The Value of Perfect Forecasting in Optimizing the Management of Energy Communities

Patrizia Beraldi[a], Luigi Gallo[b] and Alessandra Rende

Department of Mechanical, Energy and Management Engineering, University of Calabria, Italy

Keywords: Machine Learning Techniques, Solar Production, Forecasting, Renewable Energy Communities.

Abstract: The rise of Renewable Energy Communities (REC) is transforming energy systems by promoting decentralized renewable energy production, but their operational efficiency is hindered by the inherent uncertainty of production sources like photovoltaic systems. Accurate day-ahead forecasting is pivotal for optimizing REC energy management strategies, balancing production, consumption, and grid reliance. This study evaluates five machine learning (ML) models—Support Vector Regression, Random Forest, Adaptive Boosting, Gradient Boost Regression Tree, and Stacking Generalization—against standard accuracy metrics and introduces the Value of Perfect Forecast, a novel metric that quantifies the economic impact of forecast inaccuracies on REC optimization. Results indicate that, while some models perform better in standard accuracy metrics, others are more effective in reducing the economic impact of forecast errors, emphasizing the necessity of aligning forecasting approaches with optimization goals to achieve meaningful operational improvements.

1 INTRODUCTION

In recent years, Energy Communities (ECs) have gained increasing attention as a reference model for driving the energy transition. Defined as cooperative or collective groups of local stakeholders, ECs typically consist of individual renewable energy producers, such as homeowners with photovoltaic (PV) panels, as well as small and medium-sized enterprises and public institutions. If an EC's energy production comes primarily from green sources, it is classified as a Renewable Energy Community (REC).

The primary function of a REC is to facilitate the sharing of locally produced renewable energy among members of the community. This replaces the concept of self-consumption, typically considered at the individual level, with a broader concept of 'virtual' self-consumption. Ideally, energy needs should be met within the community, thereby reducing dependence on the grid, with consequent economic benefits.

In such a setting, members can achieve cost savings, by, for example, receiving surplus energy produced by a neighbour at rates significantly lower than current electricity tariffs. On the other hand, a coalition member can be incentivized to share the produced energy, receiving compensation above to net metering or feed-in tariffs. Beyond the economic benefits, social and environmental motivations are also driving the adoption and expansion of RECs. Indeed, they contribute to reducing carbon emissions in line with global climate targets and environmental stewardship. In addition, RECs promote social cohesion by strengthening local networks, making communities less vulnerable to energy price fluctuations and external supply disruptions. From the REC's perspective, achieving a high level of self-sufficiency by optimally matching production and consumption is a challenging task due to the uncertain and intermittent nature of renewable generation. In this evolving context, accurate forecasting becomes increasingly important to support the effective operation of the REC's resources. Without accurate forecasts, energy system optimization can suffer, leading to operational inefficiencies and increased reliance on the external grid. These inefficiencies can limit the benefits to coalition members and also reduce the attractiveness of REC membership.

In this paper, we focus on the optimal operation of a REC with the aim of investigating the value of the perfect forecast. The approach relies on the idea of combining predictive and prescriptive methodologies to define a robust tool to support decision-makers. The predictive analysis plays the key role of reducing

[a] https://orcid.org/0000-0002-1672-4033
[b] https://orcid.org/0009-0002-0553-4362

Beraldi, P., Gallo, L. and Rende, A.
The Value of Perfect Forecasting in Optimizing the Management of Energy Communities.
DOI: 10.5220/0013343700003893

In *Proceedings of the 14th International Conference on Operations Research and Enterprise Systems (ICORES 2025)*, pages 177-185
ISBN: 978-989-758-732-0; ISSN: 2184-4372

the uncertainty intrinsic in RES generation, whereas the prescriptive methodology benefits from accurate forecasts to provide more reliable solutions.

Most studies assess predictive techniques mainly using standard performance metrics. However, these evaluations often overlook how well these techniques contribute to the effectiveness of solving real-world optimization problems. To address this gap, we introduce a new metric, called the Value of Perfect Forecast (VPF) which quantifies the cost associated with uncertainty in the decision-making process. The VPF parallels the well-known Expected Value of Perfect Information, used in stochastic programming (Ruszczyński and Shapiro, 2003). The VPF measures the additional costs incurred when operational plans, determined by considering imperfect forecasts, must be adjusted to account for actual data. By comparing these costs to those derived from an optimal plan based on perfect information, the VPF provides a clear representation of the value decision-makers might assign to eliminating forecast uncertainty, thereby offering a more comprehensive evaluation of forecasting techniques in practical optimization contexts.

