5. **Services**: This includes environments with a winter temperature regime of 18-20 $^{\circ}C$ to be maintained for eight hours per day during the workdays. The formers are areas designated for offices, classrooms, libraries, dining rooms, warehouses, storage rooms, archives, workshops, characterized by the presence of non-medical personnel in a discontinuous manner and the absence of patients and visitors. Except for the consumption related to the production process, of sterilization centres, kitchens, and laundries must be considered and evaluated.

The need to ensure indoor air quality comfort and security standards [12] lead to the necessity of complex systems dedicated to Heating, Ventilation and Air conditioning purposes ($HVAC$) [13]. $HVAC$ (Heating, Ventilation, and Air Conditioning) systems are a critical component of any building, including hospitals. These systems are responsible for maintaining comfortable indoor temperatures, providing fresh air and ventilation, and controlling indoor air quality. However, the energy consumption associated with $HVAC$ systems can be substantial, particularly in large buildings like hospitals. Outdoor climate conditions, such as external air temperature and relative humidity, play a significant role in determining the energy consumption of $HVAC$ systems.

Indeed, the Ventilation systems require electrical energy, while the air handling process needs thermal energy in the form of hot and chilled water for the air temperature control and management, and steam for humidification purposes. Consequently, refrigeration systems (which are commonly composed of electricity-driven compression chillers) are responsible for a relevant increase in the electricity requirements of a healthcare facility, especially during the summer season [14]. Similarly, the thermal energy requirements result to be markedly driven by the hot water needed for indoor heating purposes. In general, the electrical energy demand for healthcare facilities is primarily driven by the operation of air conditioning systems and other specialized services such as surgical procedures and the use of electro-medical equipment (e.g. nuclear magnetic resonance, radiography). Additionally, the operation of refrigeration units often requires the presence of cooling towers on site to dissipate heat. This can present logistical challenges in terms of available space, leading to the adoption of space-saving forced air circulation evaporative cooling towers. However, this approach also increases electricity needs to power the cooling tower fans, leading to further increase the overall electricity requirements.

During hot summer months, the cooling load on the *HVAC* system will be higher, leading to increased energy consumption. Similarly, during cold winter months, the heating load on the *HVAC* system will be higher, again leading to increased energy consumption. Consequently, outdoor climate conditions can be often considered as one of the main building energy drivers. Effective control of the *HVAC* system is one way to minimize the impact of outdoor climate on energy consumption. For example, by implementing advanced control strategies, building operators can optimize *HVAC* operation to minimize energy consumption while still maintaining comfortable indoor temperatures and good indoor air quality. Additionally, by monitoring and controlling the ventilation rate, it is possible to reduce the amount of outdoor air that needs to be conditioned, which can lead to significant energy savings.

Hu et al. [15] report the electricity consumption of a large acute hospital in Taipei City (Taiwan), highlighting the major influence of the chillers on the overall building electricity consumption. Renedo et al. [16] analyse different possibilities for providing heating, air conditioning and hot tap water to a hospital centre, showing the high electricity consumption due to air conditioning purposes, especially during the summer. Congradac et al. [17] develop a mathematical tool for the exact calculation of room/building energy demands. Preliminarily, the paper analyses the energy balance of a hospital facility, highlighting that heating and cooling energy demand represents 48% of the overall energy consumption of the building.

A study reports a detailed correlation analysis carried out to find the main energy drivers of a healthcare facility in Italy which is the test case of this thesis work [18]. The study shows that outdoor climate conditions play a fundamental role in the electrical energy requirements behaviour. Moreover, a simplified multivariate linear regression model has been realized. These kinds of analyses are fundamental to obtain useful information on the building energy consumption, allow one to monitor the energy consumption to pursue a continuous improvement of the energy efficiency [19] and reduce environmental impact [20]. Moreover, energy efficiency can be pursued by the introduction of different layouts or technologies of the building plants and systems. A previous work by the author (reported in appendix A) develop a building energy model on the test case under analysis to carry out a techno-economic assessment of two different changes in the system layout. Moreover, some studies analyse several passive building retrofit measures which allow enhancing the buildings energy efficiency [21]. Montero et al. [22] analyse the behaviour of photovoltaic self-consumption system of a hospital in southern Europe, finding that an average of 25–30% of the annual electricity demand could be covered.

1.1.1 Energy Performance Indicators

Energy performance indicators are metrics used to measure and evaluate the energy efficiency of a system, building, or process. These indicators are typically numerical values or ratios that provide information on energy consumption, energy cost, or energy efficiency, and can be used to monitor and improve the energy performance of a facility or system. Comparing the actual energy performance indicators of the building under analysis with those of similar structures can be done to obtain valuable insights. The critical aspect of this operation lies in finding meaningful energy indicators and comparison benchmarks. It is essential to evaluate all elements that can influence the building energy consumption, in order to use appropriate adjustment factors capable of making the values of indicators comparable for structures in the same sector. Such elements for a hospital may include:

- *climatic characteristics* (climate zone, degree days, etc.);
- *building morphology* (volumes, envelope data, floors, etc.);
- *building usage* (beds, care intensity, etc.).

The following are some indicators commonly used in the healthcare sector:

- **Energy demand per bed** (kWh/bed): allows comparison between structures located in the same climate zone and with similar technology, provides the per capita energy consumption related to the number of hospitalizations. This index is heavily influenced by the climate and the internal organization of the hospital
- **Energy demand divided by useful floor area** (kWh/m^2) or **volume** (kWh/m^3): if used for buildings with similar characteristics, they allow comparison between dimensionally very different structures, may be inadequate if the usage of the environments being examined is not considered
- **Energy demand per Heating Degree Days** (kWh/HDD): The heating degree days (see section 3.2, Eq. 3.1) is a parameter useful to quantify the need for heating of a specific location. Consequently, an energy performance indicator based on this parameter can be useful to compare buildings placed in locations with marked climatic differences

The listed indicators can be easily combined to obtain reliable comparison parameters. As an example, a useful "composed" indicator can be obtained by dividing the energy demand by the product of the useful floor area and the heating degree days. This one allow to compare two buildings with different sizes and different climate conditions.

Several studies on the electrical and thermal energy demands have been conducted on thirteen Spanish Private Hospitals between 2008 and 2017, aiming to evaluate the main

hospital building energy drivers [23]. The study realizes a correlation analysis between the energy consumption of the buildings and several parameters related to the activities carried out within the analysed buildings. The results of the analysis show that both electrical and thermal energy demands are strongly influenced by the useful floor area of the facilities, while other parameters like staff in service and number of beds result have a weaker correlation with the building energy demand. A successive study [24] estimate the economic and environmental impacts of energy consumption derived from healthcare buildings and proposes several energy-saving options in the sector, finding that through an appropriate energy management of healthcare buildings it is possible to save up to 8.60 kWh/m^2 per year, for buildings of less than 5′000 m^2 (with no beds), while in healthcare buildings larger than 5′000 m^2 (with beds) it was possible to save up to 6.88 kWh/m^2 per year. The study analyse the relation of the hospitals energy consumptions in function of number of beds, staff in service, number of hospitalization, number of emergency, number of discharges and useful floor area, highlighting a marked positive linear correlation between the aforementioned parameters and the building energy demand. Moreover, while the thermal energy demand seems to be mildly affected by these parameters, their increase provoke a marked increase in electricity consumption. It is an important results since the electricity consumptions represents often the main part of the overall yearly hospitals energy consumption.

1.1.2 Regulations in healthcare field

The current Italian regulations regarding the physical-technical requirements for hospital building construction (thermo-hygrometric indoor air conditions, ventilation, and lighting) are primarily based on two reference texts (while not very recent, are still considered obligatory and holds legal validity): the *Circular of the Ministry of Public Works of 22/11/1974, no. 13011 "Physical-technical requirements for hospital building construction. Thermal, hygrometric, ventilation, and lighting properties"* and the *Decree of the President of the Republic 14/01/1997 "Minimum structural, technological, and organizational requirements for the exercise of health activities by public and private structures"*. **Circular 13011** prescribes the following points:

- Room temperature of 20 $°C \pm 2\ °C$ in all hospital departments, including services, guaranteed during the winter season through heating system.
- In the inpationts rooms, collective use rooms, and hallways during the winter season, a relative humidity value of 40% ± 5% must be guaranteed.
- The values of air exchange must be those summarized in Table 1.3 and guaranteed by a forced ventilation system in which the air is properly filtered and then introduced

into the environments at speeds not exceeding 0.15 m/s.

- In surgery rooms, delivery rooms, neonatal units, intensive care units, dialysis centers, sterile areas, and analysis laboratories, there must be an air-conditioning system that guarantees compliance with the aforementioned requirements both in summer and winter.

The **DPR 14/01/97** defines the minimum requirements that must be respected and applied by a healthcare facility, in terms of policy, objectives and activities, organizational structure, management of human and technological resources, quality assessment and improvement, and information system. The healthcare complex is divided into 13 functional environments. Table 1.4 reports the main thermo-hygrometric requirements.

There are also specific technical regulations that govern various sectors: the main ones of interest are *UNI 10339:1995* "*Aeraulic systems for well-being. Generalities, classification and requirements*" and *UNI 11425:2011* "*Ventilation and conditioning system with controlled contamination for the operating room*".

UNI 10339:1995 provides technical data on ventilation and devices to ensure good air quality. The minimum air exchanges and maximum allowable speeds required by **UNI**

Table 1.3: Prescribed air changes per hours (ACH) divided by area typology (Circular 13001)

Zone	**ACH (Vol/h)**
Inpatient rooms	2
Inpatient rooms for children	3
Diagnostic departments	6
Sanitaries	10
Special departments	12

Table 1.4: Main thermohygrometric requirements prescribed by DPR 14/01/97

Zone	T (min-max) ($°C$)	RH (min-max) (%)	ACH (Vol/h)
Surgery dep.	20-24	40%-60%	15
Obstetrics dep.	20-24	30%-60%	6
Intensive care unit	20-24	40%-60%	6
Medicines stockroom	20-26	45%-55%	2
Sterilization	20-27	40%-60%	15
Disinfection service	20-27	40%-60%	15
Morgue	18-18	55%-65%	15

10339:1995 are reported in Table 1.5. The flows reported in Table 1.5 must be corrected based on altitude, applying a specific corrective coefficient based on the height above sea level (Table 1.6). Moreover, there are also prescribed restrictions on the maximum speed at which air can be supplied into the indoor environments (see Table 1.7).

According to **UNI 11425:2011**, the requirements for the *VCCC* (Ventilation, Climate Control, Contaminant Control) system serving the operating room are defined and the environmental conditions for the rooms constituting the operating room are suggested. The use of recirculated air in individual operating rooms is allowed as long as this air is extracted and reintroduced into the same operating room without coming into contact with air flows from other rooms. In addition to Italian regulations, it is customary to also

Table 1.5: Minimum ventilation prescribed by UNI 10339:1995

Zone	**Outdoor air mass flow rate**
Inpatient rooms	11 L/s per person
Hospital wards	11 L/s per person
Steril rooms	11 L/s per person
Outpatient rooms	8.5 L/s per person
Sanitaries	Extraction of 8 Vol/h

Table 1.6: Corrective coefficient of minimum ventilation requirements based on the altitude

Altitude (m)	**Coefficient**
0	1
500	1.06
1000	1.12
1500	1.18
2000	1.25
2500	1.31
3000	1.38

Table 1.7: Maximum admitted air speed by UNI 10339:1995

Zone	**Heating (m/s)**	**Cooling (m/s)**
Inpatient rooms, Steril rooms	0.05-0.1	0.05-0.15
Obstetrics dep., Surgery rooms	0.05-0.1	0.05-0.15
Physical therapy dep.	0.1-0.2	0.15-0.25

refer to the most important foreign regulations, in particular the German **DIN 1946-4** "*Ventilation and air conditioning - Part 4: VAC systems in buildings and rooms used in the health care sector*" and the American **ASHRAE 170-2013** "*Ventilation of Health Care Facilities*". The pressure conditions prescribed by the latter for various rooms, in order to prevent possible contamination and bacterial flow leaks, are normally used in Italy as well.

1.2 Energy audit

In recent years, the topic of energy efficiency has gained increasing importance, especially considering the European goals for reducing energy consumption set for *2030* and *2050*. It is therefore important to identify the least efficient energy systems and intervene on them. The procedure for accurately analysing the system is called *energy audit*. The energy audit is often confused with the energy certification. Energy audit differs from energy certification because the latter provides a simple energy quality of the building/system in standard conditions, while the energy audit provides information regarding any energy inefficiencies of a system in its current operating conditions, identifying ad hoc solutions that are both energy and economically efficient.

The **UNI CEI EN ISO 50001:2018** (which substitutes the preceding **UNI CEI EN ISO 50001:2011**) standard entitled "*Energy management systems. Requirements and guidelines for use*" identifies energy analysis (energy audit) as an energy usage evaluation procedure, through which the interested party is able to obtain information regarding the areas of significant energy use and possible energy efficiency interventions. Until the definition of the **UNI CEI EN ISO 50001:2011**, the only document to refer to in defining the characteristics and procedure of energy audit was the **UNI CEI/TR3 11428:2011** entitled "*Energy management. Energy diagnoses. General requirements for energy diagnosis services*". This technical report provides suggestions regarding the data necessary for diagnosis, the operations to be performed, and the presentation of energy savings interventions. In *2012*, the **UNI CEI EN 16247-1** was released entitled "*Energy diagnoses - Part 1: General requirements*". The subsequent parts (16247-2,-3,-4,-5) are being drafted, including the one regarding buildings, which will complete the general framework. The **UNI CEI/TR 11428** proposes a systematic procedure aimed at:

- Providing adequate knowledge of the energy consumption profile of a building or group of buildings, an industrial activity and/or facility, or public or private services
- Identifying and quantifying energy savings opportunities from a cost-benefit perspective

- Reporting on results.

The energy audit establish standard procedure aiming to provide a description of the energy system, defining the possible efficiency improvement measures and quantifying the resulting savings. The energy audit must take into account the expectations, needs and limitations indicated by the client in order to jointly decide the scope, execution mode, objective and purpose of the diagnosis. The energy diagnosis must possess the following requirements:

1. **Completeness**: definition of the energy system comprehensive of the significant energetic aspects

2. **Reliability**: acquisition of real data in the necessary number and quality for the development of the energy inventory and survey of the structure under analysis. The responsible for the diagnosis must however verify that the energy consumption is consistent with the billing data for energy consumption or with what was detected by the measuring instrumentation

3. **Traceability**: identification and use of an energy inventory, documentation of the origin of the data and of the eventual mode of processing to support the results of the energy diagnosis, including the working hypotheses that may have been assumed

4. **Utility**: identification and evaluation under the cost/benefit profile of the interventions for improving energy efficiency. The interventions must be expressed through adequate documentation, differentiated based on the sector, the purposes, and the scope of application

5. **Verifiability**: identification of the elements that allow the client to verify the achievement of the efficiency improvements resulting from the application of the proposed interventions

1.2.1 Phases of the energy audit

The norm **UNI CEI/TR 11428** divides the energy audit process in several (standard) phases:

Preliminary communications

The energy audit responsible must agree with the client on:

- the objectives, needs, and expectations regarding the energy audit
- the purpose, limits, and degree of accuracy of the audit

- criteria for evaluating energy saving measures
- the commitment of time, resources, information, and organization required from the client
- strategic programs, management systems, or other situations that may be relevant to the energy audit and its conclusions
- opinions and constraints regarding potential energy saving measures

The person in charge of the audit must also request the client to:

- nominate staff members to communicate with
- inform the involved staff and other parties about the proposed audit and any requirements that may result from it
- ensure cooperation between the parties involved in the audit

Preliminary meeting

The purpose of the preliminary meeting is to provide all stakeholders with information regarding the objectives, scope, boundaries, and degree of detail of the energy diagnosis and to agree on all operational methods for carrying out the energy diagnosis. During the preliminary meeting, the energy audit manager must agree with the organization:

- the methods for accessing the structure
- hygiene and safety regulations in the work places affected by the diagnosis
- the resources and data, including confidential and sensitive information, that must be provided

The energy audit manager must describe the procedures and methods for planning and executing the energy diagnosis, and must also specify and justify any need for additional measurement tools. It is also necessary to obtain from the client the following data and information:

- detailed information on the characteristics of the objects of the diagnosis, including any factors that may influence consumption
- information about past events that may have influenced energy consumption in the period covered by the collected data
- information regarding unusual operating conditions, maintenance work, or other activities concurrent with the energy diagnosis that may influence its outcome

Data Collection

The energy audit manager must collect information from the client in order to obtain:

- historical data (bills, adjustment factors, operation and maintenance, etc.)
- building and systems documents (planimatries, as-builts, plants and systems data etc.)

Operative activities

It is then necessary to inspect the energy system in order to confirm the information obtained through the documentation analysis and obtain further information about the system functioning, as:

- evaluate the significant energy aspects
- identify the operational modes, user behaviours, and their influence on energy consumption and energy efficiency
- list the areas and processes that require further quantitative data to support subsequent analysis
- generate recommendations for reducing energy consumption

On the one hand, the energy audit manager must esure to have sufficient information to conduct the analysis. In particular, the energy audit manager must ensure that:

- measurements and inspections are carried out reliably and under conditions that represent normal operating conditions and in correct environmental conditions
- the previous prescription does not exclude the carrying out of inspections outside working hours, during periods of inactivity, or when no climatic load is expected
- promptly inform the client of any difficulties encountered during the field activity

On the other hand, the client must ensure to provide the required information. In particular, the client must be asked to:

- provide unrestricted and secure access to all the systems that need to be inspected
- nominate one or more people to act as guides and assistance during the in-situ activity. These people must have the necessary skills, abilities, and authority to carry out any interventions on processes and equipment, if required
- allow access to drawings, manuals, and other significant technical documentation of the energy audit systems, along with the results of any tests and measurements carried out at commissioning

Analyses

A thorough examination of the collected data and information is necessary in order to identify opportunities for improving energy efficiency. Through this analysis, the energy audit manager must derive further data and information that must consist of at least:

- a balance and flow diagram of energy divided by use and supply mode
- relationship between consumption and factors that affect variations
- actual and reference energy performance indicators
- evidence of any change in performance over time
- time-based energy demand diagram
- determination and quantification of potential energy savings and cost/benefit analysis for such interventions

1.2.2 Energy Manager

The energy audit is closely linked to a specific professional figure that is progressively becoming increasingly important and present in the environment of commercial and industrial activities: the **Energy Manager**.

The energy manager is a professional figure (which in Italy is linked to a legal obligation for energy-intensive activities) whose main tasks include the analysis, monitoring, and optimization of energy usage in companies and public or private entities, thus enabling cost, energy, environmental, and production-related savings. In recent years, the role of the energy manager has grown, integrating with the purchase of fuels and energy carriers (electricity, natural gas, etc.), the adoption of the *LCCA* (life cycle cost analysis) methodology for product and service procurement, the management of guaranteed energy performance contracts (*EPC*), the integration of resource management with the company's core business and the definition and management of informational and awareness campaigns aimed at personnel. The energy manager also plays a key role in the energy management system for organizations certified under **UNI CEI EN ISO 50001:2018**.

In other words, the Energy Manager, as suggested by the term, is the person responsible for managing energy-related matters within a company, public entity, or organization as a whole, monitoring consumption, optimizing it and promoting targeted actions towards energy efficiency and the use of renewable sources. This results in a different role depending on the dimensional characteristics of the structure in question:

- In the case of a complex organization, the energy manager will preferably be a high-level manager leading a group of predominantly technical people.

- In the case of small businesses and public entities, it will generally be an external consultant with technical competencies in energy management.
- In the case of a residence, a dedicated consultant is not feasible due to costs, and the function can be performed by support networks (e.g. consumer associations, energy points and agencies, etc.) or by those who issue the energy certification.

In the latter case, it is more properly an *energy audit manager* rather than an energy manager, lacking the management aspect. Even in the second case, the management aspect will be limited, although present (e.g. energy accounting and feasibility studies are usually the tasks of the professional, but it is unlikely that a consultant can interfere with internal company procedures), but it can make sense to define it as an energy manager. The first case is the only one in which it is not theoretically necessary for the energy manager to be a technical person, as they must essentially be a high-level manager supported by good specialists, capable of effectively influencing company energy policies. Given the nature of the energy aspects as insiders, however, a technical profile may be preferable even in this situation.

The energy manager, therefore, assesses consumption through ad hoc audits or, if available, through reports produced by telemetry, tele-control, and automation systems. He then optimize consumption through the proper regulation of facilities and their appropriate use from an energy standpoint, promote energy-conscious behaviour among employees and/or building occupants, and propose improvement investments that can potentially improve production processes or the performance of related services. Another function that often involves the energy manager is the procurement of electricity and other energy carriers. In this case, the aim is to reduce purchasing costs, possibly by promoting proper electricity load management to avoid peak power that results in higher costs. Among less common but useful options is the possibility of collaborating with the purchasing office to promote procedures that promote so-called green procurement and the purchase of machinery with low energy consumption and low operating costs (life cycle cost analysis – LCCA).

Among the major difficulties, the following can be cited:

- The need to often deal with non-technical people, both on the side of corporate decision-makers and on the side of colleagues in charge of other tasks, which requires the ability to explain technical concepts in simple terms
- The need to communicate with other corporate functions, as energy is a horizontal issue by itself, involving those who buy electricity and other fuels, but also machinery and devices, those who are responsible for the maintenance of systems, those who

design buildings and plants, those who write specifications for tenders, legal offices, and so on

To be addressed effectively, the first point requires mostly personal qualities while the second requires the involvement and commitment of corporate leaders, who must create the conditions for the energy manager to perform at their best and find the necessary collaboration. One way to respond to this is to have an energy management system (*BEMS*) certified under **UNI ISO 50001**. To detail the role of the energy manager with reference to typical actions, the following list can be considered:

- Making contact with the organization and identifying reference figures for the performance of their activities (decision-makers, procurement office, energy management experts, maintenance personnel, process line managers, administrative and accounting functions, financial functions, etc.)
- Collecting energy bills, evaluating monthly and annual consumption, verifying; identifying daily electrical and thermal load curves
- Checking existing contracts related to energy services (to ensure the services provided and to plan investment proposals properly)
- Creating a database of consumption areas, with more details for the most significant ones (features, committed and rated powers, energy/hours of operation, installation date, etc.)
- Identifying a set of *Energy Performance Indicators* (see section 1.1.1) to compare consumption between different sites and with literature
- Conducting energy diagnoses and feasibility studies (personally or with the assistance of third parties)
- Proposing interventions and feasibility studies (monitoring, waste reduction, employee awareness programs, investments in efficiency and renewables)
- Monitoring regulations and accessing incentives
- Verifying results achieved and communication programs

When analyzing only the operational aspect of an energy manager or any figure (internal or external to a company) that wants to undertake an energy analysis and auditing process, two main activities that must necessarily be carried out can be classified as follows:

- *Techno-economic assessment* of possible changes in the system layout and management strategies of the building, which may involve both changes in the management strategy of existing plants, and the evaluation of the costs incurred from the installation of new plants (including generation plants such as co-generators and photovoltaic panels)

- *Monitoring and Targeting* of the operational conditions of the facility's systems, verifying their proper functioning and identifying anomalies in energy consumption, from which possible corrective actions will result

1.3 New system management strategies and layouts assessment

The reduction of energy consumption in buildings through the adoption of energy efficiency policies is a key pillar of the European Union (EU) climate and energy strategy [25]. The continuous improvement in energy efficiency must pass through the evaluation of different solutions, which imply the modification of the building energy management strategies, or changes in the energy systems and their layout. The continuous improvement in energy efficiency is a critical aspect of ensuring sustainable and cost-effective energy consumption in buildings. This requires the evaluation and implementation of various solutions that can significantly impact energy consumption. These solutions can include modifications to the energy management strategies of a building, as well as changes to the layout and configurations of energy systems.

The implementation of energy efficiency measures in industrial and commercial buildings can be a challenging task, especially when considering the need to maintain continuity of service. These measures can involve significant changes to the building's energy systems and layouts, which can be costly and time-consuming. Additionally, it can be difficult to predict the long-term effects of these changes on energy consumption, making it challenging to justify the investment. To overcome these challenges, numerical methods can be utilized to assess the effectiveness of various operations aimed at optimizing the building's energy demand. These methods can involve the use of computer simulations, mathematical models, and optimization techniques to predict the energy consumption of the building under different scenarios. By using these techniques, it is possible to evaluate the potential energy savings, the costs of implementation, and the overall impact on the environment. This can help building owners and managers to make informed decisions about which energy efficiency measures to implement, and how to implement them in an efficient and cost-effective manner. By using numerical methods, it is possible to identify the most suitable solutions that balance the need for cost-effectiveness and continuity of service with the potential for energy savings and environmental benefits.

In hospitals and healthcare facilities, it is often impossible to carry out experiments on different solutions in terms of systems energy management strategies, due to the specific constraints in terms of service continuity and to ensure security standards of indoor air quality. Consequently, the evaluation of several operations aimed to increase the building

energy efficiency must be carried out exploiting numerical simulations able to reproduce the energy demand of the structure and its plants and systems. To this end, the dynamic models represent a powerful tool to evaluate the building energy performance or the effect of modifications on plants and systems configuration or their management strategies [26].

Process simulation is the representation of industrial processes by means of the application of mathematical and physical principles. Steady-state simulation models are widely used in industry and are fundamental to the conceptualisation, design and evaluation of processes. However, steady state is an idealistic definition used as a representation of 'design' conditions that cannot always be achieved. The most significant difference between steady-state and dynamic simulation is that the steady state assumes that the variables are constant with respect to time. This means that in the steady state there is no accumulation in the system, so the total input mass and energy correspond to the output.

In contrast, dynamic models take into account the rate of accumulation of mass and energy within the system, which allows one to determine how long it would take to reach a steady state from a specific initial state. Regarding the building energy modelling, the use of dynamic simulation is essential in order to take into account the many variables involved. For example, thermal inertia phenomena in the building can have a significant influence on the energy consumption of air conditioning, partially decoupling the indoor climatic conditions from the outdoor ones. Van der Veken et al. [27] compare the results of steady-state and dynamic building energy simulations, finding that the steady-state simulation predict an overall energy consumption about 5% higher of the dynamic counterpart, being unable to model the transient processes that take place inside the building. Horton and Braun [28] carry out a similar comparison, focussing on the load-based performance evaluation for residential air handling systems. Even in this second study, the steady-state simulation appear to overestimate the energy consumption.

1.3.1 Building Energy Modelling

A common method to perform this kind of analysis is called **Building Energy Modelling** (*BEM*). *BEM* is the practice of using computer-based simulation software to perform a detailed analysis of a building energy use and energy-using systems. The simulation software works by enacting a mathematical model that provides an approximate representation of the building. BEM includes whole-building simulation as well as detailed component analysis utilizing specialized software tools that address specific concerns, such as moisture transfer through building materials, daylighting, indoor air quality, natural ventilation, and occupant comfort.

Several software and tools have been developed to perform dynamic numerical energy simulations on buildings (e.g., *EnergyPlus* [29], *TRNSYS*, *OpenStudio*), and many studies

have been carried out over past years to perform building energy modelling and energy performance prediction. Magni et al. [30] report on the calculations carried out by these tools, analysing four of the most popular ones. They also compare the performance of the software under analysis by simulating an office building. Foucquier et al. [31] and Harish and Kumar [32], provide state-of-the-art reviews in building energy modelling and simulation of building energy systems.

As already said, despite the exploited software, the building energy modelling can be divided in two main categories: **Steady state** and **Dynamic** simulations. **Steady state simulations** base their calculations on a more representative period of the year (usually in terms of temperature). Generally, the heating or cooling thermal plant is dimensioned respectively on the basis of the hottest and coldest day, that is, the extreme weather conditions to which the building under evaluation is subject; typically, these tools use simple and stationary calculation models (in this treatment, the term "stationary" refers to a state that is not subject to change or variations towards other states; generally, it is said that in the physical model it is independent of time). Methods for designing heating or cooling systems are therefore based on sizing peak power, using simple multiplication formulas essentially based on volume and safety coefficients.

Beyond the regulatory requirements, it would be advisable to limit the exclusive use of these models to the evaluation of the energy characteristics of technologically simple buildings. Steady state models are used especially in all those tools where it is necessary to give an energy evaluation of a standard type [33]. generally, these software are used primarily to attribute an energy performance for the containment of energy consumption [34]. These methods are approximate and often cause over-sized plants that inevitably lead to increased costs and consumption. Over-dimensioning that does not only affect the upstream equipment of the plants, such as heat generators or refrigeration units, but also pipes, distribution organs, circulation organs and terminal organs. Therefore, additional costs not only related to the purchase of the plant but also to its management, which impact throughout the entire life of the building. In the following, a concise summary of the analyses for which the utilization of stationary simulation can be employed is provided:

- Estimation of peak thermal loads
- Preliminary sizing of plant and/or building
- Seasonal Performance Index
- Energy certification of buildings

A complete and versatile method for proper sizing must necessarily exploit the potential of energy modelling and in particular **dynamic simulation** (simplified scheme in figure 1.5). For technologically complex buildings both in terms of construction and plant

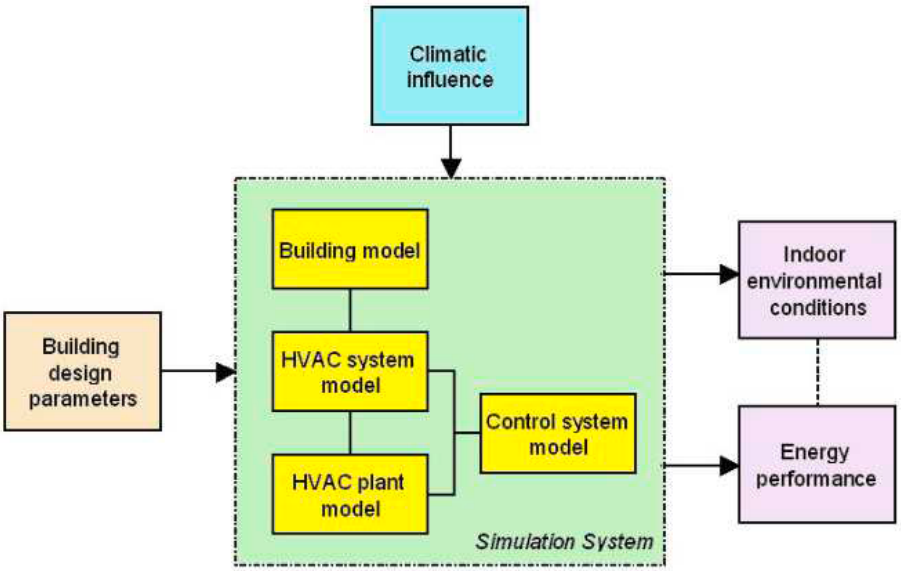

Figure 1.5: Simplified scheme for modelling of building-plant system

engineering, such as a healthcare facility, it is advisable to use a dynamic simulation model, and in any case whenever very precise and detailed information is needed about the thermodynamic performance of the system considered.

The design of highly energy-efficient buildings requires advanced tools that allow for the most correct possible evaluation of all thermodynamic phenomena. Only dynamic simulation analyses the energy performance of a building with precision and reliability, allowing for the control of the many variables present in a building-plant system. Dynamic simulation tools, unlike stationary ones, can predict the energy consumption of the plants connected to buildings over a longer period (typically annual periods). They generally consist of analysis tools for the prevention of energy consumption during design and verification, through the use of numerous physical models for the energy representation of the building, so they can be used in all stages of the structure's life and at any time considered.

A general understanding of the methods and their limitations used in simulation programs is essential if these tools are to be applied appropriately. Dynamic models predict the resolution of thermodynamic equations considering all physical phenomena as a function of time that cannot be resolved in stationary models, such as:

- Influence on loads from external shading
- Air flows from the outside environment to the inside or vice versa between indoor environments
- Estimation of seasonal thermal energy consumption
- Calculation of building CO_2 emissions

- Estimation of building management costs
- Verification of the thermal/inertial behaviour of the building envelope
- Optimal sizing of thermal systems and related subsystems
- Evaluation of energy production from renewable sources
- Evaluation and assessing of internal environment thermal comfort conditions
- Verification of natural and artificial lighting of the various rooms
- Verification of the effectiveness of natural ventilation strategies
- Analysis of pollutant generation and removal

Using a dynamic simulation program correctly, the above phenomena can be partially or completely considered and the mathematical model is constructed to represent all possible energy flows and their interactions, approaching the primary target of reality emulation. Some examples of analyses for which the use of dynamic simulation is necessary are the estimate of thermal loads based on the actual conditions of use of the building-plant system (environmental, occupational, management, etc.), the sizing of all systems and subsystems (plant, geometric, construction) with estimation of performance, seasonal thermal energy consumption, CO_2 emissions, building management costs, and optimal sizing of thermal plants and subsystems

In 2000, Kalogiorou [35] exploit *TRNSYS* software to model the building energy behaviour with a hybrid photovoltaic-thermal solar system. Several office buildings in Brazil have been simulated using *EnergyPlus* [36], providing a methodology to perform the model calibration based on the available building information. More generally, Coakley et. al. [37] analyse several numerical methods to match building energy simulation models to available measured data. Zhai et al. [38] perform building energy modelling simulation using *EnergyPlus* on three office buildings, aiming to assess the effects of ventilation on building indoor air temperature during the summer. Booten et al. [39] exploit *EnergyPlus* software to simulate the dynamic response of a residential building to different cooling strategies, aiming to shift the energy demand off-peak. They collect experimental data about the real building in which the cooling is provided by a dedicated heat pump, using it to validate the model and to assess the software capabilities in cooling energy demand modelling. Kamal et al. [40] use *OpenStudio* software to model a large office building for three different thermal energy storage systems to quantify their energy-saving capabilities. A concept of an integrated energy system for residential applications has been presented by Kotowics et al. by carrying out simulations using *EnergyPlus* [41].

EnergyPlus has been exploited by Rubeis et al. [42] to evaluate the sensitivity of the heating performance of a real residential building in Italy to different climate zones,

different climate conditions and different *HVAC* system solutions. The same building was the object of energy demand analysis, exploit the building energy model to identify energy-saving strategies [43]. *EnergyPlus* has been used to design and optimize a novel cogeneration system for an office building, allowing to perform technical and economic considerations [44]. Despite the usefulness of the *BEM* technique, simulating complex buildings can lead to a computationally expansive model which can lead to problems in the model exploitation. To partially overcome this problem, Edwards et. al. developed an *EnergyPlus* model surrogate through machine learning techniques [45]. Other fields of application of *BEM* software are represented by the urban building energy modelling. In this field, the *BEM* software is exploited to carry out city-scale energy simulations [46]. Beccali et al. [47] investigate the technical and economic feasibility of *CHP*-retrofit of existing power plants and the possible utilization of the recovered heat to supply, via a district heating and/or cooling network, the energy requests of civil users in the tertiary and residential sectors. The residential energy demand has been estimated using the *EnergyPlus* software by performing simulations for a set of reference buildings.

Focusing on the healthcare sector, accurately reproducing the energy performance of hospital buildings can be an extremely difficult task. On one hand, the complexity of the heating, ventilation, and air conditioning (*HVAC*) systems can lead to numerous problems in accurately modelling the energy consumption of these systems. The *HVAC* systems in hospitals are often more complex than those in other types of buildings due to the need for specialized equipment and the need to maintain a sterile environment.

On the other hand, the diverse needs and activities carried out within the building can also present obstacles to accurately reproducing the energy performance. Hospitals are multi-purpose facilities, with different areas dedicated to different functions such as patient care, research, and administration. These different functions and activities can make it difficult to accurately model the overall energy performance of the building. Furthermore, the presence of patients with different needs and the presence of medical equipment can lead to a high variability in the energy consumption of the building. This can make it difficult to develop a representative model of the energy consumption of the building that can be used for analysis and energy efficiency improvement. Buonomano et al. [48] present a methodology to reproduce a big hospital building in Italy. They analyse several aspects of the *BEM* realization and then exploit it to perform dynamic energy performance analysis.

A preliminar building energy model of the test case analysed in the present work has been realized. The obtained model has been analysed and exploited to evaluate changes in the building *HVAC* systems in order to reduce the building energy demand. Moreover, the optimal size of a photovoltaic generation system has been also evaluated through the

obtained model. Further details can be found in Appendix A.

1.4 Monitoring and Targeting

The technique of energy Monitoring and Targeting (M&T) is based on the management principle that effective management requires accurate measurement. Through the use of M&T techniques, energy managers are equipped with the information needed to assess the impact of their energy management practices, evaluate the results of their energy efficiency projects, and determine expected energy consumption levels over a specified time period. Furthermore, M&T techniques serve as an early warning system, alerting energy managers to instances of excessive energy consumption that may result from equipment malfunctions, operator error, unfavourable user behaviour, inadequate maintenance, and other similar factors.

The core principle of M&T involves analysing the typical correlations between energy consumption and its determinant factors such as *HVAC* equipment, production outputs, weather conditions, occupancy, and available daylight. The objective of this approach is to support business managers in achieving the following objectives:

- Identification and justification of excessive energy usage
- Detection of deviations in energy consumption patterns
- Representation of energy consumption trends over various time scales, including daily, weekly, seasonal, and operational
- Forecasting of energy usage and costs when contemplating business changes
- Identification of areas of energy wastage
- Understanding of the impact of determinant factors on energy efficiency
- Establishment of performance targets for energy management programs
- Effective management of energy consumption as a controllable expense rather than a fixed cost

The ultimate objective of this approach is to minimize energy costs through enhanced energy efficiency and effective energy management. This leads to additional benefits, including increased resource efficiency, improved production budgeting, and reduction of greenhouse gas emissions. M&T techniques are based on some fundamental steps that create a continuous feedback loop, thereby enhancing the control of energy utilization.

- **Measures**: The initial step involves the compilation of data from various meters. The availability of low-cost energy feedback davices has facilitated this process. The frequency of data compilation depends on the desired reporting interval, ranging from a few seconds to every 15 minutes or a hour. Some data can be directly obtained

from the meters, while others require calculation. These various measurements are commonly referred to as streams or channels. Relevant driving factors, such as production or degree days, also constitute streams and must be collected at intervals that align with the data from the meters.

- **Baseline definition**: The collected data must be used to define the baseline. This is essentially a mathematical model that represents the energy demand of the building based on the identified energy drivers, that is, those parameters that influence its performance. The methods used can vary, as the already mentioned Building Energy Modelling or machine learning-based ones. In the following chapters the most commonly used machine learning-based methods will be explored. The resulting model will then be an indicator (or predictor) of the building energy demand in "standard" conditions

- **Monitor variations**: The subsequent step involves monitoring the discrepancy between the predicted energy consumption and the actual consumption as recorded by meters. One commonly utilized tool for this purpose is the *CUSUM*, which stands for cumulative sum of differences. This involves calculating the difference between the predicted consumption (based on the best fit line established earlier) and the actual consumption. The *CUSUM* can then be charted over time to provide additional insights for the energy efficiency specialist. A *CUSUM* fluctuating around zero generally indicates normal process operation. However, a marked variation, either steadily increasing or decreasing, often indicates a change in the process conditions

- **Causes identification**: The energy manager, working in conjunction with building managers, are tasked with analysing the *CUSUM* graph to determine the root causes of deviations in energy consumption. Such deviations may be due to changes in occupant behaviour, process modifications, or alterations in external conditions, among others. The monitoring of these changes and the identification of the underlying causes are crucial to encourage favourable behaviours and discourage unfavourable ones, ultimately enhancing overall energy efficiency

- **Target definition**: Upon establishment of the baseline and identification of factors affecting energy consumption, it becomes possible to set meaningful targets for future consumption. Through the information obtained with the preceding steps, targets can be more accurately defined based on the actual building energy demand behaviour. The process of setting targets consists of two main components: the quantifiable reduction in energy consumption and the time needed to obtain the aforementioned reduction.

- **Monitor results**: The process of M&T then returns to the initial step of consumption measurement. The iterative nature of M&T necessitates continual feedback in order to continuously optimize performance. Upon implementation of targeted measures and setting of consumption goals, repeating the measurement process provides insight into the efficacy of the measures taken, enabling informed decisions on any necessary subsequent actions

As observed, the M&T process is primarily based on the baseline definition method, which effectively concerns the establishment of a predictive model aimed at representing the structure in its nominal conditions. This type of activity is referred to as **Building Energy Prediction**. The higher the performance of the prediction model, the more effective the building energy monitoring will be.

1.4.1 Machine Learning-based energy demand prediction

Although the *BEM* approach is a fundamental technique in the field of building energy performance evaluation, it presents a non-negligible drawback. Indeed, the data required to create a building energy digital twin are extremely detailed, starting from the characteristics of the building envelope, through the technical data of the installed systems, as well as data related to the activities performed inside the building. Such data are rarely available in adequate detail in real industrial or commercial contexts.

Depending on the specific buildings, data-driven techniques like **Machine Learning** can be valid alternatives to perform building energy forecasting and monitoring, aiming at exploiting the automatic learning process of these algorithms to find complex relationships between the building energy drivers and the building energy consumption [49, 50]. As an example, Kapp et al. [51] explore several statistical techniques and machine learning models informed by physical system parameters to predict the industrial building energy consumptions. Moreover, several studies were carried out to develop machine learning algorithms able to predict the energy demand of residential, industrial and commercial buildings. Mosavi et al. [52] classify several machine learning techniques used in the energy systems context, giving some guidelines on which method to use based on the type of application.

The choice of the machine learning model typology depends on the specific context and requirements in terms of accuracy. Moreover, the model interpretability can become a key factor in case of lack of user-experience in the specific field of machine learning. The term *model interpretability* refers to the degree to which the model itself and its predictions can be understood by a human being. As an example, the Multiple Linear Regression model (*MLR*) is basically a function in which each input features is multiplied by a coefficient.

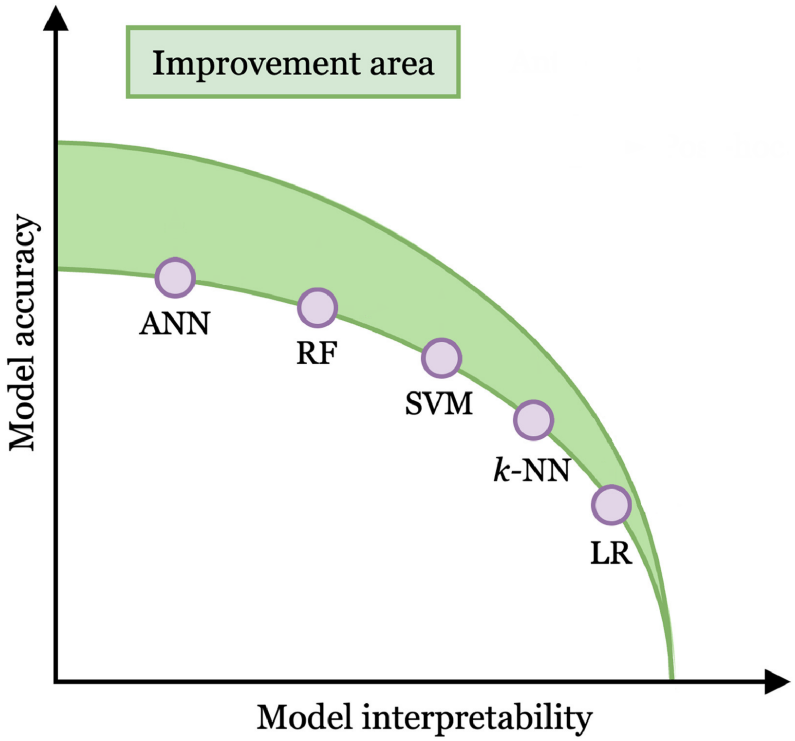

Figure 1.6: Trade-off between model accuracy and model interpretability [53]

In a certain sense, the coefficients of the individual features can be interpreted as a factor quantifying the influence of the individual feature on the target variable. In contrast, more complex models, such as Artificial Neural Networks (ANN), are mathematical models composed of a multitude of operations between large matrices, which combine the various input features in a non-linear manner. It is therefore difficult to extract information regarding the influence of individual input features on the target simply by observing the internal parameters of similar models. Chen et al. [53] review the most used methods in the building energy field, evaluating them in terms of accuracy and interpretability. The paper represent the trade-off between model accuracy (the prediction performance) and model interpretability, as shown in Figure 1.6. The figure compare several machine learning techniques, trying to graphically represent them interpretability and accuracy. As already said, the Linear Regression (LR) and the Artificial Neural Network are polar opposites in terms of the combination of accuracy and interpretability.

Artificial Neural Networks can be used to develop a system able to predict and optimize the energy consumption of a building [54–56]. Zhao et al. [57] analyse in deep the available prediction method for the building energy consumption forecasting, reviewing recently developed models, which include elaborate and simplified engineering methods, statistical methods and artificial intelligence methods. Luo et al. use a genetic algorithm to set a deep feed-forward neural network model for predicting the electricity consumption of real buildings [58]. An Artificial Neural Network model has been developed to predict the energy consumption due to air conditioning in residential buildings [59]. Hwang et al. [60]

use a deep learning approach to predict the energy performance of heating and cooling systems. The models were developed based on measured data collected in an educational building and were classified into different prediction time groups at 3-min, 15-min, 30-min, and 1-h time intervals. Wong et al. [61] have developed an artificial neural network (ANN) model to predict the energy demand of office buildings in subtropical climates. A total of nine variables were used as the input parameters, including weather, building envelope and time information.

Seyedzadeh et al. [62–64] study in deep the possibilities that machine learning models have to be applied for predicting and optimizing non-domestic building energy consumption and performance. Ngarambe et al. [65] present a review of current AI-based methodologies being used to enhance thermal comfort in indoor spaces, focussing on thermal comfort predictive models using diverse machine learning algorithms and their use in building control systems to pursue energy-saving and maintain the thermal comfort of the building. The study presents also a review of the main features used to perform predictive internal climate control. Palladino et al. [66] discuss the possibility to realize a neural network model to predict the main thermal comfort indexes defined by the international standard UNI EN ISO 7730 [67].

Yuce et al. [68] present an ANN approach to predict the energy consumption and thermal comfort level of an indoor swimming pool, evaluating the environmental and control variables that affect energy consumption and thermal comfort. The ANN is exploited to develop an optimization-based control system, which has been implemented for a specific *HVAC* system. Kalogirou et al. [69] investigate the potential of using an ANN with singular value decomposition method (*SVD*) to predict the indoor temperature to shut down the heating system controller early, achieving savings in the energy consumption. Similar methods can be applied to monitor the operation of single components of a system. Garcés et al. [70] use several Machine Learning Techniques for the estimation of energy efficiency in industrial coal boilers, while Mohanraj et al. [71] explore the possible applications of artificial neural networks (*ANN*) for energy and exergy analysis of refrigeration, air conditioning and heat pump (*RACHP*) systems. The following chapter (Ch. 2) will be devoted to analyse the machine learning methods at the basis of the building energy prediction and to carry out an in depth analysis of the methods used to develop the presented building energy monitoring method.

1.5 Motivation

This work was developed in collaboration with *Casa di Cura Villa Donatello S.p.A.* (located in Sesto Fiorentino, Florence, Italy), which made available its healthcare facility data in order to develop methods to pursue the goal of increasing the energy efficiency of the building. The objective of this thesis work is to develop methodologies for continuous monitoring of building energy demand, capable of analysing energy demand data, identifying anomalies, and generating alarms to help users take prompt corrective actions.

Starting from the physical modelling, a *preliminar* energy model of the building was created using the technique called *Building Energy Modelling* (more information about the development of the building energy model are provided in Appendix A). The obtained model was calibrated through real data of the building and its energy requirements, and then used for the techno-economic assessment of different system and energy self-production solutions. However, energy monitoring of a system can be achieved provided that a model that can correctly and precisely represent the energy performance of the building in its current state (or standard) is available. The *Building Energy Modelling* technique requires a huge amount of information related to the building envelope, the systems installed within it, the behaviour of the people inside, and so on. This aspect makes the *BEM* technique not suitable for energy monitoring purposes.

For this purpose, it is necessary to reverse the traditional point view of modelling, moving beyond the realization of a building model starting from information on its configuration, and instead starting from the final effect: the energy requirement. In this context, machine learning-based techniques offer an excellent opportunity to create a mathematical model that can predict the energy demand of a building based on various influential parameters: the so-called *Energy Drivers*. However, the biggest drawback of these methodologies lies in the need to have users with adequate experience and knowledge both in the energy field and in the field of Machine Learning, that is not always available in an industrial or commercial reality (especially for those activities whose core business is completely different).

The aim of this work is to create a systematic methodology to proceed from the raw data to a ready-to-implement monitoring method, minimizing the knowledge required to the user and, contemporary, ensure reliable performance of the monitoring method itself. Different approaches will be applied and tested, starting from methods with low computational effort that require high user experience, to method that requires significantly higher computational resources, but aiming to minimize the required background knowledge to be applied. The obtained results will be analysed and compared in terms of prediction performance and computational cost. This work is therefore also intended as a "map" that allows the end user to orient themselves on the methodology to follow based on their

specific needs.

The second part of this work regards the application of the developed method to more detailed data about the electricity consumptions of Heating, Ventilation and Air Conditioning system installed in the facility. As a result, the obtained models will be applied to perform the monitoring of the whole building and detailed systems energy consumptions, obtaining insight of the effective usefulness of the implementation of this methods in a real activity context.

Chapter 2

Background and Methods

2.1 Background

Machine Learning (*ML*) is a field of research dedicated to understanding and building methods that allow for the automated learning of a computer, i.e. that exploit data to improve performance on a certain set of tasks [72]. Machine learning algorithms build a model based on sample data, known as *training data*, in order to make predictions or decisions without being explicitly programmed to do so. Machine learning algorithms are used in a wide variety of applications, such as medicine, e-mail filtering, speech recognition, agriculture and computer vision, where it is difficult or impractical to develop conventional algorithms to perform the necessary tasks [73, 74].

There are different *ML* approaches, which are typically divided into four main categories corresponding to the different learning paradigms:

- **Supervised learning** algorithms aim to build a mathematical model through a set of data that contains both the inputs and the desired outputs. During the training process, supervised learning algorithms learn a function that can be used to predict the output associated with new inputs. Supervised learning has many applications across a wide range of fields, including natural language processing, computer vision, and financial analysis. In natural language processing, for example, supervised learning can be used to train models to classify text based on its content. In computer vision, supervised learning can be used to classify images based on their content or to identify objects within an image. In finance, supervised learning can be used to predict stock prices or to identify potential fraud cases. One of the key advantages of supervised learning is that it allows for the creation of highly accurate models that can be used to make predictions on new data. This is because the model is trained on a dataset that already contains the correct output values, allowing it to learn from this information and make accurate predictions on new data. To

implement supervised learning, a variety of algorithms can be used, including linear regression, logistic regression, decision trees, and neural networks. These algorithms differ in their complexity and performance, and the choice of algorithm often depends on the specific application and the size and complexity of the dataset.

- **Unsupervised learning**: Unsupervised learning algorithms take a data set containing only input and find a structure in the data, such as grouping or clustering data points. The algorithms, therefore, learn from test data that has not been labelled, classified or categorised. Instead of responding to feedback, unsupervised learning algorithms identify commonalities in the data and react based on the presence or absence of such points in each new datasets. One of the primary goals of unsupervised learning is to discover latent variables that can help explain the underlying structure of a dataset. This can be accomplished through techniques such as clustering, dimensionality reduction, and density estimation. Clustering is a process of grouping similar data points together based on some similarity measure, while dimensionality reduction aims to reduce the number of features in a dataset while preserving its essential structure. Density estimation, on the other hand, attempts to model the underlying probability distribution of the data. Unsupervised learning has a wide range of applications in various fields, such as natural language processing, computer vision, and bio-informatics. In natural language processing, unsupervised learning techniques such as topic modeling can be used to discover the underlying themes in large collections of text documents. In computer vision, unsupervised learning algorithms can be used to learn representations of images that can be used for tasks such as object recognition and image retrieval. In bioinformatics, unsupervised learning can be used to identify clusters of genes that are co-expressed and may be involved in common biological processes. Despite its many applications, unsupervised learning remains a challenging problem due to the absence of explicit training data. It is often difficult to evaluate the effectiveness of unsupervised learning algorithms, as there may not be a clear metric for measuring their performance. Nevertheless, recent advances in deep learning have led to significant progress in unsupervised learning, particularly in the area of deep generative models. These models can learn to generate synthetic data that closely resembles the original dataset, which has important applications in fields such as image and speech synthesis.

- **Semi-supervised learning**: Semisupervised learning is a type of machine learning that combines both labeled and unlabeled data to improve the accuracy of predictive models. While supervised learning requires a large amount of labeled data to train models, semisupervised learning can leverage a smaller amount of labeled data along with a larger amount of unlabeled data to achieve similar or even better performance.

In semisupervised learning, the model is trained on a small set of labeled data, and then uses the remaining unlabeled data to learn additional patterns and structure. This approach is especially useful in situations where obtaining labeled data is time-consuming, expensive, or otherwise difficult. By incorporating additional information from the unlabeled data, the model can generalize better and make more accurate predictions on new, unseen data. One of the key challenges in semisupervised learning is how to effectively incorporate the unlabeled data into the learning process. Many different approaches have been proposed, including self-training, co-training, and graph-based methods. Self-training involves using the labeled data to train an initial model, and then using that model to make predictions on the unlabeled data. The high-confidence predictions are then used as additional labeled data to train a new model. Co-training involves training two separate models on different subsets of the data, and then exchanging the most confident predictions between the two models to improve their performance. Graph-based methods involve constructing a graph of the data points, and using the graph structure to propagate labels from the labeled data to the unlabeled data. Semisupervised learning has many practical applications, especially in domains where obtaining labeled data is expensive or time-consuming. For example, in natural language processing, semisupervised learning can be used to improve text classification and sentiment analysis. In computer vision, semisupervised learning can be used to improve object recognition and image segmentation. In medical diagnosis, semisupervised learning can be used to improve the accuracy of disease diagnosis based on medical images. Despite its many benefits, semisupervised learning still faces several challenges. For example, it is often difficult to determine the optimal amount of labeled and unlabeled data to use in the learning process. In addition, the choice of algorithm and the quality of the labeled data can have a significant impact on the performance of the model. Nevertheless, many machine learning researchers have found that unlabelled data, when used in conjunction with a small amount of labelled data, can produce a significant improvement in learning accuracy [75].

- **Reinforcement learning**: Reinforcement learning is a type of machine learning that focuses on training agents to make sequential decisions in dynamic environments. It is inspired by the principles of behavioral psychology, which suggest that learning can be modeled as a process of receiving rewards or punishments for specific actions. In reinforcement learning, an agent interacts with an environment by taking actions and receiving rewards or penalties, and it learns to make decisions that maximize its cumulative reward over time. One of the key features of reinforcement learning is the use of trial-and-error learning. The agent explores the environment by taking

different actions, and learns from the feedback it receives in the form of rewards or penalties. The agent's goal is to learn a policy, which is a mapping from states to actions that maximizes the expected reward. This is typically accomplished using a value function, which estimates the expected cumulative reward for each state. Reinforcement learning has many practical applications in domains such as robotics, gaming, and control systems. For example, in robotics, reinforcement learning can be used to train robots to perform complex tasks such as grasping objects or navigating unfamiliar environments. In gaming, reinforcement learning has been used to develop game-playing agents that can achieve superhuman performance in games such as Go and chess. In control systems, reinforcement learning can be used to optimize the performance of complex systems such as power grids and traffic networks. Reinforcement learning has many practical applications in domains such as robotics, gaming, and control systems. For example, in robotics, reinforcement learning can be used to train robots to perform complex tasks such as grasping objects or navigating unfamiliar environments. In gaming, reinforcement learning has been used to develop game-playing agents that can achieve superhuman performance in games such as Go and chess. In control systems, reinforcement learning can be used to optimize the performance of complex systems such as power grids and traffic networks. Recent advances in reinforcement learning have led to significant progress in addressing these challenges. Deep reinforcement learning, which combines reinforcement learning with deep neural networks, has shown great promise in learning policies for high-dimensional state and action spaces. Other approaches, such as hierarchical reinforcement learning and meta-reinforcement learning, have also been developed to improve the efficiency and flexibility of reinforcement learning algorithms.

Focussing on *supervised learning* techniques, there are two common areas of application [76]:

- **Classification**: Classification is defined as the process of recognising and grouping data into predetermined categories. Classification algorithms used in machine learning use input training data in order to predict the probability that subsequent data will fall into one of the predetermined categories. In machine learning classification, the goal is to learn a mapping between input features and output labels. The training data consists of input feature vectors and their corresponding labels, and the algorithm learns to generalize from this data to make predictions on new, unseen data. The most common type of classification algorithm is the supervised learning approach, which requires a labeled training dataset to learn the mapping from input to output. The aim of the algorithm is to learn a decision boundary that can separate the input

features into different classes. There are many different types of machine learning classification algorithms, each with its own strengths and weaknesses. Some popular algorithms include decision trees, k-nearest neighbors, support vector machines, and neural networks. Decision trees are a simple and intuitive algorithm that are often used for interpretability, while k-nearest neighbors is a simple and flexible algorithm that can work well for many types of data. Support vector machines are a powerful algorithm that can be used for both linear and nonlinear classification problems, and neural networks are a highly flexible algorithm that can be used for many different types of data. One of the most common applications of classification is the filtering of e-mails into "spam" or "non-spam", as used by major e-mail service providers.

- **Regression**: Regression techniques are used to predict a *target* numeric value given a set of input data called *features* (or *predictors*). Machine learning regression is a type of algorithm used to model the relationship between a set of input variables and a continuous output variable. It is commonly used in many fields, including finance, economics, and engineering, where the goal is to make predictions based on past data. In machine learning regression, the aim is to learn a mapping between the input variables and the continuous output variable. This is done by training a model on a set of training data, where the input variables and their corresponding output values are known. The model learns to generalize from this data to make predictions on new, unseen data. There are many different types of machine learning regression algorithms, each with its own strengths and weaknesses. Similarly to the classification algorithms, some popular algorithms include linear regression, decision trees, support vector machines, and neural networks.

The present research work exploits the potentiality of the **supervised learning** to perform energy demand prediction. In particular, two different *Machine Learning* algorithms have been chosen to perform the energy predictions (see sections 2.1.1 and 2.1.1).

The training of supervised learning regression models imply the creation of a mathematical relationship between *features* and *targets* aimed at minimizing an *error estimator* able to evaluate the overall prediction performance of the model. In other words, the performance of a regression model must be reported as a deviation between the predicted and the real values. Consequently, before exploring in detail the mechanisms that characterise the various learning algorithms, it is important to specify which metrics are most commonly used in the regression field.

The **Mean Squared Error** (*MSE*) is calculated as the average of the squared differences between predicted and expected values in a dataset. It is defined as:

$$MSE = \frac{1}{N} \cdot \sum_i \left(y_i - \hat{y}_i\right)^2 \tag{2.1}$$

where y_i is the i^{th} expected value and $\hat{y}_i$ is the i^{th} predicted value. The *MSE* measures the average of the squared differences between the predicted values and the actual values. Squaring the differences has the effect of penalizing larger errors more heavily than smaller errors. The *MSE* is a non-negative value, with a value of 0 indicating a perfect prediction. The *MSE* is widely used in machine learning because it is a simple and easy-to-understand metric. However, it has some limitations. One limitation is that it is sensitive to outliers, since the squared differences of large errors can dominate the overall value of the *MSE*. Another limitation is that it does not take into account the scale of the output variable, which can make it difficult to compare the performance of models across different data sets.

The **Root Mean Squared Error** (*RMSE*) is an extension of the mean squared error. Since it is calculated as the square root of the MSE (Eq. 2.1), the $RMSE$ results conceptually equivalent. Nontheless, the main advantage of this representation lies in the fact that the units of the resulting error is the same of the *target* which the machine learning model try to predict. Consequently, the $RMSE$ results easier to interpret by the user. The $RMSE$ can be calculated as follows.

$$RMSE = \sqrt{\frac{1}{N} \cdot \sum_i \left(y_i - \hat{y}_i\right)^2} \tag{2.2}$$

The **Mean Absolute Error** (*MAE*) is a popular metric because, like RMSE, the units of the error score match the units of the target value that is being predicted. Differently from MSE and $RMSE$, the MAE does not give more or less weight to different types of errors and instead the scores increase linearly with increases in error.

$$MAE = \frac{1}{N} \cdot \sum_i \left|\left(y_i - \hat{y}_i\right)\right| = \frac{1}{N} \cdot \sum_i \left|e_i\right| \tag{2.3}$$

The MAE measures the average of the absolute differences between the predicted values and the actual values. Unlike the MSE, the MAE does not square the differences, which means that it does not heavily penalize larger errors. The MAE is a non-negative value, with a value of 0 indicating a perfect prediction. The MAE is widely used in machine learning because it is a simple and easy-to-understand metric. One of its advantages is that it is not sensitive to outliers, since the absolute differences treat all errors equally. Another advantage is that it is easy to interpret, as it represents the average magnitude of the errors. However, the MAE also has some limitations. One limitation is that it does not take into account the direction of the errors. In other words, it treats over-predictions

and under-predictions equally, which may not be desirable in some applications. Another limitation is that it is not as widely used as the MSE, which can make it difficult to compare the performance of different models.

In statistics, the **Coefficient of determination**, also known as R^2 (see Eq. 2.4), is the proportion of the variation in the dependent variable (*target*) that is predictable from the independent variables (*features*). Unlike the previous metrics, R^2 is a dimensionless value that can vary in a fixed range between 0 and 1, regardless of the absolute value assumed by the *target*. A model which perfectly predict the *target* will present $R^2 = 1$. $R^2 = 0$ is characteristic of a *baselinemodel*, which always predict the *target* average value ($\bar{y}$).

$$R^2 = 1 - \frac{\sum_i \left(y_i - \hat{y}_i\right)^2}{\sum_i \left(y_i - \bar{y}\right)^2} \tag{2.4}$$

The R^2 is widely used in machine learning because it provides a measure of the goodness of fit of the model to the data. A higher R^2 value indicates a better fit of the model to the data, while a lower R^2 value indicates a poorer fit. The R^2 can also be used to compare the performance of different models, as the model with the higher R^2 value is generally considered to be the better model. One limitation of the R^2 is that it can be affected by the number of independent variables in the model. Adding more variables to the model can increase the R^2 value, even if the variables are not significant predictors of the dependent variable. This can lead to overfitting, where the model fits the training data too closely and performs poorly on new data.

Consequently, it is a great metric to obtain a general and comparable evaluation of a predictive model, but it is not recommended to be used as a metric during a training process. Therefore, most models are trained by aiming to minimize the *MSE*, while the final evaluation of the performance is often performed using R^2.

2.1.1 Energy prediction methodologies

Multiple Linear Regression

Linear regression is a statistical approach to modelling the relationship between an independent variable known as a feature or explanatory variable and a dependent variable, also called target. The method allows one to obtain a function that takes the feature as an input variable and return the predicted value of the dependent variable. When the input features are more than one, the method is often called Multiple Linear Regression (or Multivariate Linear Regression). The searched function is the one who minimized the residuals between the prediction and the target values (model error). Generally, the process able to minimize the average error is an iterative algorithm called **Gradient**

Descent [77].

Gradient descent is an optimization algorithm used to minimize some function by iteratively moving in the direction of steepest descent as defined by the negative of the gradient. In the specific case of machine learning, the gradient descent procedure is used to update the parameters of the trained model in order to minimize the so-called *loss function*, which is one of the already described error estimation functions (Eq. 2.1, 2.2, 2.3, 2.4).

Moreover, there is a method to obtain the Multiple Linear Regression (with minimization of MSE) through matrix operations. Such method is often called *normal equation* [78, 79]. The normal equation is a mathematical formula used to find the optimal values for the parameters of a linear regression model. It is derived by setting the gradient of the cost function, which measures the difference between the predicted and actual values of the dependent variable, to zero. The solution to the normal equation provides the optimal values for the parameters that minimize the cost function. The calculation of the optimal values of the parameters can be carried out by solving the following equation:

$$\vec{c} = \left(X^T X\right)^{-1} X \vec{y} \tag{2.5}$$

Given n as the number of training examples and $k-1$ as the number of features, X is an $n \times k$ array that includes all the features values (x_1, x_2, …, x_k), is an n-dimensional vector containing the corresponding target values, and is a vector that contains the $k+1$ linear regression coefficients (c_0, c_1, c_2, …, c_k). The obtained relationship between the features and the target can be represented as:

$$y = c_0 + c_1 \cdot x_1 + \cdots + c_k \cdot x_k \tag{2.6}$$

where the coefficients (c_0, c_1, c_2, …, c_k) represent the internal parameters of the obtained model. The normal equation provides a closed-form solution to the problem of finding the best-fit line in multiple linear regression. This means that it is possible to obtain the optimal values of the parameters without using an iterative optimization algorithm. The use of normal equations can be particularly useful when the number of independent variables is small, as the computational cost of calculating the inverse of the matrix can be high when the number of independent variables is large.

