

# Replication of "Self-Supervised Learning for Large-Scale Preventive Security Constrained DC Optimal Power Flow"

February 11, 2026

## Abstract

Security-Constrained Optimal Power Flow(SCOPF) is essential for ensuring reliable power system operation under  $N-1$  contingencies. This paper proposes a preventive methodology. In this preventive DC-SCOPF with Automatic Primary Response (APR), the inclusion of contingency constraints and binary variables leads to large-scale mixed-integer formulations whose computational complexity grows rapidly with system size.

This project follows the Replicator track in which the methodology and experimental findings of the referred paper will be reproduced, which proposes a self-supervised primal-dual learning method(PDL-SCOPF) to generate near-optimal and feasible solutions. As it is self supervised learning method, it also doesn't require labeled optimal solutions.

The replication will involve re-implementing the preventive DC-SCOPF formulation with generator and line contingencies, including the APR-based contingency modeling. The core learning methodology will be followed, which includes the primal network for base-case dispatch prediction, the dual network for Lagrange multiplier estimation, the power balance repair layer ensuring nominal feasibility, and the binary search layer for contingency dispatch computation. The self-supervised training procedure inspired by the Augmented Lagrangian Method will also be reproduced.

Experimental evaluation will be conducted on power system test systems in which we verify optimality gap, power balance violations, slack variable magnitudes, and computational speedups relative to solver-based approaches. The main objective of this project is to assess the reproducibility of the reported scalability and near-optimal performance given in the paper, and to analyze the effectiveness of self-supervised primal-dual learning for large-scale constrained optimization in power systems.

# Self-Supervised Learning for Large-Scale Preventive Security Constrained DC Optimal Power Flow

Seonho Park  and Pascal Van Hentenryck 

**Abstract**—Security-Constrained Optimal Power Flow (SCOPF) plays a crucial role in power grid stability but becomes increasingly complex as systems grow. This paper introduces Primal-Dual Learning (PDL) for SCOPF (PDL-SCOPF), a self-supervised end-to-end primal-dual learning framework for producing near-optimal solutions to large-scale SCOPF problems in milliseconds. Indeed, PDL-SCOPF remedies the limitations of supervised counterparts that rely on training instances with their optimal solutions, which becomes impractical for large-scale SCOPF problems. PDL-SCOPF mimics an Augmented Lagrangian Method (ALM) for training primal and dual networks that learn the primal solutions and the Lagrangian multipliers, respectively, to the unconstrained optimizations. In addition, PDL-SCOPF incorporates a repair layer to ensure the feasibility of the power balance in the nominal case, and a binary search layer to compute, using the Automatic Primary Response (APR), the generator dispatches in the contingencies. The resulting differentiable program can then be trained end-to-end using the objective function of the SCOPF and the power balance constraints of the contingencies. Experimental results demonstrate that the PDL-SCOPF delivers accurate feasible solutions with minimal optimality gaps. The framework underlying PDL-SCOPF aims at bridging the gap between traditional optimization methods and machine learning, highlighting the potential of self-supervised end-to-end primal-dual learning for large-scale optimization tasks.

**Index Terms**—Primal-dual learning, security-constrained optimal power flow, column and constraint generation, deep learning, differentiable programming, end-to-end learning.

## NOMENCLATURE

The next list describes several symbols that will be later used. Sets and indices are in calligraphic or blackboard bold font. Bold symbols represent vectors or matrices.

### Machine learning related

$\lambda$	Dual variable estimate associated with the power balance constraint under the generator contingencies.
$D_\phi$	Dual network.

Received 29 November 2023; revised 27 April 2024 and 25 August 2024; accepted 9 November 2024. Date of publication 14 November 2024; date of current version 21 April 2025. This work was supported by NSF under Grant 2007095 and Grant 2112533. Paper no. TPWRS-01863-2023. (Corresponding author: Pascal Van Hentenryck.)

The authors are with the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: seonho.park@gatech.edu; pascal.vanhentenryck@isye.gatech.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TPWRS.2024.3498705>.

Digital Object Identifier 10.1109/TPWRS.2024.3498705

$\tilde{\mathbf{g}}_k$	Dispatch estimate under the generator contingency $k$
$P_\theta$	Primal network.
$\tilde{\eta}$	Slack variable estimate.
$\mathbf{x}$	Input parameters.
$\tilde{\mathbf{g}}$	Base case dispatch estimate.
$\tilde{\mathbf{g}}$	Base case dispatch estimate before applying the power balance repair layer.

### Parameters

$\mathbf{B}$	Generator–bus incidence matrix.
$\beta$	Threshold parameter for choosing the cuts to be added in the CCGA.
$\delta$	Relaxation parameter for the binary search layer.
$\underline{\mathbf{f}}, \bar{\mathbf{f}}$	Lower and upper bound of power flow limit.
$\gamma$	Primary response parameter.
$\hat{\mathbf{g}}$	Generator capacities.
$\underline{\mathbf{g}}, \bar{\mathbf{g}}$	Lower and upper bound of generator dispatch.
$\underline{\mathbf{d}}$	Load demands.
$\mathbf{L}$	Line outage distribution factor matrix.
$\mathbf{c}$	Cost coefficients.
$\mathbf{K}$	Power transfer distribution factor matrix.
$M_\eta$	Penalty coefficient.

### Set and indices

$\mathcal{G}, \mathcal{E}, \mathcal{N}, \mathcal{L}$	Indices of the generators, transmission lines, buses, load units.
$\mathcal{K}_g, \mathcal{K}_e$	Generator or transmission line contingency indices.
$\mathbb{K}_g$	Subset of $\mathcal{K}_g$ considered in the current CCGA iteration.
$\mathbb{U}_g, \mathbb{U}_e$	Sets of pairs of $(k, l)$ where the power flow limit constraint at line $l$ is violated under the generator or line contingency $k$

### Variables

$\mathbf{f}$	Base case power flow vector.
$\mathbf{f}_k$	Power flow vector under the generator contingency $k$
$\mathbf{g}$	Power generations (dispatches, injections)
$\mathbf{g}_k$	Power generations under the generator contingency $k$
$\mathbf{g}'_k$	Provisional variables for $\mathbf{g}_k$
$n_k$	Global signal variable under the generator contingency $k$
$\rho_k$	Binary variable indicating whether $\mathbf{g}_k$ reaches the upper limit under the generator contingency $k$
$\eta$	Slack variable for power flow limit constraints.

## I. INTRODUCTION

**P**OWER system operations must maintain the power balance between load and generation at all times. To achieve this equilibrium requires solving mathematical optimization problems in real-time and day-ahead electricity markets. These optimizations have become increasingly challenging in recent years due to the integration of renewable energy sources which are both more numerous and volatile compared to traditional synchronous generators. As a consequence, managing the reliability and the risk in the system gets increasingly important.

The Security-Constrained Optimal Power Flow (SCOPF) problem is a traditional model to ensure the stability of power systems operations under contingencies. The SCOPF guarantees that a feasible generator dispatch exists even under the failure of a single generator or transmission line, while minimizing the cost of the nominal dispatch. In actual operations, however, Independent System Operators (ISOs) do not solve the SCOPF; rather they maintain reserves that are computed outside the markets; the reserve requirements may be sub-optimal or insufficient in cases with large shares of renewables. One reason for using reserves is the high computational cost of the SCOPF compared to its equivalent Economic Dispatch (ED) with reserves, even when using a DC approximation of the power flow equations. Indeed, the DC-SCOPF is a Mixed-Integer Linear Program (MILP) whose number of binary variables increases quadratically with the number of contingencies for preventive SCOPF with Automatic Primary generator Response (APR) [1]. In contrast, the ED with the reserve constraints is a linear program that can be readily solved.

One possible approach to address the computational challenges of the SCOPF is the use of Machine Learning (ML) which has attracted significant attention in the power systems community in recent years. However, even for machine learning, the SCOPF raises fundamental challenges since generating sufficient training data is typically impractical for industrial-size power grids. Moreover, the ML models may have to predict a quadratic number of variables to ensure the power balance constraints in the contingencies.

This paper takes a different avenue. It considers the ED formulation used by the ISOs in the United States where the reserve requirements have been replaced by  $N - 1$  contingencies to produce a preventive APR-based SCOPF. The paper then proposes an *end-to-end self-supervised primal-dual Learning framework*, called PDL-SCOPF, to address the computational challenges raised by the SCOPF formulation. By virtue of being self-supervised, PDL-SCOPF does not need the optimal solutions of thousands of training instances: it just relies on the availability of an empirical distribution of the input configuration that captures future conditions of interest. Thanks to being trained end-to-end, PDL-SCOPF only needs to predict the baseline dispatch: differentiable layers are then used to restore the feasibility of the power balance for the base case and to predict the contingency dispatches using the APR. PDL-SCOPF employs a Primal-Dual Learning (PDL) framework whose training mimics the Augmented Lagrangian Method (ALM) over a set of instances: it uses a primal network to approximate unconstrained

optimization problems similar to those of the ALM and a dual network to approximate the Lagrangian multipliers used in these optimizations.

The main contribution of the paper, PDL-SCOPF, can be summarized as follows:

- PDL-SCOPF is a self-supervised method that produces near-optimal feasible solutions to preventive DC-SCOPF. Thus, it does not require a training dataset that contains optimal SCOPF solutions for a set of instances.
- PDL-SCOPF is a Primal-Dual Learning framework that mimics Augmented Lagrangian Method by predicting both primal optimal solutions (using a primal network) and their associated Lagrangian multiplier (using a dual network).
- PDL-SCOPF uses differential layers to restore the feasibility of the power balance [2] and to adapt the binary search from the Column and Constraint Generation Algorithm in [3] to compute the contingency dispatches and their violations. The predictive model and these differentiable layers can then be trained end-to-end in the framework to produce near-optimal feasible solutions.
- The scalability and accuracy of PDL-SCOPF have been validated on industry-size test cases with thousands of buses, for which it estimates solutions in milliseconds.