The idea of measuring the value of the forecasting techniques, not only in terms of accuracy, has only been partially explored in the scientific literature. (Peterssen et al., 2024) studied the impact of forecasting on the optimization of energy systems. Specifically, they compare the solutions of a linear programming problem, where the input parameters are assumed to be perfectly known, with a priority list, i.e. a heuristic strategy that does not use any forecasts at all. The results show that when no forecast is taken into account, there is a cost increase of 28%, whereas the use of limited forecasting can reduce this limit to 22%. Similarly, (Putz et al., 2023) assess the impact of forecast accuracy on local ECs. Their study finds that relying solely on conventional quality metrics for selecting a forecast approach fails to capture its true value in supporting EC operations. In their case study, where forecasts relate to electricity and thermal loads, the authors show that significant improvements in forecast quality yield only marginal gains in terms of KPIs for the EC.

Our study aims at further investigating this issue, highlighting the importance of contextualizing the forecast methods in the interplay between predictive and prescriptive methods. Specifically, we focus on predicting energy production from PV panels owned by REC members. To this end, we implement five machine learning (ML) techniques: Support Vector Regression (SVR), Random Forest (RF), Adaptive Boosting (ADA), Gradient Boost Regression Tree (GBRT), and Stacking Generalization (STK). In addition to evaluating ML techniques against traditional KPIs, we measure their practical impact by incorporating their predictions into a prescriptive optimization framework aimed at defining the best daily operating plan. The rest of paper is organized as follows: Section 2 outlines the forecasting techniques for daily PV production. Section 3 details the optimization model. Section 4 introduces the KPIs used in the evaluation process. Section 5 describes the case study, and Section 6 presents and analyzes the results. Conclusions and future reseach directions are drawn in Section 7.

2 FORECASTING METHODS

The scientific interest in the design of more and more accurate forecasting algorithms for the electricity production from renewable energy sources is evident by the very large number of publications on this subject. We refer interested readers to recent contributions that survey the main algorithms. Among the others, we cite (Sharadga et al., 2020), (Yao et al., 2019), (Ma et al., 2022), (Kodaira et al., 2021).
Here, we focus on the forecasting approaches for the day-ahead electricity production from PV panels and we briefly describe the main methods implemented in our study.

2.1 Support Vector Regression

The first technique analyzed in our study is the SVR. Let us consider a training dataset, represented by the set of pairs $\{(x_1, y_1), \dots, (x_N, y_N)\}$, where $x_i \in \mathbb{R}^n$ is a vector of n input features, and $y_i \in \mathbb{R}$ denotes the corresponding PV power generation values. The original feature space is mapped into a higher-dimensional space through a kernel function. This mapping enables the model to handle nonlinear relationships between the input features and the target variable. In this study, we have used the radial basis function, that transforms the input feature vector x_i into a new feature representation $\varphi(x_i)$, via the transformation map Φ (Ding et al., 2021). The primary goal of the SVR is to estimate a function $f(x)$ that approximates the relation between input variables and target:

$$y = f(x) = \omega\varphi(x_i) + b.$$

Here ω is the vector of weights associated with the input variables, and b is the bias term. These parameters are determined by solving the soft margin optimization problem, defined as follows:

$$min\ \frac{1}{2}\|\omega\|^2 + C\sum_{i=1}^{N}(\xi_i^+ + \xi_i^-) \quad (1)$$

$$s.t.$$

$$y_i - \omega x_i - b \leq \varepsilon + \xi_i^+ \qquad i = 1,\dots,N \quad (2)$$

$$\omega x_i + b - y_i \leq \varepsilon + \xi_i^- \qquad i = 1,\dots,N \quad (3)$$

$$\xi_i^+,\ \xi_i^- \geq 0 \qquad i = 1,\dots,N \quad (4)$$

where ξ_i^+, ξ_i^- are slack variables that allow deviations beyond the ε margin. The objective function balances complexity, controlled by the first term in (1), with penalties (second term) for errors beyond the tolerance ε, regulated by the parameter C (Vapnik, 1998).

