

Article

Impact of an ML-Based Demand Response Mechanism on the Electrical Distribution Network: A Case Study in Terni

Marco Antonio Bucarelli ^{1,*}, Mohammad Ghoreishi ¹, Francesca Santori ¹, Jorge Mira ² and Jesús Gorroñoitia ²

¹ ASM Terni S.p.A., 05100 Terni, Italy; mohammad.ghoreishi@asmterni.it (M.G.); francesca.santori@asmterni.it (F.S.)

² Atos Research & Innovation, 28042 Madrid, Spain; jorge.mira@atos.net (J.M.); jesus.gorronoitia@atos.net (J.G.)

* Correspondence: marco.antonio.bucarelli@asmterni.it

Abstract: The development of smart grids requires the active participation of end users through demand response mechanisms to provide technical benefits to the distribution network and receive economic savings. Integrating advanced machine learning tools makes it possible to optimise the network and manage the mechanism to maximise the benefits. This paper proceeds by forecasting consumption for the next 24 h using a recurrent neural network and by processing these data using a reinforcement learning-based optimisation model to identify the best demand response policy. The model is tested in a real environment: a portion of the Terni electrical distribution network. Several scenarios were identified, considering users' participation at different levels and limiting the potential with various constraints.

Keywords: demand response; distribution network; machine learning; optimisation; smart grid



Citation: Bucarelli, M.A.; Ghoreishi, M.; Santori, F.; Mira, J.; Gorroñoitia, J. Impact of an ML-Based Demand Response Mechanism on the Electrical Distribution Network: A Case Study in Terni. *Electronics* **2023**, *12*, 3948. <https://doi.org/10.3390/electronics12183948>

Academic Editor: Nikolay Hinov

Received: 4 August 2023

Revised: 10 September 2023

Accepted: 16 September 2023

Published: 19 September 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The distribution network (DN) faces a challenging transition period due to the spread of renewable energy sources, which, at the European level, are expected to reach 40% of the energy mix in 2030 to reduce greenhouse gas emissions by 55% compared with 1990 [1]. The grid must move from a static infrastructure, with large generation centres and unidirectional power transmission, to a more agile infrastructure, with many distributed energy resources (DERs) and millions of IoT devices connected and continuously exchanging data with the grid. In smart grids, end-users can actively participate in the transition with their choices (energy communities, electric vehicles, and demand shifting). Several solutions must be adopted and implemented in the grid. In this paper, the authors focus on the impact of optimised demand response (DR) systems on the DN, seeing the benefits of mass deployment.

Electrical utilities consider using DR programs to improve their networks' stability, efficiency, and reliability in response to increased power demand [2] or when the grid is affected by unplanned events [3]. The goal of DR is to allow consumers to contribute to the operation of the electric grid by reducing or shifting their electricity usage to off-peak hours, allowing the grid to relax. Advancements in automated infrastructure and DN technology have enabled residential consumers to participate in demand curtailment plans [4]. Distribution system operators (DSOs) require an advanced management system to proactively manage access to large-scale flexible resources using smart grid technologies to achieve a reliable, robust, cost-effective, and optimised network. A DSO's management system is a centralised platform capable of hosting various optimisation schemes and integrating and managing loads, generation, and other field devices at all distribution voltage levels in real time. It can determine the DR to keep the active DN within safe operating limits. As a result, in addition to financial benefits for DR participants, DSOs can address technical network issues such as voltage or thermal constraints, transformer

overloading, or cable congestion. The potential of DR has already been recognised by authorities worldwide [5], but, despite the efforts of various stakeholders over the last decade, DR programs still need to be established as tools for flexibility provision in power system operations [6].

The existing body of literature delves extensively into DR, covering a spectrum of facets including its social implications and user engagement [7–9], market strategies and economic offers [10–13], the impact on the electricity grid, IoT technologies, the automation required to implement DR mechanisms [14,15], and the role of DSO, energy retailers, and aggregators in the flexibility market [16,17]. In the context of DR's implementation and the impact on the electricity DN, a comprehensive review highlights its advantages in reducing peak loads and enhancing reliability [18]. Zakariazadeh et al. [19] introduce a method for determining load reduction to manage voltage profiles, utilising real-time data from remote units for precise control. Analysing DR's impact based on building energy efficiency, a system-wide DR management model for residential customers is presented in [20], aiming to flatten the total load profile with minimum individual cost, while an ML-based service that requires reinforcement learning (RL)-based optimisation control, namely, grid operation optimisation is underscored in [21]. DR's role as a DER enabler is highlighted in [22], where multi-objective techniques facilitate DR implementation in IoT parking stations for green energy electric vehicle charging.

Optimisation of the DN using DR mechanisms and the introduction of advanced tools for decision support within the network is a debated topic in the literature. Some review papers provide an indication of which models are most used in order to optimise load scheduling, also considering behavioural or intrinsic constraints. In [23], it is pointed out that the main methods used in the literature for studying automatic DR programmes can be divided into classical optimisation methods (linear programming, non-linear programming, and mixed integrator linear programming), heuristic optimisation (particle swarm optimisation, genetic algorithm, fuzzy logic control), and meta-heuristic optimisation (such as the firefly algorithm). The use of machine learning (ML) for grid optimisation problems is more limited, but more and more articles can be found in the literature that use these methods as a model for load management or pricing in DR mechanisms. In [24], it is reported that, considering only ML models applied to the power grid, 52% of the studies about load forecasting use artificial neural networks with one hidden layer and that supervised learning is also widely used (32%), whereas unsupervised learning and deep learning (DL) are rarer. For the problem of optimal scheduling of loads, there is more variety of models used, with 50% using swarm artificial intelligence, 22% using RL, and the remainder distributed between unsupervised, evolutionary algorithms, artificial neural networks, and mechanism design. Further statistics can be found in the review [25], mainly focusing on ML models applied to local energy communities.

Several studies focus on electric vehicles, whose recharging can be shifted in time depending on the network needs and the constraints imposed by the user, e.g., in [26], an optimal solution is proposed based on solving a stochastic problem, considering an uncertainty when the car is recharged and exploiting price variation as an enabler. In [27], three control algorithms based on linear optimisation are used in order to best manage multiple charging points and consider random consumption profiles of electric vehicles. Other types of devices typically used for DR are electrical energy storage systems, which can modulate their energy input or withdrawal according to the needs of the grid. This approach is followed in [28], where an optimal power management mechanism is presented for grid-connected PV systems with storage based on dynamic programming. In [29], a load scheduling optimisation has been realised for domestic consumers in order to implement DR mechanisms based on price differences. DR extends to microgrids, as discussed in-depth by Osman et al. [18], who propose a multi-objective darts game theory optimisation algorithm that is particularly useful post-disaster for forming microgrids and implementing emergency DR programs. In [30], both frequency and voltage issues are addressed by proposing a centralised, adaptive load-shedding algorithm, leveraging

phasor measurement units (PMUs). Zhu et al. [31] present a distributed algorithm to flatten load profiles while minimising individual customer costs.

