

Feature-Driven Strategies for Trading Wind Power and Hydrogen

Emil Helgren, Jalal Kazempour, and Lesia Mitridati

Department of Wind and Energy Systems, Technical University of Denmark, Kgs. Lyngby, Denmark

helgrenemil@gmail.com, {jalal, lemitri}@dtu.dk

Abstract—This paper develops a feature-driven model for hybrid power plants, enabling them to exploit available contextual information such as historical forecasts of wind power, and make optimal wind power and hydrogen trading decisions in the day-ahead stage. For that, we develop different variations of feature-driven linear policies, including a variation where policies depend on price domains, resulting in a price-quantity bidding curve. In addition, we propose a real-time adjustment strategy for hydrogen production. Our numerical results show that the final profit obtained from our proposed feature-driven trading mechanism in the day-ahead stage together with the real-time adjustment strategy is very close to that in an ideal benchmark with perfect information.

Index Terms—Hybrid power plants, trading decisions, feature-driven model, linear policies, real-time adjustment strategy.

I. INTRODUCTION

Hybrid power plants, comprising an electrolyzer and wind turbines, are expected to be largely installed in many countries, such as Denmark [1]. The plant operator, aiming to maximize their profit, can either allocate the entire generated wind power for hydrogen production, sell the whole or part of generated wind power to the grid, or buy power from the grid to produce more hydrogen. The plant operator should make informed decisions for trading power and hydrogen in a forward stage, when neither day-ahead and balancing prices nor the true wind power generation are realized. This requires developing a trading strategy, which is the focus of this paper.

There are many papers in the literature that develop trading strategies for wind power producers (wind power as the sole product) under price and wind uncertainty using stochastic optimization methods. Among others, [2] solves a newsvendor problem, [3] develops a two-stage stochastic programming, [4] devises a range of robust optimization models, and finally [5] develops a distributionally robust model. In addition, various data-driven approaches have been proposed that take advantage of contextual information to train a probabilistic forecast model and use its outcomes to solve a stochastic optimization problem, as summarized in [6]. These approaches include sequential training and optimization methods, in which the probabilistic forecast model is trained independently of the optimization task [7], and integrated training and optimization methods, in which the forecast model is trained in order to minimize the losses in a specific optimization task [8]. While all these methods outperform the deterministic counterpart, they require (i) knowledge of probabilistic forecasting, (ii) generating a set of scenarios, an uncertainty set, or a family of

probability distributions, and then (iii) developing a stochastic optimization model to be solved in the decision-making stage. The trading problem is even more challenging for hybrid power plants, where hydrogen is also a product in addition to wind power. Although one may expect large companies to have great expertise in making trading decisions under uncertainty, it might be more challenging for a small-scale hybrid power plant operator. Therefore, the aim of this paper is to develop a *pragmatic* trading approach, which outperforms the deterministic model by enabling the hybrid power plant to learn from historical data, without the need for the plant to use probabilistic forecasting and complex stochastic solutions.

The first contribution of this paper is to propose a novel application of a prescriptive analytics framework based on decision rules, inspired by [9], to the multi-market bidding problem of hybrid power plants. In this data-driven approach, we exploit contextual information, so-called *features*, such as historical (deterministic) forecasts of wind power and directly map them to an optimal action. In the *training stage*, the hybrid power plant operator learns feature-driven *policies* based on historical features and uncertainty realizations. In the *decision-making stage*, the learned policies are applied to new available features, leading to trading decisions without the need to solve a complex optimization problem. A few papers in the literature, e.g., [10], [11] and [12], develop feature-driven trading models for renewable power producers. To the best of our knowledge, this is the first paper that develops such a model for hybrid power plants trading both wind power and hydrogen, which is a more complicated problem. The second contribution of this paper is to extend these previous works by investigating various model architectures, in particular price- and time-dependent policies, and feature vectors, including an improved forecast feature vector. Finally, the third contribution of this paper is to develop a pragmatic rule-based adjustment strategy for real-time hydrogen production. Using an out-of-sample simulation, we show how the resulting profit from feature-driven trading in the day-ahead stage and the adjustment strategy in real time is very close to that in an ideal benchmark (oracle).

The rest of the paper is organized as follows. Section II explains the decision-making framework. Section II details the feature-driven trading strategy model in the day-ahead stage. Section IV presents the proposed real-time adjustment strategy. Section V provides numerical results. Finally, Section VI concludes the paper.

II. DECISION-MAKING FRAMEWORK

We consider a hybrid power plant comprising an electrolyzer and a wind turbine, behind the meter. The wind power generation can be exported to the grid or directed to the electrolyzer to produce hydrogen with a constant efficiency coefficient¹. The electrolyzer has also the possibility to buy power directly from the grid². The produced hydrogen is sold through a bilateral contract at a fixed hydrogen price. As part of the contract, there is a minimum amount of hydrogen that should be produced in a daily basis. In practice, this minimum daily requirement is fulfilled by filling tube trailers that are collected daily by the hydrogen off-taker. However, the proposed models could be generalized straightforwardly to other off-take structures, such as injection of hydrogen in gas pipelines. The hydrogen production thus needs to be scheduled for the whole day in order to ensure fulfilling this minimum production requirement. Any real-time adjustment should likewise account for the imposed hydrogen production quota. The power production capacity of the wind farm is assumed to be identical to the power consumption capacity of the electrolyzer.

