



HAL
open science

Prescriptive Trees for Integrated Forecasting and Optimization Applied in Trading of Renewable Energy

Akylas Stratigakos, Simon Camal, Andrea Michiorri, Georges Kariniotakis

► To cite this version:

Akylas Stratigakos, Simon Camal, Andrea Michiorri, Georges Kariniotakis. Prescriptive Trees for Integrated Forecasting and Optimization Applied in Trading of Renewable Energy. *IEEE Transactions on Power Systems*, 2022, 10.1109/TPWRS.2022.3152667 . hal-03330017v3

HAL Id: hal-03330017

<https://hal.science/hal-03330017v3>

Submitted on 11 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Prescriptive Trees for Integrated Forecasting and Optimization Applied in Trading of Renewable Energy

Akylas Stratigakos, *Student Member, IEEE*, Simon Camal, Andrea Michiorri,
and Georges Kariniotakis, *Senior Member, IEEE*

Abstract—Deriving decisions from data typically involves a sequential process with two components, forecasting and optimization. Forecasting models learn by minimizing a loss function that stands as a proxy for task-specific costs (e.g., trading, scheduling) without considering the downstream optimization, which in practice creates a performance bottleneck and obscures the impact of data on decisions. This work proposes a single data-driven module that leverages the structure of the optimization component and directly learns a policy conditioned on explanatory data. For this purpose, we describe an algorithm to train ensembles of decision trees by directly minimizing task-specific costs, and prescribe decisions via a weighted Sample Average Approximation of the original problem. We then develop a generic framework to assess the impact of explanatory data on prescriptive performance. To illustrate the efficacy of the proposed modeling approach, we consider two case studies related to trading renewable energy. First, we examine trading in a day-ahead market and propose strategies that balance optimal trading decisions and predictive accuracy. Next, we append a storage device and co-optimize the day-ahead offers and the operational policy, based on a tractable approximation using the linear decision rule approach. The empirical results demonstrate improved prescriptive performance compared to solutions derived under the standard stochastic optimization framework. Further, we provide valuable insights on how explanatory data impact optimization performance and how this impact evolves under different market designs.

Index Terms—Data-driven optimization, decision trees, energy forecasting, energy storage system, energy trading, linear decision rules, prescriptive analytics.

I. INTRODUCTION

A. Background and Motivation

ALMOST every aspect of power systems, including real-time operation, investment planning, and electricity markets, is subject to uncertainty stemming from multiple sources, such as stochastic renewable production and unknown market conditions. Facilitated by the influx of data associated with the transition towards the smart grid era, recent years have seen a burgeoning development of energy analytics tools that enable stakeholders to derive decisions in the presence of contextual information. Typically, data-driven decision-making comprises a sequential process with two components, (energy) forecasting and optimization. The first involves forecasting

uncertain quantities, e.g., renewable production or market prices, at a future time interval and in a form that ranges from deterministic (point) forecasts to probabilistic forecasts, scenarios (trajectories) and uncertainty sets. Next, these forecasts are used as inputs in an optimization problem to derive optimal decisions (prescriptions).

In the context of statistical learning, forecasting models are trained by minimizing a surrogate loss function that serves as a proxy for a task-specific cost. However, several issues arise with this approach, as moving from a prediction to a prescription might not be straightforward. Assessing the impact of forecasts on decision costs, i.e., *forecast value*, is considered to be one of the key challenges in energy forecasting in the coming years [1]. Further, directly optimizing towards forecast value rather than accuracy is identified as a high-leverage objective to employ machine learning as means of tackling climate change [2]. Studies on the economic impact of price forecasting errors [3] confirm that increased accuracy does not always translate into increased value, as the latter heavily depends on the specific task. A recently observed trend suggests moving beyond the simple statistical evaluation of prediction errors to assessing the quality of decisions obtained for different applications. For example, [4] proposes a multivariate probabilistic forecasting model and considers the economic benefits for an electricity retailer as a means of assessing its forecast value. This trend highlights two pertinent issues that serve as our current motivation. First, it is pivotal for the forecasting model to exploit the structure of the downstream optimization problem in order to maximize its value. Second, deploying multiple analytic tools in sequence increases the model chain complexity and obfuscates the impact of data on the efficacy of decisions.

This work proposes an alternative paradigm for decision-making in the presence of contextual information in power system applications, specifically in short-term trading of renewable energy. By integrating forecasting and optimization our goal is to i) improve the out-of-sample prescriptive performance in trading applications, ii) reduce the effort to model uncertainty and simplify the data-decisions pipeline, iii) quantify the impact of contextual information to optimization efficacy and enhance model explainability. To this end, we leverage recent advances in the fields of machine learning and operations research to propose a data-driven method for policy learning, that allows decisions to vary as a function of contextual information.

Parts of this research were carried in the frame of the European project RE-gions (No. 646039 and No. 775970), supported by ADEME's 'Investissement d'Avenir' program and the ERA-Net SES RegSys project, and the Smart4RES project (No. 864337), supported by the Horizon 2020 Framework Program.

B. Literature Review

Classical stochastic optimization assumes that uncertain parameters have known distributions. In practice, however, only observational data are available. The problem of single-stage *conditional* stochastic optimization in the presence of contextual information (also called explanatory data or features) has received a lot of attention recently. Three general approaches are identified: i) modifying the optimization component, ii) learning under an alternative loss function, iii) directly forecasting the decisions. For the first approach, [5] establishes a prescriptive analytics framework for data-driven decisions via a weighted Sample Average Approximation (SAA) [6] of the original problem, employing predictive machine learning algorithms, which sets the base of our approach. Several extensions exist, including robust prescriptive analytics [7], dealing with multi-stage problems [8], and considering multi-stage problems with adjustable robust optimization [9]. In [10] the residuals induced by the SAA solution are used to infer decision uncertainty. The second stream of literature considers training forecasting models under alternative loss functions in order to explicitly minimize downstream costs [11]. However, training a machine learning algorithm becomes challenging, as the alternative loss function might be non-convex and discontinuous. Gradient-based methods usually assume a smooth objective, as in an earlier work [12] that employs a specialized financial criterion as a loss function. Similarly, [13] describes an end-to-end learning approach for probabilistic forecasting, with applications for energy storage arbitrage and grid scheduling. Decision tree learning, on the other hand, can be used to directly minimize task-specific costs [14]. An alternative approach based on bilevel programming is presented in [15], where the lower problem computes the best decision given a forecast and the upper problem estimates the linear coefficients of a forecasting model that lead to minimum costs. In this framework, however, it remains unclear how to deal with multiple sources of uncertainty. Lastly, the third approach suggests applying empirical risk minimization (ERM) to directly forecast the decisions, a form of policy selection. This is appropriate when the cost function can be employed as a loss function in a learning algorithm. For example, [16] describes data-driven solutions to the newsvendor problem, which effectively result in a high-dimensional quantile regression. A significant drawback of this approach, however, is that it fails to ensure feasible solutions for out-of-sample observations.

Recent works also consider integrating forecasting and optimization and directly learning a policy from data to improve prescriptive performance. We highlight [17] and [18], which combine the framework put forward in [5] with learning under an alternative loss function, as closely related to the present work. Our approach is also based on decision trees, although we employ an alternative method to enhance tractability via a random split criterion. Furthermore, we explicitly quantify feature importance on optimization performance, which is not considered in the above-mentioned works, thus also enhancing explainability of decisions.

Regarding power systems applications and, in particular,

renewable integration in electricity markets, a vast body of work deals with decision-making in the presence of contextual information. A significant part of that work is dedicated to developing accurate forecasting models; we refer the reader to [19] and [20] for a recent treatment on renewable and electricity price forecasting. For our work, the focus is placed on deriving optimal energy offers for wind power plants (WPPs) based on probabilistic energy and price forecasts [21]–[23]. This prototypical problem has found many extensions, such as considering adjustment markets and risk-aversion [24], jointly participating in energy and reserve capacity markets [25], trading under a single-price balancing mechanism [26], and including a storage device [27].

We now highlight recent works in power systems that relate to the present methodology. Regarding learning under alternative loss functions, a task-based load forecasting model that combines deep learning with stochastic economic dispatch is proposed in [28], following the work in [13]. A closed-loop forecast-optimize module is described in [29] to solve the day-ahead unit commitment problem, employing the loss function introduced in [11]. Additionally, [30] considers wind forecasting for short-term trading applications and [31] examines load forecasting for dispatch scheduling, both relying on two-step approaches that involve first inferring a convex loss function, then training the forecasting model. However, this approach does not directly leverage the optimization component. Our previous work [32] suggests employing a risk-averse trading strategy as an alternative loss function to forecast market quantities, prior to deriving trading decisions, which still does not reduce the modeling effort and cannot handle multiple uncertainties. In [33], the framework put forward in [16] is extended by proposing linear decision rules to improve both the forecasting and trading performance of a WPP participating in a DA market. For a similar case study with photovoltaic (PV) plants, [34] describes an ERM formulation based on neural networks. Both of these works deal with variations of the newsvendor problem and proposed solutions do not guarantee feasibility of decisions for more complex problems. One way to circumvent this issue is by considering a discrete set of actions that approximate continuous decisions. For example, [35] examines mode-based control of storage posed as a supervised classification problem.