As described previously, the literature concerning the application of DR mechanisms presents numerous models and mechanisms mainly simulated on electrical datasets, often aimed at cost-effectiveness analyses and dynamic pricing programme scheduling. A lack is noted in the availability of real data and the use of innovative models based on machine learning. Indeed, by addressing these challenges and shortages, this paper aims to prepare an optimal management mechanism of network flexibility resources with DR, using RL analytics applied to a real Italian DN. Within the tool, a forecast model, i.e., a recurrent neural network (RNN) based on gated recurrent unit (GRU) layers, allows for identifying the trend in the load and generation in the next 24 h; this model has been used for online forecasting based on streaming dataset collected from real grids on a continuous basis. The optimisation and forecasting ML models have been chosen as an example of a real and effective application and do not claim to be the most effective algorithms.

In addressing the gap in the literature, the contributions of this paper can be summarised as follows:

- Modelling and testing an original ML-based forecasting model, with continuous model training, applied to the PV production trend and industrial and domestic consumption trends;
- Modelling and testing an original ML model for optimising an electricity grid, leveraging local flexibility resources;
- Integration of the models and deployment on a real power grid, analysing the benefits for the DN.

This paper is organised as follows: the introductory section presents the problem, the literature analysis, and the contributions of this article to the existing literature. Section 2 presents the ML-based forecasting models used for load and generation prediction and grid optimisation, as well as the methodological outline of this study. Section 3 presents the real case study of the Terni DN. Section 4 shows the results of applying the models in the real case, using the previously defined indicators, and finally, Section 5 reports the conclusions.

The work performed in this paper was founded by the IoT-NGIN European Project [32] and leverages the experience and contribution of some of the Consortium partners.

2. Methodology and Models

In order to evaluate the effectiveness of the ML tools developed, some energetic indicators highlighting the benefits of the DN are used. These indicators are network losses, the self-consumption rate (SCR), the self-sufficiency rate (SSR), and the reverse power flow (RPF), i.e., the energy flowing from the DN to the primary substation (PS) in the opposite direction to the one commonly followed. The equations to calculate SSR and SCR are as follows [33].

$$SCR = 1 - \frac{E_i}{E_p} \quad (1)$$

$$SSR = \frac{E_p - E_i}{E_a + E_p + E_i} \quad (2)$$

where E_i is the energy fed into the grid, E_p is the energy produced, and E_a is the energy absorbed by the grid.

2.1. Forecasting Model

The forecasting model is based on online learning, i.e., the training of ML models is performed continuously when new data are available. This paradigm acquires great importance in the context of IoT due to the large number of sensors or devices that can be present in common scenarios and the large amount of information captured by them dynamically. The online learning service supports both application programming interface

(API) REST requests and streaming data since most IoT devices generate communication flows in real time.

Forecasting the power generated by the grid over a future time frame, based on historical data collected from grid sensors, consists of predicting the following power values from the last past measured ones. For example, given the last measured 36 h of generated power, the model can predict the next 24 h. Forecasting is computed using a trained ML model that infers the prediction based on a given input tensor of measured values. This is an example of a univariate time series prediction problem addressed using ML modelling techniques. Before the ML model is ready for computing predictions, it must be trained with historically generated power data. As the data are continuously pushed by the sensor into an MQTT topic, learning is conducted using an IoT-NGIN MLaaS online learning service [21] that collects online data from the topic and trains the model after applying the following data pre-processing procedure: (i) power is computed from sensor data, (ii) the data are resampled by averaging over each hour, (iii) the data are scaled in the range [0, 1] as required by the ML algorithms applied during training, and (iv) the data are windowed across collected power data to create a dataset with the shape required by the ML training algorithms.

The ML model designed for forecasting the generated power of the grid is a DL model with layers based on RNNs [22]. RNNs have been demonstrated to work well in predicting the future behaviour of time series, although they present some disadvantages, such as the vanishing gradient problem [34]. To overcome this problem, alternative architectures for ML models, evolving from RNNs, have been proposed, including LSTM and GRUs [35]. For the power forecasting problem, the GRU is used since it solves the vanishing gradient problem suffered by the original RNN and converges faster than other types of RNNs (e.g., LSTM). After the recurrent layers, fully connected layers are added for applying linear transformations to the outputs of the GRU layer. Figure 1 depicts the layers of the implemented architecture.

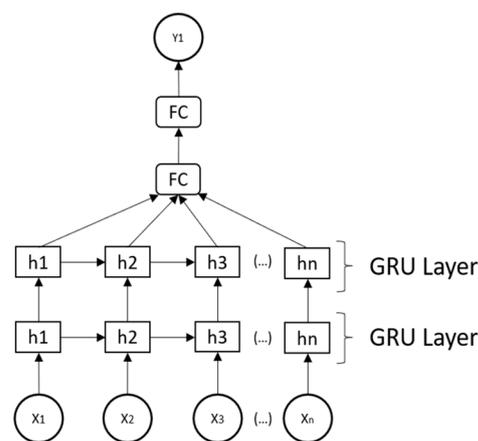


Figure 1. DL architectures for power forecasting.

This proposed model is trained to minimise the mean squared error (MSE) between the real values and the predictions. In order to understand the performance of the forecasting model, the lower the MS, the better the effectiveness. The MSE is calculated using Equation (3):

$$MSE = \frac{1}{n} \sum_{t=1}^n (x_t - \hat{x}_t)^2 \quad (3)$$

where x_t is the real value of the PV production or the building consumption time series at hourly intervals t of the evaluation period and \hat{x}_t is the produced forecast of the respective model.

A module is implemented for explainable artificial intelligence (XAI), which attempts to explain the predictions made to answer the question: Why has the model made this

prediction? XAI is a set of methods and processes that help to comprehend and trust the prediction of the ML model. Moreover, it helps to characterise model performance by providing the impact of the input data for a given prediction, adding transparency to the prediction and capacity for model bias detection.

The OL service is deployed into MLaaS using Kserve [36] with Kubeflow [37]. Kserve allows for deploying three types of components: predictor, transformer, and explainer. Each of these components exposes a REST API as an HTTP service.

The monitoring service consists of an HTTP endpoint deployed using the FastAPI framework [38], a Prometheus engine [39], and a Grafana web tool [40].

