Value-oriented price forecasting for arbitrage strategies of Energy Storage Systems through loss function tuning

Smets, Ruben; Toubeau, Jean François; Dolanyi, Mihaly; Bruninx, Kenneth; Delarue, Erik

DOI
10.1016/j.energy.2025.137112

Publication date
2025

Document Version
Final published version

Published in
Energy

Citation (APA)
Smets, R., Toubeau, J. F., Dolanyi, M., Bruninx, K., & Delarue, E. (2025). Value-oriented price forecasting for arbitrage strategies of Energy Storage Systems through loss function tuning. *Energy*, *333*, Article 137112. https://doi.org/10.1016/j.energy.2025.137112

Energy 333 (2025) 137112

Contents lists available at ScienceDirect

Energy

journal homepage: www.elsevier.com/locate/energy

Value-oriented price forecasting for arbitrage strategies of Energy Storage Systems through loss function tuning

Ruben Smets [a,e], Jean-François Toubeau [b], Mihaly Dolanyi [d], Kenneth Bruninx [c], Erik Delarue [a,e,*]

[a] Applied Mechanics & Energy Conversion, KU Leuven, Leuven, Belgium
[b] Power Systems and Markets Research Group, UMONS, Mons, Belgium
[c] Faculty of Technology, Policy and Management, TU Delft, Delft, The Netherlands
[d] Sight-E Analytics, Budapest, Hungary
[e] Energyville, Genk, Belgium

ARTICLE INFO

Keywords:
Implicit balancing
Model Predictive Control
Value-oriented forecasting
Time series forecasting
Recurrent Neural Networks

ABSTRACT

Increasing shares of renewable generation are leading to more volatile electricity prices, presenting an opportunity for Energy Storage Systems (ESS) participating in short-term electricity markets. Model Predictive Control (MPC) has been shown to be a powerful tool to leverage the latest information at the time of optimization, yet its efficacy depends on the quality of the employed price forecasts. So far, these forecasts have been developed with traditional forecasting methods instead of value-oriented approaches, which consider the downstream decision problem during the forecaster training phase. Existing value-oriented methods, however, often rely on a specific downstream problem structure. This paper addresses these shortcomings by introducing a universally applicable, value-oriented forecasting methodology that employs a generalized loss function designed to account for inter-temporal price variability, using the downstream value (i.e., profit from ESS market participation) as the selection criterion in the hyperparameter tuning step. The proposed methodology is tested on a case study considering different types of ESS participating in the Belgian balancing market through MPC. The method is benchmarked against other forecasting techniques including a neural network trained in traditional, accuracy-oriented fashion. Using real-life data over a test set of two months, we show that the methodology outperforms those traditional techniques in terms of ex-post out-of-sample profit.

1. Introduction

In recent years, net-zero emissions policies and climate awareness have driven investments in intermittent Renewable Energy Sources (RES), resulting in more volatile electricity prices in various electricity markets. These trends are strengthening the business case of Energy Storage Systems (ESSs). Indeed, arbitrage opportunities become more abundant and profitable since inter-temporal electricity price spreads tend to increase with growing intermittent RES output. One particular technology that is getting a lot of attention is the Battery Energy Storage System (BESS). Such systems tend to have high round-trip efficiencies [1], and their costs are projected to decrease with over 40% by 2030 compared to 2020 [2]. This has sparked a variety of research in finding the optimal schedule in order to maximize profits in short-term electricity markets.

1.1. ESS market participation

While bidding behavior has been described with heuristic decision rules [3], the scientific literature has widely adopted optimization techniques to model market participation. Broadly, two types of optimization strategies can be discerned. The first concerns day-ahead optimization. Many papers consider this stage due to the prominence of the day-ahead market in many energy systems, as well as the reserve markets being typically organized at that stage. Research often explores either exclusive day-ahead market participation [4] or co-optimization of energy and reserve bids [5–8]. A critical element when participating in reserve markets is the uncertainty of activation. This leads to uncertainty on the State of Charge (SoC), which is subject to box constraints. Assumptions like worst-case considerations [5,6], probabilistic guarantees [7] and only considering expected activations [8] have to be introduced to mitigate this uncertainty. These approaches,

* Corresponding author at: Applied Mechanics & Energy Conversion, KU Leuven, Leuven, Belgium.
E-mail address: erik.delarue@mech.kuleuven.be (E. Delarue).

https://doi.org/10.1016/j.energy.2025.137112
Received 2 July 2024; Received in revised form 15 May 2025; Accepted 12 June 2025
Available online 4 July 2025

Nomenclature

Sets and Indice

$\mathcal{T}_{LA}$	Set of instances in Look-Ahead horizon, indexed τ
L	Set of degradation break levels, indexed l
T	Set of instances in MPC test horizon, indexed t
X_t	Set of train data examples, indexed j
X_v	Set of validation data examples, indexed i

Decision Variables

Γ_τ	Degradation cost at instance τ [€]
θ	Neural network parameter [–]
$\tilde{s}_{\tau,l}$	BESS State of Charge at time step τ associated with degradation level l [MWh]
ξ_a	Hyperparameter related to neural network architecture [–]
ξ_l	Hyperparameter related to loss function [–]
$b_{\tau,l}$	Binary for activation of degradation level l at time step τ [–]
e^+_τ	Energy discharged at instance τ [MWh]
e^-_τ	Energy charged at instance τ [MWh]
s_τ	BESS State of Charge at time step τ [MWh]

Parameters

Δt	Duration of time step [h]
$\eta^{+/-}$	Discharge (+) and charge (−) efficiency [–]
$\hat{\lambda}_\tau(t)$	Forecasted price for lookahead instance τ at optimization time t [€/MWh]
$\lambda_\tau(t)$	Actual price for lookahead instance τ at optimization time t [€/MWh]
$\overline{P}$	Maximum (dis)charge power [MW]
$\overline{SoC}$	BESS maximum energy content [MWh]
$\underline{\text{SoC}}$	BESS minimum energy content [MWh]
C^B	Total BESS investment cost for energy capacity [€]
p^δ_l	BESS degradation break point for break level l [–]
q^δ_l	BESS energy break point for break level l [MWh]

while necessary to address the unpredictability of reserve activations, introduce compromises in terms of feasibility and economic efficiency.

The second method involves Model Predictive Control (MPC). This framework entails optimizing the schedule of the ESS with a rolling horizon, only realizing the first time step with every iteration. This allows the ESS operator to have a precise view of the SoC, as such alleviating some of the uncertainty-related shortcomings of day-ahead optimization. This has been extensively studied in US-style RT markets [9–14], and has recently been suggested for the strategy of explicit [15] and implicit [16] balancing in European markets, as well as trading on the continuous intra-day market [17].

While such optimization methodologies have been implemented considering perfect price foresight, see e.g. [18], employing price forecasts and evaluating the decisions using the actual prices improves the realism of the model. In that case, the model is classified as a predict-then-optimize problem. First, prices are predicted based on contextual information, to subsequently use them as parameters in the optimization program, also known as the "downstream problem". The success of such a predict-then-optimize method crucially depends on the forecaster's ability to accurately represent the uncertain variable, which, in this paper, corresponds to the electricity price.

1.2. Price forecasting

The literature regarding short-term price forecasting can be split into three broad categories, being (i) deep learning, (ii) statistical methods, and (iii) hybrid methods [19]. Whereas statistical methods, see e.g. [20,21], have the advantage of higher interpretability, deep learning methods gained traction since they tend to outperform their statistical counterparts [22], and will thus be the focus of this paper. Within the realm of deep learning, Recurrent Neural Networks (RNNs) are designed to accurately capture temporal dependencies. More specifically, the strength of Long-Short Term Memory (LSTM) unit-based RNNs has been shown for forecasting electricity prices [23–25], the System Imbalance (SI) [16,26] and electrical load [27].

A key component in training neural networks is the choice of the loss function, which guides the learning process, typically optimized using gradient descent. Any differentiable loss function may be chosen, and there exists a plethora of loss functions that have been implemented [28]. Among those, the Mean Squared Error (MSE) and the pinball loss are ubiquitous for deterministic and probabilistic forecasts respectively.