2.2 Random Forest

In this study, various ensemble methods are applied to model the relationship between the PV power generation and input features. Ensemble learning combines multiple basic learners to enhance predictive performance compared to individual models. The first method employed is Random Forest with out-of-bag validation (Breiman, 2001). RF is a bagging method, i.e., a technique used in regression tasks that trains multiple base learners on different subsets of the data and combines their outputs to produce a more robust prediction. In this approach, multiple subsets (called "bags") are sampled with replacement from the training dataset. Each bag is used to train an individual predictive model, resulting in a collection of models. The final prediction of PV power generation is obtained by averaging the predictions from all models. This aggregation reduces overfitting and improves generalization. In our experiments, the CART algorithm (Breiman et al., 1984) has been used to train the multiple regression trees (RTs) that serve as base learners.

2.3 Adaptive Boosting

The other method employed in this study is the Adaptive Boosting (ADA). Unlike bagging methods, which train models independently, boosting focuses on sequentially training models such that each subsequent model learns from the errors of its predecessor. In our application, ADA trains multiple RTs sequentially. The training process begins by associating an equal weight $\omega_i = \frac{1}{N}$ to N data points that compose the training set. These weights are used to compute the probability for sampling data points $p_i = \frac{\omega_i}{\sum_{i=1}^{N}\omega_i}$. Through this probability distribution, the training set is sampled with replacement. The first RT, RT_0, is trained on the sampled data set and generates the prediction $\hat{y}_i^0 = RT_0(x_i)$. A linear Loss function is computed for each data point i:

$$L_i = \frac{\hat{y}_i^0 - y_i}{D},$$

where $D = sup \mid \hat{y}_i^0 - y_i \mid$. Based on the loss values, ADA updates the weights of the data points, increasing the weights of those with higher errors. This adjustment ensures that subsequent learners focus more on data points that the previous learner struggled with. The procedure for weight updates and assigning importance to each regression tree is detailed in (Drucker, 1997). The final hourly PV generation forecast is computed as the weighted median of the predictions from all the RTs in the ensemble.

2.4 Gradient Boosting

Another technique used in our study is GBRT, where the RTs are iteratively trained to minimize the prediction error (Friedman, 2001). The algorithm begins with an initial RT, RT_0 which approximates the relationship between the target variable y and the input feature vector x. The initial prediction for each data point i are $\hat{y}_{i0} = RT_0(x_i)$. At step $t = 1$, the next RT_1 is trained to learn the residuals r_1 defined as the gradient of the loss function

$$\frac{\delta L(y, \hat{y}_0)}{\delta \hat{y}_0},$$

with respect to the previous prediction. The predictions are then updated by adding the contribution of RT_1, scaled by a learning rate η:

$$\hat{y}_1 = \hat{y}_{i0} + \eta RT_0(x_i).$$

This process continues iteratively, At each step t, the tree RT_t is trained to minimize the residual errors of the current predictions, refining the model incrementally. The final prediction y_T is given by

$$\hat{y}_T = RT_0(x_i) + \eta \sum_{t=0}^{T-1} RT_t(x_i).$$

In our study, the Mean Squared Error has been used as loss function L, minimizing the average squared differences between predicted and true values. The learning rate η is an hyper-parameter tuned during the training process with ten-fold cross validation.

2.5 Stacked Generalization

As an extension to the previously discussed methods, this work also incorporates a stacked generalization (STK) (Wolpert, 1992). Stacking combines the strengths of multiple base learners by using their predictions as inputs to a higher-level model, called the meta-learner or blender. The aim is to build a robust forecasting method by adjusting the results of the submodels and minimize the final prediction errors. Let us denote by $\mathcal{J}$ the set of tested ML models and let $\hat{y}_j$ be the corresponding forecast. Then, the final prediction is defined as:

$$\hat{y} = \sum_{j \in \mathcal{J}} \beta_j \hat{y}_j.$$

To determine the coefficient β_j a linear regression model has been used, by minimizing the error between the stacked ensemble's predictions ($\hat{y}_j$) and the actual observed values. By fitting β_j, the meta-learner assigns greater weights to models with better predictive performance.

