
Price-Aware Deep Learning for Electricity Markets

Vladimir Dvorkin

Massachusetts Institute of Technology
Cambridge, MA 02109
dvorkin@mit.edu

Ferdinando Fioretto

University of Virginia
Charlottesville, VA 22903
fioretto@virginia.edu

Abstract

While deep learning gradually penetrates operational planning of power systems, its inherent prediction errors may significantly affect electricity prices. This paper examines how prediction errors propagate into electricity prices, revealing notable pricing errors and their spatial disparity in congested power systems. To improve fairness, we propose to embed electricity market-clearing optimization as a deep learning layer. Differentiating through this layer allows for balancing between prediction and pricing errors, as oppose to minimizing prediction errors alone. This layer implicitly optimizes fairness and controls the spatial distribution of price errors across the system. We showcase the price-aware deep learning in the nexus of wind power forecasting and short-term electricity market clearing.

1 Introduction

Addressing the gap between electricity market clearing and power system operations with large shares of renewables, recent literature proposes various deep learning tools for forecasting weather-dependent generation and loads, set-points of conventional generation, and active sets of network constraints, among other applications. Deep learning models learn complex physical processes from an abundance of operational data and instantly map contextual features into informed inputs for market clearing [1], and even directly estimate market-clearing outcomes from data [2]. However well-informed and tuned these models may be, they are still prone to prediction errors that affect market outcomes for electricity producers and consumers.

This paper uncovers the propagation of prediction errors of deep learning into locational marginal prices – the ultimate indicator defining electricity market outcomes across the network – and demonstrates spatial disparity of these errors across the range of power system benchmark systems. To reduce the errors and improve fairness in electricity market applications, we propose to directly incorporate market-clearing optimization as a deep learning layer to inform predictions on pricing errors. For clarity, we use wind power forecasting as the deep learning task, but other prediction tasks may benefit from the same approach. Our results show that even small errors cause spatial disparities in congested systems, with usually a small set of buses contributing the most to the overall disparity. While transmission investments would ultimately resolve spatial disparities, informing predictions with market-clearing outcomes already improves fairness.

2 Price-Awareness for Fair Deep Learning

Consider a dataset $\{(\varphi_i, w_i)\}_{i=1}^m$ of m operational records, where the i^{th} record includes a power output w_i and the associated feature vector φ_i , collecting weather and turbine data, such as wind speeds and blade pitch angles. The dataset is used to train a deep learning model DeepWP mapping features to wind power outputs, i.e., $\text{DeepWP}(\varphi) = \hat{w}$. DeepWP is a neural network with a fully connected feedforward architecture, and we train it to minimize the prediction error $\|\hat{w} - w\|$. While standard, this training approach is unaware of the effects of prediction errors on electricity prices.

2.1 From Prediction to Electricity Price Errors

To study the impact of prediction errors on electricity prices, we work with the electricity market clearing based on a DC optimal power flow (OPF) problem [1, 2]. Given a wind power forecast $\hat{w} \in \mathbb{R}^n$, the goal is to identify the conventional generation dispatch $p \in \mathbb{R}^n$ which satisfies electric loads $d \in \mathbb{R}_+^n$ at minimum dispatch costs. Generators produce within technical limits $p, \bar{p} \in \mathbb{R}_+^n$ incurring costs, described by a quadratic function with coefficients $c \in \mathbb{R}_+^n$ and $C \in \mathbb{S}_+^n$. The power flows are computed using a matrix $F \in \mathbb{R}^{e \times n}$ of power transfer distribution factors, i.e., the flows are $f = F(p + \hat{w} - d) \in \mathbb{R}^e$ limited by $\bar{f} \in \mathbb{R}_+^e$. The DC-OPF optimization takes the form:

$$\underset{p \leq p \leq \bar{p}}{\text{minimize}} \quad p^\top C p + c^\top p \quad (1a)$$

$$\text{subject to} \quad \mathbb{1}^\top (p + \hat{w} - d) = 0 : \hat{\lambda}_b, \quad (1b)$$

$$|F(p + \hat{w} - d)| \leq \bar{f} : \hat{\lambda}_{\bar{f}}, \hat{\lambda}_{\underline{f}}, \quad (1c)$$

which minimizes total costs subject to power balance, and generation and transmission limits.