The Multiple Linear Regression technique has both advantages and drawbacks. The main advantages are:

- It can model complex relationships between a dependent variable and multiple independent variables
- It is relatively easy to implement and interpret, and can be used to test hypotheses

and make predictions

- It provides a good understanding of the contribution of each independent variable to the dependent variable

Regarding the drawbacks, they can be summarized in the following points:

- It assumes that the relationship between the dependent variable and the independent variables is linear, which may not always be the case
- It may be sensitive to outliers, which can significantly affect the model results
- It may suffer from multicollinearity, which is the problem of having two or more independent variables that are highly correlated, leading to unstable estimates of the parameters
- It assumes that the residuals of the model are normally distributed, which may not always be the case

Artificial Neural Network: Multi-Layer Perceptron

Artificial neural networks (*ANNs*) are mathematical models that aim at reproducing the learning process of the biological neural system. Indeed, an artificial neural network is a particular algorithm that can learn from input data and exploit the obtained knowledge to perform different tasks. Detailed information about the *ANNs* algorithms can be easily found in the literature [80–82]. An *ANN* is based on a collection of connected units or nodes called artificial neurons. A typical artificial neuron is shown in Figure 2.1. It has m inputs $(x1, \ldots, xm)$ that represents the input features used to predict the target y. The individual inputs are weighted by the corresponding elements $w_{1,1}, \ldots, w1, m$ and then are summed together with the bias (b). The output of a neuron is then computed by a function (known as *activation function*) that is applied to the sum of its inputs. The input data processing can be represented with Eq. 2.7 as:

$$y = f\left(\vec{W} \cdot \vec{x} + b\right) \tag{2.7}$$

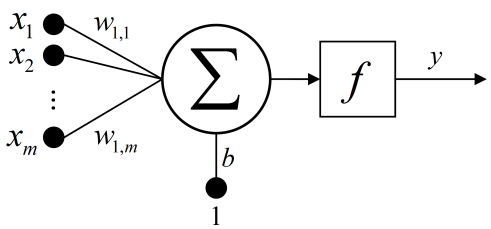

Figure 2.1: Artificial neuron representation

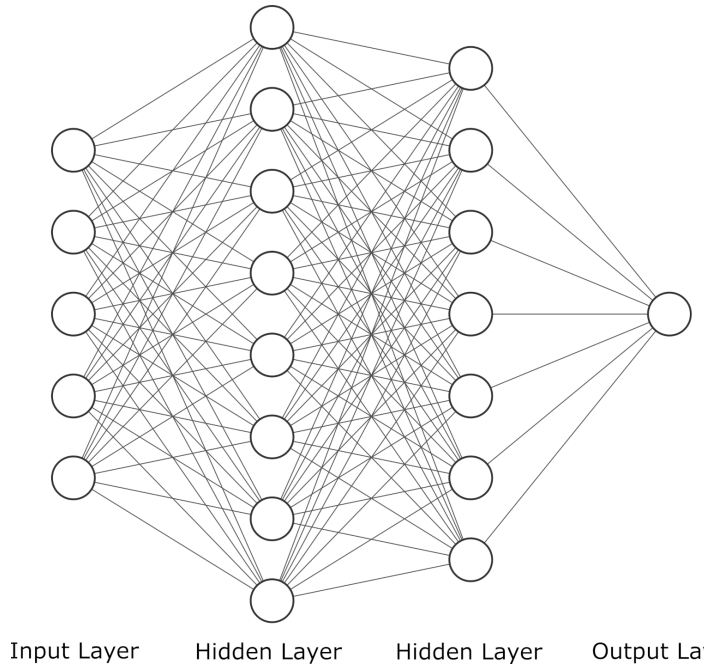

Figure 2.2: Fully connected Artificial Neural Network representation

Where $\vec{x}$ and $\vec{W}$ are the inputs and weights vectors respectively, b is the bias and f is the activation function. Neurons can be combined in artificial neural networks, allowing for complex relations between features and targets. Each neuron receives an input signal, processes it and finally sends it to other connected neurons through the connections that are often called *edges*.

Typically, neurons are aggregated into *layers* (see Figure 2.2), and each of them may perform different transformations on their inputs. One of the simplest and typical *ANN* learning processes is called feed-forward, where the signals travel from the first layer (the input layer) to the last layer (the output layer). Neural networks learn by processing examples, each of which contains known input features and a result (also known as *target*). The *ANN* adjust the weight and biases of the whole network in order to form complex relation between the input features and the target. The training of a neural network from a given example is usually conducted by determining the difference between the processed output of the network and the target output. The logic at the basis of the training algorithm is the already mentioned *gradient descent*. The Artificial Neural Network-based regression algorithm technique has both advantages and drawbacks. The main advantages are:

- It can model complex, nonlinear relationships between the input and output variables, making them a powerful tool for regression analysis
- It can handle large and complex datasets, and can learn from data that is not well-structured or contains noise
- It can perform well in cases where there are many independent variables and

interactions between them, which can be difficult to model with other regression techniques

- It can generalize well to new data and can adapt to changing environments or data

The ANN-based regression presents also several drawbacks:

- It can be computationally expensive, requiring large amounts of data and computation power to train
- It can be computationally expensive, requiring large amounts of data and computation power to train
- It is often considered as black-box model, where the internal workings are not well understood, making it difficult to interpret the results and understand the relationship between input and output variables
- It require careful selection of hyperparameters and architecture, which can be challenging for beginners

The choice of *ANN* typology should be made through proper sensitivity analyses. For example, the effectiveness of an *ANN* may be related to the number of neurons in the hidden layer [82–87] , but a larger number of neurons does not imply greater accuracy of the model. Before the choice of the network architecture and the parameters related to the learning algorithms, it is essential to focus on the available data.

A commonly used pre-processing technique is called *Feature Scaling* [88]. Feature scaling is a method used to normalize the variation range of independent variables. In data processing, it is also known as data normalization and is generally performed during the data pre-processing stage. Since the range of values assumed by the raw data can vary substantially, some machine learning techniques require their normalization in order to regularize and optimize the learning process. It is very important to test with attention the quality and the utility of the input data to the network, realizing analysis and tests aimed to obtain the best set of input features. Besides the quality of the chosen input data, it is very important to pay attention to their representation. Different techniques allow realizing a manipulation of the data (*Feature Engineering*) in order to put the data in a form that is suitable for the learning process.

The *ANN* typology used for this study is called *Feed-forward Neural Network* (*FNN*), which is the simplest type of artificial neural network devised. Since the concept at the very basis is the same of the *MLR*, the *FNN* makes it possible to obtain a "static" model that is able to well represents the facility in the first year of observation, i.e. the year that was selected as 'standard consumption'. Thereby, it will be possible to obtain a comparison between current and "standard" consumption, which will then allow variations

in the energy behaviour of the healthcare facility to be perceived. In this network, the information moves in only one direction, from the input nodes to the output nodes. If no node layer is interposed between input and output, the *ANN* is called *Single-Layer Perceptron*. If one or more layers are interposed between input and output, the resulting *ANN* is called *Multi-Layer Perceptron* and, differently from the former, is capable to create a non-linear relation between inputs and outputs.

Similarly to the *MLR* models, the artificial neural network learning process consists in a iterative variation of the model coefficients in order to pursue the minimization of an objective function (called *loss function*). Several algorithms (*optimizers*) can be used to carry out the loss function minimization. *Stochastic gradient descent* (*SGD*) is perhaps the most well-known optimization algorithm for *ANNs*. It works by iteratively updating the model parameters in the opposite direction of the gradient of the loss function with respect to the parameters. *SGD* uses a small batch of training examples to estimate the gradient, which makes it relatively efficient and scalable to large datasets. However, *SGD* can be sensitive to the learning rate and may require careful tuning in order to converge to a good solution. Another popular optimization algorithm is the *Adaptive Moment Estimation* (*Adam*) algorithm. This algorithm is an extension of *SGD* that uses an exponential moving average of the gradient and the squared gradient to adapt the learning rate on a per-parameter basis. Adam is generally considered to be more robust and easier to tune than *SGD*, and has become a popular choice for training *ANNs*. The *Adam* algorithm will be used in this study as optimizer for all the trained *ANN* models.

Assuming that the data have been correctly processed and the optimizer has been chosen,it is important to focus on the main *ANN* parameters which can markedly influence the learning process. They are often called *hyperparameters*, and they are related to different aspects of the machine learning model.

Activation functions are an important component of artificial neural networks. They are used to introduce nonlinearity into the network and are typically applied element-wise to the output of a linear combination of inputs and weights. Activation functions allow *ANN* to model complex relationships and learn features in data that may not be linearly separable. The main activation functions used for regression purposes are:

- $ReLU$: is a simple activation function that maps all negative input values to 0 and leaves positive values unchanged. It is defined as $f(x) = max(0, x)$.
- *tanh*: Hyperbolic tangent is a smooth, S-shaped curve that maps any input value to a range of -1 to 1. It is defined as $f(x) = tanh(x)$.

In terms of differences, $ReLU$ is generally faster to compute than *tanh* because it has a simpler form and does not require the use of expensive operations such as exponentiation.

However, *tanh* has a more balanced range and can be a better choice for tasks that require centered output, such as regression. Ultimately, the choice between *ReLU* and *tanh* will depend on the specific requirements of the task and the properties of the data. Both activation functions have their own unique strengths and limitations, and it is important to carefully evaluate their suitability for the task at hand.

The *learning rate* (lr) is a key hyperparameter in the training of an artificial neural network (ANN). It determines the size of the steps that the network takes in adjusting its weights and biases during the training process. A high learning rate can lead to fast convergence, but also the risk of oscillation or divergence. A low learning rate can lead to slow convergence, but also a more stable and accurate model. The choice of *learning rate* can significantly impact the performance of the trained ANN, and finding an optimal *learning rate* can be an important part of the model design process.

The number of neurons in the hidden layers (HLN) of an artificial neural network (ANN) is an important architectural hyperparameter that can significantly impact the model's ability to learn and generalize to unseen data. The hidden layers of an ANN are responsible for learning and representing the underlying patterns and relationships in the input data. Traditionally, the number of neurons in the hidden layers of an ANN has been chosen heuristically, based on the size and complexity of the input data and the desired model capacity. A common practice is to assign the number of neurons in the hidden layer based on the number of *input features* and *outputs* of the model. Various *rule-of-thumb* relationships have been proposed in the literature [56]. In this study, the choice fell on the sum of the number of input features and the number of outputs (Eq. 2.8).

$$HLN = n_i + n_o \tag{2.8}$$

This is still a choice based on the experience of the authors, and certainly does not guarantee the identification of the optimal value. This typologies of hyperparameter choices will be referred to later as the *literature-suggested*. However, recent research has shown that the choice of the number of neurons in the hidden layers can be critical for the model's ability to learn and generalize. Therefore, more advanced methods such as the *grid-search approach* will be used to fine-tune various hyperparameters.

In order to ensure the correct functioning of the learning process and to reduce the computational cost, it must be defined a condition known as *Early Stopping*. Early stopping is a technique used to prevent *overfitting* in artificial neural networks ($ANNs$). It involves monitoring the performance of the ANN on a separate validation set during training and interrupting the training process when the performance on the validation set starts to degrade. The idea behind early stopping is that if the ANN continues to be trained after the validation performance has started to decrease, it will continue to fit the

training data better and better, but at the expense of generalization to new, unseen data. By interrupting the training process before this point, it is possible to obtain a model that has lower training error but also better generalization performance. Early stopping is typically implemented by monitoring the performance on the validation set after each training epoch and interrupting training when the performance has not improved for a certain number of epochs. The number of epochs for which the performance is allowed to degrade before training is stopped is known as the *patience* of the early stopping.

Consequently, in order to correctly perform the training process, the available dataset must be divided into three parts:

- the *training set* (70%): shown to the algorithm to perform the learning.
- the *validation set* (10%): used to evaluate the performance of the model at the end of each iteration of the learning process, in order to eventually trigger the Early Stopping condition.
- the *test set* (20%): used at the end of the process to evaluate the prediction performance of the trained model.

Performance evaluation during the learning process is very important to ensure that the process goes smoothly. The main risk is that the model focuses too much on the data it is shown, failing to predict instants with different characteristics from those in the *training dataset*, losing generality. A correct learning process, consequently, must exhibit consistent variations in the loss function (i.e., the chosen error metric which, in this case, is the mean squared error) along the various Epochs (iterations). From the moment that a continued increase in validation loss occurs with continued reduction in training loss, the model will have lost its effectiveness on data not previously seen (overfitting).

2.1.2 Feature Selection

In machine learning and statistics, **feature selection** (also known as **variable selection**) is the process of selecting a subset of relevant *features* which will be used for the training of a numerical model. There are several reasons to perform an accurate procedure of feature selection before approaching the training of a machine learning model.

The guiding principle is the identification of the best subset of data which allow to obtain the best model in terms of prediction performance. Basically, identifying an useless feature can lead to several problems. Indeed, as the amount of data increase, also the computational effort necessary to complete the model training results higher. Moreover, also the effort needed to collect the data increase with the number of features. It must also be considered that a feature that is not useful for increasing the predictive performance

of a model may also be harmful. A classic example is the problem of *overfitting*. In mathematical modelling, *overfitting* is the production of an analysis that correspond too closely (or exactly) to a particular set of data, and may therefore fail to fit to additional data or predict future observation reliably. In other words, and overfitted model is a mathematical model that contains more features than can be justified by the dataset. Nontheless, the elimination of too many features can be even more harmful to the model prediction capabilities, leading to the *underfitting* problem. Indeed, an *underfitted* model is a model where the choosen subset of feature used to the training process are not enough to train a model capable to correctly represent the target behaviour. There are several methods to perform a feature selection process [89], and they can be classified in two main categories:

- **Filter methods**: The selection of feature is indipendent of any machine learning algorithms, and the features are selected based on their scores in various statistical tests for their correlation with the target variable. A high correlation with the target is generally interpreted as a high importance as an input feature of the machine learning model.
- **Wrapper methods**: These methods aim to select the features according to their actual influence on the model performance. Generally, the predictive model is trained multiple times with different features subsets, in order to obtain a **feature ranking** and, subsequently, select the **best feature subset**.

The following sections will be dedicated to provide a more detailed insight of the **feature selection methods** used for the subsequent analysis.

Filter Method

Figure 2.3: General flowchart for filter method-based feature selection

As already said, the filter methods base their functioning on various statistical tests aimed to quantify the correlations between the features and the target. In the literature, the most common statistical parameters used in the energy prediction field are the **correlation coefficients**. In particular, two correlation coefficients have been investigated.

The **Pearson's product-moment correlation coefficient** (r) can be used to assess the correlation characteristic of each pair of parameters, and it is defined as:

$$r_{xy} = \frac{\sum_{i=1}^{n}\left(x_i - \bar{x}\right)\cdot\left(y_i - \bar{y}\right)}{\sqrt{\sum_{i=1}^{n}\left(x_i - \bar{x}\right)^2}\cdot\sqrt{\sum_{i=1}^{n}\left(y_i - \bar{y}\right)^2}} \tag{2.9}$$

where x_i and y_i are the value assumed by the two variables in each sample, $\bar{x}$ and $\bar{y}$ are the sample mean values and n is the number of samples. The Pearson's coefficient ranges from -1 to 1. An absolute value of 1 implies that a linear equation describes the relationship between x and y perfectly, with all data points lying on a line. The correlation sign is determined by the regression slope, meaning that a value of $+1$ implies a total positive linear correlation, in which y increases as x increases. Conversely, a value of r equal to -1 implies a total negative linear correlation between the two analysed parameters. Finally, a value of 0 implies that there is no linear dependency between the considered variables.

Nevertheless, the Pearson coefficient is not capable to evaluate non-linear relationships between the variables. To overcome this limitation, the **Spearman's rank correlation coefficients** (ρ) were also evaluated. Spearman is a nonparametric measure of rank correlation (conceptually similar to the Pearson's coefficient), which let to assess how well the relationship between two variables can be described using a monotonic function. To obtain the spearman correlation coefficient between two features, each feature must be converted to a rank. Then, the spearman correlation can be calculated as:

$$\rho_{xy} = \frac{\sum_{i=1}^{n}\left(s_i - \bar{s}\right)\cdot\left(t_i - \bar{t}\right)}{\sqrt{\sum_{i=1}^{n}\left(s_i - \bar{s}\right)^2}\cdot\sqrt{\sum_{i=1}^{n}\left(t_i - \bar{t}\right)^2}} \tag{2.10}$$

where s and t represents the ranks of each feature.

The energy characterisation of any activity involves understanding the aspects that mainly influence the demand, i.e. finding the Energy Drivers. To do so, a detailed analysis of the correlations between electricity consumption and various parameters (e.g., calendar or climate features) must be done, defining the main building energy drivers. The correlation analysis is useful to perform a preliminary feature selection through a method often known as "filter method" [89]. In other words, the correlation analysis lets to identify features that might not be useful for improving the model prediction performance.

A machine learning algorithm is based on creating a model which exploits the correlation between one or more features and one or more targets. To carry out a good feature selection is pivotal to check the possible multicollinearities between each pair of input features. Two strongly correlated input features would provide the model with the same information. In other words, if two features result strongly dependent on each other, maintaining both features as input of the machine learning algorithm could not lead to improvement of the

model accuracy, because both features bring the same information. This type of feature is called **redundant features**.

By observing the correlations between the input features and the target, correlation analysis leads to the identification of features that have an excessively weak correlation with the target and is therefore useless for model prediction performance improvement. This type of feature is called **irrelevant features**.

Consequently, the identification of both feature typologies implies the definition of two different thresholds for the correlation coefficients: A lower bound on input feature correlations with the target (i.e., electricity demand) is used to identify the irrelevant features, while an upper bound on the correlation between each pair of input features is used to highlight redundancies in the input dataset. In the literature, some examples of threshold choices are provided. Nevertheless, these thresholds are only literature-suggested values. The thresholds used for a specific application should be chosen as a compromise between prediction accuracy and computational efficiency. The non-systematic nature of this feature selection methodology turns out to be its main drawback, as it requires a great deal of experience of the application context (building and its installations) as well as in-depth knowledge of machine learning and predictive methods. Setting thresholds incorrectly could lead to the exclusion of potentially important features, or to the inclusion of completely unnecessary features (resulting in an increased computational cost).

Wrapper method

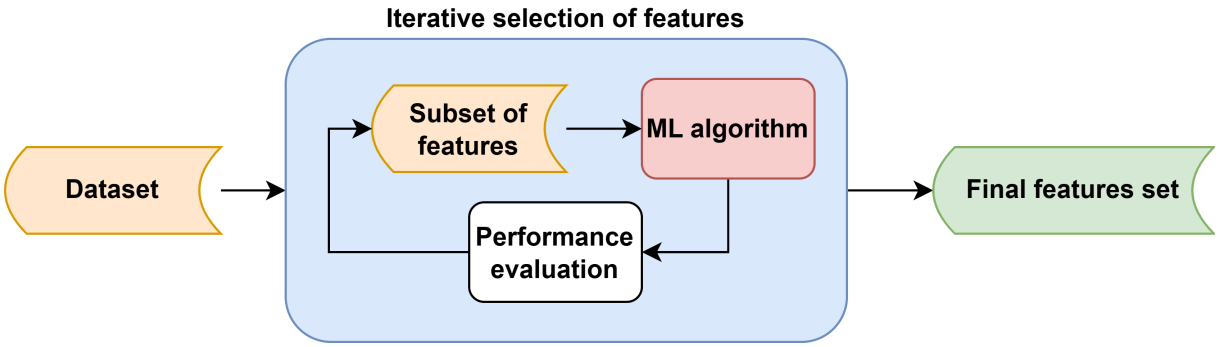

Figure 2.4: General flowchart for wrapper method-based feature selection

Differently from the *filter methods*, the **wrapper feature selection methods** aim to provide a feature renking based on the real influence that every feature has on the model prediction performance. The wrapper methods are basically iterative process that perform multiple model training with different subsets of features and measure the prediction performance based on the already mentioned performance metrics. The necessary number of training and, consequently, the overall computational cost depends on the logic used to select the various subset. Nevertheless, generally the wrapper method are markedly more

expansive than wrapper method in terms of computational costs.

Some widely used wrapper methods are:

- **Forward Selection**: Forward selection is an iterative method which is basicaly divided in a number of steps equal to the number of features. The forward feature selection procedure begins by evaluating all subsets of features that consist of a single feature. In other words, it starts by measuring the prediction performance of models trained with subsets of one components, $X_1, X_2, \ldots, X_M$ (where M is the *feature array dimensionality*). In this process, the best individual feature, $X(I)$, will be identified. The next step is to identify the best subset consisting of two components, $X(I)$ (best feature from the previous step) and another feature among the remaining $M-1$ features. Thus, there are a total of $M-1$ pairs of features to be evaluated. Once $X(II)$ (second best feature that pairs with $X(I)$) is found, these two features will be paired with the remaining features to identify the third best feature. This process will continue until each feature is evaluated, resulting in a ranking of feature importance as output.

- **Backward Elimination**: Conceptually similar to the previous one, it differs from it by starting with performance evaluation using all features, and eliminating step-by-step the one that is least useful for increasing prediction performance.

One of the main advantages of forward elimination is that it guarantees that the final feature set is optimal. This means that the algorithm will not stop until the best feature subset has been found. Forward elimination is also useful when the number of features is relatively small and the computational cost of training the model is not a concern.

On the other hand, one of the main advantages of backward elimination is that it is computationally efficient. This can be particularly useful when dealing with large datasets or when the cost of training the model is a concern. Additionally, backward elimination can handle a larger number of features compared to forward elimination. However, backward elimination has some drawbacks. One of the main drawbacks is that it is not guaranteed to find the optimal feature subset. This is because the algorithm stops when a stopping criterion is met, which may not necessarily lead to the best feature subset. Since the algorithm proceeds by iteratively removing features one by one, it is more difficult to identify the contribution of each individual feature to the overall performance of the model. Additionally, backward elimination is more prone to overfitting than forward elimination, especially when the number of features is large.

To make sure that the evaluation is general and not affected by the choice of the specific test dataset, each training process performed during the wrapper method is carried out using the procedure known as *k-fold cross validation* [90]. *K-fold cross-validation* is

a widely used method for evaluating the performance of machine learning models. In a *k-fold cross-validation*, the dataset is divided into k subdatasets of equal size (typically 3, 5 or 10). Among the subdatasets , only one subdataset is kept as test data to evaluate the model, while the remaining $k-1$ subdatasets are used as training data to determine the model parameters. The cross-validation process is then repeated k times, with each of the k subsets used sequentially as test data, while the other $k-1$ subsets are used as training sets. The performance evaluation metrics resulting from this process will then be averaged across the k trainings to determine the actual performance of the model implemented. One of the main advantages of k-fold cross-validation is that it provides a more reliable estimate of the performance of the model than a single train-test split. This is because each instance in the data is used for both training and testing, and the performance metric is averaged over multiple iterations. K-fold cross-validation also helps to reduce the impact of random variation in the data, as each instance has an equal chance of being in the test set. One of the main drawbacks is that it can be computationally expensive, especially when dealing with large datasets or complex models. This is because the model needs to be trained k times, and the time required for training can increase significantly with the number of folds and the complexity of the model.

2.1.3 Feature Engineering

Despite the chosen machine learning methodology, the input features influence more than everything else the model performance, and the *Feature Selection* is therefore a pivotal step of the propose method. Nevertheless, a practice called *Feature Engineering* can further enhance to the model performance. As already said, *Feature engineering* refers to the process by which domain knowledge is used for extracting features, properties, or attributes from raw data. The types of feature engineering used in this study mainly concern three aspects:

- *physical relations*
- *context knowledge*
- *numerical representations*

A classical example of **physics-related** feature engineering is the exploitation of *Psychrometry* to obtain new features from the measured climate features. In particular, based on dry bulb air temperature (T_{db}), relative humidity (RH) and air pressure (p_{amb}), it is possible to obtain several new features: wet bulb air temperature (T_{wb}), dew point temperature (T_{dp}), vapour pressure (p_{vap}), humidity ratio (HR), dry and moist air enthalpy (h_{dry}, h_{wet}), dry and moist air density (ρ_{dry}, ρ_{wet}). The psychrometric relation of the moist

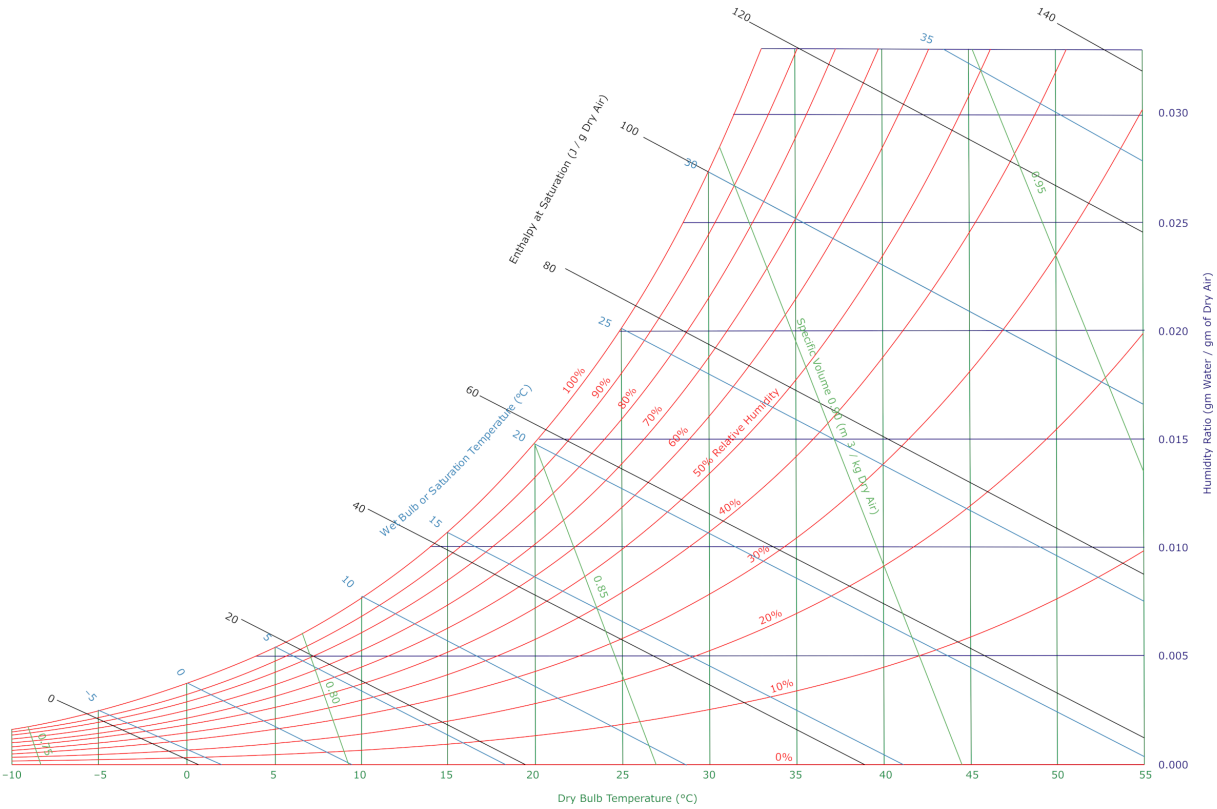

Figure 2.5: Moist air psychrometric chart at ambient pressure (sea level)

air can be resumed through the psychrometric chart, which shows the moist air properties at constant pressure (see Figure 2.5).

The whole building electricity demand is mainly guided by the requirements of *HVAC* systems. These system can be markedly influenced by the external air thrmodynamic conditions. As an example, a summer day with high humidity could lead air handling units to require more cold water to carry out de-humidification. The cold water is provided by the refrigerators, which, in turn, are cooled by evaporative cooling towers. Different climatic conditions can affect the performance of the evaporative towers and, consequently those of the refrigeration units, leading to a higher energy demand for the same amount of cooling energy required by the air handling units. These statements lead to the conclusion that different representations of external weather conditions may allow the machine learning model to better fit the system's response, allowing better predictions simply by providing it with data with a different representation.

Considering the great influence of air treatment on overall activity consumption, data processing was carried out to account for the change in the building's indoor temperature set point during the year. From *April* 15^{th} until *October* 15^{th}, the system configuration assign the set point temperature of 25 $°C$ (except some specific areas like the Nuclear Magnetic Resonance room). The rest of the year is characterized by a temperature setpoint of 20 $°C$. Consequently, two feature has been created in order to account not only the outdoor temperature, but the difference between the outdoor temperature and

the average indoor temperature requirement. For convenience, the features have been called *Heating Degree Days* (*HDD*) and *Heating Degree Hours* (*HDh*), although, unlike the classic definition of heating degree days (Eq. 3.1), they also include the summer period. In particular, the *HDD* feature assume values based on the following relations:

$$HDD = \sum_{j=1}^{n} \left(20°C - \bar{T}_j \right) \Big|_{\bar{T}_j < 20°C} \tag{2.11}$$

$$HDD = \sum_{j=1}^{n} \left(\bar{T}_j - 25°C \right) \Big|_{\bar{T}_j > 25°C} \tag{2.12}$$

where $\bar{T}$ represents the daily average outdoor air temperature. Eq. 2.11 is used in the period in which the indoor temperature set point is 20 $°C$, while eq. 2.12 is used during the rest of the year. As regards *heating degree hours* (*HDh*), it is defined base on the same periods and indoor temperature set points defined above, but its computation is not made with daily averaged values, but with the instantaneous ones.

In the field of building energy feature engineering, weather-related composite features such as *ClimateZ* can be performed. The latter is a representation of dry bulb temperature, wet bulb temperature and global solar radiation as represented which is often useful in the field of machine learning-based energy prediction (see 2.13). This feature does not really make physical meaning, but it is widely used in the field of building energy prediction, as it has often proven to be extremely useful to machine learning algorithms [89, 91]. It is a linear combination of dry bulb temperature, wet bulb temperature and global solar radiation as represented in Eq. 2.13:

$$ClimateZ = 0.97 \cdot T_{db} + 0.96 \cdot T_{wb} + 0.77 \cdot SR_{gl} \tag{2.13}$$

Regarding feature engineering related to **context knowledge**, it can be applied to different aspects. In this case, a relation can be found between the energy behavior of the building and the hour of day (see section 3). In fact, the demand for electricity turns out to have a repetitive behavior with an even 24-hour period, reaching its minimum during the night and its maximum during the day. The numerical representation of the hour of the day ($Hour$) starts from 0.00 (00:00) and goes to 23.75 (23:45). There are two main inconsistencies in its relation with the building power demand. First, the power demand at 23:45 ($Hour = 23.75$) will likely be very similar to the midnight one ($Hour = 0.00$). Consequently, representing the corresponding hour values with two completely different numbers (0.00 and 23.75) could not be the optimal solution. In other words, the hour of the day must have continuity across the midnight in order to better represent the nearly periodic power demand daily behaviour characteristic of the test case under analysis.

Moreover, as already said, the power demand presents a lower and a higher peak during a single day which are not represented by the aforementioned numerical of the Hour feature. These two aspects can be resolved by representing the hour of the day with a *properly shifted sine function*. The shift choice was based on *Spearman correlation coefficient* (Eq. 2.10), choosing the shift value which maximize the correlation with the building electricity demand. Nonetheless, this representation raises a new issue. Since the sine function is periodical, there will be two different hours of the day represented by the same value of the sine function. This fact can influence negatively the learning process, leading the *ANN* to associate two different moments of the day to the same value of the hour value. In literature, there are some works [92, 93] that resolve the problem by representing the hour of the day with two different features characterized by both sine and cosine function, so that each time-step will be characterized by a univocal pair of values ($\sin(h)$ and $\cos(h)$). A similar reasoning has been done regarding the *Month* feature, considering the yearly periodicity of the building electricity demand.

The last category of *feature engineering* concerns the exploitation of different **numerical representations** of the same feature. First of all, the concept of *categorical feature* must be defined. *Categorical features* can only take on a limited, and usually fixed, number of possible values. As an example, the *DayTyp* feature, which distinguish between workdays, pre-festivity days and festivities, has been firstly represented with a single vector which can assume values of *1*, *2* or *3*. This kind of representation is generally know as *ordinal encoding*. Nevertheless, these kind of features have not a real "order" in comparison to the target values. In this cases a viable alternative is the so-called *One-Hot encoding*. In this representation, a vector made up of elements that can assume a finite number n different values can be replaced with n binary vectors (which can therefore assume *1* or *0* values), each of which will have unit value only if the source vector assumes the corresponding value.

To ensure the correct learning process of a machine learning algorithm, it is good practice to carry out so-called feature scaling. This process involves making sure that the features in the dataset are all on the same scale [94]. In general, all real-value features of the data set should be between -1 and 1 or a range around this region. An interval that is too large or arbitrarily too small will lead to a slowing down of the learning process itself. In this case, all features with a range of variation larger than (-1,1) will be scaled using the method known as *Robust Scaler*. This Scaler removes the median and scales the data according to the quantile range (defaults to *IQR*: Interquartile Range). The *IQR* is the range between the 1^{st} quartile (25^{th} quantile) and the 3^{rd} quartile (75^{th} quantile). The *Robust Scaler* is applied as follows:

$$x_{scaled} = \frac{x_i - Q_1(x)}{Q_3(x) - Q_1(x)} \tag{2.14}$$

where x_i is the i^{th} feature sample, while $Q_1(x)$ and $Q_3(x)$ represent respectively the first and the third *quartiles*.

One of the main advantages of the robust scaler is that it can help to improve the performance of machine learning models that are sensitive to outliers. This is because the robust scaler scales the data in a way that is less influenced by extreme values, which can distort the mean and standard deviation. Another advantage of the robust scaler is that it can be used with non-Gaussian distributed data, which can be problematic for other normalization methods. This is because the median and interquartile range are more robust measures of central tendency and spread than the mean and standard deviation. On the other hand, one of the main disadvantages is that it can be computationally expensive, especially when dealing with large datasets. This is because the median and interquartile range need to be calculated for each feature, which can be time-consuming for high-dimensional data.

The *target* was also scaled to ensure uniformity of the learning process. In contrast to the features, it was simply scaled to the maximum value it takes during the training period. In this way, the output of the prediction model can be directly used, and visualised without distorting the proportions to the electrical power demand data.

2.1.4 Hyperparameter tuning: Grid-search

Hyperparameter tuning is a crucial step in the process of building a neural network, as it involves fine-tuning the parameters that govern how the network learns from data. These parameters, which are set prior to training and are not learned during the training process, are often referred to as hyperparameters. The selection of optimal hyperparameters can significantly impact the performance of a neural network. Hyperparameter tuning involves selecting the optimal values for these hyperparameters. This is typically done through a combination of manual tuning and automated search methods such as *grid search*, *random search*, and *Bayesian optimization*. The goal is to find the hyperparameters that result in the best performance on a validation set while avoiding overfitting on the training set.

Grid search is a technique used to find the best possible combination of hyperparameters for a machine learning model. The idea behind grid search is to create a grid of all possible combinations of hyperparameters and then evaluate the performance of the model for each combination. In other words, it involves constructing a grid of parameters, and then running the model for each combination of those parameters to determine which combination results in the best performance. Grid search is an exhaustive search method,

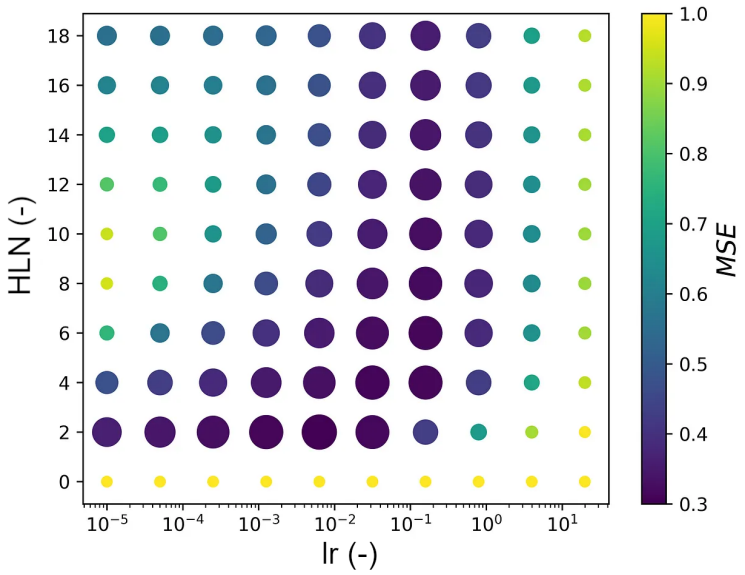

Figure 2.6: Graphical representation of the hyperparameters grid for Grid Search-based Hyperparameter tuning

meaning that it will consider all possible combinations of hyperparameters before it finds the best one. It is also relatively simple to implement and can be used to tune any machine learning model.

In general, the grid search approach can be divided in several steps:

- **Define the hyperparameters**: The first step is to define the hyperparameters to be tuned. For example, the learning rate, number of hidden layers, and number of neurons per layer
- **Define the ranges of values**: For each hyperparameter, define a range of values to explore. For example, the learning rate could be explored over the range [0.1, 0.01, 0.001]
- **Create a grid of all combinations**: Create a grid of all possible combinations of hyperparameters. For example, if there are 3 hyperparameters, each with 3 values to explore, there will be a total of $3^3 = 27$ combinations
- **Train and evaluate the model for each combination**: Train and evaluate the model for each combination of hyperparameters using a validation set. For example, if there are 27 combinations, the model will be trained and evaluated 27 times
- **Select the best combination of hyperparameters**: Select the combination of hyperparameters that performs the best on the validation set. This is typically done by selecting the combination that achieves the highest accuracy (Classification purposes) or lowest loss (e.g. *MSE* for Regression purposes)

- **Test the model with the selected hyperparameters**: Finally, test the model with the selected hyperparameters on a separate test set to evaluate its performance

While grid search is a simple and effective method for hyperparameter tuning, it can be computationally expensive, especially when the number of hyperparameters and their ranges are large.

To enhance the accuracy and generality of the ongoing evaluation, the *k-fold cross validation* methodology will be applied for each training process, as well as for feature selection. As in the previous step, the analysis will certainly be more accurate, general, and less prone to errors due to the selection of the validation and test sets. However, the computational cost will triple as a drawback. Nevertheless, compared to the feature selection procedure, the grid search process allows for working on the number of evaluated combination defined in the hyperparameter grid, thereby reducing the impact of the grid search on the overall computational cost, which, for the current study, remains significantly lower than that required by the feature selection procedure via the wrapper method.

In this study, the hyperparameters chosen for the test are three: activation function, learning rate (lr) and number of neurons in the hidden layer (HLN). These are the three parameters that most influence the machine learning of an artificial neural network: the activation function, the learning rate and the number of neurons in the hidden layer.

The *activation function* can be the linear rectifier ($ReLU$) or the hyperbolic tangent ($tanh$), because both are most commonly used for regression applications. As far as the *learning rate* (lr) is concerned, its value generally takes on similar values in different applications, being less variable depending on the input dataset. The values tested will be 5 and will vary between 0.0001 and 0.01. Values higher than 0.01 generally make it difficult to reach the model training convergence, while values lower than 0.0001 will results in a slow and costly training process, with no advantages in terms of prediction performance. As regards the *number of neurons in the hidden layer* (HLN), these parameter are more strictly related to the number of input feature. It is impossible to define a "standard grid" applicable to all the cases. Consequently, the choice of HLN parameter grid will be done based on multiples of the number of hidden layer used during the feature selection procedure, which is the sum between number of neurons of input and output layers (see Eq. 2.8).

The grid search procedure does not necessarily have to be done in a single step. Contrarily, in order to reduce the computational cost, it is highly recommended to proceed step by step, starting with a coarse grid, and then making it denser the evaluation points only where necessary.

In some cases, it will be evident that there is a minimum mean error for a well-defined combination of parameters, making the choice of final hyper-parameters evident. In other

situations, a low error point may be reached, which will persist as a parameter changes. For example, a minimum may be reached for a well-defined combination of learning rate, activation function and *HLN*, which persists even as *HLN* increases. This situation makes it clear which activation function and learning rate are optimal, but no *HLN* value can be identified to minimise the error. In that case, it is advantageous to choose the combination with the lowest possible *HLN* that allows the average error to be kept to a minimum value, as choosing a larger number of neurons would increase the computational cost without bringing any benefit in terms of performance.

In the cases that will be analysed below, the grid search procedure was developed in successive steps, making the grid denser as needed. There is no difference between this procedure and the development of the grid search in one step. Consequently, the results will be presented all together, so that the choice of the optimal set of hyperparameters is clear.

2.1.5 Statistical method for energy monitoring: CUSUM

The *CUSUM* (cumulative sum) method is based on the concept of monitoring the deviation of individual observations from a target or reference value. It is a sequential analysis technique, where the focus is on detecting changes in a monitored process behaviour as soon as possible after they occur. The *CUSUM* statistic is a running sum of the differences between each observation and the target value, with positive deviations adding to the statistic and negative deviations subtracting from it [95–97]. The *CUSUM* method is particularly useful for detecting small shifts in the process behaviour that may not be immediately apparent with traditional historical data analyses. It is also useful for detecting changes in the process that occur gradually over time, and it can be used in both univariate and multivariate settings. In the univariate setting, it is used to detect behavioural changes of a single variable, while in the multivariate setting, it is used to detect changes of multiple variables.

The entire *CUSUM* technique is based on the development of an effective predictive method that allows the prediction of the process variable under analysis (in this case, energy consumption) under standard operating conditions of the monitored system or plant. The development of accurate predictive models is therefore a fundamental requirement to achieve effective monitoring. It involves the use of a cumulative sum of the deviations from a target or reference value, which is then used to detect changes in the energy demand behaviour. The *CUSUM* method is particularly useful for detecting small or gradual changes in the monitored target behaviour, and is often used in conjunction with other statistical techniques such as control charts. The main advantage of the *CUSUM* method is its ability to detect changes in energy demand mean quickly and accurately, making it

an important tool for energy monitoring and optimization.

The *CUSUM* statistic is plotted on a control chart, with a preset reference value or threshold (known as the *decision interval*) used to determine when a shift in the target parameter behaviour has occurred. If the *CUSUM* statistic exceeds the decision interval, an alarm is triggered and the process is considered to be out of control. The decision interval is typically set based on the trade-off between the need to detect small changes in the energy demand and the need to minimize the number of false alarms. The *CUSUM* algorithm can be also used in real-time, allowing for prompt detection and response to process deviations.

The *CUSUM* method is a statistical process control technique that is widely used in the field of energy monitoring [98]. After the machine learning model training, it is used to predict the "standard" building energy demand (C_{std} in Eq. 2.15). Despite the complexity of the training algorithm, the energy characterization of a building can be difficult. Its energy demand behaviour in function of the energy drivers can change over time, due to many different causes as changes in the building activity or systems management strategy. If on one hand, a higher energy demand due to a building systems fault may be evident, other energy demand behaviour changes can be mild and difficult to find by simply analyse the historic data.

In order to better define the *CUSUM* method, Eq. 2.15 shows how the *CUSUM* can be calculated in for the j^{th} time step of a dataset.

$$S_j = \sum_{i=1}^{j} \left(C_i - C_{std,i} \right) = S_{j-1} - e_j \tag{2.15}$$

where S_j is the cumulated sum of the j^{th} time step, C_i is the current energy demand value, $C_{std,i}$ represent the aforementioned standard (predicted) energy demand for the current time step and e_j is the current residual. If the real power demand value is higher than the predicted one, S_j will be positive. Since the S value of a specific time step is computed as the cumulative sum of differences of all the preceding time steps, an increase in S will imply a higher energy consumption than the standard one. Contrarily, a continuous decrease in S implies that the building electricity requirements result constantly lower than the expected one.

The application of the *CUSUM*-based monitoring method involves the definition of "alarm" thresholds, which enable the control chart to detect and report anomalies in real time. The threshold is identified based on the concept that the model is trained with the *training dataset*. Since the training is aimed at reducing the mean squared error (*MSE*) of the model in that specific period, the fluctuations in S during the training period should be considered normal consumption behaviour. Consequently, once the method is applied

on new data, if the value of S exceeds the thresholds set based on the training period S behaviour, an alarm must be notified. Generally, the threshold is set by calculating the standard deviation of S in the training period (σ), multiplying it by a constant (k). A low value of k will result in a higher sensitivity of the method, as well as a higher frequency of alarms, while a higher value will reduce it. This threshold can also be tested, experimented and calibrated during real-time operation of the method, constituting a real "operative" control parameter. The most commonly used value of k, however, is 3.

2.2 Proposed method

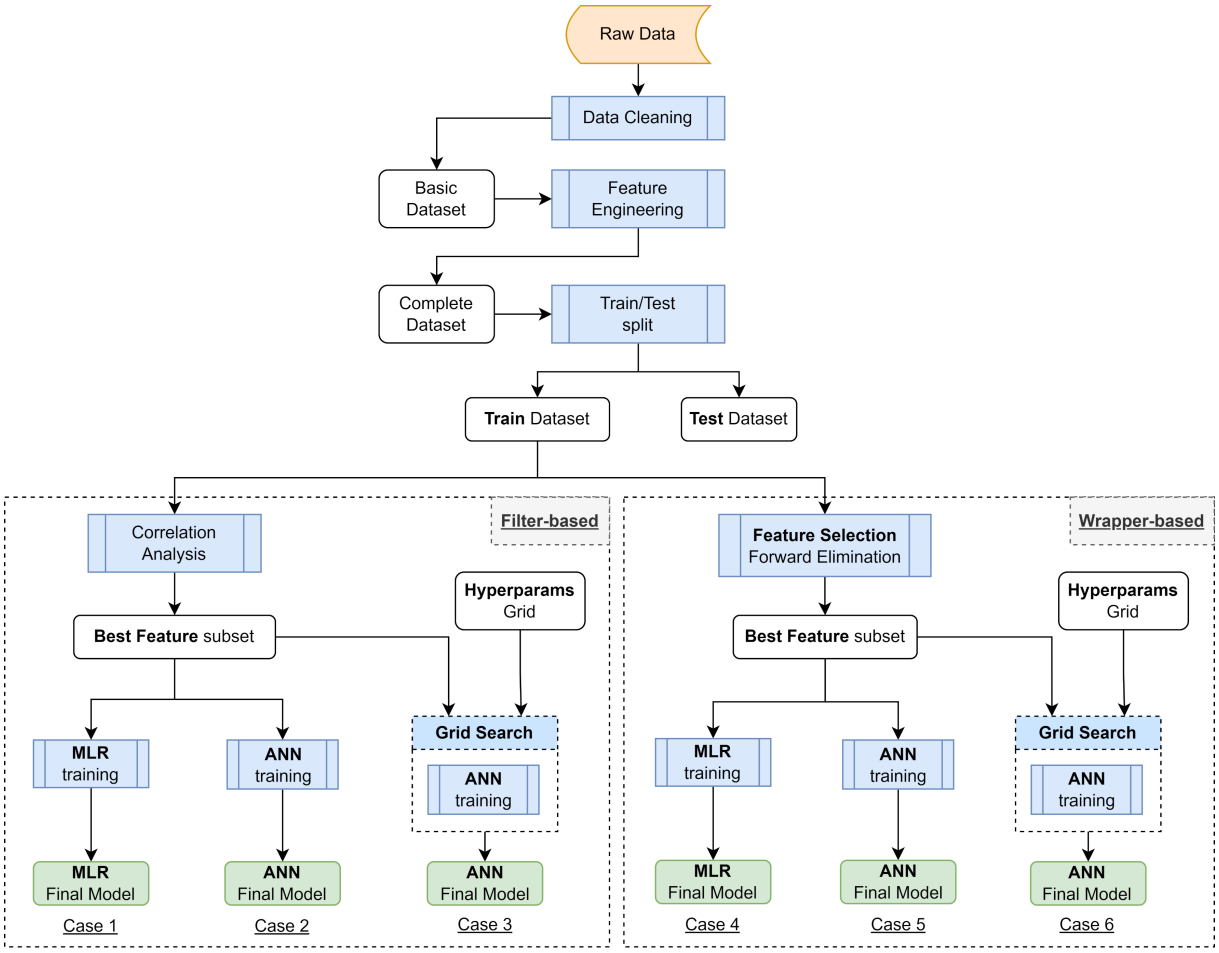

Figure 2.7: Proposed method conceptual scheme

The aim of this work is to create a systematic methodology to proceed from the raw data to a ready-to-implement monitoring method, minimizing the knowledge required to the user and, contemporary, ensure reliable performance of the monitoring method itself.

To reach this goal, the present work proposes a comparison of different methodologies to perform building energy predictions (which is the core of a machine learning-based monitoring method) aimed at building energy monitoring applications. This type of application can be extremely useful in monitoring the energy demand behaviour of a

building, as well as an important method of verifying the effect of any interventions made on the systems and their management strategies. However, depending on the method used to perform the energy prediction, these techniques can be computationally expensive, requiring computational efforts that cannot always be sustained by companies and activities. The proposed method of analysis therefore aims to perform a performance and computational cost analysis of different methodologies for training machine learning models, in order to establish which is the best in terms of cost-benefit ratio.

As already said, the entire process leads from raw data to the final application of the monitoring method. Figure 2.7 shows the flowchart of the proposed model, classifying it into several different cases (which represent different alternative approaches) and dividing it into several steps which are further described below:

1. **Data cleaning and Feature Engineering**: The raw data is analysed and cleaned, removing any implausible or obviously incorrect data on both features and target, such as too high outdoor air temperature values or negative power demand. Moreover, a data elaboration has been carried out through a procedure often called Feature Engineering (further information in sections 2.1.1 and 2.1.3). *Feature Engineering* is the practice to elaborate the available data using physical or mathematical relations aimed at extracting new information from data or to put it in different representations in order to enhance the training procedure [99]. In this first stage, the feature engineering can be performed in order to obtain as many features as possible. At the end of this stage, a complete dataset will be available to be used for the further steps of the proposed method.

2. **Feature selection**: As already explained in section 2.1.2, this step is fundamental in order to obtain the best feature subset, eliminating those features that are unnecessary or even harmful to the prediction performance of the model. This can be the most expensive step in terms of computational costs. Consequently, the choice of the feature selection method must be accurately evaluated based on the context of application. The proposed procedures exploit two different feature selection methods: Filter method (section 2.1.2) for **Cases 1**, **2** and **3** and wrapper method (section 2.1.2) in the form of *Forward Elimination* for **Cases 4**, **5** and **6**.

3. **Hyperparameters tuning**: This step is only performed for certain types of machine learning models. A simple Multiple Linear Regression model trained using the normal equation (see section 2.1.1) does not need to perform this step, as the equation found is already the one that maximises prediction performance. An artificial neural network, on the other hand, has different types of parameters to work on (see section 2.1.1). The optimal parameters are different for each dataset and, consequently, it

is important to carry out a tuning procedure in order to ensure maximisation of performance. For the present work, two types of *hyperparameters* choice will be used. The first one (**Cases 2** and **5**) consists of the choice of *hyperparameters* based on the most widely used choices in the literature, which achieves good performance while avoiding a computationally expansive procedure for parameter optimization. The second (**Cases 3** and **6**) is based on the evaluation of different combinations of *hyperparameters* through the **Grid Search** approach [100], which allows to evaluate the *hyperparameters* set that best fit the studied test case, but introducing an additional computationally onerous procedure. Grid search is a technique for hyperparameter optimization in machine learning. It is a systematic way of searching for the optimal set of hyperparameters for a given model. The approach involves specifying a range of values for each hyperparameter, and then training the model with all possible combinations of these values. The performance of the model is then evaluated using a validation set, and the best set of hyperparameters is chosen based on the highest performance metric. Grid search is a simple and widely used method for hyperparameter optimization, but it can be computationally expensive and time-consuming for models with many hyperparameters or large ranges of values. To make the grid search evaluation more general and accurate, it will be also used the k-fold cross validation already explained in section 2.1.2, which contribute to increase the required computational effort.

4. **Model training**: Once the feature selection and the hyperparameters tuning has been done, the selected model is ready to be trained. Generally, once the training dataset is defined, it is randomly divided into three parts. The first part is called *training set*, it encloses the 70% of the available data and it is shown to the model to perform the training. The second part is the *validation set* (10% of the dataset) and it is used to compute the model performance metric during the training process in order to guide the internal parameters variation aiming to the minimization of the model prediction error. Finally, the third part is the *test dataset*. It is composed by the remaining 20% of the available dataset and is used exclusively after the training is completed, in order to evaluate its prediction performance on data which have been not previously seen by the model.

5. **CUSUM**: Once the final model has been trained, it can be used to carry out building energy monitoring by comparing actual and predicted consumption over periods subsequent to the training one, exploiting the already explained *CUSUM* technique (see section 2.1.5).

Based on the above, six methods were identified, for which a prediction performance

and computational cost analysis will be performed. The six cases are summarised in Table 2.1.

Table 2.1: Analysed cases description

Case	FS	ML	HyperParams
1	Filter method	MLR	-
2	Filter method	ANN	literature-suggested
3	Filter Method	ANN	grid-search tuning
4	Wrapper Method	MLR	-
5	Wrapper method	ANN	literature-suggested
6	Wrapper method	ANN	grid-search tuning

The cases can be classified through different characteristics:

- **Feature Selection method** (*FS*): As already said, the considered feature selection methods are two: the filter method is completely based on the correlation analysis, is computationally affordable but do not consider the real effect of any feature on the prediction performance of the model. The wrapper method (Forward Elimination) provide a feature ranking based on the real impact of the features on the prediction performance of the model (see section 2.1.2), leading to a marked increase of the computational cost.

- **Machine Learning methods** (*ML*): Multiple Linear Regression (*MLR*) and Artificial Neural Network Multi-Layer Perceptron (*ANN MLP*).

- **Hyperparameter tuning method**: In the case of *Multiple Linear Regression*, there are no hyperparameters to tune, so there is no fine-tuning procedure. In applications using the *Artificial Neural Network* as a prediction model, two methodologies were tested. The first (*Cases 2* and *5*) concerns the use of hyperparameters commonly used in the literature for similar applications (**literature-suggested**), while the second concerns an iterative parameter optimization procedure that exploits the **Grid Search**.

Once the machine learning model is trained, it can be exploited to realize the statistical monitoring method through a widely used technique called Cumulative Sum of Differences (CUSUM), which has been further explained in the preceding section.

The main issues of applying machine learning methods in the context of energy monitoring in buildings involve three key aspects. Firstly, the *model prediction performance*

is highly important. The goal of the method is to highlight changes in the systems energy behaviour by comparing the values predicted by the model with the actual measured values, using the statistical technique known as *CUSUM*. As mentioned earlier, *CUSUM* is based on the cumulative sum of the obtained errors. The fluctuation of the cumulative sum during the training period is used to estimate the expected oscillation in the subsequent period, representing the normal operation of the systems. Consequently, the smaller the prediction error, the lower the *CUSUM* oscillation during the training period. This leads to increased sensitivity of the method, enabling the detection of mild and progressive changes in the building energy consumption behaviour, resulting in the identification of more changes, significantly reducing the time between the issue and its detection, and consequently, facilitating the prompt implementation of corrective measures, leading to improved energy efficiency of the building.

The second issue is directly related to the previous one and concerns the computational cost of model training. Firstly, the choice of different machine learning methodologies results in significantly different computational costs. In this thesis work, *Multiple Linear Regression* is analysed, which is characterized by simplicity of interpretation and usage, as well as a very low computational cost, at the expense of lower prediction performance and a lower capability to identify complex relationships between features and the target. The second method used (*Artificial Neural Network*) represents the polar opposite, offering significantly higher prediction performance at the cost of a considerable increase in computational cost. In addition to the basic machine learning model, the various methodologies described vary based on *feature selection* and *hyperparameters fine-tuning* procedures, proposing low computational cost procedures for both (*filter method* for feature selection and *literature-suggested hyperparameters* for hyperparameters tuning) as well as systematic procedures. As already illustrated, systematic procedures involve repeating the training process multiple times until ideal conditions are achieved for training the final model. These methods allow to reach the secondary goal of notably reduce the required user experience to perform the related tasks. Therefore, the choice of method should be based on the specific activity and the available computational effort. The following chapters will provide a detailed comparison of the various applied methods, considering prediction performance and computational cost, which effectively serves as a guide for selecting the most appropriate method.

The third issue concerns a different aspect: *data availability*. Firstly, it is necessary to have as detailed data as possible regarding the energy consumption of the building. Such data is easily obtainable for electrical energy. Energy providers make available hourly data on their websites regarding the electrical power absorbed by the entire building. However, to implement a real-time system, the installation of dedicated measurement systems within

the building may be required, leading to additional expenses. Currently, non-invasive and easily installable electrical energy measurement systems are available at affordable prices for companies. More problematic could be the monitoring of natural gas consumption, as data is not made available by energy providers. The installation of dedicated measurement systems is more challenging, as it involves installing turbine-based meters directly on the gas supply pipelines, which would require a complete shutdown of the thermal power plant, causing disruptions that can be problematic for 24/7 operational facilities. The data necessary for developing a good machine learning model depend closely on the specific energy drivers of the analysed structure. For example, a structure whose consumption is mainly driven by *HVAC* systems will have external climatic conditions as its main energy drivers, while an industrial production facility will be more influenced by production volume, the presence of operators inside the building, and the opening and closing hours of the plant. Consequently, for each entity, it will be necessary to arrange the procurement of relevant data, which is not always feasible in all contexts. For instance, a company may not be able to afford installing a dedicated weather station for measuring external climatic conditions and may have to resort to alternative methods, such as requesting data from dedicated meteorological laboratories.

Chapter 3

Test case overview

Figure 3.1: Healthcare Facility view

The historic location of the healthcare facility *Villa Donatello* was until 2018 the elegant 19^{th}-century building located near the historic city center of Florence, Italy. The building was built outside the external circle of historic walls was demolished and the avenues were created, during the Florence *Capital period* (1865-1871) and can therefore be dated around 1870. Previously a Florentine residence of the *Rospigliosi* princes until the early 20^{th} century, it then became property of the *Nevers nuns* and has been used as a care home with the name Villa Donatello since 1946. In 2014, to meet the increased demand, some activities were moved to a second building adjacent to the main location (*Villa Vittoria*), which dates back to the same time period and recently renovated. The

Figure 3.2: Villa Ragionieri during the restructuration

further growth of the customers made the old location inadequate with respect to the offer of comfort in line with the standards of the best private healthcare services, and in 2018 led to the transfer of most activities and the main location to a larger and more modern structure, built on the ashes of another historic villa (*Villa Ragionieri*) and located at the border between the municipalities of *Florence* and *Sesto Fiorentino*. The relocation facilitated the expansion of services provided by *Villa Donatello*, transforming it from a traditional private healthcare facility to a small-scale private hospital. A significant outpatient facility, equipped with the main specialized and diagnostic services, remains active at the historic location in the old building (*Villa Vittoria*), representing a convenient proximity point for those living in the south and center of Florence.

With regards to the "new headquarter" referred to as *Villa Ragionieri*, it was obtained through the restructuring of the historical building (see Figure 3.2). The project, completed in 2008, was initially intended for the conversion of the building into a large oncology center, housing operating rooms and diagnostics, hospitalization, and advanced oncological therapies, dedicating the historical part to administrative functions and the extension to healthcare services.

The building (represented in Figure 3.3), inaugurated in *October 2018*, extends itself into the surrounding lands, reaching a total useful floor area of $11'929.97\ m^2$, a total volume of $52'580.72\ m^3$ and a building exposed surface of $23'577.57\ m^2$. It is now an innovative polyfunctional clinic, strongly active on the Florentine territory. It has 8 operating rooms, 40 beds, 32 outpatient clinics, state-of-the-art equipment such as an angiograph, a robot for urologic surgery, 3 lasers for ophthalmic surgery, a latest-generation Magnetic Resonance Imaging, a digital mammograph, and an advanced ultrasound for prenatal diagnosis. Other data, referring to *December 2018*, significant of the facility's potential are:

- $55'515$ outpatient visits
- $4'680$ surgical procedures

Figure 3.3: Healthcare Facility view

- $4'506$ Magnetic Resonance Imagings (MRI) and Computerized Axial Termographies (CAT)
- $1'626$ endoscopic procedures.

Upon examination of the available documentation regarding the architectural design and internal systems of the healthcare facility, it has been determined that the facility can be divided into two distinct areas (Figure 3.4). The first area, referred to as the *historic building*, does not contain any primary healthcare functions. This area serves as administrative and support space for the healthcare facility.

The second area, referred to as the *enlargement*, is specifically designed and equipped for core healthcare activities. It comprises four floors and a rooftop terrace and is utilized for patient care, treatment, and examination. The basement of the enlargement area of the facility is primarily utilized for the housing of technical systems and ancillaries, such as transformer cabinet, power center, water stations, and the main part of the air handling units. The ground floor of the facility is designed to enclose a central area near the entrance which includes the reception area, several offices, a bar, and a small commercial space. The two hallways delimit a central area, which includes the surgery department and the intensive care unit. These areas are characterized by specific constraints in terms of ventilation and air handling and include many medical devices which require electric energy for their operation. The first floor presents two hallways similar to the ones described

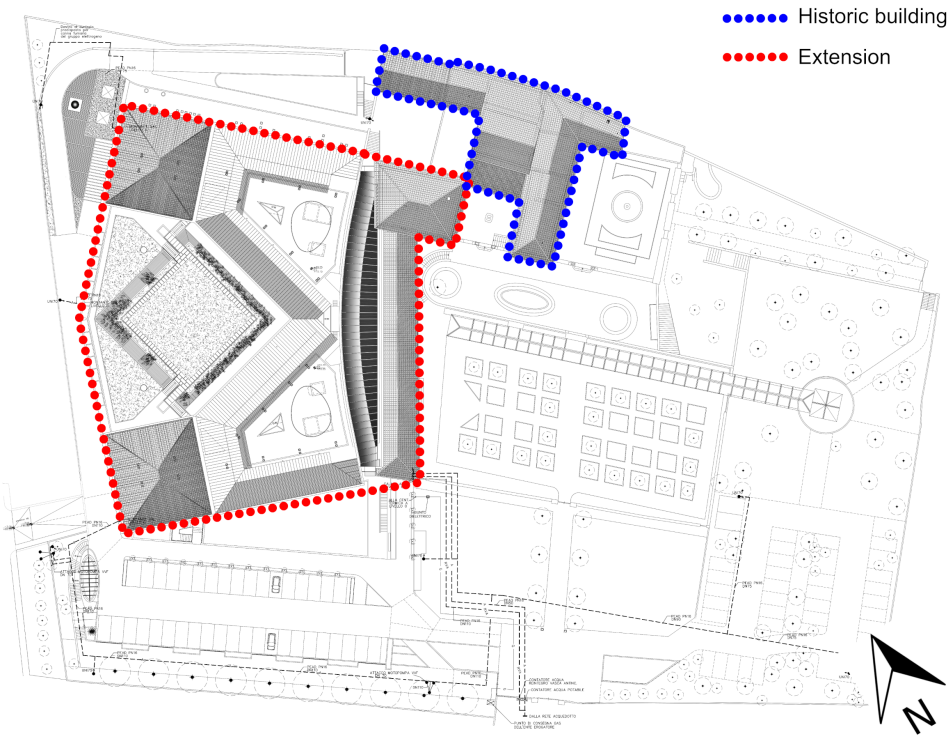

Figure 3.4: Test case planimetry

for the ground floor, but it does not have the central area. This floor hosts some offices, day hospital department and outpatient activities. The second floor is similar to the first one from a structural point of view, and it hosts the hospitalization departments. The activities performed in first and second floor need careful air handling management in order to maintain both comfort and security standards, resulting in electric energy demand for ventilation and hot water, chilled water and steam for air conditioning. Moreover, surgery and the hospitalization departments rooms need the access to compressed air and medical gasses, which lead to an additional electric energy demand for gasses and air circulation.

3.1 Plants and systems

The building thermal energy requirements (in form of hot and chilled water) are covered on-site by dedicated systems, while the electrical energy requirements are fulfilled by the electricity grid. The electrical demand, in case of national grid failure, is ensured by the presence of a 2 MW_{el} generator obtained through an internal combustion engine powered by natural gas, which is automatically activated in case of need. To ensure the continuity of essential services, various *UPS* groups are present downstream of the power center, serving three main "sections" called respectively: *Vital continuity*, which regards all the areas and equipments which mus remain active continuously, such as the surgery rooms

Figure 3.5: Refrigerators Climaveneta W3000 TECS-HF 4AS

Characteristic	Value
Cooling Power	1′300 kW_{cool}
Operative water pressure	1 MPa
Operative coolant pressure	1.9 MPa
Coolant type	R134a
Nominal Electric power	320 kW

Table 3.1: Refrigerators Climaveneta W3000 TECS-HF 4AS details

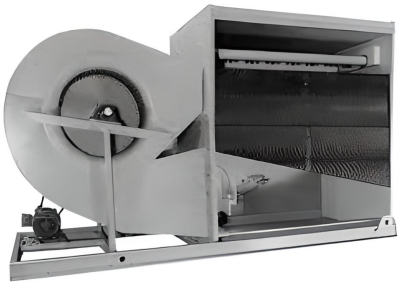

Figure 3.6: Cooling Tower Baltimore Aircoil VTL 152M

Characteristic	Value
Fan Power	18.5 kW
Water volume flow rate	36.3 L/s
Inlet water temperature	35 $°C$
Outlet water temperature	30 $°C$
Wet bulb air temperature	24.5 $°C$

Table 3.2: Cooling Tower Baltimore Aircoil VTL 152M nominal parameters

or the Magnetic Resonance, *Security*, which regards all the emergency lighting and the emergency air extractors, and *IT security*, which serve all the data centers, server and communication systems installed within the building.