3 OPTIMIZING THE REC MANAGEMENT

With the aim of evaluating the impact of perfect forecast on energy systems, we consider the problem of defining the optimal daily operational plan of a REC, where some members own generation units (prosumers), whereas others are simple consumers. The centralized management of the REC entails the definition of a shared strategy, where the energy requests and production profiles of all REC's members are considered collectively. Compared to an individual approach, where each REC's member optimizes his own resources independently, the unified one accounts for intra-community energy exchanges, thus increasing the self-sufficiency rate with a consequent maximization of the gain for trade. Electricity shortage or excess production (not used within the REC), are compensated by transactions with the power grid. We assume that the REC is managed by an aggregator that represents the interface with the power market and has to guarantee the demand satisfaction for all the community's members. The final aim is to define the most convenient operational plan. In our model, we assume that the daily demand profiles of the REC's members are known in advance, whereas the uncertain electricity production from PVs is replaced by its forecast. Since the accuracy of the supply profile directly influences the entire operational plan, deviations from the actual electricity production could significantly impact the overall costs, i.e., the aggregator might have to resort to the real-time market where rates are typically less convenient compared to the tariffs of the day-ahead electricity market.

To formally define the problem, we introduce the sets $\mathcal{N}$ and $\mathcal{T}$ associated with the REC members and the operational time horizon, respectively. In the experiments, we have considered a daily horizon, with hourly time steps. Forecasts are generated one day in advance and serve as input data in the optimization phase. The process is repeated according to a rolling horizon scheme, where each day updated forecasts and new input data are used in the optimization phase. Prosumers within the REC are supposed to be equipped with battery energy storage (BES) devices. For each $n \in \mathcal{N}$, we denote by C_n the battery size and by D_{nt} the electricity demand at time t. As for the energy production, we denote by G_{nt} the energy produced by the member n at time t. Energy produced and not directly used can be stored in the BES, if available, and used later or can, eventually, be exported to the REC. For each member $n \in \mathcal{N}$ and time $t \in \mathcal{T}$, the following decision variables are introduced in the formulation:

- SoC_{nt} state of charge of the BES;
- E^c_{nt}, E^d_{nt} amount charged in and discharged from BES, respectively;
- EC^{in}_{nt}, EC^{out}_{nt} energy amount imported from and exported to the REC, respectively;
- EG^{in}_{nt}, EG^{out}_{nt} energy amount imported from and exported to the power grid, respectively.

The aim is to define the operational plan that minimizes the total costs:

$$\min f = \sum_{t \in \mathcal{T}} \sum_{n \in \mathcal{N}} (P^G_t EG^{in}_{nt} + P^C_t EC^{in}_{nt} - R^G_t EG^{out}_{nt} - R^C_t EC^{out}_{nt}). \tag{5}$$

Here P^G_t and P^C_t denote the electricity tariffs to purchase electricity from the power market and the community, respectively, whereas R^G_t and R^C_t are the revenues obtained when selling electricity to the main grid and to the community. We assume that such data are known in advance and that $P^C_t < P^G_t$ and $R^C_t > R^G_t$ to encourage the sharing of electricity within the REC.

Below we report the main constraints introduced in our formulation. The first set of constraints (6) ensures the satisfaction of the load demand for each member n of the REC and for each time period t. Specifically, the demand can be satisfied by using energy imported from the grid and/or from the REC and, eventually, by using stored energy and the self-production. The excess can be charged to the BES and/or sold to the grid or the community.

$$EC^{in}_{nt} + EG_{nt}{}^{in} + E^d_{nt} + G_{nt} = D_{nt} + EC^{out}_{nt} + EG^{out}_{nt} + E^c_{nt} \quad \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \tag{6}$$

By (7), energy balance within the REC is guaranteed:

$$\sum_{n \in \mathcal{N}} (EC^{in}_{nt} - EC^{out}_{nt}) = 0 \qquad \forall t \in \mathcal{T} \tag{7}$$

Constraints (8) limit the total energy that can be imported to a given upper bound limit E^{max}, computed on the basis of the operation power:

$$EC^{in}_{nt} + EG^{in}_{nt} \leq E^{max} \qquad \forall n \in \mathcal{N}, \forall t \in \mathcal{T} \tag{8}$$

Constraints (9)-(12) refer to the management of the storage device:

$$SoC_{nt} = \alpha\, SoC_{nt-1} + \eta^c\, E^c_{nt} - \frac{1}{\eta^d}\, E^d_{nt} \quad \forall n \in \mathcal{N}, \forall t \in \mathcal{T} \tag{9}$$