The rest of the paper is structured as follows. Section II describes the related work in SCOPF and machine learning. Section III revisits the formulation of SCOPF and briefly describes a Column and Constraint Generation Algorithm to solve it exactly. Section IV revisits the PDL framework and Section V introduces PDL-SCOPF. Section VI presents the experimental results and compares PDL-SCOPF with a number of baselines. Section VII summarizes the paper and presents directions for future work.

## II. RELATED WORK

### A. Security-Constrained OPF

This paper focuses on the SCOPF problem with  $N - 1$  generator and line contingencies. A comprehensive review of SCOPF can be found in [4]. The SCOPF aims at determining the pre-contingency generator dispatch that minimizes operational costs while ensuring feasibility even in the event of contingencies. A  $N - 1$  contingency refers to the outage of any single component (generator or line). The SCOPF can be classified into two distinct settings: 1) the corrective case, discussed in [5], which assumes that re-scheduling is possible, and 2) the preventive case, studied for instance in [1], [3], [6], [7], where re-dispatch is not an option. This paper considers a preventive SCOPF that models an APR for generator contingencies. Such SCOPF problems are computationally challenging due to the binary variables introduced to model the APR response. Indeed, the number of binary variables increases quadratically in the number of generator contingencies. To alleviate this computational burden, various decomposition techniques have been proposed: they include Benders decomposition (e.g., [8], and the Column and Constraint Generation Algorithm (CCGA) proposed in [3], [9]. The PDL-SCOPF method proposed in this

paper uses some insights from the CCGA. Extensive reviews on AC-SCOPF are available in [10], [11]. Recent efforts to tackle AC-SCOPF, as part of the Grid Optimization Competition initiated by ARPA-E, are summarized in [12]. Both industry and academia have also explored approximations and relaxations to the AC-SCOPF, as discussed in [13], [14]. Currently, industry (e.g., ISO) predominantly employs the DC formulation for SCOPF (DC-SCOPF) [15]. Consequently, this paper focuses on the DC-SCOPF. More precisely, the formulation presented in this paper can be viewed as the economic dispatch of the ISOs where the reserve constraints have been replaced by  $N - 1$  contingencies under the APR modeling. For a broader perspective, see the recent comprehensive review on large grid optimization in [16].

### B. Machine Learning for Optimization and Power Systems

Various end-to-end Machine Learning (ML) approaches have recently emerged for optimization applications. They aim at directly estimating optimal solutions given a distribution of the input parameters. In other words, these ML approaches are used as regressions to approximate the mapping from input parameters to optimal solutions. A comprehensive overview of end-to-end optimization learning methods can be found in [17]. A ML framework to accelerate general mixed integer programming (MIP) solvers was proposed in [18]; it also offers a way to bound the optimality gap between the inference and the optimal solution. End-to-end optimization learning techniques have also founded applications in various power system optimization tasks, such as economic dispatch [2], DCOPF [19], ACOPF [20], and unit commitment [21], [22].

### C. Supervised Learning for SCOPF

Several supervised learning approaches have been applied to SCOPF. DEEPOPf [19] approximates the optimal dispatch of the SCOPF problem, taking into account line contingencies. Velloso and Van Hentenryck [23] combine a deep neural network-based mapping with the CCGA method to approximate the optimal solution of the SCOPF problem, with a focus on generator contingencies. These approaches employed supervised learning, which requires gathering pairs of input parameters and their corresponding optimal solutions for training. As demonstrated in Section VI, solving a single SCOPF instance for large-scale SCOPF is quite time-consuming, which makes the collection of such training data impractical. This paper remedies this limitation by using a self-supervised learning framework that does not necessitate the ground truth solutions for training, which, in turn, is especially effective for learning large-scale industry-sized SCOPF problems having both generator and line contingencies.

### D. Self-Supervised Learning for OPF

Self-supervised learning leverages the original optimization formulation, including the objective function, to approximate optimal solutions without relying on the ground truth data. DC3 [24] is such a self-supervised method that employs an

implicit layer to driving the learning towards feasible solutions. E2ELR [2] introduces repair layers for satisfying power balance and reserve constraints that are trained end-to-end to produce feasible and near-optimal solutions to ED problems. PDL [25] integrates primal and dual networks to approximate both the primal and dual solutions for constrained optimization problems. In this work, PDL-SCOPF leverages the E2ELR repair layer for power balance, PDL, and the binary search layer adapted from CCGA to produce near-optimal feasible solutions to industry-size SCOPF problems.

### E. Scalable Learning for OPF

Significant work has been carried out for scaling ML approach to large-scale OPF. They include Compact Learning [26], which learns a lower-rank representation of the optimal solution for ACOPF, and Spatial Decomposition [27], which proposes a two-step learning process by learning the flows between regions in a spatial decomposition before learning the optimal power flow in each region.

## III. PROBLEM FORMULATION AND OBJECTIVE

This section reviews the SCOPF problem considered in this paper. The formulation can be seen as the traditional economic dispatch of the ISOs in the United States, where the reserve constraints have been replaced by an explicit modeling of  $N - 1$  contingencies. In the formulation, the thermal limits are thus considered as soft constraints but their violations are heavily penalized in the objective function. To handle contingencies, the formulation captures the Automatic Primary Response (APR) [1], [12] used in [3], [23].

### A. The SCOPF Problem

Model 1 presents the extensive SCOPF problems with  $N - 1$  generator and line contingencies. Its primary objective is to determine the optimal active generation points for the base case while ensuring the feasibility of both generator and line contingencies. Objective (1) sums the linear cost of the base case dispatch  $\mathbf{g}$  and the penalties for violating the thermal limits in the base case and in the contingencies, with  $M_\eta$  being a large penalty coefficient.

1) *The Base Case:* The base case includes Constraints (2), (3), and (4) that should be satisfied. Constraint (2) imposes the power balance between the generation and the load demands in the nominal case. In this constraint,  $\mathbf{1}$  represents a vector whose elements are all one, and  $\mathbf{d} \in \mathbb{R}^{|\mathcal{N}|}$  is the load demand with the bus-wise representation of the load units. Constraint (3) ensures that the power flow  $\mathbf{f}$  for each transmission line falls within the predefined upper and lower limits,  $\underline{\mathbf{f}}$  and  $\bar{\mathbf{f}}$ . Any violation is penalized in the objective value. The power flow for the base case  $\mathbf{f} = \mathbf{K}(\mathbf{d} - \mathbf{B}\mathbf{g})$  is computed using a Power Transfer Distribution Factor (PTDF) matrix  $\mathbf{K} \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{N}|}$  and a generator-bus incidence matrix  $\mathbf{B}$ . Constraint (4) imposes the generation limits.

**Model 1: The Extensive SCOPF Formulation**

$$\begin{aligned}
& \min_{\mathbf{g}, \rho_k, n_k, \eta_k} \quad \mathbf{c}^\top \mathbf{g} + M_\eta \left( \sum_{k \in \{0\} \cup \mathcal{K}_g \cup \mathcal{K}_e} \|\boldsymbol{\eta}_k\|_1 \right) & (1) \\
& \text{s.t.} : \quad \mathbf{1}^\top \mathbf{g} = \mathbf{1}^\top \mathbf{d} & (2) \\
& \quad \mathbf{f} - \boldsymbol{\eta}_0 \leq \mathbf{f} = \mathbf{K}(\mathbf{d} - \mathbf{B}\mathbf{g}) \leq \bar{\mathbf{f}} + \boldsymbol{\eta}_0 & (3) \\
& \quad \mathbf{g} \leq \mathbf{g} \leq \bar{\mathbf{g}} & (4) \\
& \quad \mathbf{1}^\top \mathbf{g}_k = \mathbf{1}^\top \mathbf{d} \quad \forall k \in \mathcal{K}_g & (5) \\
& \quad \mathbf{f} - \boldsymbol{\eta}_k \leq \mathbf{f}_k = \mathbf{K}(\mathbf{d} - \mathbf{B}\mathbf{g}_k) \leq \bar{\mathbf{f}} + \boldsymbol{\eta}_k \quad \forall k \in \mathcal{K}_g & (6) \\
& \quad \underline{g}_i \leq g_{k,i} \leq \bar{g}_i \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k & (7) \\
& \quad g_{k,k} = 0 \quad \forall k \in \mathcal{K}_g & (8) \\
& \quad |g_{k,i} - g_i - n_k \gamma_i \hat{g}_i| \leq \hat{g}_i \rho_{k,i} \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k & (9) \\
& \quad g_i + n_k \gamma_i \hat{g}_i \geq \hat{g}_i \rho_{k,i} + \underline{g}_i \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k & (10) \\
& \quad g_{k,i} \geq \hat{g}_i \rho_{k,i} + \underline{g}_i \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k & (11) \\
& \quad \mathbf{f} - \boldsymbol{\eta}_k \leq \mathbf{f} + f_k \mathbf{L}_k \leq \bar{\mathbf{f}} + \boldsymbol{\eta}_k \quad \forall k \in \mathcal{K}_e & (12) \\
& \quad \boldsymbol{\eta}_k \geq 0 \quad \forall k \in \{0\} \cup \mathcal{K}_g \cup \mathcal{K}_e & (13) \\
& \quad n_k \in [0, 1] \quad \forall k \in \mathcal{K}_g & (14) \\
& \quad \rho_{k,i} \in \{0, 1\} \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k & (15)
\end{aligned}$$

2) *Generator Contingency*: Each generator contingency imposes Constraints (5), (6), and (7) to enforce the power balance, the thermal limits, and the generation bounds under the generator contingency. The only difference from the base case is Constraint (8) that specifies that generator  $k$  should remain inactive under its contingency.

