Decision-Focused Learning for Optimized Participation of an Industrial Consumer in Energy-Only and Reserve Markets

[1,2]Chloé Dupont, [1]Pietro Favaro, [1]François Vallée, [2]Bruno Francois, [1]Jean-François Toubeau
[1]*Power Systems and Markets Research Group, University of Mons, Mons, Belgium*
[2]*Univ. Lille, Arts et Metiers Institute of Technology, Centrale Lille, Junia, ULR 2697 - L2EP, F-59000 Lille, France*
Corresponding author: chloe.dupont@umons.ac.be,

***Abstract*—Many decision-making applications in modern power systems involve solving two sequential problems: (*i*) predicting unknown parameters and (*ii*) optimizing decisions under prediction uncertainty. Conventionally, the prediction model is learned by minimizing a statistical error (with respect to the true observations), without considering the impact of the predictions on the quality of decisions. In contrast, this paper aims to improve this gap by investigating *Decision-Focused Learning* wherein the optimization model is integrated into the training pipeline of the prediction model. The goal is to learn a predictor that endogenously identifies the salient features (e.g., timing of peaks, duration of valleys, etc.) that lead to the best downstream decisions. This end-to-end learning approach is applied to the day-ahead scheduling problem of a flexible consumer participating in both energy-only and reserve markets. A machine learning model is used to forecast the day-ahead energy prices, which has been trained with a value-oriented loss function, referred to as *regret* that evaluates the quality of the operation decisions. Results demonstrate the ability of *Decision-Focused Learning* to improve decision quality, leading to economic benefits for the consumer. Outcomes also reveal initializing training with a conventional least-squared-error, before using the *regret*, enhance model performance.**

***Index Terms*—Decision-Focused Learning, Optimization under Uncertainty, Power Systems, Flexibility**

I. Introduction

Energy transition and environmental objectives have raised major challenges for the actors of the energy sector. In particular, the deployment of renewable technologies has increased the uncertainty in power grids, creating the need for increasingly efficient prediction tools to optimize the operation of the system. To that end, Machine Learning (ML) is often used to predict unknown variables, which are then fed into Constrained Optimization (CO) models for decision-making, as represented in Fig. 1. Together, these tools facilitate the optimal management of power systems – including consumers, generation plants, and storage systems – by effectively handling the different sources of uncertainty (e.g., from prices and renewable generation).

This paper focuses on enhancing the training framework of the forecaster. Traditionally, a Predict-Then-Optimize (PTO) approach is adopted: the training of the ML model is performed independently from the downstream CO problem [1].

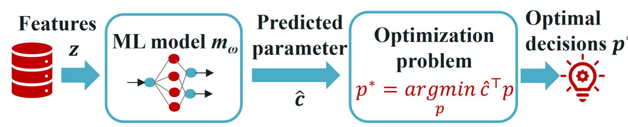

Fig. 1. Block diagram of the two stages process

However, predictions are often imperfect, and this lack of accuracy can lead to suboptimal decisions by the CO model. Traditional metrics, such as mean-squared error, assess the statistical accuracy of forecasts but often fail to capture the critical information (e.g., timing and number of peaks, delay between peaks and valleys, etc.) that lead to improved decisions. Identifying this critical information is specific to each problem and can be extremely challenging, even for experts in the field. Consequently, translating this information into an adequate statistical error metric is even more elusive.

In light of this context, we propose a Decision-Focused Learning (DFL) approach, wherein the prediction model is learned in light of the quality of the downstream decisions, which is achieved by integrating the optimization model into the training pipeline. It should be noted that, when the prediction model is trained, it can be used in the traditional two-step process of Fig. 1. This training paradigm is implemented through the definition of a specific metric called the *regret*, which directly integrates the quality of downstream decisions within the loss function. The parameters of the forecaster are learned through a gradient-based optimization.