Specifically, constraints (9) model the dynamics of the BES, linking the state of charge of a time t, to the amount in the battery at the end of the previous time and to the charged and the discharged amounts. Here α, η^c and η^d accounts for loss. We note that for the first period of the time horizon, the amount initially in the device is set equal to the energy in the battery at the last period of the previous day. Constraints (10)-(12) limit the stored and the charged and discharged amount to a percentage β of the BES capacity.

$$SoC_{nt} \leq C_n \qquad \forall n \in \mathcal{N}, \forall t \in \mathcal{T} \tag{10}$$

$$E^d_{nt} \leq \beta\, C_n \qquad \forall n \in \mathcal{N}, \forall t \in \mathcal{T} \tag{11}$$

$$E^c_{nt} \leq \beta\, C_n \qquad \forall n \in \mathcal{N}, \forall t \in \mathcal{T} \tag{12}$$

Finally, the last constraints refer to the nature of the decision variables, for every time period t and member n:

$$EC^{in}_{nt},\ EC^{out}_{nt},\ EG^{in}_{nt}, EG^{out}_{nt},\ SOC_{nt}, E^d_{nt}, E^c_{nt} \geq 0. \tag{13}$$

4 ASSESSING THE FORECAST QUALITY

We assess the quality of the forecasting methods using both traditional metrics and the new index. While the former are inherent in any forecast, regardless of the context in which they are used, the latter measure accounts for the value of prediction and evaluates the benefit of incorporating the forecast into the decision-making process.

Accuracy of the forecasting methods has been measured by traditional KPIs. In particular, in the experimental phase we have used the following traditional metrics, Root Mean Squared Errors (*RMSE*), Mean Absolute Error (*MAE*), and R-Square coefficient(R^2), defined as follows:

$$RMSE = \sqrt{\sum_{i=1}^{I} \frac{(y_i - \hat{y}_i)^2}{I}},$$

$$MAE = \sum_{i=1}^{I} \frac{|\, y_i - \hat{y}_i \,|}{I},$$

$$R^2 = 1 - \frac{\sum_{i=1}^{I} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{I} (y_i - \bar{y}_i)^2}.$$

Here index i represents the generic data point of the test-set, while y_i and $\hat{y}_i$ denote the measured and forecast values, respectively. Both RMSE and MAE measure the accuracy of the forecasting model on average: RMSE is the squared root of squared errors mean, while MAE averages absolute values of the errors. R^2 represents the proportion of variance in the dependent variable that is predictable from the independent variables. Values range from 0 to 1, with 1 indicating perfect prediction and 0 suggesting no predictive power. Each metric has its strengths: RMSE is sensitive to large errors, MAE is more robust to outliers, and R^2 indicates overall model fit.

In addition to the traditional KPIs, the value of the forecasting methods has been evaluated by the new measure, VPF. Let $\mathcal{K}$ denote the set of forecasting techniques under evaluation. Each method $k \in \mathcal{K}$ generates a supply profile for each prosumer within the REC, which then serves as input data in the optimization model. By solving the optimization problem $|\mathcal{K}|$ times− once for each forecasting method and associated production values G^k_{nt}— we evaluate the influence of each forecast on decision outcomes. Let $\mathbf{x}^k$ represent the vector of decision variables when the optimization model is run with forecast k. We then calculate the objective function value when the solution $\mathbf{x}^k$ is applied, taking into account the actual supply patterns. Forecast errors may require adjustments to the initial operational plan, which can lead to higher costs due to necessary interactions with the balancing market. If for some time periods the real PV generation is lower than the forecast one, the aggregator may be required to purchase electricity from the market at higher prices. On the contrary, if the forecast overestimates the real profile, the extra amount should be fed back to the grid. Therefore, implementing the solution $\mathbf{x}^k$ under real conditions may result in added imbalance costs. We denote the total cost after these adjustments by $z(\mathbf{x}^k)$.

To benchmark the results, we consider a scenario of perfect forecast, which yields an objective function value denoted by z^*. Thus, for each k, the VPF is defined as:

$$VPF^k = z(\mathbf{x}^k) - z^*.$$

This metric enables a direct comparison between each forecast approach and the perfect forecast scenario, highlighting the ”cost” of forecast inaccuracy in terms of suboptimal decision-making within the REC. A lower VPF indicates that the forecast method produces results closer to the ideal solution under perfect information, suggesting a higher practical utility in the decision-making process. The experiments presented in Section 6 show how the new index may complement the standard KPIs.