3) *Automatic Primary Response for Generator Contingency*: To address generator contingencies, Constraints (9)–(11) implements an APR [1], [12]. The formulation in this paper is based on the APR model used in [3], [23] where, for each generator contingency  $k$ , a system-wide signal  $n_k \in [0, 1]$  represents the level of system response required to resolve the power imbalance. Furthermore, the APR model assumes that the change in synchronized generation dispatch under a contingency is proportionate to the droop slope, which is determined by the product of generator capacity  $\hat{g}$  and the predefined parameter  $\gamma$  as in [3]. The generator capacities are defined as  $\hat{\mathbf{g}} = \bar{\mathbf{g}} - \underline{\mathbf{g}}$ . The APR constraints ensure that the generation dispatch under generator contingency remains within the generation limits. The mathematical expression for constraining the generation dispatch for synchronized generators can be defined as:

$$g_{k,i} = \min\{g_i + n_k \gamma_i \hat{g}_i, \bar{g}_i\}, \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k. \quad (16)$$

To represent the disjunctive constraint (16), binary variables  $\rho_{k,i}$  are introduced. For all  $i \in \mathcal{G}$  and  $k \in \mathcal{K}_g$  such that  $k \neq i$ , these binary variables enable the mathematical representation of the above disjunctive constraint (16). The binary variable  $\rho$  imposes

**Algorithm 1: Column-and-Constraint Generation Algorithm (CCGA) for Solving SCOPF.**


---

```

Initialize:  $\mathbb{K}_g = \emptyset, \mathbb{U}_g = \emptyset, \mathbb{U}_e = \emptyset$ 
1: for  $j = 0, 1, \dots$  do
2:   Solve Model 2 and obtain  $\mathbf{g}^{(j)}$ 
3:   Obtain  $\mathbf{g}_k^{(j)}$  using the binary search,  $\forall k \in \mathcal{K}_g$ 
4:   Compute the violations  $\alpha_{k,l}^g$  w.r.t. (26),
       $\forall l \in \mathcal{E}, \forall k \in \mathcal{K}_g$ 
5:   Compute the violations  $\alpha_{k,l}^e$  w.r.t. (27),
       $\forall l \in \mathcal{E}, \forall k \in \mathcal{K}_e$ 
6:   Break if  $\max\{\alpha_{k,l}^g\} \leq \epsilon$  and  $\max\{\alpha_{k,l}^e\} \leq \epsilon$ 
7:    $\mathbb{U}_g \leftarrow \{(k, l) \mid \alpha_{k,l}^g > \max\{\alpha_{k,l}^g\}/\beta\}$ 
8:    $\mathbb{U}_e \leftarrow \{(k, l) \mid \alpha_{k,l}^e > \max\{\alpha_{k,l}^e\}/\beta\}$ 
9:    $\mathbb{K}_g \leftarrow \mathbb{K}_g \cup \{k \mid \forall k \in \mathbb{U}_g\}$ 
10: end for

```

---

$g_k = \bar{g}$  if  $\rho = 1$ , and  $g_k = g + n_k \gamma \hat{g}$  otherwise. As a result, the extensive SCOPF problem 1, featuring the APR of generators under generator contingencies, is a MILP problem.

4) *Line Contingency*: Line contingencies impose Constraints (12). These constraints ensure that, during a line contingency, there is an immediate redistribution of power flow specified by the Line Outage Distribution Factor (LODF) [28], [29]. Under the line contingency of  $k$ ,  $k$ -th column vector, denoted as  $\mathbf{L}_k$ , of the LODF matrix  $\mathbf{L} \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{K}_e|}$ , delineates the redistribution of base case power flow at line  $k$ ,  $f_k$ , to the other lines, so as to ensure that there is no power flow at line  $k$ , i.e.,  $f_{k,k} = 0, \forall k \in \mathcal{K}_e$ .

5) *Slack Variables*: The constraints related to thermal limits (3), (6), and (12) are treated as *soft constraints*. Using soft constraints for the thermal limits is in accordance with the formulations used by ISOs for clearing the electricity markets in the United States [30], [31]. These soft constraints capture positive slack variables, as defined in (13), which are heavily penalized in the objective function.

**B. The Column and Constraint Generation Algorithm (CCGA)**

Solving the extensive SCOPF model 1 directly is extremely challenging for large networks, as the number of binary variables grows quadratically with respect to the number of generators. The optimization algorithm, CCGA, was proposed in [23] to address this computational challenge. This section reviews the CCGA briefly as the PDL-SCOPF framework proposed in this paper borrows and adapts one of its contributions.

The CCGA is summarized in Algorithm 1. It iteratively solves a master problem that contains a subset of the contingencies. Additional contingencies with violated constraints are then identified and added to the master problem after each iteration. Model 2 presents the master problem for the CCGA, which contains three types of constraints. The first set of constraints considers those for the base case that are the same as in the extensive problem 1. The second set (Constraint (19)) is concerned with the *provisional* generation dispatch  $\mathbf{g}'_k$  for each generator contingency  $k$  computed by the master problem. Constraint (19) ensures that  $\mathbf{g}'_k$  always remains within the limits of the APR.

**Model 2: The CCGA Master Problem.**

$$\min_{\mathbf{g}, \boldsymbol{\eta}_0, \{\mathbf{g}'_k\}_{k \in \mathcal{K}_g}, \{\boldsymbol{\rho}_k, n_k\}_{k \in \mathcal{K}_g}, \{\eta_{k,l}\}_{(k,l) \in \mathcal{U}_g \cup \mathcal{U}_e}} \mathbf{c}^\top \mathbf{g} + M_\eta \left( \|\boldsymbol{\eta}_0\|_1 + \sum_{(k,l) \in \mathcal{U}_g \cup \mathcal{U}_e} \eta_{k,l} \right) \quad (17)$$

$$\text{s.t. : } (2), (3), (4)$$

$$\boldsymbol{\eta}_0 \geq 0 \quad (18)$$

$$\mathbf{g}'_k - \mathbf{g} \leq \gamma \hat{\mathbf{g}} \quad \forall k \in \mathcal{K}_g \quad (19)$$

$$\underline{g}_i \leq g'_{k,i} \leq \bar{g}_i, \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k \quad (20)$$

$$g'_{k,k} = 0 \quad \forall k \in \mathcal{K}_g \quad (21)$$

$$\mathbf{1}^\top \mathbf{g}'_k = \mathbf{1}^\top \mathbf{d} \quad \forall k \in \mathcal{K}_g \quad (22)$$

$$|g'_{k,i} - g_i - n_k \gamma_i \hat{g}_i| \leq \hat{g}_i \rho_{k,i} \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k \quad (23)$$

$$g_i + n_k \gamma_i \hat{g}_i \geq \hat{g}_i \rho_{k,i} + \underline{g}_i \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k \quad (24)$$

$$g'_{k,i} \geq \hat{g}_i \rho_{k,i} + \underline{g}_i \quad \forall i \in \mathcal{G}, \forall k \in \mathcal{K}_g, i \neq k \quad (25)$$

$$\underline{f}_l - \eta_{k,l} \leq f_{k,l} \leq \bar{f}_l + \eta_{k,l} \quad \forall (k,l) \in \mathcal{U}_g \quad (26)$$

$$\underline{f}_l - \eta_{k,l} \leq f_l + (f_k \mathbf{L}_k)_l \quad (27)$$

$$\leq \bar{f}_l + \eta_{k,l} \quad \forall (k,l) \in \mathcal{U}_e \quad (27)$$

$$\eta_{k,l} \geq 0 \quad \forall (k,l) \in \mathcal{U}_g \cup \mathcal{U}_e \quad (28)$$

The third set of constraints considers a subset of generator contingencies  $\mathcal{K}_g$ , for which binary variables  $\rho_{k,i}$  and global signal variable  $n_k$  are defined. It expresses the proportional response on the provisional dispatch, as well as sets of likely active constraints for the generator ( $\mathcal{U}_g$ ) and line contingencies ( $\mathcal{U}_e$ ) respectively.

After solving the master problem 2, the CCGA computes the candidate generation dispatch  $\mathbf{g}_k^{(j)}$  under generator contingencies  $k$  (where  $j$  is the iteration number). Note that the provisional generation dispatch  $\mathbf{g}'_k$  is corrected by performing the *binary search*, introduced in [3], on  $n_k$  for each generator contingency  $k$ . This binary search procedure is designed to satisfy the power balance constraints. Also, note that the purpose of the provisional generation dispatch  $\mathbf{g}'_k$  in the master problem is to guarantee the existence of the candidate generation dispatches that satisfy the power balance constraints (as found by the binary search). Using  $\mathbf{g}_k^{(j)}$  and  $\mathbf{g}_k^{(j)}$ , the CCGA then calculates the thermal limit violations for all generator and line contingencies. If no violations are detected, it means that the current master problem has produced an optimal solution, and CCGA stops. Otherwise,  $\mathcal{U}_g$  and  $\mathcal{U}_e$  are updated by adding constraints with violations greater than the threshold  $\beta$ . CCGA converges in a finite number of iterations since it adds the generator contingencies with the violated thermal limit violations in each iteration.

The PDL framework for SCOPF proposed in this paper adapts the binary search of the CCGA in its end-to-end pipeline. The adaptation is presented in detail later in the paper.

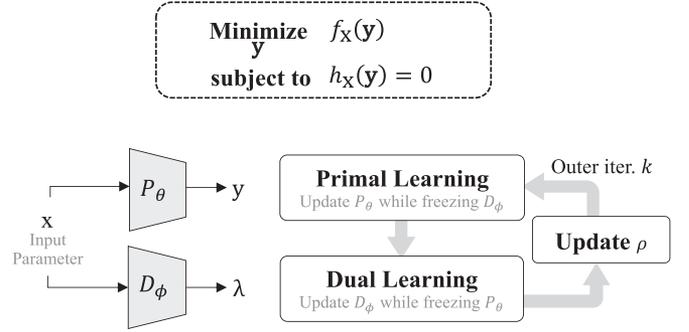


Fig. 1. Overview of the self-supervised primal-dual learning.

