

A Physics-Informed Machine Learning for Electricity Markets: A NYISO Case Study

Robert Ferrando, Laurent Pagnier, Robert Mieth, Zhirui Liang, Yury Dvorkin, Daniel Bienstock and Michael Chertkov

Abstract—This paper addresses the challenge of efficiently solving the optimal power flow problem in real-time electricity markets. The proposed solution, named Physics-Informed Market-Aware Active Set learning OPF (PIMA-AS-OPF), leverages physical constraints and market properties to ensure physical and economic feasibility of market-clearing outcomes. Specifically, PIMA-AS-OPF employs the active set learning technique and expands its capabilities to account for curtailment in load or renewable power generation, which is a common challenge in real-world power systems. The core of PIMA-AS-OPF is a fully-connected neural network that takes the net load and the system topology as input. The outputs of this neural network include active constraints such as saturated generators and transmission lines, as well as non-zero load shedding and wind curtailments. These outputs allow for reducing the original market-clearing optimization to a system of linear equations, which can be solved efficiently and yield both the dispatch decisions and the locational marginal prices (LMPs). The dispatch decisions and LMPs are then tested for their feasibility with respect to the requirements for efficient market-clearing results. The accuracy and scalability of the proposed method is tested on a realistic 1814-bus NYISO system with current and future renewable energy penetration levels.

Index Terms—Active Set Learning, Physics-Informed Machine Learning, Optimal Power Flow, Locational Marginal Prices

I. INTRODUCTION

US electricity markets are cleared in a two-stage procedure, including the day-ahead market (DAM) and the real-time market (RTM). The DAM clearing is based on security-constrained unit commitment (SCUC), which determines the least-cost generator schedule for the next day. SCUC considers day-ahead generation offers, load bids, as well as the forecast of load and renewable energy source (RES) generation [1]. The commitment of generators, with the notable exception of flexible, gas-fired units, is predominantly determined in the DAM, but the final RTM production schedule may diverge from the anticipated DAM production schedule due to changes in operating conditions. The cornerstone of both the DAM and RTM clearing is the fast solution of a DC optimal power flow (DC-OPF) problem [2] or one of its variants [3]. Naturally, RTM clears more frequently than DAM and requires the DC-OPF problem to be solved multiple times for each 5- to 15-minute operating interval. Additionally, the drastic growth of RES, especially the recent roll-out of large-scale offshore wind power plants [4], has increased the uncertainty in power system scheduling, motivating system and market operators to perform scenario-based uncertainty analyses. These analyses substantially increase the number of DC-OPF-based instances that must be solved within each 5- or 10-minute intervals. This

motivates modifying the current ISO procedures to solve DC-OPF more efficiently.

This paper proposes a learning-based algorithm named Physics-Informed Market-Aware Active Set learning OPF, or **PIMA-AS-OPF**. The PIMA-AS-OPF leverages the physical constraints in the DC-OPF model and produces market-clearing outcomes that comply with market design principles, e.g. revenue adequacy and cost recovery, to ensure both physical and economic feasibility of its outcomes. The accuracy and scalability of the proposed method is tested in a realistic 1814-bus New York Independent System Operator (NYISO) transmission network model, which accounts for the current RES and future higher RES penetration levels.

A. Literature Review

Machine learning (ML) has been leveraged to expedite the solution of many power systems optimization problems [9], including DC-OPF [5]–[8], AC-OPF [9]–[13], and security-constrained economic dispatch (SCED) [14], [15]. The ML techniques behind these approaches include logistic regression [10], random forest [9], [10], auto-encoders [14], and stacked extreme learning machine [12]. The most popular technique is the neural network (NN) framework. The NN models designed for solving the power system scheduling models include fully-connected neural networks (FCNN) [6], [15], input convex neural networks (ICNN) [5], physics-informed neural networks (PINN) [13], [16], convolutional neural networks (CNN) [11], graph neural networks (GNN) [11], [17], and deep neural networks (DNN) [14], [18].

There are two common ways to construct ML approaches for solving optimization problems. The intuitive way is to train a ML model to predict the solution of the original problem directly based on the inputs. For example, [17] proposes a GNN model which takes electricity demands as the input and gives real-time electricity prices as the output. An alternative way is using ML to solve a sub-problem of the original optimization problem, thereby reducing the computational difficulty of the original problem. A common sub-problem is *active set learning*, in which the active inequality constraints of an optimization problem are learned via classification. Reducing the number of constraints under consideration is a key step towards dimension reduction of the optimization problem, and motivates the classification of active constraints, so that optimization theory can be leveraged to alleviate the computational complexity of the problem at hand.

The active set learning technique has been widely used in solving the power system scheduling problem [6], [8], [15],

[19], [20]. For example, [19] proposes using classification algorithms, specifically neural net classifiers, to learn the mapping between uncertainty in electricity demands and the active set of constraints in the DC-OPF model, which enhances computational efficiency. Similarly, [6] constructs a NN that predicts the optimal generation cost from the load at each bus, subsequently differentiating the cost with respect to load. The resulting values dual variables (typically interpreted as LMPs) are used to learn the active constraints, from which a system of linear equations equivalent to the original DC-OPF problem can be derived. The current work in [6], [8], [19] sets a precedent for the development and implementation of the active set learning technique on larger systems and more complicated problems, such as AC-OPF and SCED. Inspired by the observation that most of the generators are dispatched to their upper or lower limits, [15] develops a “classification-then-regression” scheme to solve SCED, i.e., first identifying the generators saturated at their bounds, and then using regression to complete the primal solution with the dispatch of the “free” generators. The proposed architecture achieves relative errors below 1% in numerical experiments on a synthetic French transmission system with realistic Midwest ISO data, demonstrating its potential for real-time energy markets. Moreover, [14] treats SCED, using stacked denoising autoencoders (SDAEs) to learn the active constraints from “system operating conditions,” adapting to changes in grid topology using transfer learning. As reported in [21], one of the winning teams of the ARPA-E GO competition also employed active set learning in their proposed method. However, as the number of features (buses, lines, etc.) of the power grid increases, so does the number of active constraints and the challenge of the underlying ML problem. To address this, [20] proposes an algorithm to learn active constraints irrespective of the optimization problem formulation or its inherent uncertainty. The algorithm in [20] offers “theoretical performance guarantees” and can quickly converge when there are few active sets to be learned. However, most relevant papers, e.g., [6], [10], [17], [19], only evaluate the accuracy of the forecast results (such as the dispatch or prices in electricity markets) using the difference between the forecast and true values, while the feasibility of the forecast results in the actual market setting is ignored. For example, if the power suppliers cannot recover their power generation costs with the forecast real-time energy prices, then the prices are infeasible even if the forecast errors are small in general. On the other hand, most of the DC-OPF models studied in previous work, e.g., [5]–[8], ignore the possible curtailment in load or RES power, which is a common issue in real-world power systems. In addition, as pointed out in [15], most published case studies are for “small academic test systems,” such as IEEE 118-bus system [5], IEEE 300-bus system [22], 500-bus transmission network [9], or 661-bus utility system [14]. These systems are too small to assess scalability of the algorithms therein to larger operational transmission grids. The work in [15] does use a realistically sized test system, but with spatially inconsistent data sets to model RES.

