Spatial Network Decomposition for Fast and Scalable AC-OPF Learning

Minas Chatzos, Terrence W.K. Mak, Member, IEEE, and Pascal Van Hentenryck, Member, IEEE

Abstract—This paper proposes a novel machine-learning approach for predicting AC-OPF solutions that features a fast and scalable training. It is motivated by the two critical considerations: (1) the fact that topology optimization and the stochasticity induced by renewable energy sources may lead to fundamentally different AC-OPF instances; and (2) the significant training time needed by existing machine-learning approaches for predicting AC-OPF. The proposed approach is a 2-stage methodology that exploits a spatial decomposition of the power network that is viewed as a set of regions. The first stage learns to predict the flows and voltages on the buses and lines coupling the regions, and the second stage trains, in parallel, the machine-learning models for each region. Experimental results on the French transmission system (up to 6,700 buses and 9,000 lines) demonstrate the potential of the approach. Within a short training time, the approach predicts AC-OPF solutions with very high fidelity and minor constraint violations, producing significant improvements over the state-of-the-art. The results also show that the predictions can seed a load flow optimization to return a feasible solution within 0.03% of the AC-OPF objective, while reducing running times significantly.

Index Terms—Optimal Power Flow; Machine Learning; Neural Networks; Network Decomposition;

I. INTRODUCTION

The AC Optimal Power Flow (AC-OPF) problem is at the core of modern power system operations. It determines the least-cost generation dispatch that meets the demand of the power grid subject to engineering and physical constraints. It is non-convex and NP-hard [1], and the basic block of many applications, including security-constrained OPF [2], [3], security-constrained unit commitment [4], optimal transmission switching [5], capacitor placement [6], and expansion planning [7], among others.

Machine learning has significant potential for real-time AC-OPF applications for a variety of reasons [8]. A machinelearning model can leverage large amount of historical data and deliver extremely fast approximations (compared to an AC-OPF solver). Recent work (e.g., [9], [8]) has indeed shown that machine-learning approaches can predict AC-OPF with high fidelity and minimal constraint violations, using a combination of neural networks and Lagrangian duality. However, the training times and memory requirements of these machine-learning models can be quite significant, which limits their potential applications. Indeed, topology optimization and the stochasticity induced by renewable energy sources may lead to fundamentally different AC-OPF instances and it is unlikely that these approaches would scale to accommodate the wide variety of input configurations encoutered in practice.

This paper explores a fundamentally different avenue: It seeks a scalable machine-learning approach for predicting AC-OPF solutions that can be trained quickly. Such an approach would make it possible to train a machine-learning model quickly to accommodate a new topology or a significant change in the capacity of renewable energy sources. It would also open the possibility of training different machine-learning models for different time periods.

To achieve this goal, the paper proposes a 2-stage machinelearning approach that exploits a spatial decomposition of the power system. The power network is viewed as a set of regions, the first stage learns to predict the flows and voltages on the buses and lines coupling the regions, and the second stage trains, in parallel, the machine-learning models for each region. Experimental results on the French transmission system (up to 6,700 buses and 9,000 lines) demonstrate the potential of the approach. Within a short training time, the approach predicts AC-OPF solutions with very high fidelity and minor constraint violations, producing significant improvements over the state-of-the-art. The results also show that the predictions can seed a load flow optimization to return a feasible solution within 0.03% of the AC-OPF objective, while reducing running times significantly.

To our knowledge, the proposed approach is the first distributed training algorithm for learning AC-OPF for largescale network topology. It builds on top, and significantly extends, prior work [9], [8] combining machine learning and Lagrangian duality. Most importantly, the 2-stage approach significantly reduces the dimensionality of the learning task, allows the training to be performed in parallel for each region, and dramatically shortens training times, opening new avenues for machine learning in very large-scale system operations. It is also the first approach that can learn AC-OPF on an actual, large-scale tranmission system fast, even on reasonable hardware configurations.