1.2.1. Value-oriented forecasting

Recently, researchers argue that the forecaster should not necessarily maximize accuracy, but rather the downstream value resulting from deploying the forecaster in the downstream decision problem. This concept is known as value-oriented, or decision-focused, forecasting. For the application of ESS participation in energy markets, the downstream value corresponds to the ex-post profit.

There are two distinct approaches to such value-oriented forecasting. A first, integrated approach attempts to train the forecaster directly to maximize the downstream value. The primary challenge with this approach is that applying the gradient descent method requires calculating the derivative of the optimal decisions with respect to the forecast, i.e., a derivative "through" the optimization program. Calculating these derivatives has been accomplished for non-linear convex programs using implicit differentiation of the Karush-Kuhn–Tucker (KKT) conditions of the downstream problem [29–31]. However, (mixed)-integer linear programs exhibit zero gradients almost everywhere [32]. To address that issue, the training problem has been approximated by introducing smoothing terms in the optimization objective [32–34], or by using a surrogate value function [35]. There are two major drawbacks to this integrated approach. The first is that the training procedure is computationally expensive because it requires solving the downstream optimization problem for every example in the training set repeatedly, i.e., for every epoch in the training procedure. Secondly, although it has been implemented for stand-alone optimization, i.e., where the value function corresponds to the optimization objective, it is not obvious how it can be deployed to models with other structures such as heuristic decision rules and MPC. To date, integrated value-oriented approaches have been shown to outperform traditional forecasting-based approaches in stand-alone applications such as participation in day-ahead electricity markets [36,37]. Chapter 5 of [38] shows this can only be achieved when (i) all optimized decisions are implemented, and (ii) the optimization objective aligns directly with the realized value. In contrast, rolling horizon optimization contexts — such as MPC in real-time markets — typically involve the execution of only the first decision from each optimization problem. This leads to a misalignment between the optimization objective and the realized value, which is determined solely by the executed decisions.

In the second approach, the objective is to identify a suitable loss function, based on traditional accuracy metrics, that effectively aligns the forecaster with the downstream task. In [39], a smoothed version of a Theil accuracy metric is proposed to train day-ahead electricity prices because of its favorable properties regarding differentiablility and robustness against outliers. In the context of minimizing generation costs based on load forecasts, [40] proposes a loss function which is piece-wise linear in the load forecast error, to capture the asymmetric impact of load forecast errors on the system cost. In [41], a directionality-based penalty term is added to a Mean Absolute Error (MAE)-style loss function for equity price forecasting applied to the problem of optimal trading in financial markets. In [42], a forecaster is trained by minimizing a loss function that gives more importance to large wind power output forecast errors and it is shown that this improves the downstream profit obtained from energy trading in the day-ahead market. While these implementations of the loss-based approach to value-oriented forecasting acknowledge the subsequent use in a downstream problem, they generally do not formally integrate it in the training method of the forecaster.

1.3. Contributions

In conclusion, current loss-based approaches have not explicitly accounted for downstream value in the learning phase of the forecaster, while integrated methods are computationally expensive and unsuitable for rolling horizon optimization settings such as MPC, where only the first decision of each optimization cycle is executed. To overcome these issues, we propose the method of loss tuning — a novel, loss-based value-oriented forecasting approach. It involves defining a generalized loss function parameterized by hyperparameters, training a set of forecasters using different instances of this loss, running the downstream optimization problem with each trained forecaster, and selecting the one that yields the highest downstream value (e.g., profit from ESS market participation). The scientific contribution is threefold:

- The proposed method of loss-tuning is a loss-based value-oriented forecasting approach that systematically integrates the downstream problem that outperforms traditional forecasting methods in rolling horizon optimization settings like MPC. It is, in fact, agnostic to the specific nature of the downstream task, making it broadly applicable.
- We apply the method of loss tuning to participation of ESS in RT balancing markets and propose a parametric family of loss functions tailored to that problem. The tunable parameters of this loss function include (i) a continuously variable exponent of the error term and (ii) a variability component to accurately capture price variations over the forecast horizon.
- In a real-life case study of the Belgian balancing market using actual price data, we show that the proposed methodology can outperform traditional (accuracy-based) forecasts. The increase in ex-post out-of-sample profit ranges from 13% to 176% for different types of ESSs maximizing their profit through implicit balancing actions based on forecast-informed MPC.

The proposed methodology is of interest for any agent who deploys forecasts in a downstream (optimization) problem, and in particular for ESS operators seeking inter-temporal arbitrage opportunities in volatile short-term electricity markets. The remainder of the paper is organized as follows. Section 2 introduces European balancing markets, and describes the strategy of implicit balancing where real-time out-of-balance actions expose market participants to the imbalance price. In Section 3, we present the methodology, focusing both on the general training framework, as well as the specific implementation for an ESS maximizing its implicit balancing profits through MPC. An extensive case study considering the Belgian balancing market is presented in Section 4. Finally, we provide some concluding remarks in Section 5.

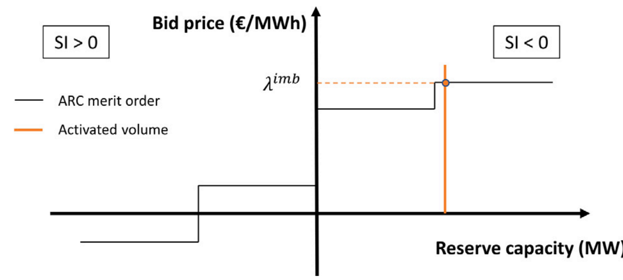

Fig. 1. Illustration of the balancing market clearing.

2. Balancing markets

In European electricity markets, the balancing market is designed to correct a mismatch in supply and demand to avoid grid frequency deviations. In the day-ahead stage, balancing service providers place bids for (i) the reservation of balancing capacity and (ii) for activation of that capacity in Real-Time (RT). Aggregating the bids for reserve activation gives rise to the Available Regulation Capacity (ARC) merit order. In the RT stage, the System Operator (SO) decides on the amount of activated reserves based on the observed and forecasted SI. As depicted in Fig. 1, the ARC merit order represents the supply side of the balancing market, while the activated reserves constitute the demand. The intersection yields the imbalance price, which is either paid by market participants in the case of excess supply in their portfolio or credited to them in the case of a shortage, effectively internalizing the cost of imbalances within the market. This pricing mechanism creates a potential opportunity: market participants can deliberately deviate from their contracted positions to capitalize on favorable imbalance prices. This strategy, known as implicit or passive balancing, has been shown to allow higher profits compared to day-ahead market participation, particularly when using MPC to optimizing decisions [16]. Importantly, this approach does not require submitting bids; exposure to the imbalance price arises solely from real-time deviations between scheduled and actual positions.

The success of such an approach relies heavily on actionable forecasting of the imbalance price, which is the central focus of this paper. Since the SI critically affects the imbalance price, a two-step approach for imbalance price forecasting has been shown to result in the highest accuracy [43]. In our previous work, we adopted a method following such a two-step procedure to forecast prices informing ESS decision-making [16,26]. The first step consists of forecasting SI quantiles using a RNN. In the second step, those quantiles, with corresponding weights, are converted to imbalance price values by minimizing the cost of activating balancing energy, which corresponds to clearing the balancing market. As such, this can be regarded as a fundamental approach for imbalance price forecasting. Notice that the second step in this procedure assumes that (i) the SO has perfect foresight of the SI and activates the exact amount of reserves to counteract the SI, and (ii) the merit order is perfectly known. In this paper, we replace the second step of the two-step forecasting approach in [16] by introducing a RNN to map data of the forecasted SI and the published ARC merit order to imbalance prices.

3. Method

This section presents our proposed loss-based value-oriented neural network training methodology. The overall approach is described in Section 3.1, while Section 3.2 details its application to an ESS participating in balancing markets using MPC.

3.1. General training framework

Here, we lay out the value-oriented NN training methodology. First, we define a new generalized loss function including a variability component in Section 3.1.1. In Section 3.1.2, the general value-oriented hyperparameter tuning framework is presented. Finally, Section 3.1.3 outlines the proposed solution strategy to the optimization problem outlined before.