2.2. Grid Optimisation Model

The grid optimisation model, aiming to maximise its SCR and SSR ratios, is designed as an RL-based optimisation model [41]. RL is an ML discipline centred on the learning of optimal behavioural policies for the decision-making of a group of agents interacting in a common environment, leading to a maximisation of a cumulative reward (i.e., expert-defined performance metric). In the context of control systems, the learned policy allows for deploying deterministic or stochastic control logic/instructions for agents interacting with the end system, such as maximising the grid SCR and SSR ratios. Some of the main concepts of RL involving the environment, agents, states, actions, rewards, observations, and policies are shown in Figure 2:

- The environment refers to the physical or simulated space that the agents interact with.
- Agents are entities that are affected by and are in a position to interact with the environment by taking action.
- Agents take a state (e.g., vector) that represents their status at every point in time. States are defined as a discrete or a continuous closed set.
- A set of actions is defined for the agents to take. This group is defined as a discrete or a continuous closed set.
- Rewards are given by the environment after the undertaking of actions by the autonomous agents.
- Observations are pre-processed snapshots (e.g., in the form of vectors) collected after each transition, which gather relevant variables from the environment as well as the previous state and the actions taken, resulting in the state and the observed reward.
- Policies, in broad terms, are the learned (deterministic or stochastic) mapping between the set of states and the set of actions.

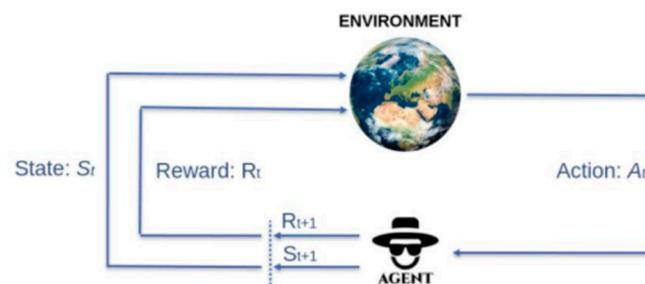


Figure 2. RL concept.

Depending on the problem to be tackled, experience is gathered with interactions in either a simulated or a real (digital or physical) environment. Based on the knowledge of the environment, it is possible to distinguish between model-free and model-based approaches [41]. For the reinforcement algorithm to learn/approximate an optimal policy, it needs to buffer enough experience from the environment. The experience collected in the model comprises transitions (steps within episodes), which are made up of observations including relevant variables from the environment, the previous agent status, the action taken, the following states, and the observed reward. In the following, the RL elements for

the grid optimisation model are described. The RL service is deployed using KServe with the Tensorforce framework [42].

2.2.1. Environment

As interactions with the real grid are not possible, a grid simulator is implemented based on the pandapower framework [43]. This framework uses the Flow Pypower power solver [44] to create a calculation network program to automate and optimise power systems. As a result, the simulator describes the same environment state as the real grid when connected with the same sources of power generation and consumption. The operation of the simulator is as follows: once the electrical network is specified in pandapower, the simulator reads the power loads demanded by all the consumer groups connected to the grid and the power generated along a day. The data have a resolution of 15 min, so there are 96 values per day of domestic and industrial consumers' power demand. Next, the simulator introduces the loads on the grid and performs a simulation. Then, the simulator outputs the grid state, consisting of the parameters to be optimised: network losses, SCR, SSR, and RPF.

2.2.2. States

The optimiser acts on the customers' energy demand: it modulates the distribution of energy demand throughout the day, so it is necessary to know the state of energy demand or, equivalently, the distribution of current energy demand. Two types of customers can be distinguished: domestic and industrial. There are 13 loads (client groups) of each. To understand better what the data look like, a pre-analysis of a dataset collected over a year was carried out. Figures 3 and 4 show the average distribution of, respectively, each domestic and industrial load for one day. The available data, provided by the grid owner, do not show the annual curve for each user but provide an averaged curve relating to the industrial cluster and one relating to the domestic cluster, as well as the photovoltaic production curve per unit of installed power. Furthermore, the capacities of each load and each photovoltaic system are known. In fact, the user curves are proportional to each other according to the installed power. To simplify the optimisation model, it only acts on higher loads, i.e., those having a stronger influence on the grid (the top eight loads, combining four domestic and four industrial loads, are selected since they cover, respectively 70% and 77% of the total cluster load). It is considered that the flexibility of users, i.e., the ability to vary their consumption curve, is proportionally the same for all users, with the difference being that larger users are able to provide higher-performance services for the entire DN. Reducing the number of actors simplifies the computational aspects of the problem and speeds up the solution.

It can be seen that the trend in domestic users peaks in the morning hours, and the two main peaks are around lunchtime and dinnertime. This is an averaged curve, which therefore considers full- and part-time workers, students, and pensioners. The curve for industrial and commercial users presents a peak in the early hours of the day, corresponding to the start-up of engines and other consumer devices, and a slow decrease in consumption, interrupted only by a recovery around noon.

Discretisation of States

As the power demand continuously changes, it must be discretised, as a discrete state space is needed for training the optimiser in a reasonable time. The state space should be as small as possible in order to be completely explored by the optimiser so that the complete policy is learned. However, an excessive reduction of the time interval would increase the state space, imposing higher exploration time on the optimiser training process. After initial tests with quarter-hourly discretisation, it was noted that an extremely large number of episodes were required by the optimiser to explore the state space, despite running on a very high-performance HPC cluster. Hence, it was decided to carry out hourly load discretisation. An approach with much lower discretisation times, such as

15 s or less, would be desirable for electrical aspects (overcurrents, voltage issues. . .), for energy aspects, such as SCR or SSR, as the objective of the grid optimiser, there is no need to use short analysis times. In this respect, the Italian legislation also uses hourly resolution times for the calculation of the SCR for the calculation of incentives for renewable energy communities [45].

Therefore, the states consist of the sum of all the loads, discretised by the hour. This approach ended with 48 states, consisting of two vectors of 24 elements each (one for domestic consumers and another for industrial consumers). Each state could take a continuous value in the range [0, 1], that is, the percentage over the total daily demand for each daily hour. Therefore, the range [0, 1] was also discretised into 10 bins. Furthermore, it is possible to apply a further normalisation by dividing the state value by a “load threshold” so that the possible states of demand in the range [0, 1] are constraints (to reduce the state space dimension) to those values with a higher likelihood of being explored and applied by the optimiser. In this way, the simulator can explore more states in a reasonable time.