II. RELATED WORK

Machine learning has attracted significant attention in the power systems community: recent overviews of the various approaches and applications can be found in [10], [11]. In the context of AC-OPF, various approaches have been proposed for learning the active set of constraints [12], [13], [14], [15], [16], imitating the Newton-Raphson algorithm [17], or learning warm-start points for speeding-up the optimization process [18], [19]. Several approaches aim to predict optimal dispatch

The authors are affiliated with the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. E-mail: minas@gatech.edu, wmak@gatech.edu, pvh@isye.gatech.edu.

$$\mathbf{minimize} \quad \sum_{i=1}^{|\mathcal{N}|} c_i(p_i^g) \tag{1}$$

subject to:

 $\underline{v}_i \leq v_i \leq \overline{v}_i$ $\forall i \in \mathcal{N}$ (2) $\underline{p}_i^g \le p_i^g \le \overline{p}_i^g$ $\forall i \in \mathcal{N}$ (3r) $q_i^g \leq q_i^g \leq \overline{q}_i^g$ $\forall i \in \mathcal{N}$ (3i) $(p_{ij}^f)^2 + (q_{ij}^f)^2 \le |\overline{S}|_{ij}^2$ $\forall (i,j) \in \mathcal{E} \quad (4)$ $p_{ij}^f = g_{ij}v_i^2 - v_iv_j(b_{ij}\sin(\theta_i\textbf{-}\theta_j) + g_{ij}\cos(\theta_i\textbf{-}\theta_j))$ $\forall (i,j) \in \mathcal{E} \ (5r)$ $q_{ij}^{\tilde{f}} = -b_{ij}v_i^2 - v_iv_j(g_{ij}\sin(\theta_i \cdot \theta_j) - b_{ij}\cos(\theta_i \cdot \theta_j)) \quad \forall (i,j) \in \mathcal{E} \quad (5i)$ $p_i^g - p_i^d = \sum_{(i,j) \in \mathcal{E}} p_{ij}^f$ $\forall i \in \mathcal{N}$ (6r) $q_i^g - q_i^d = \sum_{(i,j)\in\mathcal{E}} q_{ij}^f$ $\forall i \in \mathcal{N}$ (6*i*)

Fig. 1: The OPF Formulation.

decisions [20], [21], [22] but these were limited to small-case studies. In the context of DC-OPF, it is worth mentioning the results of [23], [24] that provide formal guarantees on the predictions of neural networks. The application of machine learning to the security-constrained extension of the DC-OPF is presented in [25], [26].

III. PRELIMINARIES

a) The AC Optimal Power Flow Problem: A power network is modeled as an undirected graph $(\mathcal{N}, \mathcal{E})$ where \mathcal{N} and \mathcal{E} are the set of buses and transmission lines. The set of generators and loads are denoted by \mathcal{G} and \mathcal{L} . The goal of the OPF is to determine the generator dispatch of minimal cost that satisfies the load. The OPF constraints include engineering and physical constraints. The OPF formulation is shown in Figure 1. The power flow equations are expressed in terms of complex power of the form S = (p+jq), where p and q denote the active and reactive powers, admittances of the form Y = (q + ib), where q and b denote the conductance and susceptance, and voltages of the form $V = (v \angle \theta)$, with magnitude v and phase angle θ . The formulation uses v_i, θ_i, p_i^g , and q_i^g to denote the voltage magnitude, phase angle, active power generation, reactive power generation at bus *i*. Moreover, p_{ij}^{f} and q_{ij}^{f} denote the active and reactive power flows associated with line (i, j). The OPF receives as input the demand vectors p_i^d and q_i^d for each bus *i*. The objective function captures the cost of the generator dispatch. Typically, $c_i(\cdot)$ is a linear or quadratic function. Constraints (2), (3r), and (3i) capture operating bounds for the associated variables. The thermal limit for line (i, j) is captured via constraint (4). Constraints (5r) and (5i) capture Ohm's Law. Constraints (6r) and (6i)capture Kirchhoff's Law.