## IV. PRIMAL-DUAL LEARNING

This section reviews Primal-Dual Learning (PDL) [25], the machine learning framework used to approximate the SCOPF.

## A. The Augmented Lagrangian Method

Consider the following optimization problem

$$\min_{\mathbf{y}} f_{\mathbf{x}}(\mathbf{y}) \text{ subject to } \mathbf{h}_{\mathbf{x}}(\mathbf{y}) = 0. \quad (29)$$

where  $\mathbf{x}$  represents instance parameters that determine the objective function  $f_{\mathbf{x}}$  and the equality constraint  $\mathbf{h}_{\mathbf{x}}$ . The *penalty method* addresses this problem by solving a sequence of unconstrained optimization problems of the form

$$f_{\mathbf{x}}(\mathbf{y}) + \rho \mathbf{1}^\top \nu(\mathbf{h}_{\mathbf{x}}(\mathbf{y})), \quad (30)$$

where  $\nu(\cdot)$  is the element-wise violation penalty function, i.e.,  $\nu(x) = x^2$ , and  $\rho$  is the penalty coefficient. The *Augmented Lagrangian Method* (ALM) [32], [33], [34], [35] extends the penalty method and solves unconstrained optimization problems of the form

$$f_{\mathbf{x}}(\mathbf{y}) + \boldsymbol{\lambda}^\top \mathbf{h}_{\mathbf{x}}(\mathbf{y}) + \frac{\rho}{2} \mathbf{1}^\top \nu(\mathbf{h}_{\mathbf{x}}(\mathbf{y})) \quad (31)$$

where  $\boldsymbol{\lambda}$  are the Lagrangian multiplier approximations. These multipliers are updated using the rule

$$\boldsymbol{\lambda} \leftarrow \boldsymbol{\lambda} + \rho \mathbf{h}_{\mathbf{x}}(\mathbf{y}). \quad (32)$$

## B. Self-Supervised Primal-Dual Learning

PDL is a self-supervised method for training deep neural networks for learning constrained optimization problems of the form (29). PDL mimics the ALM on a set of training instances, with the aim of producing approximations to the primal and dual optimal solutions for unseen problems from the same distribution. An overview of PDL is shown in Fig. 1. PDL is composed of a primal network and a dual network that are trained iteratively in sequence to predict the primal solutions and the Lagrangian multipliers of the ALM.

At each iteration, the *primal learning* step updates the parameters  $\theta$  of the primal network  $P_\theta$  while keeping the dual network  $D_\phi$  fixed. Given the output of the frozen dual network  $\boldsymbol{\lambda}$  and the

**Algorithm 2:** Primal-Dual Learning (PDL) [25].

**Parameter:** Initial penalty coefficient  $\rho$ , Maximum outer iteration  $K$ , Maximum inner iteration  $L$ , Penalty coefficient updating multiplier  $\alpha$ , Violation tolerance  $\tau$ , Upper penalty coefficient safeguard  $\rho_{\max}$

**Input:** Input Distribution  $\mathcal{J}$

**Output:** learned primal and dual nets  $P_\theta, D_\phi$

```

1: for  $k \in \{1, \dots, K\}$  do
2:   for  $l \in \{1, \dots, L\}$  do  $\triangleright$  Primal Learning
3:     Update  $P_\theta$  using  $\nabla_\theta \mathcal{L}_p$  (See (33))
4:   end for
5:   Calculate  $v_k$  as (35)
6:   Define  $D_{\phi_k}$  by copying  $D_\phi$ 
7:   for  $l \in \{1, \dots, L\}$  do  $\triangleright$  Dual Learning
8:     Update  $D_\phi$  using  $\nabla_\phi \mathcal{L}_d$  (See (34))
9:   end for
10:  Update  $\rho$  using (36)
11: end for
12: return  $P_\theta$  and  $D_\phi$ 

```

penalty coefficient  $\rho$ , the primal learning uses the loss function

$$\mathcal{L}_p(\mathbf{y}|\lambda, \rho) = f_{\mathbf{x}}(\mathbf{y}) + \lambda^T \mathbf{h}_{\mathbf{x}}(\mathbf{y}) + \frac{\rho}{2} \mathbf{1}^T \nu(\mathbf{h}_{\mathbf{x}}(\mathbf{y})), \quad (33)$$

which is the direct counterpart of the unconstrained optimization (31). After completion of the primal learning, PDL applies a *dual learning* step that updates the parameters  $\phi$  of the dual network  $D_\phi$ . The dual learning training uses the loss function

$$\mathcal{L}_d(\lambda|\mathbf{y}, \lambda_k, \rho) = \|\lambda - (\lambda_k + \rho \mathbf{h}_{\mathbf{x}}(\mathbf{y}))\|, \quad (34)$$

which is the direct counterpart of the update rule for the Lagrangian multipliers of the ALM (32). Unlike the ALM, the Lagrangian multipliers are the output returned by the dual network. In order to update the Lagrangian multipliers, PDL employs the *copied dual network*  $D_{\phi_k}$  to obtain the Lagrangian multipliers  $\lambda_k$  of the current outer iteration  $k$  and update the dual network  $D_{\phi_k}$  using the loss function (34). This approach maintains the scalability of the learning process by avoiding to store the dual values for all training instances.

To handle severe violations, each iteration may increase the penalty coefficient  $\rho$ . At iteration  $k$ , this update uses the maximum violation  $v_k$  defined as

$$v_k = \max_{\mathbf{x} \sim \mathcal{D}} \{\|\mathbf{h}_{\mathbf{x}}(\mathbf{y})\|_\infty\}, \quad (35)$$

where  $\mathcal{D} = \{\mathbf{x}^{(i)}\}_{i=1}^N$  is the training dataset. The penalty coefficient increases when the maximum violation  $v_k$  is greater than a tolerance value  $\tau$  times the maximum violation from the previous iteration  $v_{k-1}$ , i.e.,

$$\rho \leftarrow \min\{\alpha\rho, \rho_{\max}\} \text{ if } v_k > \tau v_{k-1}, \quad (36)$$

where  $\tau \in (0, 1)$  is the tolerance,  $\alpha > 1$  is an update multiplier, and  $\rho_{\max}$  is an upper bound on the penalty coefficient. In other words, this update process increases the penalty coefficient only when the violation exceeds the specified tolerance, up to the maximum value  $\rho_{\max}$ .

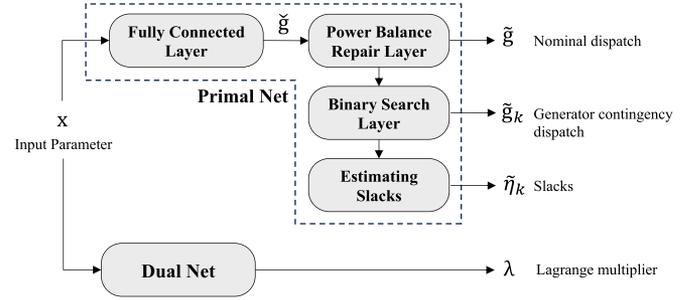


Fig. 2. The primal and dual networks of PDL-SCOPE.

The overall PDL procedure is specified in Algorithm 2, and alternates between updating the primal and dual networks during training. This iterative process resembles the ALM for constrained optimization and replaces the unconstrained optimizations and the Lagrangian updates in each outer iteration by the training of the primal and dual networks. At each iteration, these networks approximate the solutions of the unconstrained optimizations and the Lagrangian multipliers for all instances in the training set. At inference time, these networks return approximations of the primal and dual optimal solutions for unseen problem instances.

## V. PRIMAL-DUAL LEARNING FOR SCOPE

This section describes PDL-SCOPE, the framework that applies PDL to learn large-scale SCOPFs. The primal variables to approximate for the SCOPFs are  $\mathbf{y} := \tilde{\mathbf{g}}, \{\tilde{\mathbf{g}}_k, \tilde{n}_k, \tilde{\rho}_k\}_{k \in \mathcal{K}_g}, \{\tilde{\eta}_k\}_{0\} \cup \mathcal{K}_g \cup \mathcal{K}_e$ , the objective function  $f_{\mathbf{x}}(\mathbf{y})$  is the original objective function (1) of Model 1, and the constraints  $\mathbf{h}_{\mathbf{x}}(\mathbf{y})$  capture the power balance (5) for the generator contingencies as explained shortly. Fig. 2 provides a schematic representation of the primal and dual networks which, given the input configuration vector  $\mathbf{x}$ , estimate the primal and dual solutions for SCOPF, respectively. There are two key innovations in the design of the primal learning network of PDL-SCOPE: (1) the use of a repair layer for restoring the power balance of the base case and (2) the use of a binary search layer to estimate the generator dispatches in the generator contingencies.