The locational marginal prices (LMPs), induced on a particular wind power forecast $\hat{w}$, are derived from the optimal dual variables associated with power balance and power flow constraints:

$$\pi(\hat{w}) = \hat{\lambda}_b \cdot \mathbb{1} - F^\top (\hat{\lambda}_{\bar{f}} - \hat{\lambda}_{\underline{f}}) \in \mathbb{R}^n, \quad (2)$$

where the first term is the system-wide price adjusted by the second term due to congestion.

Remark 1 (Uniqueness). *If convex problem (1) is feasible for some forecast $\hat{w}$, then its primal and dual solutions are unique thanks to strictly monotone objective function (1) [3]. Hence, electricity price $\pi(\hat{w})$ is also unique w.r.t. forecast $\hat{w}$.*

Hence, the LMP error, defined as $\delta\pi = \pi(\hat{w}) - \pi(w) \in \mathbb{R}^n$, is also unique w.r.t. forecast $\hat{w}$.

Two key observations arise from analyzing Eq. (2). The first relates to disparities due to congestion.

Observation 1 (Spatial disparity). *In congested networks, for which $\mathbb{1}^\top (\hat{\lambda}_{\bar{f}} + \hat{\lambda}_{\underline{f}}) > 0$, the price error at bus i is proportional to the i^{th} column of matrix F of power transfer distribution factors.*

Hence, the same prediction error has a disparate effect on electricity prices, depending on location. This effect under DeepWP predictions is illustrated in Fig. 2 (left plane), where LMP errors vary in the range from -4.2 \$/MWh to $+0.9$ \$/MWh on average.

The second observation identifies a unique system bus, reference bus r , with small price errors.

Observation 2 (Reference bus). *Since the r^{th} column of F is all zeros, the price error at the reference bus only includes the error of the system-wide term in (2).*

The spatial disparities is measured using the notion of α -fairness [4]:

$$\alpha = \max_{i \in \{1, \dots, n\}} \|\mathbb{E}[\|\delta\pi_i\|] - \mathbb{E}[\|\delta\pi_r\|]\|. \quad (3)$$

where the expectation is w.r.t. the dataset distribution. Parameter α is the fairness bound, with smaller values denoting stronger fairness. While the original definition takes the maximum across all pairs of individuals (buses), we use bus r as a reference, as it only includes the error of the system-wide price.

2.2 Electricity Market Clearing as a Neural Network Layer

To minimize the impact of prediction errors on electricity prices, we develop a new deep learning architecture DeepWP+, depicted in Fig. 1. Relative to DeepWP, the new architecture includes a market-clearing optimization layer which computes electricity prices in response to predictions and passes them on to the loss function, including both wind power prediction and price errors.

Directly acting on problem (1), however, is computationally challenging: it requires solving the problem and then differentiating through the large set of Karush–Kuhn–Tucker conditions. We propose to work with the dual optimization instead. For brevity, we rewrite problem (1) as

$$\underset{p}{\text{minimize}} \quad p^\top C p + c^\top p \quad (4a)$$

$$\text{subject to} \quad A p \geq b(\hat{w}) : \hat{\lambda}, \quad (4b)$$




Figure 1: The standard DeepWP and proposed DeepWP+ learning architectures.