3.1.1. Generalized loss function

This paper revolves around the premise that when training a forecaster, the modeler is unaware of the optimal (statistical) loss function to be implemented such that the resulting forecaster yields favorable downstream performance. To overcome this issue, we propose to consider a generalized loss function, defined as a parametric family:

$$\mathcal{L} = f(\hat{\lambda}, \lambda; \xi_l), \tag{1}$$

with ξ_l the parameters defining the specific loss function, $\hat{\lambda}$ the forecasted variable and λ its corresponding ground truth. Depending on the problem at hand, λ and $\hat{\lambda}$ can be scalars, vectors, multi-dimensional arrays or even probability distributions. The form of the parametric family should be defined based on an understanding of the problem. In Section 3.2.2, we propose a parametric family for time series forecasting of prices in the context of an agent seeking inter-temporal arbitrage opportunities.

3.1.2. Value-oriented hyperparameter tuning

The traditional training procedure of a NN considers different sets of HyperParameters (HPs) which can be related, but are not limited, to the NN architecture (e.g. activation functions, number of hidden layers, etc.). Because of their unknown impact on the forecast accuracy, an iterative procedure is adopted in which different NNs with different values of those HPs are trained to minimize the pre-defined loss function $\mathcal{L}$ on the training set. In the validation step, the NN that has the lowest loss value on the validation set, which is composed of data that is not used in calculating the gradients in the training step, is selected. This procedure can be translated mathematically to:

$$\underset{\xi_a}{\text{minimize}} \sum_{i \in X_v} \mathcal{L}\left(\pi_{\theta^*}^{\xi_a}(x_i), \lambda_i\right) \tag{2a}$$

subject to:

$$\theta^* = \arg\min_{\theta} \sum_{j \in X_t} \mathcal{L}\left(\pi_{\theta}^{\xi_a}(x_j), \lambda_j\right), \tag{2b}$$

where x_i represents the i^{th} example of input features fed to the NN, and λ_i the actual values of the quantity to be predicted, associated to that example. The predicted quantity is the output of the NN, i.e., $\hat{\lambda}_i = \pi_{\theta}^{\xi_a}(x_i)$, where ξ_a are the architectural HPs of the NN, and θ the NN parameters, also referred to as "weights and biases". In Problem (2), the upper level optimization tries to find the optimal hyperparameters ξ_a, that upon training of the NN parameters θ using the train set X_t, minimizes the loss on the validation set X_v.

Here, we propose a new training framework that extends Problem (2) by (i) using a variable loss function and considering its associated unknown parameters ξ_l as HPs to be tuned, and (ii) evaluating the validation performance using the downstream value. This forecaster training problem can then be written as:

$$\underset{\xi_a, \xi_l}{\text{maximize}} \sum_{i \in X_v} V\left(d_i^*, \lambda_i\right) \tag{3a}$$

subject to:

$$d_i^* = DP\left(\pi_{\theta^*}^{\xi_a}(x_i)\right) \tag{3b}$$

$$\theta^* = \arg\min_{\theta} \sum_{j \in X_t} \mathcal{L}_{\xi_l}\left(\pi_{\theta}^{\xi_a}(x_j), \lambda_j\right). \tag{3c}$$

Here, DP represents the downstream problem which is informed by the predicted values. The output of this problem is the optimal decision d^*. The value function V is typically the objective function of DP, evaluated using the ground truth λ, if it is an optimization program.

Within this training procedure (3), there are no constraints to the type of forecaster nor the form of the downstream problem, in contrast to existing integrated methods to value-oriented forecasting [33,35]. Therefore, the forecaster can leverage the predictive power of, e.g., deep learning models, while the downstream problem can be anything ranging from optimization to MPC and a heuristic decision rule.

While Problem (3) formally describes the objective of maximizing downstream value on the validation set, using forecasters trained on the training data, implementing this as an integrated optimization is often computationally intractable, especially in settings where the downstream problem corresponds to an MPC problem. In practice, and as we propose in this work, it is more efficient to decouple the training and validation procedures. This separation allows for scalable use of powerful forecasting models while still guiding hyperparameter selection based on downstream performance, as discussed in Section 3.1.3.

3.1.3. Solution strategy

Algorithm 1 Value-Oriented Hyperparameter Tuning

Input: Training data X_t, validation data X_v, number of epochs K and HP combinations M
Initialize: HP spaces Ξ^a (architecture), Ξ^l (loss)
Sample M combinations $\{(\xi_m^a, \xi_m^l)\}_{m=1}^{M}$ from $\Xi^a \times \Xi^l$ ▷ (A1)

for $m = 1$ to M **do**
  $\theta_m^* \leftarrow \text{TrainModel}(\xi_m^a, \xi_m^l, X_t, X_v)$ ▷ (A2)
end for

for $m = 1$ to M **do**
  for all $i \in X_v$ **do**
    Forecast: $\hat{\lambda}_i^m \leftarrow \pi_{\theta_m^*}^{\xi_m^a}(x_i)$
    Optimize: $d_i^{*,m} \leftarrow DP(\hat{\lambda}_i^m)$ ▷ (B1)
  end for
  Evaluate: $V_m \leftarrow \sum_{i \in X_v} V(d_i^{*,m}, \lambda_i)$ ▷ (B2)
end for
Select: $m^* \leftarrow \text{argmax}_m V_m$ ▷ (B3)
Output: Trained model $\pi_{\theta_{m^*}^*}^{\xi_{m^*}^a}$

function $\text{TrainModel}(\xi^a, \xi^l, X_t, X_v)$
  Define model $\pi_{\theta}^{\xi^a}$ and loss $\mathcal{L}_{\xi^l}$
  Initialize parameters θ
  Initialize $\mathcal{L}_{\text{best}} \leftarrow \infty$
  for $k = 1$ to K **do** ▷ Gradient descent
    $\theta \leftarrow \theta - \eta \cdot \nabla_{\theta} \mathcal{L}_{\xi^l}$
    $\mathcal{L}_{\text{val}} \leftarrow \frac{1}{|X_v|} \sum_{i \in X_v} \mathcal{L}_{\xi^l}(\pi_{\theta}^{\xi^a}(x_i), \lambda_i)$
    if $\mathcal{L}_{\text{val}} < \mathcal{L}_{\text{best}}$ **then**
      $\mathcal{L}_{\text{best}} \leftarrow \mathcal{L}_{\text{val}}$
      $\theta^* \leftarrow \theta$
    end if
  end for
  return θ^*
end function

Because the lower level in Problem (2) is usually highly non-convex, the main method for solving it is the gradient descent algorithm [44]. As solving the lower level is computationally intensive, the upper level is often solved by trying a limited number of HP samples. Even though algorithms for finding the optimal HP values have been introduced, randomly sampling the HPs can yield similar performance [45]. The solution strategy for the newly proposed training procedure is outlined

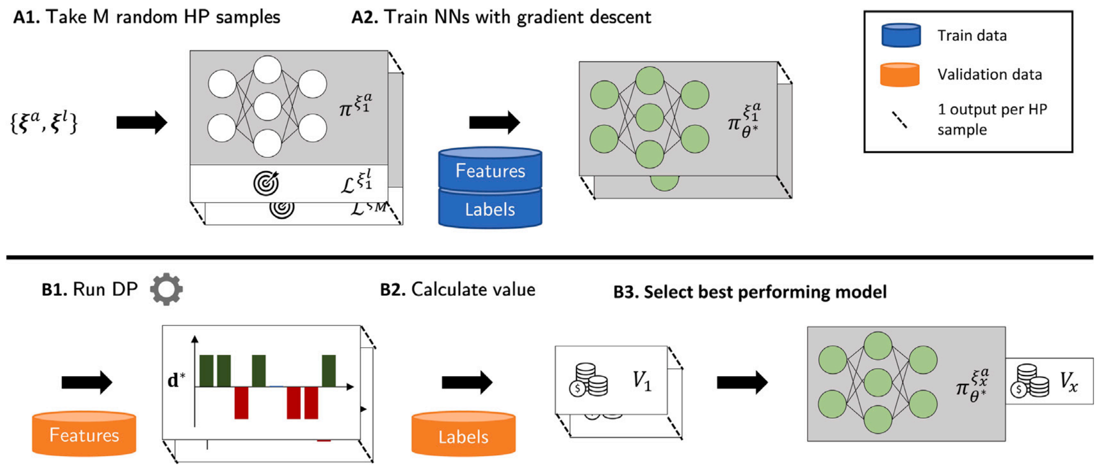

Fig. 2. Solution strategy of the value-oriented HP tuning problem (3). Starting from loss and architectural HPs, M loss functions and randomly initialized NNs are defined. Those NNs are trained to minimize their respective loss function with the train data. The downstream problem is run M times on the validation set with forecasts from the different trained models, yielding M different decision sets, from which the downstream value can be calculated using the ground truth of the forecasted value. The model yielding the highest downstream value is selected and deployed on the test set.