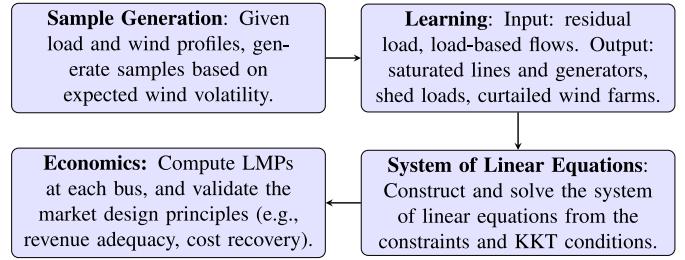


Fig. 1. A summary of the PIMA-AS-OPF procedure.

B. Outline of the paper

This paper proposes a ML-based optimization algorithm, PIMA-AS-OPF, to expedite the solution of DC-OPF in a principled manner. A schematic for the proposed PIMA-AS-OPF approach is shown in Fig. 1. First, the day-ahead security constrained unit commitment (SCUC) problem is solved. The real-time DC-OPF problem is then solved to generate samples, with the SCUC solution assumed throughout. For each hour of the next day, the second stage trains specialized ML models (fully-connected neural networks). The hourly resolution matches market operating practices and aids the reduction of computational complexity, as all real-time problems solved during an hour rely on the same SCUC solution [15]. The NN input is the net load (i.e., demand minus RES injection) and the power flow across transmission lines. The NN outputs include the set of active constraints (i.e., constraints that are active in the optimal solution, this identifying saturated generators and transmission lines), as well as the non-zero load and wind curtailments. Third, using the NN outputs, the original DC-OPF problem is reduced to a linear system of equations, which can be solved efficiently. Fourth, based on the solution from the last step, the LMPs can be computed easily. The resulting LMPs are tested with respect to the requirements of efficient market clearing results to ensure their feasibility. The training process is performed day-ahead so that the trained models can be used during the operating day in real time to solve DC-OPF based on a large number of possible configurations given updated forecasts and scenarios. Table I compares the proposed PIMA-AS-OPF with three relevant ML approaches used to solve DC-OPF. The proposed approach aligns with the theme of *physics-informed* machine learning [23] by directly taking advantage of the physics of the transmission grid and structure of the presently formulated real-time optimization procedure to construct an ML algorithm. Moreover, by testing the results against the crucial market properties such as revenue adequacy and cost recovery, the proposed method is also feasible for market implementations.

II. PRELIMINARIES

This section describes the DC-OPF formulation, presents the Karush-Kuhn-Tucker (KKT) conditions of the DC-OPF formulation as used in the PIMA-AS-OPF algorithm, derives the expression for LMPs, and then clarifies the crucial properties of an efficient market.

TABLE I
COMPARISON OF EXISTING APPROACHES AND THE PROPOSED PIMA-AS-OPF APPROACH.

Article	Task	Methodology	Extension in PIMA-AS-OPF
Deka et al. (2019) [19]	Efficient solution of DC-OPF.	Learn the active set of constraints for the optimal DC-OPF solution.	Substitute ML with solving a system of linear equations when possible, and solve dual problem.
Liu et al. (2021) [17]	Topology-aware solution of dual DC-OPF problem with adherence to constraints.	Solve DC-OPF using graph NN in order to predict LMPs.	Solve both primal and dual problem, ensure LMPs adhere to financial coherency.
Chen et al. (2022) [6]	Efficient learning of dual DC-OPF solution and primal objective.	Use NN to learn dual variables and LMPs.	Recover objective algebraically, check market conditions, validate on large, realistic system.

A. DC-OPF Formulation

Consider the following single-period DC-OPF model, which allows for load shedding and wind curtailment and computes power flows using the Power Transfer Distribution Factor (PTDF) matrix Φ :

$$\min_{p_g, s_i, c_w} \sum_g \left(c_g^{(2)} p_g^2 + c_g^{(1)} p_g \right) + \sum_i c_i^{\text{LNS}} s_i + \sum_w c_w^{\text{RES}} c_w, \quad (1a)$$

$$\text{s.t. } \sum_i (\ell_i - s_i) = \sum_g p_g - \sum_w c_w, \quad (1b)$$

$$\forall k \in L: -f_k^{\max} \leq \sum_i \Phi_{ki} \left[\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right] \leq f_k^{\max}, \quad (1c)$$

$$\forall g \in G: p_g^{\min} \leq p_g \leq p_g^{\max}, \quad (1d)$$

$$\forall i \in B: s_i \geq 0, \quad (1e)$$

$$\forall w \in W: c_w \geq 0, \quad (1f)$$

where B , L , G , and W are the sets of nodes, lines, conventional generators, and wind farms in the system, respectively, G_i and W_i are the sets of generators and wind farms connected to node i , ℓ_i is the forecast net load (i.e., forecast demand minus forecast renewable injections) at node i , f_k^{\max} is the thermal capacity of line k , p_g is the (active) power output of generator g , p_g^{\min} and p_g^{\max} are the upper and lower bounds of p_g , and s_i and c_w are slack variables representing the curtailed load at node i and the curtailed wind power at wind farm w . Objective (1a) minimizes the cost of power generation plus the penalties for load shedding and wind curtailment, where $c_g^{(2)}$ and $c_g^{(1)}$ are parameters in the quadratic cost function of generator g , and c_i^{LNS} and c_w^{RES} are the unit cost of curtailed load and wind, respectively. Note that the slack variables s_i and c_w should be capped by the forecast load and wind power, respectively, but those constraints are omitted for simplicity of the following manipulations since they are unlikely to be active. Adding these linear constraints will not affect the underlying method, as well as its computational complexity and applicability.

Solving the formulation (1) effectively requires assuming that the unit commitment decisions on all conventional generators are known.

The proposed PIMA-AS-OPF algorithm can solve a number of real-time DC-OPF instances, each with their own configura-

tions of wind power generation, therefore providing stochastic insights from a set of resulting solutions.

The DC-OPF formulation in (1) can be re-stated as the following Lagrangian:

$$\min_{\mathbf{p}} \max_{\mathbf{d}} \mathcal{L}(\mathbf{p}, \mathbf{d}) \quad \text{s.t. Eq. (1e) and (1f), } \mathbf{d} \setminus \lambda \geq 0, \quad (2)$$

$$\begin{aligned} \mathcal{L}(\mathbf{p}, \mathbf{d}) := & \sum_g \left(c_g^{(2)} p_g^2 + c_g^{(1)} p_g \right) + \sum_i c_i^{\text{LNS}} s_i \\ & + \sum_w c_w^{\text{RES}} c_w + \lambda \left[\sum_i (\ell_i - s_i) - \sum_g p_g + \sum_w c_w \right] \\ & + \sum_k \bar{\mu}_k \left[\sum_i \Phi_{ki} \left(\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right) - f_k^{\max} \right] \\ & - \sum_k \underline{\mu}_k \left[f_k^{\max} + \sum_i \Phi_{ki} \left(\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right) \right] \\ & + \sum_g \left[\bar{\alpha}_g (p_g - p_g^{\max}) + \underline{\alpha}_g (p_g^{\min} - p_g) \right], \end{aligned} \quad (3)$$

where $\mathbf{p} := (p_g, s_i, c_w)$ denotes the vector of primal variables, and $\mathbf{d} := (\lambda, \bar{\mu}_k, \underline{\mu}_k, \bar{\alpha}_g, \underline{\alpha}_g)$ denotes the vector of Lagrange multipliers associated with respective primal constraints.