### A. The Primal Network

The primal network estimates the nominal dispatch, the contingency dispatches, and the slacks of the thermal constraints. As shown in Fig. 2, it uses three main components: (1) a fully connected layer that produces a first approximation  $\hat{\mathbf{g}}$  of the nominal dispatch; (2) a power balance feasibility layer  $\tilde{\mathbf{g}}$  that produces a second approximation of the nominal dispatch that is guaranteed to satisfy the power balance constraints; and (3) a binary search layer that computes an approximation  $\tilde{\mathbf{g}}_k$  to the contingency dispatches by mimicking the binary search of the CCGA. *These three components, and the computation of the constraint slacks, constitutes a differentiable program for the primal learning step that is trained end-to-end.*

**Algorithm 3:** The Binary Search Layer  $\text{BSLayer}(\mathbf{g}, k)$ .**Parameter:** Maximum iteration  $t$ Initialize:  $n_k = 0.5, n_{\min} = 0, n_{\max} = 1$ 


---

```

1: for  $j = 0, 1, \dots, t$  do
2:    $g_{k,i}^{(j)} \leftarrow \min\{g_i + n_k \gamma_i \hat{g}_i, \bar{g}_i\}, \forall i \in \mathcal{G}$ 
3:    $g_{k,k}^{(j)} \leftarrow 0$ 
4:    $e_k \leftarrow \mathbf{1}^\top \mathbf{g}_k^{(j)} - \mathbf{1}^\top \mathbf{d}$ 
5:   if  $e_k > 0$  then:  $n_{\max} \leftarrow n_k$ 
6:   else:  $n_{\min} \leftarrow n_k$ 
7:    $n_k \leftarrow 0.5(n_{\max} + n_{\min})$ 
8: end for
9: for  $i \in \mathcal{G}$  do
10:  if  $g_i + n_k \gamma_i \hat{g}_i > \bar{g}_i$  then:  $\rho_{k,i} \leftarrow 1$ 
11:  else:  $\rho_{k,i} \leftarrow 0$ 
12: end for
13: return  $\mathbf{g}_k^{(j)}, n_k, \boldsymbol{\rho}_k$ 

```

---

1) *The Fully Connected Layers:* The fully connected layer produces an approximation  $\tilde{\mathbf{g}}$  of the nominal dispatch that satisfies the generator bound constraints (4). This can be achieved by applying an element-wise sigmoid function to the last layer of the fully connected neural network [19], [25].

2) *Power Balance Repair Layer for Base Case:* The power balance repair layer is borrowed from [36] and guarantees that the primal network generates a nominal dispatch that satisfies the power balance constraint (2). It receives  $\tilde{\mathbf{g}}$  as an input and generates a new nominal dispatch  $\hat{\mathbf{g}}$  by applying a proportional scaling of all generators as follows:

$$\hat{\mathbf{g}} = \begin{cases} (1 - \zeta^\uparrow)\tilde{\mathbf{g}} + \zeta^\uparrow \bar{\mathbf{g}} & \text{if } \mathbf{1}^\top \tilde{\mathbf{g}} < \mathbf{1}^\top \mathbf{d}, \\ (1 - \zeta^\downarrow)\tilde{\mathbf{g}} + \zeta^\downarrow \underline{\mathbf{g}} & \text{otherwise,} \end{cases} \quad (37)$$

where  $\zeta^\uparrow$  and  $\zeta^\downarrow$  are defined as,

$$\zeta^\uparrow = \frac{\mathbf{1}^\top \mathbf{d} - \mathbf{1}^\top \tilde{\mathbf{g}}}{\mathbf{1}^\top \bar{\mathbf{g}} - \mathbf{1}^\top \tilde{\mathbf{g}}}, \quad \zeta^\downarrow = \frac{\mathbf{1}^\top \tilde{\mathbf{g}} - \mathbf{1}^\top \mathbf{d}}{\mathbf{1}^\top \tilde{\mathbf{g}} - \mathbf{1}^\top \underline{\mathbf{g}}}. \quad (38)$$

In the case where the total generation is smaller than the total loads, i.e.,  $\mathbf{1}^\top \tilde{\mathbf{g}} < \mathbf{1}^\top \mathbf{d}$ , the coefficient  $\zeta^\uparrow$  proportionally increases  $\tilde{\mathbf{g}}$  so that the power balance for the base case is satisfied, i.e.,  $\mathbf{1}^\top \hat{\mathbf{g}} = \mathbf{1}^\top \mathbf{d}$ . The direct output from the fully connected layer  $\tilde{\mathbf{g}}$  is corrected to  $\hat{\mathbf{g}}$ , that is

$$\begin{aligned} \mathbf{1}^\top \hat{\mathbf{g}} &= (1 - \zeta^\uparrow)\mathbf{1}^\top \tilde{\mathbf{g}} + \mathbf{1}^\top \bar{\mathbf{g}} \\ &= \mathbf{1}^\top \tilde{\mathbf{g}} + \zeta^\uparrow(\mathbf{1}^\top \bar{\mathbf{g}} - \mathbf{1}^\top \tilde{\mathbf{g}}) \\ &= \mathbf{1}^\top \tilde{\mathbf{g}} + \mathbf{1}^\top \mathbf{d} - \mathbf{1}^\top \tilde{\mathbf{g}} \\ &= \mathbf{1}^\top \mathbf{d}. \end{aligned} \quad (39)$$

Similarly, when  $\mathbf{1}^\top \tilde{\mathbf{g}} \geq \mathbf{1}^\top \mathbf{d}$ , the power balance layer is designed to satisfy the power balance constraint for the base case. For more details, please refer to [36]. This power balance layer is differentiable almost everywhere and can thus be naturally integrated into the whole training process.

3) *The Binary Search Layer for Generator Contingencies:* Estimating the dispatches for all  $N - 1$  generator contingencies

presents a computational challenge, as the number of binary variables grows quadratically with the number of generator contingencies. The learning problem becomes impractical when dealing with industry-sized power grids. To address this challenge, the primal network only predicts the nominal dispatch and uses a binary search layer, inspired by the CCGA algorithm, to compute the generator dispatches under all contingencies. The *binary search layer* is an adaption of its CCGA counterpart and is described in Algorithm 3. It estimates the dispatches under the generator contingencies ( $\mathbf{g}_k$ ) and the APR-related variables ( $n_k, \boldsymbol{\rho}_k$ ) from the nominal dispatches. The algorithm performs a binary search on the global signal  $n_k$  in order to try to find contingency dispatches that satisfy the power balance constraint. Contrary to its CCGA counterpart, Algorithm 3 may not always satisfy the power balance constraint in the contingencies because  $\tilde{\mathbf{g}}_k$  does not consider constraint (19), which provides such a guarantee. In the experiment, it is observed that the provisional constraint (19) is violated in some cases in the early stage of the training, but it is satisfied when the training is complete. Algorithm 3 is described for a single training instance but it can be easily adapted to simultaneously compute these solutions for all instances and all generator contingencies in a minibatch. Note also that while the forward process of the binary search layer is conducted using Algorithm 3, the backpropagation of it to compute the first derivatives of the contingency dispatches with respect to the base case dispatches is obtainable through the expression in terms of  $\tilde{\mathbf{g}}$  and  $n_k$  (see (16)) which is efficient and differentiable almost everywhere.

4) *Retrieving the Slack Estimates for Thermal Limits:* Once the generation dispatches for the base case and generator contingencies are estimated, it is possible to calculate the slack variables for the base case, the generator contingencies (6), and the line contingencies (12). For the base case, the slack variables can be estimated by

$$\tilde{\eta}_0 = \max\{0, \mathbf{f} - \bar{\mathbf{f}}, \underline{\mathbf{f}} - \mathbf{f}\}. \quad (40)$$

The slack variables associated with generator contingency  $k \in \mathcal{K}_g$  are given by

$$\tilde{\eta}_k = \max\{0, \mathbf{K}(\mathbf{d} - \mathbf{B}\mathbf{g}_k) - \bar{\mathbf{f}}, \underline{\mathbf{f}} - \mathbf{K}(\mathbf{d} - \mathbf{B}\mathbf{g}_k)\}. \quad (41)$$

Element  $l$  of the slack variables associated with the line contingency  $k \in \mathcal{K}_e$  is given by

$$\tilde{\eta}_k = \max\{0, \mathbf{f} + f_k \mathbf{L}_k - \bar{\mathbf{f}}, \underline{\mathbf{f}} - \mathbf{f} - f_k \mathbf{L}_k\}, \quad (42)$$

where  $l \neq k$ . If  $l = k$ ,  $\eta_{k,k}^e = 0$ .

5) *The Relaxed Constraints:* The modeling of the original SCOPF problem (1) offers the advantageous property that several constraints are implicitly satisfied. As previously described, the power balance constraint for the nominal case (2) and the generation dispatch bound (4) are fulfilled through the power balance repair layer. Additionally, by using the binary search layer (Algorithm 3, all the APR-related constraints for generator contingencies (5), (7)–(11), (14), (15) are satisfied. Furthermore, the lower bounds for the slack variables (13) are also satisfied through (40), (41), and (42). Note that the thermal limit constraints (3), (6), and (12) are soft constraints. Hence, only the power balance constraints in the generator contingencies may

TABLE I  
SPECIFICATIONS OF THE SCOPF TEST CASES

Test Case	$ \mathcal{N} $	$ \mathcal{G} $	$ \mathcal{L} $	$ \mathcal{E} $	$ \mathcal{K}_g $	$ \mathcal{K}_e $	$\dim(\mathbf{x})$
300_ieee	300	69	201	411	57	322	339
1354_peg	1354	260	673	1991	193	1430	1193
1888_rte	1888	290	1000	2531	290	1567	1580
3022_goc	3022	327	1574	4135	327	3180	2228
4917_goc	4917	567	2619	6726	567	5066	3753
6515_rte	6515	684	3673	9037	657	6474	5041

be violated. It is precisely those constraints that are captured in the primal loss function of PDL-SCOPF, i.e.,

$$\mathbf{h}_x(\mathbf{y})_k = \mathbf{1}^\top \bar{\mathbf{g}}_k - \mathbf{1}^\top \mathbf{d} \quad (k \in \mathcal{K}_g).$$

### B. The Dual Network

Since all the other constraints are satisfied, the dual network produces optimal dual estimates  $\lambda = D_\phi(\mathbf{x})$  for the generator contingency power balance constraints. Experiments were run to evaluate whether adding the output of the primal network to the input of the dual network would yield performance improvements. The results were inconclusive, so the PDL-SCOPF follows the schema from Section IV.