in Algorithm 1 and visualized in Fig. 2. The upper part (A) depicts the training step whereas the lower part (B) depicts the step related to selecting the best trained NN. In (A1), M sets of values are randomly sampled for the predefined architectural and loss-related HPs, being denoted ξ^a and ξ^l respectively. This results in M distinct NNs and loss functions $\{\{\pi^{\xi^a_m}, \mathcal{L}^{\xi^l_m}\}|m = 1, \dots, M\}$. In (A2), those NNs are trained to minimize their respective loss functions on the training data with gradient descent, using early stopping on the validation set. In (B1), the resulting trained NNs are deployed on the same validation set to yield forecasts $\hat{\lambda}$, i.e. imbalance prices in this paper. These are fed to the Downstream Problem DP leading to M (optimized) decision sets d^*. In (B2), the downstream value is calculated from the optimized decision sets and the validation labels, i.e., the ground truth λ. Finally, in (B3), the trained NN with the highest downstream value is selected and can be deployed on the test set. Note that the upper and lower part are decoupled. Indeed, the choice of HP values is pre-defined and as such does not depend on intermediate training outcomes. The same M trained NNs can therefore be tested on different downstream problems. This can be useful for agents with different tasks that require similar forecasting capabilities, such as an owner of multiple assets with distinct technical characteristics simultaneously participating in electricity markets.

3.2. Implementation for ESS implicit balancing

In this section, we describe the ESS profit maximization strategy in RT balancing markets that constitutes the downstream problem in this paper. In Section 3.2.1, we write down the MPC algorithm determining the ESS actions. Section 3.2.2 presents a generalized loss function tailored for arbitrage-seeking agents, such as ESSs. Finally, Section 3.2.3 details how we reduce the amount of hyperparameters arising from that generalized loss function.

3.2.1. ESS model predictive control problem

When considering ESS participation in electricity markets, the European balancing markets are an interesting case because the regime-switching of prices leads to high and frequent price spreads [46,47]. At the same time, it is notoriously difficult to accurately predict the SI, and therefore imbalance prices [48]. A strategy that is increasingly investigated is applying Model Predictive Control (MPC) to constantly re-optimize the ESS schedule in RT markets. By doing that, the difficulty of forecasting is partly overcome as the latest information can be leveraged to produce the best possible forecasts. A second merit of MPC is that it exploits exact knowledge of the SoC of the ESS, contrasting decision-making in day-ahead which is subject to uncertainty on the activation of reserves. In our previous work [16], we have adopted MPC for participation in RT European balancing markets through implicit (or passive) balancing. With this strategy, the ESS takes RT out-of-balance positions, and as such is exposed to the imbalance price without the need of decisions in day-ahead. The optimization program governing the MPC reads:

$$\max \sum_{\tau \in \mathcal{T}_{LA}} \left(\hat{\lambda}_\tau(t) \left(e^+_\tau \eta^+ - \frac{e^-_\tau}{\eta-} \right) - \Gamma_\tau \right) \tag{4a}$$

subject to:

$$0 \le e^+_\tau + e^-_\tau \le \overline{P} \Delta t \qquad \forall \tau \tag{4b}$$

$$s_{\tau+1} = s_\tau + e^-_\tau \eta^- - e^+_\tau / \eta + \qquad \forall \tau \tag{4c}$$

$$\underline{\text{SoC}} \le s_\tau \le \overline{SoC} \qquad \forall \tau \tag{4d}$$

$$s_1 = SoC(t) \tag{4e}$$

$$0 \le \Gamma_\tau \ge C^B \left(\delta_{\tau+1} - \delta_\tau \right) \qquad \forall \tau \tag{4f}$$

$$\delta_\tau = \sum_{l \in L} b_{\tau,l} p_l + \frac{p_{l+1} - p_l}{q_{l+1} - q_l} \left(\tilde{s}_{\tau,l} - b_{\tau,l} q_i \right) \qquad \forall \tau \tag{4g}$$

$$b_{\tau,l} q_l \le \tilde{s}_{\tau,l} \le b_{\tau,l} q_{l+1} \qquad \forall \tau, l \tag{4h}$$

$$\sum_{l \in L} \tilde{s}_{\tau,l} = s_\tau \qquad \forall \tau \tag{4i}$$

$$\sum_{l \in L} b_{\tau,l} = 1 \qquad \forall \tau \tag{4j}$$

$$b_{\tau,l} \in \{0, 1\}, \qquad \forall \tau, l \tag{4k}$$

with decision variables $\{\boldsymbol{e}^+, \boldsymbol{e}^-, \boldsymbol{s}, \tilde{\boldsymbol{s}}, \boldsymbol{b}, \boldsymbol{\Gamma}\}$, adopting the convention that bold symbols represent vectors. In the objective, Eq. (4a), the BESS maximizes the expected profit from the implicit balancing actions based on imbalance price forecasts $\hat{\boldsymbol{\lambda}}$, corrected with the degradation costs from cycling the battery. Eqs. (4b)–(4d) are typical ESS constraints limiting the (dis)charge power and the SoC, and dictating the SoC evolution from the (dis)charge decisions. Eq. (4e) sets the initial SoC. Finally, Eqs. (4f)–(4k) account for the (non-convex) battery degradation costs through a mixed-integer linear approximation, see [16,49].

The MPC approach operates in a rolling horizon fashion: at each time step t, the optimization problem (4) is solved over a finite lookahead horizon $\mathcal{T}_{LA}$ using the latest forecasts $\hat{\lambda}_\tau(t)$. However, only the first control action, i.e., charging or discharging at the current time, is implemented. The SoC is then updated according to the following equation:

$$SoC(t+1) = SoC(t) + e_1^{-*}(t)\eta^- - \frac{e_1^{+*}(t)}{\eta^+}, \tag{5}$$

and the horizon rolls forward by one step. This process repeats sequentially over the entire MPC horizon T, continuously incorporating updated forecasts and system states, thereby allowing the ESS to adapt to changing market conditions and forecast uncertainties in real time.

The downstream value function is the same as the optimization objective function (4a), but the predicted price is replaced with the actual value. Additionally, the sum runs over the MPC horizon instead of the optimization horizon, taking for every time step the optimized values of the first instance:

$$V = \sum_{t\in T} \lambda_1(t)\left(e_1^{+*}(t)\eta^+ - \frac{e_1^{-*}(t)}{\eta-}\right) - \Gamma_1^*(t). \tag{6}$$

3.2.2. Generalized loss for arbitrage

In this section, we define a parametric family of loss functions for time series forecasting in the context of capturing inter-temporal arbitrages. It generalizes three dynamics having an a priori unknown impact on the downstream value (6). The generalized loss of a vector prediction $\hat{\boldsymbol{\lambda}}$ over N lookahead instances, provided the ground truth $\boldsymbol{\lambda}$, is given by:

$$\mathcal{L} = \sum_{n=1}^{N} \Bigg(w_{n,0}|\lambda_n - \hat{\lambda}_n|^p + \sum_{k=1}^{N-n} w_{n,k}|(\lambda_{n+k} - \lambda_n) - (\hat{\lambda}_{n+k} - \hat{\lambda}_n)|^p \Bigg), \tag{7}$$

where the hyperparameters[1] $w_{n,k}$ are the weights of the loss function pertaining to the n^{th} lookahead instance and k^{th} lookahead time difference. Note that setting $w_{n,k} = 0$ for nonzero k and $w_{n,0} = 1/n$ for all n is equivalent to using the MAE and MSE when $p = 1$ and $p = 2$ respectively. The first generalization of (7) with respect to traditional loss functions relates to p. When using the MSE as loss function, one could expect that the variability over the forecast lookahead period may be lower compared to a NN trained using the MAE, as large forecast errors are penalized more severely. Since the variability of the predicted price is expected to affect the quality of decision-making, but the exact effect is a priori unknown, p is considered a continuously variable hyperparameter. The second and third generalizations relate to the weights array $\{w_{n,k}|n \in [1,\ldots,N], k \in [1,\ldots,N-n]\}$. The first dimension of the weights array assigns different importance to different lookahead instances. It may, for example, be beneficial to assign lower weights to instances further in the future when it is known that the input features for those instances are less accurate. Finally, if the weights array is nonzero for $k > 0$, the generalized loss function exhibits a variability component. While traditional loss functions penalize the static difference between the forecasted and actual values, here the difference in variability between forecasted and actual values is accounted for. As correctly capturing price differences is crucial for agents seeking to exploit inter-temporal arbitrages, such a variability component in the loss function is expected to increase the downstream value of the price forecast.