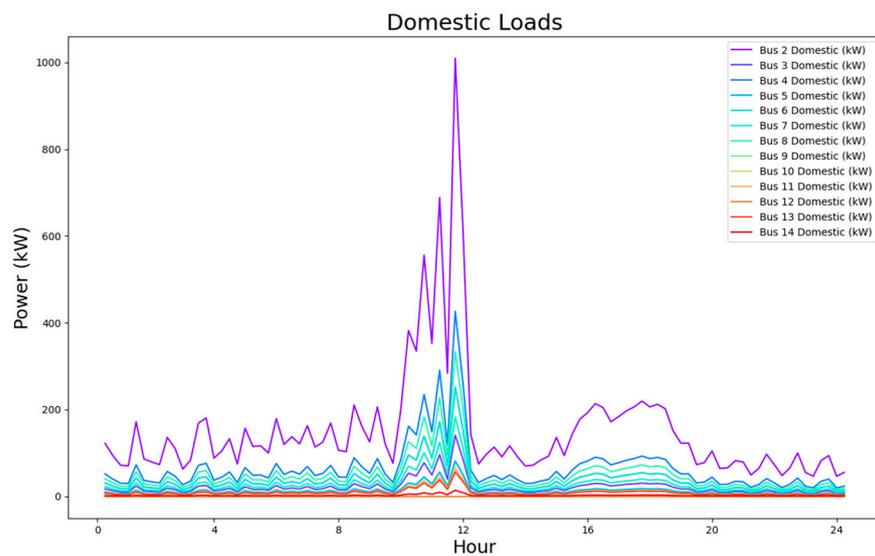


Figure 3. Energy demand for domestic clusters.

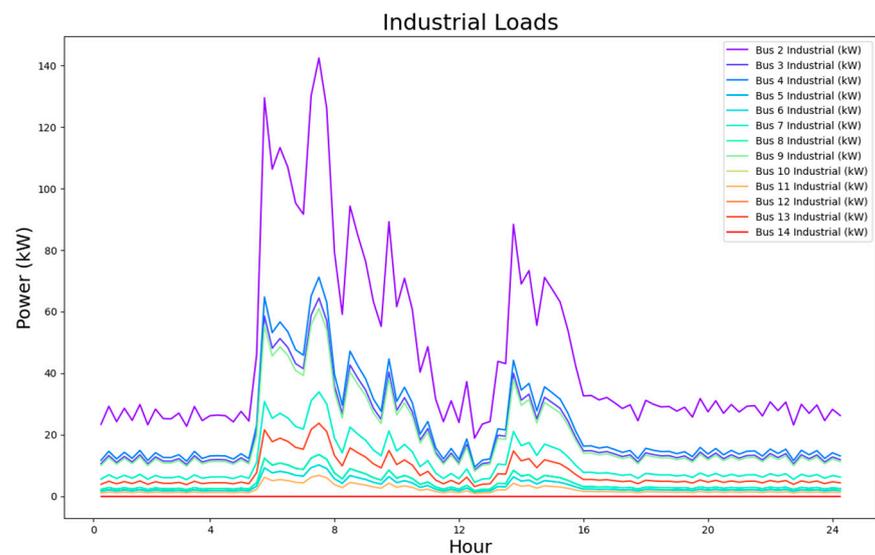


Figure 4. Energy demand for industrial clusters.

2.2.3. Actions

The RL-based optimiser seeks to optimise the energy demand of the grid. To do so, the distribution of energy demand of each load is shifted throughout the day. This optimiser uses four sets of discrete actions (see Table 1).

Table 1. Action subset.

Action Subset	Length
Domestic demand cluster selection	4
Industrial demand cluster selection	4
Domestic initial time slot	24
Industrial initial time slot	24
Final domestic time slot	2
Final industrial time slot	2
Demand shift for domestic cluster	2
Industrial shift for domestic cluster	2

The first action set is the load selection. This action set includes eight possible actions (select one load from the eight available). The second action set determines the start and end time of the energy displacement. The third action set determines the percentage of energy that is shifted at different times. Three different energy shifts were considered: 1%, 5%, and 10%. For example, if from 8:00 a.m. to 9:00 a.m. there is an energy demand of 10 kW and it is chosen to make a 10% shift to the time slot between 9:00 a.m. and 10:00 a.m., the second time slot will increase the power demand by 1kW, and the energy demand from 8:00 a.m. to 9:00 a.m. will decrease it by the same amount. This resulted in a set of 147,456 possible actions.

The energy demand shift presents several restrictions: (i) the total energy demanded by a load must remain constant throughout the day, (ii) the shift can only occur within two contiguous time slots, i.e., the displacement cannot be applied to any arbitrary distant time slots, and (iii) the demand cannot be negative.

The power demand load describes the daily demand, with a 15 min interval, resulting in a 96-value vector for a full day. In order to be aligned with the state discretisation of 1 h (as described in Section 2.2.2), the optimiser resamples the loads by averaging the samples within each hour.

2.2.4. Rewards

The objective of the system is to optimise the operation of the grid. This version tries to maximise the SSR and SCR. The reward function is the mean of both parameters. In order to obtain greater flexibility during the training process, the reward is computed as a weighted average:

$$R = \frac{1}{2}(\alpha \cdot SCR + (1 - \alpha) \cdot SSR) \quad (4)$$

By tuning the α parameter, it is possible to address the optimisation needs and promote SSR over SCR (or vice versa). In this way, we can force SCR and SSR to evenly contribute to the reward, by setting $\alpha = \frac{1}{2}$, or make the reward equal to the SCR, by setting $\alpha = 1$.

2.2.5. Agents

The optimisation model includes agents for the training of PPO, DQN, and DDQN [41]. The optimisation model was implemented using the Tensorforce RL framework. For training, some improvements were conceived and implemented, including:

- Variable number of steps per episode: some experiments show that the reward fell down after some steps and was incapable of going up for the rest of the episode. To avoid wasting time in the training process, the episode is concluded when the obtained rewards go down a customisable threshold of 20%.

- Variable exploration time: the training process supports the configuration of a variable exploration rate, which can be diminished as the learning process progresses over more learning episodes.
- State density matrix: the optimisation training process registers all the states visited by the agent, intended to give a clear vision of the agent's preferable combination of states, aiming to understand the reasons for the agent to choose such a state combination as optimal. This state's matrix consists of 48 columns, corresponding to the 24 h of both domestic and industrial consumers, and 10 rows, corresponding to the 10 state bins available for each state.