## VI. EXPERIMENTS

### A. The Experimental Settings

1) *Test Cases*: The effectiveness of PDL-SCOPF is assessed on six specific cases from the Power Grid Library (PGLIB) [37] given in Table I. Note that the sizes of  $\mathcal{K}_g$  and  $\mathcal{K}_e$  may differ from the actual number of generators and transmission lines. Indeed, the experiments exclude contingency scenarios where the generator capacity is zero or the lower limit  $\underline{g}$  is negative, which indicate the possible presence of dispatchable loads. Regarding line contingencies, the experiments exclude transmission lines that disconnect the network completely. These lines can be identified using the LODF matrix  $\mathbf{L}$  [28].

2) *Instance Generation*: Training and testing instances obtained by perturbing the load demands, the cost coefficients, and the upper bounds of the generation dispatch, i.e.,  $x := \{\mathbf{d}, \mathbf{c}, \bar{\mathbf{g}}\}$ , in the PGLIB configuration. This generalizes earlier settings [20], [24], [25], where the only load demand  $\mathbf{d}$  is perturbed. As in [20], [24], [25], the load demands were sampled from the input distribution  $\mathcal{J}$ , which is defined as the truncated multivariate Gaussian distribution. Thus, the samples are generated as,

$$\mathbf{d} \sim \mathcal{TN}(\mathbf{d}_0, \Sigma, (1 - \mu)\mathbf{d}_0, (1 + \mu)\mu\mathbf{d}_0), \quad (43)$$

where  $\mathbf{d}_0$  is the base load demands defined in the PGLIB and  $\Sigma$  is the covariance matrix. Element  $\Sigma_{ij}$  of  $\Sigma$  is defined as  $\Sigma_{ij} = \alpha \sigma_i \sigma_j$  where  $\sigma_i$  and  $\sigma_j$  are proportional to  $d_i$  and  $d_j$  (i.e., scaled by the Z-score representing the 95th percentile of a normal Gaussian with a unit standard deviation), and the correlation coefficient  $\alpha$  is 1 when  $i = j$  and 0.5 otherwise.  $\mu$  is set to 0.5, meaning that the load demands were set to be perturbed by  $\pm 50\%$  at most. To perturb the cost coefficients and the dispatch

TABLE II  
THE NUMBER OF BINARY AND CONTINUOUS VARIABLES DENOTED AS #BV AND #CV AND CONSTRAINTS (#Cnst) IN EXTENSIVE SCOPF PROBLEM BEFORE AND AFTER PRESOLVE

Test Case	Before			After		
	#CV	#BV	#Cnst	#CV	#BV	#Cnst
300_ieee	160.2k	3.2k	327.9k	44.5k	3.2k	82.0k
1354_peg	3.3m	50.0k	6.7m	642.3k	50.0k	642.3k
1888_rte	4.8m	83.8k	9.7m	503.4k	83.8k	649.7k
3022_goc	14.6m	106.6k	29.4m	-	-	-
4917_goc	38.2m	321.5k	77.1m	-	-	-
6515_rte	64.9m	449.4k	130.7m	-	-	-

<sup>-</sup> Indicates that presolve runs out of memory.  $k$  and  $m$  signify  $10^3$  and  $10^6$ .

upper bounds, the experiments use the base values already provided in PGLIB. These base values are then adjusted by multiplying by factors specific to each instance. Specifically, per instance, two independent factors were sampled from the multivariate Gaussian distributions  $\mathcal{N}(\mathbf{1}, \Sigma)$ , where  $\Sigma$  uses a correlation coefficient  $\alpha$  of 0.8. The cost coefficients  $\mathbf{c}$  are safeguarded to be nonnegative. The dispatch upper bounds are truncated to be greater than its lower bounds as  $\bar{\mathbf{g}} \leftarrow \max\{\bar{\mathbf{g}}, \underline{\mathbf{g}} + 0.01\hat{\mathbf{g}}\}$ . The dimension of the input parameter  $\dim(\mathbf{x})$  is the same as  $2|\mathcal{G}| + |\mathcal{L}|$ . PDF-SCOPF is trained by generating instances on the fly for each update of the primal and dual networks. Note once again that there is no need to compute the optimal solutions for the instances generated during the training of PDF-SCOPF.

3) *Testing Data and Ground Truth*: To evaluate PDF-SCOPF, 1,000 instances are sampled, using the predefined distribution. Clearly, these testing instances are unseen during the training process and are generated after training. To obtain the optimal solution, testing instances were solved using the CCGA using the parameter settings suggested in [3]. The penalty coefficient for slacks  $M_\eta$  is set to 1500 as in [36]. The code is implemented based on Pyomo [38], and solved using Gurobi 10.0 [39] on 24 physical cores of a machine with Intel Xeon 2.7 GHz with 256 GB RAM. Table II highlights the size of the extensive SCOPF formulation (1): it gives the numbers of variables and constraints before and after applying the Presolve (with the default setting) in Gurobi. Presolve cannot even complete due to out of memory for the larger test cases. For the 3022\_goc, there are about 14.6 million continuous variables and 106,600 binary variables. In the largest case 6515\_rte, there are about 64.9 million continuous variables, 449,400 binary variables, and 130.7 million constraints. Table III presents the computing times and the number of iterations of the CCGA. Observe the challenging 6515\_rte case, which requires at least 4648 seconds to solve.

4) *Baselines*: PDL-SCOPF is evaluated against three machine learning based baselines to assess its performance. *It is important to stress that all baselines use the same primal neural architecture, including the repair layer for the power balance in the base case.* Only PDL-SCOPF also uses a dual network. The first baseline, denoted as *Penalty*, is a self-supervised framework that uses the penalty function (30) as a loss function. The other two baselines are supervised learning (SL) frameworks. The first supervised learning framework, *Naive*, uses a loss function

TABLE III  
ELAPSED TIME AND ITERATIONS FOR SOLVING SCOPF TEST INSTANCES USING CCGA

Test Case	Solving Time (s)			#Iteration		
	min	mean	max	min	mean	max
300_ieee	7.72	32.06	99.10	4	5.06	8
1354_peg	42.66	97.10	1125.91	3	4.79	7
1888_rte	171.92	654.39	2565.6	3	4.93	6
3022_goc	628.67	6932.20	15844.52	8	10.71	17
4917_goc	3720.33	8035.86	15316.95	7	10.98	16
6515_rte	4648.31	9560.78	92430.47	4	6.65	10

that minimizes the distance between the base case generation dispatch estimates and the ground truth, defined as  $\|\tilde{\mathbf{g}} - \mathbf{g}^*\|_2$ . The second supervised learning framework is inspired by the *Lagrangian Duality* (LD) approach from [20]: it combines the Naïve loss function with a penalty for constraint violations, i.e., loss function

$$\|\tilde{\mathbf{g}} - \mathbf{g}^*\|_2 + \rho \mathbf{1}^\top \nu(\mathbf{h}_x(\mathbf{y})). \quad (44)$$

Contrary to [20], the penalty coefficient  $\rho$  is not updated and is set to  $1e3$ , which was chosen after hyperparameter tuning. Note that to test supervised frameworks, a significant number of training instances need to be accumulated. For the supervised learning baselines, 10,000 training instances for the six test cases were accumulated. The results will show that solving these SCOPF problems is time-consuming, making it prohibitive for large-scale cases. However, it was felt important to compare the supervised and self-supervised frameworks and demonstrate that PDL-SCOPF can achieve a comparable or better performance without the need to solve training instances. The performance of the proposed approach as well as the other three baselines are evaluated by measuring the optimality gap and constraint violations. The optimality gap in percentage is calculated as the average value of the optimality gaps over the test instances, i.e.,  $\frac{|\text{obj}(\mathbf{y}^*) - \text{obj}(\mathbf{y})|}{|\text{obj}(\mathbf{y}^*)|}$ , where  $\text{obj}(\cdot)$  refers to the objective function of the problem,  $\mathbf{y}^*$  is the optimal solution obtained from Gurobi, and  $\mathbf{y}$  is the optimal solution estimate from the proposed method or one of the baselines. This optimality gap evaluates how much the optimal solution estimate deviates from the ground truth optimal solution obtained by Gurobi.

5) *Architectural Details*: Both the primal and dual networks consist of four fully-connected layers, each followed by Rectified Linear Unit (ReLU) activations. Layer normalization [40] is applied before the fully-connected layers for the primal network only. The number of hidden nodes in each fully-connected layer is proportional to the dimension of the input parameter, and is set to be  $1.5 \dim(\mathbf{x})$ .

6) *Training Setting*: The training uses a mini-batch size of 8 and a maximum of 1,000 epochs for supervised baselines. The models are trained using the Adam optimizer [41] with a learning rate of  $1e-4$ , which is reduced by a factor of 0.1 at 90% of the total iterations. For PDL-SCOPF, the number of outer iterations ( $K$ ) is set to 20, and the maximum number of inner iterations ( $L$ ) is set to 2,000 for both the primal and dual

TABLE IV  
MAXIMUM VIOLATIONS ON POWER BALANCE CONSTRAINTS FOR GENERATOR CONTINGENCY (IN P.U.)

Test Case	Naïve (SL)	LD (SL)	Penalty (SSL)	PDL-SCOPF (SSL)
300_ieee	0.008	0.001	0.003	0.000
1354_peg	0.022	0.003	0.000	0.000
1888_rte	0.044	0.003	0.000	0.000
3022_goc	0.088	0.006	0.000	0.000
4917_goc	0.067	0.002	0.000	0.000
6515_rte	0.001	0.000	0.000	0.000

networks, resulting in a total of 80,000 iterations. This same iteration count is applied to the self-supervised learning baseline. For PDF-SCOPF, total  $8 \times 80,000$  instances (i.e., inputs) were drawn from the input distribution  $\mathcal{J}$  on the fly for training the primal and dual networks in lines 3 and 8 of Algorithm 2. Also, for the self-supervised baseline, the same number of instances is used for training the primal network. The implementation uses PyTorch, and all models were trained on a machine equipped with a NVIDIA Tesla V100 GPU and an Intel Xeon 2.7 GHz CPU. Averaged performance results based on five independent training processes with different seeds are reported.