3.2.3. Reduction of loss-related hyperparameters

An issue arising from generalized loss (7) is the large amount of parameters, and hence loss-related HPs. Indeed, the amount of weights equals $|\mathcal{T}_{LA}| \cdot (|\mathcal{T}_{LA}| + 1)/2$ with $|\mathcal{T}_{LA}|$ the length of the lookahead horizon. This results in a 55-dimensional HP space when a lookahead horizon of 10 time steps is implemented. A comprehensive search in this space requires training a large amount of NNs which is computationally inefficient. To mitigate this problem, we propose three distinct expressions for the weights in (7) to reduce the amount of HPs:

[1] Corresponding to the loss-related hyperparameters ξ_l in Problem (3)

$$\text{VOa:}\quad w_{n,k} = \begin{cases} e^{-\alpha n} & \text{if } k = 0 \\ 0 & \text{otherwise} \end{cases} \tag{8a}$$

$$\text{VOb:}\quad w_{n,k} = \begin{cases} 1 - A & \text{if } k = 0 \\ A & \text{otherwise} \end{cases} \tag{8b}$$

$$\text{VOc:}\quad w_{n,k} = \begin{cases} (1-A)e^{-\alpha n} & \text{if } k = 0 \\ Ae^{-(\alpha n + \beta k)} & \text{otherwise.} \end{cases} \tag{8c}$$

The codes (VOa, VOb and VOc) associated with the expressions are used to refer to models trained with that specific HP reduction. “VO” here refers to the fact that these models are trained with the Value-Oriented methodology. Expression VOa disregards the variability component in (7) and considers an exponentially decaying importance of accurately forecasting the price further in the future. Such exponentially decaying weights are chosen, since it has been shown [16] that the input features to the model (SI forecast) are significantly less accurate for further away lookahead time steps compared to those in the near future. The VOa model has 2 loss-related HPs: $\{\alpha, p\}$, where p is the exponent in (7) and α the coefficient of exponential decay of the weight w.r.t. the lookahead. VOb does consider the variability component, and gives every instance in the lookahead horizon equal importance. It also has 2 loss-related HPs: $\{A, p\}$, where A represents the weight of the variability component in the loss function. Finally, VOc considers a combined exponential decaying importance of both future lookahead instances and time differences in the variability component. Its 4 loss-related HPs are: $\{A, \alpha, \beta, p\}$, where β is the coefficient of exponential decay of the weight w.r.t. the time difference.

4. Case study

In this section, the presented methodology outlined in Section 3 is tested on a case study of the Belgian Balancing market. First, Section 4.1 describes the design of the case study. Section 4.2 discusses the downstream profit performance results when implementing the forecasters in the proposed MPC algorithm. Section 4.3 gives an overview of the forecasting errors produced by the different trained models. Finally, an analysis of the impact from the assumptions in the generalized loss function is given in Section 4.4.

4.1. Case study design

The proposed methodology is validated on the Belgian balancing market for a test set comprising the months November and December 2020. The train and validation sets comprise the period January 2019–October 2020, applying a 70/30 train/validation split. The input features for the imbalance price forecaster are forecasted SI quantiles and the ARC merit order. The merit order consists of 20 price levels of balancing power activation ranging from −1000 MW to +1000 MW and is found on the Elia Data Download page [50]. Quarter hourly actual imbalance price values are used for the ex-post analysis and are also found from [50]. The forecasted SI quantiles are generated from a separate RNN forecaster proposed in [26]. The SI forecaster was trained with data from 2015–2018. The SI forecaster uses two categories of input features. The first consists of historical observations, including past SI values, load, and generation from wind, solar, gas, and nuclear sources, and inter-zonal power flows. The second includes forecasted values, which are provided by the TSO [50], for the same variables: load, generation from different sources, and inter-zonal flows. All features exhibit a quarter-hourly granularity.

To test the proposed value-oriented methodology, training procedure (3) is performed for the three expressions of loss-related HPs from (8) separately. In each run, the outer optimization (HP tuning) consists of evaluating 100 random samples of loss-related HPs. For all the trained models, the architecture is unchanged, such that we single

Table 1
Overview of different MPC applications.

Application code →	B1	G1	G4
Power rating (MW)	1	1	1
Energy rating (MWh)	1	1	4
Rountrip efficiency	0.9	0.81	0.81
Inv. cost (€/MWh)	150,000	–	–

out the effect of a varying loss function. The chosen architecture is an encoder–decoder Long-Short Term Memory (LSTM) RNN with 64 hidden units. The encoder has a lookback horizon of 4 quarter hours and the decoder a lookahead horizon of 10 quarter hours. All the models are implemented with the *PyTorch* library, and trained using the *ADAM* optimization procedure. For the downstream problem in (3), we employ MPC algorithm (4) maximizing ESS profits. Three different implementations of the MPC are used to showcase how this method can be generalized. In the first application, B1, the investment cost determining the battery degradation costs in Eq. (4f) amounts 150,000€/MWh and a roundtrip efficiency of 90% is assumed, resembling a large-scale BESS [16]. In the other applications, the degradation costs are neglected and the ESS has a roundtrip efficiency of 81%, representing a general non-specified storage system. The two applications have an Energy-to-Power (E/P) ratio of 1 and 4 and are referred to as G1 and G4 respectively. An overview of the different applications is given in Table 1.

Two distinct benchmarking methods are used to show the effectiveness of the proposed value-oriented methodology. The first benchmark is an Optimization Program (OP) that minimizes the cost of activating balancing energy from the ARC merit order, elaborated upon in Section 2. This is also the only benchmark available in the literature for imbalance price forecasting applied to ESS profit maximization, see [16,26]. A second benchmark is the Traditional (T) forecast training procedure (2). Again, 100 forecasters are trained using three loss functions: T1 refers to model trained with the MAE. T2 refers to model trained with the MSE. Finally, T3 refers to the model trained with the cubed absolute error between the forecast and actual prices.

4.2. Comparing forecasters with downstream value

Table 2 summarizes the results obtained from using the different forecasting models in the different MPC applications. When deploying these models in real-life, the agent would in principle select the forecaster with the best performance on the validation set. Note that the metric measuring that performance is different for the traditional models (being the validation loss) than for the value-oriented models (being the MPC profit on the validation set). The MPC profit on the test set attained by the selected forecaster is then the true ex-post out-of-sample performance of the different models. To showcase the robustness of the results, the range of the validation profit for the 5 best performing forecasters per model is given, together with the profit range of the same forecasters in the test set. Even though the traditional models (T1,T2,T3) did not have varying HPs in their training procedure, there is still a range in profit outcomes, typically amounting around 10% of the maximum value for the G1 and G4 applications, and even 30% for the B1 application. This profit range results from the sensitivity to randomness in the NN training procedure, arising from the random initialization of weights and biases, and random shuffling of training examples. Indeed, as the gradient descent method for training the neural networks does not guarantee a globally optimal solution, different initial sampling of the forecaster parameters leads to different solutions of the training procedure, yielding different ESS decisions and profits. The Perfect Foresight (PF) benchmark considers the actual imbalance prices within the optimization and serves as an upper bound to the attainable profits per application. Note that in general, the agent with perfect price foresight strongly outperforms the one taking decisions based on uncertain forecasts in terms of profit. This is due to the highly uncertain nature of the balancing market. This effect increases with decreasing ESS energy content, as this results in a higher importance of timing the optimal (dis)charge moments exactly right.