The hybrid power plant aims to maximize the total profit from electricity trading and hydrogen sales. This leads to two stages for decision making:

Day-ahead decision making: Given the available (deterministic) wind power forecast, the hybrid power plant must decide how much electricity to trade (buy or sell) in the day-ahead market and how to schedule the operation of the electrolyzer for each hour of the following day, depending on hourly day-ahead market prices and the fixed hydrogen price. Recall the minimum daily hydrogen production requirement needs to be accounted for. In Nord Pool, all these decisions for every hour $t \in \mathcal{T}$ of day D should be made before noon of day D-1.

Real-time decision making: The hybrid power plant must settle imbalances arising from deviations in wind power production (compared to the day-ahead schedule) in the balancing market. Therefore, at the time of delivery, i.e., in every hour t of day D, the plant must decide how to adjust the schedule of the electrolyzer based on the realized wind power production and electricity balancing prices, while ensuring the minimum daily hydrogen production requirement is fulfilled.

III. DAY-AHEAD BIDDING

We first define a feature-driven model for making power and hydrogen trading decisions in the day-ahead stage. Next, we explain feature vectors, and then elaborate on architecture of linear policies. We eventually provide an optimization model for training the proposed feature-driven model.

¹The efficiency of electrolyzers in practice is variant, depending on their power consumption level following a non-linear curve. We refer the interested reader to [13] for alternative convexified solutions.

²According to the ‘‘Renewable Fuels of Non-Biological Origin (RFNBO)’’ Delegated Act of the European Commission published in 2023 [14], the produced hydrogen in this case will not be seen green except for certain operational conditions. We do not consider this legislation in our paper.

A. Model Definition

The main challenge for the hybrid power plant in the day-ahead stage is to account for uncertainty on day-ahead and balancing electricity prices and wind power production.

Aiming to develop a feature-driven trading strategy in the day-ahead stage, we first define the vector of available features with N elements corresponding to hour t as

$$\mathbf{X}_t = [X_t^1, \dots, X_t^N] \quad \forall t \in \mathcal{T}. \quad (1)$$

The features should be those contextual information for hour t of the following day that the plant has access to when making day-ahead decisions. Examples of available features for the hybrid power plant at the day-stage stage are (i) deterministic forecasts of wind power generation of the underlying farm at different points before the day-ahead market gate closure at the noon of day D-1, (ii) aggregate wind power forecast in the area (or bidding zone) where the farm is located, (iii) aggregate wind power forecast in the neighboring areas, (iv) realized wind power production in day D-1.

Let variable p_t^{DA} denote the electricity traded (bought or sold) in the day-ahead market for hour t of the next day. In addition, variable p_t^{H} is the electricity consumption schedule of the electrolyzer for the same hour. We exploit linear decision rules [15], so called *linear policies*, to define p_t^{DA} and p_t^{H} as a function of features \mathbf{X}_t , such that

$$p_t^{\text{DA}} = \mathbf{q}^{\text{DA}} \mathbf{X}_t^\top \quad \forall t \in \mathcal{T} \quad (2a)$$

$$p_t^{\text{H}} = \mathbf{q}^{\text{H}} \mathbf{X}_t^\top \quad \forall t \in \mathcal{T}, \quad (2b)$$

where $(\cdot)^\top$ is the transpose operator. Vectors $\mathbf{q}^{\text{DA}} \in \mathbb{R}^N$ and $\mathbf{q}^{\text{H}} \in \mathbb{R}^N$ are the policies, so called *q-policies*, to be learned from historical data in the training stage.

If the model is designed and trained appropriately, \mathbf{q}^{DA} and \mathbf{q}^{H} will reflect relations between the input features and the decision variables p_t^{DA} and p_t^{H} that persist into the future, such that the model will produce well-performing trading decisions for new feature vectors.

Feature Vectors: The day-ahead market allows for simultaneously placing multiple independent price-quantity bids $\{\lambda_b, p_b\}_{b=1, \dots, B}$ for each hour of the following day. Therefore, we can model the electricity traded in the day-ahead market p_t^{DA} as a function of the day-ahead price at hour t . This is achieved by introducing an unknown price variable $\lambda \in \mathbb{R}$ in the feature vector \mathbf{X}_t . Since all other values in the feature vector are known at the time of decisions, this results in p_t^{DA} being a linear function of this unknown variable λ . In practice, the day-ahead price-quantity bids are then derived by discretizing this linear function and approximating it as a stepwise function with arbitrary step size, as illustrated in Fig. 1. Each step b of the resulting stepwise function will represent one pair of price-quantity bids $\{\lambda_b, p_b\}_{b=1, \dots, B}$. Since the electricity traded in the day-ahead market and the power consumption p_t^{H} of the electrolyzer are interdependent, p_t^{H} is also expressed as a linear function of the day-ahead electricity price at hour t .

All feature vectors investigated in this paper will have the form $\mathbf{X}_t = [\tilde{\mathbf{X}}_t, \lambda, 1]$, where the constant feature 1 provides an

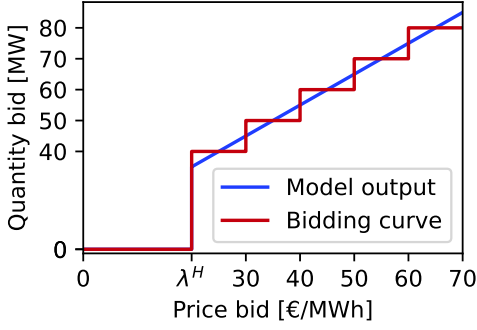


Fig. 1: The discretization of $p_t^{\text{DA}} = \mathbf{q}^{\text{DA}} \mathbf{X}_t^\top$ to create piece-wise price-quantity bids $\{\lambda_b, p_b\}_{b=1, \dots, B}$.

intercept and $\tilde{\mathbf{X}}_t \in \mathbb{R}^{N-2}$ is the vector of remaining features. The remaining features can include any relevant data that are expected to be related with the uncertain parameters including the wind power production (e.g., wind power forecasts) and the day-ahead and balancing prices (e.g., electricity price and demand forecasts and imbalance forecasts). Note that the plant has usually access to an extensive database, including *historical* values of these features and corresponding *realized* values of the uncertain parameters.