7) *Parameter Settings*: The PDL-SCOPF parameters are configured following the recommendations in [25]. A minor adjustment relates to updating the dual net. When the penalty coefficient becomes too high, updating the dual net can be unstable due to the substantial gap between the desired dual estimates and the current estimates in (34). To address this, the penalty coefficient used in the loss function (34) is fixed at  $1e-1$ , regardless of how the penalty coefficient is updated (36). The initial penalty coefficient used in the primal loss function (33) is set to 0.1 and may be increased up to  $\rho_{\max} = 1e8$ . For updating the penalty coefficient,  $\tau$  and  $\alpha$  are set to 0.9 and 2.0, respectively. To balance the objective function value and the constraint violations in the primal loss function (33), the objective function is divided by  $1e5$ .

## B. Numerical Results

The numerical results compare PDL-SCOPF with the ground truth and the baselines by evaluating the optimality gap and constraint violations. First, to check whether the constraints are satisfied by the optimal solution estimates from ML-based methods, Tables IV–V show the violations of the hard constraints, that is, the power balance constraint for generator contingencies (5). Table IV reports the averaged maximum violation of the generator contingency power balance constraints on testing instances. PDL-SCOPF and the SSL penalty method have negligible violations of the power balance constraints on all instances. In contrast, the supervised methods exhibit violations that can be significant, especially for the Naïve SL baseline. Table V shows the corresponding number of generator contingencies where the power balance constraint is violated and its percentages over the whole number of generator contingencies in parenthesis. The violation are calculated with a tolerance of  $1e-4$ . This table also

TABLE V  
THE NUMBER OF GENERATOR CONTINGENCIES WITH VIOLATED POWER BALANCE CONSTRAINTS (PERCENTAGES IN PARENTHESIS)

Test Case	Naïve (SL)	LD (SL)	Penalty (SSL)	PDL-SCOPF (SSL)
300_ieee	5(8.77%)	2(3.51%)	2(3.51%)	0(0.00%)
1354_peg	14(7.25%)	4(2.07%)	0(0.00%)	0(0.00%)
1888_rte	31(10.68%)	7(2.41%)	0(0.00%)	0(0.00%)
3022_goc	57(17.43%)	13(3.98%)	0(0.00%)	0(0.00%)
4917_goc	42(7.41%)	12(2.12%)	0(0.00%)	0(0.00%)
6515_rte	2(0.30%)	0(0.00%)	0(0.00%)	0(0.00%)

TABLE VI  
MEAN OPTIMALITY GAP (%)

Test Case	Naïve (SL)	LD (SL)	Penalty (SSL)	PDL-SCOPF (SSL)
300_ieee	1.021	<b>0.908</b>	3.867	2.805
1354_peg	13.447	2.700	2.533	<b>0.856</b>
1888_rte	22.115	2.436	4.969	<b>1.960</b>
3022_goc	159.116	8.008	1.312	<b>0.983</b>
4917_goc	47.212	2.096	0.454	<b>0.210</b>
6515_rte	2.419	1.292	2.069	<b>0.815</b>

Best values in bold.

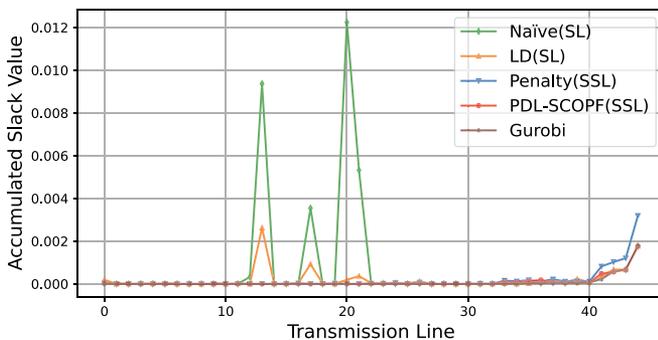


Fig. 3. Averaged total power balance constraint slack values for each transmission line in 1354\_pegase.

underscores PDL-SCOPF satisfies the hard constraints, which was a primary goal of the proposed training procedure.

Second, Table VI reports the mean optimality gap in percentage, providing a comparison between the CCGA algorithm and its learning counterparts. The results show that PDL-SCOPF is almost always the strongest method with optimality gaps often below 1%. It is only dominated on the IEEE 300-bus system by LD but this result must be treated carefully since LD violates some of the power balance constraints. Note that PDL-SCOPF significantly dominates the Penalty (SSL) method. Also, for the thermal limits for the base case and contingencies, it is desirable to have smaller slack variable values. For instance of 1354\_pegase, Fig. 3 represents the the sum of the slack values for the base case, the 193 generation contingencies, and the 1430 transmission contingencies. It only shows the slack variable values for the transmission lines in which only positive values are observed. The transmission lines are sorted by the total slack values for Naïve SL baseline in ascending order.

TABLE VII  
AVERAGED TRAINING TIME (IN GPU) FOR PDL-SCOPF AND ACCUMULATED CPU TIME TO PREPARE THE SUPERVISED TRAINING DATASET

Test Case	Training	Sampling
300_ieee	26min	89hr 3min
1354_peg	36min	269hr 43min
1888_rte	43min	1817hr 45min
3022_goc	1hr 7min	19256hr 7min
4917_goc	1hr 59min	22321hr 50min
6515_rte	2hr 54min	26557hr 43min

TABLE VIII  
AVERAGED INFERENCE TIME OF PDL-SCOPF WITH 1 INSTANCE OR 8 INSTANCES ON CPU OR GPU IN MILLISECONDS AND AVERAGED SPEEDUP OVER GUROBI ON 1 CPU

Test Case	PDL-SCOPF Inference Time (ms)				Speedup
	8@GPU	1@GPU	8@CPU	1@CPU	
300_ieee	5.163	5.117	12.392	3.527	9088.23×
1354_peg	6.917	5.146	290.376	25.745	3771.60×
1888_rte	7.956	5.145	455.846	46.876	13960.02×
3022_goc	13.486	5.823	1484.019	183.879	37699.81×
4917_goc	31.775	8.043	3975.468	479.881	16745.53×
6515_rte	51.9370	10.576	6767.474	823.016	11616.76×

PDL-SCOPF exhibits the smallest slack values when compared with the baseline and ground truth (Gurobi). Note that positive slack values are quite sparsely observed only on a few transmission lines for PDL-SCOPF, meaning that the transmission thermal limits are satisfied on most transmission lines for the contingencies. This result also contributes to the small optimality gaps of PDL-SCOPF (Table VI) because non-zero slack values lead to suboptimal solutions given their high penalties in the objective.

Table VII reports the elapsed GPU times for PDL-SCOPF training and the accumulated CPU time for generating samples for the supervised baselines. The time to generate solutions offline is extremely high, even with the CCGA. The self-supervised learning including PDL-SCOPF does not necessitate the ground truth sampling process, thus brings substantial benefits in training time. It is remarkable that, for the largest test case 6515\_rte, the training time of 2 hours and 54 minutes is significantly smaller than the time required to solve the single worst-case SCOPF instance. The latter takes approximately 25 hours and 40 minutes (92430.47 seconds equivalently, as shown in Table III) in the test dataset. This efficiency in training can be attributed to the scalable modeling and the self-supervised nature of PDL-SCOPF.

Table VIII reports the inference time in milliseconds of PDL-SCOPF on CPUs and GPUs for a single instance or a batch of 8 instances. Even for the largest test case, PDL-SCOPF provides a high-quality approximation to a single instance in about 10 milliseconds and to a batch of instances in about 50 milliseconds. Note also the significant benefits of using GPUs. These results show that PDL-SCOPF is 4 orders of magnitude faster than Gurobi on these instances.

Altogether, these results show that PDL-SCOPF, once trained, provides near-optimal solutions to SCOPF in milliseconds without the need to solve any instance offline. It provides operators with a new tool to inform real-time decision making.

## VII. CONCLUSION

This paper introduces a self-supervised primal-dual learning framework PDL-SCOPF to approximate the optimal solutions for large-scale SCOPF problems. The framework produces near-optimal feasible solutions in just milliseconds, and is self-supervised, precluding the need to generate optimal solutions to the training instances. PDL-SCOPF is a primal-dual learning method that mimics the ALM framework, alternating between approximating unconstrained optimization problems using a primal network and estimating the Lagrangian multipliers with a dual networks. The primal network is trained with a loss function that combines the original objective of the SCOPF with constraints capturing the power balance of the generator contingencies. It features a repair layer to ensure the feasibility of the power balance of the nominal case, and a binary search layer (inspired by the CCGA) that computes the contingency dispatches from the nominal dispatch. The whole framework is trained end to end, in the ALM style.

Experiment results conducted on a range of test cases, featuring systems with up to approximately 6,500 generators, demonstrated the effectiveness of the proposed methodology. By combining self-supervised learning, primal dual learning, and implicit feasibility and completion layers, PDL-SCOPF produces predictions with no constraints violations and small optimality gaps that compare dominantly with existing supervised and self-supervised methods.

This research represents an important step in demonstrating the scalability of self-supervised end-to-end optimization learning frameworks. These frameworks have the potential to transform how to learn large-scale optimization problems. In the context of SCOPF problems, potential applications include corrective SCOPF [5], stochastic SCOPF [42], and chance constraint SCOPF problems [43]. Self-supervised end-to-end learning may also be an avenue for large-scale optimization challenges in expansion planning and topology optimization.