B. KKT Conditions of DC-OPF

The PIMA-AS-OPF algorithm, described in Section III, relies on the KKT conditions of the convex optimization in (2). The KKT conditions, which consist of *complementary slackness* and *stationarity at the optimum*, are necessary and sufficient conditions for the solution of a convex optimization problem to be optimal [24]. Complementary slackness states that a non-zero Lagrange multiplier exists only if the corresponding inequality constraint is active, i.e., satisfied with equality. Stationarity at the optimum refers to the partial derivative of the Lagrangian being zero when evaluated at the optimal solution. Accordingly, the KKT conditions of (2) are given as follows:

Complementary Slackness:

$$\forall k: \bar{\mu}_k \left[\sum_i \Phi_{ki} \left(\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right) - f_k^{\max} \right] = 0, \quad (4a)$$

$$\forall k: \underline{\mu}_k \left[f_k^{\max} + \sum_i \Phi_{ki} \left(\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right) \right] = 0, \quad (4b)$$

$$\forall g: \bar{\alpha}_g (p_g - p_g^{\max}) = 0, \quad (4c)$$

$$\forall g: \underline{\alpha}_g (p_g^{\min} - p_g) = 0, \quad (4d)$$

Stationarity:

$$\forall g: 2c_g^{(2)}p_g + c_g^{(1)} - \lambda + \sum_k \Phi_{ki}(\bar{\mu}_k - \underline{\mu}_k) + \bar{\alpha}_g - \underline{\alpha}_g = 0, \quad (4e)$$

$$\forall i: c_i^{\text{LNS}} - \lambda + \sum_k \Phi_{ki}(\bar{\mu}_k - \underline{\mu}_k) = 0, \quad (4f)$$

$$\forall w: c_w^{\text{RES}} + \lambda + \sum_k \Phi_{ki}(\underline{\mu}_k - \bar{\mu}_k) = 0, \quad (4g)$$

C. LMPs and Market Design Properties

LMPs in electricity markets reflect the cost of producing, as well as the price of consuming, an additional MWh of electricity at a given bus. The LMP at node i can be derived by differentiating the Lagrangian in Eq. (3) with respect to net load ℓ_i , as:

$$\forall i: \text{LMP}_i := \frac{\partial \mathcal{L}}{\partial \ell_i} = \lambda + \sum_k \Phi_{ki}(\underline{\mu}_k - \bar{\mu}_k). \quad (5)$$

Specifically, λ is the LMP at the reference node, and it is also the LMP at all nodes when the transmission capacity is unlimited, i.e., when $\underline{\mu}_k$ and $\bar{\mu}_k$ are zero.

LMPs should have the following critical properties to ensure sound electricity market operations [25]:

1) *Revenue Adequacy*: From the perspective of the system operator, the total payment made by the market to the producers must be covered by the total payment made to the market by consumers. In addition, recall the system operator must pay the designated penalties for load shedding and wind curtailment. While not directly enforced during real-time operations, these penalties represent the unanticipated cost to the system operator, e.g., of having to rely more heavily on traditional generation technologies when wind power is curtailed. Thus, the revenue adequacy condition is formulated as:

$$\begin{aligned} & \sum_g \text{LMP}_{i(g)} p_g + \sum_i c_i^{\text{LNS}} s_i + \sum_w c_w^{\text{RES}} c_w \\ & \leq \sum_i \text{LMP}_i (\ell_i - s_i). \end{aligned} \quad (6)$$

2) *Cost Recovery*: From the perspective of individual generators, market participants must recover their operating costs based on the resulting LMPs, i.e., the payment made to each generator must be sufficient to recover their operating costs. Therefore, the cost recovery condition is formulated as follows:

$$\forall g: 2c_g^{(2)}p_g^2 + c_g^{(1)}p_g \leq \text{LMP}_{i(g)} p_g. \quad (7)$$

For generators which do not reach their upper or lower limits (so-called marginal generators), Eq. (7) holds with equality.

3) *Efficiency via Strong Duality*: Since DC-OPF as stated in Eq. (1) is a quadratic, convex optimization problem with all affine constraints, the strong duality is guaranteed, i.e., the primal optimal objective and the dual optimal objective are equal. In the unlikely event that $s_i = \ell_i$ or $c_w = \omega_i$,

in principle, an additional nonzero Lagrangian multiplier is activated and must be accounted for in the Lagrangian and strong duality condition.

III. PIMA-AS-OPF ALGORITHM

In this section, an ML approach (PIMA-AS-OPF) is proposed to enhance the computational efficiency of solving the DC-OPF given in (1). The proposed PIMA-AS-OPF procedure includes four steps:

- 1) Generate multiple DC-OPF instances using a standard DC-OPF optimization solver and use these as ground truth for the subsequent machine learning steps.
- 2) Train a Neural Network (NN) to identify the active constraints and the non-zero slack variables in DC-OPF. The loss function for NN training is designed in Section III-B, while the NN inputs and outputs are presented in Section III-A.
- 3) Build an equivalent system of equations based on the NN outputs, which provides the same optimal solution as the DC-OPF problem, but is more computationally amenable. This equivalent system of equations is derived in Section III-C.
- 4) Use the derived relationship in (5) to compute the resulting LMPs. Test the resulting LMPs against the market properties in Section II-C such as revenue adequacy and cost recovery to validate their feasibility. This result validation procedure is discussed in Section III-D.

This approach exploits the fact that the number of active constraints in DC-OPF is relatively small, which is supported by NYISO's findings in [26] that only about 15 out of the 2203 lines in the system are regularly operated at their nominal thermal limits.

A. NN Inputs and Outputs

Active set learning can be interpreted as a supervised classification problem where an ML model, e.g., a NN, is trained to assign a binary label (0/1) to each relevant variable. The physical meaning of labels varies based on each specific application.

The NN inputs consist of the information accessible to ISOs during the real-time scheduling process. This information includes the unit commitment results and the forecast demand and renewable energy generation, which can be used to calculate the net load (ℓ_i). This information, however, requires pre-processing before being provided to the NN. Specifically, the "load flows" are computed as explained below based on the given net load and used as an additional input to the NN. Recall that the power flow on line k can be expressed as:

$$f_k = \sum_{i \in k_i} \Phi_{ki} \left[\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right]. \quad (8)$$

Although the values of p_g , s_i , and c_w in (8) are unknown before the DC-OPF is solved, a proxy for the power flows in the system can be computed based on ℓ_i and the topology of

the system as encoded in the PTDF matrix Φ . Therefore, the *load flow* on line k is defined as:

$$f_k^{\text{load}} := \sum_{i \in k_i} \Phi_{ki} \ell_i. \quad (9)$$

Both the load at each node and the load flow on each line are provided as inputs to the NN. In Section IV, it is demonstrated that the addition of load flows leads to improved replication of dispatch decisions and LMPs.

Classic active set learning methods for solving DC-OPF can identify the saturated lines and generators, see e.g., [19], [20]. However, our approach goes further by also determining if pre-arranged adjustments, such as load shedding or wind curtailment, are necessary to make the DC-OPF problem feasible. As a result, the desired NN outputs include:

- 1) The committed generators that reach their upper or lower output limits,
- 2) The transmission lines that reach their thermal limits,
- 3) The buses with non-zero load shedding, and
- 4) The wind farms with non-zero wind power curtailment.