The central water supply originates from a water meter placed outside the building and provides water requirements to the thermal power plant and refrigeration system.

The water cooling is carried out through three water-cooled refrigeration units (see Figure 3.5 and Table 3.1). Each unit is characterized by a nominal electric power of 320 kW and a nominal *COP* of 3.92 and is equipped with four centrifugal compressors with digital speed control of the impellers and magnetic levitation bearings. As shown by Figure 3.9, the the refrigerators are connected to the circuit of chilled water, that flows in each refrigeration units' evaporators. After the cooling process, the primary water from all the refrigeration units is led to a chilled water collector by the primary pumps (three for each unit). Then secondary pumps draw the water from the aforementioned collector to supply all the different areas of the building (primarily the *AHUs*). Another water circuit connected to the condenser of the refrigerators led the cooling water to the evaporative towers located on the roof.

Indeed, to achieve the refrigeration of the coolant, six evaporative towers (located on the extension building roof) serve the refrigeration units. The total cooling capacity is

Figure 3.7: Hot Water generator ICI Caldaie GREENOx-e70

Characteristic	Value
Thermal power	700 kW
Hot water temperature	80 $°C$
Max operations pressure	5 bar
Stack gases pressure losses	5.6 mbar
Full load efficiency	95.37%

Table 3.3: Hot Water generator ICI Caldaie GREENOx-e70 nominal parameters

Figure 3.8: Steam generator ICI Caldaie PX300

Characteristic	Value
Nominal thermal power	230 kW
Maximum thermal power	233 kW
Steam production	300 kg/h
Operative pressure	4.5 bar

Table 3.4: Steam generator ICI Caldaie PX300 nominal parameters

4′560 kW (6 × 760 kW). Further detail are reported in Table 3.2.

Four gas-powered hot water generators are installed within the thermal power plant, for a nominal thermal power of 2′800 kW_{th} (4 × 700 kW_{th}). As previously stated, the hot water generators fulfil the entire hot water demand of the facility. This demand is mainly composed of the thermal energy requirement for heating, and therefore the demand derived from the pre and post-heating coils of the *AHUs*. Since in an *AHU* the post-heating coil operates after the cooling coil (in order to regulate the temperature of the air entering the rooms), a portion of the demand for hot water for air heating purposes will take place also during the summer months. A significantly lower portion is composed of the generation for domestic hot water.

The thermal power plant of the healthcare facility encloses also three natural gas-powered steam generators, for a total nominal thermal power of 690 kW_{th} (3 × 230 kW_{th}), ensuring all the steam needed for air treatment purposes. The generation of steam for the sterilization of medical equipment and beds is instead delegated to electric steam generators installed in the rooms used for that activity. Further detail about the three steam generators are resumed in Table 3.4.

Table 3.5: Air handling units installed within the facility

AHU	**Area**	**Mass flow rate** (m^3/h)
CDZ1	Radiotherapy	10′800
CDZ2	Laboratories	5′950
CDZ3	Locker rooms	2′250
CDZ4	Locker rooms	2′250
CDZ6	Diagnostic dep.	18′200
CDZ11	Hall	12′500
CDZ12	Day Hospital dep.	3′180
CDZ13	Inpatient dep.	4′200
CDZ14	Inpatient dep.	4′200
CDZ15	Inpatient dep.	4′200
CDZ16	Inpatient dep.	4′200
CDZ17	Conference room	3′180
CDZ18	MRI	2′150
CDZ19	Bar	1′300
CDZ20	Commercial area	1′550
CDZ21	Laboratories	3′180
CDZ22	Offices	4′825
CDZ23	Antiblastic lab.	3′150
CDZ24	Fisiotherapy dep.	4′200
CDZ C1	Kitchen	1′100
CDZ C2	Cafeteria	1′100
CTA1	Surgery room 1	2′800
CTA2	Surgery room 2	2′800
CTA3	Surgery room 3	2′800
CTA4	Surgery room 4	2′800
CTA5	Surgery room 5	2′800
CTA6	Surgery room 6	4′200
CTA7	Surgery room 7	3′400
CTA8	Surgery room 8	11′750
CTA9	Recovery room	18′000
CTA10	Decontamination area	11′250
CTA11	Intensive care unit	15′500

The *HVAC* system is composed of Thirty-two air handling units. Two air handlers are dedicated to the *historic building* (one of these two serve the conference room and is normally switched off), while the other Thirty are used to maintain the air quality standards in the *enlargement area*. Table 3.5 resume all the installed *AHU*, specifying both the nominal air mass flow rate and the building area served by the specific unit.

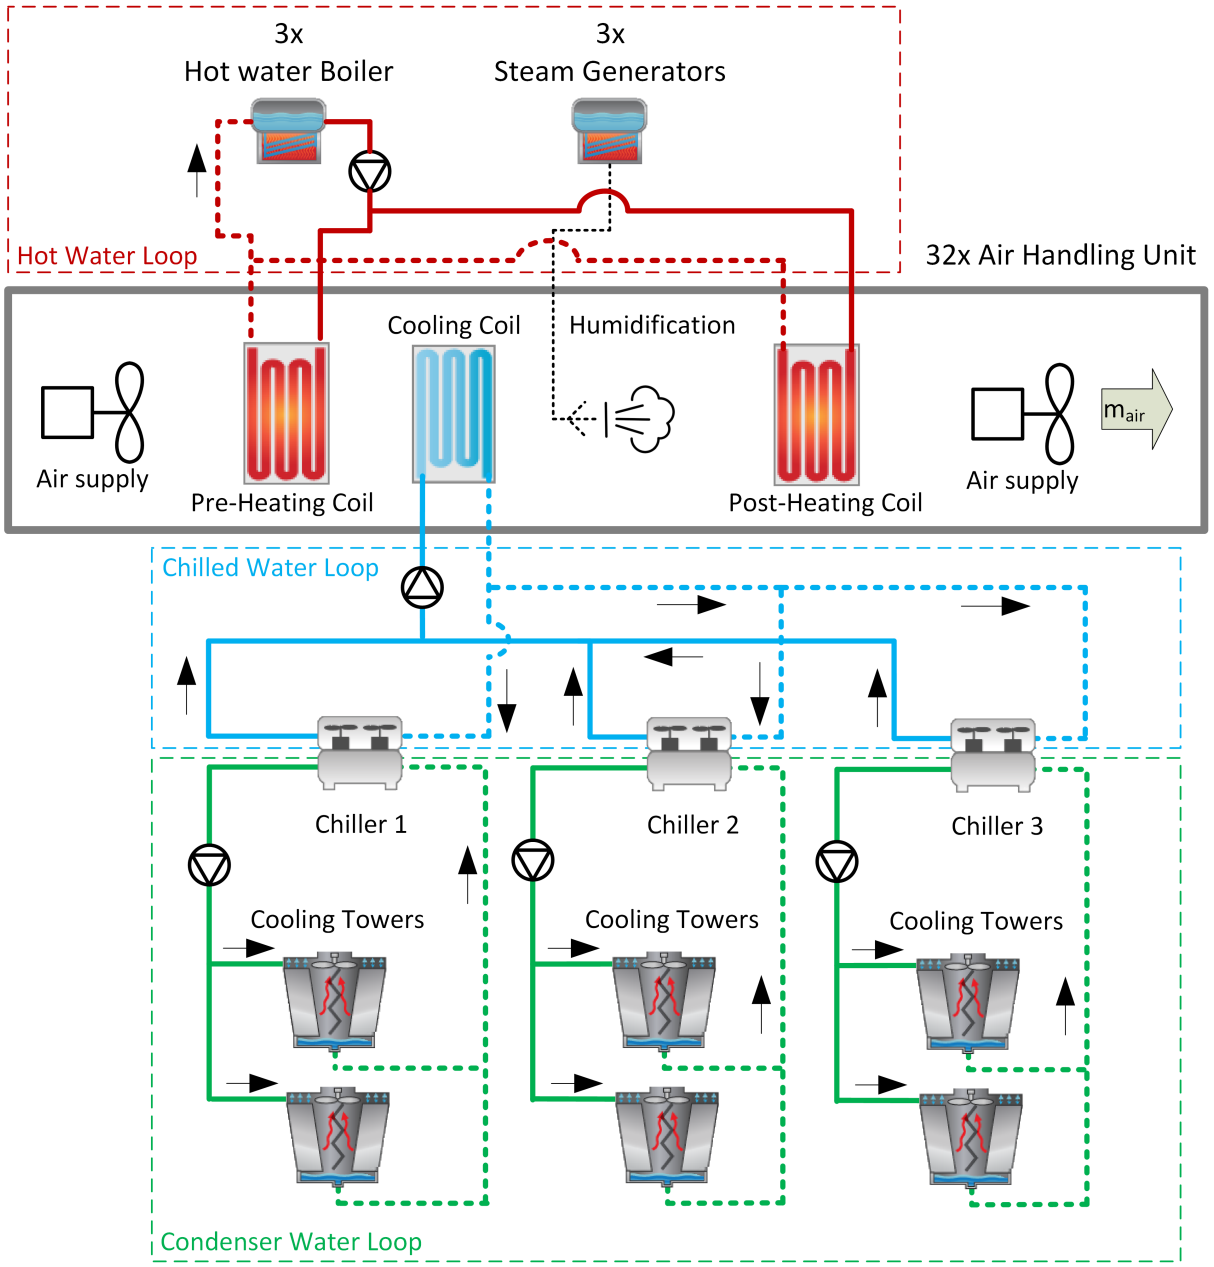

Figure 3.9: Healthcare Facility HVAC systems scheme

The electrical supply of the facility is at medium voltage (15′000 V). From the medium voltage cell, the electrical line reaches the transformer station, where three power transformers provide for the transition to low voltage (380 V). Downstream of these, the power center of the healthcare facility is supplied, which is divided into two sections. The normal section is exclusively supplied by the national network, while the preferential section has the possibility, in case of service interruption, to be supplied by a generator set, in order to guarantee continuity of service. The two sections are composed of a total of 75 lines that, starting from the power center, branch out throughout the facility, supplying the different systems and utilities.

In conclusion, all the aforementioned systems are fully controlled and managed automatically by a centralized Building Energy Management System (BEMS). As seen, the

generation of cold and heat are separate systems, located in completely different places within the building. However, the energy generated in both cases is conducted to the air handling system, which is fractionated into various sub-groups positioned in different technical rooms. Both the heating and cooling systems therefore find a sort of junction in the air handling units that, through automated controls, heavily influence their energy demand. To better visualize the configuration of the entire system, a simplified diagram has been created and is shown in Figure 3.9.

3.2 Climate

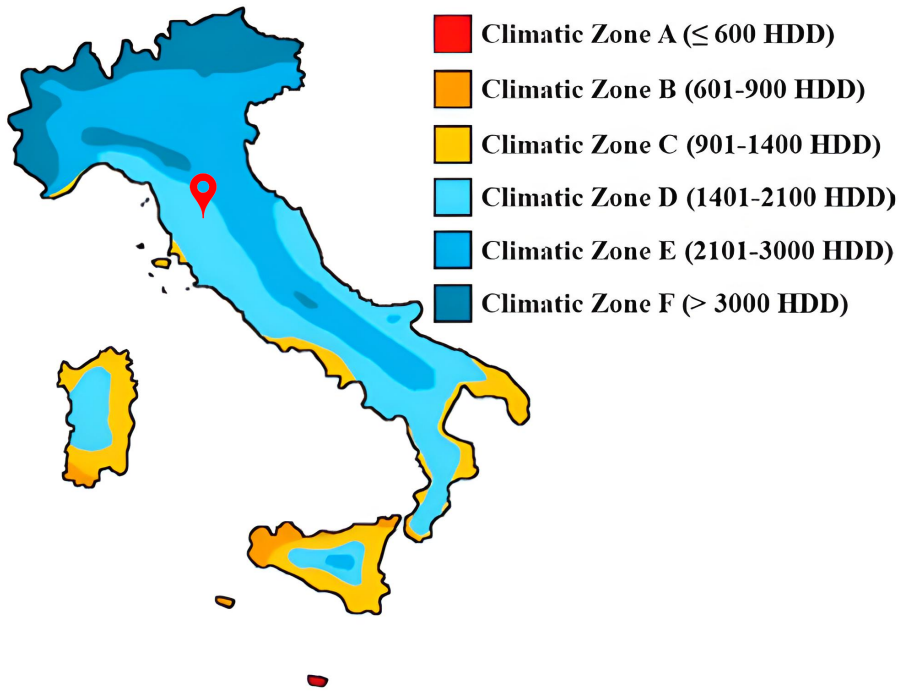

Figure 3.10: Italian Climatic Zones classification

Healthcare and Hospitals are generally characterized by a strong correlation between energy consumption and external climate conditions, due to their extensive use of *HVAC* systems to ensure comfort and security indoor air quality standards. Consequently, to describe the building and its activities is pivotal to first contextualize it in the corresponding climate conditions. As depicted in Figure 3.10, Italy is divided into six climatic zones, which are basically used to define the period of the year and the number of hour per day of permitted activation of heating systems for private and public buildings. The climatic zones are defined based on a simple index called heating degree day (HDD) generally used as an indicator of the expected heating energy requirements for a building in different climate conditions. This parameter is defined as the cumulative sum of the differences between the daily averaged outdoor air temperature ($T_{e,j}$) and a reference temperature called T_0 (which represent the expected indoor air temperature set-point), summed over a defined period (usually a year) and only when $T_{e,j} < T_0$:

$$HDD = \sum_{j=1}^{n} \left(T_0 - T_{e,j} \right) \Big|_{T_{e,j} < T_0} \tag{3.1}$$

Where n represents the number of days in the conventional heating period and j is the index which represents the day of the considered period. $T_{e,j}$ is the daily mean temperature value, while T_0 (expressed in $°C$) is a reference temperature value selected based on the required internal temperature. The Italian regulation (*D.P.R. 412/1993*)

provides a reference temperature of 20 $°C$. The *HDD* assumes a non-zero value only in the days characterized by a mean temperature (T_e) lower than the conventional one (T_0). If on one hand the *HDD* are defined to quantify the need for heating during a specific period, on the other hand a second parameter can be defined in order to quantify the need for cooling. This parameter is called *Cooling Degree Days* (*CDD*) and is defined as:

$$CDD = \sum_{j=1}^{n} \Big(T_{e,j} - T_0 \Big) \Big|_{T_{e,j} > T_0} \tag{3.2}$$

UNI 10339 and **UNI 10349** provides a reference temperature for *CDD* which vary based on the locations, and for the test case under analysis is 25 $°C$. Nevertheless, it is important to remark that this parameter is not used in the Italian regulation to define any constraint or best practice for choosing the summer indoor air temperature set-points.

The healthcare facility under analysis is located in *Sesto Fiorentino*, near *Florence*, Italy (latitude: 43.82426335607834, longitude: 11.223063227171048). The location is classified by the Italian regulation as a **climatic zone D**, assigning a conventional Heating Degree Day value equal to 1′772 *HDD*. For the municipality of *Sesto Fiorentino*, the natural gas powered heating systems can be activated from *November 8* until *April 7*. Moreover, the heating systems can operate 11 hours per day, with prescribed indoor air temperature of 17 ± 2 $°C$ for the industrial activities and 19 ± 2 $°C$ for all the other building typologies. For other heating system typologies (not natural gas-powered), the admitted period extends itself from *November 1* until *April 15*, with a limit of 12 hours per day and prescribed indoor air temperature set-point of 18 ± 2 $°C$ for industrial buildings and 20 ± 2 $°C$ for all the others.

The presented analysis has been carried out through real climate data provided by *"Consorzio LaMMA"* related to the year *2020* that presents 1′946.7 *HDD*, resulting in

Table 3.6: Climate feature description

Feature	**Description**	**Unit**	**min**	**MAX**	μ	σ
T_{db}	Dry bulb temp.	K	270.85	313.25	289.68	7.98
RH	Relative Humidity	%	9.00	100.00	63.06	22.78
SR_{gl}	Global Solar Rad.	W/m^2	0.00	1127.20	170.07	256.89
SR_{diff}	Diffusive Solar Rad.	W/m^2	0.00	454.00	50.77	75.06
I	Illuminance	Lx	0.00	123894.00	16102.29	25119.33
p_{amb}	Ambient air pressure	bar	1.0100	0.9780	1.0380	0.0076
WS	Wind speed	m/s	0.10	11.30	1.75	1.36
WD	Wind direction	deg	2.00	360.00	202.82	94.35
$Prcp$	Precipitation	mm	0.000	18.200	0.024	0.206

a colder year in comparison to the conventional one (1′772 *HDD*). Moreover, regarding the Cooling Degree Days (*CDD*), the healthcare facility location in the *2020* present $CDD = 129.3$. The exploited climatic data have been obtained through real measures carried out exploiting a climatic station placed near the healthcare building. All the available parameters are shown in Table 3.6 and Figure 3.11. In particular, the figure reports the outdoor air dry bulb temperature in $°C$ (although the data will be exploited in K) in order to help the temperature behaviour understanding. The statistical characterization of the available climate data is succinctly presented in Table 3.6, providing a comprehensive overview of the key metrics derived from the dataset.

As explained in section 2.1.3, the available data can be further elaborated through the *Feature Engineering* process, in order to obtain further data which could improve the machine learning model training process. In particular, through the combination of outdoor air dry bulb temperature (T_{db}), relative humidity (RH) and air pressure (p_{amb}), by applying the psychrometric relations, it is possible to obtain: wet bulb air temperature (T_{wb}), dew point temperature (T_{dp}), vapour pressure (p_{vap}), humidity ratio (HR), dry and moist air enthalpy (h_{dry}, h_{wet}), dry and moist air density (ρ_{dry}, ρ_{wet}). The psychrometric relations were used exploiting an open source third party python package called *Psychrolib* [101], which is based on the relation prescribed by ASHRAE [102].

Parameters related to the heating degree days (namely, *HDD* following Eq. 2.11 and 2.12, and *HDh* as specified in section 2.1.3) have been obtained. The "Wind" features, which have been originally represented through the direct measures as two separate vectors

Table 3.7: Engineered climate feature description

Feature	**Description**	**Unit**	**min**	**MAX**	μ	σ
T_{wb}	Wet bulb temp.	K	268.60	296.08	284.61	5.52
T_{dp}	Dew point temp.	K	257.41	293.88	281.01	6.06
p_{vap}	Vapour pressure	bar	0.0015	0.0245	0.0114	0.0043
HR	Humidity Ratio	-	0.0009	0.0154	0.0071	0.0027
h_{dry}	Dry air ent.	J/kg	100.60	38328.60	16157.58	8032.14
h_{wet}	Moist air ent.	J/kg	1890.53	67637.83	34041.05	13537.50
ρ_{dry}	Dry air dens.	kg/m^3	1.1279	1.3146	1.2202	0.0352
ρ_{wet}	Moist air dens.	kg/m^3	1.1214	1.3127	1.2151	0.0364
w_x	E-W wind speed	m/s	-6.8612	11.4000	-0.0584	1.7045
w_y	N-S wind speed	m/s	-8.0369	8.1351	0.0017	1.4601
HDD	Degree Day	HDD	-4.94	18.10	6.59	5.48
HDh	Degree Hour	HDh	-13.10	24.50	6.59	6.43
$ClimateZ$	Climate Z	-	254.07	320.91	288.85	12.92

"speed" (WS) (expressed in m/s) and "direction" (WD) (expressed in degrees), have been converted in two speed components: the *East-West* velocity component (w_x) and the *North-South* velocity component (w_y). The final feature engineering procedure applied to the climate data allow to obtain the already mentioned *ClimateZ*, which is expressed in Eq. 2.13. All the climate feature that will be used in the following chapters are resumed and statistically analysed in Tables 3.6 and 3.7.

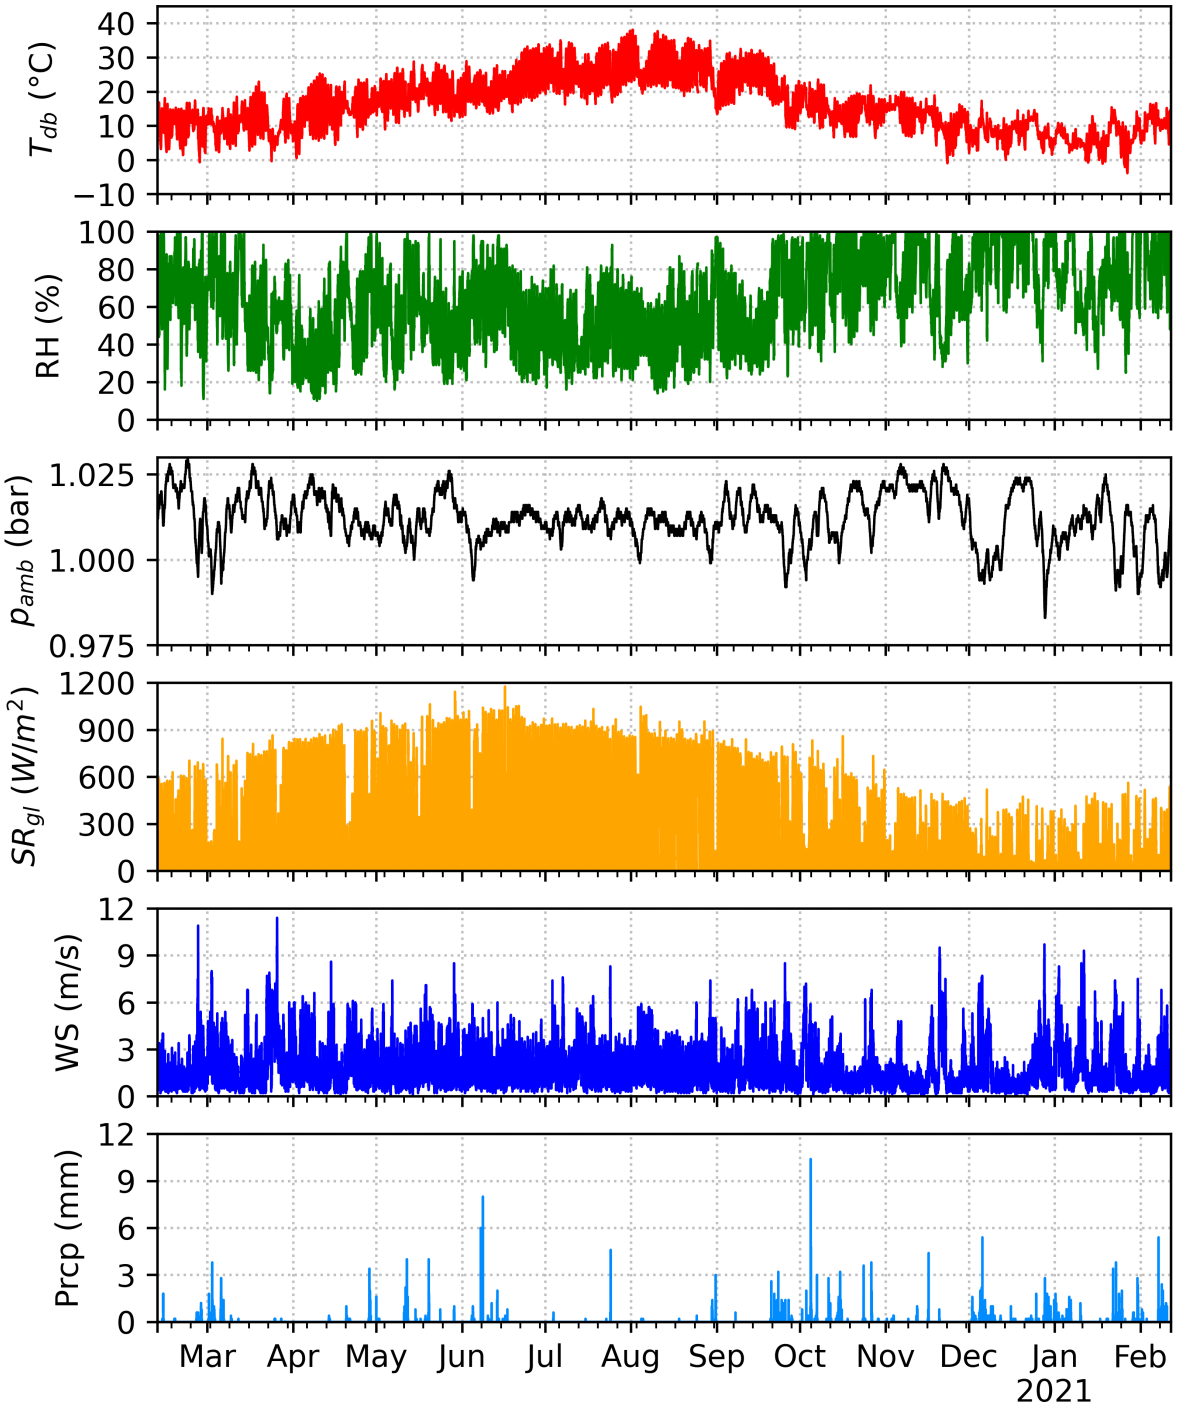

Figure 3.11: Main climate features behaviour at the facility location in 2020

3.3 Energy consumptions

First of all, it is important to notice that all the data presented in this section have been scaled by their maxumim value found in the analysed period, according to what was said talking about the *feature and target scaling* in section 2.1.1. The monthly thermal energy demand have been divided by its maximum monthly thermal energy consumption, as well as the monthly electricity demand. As regards the detailed data about the building electric power demand (that has a 15-minutes time step), all the data have been scaled by the maximum value of instantaneous absorbed power related to the whole building data. The detailed about *HVAC* subsystems (refrigerators, air handling units and pumps), have been also divided by the maximum whole building power demand, in order to better visualize the proportions between the overall and specific consumptions. Indeed, all the data will vary between 0 and one, but maintaining values proportional to the percentage of the instantaneous whole building energy demand. Consequently, the percentage of a specific subsystem can be obtained directly by looking at the specific values shown in the following figures.

Upon analyses of the energy provider bills, it was obtained the monthly energy consumption related to the year *2020* of the test case under analysis.

The total consumption of the healthcare facility is divided into different portions depending on the energy source. In particular, in *2020*, the electricity requirements are almost 59% of the overall energy demand, while the thermal energy requirements account for the remaining 41%. 2020 was a unique year due to the pandemic. Nonetheless, the activities carried out within the facility did not undergo any changes compared to previous

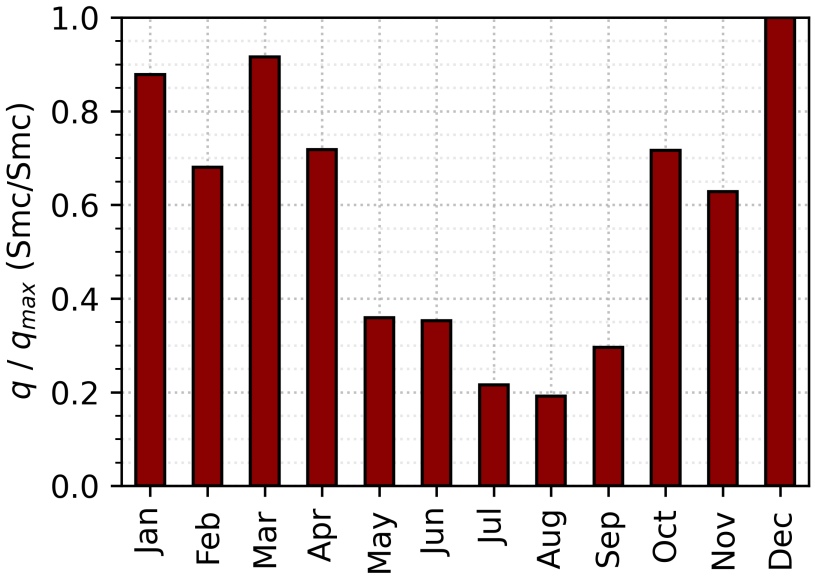

Figure 3.12: Monthly thermal energy consumption scaled by maximum monthly consumptions

years, both in terms of healthcare services and the use of energy facilities. In Figure 3.12, the dimensionless monthly energy demand can be observed. It can be seen that there is a strong difference in energy demand between winter months and summer months. In particular, summer months define a baseline for the building thermal energy requirements, mainly consisting of the need for hot water generation and a small part of air heating downstream from the *AHU* cooling coils (post-heating coils). The rest is tightly linked with the outdoor air temperature, resulting in high thermal energy demand during winter and minimal needs during summer. In intermediate periods, a gradual variation of the energy consumption can be noticed, which tends to follow the monthly average temperature trend. The monthly billing of electrical energy by the energy provider is calculated by quantifying the energy consumed across three distinct time intervals, which are markedly based on the national energy demand patterns throughout a weekly cycle. These time slots are represented in figure 3.13 and are defined as follows:

- **F1** represents the periods that go from Monday until Friday, from 08:00 until 19:00 (excluding national festivity days)
- **F2**, represents the period from Monday until Friday, at 07:00 - 08:00 and 19:00 - 23:00. time slots. Moreover, it includes Saturday from 07:00 until 23:00
- **F3** represents the periods between Monday and Saturday, from 23:00 until 07:00. Every Sunday and all national festivities are considered as part of the *F3* time slot

Figure 3.14 shows the obtained data about the building electricity demand. Electricity consumption has been scaled by dividing it by the maximum monthly consumption found in the year under analysis (E_{max}). It presents a marked increase in the summer. July results to be the most energy-intensive month, presenting a +50% of February energy consumption due to the massive air handling requirements characteristic of the summer period. This preliminary energy demand analysis shows an energy demand dependency on the seasonal changes of the outdoor climate conditions. Nevertheless, aiming at defining the main building energy drivers, the analysed data are not sufficiently detailed to obtain reliable information. Consequently, more detailed electricity demand data were obtained

	0 - 7	7 - 8	8 - 19	19 - 23	23 - 24
Workdays	**F3**	**F2**	**F1**	**F2**	**F3**
Saturday	**F3**	**F2**	**F2**	**F2**	**F3**
Sunday	**F3**	**F3**	**F3**	**F3**	**F3**

Figure 3.13: Time slot definition table

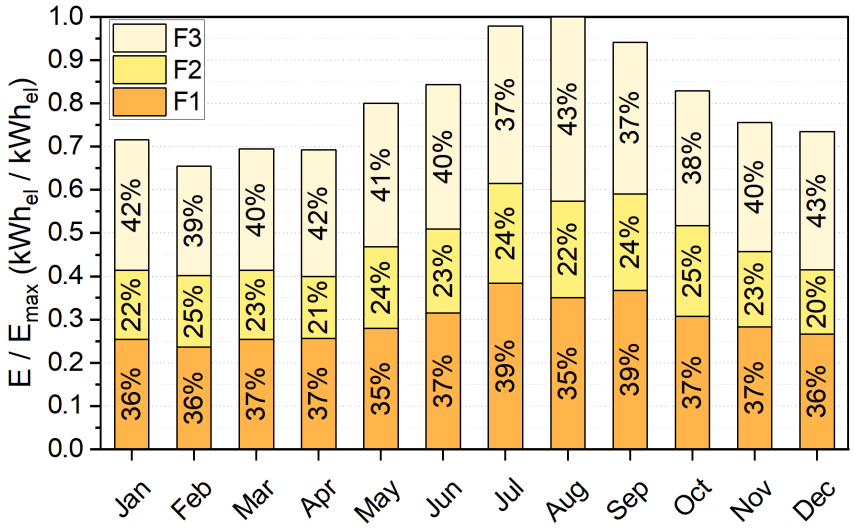

Figure 3.14: Monthly electrical energy consumption divided by timeslot

from the energy provider, who made available the data set from February 2020 to October 2022 (15-minutes time step).

Energy Performance Indexes

Based on the data collected so far, it is possible to calculate some energy performance indices. To compare with other similar structures, a work by "IRES Piemonte" was found that addresses the *2016* energy consumption analysis of all hospital structures in the Italian region of "Piemonte". In this study, three main indices were used:

- **Primary energy consumption index** [toe/m^2]: Basically calculates the overall yearly energy consumption of a building converting it in *toe*, dividing it by the useful floor area

- **Thermal energy consumption index** [$kWh_{th}/(m^2 \times HDD)$]: The thermal energy consumption (expressed in kWh) must be divided by the useful floor area of the hospital in order to compare structures with different sizes. However, the specific degree days of the location play an important role in heating. To make the quantity comparable, the obtained index is further divided by the degree days that were measured in the specific location in the analysed year.

- **Electrical energy consumption index** [kWh_{el}/m^2]: The "IRES Piemonte" study calculates this index by dividing the electricity consumption only by the value of the useful floor area. An index that is even more comparable with structures located in different places could have followed the same logic as the index seen previously, further dividing the index by the cooling degree days found for the specific year in the specific location. In any case, the summer temperatures of the healthcare

facility under analysis are not that different from the ones found in Piemonte, making comparison possible.

Regarding the **Primary energy consumption index** the values found for Hospitals in *Piemonte*, vary from 0.02 to 0.08 toe/m^2. Regarding the healthcare facility under analysis, its value, calculated for the year 2021, is 0.079066 toe/m^2, thus being just below the highest value found in the *IRES Piemonte* study. The **Thermal energy consumption index** of *IRES Piemonte* study ranges between 0.05 to 0.2 $kWh_{th}/(m^2 \times HDD)$, while the present test case obtain a value of 0.264 $kWh_{th}/(m^2 \times HDD)$, a value 25% higher than the maximum found in the literature study. Finally, the **Electrical energy consumption index** of Hospitals in *Piemonte* vary between 45 and 200 kWh_{el}/m^2, while the analysed test case Index is 404.64 kWh_{el}/m^2, doubling that found in literature.

The analysis of the obtained energy performance indexes highlight that the healthcare facility under analysis appears to be significantly more energy-intensive than the structures analysed in the literature. However, it must be considered that the structure itself was derived from the extension of a historic building. The design was certainly influenced by the landscaping constraints imposed by the municipality, resulting in the creation of a hospital that have markedly higher "energy-density" than the standards, where every available space has been extensively exploited to perform the healthcare core activities or to enclose the energy systems. Nevertheless, it is also true that significant improvement margins have been identified regarding the management strategies of the energy systems. During the data collection and documentation analysis activities, which continued throughout the entire work presented in this thesis, many improvements have been made, leading to non-negigible energy savings especially regarding the electricity demand (the latest analyses indicates a decrease of 14% of the overall yearly energy consumption in comparison to the predicted ones in 2022).

Detailed building electricity demand analysis

The electrical supply of the facility is at medium voltage (15′000 V). From the medium voltage cell, the electrical line reaches the transformer station, where three power transformers provide for the transition to low voltage (380 V). Downstream of these, the power center of the healthcare facility is supplied, which is divided into two sections. The normal section is exclusively supplied by the national network, while the preferential section has the possibility, in case of service interruption, to be supplied by a generator set, in order to guarantee continuity of service. The two sections are composed of a total of 75 lines that, starting from the power center, branch out throughout the facility, supplying different utilities.

Figure 3.15: Wireless measuring system installed within the building power center

On *February 12, 2020*, each line was equipped with a self-powered wireless sensor (see Figure 3.15) that allows to carry out real time measurements of the electric power absorbed by each line. The data obtained are transferred in real-time to an online platform where the data is stored, constituting in fact a history of the consumption of the facility divided by the different utilities. Due to the complexity of the electrical system, it has been considered appropriate to aggregate the data obtained in different categories, in order analyse the electricity demand distribution among the different building systems.

As a first step, a manual data cleaning was carried out, which allowed for the detection and elimination of data related to failures of the electricity service. In addition, it was necessary to remove maintenance-related periods when the facility was disconnected from the national electricity grid. This procedure is fundamental to obtain a dataset which well represent the standard energy demand of the healthcare building.

Figure 3.16 reports the behaviour of the whole building electric power demand from *February* 12^{th} *2020* until *February* 11^{th} *2021*. The figure highlight a strong variation of the power load during the year. In particular, the electricity demand presents a marked increase during the summer season, coherently with the information provided by the energy

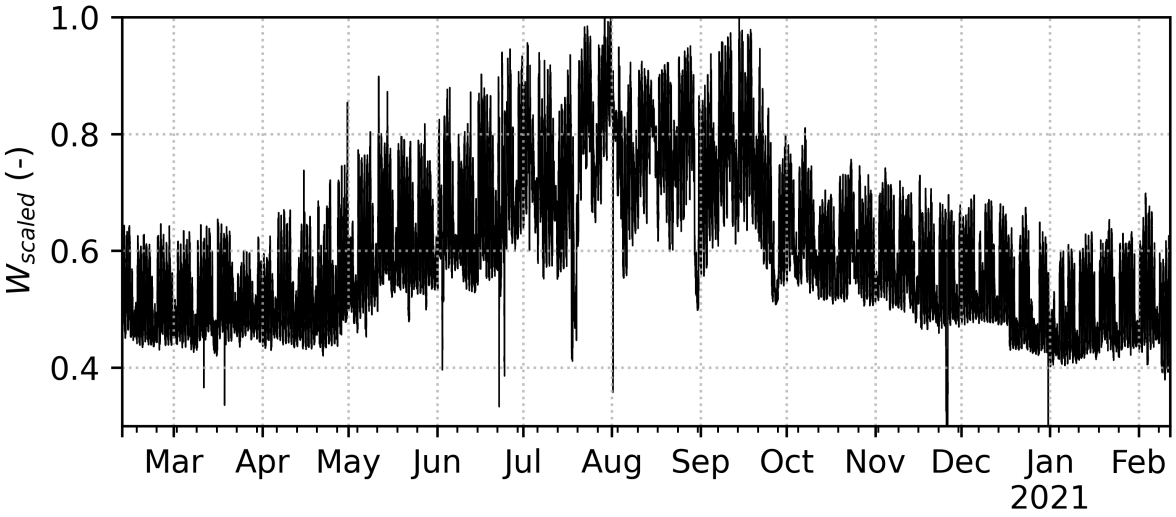

Figure 3.16: Whole building electric power demand

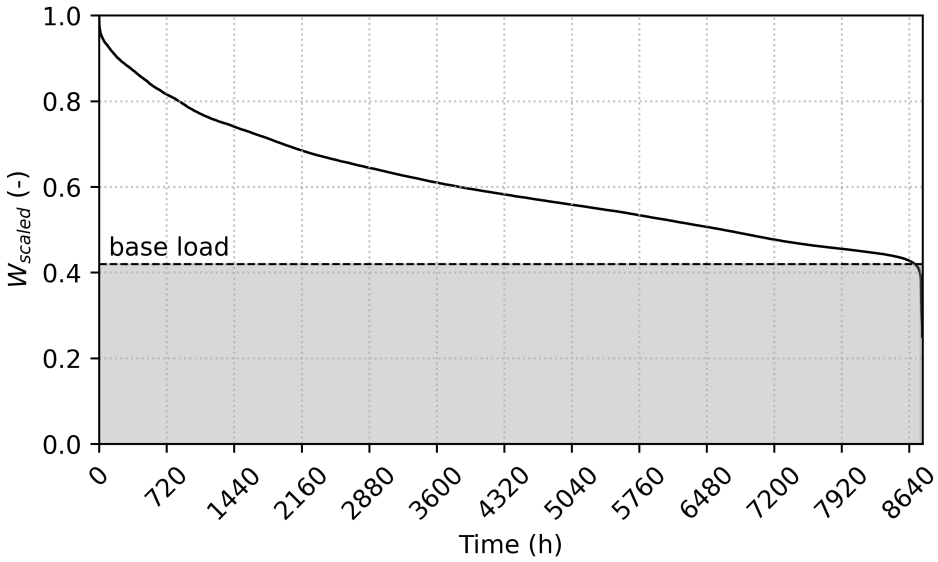

Figure 3.17: Whole building electric power demand cumulated curve

provider bills (Figure 3.14). By comparing the T_{db} behaviour in Figure 3.11 and the whole building energy demand behaviour (Figure 3.16), several similarities can be noticed, a sign that the climate conditions have strong influences on the building energy demand.

Moreover, the cumulated power curve of the whole building electrical power demand have been drawn in Figure 3.17. The figure allow to better visualize the *base load*, which is the minimum consumption that the structure maintains for all hours of the year. Regarding the whole building electricity demand, the base load is above the 40% of the maximum absorbed power during the analysed year, confirming the high level of service continuity characteristic of the test case under analysis. The figure show also that the power demand remain above the 80% of the maximum absorbed power for almost 720 hours, equal to an entire month of activity.

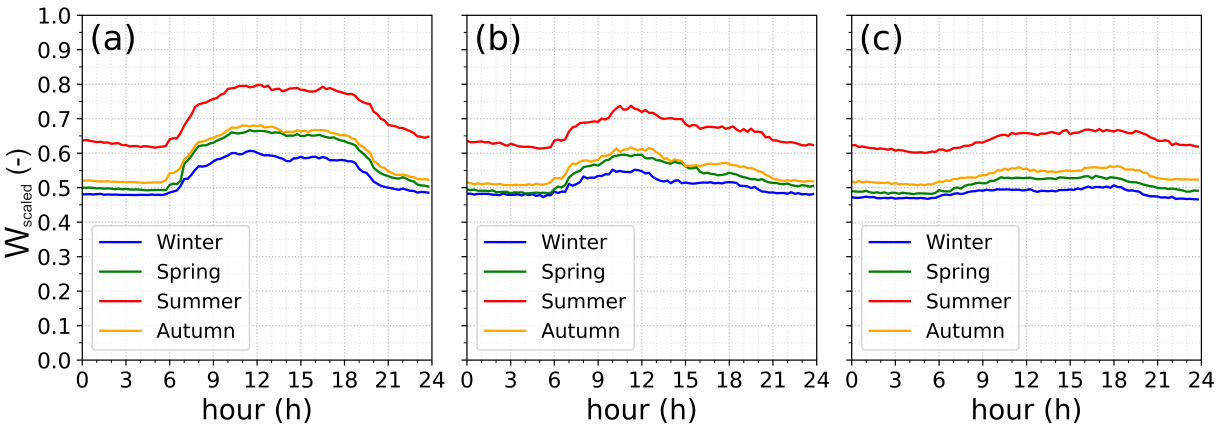

Figure 3.18: Average power load curve behaviour during the four seasons for workdays (a), Saturday (b) and Sunday/Festivity (c)

Figure 3.18 depicts the analysed data as daily average power curves, one for each season. The electricity power values have been scaled by dividing them by the maximum instantaneous electrical power consumption found in the year under analysis (W_{max}). If on the one hand, marked differences between winter, autumn and summer can be found, on the other hand, the spring power curve tends to have a behaviour similar to the autumn one. Focussing on the daily energy demand behaviour, the figure shows that workdays present similar characteristics regardless of the season. From midnight onwards, the electricity demand remains constant until 6:00 a.m. However, summer presents a nearly linear electricity demand decrease, which can be related to a lower operation intensity of the air handling system due to the gradual external temperature drop. Moving from 06:00 a.m. up to 08:30 a.m., progressive growth of the electrical power demand takes place.

All seasons present similar behaviours because of the increase in healthcare activity intensity. The most energy-intensive period takes place from 08:30 a.m. until 6:00 p.m, which is the most intense period for the building activities, leading to high energy demand for *HVAC* systems and lighting, and intense medical equipment exploitation. Morning and afternoon (10:30 a.m. - 1:00 p.m. and 4:00 p.m. - 6:00 p.m.) present higher electricity demand. Then, contrary to the morning, there is a progressive energy demand decrease, reaching typical night-time values by the end of the day. Sundays are characterized by smoother power curves. The electrical power demand night behaviour is similar to the one found during workdays.

Conversely, during the day, the electricity demand does not present the sharp increase typical of working days. Moreover, the day/night transition results longer than the one found in workdays, since the healthcare core activities are less intense and the medical staff on duty is minimum. Saturdays present intermediate characteristics between workdays and Sundays. An increase in energy consumption is apparent throughout the day, but reaches lower values compared to working days, and shows a progressive decrease from

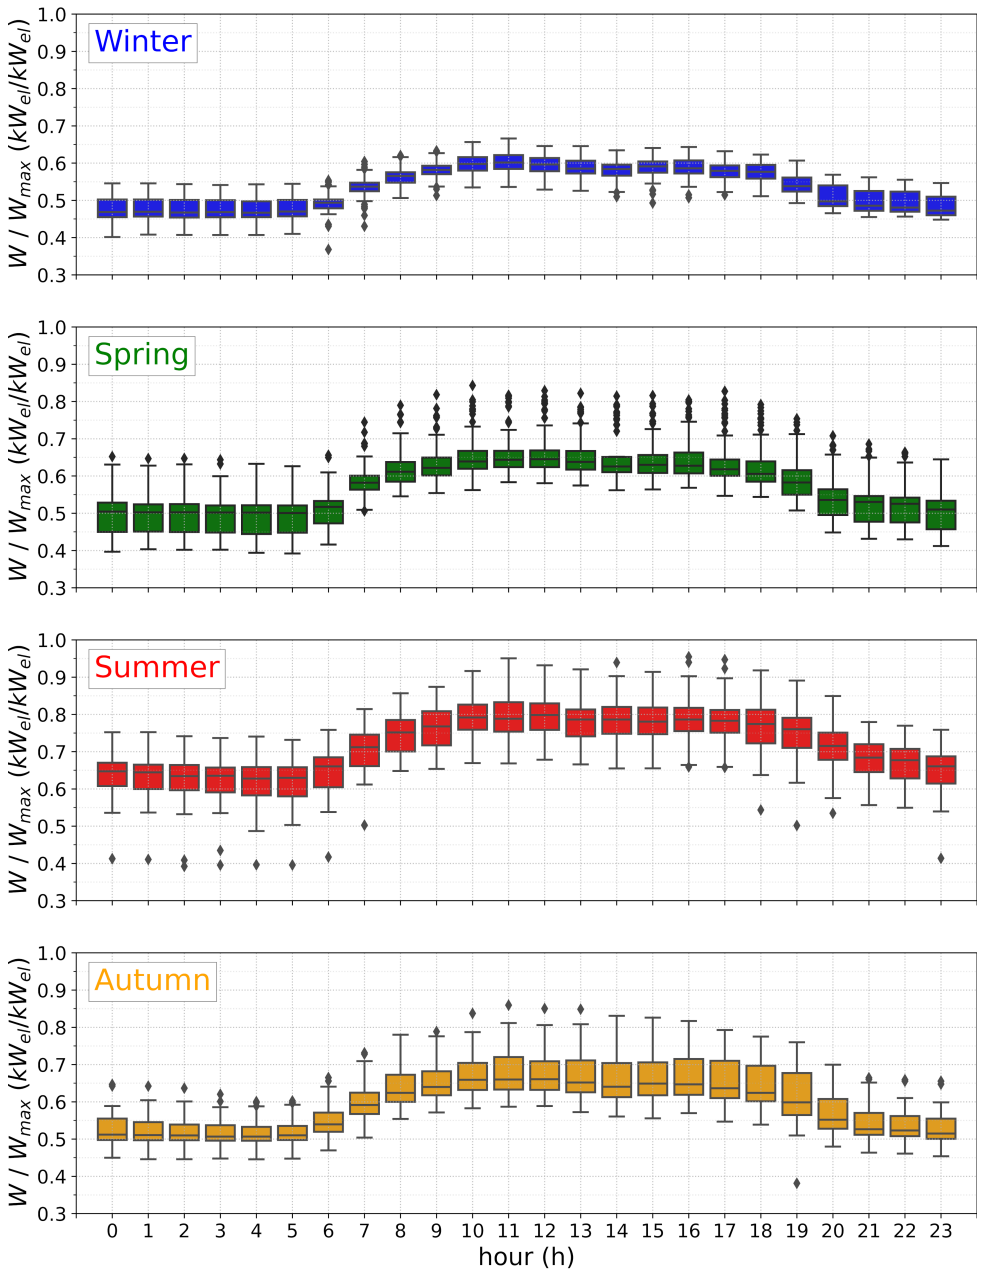

Figure 3.19: Box plots of the workday building energy demand divided by season

15:00, corresponding to the reduction in the intensity of healthcare activity.

Figure 3.19 shows the distribution of the hourly averaged electrical power demand through a box plot. The deviation of the power curves changes markedly based on the season but remains relatively constant throughout the day. The value of the standard deviation becomes larger as one moves from winter to summer. This behaviour can be explained considering that the summer electricity consumption is markedly affected by the external temperature due to the massive use of *HVAC* systems. Consequently, the temperature variation influence markedly more the summer electricity demand in comparison to the winter one. Indeed, the winter electrical power demand results weakly influenced by the external temperature, because the energy required for air handling purposes is mainly constituted by thermal energy. The electrical energy that is consumed

during winter (which includes air treatment-related applications - e.g., water and air circulation) is in any case less related to the outdoor air temperature in comparison to warmer seasons.

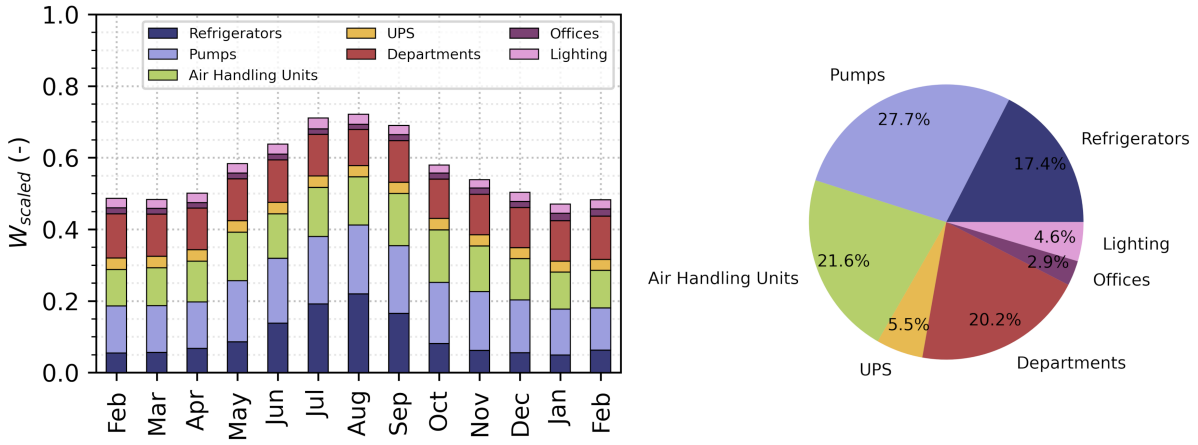

Figure 3.20: Healthcare facility energy demand repartition among different energy systems and applications

The electricity demand is composed of various usage categories, each of which may have different characteristics and requirements. Figure 3.20 shows the energy demand repartition among several different classes on monthly and yearly bases. The voice *Refrigerators* refers to the three electricity-driven chillers devoted to the water cooling process, which will be mainly exploited for air conditioning purposes. The cooled water will be send to the *Air Handling Units*, which perform the air treatments with hot water, chilled water and steam. The treated air circulation is provided by the *AHU* fans which exploits electricity to send the air to the various building areas. The chilled water circulation is ensured by several pumps, accounting for both refrigerators water circuits (primary and secondary) and the ducts which lead the chilled water to the air handling units. Moreover, *Pumps* include the electricity demand of all the pumps which compose the water circulation system. *Departments* encloses all the healthcare department and electro-medical equipment and is indeed one of the most energy intensive categories. On the other hand, *Offices*, *Lighting*, and *UPS* results less relevant in shaping overall electricity demand.

Indeed, Figure 3.20 clearly shows that 72.1% of the yearly energy requirement (referring to the period which ranges from *February 12, 2020* and *February 11, 2021*) is due to the *HVAC* system operations (which mainly consists in *Refrigerators*, *Air Handling Units* and *Pumps*). If on one hand *AHU* and *Pumps* present low electricity requirements variations throughout the analysed year, the refrigerators show a marked relation with the outdoor air temperature, presenting a marked increase in electric energy needs as the summer season approaches, due to the massive energy requirements for air conditioning purposes.

Figure 3.21 shows the refrigerators electric power demand form *February 12, 2020* to

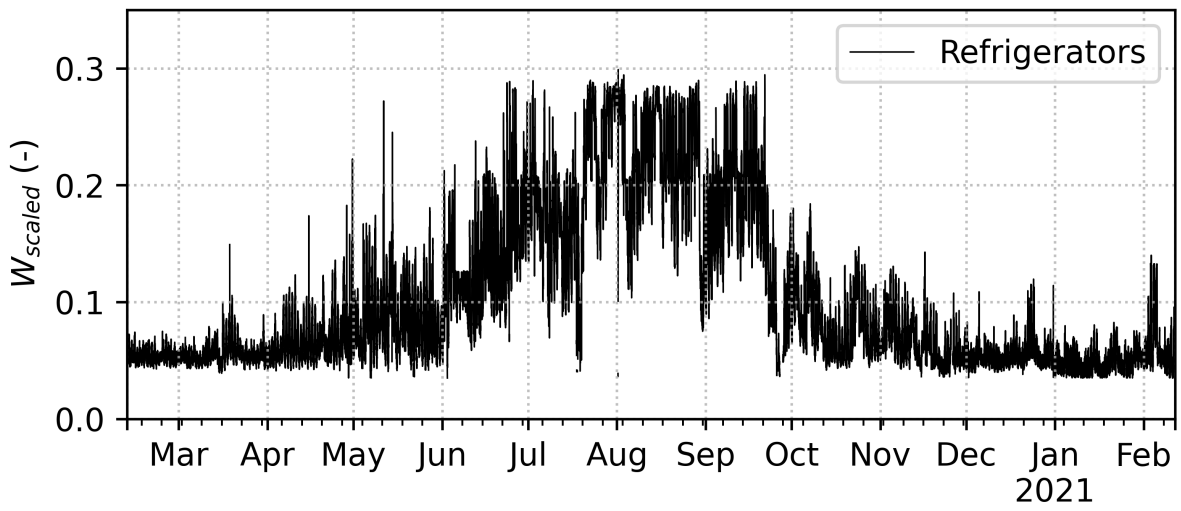

Figure 3.21: Refrigerators electric power demand form *February 12, 2020* to *February 11, 2021*

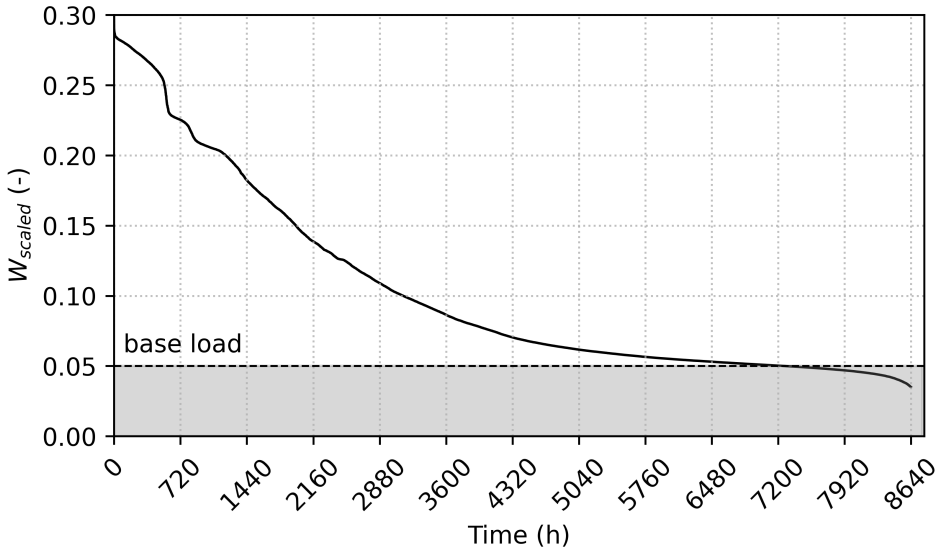

Figure 3.22: Refrigerators electric power demand cumulated curve

February 11, 2021, highlighting the notable increase of the energy requirements during the summer due to the massive use of air conditioning systems. Indeed, the refrigerators cumulated power curve (Figure 3.22) highlight that 50% of the refrigerators operative time osculates around tde "base load" value (which is the 5% of the maximum whole building power demand). The rest of the time, the refrigerators shows an increase in the absorbed power with the maximum one that reach almost 30% of the maximum whole building power demand, confirming the marked influence of the seasons on the electricity needs due to refrigeration purposes.

This aspect is better highlighted by observing the average power curves depicted in Figure 3.23. Indeed, spring, autumn and winter seasons tents to have similar behaviours, even though, during the daytime, spring shows a higher demand for chilled water. This

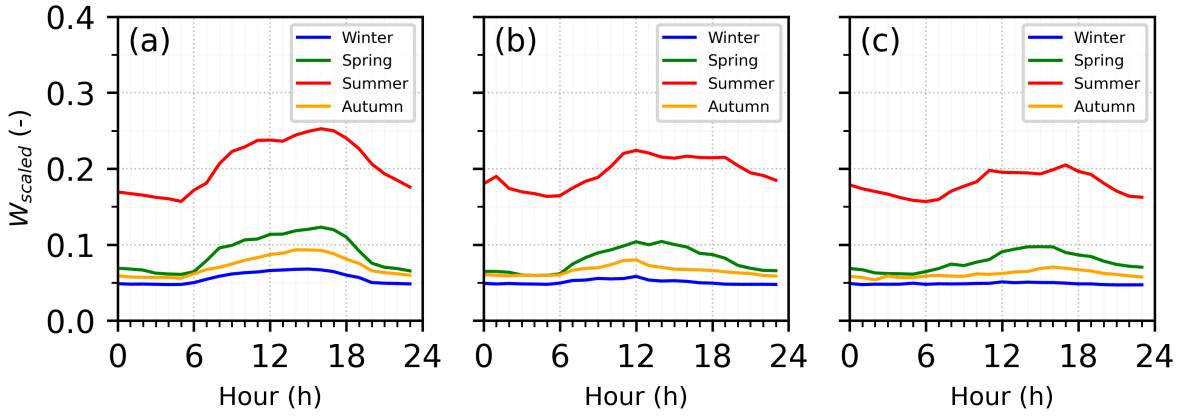

Figure 3.23: Refrigerators average power load curve behaviour during the four seasons for workdays (a), Saturday (b) and Sunday/Festivity (c)

phenomenon may be due to the psychological impact that the transitional seasons have on individuals. During the spring season, coming from winter, where people are accustomed to cold environments, they may be more inclined to turn on air conditioning. Conversely, in autumn, coming from the hot summer, it is more common to perceive the cold more acutely at the similar air temperature and humidity, leading to a lower use of air conditioning systems. Nevertheless, the summer results unequivocally higher than the previous three, resulting in a massive requirements for electricity during the whole day. Focussing on the weekly time scale, it can be seen that the differences noted for the whole building energy demand between workdays and weekends (see Figure 3.18) can be found also of the refrigerators energy demand.

As already noticed through the monthly energy demand reported in figure 3.20, the air handling units electric power demand (Figure 3.24) shows a more regular behaviour throughout the year in comparison to the refrigerators one, being mostly influenced by the daily cycle in comparison to the seasonal one. Also the cumulated power curve (Figure 3.25) shows this behaviour. The base load is not clear even if the minimum electrical absorbed power is above 7.5% of the maximum building electrical power need. The cumulated curve shows a linear increase while approaching its maximum value, which results to be 17.5% of the whole building maximum power demand.

By observing the average power curves (Figure 3.26), it can be confirmed that there is a marked variation of the energy demand during the daily time scale, leading to higher requirements during the daytime, when the healthcare activities are more intense. The seasonal behaviour shows that summer and autumn present higher power demands, followed by the spring and, then, by the winter. By comparing autumn and spring, it can be seen that the behaviour is in contrast with the one found for the refrigerators. Nevertheless, the variation due to the people perception influence only the temperature in which the treated air is provided to the internal environments, since the local thermal control allow

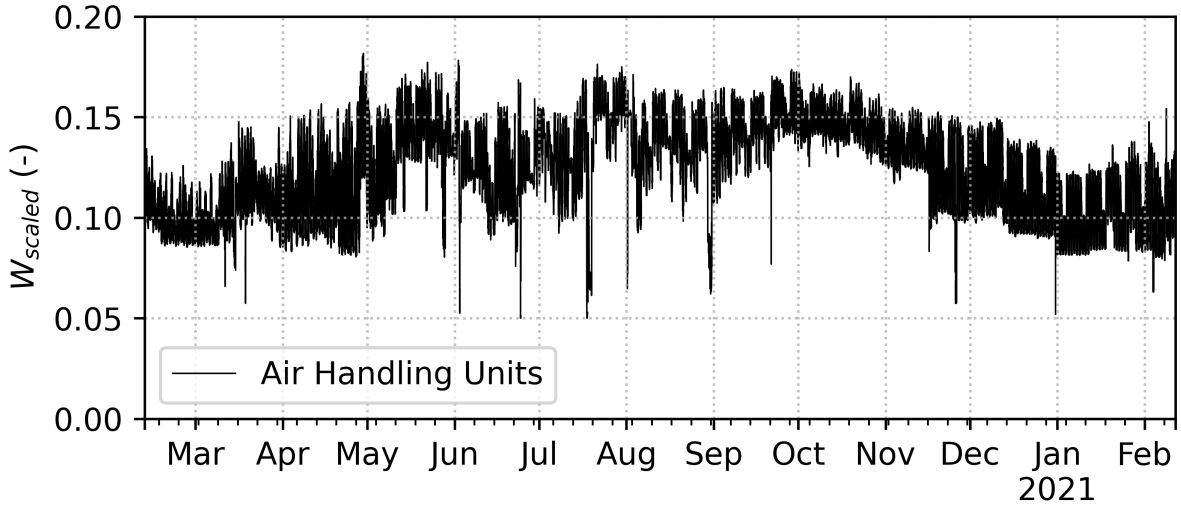

Figure 3.24: Air handling units electric power demand form *February 12, 2020* to *February 11, 2021*

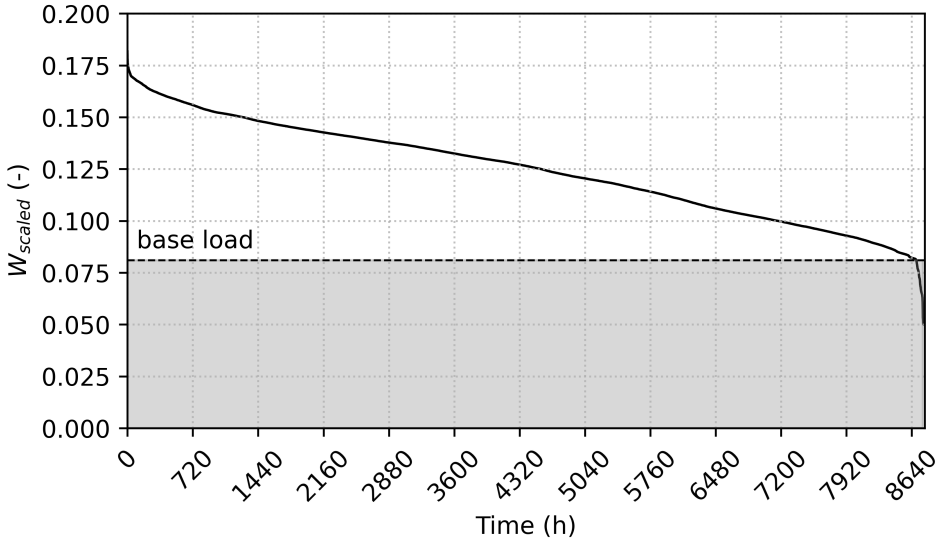

Figure 3.25: Air handling units cumulated electric power curve

the user to modify only the temperature set-point and not the air mass flow rate.

It is logical to assume that the energy consumption of the air handling unit, which utilizes chilled water, mimics the increased demand for energy observed in refrigeration systems during the spring season in comparison to the autumn. The opposite behaviour exhibited by air handling units is due to the significant influence that relative humidity has on the performance of evaporative towers, which serve refrigeration systems. High relative humidity results in reduced cooling capacity of the towers, and thus, a decrease in the performance of refrigeration units. Therefore, despite the slightly lower demand for chilled water by air handling units during the spring season, the higher rate of external air humidity leads to a greater consumption of electricity by refrigeration units.