Architecture of Linear Policies: Several approaches can be used regarding the architecture of the \mathbf{q} -policies. The simplest but restrictive architecture, so-called *general architecture*, considers *constant* policies $\mathbf{q}^{\text{DA}} \in \mathbb{R}^N$ and $\mathbf{q}^{\text{H}} \in \mathbb{R}^N$ for every hour t .

It is widely acknowledged that daily patterns exist in the day-ahead prices. Accounting for these patterns could remarkably increase the performance of the model. This can be achieved by introducing an *hourly architecture*, where the \mathbf{q} -policies in (2) are specific to the hour of the day and denoted as $\mathbf{q}_{h_t}^{\text{DA}} \in \mathbb{R}^N$ and $\mathbf{q}_{h_t}^{\text{H}} \in \mathbb{R}^N$ for $h_t = t \bmod 24$, with \bmod representing the modulo operator. This hourly architecture allows the model to learn daily patterns. However, it results in 24 times as many \mathbf{q} -policies to train as in the general architecture, and thus requires a larger training dataset. The trade-off between these factors will be empirically studied later in our case study.

Besides, in order to capture more complex and non-linear dependencies between the features \mathbf{X}_t and decision variables p_t^{DA} and p_t^{H} , multiple sets of model parameters can be introduced, each covering one domain of the features. In particular, we introduce M price domains $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_M\}$, each defined as a range of prices $\mathcal{D}_i = [\lambda_{i-1}, \lambda_i]$ with $\lambda_{i-1} < \lambda_i$ for all $i \in \{1, \dots, M\}$. Over each price domain, we define \mathbf{q} -policies as \mathbf{q}_i^{DA} and \mathbf{q}_i^{H} . As a result, the decision variables p_t^{DA} and p_t^{H} are expressed as piece-wise linear functions of the day-ahead prices, such that

$$p_t^{\text{DA}} = \mathbf{q}_i^{\text{DA}} \mathbf{X}_t^\top \quad \text{if } \lambda \in \mathcal{D}_i, \quad \forall t \in \mathcal{T} \quad (3a)$$

$$p_t^{\text{H}} = \mathbf{q}_i^{\text{H}} \mathbf{X}_t^\top \quad \text{if } \lambda \in \mathcal{D}_i, \quad \forall t \in \mathcal{T}. \quad (3b)$$

Note that a trade-off between complexity and flexibility can be achieved, if optimal domains can be chosen based on

statistical characteristics of the data. For example, a natural threshold where it might be beneficial to divide the model into different trading strategies is when the day-ahead electricity price is equal to the hydrogen price λ^{H} .

B. Optimization Model for Training Stage

In the training stage we determine the adequate values of the \mathbf{q} -policies using a batch learning mechanism. The training of the \mathbf{q} -policies can be represented as an optimization problem, in which the profit of the plant is maximized over a *historical* dataset, containing values of the feature vectors along with corresponding realizations of the uncertain parameters, e.g., market prices, for all $t \in \mathcal{T}^{\text{hist}}$. This optimization problem, formulated as a mixed-integer linear program (MILP), reads as

$$\max_{\Omega} \sum_{t \in \mathcal{T}^{\text{hist}}} \left[\lambda_t^{\text{DA}} \mathbf{q}^{\text{DA}} \mathbf{X}_t^\top + \lambda^{\text{H}} \rho^{\text{H}} \mathbf{q}^{\text{H}} \mathbf{X}_t^\top + \lambda_t^{\text{UP}} o_t - \lambda_t^{\text{DW}} u_t \right] \quad (4a)$$

$$\text{s.t. } o_t - u_t = P_t^{\text{W}} - \mathbf{q}^{\text{DA}} \mathbf{X}_t^\top - \mathbf{q}^{\text{H}} \mathbf{X}_t^\top \quad \forall t \in \mathcal{T}^{\text{hist}} \quad (4b)$$

$$0 \leq o_t \leq M(1 - b_t) \quad \forall t \in \mathcal{T}^{\text{hist}} \quad (4c)$$

$$0 \leq u_t \leq M b_t \quad \forall t \in \mathcal{T}^{\text{hist}} \quad (4d)$$

$$0 \leq \mathbf{q}^{\text{H}} \mathbf{X}_t^\top \leq \bar{P}^{\text{H}} \quad \forall t \in \mathcal{T}^{\text{hist}} \quad (4e)$$

$$-\bar{P}^{\text{H}} \leq \mathbf{q}^{\text{DA}} \mathbf{X}_t^\top \leq \bar{P}^{\text{W}} \quad \forall t \in \mathcal{T}^{\text{hist}} \quad (4f)$$

$$\sum_{t=24(d-1)+1}^{24d} \rho^{\text{H}} \mathbf{q}^{\text{H}} \mathbf{X}_t^\top \geq \underline{H} \quad \forall d \in D^{\text{hist}}, \quad (4g)$$

where the set of variables $\Omega = \{\mathbf{q}^{\text{DA}}, \mathbf{q}^{\text{H}}, o_t, u_t, b_t\} \forall t \in \mathcal{T}^{\text{hist}}$ includes the \mathbf{q} -policies \mathbf{q}^{DA} and \mathbf{q}^{H} , as well as the real-time decisions, including the over-production o_t and under-production u_t settled in real time, and the auxiliary binary variable $b_t \in \{0, 1\}$ indicating the state of over- or under-production. Note that the \mathbf{q} -policies in (4) can take index h_t and/or index i as discussed in Section III-A