## REFERENCES

- [1] Y. Dvorkin, P. Henneaux, D. S. Kirschen, and H. Pandžić, "Optimizing primary response in preventive security-constrained optimal power flow," *IEEE Syst. J.*, vol. 12, no. 1, pp. 414–423, Mar. 2018.
- [2] W. Chen, S. Park, M. Tanneau, and P. Van Hentenryck, "Learning optimization proxies for large-scale security-constrained economic dispatch," *Electric Power Syst. Res.*, vol. 213, 2022, Art. no. 108566.
- [3] A. Velloso, P. Van Hentenryck, and E. S. Johnson, "An exact and scalable problem decomposition for security-constrained optimal power flow," *Electric Power Syst. Res.*, vol. 195, 2021, Art. no. 106677.
- [4] F. Capitanescu et al., "State-of-the-art, challenges, and future trends in security constrained optimal power flow," *Electric Power Syst. Res.*, vol. 81, no. 8, pp. 1731–1741, 2011.
- [5] Q. Wang, J. D. McCalley, T. Zheng, and E. Litvinov, "Solving corrective risk-based security-constrained optimal power flow with Lagrangian relaxation and benders decomposition," *Int. J. Elect. Power Energy Syst.*, vol. 75, pp. 255–264, 2016.
- [6] Y. Li and J. D. McCalley, "Decomposed SCOPF for improving efficiency," *IEEE Trans. Power Syst.*, vol. 24, no. 1, pp. 494–495, Feb. 2009.
- [7] A. J. Ardakani and F. Bouffard, "Identification of umbrella constraints in DC-based security-constrained optimal power flow," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 3924–3934, Nov. 2013.
- [8] D. Bertsimas, E. Litvinov, X. A. Sun, J. Zhao, and T. Zheng, "Adaptive robust optimization for the security constrained unit commitment problem," *IEEE Trans. Power Syst.*, vol. 28, no. 1, pp. 52–63, Feb. 2013.
- [9] B. Zeng and L. Zhao, "Solving two-stage robust optimization problems using a column-and-constraint generation method," *Operations Res. Lett.*, vol. 41, no. 5, pp. 457–461, 2013.
- [10] L. Platbrood, F. Capitanescu, C. Merckx, H. Crisciuciu, and L. Wehenkel, "A generic approach for solving nonlinear-discrete security-constrained optimal power flow problems in large-scale systems," *IEEE Trans. Power Syst.*, vol. 29, no. 3, pp. 1194–1203, May 2014.
- [11] F. Capitanescu, "Critical review of recent advances and further developments needed in ac optimal power flow," *Electric Power Syst. Res.*, vol. 136, pp. 57–68, 2016.
- [12] I. Aravena et al., "Recent developments in security-constrained AC optimal power flow: Overview of challenge 1 in the ARPA-E grid optimization competition," *Operations Res.*, vol. 71, no. 6, pp. 1997–2014, 2023.
- [13] C. Coffrin and P. Van Hentenryck, "A linear-programming approximation of ac power flows," *INFORMS J. Comput.*, vol. 26, no. 4, pp. 718–734, 2014.
- [14] C. Coffrin, H. L. Hijazi, and P. Van Hentenryck, "The QC relaxation: A theoretical and computational study on optimal power flow," *IEEE Trans. Power Syst.*, vol. 31, no. 4, pp. 3008–3018, Jul. 2016.
- [15] A. Marano-Marcolini, F. Capitanescu, J. L. Martinez-Ramos, and L. Wehenkel, "Exploiting the use of DC SCOPF approximation to improve iterative ac SCOPF algorithms," *IEEE Trans. Power Syst.*, vol. 27, no. 3, pp. 1459–1466, Aug. 2012.
- [16] A. Pandey, M. R. Almassalkhi, and S. Chevalier, "Large-scale grid optimization: The workhorse of future grid computations," *Curr. Sustainable/Renewable Energy Rep.*, vol. 10, no. 3, pp. 139–153, 2023.
- [17] J. Kotary, F. Fioretto, P. Van Hentenryck, and B. Wilder, "End-to-end constrained optimization learning: A survey," 2021, *arXiv:2103.16378*.
- [18] V. Nair et al., "Solving mixed integer programs using neural networks," 2020, *arXiv:2012.13349*.
- [19] X. Pan, T. Zhao, M. Chen, and S. Zhang, "DeepOPF: A deep neural network approach for security-constrained DC optimal power flow," *IEEE Trans. Power Syst.*, vol. 36, no. 3, pp. 1725–1735, May 2021.
- [20] F. Fioretto, T. W. Mak, and P. Van Hentenryck, "Predicting AC optimal power flows: Combining deep learning and Lagrangian dual methods," in *Proc. AAAI Conf. Artif. Intell.*, 2020, vol. 34, no. 01, pp. 630–637.
- [21] A. S. Xavier, F. Qiu, and S. Ahmed, "Learning to solve large-scale security-constrained unit commitment problems," *INFORMS J. Comput.*, vol. 33, no. 2, pp. 739–756, 2021.
- [22] S. Park, W. Chen, D. Han, M. Tanneau, and P. Van Hentenryck, "Confidence-aware graph neural networks for learning reliability assessment commitments," *IEEE Trans. Power Syst.*, vol. 39, no. 2, pp. 3839–3850, Mar. 2024.
- [23] A. Velloso and P. Van Hentenryck, "Combining deep learning and optimization for preventive security-constrained DC optimal power flow," *IEEE Trans. Power Syst.*, vol. 36, no. 4, pp. 3618–3628, Jul. 2021.
- [24] P. L. Donti, D. Rolnick, and J. Z. Kolter, "DC3: A learning method for optimization with hard constraints," 2021, *arXiv:2104.12225*.
- [25] S. Park and P. Van Hentenryck, "Self-supervised primal-dual learning for constrained optimization," in *Proc. AAAI Conf. Artif. Intell.*, 2023, vol. 37, no. 4, pp. 4052–4060.
- [26] S. Park, W. Chen, T. W. Mak, and P. Van Hentenryck, "Compact optimization learning for AC optimal power flow," *IEEE Trans. Power Syst.*, vol. 39, no. 2, pp. 4350–4359, Mar. 2024.
- [27] M. Chatzos, T. W. Mak, and P. Van Hentenryck, "Spatial network decomposition for fast and scalable AC-OPF learning," *IEEE Trans. Power Syst.*, vol. 37, no. 4, pp. 2601–2612, Jul. 2022.
- [28] J. Guo, Y. Fu, Z. Li, and M. Shahidepour, "Direct calculation of line outage distribution factors," *IEEE Trans. Power Syst.*, vol. 24, no. 3, pp. 1633–1634, Aug. 2009.
- [29] D. A. Tejada-Arango, P. Sánchez-Martin, and A. Ramos, "Security constrained unit commitment using line outage distribution factors," *IEEE Trans. Power Syst.*, vol. 33, no. 1, pp. 329–337, Jan. 2018.
- [30] X. Ma, H. Song, M. Hong, J. Wan, Y. Chen, and E. Zak, "The security-constrained commitment and dispatch for midwest ISO day-ahead co-optimized energy and ancillary service market," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, 2009, pp. 1–8.
- [31] MISO, "Real-time energy and operating reserve market software formulations and business logic," 2023, business Practices Manual, Energy Operating Reserve Markets Attachment D., BPM-002-r24.

- [32] M. J. Powell, "A method for nonlinear constraints in minimization problems," *Optimization*, pp. 283–298, 1969.
- [33] R. T. Rockafellar, "Augmented lagrange multiplier functions and duality in nonconvex programming," *SIAM J. Control*, vol. 12, no. 2, pp. 268–285, 1974.
- [34] R. Andreani, E. G. Birgin, J. M. Martínez, and M. L. Schuverdt, "On augmented lagrangian methods with general lower-level constraints," *SIAM J. Optim.*, vol. 18, no. 4, pp. 1286–1309, 2008.
- [35] D. P. Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods*. New York, NY, USA: Academic Press, 2014.
- [36] W. Chen, M. Tanneau, and P. Van Hentenryck, "End-to-end feasible optimization proxies for large-scale economic dispatch," *IEEE Trans. Power Syst.*, vol. 39, no. 2, pp. 4723–4734, Mar. 2024.
- [37] S. Babaeinejadсарookolae et al., "The power grid library for benchmarking AC optimal power flow algorithms," 2019, *arXiv:1908.02788*.
- [38] W. E. Hart et al. *Pyomo-Optimization Modeling in Python*. Berlin, Germany: Springer, 2017, vol. 67.
- [39] Gurobi Optimization, LLC, "Gurobi optimizer reference manual," 2023. [Online]. Available: <https://www.gurobi.com>
- [40] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, *arXiv:1607.06450*.
- [41] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [42] H. Sharifzadeh and N. Amjady, "Stochastic security-constrained optimal power flow incorporating preventive and corrective actions," *Int. Trans. Elect. Energy Syst.*, vol. 26, no. 11, pp. 2337–2352, 2016.
- [43] L. Roald, F. Oldewurtel, B. Van Parys, and G. Andersson, "Security constrained optimal power flow with distributionally robust chance constraints," 2015, *arXiv:1508.06061*.



**Seonho Park** received the Ph.D. degree in industrial and systems engineering from the University of Florida, Gainesville, FL, USA. He is currently a Postdoctoral Fellow with the Georgia Institute of Technology, Atlanta, GA, USA. His research interests include optimization and machine learning, with a current focus on leveraging machine learning techniques to accelerate optimization processes in power system applications.



**Pascal Van Hentenryck** is currently the A. Russell Chandler III Chair and a Professor with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA. He is also the Director of the NSF AI for Advances in Optimization, Atlanta, and the Director of Tech AI, AI Hub, Georgia Tech, Atlanta. His research interests include artificial intelligence, optimization, and applications in energy, supply chains, manufacturing, and supply chains. He was the recipient of two honorary Doctoral degrees, and teaching excellence awards with Brown University, Providence, RI, USA, and Georgia Tech. Many of his optimization systems have been in commercial use for more than 20 years. He is an INFORMS Fellow and a Fellow of the Association for the Advancement of Artificial Intelligence.