Table 1: Wind power prediction and LMP errors under conventional (DeepWP) and price-aware (DeepWP+) deep learning models

case	wind power data			DeepWP						DeepWP+					
	bus	capacity	$\bar{f}$ -scale	RMSE($\hat{w}$)	RMSE($\hat{\pi}$)	CVaR($\hat{\pi}$)	α -value			RMSE($\hat{w}$)	RMSE($\hat{\pi}$)	CVaR($\hat{\pi}$)	α -value		
		MW	[p.u.]	MWh	\$/MWh	\$/MWh	\$/MWh	MWh	gain	\$/MWh	gain	\$/MWh	gain	\$/MWh	gain
14_ieee	14	100	1.00	0.35	0.62	1.52	0	0.35	+0.6%	0.61	-0.6%	1.50	-0.8%	0	—
57_ieee	38	600	0.60	2.31	11.03	34.64	32.08	2.60	+11.2%	10.72	-2.9%	33.59	-3.1%	30.92	-3.8%
24_ieee	15	1,000	0.75	4.08	8.62	37.70	27.48	4.51	+9.6%	8.33	-3.5%	36.35	-3.7%	26.26	-4.6%
39_epri	6	1,500	0.70	5.94	11.15	31.21	17.53	6.43	+7.6%	10.19	-9.4%	28.02	-11.4%	15.84	-10.7%
73_ieee	41	1,000	0.80	4.02	5.12	16.21	32.83	5.51	+26.9%	4.24	-20.8%	13.41	-20.9%	26.63	-23.3%
118_ieee	37	500	0.75	2.29	3.59	11.32	17.91	2.60	+12.1%	2.88	-24.7%	9.06	-25.0%	14.09	-27.2%

where the feasible region is parameterized by forecast $\hat{w}$. The dual problem takes the form:

$$\underset{\lambda \geq 0}{\text{maximize}} \quad q(\hat{w})^\top \lambda - \lambda^\top Q \lambda, \quad (5)$$

where $q(\hat{w}) = AC^{-1}c + b(\hat{w})$ and $Q = AC^{-1}A^\top$. Since the two problems are strictly convex and concave, respectively, the strong duality holds, such that the LMPs can be extracted from the dual problem using expression (2). A virtue of the dual problem (5) is that it only includes the non-negativity constraints; it is thus simpler to solve and differentiate through than its primal counterpart in (4).

The market clearing as an optimization layer is implemented using `DiffOpt.jl` – a library for differentiating through the solution of optimization problems in Julia [5].

3 Numerical Results and Discussion

For numerical experiments, we use power system benchmark systems from [6]. In each system, we install one wind farm and uniformly scale transmission capacity to provoke congestion, as shown in Table 1. To ensure fair comparison, we use the same wind power forecasting data¹ across all systems. The data includes wind power output as a function of wind speed, wind direction and blade pitch angle features. We independently sample 1,000 scenarios for training and testing.

The DeepWP architecture consists of four hidden layers with 30 neurons each, all using ReLU as activation functions. The training consists of three stages, all using ADAM optimizer: pre-training with 500 epochs and learning rate $1e-4$, then 1000 epochs with learning rate $5e-5$, and final 100 epochs with learning rate $5e-6$. The training of DeepWP+ starts from the 501st epoch using the parameters of the pre-trained DeepWP model obtained at the first stage. As neural network parameters are initialized at random, we report the average results across 100 trained models.

We discuss the results using the root mean square error of wind power prediction, denoted $\text{RMSE}(\hat{w})$, that of prices, denoted $\text{RMSE}(\hat{\pi})$, the root mean square error across 10% of the worst-case scenarios, denoted $\text{CVaR}(\hat{\pi})$, and α -fairness bound.

¹<https://www.kaggle.com/datasets/theforcecoder/wind-power-forecasting>




Figure 2: Projection of DeepWP and DeepWP+ wind power predictions errors on locational marginal price errors in the IEEE 118-bus system.