The objective function (4a) maximizes the total profit of the hybrid power plant over the historical dataset, including four terms. The first term $\lambda_t^{\text{DA}} \mathbf{q}^{\text{DA}} \mathbf{X}_t^\top$ corresponds to the revenue/cost from electricity trading in the day-ahead market, where $\lambda_t^{\text{DA}} \in \mathbb{R}$ is the historical day-ahead price reported by the market operator. Note that $\mathbf{q}^{\text{DA}} \mathbf{X}_t^\top$, reflecting the power quantity, could be positive or negative, indicating whether the plant sells or buys power. The second term $\lambda^{\text{H}} \rho^{\text{H}} \mathbf{q}^{\text{H}} \mathbf{X}_t^\top$ is the revenue from hydrogen sales, where $\lambda^{\text{H}} \in \mathbb{R}_+$, in €/Kg, is the fixed hydrogen price. In addition, $\rho^{\text{H}} \in \mathbb{R}_+$, in Kg/MWh, is the constant power-to-hydrogen efficiency of the electrolyzer, whereas $\mathbf{q}^{\text{H}} \mathbf{X}_t^\top$ reflects the power consumed by the electrolyzer. Finally, the third and fourth terms in (4a) refer to the settlement of over- and under-production, respectively, in the balancing market with a dual-price imbalance settlement scheme. Note that $\lambda_t^{\text{UP}} \in \mathbb{R}$ is the upward regulation price, which is lower than or equal to the day-ahead price λ_t^{DA} , reflecting a lost opportunity cost for over-production. In addition, $\lambda_t^{\text{DW}} \in \mathbb{R}$ is the downward regulation price, which is higher than or equal to the day-ahead price λ_t^{DA} , reflecting a

penalty cost for under-production. By this, it is always costly for the hybrid power plant to create any imbalance either over- or under-production.

Constraint (4b) defines the power imbalance $o_t - u_t$ to be settled in real time as the difference between the realized wind power generation P_t^W and the power scheduled in the day-ahead stage. The set of disjunctive constraints (4c)-(4d), with M a large enough positive constant, enforces over- and under-production to not happen simultaneously at each hour t . If $b_t = 1$, then (4c) enforces the over-production o_t to be zero. On the other hand, if $b_t = 0$, then (4d) sets the under-production u_t to zero. Constraint (4e) enforces the power consumption of the electrolyzer to lie within zero and its capacity \overline{P}^H . Note that the lower bound being set to zero means that the electrolyzer cannot operate as a fuel cell, converting the hydrogen back to electricity. In order to avoid imbalance settlement costs, the power traded in the day-ahead market is restricted by (4f) to lie within the minus consumption capacity of the electrolyzer and the nominal capacity of the wind turbine \overline{P}^W . Finally, the minimum daily hydrogen production requirement \underline{H} is enforced by (4g) for each day $d \in D^{\text{hist}}$ in the historical dataset. This requirement could be alternatively implemented using a slack variable penalized in the objective function. However, a value for such a penalty is unknown. One may assume the penalty would always be higher than the revenue that could be generated by not meeting the requirement.

C. Decision-Making Stage and Retraining

For each day of the testing period, the trained \mathbf{q} -policies are applied to a new feature vector \mathbf{X}_t for each hour t of the following day, to compute the linear functions representing the electricity traded in the day-ahead market and consumed by the electrolyzer as functions of the future day-ahead prices. These functions are then discretized, and stepwise price-quantity bids are placed in the day-ahead market. The realized day-ahead price will then determine the electricity traded in the day-ahead market and consumed by the electrolyzer. For the sake of simplicity, we assume none of the technical constraints of the hybrid power plant is violated in the testing period. This is in practice likely as long as the testing data does not differ significantly from the training one. If it is the case, one should restore feasibility of those constraints by adjusting trading decisions at the minimum cost, or projecting the decisions onto the feasible space.

Once the uncertain parameters, i.e., day-ahead prices, balancing prices, and wind power production, are realized, the new feature vector and the associated realizations are added to the historical dataset. Retraining the model considering more recent historical data points is thus performed at regular intervals. One can use a sliding window approach in order to keep the length of the training dataset constant, as illustrated in Fig. 2. Note that, at each retraining interval, the blue part of the window represents the previous testing data points recently added to the training dataset. The length of this window




Fig. 2: Illustration of sliding window approach to determine training data in each retraining interval.

is a hyperparameter of the model that should be selected appropriately, as further discussed in our numerical analysis.

IV. REAL-TIME ADJUSTMENTS

We assume the electrolyzer is subject to negligible ramping constraints. Therefore, the hybrid power plant is able to adjust the hydrogen production schedule in real time based on the realized wind power production and balancing prices. This section introduces a rule-based adjustment algorithm for the real-time operation of the electrolyzer that accounts for the minimum daily hydrogen production requirement.

A. Decision Rule for a Single Hour

In each hour, the decision rule is dependent on whether balancing prices are higher or lower than the hydrogen price. Hereafter, to make electricity and hydrogen prices comparable in terms of €/MWh, we refer to $\rho^H \lambda^H$ as the hydrogen price. Balancing prices are unknown until after the balancing settlements, however it is assumed that they can be approximated to a sufficient degree for the current hour by predicting the system status [16] and observing the intraday market. This will be explained in further detail after the introduction of the algorithm and required parameters.