Regarding pumps, observing the annual trend of absorbed power, a nearly constant

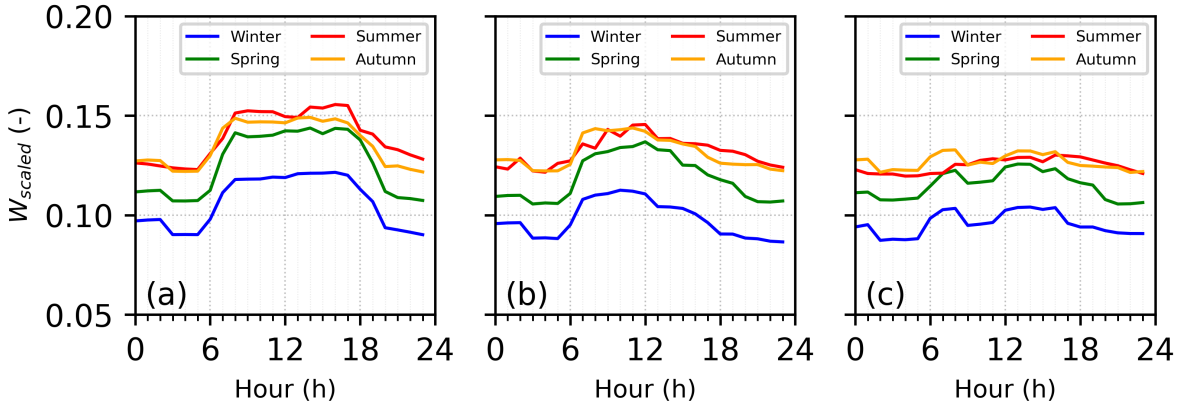

Figure 3.26: Air handling units average power load curve behaviour during the four seasons for workdays (a), Saturday (b) and Sunday/Festivity (c)

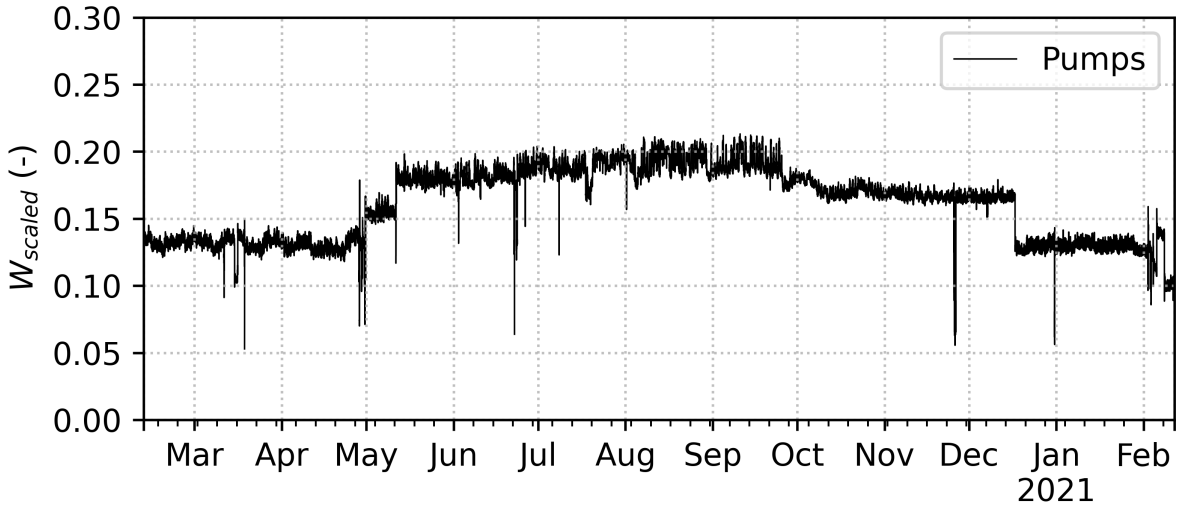

Figure 3.27: Air handling units electric power demand form *February 12, 2020* to *February 11, 2021*

behaviour can be identified during different periods of the year. There are mainly sudden changes in certain periods of the year, from which the step-like behaviour visible in the Figure 3.27 derives. Also observing the average power load curves (Figure 3.29), a constant trend of energy requirements is identified on the daily time scale, regardless of the day typology (workday, Saturday or Festivity), with slight increases during the daytime of summer workdays.

Indeed, looking at Figure 3.28, the "step-like" behaviour of the cumulated curve is apparent. The minimum load is found as almost 9% of the maximum building electricity need. A first step which oscillates around 12.5% take place for almost 3′000 hours per year. The highest energy needs take place for more than 5′000 hours, oscillating between 15% and 20% of the maximum whole building power demand.

This peculiar behaviour can be explained by considering that the primary circuit of the refrigeration units is responsible for a considerable part of the consumption related to

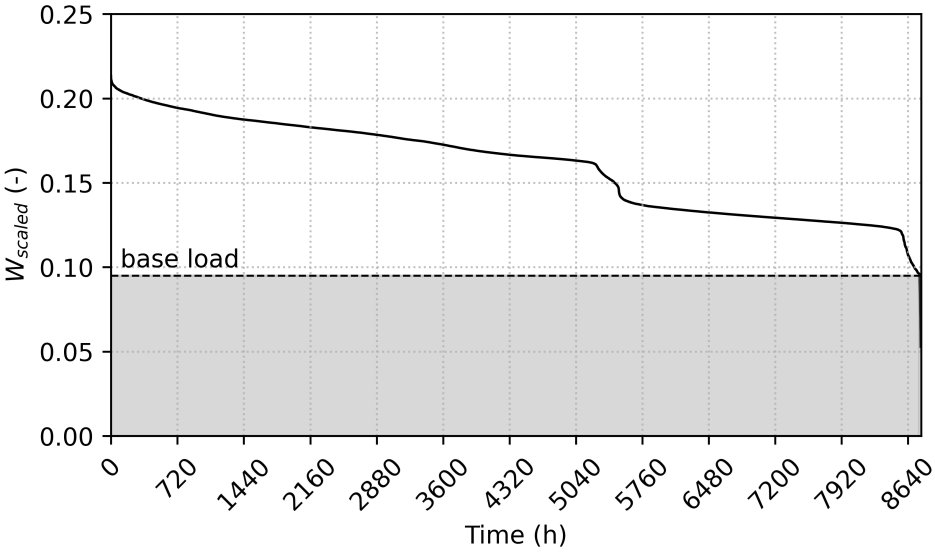

Figure 3.28: Air handling units cumulated electric power curve

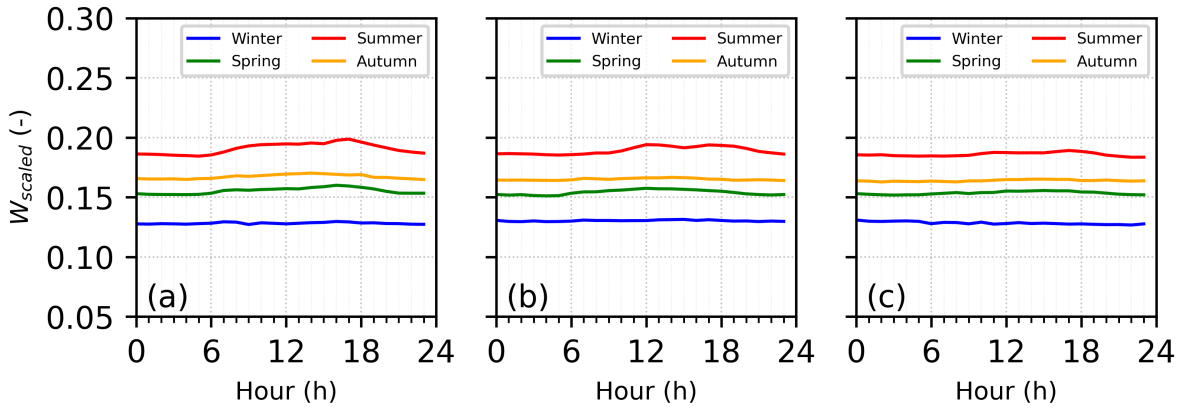

Figure 3.29: Air handling units average power load curve behaviour during the four seasons for workdays (a), Saturday (b) and Sunday/Festivity (c)

the pumps. During the summer season, where the need for chilled water is higher than the rest of the year, generally two or three refrigerators are operative, resulting in a significant increase in electrical energy related to the pumps. In other seasons, the need for cold water is generally satisfied by a single refrigeration unit, resulting in a sudden reduction of the power demand of the circulation systems. Furthermore, during winter, the demand for cold water from the air handling units is at its minimum, reducing the need for circulation of cold water in the ducts that branch out throughout the building. This results in the additional reduction that can be observed during the winter season.

Chapter 4

Results

Following the guidelines outlined in section 2.2, the present chapter is dedicated to the presentation and analysis of the obtained results. The first part provides a detailed analysis of the application of all cases procedures included in the proposed methodology. Each step defined in the proposed method is presented and commented. In terms of *feature selection*, both the filter method (applied to **Cases 1**, **2**, and **3**) and the wrapper method applied to **Cases 4**, **5**, and **6** are analysed, and the resulting feature rankings are compared. Subsequently, the results obtained from the final models training procedures are analysed, comparing the resulting performance metrics and the time required to complete the entire computation. Since the hyperparameter tuning processes are only performed for cases **3** and **6**, the related results will be included in the respective sections, analysing both the grid search process and the final model results. As a result, the most accurate procedure will be defined. Through this information, the chosen case is applied to the electricity demand data for the *HVAC* system components described in section 3.3. The results obtained are analysed and compared, Similarly to what was done for the whole building electricity demand.

The second and final part of this chapter focuses on utilizing the models obtained through the previous analyses to finalize the monitoring methods using the *CUSUM* technique. In particular, the models are applied to the part of the dataset related to the period which follow the training one, analysing the results through the realization of the the *CUSUM* control charts and comparing them with the actual events of the building and its systems during the monitoring period.

4.1 Whole Building energy prediction

4.1.1 Feature Selection

This section is devoted to the application of the two feature selection methods illustrated in section 2.1.2. Referring to table 2.1, the analysed cases are divided into the following cases: **Case 1** (*MLR*) and **Cases 2** and **3** (*ANN*) exploit the methodology called *Filter Method* (see section 2.1.2), while **Case 4** (*MLR*) and **Cases 5** and **6** (both *ANN*) exploit the *Wrapper Method* (see section 2.1.2).

Filter Method feature selection: Cases 1, 2 and 3

Due to the large number of features available, the correlation analysis will be shown by dividing them into several categories in order to facilitate the data visualization. It is important to specify that features represented through *One-Hot encoding* (see section 2.1.3) will not be included in the correlation analysis. Indeed, the correlation coefficient between the electricity demand and a binary vector will be low and not indicative of the effect this will have on the final model. Since literature analysis has shown that it is often advisable to represent categorical variables through that encoding method, categorical features that will be included in the following steps of the method will automatically be represented in *One-Hot encoding*. As already explained in section 2.1.2, the *Filter Method* imposes the definition of two different *thresholds* in order to identify both **irrelevant** and **redundant** features among the available ones. These thresholds have no standard values. They indeed strongly depend on the specific available dataset. Although there are no standard rule to define the thresholds, some common values can be found in similar applications. Setting $\tau_i = 0.6$ for irrelevant features and $\tau_r = 0.9$ for redundant features seems to work in several cases. However, by analysing the correlation coefficients of the features with the target

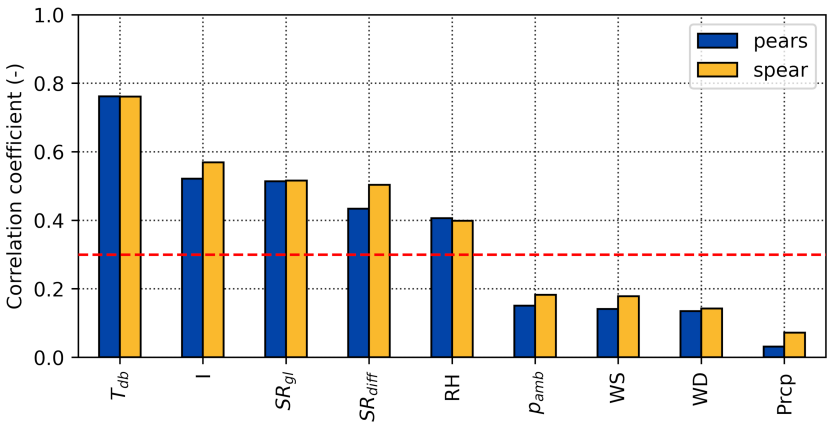

Figure 4.1: Basic climate features ranking based on their correlation coefficients with the power demand

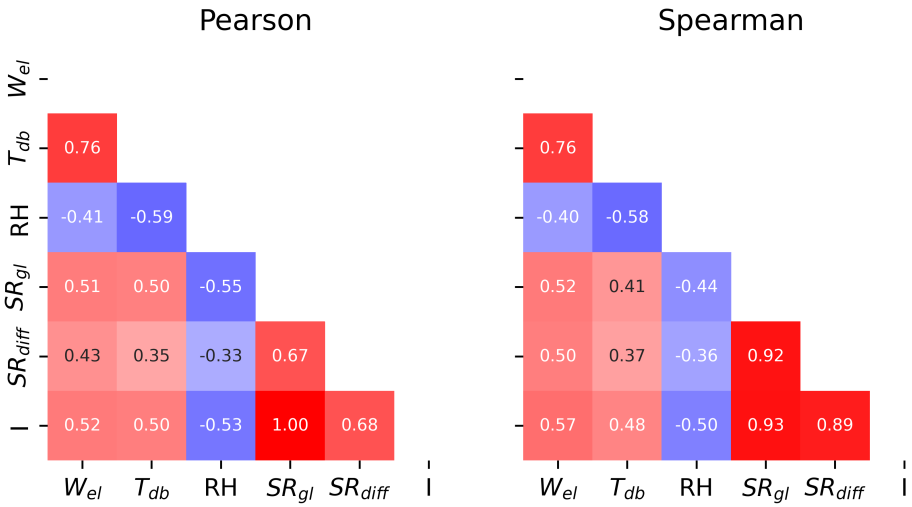

Figure 4.2: Correlation coefficient matrix for standard climate features

(Figures 4.1, 4.4, 4.6), a fairly sharp step between highly and less correlated features can be identified. This value is 0.3, and was therefore taken as a threshold for the identification of less relevant features. As regards the redundant threshold, instead, the value found in literature was deemed suitable for the given dataset. To resume, the threshold used to identify the irrelevant feature is $\tau_i = 0.3$, the one related to the redundant feature identification is $\tau_r = 0.9$.

Looking at the climate features available(Figure 4.1), i.e. features from direct measurements (see Table 3.6 for more information), it is possible to observe their correlation coefficients with electrical power demand (W_{el}). The most correlated feature with electrical power demand is T_{db}, (Outdoor air dry bulb temperature, $r = 0.76$ and $\rho = 0.76$). Considering that the healthcare facility has large air conditioning systems, the temperature of the external environment will necessarily have a great influence on its utilisation and, consequently, on the resulting energy demand. The cooling systems (refrigerators) exploit

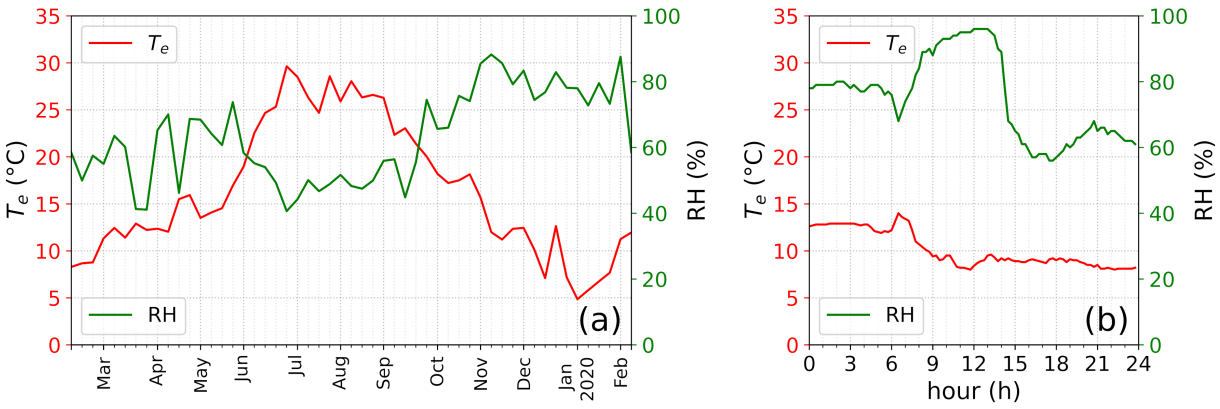

Figure 4.3: Outdoor air temperature and relative humidity behaviour: Weekly averaged 2019 (a), 05-05-2019 (b)

electricity for its operations, while the hot water generation is obtained through thermal energy (Natural Gas). Indeed, the increase in electricity demand of the building increase with the increase of outdoor air, resulting in positive correlation coefficients.

Other highly correlated features are those related to the solar radiation, respectively I (Illuminance, $r = 0.52$ and $\rho = 0.57$), SR_{gl} (Global solar radiation, $r = 0.51$ and $\rho = 0.52$) and SR_{diff} (Diffusive solar radiation, $r = 0.43$ and $\rho = 0.50$). All the correlation coefficients are positive, meaning that the electricity demand increase with the solar radiation. This can be explained considering that the main healthcare activities are performed during the day, even if the building remain operative continuously. Consequently, it could be assumed that the high correlation coefficient is not due to the direct effect of solar radiation, but

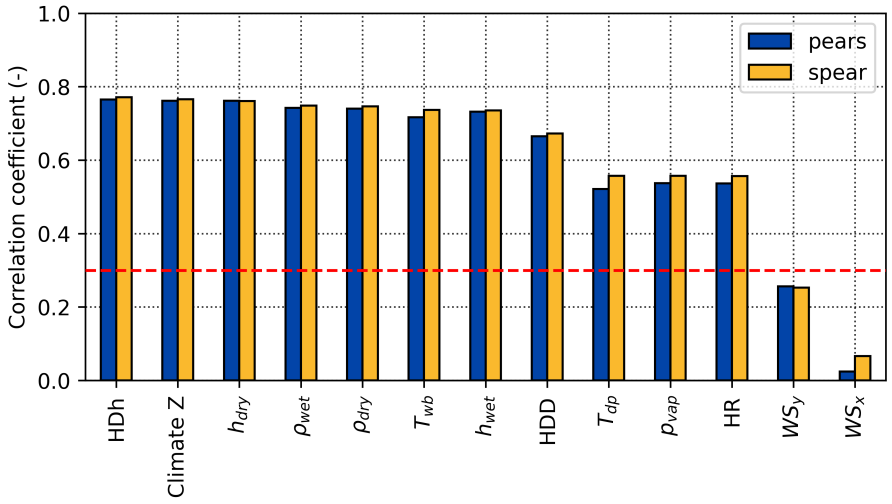

Figure 4.4: Elaborated climate features ranking based on their correlation coefficients with the power demand

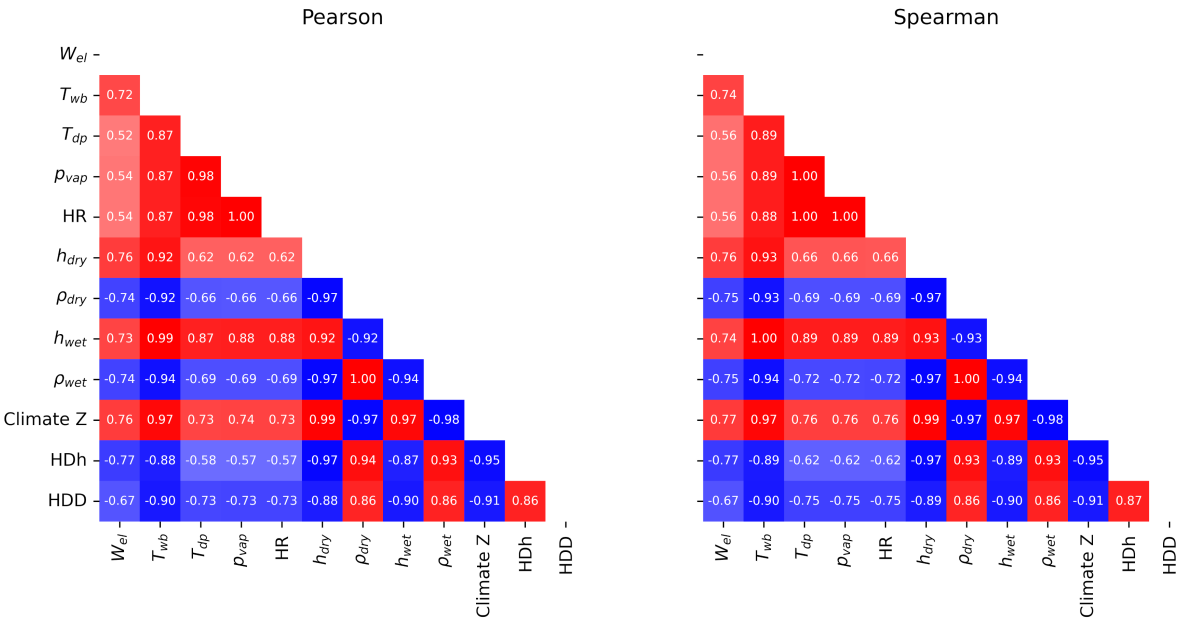

Figure 4.5: Correlation coefficient matrix for elaborated climate features

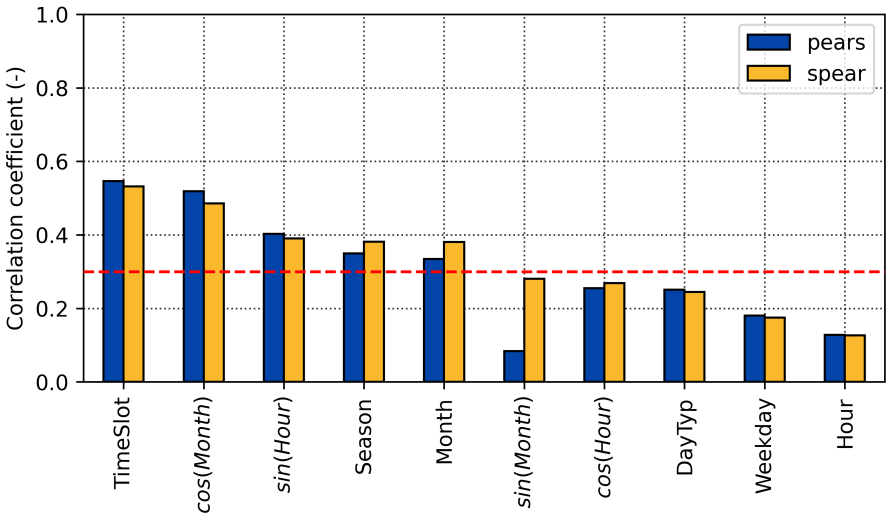

Figure 4.6: Calendar features ranking based on their correlation coefficients with the power demand

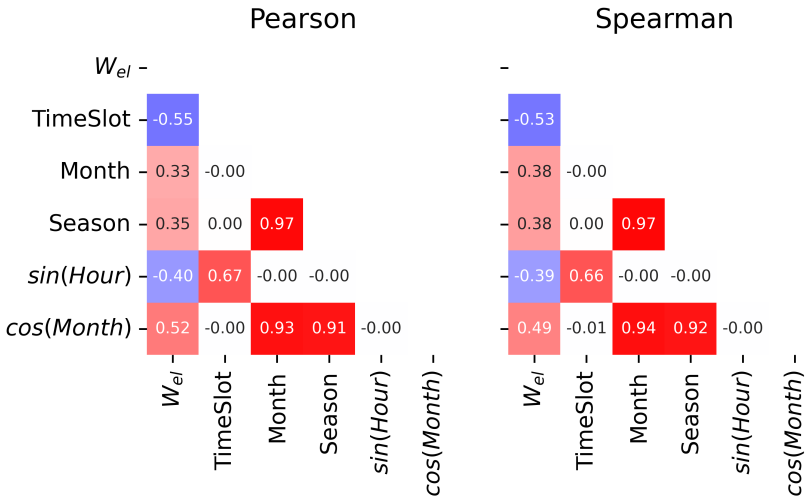

Figure 4.7: Correlation coefficient matrix for Calendar features

to a coincidence with increased activity during the day. Nevertheless, the facility under analysis present large windowed surfaces on the roof. This characteristic lead to a marked increase of cooling requirements during the summer days, with a consequent increase in refrigerators and air handling units exploitation.

The last accepted feature of this subset is RH ($r = -0.41$, $\rho = -0.40$), which is the outdoor air relative humidity. This feature might appear to be of minor importance. However, the fact that cooling units are cooled by evaporative towers must be taken into account. Consequently, the efficiency of the towers, and therefore of the chillers, will be very much related to humidity.

Ambient air pressure (p_{amb}), Wind speed and direction (WS, WD) and precipitation ($Prcp$), present low correlation with the power demand ($r, \rho < \tau_i$) and can be discarded.

Figure 4.2 highlight a high correlation between I and SR_{gl} ($r = 1.00$, $\rho = 0.93$), and

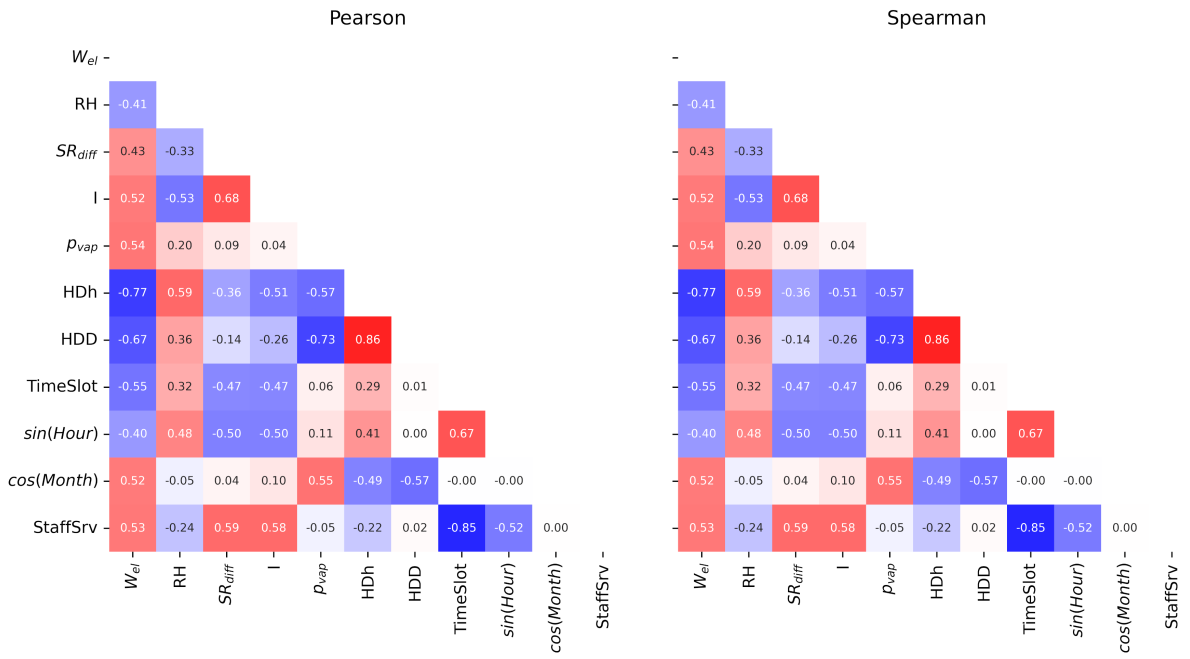

Figure 4.8: Correlation coefficient matrix for remaining features

SR_{gl} results highly correlated also with SR_{diff} ($r = 0.67$, $\rho = 0.92$). Contrarily, I and SR_{diff} present correlation coefficients lower than the redundant threshold τ_r. Consequently, the SR_{gl} can be identified as a **redundant** feature and can be discarded from the dataset.

It is important to notice that other features present intermediate correlations between each other. As an example, T_{db} shows a strong correlation with the RH ($r = -0.59$, $\rho = -0.58$). In particular, RH tends to decrease corresponding to a temperature increase. Even that the correlation coefficient between outdoor air temperature and relative humidity highlight a negative correlation that reflects the yearly behaviour of these two variables (Figure 4.3a), looking at a daily time scale the specific behaviour can be completely different. Figure 4.3a shows the weekly averaged behaviour of both outdoor air temperature and relative humidity for the year *2020*. It seems to confirm the information obtained by simply looking at the correlation coefficient between the two variables, reported in Figure 4.2. Indeed, relative humidity tends to decrease as the outside air temperature rises, resulting in a negative linear correlation. Nevertheless, focusing on some peculiar daily behaviour as $May5th, 2019$ (see Figure 4.3b), it can be noticed a different correlation which leads to a contemporary decrease of the two variables, a trend that is opposite to the one indicated by the related correlation coefficient. Consequently, to correctly characterize the daily time-scale building electricity demand, it is important to take into account both features as input of the machine learning model.

Following section 2.1.3, the available climate features have been elaborated through psychrometric relations. Morover, the *Climate Z* feature has been added according to

equation 2.13. Looking at Figure 4.4, appear evident that there are two features which can be considered as **irrelevant**: namely WS_y ($r = -0.26$, $\rho = -0.25$) and WS_x ($r = -0.025$, $\rho = -0.067$), confirming the information obtained through the basic climate features analysis explained before.

Figure 4.5 highlight several high correlations between the input features. The redundant couples of features are:

- T_{wb} and h_{wet} ($r = 0.99$, $\rho = 1.00$)
- T_{dp} and p_{vap} ($r = 0.98$, $\rho = 1.00$)
- p_{vap} and HR ($r = 1.00$, $\rho = 1.00$)
- h_{dry} and *ClimateZ* ($r = 0.99$, $\rho = 0.99$)
- ρ_{dry} and ρ_{wet} ($r = 0.99$, $\rho = 1.00$)
- h_{wet} and *ClimateZ* ($r = 0.97$, $\rho = 0.97$)
- ρ_{wet} and *ClimateZ* ($r = 0.98$, $\rho = 0.98$)
- *ClimateZ* and *HDh* ($r = 0.95$, $\rho = 0.95$)

Given the high degree of redundancy of these data, by referring to the correlations between features and targets, it can be established that the **redundant** features are T_{wb}, T_{dp}, HR, h_{dry}, ρ_{dry}, h_{wet}, ρ_{wet} and *ClimateZ*.

Moving to the **calendar features**, Figure 4.6 shows that five features present weak correlation with the electricity demand. Nevertheless, the sine functions (see section 2.1.3) applied to *Month* and *Hour* are couples of feature which are necessary to correctly represent the "original" feature. As an example, the hour of the day represented with a single sine function, will present the same value in two different moments of the day, leading to the risk of compromising the prediction performance of the *ML* model. Consequently, the presence of high correlation between only one of the two features and electrical power, elects both as relevant features. The **irrelevant** calendar features are therefore: *DayTyp*, *Weekday* and *Hour* (ordinal encoding). It is interesting to notice that both representation of *Month* feature (ordinal encoding and sine functions) are potential input features for the subsequent *ML* model.

Nevertheless, looking at Figure 4.7, it can be seen that *Month* and *Season* feature present high correlations between each other ($r = 0.97$, $\rho = 0.97$). This fact is reasonable, given that *Season* and *Month* parameters have similar time-scales. Since *Season* results to have higher correlation with the power demand (target), the *Month* feature will be marked as redundant. Nevertheless, *Season* present high correlation also with a different

Table 4.1: Feature rejected by filter method: Case 1 and 3

Subset	**Irrelevant**	**Redundant**
Standard climate	p_{amb}, WS, WD, $Prcp$	SR_{gl}
Extracted climate	WS_x, WS_y	T_{wb}, T_{dp}, HR, h_{wet}, h_{dry}, ρ_{wet}, ρ_{dry} *ClimateZ*
Standard calendar	$DayTyp$, $Weekday$, $Hour$	$Month$, $Season$
Remaining features	-	T_{db}

numerical representation of the *Month* feature: $\cos(M)$ ($r = 0.91$, $\rho = 0.92$). In this case, *Season* have a weaker correlation with the target, and will be consequently be excluded from the following analysis.

Once all the features subgroups correlations were analysed, a final check on the remaining features must be done in order to ensure the removing of further redundancy in the dataset (given that all the remaining features have been selected as *relevant features*). Figure 4.8 shows the correlation coefficients grids of the remaining input features. In fact, there is one last pair of features that shows correlation coefficients greater than the redundancy threshold (τ_r): HDh and T_{db} ($r = 0.97$, $\rho = 0.97$). Between these two, the most correlated is HDh, showing that the "difference" between outdoor and indoor air temperature have greater influence on the overall consumptions. Consequently, T_{db} can be

Table 4.2: Final features ranking obtained by appllying filter method

Rank	**Feature**
1	*HDh*
2	*HDD*
3	*TimeSlot (One-Hot)*
4	p_{vap}
5	*StaffSrv*
6	*I*
7	*Month (Sine functions)*
8	SR_{diff}
9	*RH*
10	*Hour (Sine functions)*

marked as **redundant**, and will be removed from the final dataset.

The correlation matrices were thoroughly examined at this stage. To summarize the results of this analysis, in Table 4.1 all the removed feature are shown. They are divided in *Irrelevant* and *Redundant* features and also by their subgroup. Moreover, also a ranking of the importance of accepted features can be drawn up (Table 4.2) based on the correlation coefficients that they have with the electrical power demand. The *Final Dataset* obtained with the *Filter Method* is composed of 12 features.

The analysis of the correlation matrices revealed that the features with the most significant impact on the building electrical energy demand are heating degree days (HDD) and heating degree hours (HDh). These features provide insight into the local climate conditions, particularly dry bulb temperature (T_{db}). This underscores the significant influence of the $HVAC$ system on the overall energy consumption of the building. Other climate features are the vapour pressure (p_{vap}) and relative humidity (RH), which highlight the influence of the air humidity on the systems performance. A higher humidity leads to increase the electricity used for dehumidification purposes.

Moreover, as already said, a higher humidity leads to worse performance of the evaporative towers, resulting in reduced performance of the entire cooling system, increasing consumption due to the refrigeration units and primary circuit pumps (refrigeration units - evaporative towers) and secondary circuit (refrigeration units - air handling units). Illuminance (I) and diffusive solar radiation (SR_{diff}) also contribute to the energy demand behaviour. The large windowed surfaces of the building contribute to increase the building heating on sunny days, which, during the summer period, leads to higher electricity demand for air conditioning. Indeed, the combination of the two features provide information about cloud cover at a specific time step.

Calendar features include three variables related to different time scales. *Hour* captures daily power demand variation, *TimeSlot* reflects weekly rhythms, and *Month* captures seasonal trends.

Finally, the number of staff in service (*StaffSrv*) appears to also influence power demand behaviour, providing insight about the intensity of healthcare activity and its influence on the overall electricity demand.

Wrapper method feature selection applied to MLR: Case 4

Table 4.3: Accepted features for wrapper method with MLR

Rank	Feature	R^2	MSE	RMSE	MAE
1	*ClimateZ*	0.6616	0.0055589	0.074558	0.060589
2	*TimeSlot* (One-Hot)	0.7965	0.0033433	0.057821	0.045801
3	*Month* (One-Hot)	0.8658	0.0022039	0.046945	0.034762
4	*Hour* (One-Hot)	0.8762	0.0020333	0.045092	0.033093
5	*HDD*	0.8794	0.0019807	0.044505	0.032723
6	h_{wet}	0.8822	0.0019353	0.043992	0.032632
7	T_{wb}	0.9102	0.0014756	0.038414	0.027534
8	*StaffSrv*	0.9121	0.0014440	0.038000	0.027211
9	*Weekday* (One-Hot)	0.9131	0.0014267	0.037771	0.026965
10	*RH*	0.9137	0.0014175	0.037650	0.026703
11	p_{amb}	0.9141	0.0014111	0.037564	0.026634
12	*DayTyp* (One-Hot)	0.9143	0.0014067	0.037506	0.026593
13	rho_{dry}	0.9146	0.0014026	0.037451	0.026564
14	T_{dp}	0.9152	0.0013929	0.037321	0.026458
15	h_{dry}	0.9159	0.0013811	0.037163	0.026290
16	*HR*	0.9163	0.0013753	0.037086	0.026144
17	*Hour* (Sine functions)	0.9164	0.0013722	0.037043	0.026116
18	*I*	0.9166	0.0013704	0.037019	0.026108
19	w_x, w_y	0.9166	0.0013690	0.037001	0.026092
20	ρ_{wet}	0.9167	0.0013686	0.036994	0.026054
21	p_{vap}	0.9170	0.0013636	0.036928	0.026018
22	SR_{gl}	0.9170	0.0013635	0.036926	0.026013
23	*TimeSlot* (Ordinal)	0.9170	0.0013634	0.036924	0.026013
24	*HDh*	0.9170	0.0013634	0.036925	0.026022
25	*Hour* (Ordinal)	0.9170	0.0013633	0.036922	0.026019

Applying the wrapper method to the feature selection process leads to different results in comparison to the preceding method. In particular, the application of this method on **Case 4**, which apply the *MLR* method on the same dataset, the accepted features

pass from the 10 accepted by the filter method, to 25. The final ranking of the accepted features is reported in Table 4.3, and the metrics obtained for each steps are also reported in the table.

The feature ranking results markedly different from those obtained using the filter method shown in Table 4.2. In the present case, the most important feature is *ClimateZ*, whose representation was discussed in section 2.1.3, Equation 2.13). *ClimateZ* combine three of the main parameters which are useful to characterize the outdoor climate conditions: dry bulb air temperature (T_{db}), air humidity (RH) and global solar radiation (SR_{gl}). The impact of weather conditions on the energy consumption of a building has been thoroughly examined and discussed in previous chapters, but *ClimateZ* seems to present climate information in a way that facilitates the learning process. Indeed, the *ClimateZ* alone is able to obtain a model with a relatively high prediction performance ($R^2 = 0.6616$).

As regards the calendar features, *TimeSlot*, *Month* and *Hour* (represented with *One-Hot encoding*) are respectively 2^{nd}, 3^{rd} and 4^{th} in the features ranking (Table 4.3), indicating that the building energy demand is significantly affected by daily, weekly and seasonal time-scales. The combination of these three features and the previously accepted *ClimateZ* let to enhance the prediction performance, reaching an R^2 of 0.8762. Some additional consideration can be done regarding calendar features that have been accepted by the wrapper method: $TimeSlot$ most important representation is the *One-Hot encoding*, but also the *Ordinal encoding* play a role in the model learning process, even if in a less marked way. $Hour$ feature is accepted in three different representation. *One Hot encoding* allows for differentiation between different time steps, *Sine functions* (defined in section 2.1.3) carries with it information about the cyclical nature of each day. *Ordinal encoding* has a minimal impact on prediction performance, but when combined with the other two methods, it still allows for the establishment of more effective relationships between energy demand and hour of day. Moreover, also the $Weekday$ and $DayTyp$ features have been accepted. This suggests that the different days of the week have unique characteristics that are repeated from week to week, which can be associated with the weekly cycle of activities that take place within the facility, and differences between each day of the week can be systematically found (not only the already noticed differentiation between the day typology). Consequently, the designation of different days of the week as an input feature allows for a better representation of the weekly cycle of the building electricity demand.

Regarding climate features, apart from *ClimateZ*, others have proven to make a significant contribution to the model learning process. heating degree days (HDD) is 5^{th} in features ranking, a sign that the daily average temperature difference between outdoor and indoor environments significantly influence the air handling systems exploitations. Moreover, moist air enthalpy (h_{wet}) and wet bulb temperature (T_{wb}) are respectively 6^{th}

and 7^{th} in the feature ranking, as they carry with them information about the humidity of the outdoor air. Several other climate features have been accepted, even if the highest contributes come from the first seven. The complete process of feature selection have required 76 seconds leading on a preliminary R^2 of 0.9170.

It is important to remark that the metrics shown at the end of the feature selection process do not necessarily be equal to the one obtained by evaluating the final model. Indeed, while the wrapper feature selection exploits the *k-fold cross validation* (see section 2.1.2), evaluating the model on 33% of the available training dataset, the final model will be trained using 70% of the dataset for the training process, 10% for internal validation purposes, and 20% for the final performance evaluation. Moreover, each part of the dataset has been randomly extracted in order to ensure the maximum generality of the obtained models and, consequently, does not necessarily use the same samples to train the model, leading to evaluate the metrics on different sub-datasets.

Wrapper method feature selection applied to ANN: Case 5 and 6

Table 4.4: Accepted features for wrapper method using ANNs

Rank	**Feature**	R^2	**MSE**	**RMSE**	**MAE**
1	*ClimateZ*	0.7250	0.0045174	0.067212	0.053976
2	*StaffSrv*	0.8718	0.0021057	0.045888	0.034359
3	*Month* (One-Hot)	0.9140	0.0014129	0.037589	0.026064
4	*Hour* (Sine functions)	0.9293	0.0011620	0.034088	0.023618
5	*Weekday* (One-Hot)	0.9425	0.0009438	0.030722	0.021352
6	T_{dp}	0.9535	0.0007642	0.027644	0.019498
7	p_{amb}	0.9603	0.0006521	0.025536	0.017885
8	*HDD*	0.9624	0.0006167	0.024834	0.017514
9	*HR*	0.9638	0.0005949	0.024391	0.017185
10	*Hour* (Ordinal)	0.9639	0.0005916	0.024323	0.017382
11	*DayTyp* (Ordinal)	0.9642	0.0005883	0.024255	0.017143
12	*RH*	0.9639	0.0005931	0.024353	0.017207
13	ρ_{wet}	0.9642	0.0005880	0.024249	0.017142

This particular case imply the application of the feature selection through wrapper method applied to *ANNs*. Consequently, it is important to define a starting network architecture (including all hyperparameters) that allows for training in the different steps of the method. As previously mentioned, the choice of hyperparameters is based on those most commonly used in literature for this type of application (literature-suggested in Figure 2.7). Below, the parameters used for the specific goal of feature selection will be described. Considering all the available features, the maximum number of inputs (n_i) reaches 86 (considering that several features are composed of multiple vectors), with a single output ($n_o = 1$). Following the altrady discussed Equation 2.8, the number of neurons in the hidden layer (HLN) have been set to 87. As regards the other hyperparameters, the *learning rate* has been set to 0.01, while the chosen activation function is the *ReLu*, which is the standards for regression applications.

Using the wrapper method for feature selection on artificial neural networks leads to different outcomes compared to previous cases. The number of accepted features, as shown in Table 4.4, is fewer than in the previous case, with a total of 13.

ClimateZ still holds the top spot in the ranking, followed by the staff in service (*StaffSrv*) as the second highest-ranked feature in the ranking. It is plausible to state that it is not the presence of a larger number of people in the facility that results in higher consumption. The opposite is more likely true, considering that periods of increased occupancy coinciding with periods of greater healthcare activity and, therefore, more extensive use of the facility systems. In other words, if the regulation of the systems were kept in their standard configuration and all the people present left, the consumption would remain the same. However, it is true that the management strategies (especially for $HVAC$ systems) are designed based on the activities performed within the building, which are directly related to the presence of people in the building. So, this feature can be considered an "indirect" cause of increased consumption, just like the hour of day. In fact, both influence the daily cycle of the activity.

Calendar features related to all time scales remain very important. *Month* is only accepted in *One-Hot encoding* form, as is *Weekday*, while the *Hour* is accepted in both *sine function* form and *ordinal encoding*. Exhausting the calendar features, in addition to *Weekday*, the type of day (*DayTyp*) in *Ordinal encoding* also proves useful for model learning.

Regarding the remaining climate-related features, many of the considerations made in previous sections are valid. Wind speed (WS) is no longer accepted and the information derived by solar radiation and ambient air temperature is included exclusively through the *ClimateZ* feature. The features T_{dp}, p_{amb}, HDD, HR, RH and ρ_{wet}, which provide important information about air humidity, are accepted. As previously observed, air humidity is

important for the performance of indoor air conditioning systems and, consequently, has a significant influence on the energy behaviour of the building.

Comparing the prediction performance of the current case with that of the wrapper method applied to multiple linear regression, it can be seen that the use of artificial neural networks leads to significantly better performance. The increase in prediction accuracy from one step of the wrapper method to the next one is more marked, resulting in significantly higher prediction performance with a smaller number of input features.

4.1.2 Model training

Once the *Feature Selection* phases has been carried out, the obtained best feature subset is used to train the model. The desired model must represent the building energy demand behaviour of a specific year. This is because the model (which represent the first year of data), will be compared with new data about the subsequent years in order to highlight any changes in the energy demand behaviour. In other words, the *ML* model will be used to predict the *expected* energy demand. If any change in the building energy system layout or management strategies take place after the training period, the real energy demand will be different than the expected one, allow the monitoring method to highlight this anomaly. The chosen training period goes from 12^{th} *February* 2020 until 11^{th} *February* 2021.

It is important to remark that a further operation must be done to the dataset. If the dataset used to train the model were also used to evaluate its prediction performance, the model would always be well trained. To ensure the overall reliability of the obtained model, it is important to evaluate its performance on data not previously shown to it. However, it would not be correct to show him data from years after the training period, as these will be used as a dataset to apply the energy monitoring method. In addition, using data after the training period would imply the need to have data for a period longer than one year to correctly carry out the training process, thus delaying the possibility of applying the proposed method. However, a building electricity demand dataset with a time step of 15 minutes will certainly be highly redundant. There will be many samples related to different day types, months, seasons, temperatures and hours. It is therefore possible to randomly subtract a small amount of data from the available training data to be able to test the prediction performance of the obtained model.

In this case, 20% of the available samples in the dataset were randomly extracted, which will therefore not be shown to the model during training, but will instead be used to evaluate its performance. A *test dataset* will then be formed. Once the model is trained, it will be applied to the features contained in the *test dataset*. The values predicted by the model itself will be compared with the real energy consumption data, and then the main evaluation metrics will be evaluated on the test dataset (see Eq. 2.1, 2.2, 2.3, 2.4). All

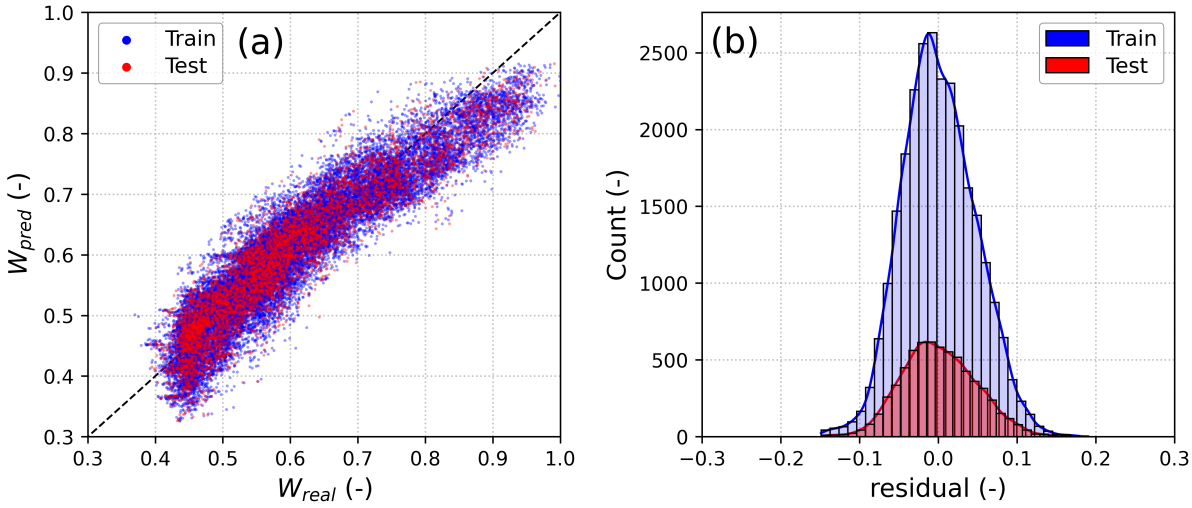

Figure 4.9: Real data VS. prediction (Case 1): (a) Comparison of real and predicted whole building power demand. (b) Residual statistical distribution

the training processes will be analysed separately, analysing and comparing the prediction and the real values. Then, all the results will be compared together in terms of prediction performance and computational cost.

Case 1: Multiple Linear Regression with Filter Method

This section is devoted to analyse the results of the multiple linear regression training (*Case 1*). In particular, this section will present the results of the *MLR* model trained through the filter method-based feature selection. This case represents the simplest and less systematic method among all the analysed cases. The *MLR* allow to realize a simple and interpretable model, which is not able to model complex relationships between features and target. Moreover, the feature selection performed through the filter method must be performed through a proper correlation analysis, requiring a high previous knowledge about the context and the building requirements.

As mentioned above, the first year of data was randomly divided into two datasets. 80% percent of the data constitutes the *training dataset* to train the model. The remaining 20% constitutes the *test dataset* to evaluate the model performance.

Figure 4.9a exploit a scatter plot to compare real and predicted values for both *training dataset* (blue dots) and *test dataset* (red dots). First of all, it can be noted that the residual (differences between real and predicted data) behaviour of the model prediction is comparable for both training and test data. This aspect is better highlighted in Figure 4.9b, which shows the statistical distribution of the prediction residuals. Moreover, Figure 4.9a shows that the points tend to lie over the graph diagonal, a sign that most of the instantaneous power demands have been predicted with acceptable accuracy. Indeed, a perfect prediction leads to having all points located on the graph diagonal. Thus, the

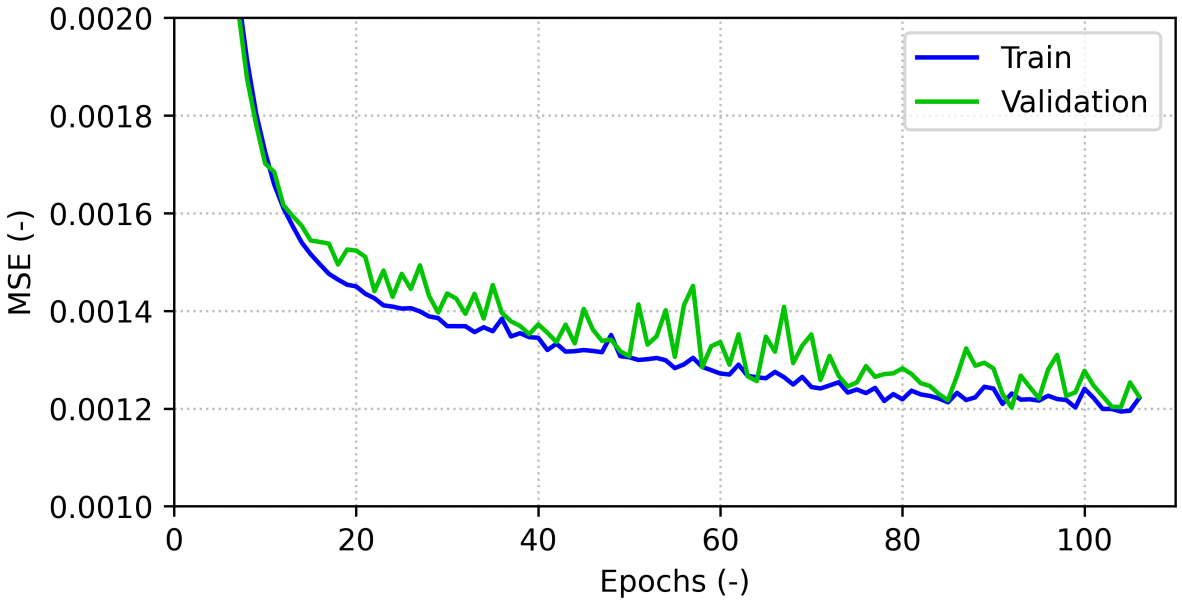

Figure 4.10: Training and validation loss behaviour (Case 2)

scattering of points around the diagonal is due to imperfect model prediction performance. To quantify it, the metrics previously explained have been used. All the calculated metrics are summarized in Table 4.6. The performance of the model is thus summarized by an R^2 value of 0.8555, an indication of the trained model's good performance.

Nevertheless, it can be seen that for low power demand (bottom left of Figure 4.9a), the model tend to increase its dispersion around the diagonal, struggling to predict the target correctly and alternating between overestimates and underestimates. At high power demand (top right of Figure 4.9a) the model systematically underestimate the power demand. To better show these tendencies, Figure 4.23 shows the comparison between real and predicted values for two weeks representative respectively of highest and lowest electricity demand. The first is a summer week (July 2020) and represent a period in which the electricity demand reach its maximum values, with an intensive use of $HVAC$ systems. The second is a winter week (December 2020), where the electricity demand is at its minimum. Looking at the summer week in Figure 4.23, it can be seen that most of the underestimation happen during the afternoon of the workdays. The winter week highlight that during the workdays nights, the model alternates underestimation and overestimation. All of these behaviours lead to the fact that the model cannot perfectly find the relationship between the building consumption and the outdoor air temperature (T_{db}) during periods of very high and very low consumption, while it predicts the consumption trends very well during intermediate periods.

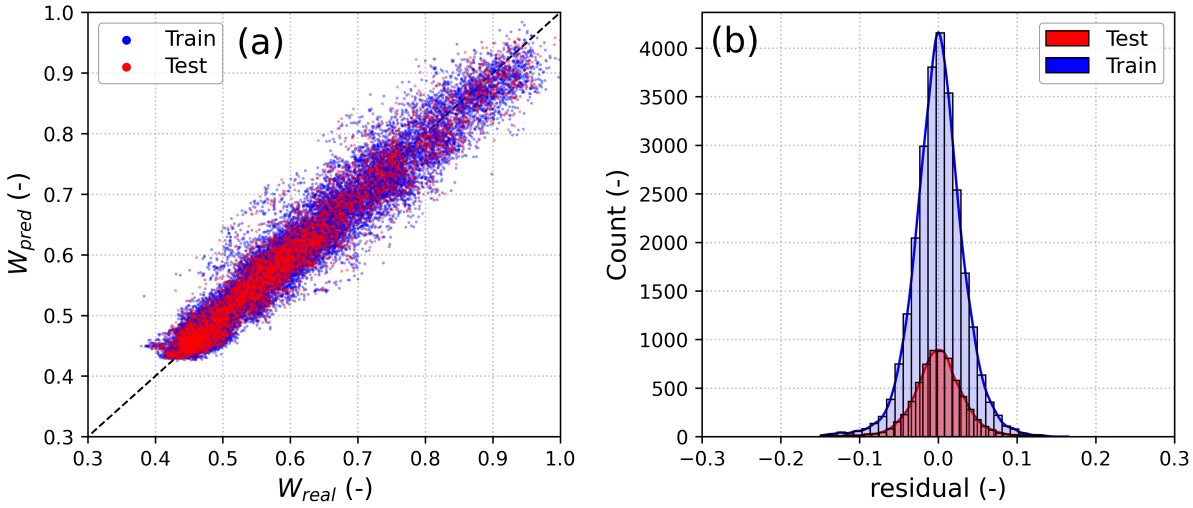

Figure 4.11: Real data VS. prediction (Case 2): (a) Comparison of real and predicted whole building power demand. (b) Residual statistical distribution

Case 2: ANN with filter method

In this section, the results of the implementation of an Artificial Neural Network Multi-Layer Perceptron (*ANN-MLP*) with feature selection using the *filter method* for **Case 2** are presented. As the feature selection is based on the filter method, the number of input features will be 10, just as in the previous case. As previously stated in previous sections, the hyperparameters of the network were selected based on commonly used parameters in literature and reference libraries (e.g., *scikit learn*). The learning rate was set to 0.01, and the chosen activation function was the rectified linear unit (*ReLU*). With regards to the number of neurons in the hidden layer (*HLN*), they were derived from the equation 2.8. Given that some of the accepted input features consist of multiple vectors (specifically, *Month* and *Hour* expressed through sine and cosine functions and the *TimeSlot* expressed through *One-Hot encoding*), the number of nodes in the input layer of the network will be higher than 10, reaching a total of 14 input layer nodes. Consequently, the chosen *HLN* results equal to 15. Figure 4.10 shows that, apart from a fluctuation in validation loss, both loss functions shows a continuous reduction, and the learning process was therefore successfully concluded after *109 Epochs*. It can be noticed that the training and validation loss assumes similar values during the whole training, even if the validation loss tents to have higher oscillation. This is due to the fact that the validation data are not used to adjust the internal parameters of the *ANN* during the training, but only to evaluate the prediction performance at the end of each epoch and, consequently, shows a mildly reduced prediction capabilities of the model on unseen data. This fact is useful to better explain the importance to evaluate the model on new data (test dataset) at the end of the training, but is not useful to evaluate the model validity.

Looking at Figure 4.11a, it can be seen that the predictions, for both training and

test datasets, turn out to be more uniform around the diagonal of the graph, not showing the irregularities for lower and upper extremes of power input that were found in the previous case. The distribution of residuals (Figure 4.11b) is also very similar for the two datasets, further confirming the validity of the trained model. The error evaluation parameters are summarized in Table 4.6. It can be seen a reduction in *RMSE*, *MSE* ans *MAE*, resulting in an increase of R^2 up to 0.9019. Indeed, the *ANN* allow to include non-linearities in the relation which lead from the input to the output, helping to better model complex behaviour like the current one. However, by observing the region in which the dimensionless power demand ranges between 0.4 and 0.5, a systematic prediction defect can be observed. Indeed, for W_{real} lower than 0.5, the model's tendency to settle on a constant predicted value of about 0.42 becomes increasingly evident. This means that, for low power absorbed, the model tends to overestimate consumption, failing to follow the lower power absorption peaks. This behaviour is similar to the one found for the preceding case, even if in this case the error is less marked due to the increase in overall prediction performance. This lead to think that, despite the mathematical model complexity, the used input features are not sufficient to correctly model the whole building electricity demand.

Nevertheless, as illustrated in Figure 4.23, the model for **Case 2** exhibits improved performance in tracking the high and low peaks of the building power demand when compared to that of **Case 1**, even if visible prediction errors at the peaks continue to be observed. In addition, it is still difficult for the model to correctly predict the energy behaviour of the structure on weekends and holidays. In conclusion, as summarized by the determination coefficient ($R^2 = 0.908$, Table 4.6), the model generally has greater prediction capabilities compared to the previous one, confirming the improved performance associated with the use of *ANNs*.

Case 3: ANN with filter method and hyperparameters tuning

To further enhance the prediction performance of the *ANN* model, a grid search approach (see 2.2) has been carried out on three model hyperparameters in order to perform the *Hyperparameters tuning* and, consequently, provide a model trained with the best feature subset and the best model architecture. The tuned hyperparameters are the learning rate (lr), the number of neurons in the hidden layer (HLN) and the activation function. The learning rate (lr) ranges between 0.0001 and 0.01, while the hidden layer neurons (HLN) vary between 8 and 750. Focussing on the hidden layer neurons number, the studied values have been obtained as multiples of the base value, which have been used in the previous case (**Case 2**, see equation 2.8).

The analyzed activation functions are the *Rectifier Linear Unit* ($ReLU$) and the *hyper-*

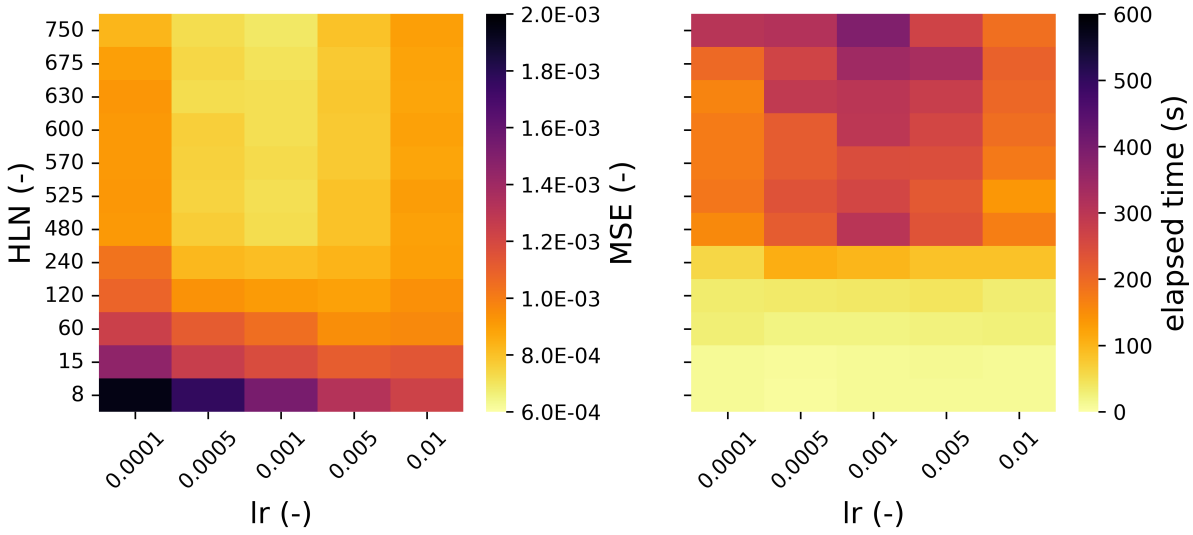

Figure 4.12: Hyperparameter tuning results (Case 3) with "ReLU" activation function

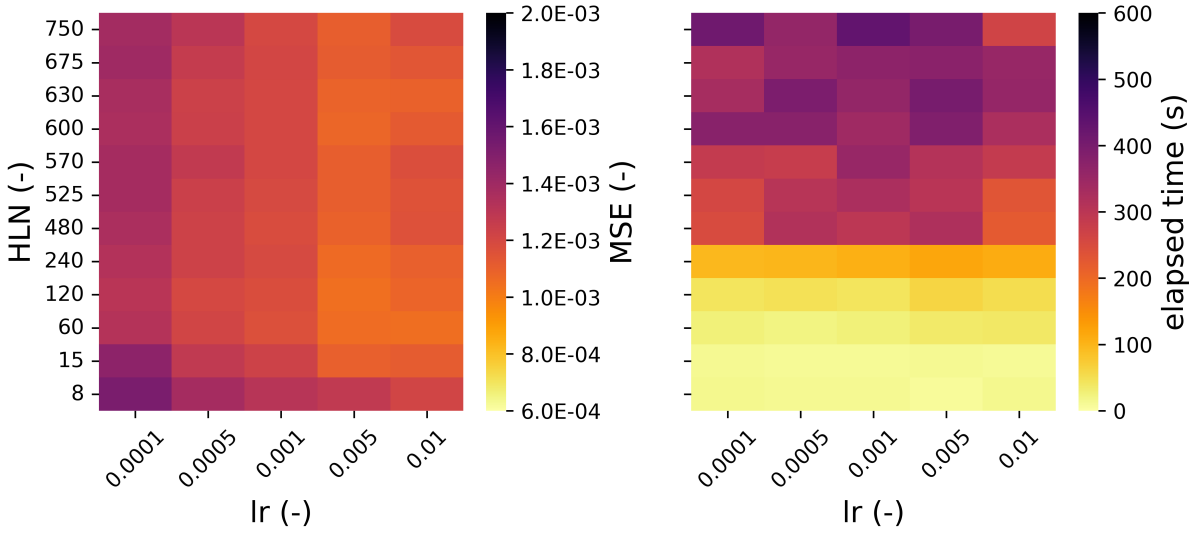

Figure 4.13: Hyperparameter tuning results (Case 3) with "tanh" activation function

bolic tangent (*tanh*), which are the most common activation functions used in regression applications. Once the hyperparameter values has been decided (*Hyperparameters grid*), the grid search approach involves training the model with all possible combinations of these values. The performance of the models have been then evaluated using a validation set, and the best set of hyperparameters is chosen based on the highest performance metric. Figures 4.12 and 4.13 depict the heatmaps of the obtained results, analysing the variation of both *MSE* and the elapsed time of each performed training procedure, divided by activation function.

The results of this study demonstrate that the use of *ReLu* activation functions in training models results in higher prediction performance compared to *tanh* activation function. Furthermore, this is achieved with a general decrease in training time. Additionally, as the number of hidden layer neurons (*HLN*) is increased, the prediction performance of

the *ReLU*-based models also increases. However, this is accompanied by an increase in training time.

A learning rate of 0.001 was found to be optimal for models with 480 or more *HLN*, as it resulted in the highest prediction performance. When using a learning rate of 0.001, an increase in prediction performance can be observed with increasing *HLN* until a value of 480 is reached, after which the mean square error (*MSE*) value does not significantly change. The grid search procedure allow to define the best values of the three investigated hyperparameters, leading to the choice of *ReLU* activation function, $HLN = 480$ and $lr = 0.001$. The selected hyperparameters have been used to perform the final training. The loss and validation loss behaviour during the training are valid, confirming the generality of the trained model (Figure 4.14). Comparing their behaviour with the ones found for the preceding case, it can be seen that the values of the two functions present more marked differences. This mean that the prediction performance of the obtained model on data that do not partecipate directly to the training have reach its maximum, while the prediction capabilities on the training data could be further increase. Nevertheless, increasing the performance on the training data without enhance the ones found on new data is not useful. Moreover, with the training process after this point could lead to a change of the validation loss behaviour, which could start to increase despite a further decrease of the training loss. This phenomena is the already mentioned overfitting, which is the main reason behind the implementation of training stopping criteria based on the loss and validation loss function behaviours (*Early Stopping*).