Given the commitment status of generators, all desired outputs can be determined given the net load and the topology of the system, which ensures the feasibility of the NN model. The NN output is a matrix that consists of binary values. Since a Lagrange multiplier is zero if and only if the corresponding constraint is inactive, the NN outputs will include the Lagrange multipliers associated with constraints of generator outputs and transmission capacities, i.e., Eqs. (1d) and (1c). Meanwhile, the NN outputs should also include binary variables indicating whether load or wind power is curtailed, where “1” denotes “with curtailment” and “0” denotes “without curtailment.”

B. Loss Function for NN Training

In the PIMA-AS-OPF algorithm, the NN is used to identify the active constraints of generators and transmission lines in the DC-OPF formulation, as well as the occurrences of load shedding and wind curtailment. Accordingly, the label of “1” means a constraint is active or a slack variable is non-zero, and vice versa. Following conventional classification models such as logistic regression, the following loss function is designed for NN training:

$$L := \sum_{s,v} \left| \max_s d_v^s \right| L_{\text{reg}} \left(\mathbf{y}_v^s, \mathcal{B}(d_v^s) \right). \quad (10)$$

In Eq. (10), vector d_v denotes the value of variable v , which contains all of the variables of interest, such as the Lagrange multipliers of relevant constraints in DC-OPF. The values of d_v are examined at each sample of the testing set s . The function $\mathcal{B}(d)$ is a binary mapping derived from d . For example, if $\bar{\mu}_k \neq 0$, then $\mathcal{B}(\bar{\mu}_k) = 1$; else $\mathcal{B}(\bar{\mu}_k) = 0$. Meanwhile, matrix \mathbf{y} denotes the neural network outputs, which are between 0 and 1, and are compared to the ground truth values given by $\mathcal{B}(d_v^s)$. The components of \mathbf{y} are assessed using a pre-set threshold, e.g., 0.5, to determine whether or not the corresponding constraints are active. The function

$L_{\text{reg}}(x, y)$ measures the difference between x and y , as given by:

$$L_{\text{reg}}(x, y) = \begin{cases} -\log(x), & y = 1; \\ -\log(1-x), & y = 0. \end{cases}$$

Therefore, $L_{\text{reg}} \left(\mathbf{y}_v^{d(s)}, \mathcal{B}(d_v^s) \right)$ represents the difference between the NN outputs and its true value with respect to the s^{th} sample in the training set and the v^{th} variable of interest. In addition, this difference is weighted by the maximum value of d_v across the training data set. These weights place greater emphasis on penalizing the NN for incorrect identification of active constraints with large Lagrange multipliers, as this can cause significant errors in LMPs whose values are computed based on the Lagrange multipliers (see Eq. (5)). Meanwhile, more weight will be placed on variables that are not consistently zero or non-zero, as they display less discernible patterns are therefore more challenging for the NN to learn. This weighting method is specifically designed for the task of identifying active constraints since it takes into account the underlying structure of the DC-OPF formulation.

C. Equivalent System of Equations

In what follows, the “free” generators are defined as the generators with none of their output limits saturated. Similarly, the “free” lines are the transmission lines with none of their thermal limits are saturated.

Based on the NN outputs, all the zero elements in the vectors of Lagrange multipliers $\{\bar{\alpha}_g, \underline{\alpha}_g\}_{g \in G}$ and $\{\bar{\mu}_k, \underline{\mu}_k\}_{k \in L}$, as well as the zero elements in the vectors of slack variables $\{s_i\}_{i \in B}$ and $\{c_w\}_{w \in W}$, can be obtained. Therefore, the remaining unknown variables include the *non-zero* primal variables p_g, s_i^*, c_w^* and the *non-zero* Lagrange multipliers $\bar{\mu}_k^*, \underline{\mu}_k^*, \bar{\alpha}_g^*$ and $\underline{\alpha}_g^*$, where superscript * denotes non-zero values. With this information, the original DC-OPF formulation can be reduced to the following equivalent system of equations (ESE):

$$\forall \bar{\mu}_k^* : \sum_i \Phi_{ki} \left[\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right] = f_k^{\max}, \quad (11a)$$

$$\forall \underline{\mu}_k^* : \sum_i \Phi_{ki} \left[\sum_{g \in G_i} p_g - \sum_{w \in W_i} c_w - \ell_i + s_i \right] = -f_k^{\max}, \quad (11b)$$

$$\forall \bar{\alpha}_g^* : p_g = p_g^{\max}, \quad (11c)$$

$$\forall \underline{\alpha}_g^* : p_g = p_g^{\min}, \quad (11d)$$

$$\forall g \in G_{\text{free}} : 2c_g^{(2)} p_g + c_g^{(1)} - \lambda + \sum_k \Phi_{ki} (\bar{\mu}_k - \underline{\mu}_k) = 0, \quad (11e)$$

$$\forall s_i : c_i^{\text{LNS}} - \lambda + \sum_k \Phi_{ki} (\bar{\mu}_k - \underline{\mu}_k) = 0, \quad (11f)$$

$$\forall c_w : c_w^{\text{RES}} + \lambda + \sum_k \Phi_{ki} (\underline{\mu}_k - \bar{\mu}_k) = 0, \quad (11g)$$

$$\sum_i (\ell_i - s_i) = \sum_g p_g - \sum_w c_w, \quad (11h)$$

where Eqs. (11a)-(11d) are the complementary slackness conditions of the identified active constraints, Eq. (11e) is the stationarity of free generators, Eqs. (11f) and (11g) are the same as the the stationarity in Eqs. (4f) and (4g), and Eq. (11h)

is the same as the power balance constraint in Eq. (1b). Since the objective of the PIMA-AS-OPF algorithm is to calculate the LMPs, the non-zero Lagrange multipliers that are not used in the expression of LMPs in Eq. (5), such as $\bar{\alpha}_g^*$ and $\underline{\alpha}_g^*$, can also be eliminated from the ESE to reduce the dimensionality of the problem. In summary, the dimensionality of the ESE is $N \times N$, where $N = |k : \bar{\mu}_k \neq 0| + |k : \underline{\mu}_k \neq 0| + |G| + |i : s_i > 0| + |w : c_w > 0| + 1$. This is a significant reduction from the dimensionality of the original DC-OPF problem, which is $3|G| + 2|L| + |B| + |W| + 1$, with $|G|, |L|, |B|, |W|$ denoting the total number of generators, lines, buses, and wind farms, respectively.

D. LMP Validation Based on Market Properties

For a ML approach to be adopted by real-world system operators, it is important to ensure the feasibility of its output. As discussed in Section II-C, efficient market clearing results should have the properties of revenue adequacy, cost recovery, and efficiency. Therefore, to validate the effectiveness of the resulting LMPs, they need to be tested against selected market properties, such as Eqs. (6) and (7). In the proposed PIMA-AS-OPF algorithm, these conditions follow as direct consequences of the DC-OPF formulation. However, the learning algorithm is designed in such a way as to be robust to “small errors”. Thus, it is anticipated that, if a saturated line, generator, or another constraint is misclassified, and the associated Lagrange multiplier is small in magnitude, the market design properties are still preserved. This claim is verified in the case study presented in Sec. IV.