3. Case Study

ASM Terni S.p.A. is a multi-utility in Terni municipality, an Italian city in the Umbria area, which manages operations for the administration of electricity, gas, and water networks and is responsible for waste collection and disposal. ASM, through its business unit Terni Distribuzione Elettrica, owns and operates the DN serving 65,000 users. ASM manages 65,000 smart meters, 700 secondary substations, and three PS. In total, 1276 LV solar plants and 53 MV PV make up Terni's distributed generation's installed power of 63.4 MVA; 40% of the DN's yearly energy consumption is met by this DER. Additional information about ASM DN can be found in [46–48], which, respectively, exploited the digital twin concept in the power DNs, investigated false data impact, and analysed the integration of a hydrogen supply chain for feeding vehicles based on fuel cell technology.

In order to evaluate the effectiveness of the ML tools developed, a portion of the Terni network in MV was used as a case study, as depicted in Figure 5, consisting of 14 nodes, one of which represents a PS, while the others represent secondary substations and 30 power lines. Each of the 13 secondary substations feeds a different capacity of industrial loads, for a total of 650 kW of installed power and domestic loads (3374 kW), and some of them have photovoltaic systems (3663 kW). Regarding the historical data from the last 5 years, a yearly average curve of industrial and domestic loads and PV production is available, with a granularity of 15 min. These curves are used as a baseline in the optimisation tool. A PMU is located in the PS, which is capable of monitoring voltage, active and reactive power, and other electrical parameters in real time; this sensor is used for forecasting analysis.

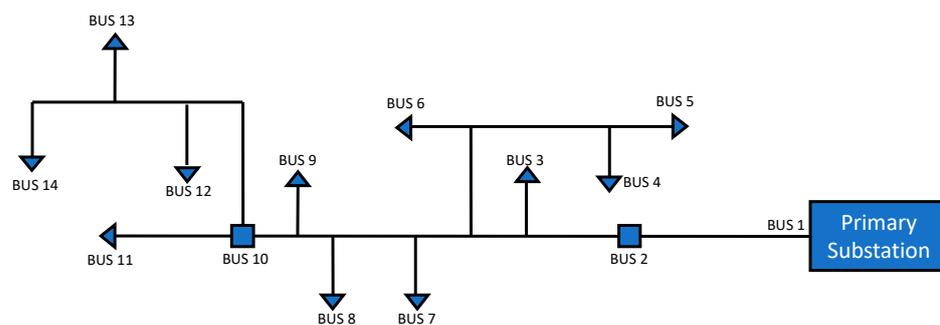


Figure 5. Schematic diagram showing the portion of the 20 kV grid used as a case study. Some nodes are connected via different sections of the line, for example, the overhead and cable types or with different typologies. For reasons of graphic clarity, they were not distinguished here.

Experiments were conducted using an Atos HPC cluster of 12 nodes with 720 cores each and a total memory of 2968 Gb. Experiments were conducted on 16 cores and 128 Gb RAM, which took days to complete.

4. Results

4.1. Forecasting Results

To verify that the proposed ML architectures are valid, it was used to predict the voltage trend and the active and reactive power of the PMU, located in a PS, a smart meter, and two power quality analysers. Forecasting was carried out using a dataset collected for

18 months through the MQTT topic broadcast by sensors, by training the proposed models. The predicted data were chosen to have a time horizon of 24 h and a resolution of one hour, whose forecasting was based on the last 36 h of data averaged every hour, resulting in an input tensor of 36 values. The proposed models were trained with the hyper-parameters listed in Table 2. The optimiser, called Adam, is an algorithm for first-order gradient-based optimisation of stochastic objective functions, based on adaptive estimates of lower-order moments [49].

Table 2. Hyper-parameters for power forecasting models.

Hyper-Parameters	Value
Epochs	50
Learning Rate	0.005
Optimiser	Adam
Loss Function	Mean squared error
Batch size	128

After training the model, inferences for forecasting the generated power were computed, comparing predictions with actual measure power. Figure 6 shows the actual power data (orange line), inferences performed using the ML model (blue points), and the forecasting intervals with 90% confidence (blue area). The forecasting intervals can be computed when the errors between the actual data and the model predictions present a distribution that can be considered Gaussian. To confirm the errors came from a Gaussian distribution, they were subjected to normality tests: Shapiro–Wilk [50], Anderson–Darling [51], and D’Agostino–Pearson [52]. These tests consist of statistical hypothesis tests that check whether the data contain specific properties. Thus, two hypotheses are defined: the null and alternative hypotheses. The null hypothesis supports that the data probably come from a normal distribution, while the alternative hypothesis suggests that the data present a different distribution. The statistical test returns a probability known as the *p*-value. If the result presents a value lower than the defined significance level (0.05 in this case), then the null hypothesis must be rejected so that the data distribution can be assumed to be normal. Table 3 lists the *p*-values obtained.

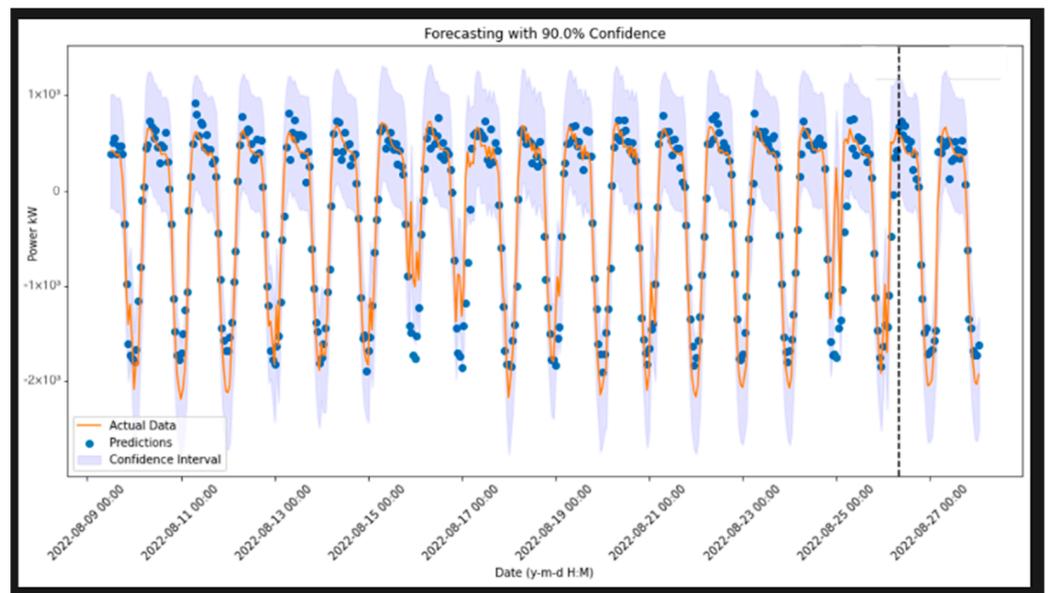


Figure 6. Training results for generated power forecasting for 20 days of analysis in August 2022.