Table I groups relevant applications by methodology, and provides a qualitative comparison. Our main contributions are summarized as follows:

- We propose and validate an alternative data-driven modeling approach that leverages contextual information to improve prescriptive performance in renewable trading applications, which also reduces the modeling effort, handles multiple sources of uncertainty, and guarantees feasible decisions.
- Methodological contributions include a novel prescriptive tree algorithm, which shows significant reductions in computational costs, and adapting well-known feature importance metrics from the machine learning literature to a prescriptive analytics context, departing from the classical regression setting.
- We illustrate the efficacy of the proposed approach in two

TABLE I: Summary of relevant research in power systems. Superscript * denotes learning under an alternative loss function, † denotes directly forecasting decisions, and †† denotes integrated forecasting-optimization approaches.

Ref.	Feasibility guarantees	Multiple types of uncertainty	Reduced modeling effort
[28], [30], [31], [37]*	✓	-	-
[29]*, ††	✓	-	✓
[33], [34]†	-	✓	✓
This work*, ††	✓	✓	✓

case studies of increasing complexity related to renewable trading. First, we examine trading in a DA market under different pricing mechanisms and propose strategies that balance trading cost and predictive accuracy. Next, we append a storage device and jointly optimize the DA offering strategy and operational control policy; for the latter, we employ the linear decision rule approach [36] to provide a tractable approximation. To our knowledge, embedding the control policy within the learning component presents a novel application of integrated forecasting-optimization.

The rest of the paper is organized as follows. Section II presents the mathematical background and the proposed methodology. Section III formulates the renewable trading problem. Results are presented in Section IV. Finally, we draw conclusions and provide directions for future research in Section V.

II. MATHEMATICAL BACKGROUND AND PROPOSED METHODOLOGY

A. Conditional Stochastic Optimization

We study single-stage stochastic optimization problems with $Y \in \mathcal{Y} \subseteq \mathbb{R}^{d_y}$ being uncertain parameters of interest (e.g., renewable production, market prices) and $X \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$ a set of associated features (e.g., weather conditions), following a joint probability distribution $(X, Y) \sim \mathbb{Q}$. We are interested in approximating the conditional stochastic optimization problem

$$v = \min_{z \in \mathcal{Z}} \mathbb{E}_{\mathbb{Q}}[c(z; Y) | X = \bar{x}] = \min_{z \in \mathcal{Z}} \mathbb{E}_{y \sim \mathbb{Q}_{\bar{x}}}[c(z; Y)], \quad (1)$$

where v is the objective value, $z \in \mathbb{R}^{d_z}$ is the decision vector, \mathcal{Z} is a convex set of feasible solutions, $c(\cdot)$ is a cost function, \bar{x} is a new observation of X , and $\mathbb{Q}_{\bar{x}}$ is the marginal distribution of Y conditioned on \bar{x} . In place of the true distributions, we have access to a training data set $\{(y_i, x_i)\}_{i=1}^n$ of n observations and aim at learning a *policy* \hat{z} that varies as a function of X .

The fundamental method of approximating (1) given a set of observations (empirical or sampled) y_i of Y is via SAA [6]:

$$\hat{z}^{SAA} = \arg \min_{z \in \mathcal{Z}} \sum_{i=1}^n \frac{1}{n} c(z; y_i). \quad (2)$$

While the SAA enjoys several nice theoretical properties, such as consistency and asymptotic optimality, it does not leverage

the available contextual information. The standard modeling approach to do so is to first employ a forecasting model $f : \mathcal{X} \rightarrow \mathcal{Y}$ that maps observations of X to Y , and then solve an optimization problem, which we denote “Forecast, then Optimize” (FO). Forecasting model f is typically selected from a class of machine learning or statistical models. We differentiate between the case of deterministic forecasts

$$\hat{z}^{EV} = \arg \min_{z \in \mathcal{Z}} c(z; \mathbb{E}[Y | X = \bar{x}]), \quad (3)$$

where EV stands for Expected Value, and the generalized case of probabilistic forecasts

$$\hat{z}^{FO} = \arg \min_{z \in \mathcal{Z}} \mathbb{E}_{y \sim \hat{\mathbb{Q}}_{\bar{x}}} [c(z; Y)], \quad (4)$$

where $\hat{\mathbb{Q}}_{\bar{x}}$ denotes the inferred conditional distribution of Y . In turn, (4) is solved with standard stochastic optimization techniques. While in most applications, the FO approach constitutes the current state of the art, note that an optimal solution to (4) is not an optimal policy, as various implicit assumptions are in place (e.g., calibration of predictive density). Therefore, the decision-maker cannot effectively hedge against miscalibrated forecasting models.

The framework described in [5] integrates predictive and prescriptive analytics by forming a weighted SAA of (1) to derive decisions conditioned on contextual information, termed *predictive prescriptions*. These prescriptions, which retain consistency and asymptotic optimality, are defined as:

$$\hat{z}(\bar{x}) = \arg \min_{z \in \mathcal{Z}} \sum_{i=1}^n \omega_{n,i}(\bar{x}) c(z; y_i), \quad (5)$$

where $\omega_{n,i}(\bar{x})$ denotes weights obtained from local learning algorithms, e.g., nearest neighbors and decision trees. In the original work [5], $\omega_{n,i}(\bar{x})$ are derived by training the various algorithms for prediction. In this work, we derive $\omega_{n,i}(\bar{x})$ by directly minimizing decision costs, thus leveraging the structure of the downstream optimization problem and providing a more informed approximation of the decisions. This is further motivated by the fact that in Section III we consider uncertainties from different sources (renewable production and market quantities), which in turn depend on a different set of features. Training a local learning algorithm to predict both of these, would inevitably lead to suboptimal performance. Instead, the proposed approach enables the model to assess the relative impact of each uncertain parameter on the downstream decision costs and weight associated features accordingly during learning, while also exploiting possible cross-dependencies.

Following [18], we formally define the problem of searching over functions $f : \mathcal{X} \rightarrow \mathcal{Y}$ that improve prescriptive performance as

$$\begin{aligned} \min_{f \in \mathcal{F}, z^f(x_i) \in \mathcal{Z}} \quad & \sum_{i \in [n]} c(z^f(x_i); y_i) \\ \text{s.t.} \quad & \end{aligned} \quad (6a)$$

$$z^f(x_i) = \arg \min_{z \in \mathcal{Z}} \sum_{j \in [n]} \omega_{n,j}^f(x_i) c(z; y_j) \quad \forall i \in [n], \quad (6b)$$

where $[n] := \{1, \dots, n\}$. In the following, we focus exclusively on decision trees, which typically outperform other local learning algorithms, trained to directly output prescriptions instead of predictions (*prescriptive trees*). Fig. 1 presents an overview of the different modeling approaches.

B. Prescriptive Trees with Random Splits

Decision tree learning is a widely popular machine learning algorithm, employed both for classification and regression tasks. The proposed method follows the classification and regression trees (CART) [38] approach, that recursively applies locally optimal binary splits to partition feature space \mathbb{R}^{d_x} , resulting in a set of L leaves. A node split separates a region $R \subseteq \mathbb{R}^{d_x}$ at feature $j \in d_x$ and point s into two disjoint partitions $R = R_l \cup R_r$, such that $R_l = \{i \in [n] : x_{ij} < s\}$ and $R_r = \{i \in [n] : x_{ij} \geq s\}$, with scalar x_{ij} denoting the i -th observation of the j -th feature. Thus, observations that satisfy $x_{ij} < s$ fall to the left of the node, while the rest fall to the right. For brevity of exposition we focus exclusively on quantitative features, although it is straightforward to also include categorical features.

Decision trees are prone to overfitting, i.e., they suffer from high variance, which significantly hinders their predictive capacity. Randomization-based ensemble methods address this issue and lead to impressive predictive performance. Popular methods include bootstrap aggregation (bagging), Random Forests [39] and Extremely Randomized Trees (ExtraTrees) [40].

In this work, we employ the cost function of (1) as an alternative loss function that determines the node split criterion. For each tree node, the locally optimal split is derived from:

$$\min_{j,s} \left[\min_{z_l \in \mathcal{Z}} \sum_{i \in R_l} c(z_l; y_i) + \min_{z_r \in \mathcal{Z}} \sum_{i \in R_r} c(z_r; y_i) \right]. \quad (7)$$

The inner minimization problems correspond to the SAA solution of each partition, with \hat{z}_l, \hat{z}_r being the estimated locally constant decisions of the left and right child node. Problem (7) is of discrete nature and must be solved once per each tree node. When decision trees are trained for prediction (i.e., regression), the standard approach is to order all observations per selected feature j , evaluate each candidate split point, and select the best one. This approach relies on the existence of an analytical solution for the internal minimization problems. In the regression setting, for example, the SAA solution equals the within leaf average, which can be updated recursively for all candidate splits. Unfortunately, this does not apply to general constrained problems. In that case, we need to call a general-purpose convex solver for each of the two SAA problems per each candidate split, which, depending on the underlying problem, could lead to a significant increase in computation time. To this end, we propose employing a randomized split criterion, following the paradigm of the ExtraTrees algorithm [40], which significantly decreases the number of candidate splits evaluated per node. We refer to an ensemble of prescriptive trees as *Prescriptive Forest (PF)*.

For a single prescriptive tree, we start from the top with a full data set and recursively partition the feature space until

no further improvements are possible or a stopping criterion is met. Typical stopping criteria include the maximum tree depth Δ^{max} , the minimum number of observations n_{min} that fall at each leaf, and a predefined threshold for cost reduction. At each node of each tree, we randomly select a subset of K features from X and for each feature randomly select a candidate split point within its range. Next, we estimate the aggregated cost of (7) for each candidate split and compare it with the cost at its root node, updating the tree structure accordingly. The process of training a prescriptive tree is detailed in Algorithm 1.