As mentioned earlier, we assume that the balancing market follows a dual-pricing mechanism, implying that $\lambda_t^{\text{UP}} \leq \lambda_t^{\text{DA}} \leq \lambda_t^{\text{DW}}$ for each hour $t \in \mathcal{T}$. The following three situations can thus occur regarding the relation between the hydrogen price and balancing prices: (i) $\rho^H \lambda^H \leq \lambda_t^{\text{UP}} \leq \lambda_t^{\text{DW}}$; (ii) $\lambda_t^{\text{UP}} \leq \lambda_t^{\text{DW}} \leq \rho^H \lambda^H$; and (iii) $\lambda_t^{\text{UP}} \leq \rho^H \lambda^H \leq \lambda_t^{\text{DW}}$. In the first situation where both up and down balancing prices are higher than the hydrogen price, producing hydrogen would result in less revenue than what is compensated for over-production in the balancing market. This means that the producer would benefit from adjusting the electrolyzer's consumption down and exporting as much electricity to the grid as possible. In the second situation where both balancing prices are lower than the hydrogen price, producing hydrogen results in higher revenue than the cost incurred from under-production in the balancing market. This means that the hybrid power plant would benefit from adjusting the electrolyzer's consumption up and importing as much electricity from the grid as possible. In the final situation, over-production results in lower revenues than hydrogen production, and under-production results in higher costs than revenues from hydrogen production. In this situation, the electrolyzer

should be adjusted in order to minimize any deviation from the day-ahead market bid. This decision rule that maximizes the profit for a single hour can be formulated as a piece-wise linear function $\Pi(\cdot)$ that returns the adjusted electricity consumption level of the electrolyzer, such that

$$\Pi(\delta_t, \lambda_t^H, \lambda_t^{\text{UP}}, \lambda_t^{\text{DW}}) = \begin{cases} 0 & \text{if } \lambda_t^{\text{UP}} > \rho^H \lambda^H \\ \bar{P}^H & \text{if } \lambda_t^{\text{DW}} < \rho^H \lambda^H \\ \alpha^H[\delta_t] & \text{otherwise,} \end{cases} \quad (5)$$

where $\delta_t = P_t^{\text{W}} - p_t^{\text{DA}}$ is the surplus/deficit of wind power production compared to the day-ahead market bid, and $\alpha^H[\cdot]$ ensures that the electrolyzer production is feasible:

$$\alpha^H[\delta_t] = \begin{cases} 0 & \text{if } \delta_t < 0 \\ \bar{P}^H & \text{if } \delta_t > \bar{P}^H \\ \delta_t & \text{otherwise.} \end{cases} \quad (6)$$

Adjusting the electricity consumption of the electrolyzer according to the decision rule (5) will maximize the generated revenues for a single hour, when the estimated balancing price is the only source of uncertainty. Note that (5) does not require an accurate estimation of balancing prices, but only whether balancing prices will be higher or lower than the hydrogen price, which is a more straightforward task. Whether the balancing price will be higher or lower than the known day-ahead price will depend on the system status (surplus or deficit energy in the system).

B. Real-time Adjustment Algorithm

Implementing the hourly decision rule (5) independently for each single hour of the optimization period would not guarantee that the minimum daily hydrogen production requirement \underline{H} is satisfied. Instead, we introduce an algorithm which only adjusts the electrolyzer's schedule according to the hourly decision rule in a given hour t if this requirement can still be met during the day. Therefore, the electrolyzer's consumption can always be adjusted upward from the day-ahead schedule at any given hour. However, in order to perform a downward adjustment in a given hour t of the day d , the adjusted hydrogen production in this hour $\rho^H p_t^{\text{adj}}$, plus the cumulative realized production in previous hours $\rho^H p_t^{\text{real}} = \sum_{i=24(d-1)+1}^{t-1} \rho^H p_i^{\text{adj}}$, plus the cumulative production schedule for the remaining hours of the day $\sum_{i=t+1}^{24d} \rho^H p_i^H$ must be higher than the minimum daily hydrogen production requirement \underline{H} . For example, let us consider a hybrid power plant with a certain minimum daily hydrogen production requirement, equivalent to the consumption of 15 MWh. If the realized cumulative hydrogen production at hour 15 is equal to 10 MWh and the cumulative planned hydrogen production from hour 15 to 23 is equal to 7 MWh, then the electrolyzer has a scheduled hydrogen surplus of 2 MWh, which can be adjusted down at hour 15. If the hourly decision rule (5) outputs a lower set-point for the electrolyzer than its schedule for this hour, the electrolyzer can then be turned down as much as the surplus allows. On the contrary, if the realized cumulative hydrogen production is only equal to 8 MWh, then the electrolyzer

Algorithm 1 Real-time upward and downward adjustment

```

1: for each day  $d$  do
2:   Initialize  $p_{24(d-1)+1}^{\text{real}} = 0$ 
3:   for each hour  $t \in [24(d-1)+1, 24d]$  do
4:     Receive model output:  $p_t^{\text{DA}}, \mathbf{p}^H$ 
5:     Receive balancing price forecasts:  $\lambda_t^{\text{UP}}, \lambda_t^{\text{DW}}$ 
6:     Realize production:  $P_t^{\text{W}}$ 
7:     Compute real-time deviation  $\delta_t = P_t^{\text{W}} - p_t^{\text{DA}}$ 
8:     Compute  $p_t^{\text{H}*} = \Pi(\delta_t, \lambda_t^{\text{UP}}, \lambda_t^{\text{DW}})$ 
9:     if  $p_t^{\text{H}*} > p_t^H$  then
10:        $p_t^{\text{adj}} = p_t^{\text{H}*}$ 
11:     else
12:        $p_t^{\text{adj}} = \max\left(p_t^{\text{H}*}, \frac{\underline{H}}{\rho^H} - \left(p_t^{\text{real}} + \sum_{i=t+1}^{24d} p_i^H\right)\right)$ 
13:     end if
14:     Receive revenue from allocations  $(p_t^{\text{DA}}, p_t^{\text{adj}})$ 
15:      $p_{t+1}^{\text{real}} = p_t^{\text{real}} + p_t^{\text{adj}}$ 
16:   end for
17: end for

```

has no scheduled hydrogen surplus and cannot be adjusted down in this hour, regardless of the outputs of the hourly decision rule (5). The size of the production requirement compared to the production capacity of both the wind farm and the electrolyzer thus directly affects the possible space of adjustments, with a higher requirement reducing the amount of adjustment possible. The complete adjustment algorithm is provided in Algorithm 1.