Table 3. Normality test results (*p*-values) for power forecasting models.

Normality Test	Value
Shapiro–Wilk	0.47
Anderson–Darling	0.76
Agostino–Pearson	0.10

Within the confidence intervals, the model learns the seasonal variations in the generated power quite well, as depicted in Figure 6, obtaining an MSE of 0.009 during training. As shown in Figure 6, there is good matching between the observed and predicted values within the confidence interval, which captures both the trend and seasonality in the future series behaviour.

4.2. Results of Grid Optimisation Using the Demand Response

The grid optimisation model was trained with the agent’s hyperparameters listed in Table 4. In our experiment, $\alpha = 1$ was chosen because, from the DSO’s point of view, it is more convenient to implement a DR mechanism that favours the self-consumption of local DER to reduce RPF and reduce demand from the grid. A high SCR means low grid losses. SSR, on the other hand, is more related to consumers and the possibility of becoming grid-independent.

Table 4. Hyper-parameters for grid optimisation models.

Hyper-Parameter	Value
RL Algorithm	PPO
Learning Rate	0.001
Optimiser	Adam
Multistep	10
Batch size	1
Reward discount	10%
α	1
Load threshold	30%

Initial experiments were conducted to determine each episode’s optimal duration (i.e., #steps), which resulted in 2000 steps. Also, as described before, better optimisation results were obtained by applying a normalised discretisation of the state space. The results show the agent was exploring the state space much better than if normalised discretisation had not been applied. Figure 7 shows the matrix space of visited loads after completing the experiment.

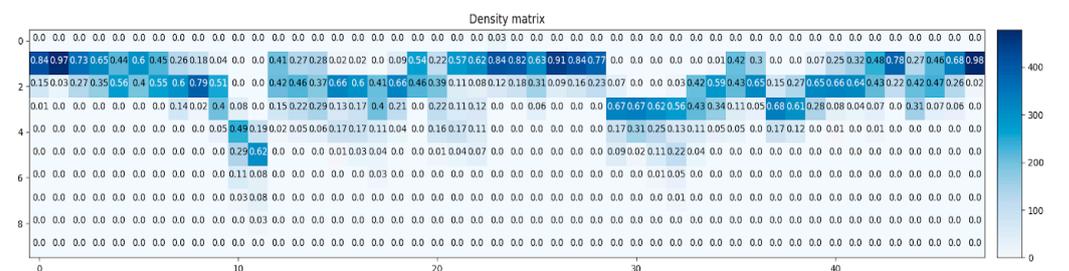


Figure 7. Load state space.

As a result of training, the reward was improved (Figure 8) after applying learnt actions on the user’s demand load, resulting in better SCR. The achieved SCR of around 52% is very high, considering that the portion of the grid in question has an average daily production of 23.6 MWh and a consumption of 12.9 MWh. Thus, the maximum SCR that can theoretically be achieved is 55%, which occurs in the case of completely zero energy

absorption from the grid. In Figure 8, five different episodes of the training have been depicted.



Figure 8. Reward evolution over episodes considering $\alpha = 1$.

In Figure 9, domestic loads are shown before (i.e., left) and after being optimised (i.e., right) along hours in the day on the x -axis. A similar trend is recognised for industrial loads. The optimised load shows a shape closer to the theoretical optimal load, with a bell shape, where the load is concentrated at noon. For the optimiser to reach the optimal shape, an extremely longer episode duration, beyond our computing capability, is required to exhaust the exploration of the state space in order to compute the optimal policy. However, the initial training results, within our computational constraints, show progress in the right direction.

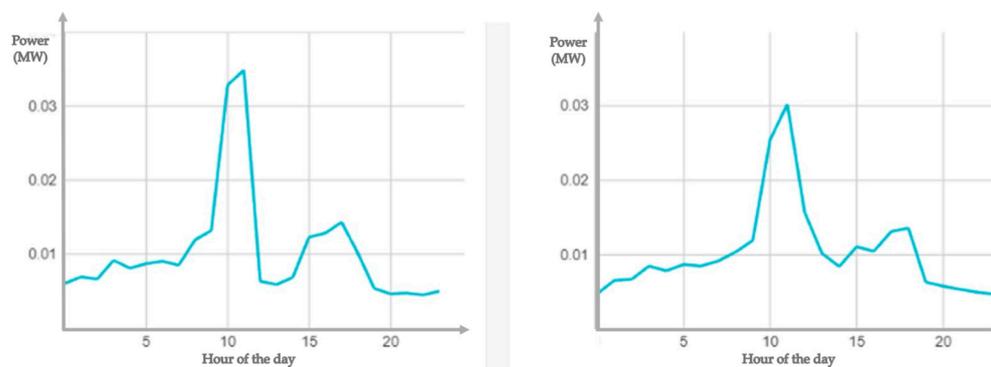


Figure 9. Domestic demand load before (on the left) and after (right) the training.

Further experiments are required to improve the optimised SSR and SCR values to become close to the theoretical ones. This can be accomplished by reducing the dimension of the state space so that within the affordable computation time, the agent can exhaust the state space when searching for optimal user demand loads.

5. Conclusions

This paper shows some of the services developed within the European IoT-NGIN project and their integration within the electricity infrastructure of ASM Terni. This paper aimed to show the application of ML tools within the DN and highlight the added value that can be obtained. Analytics for forecasting and grid optimisation were presented. The GRU-based forecasting applied to a PMU, a smart meter, and two power quality analysers indicated the active and reactive power and voltage values in the next 24 h. At the same time, the grid optimisation tool, which consists of an electrical simulation part using pandapower and an RL-based optimisation part for the consumer curves, managed the decisions of the DR mechanisms on a portion of the medium voltage network (14 nodes) of ASM Terni. Technical details of the services were provided, accompanied by some results.

A limitation of this paper is its simplified approach to user participation in DR mechanisms, which does not take into account the real availability of flexibility resources (storage, loads that can be shifted over time, etc.), which was assumed to be a portion of the total load.

In future research, we plan to test additional ML models for application in these topics, compare the results, and apply the models to a more significant portion of the electrical network. Furthermore, we plan to evaluate how the integration between the forecasting service and the optimisation service can be strengthened using single analytics that exploit the results provided by both services in a cascade.