Algorithm 1 PrescriptiveTree

Input: Data $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, current partition R , current depth Δ , hyperparameters $\{n_{min}, K, \Delta^{max}\}$

Output: Prescriptive tree b

```

1: Determine cost  $v = \min_{z \in \mathcal{Z}} \sum_{i \in R} c(z; y_i)$ 
2: Set  $v_{min} \leftarrow v$ ,  $split \leftarrow \text{False}$ 
3: if  $\Delta < \Delta^{max}$  and  $n \geq 2n_{min}$  then
4:   for  $\kappa = 1, \dots, K$  do
5:     Sample feature  $j \in d_x$  without replacement
6:     Sample split point  $s$  from range of feature  $j$ 
7:     Left child node:  $R_l = \{i \in [n] : x_{ij} < s\}$ 
8:     Right child node:  $R_r = \{i \in [n] : x_{ij} \geq s\}$ 
9:     if  $|R_l| \geq n^{min}$  and  $|R_r| \geq n^{min}$  then
10:       $v = \min_{z_l \in \mathcal{Z}} \sum_{i \in R_l} c(z_l; y_i) + \min_{z_r \in \mathcal{Z}} \sum_{i \in R_r} c(z_r; y_i)$ 
11:      if  $v < v_{min}$  then
12:        Set  $j^* \leftarrow j$ ,  $s^* \leftarrow s$ 
13:        Update  $v_{min} \leftarrow v$ ,  $split \leftarrow \text{True}$ 
14:      end if
15:    end if
16:  end for
17: if  $split == \text{True}$  then
18:    $\mathcal{D}^l = \{(x_i, y_i) : i \in R_l\}$ 
19:    $\mathcal{D}^r = \{(x_i, y_i) : i \in R_r\}$ 
20:    $b^l = \text{PrescriptiveTree}(\mathcal{D}^l, R_l, \Delta + 1)$ 
21:    $b^r = \text{PrescriptiveTree}(\mathcal{D}^r, R_r, \Delta + 1)$ 
22:   Update tree structure  $b$ 
23: end if
24: end if
25: return  $b$ 

```

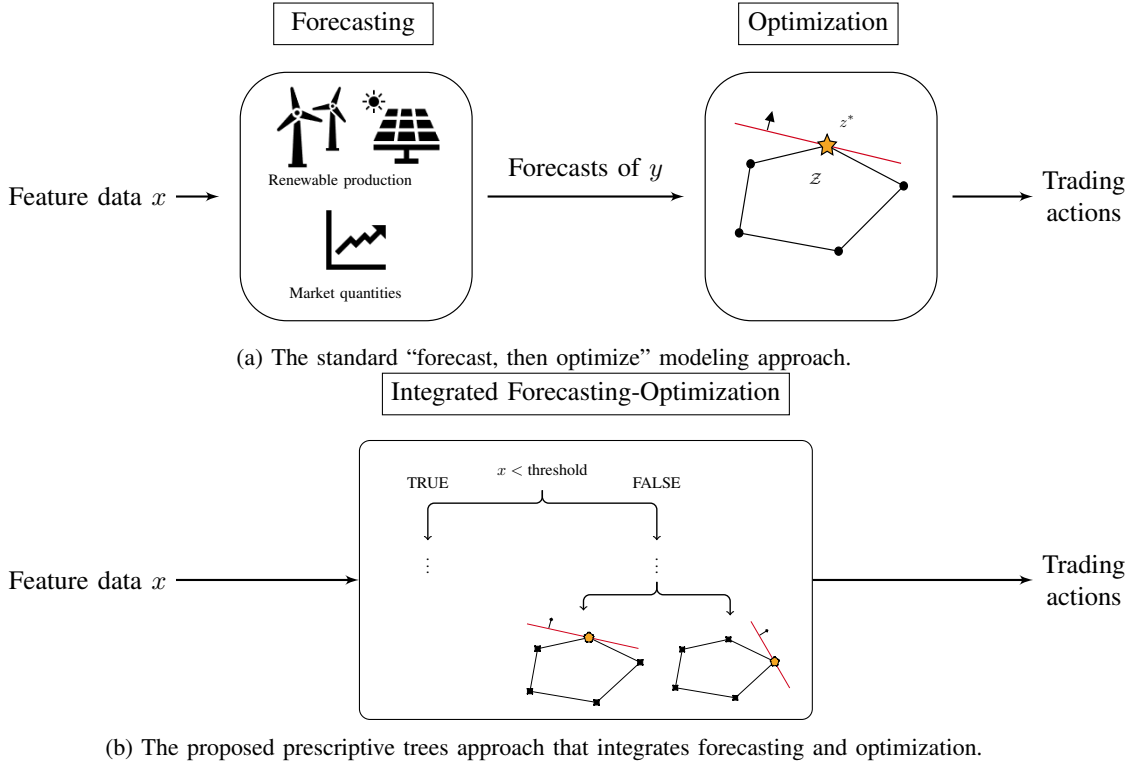
For a single prescriptive tree, the corresponding weights $\omega_{n,i}(\bar{x})$ for a new query \bar{x} are obtained as

$$\omega_{n,i}(\bar{x}) = \frac{\mathbb{I}[R(x_i) = R(\bar{x})]}{|R(\bar{x})|}, \quad (8)$$

where $R(\bar{x})$ is the leaf that \bar{x} falls into, $|\cdot|$ the leaf cardinality and $\mathbb{I}[\cdot]$ an indicator function that checks whether training observation x_i falls into $R(\bar{x})$. For an ensemble of B trees the weights are obtained from

$$\omega_{n,i}(\bar{x}) = \frac{1}{B} \sum_{b=1}^B \frac{\mathbb{I}[R^b(x_i) = R^b(\bar{x})]}{|R^b(\bar{x})|}. \quad (9)$$

From Algorithm 1 we observe that a single tree is fully compiled, i.e., each leaf outputs a prescription, thus providing a direct mapping of input data to decisions, while also ensuring



(a) The standard "forecast, then optimize" modeling approach.

(b) The proposed prescriptive trees approach that integrates forecasting and optimization.

Fig. 1: Deriving trading decisions from contextual information. In both cases the objective is uncertain and must be inferred from available data.

feasibility. For an ensemble of B trees, an additional step is required to approximate the solution from (5), which is detailed in Algorithm 2.

Algorithm 2 PredictivePrescription

Input: New query \bar{x}

Output: Prescription \hat{z}

- 1: Initialize $\omega_{n,i}(\bar{x}) = 0$
 - 2: **for** tree $b = 1, \dots, B$ **do**
 - 3: $\omega_{n,i}(\bar{x}) = \omega_{n,i}(\bar{x}) + \frac{1}{B} \frac{\mathbb{I}[R^b(x_i) = R^b(\bar{x})]}{|R^b(\bar{x})|}$
 - 4: **end for**
 - 5: **return** $\hat{z}(\bar{x}) = \arg \min_{z \in Z} \sum_{i \in [n]} \omega_{n,i}(\bar{x}) c(z; y_i)$
-

As discussed, the main computational cost of Algorithm 1 occurs during the evaluation of candidate splits. The motivating factor behind the random split criterion lies in the expected reduction in computation time, as only a small number of splits are evaluated at each node. Computational experiments between the ExtraTrees and the Random Forest algorithm [40] suggest an average reduction in training time by a factor of 3 for $K = \sqrt{d_x}$, which can rise up to a factor of 10 for wider data sets (larger d_x). Regarding the ensemble size B , the generalization error is expected to monotonically decrease as B increases, thus the computation time is the main consideration for its selection. Note that the task of training an ensemble is trivially parallelizable. Similarly, the rest of the hyperparameters K, n_{min} represent an inherent trade-off between model capacity and computational costs (single trees are maximally grown, thus Δ^{max} is set at infinity). The

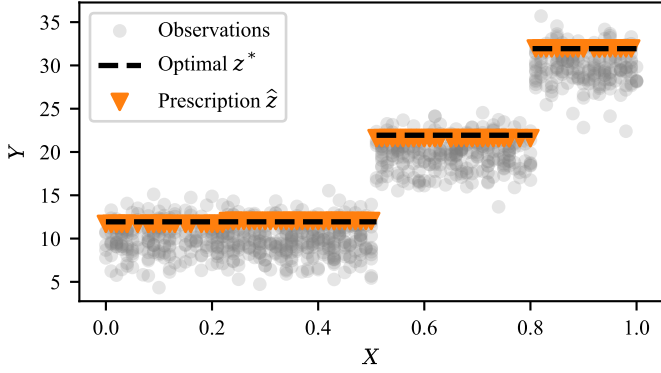
number of candidate splits K controls how strong individual splits are (for $K = 1$ splits are completely random), while regarding n_{min} , larger values result in shallower trees (and reduced computations), with higher bias and lower variance. Section IV presents detailed results on how hyperparameters B, K and n_{min} affect model performance under different loss functions and an empirical comparison for different node splitting criteria.