However, the partial derivative of the *regret* with respect to the prediction model parameters is not straightforward as linear and mixed-integer linear programs have a zero gradient almost everywhere. This results in a more complex back-propagation phase. Various methods were developed to overcome the computational challenges posed by DFL and the *regret* loss function, each involving specific assumptions about the associated optimization problem [2]. Some works have managed the discontinuity introduced by the regret in the training phase through analytical smoothing for convex optimization problems [3] or by using unrolled solvers [4]. Further research implemented a smoothing by adding random perturbations and used blackbox combinatorial solvers [5] or a perturbed optimizer [6]. The above-described methods can be associated with the implementation of differentiable

979-8-3503-9042-1/24/$31.00 ©2024 IEEE

optimization layers and they can thus be introduced anywhere in a neural network architecture.

The approach adopted in this work is based on the definition of a surrogate function of the *regret*: the *Smart Predict, then Optimize* (SPO+) function defined in [7]. The SPO+ loss applies to any CO problems if the predicted parameter appears linearly in the objective function which is the case of the day-ahead prices in the scheduling problem, as detailed in Section II. Among the real-world applications of DFL, little research has been conducted in the field of power systems, and new research is therefore needed to fully investigate its potential. DFL was applied to wind power predictions for energy system cost minimization in [8]. Then, the SPO+ framework was used to learn an electricity price prediction model dedicated for energy storage system arbitrage in [9].

This paper investigates the benefits of the DFL approach for a day-ahead scheduling problem of a consumer minimizing its energy costs. A ML model is used to forecast the electricity prices of the day-ahead market. Then, a linear optimization problem is solved to schedule the energy consumption and the reserve participation of the consumer. The SPO+ loss is used to train the ML model in an end-to-end process, considering the impact of forecast inaccuracies on the downstream decisions. This paper brings the following contributions :

1) We propose a DFL framework for the day-ahead scheduling of a consumer optimizing jointly its participation to the day-ahead energy-only market and the reserve capacity market. The prices of the energy-only market are predicted by a Recurrent Neural Network (RNN) using Long-Short Term Memory (LSTM) units. In contrast to the conventional PTO approach, the forecaster, i.e., the RNN, considers the downstream optimization problem during its training thanks to the SPO+ loss, which represents an upper-bound of the regret, to lead to better decisions.
2) We show that initializing the predictor training with a conventional least-squared-error approach, before starting the Decision-Focused Learning (DFL) method, enhances model performance by reducing computation time and increasing forecast accuracy while maintaining the quality of decision-making improvements.

The remaining of the paper is organized as follows. Section II presents the optimization problem and the implementation of DFL. Then, the performances of DFL are evaluated in Section III on a realistic case-study using historical data from the Belgian electricity markets. Section IV concludes the paper and gives some perspectives for future work.

II. Methodology

The end-to-end model used by a factory owner for the day-ahead scheduling of its plant is made up of two parts: (*i*) the ML architecture used to forecast the prices for the energy-only spot market and (*ii*) the CO problem which optimizes the consumption profile and the reserve participation. We start by describing in Subsection II-A the formulation of the optimization problem. Then, in Subsection II-B, we explain how it is embedded into the learning of the forecaster thanks to the SPO+ loss.

A. Optimization problem

A factory owner wants to minimize its electricity cost by optimizing the power demand of the company and the power capacity allocated to the secondary reserve, also referred to as automatic Frequency Restoration Reserve (aFRR). The objective function of the scheduling problem (1) is made up of (*i*) the cost of purchasing electrical power p_t^{DA} for a time interval Δ_t at a price c_t^{DA} as predicted by the forecaster; (*ii*) the benefits earned by providing upward and/or downward power capacity $p_t^{\mathrm{RP}u,d}$ at a price $c_t^{\mathrm{RP}u,d}$ for the reserve. The company is assumed to be a price taker, i.e., its actions do not impact the market clearing prices. An hourly time step (t) is used to align with the market granularity and the time horizon of the problem ($\mathcal{T}$) is 24 hours.

$$\Pi = \underset{\Theta}{\mathrm{Min}} \sum_{t\in\mathcal{T}} [\underbrace{c_t^{\mathrm{DA}} p_t^{\mathrm{DA}} \Delta_t}_{(i)} - \underbrace{c_t^{\mathrm{RP}u} p_t^{\mathrm{RP}u} - c_t^{\mathrm{RP}d} p_t^{\mathrm{RP}d}}_{(ii)}] \ \forall t \quad (1)$$

The objective function (1) involves several assumptions. First, the day-ahead energy and reserve markets are jointly cleared for the 24 hours of the next day. While the day-ahead electricity prices c_t^{DA} are predicted with a forecaster, the reserve capacity procurement prices $c_t^{\mathrm{RP}u,d}$ are assumed to be perfectly known. Lastly, the reserve activation remuneration is supposed to offset exactly the operating costs associated with this reserve activation.