V. NUMERICAL ANALYSIS

This section provides a numerical analysis of the proposed feature-driven day-ahead trading strategy and rule-based real-time adjustment. All source codes are available in [17].

A. Case Study Data

We use historical data for the year 2019 for training, and the year 2020 for testing, ensuring that the models are evaluated through all seasonal variations. It is essential to evaluate the models with out-of-sample testing in this way, using more recent data for testing than was used for training, to ensure the performance of the models is due to them generalizing observed patterns in the features and not overfitting the training data. All price data is publicly available on ENTSO-e [18], along with production data for specific wind farms, and aggregated forecasts of offshore/onshore wind production in the Danish bidding zones DK1 and DK2. The production from the wind farm in Roedsand, Denmark is used to model the realized wind production of the hybrid power plant.

Realistic forecasts of day-ahead price and wind power production have been provided by Siemens Gamesa Renewable Energy for five consecutive months. Two years of synthetic forecasts are created for both price and wind production, by fitting the statistical characteristics of the forecast error of Siemens Gamesa forecasts to a theoretical distribution, and then sampling from this distribution to generate new

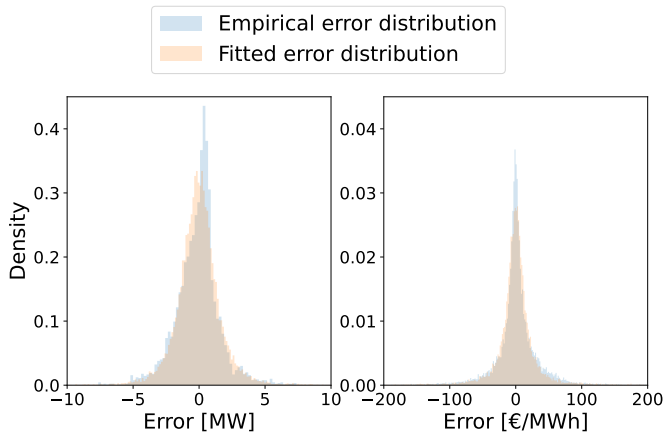


Fig. 3: Distribution of day-ahead market price (right) and wind power generation (left) forecast errors.

forecast errors with the same characteristics. Fig. 3 provides a histogram of both original and generated forecast errors for both day-ahead market price and wind production forecasts.

Forecasts are produced by sampling forecast error values from the fitted distributions, and then adding these errors to the realized production data to create a forecast. The entire training and testing datasets thus consist of realized day-ahead prices, upward and downward balancing prices, and production for the wind farm in Roedsand, and forecasts for day-ahead prices and wind production, for 2019 and 2020.

B. Models to be Compared

This numerical analysis compares several trained models with various combinations of (i) linear policies architecture as introduced in Section III-A and (ii) feature vectors. These trained models are compared to a deterministic (**Det.**) optimization approach as a base-case, and a benchmark with perfect information (**Hindsight**). In addition, we explore the impact of the length of the training period, ranging between 1 and 12 months, on the outcomes of different models.

We first compare models with **General Architecture (GA)** and **Hourly Architecture (HA)** for \mathbf{q} -policies. Recall that \mathbf{q} -policies take index t in the **HA** model, which is not the case in the **GA** model. Next, we compare the **GA** and **HA** in the cases where a single set of policies is defined over all values of prices, and a case where the policies are split into multiple **Price Domains (PD)**. When considering multiple price domains, we refer to models as **GA+PD** and **HA+PD**. Three price domains are chosen: one below the hydrogen price $\rho^H \lambda^H$, one above the 90% percentile of realized day-ahead prices in the training data, and one in between. The hydrogen price is a natural threshold to separate the trading decisions because producing hydrogen is necessarily more profitable than selling power for day-ahead prices below the hydrogen price. The 90% percentile is chosen to reduce the influence of extreme outliers on the \mathbf{q} -policies.

Additionally, we investigate three different types of feature vectors \mathbf{X}_t . Firstly, the simplest **Reduced Feature (RF)** vec-