C. An Illustrative Example

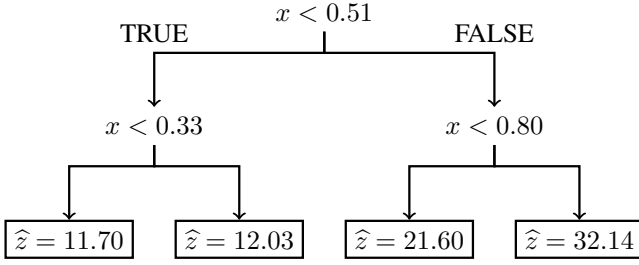
To illustrate the proposed method, we study a toy news vendor problem. Consider an uncertain demand generated from $Y = 10 + 10\mathbb{I}[X > 0.5] + 10\mathbb{I}[X > 0.8] + \epsilon$, with X being a single feature following a uniform distribution $U[0, 1]$ and ϵ random noise following a normal distribution $N(0, 2)$. The cost function is set as $c(z; Y) = 2(Y - z)^- + 10(z - Y)^+$. We sample 1000 observations and train a single prescriptive tree with $\Delta^{max} = 2$. For the purpose of illustration, instead of random splits, we considered evaluating splits at 100 equally spaced quantiles of X . Figure 2a presents a scatterplot with the in-sample fit, while Fig. 2b presents the tree structure, both of which highlight the fact that the tree does a good job of approximating the optimal decision.

D. Measuring the Prescriptiveness

Explainability is pivotal to disseminating the results to industry stakeholders and enabling large scale adoption of analytics tools in real-life applications. Here, we are interested in a quantitative assessment of the impact of the various



(a) In-sample fit and derived decisions.



(b) Prescriptive tree structure.

Fig. 2: Illustrative example.

features on the efficacy of decisions (*prescriptiveness*). This is especially important in cases where obtaining explanatory data incurs in and of itself additional costs, e.g., acquiring weather forecasts for multiple locations. We propose adapting the well-known Mean Decrease Impurity (MDI) in a prescriptive analytics context. Provided a scoring rule, the MDI measures the total decrease in node impurity (dissimilarity) weighted by the probability of reaching a specific node averaged over the ensemble [41]. Considering a prescriptive tree node R_0 partitioned at (j, s) into R_1, R_2 , the decrease in aggregated cost is measured as:

$$\Delta v(j, s) = v(R_0) - v(R_1) - v(R_2). \quad (10)$$

For an ensemble of B trees, the importance of feature j in terms of prescriptiveness, $Imp(j)$, is measured as the aggregated cost decrease over all the nodes that j defines the split variable, over all trees B in the ensemble:

$$Imp(j) = \frac{1}{B} \sum_{b=1}^B \sum_{\ell \in R_{1:L}: j_{\ell}=j} p(b) \Delta v(j_{\ell}, s), \quad (11)$$

with $p(b) = \frac{|R_{\ell}^b|}{n}$ being the proportion of observations reaching node R_{ℓ} in tree b and j_{ℓ} the feature used for splitting that node. The MDI is estimated internally during training, and is thus obtained without additional computational costs.

We also consider measuring prescriptiveness by adapting the permutation importance technique proposed in [39]. First, we estimate aggregated costs with respect to the selected objective function over a hold-out set, which defines the base score. Next, we iterate over all the features, permute (re-shuffle) each one, and derive new prescriptions, repeating the process

a number of times. The permutation importance is then defined as the *expected cost increase* compared to the base score. Preliminary analysis indicated that this approach leads to a significant increase in computational costs, as prescriptions need to be re-optimized at each query. Therefore, we omit it from the experimental results, but note that it presents an attractive alternative for the case of a single prescriptive tree.

III. RENEWABLE TRADING

This section applies the proposed methodology to renewable energy trading. We consider an aggregation of plants participating as a single entity in short-term markets and present two case studies of increasing complexity. First, we derive optimal offers for participating in a DA market. Next, we append a generic storage device and jointly optimize the DA offers and the operational control policy of the storage.

A. Case Study 1: Trading in a DA market

First, we examine trading in a DA market as a price-taker under different balancing mechanisms. Prior to market closure, the producer submits an energy offer p_t^{offer} for each clearing period t of the DA market. As temporal constraints do not apply, subscript t is dropped from the formulation. During real-time (RT) operation, the system operator activates balancing reserves to maintain the demand-supply equilibrium and stabilize the system frequency. The system assumes two states, namely *short*, i.e., demand exceeds supply and upward regulation is required, and *long*, i.e., supply exceeds demand and downward regulation is required. Based on RT production, the producer buys back (sells) the amount of energy shortage (surplus) in order to balance its individual position. In the following, we present problem formulations that apply to different balancing market designs.

Let p^E denote the stochastic renewable production, π^{da} the clearing price of the DA market and $\pi^{\uparrow/\downarrow}$ the marginal cost of activating upward/downward regulation services. Under the assumption of individual rationality, shortage of supply leads to increased RT marginal costs. Thus, we assume that if the system is short, it holds that $\pi^{\uparrow} \geq \pi^{da}$ and $\pi^{\downarrow} = \pi^{da}$, while if the system is long, then $\pi^{\downarrow} \leq \pi^{da}$ and $\pi^{\uparrow} = \pi^{da}$. Let us further define $\lambda^{\uparrow} = \max\{0, \pi^{\uparrow} - \pi^{da}\}$ and $\lambda^{\downarrow} = \max\{0, \pi^{da} - \pi^{\downarrow}\}$ as the respective upward and downward unit regulation costs. Evidently, it holds that $\lambda^{\uparrow} \cdot \lambda^{\downarrow} = 0$, i.e., only one of them (at most) assumes a value greater than zero for a given settlement period.

Under a *single-price* balancing mechanism, the profit for each settlement period is defined as:

$$\begin{aligned} \rho^{\text{single}} &= \pi^{da} p^{offer} + \pi^{\uparrow} (p^E - p^{offer}) + \pi^{\downarrow} (p^E - p^{offer}) \\ &= \pi^{da} p^E - \underbrace{[-\lambda^{\uparrow} (p^E - p^{offer}) + \lambda^{\downarrow} (p^E - p^{offer})]}_{\text{imbalance cost}}. \end{aligned} \quad (12)$$

Here, $\{p^E, \lambda^{\uparrow}, \lambda^{\downarrow}\}$ defines the uncertain problem parameters (i.e., parameter Y). Since profit is affine with respect to the contracted energy, it is trivial to derive the optimal energy offer analytically as:

$$p^{offer*} = \begin{cases} p^{\min}, & \text{if } -\hat{\lambda}^{\uparrow} + \hat{\lambda}^{\downarrow} \leq 0 \\ p^{\max}, & \text{if } -\hat{\lambda}^{\uparrow} + \hat{\lambda}^{\downarrow} > 0, \end{cases} \quad (13)$$

where $\hat{\cdot}$ denotes expected (forecast) values, see [22, Section II] for proof. We interpret (13) as follows: the optimal offer equals zero if the system is expected to be short (typically $p^{min} = 0$) and the nominal capacity if the system is expected to be long. The case of zero costs is merged with the former without loss of generality. In practice, however, following this policy incurs great risks and could constitute market abuse, which motivates designing a strategy that does not lead to excessive imbalances.

Alternatively, if the balancing market operates under a *dual-price* balancing mechanism, (12) is modified to impose non-arbitrage between the DA and the balancing market. The single period profit is now defined as:

$$\rho^{dual} = \pi^{da} p^E - \underbrace{[-\lambda^\uparrow (p^E - p^{offer})^- + \lambda^\downarrow (p^E - p^{offer})^+]}_{\text{imbalance cost}}, \quad (14)$$

where $(\cdot)^- = \min\{\cdot, 0\}$ and $(\cdot)^+ = \max\{\cdot, 0\}$. Hence, the term defining the imbalance cost is non-negative, which in turn means that no additional profit can be attained in the balancing market. This contrasts the single-price market design, where deviations that help restore the system frequency result in negative imbalance costs, i.e., additional profit. Provided probabilistic forecasts for all the uncertain parameters, the optimal offer is derived analytically as:

$$p^{offer*} = \hat{F}^{-1}\left(\frac{\hat{\lambda}^\downarrow}{\hat{\lambda}^\downarrow + \hat{\lambda}^\uparrow}\right), \quad (15)$$

where \hat{F}^{-1} is the predicted inverse cumulative distribution function (c.d.f.) of p^E . The above solution holds without assuming independence between energy production and unit regulation costs, see [22, Section III] for details. The optimal offering strategy thus requires probabilistic forecasts of energy production and deterministic forecasts of unit regulation costs.

For both cases, we propose a hybrid policy that balances profit maximization and energy forecasting accuracy. Specifically, the problem is formulated as:

$$\min_{p^{offer}} \mathbb{E} \left[(1-k)(-\rho^{\text{single/dual}}) + k \|p^E - p^{offer}\|_2^2 \right] \quad (16a)$$

$$\text{s.t.} \quad p^{min} \leq p^{offer} \leq p^{max}. \quad (16b)$$

The objective function (16a) minimizes a convex combination of (normalized) trading cost and prediction error, which sets the loss function of the tree algorithm. We can interpret this objective as adding a regularization term that penalizes excessive deviations from the expected energy production; contrary to other risk-averse formulations, this provides a more intuitive trade-off to stakeholders. This trade-off is controlled from design parameter k . Specifically, for $k = 0$ we retrieve a purely prescriptive task, while for $k = 1$ we obtain a purely predictive task with a standard regression loss.