IV. CASE STUDY IN NYISO 1814-BUS SYSTEM

The effectiveness of the proposed PIMA-AS-OPF algorithm is demonstrated on the 1814-bus NYISO system, which is referred to as NYISO-1814 in this section. The NYISO-1814 system, as shown in Fig. 2, consists of 1814 buses, 2203 lines, 362 conventional generators, and 33 wind farms, and it is sufficiently large to test the PIMA-AS-OPF scalability.

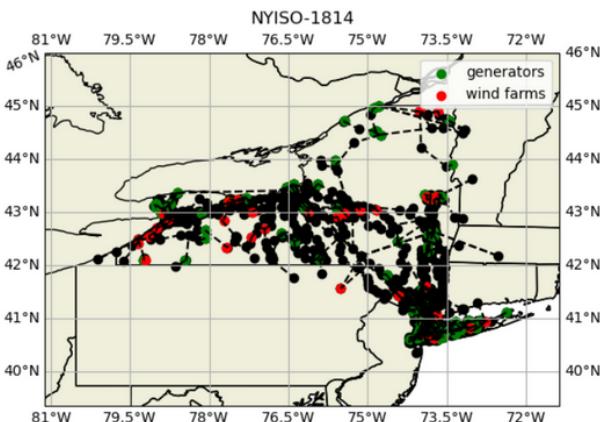


Fig. 2. The NYISO-1814 system.

A. Simulation Setup

The ground truth DC-OPF solutions are generated using Gurobi 10.0 [27] via JuMP [28] in the Julia programming language. The neural networks for classification are constructed and trained using the Flux.jl package [29], [30]. All other computations are performed using the native functions of Julia, or with the LinearAlgebra.jl package.

All the NN models in this case study share the same architecture, consisting of an input layer, three hidden layers with 30 neurons each and a ReLU activation function, and an output layer with a sigmoid activation function. The training of the NN models is performed on a high-performance computer, and it takes approximately 15-20 min to train each NN with a maximum training duration of 1000 epochs. In practice, the real-time system operators will be provided with a ready-to-use NN trained in advance based on the unit commitment results of day-ahead scheduling, and thus will only be required to perform the “system of linear equations” phase (as in Fig. 1) of the PIMA-AS-OPF approach in real time.

The testing of the NN is performed on a Lenovo IdeaPad with 12 GB RAM and an Intel Core i5 10 Gen processor. In summary, the testing results show that DC-OPF is solved by Gurobi in 0.11-0.13 sec, while the equivalent system of equations is solved in only 0.04-0.05 sec.

B. Six Basic Cases

To assess the effectiveness of the scheme, DC-OPF is solved for three representative combinations of load and wind injections obtained from two days in February and August, respectively, at different times. Furthermore, for each time, a *base* and *high* wind profile is considered to capture both current and projected wind power outputs under increased renewable penetration. Accordingly, Table II summarizes the six unit commitment (UCs) instances used to solve Eq. (2) for the NYISO system. Each unit commitment instance is solved *a priori*, resulting in a schedule of generators to be *on* and *off* when DC-OPF is solved in real-time. The three selected times constitute conditions under which the number of committed generators is small, moderate, and large. Cases A7B and A7H represent typical early-morning conditions, with 273 out of 362 generators committed. By contrast, the evening commute brings greater demand and necessitates greater generation capacity throughout the NYISO system, leading to 313 out of 362 commitments for A17B and A17H. In addition, an unanticipated thunderstorm occurred in the vicinity of Long Island at this time, presenting additional challenges with respect to uncertainty from wind farms. Finally, demand is relatively low in the early morning, so relatively few generators are online for F0B, with even fewer required with a higher wind power output level in F0H.

C. Sample Generation Procedure

Training input data samples are constructed from the given wind injections by adding tuned Gaussian noise. For each wind farm i , wind power injection w_i^0 is perturbed as:

$$w_i = w_i^0(1 + \eta \xi_i), \quad \xi_i \sim \mathcal{N}(0, 1), \quad (12)$$

TABLE II
OVERVIEW OF SIX BASIC CASES.

Case	Time	Wind	On ¹	P _{min} ²	P _{max} ³	P _{sched} ⁴
A7B	Aug.28, 7am	Base	273	87.49	289.41	231.10
A7H	Aug.28, 7am	High	273	87.49	289.41	231.10
A17B	Aug.28, 5pm	Base	313	102.53	336.59	312.79
A17H	Aug.28, 5pm	High	313	102.53	336.59	312.79
F0B	Feb.13, 12am	Base	187	55.78	184.71	142.67
F0H	Feb.13, 12am	High	141	50.81	158.43	123.89

¹ Number of committed generators in this case

² Minimum output of all committed generators (unit: MW)

³ Maximum output of all committed generators (unit: MW)

⁴ Scheduled output of all committed generators (unit: MW)

Noise level η is set to 0.01, 0.05, 0.1, and 0.15, increasing from a negligible level to one consistent with typical wind power variance [31], [32]. Note that load is not varied for this particular phase of the case study, as most of the variability arises from wind. The combination of changing unit commitment, wind penetration, and noise configurations introduced to the wind farms on top of the pre-set conditions presents varying levels of challenge to the neural network with respect to learning saturated generators, saturated lines, shed loads, and curtailed wind injections. Hence, a broad range of operating decisions can be captured. ML tools will be exposed to these decisions in real-world operations.

For each case in Table II, 1000 samples are generated, half of which are allocated for training and the rest for testing. A separate neural network is trained for each configuration. However, all neural networks share the same architecture from Section IV-A.

Remark 1. The complementary slackness conditions for the generator constraints are omitted, as saturated generators are fixed during the DC-OPF solution. As such, stationarity conditions are only written for free generators $g \in G_{\text{free}}$. Generators which reach their upper or lower limits in real-time are learned by the NN. Since these limits are known, the corresponding p_g values are assumed fixed for the purpose of constructing the system of linear equations, and are subtracted from the net load ℓ_i at the bus where they are located prior to solving the system.

D. Accuracy of Classification

The success of the proposed ML-driven OPF scheme depends on the correct classification of saturated generators, saturated lines, shed loads, and curtailed wind farms by the neural network. For example, misidentifying a saturated line with a high-magnitude Lagrange multiplier as below its limit by the neural network will result in large LMP errors at the attached and possibly nearby buses. On the other hand, by analyzing the results of the neural network classification, LMP and dispatch errors may be localized to particular buses, lines, generators, and wind farms, informing strategic improvements to the machine learning scheme to target the most vulnerable components of the power grid.

TABLE III
PERCENTAGE OF MISCLASSIFICATIONS IN DIFFERENT TEST CASES.