Author Contributions: Conceptualisation, M.A.B., J.M. and J.G.; methodology, M.A.B. and J.G.; software, J.M. and J.G.; validation, M.A.B., M.G. and F.S.; formal analysis, M.A.B., J.G. and F.S.; investigation, M.G. and J.M.; resources, F.S.; data curation, M.A.B. and M.G.; writing—original draft preparation, M.G., J.M. and J.G.; writing—review and editing, M.A.B., J.M. and J.G.; visualisation, J.G. and J.M.; supervision, F.S. and J.G.; project administration, M.A.B.; funding acquisition, F.S. and J.G. All authors have read and agreed to the published version of the manuscript.

Funding: The work presented is based on research conducted within the framework of the H2020 European Commission project IoT-NGIN (grant agreement no. 957246). The content of this paper is the sole responsibility of its author and does not necessarily reflect the views of the EC.

Data Availability Statement: Restrictions apply to the availability of these data. Data was obtained from ASM Terni S.p.A. and are available from the author Francesca Santori with the permission of ASM Terni S.p.A.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

API	application programming interface
DER	distributed energy resources
DL	deep learning
DN	distribution network
DR	demand response
DSO	distribution system operator
GRU	gated recurrent unit
LSTM	long short-term memory
ML	machine learning
MQTT	message queuing telemetry transport
MSE	mean squared error
PMU	phasor measurement unit
PS	primary substation
RL	reinforcement learning
RNN	recurrent neural network
RPF	reverse power flow
SCR	self-consumption rate
SSR	self-sufficiency rate

References

1. European Environment Agency. Total Greenhouse Gas Emission Trends and Projections in Europe. 2022. Available online: <https://www.eea.europa.eu/ims/total-greenhouse-gas-emission-trends> (accessed on 10 September 2023).
2. Albadi, M.H.; El-Saadany, E.F. Demand response in electricity markets: An overview. In Proceedings of the 2007 IEEE Power Engineering Society General Meeting, Tampa, FL, USA, 24–28 June 2007; pp. 1–5. [CrossRef]
3. Chen, C. Demand response: An enabling technology to achieve energy efficiency in a smart grid. In *Application of Smart Grid Technologies*; Lamont, L.A., Sayigh, A., Eds.; Academic Press: Cambridge, MA, USA, 2018; pp. 143–171. [CrossRef]
4. Haider, H.T.; See, O.H.; Elmenreich, W. A review of residential demand response of smart grid. *J. Renew. Sustain. Energy Rev.* **2016**, *59*, 166–178. [CrossRef]
5. European Commission. *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee, the Committee of the Regions and the European Investment Bank—Clean Energy for All Europeans*; European Commission: Luxembourg, 2016.

6. Willems, B.; Zhou, J. The clean energy package and demand response: Setting correct incentives. *Energies* **2020**, *13*, 5672. [CrossRef]
7. Srivastava, A.; Passel, S.V.; LAes, E. Assessing the success of electricity demand response programs: A meta-analysis. *Energy Res. Soc. Sci. J.* **2018**, *40*, 110–117. [CrossRef]
8. Larsen, S.P.; Johra, H. User engagement with smart home technology for enabling building energy flexibility in a district heating system. *IOP Conf. Ser. Earth Environ. Sci.* **2019**, *352*, 012002. [CrossRef]
9. Eguiarte, O.; de Agustín-Camacho, P.; Garrido-Marijuán, A.; Romero-Amorrortu, A. Domestic space heating dynamic costs under different technologies and energy tariffs: Case study in Spain. *Energy Rep.* **2020**, *6*, 2205. [CrossRef]
10. Aalami, H.; Moghaddam, M.P.; Yousefi, G. Demand response modeling considering interruptible/curtailable loads and capacity market programs. *Appl. Energy* **2010**, *87*, 243–250. [CrossRef]
11. Mohsenian-Rad, A.; Wong, V.W.S.; Jatskevich, J.; Schober, R.; Leon-Garcia, A. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Trans. Smart Grid.* **2010**, *1*, 320–331. [CrossRef]
12. Nguyen, D.T.; Nguyen, H.T.; Le, L.B. Dynamic pricing design for demand response integration in power distribution networks. *IEEE Trans. Power Syst.* **2016**, *31*, 3457–3472. [CrossRef]
13. Rahmani-andebili, M. Modeling nonlinear incentive-based and price-based demand response programs and implementing on real power markets. *Electr. Power Syst. Res.* **2016**, *132*, 115–124. [CrossRef]
14. Kathirgamanathan, A.; Péan, T.; Zhang, K.; De Rosa, M.; Salom, J.; Kummert, M.; Finn, D.P. Towards standardising market-independent indicators for quantifying energy flexibility in buildings. *Energy Build.* **2020**, *220*, 110027. [CrossRef]
15. Lu, R.; Hong, S.H.; Yu, M. Demand response for home energy management using reinforcement learning and artificial neural network, *IEEE Trans. Smart Grid.* **2019**, *10*, 6629–6639.
16. Gao, N.; Ge, S.; Tian, Y.; You, C. A review of decision-making strategies of profit seeking demand response aggregators. In Proceedings of the 2020 IEEE sustainable power and energy conference (ISPEC), Chengdu, China, 23–25 November 2020; pp. 2135–2140.
17. Lu, X.; Li, K.; Xu, H.; Wang, F.; Zhou, Z.; Zhang, Y. Fundamentals and business model for resource aggregator of demand response in electricity markets. *Energy* **2020**, *204*, 117885. [CrossRef]
18. Osman, S.R.; Sedhom, B.E.; Kaddah, S.S. Impact of implementing emergency demand response program and tie-line on cyber-physical distribution network resiliency. *Sci. Rep.* **2023**, *13*, 3667. [CrossRef]
19. Zakariazadeh, A.; Homaei, O.; Jadid, S.; Siano, P. A new approach for real time voltage control using demand response in an automated distribution system. *Appl. Energy* **2014**, *117*, 157–166. [CrossRef]
20. Safdarian, A.; Fotuhi-Firuzabad, M.; Lehtonen, M. A distributed algorithm for managing residential demand response in smart grids. *IEEE Trans. Ind. Inf.* **2014**, *10*, 2385–2393. [CrossRef]
21. Borou, S.; Anastasakis, Z.; Voulkidis, A.; Velivassaki, T.; Mira, J.; Moreno, I.; Gorroñoigoitia, J.; Bardisbanian, H. Enhanced IoT Federated Deep Learning/Reinforcement ML. IoT NGIN Report. Available online: https://iot-ngin.eu/wp-content/uploads/2023/01/IOT-NGIN_D3.3_V1.0_PENDING_EC_APPROVAL.pdf (accessed on 10 September 2023).
22. Rumelhart, D.E.; Hinton, G.E.; Williams, R.J. *Learning Internal Representations by Error Propagation*; Institute for Cognitive Science, University of California: San Diego, CA, USA, 1985.
23. Tiwari, A.; Pindoriya, N.M. Automated Demand Response in Smart Distribution Grid: A Review on Metering Infrastructure, Communication Technology and Optimization Models. *Electr. Power Syst. Res.* **2022**, *206*, 107835. [CrossRef]
24. Antonopoulos, I.; Robu, V.; Couraud, B.; Kirli, D.; Norbu, S.; Kiprakis, A.; Flynn, D.; Elizondo-Gonzalez, S.; Wattam, S. Artificial intelligence and machine learning approaches to energy demand-side response: A systematic review. *Renew. Sustain. Energy Rev.* **2020**, *130*, 109899. [CrossRef]
25. Hernandez-Matheus, A.; Löschenbrand, M.; Berg, K.; Fuchs, I.; Aragüés-Peñalba, M.; Bullich-Massagué, E.; Sumper, A. A systematic review of machine learning techniques related to local energy communities. *Renew. Sustain. Energy Rev.* **2022**, *170*, 112651. [CrossRef]
26. Mohsenian-Rad, H.; Ghamkhari, M. Optimal charging of electric vehicles with uncertain departure times: A closed-form solution. *IEEE Trans. Smart Grid* **2014**, *6*, 940–942. [CrossRef]
27. Van Der Kam, M.; van Sark, W. Smart charging of electric vehicles with photovoltaic power and vehicle-to-grid technology in a microgrid; a case study. *Appl. Energy* **2015**, *152*, 20–30. [CrossRef]
28. Riffonneau, Y.; Bacha, S.; Barruel, F.; Ploix, S. Optimal power flow management for grid connected pv systems with batteries. *IEEE Trans. Sustain. Energy* **2011**, *2*, 309–320. [CrossRef]
29. Haider, H.T.; Muhsen, D.H.; Al-Nidawi, Y.M.; Khatib, T.; See, O.H. A novel approach for multi-objective cost-peak optimization for demand response of a residential area in smart grids. *Energy* **2022**, *254*, 124360. [CrossRef]
30. Gerke, B.F.; Zhang, C.; Murthy, S.; Satchwell, A.J.; Present, E.; Horsey, H.; Wilson, E.; Parker, A.; Speake, A.; Adhikari, R.; et al. Load-driven interactions between energy efficiency and demand response on regional grid scales. *Adv. Appl. Energy* **2022**, *6*, 100092. [CrossRef]
31. Zhu, J.R.; Jin, Y.; Zhu, W.; Lee, D.-K.; Bohlooli, N. Multi-objective planning of micro-grid system considering renewable energy and hydrogen storage systems with demand response. *Int. J. Hydrogen Energy* **2023**, *48*, 15626–15645. [CrossRef]
32. IoT-NGIN Project Website. Available online: <https://iot-ngin.eu/> (accessed on 10 September 2023).
33. Energy Community. *POLICY GUIDELINES by the Energy Community Secretariat on the Grid Integration of Prosumers*; Energy Community: Vienna, Austria, 2018.