The feasible region Θ of the decision variables is subject to a set of market-based and technical constraints.

$$\sum_{t\in\mathcal{T}} \left(p_t^{\mathrm{DA}} - P_t^{\mathrm{exp}}\right) = 0 \quad (2)$$

$$\sum_{t\in\mathcal{T}} \left(p_t^{\mathrm{RPu}} - p_t^{\mathrm{RPd}}\right) = 0 \quad (3)$$

$$0 \le p_t^{\mathrm{RPu,d}} \le P^{\mathrm{RPmax}} \quad \forall t \quad (4)$$

$$P_t^{\mathrm{min}} + p_t^{\mathrm{RPu}} < p_t^{\mathrm{DA}} < P_t^{\mathrm{max}} - p_t^{\mathrm{RPd}} \quad \forall t \quad (5)$$

$$p_{t-1}^{\mathrm{DA}} - p_t^{\mathrm{DA}} \le \Delta^{\mathrm{down}} \quad \forall t \quad (6)$$

$$p_t^{\mathrm{DA}} - p_{t-1}^{\mathrm{DA}} \le \Delta^{up} \quad \forall t \quad (7)$$

First, the total power consumption over the day p_t^{DA} must match the expected consumption P_t^{exp} (2) which is associated with a production target of the factory that must be fulfilled. In the same perspective, the symmetrical participation in both reserve products is enforced by (3). The capacity allocated to the reserve markets is limited by P^{RPmax} (4), ensuring that the power offer complies with the technical requirements of the factory (i.e., 7.5 min for the secondary reserve). Then, the power demand p_t^{DA} must respect the flexibility bounds (P_t^{max} and P_t^{min}) while considering the reserve participation (5). Finally, the ramping capabilities (Δ^{up} and Δ^{down}) of the factory limit the changes in power demand between two consecutive time steps (6) and (7).

B. DFL implementation

In order to predict the prices of the spot market c^{DA} for the following day, a ML tool m with parameters ω is used. From a DFL perspective, any ML architecture can be used as long as the relationship between its output and its parameters is differentiable. This is not limiting since it includes all the neural networks. The general form of the prediction model can be written as :

$$\hat{c}^{\mathrm{DA}} = m_\omega(z) \tag{8}$$

The predictions $\hat{c}^{\mathrm{DA}}$ are made using past day-ahead prices z as input features. The ML model m can be trained and its parameters ω tuned with an appropriate loss function and a stochastic gradient descent. The key change introduced by DFL — compared to conventional PTO methods — lies in the definition of the loss function, sometimes referred to as the regret *r* (9). It incorporates the influence of the predicted parameters $\hat{c}^{\mathrm{DA}}$ on the decisions of the optimization problem. For the sake of simplicity, $\hat{c}$ denotes the parameters of the CO objective, including the predicted electricity prices $\hat{c}^{\mathrm{DA}}$ and the other known costs $c^{\mathrm{RPu,d}}$, while c represents the true prices of both markets. Then, $p^*(\hat{c})$ represents the optimal value of the decision variables (p_t^{DA}, $p_t^{\mathrm{RPu,d}}$) based on predicted prices and $p^*(c)$ corresponds to the optimal decisions based on true prices. Finally, $f(\cdot)$ is the objective function minimized in the optimization problem (1).

$$r\left(p^*(\hat{c}), c\right) = f\left(p^*(\hat{c}), c\right) - f\left(p^*(c), c\right) \tag{9}$$

The *regret* is defined as the degradation of the objective due to the inaccuracies on the predicted parameters $\hat{c}$ compared to the fully informed case $f(p^*(c), c)$. Fig. 2 illustrates the use of the regret *r* in a DFL approach (red) to train the ML model compared to PTO (orange), which is based on a MSE loss.