tor contains only the deterministic wind production forecast \hat{P}_t^W , such that $\mathbf{X}_t^{\text{RF}} = [\hat{P}_t^W, \lambda, 1]$. Secondly, the **Augmented Feature (AF)** vector includes additional features, namely the aggregated onshore/offshore wind production forecasts in the two bidding zones DK1 and DK2, released by the Danish transmission system operator, Energinet, before the time of bidding, such that $\mathbf{X}_t^{\text{AF}} = [\tilde{\mathbf{X}}_t^{\text{AF}}, \lambda, 1]$ with $\tilde{\mathbf{X}}_t^{\text{AF}} = [\hat{P}_t^W, \hat{P}_t^{\text{on,DK1}}, \hat{P}_t^{\text{on,DK2}}, \hat{P}_t^{\text{off,DK1}}, \hat{P}_t^{\text{off,DK2}}]$. This augmented feature vector may capture additional information related to the uncertainty sources at the cost of training additional \mathbf{q} -policies and requiring a larger amount of training data. Increasing the length of the historical dataset raises issues of stationarity of the environment, in particular of electricity prices. Therefore, a so-called **Forecast Model (FM)** feature vector is introduced, that accounts for the same additional features $\tilde{\mathbf{X}}_t^{\text{AF}}$, while minimizing the added complexity of the model. This is achieved by first training a separate feature-driven forecast model that provides a single improved wind power generation forecast feature $\tilde{X}_t^{\text{FM}} = \tilde{\mathbf{q}}^{\text{FM}} [\tilde{\mathbf{X}}_t^{\text{AF}}, 1]^\top \in \mathbb{R}$, where $\tilde{\mathbf{q}}^{\text{FM}}$ represents a vector of policies that is trained using a historical dataset. As the environment related to wind power generation is expected to be stationary, to a large extent, this feature-driven forecast model can be trained on a longer historical dataset. Then, the **FM** feature vector is defined as $\mathbf{X}_t^{\text{FM}} = [\tilde{X}_t^{\text{FM}}, \lambda, 1]$.

C. Results

The results presented are the profits generated by each model from electricity trades in the day-ahead market, hydrogen production, and settlements in real time, evaluated ex-post for the entire year 2020.

Fig. 4a illustrates the comparison between the different policy architectures, showing the best performing feature vector and training period for each architecture. This plot shows that introducing multiple price domains increases the performance radically for both types of architectures, resulting in the **HA+PD** model outperforming the **GA+PD** model with price domains included. Additionally, we note that the training period yielding the best results for the **GA** model is 5 months, and the model performed within 99% of this performance after only 3 months. With a training period of 3 months, the **GA** model generates around 5% higher revenues than the **HA** model. This indicates that for time periods with a more non-stationary environment than the one used in the study, or with limited available historical data, the **GA** model might be a more appealing choice. Since Fig. 4a shows that **HA+PD** is the model that outperforms others, we further analyze this model in the rest of this section.

Fig. 4b depicts the performances of **HA+PD** for different feature vectors and varying lengths of the training dataset, from one to 12 months. First, we observe that the improvement of the model by having access to a training dataset longer than 7 months is limited, irrespective of the feature vectors chosen. For example, the profit when models trained on a 12-month dataset is only 1-2% higher than when trained on a 7-month dataset. The second observation is that the model




Fig. 4: Ex-post profit of the hybrid power plant with (a) four policy architectures (**GA**, **HA**, **GA+PD**, **HA+PD**); (b) the **HA+PD** model for various training dataset lengths and available feature vectors (**AF**, **RF**, **FM**), compared to the **Deterministic** and **Hindsight** models; and (c) the **HA+PD** model before and after the real-time adjustment, compared to the deterministic (**Det.**) and **Hindsight** model, and the **Optimal Adjustment** strategy.

with the **FM** feature vector setting usually exhibits a better performance than the ones with the **AF** and **RF** feature vector settings, especially when the length of the training dataset is lower than 12 months. However, these three models show a similar performance when the plant uses a dataset of 12 months. The third observation is the performance of our models in comparison to the deterministic model, where the hybrid power plant makes day-ahead bidding decisions by using deterministic forecasts only. We notice that, by using a training dataset of at least 4 months, our model outperforms the deterministic model. For example, when using the 12-month dataset, the plant makes around 5.9% more profit by using the **HA+PD** model with either of three feature vector settings. The final observation corresponds to the comparison of the **HA+PD** model with respect to the hindsight model, in which the plant has perfect information on the prices and wind generation. We observe that the profit earned by our proposed **HA+PD** model is only 3.7% lower than the profit in hindsight, whereas it is 9.8% for the deterministic model. For the rest of this section, we further analyze the best-performing model, i.e., the **HA+PD** model with a 12-month training dataset and (**AF**) feature vector setting.

Fig. 4c depicts the profit across various models before and after real-time adjustments. Obviously, the hindsight model does not need any real-time adjustment. The *final* profit, i.e., after in the day-ahead stage plus real time, earned by the deterministic (**Det.**) and the proposed **HA+PD** model is computed using the rule-based adjustment strategy proposed in Section IV-B. This profit is compared to the final profit obtained by the **HA+PD** model with an *optimal adjustment*, implemented as an optimization problem over 24 hours with perfect foresight, which provides an upper bound on the benefits of any adjustment policy. We observe that the the proposed real-time adjustment strategy achieves a higher increase in the profit of the deterministic model compared to that of the **HA+PD**

model. Indeed, due to a less efficient day-ahead scheduling, the deterministic model has more potential for improvement in the real-time stage than the **HA+PD** model. However, the final profit of the deterministic model is still slightly lower than that of the proposed **HA+PD** model. Yet, it is important to note that its result relies on having access to accurate upward and downward balancing price forecasts for the the real-time adjustment strategy, which is not a straightforward task. As the final profits of the deterministic model are highly dependent on the performance of the real-time adjustment policy, in the absence of accurate price forecasts, the deterministic model risks achieving significantly lower final profits than the **HA+PD** model. Additionally, we observe that the proposed rule-based adjustment strategy achieves almost as much profit as the optimal adjustment one, under the assumption of accurate price forecasts. This shows that, owed to the proposed real-time adjustment strategy, the final profit of the **HA+PD** model gets even closer to the profit in hindsight, making it an even more attractive model.