34. Pascanu, R.; Mikolov, T.; Bengio, Y. On the difficulty of training recurrent neural networks. In Proceedings of the International Conference on Machine Learning, Atlanta, GA, USA, 16–21 June 2013; pp. 1310–1318.
35. Dey, R.; Salem, F.M. *Gate-Variants of Gated Recurrent Unit (GRU) Neural*; Department of Electrical and Computer Engineering: Vancouver, BC, Canada, 2017.
36. Kserve. Available online: <https://kserve.github.io/website/0.10/> (accessed on 10 September 2023).
37. Kubeflow. Available online: <https://www.kubeflow.org/> (accessed on 10 September 2023).
38. FastAPI. Available online: <https://fastapi.tiangolo.com> (accessed on 10 September 2023).
39. Prometheus. Available online: <https://prometheus.io/> (accessed on 10 September 2023).
40. Grafana. Available online: <https://grafana.com/> (accessed on 10 September 2023).
41. Sutton, R.S.; Barto, A.G. Reinforcement learning: An introduction. MIT Press: Cambridge, MA, USA, 2018.
42. Tensorforce. Available online: <https://github.com/tensorforce/tensorforce> (accessed on 10 September 2023).
43. Pandapower. Available online: <http://www.pandapower.org/> (accessed on 10 September 2023).
44. Pypower. Available online: <https://pypi.org/project/PYPOWER/> (accessed on 10 September 2023).
45. GSE, Gruppi di Autoconsumatori di Energia Rinnovabile che Agiscono Collettivamente e Comunità di Energia Rinnovabile. Available online: <https://www.gse.it/servizi-per-te/autoconsumo/gruppi-di-autoconsumatori-e-comunita-di-energia-rinnovabile/documenti> (accessed on 10 September 2023). (In Italian).
46. Bragatto, T.; Bucarelli, M.A.; Carere, F.; Cresta, M.; Gatta, F.M.; Geri, A.; Maccioni, M.; Paulucci, M.; Poursoltan, P.; Santori, F. Near real-time analysis of active distribution networks in a Digital Twin framework: A real case study. *Sustain. Energy Grids Netw.* **2023**, *35*, 101128. [[CrossRef](#)]
47. Bragatto, T.; Bucarelli, M.A.; Bucarelli, M.S.; Carere, F.; Geri, A.; Maccioni, M. False Data Injection Impact on High RES Power Systems with Centralized Voltage Regulation Architecture. *Sensors* **2023**, *23*, 2557. [[CrossRef](#)] [[PubMed](#)]
48. Bragatto, T.; Bucarelli, M.A.; Carere, F.; Cavadenti, A.; Santori, F. Optimization of an energy district for fuel cell electric vehicles: Cost scenarios of a real case study on a waste and recycling fleet. *Int. J. Hydrogen Energy* **2022**, *47*, 40156–40171. [[CrossRef](#)]
49. Kingma, D.P.; Ba, J.L. ADAM: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
50. Shapiro, S.S.; Wilk, M.B. An Analysis of Variance Test for Normality. *Biometrika* **1965**, *52*, 591–611. [[CrossRef](#)]
51. Anderson, T.W.; Darling, D.A. Asymptotic Theory of Certain “Goodness of Fit” Criteria Based on Stochastic Processes. *Ann. Math. Stat.* **1952**, *23*, 193–212. [[CrossRef](#)]
52. D’Agostino, R.B.; Pearson, E.S. Tests for Departure from Normality. Empirical Results for the Distributions of b_2 and $\sqrt{b_1}$. *Biometrika* **1973**, *60*, 613–622. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.