The results analysis shown in Figure 4.15 highlight an increase in the prediction performance in comparison to the previous cases, achieving an R^2=0.932 (Table 4.6).

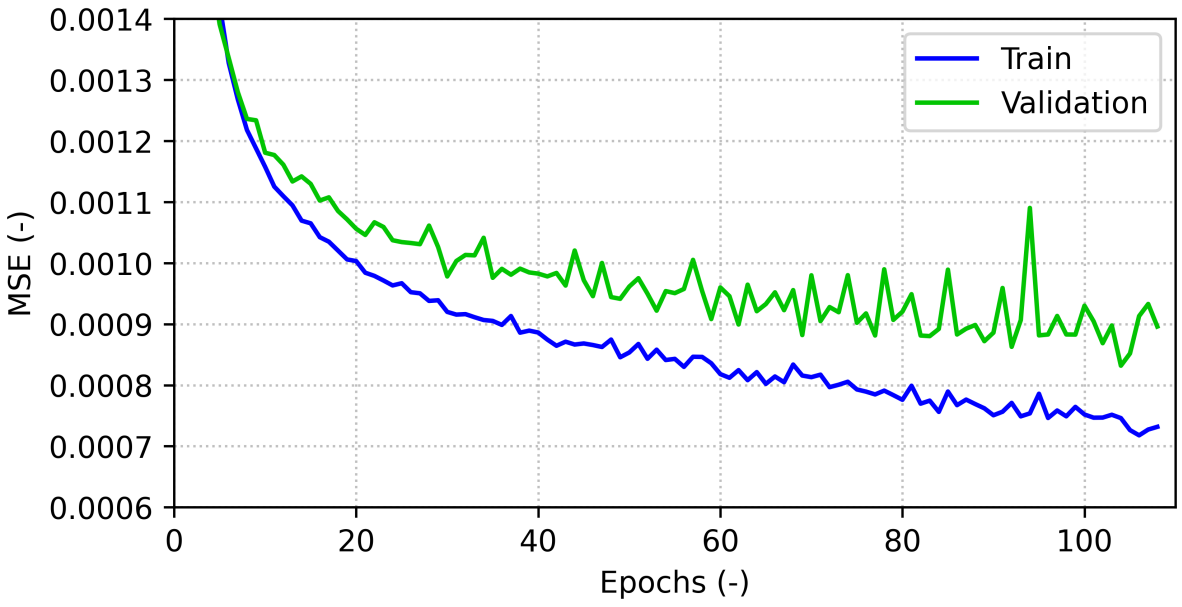

Figure 4.14: Training and validation loss behaviour (Case 3)

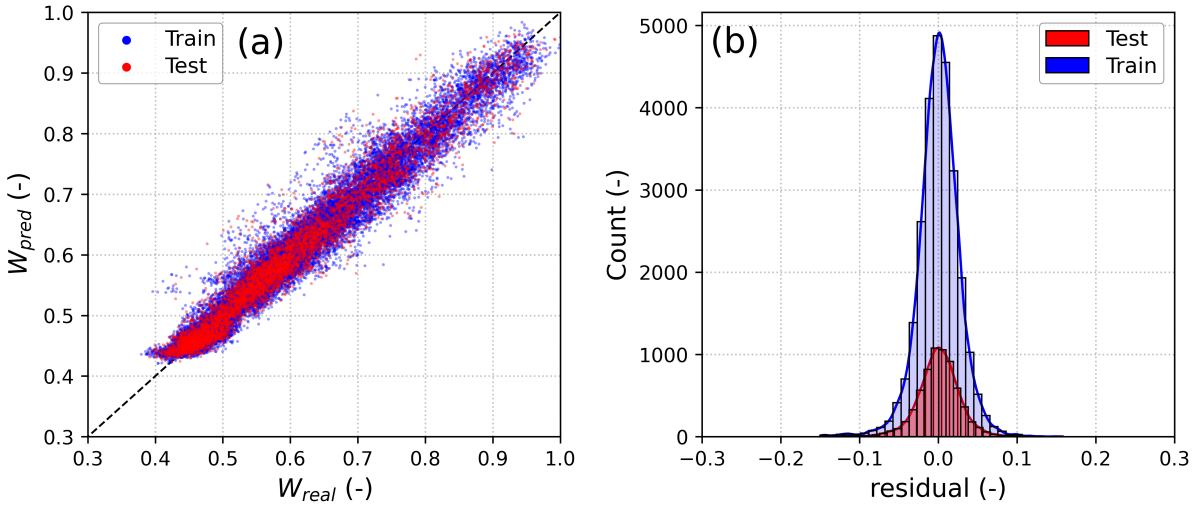

Figure 4.15: Real data VS. prediction (Case 3): (a) Comparison of real and predicted whole building power demand. (b) Residual statistical distribution

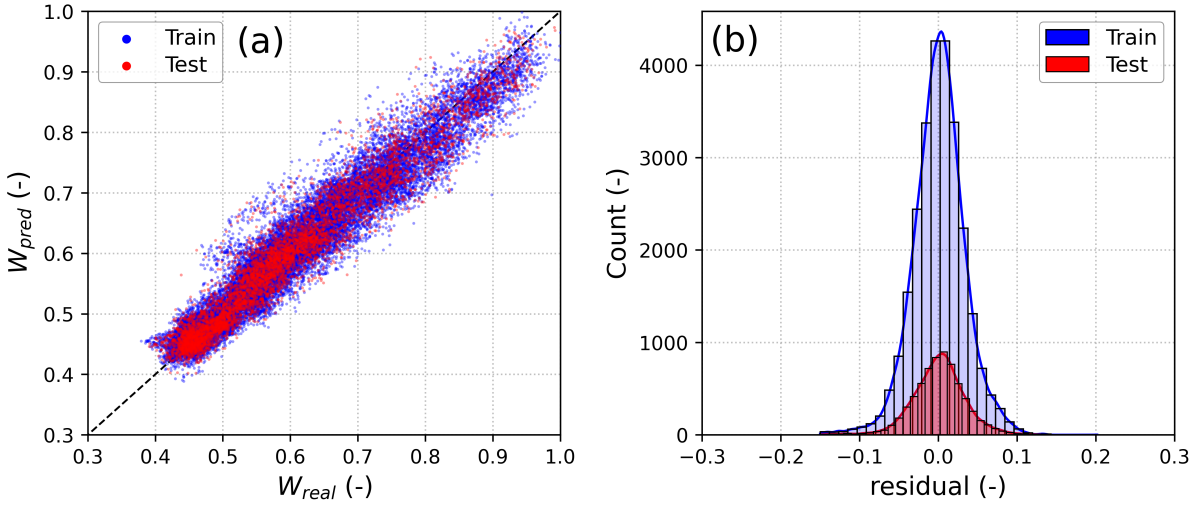

Figure 4.16: Real data VS. prediction (Case 4): (a) Comparison of real and predicted whole building power demand. (b) Residual statistical distribution

despite that, a mild overestimation of low consumption peaks continue to be observed. Figure 4.23 shows that the obtained model is able to better predict the energy consumption at higher and lower peaks, achieving also an increase in prediction performance during the weekends.

Case 4: MLR with Wrapper method

As already seen in section 4.1.1, the wrapper method applied to the Multiple Linear Regression method allow to accept a higher number of features for the final training (25), allowing for a significant increase in the model prediction performance, maintaining the advantages of *MLR*. In particular, since the *MLR* model presents a low computational cost, its increase due to the application of the wrapper method (multiple trainings) results

markedly affordable. The systematic nature of this feature selection method leads to a sensible decrease in context knowledge required for the feature selection. Differently from the *filter method*, the features are selected by directly applying the *ML* model on the various feature subsets. Consequently, this procedure allows to decrease the possibilities of excluding useful features or including harmful ones.

Comparing the scatter plots of **Case 1** (Figure 4.9a) and **Case 4** (Figure 4.16a), it is clear that there is an improvement in prediction performance. The errors at the peaks are significantly reduced, as well as the scatter around the diagonal of the graphs. In terms of metrics, the improvement is reflected in an increase in the R^2 from 0.856 in **Case 1** to 0.919 in **Case 4**, a slightly better performance than that obtained by applying the neural network to the features accepted through the filter method (**Case 2**). This fact makes it evident the importance of the feature selection phase in the overall learning process. Indeed, by simply applying a systematic procedure to the feature selection, it was possible to obtain a much simpler and computationally less onerous model with the available data, achieving prediction performance comparable with an *ANN* with a non-optimized input data. The computational cost that is still higher than the latter, requiring 76.4 s to complete the process (**Case 2** require almost 26 seconds less to complete the training).

However, observing Figure 4.23, it can be seen that multiple linear regression, although with significant average performance, presents more accentuated problems compared to neural networks in modelling consumption during weekends, therefore struggling to distinguish the day typology. In any case, regardless of the method of feature selection, the inability of multiple linear regression models to capture non-linear relationships can only be overcome by changing the type of predictive algorithm. The non-linearity introduced by artificial neural networks through the *activation function* can in no way be reproduced by keeping the model simple. This is the greatest drawback observable for linear regression models. Nevertheless, it is clear that in applications where reducing computational cost is a key factor, this method can be a valid way to obtain a systematic monitoring method, giving up better performance in favour of low computational cost and model interpretability, and completely eliminating the need for prior knowledge in the field of artificial neural networks.

Case 5: ANN with Wrapper method

In this section, we discuss the fifth case of training a neural network using the accepted features obtained through the wrapper method (which we discussed in Section 4.1.1). The results obtained from this training are similar to those obtained in previous cases, with some notable differences. The behavior of loss and validation loss, as shown in Figure 4.17, is similar to previous cases. Additionally, the distributions of the residuals on the

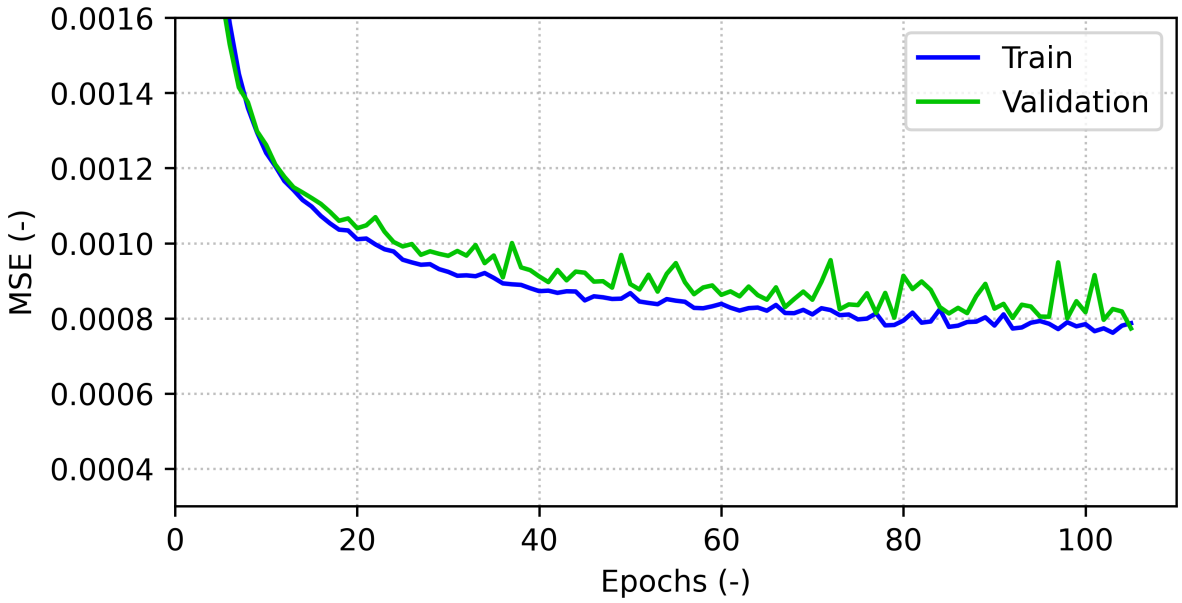

Figure 4.17: Training and validation loss behaviour (Case 5)

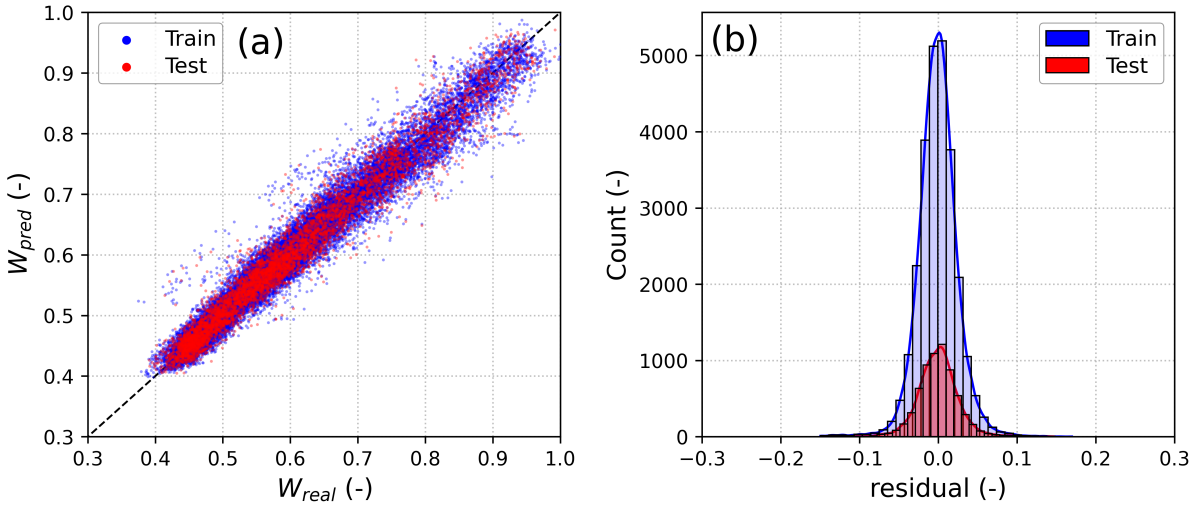

Figure 4.18: Real data VS. prediction (Case 5): (a) Comparison of real and predicted whole building power demand. (b) Residual statistical distribution

training and test datasets, as shown in Figure 4.18, are also comparable to previous cases. However, the prediction performance of the neural network in this case is superior, with an R^2 value of 0.938. However, this increase in prediction performance comes at the cost of a significant increase in computational cost. Specifically, the use of the wrapper method as a feature selection procedure for the neural network increases the computational cost from 76.4 seconds in Case 4 to 8.91 hours in Case 5. Depending on the application context, this increase in computational cost may make the use of this method impractical. Therefore, while the wrapper method can improve the prediction performance of the neural network, it comes with a significant increase in computational cost that must be considered in practical applications.

Case 6: ANN with Wrapper method and hyperparameters tuning

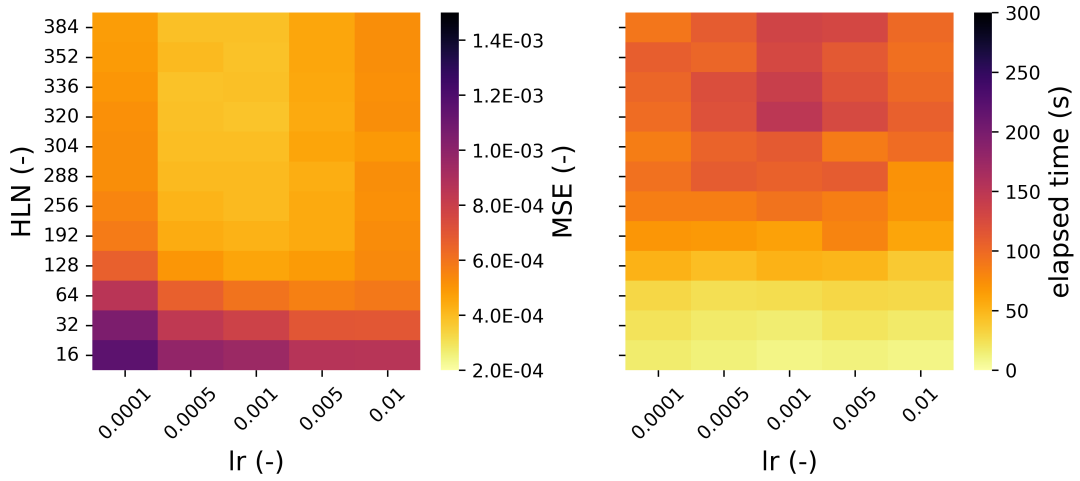

Figure 4.19: Hyperparameter tuning results (Case 6) with "ReLU" activation function

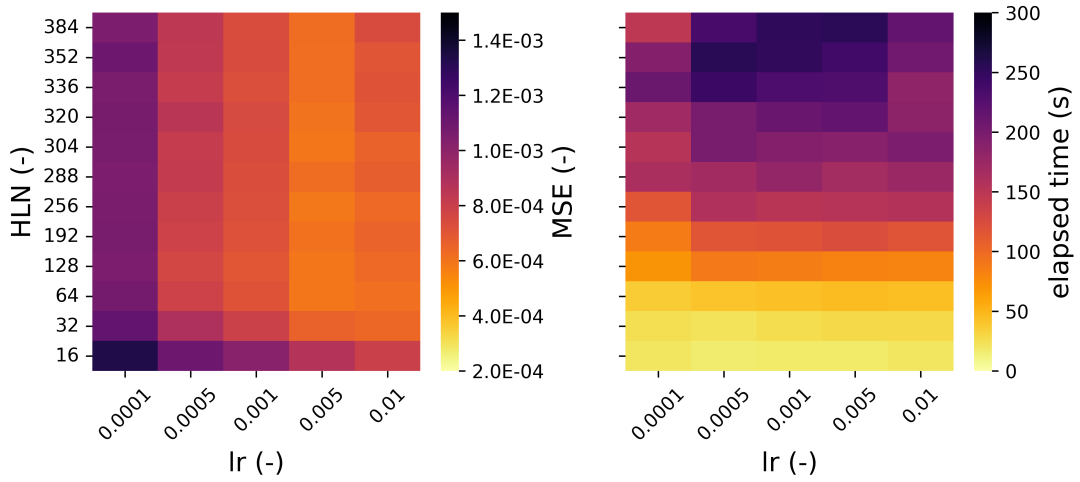

Figure 4.20: Hyperparameter tuning results (Case 6) with "tanh" activation function

The last case to be analyzed consists of feature selection using the wrapper method and subsequent optimization of hyperparameters through grid-search. The hyperparameter grid was chosen using the same logic as in **Case 3**. Despite the selected features being numerically superior to those identified in **Case 3**, in this case the optimal value of neurons in the hidden layer is significantly lower.

This can be interpreted in the following way: the wrapper method allows for a feature selection based on their actual contribution to the training of a model, something that does not occur if the exploited method is based on correlation coefficients (*Filter method*). As a result, the features selected in this case are extremely effective in helping the learning process, allowing for excellent results with a lower model complexity.

As found for **Case 3**, the best activation function results the *ReLU*, while the *tanh* cause a decrease in prediction performance and an increase in computation time required

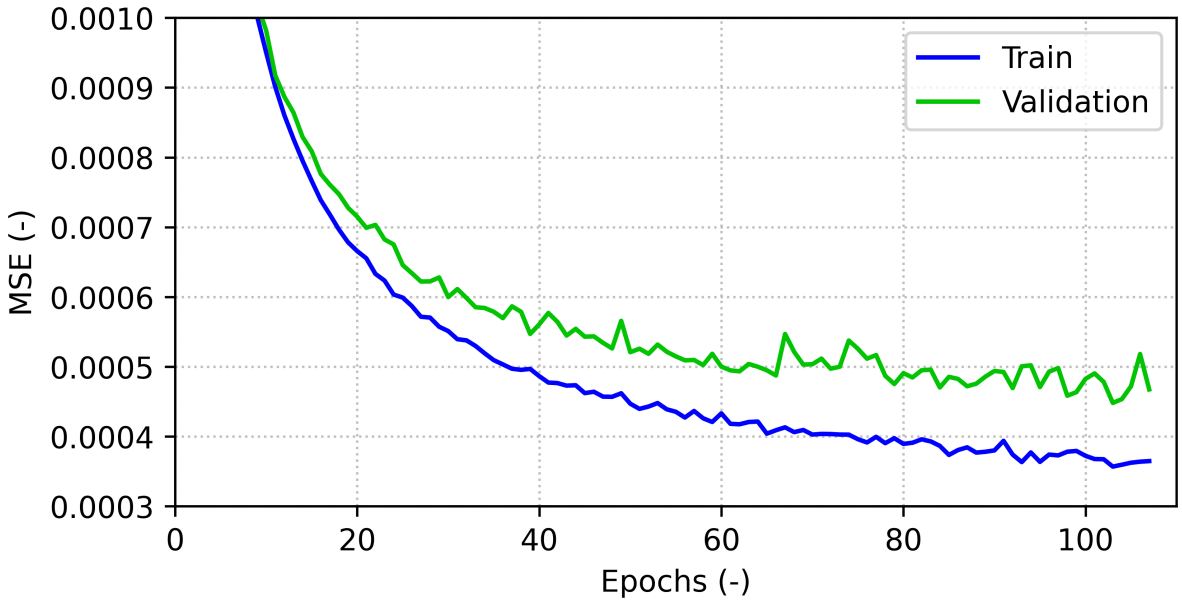

Figure 4.21: Training and validation loss behaviour (Case 6)

for each training. The optimal hyperparameter set is composed of a number of neurons in hidden layer equal to 320, while the best learning rate is confirmed to be 0.001.

The validation of the model is confirmed by the behaviour of the loss and validation loss parameters (Figures 4.21), which fall to values that are markedly lower than those seen in the previous cases, showing behaviour that is comparable between the training set and the validation set, and thus confirming the absence of overfitting. Looking at the figure 4.22, it is possible to notice a distribution of residuals that is significantly more concentrated around the zero value, a sign that the prediction performance is superior to the previous cases. The average metrics presented in table 4.6 confirm the result, raising the value of R^2 to a value of 0.954. Finally, observing figure 4.23, the model's greater ability to predict

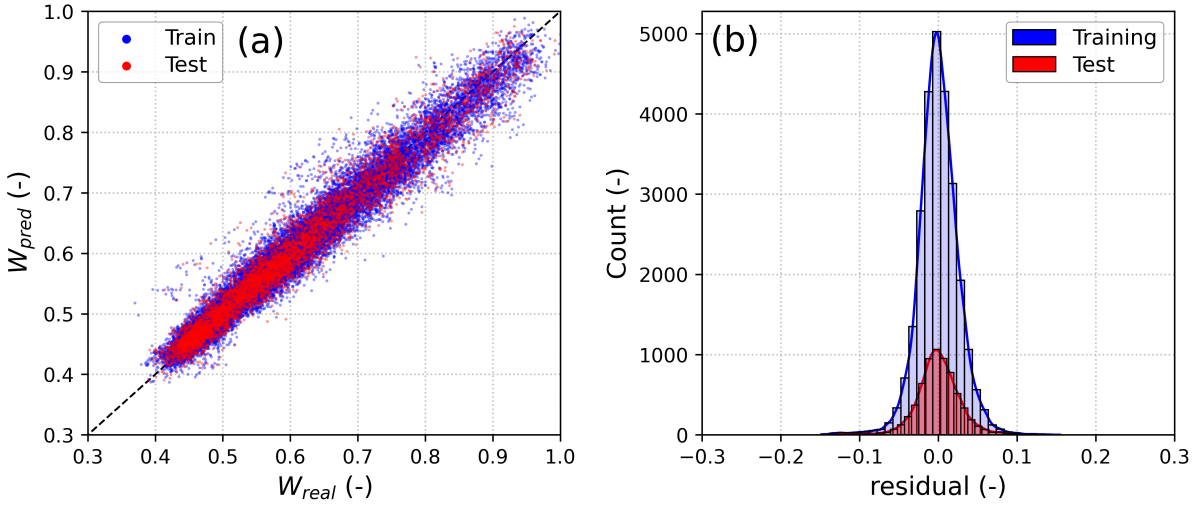

Figure 4.22: Real data VS. prediction (Case 6): (a) Comparison of real and predicted whole building power demand. (b) Residual statistical distribution

the structure's consumption is apparent, showing a curve almost completely superimposed on actual consumption during both the summer and winter periods. The discernment problems between weekdays and weekends have almost completely disappeared, like the prediction problems for high and low peaks during the daily cycle.

Discussions

Table 4.5: ANNs hyperparameters summary

Hyperparameter	*Activation function*	*lr*	*Input features*	*HLN*
Case 2	ReLU	0.01	14	15
Case 3	ReLU	0.001	14	480
Case 5	ReLU	0.01	13	32
Case 6	ReLU	0.001	13	320

Table 4.5 summarizes the selected hyperparameters for all cases involving neural networks. The main observation that can be made is that for this particular dataset, the learning rate (lr) suggested by the literature appears to be overestimated. The models seem to learn better with a lower learning rate (0.001), which allows for more regular learning and facilitates the convergence of the loss function. However, reducing the learning rate appears to benefit from a significant increase in the number of neurons in the hidden layer (HLN). A higher HLN effectively increases the number of linear combinations between the input features within the model, forming more complex relationships that can better adjust and predict consumption more effectively. As regards the *activation function*, the rectified linear unit ($ReLU$) often outperforms the hyperbolic tangent ($tanh$) and is in fact the most widely used function for regression applications.

Table 4.6 resume the performance metrics of the final models obtained by the application of each case. As previously mentioned, the prediction performance obtained for **Case 4** ($R^2 = 0.919$), that is the MLR with the wrapper method as a feature selection procedure, ensures clearly superior performance compared to the corresponding case that uses the filter method instead (**Case 1**, $R^2 = 0.856$). It also shows comparable performance to the other cases that use the filter method to perform feature selection (**Case 2** and **Case 3**), with a computational cost that is comparable to the first and clearly lower than the second. The advantage of **Case 4**, furthermore, is that it uses a systematic procedure for selecting input data, avoiding to exclude useful features from the final dataset and significantly reducing the amount of experience and knowledge required by the user to

Table 4.6: Prediction performance metrics

Metric	**Case 1**	**Case 2**	**Case 3**	**Case 4**	**Case 5**	**Case 6**
RMSE	0.0494	0.0390	0.0338	0.0371	0.0321	0.0279
MSE	0.00244	0.00150	0.000115	0.00138	0.00104	0.00078
MAE	0.0386	0.0280	0.0233	0.0261	0.0222	0.0193
R^2	0.856	0.908	0.932	0.919	0.938	0.954
elapsed time	0.007 s	40.5 s	2.5 h	76.4 s	8.9 h	10.4 h

create the final model. In an industrial context, therefore, the use of this method could be considered sufficient to achieve a solid and reliable monitoring tool.

From the analysis of **Cases 5** and **6**, which use artificial neural network as a predictive model, an additional conclusion can be drawn. The improvement in performance is evident only in **Case 6** ($R^2 = 0.954$), where the feature selection procedure using the wrapper method is followed by the hyperparameter tuning process. The total execution time of the *wrapper method* is predominant compared to all other phases of the process, and it is therefore unnecessary to perform this process on a neural network model without also performing a fine tuning procedure (**Case 5**). In terms of prediction performance, the latter case is clearly ahead of the others, providing in conclusion a model trained with the best subset of features and with optimal hyperparameters, and using a procedure that reduces the user's knowledge required to the mere processing of features through the feature engineering procedure.

It is clear that the distribution of the total computation time required for the process is closely dependent on the overall number of features to be analysed and the number of values entered into the hyperparameter grid that defines the fine-tuning procedure. It is also true that a reduction in the number of features or the number of combination in the hyperparameter grid aimed to reduce computational cost, would likely result in a reduction in the performance of the final model, making it increasingly similar to those obtained with less burdensome procedures, and consequently reducing the benefits provided by the "complete" method. As a result, the choice of method to use must certainly be a compromise between the need for model accuracy and the available computational resources. The analysis carried out in this section therefore constitutes a useful map to guide the user in choosing the most appropriate method based on the specific context of application.

The burden of computational cost is definitely an aspect that allows for further future developments. Contemporary to the activities described in this work, an novel learning

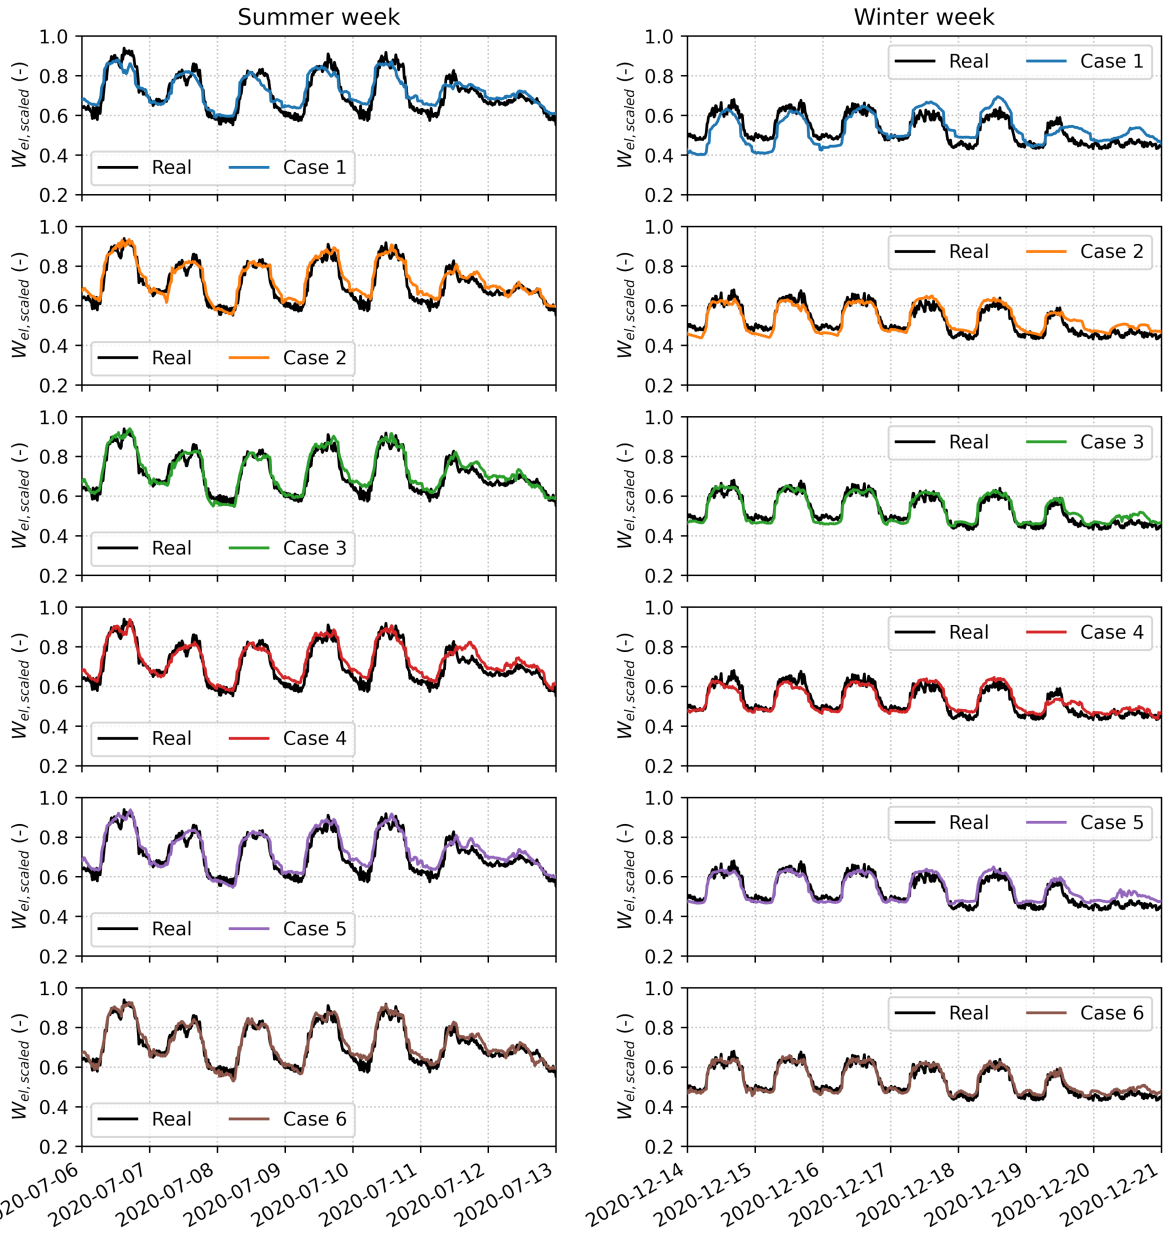

Figure 4.23: (a) Comparison of real and predicted whole building power demand. (b) Residual statistical distribution

algorithm that could provide significant advantages has been developed. The algorithm involves not using the entire available dataset at each epoch of training as the traditional algorithms, but rather randomly extracting a small percentage (*batch size*) and evaluating prediction error on it. This significantly reduces the number of elementary mathematical operations performed for each iteration of the learning process. The batch size evaluation is automatically chosen by the developed algorithm, which has the ability to vary it during the learning process according to an appropriate control logic. In addition, the value of the learning rate, a crucial parameter for correctly perform the training process and often difficult to choose, is automatically calculated based on the prediction error calculated at the end of each iteration.

These two aspect allow for:

- Reduce the computational cost of the feature selection phase thanks to the dynamic batch size and learning rate adjustments, leading to decrease the training time of each single training. Since several training processes are involved in this phases, even a small reduction in computational cost lead to non-negligible reduction of the overall elapsed time
- Removing the learning rate from the tunable hyperparameters allow for a significative reduction of the number of hyperparameters combination considered in the hyperparameter grid of the fine tuning using grid search approach. This, combined with the previously described effect, allow for markedly reduce the computational effort of the hyperparameter fine tuning phases.

This algorithm is better presented in appendix B and, although it is still under development, it shows promising results.

4.2 HVAC system energy prediction

The previous sections highlighted that the best results in terms of model prediction performance have been obtained following the **Case 6** procedure. This includes the exploitation of *Wrapper method* for the feature selection process, followed by the grid search method for the hyperparameters tuning. The comparison between model prediction performance and computational cost have been reported in Table 4.6. The choice of the method must be taken as a compromise between these two factors. It depends, in conclusion, by the requirements in terms of prediction performance and by the available computational resources of the specific application. However, it should not be forgotten that a reduction in computational cost entails an increase in the knowledge and experience required by the user, especially moving towards the *filter method*.

This section will be devoted to the application of the best training procedure resulting from the previous analyses (*Case 6*) to further data about the electricity demand of the healthcare facility under analysis. In particular, exploiting the measuring system installed within the building power center (see Section 3.3), more detailed electricity demand data regarding the *HVAC* systems have been acquired referring to the same period of the previously analysed data about the whole building energy demand.

As already described in the "Test case overview" section (3.3), three data series were analysed. In particular:

- **Refrigerators**: This measure include the electricity demand of all the three refrigerator units installed within the building, which fulfil all the requirements in terms of chilled water.
- **Air Handling Units**: This voice includes all the *32* air handling units installed within the building. Unfortunately, only the aggregated consumption of all the *AHU* were available.
- **Pumps**: This voice include all the pumps used to carry out the water circulation. Upon analysis of the system documentations, it can be seen that the demand is mainly constituted by the primary circuits of the refrigerators and to the pumps which lead the chilled water to the *AHU*.

Due to the high degree of aggregation of the data, it was only possible to study these three important parts of the overall consumption in aggregate form. it is therefore obvious that the installation of a larger number of sensors within the building would allow a more detailed monitoring of individual components. However, a more detailed monitoring of individual components may require more input data. In other words, if the objective were to monitor a single pump, it is highly probable that climate and calendar data would not

be sufficient. It would be necessary to have histories of the setpoints and control logics set, activation schedules of the systems connected to the individual component, and so on.

In other words, the degree of details for the monitoring methods must be a compromise between the necessity to monitor detailed system components and the capability to provide further operative data.

However, the aim of this final section is to apply the method to the three parts listed above. In fact, by following the same steps described in the previous section, it is possible to create a monitoring method for each of these three items from scratch. This highlights the high degree of reproducibility of the method developed, which can then be applied to completely different contexts, with the only requirement of having sufficient data available.

4.2.1 Refrigerators

Correlation analysis

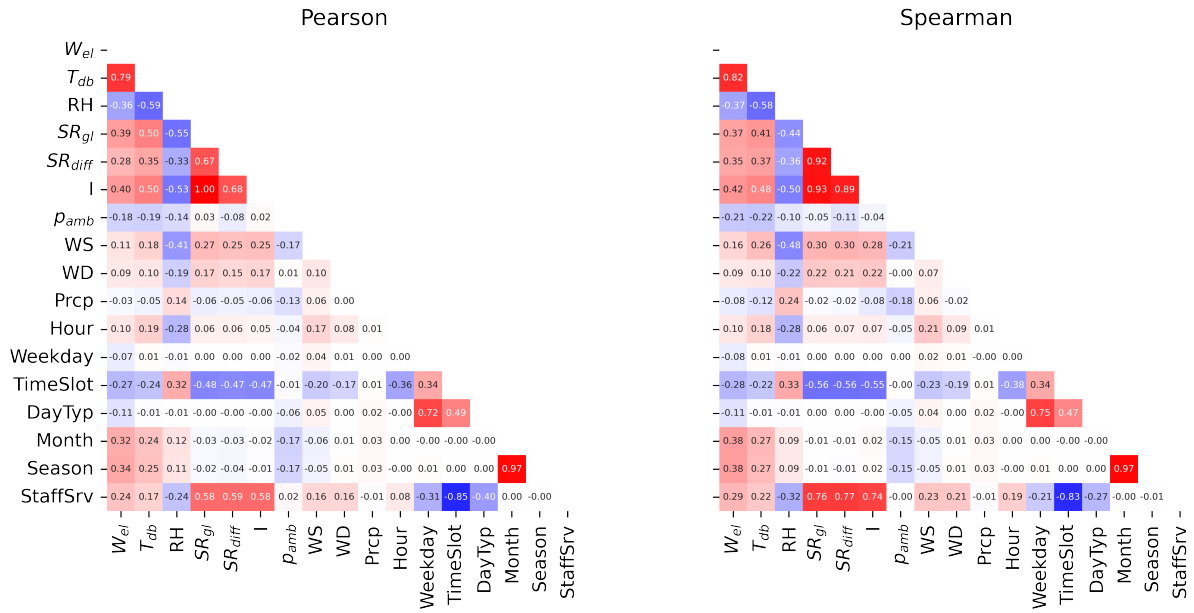

Figure 4.24: Correlation matrices for Refrigerators electricity demand

Although correlation coefficients are not useful for the application of the method, studying their values nevertheless helps to better understand which aspects most influence consumption. Figure 4.24 represents the correlation matrices of the "basic" available feature. As expected, the most correlated feature is the air dry bulb temperature T_{db} ($r = 0.79$, $\rho = 0.82$). Figure 4.25 is a scatter plot in which each point represents a sample of the available dataset. In particular, the outdoor air dry bulb temperature have been compared with the instantaneous power demand. It can be immediately noticed that there is a nearly quadratic relation between the two data series. Indeed, the increase in the refrigerators power demand in correspondence to an increase of the air temperature

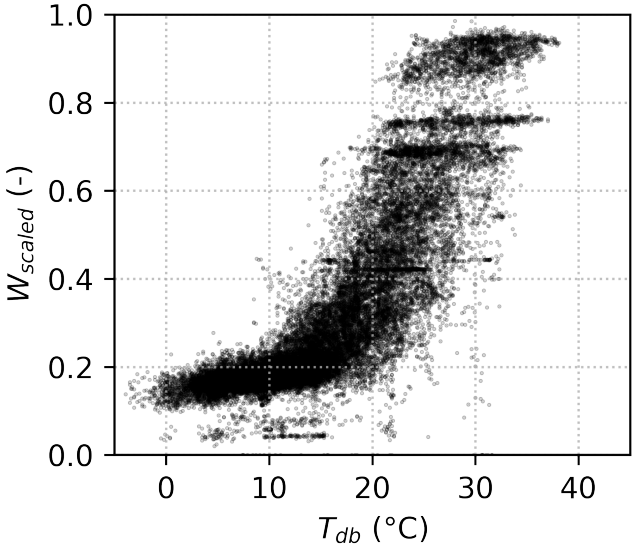

Figure 4.25: Dry bulb outdoor air temperature VS refrigerators power demand

become more marked with higher temperatures, corresponding to the summer period and, consequently, with higher exploitation of the air conditioning systems.

It is interesting to notice that the quadratic relation between these two variables is also highlighted by the correlations, leading to a spearman coefficient higher than the corresponding pearson's one. Indeed, while pearson is usefull to quantify *linear* correlations, spearman coefficient is made to highlight also non-linear correlations (see section 2.1.2).

Other climate features highly correlated with the power demand are the solar radiation-related ones (SR_{gl} and SR_{diff}). Indeed, the air conditioning needs are influenced by the increase of the indoor air temperature caused by high exposition to solar radiation of the external building surfaces (especially the windowed ones). Also the air humidity (RH) play a relevant role in the refrigerators consumptions, given that their performance are related to the performance of the evaporative cooling towers. As regards the calendar features, the seasonal and weekly time scales results the most correlated, with *Month*, *Season* and *TimeSlot* features presenting an intermediate correlation coefficient. Also the *StaffSrv* feature, which represent the number of employees present within the building in a given time step, present intermediate correlation. This fact can be in part related with the daily activities cycle, and not directly with the number of people.

Feature Selection

To carry out the feature selection exploiting the *wrapper method* applied to an *ANN*, the same data elaboration previously done for the whole building energy demand have been carried out, performing the same feature engineering processes. Consequently, the input features dataset is the same already used in previous cases. The starting model architecture and hyperparameters used to carry out the wrapper method feature selection

Table 4.7: Accepted features (Refrigerators)

Rank	**Feature**	R^2	**MSE**	**RMSE**	**MAE**
1	*ClimateZ*	0.8560	0.0075304	0.086778	0.056539
2	*Month* (One-Hot)	0.8925	0.0056222	0.074981	0.049158
3	*Weekday* (One-Hot)	0.9158	0.0044022	0.066349	0.042259
4	*Hour* (Sine functions)	0.9443	0.0029137	0.053978	0.033869
5	T_{dp}	0.9595	0.0021208	0.046052	0.029102
6	*HDD*	0.9648	0.0018413	0.042911	0.027915
7	*Season* (One-Hot)	0.9650	0.0018305	0.042785	0.027661
8	p_{amb}	0.9672	0.0017132	0.041391	0.026975
9	h_{wet}	0.9682	0.0016615	0.040762	0.026721
10	*Month* (Ordinal)	0.9685	0.0016458	0.040569	0.026428
11	*WD*	0.9681	0.0016702	0.040868	0.026607
12	*Hour* (One-Hot)	0.9680	0.0016729	0.040901	0.026834
13	*StaffSrv*	0.9684	0.0016511	0.040634	0.026877
14	*DayTyp* (One-Hot)	0.9687	0.0016365	0.040454	0.026467
15	ρ_{wet}	0.9697	0.0015865	0.039831	0.026275
16	T_{wb}	0.9697	0.0015960	0.039825	0.025924
17	*Month* (Sine functions)	0.9697	0.0015857	0.039821	0.025903
18	*HR*	0.9698	0.0015853	0.039817	0.025888

are chosen using literature-suggested hyperparameters. The number of neuron in the hidden layer is set to 87, the learning rate is set to 0.01 and the selected activation function is the rectifier unit ($ReLU$).

Table 4.7 shows the feature ranking resulting from the application of the feature selection process. Some similarities can be found between the features ranking and the information obtained through the exploratory correlation analysis. Indeed, the most important feature is *ClimateZ*, which is a combination of solar radiation, relative humidity and dry bulb temperature. These are the three feature which have higher correlations with the power demand of the refrigerators. Also the Heating Degree Days feature (HDD) results important for the model training process, a sign that also the difference between outdoor and indoor air temperature play a fundamental role in the refrigerators electricity demand. This information was expected, since the refrigerators are mainly exploited for

air conditioning purposes. The wrapper method shows that their representation with the linear combination allowed by *CLimateZ* feature results more efficient aiming to enhance the automatic learning process than using the three features individually.

The following three positions are taken by *Month*, *Weekday* and *Hour*, representing the seasonal, weekly and daily cycles respectively. The seasonal time scale variation of the refrigerators power demand is mainly influenced by the air cooling requirements which are magnified during the summer due to the increase of the outdoor air temperature and solar radiation. Weekly and daily time scale, are instead more influenced by the activity intensity, which decrease during the low activity periods (nights and weekends). As regards the seasonal time-scale, also *Season* and *Month* (in Ordinal encoding representation) are useful to increase the prediction performance.

Moreover, three parameters which help to consider the effect of the evaporative towers result to be important to predict the electricity demand of the refrigerators: namely dew point temperature (T_{dp}), ambient pressure (p_{amb}) and moist air enthalpy (h_{wet}). As already said, the outdoor air humidity can influence the cooling towers performance and, consequently, the evaporator water inlet temperature. The refrigerators chillers *COP* is indeed related to the coolant condensing pressure [103], which is strictly related to the water temperature coming from the cooling towers. The maximum chiller performance can be achieved by maintaining the nominal coolant condensing pressure. A decrease in the cooling tower performance can negatively affect the chiller *COP* and, consequently, the required electrical power demand to provide the same amount of cooling power to the primary water.

Is important to notice that the increase in the prediction performance occurs until the feature *Month (Ordinal)* is considered. Then, the performance start to decrease when Wind Direction (*WD*) and *Hour (Sine functions)* are considered. It would be possible to think that this is the time to stop accepting input features. However, proceeding with the analysis, it can be seen that performance starts to increase again for the next 8 features, until it reaches its maximum value at the input of the Humidity Ratio (*HR*). From this it can be deduced that, although it is possible to stop at feature number 10 in order to make the model lighter (which would require less input data), the combination of the following features which on their own would not bring an advantage in terms of model accuracy, when considered together can lead to a further increase in model prediction performance.

Consequently, wishing to test the best possible model, 18 features were considered (as shown in the Table 4.7), which were used to proceed with the analysis.

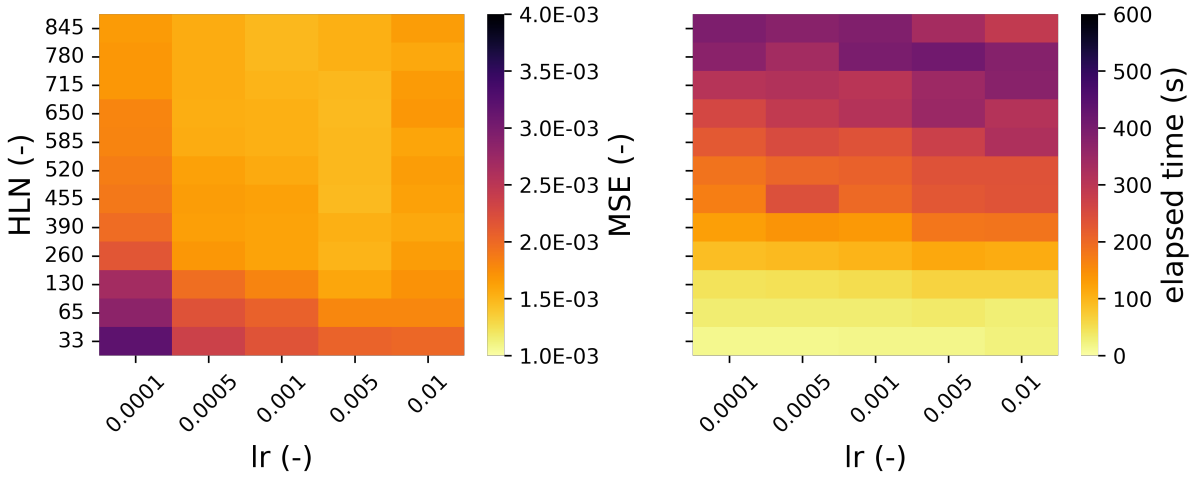

Figure 4.26: Hyperparameter tuning results with "ReLU" activation function for Refrigerators model

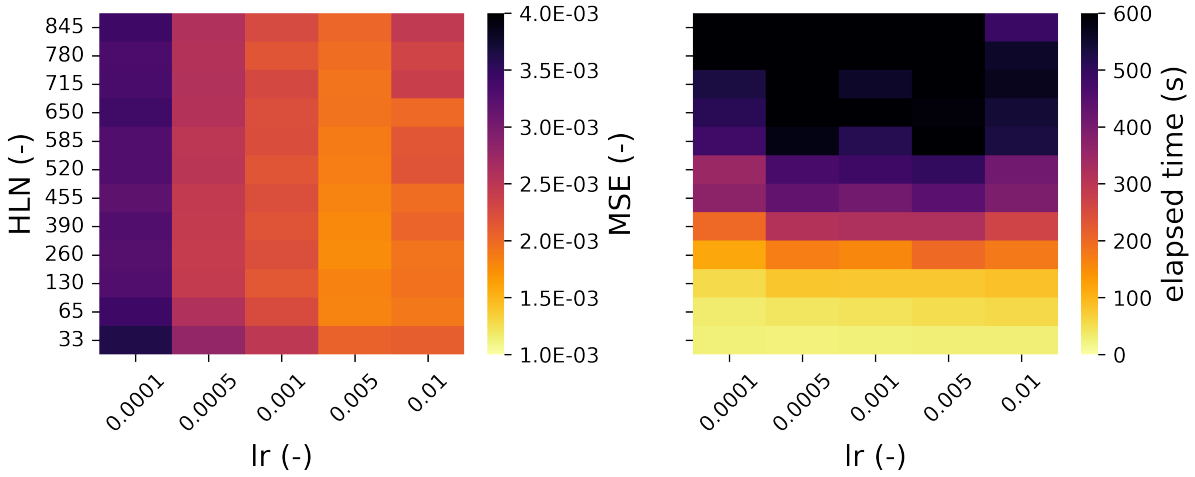

Figure 4.27: Hyperparameter tuning results with "tanh" activation function for Refrigerators model

Hyperparameters Tuning

Once the accepted features have been defined, it is possible to proceed with the hyperparameter tuning phase. As already done for the whole building energy demand, three hyperparameters was selected for the optimization procedure. Figure 4.26 and 4.27 show the model performance and elapsed time heatmaps for all the possible combinations of the selected hyperparamenters. The hyperbolic tangent results not convenient, leading to worst prediction performance and increasing the time needed to perform the trainings. Indeed, the *ReLU* lead to a general increase of performance and a decrease of the needed computational time. Looking at Figure 4.26, it is possible to see that the better performance are obtained with the learning rate set to 0.005, a lower value in comparison to the one obtained for the whole building energy demand ($lr = 0.01$), confirming the importance of performing the a hyperparameter tuning specific for the used dataset.

Concluding, the best performance were obtained with the $ReLU$ activation function, $lr = 0.005$ and a number of neurons in the hidden layer equal to 455, a combination which allow to obtain good prediction performance and to contain the computational cost (in comparison to higher number of neuron in the hidden layer).

Final model training

At this point, the best feature subset has been defined through the feature selection procedure, while the best set of hyperparameters has been obtained at the conclusion of the hyperparameters tuning process through the grid-search approach. It is now possible to proceed to the final phase of model developing, using the obtained results to perform the final training. As it can be seen from Figure 4.28, which represents the trend of the loss function on the training and validation datasets during the training process, both curves show good regularity throughout the process. There are clear generalization limitations, as the reduction of the loss function on the training dataset shows a reduction significantly lower than what happens for the validation dataset. However, this aspect is not of great interest, as it does not compromise the validity of the training process. Being both curves monotonically decreasing (on average) it is possible to exclude the possibility of overfitting.

The training concludes after 108 iterations, reaching an excellent R^2 of 0.97 on the test dataset, which is composed of randomly extracted portion of data not shown to the model (Table 4.8 reports all the computed metrics).

Moving on to the residual analysis, in Figure 4.29a it is possible to observe the scatter plot that compares actual and predicted values for both the training dataset (data already "seen" by the model during training) and the test dataset. The graph tends to show a concentration of points in certain areas. This is due to the fact that refrigeration groups,

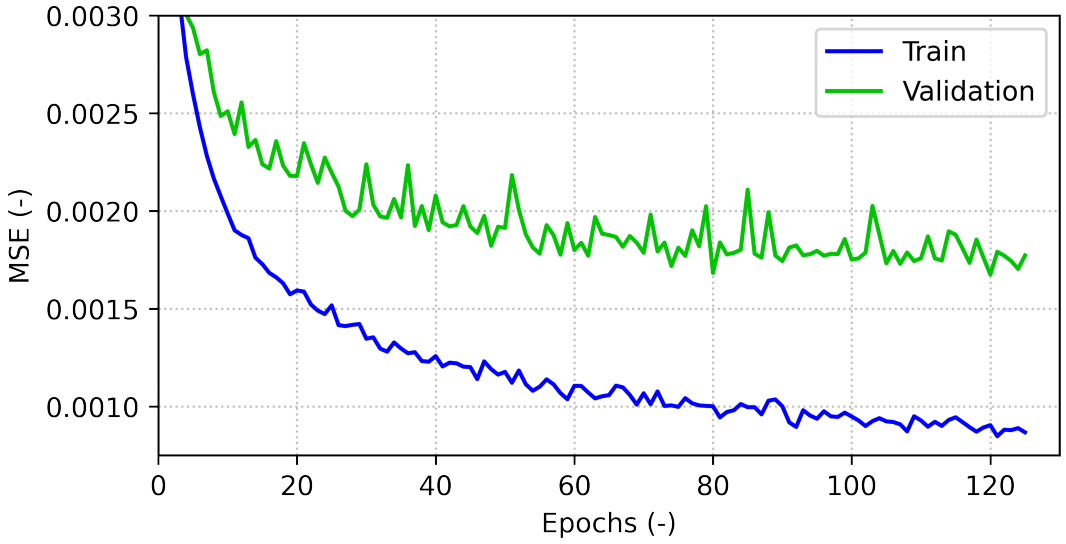

Figure 4.28: Loss and validation loss of the Refrigerators model training process

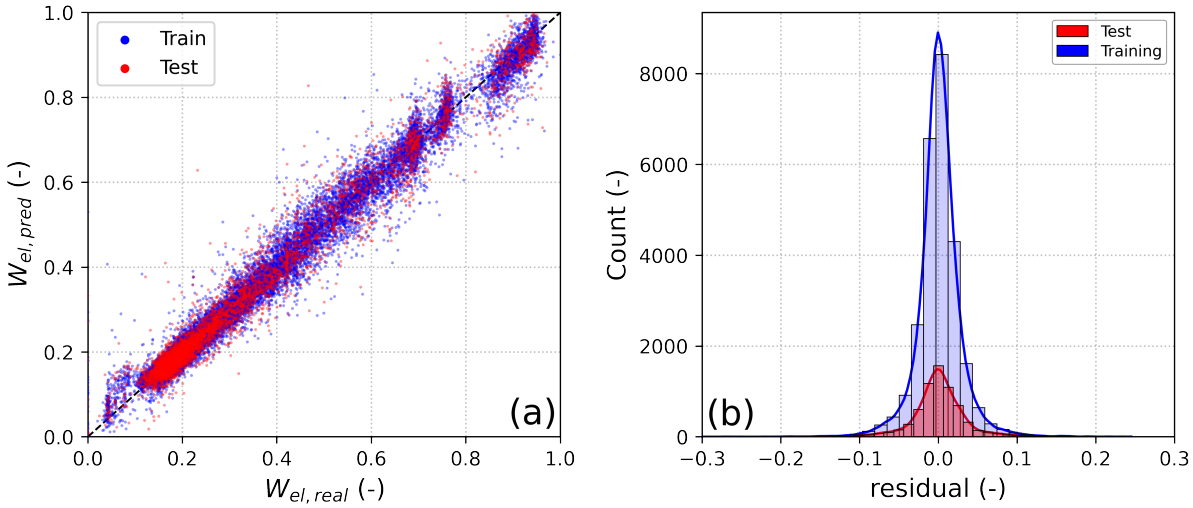

Figure 4.29: Real data VS. prediction (Refrigerators model): (a) Comparison of real and predicted power demand. (b) Residual statistical distribution

although finely tuneable, are inserted in series based on the instantaneous refrigeration energy demand. Consequently, the absorbed power levels could be (to a first approximation) "discretized". Hence, the clusters represent areas where, commonly, a new refrigeration group is activated to meet the refrigeration demand.

In any case, as can be seen, the dispersion of points around the diagonal is practically uniform over the entire operating range, presenting a small cloud of points in the low power area (lower left) that is slightly overestimated compared to the real counterpart. However, it must be said that the time in which the refrigerators works at such low loads is truly negligible in comparison to the entire standard year of operation, leading to a negligible error on the overall result and to a invisible effect on the subsequent monitoring method. As can be seen from Figure 4.29b, the distribution of residuals is markedly concentrated around zero for both datasets, maintaining a Gaussian behaviour and, consequently, confirming the validity and generality of the obtained model.

Table 4.8: Final model performance metrics (Refrigerators)

Metric	**Value**
R^2	0.9702
MSE	0.0015673
RMSE	0.039480
MAE	0.026494

4.2.2 Pumps

The water circulation system within the building is ensured by a complex pump system. The measuring system installed within the building power center has allowed us to obtain data related to the aggregated consumption of all the pumps. With regards to the water cooling system, the pumps of the primary water cooling circuit (between the collector and the refrigeration unit), the secondary circuits that allow the chilled water to be conducted to the various air handling units, and finally the pumps that ensure the circulation of the cooling water, i.e. the ducts which connect the chillers condensers (basement level) and the evaporative cooling towers on the roof, are all considered. A similar analysis can be performed for the hot water circuits, as the measured data includes the pumps used to deliver hot water to the pre and post-heating batteries of the air handling units. Nevertheless, the power demand is mainly shaped by the chilled water pumps circulation.

Correlation analysis

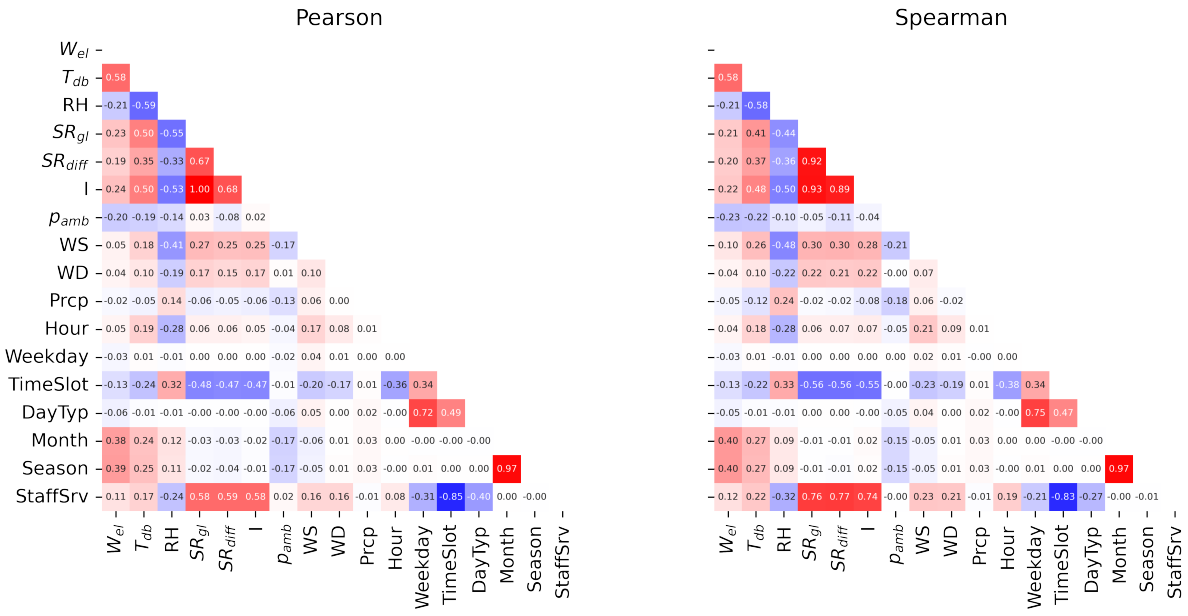

Figure 4.30: Correlation matrices for Pumps electricity demand

In order to achieve the goal of creating a predictive model for the pumps electricity demand, it is important to perform a preliminary analysis of the correlations between the available features (already analyzed in previous chapters) and the electricity demand of the water circulation system (Figure 4.30). In this case as well, a clear correlation between absorbed power and the outdoor air dry bulb temperature (T_{db}) can be observed ($r = 0.58$, $\rho = 0.58$). Since the pumps are dedicated to the air conditioning system, the strong correlation between their energy consumption and the outdoor weather conditions is not surprising. Unlike what was seen for the refrigeration units, the variation on a daily

time scale is less marked, resulting in long periods of almost constant consumption that take on different levels depending on the season. This is because during the night period at least one refrigerator is operative. The primary pumps are regulated through an ON/OFF management strategy and, consequently, a refrigerator operating at its base load will lead to the activation of the primary pumps at their nominal load. During the winter, the most used pumps are those related to the primary circuit of a single refrigerator. This lead to marked seasonal variations. In fact, the only two other features with noteworthy correlations are *Month* ($r = 0.38$, $\rho = 0.40$) and *Season* ($r = 0.39$,$\rho = 0.40$), which effectively mark the different regimes of use of this plant.

Feature Selection

The available features are obtained using the same Feature Engineering procedures described earlier and used for previous models. The wrapper method is initialized with all 86 obtained features. As a result, the preliminary neural network used for the feature selection procedure was initialized with the hyperparameters suggested by the literature ($lr = 0.01$, $ReLU$ as activation function and $HLN = 87$). As per procedure, all input features were scaled using the *RobustScaler* defined in Equation 2.14, while the absorbed power was restored to a unit variation interval by dividing it by the maximum value found

Table 4.9: Accepted features (Pumps)

Rank	**Feature**	R^2	**MSE**	**RMSE**	**MAE**
1	*Month* (Sine functions)	0.8448	0.0025593	0.050589	0.035062
2	T_{wb}	0.8699	0.0021456	0.04632	0.030192
3	p_{amb}	0.8996	0.001654	0.040669	0.025928
4	*Weekday* (One-Hot)	0.9307	0.0011423	0.033798	0.02208
5	*Hour* (Sine functions)	0.9443	0.0009152	0.030253	0.020029
6	*Month* (One-Hot)	0.9539	0.0007604	0.027576	0.018197
7	*HDD*	0.9622	0.0006229	0.024957	0.016263
8	*Hour* (One-Hot)	0.9671	0.0005425	0.023292	0.015785
9	*HR*	0.9688	0.0005138	0.022668	0.015492
10	p_{vap}	0.9694	0.0005048	0.022467	0.015256
11	*StaffSrv*	0.9696	0.0005001	0.022363	0.015581
12	*DayTyp* (One-Hot)	0.9698	0.000497	0.022292	0.015271
13	T_{dp}	0.9699	0.0004956	0.022263	0.015237

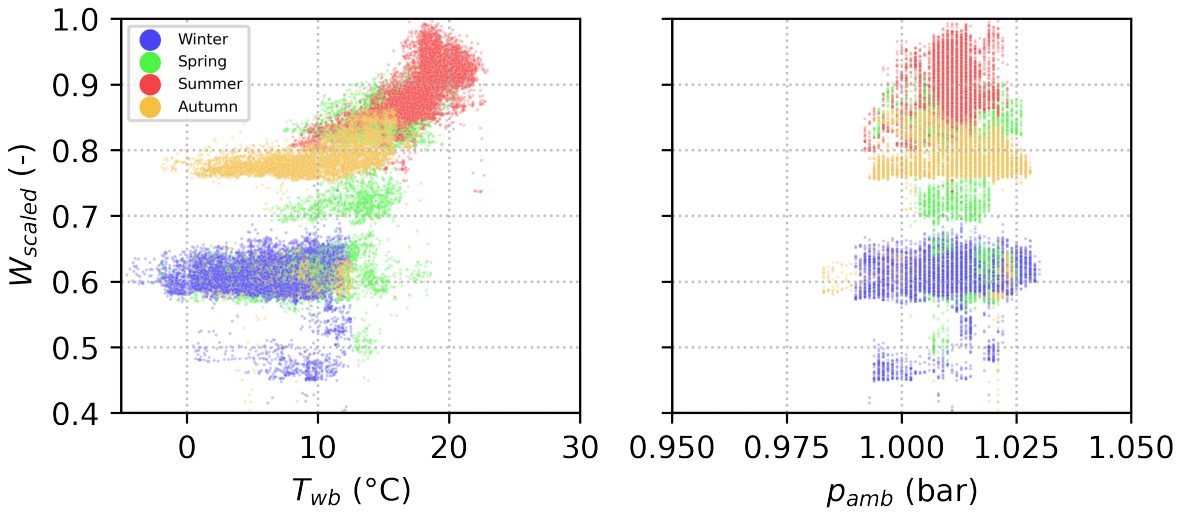

Figure 4.31: Outdoor air wet bulb temperature and ambient pressure VS. pumps electricity demand

in the training dataset.