B. Case Study 2: Including Storage

Energy storage presents a promising avenue to support the participation of renewable plants in electricity markets and

enhance their profit [27], offering functions such as arbitraging in DA markets and compensating for deviations from the submitted schedule during real-time operation. Here, we extend the previous problem by appending a generic storage device to the aggregation of renewable plants and jointly optimize the DA offers, considering a closed system, and the operational control policy of the storage. We maintain a similar setting as before, i.e., the aggregation participates in a DA market subject to imbalance penalties, considering a dual-price balancing mechanism. Participating in additional markets, such as intraday, or offering balancing services is not examined. Optimizing over the operational policy of the storage means that we allow recourse actions based on the realization of uncertainty. This defines a multi-stage dynamic optimization problem; a tractable reformulation is provided by applying the linear decision rule (LDR) approach [36], modeling real-time decisions as an affine function of uncertainty, in this case energy forecast error. Throughout, we use $\tilde{\cdot}$ to denote decisions that depend on the realization of uncertainty. Index t is used to define a specific time period (scalar), while absence of t defines a vector over the DA horizon of length T .

Let $\xi \in \Xi \subseteq \mathbb{R}^T$ define the energy forecast error for the DA horizon, i.e., a sample path of errors of length T , and Ξ define an uncertainty set. The uncertain renewable production is defined as $p^E = \hat{p}^E + \xi$, i.e., the expected value (forecast) $\hat{p}^E \in \mathbb{R}^T$ plus the error term ξ . We define recourse actions of the storage as an affine function of uncertainty. For example, the decision vector for charging is defined $\tilde{p}^{ch} = \hat{p}^{ch} + D^{ch}\xi$; here, $\hat{p}^{ch} \in \mathbb{R}^T$ denotes the scheduled DA decisions and $D^{ch} \in \mathbb{R}^{T \times T}$ is a linear coefficient matrix that determines the operational policy mapping realizations of uncertainty ξ to recourse actions. Note that the whole error history is considered; to retain non-anticipativity we require D^{ch} to be lower-triangular.

We consider a modified version of [27] and design a control policy that aims at minimizing the imbalance volume, without considering the balancing mechanism in the objective function. The results presented in Section IV show that this represents a realistic application for a dual-price balancing mechanism. Following our previous formulation, we optimize over a convex combination of trading performance and deviations from the contracted energy. For simplicity, we assume DA prices are known. The problem is defined as

$$\min_{\mathcal{P}} \mathbb{E} \left[\sum_{t=1}^T -(1-k)\pi_t^{da} p_t^{offer} + k \|p_t^{output} - p_t^{offer}\|_2^2 \right] \quad (17a)$$

$$\text{s.t.} \quad p^{min} \leq p^{offer} \leq p^{max}, \quad (17b)$$

$$p_t^{soc} = p_{t-1}^{soc} + \eta^{ch} \tilde{p}_{t-1}^{ch} + \frac{1}{\eta^{dis}} \tilde{p}_{t-1}^{dis} \quad \forall t \in [T], \quad (17c)$$

$$p_1^{soc} = p_T^{soc} = p_0, \quad (17d)$$

$$p^{output} = \hat{p}^E + \xi + \tilde{p}^{dis} - \tilde{p}^{ch}, \quad (17e)$$

$$0 \leq \tilde{p}^{dis} \leq c^{dis} \quad \forall \xi \in \Xi, \quad (17f)$$

$$0 \leq \tilde{p}^{ch} \leq c^{ch} \quad \forall \xi \in \Xi, \quad (17g)$$

$$0 \leq p^{soc} \leq B^{max} \quad \forall \xi \in \Xi, \quad (17h)$$

where $\mathcal{P} = \{p^{offer}, \hat{p}^{ch}, \hat{p}^{dis}, D^{ch}, D^{dis}\}$ is the set of decision variables, and p^{soc}, p^{output} are auxiliary variables for the induced state of charge in the storage and the actual output of the plant-storage system. The objective (17a) minimizes a convex combination of trading profit from the DA market and deviations between actual output and the contracted energy. The trade-off is controlled from parameter k . For $k = 0$ the main function of the storage is to arbitrage in the DA market, while for $k = 1$ the focus is placed on compensating deviations from the schedule during real-time operation. Problem constraints include the limits for contracted energy (17b), the state transition equation of the storage (17c), initial and terminal conditions for the state of charge (17d), and technical limits of the storage (17f)-(17h). In a data-driven setting, we assume that ξ belongs to a finite set $\Xi = \{\xi_i\}_{i=1}^n$, which we use to approximate the objective (17a) (remark that the i -th observation ξ_i is a sample path of length T). Further, we employ duality theory and standard techniques from robust optimization, to reformulate (17f)-(17h). For reference, constraint (17f) is reformulated as follows. First, we define a polyhedral uncertainty set $\Xi' = \{\xi : H\xi \leq h\}$, where $H = [I, -I]^T \in \mathbb{R}^{2T \times T}$, with I defining an identity matrix, and $h \in \mathbb{R}^{2T}$ containing the upper and lower bound for each period t . Next, we define the inner max problem

$$\left\{ \begin{array}{l} \max_{\xi} \quad \hat{p}^{dis} + D^{dis} \xi \\ \text{s.t.} \quad H\xi \leq h : \mu \end{array} \right\} \leq c^{dis} \quad (18)$$

where μ denotes the dual variables. From duality, we derive

$$\left\{ \begin{array}{l} \max_{\mu} \quad -h^T \mu \\ \text{s.t.} \quad H^T \mu = D^{dis} \end{array} \right\} \leq c^{dis} - \hat{p}^{dis} \quad (19)$$

which finally leads to

$$\exists \mu, \text{ with } h^T \mu \leq c^{dis} - \hat{p}^{dis}, H^T \mu = D^{dis}, \mu \geq 0. \quad (20)$$

The rest of the constraints are reformulated in a similar fashion. If the uncertainty set Ξ' is too wide, no control will take place during real-time operation, while if it is too tight, it is possible to get infeasible actions. During the actual implementation, we add a saturation block to ensure feasible recourse actions. Therefore, the maximum charge is set as $\min\{c^{in}, \frac{B_{max} - p^{soc}}{\eta^{ch}}\}$, while the maximum discharge is $\min\{c^{out}, p^{soc} \eta^{dis}\}$. Lastly, we note that Ξ' varies on an hourly and daily basis. To determine h we use the intervals derived from the underlying samples ξ_i . For reference, consider the example forecasts shown in Fig. 3. At 00:00 scenarios show small dispersion (i.e., less uncertainty), which results in tighter upper and lower bounds in h . On the other hand, at 12:00 the derived bounds in h are wider, due to the larger dispersion of the underlying scenarios.

C. Energy and Price Forecasting

As discussed, the standard FO approach requires forecasting the uncertain quantities prior to solving a stochastic optimization problem. The problem described above requires estimating the conditional distribution of p^E and the conditional

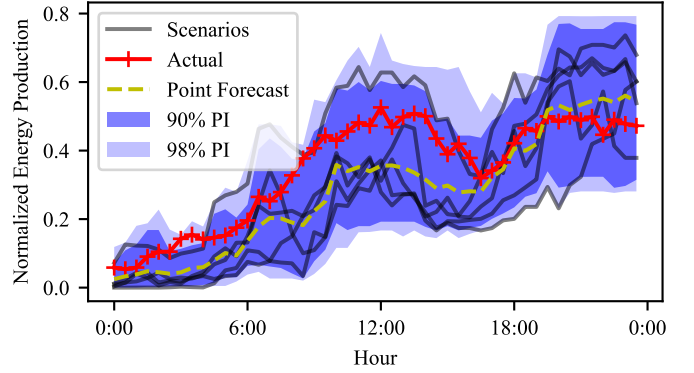


Fig. 3: Day-ahead renewable energy forecasts: point forecasts, probabilistic forecasts (prediction intervals or PI), and scenarios.

expectations of $\lambda^{\uparrow/\downarrow}$. To this end, the producer employs feature vectors x^E and x^{market} that include information associated with the uncertain parameters, e.g. weather predictions, historical energy production, and historical market prices, among others. As our purpose is not to provide a comprehensive analysis of forecasting models, we employ established benchmarks. For probabilistic energy forecasting, we select the Quantile Regression Forests model, a machine learning model with state-of-the-art performance in energy forecasting [42]. For the second case study, the standard modeling approach requires as input temporally correlated scenarios (trajectories) of renewable production. To this end, we employ a Gaussian copula function, given the conditional marginal predictive densities obtained from the QRF, following the procedure detailed in [43]. For reference, different types of forecasts are shown in Fig. 3.

Contrary to energy forecasting, the literature on forecasting unit regulation costs is relatively scarce. A standard practice is to partition the problem into three prediction tasks, namely upward and downward regulation cost and direction prediction. Finally, individual components are combined according to the requirements of the specific policy by the law of total expectation. Specifically, the individual components are:

$$\hat{\phi} = \mathbb{P}(\lambda^{\uparrow} > 0), \quad (21)$$

$$\hat{\lambda}^{\uparrow} = \hat{\phi} \mathbb{E}[\lambda^{\uparrow} | \lambda^{\uparrow} > 0], \quad (22)$$

$$\hat{\lambda}^{\downarrow} = (1 - \hat{\phi}) \mathbb{E}[\lambda^{\downarrow} | \lambda^{\downarrow} > 0], \quad (23)$$

where $\hat{\phi}$ is the estimated probability of the system being short. Therefore, the prediction for the upward unit regulation cost λ^{\uparrow} equals the expectation of a regression model trained conditionally on the system being short, weighted by probability $\hat{\phi}$. Following [44], we apply exponential smoothing to model the individual components.

IV. EVALUATION AND RESULTS

A. Experiment Setup

1) *Data*: For the computational experiments, we consider an aggregation of 3 WPPs and 1 PV plant, with a total

TABLE II: Storage parameters (normalized).