Table 2 shows that the value-oriented models strongly outperform both the optimization program and the traditional models. The profit improvement ranges from 13% for application G1 to 176% for application B1. For VOb and VOc, the profit of all 5 best performing models is higher than the attained profit by the OP, demonstrating the merit of including the variability component in (7), which is nonzero only for VOb and VOc. Recall that the profits for the B1 application include the degradation costs, therefore covering the investment cost of the battery, leading to significantly lower profit values compared to the applications without degradation costs.

Fig. 3 gives an in-depth comparison of the different models for application G1. It shows a snapshot of predicted imbalance prices and resulting forward discharge schedules. At the start of the lookahead horizon, all the models are at the maximum SoC. The comparison is limited to the OP, traditional models T1 and T2, and the value-oriented VOc for the sake of clarity. The most significant observation is that the only forecaster correctly predicting the qualitative short-term price evolution is VOc. Indeed, whereas the other models believe the price in the first time step is a local minimum, the forecast by VOc correctly identifies it as a local maximum. For that reason, only the agent acting by the VOc forecasts will discharge in the first time step. Since in the MPC algorithm re-optimization occurs after every instance realization, the short-term behavior will have the highest impact on the final result. Secondly, the low price forecasts by the OP can be attributed to missing information in the merit order. These negative price forecasts lead to the peculiar behavior of systematic charging decisions, resulting from the optimization program in (4) allowing for both charging and discharging within a quarter hour, resulting from constraint (4b).[2] Even though the impact on the SoC is zero, this behavior leads to a net charge from the grid due to (dis)charge inefficiencies. As the ESS is fully charged at the start of the snapshot in Fig. 3 such combined charging/discharging is indeed the most profitable strategy given the negative price forecasts. Throughout the test set, all models exhibit such behavior, with an impact of less than 0.5% on their total profit.

4.3. Comparing forecasters with accuracy metrics

Fig. 4 shows the MAE and the Root Mean Squared Error (RMSE) of the imbalance price predictions over the test set, comparing the different models introduced in Section 4.1. The first observation is that models trained to minimize different loss functions (unsurprisingly) perform differently when considering different, but equally reasonable, accuracy metrics. This is the basis for the philosophy adopted in training procedure (3) of varying the loss function and evaluating its impact on the downstream problem.

A second observation lies in the statistical accuracy of models trained from different training strategies. The traditional models (T1,T2,T3), trained according to procedure (2) tend to outperform the value-oriented models with non-zero variability component in (7), VOb and VOc. This is a sensible result as those value-oriented models are partly trained to predict the variability of the imbalance prices, which is not captured in the MAE and RMSE accuracy metrics. Finally, the benchmark of estimating the imbalance price with an optimization program generally exhibits a poor statistical accuracy.

Comparing Fig. 4 to Table 2 leads to the observation that even though the traditional models do outperform the OP in terms of traditional forecast accuracy, they do not generally outperform the OP

[2] This is indeed physically possible since the storage operator could decide to first discharge 7.5 min, and then charge for the remainder of the quarter hour.

Table 2
Overview of attained profit by different forecasting models and a Perfect Foresight (PF) benchmark for the three applications summarized in Table 1. The bold numbers indicate the profit by the model that would be selected (being the one with the highest validation profit) per application. All numbers in €/MW/qh.

App.	Metric	OP	T1	T2	T3	VOa	VOb	VOc	PF
B1	Validation profit	0.10	0.18	0.24	0.09	0.33	**0.37**	0.34	1.87
	Test profit	0.17	0.25	0.18	0.12	0.30	**0.47**	0.38	2.08
	Top 5% val. profit range	–	[0.17,0.22]	[0.22,0.25]	[0.09,0.14]	[0.31,0.33]	[0.34,0.37]	[0.32,0.34]	–
	Top 5% test profit range	–	[0.19,0.25]	[0.17,0.25]	[0.09,0.15]	[0.27,0.39]	[0.45,0.47]	[0.36,0.47]	–
G1	Validation profit	1.57	1.30	1.48	1.35	1.69	1.66	**1.70**	4.19
	Test profit	1.60	1.32	1.58	1.35	1.66	1.84	**1.80**	4.28
	Top 5% val. profit range	–	[1.25,1.31]	[1.47,1.55]	[1.33,1.39]	[1.64,1.69]	[1.62,1.66]	[1.66,1.70]	–
	Top 5% test profit range	–	[1.31,1.33]	[1.44,1.60]	[1.25,1.46]	[1.60,1.73]	[1.64,1.84]	[1.68,1.80]	–
G4	Validation profit	2.11	1.62	1.99	1.93	2.21	**2.66**	2.63	4.77
	Test profit	2.16	1.63	2.18	1.94	2.16	**2.62**	2.84	4.86
	Top 5% val. profit range	–	[1.54,1.66]	[1.99,2.10]	[1.85,2.00]	[2.16,2.21]	[2.19,2.66]	[2.22,2.63]	–
	Top 5% test profit range	–	[1.57,1.63]	[1.92,2.18]	[1.86,2.07]	[2.16,2.28]	[2.21,2.62]	[2.20,2.84]	–

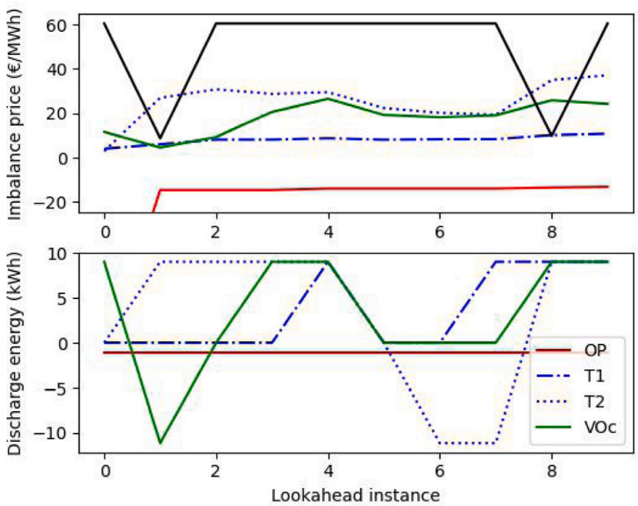

Fig. 3. Imbalance price forecast (top) and optimized forward discharge energy decisions (bottom) for different forecasting models applied to an ESS with $E/P = 1$ not considering degradation costs at time step 01/11/2020 07:45. The black curve represents the actual price.

in terms of realized profit. In fact, the OP significantly outperforms T1 and T3, whereas it achieves similar results as T2. This observation strengthens the insight that improving on traditional accuracy metrics does not necessarily correspond to improving the downstream value, and therefore that value-oriented methodologies should be adopted when price forecasts are deployed in downstream decision problems.

4.4. Analysis of loss-related hyperparameters

In this section, a concrete example is used to showcase the relevance of allowing the loss function to have varying hyperparameters. Fig. 5 shows the validation profit of all the trained forecasters within the VOb model for the three different MPC applications, as a function of the exponent p in the generalized loss function (7). The validation profit is expressed in a dimensionless score calculated as the validation profit divided by the median validation profit for that application, which facilitates the comparison across applications. As a single forecaster typically results in different validation profit scores for different applications (seen in the vertical spread per value of p), it can be concluded that it is indeed sensible to consider application-specific loss functions. A second observation is the trend of lower available energy content requiring lower p. Indeed, on the left side of Fig. 5, B1 tends to perform best and G4 worst, whereas it is the reversed order on the right side of the graph. One could regard B1 as the model with the lowest available energy, as the agent is reluctant to access low SoC states because of the relatively high degradation costs in that regime. These results are in line with the relative performance of models T1 and T3 in Table 2. This effect may be explained by considering the variability of the forecasts. A forecaster trained with low p will yield more variable forecasts than a forecaster trained with high p. The low variability of the predictions of a forecaster with high p inhibits it from taking decisions with limited energy availability (especially when degradation costs are also taken into account), whereas predictions with high variability may lead to erroneous (dis)charge decisions for a forecaster trained with low p. Finally, the variation in validation profit is much higher for the application considering degradation costs than for those without. This can be explained through the degradation costs amounting roughly 80% of the arbitrage profits. Consequently, small changes in attained arbitrage profits can have a comparatively large effect on the total profit. This effect also contributes to the larger profit improvement, of the value-oriented methodology w.r.t. the benchmarks, for application B1 compared to the other applications.