Table 4.9 reports the ranking of the features accepted by the wrapper method. It is interesting to note that the two of the most important feature for predicting the demand for electricity is the outdoor air wet bulb temperature (T_{wb}), while the second one is the ambient pressure (p_{amb}). However, it should be considered that most of the pumps operate between refrigeration units (and therefore evaporative towers) and air handling units. These two macro-areas of the system are significantly influenced by weather conditions. It can also be observed that the combination of the first two features is, in a certain way, an indicator of the humidity of the ambient air, which on one hand influences the demand for cold water by the *AHUs* for the purpose of dehumidifying and cooling indoor environments, and on the other hand heavily influences the effectiveness of the evaporative towers.

The efficiency of the refrigerators depends on these latter ones, which will vary the temperature of the produced cold water, potentially leading to an increase in the demand for cold water. In Figure 4.31 the relationships between the two mentioned features and the demand for electrical energy are graphically represented. While for the wet bulb temperature (T_{wb}) a dependence can be observed, the ambient pressure seems almost completely uncoupled from the target, as confirmed by the correlation coefficients illustrated in Figure 4.30 ($r = 0.20$, $\rho = 0.23$). Seems that the "representation" of relative humidity as a combination of these two climate features positively influences the model training. Further confirmation of what has been said can be noted by the presence of *HR* in ninth position in the ranking, of p_{vap} in tenth position, and T_{dp} in thirteenth position.

Regarding the *calendar features*, the most important feature (which is the first in the feature ranking) is the *Month* represented by the sine functions and, after, *Month* in the form of One-Hot encoding. Both features provides information about the seasonal

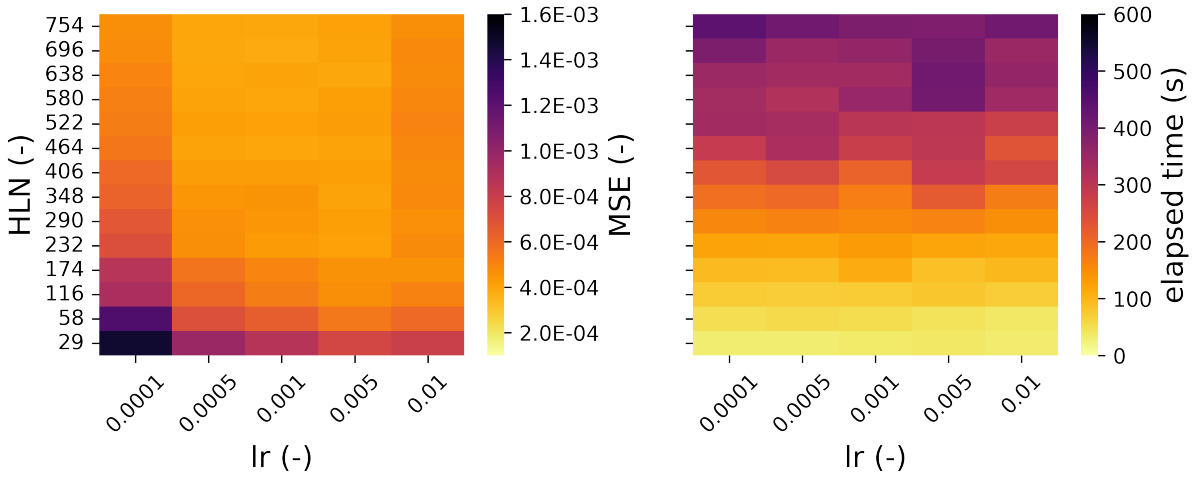

Figure 4.32: Hyperparameter tuning results with "ReLU" activation function for Pumps model

cycle. The influence of the seasonal cycle is evident by observing Figure 4.31, in which the four seasons have been coloured differently. It is clear that winter consumption and summer consumption are at opposite ends, while the both Autumn and Spring are split in half, taking on either summer or winter behaviours. What actually changes the behaviour significantly is the activation of the summer air conditioning system, which is predominant in terms of the pumps energy demand. Then, it can be seen that *Weekday* expressed through One-Hot encoding constitutes one of the most important features, in accordance with the weekly cycle that the entire air treatment system follows, being clearly underutilized during weekends. For the same reason, but relating to the daily cycle of activities in the structure, we find *Hour* expressed through the two sine functions.

Hyperparameters Tuning

The hyperparameters tuning procedure was performed, as previously, using the grid search approach. The investigated activation functions are the rectifier linear unit ($ReLU$) and the hyperbolic tangent ($tanh$). The learning rate (lr) takes the same five values used previously. Regarding the number of neurons in the hidden layer (HLN), the series of investigated values was obtained by considering multiples and submultiples of the "literature-suggested" number of neurons, which is the sum of the number of inputs (58) and outputs (1).

Figures 4.32 and 4.33 represent the results of the hyperparameter tuning process using heatmaps, showing on the left the mean squared error of each step of the method, and on the right the corresponding time taken to perform the trainings. It should be noted that, as explained in chapter 2, each step of this method consists of three separate trainings, following the *k-fold cross validation* method to ensure the generalization of the process.

As has already happened in previous cases, the trainings performed with *tanh* as the

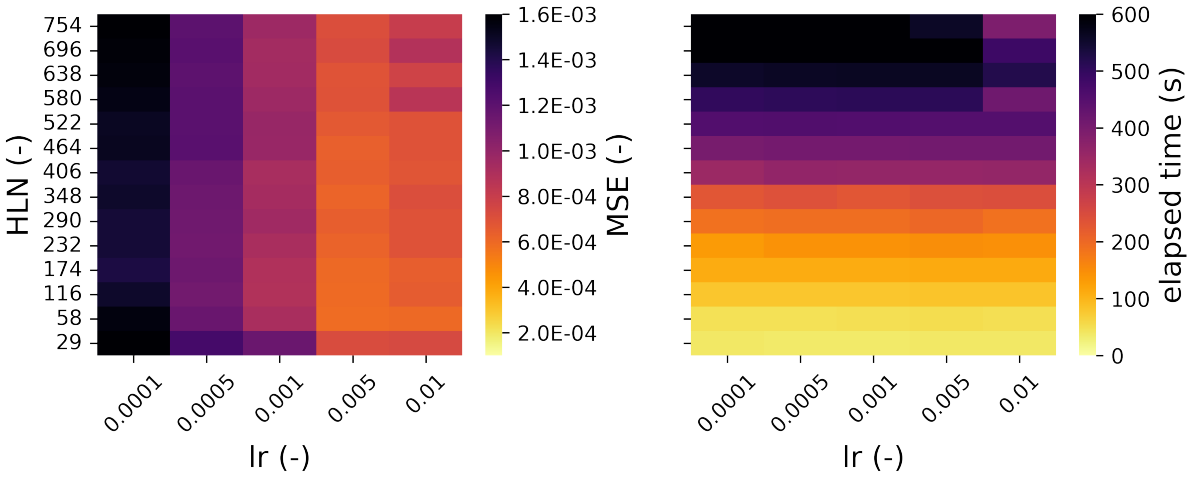

Figure 4.33: Hyperparameter tuning results with "tanh" activation function for Refrigerators model

activation function have significantly lower performances compared to those performed using the rectifier linear. Therefore, by observing Figures 4.32 in more detail, it can be seen that the computing times are constant as the learning rate varies, while the prediction performance has minimums for intermediate *lr* values. As previously observed, as the number of neurons in the hidden layer varies, the performance increases significantly, together with the computation times. The best compromise is constituted by the training with $ReLU$ as the activation function, 0.001 as the learning rate, and a number of hidden layer neurons (HLN) equal to 464.

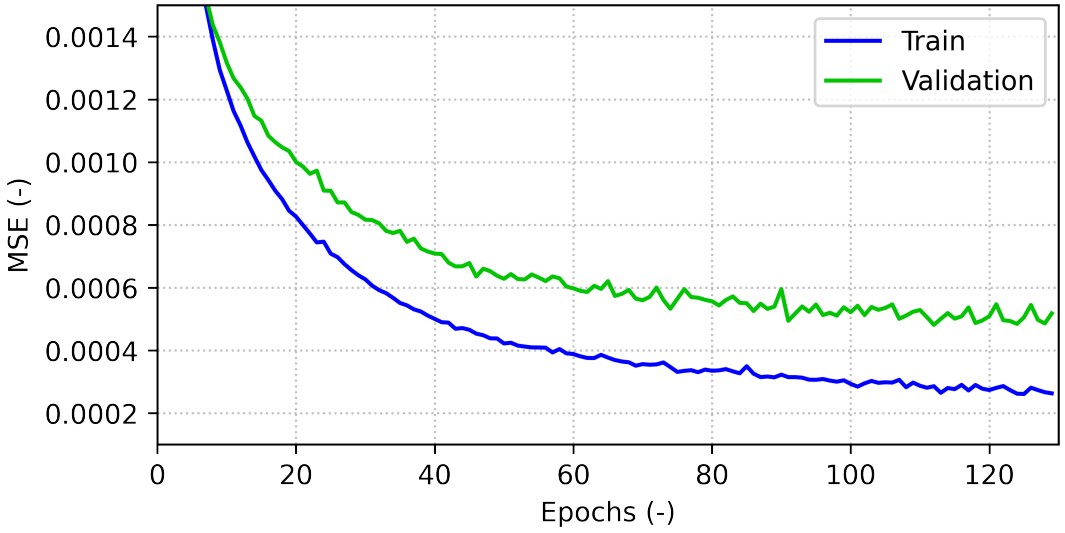

Figure 4.34: Loss and validation loss of the Pumps model training process

Table 4.10: Final model performance metrics (Pumps)

Metric	Value
R^2	0.9551
MSE	0.0007420
$RMSE$	0.027076
MAE	0.017656

Final model training

Once the best feature subset and the best hyperparamenters set were determined, the final model training has been carried out. Figure 4.34 shows the loss and validation loss functions behaviour along the entire model training process. The comparison shows that the model generally performs better on the model training, resulting in a lower value of loss function in comparison to the validation loss one during all the epochs. Nontheless, the validation loss behaviour present a markedly similar behaviour, meaning that the model present good prediction performance also for data which were not previously seen by the model. The loss and validation loss behaviour show low oscillations, further confirming the validity of the model. The training process reach the convergence after 124 epochs. Further analyses about the prediction performance of the model have been carried out on the test dataset (20% of the entire training data). First of all, the main metrics defined in section 2.1 were calculated and they is reported in Table 4.10. The model shows high prediction performance, reaching an R^2 score of 0.9551 on the test dataset (which were not involved in the model training process).

Tu further analyse the model prediction capabilities, an in-depth analysis of the prediction on both training and test dataset have been carried out. Similarly to the preceding cases, the real and predicted values have been compared in Figure 4.35a exploiting a scatter plot. The figure shows that a great part of available data-prediction couples tend to cluster in two areas. The first is characterized by a "low" power demand, while the second cover a greater span in correspondence to higher power demand. This can be explained by observing the pumps electricity demand behaviour along the training year (Figure 3.27), characterized by long periods at constant power demand separated by sudden and marked changes in correspondence to the change between winter and summer management strategies of the *HVAC* systems. In those areas, the training and test predictions results to have similar behaviours, confirming the "goodness of fit" of the analysed model. Nevertheless, there are other samples which are not included in these two cluster which present higher dispersion in comparison to the previous ones ($0.4 < W < 0.58$ and $0.7 < W < 0.75$ in Figure 4.35). The samples belonging to these areas are characteristic of periods of the year when the transition from the summer to

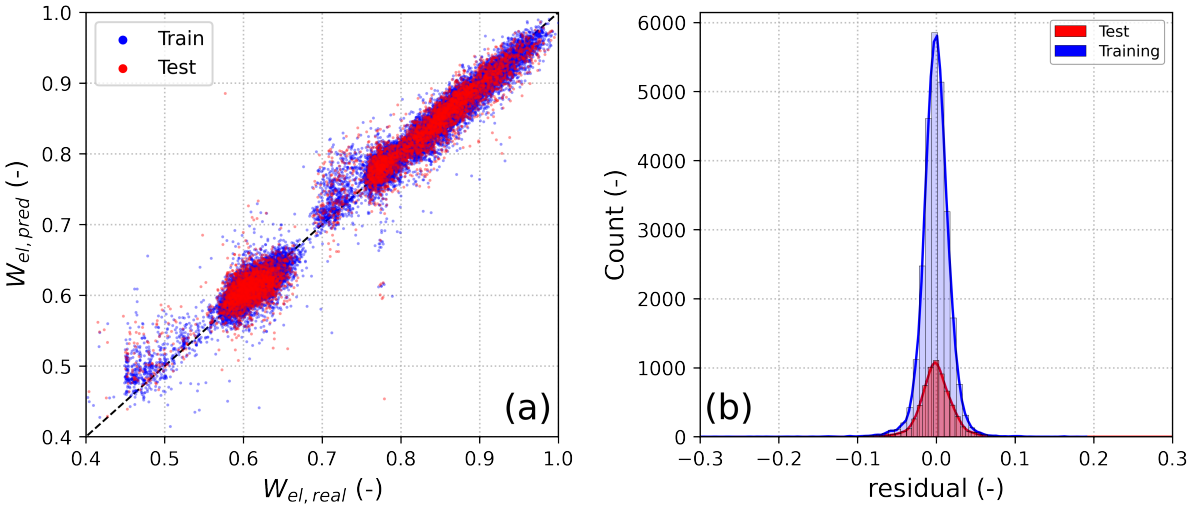

Figure 4.35: Real data VS. prediction (Pumps model): (a) Comparison of real and predicted power demand. (b) Residual statistical distribution

the winter configuration (or vice versa) has taken place gradually (see in particular the periods at the beginning of May 2020 and the end of February 2021) through hand-made system managemetn strategies tuning performed by the technical staff of the facility. Since the 15-minute time-step datasets are highly redundant, it is possible to keep these two periods within the starting dataset, as they have characterized the overall behaviour of the facility. In fact, although the average prediction performance is lower in these two periods, they are still well modelled by the trained neural network, which is able to predict their behaviour even on the test dataset.

Finally, the residual histogram shown in Figure 4.35b let to further confirm the validity of the obtained model. Indeed, the residual distributions of the test dataset is highly comparable with the training dataset one. The residuals are markedly concentrated around the zero, presenting similar distributions for negative and positive values. Given the preceding results analyses, the model is considered valid for its exploitation in the monitoring method, exhibiting exceptional prediction capabilities.

4.2.3 Air handling units

Correlation analysis

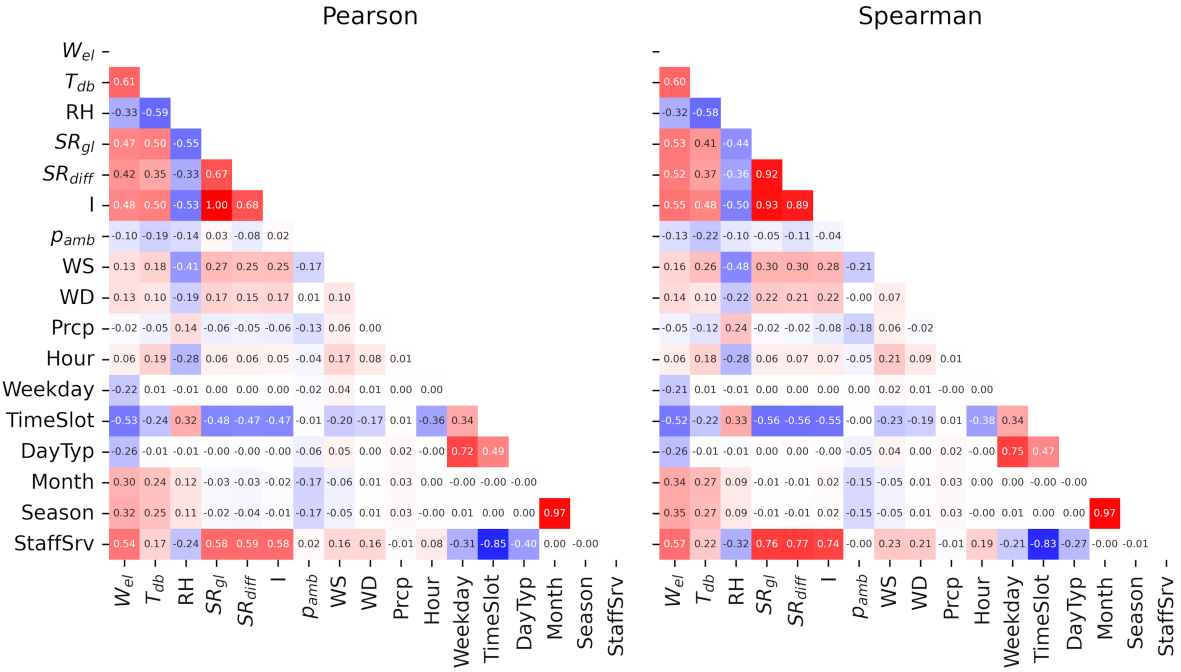

Figure 4.36: Correlation matrices for Air Handling Units electricity demand

As already said, the air handling unit are operative during the whole year to ensure the correct indoor environment ventilation and indoor climate conditions. Their electricity demand, (which is mainly guided by the fans operations) is less variable on a seasonal scale, being mainly affected by the daily cycle and the specific daily and weekly schedules assigned to each unit.

Figure 4.36 shows the correlation coefficients of the air handling units electricity demand with the basic available features. The outdoor air dry bulb temperature (T_{db}) is the most correlated features, with a Pearson's correlation coefficient (r) of 0.61 and a Spearman correlation coefficient (ρ) of 0.60. Confirming the major influence of the daily and weekly time scales, the other highly correlated features are:

- the ones related to the solar radiation; SR_{gl} ($r = 0.47$, $\rho = 0.53$), SR_{diff} ($r = 0.42$, $\rho = 0.52$), I ($r = 0.48$, $\rho = 0.55$)
- *TimeSlot* ($r = -0.53$, $\rho = -0.52$), which is mainly related to the weekly time scale
- *Staff Srv* ($r = 0.54$, $\rho = 0.57$), related to the daily and weekly building activities cycle

Other weekly parameters present an intermediate correlation with the *AHUs* electricity demand, namely *Weekday* ($r = 0. - 0.22$, $\rho = 0.21$) and *DayTyp* ($r = 0. - 0.26$, $\rho = 0.26$).

Nevertheless, intermediate correlations can be found for the seasonal time scale features as *Month* ($r = 0.30$, $\rho = 0.34$) and *Season* ($r = 0.0.32$, $\rho = 0.35$). Indeed, the seasonal time scale, even if in a less marked way, influence the air handling unit exploitation because of the specific adjustment made on the air flow rate set point during the workdays based on the people feedback, leading to mildly increase the ventilation requirements during the summer period.

Relative humidity (RH) present intermediate correlation with the electricity demand ($r = -0.33$, $\rho = -0.32$). The relation between W and RH can be resumed in two main points:

- The variation of the air density based on the air humidity, but also to the reduced performance of the evaporative cooling towers during periods with high air humidity. A lower cooling tower performance lead to a decrease in the refrigerators COP and, consequently, to an increase of the chilled water temperature or, in extreme cases, to a reduction in the chilled water mass flow rate deliverable by the chillers. Consequently, the capabilities of the air handling units to reduce the external air temperature will decrease. This fact often lead to manually adjust the air mass flow rate treated by the *AHUs*. It is important to remark that these aspects concern the people perception and behaviour, aspects that are difficult to implement in a physical model of a building

- The correlation coefficient are both negatives, meaning that an increase of the air humidity leads (generally) to a decrease in the *AHUs* energy demand. This lead to think that the relation between RH and W is more related to the seasonal temperature variation. Indeed, RH and T_{db} have high negative correlation coefficients between each other ($r = -0.59$, $\rho = -0.58$). Since the energy demand of the air handling unit present a marked positive linear relation with T_{db} and higher values of the latter tents to correspond to low RH values, the relation between W and RH is probably a sort of coincidence.

Nevertheless one possibility does not exclude the other. The two phenomena are both present, but given that they have opposite effects on the correlation coefficients, they tend to compensate each other. This leads to the low correlation coefficients found between W and RH.

The relation between W, T_{db} and RH is graphically represented in figure 4.37. As it can be seen, the relative humidity values results lower in correspondence to high values of outdoor air temperature. Contemporary, the range in which the absorbed power varies seems to reduce as the outdoor temperature increase, leading to higher electricity demand in correspondence of the lower relative humidity values. The effect of RH on the energy

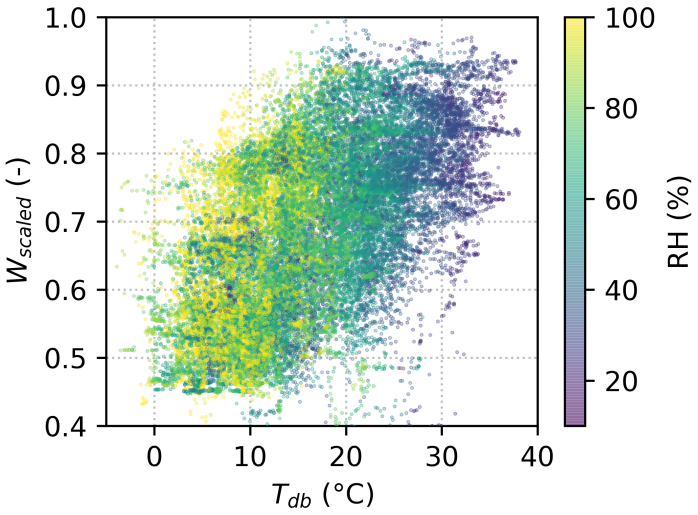

Figure 4.37: Dry bulb outdoor air temperature VS AHU power demand

demand can be seen by removing the temperature variation. As an example, it can be extracted the data related to a small range of high temperature (temperatures in which the evaporative towers are highly exploited), referring only to the daytime of workdays, in order to exclude the effect of the activity intensity. So, it is possible to observe data sets that are minimally affected by factors other than relative humidity using this method.

Figure 4.38 shows that, considering only the effect of the relative humidity, the power demand seems to increase with the latter. This logic makes it possible to further emphasise the importance of having both features (or features derived from both through the feature engineering process) as input to machine learning models, so as to be able to predict energy demand in more detail.

Finally, in Figure 4.36 can be seen that the p_{amb} feature results to be weakly correlated with the $AHUs$ power demand ($r = -0.10$, $\rho = -0.13$). This phenomenon is peculiar,

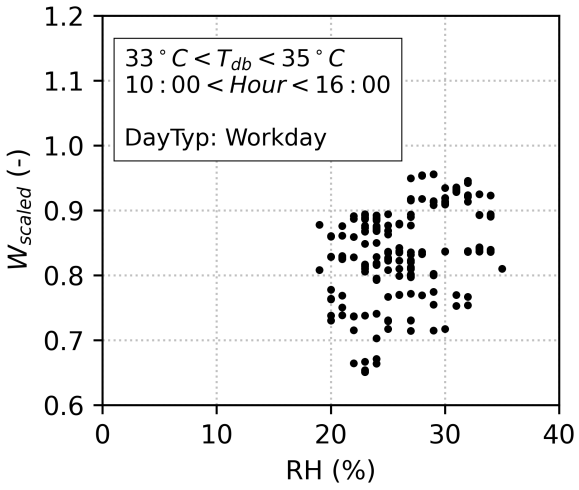

Figure 4.38: Outdoor air relative humidity VS AHU power demand

as ambient air pressure should have a significant influence on the performance of the supply and exhaust fans, which directly interact with the outdoor air. Nevertheless, as already discussed, a low correlation coefficient does not necessarily imply low influence on consumption, nor that the feature in question cannot contribute to the training of predictive models.

Feature Selection

Table 4.11: Accepted features for wrapper method (Air Handling Units)

Rank	**Feature**	R^2	**MSE**	**RMSE**	**MAE**
1	*Month* (One-Hot)	0.4247	0.0093662	0.096779	0.076176
2	*StaffSrv*	0.6929	0.0050001	0.070711	0.050046
3	h_{wet}	0.774	0.0036793	0.060657	0.042465
4	*Weekday* (One-Hot)	0.8124	0.0030538	0.055261	0.039074
5	*Hour* (One-Hot)	0.8637	0.0022188	0.047104	0.032875
6	*HDD*	0.9043	0.0015584	0.039477	0.028161
7	p_{amb}	0.9183	0.0013308	0.036481	0.025853
8	HR	0.9288	0.0011584	0.034036	0.024451

Following the same procedure used for the preceding cases, the feature selection has been carried out through the wrapper method. The results highlight that eight features results to enhance the artificial neural network model learning process. The selected hyperparameters are the same used for the preceding cases ($lr = 0.01$, $ReLU$ as *activation function*, $HLN = 86$). Table 4.11 reports the feature ranking obtained at the end of the feature selection. First of all, it can be observed that outdoor air temperature (T_{db}) is not present among the selected features. Following the reasoning made during the correlation analysis, it was stated that the instantaneous temperature have a weak influence on the electricity consumption of AHU fans. Therefore, the high correlation coefficient was mainly due to the influence of temperature on a seasonal scale, and consequently, that data was mainly indicative of the current season and how the air treatment is used (in summer, the processed air flow rate is generally increased to deal with the lower cooling capacity of the system during the hottest hours). This information can be summarized by a combination of heating degree days (HDD) and month of the year ($Month$), which are both present in the accepted feature ranking.

The correlation analysis had highlighted intermediate correlations regarding some features at the weekly and daily scale. In fact, among the accepted features, there are the

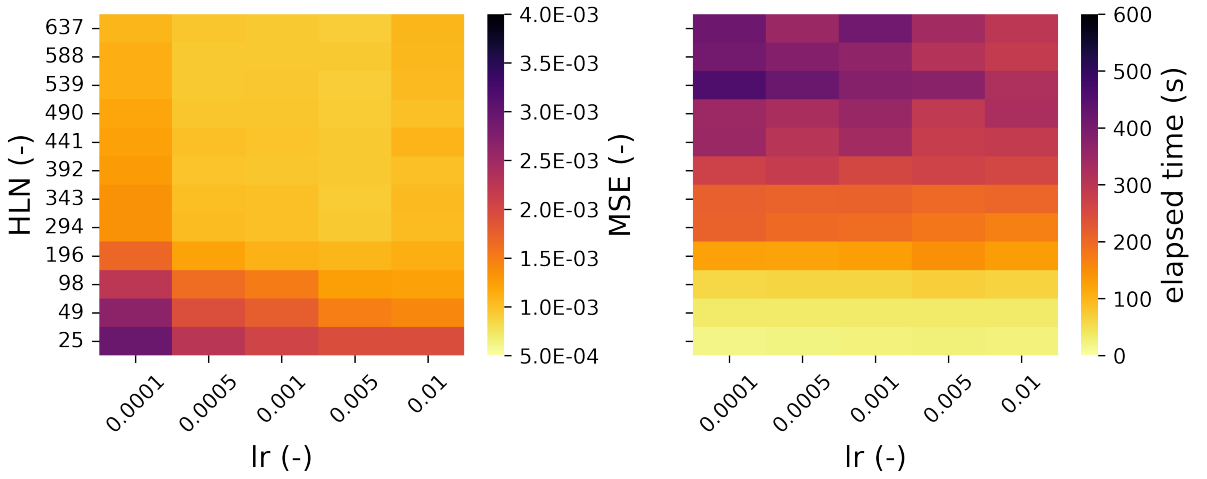

Figure 4.39: Hyperparameter tuning results with "ReLU" activation function for Air Handling Units model

staff in service (*StaffSrv*) which results second in the feature ranking and regards both weekly and daily cycles, and both *Weekday* for the weekly time scale and *Hour* for the daily one. During the correlation analysis, the effect of the relative humidity have been thoroughly analysed. Although the specific parameter is not an accepted feature, there are two features in the ranking that bring information about the air humidity, demonstrating that this is also a determining factor for accurate consumption prediction. These two feature are the moist air enthalpy (h_{wet}), which is third in the feature ranking and results a pivotal feature for the learning process, and the Humidity Ratio (HR) at the bottom of the feature ranking, which only marginally improves the prediction performance.

Among the accepted features, there is a last feature in seventh position, the ambient pressure (p_{amb}). This feature had a low correlation coefficient with the electricity demand, seeming to not affect the consumption of air handling unit fans. This fact seemed peculiar, given that the fans take in and discharge air from the external environment. In fact, proceeding with the systematic method of feature selection, this feature results instead to have a positive impact on the model prediction performance, although it cannot be compared to the top features in the ranking. In any case, this is information that could not have been possible to obtain using simpler methods such as selecting the feature based on correlation coefficient analysis.

Hyperparameters Tuning

The hyperparameters tuning procedure was performed, as previously, using the grid search approach. The investigated activation functions are the rectifier linear unit ($ReLU$) and the hyperbolic tangent ($tanh$). The learning rate (lr) takes the same five values used previously. Regarding the number of neurons in the hidden layer (HLN), the series of investigated

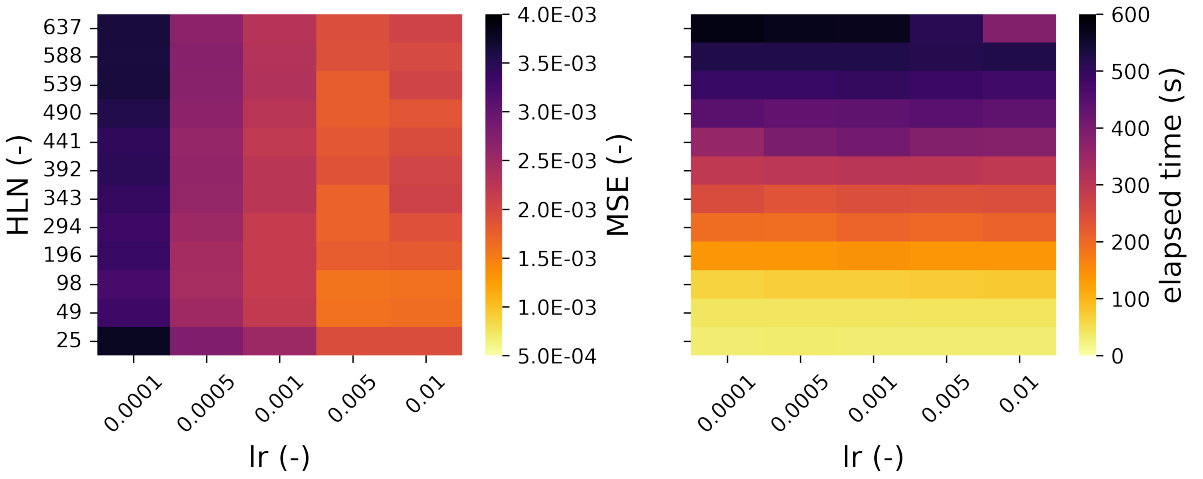

Figure 4.40: Hyperparameter tuning results with "tanh" activation function for Air Handling Units model

values was obtained by considering multiples and submultiples of the "literature-suggested" number of neurons, which is the sum of the number of inputs (48) and outputs (1).

The behaviours of prediction performance and elapsed time corresponding to the variation of hyperparameters are represented in Figures 4.39 and 4.40. The *ReLU* has been confirmed to be superior to the *tanh* both in terms of prediction performance and computational cost. Regarding the learning rate, Figure 4.39 shows a clear minimum in the prediction error (MSE) in correspondence of $lr = 0.005$, especially for high number of neurons in the hidden layer. Nevertheless, the aim to choose best combination of prediction performance and computational cost (elapsed time) lead to select $HLN = 343$.

Final model training

Table 4.12: Final model performance metrics (AHU)

Metric	Value
R^2	0.9290
MSE	0.0011522
$RMSE$	0.033832
MAE	0.024403

Finally, a training has been carried out with the information obtained through the preceding steps. Figure 4.41 shows the loss function and validation loss function behaviour along the training process. The two functions present similar trends, even if non-negligible oscillation. Nevertheless the training process ends after 122 epochs, reaching an R^2 of 0.929 on the test dataset. The other metrics are presented in Table 4.12.

The scatter plot used to compare the real consumption values and the predicted one is

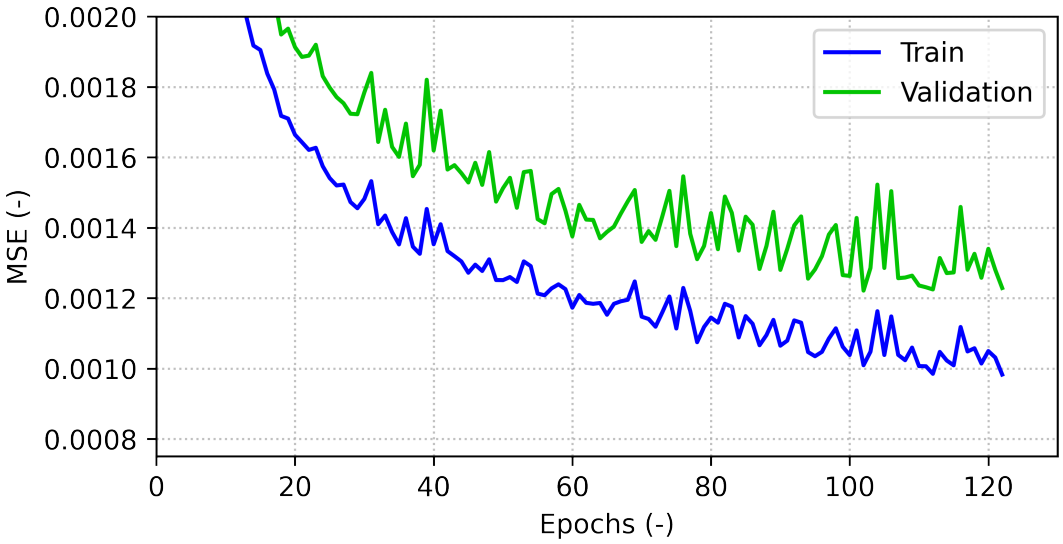

Figure 4.41: Loss and validation loss of the Air Handling Units model training process

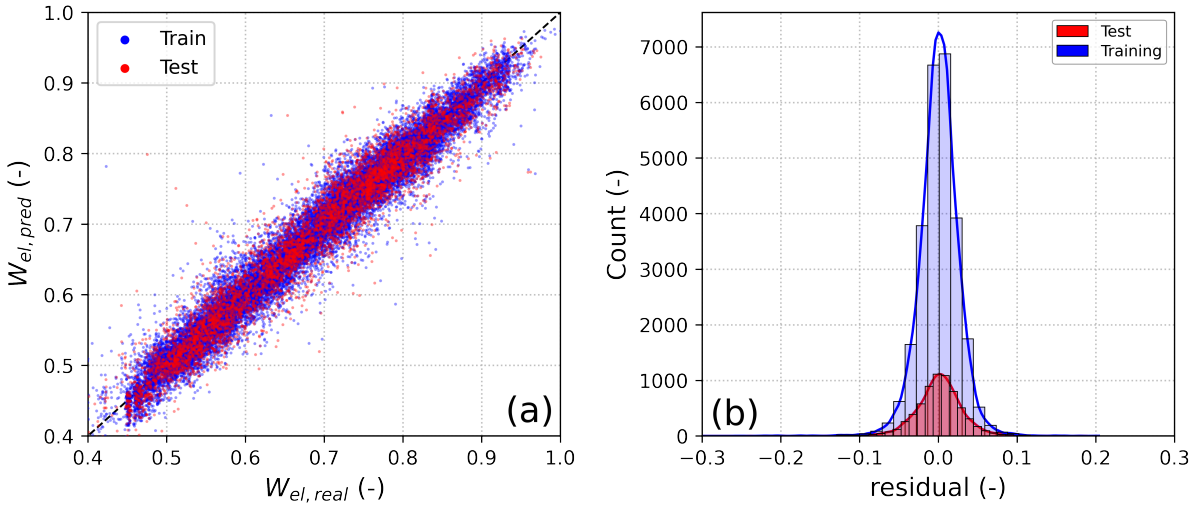

Figure 4.42: Real data VS. prediction (AHU model): (a) Comparison of real and predicted power demand. (b) Residual statistical distribution

reported in Figure 4.42a. The points results uniformly and densely distributed around the graph diagonal, a sign that demonstrate the goodness of fit of the obtained model in every situation. Figure 4.42b shows the statistical distribution of the model prediction residual (errors between real and predicted values). The residuals oscillates between -0.1 and 0.1, with a gaussian-like distribution which is similar of both training and test datasets. The results have been further analysed and it can be concluded that the obtained model has sufficient prediction performance and the capability to generalise to previously unseen data.

4.3 Monitoring method

In the previous sections, extensive tests were conducted on various methodologies to exploit raw data to realize predictive models based on machine learning. The tests were performed on the whole building electricity of the Healthcare Facility, comparing prediction performance, computational costs, and the amount of knowledge and experience required to apply the specific methodology. As a result, the selected method (**Case 6**) allow to simultaneously obtain the best prediction performance and lower prior knowledge required, with the main drawback of an increased computational cost. This method was then applied to three data series related to the *HVAC* systems of the test case under analysis. As a result, four predictive models (Whole building, Refrigerators, Pumps, Air Handling Units) have been trained, validated, and are ready for use. In the next chapter, these models will be used to implement the *CUSUM* method, analysing in detail the information obtained through its application.

This section is focused on the application of the obtained models to realize the monitoring method. The latter exploit the technique called Cumulative Sum of Differences (*CUSUM*) to perform the *building energy monitoring*, which follows the procedure explained in section 2.1.5 (Eq. 2.15).

Figure 4.43 shows the *CUSUM* chart for two cases. The first (**Case 1**) exploit Multiple Linear Regression with filter-based feature selection, which present the lower computational cost and the worst prediction performance. The second analysed model is **Case 6**, in which the *ANN* model is trained using the best feature subset (defined through wrapper-based feature selection) and the optimal hyperparameters (obtained through the grid search fine tuning procedure.

First of all, the behaviour of the *CUSUM* test period can be analysed (see red line in Figure 4.43. As expected, both cases have roughly the same behaviour.

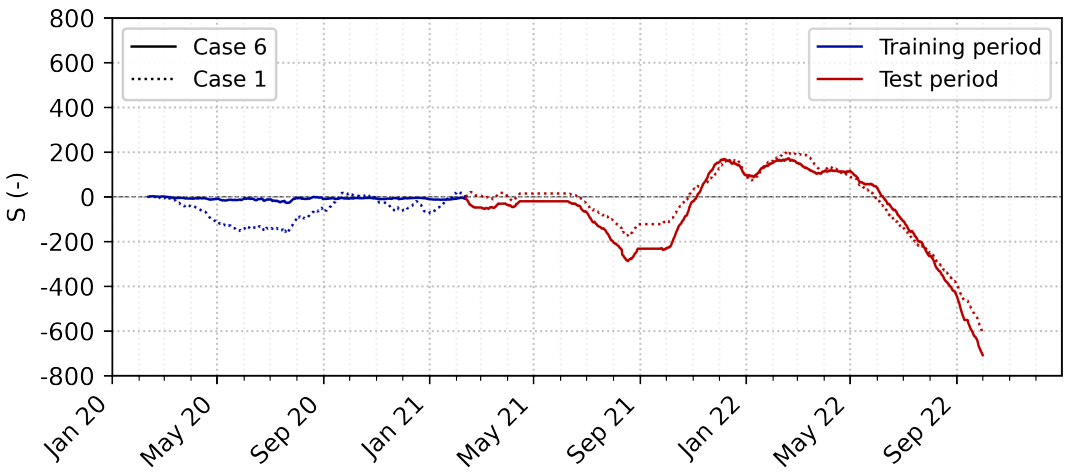

Figure 4.43: Cumulative Sum of Differences using for Case 1 and Case 6 models

Starting from *February 2021*, a general maintenance of the *Pumps* system have been carried out, leading to a mild reduction in the overall energy consumptions. Maintenance operations led to interruptions of the *HVAC* systems in several areas of the facility, leading to a physiological reduction in consumption due to abnormal activity rhythms. In fact, the main reason that led to schedule the maintenance operations during the winter period was the need to minimise discomfort to people in the various departments and areas. An interruption of the air conditioning in the middle of a hot summer period would in fact have been much more difficult to tolerate. Consequently, starting from *February 2021*, a decrease can be identified in the *CUSUM* curve of **Case 6**. It is important to notice that the monitoring method based on Multiple Linear Regression (**Case 1**) is unable to highlight this reduction.

On *February 23, 2021*, the pumps system was fully functioning. Moreover, in the same period the air handling units related to the surgery department were regulated in order to increase the overall handled air mass flow rate, leading the overall energy demand to return similar to the expected one. As the summer season approaches, the impact that the Air Handling Unit and refrigerators electricity demand has on the overall energy consumption increase in comparison to the pumps, which instead maintain nearly constant energy consumptions (see Figure 3.20). This leads the *CUSUM* curve to gradually decrease the descending trend, proceeding horizontally. In other words, the reduction of the pumps electricity demand become proportionally too low to be sensed by the monitoring methods, which, instead, highlight an energy consumption aligned to the standard one (which is related to the previous year) after *February 23, 2021*.

In *June 10, 2021* the installation of new Air Handling Unit and new *UPSs* have been carried out. Nevertheless, a general maintenance of the three refrigerators have been carried out (also visible in Figure 4.46), leading to a sensible reduction in their energy demand for the same amount of cooling energy supplied. The *CUSUM* curves instantaneously start to decrease (despite the new systems installation) from this moment until *August 2021*, when an interruption of the electrical service caused some problems to the centralised building energy management system. Automatic control of some air handling units was no longer possible and they were left in manual mode until *August 29*, when the control system was restored.

At the beginning of *October 2021*, an important inefficiency has been found regarding the management strategy of the refigerator's primary circuit pumps: they were set to remain operative even when the chillers were switched off. A prompt correction of the control logic has been implemented on *October 2, 2021*, leading to a reduction in the pumps overall electricity demand.

On *October 4 2021*, a new abrupt interruption of the facility's power supply led to

extensive damage to the centralised control system.

This issue led to a general reset in many of the refrigeration units operating parameters which loss the their optimal values, with a consequent increase in their energy demand. The issue was not detected until *December 10, 2022*, leading to several months of wasted energy. The presence of an advanced monitoring method, integrated with the systems control system, would have allowed an early warning of the problem, allowing for timely preparation of corrective measures.

By the end of December, several actions was carried out. First, an optimization campaign of the pressure set point of the secondary pump (which allow the water circulation from the refrigerators to the air handling units) was done. The refrigerators operative and control parameters issue was found, leading to restore the correct control strategy. Moreover, the refrigerators chilled water temperature set point has been increased in order to optimize the refrigerators operations.

In *January 2022*, a significant change in the processed air flow rate was made on the air handling unit of the radiotherapy department, which had long been used as a storage room, effectively halving the consumption of the corresponding *AHU*. The effectiveness of these actions is confirmed by the decrease in the *CUSUM* curve over that period, which shows a clear reduction in consumption compared to expectations.

On *10 January 2022*, new *AHU* dedicated to the in-patient rooms were installed, leading to a physiological increase in the building's energy demand. In addition, the reduced consumption of the *HVAC* systems in the winter period makes the new *AHUs* and *UPSs* installed in *May 2021* weigh much more heavily than before. The combination of these factors is detected by the energy monitoring method, which again shows an increasing curve, a sign that actual consumption is higher than predicted by the predictive model.

Starting from *February 20, 2022*, an optimization campaign was carried out on the air handling units operative parameters. The processed air flow rates have been set according to the current regulations and some climate sensors were installed in order to dynamically adjust the air temperature set point at the cooling coils, reducing the chilled water needs for air handling purposes. The mass flow rate of the surgery department air handling units have been increased in order to ensure the correct ventilation, which is a critical factor in these type of activities. The *CUSUM* shows a mild decrease. Even if the overall energy consumption results slightly lower than the expected one, the energy consumption due to the air circulation results higher than before. These two aspects almost counterbalance each other, leading the curve to proceed almost horizontally.

Between *March* and *May 2022*, the position of the dampers inside the ventilation ducts that allow the distribution of the *AHU* processed air among the individual rooms was checked. Especially in large departments served by a single *AHU*, many incorrect air

distributions were identified. Office rooms received a much higher air mass flow rate in proportion to more critical rooms, such as those containing electro-medical equipment. A repositioning of the dampers inside the air ducts allowed a more balanced air distribution, leading to a reduction in the total air flow rate processed by larger *AHUs*, such as the one serving the diagnostics department.

On the beginning of *May 2022*, nevertheless, the increase of external temperatures lead to carry out some adjustments in order to ensure the correct air treatment for the healthcare facility departments. In particular, the refrigerators chilled water set point were decreased in order to ensure the correct functioning of the Air Handling Units cooling coils. This leads to an increase of the refrigerators loads and, consequently, of the related energy consumptions.

Between the end of *May* and *June 2022*, two contemporary action were carried out on the refrigerators. First, a thorough review was carried out on the control logics of the refrigeration units and, consequently, on the activation of the primary pumps. In particular, the simultaneous activation of several refrigeration units was reduced to the minimum possible, thus significantly reducing the electrical energy consumption relative to the circulation of water in the primary circuit and the activation of the relative evaporation towers. Moreover, an in-depth maintenance has been done on the refrigerators, obtaining a significant increase in their efficiency. As confirmed by the *CUSUM* chart, the implemented measure resulted in an apparent energy saving that has been maintained over time, and continues until the end of the analysed period.

The difference in the standard deviation of S during the training period (blue lines) between **Case 1** and **Case 6** is clearly visible. Furthermore, the lower precision of the *MLR* model leads to a reduction in oscillations during the test period, further reducing the possibility of highlighting anomalies in the energy demand behaviour.

Indeed, the model training process aims at minimizing the mean square error over the training period, leading the S values to oscillate around zero. The standard deviation σ of these oscillations (training period) provides a range within which variations in S can be considered normal. In other words, the training period has been taken as a reference for the S variability. The determination of the acceptable range of S can vary depending on the specific application requirements. A narrower range leads to more frequent alarms, as the value of S during the testing period would fall outside of the acceptable range more often, resulting in a higher number of alarms.

It can be concluded that a more accurate model will result in lower oscillations of S during the training period, and, as a consequence, an improved ability to detect anomalies. Specifically, a smaller oscillation during the test period will be required to trigger the anomaly condition, resulting in a more sensitive and prompt monitoring tool.

The ultimate goal of the *CUSUM*-based monitoring method is to establish a system capable of real-time monitoring of building or system consumption, to detect anomalies, and upon exceeding a certain deviation considered "standard", alert the user, ensuring the possibility to analyse the issue and prepare prompt corrective actions.

In order to simulate the implementation of the monitoring method in the Building Energy Management System, the *CUSUM* chart described above, have been further elaborated. Since one of the aim of the monitoring method is to provide alerts and alarms in case of anomalies detected in the building energy demand behaviour, it is important to establish a basic criteria. Therefore, as explained in section 2.1.5, it is necessary to establish threshold levels beyond which the method will trigger an alarm. The adjustment of these thresholds is entirely independent of previous steps and must be chosen based on specific context considerations, as well as subject to modifications and calibrations during normal operation of the structure. In general, the threshold is selected based on the values that the *CUSUM* assumes during the predictive model training period, thus, in this case, on the first year of available data.

In this work, the threshold was established by utilizing the standard deviation of the *CUSUM* curve for each model created. Specifically, the threshold is equal to three times the standard deviation. In practice, a "standard" building consumption in the years following the training period will oscillate around zero within a range of $\pm 3\ times\sigma$. Exceeding the upper threshold will indicate that the actual building consumption has been consistently higher than expected for a consecutive period of time. Similarly, if the lower threshold is exceeded, it will indicate a systematically lower actual consumption than expected. Assuming real-time implementation of the method, following an alarm, there will likely be an intervention to restore normal system operation, which should bring the curve back to proceeding horizontally, at least until new anomalies arise. Assuming the alarm was generated by exceeding the upper threshold, from that point on the curve will proceed horizontally, remaining outside the threshold. Therefore, it will continue to generate alarms with each update of the curve (in this case, every 15 minutes). As a result, following a corrective intervention for a malfunction, the *CUSUM* should be reset to the zero value in order to bring the *CUSUM* curve back to its initial conditions.

For example, Figure 4.44 shows the *CUSUM* curve obtained through **Case 6**, resetting its value at each "known" intervention performed on the plant, as previously reported during the analysis of the curve behaviour (red line in Figure 4.43). The normal *CUSUM* behaviour (gray line) has also been reported for clarity. The blue shaded region represents the region where the *CUSUM* variation is considered normal, obtained therefore through the standard variation of the CUSUM during the year preceding the period represented in the figure. It is also possible to notice dotted vertical lines symbolizing the first alarms

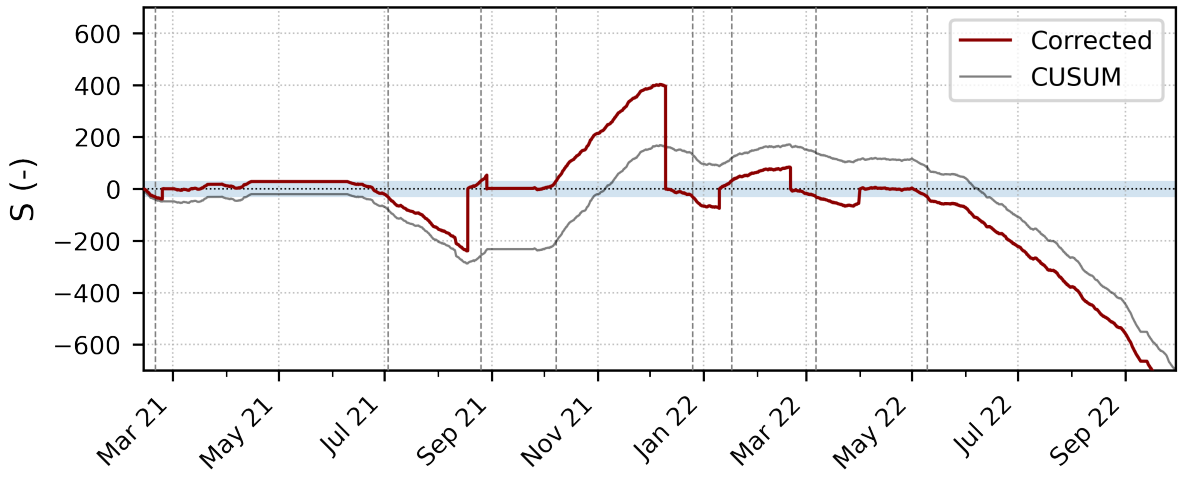

Figure 4.44: CUSUM alarms representation for Case 6

generated by the curve following the overcoming of the alarm threshold. Since the *CUSUM* was not implemented during the analysed period, resolution interventions are in no way related to the generation of the alarm, and a noticeable delay in intervention can be observed in certain cases.

Following what has been stated, the advantage of creating a predictive model that is as accurate and versatile as possible has not yet been highlighted. The main advantage becomes evident by comparing **Case 1** (*MLR* with filter method) and **Case 6** (*ANN* with wrapper method and hyperparameter tuning). Indeed, a more accurate model reduces the fluctuations of the *CUSUM* during the training period due to the reduction of the magnitude of the residuals at individual moments. This results in a significant reduction of the alarm threshold that will be applied during operations and, as a consequence, the readiness of the system will in detecting any anomalies. A perfect example of this situation is represented by the power interruption that occurred on *October 4, 2021*, which led to serious problems in the automatic management of the plants.

Figure 4.45 shows the trends of the two cases *CUSUM* curves, from the power service interruption until the complete resolution of the problem (where the *CUSUM* curve is then "artificially" reset). In addition to the two curves, other information has been represented. The coloured areas represent the *CUSUM* alarm thresholds respectively for **Case 1** (red area) and *Case 6* (blue area). The substantial difference in the ranges assumed by the two areas is evident. The black dotted vertical lines represent respectively the moment of the service disruption (and therefore the damage) and the moment when the restoration interventions were completed. Observing the blue curve, it can be observed how the change in slope occurs before the actual conclusion of the interventions. This is due to the fact that the interventions were carried out over a period that included several days, leading to a gradual restoration of normal operation. At the conclusion of all interventions (exactly

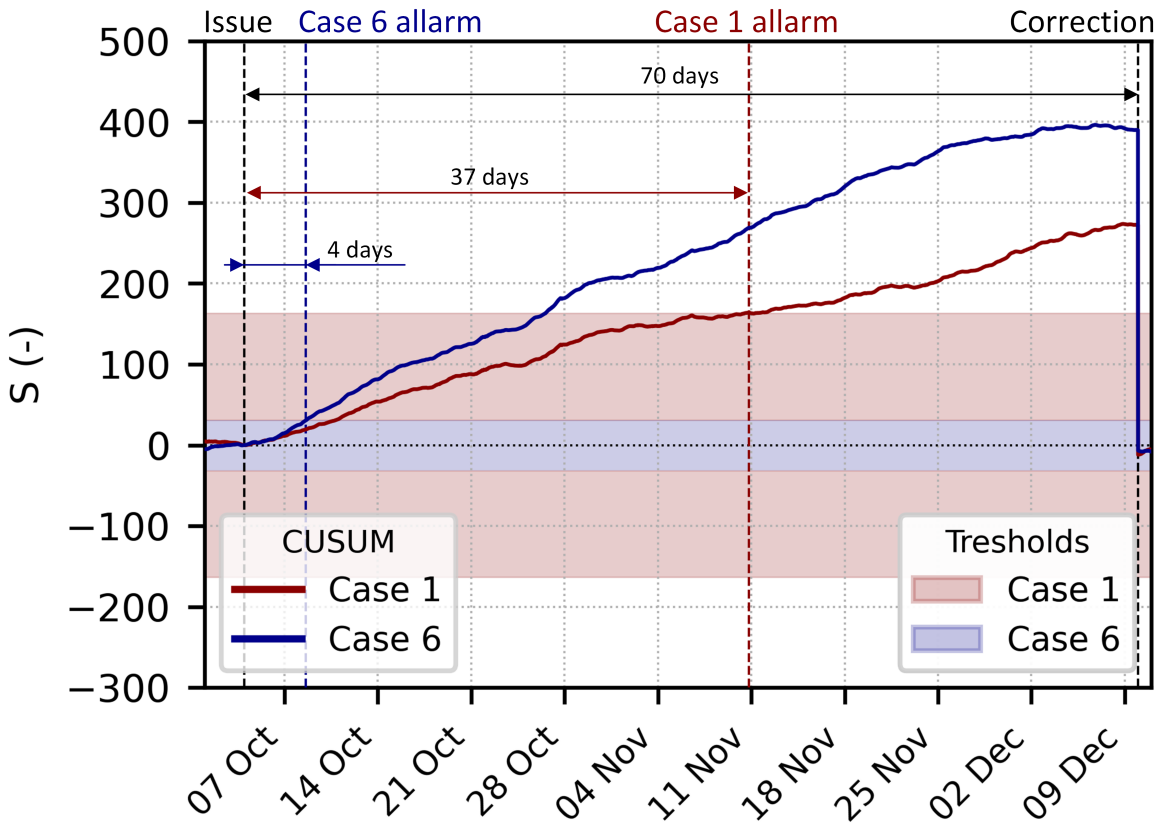

Figure 4.45: Graphical visualization of time needed by the two monitoring method to identify the anomaly in the building energy demand

as would occur in a real-world use of the method), the *CUSUM* is reset.

The latest information reported on the graph concerns the two coloured vertical dashed lines, which represent the moment the threshold is exceeded (and, as a result, the automatic notification of the alarm to users). These lines are coloured in a consistent manner with the model they refer to, red for *Case 1* and blue for *Case 6*. The monitoring method based on Multiple Linear Regression is able to detect the anomaly after about 37 days, while the one based on the optimized neural network requires only 4 days. As already stated, this is mainly due to the significantly higher prediction performance of the second case. The implementation of the second method would therefore have allowed for corrective actions to be taken over a month earlier than the first method. It should also be considered that no monitoring method was implemented during that period. It can in fact be noted that corrective actions were taken about 70 days after the advent of the problems, leading to entire months energy waste which was potentially avoidable with the use of the monitoring method.

The potential of this method lies in the fact that its applicability does not strictly depend on the context of application, but mainly on the availability of data. In industrial and commercial sectors, plants are generally managed by a centralized Building Energy

Management System (*BEMS*) equipped with various sensors to obtain useful operating data for implementing automatic regulation strategies. In these contexts, therefore, data accessibility is rarely an issue. Furthermore, detailed data on the consumption of a specific point of energy delivery are made available on the user personal web page of most of the italian energy provider, avoiding the need for the autonomous installation of dedicated measurement systems.

However, technological advancements now make it possible to create dedicated, non-invasive and modular measurement systems at a relatively low or certainly affordable cost for most of the commercial and industrial activities. As a result, in addition to monitoring the overall energy consumption of the structure by upgrading the installed systems, it would be possible to perform a detailed monitoring of individual areas or components of the building energy systems, providing the possibility of obtaining increasingly specific and interpretable alarms, further facilitating the possibility of targeted corrective actions.

The healthcare facility under analysis installed a system for measuring energy consumption. The next section will be dedicated to applying the proposed monitoring method to more specific Heating, Ventilation, and Air Conditioning (*HVAC*) subsystems, whose results will be compared to those obtained by applying the method to the whole building energy demand.

In the preceding, it was proved that the energy consumption monitoring of the entire facility energy demand can be crucial in practical application, providing highly accurate alarms and allowing corrective actions to be taken. However, it does not provide information regarding the specific issue that has arisen. This task is the responsability of the facility Energy Manager (or Energy Team), who, through knowledge of the installed systems, will be able to identify a problem based on the symptom reported by the energy monitoring method.

However, the presence of dedicated area measurement systems in the analysed test case (see Section 4.2) made it possible to create predictive models dedicated to the *HVAC* system. This allowed the developing monitoring methods dedicated to more detailed sections of the structure as the measurement system becomes more detailed. The information obtained through the *CUSUM* chart related to the whole building energy demand (analyzed in the preceding sections) will then be compared with the information obtainable by monitoring individual parts of the system, whose models were created in the previous section. Using the four predictive models, the *CUSUM* control charts were created, and are shown in Figure 4.46.

On *June 10, 2021*, new *AHUs* and new *UPSs* were installed in the structure. However, the increase in consumption was compensated by the maintenance of the refrigeration units. In fact, it can be seen how the curve related to the whole building energy consumption

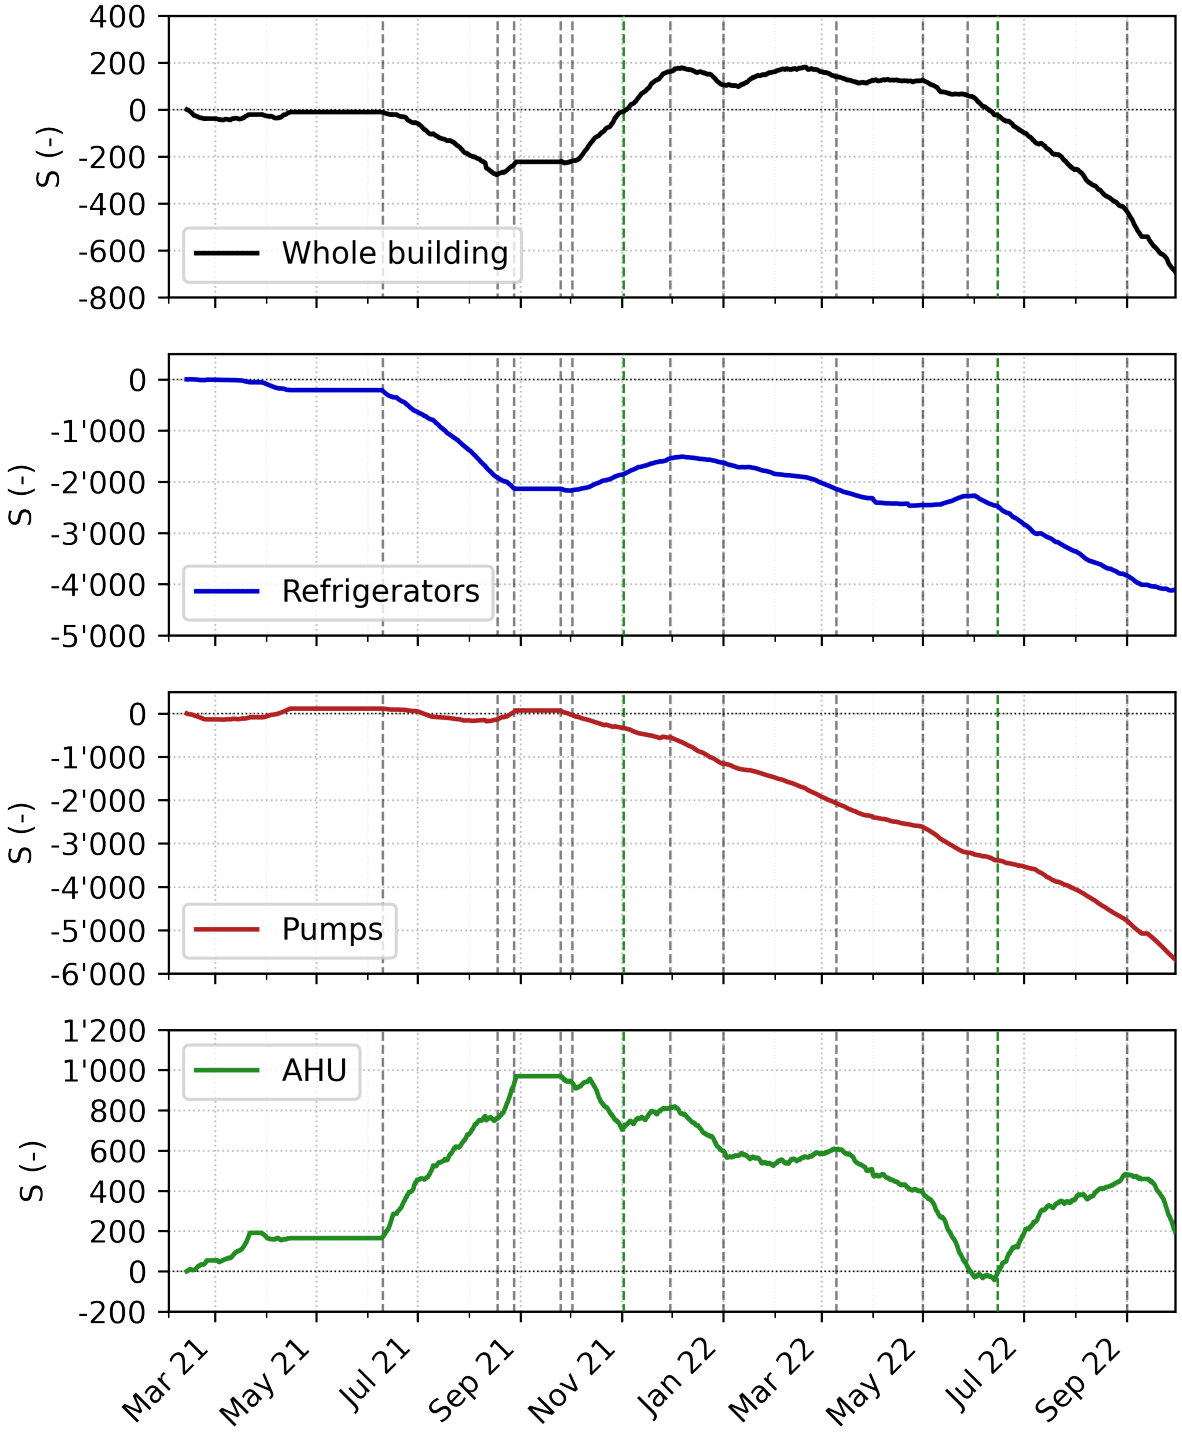

Figure 4.46: CUSUM chart for all the definitive models

is decreasing, similarly to that of the refrigeration units which in fact constitute the majority of the energy savings obtained during this period. Since the activity of the refrigeration units (in terms of hours of operation) remained unchanged, the consumption of the pumps does not undergo particular deviations compared to normal consumption. Indeed, it can be seen how the pump *CUSUM* curve remains horizontal. By monitoring the energy consumption of the *AHUs* alone, it can be instead noticed a marked increase in consumption. Nevertheless, this behaviour have been compensated by the energy demand

reduction related to maintenance of the refrigeration units, and cannot be detected by the *CUSUM* of the whole building energy demand.