Parameter	Value
B^{max}	0.5
c^{ch}	$0.5B^{max}$
c^{dis}	$0.2B^{max}$
η^{ch}	0.8
η^{dis}	0.9

capacity of 49 MW (16% PV share), respectively located in northern and southern France. For the storage case study, a typical set of parameters is shown in Table II. Contextual information x comprises features typically used as input in forecasting applications. A forecast horizon of 12-36 hours ahead is considered. In order to deal with possible temporal correlations, we conduct a preliminary analysis by examining the partial autocorrelation function (PACF) of target variables (energy and prices) and include relevant lags as additional features in x . By sufficiently enlarging the feature space x with historical lags, we treat training data (y_i, x_i) as i.i.d.

Regarding renewable production, the feature vector x^E comprises weather forecasts from a numerical weather predictions (NWP) model, namely wind speed, wind direction, temperature, cloud coverage, and solar radiation forecasts for each plant location. The NWP forecasts are issued at 00:00 on day D-1 spanning a horizon of 24-48 hours ahead. Examining the PACF did not reveal any lags to be important, thus we do not include any in x^E ; this result is standard in renewable forecasting for horizons larger than a couple of hours ahead.

Regarding market data, we employ data for the French electricity market for the same period, downloaded from [45]. Market-related contextual information x^{market} include historical lags (as indicated from the PACF) for DA prices (one day and one week prior), historical lags for system imbalance volumes (two days prior), and DA forecasts for available thermal generation, electricity demand, and renewable generation at transmission level. The system-wide forecasts issued from the operator are processed to determine a net load series, by subtracting the expected renewable production from the expected electricity demand, and a system margin series, defined as the ratio of net load to available thermal generation. In addition, we include categorical variables for the calendar effect, namely day of the week and hour of the day. For the tree algorithm, feature vectors x^E and x^{market} are concatenated, resulting in a total of $d_x = 20$ features. Models are trained on one year of historical data spanning 2019 and evaluated on the first 4 months of 2020. Lastly, a half-hour settlement period is assumed for the DA and balancing market.

2) *Benchmarks*: Our goal is to showcase the ability of the proposed approach to provide informed decisions under different objectives using a single data-driven model, without the need to deploy multiple forecasting models. The following approaches are compared:

- *FO*: The standard sequential modeling approach. This involves i) deriving probabilistic energy forecasts using the QRF model, ii) forecasting the unit regulation costs with exponential smoothing, iii) solving a stochastic

optimization problem.

- *PF*: Predictive prescriptions with weights derived from a prescriptive forest with random splits.

For the sake of comparison, the naive SAA solution and the perfect-foresight solution are also estimated. Optimization problems are solved either analytically, when applicable, or numerically. As mentioned, different values of design parameter k define different objectives; thus, a different prescriptive forest is trained for each value of k . Note that we consider offline (batch) learning, thus we implicitly assume stationarity. If the underlying processes are non-stationary, then offline learning will not suffice and online learning should be considered. This is outside the scope of the current work but presents an interesting research direction.

3) *Evaluation*: The proposed approaches are examined in terms of out-of-sample prescriptive performance and trading results. For the former, we employ the coefficient of prescriptiveness P [5], a unitless metric that measures relative optimization performance against the naive SAA and the perfect-foresight solution. Specifically, for each i in $\{FO, PF\}$ and different values of design parameter k the coefficient P is estimated as:

$$P_{i,k} = 1 - \frac{\widehat{v}_k^i - \widehat{v}_k^*}{\widehat{v}_k^{SAA} - \widehat{v}_k^*}, \quad (24)$$

where $\widehat{v}_k^i, \widehat{v}_k^{SAA}, \widehat{v}_k^*$ are the aggregated cost over the test set under the i , SAA, and perfect-foresight method. Note that the SAA approach optimizes over the empirical distributions without leveraging contextual information. The coefficient P is bounded above by one, while negative values indicate a failure to outperform the SAA. Regarding trading results, we estimate aggregated profit and risk. The conditional value at risk at 5% level ($CVaR_{5\%}$) is used as a proxy for trading risk, defined as the expected profit over the 5% worst returns.

B. Effect of Hyperparameters and Split Criterion

This subsection examines the performance of the tree algorithm with respect to hyperparameters $\{B, K, n_{min}\}$ in a controlled setting. We consider the problem of trading in a DA market under a single-price balancing mechanism as a test bed and examine prescriptive performance for values of $k = \{0, 0.5, 1\}$ by randomly sampling $n = 1000$ training and validation observations and estimating the respective coefficient of prescriptiveness P . The process is repeated 10 times.

Fig. 4a plots the prescriptive performance as a function of the ensemble size B for the different values of k . The performance appears to be insensitive to the size of the ensemble, with similar results across the different tasks. Note that the discrepancy across the levels of coefficient P for the different values of k is attributed to the relative difficulty of the underlying problem; for example, for $k = 1$ we retrieve higher values of P , which means that the regression task is relatively easier. We discuss this in more detail in the next subsection. Next, we examine the effect of number of splits evaluated per node K , which controls the model capacity. For $K = 1$ node splits are completely random (requiring minimum computations), while for $K = d_x$ all features are considered. From Fig. 4b a significant discrepancy across tasks is evident.

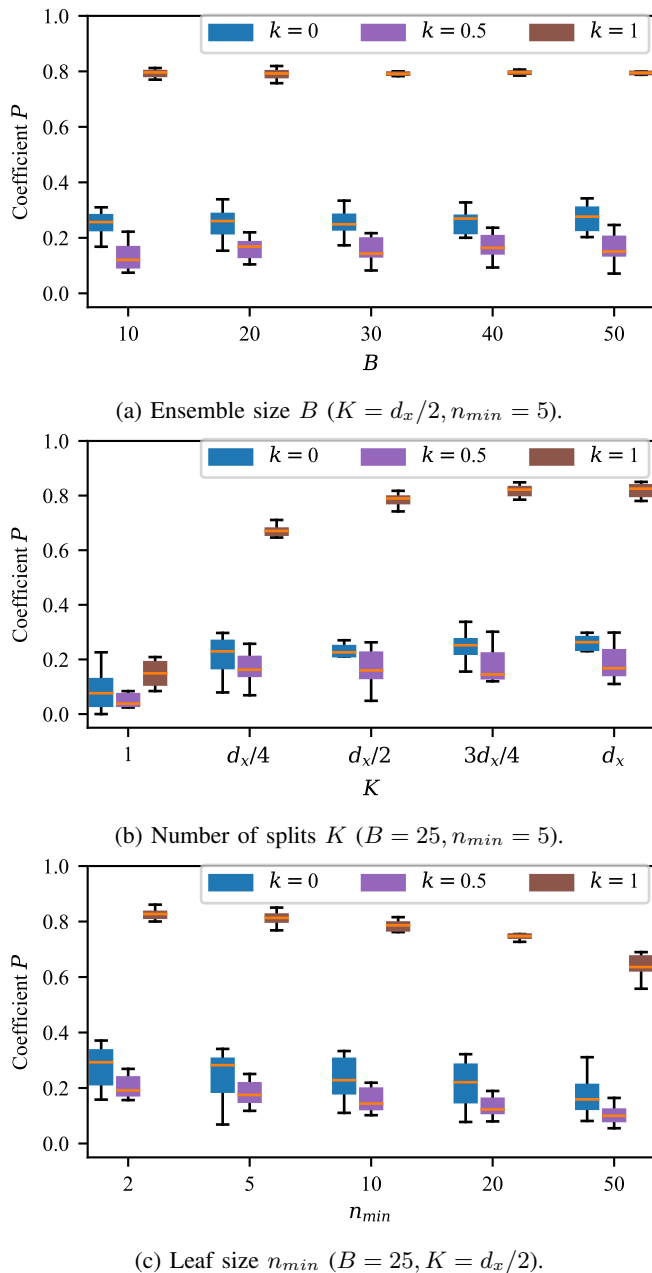


Fig. 4: Effect of hyperparameters B , K , and n_{min} .

Specifically, the selection of K has a notable effect on the performance of the predictive task, with a significant decrease for lower values of K , while this effect is less pronounced for the other tasks. Overall, higher values of K lead to increased prescriptive performance for all tasks. Finally, we examine the impact of the minimum leaf size n_{min} . Generally, smaller values of n_{min} result in lower bias, while larger values provide a smoothing effect. Fig. 4c indicates a decrease in performance for values of leaf size greater than 10, with the effect being more pronounced for the predictive task ($k = 1$). For the rest of this section, the results presented are estimated with hyperparameters $K = 3d_x/4, B = 50$, and $n_{min} = 10$.

Lastly, we repeat the experiment for $k = 0.5$ and examine different node splitting criteria. Specifically, we consider

ordering observations and evaluating all splits as in Random Forests (RF), evaluating on 10 equally spaced quantiles of the empirical distribution of each feature (RF-Q), and random splits as in ExtraTrees (ET). We remark that the effect of the hyperparameters varies for the different splitting criteria. Thus, we are not primarily interested in an exhaustive comparison in terms of prescriptive performance, but rather want to highlight the effect of the selected criterion on computational costs for a specific set of hyperparameters. Table III presents results in terms of prescriptive performance and average CPU time train a single tree over 10 iterations using a standard machine with an Intel Core i7 CPU with a 2.3GHz clock rate and 32GB of RAM. We observe that the random split criterion shows a significant reduction in computation time, both against the RF and the RF-Q approaches, without compromising prescriptive performance. Note that the training time depends on the structure of the underlying problem. In this experiment the problem is relatively simple; for larger problems (e.g., including storage) the RF criterion becomes intractable.