Case	Noise η	Generators ¹	Lines ²	Load ³	Wind ⁴
A7B	0.01	1.10%	0	0	0
	0.05	1.00%	0	0	0
	0.1	0.88%	0	0	0
	0.15	0.82%	0	0	0
A7H	0.01	0	0.01%	0	0
	0.05	0.42%	0.02%	0	0
	0.1	1.14%	0.02%	0	0.04%
	0.15	1.90%	0.03%	0	0.29%
A17B	0.01	0.05%	0.001%	0	0
	0.05	0.42%	0.02%	0	0
	0.1	0.72%	0.02%	0	0
	0.15	0.92%	0.02%	0	0
A17H	0.01	0	0	0	0
	0.05	0.10%	0.0005%	0	0.03%
	0.1	0.47%	0.02%	0	0.57%
	0.15	1.21%	0.01%	0	1.18%
F0B	0.01	0.04%	0.02%	0	1.18%
	0.05	0.41%	0.02%	0	1.41%
	0.1	0.89%	0.02%	0	1.45%
	0.15	1.49%	0.03%	0	1.45%
F0H	0.01	1.43%	0.002%	0	0.08%
	0.05	5.79%	0.02%	0	1.42%
	0.1	8.45%	0.03%	0	2.47%
	0.15	9.43%	0.04%	0	3.42%

¹ Misclassification of committed generators being free or non-free

² Misclassification of transmission lines being free or non-free

³ Misclassification of load shedding being zero or non-zero

⁴ Misclassification of wind curtailment being zero or non-zero

Table III displays the percentage of each quantity of interest which are misclassified by the neural network itemized for each configuration and noise level. The following observations are highlighted:

- 1) Across all configurations, as the noise associated with the wind perturbation increases, the number of misclassified lines and wind farms increases. This is the result of more volatile line congestion patterns and of increased wind power at farm locations with curtailment.
- 2) The neural network is more successful for the base wind than for the high wind configurations, which results in more frequent curtailment in the latter case.
- 3) Overall performance of the neural network is notably worse for F0B and F0H than for the other configurations. These two configurations correspond to a time of 12:00 midnight, when wind is typically higher. This results in a surplus production of wind power which, especially when combined with the higher levels of volatility, leads to less predictable patterns in wind curtailment that are missed by the neural network. Indeed, misclassifications are evident by the presence of infeasible wind curtailment values in the F0H case. Furthermore, the effect of adding η level of noise to the wind profile in F0B and F0H results in a higher level of relative noise to the net load, which is provided as an input to the neural network and thus affects the accuracy of classifications. Combining the high wind profile with the highest level of noise, performance of the neural network, and consequently the system of equations, is

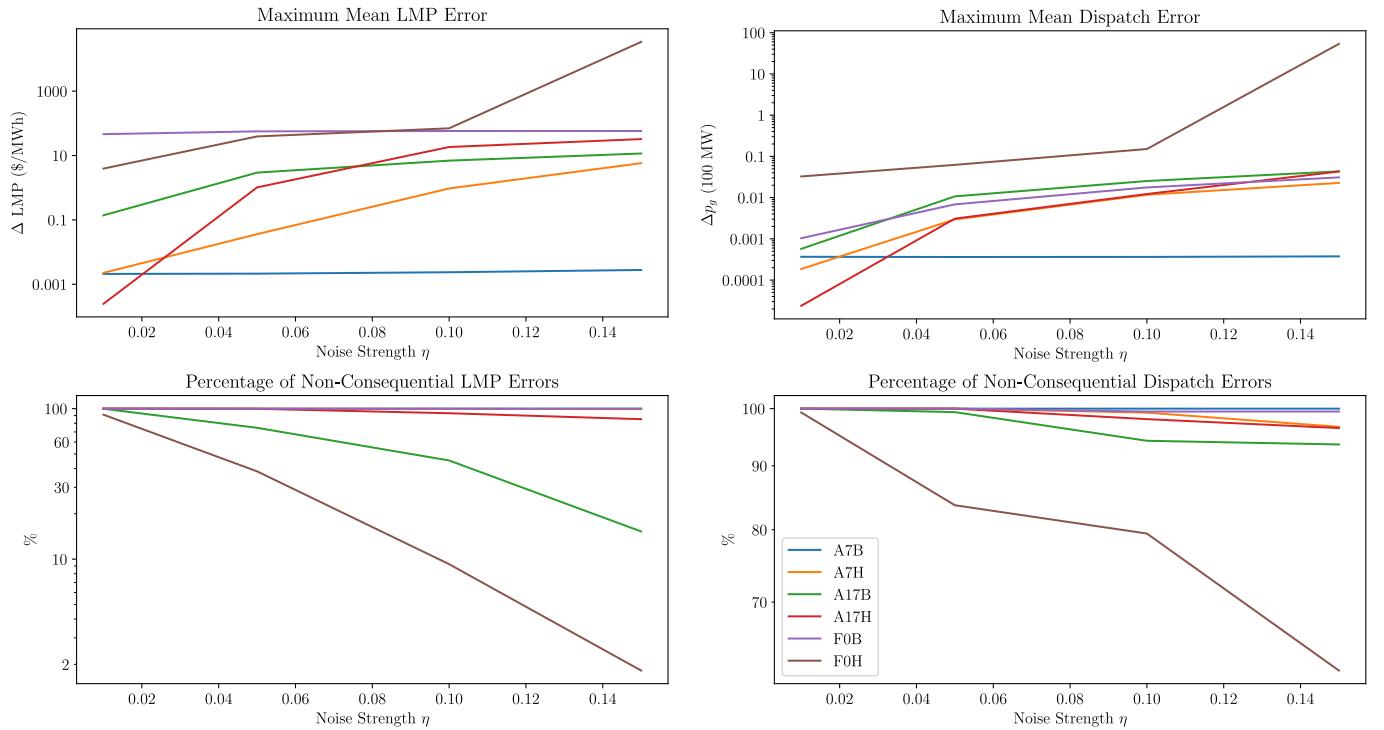


Fig. 3. Errors in reproduced LMPs (ΔLMP) and dispatch of free generators (Δp_g). The errors are defined as consequential when $|\Delta \text{LMP}| > 0.1$ \$/MWh and $|\Delta p_g| > 0.01$ p.u., respectively.

unsatisfactory for these configurations.

- 4) Misclassification of saturated generators is less consistent than for the other quantities. This is not necessarily damaging to the ultimate OPF solution, as the system of equations may still compute a dispatch near the generator capacity if e.g., a saturated generator is considered free. Similarly, classifying a generator near, but not exactly equal to its limit as saturated is not directly harmful to the quality of the LMP and dispatch replication. For instance, in testing trial 98 in the A17H test case, all of the p_g discrepancies are below 0.25 MW p.u., and at least one correctly classified generator exhibits a larger error than the misclassified generators do.
- 5) Load shedding is identified with complete accuracy across the board for all configurations. The load shedding penalty c_i^{LNS} is chosen such that it is unique for every bus to avoid ambiguous DC-OPF solutions, which results in more consistent load shedding patterns to be learned by the neural network.