Table 2: Average CPU time per training epoch of DeepWP+ architecture

case	14_ieee	24_ieee	39_epri	57_ieee	73_ieee	118_ieee
time [min.]	0.06	0.26	0.22	0.48	2.03	6.28

Table 1 provides a statistical summary for different systems. Under prediction errors of the price-agnostic DeepWP model, LMP errors vary between 0.62 and 11.15 \$/MWh. This range, obtained using the *same* wind power dataset, is explained by at least two factors. First, all systems have unique cost structures: the average LMP spans the range from 42.8 (14_ieee) to 149.3 (39_epri) \$/MWh, for the median wind power output. Second, congested networks prone to larger errors. This is particularly evident from the 14_ieee system, where no wind power output scenario causes congestion. The zero fairness bound α for this system also supports this observation.

The application of the price-aware DeepWP+ model, on the other hand, demonstrates the reduction of LMP errors, varying from 0.6 to 24.7% relative to the DeepWP model. Even larger reductions are observed across 10% of the worst-case wind scenarios, measured by $\text{CVaR}(\hat{\pi})$. For example, in the 39_epri case, while the average LMP error is reduced by 0.96 \$/MWh, the worst-case error decreases by 3.19 \$/MWh. Although the focus of DeepWP+ is on price errors, it implicitly minimized the disparity of LMP errors, as shown by the reduced fairness bound α across all congested systems. Figure 2 depicts this effect for the most geographically distributed 118_ieee system, where the fairness bound α is reduced by 27.2% thanks to price-aware predictions. Observe, while the majority of buses demonstrate near zero price errors, it is a small set of buses contributing to α —fairness statistic the most. Importantly, the benefits of DeepWP+ do come at the expense of increasing prediction errors. However, all DeepWP+ predictions are feasible and yield competitive and fairer—in the α —fairness sense—electricity prices.

The price error reduction and fairness improvements are due to the bias which DeepWP+ introduces during the training procedure. Figure 3 visualizes how the bias is introduced starting from the 501th epoch, which diverts the training towards the desired result, i.e., to minimal prediction or minimal price error. Finally, we remark that incorporating the electricity market as an optimization layer in deep learning increases the computational burden of the training procedure. Table 2 reports CPU times per training epoch, which tends to increase in the size of the system.

Overall, our results show that embedding market clearing as a deep learning layer informs predictions on market outcomes and improves algorithmic fairness in electricity markets.

Acknowledgements

Vladimir Dvorkin is supported by the Marie Skłodowska-Curie Actions COFUND Postdoctoral Program, Grant Agreement No. 101034297 – project Learning ORDER; Ferdinando Fioretto is supported by National Science Foundation grants No. 2242931, 2232054, 2143706, and 2128426.




Figure 3: Prediction (top) and locational marginal price (bottom) errors in the 39-bus EPRI system during the training of DeepWP and DeepWP+ models.

References

- [1] Robert Ferrando, Laurent Pagnier, Robert Mieth, Zhirui Liang, Yury Dvorkin, Daniel Bienstock, and Michael Chertkov. A physics-informed machine learning for electricity markets: A NYISO case study. *arXiv preprint arXiv:2304.00062*, 2023.
- [2] Shaohui Liu, Chengyang Wu, and Hao Zhu. Topology-aware graph neural networks for learning feasible and adaptive ac-opf solutions. *IEEE Transactions on Power Systems*, 2022.
- [3] Vladimir Dvorkin Jr, Jalal Kazempour, and Pierre Pinson. Electricity market equilibrium under information asymmetry. *Operations Research Letters*, 47(6):521–526, 2019.
- [4] Cuong Tran, Ferdinando Fioretto, Pascal Van Hentenryck, and Zhiyan Yao. Decision making with differential privacy under a fairness lens. In *IJCAI*, pages 560–566, 2021.
- [5] Akshay Sharma, Mathieu Besançon, Joaquim Dias Garcia, and Benoît Legat. Flexible differentiable optimization via model transformations. *arXiv preprint arXiv:2206.06135*, 2022.
- [6] Carleton Coffrin, Russell Bent, Kaarthik Sundar, Yeesian Ng, and Miles Lubin. Powermodels. jl: An open-source framework for exploring power flow formulations. In *2018 Power Systems Computation Conference (PSCC)*, pages 1–8, 2018.