TABLE III: Average performance (\pm one standard deviation) for sample size $n = 1000$.

Split criterion	RF	RF-Q	ET
Coefficient P	0.16 \pm 0.08	0.18 \pm 0.05	0.16 \pm 0.04
Single tree CPU time (sec)	650.58 \pm 103.84	26.43 \pm 1.80	2.15 \pm 0.24

C. Case Study 1 Results

1) *Single-price market*: First, we examine results for a single-price balancing market. Table IV presents aggregated results for $k = \{0, 0.25, 0.5, 0.75, 1\}$, with the PF approach leading to an expected profit increase of 3.82% across all values of k , with a maximum of 7.44% for $k = 0.25$. As expected, larger values of k lead to more conservative offers and thus to a higher $CVaR_{5\%}$, as the minimization of the imbalance volume is weighted more heavily in the objective. This is evident in Fig. 5, with offers showing larger deviations from actual production as k decreases. Fig. 6 further highlights the improved risk-reward trade-off of the PF, as it sets the efficient frontier, attaining higher revenue for a given level of risk and vice versa. The above observations regarding trading performance are further validated by examining the coefficient of prescriptiveness P , with the PF showing improved performance for all k smaller than 1. Interestingly, the FO approach fails to outperform the naive SAA solution for lower values of k (P is negative). Both approaches converge to similar performance for $k = 1$, which effectively means that the proposed method achieves similar predictive performance with the QRF algorithm. Lastly, we observe significantly lower values of coefficient P for k less than 1 compared to the regression task, highlighting the fact that trading is relatively more demanding than simple energy forecasting, as the relative distance from the perfect-foresight solution is greater. This could be attributed to the fact that trading tasks require an estimation of unit regulation costs in the DA horizon, which in practice is known to be difficult. Nonetheless, the results

TABLE IV: Results For Single-Price Market

k	0		0.25		0.5		0.75		1	
Method	FO	PF	FO	PF	FO	PF	FO	PF	FO	PF
Total Profit (10^3 EUR)	1 191	1 250	1 170	1 257	1 170	1 225	1 182	1 212	1 184	1 178
$CVaR_{5\%}$ (EUR)	-442.44	-353.51	-403.50	-281.84	-243.29	-228.68	-119.56	-132.32	-92.45	-105.12
Coefficient P	0.06	0.15	-0.01	0.13	-0.01	0.08	0.11	0.17	0.85	0.85

TABLE V: Results For Dual-Price Market

k	0		0.25		0.5		0.75		1	
Method	FO	PF	FO	PF	FO	PF	FO	PF	FO	PF
Total Profit (10^3 EUR)	1 130	1 137	1 130	1 140	1 130	1 140	1 130	1 141	1 141	1 138
$CVaR_{5\%}$ (EUR)	-97.29	-106.30	-97.30	-99.98	-97.27	-104.14	-97.30	-104.98	-99.94	-107.46
Coefficient P	0.62	0.66	0.63	0.68	0.64	0.69	0.66	0.72	0.86	0.86

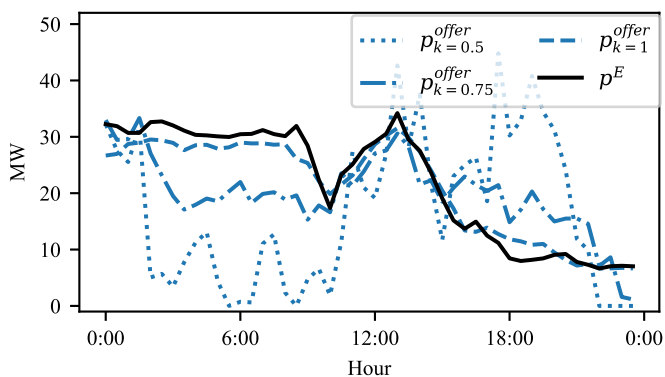
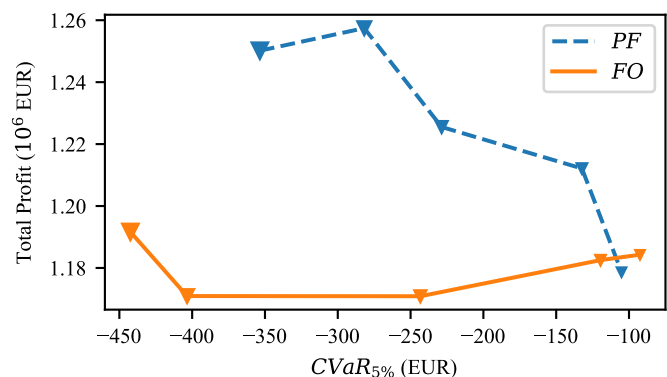


Fig. 5: Illustration of actual production and different DA offers for a single day.

provide a quantitative estimation of this empirical knowledge, which we believe to be of use to both researchers and other stakeholders.

2) *Dual-price market*: Table V presents aggregated results for a dual-price balancing mechanism. Contrary to the single-price case, trading performance is rather insensitive to the selection of parameter k , with the PF resulting in an expected profit increase of 0.62%. In fact, we do not observe significant differences between the optimal trading strategy ($k = 0$) and offering the expected energy production ($k = 1$). This could be attributed to the non-arbitrage condition imposed by the market design. This notion is further supported by examining the evolution of P , which for lower values of k is significantly greater than the single-price market. Overall, the PF consistently leads to similar or improved prescriptive performance as shown by the values of coefficient P . We highlight the fact that this improved performance also comes at a reduced modeling effort, as we employ a single data-driven model and thus avoid developing multiple forecasting models.

3) *Feature Importance*: Next, we investigate how different features affect prescriptive performance, as measured by the adapted MDI method. A subset of the most important features is plotted in Fig. 7, with feature importances normalized to add up to one. Considering a single-price market, we observe that for lower values of k , variables pertaining to estimating

Fig. 6: Risk versus reward for trading in a single-price market. Marker size is analogous to k . Values towards top and right are preferred.

unit regulation costs assume greater weight. This is attributed to the PF placing more weight on the trading cost term in the objective. Specifically, the aggregation of the expected system margin and net load, expected temperature at the WPP site, and lag observations for system imbalance volume assume approximately 65% of the total feature importance for $k = \{0, 0.25, 0.5\}$. Note that the WPPs are located in close proximity to large metropolitan areas and interconnections with neighbouring countries, thus the expected temperature effectively serves as a proxy for electricity demand. As k increases, the importance of features related to energy forecasting gradually increases, with the expected wind speed at the WPP site reaching approximately 75% of the total feature importance for $k = 1$.

Under a dual-price mechanism, we observe significantly fewer variations in feature importance across the different values of k , resembling the results in Table V. Specifically, the expected wind speed at the WPP location is consistently the most important variable throughout, with its importance ranging from 60% to 78%, followed by the expected system margin. Previous works on similar case studies mention that renewable forecasting is relatively more important than price forecasting [33]. The results presented in Table V and Fig. 7

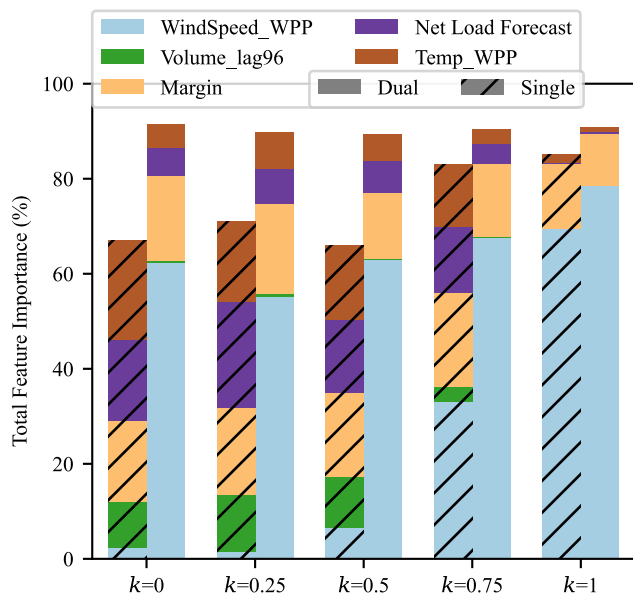


Fig. 7: Normalized feature importance based on prescriptiveness.

provide quantitative evidence for these assertions by jointly considering the two sources of uncertainty in the problem formulation and measuring the impact of different features.

Finally, examining results across the different market designs, enables us to conclude that forecasting regulation costs is relatively more important if a single-price balancing mechanism is in place, while renewable energy forecasting should be the primary focus for participants in dual-price markets.

D. Case Study 2 Results

Next, we present results for the storage case study. To examine trading performance, we assume participation in a market with a dual-price balancing mechanism. As indicated by the previous results in Section IV-C2, in practice there is no significant difference between the optimal offering strategy and offering the expected energy production under this market design. Therefore, the operational control policy implemented, i.e., using storage to minimize deviations from the submitted schedule, also makes sense from an economic perspective. Table VI shows results for a value of $k = 0.75$. In terms of trading performance, the proposed method attains a 3.07% profit increase, accompanied with a decrease in $CVaR_{5\%}$. The prescriptive performance is also improved as evident by the increased value of coefficient P . Lastly, compared to the previous case study, significantly higher profits are observed. This highlights the storage’s ability to support renewables in market applications and validates the applicability of the proposed control policy.