5 TEST CASE AND DATA SETTING

To evaluate the impact of forecasting on the operation of RECs, we have considered a simple, yet meaningful, case study related to a small community composed of three types of members. Two of them are prosumers, equipped with PV panels and BESs, whereas the other is a simple consumer. Table 1 reports the energy assets of the REC members, in terms of number of PV panels and BES capacity, along with the maximum operationing limits.

Table 1: Users' energy assets.

Member	Number panels	Storage Capacity (kWh)	Grid Operation limit (kW)
1	16	6	4.5
2	32	12	6
3	-	-	3

Demand and tariffs have been assumed to be known. The former has been derived using the data from (ARERA, 2024) and (Giordano et al., 2020), whereas the electricity tariffs have been derived using the Italian Single National Price as reference (Gestore Mercati Energetici, 2023). The model has been coded in Python and solved by the commercial solver (Gurobi Optimization, LLC, 2024).

As for the forecasting, experiments have been carried out by considering a three-year data set, including 23832 observations, with a resolution of 1-hour. The data have been provided by the University of Calabria, Italy, and the following features have been considered: Hour of the day (ranging from 1 to 24), Relative Humidity (%), Temperature (°C), and Wind Speed (km/h). The target value is the generation of a module installed on the rooftop of one of the university buildings[3]. Before running the tests, a cleaning phase has been performed. Data points, for which not all the features were available, have been removed (735 data points). Night observations have been also excluded from the dataset for preventing bias. At the end of the cleaning phase, the final number of data records was reduced to 13462.

To improve the performance of the ML models, a feature engineering process has been implemented. This step focuses on deriving new, informative input variables by transforming and processing the original ones. This step starts by examining the correlation between the input variables and the target variable using a correlation map, shown in Figure 1. This visualization helps to identify the most relevant features and potential redundancies or collinearities in the dataset.

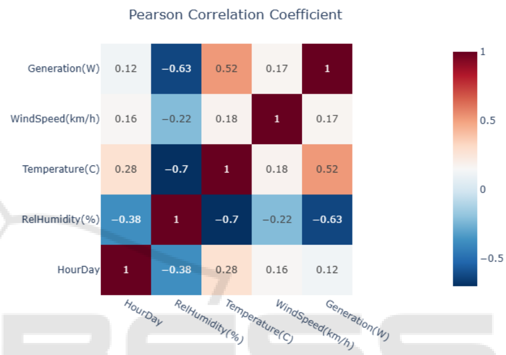

Figure 1: Correlation matrix of the original features.

Analyzing the data, it appears that the feature *HourDay* is one of the less significant variables for predicting power generation. However, a closer look reveals that the relationship between power generation and the hour of the day is similar to a quadratic function. To account for this, the squared values of the hour of the day have been introduced as an additional input feature, improving the model's ability to capture this non-linear relationship.

Furthermore, as highlighted in the results of (Nicoletti and Bevilacqua, 2024), the relationship between power generation and the meteorological variables— temperature and relative humidity— can be more accurately modeled by incorporating temporal variations. Specifically, the current values of these variables, as well as their values one hour before and one hour after, are included as input features. These additional variables allow the model to better capture the predictive power of short-term fluctuations, which are crucial for accurate forecasting. Finally, the most

[3]The PV module, employing polycrystalline silicon technology is characterized by latitude: 39°21'N and Longitude: 16°13'E. The module's surface presents an area of 1663 mm x 998 mm, and a tilt angle of 30°. It faces south-east, with a nominal power of 245 W.

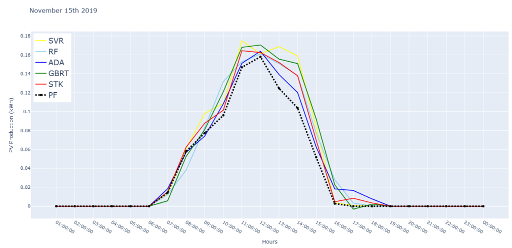

Figure 2: Power production for an autumn day: predictions versus perfect forecast.

recent historical values of power generation are included as a critical input feature. At the time of forecasting, the most recent available data for each hour comes from 48 hours prior. Incorporating this historical information ensures that the model exploits temporal continuity and accounts for persistent trends or patterns in power generation.

Data processing, including feature engineering and the application of ML has been implemented in Python. Data have been split in the training and the test sets with a ratio of 15%.