Finally, Fig. (5) provides two examples of feature-driven bidding strategies in the day-ahead market, obtained by the **HA+PD** model. The left plot shows the bidding strategy in a representative hour during which the hybrid power plant sells power to the grid. In this hour, the plant sells the first 10 MW at price $\rho^H \lambda^H$ and the rest at the 90% percentile of the realized day-ahead prices in the training data ($\lambda^{90\%}$). The right plot shows a representative hour during which the hybrid power plant buys power from the grid. The first 10 MW is bought at price $\lambda^{90\%}$ and the rest at price $\rho^H \lambda^H$.

We applied a retraining procedure to all models with a monthly retraining schedule. This retaining procedure is expected to improve the performance of the algorithm over time, for instance, adapting to seasonal wind patterns and changes in electricity prices distribution. However, over the testing period considered, retraining provided minimal improvements, and is




Fig. 5: Example of resulting price-quantity bidding curves in the day-ahead market, when the hybrid power plant buys (**right**) and sells (**left**) electricity. Recall $\rho^H \lambda^H$ is the hydrogen price, and $\lambda^{90\%}$ is the 90% percentile of realized day-ahead prices in the training data.

therefore not considered adequately to draw any conclusions about its impact.

VI. CONCLUSION

This paper explains how hybrid power plants with co-located wind turbines and an electrolyzer can implement efficient feature-driven models to learn from historical data and make informed day-ahead trading decisions. The proposed feature-driven model derive trading (selling or buying) decisions in the day-ahead electricity market, as well as a hydrogen production schedule for the next day, fulfilling the minimum daily hydrogen quota. This is a pragmatic solution, which properly accounts for wind power and price uncertainty without the need to generate probabilistic forecasts or solve complex stochastic optimization problems. Our numerical analysis shows that the proposed feature-driven models outperform the deterministic model, and require less adjustment in real time. In addition, they result in a final profit which is close to that in hindsight.

For future work, it is of interest to explore the performance of the model in a non-stationary environment. This could be tackled by developing online decision-making methods, similar to the one proposed in [19] for wind power trading only. Another valuable potential of implementing an online decision-making method is an expansion of the adjustment model to account for the uncertainty in the prediction of the balancing price and thereby improve the value of real-time adjustment. It is also interesting to include more complex physical characteristics of the electrolyzer, such as non-linear efficiency and degradation curves, and add more relevant assets to the hybrid power plant, such as battery and hydrogen storage.

ACKNOWLEDGEMENT

We would like to thank Siemens Gamesa Renewable Energy for providing us with historical forecast data. We also thank the Danish Energy Technology Development and Demonstration Programme (EUDP) for supporting this research through HOMEY (Grant number: 64021-7010) and ViPES2X (Grant number: 640222-496237) projects.

REFERENCES

- [1] Danish Ministry of Climate, Energy and Utilities, “The government’s strategy for power-to-x,” p. 5, 2021, https://ens.dk/sites/ens.dk/files/ptx/strategy_ptx.pdf.
- [2] 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.
- [3] J. Morales, A. 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.
- [4] A. Sun and A. J. Conejo, *Robust Optimization in Short-Term Power System Operations*. Springer, 2021.
- [5] P. Pinson, “Distributionally robust trading strategies for renewable energy producers,” *IEEE Transactions on Energy Markets, Policy and Regulation*, vol. 1, no. 1, pp. 37–47, 2023.
- [6] U. Sadana, A. Chenreddy, E. Delage, A. Forel, E. Frejinger, and T. Vidal, “A survey of contextual optimization methods for decision making under uncertainty,” *arXiv preprint arXiv:2306.10374*, 2023.
- [7] D. Bertsimas and N. Kallus, “From predictive to prescriptive analytics,” *Management Science*, vol. 66, no. 3, pp. 1025–1044, 2019.
- [8] A. N. Elmachtoub and P. Grigas, “Smart “predict, then optimize,”” *Management Science*, vol. 68, no. 1, pp. 9–26, 2022.
- [9] L. H. Liyanage and J. G. Shanthikumar, “A practical inventory control policy using operational statistics,” *Operations Research Letters*, vol. 33, no. 4, pp. 341–348, 2005.
- [10] M. Muñoz, 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.
- [11] A. Stratigakos, S. Camal, A. Michiorri, and G. Kariniotakis, “Prescriptive trees for integrated forecasting and optimization applied in trading of renewable energy,” *IEEE Transactions on Power Systems*, vol. 37, no. 6, pp. 4696–4708, 2022.
- [12] K. Parginos, R. Bessa, S. Camal, and G. Kariniotakis, “Interpretable data-driven solar power plant trading strategies,” in *2022 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, 2022, pp. 1–5.
- [13] E. Raheli, Y. Werner, and J. Kazempour, “A conic model for electrolyzer scheduling,” *Computers & Chemical Engineering*, vol. 179, p. 108450, 2023.
- [14] European Commission, “Delegated act on renewable fuels of non-biological origin,” 2023, https://energy.ec.europa.eu/system/files/2020-07/hydrogen_strategy_0.pdf.
- [15] 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.
- [16] A. Skajaa, K. Edlund, and J. M. Morales, “Intraday trading of wind energy,” *IEEE Transactions on Power Systems*, vol. 30, no. 6, pp. 3181–3189, 2015.
- [17] “Feature-driven-trading-of-wind-energy-and-hydrogen,” 2023. [Online]. Available: <https://github.com/emilhelgren/feature-driven-trading-of-wind-power-and-hydrogen>
- [18] “ENTSO-e transparency platform,” 2023. [Online]. Available: https://transparency.entsoe.eu/?site_preference=normal
- [19] M. A. Muñoz, P. Pinson, and J. Kazempour, “Online decision making for trading wind energy,” *Computational Management Science*, vol. 20, pp. 1–31, 2023.