To update the ML model parameters ω via stochastic gradient descent, the partial derivative of the regret with respect to ω must be computed. As depicted in equation (10), this derivative can be expressed using the chain rule of differentiation.

$$\frac{\partial\, r\left(p^*(\hat{c}), c\right)}{\partial \omega} = \frac{\partial\, r\left(p^*(\hat{c}), c\right)}{\partial p^*(\hat{c})} \frac{\partial p^*(\hat{c})}{\partial \hat{c}} \frac{\partial \hat{c}}{\partial \omega} \tag{10}$$

The challenge of this operation lies in the second factor of the right-hand side. The derivation of the optimal decision $p^*(\hat{c})$ with respect to the predicted parameter $\hat{c}$ is not straightforward. The optimization problem generates this mapping $\hat{c} \rightarrow p^*(\hat{c})$, which may be discontinuous or non-differentiable, for instance, in linear programs. One suggestion to overcome this difficulty is to differentiate a surrogate loss function which constitutes an upper bound of the regret. Elmachtoub and Grigas defined such a function entitled "Smart Predict, then Optimize" (SPO+) [7] :

$$SPO^+(\hat{c}, c) = \max_{p \in \Theta}\{c^{\top} p - 2\hat{c}^{\top} p\} + 2\hat{c}^{\top} p^*(c) - c^{\top} p^*(c) \tag{11}$$

Minimizing the SPO+ loss corresponds to minimizing an upper bound of the *regret*. This framework applies to any CO problem whose predicted parameters appear linearly in the objective function. In this definition, Θ is a compact (i.e., closed and bounded) and convex feasible region. The SPO+ definition relies on the fact that the *regret* is impervious to the scaling of $\hat{c}$, which can be replaced by $\alpha\hat{c}\ \forall\alpha$. To update the ML model parameters, the derivative of the SPO+ function is required. To that end, a subgradient of the SPO+ loss is used :

$$-2\, p^*(2\hat{c} - c) + 2\, p^*(c) \in \frac{\partial SPO^+(\hat{c}, c)}{\partial \hat{c}} \tag{12}$$

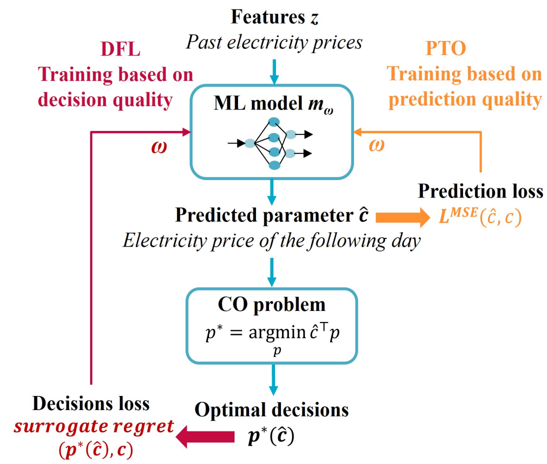

Fig. 2. Training process as decision-focus learning (red) or predict-then-optimize (orange).

III. Case Study

In this paper, a large electricity consumer aims at optimizing its power demand by predicting the electricity prices and solving a day-ahead scheduling problem. The company has a daily production target which requires a precisely defined amount of energy. Nevertheless, the hourly load can be shifted throughout the day as long as it lies within the flexibility range centered around the expected consumption profile. Subsection III-A first presents the key data of the study, then the performances of the proposed DFL framework relying on SPO+ are discussed in Subsection III-B. The conventional PTO method is used as benchmark. Finally, the relevance of warm-starting the training of the DFL model with conventional error minimization training is demonstrated empirically in Subsection III-C.

A. Data Description

The day-ahead market prices are retrieved from EPEX (European Power Exchange) SPOT for the year 2023. Over that period, electricity prices fluctuate hourly from -120 €/MWh to 330.36 €/MWh with an average of 97.27 €/MWh. Remuneration for allocating upward or downward reserve capacity is retrieved from the Belgian TSO website, Elia [10]. The upward reserve prices range from 0 €/MW to 172.11 €/MW, hovering around 39.01 €/MW on average. The downward reserve prices fluctuate between 0 €/MW and 153.17 €/MW with an average of 11.46 €/MW. The maximum power demand of the company is 50 MW and the consumption flexibility involves reducing or increasing planned consumption within a range of 20 to 30%.