E. Accuracy of DC-OPF Solution

Both the primal (dispatch of generators) and dual (LMPs) solution of DC-OPF are of interest. Figure 3 depicts the error in learned LMPs and dispatch using the mean absolute error across the testing set. Subfigures 3(a) and 3(b) show the maximum mean error, i.e., the error at the bus or (free) generator experiencing the largest discrepancy, while subfigures 3(c) and 3(d) show the percentage of buses or generators with mean errors below a specified threshold. These measures of success

are consistent with those for the neural network. Namely, as neural network performance improves, so does the accuracy of computation of dispatch and LMPs. Indeed, for most of the inspected configurations, over 90% of buses and generators exhibit mean errors of an acceptable magnitude, while maximum errors may be identified and explained in an interpretable way. Notably, the F0H configuration, at high levels of noise, exhibits unsatisfactory results by all measures, with some LMP errors upwards of \$30,000/MWh, which coincides with misclassification of curtailment in this case. As previously mentioned, it is possible to interpret the buses and generators experiencing large errors. Figure 4 provides heat maps for the A17H configuration at all levels of noise, with frequently misclassified lines, wind farms, and generators emerging at high levels of noise and indicated on the maps. The magnitudes of errors increase as noise increases. In addition, the largest LMP errors are concentrated in “clusters” of buses, and each cluster may be associated with a misclassified line or wind farm. Since the entire set of misclassified entities is small, NNs can be adapted by introducing penalty terms to the loss function to target misclassification of those particular lines and wind farms. This aligns with the notion of “physics-informedness,” which, in physics-informed neural networks, typically entails adding terms to the loss function to enforce constraints derived from the underlying physics of the system (e.g. satisfaction of an ODE) [23].

F. LMP Validation Results

To conclude the case study, the fidelity of the testing set with respect to selected market design conditions is assessed. The

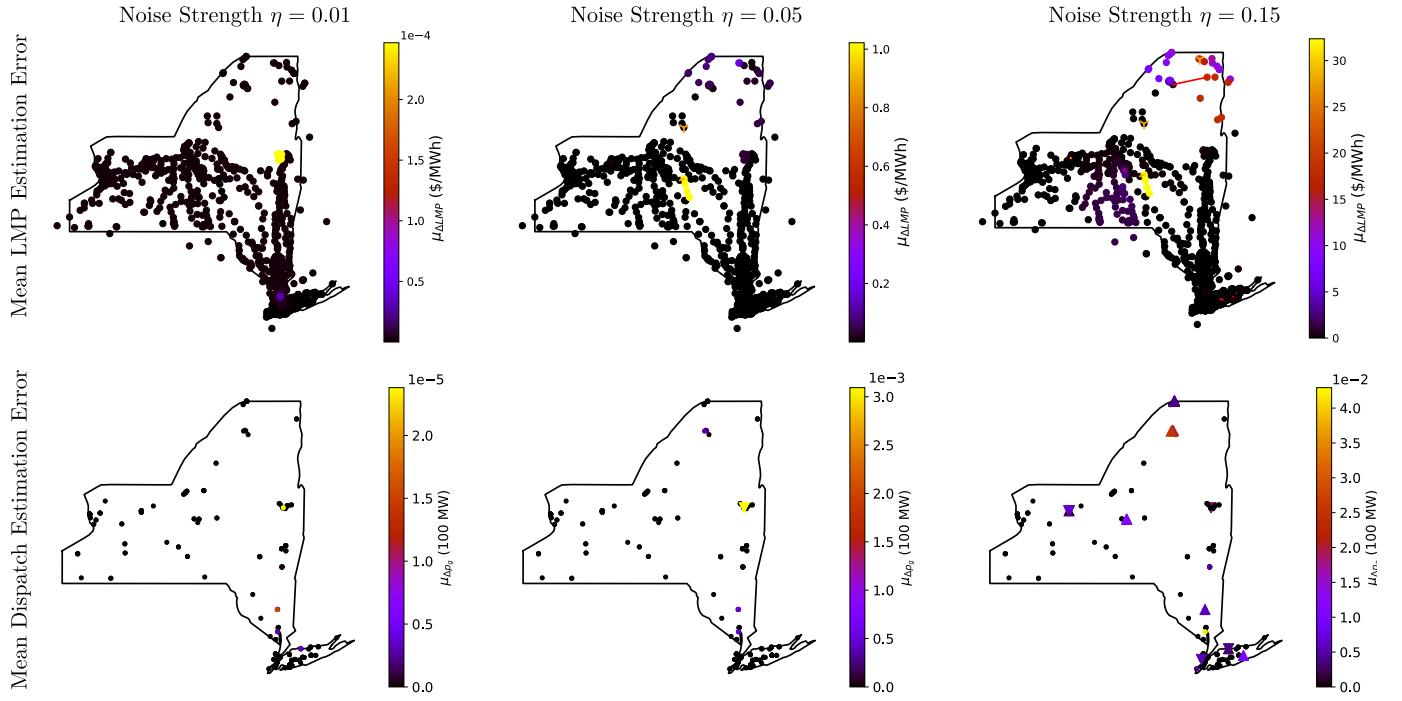


Fig. 4. Mean LMP and p_g errors for the A17H configuration. Misclassified lines and wind farms indicated are display, in top panels, in red for false negatives and in orange for false positives. Consequential misclassified generators (i.e. when dispatch error is greater than 0.001 p.u.) are displayed, in bottom panels, as \blacktriangledown for false negatives and as \blacktriangle for false positives.

goal of PIMA-AS-OPF is to replicate the ground-truth DC-OPF in all respects as accurately as possible. As the ground-truth DC-OPF is convex, efficiency via strong duality holds by construction, as long as the neural network classification step is successful. Thus, empirical results are provided for revenue adequacy and cost recovery. Furthermore, cost recovery is only assessed for the “free” generators in each sample. Considering all sample configurations except for F0H, revenue adequacy is satisfied in all samples. This demonstrates that, under unit commitments for which DC-OPF solution via PIMA-AS-OPF is a reasonable approach, even when the neural network does not perform perfectly, the desired market design conditions are still satisfied. In the F0H configuration, which exhibited the least satisfactory results by far, all samples satisfy revenue adequacy by (6). In addition, considering all free generators across the testing set, 195 out of 23767 (0.82%) do not satisfy cost recovery with equality. Revenue adequacy in both forms is satisfied for all samples in F0B, with only 3 out of 240000 (0.0125%) of free generators failing to satisfy cost recovery.

V. CONCLUSION AND FUTURE WORK

This paper bridges efforts in the power systems community to accelerate optimization using machine learning with the physical constraints and economic properties of electricity markets. The case study on NYISO-1814 demonstrates that the NN in the PIMA-AS-OPF pipeline can accurately identify the saturated lines, saturated generators, shed loads, and curtailed wind farms, for different sets of unit commitments and various levels of wind volatility. The outputs of the NN can then be used to reproduce the dispatch and LMPs using a system

of linear equations, and the resulting LMPs adhere to the principle of revenue adequacy and cost recovery. However, some particular sample configurations may not be replicated effectively. Nonetheless, for these cases, specific features of the NYISO system that contribute to these deficiencies can be identified.

The optimization procedure in the DC-OPF problem is not completely replaced by ML in this study. Instead, ML is used to leverage specific properties of the optimization constraints, such as discrete degrees of freedom (such as saturated lines and generators, loads shed, and renewable generators curtailed) and sparsity, while preserving the physics underlying power system operations. In other words, the ML algorithm is only used to learn the infrequent discrete decisions. The remaining DC-OPF solution is computed using a system of linear equations, which is an exact method and more computationally efficient than optimization.

Future work will involve modifying the NN architecture and adjusting hyper-parameters, particularly when the current NN is not performing satisfactorily. Additionally, a similar ML approach for the unit commitment problem will be developed, which poses unique challenges due to its mixed-integer programming nature. The goal is to create an integrated learning-based pipeline for the joint UC-OPF problem that can be used for both day-ahead and real-time market operations.