On *August 18, 2021*, an interruption in the electrical service causes issues in the management system, particularly related to the *AHU* operations. This issue has been resolved ten days later. Since the main subject of the issue are the air handling units, the refrigerators operations remain unchanged, while the mass flow rate of the air handling units and consequently the chilled water exploitation increase. Indeed, looking at figure 4.46, it can be seen that the refrigerators curve presents a mild slope increase due to the increase in chilled water production, but their consumption continue to be lower in comparison to the ones found during the training period, confirming the effectiveness of the previous maintenance operations. The pumps *CUSUM* start to increase because of higher need of chilled water circulation toward the air handling units, even if the consumption increase remain almost imperceptible. The main effect can be seen by observing the *AHU CUSUM*. Indeed, the *AHU* curve present a marked slope increase in this period, due to the increase in *AHU* operations caused by the issue in the *BEMS* system.

As already said in the preceding, on *October 2, 2021*, a new abrupt interruption of the national electrical service cause some issues in the refrigerators control parameters, leading to an increase in the energy consumption. Consequently, the whole building and the refrigerators *CUSUM* start to increase. Indeed, the chilled water temperature set point is too low in these period. An interesting effect take place on pumps and *AHUs*. The lower temperature of the chilled water lead to decrease the chilled water mass flow rate need due to the automatic *HVAC* management, decreasing the energy demand related to the pumps. Moreover, it increase the refrigeration capabilities of the air handling units, leading to decrease their operative load for several areas, reducing the overall *AHUs* energy demand.

On *November 2, 2021*, the air mass flow rates of the *AHUs* related to diagnostic department and magnetic resonance were increased in order to ensure the control of the indoor climate condition. This change was made for specific needs on that day and not for the normal operation of the structure. However, the change was mistakenly maintained, leading to a significant increase in air treatment related consumption. Looking at figure 4.46, it can be seen that this anomaly has been highlighted only by the *AHU CUSUM*, while the whole building curve does not highlight any changes in its behaviour. This highlights the significance of developing highly specific monitoring methods. The alarm generated by the *CUSUM* for the *AHUs* not only would have revealed an anomaly that would have gone unnoticed, but it would have also provided information that the issue was caused by the air handling units, significantly reducing the scope of the problem and allowing for targeted, prompt corrections to be made, thereby avoiding costs associated with inspecting systems and facilities that are unrelated to the problem.

On *November 30, 2021*, several action were taken to optimize the energy demand of the *HVAC* system. First of all, the issue introduced on *November 2* was acknowledged and resolved, restoring the normal *AHU* operations that, indeed, return similar to the preceding slope. Moreover, it was noticed that the processed air by the *AHU* of to the radiotherapy department (currently used as a storage room) much higher than what was necessary. The *AHU* control parameters were corrected, leading to half its energy demand. Then, the operational parameter of the refrigerators were restored to their correct values, as testified by the refrigerators *CUSUM* which stop to increase. These should cause a variation of the *CUSUM* slope, leading it to proceed horizontally. Nevertheless, since the the overall processed air flow from the *AHU* was reduced, the demand for chilled water also decreased accordingly. Indeed, the *CUSUM* curve relative to the chillers assumes a slightly negative slope.

On *May 28, 2022* an in-depth maintenance of the refrigerators unit were performed. Moreover, as already said, their management strategies was changed to reduce the contemporary functioning of multiple units. However, the unusual hot summer forced users to manually set the *AHUs* to their maximum potential. Indeed, while the *CUSUM* related to the entire building, the refrigeration groups, and the pumps show a reduction in consumption in the subsequent period, that of the *AHUs* highlights an increase, bringing the curve from a clearly negative slope to proceed horizontally, bringing consumption in line with what was obtained in the training period (before optimization operations).

On *June 15, 2022*, several other *AHUs* were manually set to 100% load, leading the *AHUs CUSUM* to markedly increase until *September 1, 2022*, when their normal functioning was restored. It is important to note that this increase in the *AHU* energy demand is highlighted only by the *AHU* monitoring method, while the whole building *CUSUM* do not sense anything because of the contemporary marked decrease in chillers electricity demand.

4.3.1 Discussions

Beyond the description of individual modifications and anomalies detected, it is noteworthy to observe that the four curves depicted in Figure 4.46 are consistent. The high impact of the *HVAC* systems on the whole building energy demand is also testified by the fact that changes found in the cumulative sum of the entire building are almost always attributed to modifications of one of these systems. The variations of the four cumulative sums are in fact very "coordinated."

On the one hand, monitoring the electrical energy of the entire building is important and, most importantly, achievable without particular efforts like the installation of dedicated measuring systems, allowing the implementation of a real-time diagnostic system that detects "generic" symptoms that require a cause to be found. On the other hand, especially for industrial activities, it is becoming increasingly convenient to implement advanced and complex measurement systems, which would allow the implementation of complex monitoring systems for various areas of the building, thus providing more precise indications regarding the detected anomalies. As already seen, if not enough detail is given, it is possible to miss anomalies in the consumption of a component if they are compensated by other anomalies. The perfect example is the increase in *AHU* consumption in mid-June 2022, which is not detected by the monitoring of overall consumption, as it is compensated by the net reduction in refrigeration group consumption.

Secondarily, monitoring based on complex and distributed measurement systems within the building, reduces the needed user's experience in interpreting alarms, by providing increasingly specific indications about the source of the anomaly as the measurement system becomes more detailed. This aspect produces a synergy with one of the goals of the method proposed in this work, namely reducing the knowledge required in the engineering, mathematical, and computer science fields in order to develop and apply the monitoring method by adopting a systematic approach. With a detailed measurement system and using the proposed method, it will be possible to proceed almost automatically from input data to the final model, provided sufficient computational resources are available. In any case, the cases analysed in this section allow for the evaluation of the compromise between the required knowledge and computational cost that best fits the specific application.

4.4 Energy saving evaluation

The *CUSUM* curves obtained in the previous section are useful for identifying deviations in the energy behaviour of the building, thereby detecting abnormal consumption. Consumptions exceeding the expected values will highlight the occurrence of issues causing energy inefficiency. Conversely, consumptions lower than expected could indicate periods of higher energy efficiency, enabling the verification of the effectiveness of any actions taken for this purpose.

The predictive models developed in this study can be exploited to quantify energy waste or savings during these periods. As for the analysed test case, three main periods are highlighted by the *CUSUM* curves:

- From October 5^{th} to December 9^{th}, 2021, a steady increase in consumption is observed following a national electrical system interruption, which led to malfunctions in the facility management system that were not immediately recognized and corrected. The development of *CUSUM* curves for individual components comprising the air handling system lead to attribute the increase in consumption to the three installed refrigeration units

- On September 25^{th}, 2021, a modification in the logic of the water circulation system management resulted in significant energy savings from the pumps installed in the building. This savings is maintained until the end of the analysed period. The predictive consumption model of the entire building would make it impossible to evaluate such savings since other events and modifications occurred in other parts of the structure in the subsequent period. However, by using the predictive consumption model specifically for the pumps energy consumptions, it is possible to distinguish the analysed variation from those that occurred later, allowing for the assessment of the actual energy savings resulting from the pump intervention

- On May 28^{th}, 2022, a maintenance operation was carried out on the refrigeration units, significantly reducing their electrical consumption. Following the same reasoning as in the previous point, it is possible to leverage the predictive consumption model for the refrigeration units to estimate the energy savings achieved through the maintenance operations, separating it from the concurrent and anomalous increases in the air handling units' consumption due to manual interventions on their usage schedules (see previous section)

This section will be dedicated to estimating the variations in energy consumption for the aforementioned cases, which are considered the best and clearest examples to

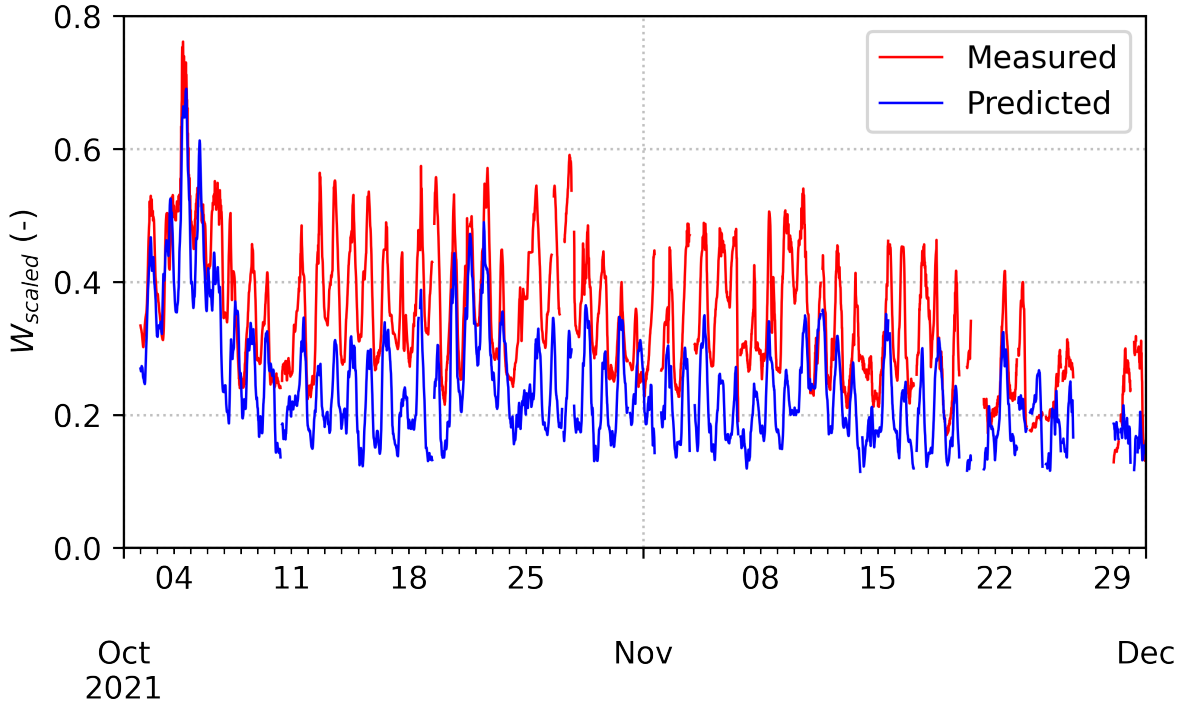

Figure 4.47: Measured Vs. Predicted refrigerators energy consumption line plot (Oct. 5 - Dec. 9 2021)

demonstrate the potential of the proposed methodology for conducting energy consumption analyses of the structure and its systems.

The first analysed case pertains to the period from October 5^{th} to December 9^{th}, 2021, during which the consumption of the refrigeration units consistently exceeded the standard. Considering that the analysed period follows the training period, precise prediction performance is not expected, as the purpose of the method is precisely to highlight a constant deviation between the predicted and measured values. The difference between the two behaviours can be observed in Figure 4.47, which depicts the curves of measured consumption (red curve) and predicted consumption (blue curve). It can be noted that for almost the entire analysed period (as confirmed by the *CUSUM*), the predicted consumption remains significantly below the measured consumption.

It is possible to further analyse the comparison between the values taken by the two data series by examining them instant by instant through a scatter plot (Figure 4.48). Each point in the graph represents an observation instant. As can be seen, two distinct zones are identified, separated by the diagonal (dashed red line). The upper part is a zone where the predicted values are higher than the measured values. In other words, if a point resides in this zone, the considered system is consuming less than expected, leading to a decrease in the *CUSUM* curve. The lower zone represents the exact opposite, and each point residing in this area exhibits real consumption higher than the predicted consumption. As already evident from the simple curves, it can be observed that the majority of instants have real

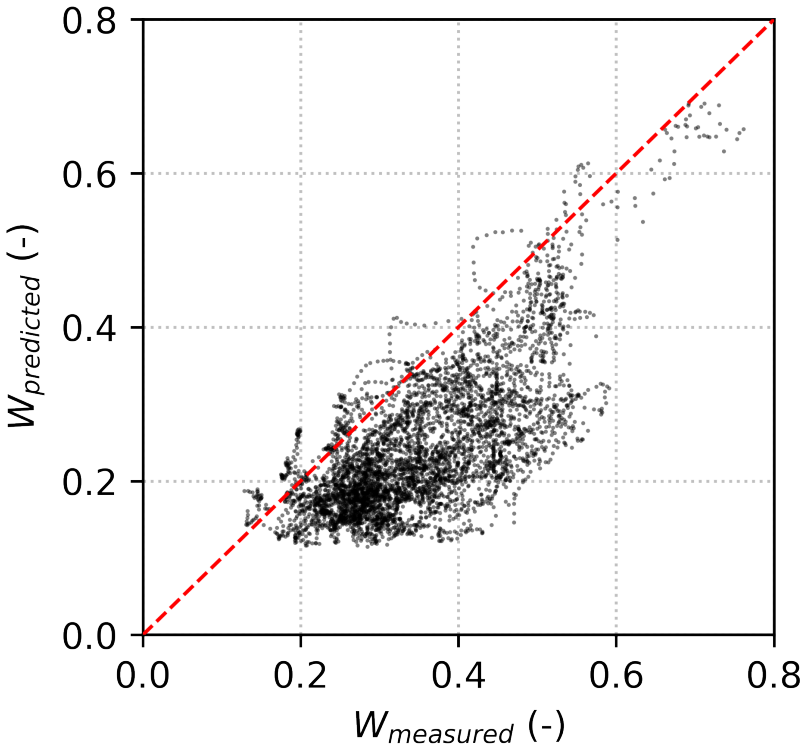

Figure 4.48: Measured Vs. Predicted refrigerators energy consumption scatter plot (Oct. 5 - Dec. 9 2021)

consumption significantly higher than the predicted consumption, with the cloud of points primarily residing in the lower zone.

Referring back to Figure 4.47, the two trends represent the instantaneous power (consumed or predicted) of the refrigeration units. Following this reasoning, it is possible to calculate the total energy quantity consumed in the analysed period, as well as the energy that the system would have consumed under normal conditions. In fact, the area between the two curves represents the amount of wasted energy in the considered time period, allowing for an estimation of energy waste and, consequently, the potential benefit that could be obtained through prompt corrective action if the monitoring method were implemented in the *BEMS*. In this case, the issue with the management system led to a 29.83% increase in consumption for the refrigeration units. By weighing this energy excess against the overall building consumption during the analysed period, it can be concluded that the unresolved issue resulted in a 4.69% increase in overall consumption.

The period starting on September 25^{th}, 2021, following the improvement of the pump management system, allows us to observe a phenomenon opposite to the previous one. From that moment, the pump-related consumption is lower than expected until the end of the considered period. Figure 4.49 allows comparing the actual consumption curve with the predicted one for the considered period. Several observations can be made compared to the previous case. Firstly, it is evident that the range of variation in measured powers

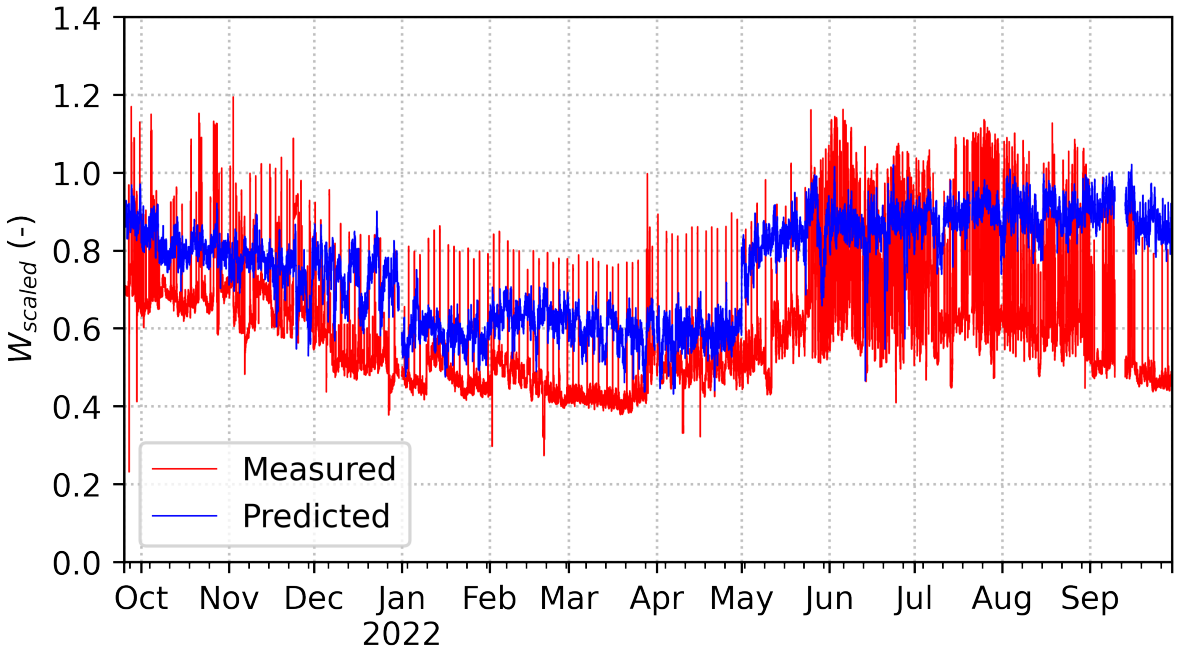

Figure 4.49: Measured Vs. Predicted pumps energy consumption line plot (Sep. 25 2021 - Sep. 30 2022)

compared to the predicted powers is different. The predicted curve (obtained using the model from the previous year) shows that with the boundary conditions of the considered period, the management system would have maintained higher levels, especially regarding the minimum absorbed power. This can be explained by considering that during the training period, the pumps related to the primary circuit of the refrigeration units were kept active even for those refrigeration units that were not in operation, generating a constant and easily avoidable baseline consumption. In fact, the measured data show significantly lower minimum powers, as the pumps were deactivated consistently with their corresponding refrigerators.

Furthermore, it can be noted that the upper peaks of the measured consumption often exceed the predicted ones, especially during the summer periods. This phenomenon is due to the fact that the abnormal temperatures during the summer of 2022 led to manual adjustments of many air treatment systems, resulting in an increased cooling demand for the entire system and, consequently, for the pumps in the primary circuit. As will be observed in the analysis of the subsequent period related to the refrigerators, this increase in consumption is not captured by the energy monitoring method specific to the refrigerators themselves, as it is completely offset by the maintenance work carried out just before the summer period. However, it is possible to highlight that if the automatic regulation of the air treatment system had been maintained during the summer of 2022, the savings comparable to the maintenance of the refrigeration units would have been even more significant.

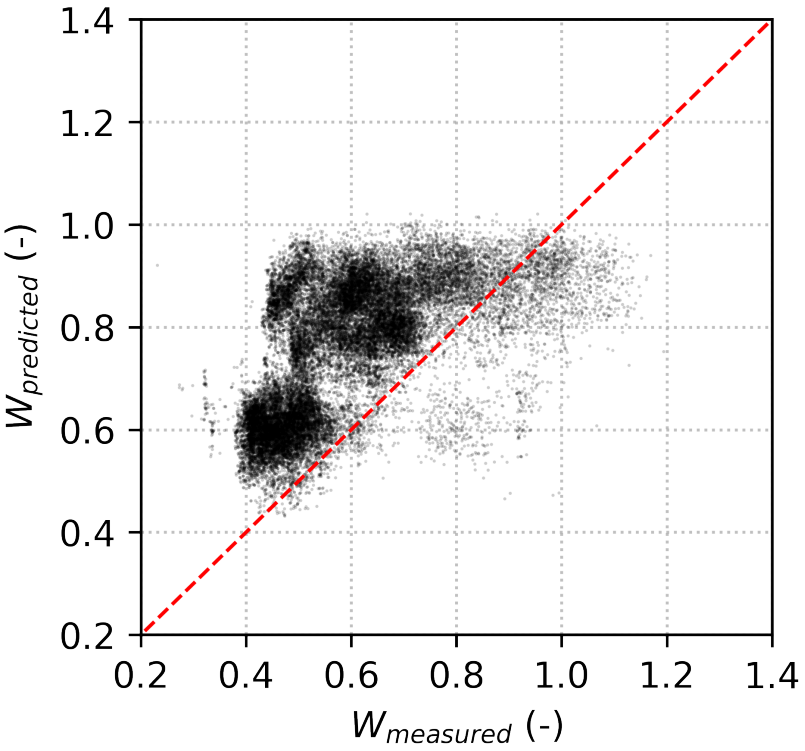

Figure 4.50: Measured Vs. Predicted pumps energy consumption scatter plot (Sep. 25 2021 - Sep. 30 2022)

Figure 4.50 allows for a clearer visualization of the comparison between the actual pump consumption and the value predicted by the machine learning model. From this graph, it can be observed that the majority of data points show a higher predicted consumption compared to the corresponding measured consumption, which is the main reason for the negative slope of the *CUSUM* curve. A smaller portion of points, however, exhibits the opposite trend (below the diagonal line on the graph). These points correspond to the summer periods mentioned earlier. Considering that the values shown represent normalized absorbed powers using the maximum pump consumption value during the training period, it can be noticed that a small portion of points exceeds 1, indicating measured values higher than the maximum power absorption during the training period. This can be easily attributed to the abnormal use of the air treatment system during the summer of 2022.

For the analysed period, it is possible (as in the previous period) to calculate that the modifications to the pump management strategy have led to a reduction in their consumption by -27.78%. This savings has had a significant impact on the overall building consumption during this period, resulting in an energy saving of -5.88%.

As previously mentioned, on May 28^{th}, 2022, a maintenance operation was conducted on the refrigeration units, effectively reducing their energy requirements. Similarly to the previous cases, it is possible to analyze in more detail the period affected by this change, which extends until the end of the analysed period. As mentioned earlier, abnormal use of

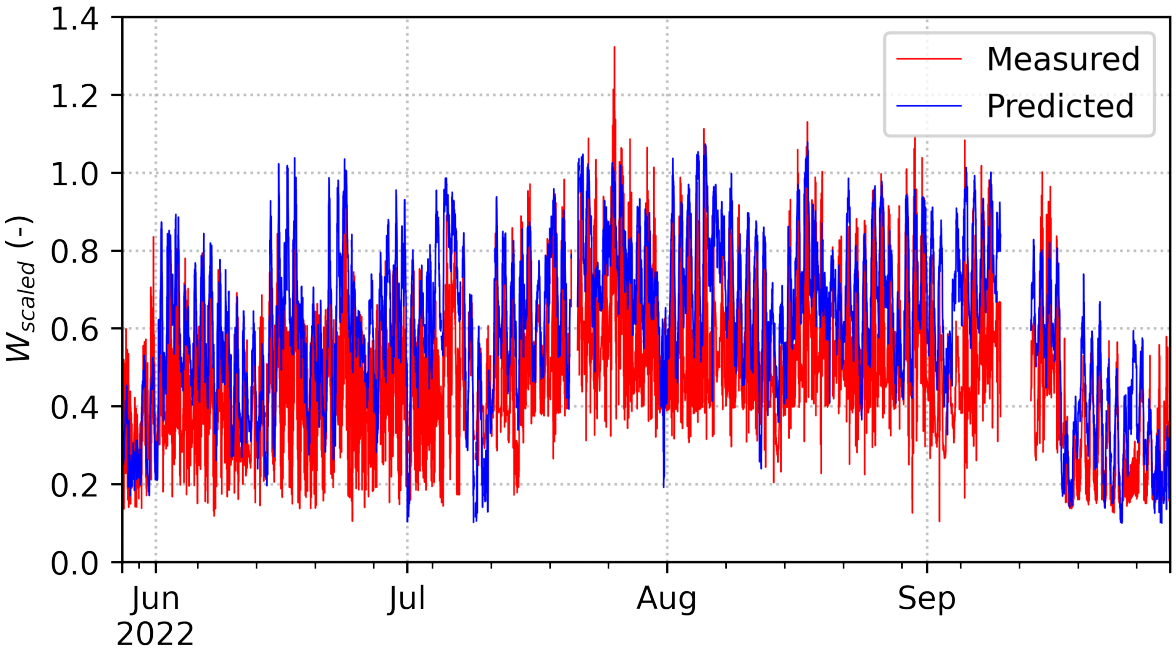

Figure 4.51: Measured Vs. Predicted refrigerators energy consumption line plot (May 28 2022 - Sep. 30 2022)

the air treatment system resulted in significant increases in energy consumption, which were offset by the efficiency operations. Through this analysis, it is therefore only possible to estimate the combined effect of these two events, as it is not possible to estimate the energy savings achieved solely through the maintenance operation.

Figure 4.51 illustrates the measured and predicted consumption curves, as seen in previous cases. It can be easily observed that the largest differences occur during off-peak periods. While the high consumption peaks remain similar, it is noticeable that in actual operation, the refrigeration units are able to operate more frequently at lower loads, especially during times when the refrigeration demand is not particularly high. The maintenance operation has thus allowed for improved performance, particularly for loads that deviate from the nominal load of the refrigerators.

Through figure 4.52, it is easier to compare the two data series. Firstly, it can be observed that the majority of data points are positioned above the diagonal line of the graph, confirming that most of the analysed instances have a measured power value lower than the predicted power. A smaller portion of points is also located below the diagonal. By also referring to Figure 4.51, it can be noticed that the majority of points below the diagonal are caused by periods further away from the summer period (at the beginning and end of the analysed period). As mentioned earlier, this is mainly due to the installation of new air treatment units, activated after the training period, which contributed to increasing the cooling demand of the facility.

Furthermore, it can be noticed that a portion of the analysed instances has measured

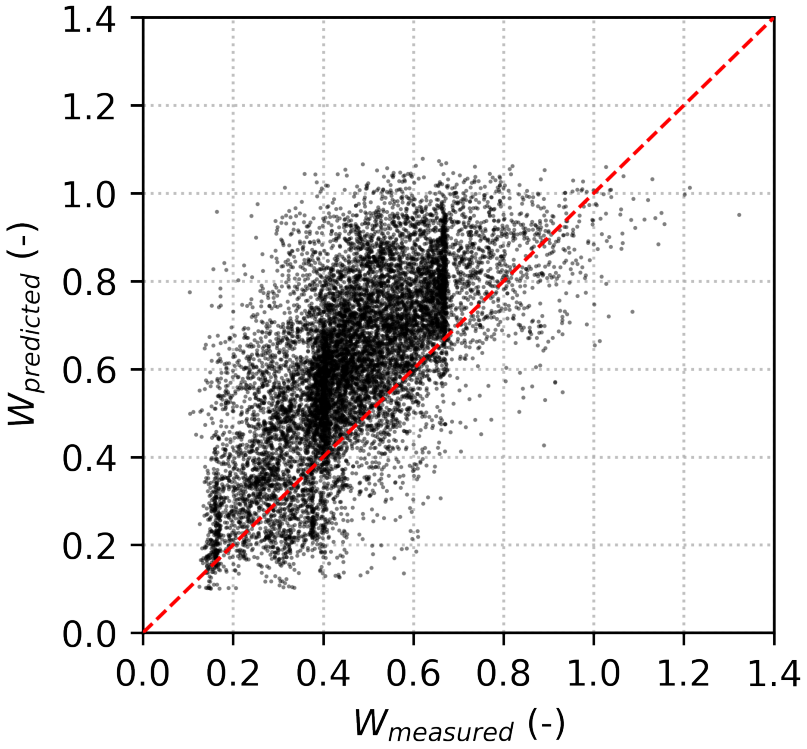

Figure 4.52: Measured Vs. Predicted refrigerators energy consumption scatter plot (May 28 2022 - Sep. 30 2022)

W_{scaled} values exceeding unity, corresponding to instances where the measured electrical power surpasses the maximum power absorbed by the refrigerators during the training period. This is caused by a net increase in the utilization of air treatment units (compared to normal management strategy), resulting in higher loads due to the increased cooling demand during the hottest hours. In support of this, the *CUSUM* curve related to the air handling units (*AHU*) in Figure 4.46 highlights the significant increase in energy consumption of the *AHUs* during the summer months of 2022.

By comparing the measured and predicted values, it can be concluded that the consumption is reduced by -29.70% compared to the predicted consumption. This energy savings is significantly reflected in the overall energy consumption of the entire building, which is reduced by -6.19%. It is important to emphasize that these results were achieved during a period when two opposing trends emerged. On one hand, the energy efficiency generated by the maintenance operations greatly reduced consumption. On the other hand, the increased cooling demand due to the different management of air treatment during the summer (air handling units no longer automatically regulated but manually set to 100% load) led to an increase in consumption. It is therefore reasonable to assume that under normal conditions, the maintenance work would have allowed for even greater energy savings.

Chapter 5

Conclusions

The aim of this research was to explore machine learning-based methodologies for creating tools to monitor the energy consumption of a building. Specifically, the method was designed to be as "systematic" as possible, meaning that adjustments to the procedure required to apply the method to a specific test case were minimized. As a result, the user's experience required to implement the method in their activity's energy consumption analysis procedures or centralized building energy management system was also minimized.

To execute this study, six distinct methodologies were developed, each utilizing different machine learning methods and varying techniques for feature selection and hyperparameters tuning. Some key information emerged regarding the selected features. The outdoor air temperature was found to be the most important factor. While the *Filter Method* place *heating degree days*-related features to the first ranks, the *Wrapper Method* place *ClimateZ* first in all cases where it was applied. A second important feature was the staff in service, which affects the activation of air conditioning systems and the daily and weekly cycle of activities. However, this parameter results less important for the *Filter Method* and for the *Wrapper Method* applied to *Multiple Linear Regression* models, but highly significant for the *Wrapper Method* based on *Artificial Neural Network* models. Finally, the rankings always included calendar features regarding the daily, weekly, and seasonal time scales (e.g., Hour, Weekday, Month). The application of the methodology revealed the general superiority of cases based on *Artificial Neural Networks* compared to those based on *Multiple Linear Regression*, with prediction performance ranging from $R^2 = 0.856$ in the case of Multiple Linear Regression model combined with the Filter Method for the feature selection (**Case 1**), to $R^2 = 0.954$ in the case of Artificial Neural Network combined with the Wrapper Method and Grid Search (**Case 6**).

Pointwise analysis of the deviations between real and predicted energy consumption, indeed, reveals a tendency of **Case 1** to not accurately follow the building's energy demand during daily peaks (highs and lows), a factor that is almost entirely resolved by using

Case 6. This analysis strongly highlights the lack of adaptability to new data of Multiple Linear Regression models, a crucial aspect for the application of a predictive model for energy monitoring purposes.

It is interesting to note how the application of the Wrapper Method for feature selection on the Multiple Linear Regression-based method (**Case 4**) is able to significantly increase the accuracy of the model, while only modestly increasing the computational cost, making it a great compromise between prediction performance and computational cost.

The method that provides the best results is the one obtained through **Case 6**. Therefore, this method was also applied to data related to the electrical power absorbed by individual subsystems of the Heating, Ventilation and Air Conditioning system.

Each system shows presents different features rankings, maintaining *ClimateZ* and calendar feature related to different time scales (daily, weekly, seasonal) at the top ranks. The obtained models show promising results ($R^2 > 0.92$).

The developed models were then used to generate the *CUSUM* control charts, which allowed for observing trends and anomalies in energy consumption. The addition of specific *HVAC* system models made it possible to make it easier to relate the detected anomaly to its cause and identify further anomalies that compensated for each other, thus being undetectable by the model of the whole building electricity demand.

In conclusion, the presented work has allowed obtaining a method to proceed systematically from raw data related to a specific context to the implementation of a complete and functional monitoring method. It is possible to see the work as a map that helps the user to choose the method that represents the optimal compromise between computational cost and prior knowledge, adapting as best as possible to the specific application context. This methodology for energy monitoring can also be easily implemented in a Building Energy Management System, providing an alert and diagnostic method capable of effectively and quickly highlighting anomalies in the energy demands of a building, thus allowing for prompt corrective actions and contributing to the constant search for increasing the energy performance of the building and its systems.

Due to its nature, the method can be easily applied to different application contexts and different energy vectors, as the only requirement is the availability of a sufficient amount of data. Its greatest strength is also its greatest limitation. The intrinsic dependence on the availability of historical data related not only to energy consumption but also to climatic conditions, occupancy, and other factors, makes it difficult to use in contexts without adequate measurement and data handling systems. However, with the ongoing shift towards the Industry 4.0 paradigm and the increasing availability of IoT technologies that are less invasive and more accessible, the presented method is widely applicable as of today and it expected to become even more applicable in the future.

Appendix A

Building energy modelling

Building energy modelling is a widely used technique that allows realizing building digital twins capable to reproduce the energy behaviour of a building. Several tools allow to carry out BEM analysis (e.g., TRNSYS©). In the present study, it was decided to explore the potentiality of open-source software called EnergyPlus™ [104]. EnergyPlus™ is a whole building energy simulation program developed by the U.S. Department of Energy's (DOE) Building Technologies Office (BTO). It allows to model both energy consumption (e.g., heating, cooling, ventilation, lighting and plug and process loads) and water use in buildings. Moreover, there are third-party tools that expand the energy plus potentiality by providing graphical interfaces useful to develop the building model and its plants. Indeed, this procedure could result markedly difficult and prone to human errors if realized directly in EnergyPlus™ environment, depending on the building complexity. To realize the 3D model of the structure, it has been used Sketchup©. To connect the realized model to EnergyPlus™, it has been exploited an open-source software called OpenStudio™ [105]. Through the combination of SketchUp© and OpenStudio™, it has been possible to realize the 3D model of the structure and to set up its plant and system coherently to the real structure.

A.1 Building energy model description

The 3D model (represented in Figure A.1) has been designed trying to replicate the enlargement area of the healthcare facility, which results to be the most energy-intensive area of the building. The model, composed of 610 surfaces (internal and external walls, floors) and 160 windows, has been realized based on the structure planimetry, even though the internal environment does not replicate exactly the real one, ignoring part of the internal walls. Indeed, since the focus was to obtain a good representation of the facility trying to keep the model as simple as possible, part of the internal walls has been not

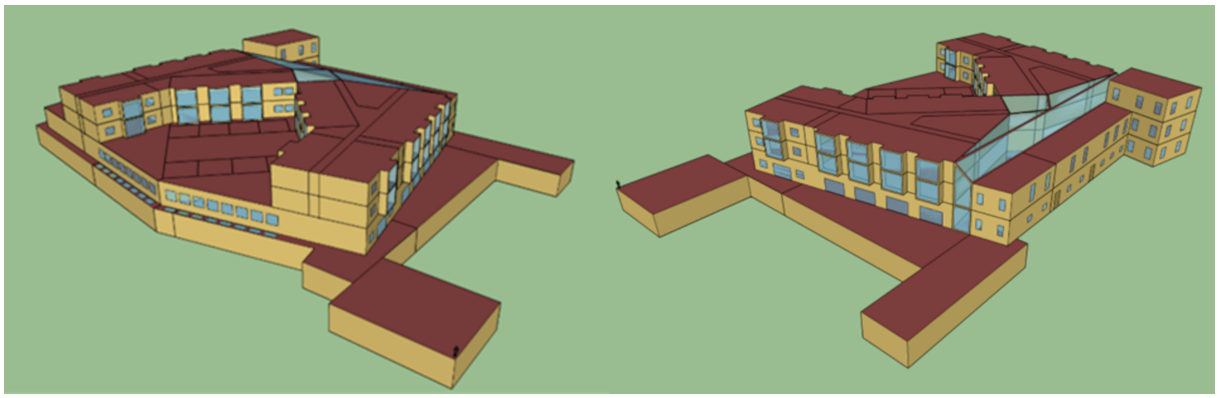

Figure A.1: Multiple views for the 3D model of the healthcare facility under analysis

represented.

A Building Energy Model (**BEM**) must be divided into a finite number of thermal zones, which are collections of spaces having similar space-conditioning requirements, the same heating and cooling set-point, and consequently, they are treated as a single space to carry out the simulation. This subdivision has been carried out considering the activity performed in each area and its internal climate control requirements (it can be seen in Figure A.2). Moreover, there are specific rooms characterized by peculiar needs in terms of internal climate control (e.g., Nuclear Magnetic Resonance, Surgery rooms). These areas have been maintained as separated thermal zones to better replicate their needs. The ambulatories and the inpatient departments ($4'000\ m^3$) are located on the first and second floor of the building, around the hallways branching out from the Entrance/Hall ($3'000\ m^3$). Each of the four considered areas is composed of twenty rooms facing the relative hallway. Considering that they are served by two dedicated Air Handling Units, and the internal climate requirements do not change among the different rooms, the two inpatient departments have been represented as four thermal zones (without internal walls).

The ground floor is composed of several areas which are characterized by completely different activities and internal climate requirements. The hall, the entrance and the commercial area have been represented through three separated thermal zones since they are served by three different air handling units. The central area of the ground floor is dedicated to healthcare activities. The seven surgery rooms ($230\ m^3$) and their central passage area ($1'400\ m^3$) are treated as separate thermal zones because they are served by eight different air handling units. Moreover, since not all the surgery rooms are operative, this solution allows the user to better control their functioning. Another thermal zone has been realized to model the surgery locker room.

The Nuclear Magnetic Resonance ($300\ m^3$) and the Angiography department ($300\ m^3$) equipment presents more restrictive requirements in terms of temperature and relative humidity in comparison to other areas. Moreover, the real building presents two dedicated

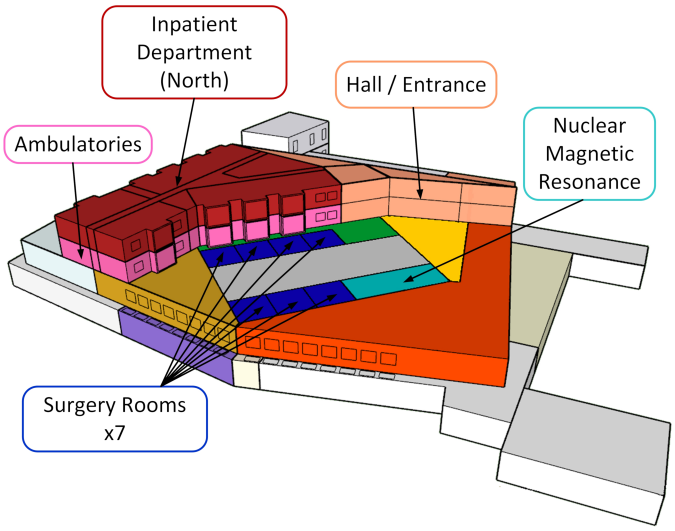

Figure A.2: Thermal zones subdivision for the realized 3D model of the healthcare facility

air handling units which serve these areas. Consequently, they are modelled as two thermal zones. Following the same logic, also Diagnostic and Ophthalmology departments (3′000 m^3) have been modelled separately. A single thermal zone has been realized to model all the offices since they are characterized by similar indoor climate conditions and people occupancy schedules. The basement is mainly dedicated to the technical rooms, except a medicine storage room, laboratory and two locker rooms. Consequently, it has been divided into four thermal zones to better model the different activities carried out in each area. A previous study carried out by the authors on the same building [18] highlights that the external climate conditions, particularly the outdoor air temperature, are the main factors that influence electrical and thermal energy consumption.

Consequently, considering that the core activities do not have a marked influence on global consumption, and since detailed data about the core activities carried out within the building are not available, the activity schedule and related energy needs have been assigned using standard hospital settings and schedules available on the used software, particularly through the Building Component Library (BCL), which allows the user to access to standard characteristics of hospitals in terms of intensity of medical exams and surgery activity, people occupancy and other operative details. Regarding the internal climate conditions, it has been set by creating several schedules to fulfil the requirements of the different thermal zones. Through the graphical interface provided by OpenStudio, it has been possible to model the plants and systems installed within the healthcare building. As regards the plants and systems realization, it has been made two fundamental considerations.

First of all, some plant components are meant to work as a backup (one hot water boiler, one refrigerator and its two cooling towers) that are not included in the realized model. Moreover, the realized 3D model does not perfectly replicate the real structure

and many areas have been merged. Consequently, it was not possible to recreate all the air handling units, but they are merged to obtain larger units that cover the entire requirements of the structure in terms of ventilation and air conditioning. The realized plants and systems can be graphically represented with the scheme shown in Figure **??**. They have been realized to reproduce the existent plants and systems of the healthcare facility, and they are listed below:

- Three hot water boilers (Thermal power of $3x700\ kW_{th}$)
- Two water-cooled refrigerators (Electrical power of $2x320\ kW_{el}$)
- Four evaporative cooling towers (Cooling capacity of $4x760\ kW_{cool}$)
- Fifteen Air Handling Units

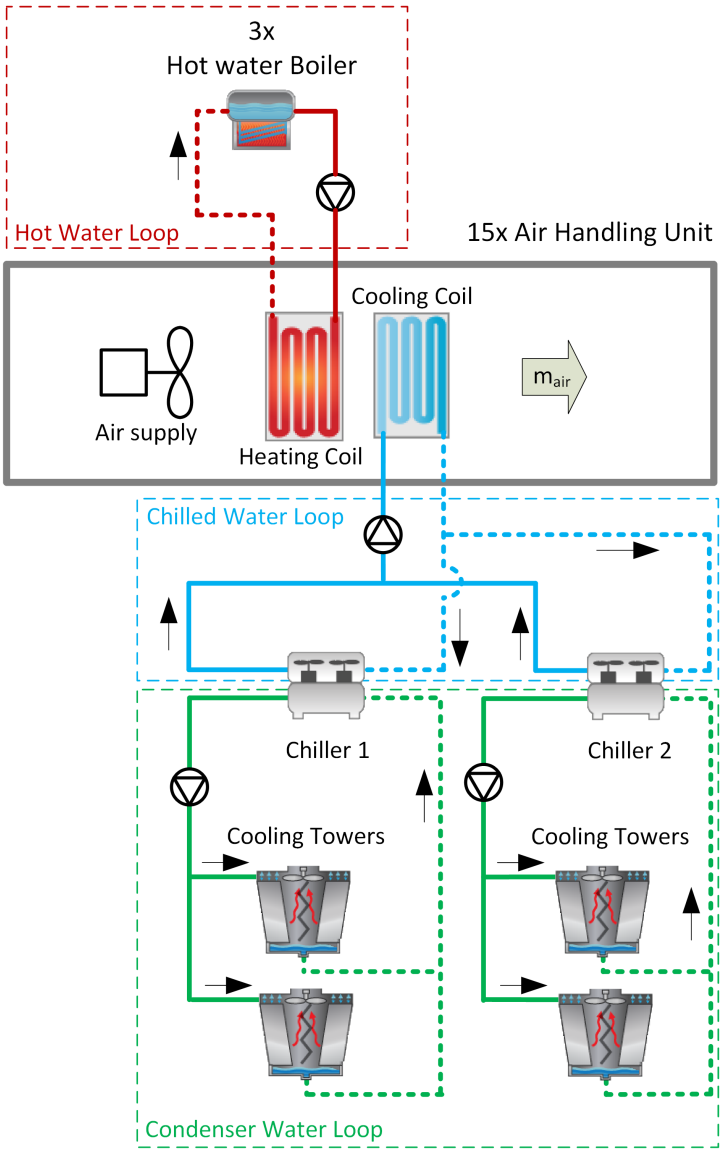

Figure A.3: HVAC system scheme for the proposed energy model

The air handling units are modelled in order to lead the outdoor air to two air-to-water heat exchangers (cooling and heating coils). The handled air is then led to the various

building areas to carry out the air conditioning. Each thermal zone has been characterized by an appropriate temperature set-point schedule based on the real requirements of the healthcare facility and, consequently, the water flow rate needed by the heating and cooling coils varies aiming to maintain the indoor air temperature to the desired value.

The hot water boilers are modelled as a part of the Hot Water Loop shown in Figure A.3. The water coming from the inlet water system is divided between the operative boilers, which are responsible to maintain the outlet water temperature to a set-point value equal to 75 $°C$. The hot water temperature set-point has been chosen according to the real building plant control logic. Since the hot water flow rate is regulated to match the healthcare requirements, the boilers are activated in sequence to maintain the water temperature to the desired set-point. Then, the hot water is led to the heating coils on the already described air handling units, which are responsible to perform the air-water heat exchange.

The water refrigeration is carried out by means of two water-cooled refrigerators (called chillers in Figure A.3) The whole refrigeration system can be divided into two separate loops. The Condenser Water Loop is related to the water that circulates through the condenser, responsible to remove the heat from the refrigerant vapour. Figure 3 shows that each refrigerator is part of a dedicated loop that brings cooling water to the related evaporative cooling towers, which attempt to maintain their outlet water temperature at (or below) the set-point of 25 $°C$. As regards the chilled water loop, the refrigerators aim to carry out the water cooling to a water temperature set-point of 7.5 $°C$, regulating their load based on it. Once the water refrigeration has been carried out, the chilled water is led to all the cooling coils of the active air handling unit to perform the air conditioning.

A BEM model must receive in input also the climate data specific of the area in which the building is located. EnergyPlus requires an energy plus weather file that must contain hourly weather data for the entire simulation period. Since data about energy and water consumption related to 2019 was available, climate data related to the same year has been collected. "*Consorzio LaMMa*" provided the climate data related to the year 2019. In particular, the parameters that have been used as input for the realized model are external temperature, relative humidity, ambient pressure, global, diffusive and normal solar radiation, wind speed, wind direction and precipitation. Once the model and the input data were ready, the simulation parameters have been set. The *ShadowCalculation* is used to control the details of the solar, shadowing and daylighting models available on the software. There are two basic methods available for the calculations. Shadowing calculations are performed over a period of twenty days by default settings. However, this value can be pivotal to determine the amount of solar radiation entering the building and, influencing the amount of cooling or heating load needed to maintain the building

internal climate requirements. Even though the weather file contains data related to solar radiation, the internal calculation of the sun position will govern how it affects the various parts of the building. The proposed model performs shadow calculation every two hours to keep the calculation cost acceptable and obtain a good representation of the effect of solar radiation on the building.

EnergyPlus let the user choose the model used for surface convection at the inside face of all the heat transfer surfaces in the model. The selected algorithm is called *TARP* and it correlates the heat transfer coefficient to the temperature difference for various orientations. As regards the outside convection, for the external walls it has been used the *DOE-2* algorithm, which is derived from field measurements on rough surfaces. Contrarily, for the windows, it has been set the *MoWiTT* algorithm, that is similar to the previous one but derived from measurement on smooth surfaces. The heat balance algorithm regards what type of heat and moisture transfer algorithm will be used for calculating the building's surfaces behaviour. In this case, the *ConductionTransferFunction* was selected. This algorithm does not take into account the moisture storage or diffusion in the construction elements. The Zone air heat balance algorithm is used to calculate zone air temperatures and humidity ratios. The selected method uses the third-order finite difference approximation to solve the zone air energy and moisture balance equations.

A study has been conducted using EnergyPlus aiming to evaluate the impact of the time-step size on the building energy model results [106]. The study highlights that for test cases characterized by HVAC systems that play a relevant role in the total building energy consumption, the time-step size should be lower than 10 minutes to appropriately simulate the building behaviours. The model presented in this study has been set using a time-step size of one minute, which is the minimum time-step size that can be set on the software. The presented simulations have been carried out on a computer with an Intel® Core™ i5-8265U CPU @ 1.60 GHz and 12 GB of installed RAM. The simulations required a computational time of 63 minutes.

A.2 Model results

to provide an insight into predictive capabilities of the realized model, the results of the analysis will be compared with the real monthly requirements of energy and water demands. Subsequently, since more detailed data about real electrical energy consumption were available, the results of the simulation will be compared with the real electrical energy consumption of the healthcare facility. It is important to remark that the presented model is not intended to perfectly reproduce the real building in terms of energy demand. The aim is that to create a model able to qualitatively reproduce the healthcare energy needs, allowing to evaluate the effect that any modification of the plants and systems or to their management strategy have on the global healthcare energy requirements. The building thermal energy requirements are mainly due to the heating and air conditioning needs and to domestic hot water generation. Both requirements are fulfilled through the already described hot water boilers. The air handling needs hot water during both summer and winter operations.

During the winter, the air is pre-heated through a pre-heating coil, which is invested by airflow before the latter pass through the fan. The air temperature is then regulated through post-heating coils to match the internal temperature set point. During the summer period, conversely, the pre-heating coil is generally deactivated, while a cooling coil carries out the air refrigeration and humidity control. The air must be post-heated to match the internal temperature needs. Consequently, the building presents a non-negligible requirement of thermal energy during the summer period.

As Figure A.4a shows, the thermal energy demand presents a minimum during the summer period. Then, it gradually increases in correspondence to the diminishing of monthly averaged external temperature toward the winter period. The monthly consumption calculated through the Energy Plus model seems to follow the same behaviour but systematically tends to underestimate the thermal energy consumption during the winter, while overestimating it during the summer period. This behaviour can be explained considering that the available data about the structure and its plants and systems are not enough to perfectly represents the building needs. For thermal energy and water, the available data are the monthly consumption.

Consequently, it has been not possible to draw conclusions about their daily behaviour to perform a detailed numerical calibration. Moreover, the healthcare facility started its operation in 2018, resulting in a building that is prone to several changes in terms of spaces employment and activities carried out within the structure.

Figure A.4b shows the monthly building water consumption. The water needs are mainly caused by Domestic hot water use. Nevertheless, it was estimated that almost 30% of the yearly building water needs can be attributed to the evaporative tower functioning.

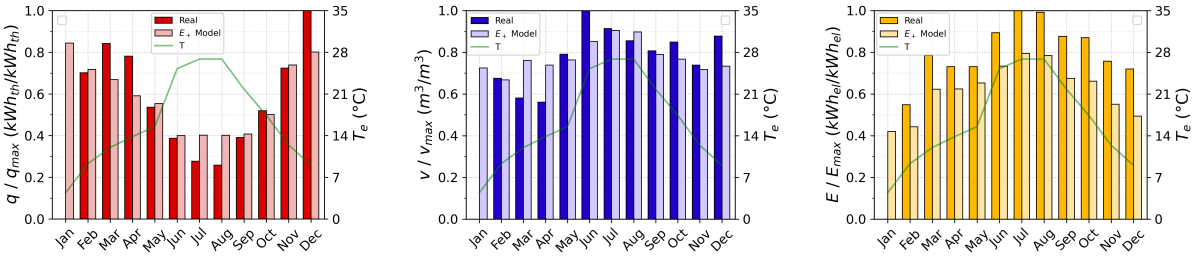

(a) Thermal energy (b) Water (c) Electricity

Figure A.4: Monthly building energy and water needs: Measures VS. Simulation

There are no available real data that allow to carry out detailed analyses on the building water demand, but it can be noticed a correlation between the external temperature and the water consumption. Indeed, during the summer, the cooling system exploitation increase to fulfil air handling purposes. Consequently, the evaporative towers exploitation tends to increase, leading to a growth in water consumption. The water needs during the rest of the year seem to have a weak correlation with the climate condition. This behaviour is compatible with the physiological variations of domestic hot water requirements.

Figure A.4c shows the monthly electrical energy consumption of the Healthcare Facility. Moreover, the monthly averaged outdoor temperature has been reported (green line) to graphically highlight the strong dependency of the electric energy demand on climate conditions. The energy demand presents marked peaks during the summer and reaches its minimum during the winter. This is due to the higher exploitation of refrigerators and HVAC systems during the summer. June and July result to be the most electric energy-intensive months, coherently with the high level of monthly averaged external temperature. August, despite the high outdoor air temperature, presents lower electricity consumption in comparison to the other summer months. This is due to the reduction of the healthcare core activities, which let to switch off unused departments and the related HVAC systems. Nevertheless, during the rest of the year, the monthly thermal energy consumption maintains itself to values of q/q_{max} between 0.5 and 0.7. Looking at the simulation results about the electrical energy consumption, they follow the same behaviour found in real data.

Figure A.5 and Figure A.6 compare the real instantaneous absorbed electrical power with the simulation results. Figure A.6 shows a comparison of the power demand behaviour over five days of May 2019. The two lines highlight the similarity between the power demand behaviours of the real building and the simulation. Figure A.5 shows a comparison of the instantaneous total electric power demand of the analysed building through a scatter plot. Each point compares the real and the simulated electrical power demand, and its colour change depending on the external temperature. The diagonal (black dashed line) represents the positions in which every point should be located if the realized model can

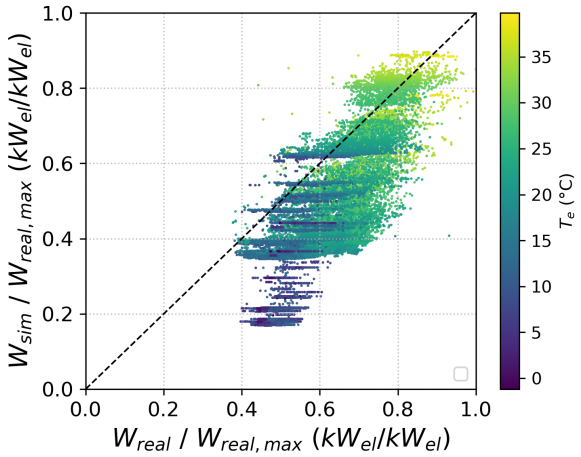

Figure A.5: Comparison of real and simulated instantaneous electrical power demands

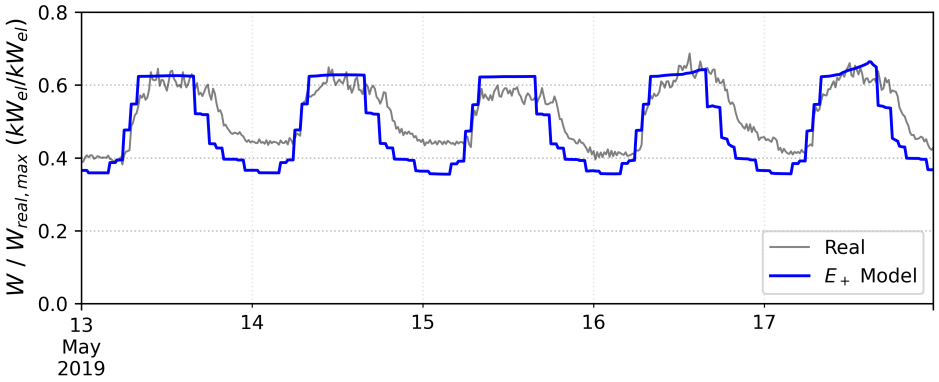

Figure A.6: Comparison between real and simulated building electrical power demand (from May 13^{th} until May $17th$)

perfectly predict electricity consumption.

As Figure A.5 depicts, the simulation performs a good prediction during the energy-intensive period, which generally corresponds to high external temperature periods. Contrarily, tends to underestimate the electric energy consumption during the rest of the year. It is apparent that different time-steps characterized by similar outdoor air temperature values present, a value of simulated electrical power corresponds to a range of real electrical power demand. This behaviour indicates that some aspects that influence the energy demand are not considered in the realized BEM model. Since detailed data about core activities schedules and people occupancy are not available, these have been modelled referring to standard hospital schedules provided by the exploited software. Consequently, the simulation performs better when the HVAC system strongly influences the total building electrical energy needs. Nevertheless, the lower electricity consumption showed by simulation results in comparison to the real building ones could highlight inefficiencies in plant and systems operation or control strategies.

A.3 Techno-economic assessment of alternative scenarios

A.3.1 Cooling Towers

Through the calibration procedure described in the previous sections, it has been possible to obtain a dynamic model of the healthcare facility that qualitatively represents the real building. Although the absence of more precise data regarding the activities of the healthcare facility did not allow the authors to perfectly model the building and its plants and systems, the obtained model allowed to perform some further simulations to evaluate alternative scenarios. In particular, the present work is aimed to evaluate several scenarios in which alternative technologies have been implemented.

- ECT_c (Current Evaporative Cooling Towers): It represents the current building plant layout
- ECT_{new} (New Evaporative Cooling Towers): Substitution of the evaporative cooling towers with more efficient ones of the same typology
- DCT (Dry Cooling Towers): Substitution of the evaporative cooling towers with dry cooling towers

ECT_{new} and DCT imply the substitution of the existent evaporative cooling towers. The first case (ECT_{new}) will evaluate the technical and economic feasibility of the substitution of the current cooling towers with more modern and efficient ones. Table A.1 reports the main characteristic of the proposed cooling tower. It is apparent the sensible reduction of electrical power needed for the air ventilation (about 60%), maintaining the cooling capacity of the current system. The second case (DCT) involves the substitution of the cooling towers with dry cooling towers. The dry cooling towers are heat exchangers, which carry out the water cooling by using outdoor air as a secondary fluid. The working fluid (water) is refrigerated in a closed circuit which let to avoid water evaporation, leading to a decrease of the building water needs in comparison to an evaporative cooling plant.

Table A.1: Current VS. New cooling towers characteristics

	Unit	$\boldsymbol{ECT_c}$	$\boldsymbol{ECT_{new}}$	$\boldsymbol{DCT}$
Cooling capacity	kW_{cool}	760	762	729
Water flow rate	L/s	36.3	36.4	–
Fan nominal electric power	kW_{el}	18.5	7.5	56

Removing the contribution of the water evaporation from the cooling process causes an increase in the air mass flow rate needed by the cooling tower itself, resulting in an evident increase in electrical energy consumption. The new components have been selected aiming to maintain the total cooling capacity of the plant. The four evaporative cooling towers (two for each refrigerator) have been substituted by four dry cooling towers, maintaining the cooling capacity of the current plant (4×760 kW). Since the cooling towers are responsible for a non-negligible amount of total yearly water consumption and since they can influence the efficiency of the chiller, the present analysis will be focussed on the evaluation of the effect of the new plant configuration on the water and electric energy consumptions.

First, the two simulations using evaporative cooling tower technology have very similar water consumption. The new evaporative towers produce annual water savings of about 0.35%. As Table A.1 reports, the new evaporative cooling tower presents the same cooling capacity as the current system. Conversely, the electrical power needed by the fan to ensure the air circulation inside the tower results markedly lower. Indeed, the ECT_{new} test case presents yearly electrical energy consumptions related to the cooling towers reduced by 59.87% in comparison to the ones obtained in the ECT_c simulation. Focussing on the dry cooling towers (DCT), the simulation presents a drastic increase in building electric energy requirements (+14.81%). This behaviour is mainly due to the marked increase of electrical power needed by the dry cooling tower fans in comparison to the ECT test cases (see Table A.1).

Conversely, the water use related to the cooling towers is completely erased, resulting in a -10.61% of the total building water consumption. By comparing the two cases with the evaporative cooling towers, they maintain similar behaviour in terms of electric energy consumptions. The new evaporative cooling towers let to obtain an almost uniform reduction in electrical energy consumption. Instead, as expected, the simulation with the dry cooling tower results in higher consumptions in comparison to the ones obtained with the evaporative cooling towers.

The difference between the two configurations become more marked during the summer period, in which the exploitation of the cooling plant results more intensive. Figure A.7 depicts the scatter plot of the instantaneous electrical energy consumption of the entire healthcare facility by comparing the current plant configuration (ECT_c), with the two presented alternative cases (ECT_{new}, DCT). These plots confirm that case ECT_{new} allows obtaining a general reduction in electrical energy consumptions because it presents all the markers below the diagonal (black dashed line). Contrarily, the dry cooling towers cause the already explained electrical energy increase. Nevertheless, the drastic increase in electrical energy consumption can be due to several different phenomena.

First of all, the nominal electric power of the cooling tower fans results significantly higher for the dry cooling tower case in comparison to the evaporative one. Indeed, the former requires a greater airflow to ensure water refrigeration, leading to a marked increase of electrical energy requirements of the cooling tower fans. Moving from winter to summer the difference between case *DCT* and the other simulations become more marked in terms of energy demand, a phenomenon related to the increase of *HVAC* systems exploitation with the external temperature raise, a behaviour that is highlighted in Figure A.7b.

This figure shows four separated areas which seem to be mainly divided by temperature level. This behaviour can be explained by considering that the dry cooling tower model present inside the *EnergyPlus* software is not controlled as a variable-speed fan, but it can be regulated only through *ON/OFF* control logic. Consequently, the four areas highlighted in the figure are due to the simultaneous activation of one or more cooling towers. Particularly, it seems that the increase in electrical power demand characteristic of a particular zone corresponds to a coherent increase of the outdoor temperature. Indeed, the zone characterized by the highest electrical power demand corresponds to the highest temperature period which requires massive exploitation of the cooling towers, while the coldest temperature can be found in correspondence to the lower electrical power demand area.

Focusing on the comparison between ECT_c and *DCT* simulation, Figure A.7 shows a different phenomenon that plays a role in the general increase of electric energy consumptions. The lower efficiency of the dry cooling towers leads to higher tower exploitation during the whole year. Figure A.8a shows that the evaporative cooling towers often outperforms the dry one, bringing the water to temperatures significantly lower in comparison to the one reached with the dry cooling towers. This behaviour becomes more evident with the increase of the outdoor air temperature.

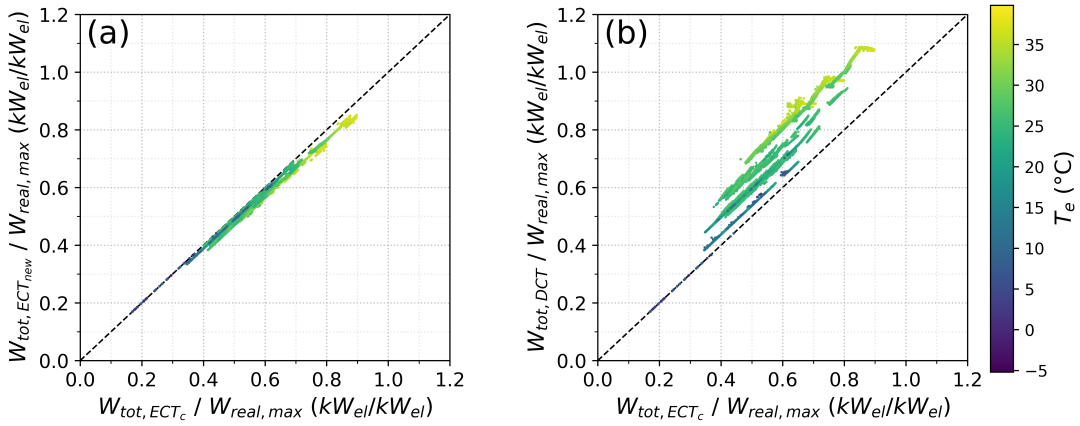

Figure A.7: Comparison between instantaneous facility power demand between *ECT* and *DCT* simulations

The increase of the Evaporator Inlet Water Temperature leads to performance issues in

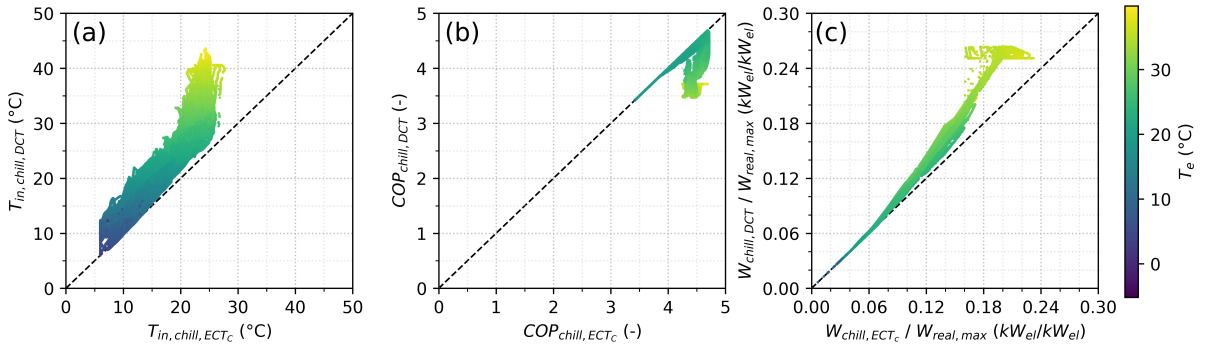

Figure A.8: Comparison between ECT_c and DCT simulations: Main chiller inlet water temperature coming from the cooling towers (a), Instantaneous COP of the main chiller (b), Instantaneous electrical power demand of the main chiller (c).

the chillers functioning. Figure A.8b shows that the COP of the main chiller served by dry towers results lower than the evaporative tower test case for high external temperature.

For lower external temperatures, the two chillers seem to work in the same conditions because the cooling towers let to match the inlet water temperature set-point. The behaviours analysed above can be resumed by observing Figure A.8c, in which the instantaneous absorbed power of the main chiller has been reported. As expected, the chiller electrical power demand is comparable between the two analysed cases in case of low external temperature. When the outdoor climate reach typical summer values, the absorbed power of the dry cooling tower test case start to become greater in comparison to the one found in the evaporative cooling tower test case. This difference results magnified with the increase of the external temperature.

In other words, the substitution of the current evaporative cooling towers with dry cooling towers leads to two different effects on the total building electricity consumption. The main one is the direct consumption increase due to the greater nominal power characteristic of the dry cooling towers. It is responsible for about 90% of the total electrical energy consumption increase. Moreover, there is an indirect effect in which the lower cooling efficiency of the dry cooling towers has repercussions on the chiller functioning. About 10% of the total building electrical energy consumption is caused by the COP decrease of the chillers due to higher water temperature at the evaporator inlet. It is important to notice that, if the first effect can be roughly estimated with relatively simple calculations once the electrical power demand of the current cooling tower is known (e.g., through dedicated measures), the indirect effect on the chiller can be estimated only through dynamic building energy models.