The forecaster is an LSTM network, built with PyTorch and trained via stochastic gradient descent using the *Adam* optimizer [11]. The network is made up of one hidden layer of 20 neurons. The ML tool aims at predicting the electricity prices for the coming day based on observation data available at the time of prediction (i.e., past day-ahead prices), using a look-back window of 3 days. In this application, LSTM is used as it can efficiently capture temporal dependencies [12]. The optimization model is implemented in Python, using Pyomo, and is solved with *Gurobi*. The PyEPO library is used to implement SPO+ loss [13]. The simulations were conducted on a personal computer setup running Windows 11 with RAM 32Go, CPU Intel Core I7-13700HX (2.1GHz).

B. Model Performances

The performances of the model trained in a DFL fashion with the SPO+ loss are compared to a PTO approach. The latter uses a MSE loss to train the forecaster without considering the downstream optimization problem. For both methods, the dataset is randomly split into 80% of training data and 20% of test data. The quality of the decisions made by the optimization model is evaluated over the 76-day test set using a relative regret score. The regret assesses the degradation of the objective due to inaccuracies in predicting the parameter $\hat{c}$ compared to the fully informed case (i.e., perfect forecast) where the actual parameter c is known.

$$Rel\ regret(\hat{c}, c) = \frac{c^\top p^\star(\hat{c}) - c^\top p^\star(c)}{c^\top p^\star(c)} \tag{13}$$

The statistical accuracy of the predictions $\hat{c}$ made by the ML model is evaluated with the relative mean-absolute-error (Rel MAE) while the quality of the decisions p is evaluated with the relative regret. The results are reported in Table I.

First, PTO and DFL approaches are compared at their best training duration, 28 epochs for PTO and 35 for DFL. As expected, the accuracy of the DFL model, with a Rel MAE of 34.01 %, is lower than the PTO approach (20.58 %). However, the DFL generates the smallest relative regret (1.61%), reflecting enhanced decision-making compared to the PTO approach (1.74%). Therefore, DFL allows to take better decisions at the expense of the statistical accuracy of the uncertain parameter forecast. Regarding training time, the DFL is significantly longer (242 s) than the PTO (6 s). This is due to the higher complexity of the DFL since, for each sample, the optimization problem must be solved.

TABLE I
COMPARISON OF DFL AND PTO APPROACHES

Epochs	Rel MAE [%]	Rel regret [%]	Time [s]
	PTO		
5	23.68	2.65	1
15	21.54	2.14	3
25	20.87	1.85	5
28	**20.66**	**1.74**	**6**
35	20.58	1.80	7
	DFL		
5	34.65	2.82	30
15	34.68	1.86	171
25	34.28	1.65	209
35	**34.01**	**1.61**	**242**

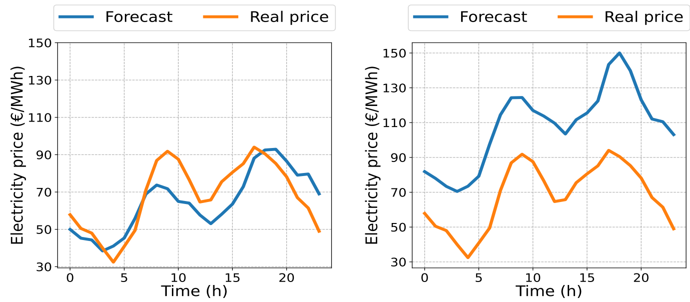

Fig. 3. Electricity price predictions using PTO (left) and DFL (right)

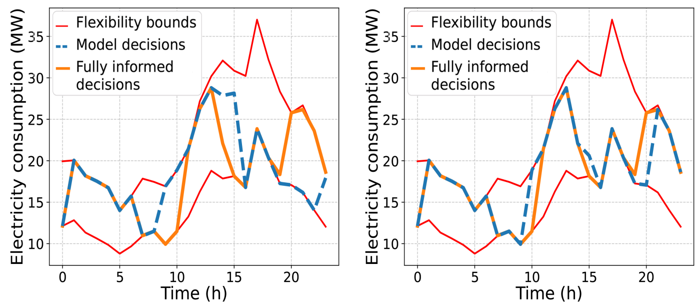

Fig. 4. Optimization problem decisions using PTO (left) and DFL (right)