Note that in Fig. 5, the exact same 100 trained NNs were evaluated in three different applications. In theory, different agents with similar applications (that require similar generalized loss functions) could combine their computing power to train the models, and separately investigate which of the models performs best on their own distinct application. This is not possible for existing integrated value-oriented forecasting methodologies [33,35], which require a specific training technique for specific types of applications. This highlights the advantage of the proposed generalizable methodology.

4.5. Computational aspects

Table 3 presents the average training times for the benchmark model (T2) and the proposed loss-tuning methods (VOa, VOb, VOc). As expected, NN training takes longer for the value-oriented models, due to the added complexity introduced by our proposed parametric family of loss functions including a variability component. Adding the optimization on the validation set to the NN train times, loss-tuning leads to training procedures that are 28% to 48% higher than those of the traditional approach. However, this remains two orders of magnitude faster than integrated value-oriented methods, which require solving the downstream optimization problem during every forward pass of training. Given that (i) our training set is 2.3 times larger than the validation set, (ii) each optimization problem takes around 630 s to solve and (iii) models are trained over approximately 150 epochs, the cumulative cost of optimization in an integrated setup would exceed 200,000 s (roughly 55 h). For this reason, we did not implement integrated methods in this study.

5. Conclusion

With increasing volatility of RT electricity prices and declining costs of storage solutions, ESSs should look into novel ways to exploit

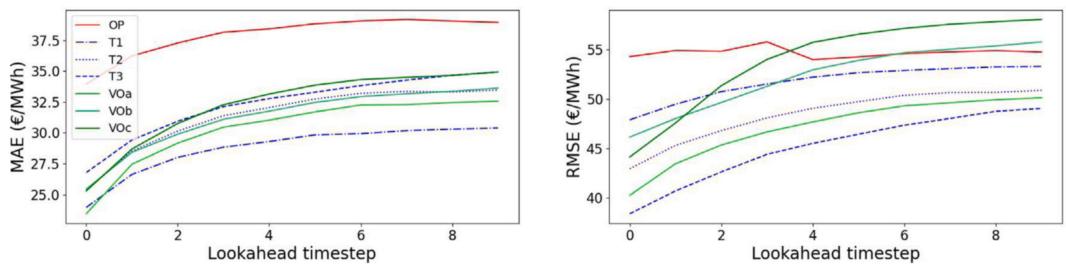

Fig. 4. Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) of predicted imbalance prices comparing seven forecasting models. The value-oriented models were selected for application G1. The evaluation comprises the test set of November and December 2020.

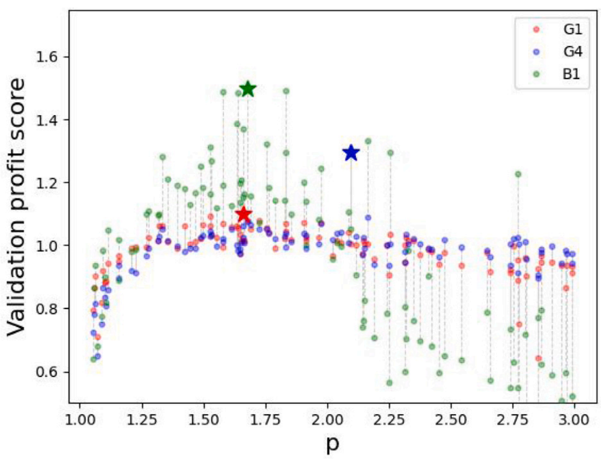

Fig. 5. Validation profit score in function of exponent p in the generalized loss function evaluated for the model VOb, comparing different applications. The forecaster with the highest performance per application is indicated with a star.

Table 3
Summary of total training time per forecaster, divided into neural network training and validation set optimization. All times are reported in seconds.

Model	T2	VOa	VOb	VOc
Avg. NN train time	2,366	2,401	2,497	2,878
Avg. optimization time	0	628	631	627
Avg. total train time	2,366	3,029	3,128	3,505

these price spreads. We have shown that traditionally-trained imbalance price forecasters do not necessarily improve ESS arbitrage profits compared to a fundamental approach to price forecasting, despite their higher statistical accuracy. Based on that insight, this paper presents a universally applicable value-oriented methodology for training time series forecasters by defining a new generalized loss function tailored to better capture the variability of electricity prices, and putting forth a hyperparameter tuning procedure that considers downstream value as a selection criterion instead of the traditional forecast accuracy. A case study of the Belgian balancing market, using actual imbalance prices, shows that the proposed model outperforms existing benchmarks in terms of ex-post out-of-sample profit from 13% to 176% depending on the specific application. Building on the findings of this study, future research could investigate the application of the loss tuning method in other rolling horizon optimization contexts, such as intra-day market participation and other energy applications involving real-time control.

CRediT authorship contribution statement

Ruben Smets: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Conceptualization. **Jean-François Toubeau:** Writing – review & editing, Validation, Supervision, Methodology, Conceptualization. **Mihaly Dolanyi:** Writing – review & editing, Investigation, Conceptualization. **Kenneth Bruninx:** Writing – review & editing, Supervision, Methodology, Investigation, Conceptualization. **Erik Delarue:** Writing – review & editing, Supervision, Resources, Project administration, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] Steilen Mike, Jörissen Ludwig. Chapter 10 - hydrogen conversion into electricity and thermal energy by fuel cells: Use of H2-systems and batteries. In: Electrochemical energy storage for renewable sources and grid balancing. Amsterdam: Elsevier; 2015, p. 143–58.
[2] Cole Wesley, Frazier A Will, Augustine Chad. Cost Projections for Utility-Scale Battery Storage: 2021 Update. Tech. rep., National Renewable Energy Laboratory; 2021.
[3] Hassler Michael. Heuristic decision rules for short-term trading of renewable energy with co-located energy storage. Comput Oper Res 2017;83:199–213.
[4] Liu Jian, Sun Xin-yue, Bo Rui, Wang Siyuan, Ou Meng. Economic dispatch for electricity merchant with energy storage and wind plant: State of charge based decision making considering market impact and uncertainties. J Energy Storage 2022;53:104816.
[5] Pandzic K, Bruninx K, Pandzic H. Managing risks faced by strategic battery storage in joint energy-reserve markets. IEEE Trans Power Syst 2021;36(5):4355–65.
[6] Khojasteh Meysam, Faria Pedro, Vale Zita. A robust model for aggregated bidding of energy storages and wind resources in the joint energy and reserve markets. Energy 2022;238:121735.
[7] Toubeau Jean-Francois, Bottieau Jeremie, De Greeve Zacharie, Vallee Francois, Bruninx Kenneth. Data-driven scheduling of energy storage in day-ahead energy and reserve markets with probabilistic guarantees on real-time delivery. IEEE Trans Power Syst 2021;36(4):2815–28.
[8] Nasrolahpour Ehsan, Kazempour Jalal, Zareipour Hamidreza, Rosehart William D. A bilevel model for participation of a storage system in energy and reserve markets. IEEE Trans Sustain Energy 2018;9(2):582–98.
[9] Wang Yubin, Zheng Yanchong, Yang Qiang. Optimal energy management of integrated energy systems for strategic participation in competitive electricity markets. Energy 2023;278:127883.
[10] Arteaga Juan, Zareipour Hamidreza. A price-maker/price-taker model for the operation of battery storage systems in electricity markets. IEEE Trans Smart Grid 2019;10(6):6912–20.
[11] Khatami Roohallah, Oikonomou Konstantinos, Parvania Masood. Look-ahead optimal participation of compressed air energy storage in day-ahead and real-time markets. IEEE Trans Sustain Energy 2019;11(2):682–92.
[12] Shi Jie, Yu Nanpeng, Gao H Oliver. Bidding strategy for wireless charging roads with energy storage in real-time electricity markets. Appl Energy 2022;327.
[13] Chen Yize, Shi Yuanyuan, Zhang Baosen. Optimal control via neural networks: A convex approach. 2018.