V. CONCLUSIONS

This work presented an integrated forecasting and optimization approach to improve prescriptive performance in renewable trading applications. We proposed a prescriptive tree

TABLE VI: Results-Storage ($k = 0.75$)

Method	FO	PF
Total Profit (10^3 EUR)	1 628	1 678
$CVaR_{5\%}$ (EUR)	-8.88	-6.12
Coefficient P	0.89	0.92

algorithm that minimizes task-specific costs for conditional stochastic optimization problems, employing a random split criterion to speed up computations. Further, we provided a framework to measure feature importance in terms of the impact on optimization efficacy under different objective functions.

We validated the proposed approach in two case studies related to short-term trading of renewable generation and examined performance under different balancing mechanisms. The proposed approach led to an average increase in aggregate profit of 3.82% and 0.62% for trading in a DA market under a single- and dual-price balancing mechanism, respectively. In an additional case study that included a generic storage in the dual-price setting, a 3.07% profit increase was shown. Overall, we observed consistently better or similar prescriptive performance against the current state of the art, which also came at a reduced modeling effort. Lastly, we examined feature importance under different objectives and across different market designs, demonstrating the capability of the proposed solution to measure the impact of data on decision-making, and provided insights on trading under different regulatory frameworks. Future work could focus on learning in an adaptive (online) setting and enhancing model interpretability.

ACKNOWLEDGMENT

The authors wish to thank ENGIE GREEN for providing historical production data.

REFERENCES

- [1] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, and H. Zareipour, “Energy forecasting: A review and outlook,” *IEEE Open Access Journal of Power and Energy*, 2020.
- [2] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown *et al.*, “Tackling climate change with machine learning,” *arXiv preprint arXiv:1906.05433*, 2019.
- [3] H. Zareipour, C. A. Canizares, and K. Bhattacharya, “Economic impact of electricity market price forecasting errors: A demand-side analysis,” *IEEE Transactions on Power Systems*, vol. 25, no. 1, pp. 254–262, 2009.
- [4] J.-F. Toubeau, J. Bottieau, F. Vallée, and Z. De Grève, “Deep learning-based multivariate probabilistic forecasting for short-term scheduling in power markets,” *IEEE Transactions on Power Systems*, vol. 34, no. 2, pp. 1203–1215, 2018.
- [5] D. Bertsimas and N. Kallus, “From predictive to prescriptive analytics,” *Management Science*, vol. 66, no. 3, pp. 1025–1044, 2020.
- [6] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, 2014.
- [7] D. Bertsimas and B. Van Parys, “Bootstrap robust prescriptive analytics,” *Mathematical Programming*, pp. 1–40, 2021.
- [8] D. Bertsimas and C. McCord, “From predictions to prescriptions in multistage optimization problems,” *arXiv preprint arXiv:1904.11637*, 2019.
- [9] D. Bertsimas, C. McCord, and B. Sturt, “Dynamic optimization with side information,” *arXiv preprint arXiv:1907.07307*, 2019.

- [10] R. Kannan, G. Bayraktan, and J. R. Luedtke, "Data-driven sample average approximation with covariate information," *Optimization Online*. URL: http://www.optimization-online.org/DB_HTML/2020/07/7932.html, 2020.
- [11] A. N. Elmachtoub and P. Grigas, "Smart "predict, then optimize,"" *Management Science*, 2021.
- [12] Y. Bengio, "Using a financial training criterion rather than a prediction criterion," *International journal of neural systems*, vol. 8, no. 04, pp. 433–443, 1997.
- [13] P. L. Donti, B. Amos, and J. Z. Kolter, "Task-based end-to-end model learning in stochastic optimization," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 5490–5500.
- [14] A. N. Elmachtoub, J. C. N. Liang, and R. McNellis, "Decision trees for decision-making under the predict-then-optimize framework," *arXiv preprint arXiv:2003.00360*, 2020.
- [15] M. A. Muñoz, S. Pineda, and J. M. Morales, "A bilevel framework for decision-making under uncertainty with contextual information," *arXiv preprint arXiv:2008.01500*, 2020.
- [16] G.-Y. Ban and C. Rudin, "The big data newsvendor: Practical insights from machine learning," *Operations Research*, vol. 67, no. 1, pp. 90–108, 2019.
- [17] N. Kallus and X. Mao, "Stochastic optimization forests," *arXiv preprint arXiv:2008.07473*, 2020.
- [18] N. Mundru, "Predictive and prescriptive methods in operations research and machine learning: an optimization approach," Ph.D. dissertation, Massachusetts Institute of Technology, 2019.
- [19] G. Kariniotakis, *Renewable Energy Forecasting: From Models to Applications*. Woodhead Publishing, 2017.
- [20] J. Nowotarski and R. Weron, "Recent advances in electricity price forecasting: A review of probabilistic forecasting," *Renewable and Sustainable Energy Reviews*, vol. 81, pp. 1548–1568, 2018.
- [21] P. Pinson, C. Chevallier, and G. N. Kariniotakis, "Trading wind generation from short-term probabilistic forecasts of wind power," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1148–1156, 2007.
- [22] C. J. Dent, J. W. Bialek, and B. F. Hobbs, "Opportunity cost bidding by wind generators in forward markets: Analytical results," *IEEE Transactions on Power Systems*, vol. 26, no. 3, pp. 1600–1608, 2011.
- [23] E. Y. Bitar, R. Rajagopal, P. P. Khargonekar, K. Poolla, and P. Varaiya, "Bringing wind energy to market," *IEEE Transactions on Power Systems*, vol. 27, no. 3, pp. 1225–1235, 2012.
- [24] J. M. Morales, A. J. Conejo, and J. Pérez-Ruiz, "Short-term trading for a wind power producer," *IEEE Transactions on Power Systems*, vol. 25, no. 1, pp. 554–564, 2010.
- [25] S. Camal, A. Michiorri, and G. Kariniotakis, "Optimal offer of automatic frequency restoration reserve from a combined pv/wind virtual power plant," *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 6155–6170, 2018.
- [26] J. Browell, "Risk constrained trading strategies for stochastic generation with a single-price balancing market," *Energies*, vol. 11, no. 6, p. 1345, 2018.
- [27] H. Ding, P. Pinson, Z. Hu, and Y. Song, "Optimal offering and operating strategies for wind-storage systems with linear decision rules," *IEEE Transactions on Power Systems*, vol. 31, no. 6, pp. 4755–4764, 2016.
- [28] J. Han, L. Yan, and Z. Li, "A task-based day-ahead load forecasting model for stochastic economic dispatch," *IEEE Transactions on Power Systems*, vol. 36, no. 6, pp. 5294–5304, 2021.
- [29] L. Wu, X. Chen, Y. Yang, and Y. Liu, "Feature-driven economic improvement for network-constrained unit commitment: A closed-loop predict-and-optimize framework," *IEEE Transactions on Power Systems*, pp. 1–1, 2021.
- [30] G. Li and H.-D. Chiang, "Toward cost-oriented forecasting of wind power generation," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 2508–2517, 2016.
- [31] J. Zhang, Y. Wang, and G. Hug, "Cost-oriented load forecasting," *arXiv preprint arXiv:2107.01861*, 2021.
- [32] A. Stratigakos, A. Michiorri, and G. Kariniotakis, "A value-oriented price forecasting approach to optimize trading of renewable generation," in *14th IEEE PowerTech*, 2021.
- [33] M. Munoz, J. M. Morales, and S. Pineda, "Feature-driven improvement of renewable energy forecasting and trading," *IEEE Transactions on Power Systems*, vol. 35, no. 5, pp. 3753–3763, 2020.
- [34] T. Carriere and G. Kariniotakis, "An integrated approach for value-oriented energy forecasting and data-driven decision-making application to renewable energy trading," *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6933–6944, 2019.
- [35] G. Henri and N. Lu, "A supervised machine learning approach to control energy storage devices," *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 5910–5919, 2019.
- [36] A. Georghiou, D. Kuhn, and W. Wiesemann, "The decision rule approach to optimization under uncertainty: methodology and applications," *Computational Management Science*, vol. 16, no. 4, pp. 545–576, 2019.
- [37] A. Stratigakos, A. Michiorri, and G. Kariniotakis, "A Value-Oriented Price Forecasting Approach to Optimize Trading of Renewable Generation," in *2021 IEEE Madrid PowerTech*. Madrid, Spain: IEEE, Jun. 2021. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03208575>
- [38] L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen, *Classification and regression trees*. CRC press, 1984.
- [39] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [40] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Machine learning*, vol. 63, no. 1, pp. 3–42, 2006.
- [41] G. Louppe, L. Wehenkel, A. Suter, and P. Geurts, "Understanding variable importances in forests of randomized trees," *Advances in neural information processing systems* 26, 2013.
- [42] K. Bellinguer, V. Mahler, S. Camal, and G. Kariniotakis, "Probabilistic forecasting of regional wind power generation for the eem20 competition: a physics-oriented machine learning approach," in *2020 17th International Conference on the European Energy Market (EEM)*. IEEE, 2020, pp. 1–6.
- [43] P. Pinson, H. Madsen, H. A. Nielsen, G. Papaefthymiou, and B. Klöckl, "From probabilistic forecasts to statistical scenarios of short-term wind power production," *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, vol. 12, no. 1, pp. 51–62, 2009.
- [44] T. Jónsson, P. Pinson, H. A. Nielsen, and H. Madsen, "Exponential smoothing approaches for prediction in real-time electricity markets," *Energies*, vol. 7, no. 6, pp. 3710–3732, 2014.
- [45] ENTSO-E. Transparency platform. [Online]. Available: <https://transparency.entsoe.eu/>