Fig. 3 and 4 illustrate these observations for a particular day of the test set, i.e., December 20, 2023. On the one hand, the electricity price predictions (blue line) are more accurate with the PTO approach in Fig. 3. On the other hand, DFL better captures the changing trend (peaks and valleys) of the real price (orange line), which is essential to proper decision-making. Fig. 4 confirms that the model decisions (blue line) are closer to the optimal ones (orange line) when DFL is applied. The company can adapt its consumption within flexibility limits (red lines). For example, the day-ahead electricity price peaks around 9 am that day. This is precisely identified by the prediction tool using DFL in Fig. 3, and the downstream optimization model minimizes the company's consumption at this time, as shown in Fig. 4. This is not the case with the PTO approach, which, due to a prediction error, anticipated this decrease in the electricity price. On that day, the electricity bill associated with a perfect price forecast (i.e., oracle) would be 25,829 €. The classical PTO method results in a 513 € higher bill (1.95%) while DFL leads to 259 € (1%) more. The DFL would save 254 €, i.e. almost 1% of the PTO bill.

Table I also displays the performances of both approaches along the training process. This highlights the limited performances of DFL in the early training steps. Indeed, after 5 epochs, in addition to providing more accurate predictions, the PTO approach gives better decisions. The regret stands at 2.65 % compared to 2.82 % for DFL. Fig. 5 displays the evolution of the two training metrics (MSE for the PTO and regret for the DFL) during the training (red) and testing (blue) phases as a function of the number of epochs. These graphs illustrate the slower convergence of the regret in the early stages of training compared to the MSE loss of the

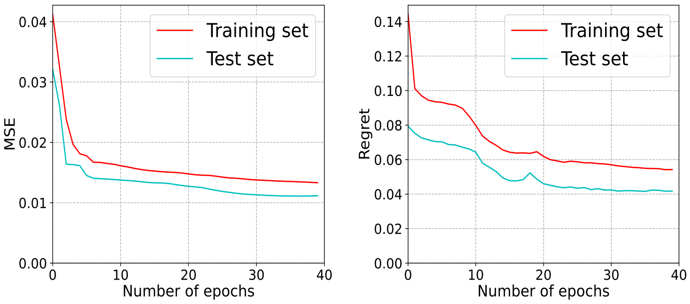

Fig. 5. Loss of PTO training (left) and regret of DFL training (right)

traditional PTO method. Table I also highlights the optimal number of epochs necessary in the training phase. The DFL method requires more epochs (35) than the PTO approach (28). In addition, as this integrated approach is more complex, the training time is considerably longer (242 s) than when training is based on a classic loss (6 s). These unfavorable indicators to DFL (Rel MAE, number of epochs and training time) pave the way for further improving the training strategy, which is discussed in Subsection III-C.

C. Impact of a Warm Start

Even though better decisions can be achieved thanks to the DFL, it suffers major drawbacks such as a decrease in the forecast precision of the uncertain parameters and a higher training time. A potential strategy to improve the model training is to perform a Warm Start (WS). Rather than randomly initializing the parameters of the ML model at the beginning of the DFL, an initial *pre-training* phase using a MSE loss can be employed to initialize the parameters before a fine-tuning performed by the DFL. To evaluate the influence of a WS on the performance of the DFL model, Table II shows the evolution of the training metrics depending on the number of epochs spent in conventional MSE-based training. The training phase using the SPO+ loss is proportionally reduced to prevent overfitting. The first contribution of the WS is a reduction in the relative MAE from an initial value of 34.01 % to 22.12 % when allocating 25 epochs to the pre-training phase. This improvement is achieved at the expense of a very slight increase (from 1.61% to 1.65%) in the regret compared to the best model using the SPO+ loss alone. However, this increase is minor and the newly obtained regret (1.65%) remains clearly below the values of the PTO approach (1.74 %). WS also positively impacts the computation times, decreasing it up to one-fourth of the original value (from 242 s to 62 s). Indeed, the WS combines the advantages of both losses: keeping a low regret reflecting optimal decisions, while improving the prediction accuracy and reducing the computation time.