6 NUMERICAL RESULTS

This section is devoted to the presentation and discussion of the numerical results. Figure 2 shows the forecast power production obtained considering the different ML techniques for a given day in November. In the same Figure, the perfect forecast, representing the actual power production, is also included. As shown, all the techniques provide predictions that capture the general trend of the real production values; however, differences in accuracy and behavior are apparent among the methods. Some techniques closely follow the actual production, especially during peak production periods, while others exhibit more significant, even though limited, deviations. Quantitative performance metrics further highlight the accuracy of each technique. As summarized in Table 2, RF method consistently outperforms the other methods, achieving the lowest values for RMSE and MAE, as well as the highest of R^2. The STK method demonstrates comparable performance to the RF, suggesting that it effectively combines the strengths of multiple models, although it does not significantly outperform the best individual technique (RF in this case). On the other hand, the GBRT method performs the worst among the tested techniques underlining the limitations in its ability to capture the underlying patterns.

Forecast data have been used as input for the optimization problem presented in Section 3. The problem is solved iteratively, using each time the forecast

Table 2: Accuracy evaluation via traditional KPIs.

Methods/KPIs	RMSE	MAE	R^2
SVR	23.16	13.20	0.89
RF	20.06	12.15	0.91
ADA	22.76	15.70	0.89
GBRT	23.44	14.95	0.88
STK	21.91	13.29	0.90

obtained with a different technique. The resulting daily strategies, represented in terms of energy transaction with the main grid, have been used to determine the VPF. Imbalances, evaluated with the respect to the strategy suggested when considering the perfect forecast, are corrected by recurring to the balancing market. To illustrate this process, Figure 3 depicts how imbalances are managed when forecasts from the RF technique are used. The figure highlights the initial transaction strategy based on RF predictions and the subsequent adjustments required when real values are observed.

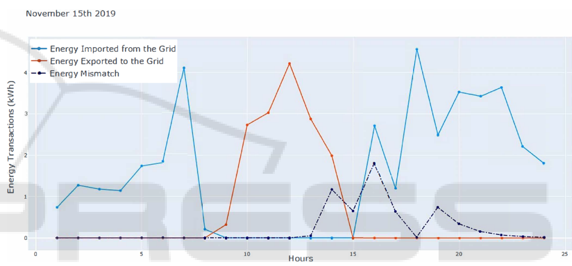

Figure 3: REC transactions with the grid.

To gain further insights, the analysis was extended by running the model for representative days across different seasons. This seasonal analysis revealed the impact of weather and environmental variability on forecasting accuracy. Table 3 reports the values of the VPF measure (in percentage) for the different techniques and the different days.

Table 3: VPF (%) for the different ML techniques for different days.

	Winter	Spring	Summer	Fall	Mean
SVR	0.52	1.22	4.74	7.97	3.61
RF	0.16	1.12	2.43	4.78	2.12
ADA	0.00	0.98	2.83	2.25	1.51
GBRT	0.76	1.09	3.30	7.60	3.19
STK	0.27	1.08	3.49	4.51	2.34

Looking at the results some considerations can be drawn. First of all, we may observe that for all the techniques, ”Fall” presents the higher VPF values, indicating higher variability potentially related to instability in weather conditions. This may be due to unpredictable weather patterns, such as rapid temperature changes or varying cloud cover, which increase forecasting difficulty. On the contrary, ”Winter” shows the best values reflecting very likely more stable conditions easier to predict. When comparing the techniques for a given reference day, we may note that ADA provides the best results followed by the RF and the STK. Thus, while ADA is not competitive when evaluated via traditional KPIs, when considering the VPF seems to provide the best results.

The results presented further emphasize an important observation: the forecasting technique that performs best under conventional metrics may not always be the most effective for prescriptive purposes, especially in contexts where accurate adjustments and decision-making are critical. This distinction highlights the need to tailor performance evaluations to the specific application or domain requirements.

7 CONCLUSIONS

This paper focuses on the evaluation of forecasting techniques from a prescriptive perspective. Specifically, the study applies five ML techniques to train predictive models for PV generation forecasting, which are then used as input parameters for an optimization problem aimed at defining the optimal daily operational strategy for a REC. The different techniques are evaluated by using both standard accuracy metrics and a new measure, the VPF, used to measure the cost incurred to adjust operational plan for compensate for the deviations between forecast and actual values. The preliminary results seem to point out that the ML model with the highest score on the standard statistical metrics is not necessarily the most effective for optimization purposes. This study underscores the need for a more holistic approach to evaluating forecasting models, especially when they are integrated into optimization workflows. By considering both standard metrics and application-specific indices like VPF, stakeholders can select models that are not only accurate, but also cost-effective for operational decision-making.