[14] Deng Zhuofu, Liu Chenxu, Zhu Zhiliang. Inter-hours rolling scheduling of behind-the-meter storage operating systems using electricity price forecasting based on deep convolutional neural network. Int J Electr Power Energy Syst 2021;125:106499.

[15] Finhold Elisabeth, Gärtner C, Grindel R, Heller T, Leithäuser N, Röger E, et al. Optimizing the marketing of flexibility for a virtual battery in day-ahead and balancing markets: A rolling horizon case study. Appl Energy 2023;352:121667. http://dx.doi.org/10.1016/j.apenergy.2023.121667.

[16] Smets Ruben, Bruninx Kenneth, Bottieau Jérémie, Toubeau Jean-François, Delarue Erik. Strategic implicit balancing with energy storage systems via stochastic model predictive control. IEEE Trans Energy Mark Policy Regul 2023;1–14. http://dx.doi.org/10.1109/TEMPR.2023.3267552.

[17] Semmelmann Leo, Dresselhaus Jannik, Miskiw Kim K, Ludwig Jan, Weinhardt Christof. An algorithm for modelling rolling intrinsic battery trading on the continuous intraday market. ACM SIGENERGY Energy Inform Rev 2025;4(4):163–74.

[18] Brijs Tom, Geth Frederik, De Jonghe Cedric, Belmans Ronnie. Quantifying electricity storage arbitrage opportunities in short-term electricity markets in the CWE region. J Energy Storage 2019;25:100899.

[19] Lago Jesus, Marcjasz Grzegorz, De Schutter Bart, Weron Rafał. Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. Appl Energy 2021;293:116983.

[20] Li Ji, Wang Weiqing, Yuan Zhi, Chen Jun, Xu Lei. Optimal multi-market scheduling of a sustainable photovoltaic-oriented distribution network hosting hydrogen vehicles and energy storages: A regret assessment optimization. J Energy Storage 2023;66:107489.

[21] Lin Yu, Lu Qin, Tan Bin, Yu Yuanyuan. Forecasting energy prices using a novel hybrid model with variational mode decomposition. Energy 2022;246:123366.

[22] Lago Jesus, De Ridder Fjo, De Schutter Bart. Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. Appl Energy 2018;221:386–405.

[23] Meng Anbo, Wang Peng, Zhai Guangsong, Zeng Cong, Chen Shun, Yang Xiaoyi, et al. Electricity price forecasting with high penetration of renewable energy using attention-based LSTM network trained by crisscross optimization. Energy 2022;254:124212.

[24] Zhou Siyu, Zhou Lin, Mao Mingxuan, Tai Heng-Ming, Wan Yihao. An optimized heterogeneous structure LSTM network for electricity price forecasting. IEEE Access 2019;7:108161–73.

[25] Brusaferri Alessandro, Matteucci Matteo, Portolani Pietro, Vitali Andrea. Bayesian deep learning based method for probabilistic forecast of day-ahead electricity prices. Appl Energy 2019;250:1158–75.

[26] Bottieau Jeremie, Hubert Louis, De Greve Zacharie, Vallee Francois, Toubeau Jean-Francois. Very-short-term probabilistic forecasting for a risk-aware participation in the single price imbalance settlement. IEEE Trans Power Syst 2020;35(2):1218–30.

[27] Ireshika Muhandiram Arachchige Subodha Tharangi, Kepplinger Peter. Uncertainties in model predictive control for decentralized autonomous demand side management of electric vehicles. J Energy Storage 2024;83:110194.

[28] Jadon Aryan, Patil Avinash, Jadon Shruti. A comprehensive survey of regression based loss functions for time series forecasting. 2022.

[29] Amos Brandon, Kolter J Zico. OptNet: Differentiable Optimization as a Layer in Neural Networks. Ithaca: Cornell University Library; 2021, arXiv.org.

[30] Agrawal Akshay, Amos Brandon, Barratt Shane, Boyd Stephen, Diamond Steven, Kolter J Zico. Differentiable convex optimization layers. In: Wallach H, Larochelle H, Beygelzimer A, d' Alché-Buc F, Fox E, Garnett R, editors. In: Advances in neural information processing systems, vol. 32, Curran Associates, Inc.; 2019.

[31] Han Jiayu, Yan Lei, Li Zuyi. A task-based day-ahead load forecasting model for stochastic economic dispatch. IEEE Trans Power Syst 2021;36(6):5294–304.

[32] Wilder Bryan, Dilkina Bistra, Tambe Milind. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. 2018.

[33] Donti Priya L, Amos Brandon, Kolter J Zico. Task-based end-to-end model learning in stochastic optimization. 2017.

[34] Ferber Aaron M, Wilder Bryan, Dilkina Bistra, Tambe Milind. MIPaaL: Mixed integer program as a layer. 2019, CoRR, arXiv:1907.05912.

[35] Elmachtoub Adam N, Grigas Paul. Smart "predict, then optimize". Manag Sci 2022;68(1):9–26.

[36] Sang Linwei, Xu Yinliang, Long Huan, Hu Qinran, Sun Hongbin. Electricity price prediction for energy storage system arbitrage: A decision-focused approach. IEEE Trans Smart Grid 2022;13(4):2822–32.

[37] Huang Xiaoge, Zhao Tianqiao, Huang Bin, Zhang Ziang, Yue Meng. Advancing energy system optimization via data-centric task-oriented forecasting: An application in PV-battery operation. Appl Energy 2025;378:124753.

[38] Smets Ruben. Participation of Energy Storage Systems in Short-Term Electricity Markets: Exploring the Interaction Between Optimization and Machine Learning [Ph.D. thesis], KU Leuven; 2024.

[39] Loutfi Ahmad Amine, Sun Mengtao, Loutfi Ijlal, Solibakke Per Bjarte. Empirical study of day-ahead electricity spot-price forecasting: Insights into a novel loss function for training neural networks. Appl Energy 2022;319:119182.

[40] Zhang Jialun, Wang Yi, Hug Gabriela. Cost-oriented load forecasting. Electr Power Syst Res 2022;205:107723.

[41] Matuozzo Alberto, Yoo Paul D, Provetti Alessandro. A right kind of wrong: European equity market forecasting with custom feature engineering and loss functions. Expert Syst Appl 2023;223.

[42] Wang Jiale, Zhou Yidan, Zhang Yao, Lin Fan, Wang Jianxue. Risk-averse optimal combining forecasts for renewable energy trading under CVaR assessment of forecast errors. IEEE Trans Power Syst 2023;1–13.

[43] Klæ boe Gro, Eriksrud Anders Lund, Fleten Stein-Erik. Benchmarking time series based forecasting models for electricity balancing market prices. Energy Syst (Berl Periodical) 2013;6(1):43–61.

[44] Ruder Sebastian. An overview of gradient descent optimization algorithms. 2016.

[45] Alibrahim Hussain, Ludwig Simone A. Hyperparameter optimization: Comparing genetic algorithm against grid search and Bayesian optimization. In: 2021 IEEE congress on evolutionary computation. IEEE; 2021, p. 1551–9.

[46] Bottieau Jeremie, Wang Yi, De Greve Zacharie, Vallee Francois, Toubeau Jean-Francois. Interpretable transformer model for capturing regime switching effects of real-time electricity prices. IEEE Trans Power Syst 2023;38(3):1–14.

[47] Bunn Derek W, Inekwe John N, MacGeehan David. Analysis of the fundamental predictability of prices in the british balancing market. IEEE Trans Power Syst 2021;36(2):1309–16.

[48] Garcia Maria P, Kirschen Daniel S. Forecasting system imbalance volumes in competitive electricity markets. IEEE Trans Power Syst 2006;21(1):240–8.

[49] Xu Bolun, Zhao Jinye, Zheng Tongxin, Litvinov Eugene, Kirschen Daniel S. Factoring the Cycle Aging Cost of Batteries Participating in Electricity Markets. Ithaca: Cornell University Library; 2017, arXiv.org.

[50] Elia. Downloaden van data. 2022, URL https://www.elia.be/nl/grid-data/data-download-page.