An economic analysis has been conducted using the NPV method (Net Present Value), to assess the economic feasibility of the proposed solution. The NPV method is a cumulative indicator updated on a yearly basis, which considers the investment and its effect along with the whole plant operative life. The calculation method has been reported

below:

$$NPV = \sum_{j=0}^{n} \left(\frac{CF_j}{(i+1)^j} \right) - I_0 \tag{A.1}$$

where n is the operative life of the component (expressed in number of years), CF_j is the cash flow cases I_0 is the investment and i is the annual discount rate, which has been considered equal to 7%.

The $CAPEX$ (I_0) of the two cooling tower typologies has been obtained directly from the provider, while the $OPEX$ has been estimated based on the current scenario, considering the ordinary and extraordinary maintenance and also the costs derived from the water handling and electricity consumption. The cash flow is then obtained as the difference between the operational expenditure of the current configuration and that of the new one. Figure 11 depicts the NPV behaviours of the two proposed solutions. The dry cooling tower solution presents a diminishing NPV. Savings related to water consumption and reduced maintenance costs do not compensate for the increase in electricity consumption, which causes an overall increase in $OPEX$ for the dry cooling tower scenario, making the investment not worthwhile.

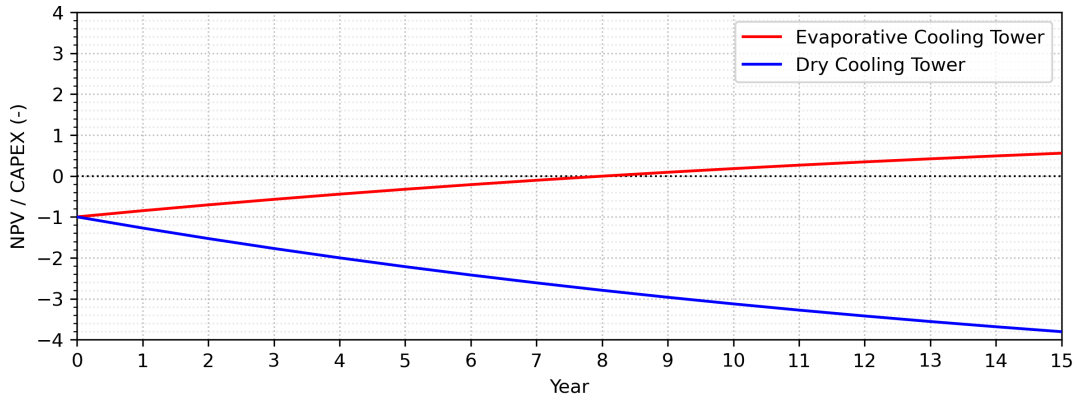

Figure A.9: Net Present Value related to the new evaporative and adiabatic cooling towers

The ECT_{new} scenario presents a positive cash flow, given that the new evaporative cooling towers allow to maintain the plant cooling capacity and obtain a reduction in terms of electric energy and water needed for their functioning. The solution presents a PBP of about 7 years with a profit index equal to 70% of the initial investment. The Building Energy Modelling technique allows to carry out preliminary evaluations of plant and systems change. The substitution components like evaporative cooling towers influences the global consumption not only in terms of the electrical energy needs of the specific component itself but also by influencing the operative conditions of connected equipment (e.g., refrigerators). The "side-effect" evaluation of plants and systems modification can be evaluated only by means of dynamic simulations, making them a key technique to increase

the building energy efficiency.

A.3.2 Photovoltaic generation system

BEM methodology can be applied to assess the impact of different energy generation systems on the energy demand of a building. Since the available climate input data include instantaneous solar radiation, the model has also been used to carry out a techno-economic analysis aimed to determine the optimal configuration of a photovoltaic generation system.

A study conducted by the International Renewable Energy Agency called "Power Generation Cost in 2019 [107] reports detailed data regarding the installation and maintenance costs of solar PV panels in different countries from 2010 to 2019, showing a general and marked decrease in installation costs in recent years Specifically, the cost of installation for commercial activity in Italy was around $4'500$ €$/kW_p$ in 2010. Subsequently, it has gradually decreased, reaching the current average cost of 940 €$/kW_p$. Consequently, photovoltaic systems for power generation could be a convenient investment for activities characterized by high electricity consumption. However, the available surface area of the specific site can be a key factor regarding the technical and economic feasibility of the installation. The Building Energy Model has been exploited to evaluate the performance of a photovoltaic system installed on the roof of the building. First of all, a study has been carried out to establish the optimal tilt and azimuthal angles of the panels to maximize the energy production of the system. The orientation of the panels has been represented through two different angles: Azimuth angle (α) is the angle between the current panel orientation and the south. Tilt angle (θ) is the angle between the horizontal position and the surface of the panel.

Table A.2: Photovoltaic panel characteristics and costs

	Unit	**Low efficiency**	**High efficiency**
Nominal Power	kW_p/m^2	0.1959	0.2348
Panel Efficiency	–	0.1830	0.2210
Installation Cost	€$/kW_p$	979.68	$1'300.00$
Operation and Maintenance Cost	€$/kW/y$	8.50	8.50
Life	y	25	25

Azimuth optimization has been conducted by placing a single panel on the roof of the model. The first orientation was chosen consistently with most real installations, taking the south direction for the azimuth angle ($\alpha = 0°$) and 30° for the tilt angle ($\theta =$

30°). Then, α has been varied around the south direction to find the optimal azimuthal orientation. The results show that the optimal orientation occurs with south-facing panels tilted by 31 degrees.

Once the optimal panel orientation was found, the model has been exploited to evaluate the economic feasibility of different photovoltaic systems. Two different monocrystalline photovoltaic panels have been selected from those available on the market. The objective was to compare two different systems, one consisting of low-efficiency panels and the other consisting of high-efficiency panels that are more expensive than the former. Moreover, to consider the performance decay of the installed generation system, it has been considered an efficiency reduction equal to 1% of the panel efficiency every year. In Table A.2 the main characteristics of the two selected typologies are resumed.

The present study aims to evaluate the economic feasibility of the generation systems and to find the optimal solution for the healthcare facility under analysis. To this end, several simulations have been carried out varying the total installed surface from 0 to 15′000 m^2 to find the economically optimal configuration. It is important to consider that the available surface of the healthcare facility reaches a maximum of 1′000 m^2. Nevertheless, finding the optimal surface of the generation system let to obtain further knowledge about the structure needs, useful to evaluate future investments. Three economic parameters have been used to perform the economic feasibility evaluation of the considered PV systems. The first one is the NPV, already defined in Equation A.1. In this scenario, the Cash Flow is composed by:

- The costs avoided due to the self-consumption of electricity produced by the power generation system
- The income derived from the sale of the excess electricity that the building is unable to use
- The annual operating and maintenance costs of the photovoltaic generation systems, which are proportional to the size of the system itself (see Table A.2)

The second one is the Payback Period (PBP), which represents the amount of time needed to recover the cost of an investment. The third parameter is called "Profit Index" represents the relationship between costs and benefits of a proposed investment, is defined as:

$$PI = \frac{NPV_{25y}}{I_0} \tag{A.2}$$

The three parameters have been calculated for different employed surfaces between 100 and 15′000 m^2. Figure A.10a a reports the NPV at the end of panel life cycle (25

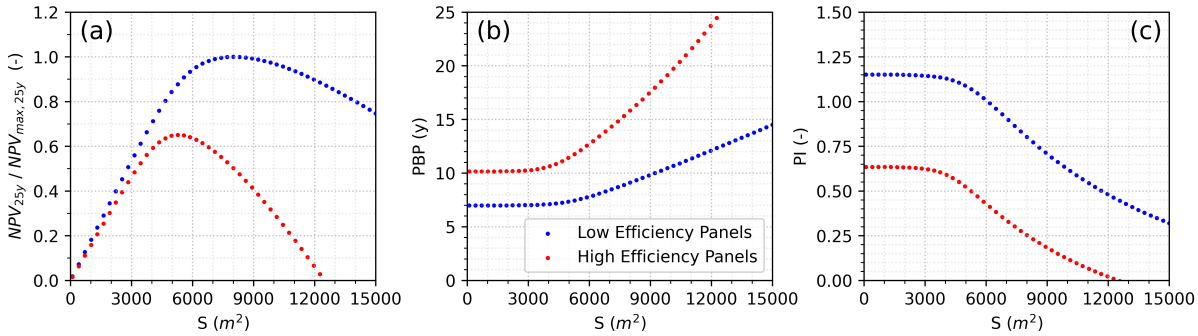

Figure A.10: Economic parameters as a function of panel total installed surface

years) behaviour as a function of the employed surface for both panel typologies. The low-efficiency panel presents a NPV similar to the high-efficiency for surfaces lower than 3′000 m^2. After this point, the high-efficiency panel reach its maximum NPV with a total installed surface of 5′000 m^2. Over this surface value, the high-efficiency panel start to decrease its NPV, while the low-cost one keeps increasing its NPV until the surface value of 7′500 m^2. By observing the life-cycle end NPV only, the optimal surface seems to be 7′500 m^2, which allow obtaining higher benefits in terms of NPV. Nonetheless, higher employed surfaces lead to greater Capital Expenditures, and consequently to a larger payback period (see Figure A.10b). Consequently, the real benefits derived from the analysed investments can be summarized by observing Figure A.10c, which depicts the Profit Index behaviour over the system surface. As an example, looking at the low-efficiency panels (blue dotted line in Figure A.10) it can be noticed that both PBP and PI turn out to be almost constant until the installed surface reaches 3′000 m^2.

This behaviour can be explained by considering the different contributions that self-consumption and excess energy production have on the cash flow. Looking at Figure A.11 it can be noticed that 3′000 m^2 is the maximum installed surface that allows for full self-consumption electricity produced by the PV generation system.

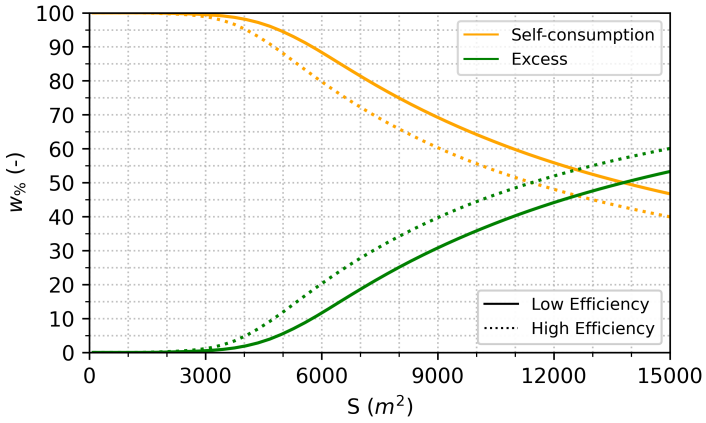

Figure A.11: Electrical energy distribution between self-consumption and surplus as a function of installed surface

Larger surfaces imply that energy production becomes greater than the healthcare facility energy demand. Figure A.11 divide energy production into two different classes, Self-consumption and Excess. It should be noted that their impact on the resulting cash flow is very different. On one hand, self-consumption results in a reduction of electricity purchased from the energy provider, increasing the cash flow as full-price electrical energy.

On the other hand, the electricity production that exceeds the building's energy demand must be sold to the energy supplier at a significantly lower price. As a result, self-consumed electricity has a greater impact on cash flow than excess electricity. Based on these premises, it can be noticed that in Figure A.10c the *PI* related to the low-efficiency panels starts from a value of 1.15 with a photovoltaic system surface equal to 100 m^2. Then, it decreases by 0.65% from 100 until 3′000 m^2, when the photovoltaic system electricity production becomes greater than the building demand. In this range of surfaces, both *CAPEX* and *CF* growth proportionally to the installed surface.

Then, a sharp decrease takes place, leading the *PI* to a value equal to 27.8% of the initial one, corresponding to a surface of 15′000 m^2. As already said, because of its lower economic value, surplus electrical energy has a lower impact on cash flow in comparison to self-consumed one. Consequently, an increase of excess electrical energy will lead the *CAPEX* to increase more sharply than the annual cash flow, causing a progressive decrease of *PI* and an increase in *PBP*. All the analysed parameters highlight the unfeasibility of the high-efficiency panels installation. Indeed, the gain in energy production efficiency does not compensate for the related *CAPEX* increase, leading to lower *NPV* and *PI*, corresponding to higher *PBP*. On the other hand, the low-cost panels present a range of surfaces that are optimal for maximizing the economic benefits.

Surfaces between 4′000 m^2 and 6′000 m^2 allow obtaining an acceptable *NPV* and maintaining low *PBP* and high *PI*, granting the economic feasibility of the proposed investment. Moreover, choosing a surface lower than the optimal one will result in a feasible investment because the payback and the profit index do not change in comparison to the optimal ones. Consequently, the search for the optimal surface of the photovoltaic generation system act as an upper limit for the installed surface and, consequently, for the installed power.

A.4 Discussions

The aim of continuously improving the energy efficiency of a building can therefore be divided into two broad macro-categories. On the one hand, it is necessary to optimise the management strategies of the building's energy systems. On the other hand, there is a need to evaluate solutions for modifying the layout of the building's systems. Often new

solutions cannot be easily tested on the facility due to the needs of the activities carried out within it (e.g., service continuity). In addition, the implementation of these solutions can be costly. In this context, building energy modelling is an important instrument. It makes it possible to carry out simulations aimed at preliminarily estimating the benefits that a given plant modification could bring to the energy demand of the building, making it possible to preliminarily exclude solutions or modifications that are not viable.

Healthcare facilities are activities with peculiar needs in terms of energy demand and service continuity. This makes it impossible to carry out experiments aimed at improving the energy efficiency of buildings. Consequently, healthcare facilities and other similar activities have several advantages in using alternative methods to explore different strategies to pursue continuous improvement of energy efficiency. Building Energy Modelling is a valid way to assess building energy efficiency and to select cost-effective actions to optimise the energy requirements of buildings. However, these techniques have limitations. Building a model requires extensive knowledge of the structure and its systems. Detailed information about the building envelope characteristics is necessary to accurately reproduce the interaction between outdoor and indoor areas, as well as the components of the systems require accuracy in modelling their operational behaviour.

The presented study provides an energy model able to qualitatively reproduce the system and plants requirements of the building under analysis. The numerical results have been compared with real data about energy and water consumption related to the year 2019, highlighting that the behaviour of the facility water and energy demands results coherent with the ones which characterize the real building. Then, the realized building energy model has been exploited to evaluate the effects that two different systems changes had on building energy needs. The first involves the substitution of the evaporative towers installed on the building roof. Since the building presents high electric energy and water consumption, they have been considered two candidates:

- Dry cooling towers installation led to cancel the cooling water needs but cause a non-negligible increase of the electrical energy requirements.
- Modern evaporative cooling towers, contrarily, lead to reduce the cooling towers electrical energy needs, maintaining their water needs almost unchanged in comparison to the current ones.

The techno-economic study of the obtained results highlights that the dry cooling tower solution does not bring advantages because the increase of electricity consumption overcome the obtained water-saving benefit. Moreover, the dry cooling tower led the refrigerators to work in worse conditions, causing an additional increase in electricity demand. It is important to remark that this kind of side-effects can be effectively evaluated only by

exploiting dynamic energy simulations.The second test case involves the techno-economic evaluation of a photovoltaic generation system. The study was conducted in order to evaluate the economic optimum in terms of generation system total employed surface. Two different kinds of panels have been evaluated:

- High-efficiency and expansive panel
- Low-efficiency and low-cost panel

The results of the techno-economic comparison highlight that the high-efficiency panel results to be a worse investment in comparison to the low-cost one. The higher efficiency of the former is not enough to justify the higher *CAPEX*, regardless of the system surface. Moreover, it has been calculated that the more convenient employed surface is around $6'000\ m^2$, a surface unavailable for the studied test case. Nevertheless, this information let to conclude that the photovoltaic panel could be a feasible investment.

Concluding, the study led to the *BEM* application on a peculiar test case, characterized by complex energy systems and internal core activities, exploiting it to draw conclusions about possible future energy efficiency solutions.

Appendix B

A stochastic first-order trust-region method with inexact restoration (SIRTR)

In this section, we describe a stochastic iterative method first proposed in [108] and suited for training regression models arising in the context of Building Energy Prediction. The proposed method solves the large-scale minimization problem involved in the training phase by iteratively performing a stochastic gradient-based step, where the full gradient of the objective function is subsampled, and both the batch size and the learning rate are updated through an automatic acceptance criterion. The adaptive tuning of the hyper-parameters is one of the key features that makes our method particularly attractive for the efficient prediction of building energy demands.

Throughout the section, we let $\{(a_i, b_i)\}_{i=1}^N$ denote a generic training set of size N, where $a_i \in \mathbb{R}^n$ represents the i^{th} feature vector, and $b_i \in \mathbb{R}$ stands for the target variable of the i^{th} example. The goal is to train the prediction function $h(\cdot; x) : \mathbb{R}^n \to \mathbb{R}$, whose parameters are encapsulated in the vector $x \in \mathbb{R}^n$. One common approach is to compute the parameters vector x as the solution of the following minimization problem

$$\min_{x \in \mathbb{R}^n} f_N(x) \equiv \frac{1}{N} \sum_{i=1}^{N} \ell(h(a_i; x); b_i), \tag{B.1}$$

where $\ell : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ is a continuously differentiable loss function. Intuitively, the function ℓ measures the distance between a target variable b_i and its predicted value $h(a_i; x)$. Hence, x is computed as the vector that minimizes the average loss measured by the model on the training set.

The proposed method, called Stochastic Inexact Restoration TRust-region (SIRTR), is a first-order iterative method for solving B.1. Starting from an initial guess $x_0 \in \mathbb{R}^n$,

SIRTR refines progressively the approximation x_k of the models parameters by iteratively applying a stochastic gradient-based step with *adaptive* batch size and learning rate. At each iteration k, SIRTR computes a stochastic approximation g_k of the full gradient $\nabla f_N(x_k)$, which is obtained by randomly selecting a batch $I_{N_{k+1,g}} \subseteq \{1, \dots, N\}$ of size $N_{k+1,g}$ and computing the following sample average approximation

$$g_k = \frac{1}{N_{k+1,g}} \sum_{i \in I_{N_{k+1,g}}} \nabla \ell(h(a_i; x_k); b_i). \tag{B.2}$$

Then, SIRTR computes the search direction p_k according to a first-order trust-region approach [109, 110], and accepts or rejects the trial point $x_k + p_k$ according to an acceptance criterion that depends on the evaluation of the following two function approximations

$$f_{N^t_{k+1}}(x_k) = \frac{1}{N^t_{k+1}} \sum_{i \in I_{N^t_{k+1}}} \ell(h(a_i; x_k); b_i) \tag{B.3}$$

$$f_{N^t_{k+1}}(x_k + p_k) = \frac{1}{N^t_{k+1}} \sum_{i \in I_{N^t_{k+1}}} \ell(h(a_i; x_k + p_k); b_i), \tag{B.4}$$

where again $I_{N^t_{k+1}} \subseteq \{1, \dots, N\}$ is a batch of size N^t_{k+1} randomly selected. The sample size N^t_{k+1} is fixed through an adaptive procedure inspired by the *Inexact Restoration* approach for constrained differentiable optimization [111, 112]. Such a procedure is based on the trivial reformulation of problem (B.1) as the constrained problem

$$\begin{aligned} \min \quad & f_M(x) \equiv \frac{1}{M} \sum_{i \in I_M} \ell(h(a_i; x); b_i), \\ x \in \mathbb{R}^n \\ \\ M = N \end{aligned} \tag{B.5}$$

where f_M is a sample average approximation obtained by selecting a batch $I_M \subseteq \{1, \dots, N\}$ of size M. The proposed method addresses (B.5) by improving *optimality* with respect to the function approximation f_M, and *feasibility* with respect to the constraint $M = N$, in a modular way. Specifically, SIRTR accepts the point $x_k + p_k$ if and only if it provides a sufficient decrease for the following merit function

$$\Psi(x, M, \theta) = \theta f_M(x) + (1 - \theta) \left(\frac{N - M}{N} \right), \quad \theta \in (0, 1), \tag{B.6}$$

which is a convex combination of the function approximation f_M and the *infeasibility measure* $\varphi(M) = (N - M)/N$. Intuitively, Ψ balances optimality with feasibility, and minimizing Ψ in place of f may have some advantages in terms of computational savings

during the iterations; indeed, we may allow for a decrease in M, provided that the function f_M decreases enough from one iteration to the other.

We describe the $k-$th iteration of SIRTR applied to problem (B.1) by dividing it into four steps. Note that, at each iteration k, we have at our disposal the previous iterate x_k, the function sample size N_k, the trust-region radius $\delta_k \in (0, \delta_{\max})$, a penalty parameter $\theta_k \in (0,1)$, and we let either $\mathcal{F}_k = 1$ if $N_k = N$, or $\mathcal{F}_k = 0$ if $N_k < N$. Furthermore, the parameters $\theta_0 \in (0,1)$, $r \in (0,1)$, $\eta_1 \in (0,1)$, $\eta_2 > 0$, $\mu > 0$, $\delta_0 \in (0, \delta_{\max})$, $\gamma > 1$ are chosen by the user.

- Restoration phase (Improve feasibility). If $N_k < N$, a new sample size $\widetilde{N}_{k+1} \leq N$ is determined such that the infeasibility measure is sufficiently decreased via the inequality
$$\widetilde{N}_{k+1} \geq rN_k + (1-r)N.$$
If $N_k = N$, we just set $\widetilde{N}_{k+1} = N$.

- Restoration phase (Reduce complexity). Once $\widetilde{N}_{k+1}$ is computed, we attempt to reduce such a a value by computing a trial sample size N^t_{k+1} that satisfies the following inequalities
$$N^t_{k+1} \leq \widetilde{N}_{k+1} \quad \text{and} \quad \frac{\widetilde{N}_{k+1} - N^t_{k+1}}{N} \leq \mu\delta_k^2.$$
Note that $N^t_{k+1} = \widetilde{N}_{k+1}$ for δ_k small enough.

- Trust-region. The function approximation $f_{N^t_{k+1}}(x_k)$ and gradient approximation g_k are computed according to (B.3) and (B.2), respectively. The batch size $I_{N_{k+1,g}}$ is chosen as a subset of $I_{N^t_{k+1}}$, with $N_{k+1,g} \leq N^t_{k+1}$. The search direction $p_k \in \mathbb{R}^n$ is then computed by a first-order trust-region approach, i.e., minimizing a linear approximation $m_k(p) = f_{N^t_{k+1}}(x_k) + g_k^T p$ of the objective function f_N over a ball centered at x_k with radius δ_k, which gives
$$p_k = \underset{\|p\| \leq \delta_k}{\operatorname{argmin}}\ m_k(p) \ = -\delta_k \frac{g_k}{\|g_k\|}.$$
Note that p_k can be interpreted as a stochastic gradient step with adaptive learning rate $\delta_k / \|g_k\|$.

- Step acceptance. The acceptance criterion for the trial point $x_k + p_k$ requires the computation of a function approximation $f_{N^t_{k+1}}(x_k + p_k)$ of the form (B.4), and a penalty parameter θ_{k+1} such that the relation $\text{Pred}(\theta_{k+1}) \geq \eta_1 \left(\frac{\widetilde{N}_{k+1} - N_k}{N}\right)$ holds. SIRTR accepts or rejects $x_k + p_k$ according to the ratio between the *actual reduction*

of the merit function Ψ

$$\begin{aligned}\mathrm{Ared}_k(\theta_{k+1}) &= \Psi(x_k, N_k, \theta_k) - \Psi(x_k + p_k, N_{k+1}^t, \theta_{k+1}) \\ &= \theta_{k+1}(f_{N_k}(x_k) - f_{N_{k+1}^t}(x_k + p_k)) + (1 - \theta_{k+1}) \left(\frac{N_{k+1}^t - N_k}{N} \right),\end{aligned}$$

and the *predicted reduction* according to the model m_k and the sample size $\widetilde{N}_{k+1}$, i.e.,

$$\mathrm{Pred}_k(\theta_{k+1}) = \theta_{k+1}(f_{N_k}(x_k) - m_k(p_k)) + (1 - \theta_{k+1}) \left(\frac{\widetilde{N}_{k+1} - N_k}{N} \right).$$

If $\frac{\mathrm{Ared}_k(\theta_{k+1})}{\mathrm{Pred}_k(\theta_{k+1})} \geq \eta_1$ (which requires a sufficient reduction of Ψ) and $(\|g_k\| - \eta_2\delta_k)\mathcal{F}_k \geq 0$ (which activates only when $N_k = N$), then the iteration k is declared successful, the trial point $x_k + p_k$ is accepted, and the trust-region radius is increased, so that x_{k+1} and δ_{k+1} are updated as

$$x_{k+1} = x_k - \frac{\delta_k}{\|g_k\|} g_k \tag{B.7}$$

$$\delta_{k+1} = \max\{\gamma\delta_k, \delta_{\max}\}. \tag{B.8}$$

Otherwise, the iteration k is declared unsuccessful, the trial point $x_k + p_k$ is rejected, and the trust-region radius is decreased, leading to the following update

$$x_{k+1} = x_k \tag{B.9}$$

$$\delta_{k+1} = \frac{\delta_k}{\gamma}. \tag{B.10}$$

The algorithm is reported in its entirety in the box below.

The presented algorithm has been tested on an energy demand dataset related to the test case described in this thesis. Nevertheless, the training data are related to the year 2019 (while this thesis work use the 2020 one). The objective of this work was to develop a novel training algorithm in aiming to reduce the computational effort related to the training phase. Indeed, the batch size automatic variation allow to reduce the computational effort of each epoch, consequently reducing the computational cost of a single training. Moreover, the "dynamic" learning rate lead to two main advantages. On the one hand, the training procedure is capable to adjust the learning rate during the training process, choosing the best one to perform each epoch and refusing any wrong

Algorithm : SIRTR

Given $x_0 \in \mathbb{R}^n$, N_0 integer in $(0, N]$, $\theta_0 \in (0,1)$, $0 < \delta_0 < \delta_{\max}$, $\gamma > 1$, $r, \eta_1 \in (0,1)$, μ, $\eta_2 > 0$.

0. Set $k = 0$, `iflag=succ`.
1. If `iflag=succ`
   If $N_k < N$ set $\mathcal{F}_k = 0$, else set $\mathcal{F}_k = 1$. Find $\widetilde{N}_{k+1}$ such that

$$N_k \leq \widetilde{N}_{k+1} \leq N \quad \text{and} \quad \widetilde{N}_{k+1} \geq rN_k + (1-r)N.$$

   Else set $\widetilde{N}_{k+1} = \widetilde{N}_k$, $\mathcal{F}_k = \mathcal{F}_{k-1}$.
2. Find N_{k+1}^t such that

$$N_{k+1}^t \leq \widetilde{N}_{k+1} \quad \text{and} \quad \frac{\widetilde{N}_{k+1} - N_{k+1}^t}{N} \leq \mu\delta_k^2.$$

3. Choose g_k as in (B.2), build $m_k(p) = f_{N_{k+1}^t}(x_k) + g_k^T p$, and set $p_k = -\delta_k g_k / \|g_k\|$.
4. If $N_k = N$, $N_{k+1}^t < N$ and $f_N(x_k) - m_k(p_k) < \delta_k \|g_k\|$, then take $\delta_k = \delta_k/\gamma$ and go to Step 2.
5. Compute the penalty parameter θ_{k+1}

$$\theta_{k+1} = \begin{cases} \theta_k & \text{if } \ \mathrm{Pred}_k(\theta_k) \geq \eta(\varphi(N_k) - \varphi(\widetilde{N}_{k+1})) \\ \dfrac{(1-\eta)(\varphi(N_k) - \varphi(\widetilde{N}_{k+1}))}{m_k(p_k) - f_{N_k}(x_k) + \varphi(N_k) - \varphi(\widetilde{N}_{k+1})} & \text{otherwise.} \end{cases} \tag{B.11}$$

6. If $\mathrm{Ared}_k(x_k + p_k, \theta_{k+1}) \geq \eta_1 \mathrm{Pred}_k(\theta_{k+1})$ and $(\|g_k\| - \eta_2\delta_k)\mathcal{F}_k \geq 0$ define

$$x_{k+1} = x_k + p_k, \qquad \delta_{k+1} = \max\{\gamma\delta_k, \delta_{\max}\}. \tag{B.12}$$

   set $N_{k+1} = N_{k+1}^t$, $k = k+1$, `iflag=succ` and go to Step 1.
   Else set $x_{k+1} = x_k$, $\delta_{k+1} = \delta_k/\gamma$, $N_{k+1} = N_k$, $k = k+1$, `iflag=unsucc` and go to Step 1.

choice through the *Trust-Region* approach. On the other hand, it allow to completely remove the learning rate from the list of "tunable" hyperparameters.

These two aspect allow for:

- Reduce the computational cost of the feature selection phase thanks to the dynamic batch size and learning rate adjustments, leading to decrease the training time of each single training. Since several training processes are involved in this phases, even a small reduction in computational cost lead to non-negligible reduction of the overall elapsed time

- Removing the learning rate from the tunable hyperparameters allow for a significative reduction of the number of hyperparameters combination considered in the

hyperparameter grid of the fine tuning using grid search approach. These, combined with the previously described effect, allow for markedly reduce the computational effort of the hyperparameter fine tuning phases.

The developed algorithm is still in the testing phase, and is therefore still work in progress. However, comparisons have already been made of the behaviour of validation loss functions obtained with *SIRTR* and those obtained with the other two most commonly used algorithms: Stochastic Gradient Descent (*SGD*, see Figure B.1) and *Adam* (see Figure B.2).

Since the *SIRTR* algorithm do not use the full dataset at each epoch, but randomly extract a portion of data quantified by the dynamic batch size, the loss function behaviour have been observed by comparing the cumulative number of data observed during the training process. The latter have been quantified through the parameter fev/N, where fev is the number of sample used to calculate the model error at the end of each epoch, and N is the total number of samples which compose the training dataset.

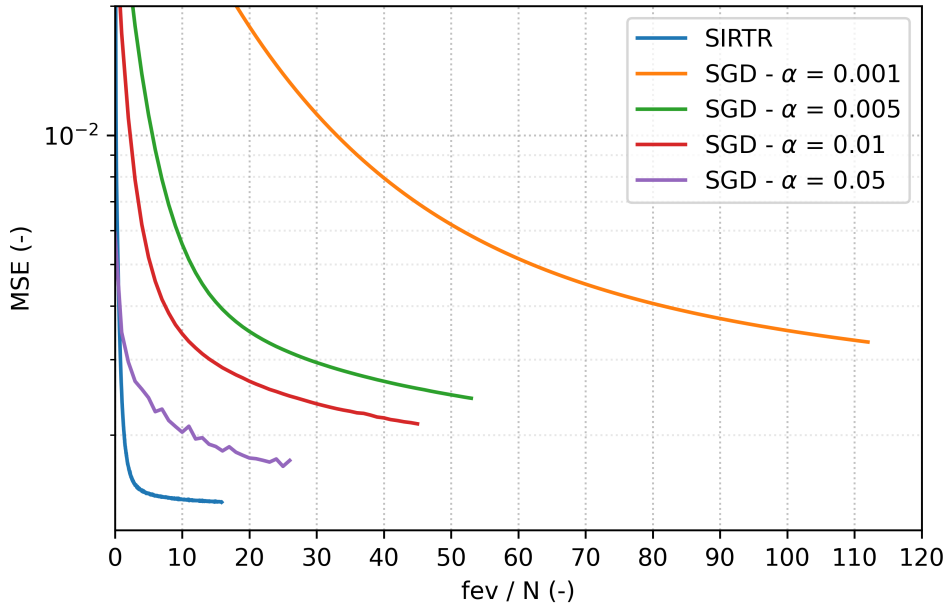

Figure B.1: Loss and Validation loss function behaviours for SIRTR and SGD with different learning rate values

The *SIRTR* approach outperforms the *SGD* algorithm regardless the learning rate value. Decreasing the learning rate of the *SGD* lead to obtain better performance, but yet incomparable with the one obtained using *SIRTR*, which results faster to converge and allow to markedly lower *mean squared error* at the end of the training.

The *Adam* algorithm results comparable in terms of prediction performance, capable to preform better than *SIRTR* for learning rate (α) of 0.005 with a similar convergence speed. Nevertheless, it is strongly sensible to the learning rate value. Moreover, the validation loss function behaviour appear irregular (similarly to what have been found for the training

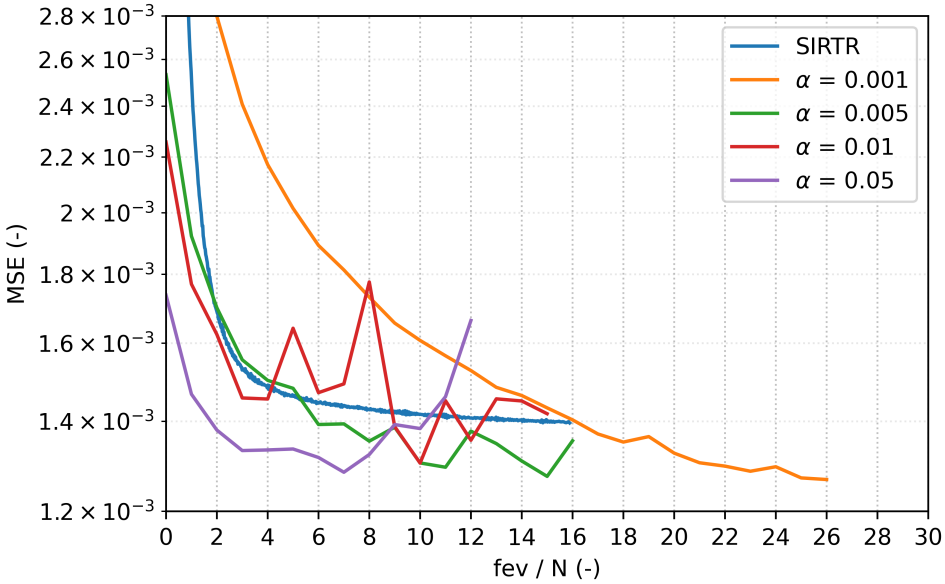

Figure B.2: Loss and Validation loss function behaviours for SIRTR and Adam with different learning rate values

procedures analysed in this thesis) and less stable in comparison to the *SIRTR* one. Finally, the choose of higher learning rate value ($lr = 0.05$) lead the validation function to start to increase after few epochs, a fact that highlight that the training process is prone to the overfitting problem.

Concluding, the *SIRTR* algorithm shows promising results, allowing to markedly reduce the computational cost of the systematic monitoring method proposed in this thesis work and being able to make the training process more stable and controllable, avoiding training issues as the overfitting.

Bibliography

[1] C. A. Balaras, E. G. Dascalaki, I. Psarra and T. Cholewa, 'Primary energy factors for electricity production in europe', *Energies*, vol. 16, no. 1, 2023, ISSN: 1996-1073. DOI: 10.3390/en16010093.

[2] *Energy, transport and environment statistics: 2020 edition*, eng, 2020 edition. Luxembourg: Publications Office of the European Union, 2020, OCLC: 1230562757, ISBN: 978-92-76-20736-8.

[3] Terna and Snam, *Scenario national trend italia*, 2021.

[4] *Piano nazionale integrato per l'energia e per il clima*, Dec. 2019.

[5] E. Commission and D.-G. for Energy, *EU energy in figures : statistical pocketbook 2022*. Publications Office of the European Union, 2022. DOI: doi/10.2833/334050.

[6] D. Helm, 'The european framework for energy and climate policies', *Energy Policy*, vol. 64, pp. 29–35, 2014, ISSN: 0301-4215. DOI: https://doi.org/10.1016/j.enpol.2013.05.063.

[7] S. Papadopoulos, C. E. Kontokosta, A. Vlachokostas and E. Azar, 'Rethinking hvac temperature setpoints in commercial buildings: The potential for zero-cost energy savings and comfort improvement in different climates', *Building and Environment*, vol. 155, pp. 350–359, 2019.

[8] *Annuario statistico del servizio sanitario nazionale*, 2020.

[9] C. Shen, K. Zhao, J. Ge and Q. Zhou, 'Analysis of Building Energy Consumption in a Hospital in the Hot Summer and Cold Winter Area', en, *Energy Procedia*, Innovative Solutions for Energy Transitions, vol. 158, pp. 3735–3740, Feb. 2019, ISSN: 1876-6102. DOI: 10.1016/j.egypro.2019.01.883.

[10] M. J. Eckelman and J. Sherman, 'Environmental Impacts of the U.S. Health Care System and Effects on Public Health', en, *PLOS ONE*, vol. 11, no. 6, e0157014, Jun. 2016, Publisher: Public Library of Science, ISSN: 1932-6203. DOI: 10.1371/journal.pone.0157014.

[11] I. S. Dobosi, C. Tanasa, N.-E. Kaba, A. Retezan and D. Mihaila, 'Building energy modelling for the energy performance analysis of a hospital building in various locations', en, *E3S Web of Conferences*, vol. 111, p. 06 073, 2019, Publisher: EDP Sciences, ISSN: 2267-1242. DOI: 10.1051/e3sconf/201911106073.

[12] R. E. Stockwell, E. L. Ballard, P. O'Rourke, L. D. Knibbs, L. Morawska and S. C. Bell, 'Indoor hospital air and the impact of ventilation on bioaerosols: A systematic review', en, *Journal of Hospital Infection*, vol. 103, no. 2, pp. 175–184, Oct. 2019, ISSN: 0195-6701. DOI: 10.1016/j.jhin.2019.06.016.

[13] J. Skoog, N. Fransson and L. Jagemar, 'Thermal environment in Swedish hospitals: Summer and winter measurements', en, *Energy and Buildings*, vol. 37, no. 8, pp. 872–877, Aug. 2005, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2004.11.003.

[14] A. Viloria, W. Osal, C. Vásquez, C. Gonzalez, N. Varela and M. Gaitán-Angulo, 'Energy Efficiency Index of Ambulatories and Hospitals', *International Journal of Control Theory and Applications*, vol. 9, pp. 59–64, Dec. 2016.

[15] S.-C. Hu, J. Chen and Y. Chuah, 'Energy cost and consumption in a large acute hospital', *International Journal On Architectural Science*, vol. 5, Jan. 2004.

[16] C. J. Renedo, A. Ortiz, M. Mañana, D. Silió and S. Pérez, 'Study of different cogeneration alternatives for a Spanish hospital center', en, *Energy and Buildings*, vol. 38, no. 5, pp. 484–490, May 2006, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2005.08.011.

[17] V. Čongradac, B. Prebiračević, N. Jorgovanović and D. Stanišić, 'Assessing the energy consumption for heating and cooling in hospitals', en, *Energy and Buildings*, vol. 48, pp. 146–154, May 2012, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2012.01.022.

[18] A. Bianchini, C. Carcasci, G. Manfrida and M. Zini, 'Reconstruction and Analysis of the Energy Demand of a Healthcare Facility in Italy', en, *E3S Web of Conferences*, vol. 197, p. 02 009, 2020, ISSN: 2267-1242. DOI: 10.1051/e3sconf/202019702009.

[19] B. Singer and W. Tschudi, *High Performance Healthcare Buildings: A Roadmap to Improved Energy Efficiency*, Place: Berkeley Publisher: Lawrence Berkeley National Laboratory Volume: 515, 2009.

[20] M. Zini, R. Sodini and C. Carcasci, 'Modeling and Optimization of a Hospital Gas Turbine-Based Cogeneration System', *Journal of Engineering for Gas Turbines and Power*, vol. 144, no. 11, Sep. 2022, ISSN: 0742-4795. DOI: 10.1115/1.4055418.

[21] K. Qu, X. Chen, Y. Wang, J. Calautit, S. Riffat and X. Cui, 'Comprehensive energy, economic and thermal comfort assessments for the passive energy retrofit of historical buildings - A case study of a late nineteenth-century Victorian house renovation in the UK', en, *Energy*, vol. 220, p. 119 646, Apr. 2021, ISSN: 0360-5442. DOI: 10.1016/j.energy.2020.119646.

[22] I. Montero, M. Miranda, F. Barrena, F. Sepúlveda and J. Arranz, 'Analysis of photovoltaic self-consumption systems for hospitals in southwestern europe', *Energy and Buildings*, vol. 269, p. 112 254, 2022, ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2022.112254.

[23] J. García-Sanz-Calcedo, M. Gómez-Chaparro and G. Sanchez-Barroso, 'Electrical and thermal energy in private hospitals: Consumption indicators focused on healthcare activity', en, *Sustainable Cities and Society*, vol. 47, p. 101 482, May 2019, ISSN: 2210-6707. DOI: 10.1016/j.scs.2019.101482.

[24] J. García-Sanz-Calcedo, A. Al-Kassir and T. Yusaf, 'Economic and environmental impact of energy saving in healthcare buildings', *Applied Sciences*, vol. 8, no. 3, 2018, ISSN: 2076-3417. DOI: 10.3390/app8030440.

[25] M. Economidou, V. Todeschi, P. Bertoldi, D. D'Agostino, P. Zangheri and L. Castellazzi, 'Review of 50 years of EU energy efficiency policies for buildings', *Energy and Buildings*, vol. 225, p. 110 322, 15th Oct. 2020, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2020.110322.

[26] C. Buratti, E. Moretti, E. Belloni and F. Cotana, 'Unsteady simulation of energy performance and thermal comfort in non-residential buildings', en, *Building and Environment*, vol. 59, pp. 482–491, Jan. 2013, ISSN: 0360-1323. DOI: 10.1016/j.buildenv.2012.09.015.

[27] J. Van der Veken, D. Saelens, G. Verbeeck and H. Hens, 'Comparison of steady-state and dynamic building energy simulation programs', *Proceedings of the international Buildings IX ASHREA conference on the performance of exterior envelopes of whole buildings*, 2004.

[28] P. Dhillon, W. T. Horton and J. E. Braun, 'Comparison of Steady-State and Dynamic Load-Based Performance Evaluation Methodologies for a Residential Air Conditioner', *International Refrigeration and Air Conditioning Conference*, May 2021.

[29] L. Berkeley et al., *EnergyPlus Essentials*, en, 2019.

[30] M. Magni, F. Ochs, S. de Vries, A. Maccarini and F. Sigg, 'Detailed cross comparison of building energy simulation tools results using a reference office building as a case study', en, *Energy and Buildings*, vol. 250, p. 111 260, Nov. 2021, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2021.111260.

[31] A. Foucquier, S. Robert, F. Suard, L. Stéphan and A. Jay, 'State of the art in building modelling and energy performances prediction: A review', en, *Renewable and Sustainable Energy Reviews*, vol. 23, pp. 272–288, Jul. 2013, ISSN: 1364-0321. DOI: 10.1016/j.rser.2013.03.004.

[32] V. S. K. V. Harish and A. Kumar, 'A review on modeling and simulation of building energy systems', en, *Renewable and Sustainable Energy Reviews*, vol. 56, pp. 1272–1292, Apr. 2016, ISSN: 1364-0321. DOI: 10.1016/j.rser.2015.12.040.

[33] UNI, *Uni/ts 11300-1:2008, "prestazioni energetiche degli edifici. parte 1: "determinazione del fabbisogno di energia termica dell'edificio per la climatizzazione estiva e invernale"*, May 2008.

[34] *Directive 2002/91/ec of the european parliament and council of december 16, 2002 on energy performance in buildings*, 2002.

[35] S. A. Kalogirou, 'Use of TRNSYS for modelling and simulation of a hybrid pv–thermal solar system for Cyprus', en, *Renewable Energy*, vol. 23, no. 2, pp. 247–260, Jun. 2001, ISSN: 0960-1481. DOI: 10.1016/S0960-1481(00)00176-2.

[36] A. Pedrini, F. S. Westphal and R. Lamberts, 'A methodology for building energy modelling and calibration in warm climates', en, *Building and Environment*, vol. 37, no. 8, pp. 903–912, Aug. 2002, ISSN: 0360-1323. DOI: 10.1016/S0360-1323(02)00051-3.

[37] D. Coakley, P. Raftery and M. Keane, 'A review of methods to match building energy simulation models to measured data', en, *Renewable and Sustainable Energy Reviews*, vol. 37, pp. 123–141, Sep. 2014, ISSN: 1364-0321. DOI: 10.1016/j.rser.2014.05.007.

[38] Z. (Zhai, M.-H. Johnson and M. Krarti, 'Assessment of natural and hybrid ventilation models in whole-building energy simulations', en, *Energy and Buildings*, vol. 43, no. 9, pp. 2251–2261, Sep. 2011, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2011.06.026.

[39] C. Booten and P. C. Tabares-Velasco, 'Using EnergyPlus to Simulate the Dynamic Response of a Residential Building to Advanced Cooling Strategies', en, p. 10, Aug. 2012.

[40] R. Kamal, F. Moloney, C. Wickramaratne, A. Narasimhan and D. Y. Goswami, 'Strategic control and cost optimization of thermal energy storage in buildings using EnergyPlus', en, *Applied Energy*, vol. 246, pp. 77–90, Jul. 2019, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2019.04.017.

[41] J. Kotowicz and W. Uchman, 'Analysis of the integrated energy system in residential scale: Photovoltaics, micro-cogeneration and electrical energy storage', en, *Energy*, vol. 227, p. 120 469, Jul. 2021, ISSN: 0360-5442. DOI: 10.1016/j.energy.2021.120469.

[42] T. de Rubeis, S. Falasca, G. Curci, D. Paoletti and D. Ambrosini, 'Sensitivity of heating performance of an energy self-sufficient building to climate zone, climate change and HVAC system solutions', en, *Sustainable Cities and Society*, vol. 61, p. 102 300, Oct. 2020, ISSN: 2210-6707. DOI: 10.1016/j.scs.2020.102300.

[43] T. de Rubeis, I. Nardi, D. Ambrosini and D. Paoletti, 'Is a self-sufficient building energy efficient? Lesson learned from a case study in Mediterranean climate', en, *Applied Energy*, vol. 218, pp. 131–145, May 2018, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2018.02.166.

[44] E. Amiri Rad, S. Maddah and S. Mohammadi, 'Designing and optimizing a novel cogeneration system for an office building based on thermo-economic and environmental analyses', en, *Renewable Energy*, vol. 151, pp. 342–354, May 2020, ISSN: 0960-1481. DOI: 10.1016/j.renene.2019.11.024.

[45] R. E. Edwards, J. New, L. E. Parker, B. Cui and J. Dong, 'Constructing large scale surrogate models from big data and artificial intelligence', en, *Applied Energy*, vol. 202, pp. 685–699, Sep. 2017, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2017.05.155.

[46] F. Johari, G. Peronato, P. Sadeghian, X. Zhao and J. Widén, 'Urban building energy modeling: State of the art and future prospects', en, *Renewable and Sustainable Energy Reviews*, vol. 128, p. 109 902, Aug. 2020, ISSN: 1364-0321. DOI: 10.1016/j.rser.2020.109902.

[47] M. Beccali, G. Ciulla, B. Di Pietra, A. Galatioto, G. Leone and A. Piacentino, 'Assessing the feasibility of cogeneration retrofit and district heating/cooling networks in small Italian islands', en, *Energy*, vol. 141, pp. 2572–2586, Dec. 2017, ISSN: 0360-5442. DOI: 10.1016/j.energy.2017.07.011.

[48] A. Buonomano, F. Calise, G. Ferruzzi and A. Palombo, 'Dynamic energy performance analysis: Case study for energy efficiency retrofits of hospital buildings', en, *Energy*, vol. 78, pp. 555–572, Dec. 2014, ISSN: 0360-5442. DOI: 10.1016/j.energy.2014.10.042.

[49] S. Fathi, R. Srinivasan, A. Fenner and S. Fathi, 'Machine learning applications in urban building energy performance forecasting: A systematic review', *Renewable and Sustainable Energy Reviews*, vol. 133, p. 110 287, 1st Nov. 2020, ISSN: 1364-0321. DOI: 10.1016/j.rser.2020.110287.

[50] S. Seyedzadeh, F. P. Rahimian, I. Glesk and M. Roper, 'Machine learning for estimation of building energy consumption and performance: A review', *Visualization in Engineering*, vol. 6, no. 1, p. 5, 2nd Oct. 2018, ISSN: 2213-7459. DOI: 10.1186/s40327-018-0064-7.

[51] S. Kapp, J.-K. Choi and T. Hong, 'Predicting industrial building energy consumption with statistical and machine-learning models informed by physical system parameters', *Renewable and Sustainable Energy Reviews*, vol. 172, p. 113 045, 2023, ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2022.113045.

[52] A. Mosavi, M. Salimi, S. Faizollahzadeh Ardabili, T. Rabczuk, S. Shamshirband and A. R. Varkonyi-Koczy, 'State of the art of machine learning models in energy systems, a systematic review', *Energies*, vol. 12, no. 7, p. 1301, Jan. 2019. DOI: 10.3390/en12071301.

[53] Z. Chen, F. Xiao, F. Guo and J. Yan, 'Interpretable machine learning for building energy management: A state-of-the-art review', *Advances in Applied Energy*, vol. 9, p. 100 123, 2023, ISSN: 2666-7924. DOI: https://doi.org/10.1016/j.adapen.2023.100123.

[54] M. Ilbeigi, M. Ghomeishi and A. Dehghanbanadaki, 'Prediction and optimization of energy consumption in an office building using artificial neural network and a genetic algorithm', *Sustainable Cities and Society*, vol. 61, p. 102 325, 1st Oct. 2020, ISSN: 2210-6707. DOI: 10.1016/j.scs.2020.102325.

[55] A.-D. Pham, N.-T. Ngo, T. T. Ha Truong, N.-T. Huynh and N.-S. Truong, 'Predicting energy consumption in multiple buildings using machine learning for improving energy efficiency and sustainability', *Journal of Cleaner Production*, vol. 260, p. 121 082, 1st Jul. 2020, ISSN: 0959-6526. DOI: 10.1016/j.jclepro.2020.121082.

[56] J. Runge and R. Zmeureanu, 'Forecasting energy use in buildings using artificial neural networks: A review', *Energies*, vol. 12, no. 17, p. 3254, Jan. 2019. DOI: 10.3390/en12173254.

[57] H.-x. Zhao and F. Magoulès, 'A review on the prediction of building energy consumption', *Renewable and Sustainable Energy Reviews*, vol. 16, no. 6, pp. 3586–3592, 1st Aug. 2012, ISSN: 1364-0321. DOI: 10.1016/j.rser.2012.02.049.

[58] X. J. Luo *et al.*, 'Genetic algorithm-determined deep feedforward neural network architecture for predicting electricity consumption in real buildings', *Energy and AI*, vol. 2, p. 100 015, 1st Nov. 2020, ISSN: 2666-5468. DOI: 10.1016/j.egyai.2020.100015.

[59] L. Yan and M. Liu, 'A simplified prediction model for energy use of air conditioner in residential buildings based on monitoring data from the cloud platform', *Sustainable Cities and Society*, vol. 60, p. 102 194, 1st Sep. 2020, ISSN: 2210-6707. DOI: 10.1016/j.scs.2020.102194.

[60] J. K. Hwang, G. Y. Yun, S. Lee, H. Seo and M. Santamouris, 'Using deep learning approaches with variable selection process to predict the energy performance of a heating and cooling system', *Renewable Energy*, vol. 149, pp. 1227–1245, 1st Apr. 2020, ISSN: 0960-1481. DOI: 10.1016/j.renene.2019.10.113.

[61] S. L. Wong, K. K. W. Wan and T. N. T. Lam, 'Artificial neural networks for energy analysis of office buildings with daylighting', *Applied Energy*, vol. 87, no. 2, pp. 551–557, 1st Feb. 2010, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2009.06.028.

[62] S. Seyedzadeh, F. P. Rahimian, S. Oliver, I. Glesk and B. Kumar, 'Data driven model improved by multi-objective optimisation for prediction of building energy loads', *Automation in Construction*, vol. 116, p. 103 188, 1st Aug. 2020, ISSN: 0926-5805. DOI: 10.1016/j.autcon.2020.103188.

[63] S. Seyedzadeh, F. Pour Rahimian, P. Rastogi and I. Glesk, 'Tuning machine learning models for prediction of building energy loads', *Sustainable Cities and Society*, vol. 47, p. 101 484, 1st May 2019, ISSN: 2210-6707. DOI: 10.1016/j.scs.2019.101484.

[64] S. Seyedzadeh, F. Pour Rahimian, S. Oliver, S. Rodriguez and I. Glesk, 'Machine learning modelling for predicting non-domestic buildings energy performance: A model to support deep energy retrofit decision-making', *Applied Energy*, vol. 279, p. 115 908, 1st Dec. 2020, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2020.115908.

[65] J. Ngarambe, G. Y. Yun and M. Santamouris, 'The use of artificial intelligence (AI) methods in the prediction of thermal comfort in buildings: Energy implications of AI-based thermal comfort controls', *Energy and Buildings*, vol. 211, p. 109 807, 15th Mar. 2020, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2020.109807.

[66] D. Palladino, I. Nardi and C. Buratti, 'Artificial neural network for the thermal comfort index prediction: Development of a new simplified algorithm', *Energies*, vol. 13, no. 17, p. 4500, Jan. 2020. DOI: 10.3390/en13174500.

[67] A. Albatayneh, D. Alterman, A. Page and B. Moghtaderi, 'The impact of the thermal comfort models on the prediction of building energy consumption', *Sustainability*, vol. 10, no. 10, p. 3609, Oct. 2018. DOI: 10.3390/su10103609.

[68] B. Yuce, H. Li, Y. Rezgui, I. Petri, B. Jayan and C. Yang, 'Utilizing artificial neural network to predict energy consumption and thermal comfort level: An indoor swimming pool case study', *Energy and Buildings*, vol. 80, pp. 45–56, 1st Sep. 2014, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2014.04.052.

[69] S. A. Kalogirou, 'Artificial neural networks in energy applications in buildings', *International Journal of Low-Carbon Technologies*, vol. 1, no. 3, pp. 201–216, 1st Jul. 2006, ISSN: 1748-1317. DOI: 10.1093/ijlct/1.3.201.

[70] H. O. Garcés *et al.*, 'Energy efficiency monitoring in a coal boiler based on optical variables and artificial intelligence', *IFAC-PapersOnLine*, 20th IFAC World Congress, vol. 50, no. 1, pp. 13 904–13 909, 1st Jul. 2017, ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2017.08.2209.

[71] M. Mohanraj, S. Jayaraj and C. Muraleedharan, 'Applications of artificial neural networks for refrigeration, air-conditioning and heat pump systems—a review', *Renewable and Sustainable Energy Reviews*, vol. 16, no. 2, pp. 1340–1358, 1st Feb. 2012, ISSN: 1364-0321. DOI: 10.1016/j.rser.2011.10.015.

[72] T. M. Mitchell, *Machine Learning*. McGraw-Hill, 1997, 414 pp., Google-Books-ID: EoYBngEACAAJ, ISBN: 978-0-07-115467-3.

[73] M. Yoosefzadeh-Najafabadi, H. J. Earl, D. Tulpan, J. Sulik and M. Eskandari, 'Application of machine learning algorithms in plant breeding: Predicting yield from hyperspectral reflectance in soybean', *Frontiers in Plant Science*, vol. 11, p. 624 273, 12th Jan. 2021, ISSN: 1664-462X. DOI: 10.3389/fpls.2020.624273.

[74] J. Hu, H. Niu, J. Carrasco, B. Lennox and F. Arvin, 'Voronoi-based multi-robot autonomous exploration in unknown environments via deep reinforcement learning', *IEEE Transactions on Vehicular Technology*, vol. 69, no. 12, pp. 14 413–14 423, Dec. 2020, Conference Name: IEEE Transactions on Vehicular Technology, ISSN: 1939-9359. DOI: 10.1109/TVT.2020.3034800.

[75] O. Chapelle, B. Scholkopf and A. Zien Eds., 'Semi-supervised learning (chapelle, o. et al., eds.; 2006) [book reviews]', *IEEE Transactions on Neural Networks*, vol. 20, no. 3, pp. 542–542, Mar. 2009, ISSN: 1941-0093. DOI: 10.1109/TNN.2009.2015974.

[76] A. Géron, *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: concepts, tools, and techniques to build intelligent systems*, Second edition. Beijing [China] ; Sebastopol, CA: O'Reilly Media, Inc, 2019, 819 pp., ISBN: 978-1-4920-3264-9.

[77] S. Bellavia, T. Bianconcini, N. Krejić and B. Morini, 'Subsampled first-order optimization methods with applications in imaging', in *Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging: Mathematical Imaging and Vision*, K. Chen, C.-B. Schönlieb, X.-C. Tai and L. Younces, Eds. Cham: Springer International Publishing, 2021, pp. 1–35, ISBN: 978-3-030-03009-4. DOI: 10.1007/978-3-030-03009-4_78-1.

[78] D. C. Montgomery, *Design and Analysis of Experiments*. John Wiley & Sons, 2017, 752 pp., Google-Books-ID: Py7bDgAAQBAJ, ISBN: 978-1-119-11347-8.

[79] M. Arioli and S. Gratton, 'Linear regression models, least-squares problems, normal equations, and stopping criteria for the conjugate gradient method', *Computer Physics Communications*, vol. 183, no. 11, pp. 2322–2336, 2012, ISSN: 0010-4655. DOI: https://doi.org/10.1016/j.cpc.2012.05.023.

[80] M. W. Ahmad, M. Mourshed and Y. Rezgui, 'Trees vs neurons: Comparison between random forest and ANN for high-resolution prediction of building energy consumption', *Energy and Buildings*, vol. 147, pp. 77–89, 15th Jul. 2017, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2017.04.038.

[81] O. I. Abiodun, A. Jantan, A. E. Omolara, K. V. Dada, N. A. Mohamed and H. Arshad, 'State-of-the-art in artificial neural network applications: A survey', *Heliyon*, vol. 4, no. 11, e00938, 1st Nov. 2018, ISSN: 2405-8440. DOI: 10.1016/j.heliyon.2018.e00938.

[82] M. T. Hagan, H. B. Demuth, M. H. Beale and O. De Jésus, *Neural network design*, 2nd edition. s.L: Martin T. Hagan, 2014, 1 p., ISBN: 978-0-9717321-1-7.

[83] C. Buratti, E. Lascaro, D. Palladino and M. Vergoni, 'Building behavior simulation by means of artificial neural network in summer conditions', *Sustainability*, vol. 6, no. 8, pp. 5339–5353, Aug. 2014. DOI: 10.3390/su6085339.

[84] J. Yao and J. Xu, 'Research on the BPNN in the prediction of PMV', *Applied Mechanics and Materials*, vol. 29-32, 1st Aug. 2010. DOI: 10.4028/www.scientific.net/AMM.29-32.2804.

[85] C. Buratti, M. Barbanera and D. Palladino, 'An original tool for checking energy performance and certification of buildings by means of artificial neural networks', *Applied Energy*, vol. 120, pp. 125–132, 1st May 2014, ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2014.01.053.

[86] P. Anand, D. Cheong, C. Sekhar, M. Santamouris and S. Kondepudi, 'Energy saving estimation for plug and lighting load using occupancy analysis', *Renewable Energy*, vol. 143, pp. 1143–1161, 1st Dec. 2019, ISSN: 0960-1481. DOI: 10.1016/j.renene.2019.05.089.

[87] S. Haykin, *Neural Networks: A Comprehensive Foundation*. Upper Saddle River, NJ: Prentice Hall, 21st Jul. 1998, 842 pp., ISBN: 978-0-13-908385-3.

[88] P. Juszczak, D. Tax and R. P. Duin, 'Feature scaling in support vector data description', pp. 95–102, 18th May 2002.

[89] L. Zhang and J. Wen, 'A systematic feature selection procedure for short-term data-driven building energy forecasting model development', *Energy and Buildings*, vol. 183, pp. 428–442, 15th Jan. 2019, ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2018.11.010.

[90] R. Kohavi *et al.*, 'A study of cross-validation and bootstrap for accuracy estimation and model selection', in *Ijcai*, Montreal, Canada, vol. 14, 1995, pp. 1137–1145.

[91] J. C. Lam, K. K. W. Wan, T. N. T. Lam and S. L. Wong, 'An analysis of future building energy use in subtropical Hong Kong', en, *Energy*, vol. 35, no. 3, pp. 1482–1490, Mar. 2010, ISSN: 0360-5442. DOI: 10.1016/j.energy.2009.12.005.

[92] K. Li, C. Hu, G. Liu and W. Xue, 'Building's electricity consumption prediction using optimized artificial neural networks and principal component analysis', *Energy and Buildings*, vol. 108, pp. 106–113, 2015.

[93] X. Luo, L. O. Oyedele, A. O. Ajayi, O. O. Akinade, H. A. Owolabi and A. Ahmed, 'Feature extraction and genetic algorithm enhanced adaptive deep neural network for energy consumption prediction in buildings', *Renewable and Sustainable Energy Reviews*, vol. 131, p. 109 980, 2020.

[94] E. Bisong, 'Optimization for machine learning: Gradient descent', in *Building Machine Learning and Deep Learning Models on Google Cloud Platform*, Springer, 2019, pp. 203–207.

[95] E. S. Page, 'Continuous inspection schemes', *Biometrika*, vol. 41, no. 1/2, pp. 100–115, 1954.

[96] L. C. Alwan, 'Cusum quality control-multivariate approach', *Communications in Statistics-Theory and Methods*, vol. 15, no. 12, pp. 3531–3543, 1986.

[97] D. C. Montgomery, *Introduction to statistical quality control*. John Wiley & Sons, 2020.

[98] V. S. Puranik, 'Cusum quality control chart for monitoring energy use performance', in *2007 IEEE International Conference on Industrial Engineering and Engineering Management*, 2007, pp. 1231–1235. DOI: 10.1109/IEEM.2007.4419388.

[99] Z. Wang, L. Xia, H. Yuan, R. S. Srinivasan and X. Song, 'Principles, research status, and prospects of feature engineering for data-driven building energy prediction: A comprehensive review', *Journal of Building Engineering*, vol. 58, p. 105 028, 2022, ISSN: 2352-7102. DOI: https://doi.org/10.1016/j.jobe.2022.105028.

[100] F. J. Pontes, G. Amorim, P. P. Balestrassi, A. Paiva and J. R. Ferreira, 'Design of experiments and focused grid search for neural network parameter optimization', *Neurocomputing*, vol. 186, pp. 22–34, 2016.

[101] D. Meyer and D. Thevenard, *PsychroLib: a library of psychrometric functions to calculate thermodynamic properties of air*, version 2.5.0, [Online; accessed 22-Feb-2023], Apr. 2020. DOI: 10.5281/zenodo.3748874.

[102] ASHRAE, *Fundamentals, ASHRAE–American society of heating (2017 edition*. 2017.

[103] A. Kumbhar, N. Gulhane and S. Pandure, 'Effect of various parameters on working condition of chiller', *Energy procedia*, vol. 109, pp. 479–486, 2017.

[104] *EnergyPlus | EnergyPlus*.

[105] *OpenStudio | OpenStudio*, Jul. 2021.

[106] P. C. Tabares-Velasco, 'Time Step Considerations when Simulating Dynamic Behavior of High Performance Homes', English, Oak Ridge, TN: Oak Ridge National Laboratory, Tech. Rep. NREL/CP-5500-68486, Sep. 2016.

[107] IRENA. 'Renewable power generation costs in 2019', /publications/2020/Jun/Renewable-Power-Costs-in-2019. (2019), [Online]. Available: /publications / 2020 / Jun / Renewable-Power-Costs-in-2019 (visited on 07/06/2021).

[108] S. Bellavia, N. Krejic, B. Morini and S. Rebegoldi, 'A stochastic first-order trust-region method with inexact restoration for finite-sum minimization', *Computational Optimization and Applications*, vol. 84, no. 1, pp. 53–84, 2023.

[109] A. S. Bandeira, K. Scheinberg and L. N. Vicente, 'Convergence of trust-region methods based on probabilistic models', *SIAM J. Optim.*, vol. 24, no. 3, pp. 1238–1264, 2014.

[110] R. Chen, M. Menickelly and K. Scheinberg, 'Stochastic optimization using a trust-region method and random models', *Math. Program.*, vol. 169, no. 2, pp. 447–487, 2018.

[111] G. E. Birgin, N. Krejic and J. M. Martinez, 'On the employment of inexact restoration for the minimization of functions whose evaluation is subject to programming errors', *Math. Comput.*, vol. 87, no. 311, pp. 1307–1326, 2018.

[112] J. M. Martinez and E. A. Pilotta, 'Inexact restoration algorithms for constrained optimization', *J. Optim. Theory Appl.*, vol. 104, pp. 135–163, 2000.