In this sense, the present work constitutes a preliminary investigation within this field of application. Firstly, it is recommended that the insights provided by the new index introduced in this study should be further replicated through a range of different computational experiments and application settings. This will demonstrate the practical utility and generalisability of the proposed index in different contexts. Future research directions could focus on achieving a deeper integration between predictive and prescriptive processes. One promising approach may be embedding the prediction process within optimization models by leveraging innovative approaches like constraint learning. In alternative the training process of predictive models may be carried out, taking in account the structure of the optimization problem, like in ’Smart Predict, Then Optimize’ framework.

ACKNOWLEDGMENTS

We acknowledge the financial support from: PNRR MUR project PE0000013-FAIR.

REFERENCES

ARERA (2024). Analisi dei consumi domestici, prelievo medio orario per provincia, 2021- cosenza. https://www.arera.it/dati-e-statistiche/dettaglio/analisi-dei-consumi-dei-clienti-domestici [Accessed: (31/10/2024)].

Breiman, L. (2001). Random forests. *Machine Learning*, 45.

Breiman, L., Friedman, J., Stone, C., and Olshen, R. (1984). *Classification and Regression Trees*. Taylor & Francis.

Ding, X., Liu, J., Yang, F., and Cao, J. (2021). Random radial basis function kernel-based support vector machine. *Journal of the Franklin Institute*, 358(18):10121–10140.

Drucker, H. (1997). Improving regressors using boosting techniques. In *Icml*, volume 97, page e115. Citeseer.

Friedman, J. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29:1189–1232.

Gestore Mercati Energetici (2023). National hourly electricity selling prices. https://gme.mercatoelettrico.org/it-it/Home/Esiti/Gas/IGIndex/Statistiche/Sintesi [Accessed: (19/10/2024)].

Giordano, A., Mastroianni, C., Scarcello, L., and Spezzano, G. (2020). An optimization model for efficient energy exchange in energy communities. In *2020 Fifth International Conference on Fog and Mobile Edge Computing (FMEC)*, pages 319–324.

Gurobi Optimization, LLC (2024). Gurobi Optimizer Reference Manual.

Kodaira, D., Tsukazaki, K., Kure, T., and Kondoh, J. (2021). Improving forecast reliability for geographically distributed photovoltaic generations. *Energies*, 14(21).

Ma, M., He, B., Shen, R., Wang, Y., and Wang, N. (2022). An adaptive interval power forecasting method for photovoltaic plant and its optimization. *Sustainable Energy Technologies and Assessments*, 52:102360.

Nicoletti, F. and Bevilacqua, P. (2024). Hourly photovoltaic production prediction using numerical weather data and neural networks for solar energy decision support. *Energies*, 17(2).

Peterssen, F., Schlemminger, M., Lohr, C., Niepelt, R., Hanke-Rauschenbach, R., and Brendel, R. (2024). Impact of forecasting on energy system optimization. *Advances in Applied Energy*, 15:100181.

Putz, D., Gumhalter, M., and Auer, H. (2023). The true value of a forecast: Assessing the impact of accuracy on local energy communities. *Sustainable Energy, Grids and Networks*, 33:100983.

Ruszczyński, A. and Shapiro, A. (2003). *Stochastic Programming, Handbook in Operations Research and Management Science*. Elsevier Science, Amsterdam.

Sharadga, H., Hajimirza, S., and Balog, R. S. (2020). Time series forecasting of solar power generation for large-scale photovoltaic plants. *Renewable Energy*, 150:797–807.

Vapnik, V. (1998). *The Support Vector Method of Function Estimation*, pages 55–85. Springer US, Boston, MA.

Wolpert, D. H. (1992). Stacked generalization. *Neural Networks*, 5(2):241–259.

Yao, S., Pan, L., Yu, Z., Kang, Q., and Zhou, M. (2019). Hierarchically non-continuous regression prediction for short-term photovoltaic power output. In *Proceedings of the 2019 IEEE 16th International Conference on Networking, Sensing and Control, ICNSC 2019*, pages 379–384.