ACKNOWLEDGMENT

The authors would like to collectively acknowledge with gratitude the ARPA-E PERFORM grant from which this work arose. The UArizona team acknowledges complementary

funding from the office of the Vice President for Research. The work of RF and MC on the project was partially supported by the UArizona NSF RTG in Data-Driven Discovery.

REFERENCES

- [1] “Manual 11: Day-ahead scheduling manual,” NYISO Energy Market Operations, 2022. [Online]. Available: https://www.nyiso.com/documents/20142/2923301/dayahd_schd_mnl.pdf/0024bc71-4dd9-fa80-a816-f9f3e26ea53a
- [2] “Manual 12: Transmission and dispatch operations manual,” NYISO Energy Market Operations, 2022. [Online]. Available: https://www.nyiso.com/documents/20142/2923301/trans_disp.pdf/9d91ad95-0281-2b17-5573-f054f7169551
- [3] F. Li and R. Bo, “Dcopf-based lmp simulation: algorithm, comparison with acopf, and sensitivity,” *IEEE Transactions on Power Systems*, vol. 22, no. 4, pp. 1475–1485, 2007.
- [4] “Offshore wind market report: 2022 edition,” US Department of Energy, 2022. [Online]. Available: <https://www.energy.gov/sites/default/files/2022-09/offshore-wind-market-report-2022-v2.pdf>
- [5] L. Zhang, Y. Chen, and B. Zhang, “A convex neural network solver for dcopf with generalization guarantees,” *IEEE Transactions on Control of Network Systems*, vol. 9, no. 2, pp. 719–730, 2021.
- [6] Y. Chen, L. Zhang, and B. Zhang, “Learning to solve dcopf: A duality approach,” *Electric Power Systems Research*, vol. 213, p. 108595, 2022.
- [7] Y. Ng, S. Misra, L. A. Roald, and S. Backhaus, “Statistical learning for dc optimal power flow,” in *2018 Power Systems Computation Conference (PSCC)*. IEEE, 2018, pp. 1–7.
- [8] L. Pagnier, R. Ferrando, Y. Dvorkin, and M. Chertkov, “Machine learning for electricity market clearing,” *arXiv:2205.11641*, 2022, presented at the 11th Bulk Power Systems Dynamics Symposium (IREP 2022), July 25-30, 2022, Banff, Canada. [Online]. Available: <https://arxiv.org/abs/2205.11641>
- [9] J. Rahman, C. Feng, and J. Zhang, “Machine learning-aided security constrained optimal power flow,” in *2020 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2020, pp. 1–5.
- [10] R. Canyasse, G. Dalal, and S. Mannor, “Supervised learning for optimal power flow as a real-time proxy,” in *2017 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*. IEEE, 2017, pp. 1–5.
- [11] T. Falconer and L. Mones, “Leveraging power grid topology in machine learning assisted optimal power flow,” *IEEE Transactions on Power Systems*, 2022.
- [12] X. Lei, Z. Yang, J. Yu, J. Zhao, Q. Gao, and H. Yu, “Data-driven optimal power flow: A physics-informed machine learning approach,” *IEEE Transactions on Power Systems*, vol. 36, no. 1, pp. 346–354, 2020.
- [13] R. Nelliikkath and S. Chatzivasileiadis, “Physics-informed neural networks for ac optimal power flow,” *Electric Power Systems Research*, vol. 212, p. 108412, 2022.
- [14] Y. Yang, Z. Yang, J. Yu, K. Xie, and L. Jin, “Fast economic dispatch in smart grids using deep learning: An active constraint screening approach,” *IEEE Internet of Things Journal*, vol. 7, no. 11, pp. 11 030–11 040, 2020.
- [15] W. Chen, S. Park, M. Tanneau, and P. Van Hentenryck, “Learning optimization proxies for large-scale security-constrained economic dispatch,” *Electric Power Systems Research*, vol. 213, p. 108566, 2022.
- [16] R. Nelliikkath and S. Chatzivasileiadis, “Physics-informed neural networks for minimising worst-case violations in dc optimal power flow,” in *2021 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 2021, pp. 419–424.
- [17] S. Liu, C. Wu, and H. Zhu, “Graph neural networks for learning real-time prices in electricity market,” *arXiv preprint arXiv:2106.10529*, 2021.
- [18] M. K. Singh, S. Gupta, V. Kekatos, G. Cavraro, and A. Bernstein, “Learning to optimize power distribution grids using sensitivity-informed deep neural networks,” in *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 2020, pp. 1–6.
- [19] D. Deka and S. Misra, “Learning for DC-OPF: Classifying active sets using neural nets,” in *2019 IEEE Milan PowerTech, arXiv/1902.05607*, 2019, pp. 1–6.
- [20] S. Misra, L. Roald, and Y. Ng, “Learning for constrained optimization: Identifying optimal active constraint sets,” *INFORMS Journal on Computing*, vol. 34, no. 1, pp. 463–480, 2022.
- [21] F. E. Curtis, D. K. Molzahn, S. Tu, A. Wächter, E. Wei, and E. Wong, “A decomposition algorithm for large-scale security-constrained ac optimal power flow,” *arXiv:2110.01737*, 2021.
- [22] X. Pan, T. Zhao, M. Chen, and S. Zhang, “Deepopf: A deep neural network approach for security-constrained dc optimal power flow,” *IEEE Transactions on Power Systems*, vol. 36, no. 3, pp. 1725–1735, 2020.
- [23] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, “Physics-informed machine learning,” *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.
- [24] A. Beck, *Introduction to nonlinear optimization: Theory, algorithms, and applications with MATLAB*. SIAM, 2014.
- [25] F. A. Wolak, *Market Design and Price Behavior in Restructured Electricity Markets: An International Comparison*. Boston, MA: Springer US, 2000, pp. 127–152. [Online]. Available: https://doi.org/10.1007/978-1-4615-4529-3_8
- [26] “Guide 01: Market participants user’s guide,” NYISO Stakeholder Services, 2021. [Online]. Available: <https://www.nyiso.com/documents/20142/3625950/mpug.pdf/>
- [27] Gurobi Optimization, LLC, “Gurobi Optimizer Reference Manual,” 2023. [Online]. Available: <https://www.gurobi.com>
- [28] I. Dunning, J. Huchette, and M. Lubin, “Jump: A modeling language for mathematical optimization,” *SIAM Review*, vol. 59, no. 2, pp. 295–320, 2017.
- [29] M. Innes, E. Saba, K. Fischer, D. Gandhi, M. C. Rudiloso, N. M. Joy, T. Karmali, A. Pal, and V. Shah, “Fashionable modelling with flux,” *CoRR*, vol. abs/1811.01457, 2018. [Online]. Available: <https://arxiv.org/abs/1811.01457>
- [30] M. Innes, “Flux: Elegant machine learning with julia,” *Journal of Open Source Software*, 2018.
- [31] N. Lipman, P. Dunn, P. Musgrave, B. Sexon, and G. Slack, “Wind/diesel integration for electricity supply to isolated communities,” in *Energy for Rural and Island Communities*, J. TWIDELL, Ed. Pergamon, 1982, pp. 149–164. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9780080276069500210>
- [32] Y. Dvorkin, M. Lubin, S. Backhaus, and M. Chertkov, “Uncertainty sets for wind power generation,” *IEEE Transactions on Power Systems*, vol. 31, no. 4, pp. 3326–3327, 2016.