TABLE II
IMPACT OF A WARM START ON DFL

WS Epochs	DFL Epochs	Rel MAE [%]	Regret [%]	Time [s]
0	35	34.01	1.61	242
5	30	29.48	1.69	181
10	25	28.37	1.66	154
15	20	25.57	1.66	128
20	15	22.89	1.64	92
25	10	22.12	1.65	62

IV. CONCLUSION

DFL is a promising concept for power systems applications. This paper presents a DFL approach for an electricity price forecaster feeding the day-ahead scheduling of a consumer participating in both energy-only and reserve markets. The company's electricity demand is flexible and, based on predictions of the day-ahead market, the factory owner can optimize its consumption and reserve participation. We use the SPO+ loss to train the ML model while considering the impact of forecast inaccuracies on the scheduling tool decisions. Compared to a classical PTO approach, the end-to-end training introduced by DFL enhances the decisions of the optimization problem to the detriment of prediction accuracy. With DFL, the forecaster focuses on the overall trend of the predicted parameters and captures its peaks and valleys accurately thus leading to better decisions. We demonstrate that warm-starting the training of the DFL model with conventional MSE training improves the model performances in terms of prediction quality and computation speed with a minor impact on the decision enhancement. Our future research will investigate other methods combining prediction and optimization tools to overcome the limiting assumptions of SPO+ loss such as the linearity of the objective function and the inability to have predicted parameters in the constraints.

REFERENCES

[1] D. den Hertog and K. Postek, "Bridging the gap between predictive and prescriptive analytics-new optimization methodology needed," 2016. [Online]. Available: https://optimization-online.org/2016/12/5779/
[2] J. Mandi, J. Kotary, et al., "Decision-focused learning: Foundations, state of the art, benchmark and future opportunities,", Report, arXiv preprint arXiv:2307.13565, 2023.
[3] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond and J. Z. Kolter, "Differentiable convex optimization layers," *Advances in neural information processing systems*, 32, 2019.
[4] J. Kotary, M. Dinh and F. Fioretto, "Backpropagation of unrolled solvers with folded optimization," Report, arXiv preprint arXiv:2301.12047, 2023.
[5] M. Vlastelica, A. Paulus, V. Musil, G. Martius and M. Rolinek, "Differentiation of blackbox combinatorial solvers", *International Conference on Learning Representations*, p.1, 2019.
[6] Q. Berthet, M. Blondel, O. Teboul, M. Cuturi, J.PI Vert, F. Bach, "Learning with differentiable perturbed optimizers," *Advances in neural information processing systems 33*, p.9508–9519, 2020
[7] A. N. Elmachtoub and P. Grigas, "Smart "predict, then optimize," *Management Science*, 68(1), p.9–26, 2022.
[8] D. Wahdany, C. Schmitt and J.L. Cremer, "More than accuracy: end-to-end wind power forecasting that optimises the energy system," *Electric Power Systems Research*, 221:109384, 2023.
[9] L. Sang, Y. Xu, H. Long, Q. Hu and H. Sun, "Electricity price prediction for energy storage system arbitrage: a Decision-Focused approach," *IEEE Transactions on Smart Grid*, vol. 13, no. 4, p.2822-2832, 2022.
[10] Open Data Elia, 2024. [Online], https://opendata.elia.be/pages/home/ [Accessed: 19-Feb-2024].
[11] D.P. Kingma and J. Ba, "Adam: a method for stochastic optimization," Report, arXiv preprint arXiv:1412.6980, 2014.
[12] E. Popovska and G. Georgieva-Tsaneva, "Day-ahead electricity price forecasting using long-short term memory recurrent neural network," *Innovative STEM Educ*, 4: 139–148, 2022.
[13] B. Tang and E.B. Khalil, "Pyepo: A pytorch-based end-toend predict-then-optimize library for linear and integer programming," Report, arXiv preprint arXiv:2206.14234, 2022.