b) Neural Network Architectures: Neural networks have achieved tremendous success in approximating highly complex, nonlinear mappings in various domains and applications. A Neural Network (NN) consists of a series of layers, the output of each layer being the input to the next layer. The NN layers are often fully connected and the function connecting the layers is given by $o = \pi(Wx + b)$, where $x \in \mathbb{R}^n$ is the input vector, $\boldsymbol{o} \in \mathbb{R}^m$ the output vector, $\boldsymbol{W} \in \mathbb{R}^{m \times n}$ a weight matrix, and $\boldsymbol{b} \in \mathbb{R}^m$ a bias vector. The function $\pi(\cdot)$ is non-linear (e.g., a rectified linear unit (ReLU)).

c) Notations: The cardinality of set \mathcal{X} is denoted by $|\mathcal{X}|$. [N] represents the set $\{1, 2, \ldots, N\}$. Vectors are displayed using bold letters and $\boldsymbol{x} = [x_1, x_2, \ldots, x_n]^{\top}$. The elementwise lower (resp. upper) bound of the vector \boldsymbol{x} is denoted by $\underline{\boldsymbol{x}}$ (resp. $\overline{\boldsymbol{x}}$). In learning algorithms, the prediction for \boldsymbol{x} is denoted by $\hat{\boldsymbol{x}}$.

IV. LEARNING AC-OPF

A. OPF Learning Goals

Given loads (p^d, q^d) , the learning goal is to predict the optimal control setpoints (p^g, q^g) of the generators, the bus voltage v, and the phase angle difference $\Delta \theta$ of the lines. This task is equivalent to learning the complex, nonlinear, high-dimensional mapping:

$$\mathcal{O}: \mathbb{R}^{2|\mathcal{L}|} \to \mathbb{R}^{|\mathcal{N}| + |\mathcal{E}| + 2|\mathcal{G}|} \tag{7}$$

which maps the loads onto an optimal AC-OPF solution. The input to the learning task is a dataset

$$\mathcal{D} = \{(oldsymbol{p^d},oldsymbol{q^d})^t, (oldsymbol{v},oldsymbol{\Delta heta},oldsymbol{p^g},oldsymbol{q^g})^t\}_{t=1}^T$$

consisting of T instances specifying the inputs and outputs.

B. A Lagrangian Dual Model for Learning AC-OPF

One of the challenges of learning mapping \mathcal{O} is the presence of physical and engineering constraints. Ideally, given a NN $\mathcal{O}[w]$ parameterized by weights w, the goal is to find the optimal solution w^* of the problem:

$$\min_{\boldsymbol{w}} \quad \mathbb{L}_{0}(\hat{\boldsymbol{v}}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{p}}^{\boldsymbol{g}}, \hat{\boldsymbol{q}}^{\boldsymbol{g}}) \tag{10}$$
s.t. $(\hat{\boldsymbol{v}}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{p}}^{\boldsymbol{g}}, \hat{\boldsymbol{q}}^{\boldsymbol{g}}) = \mathcal{O}[\boldsymbol{w}](\boldsymbol{p}^{\boldsymbol{d}}, \boldsymbol{q}^{\boldsymbol{d}})$
 $(\hat{\boldsymbol{v}}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{p}}^{\boldsymbol{g}}, \hat{\boldsymbol{q}}^{\boldsymbol{g}}, \hat{\boldsymbol{p}}^{\boldsymbol{f}}, \hat{\boldsymbol{q}}^{\boldsymbol{f}}) \text{ satisfy (2)-(6i)}$

where \mathbb{L}_0 denotes the average norm of the difference between the ground truth and the predictions $\mathbb{L}_0(\hat{x}) = \frac{1}{T} \sum_{t=1}^{T} ||x^t - \hat{x}^t||$ over all training instances, and (\hat{p}^f, \hat{q}^f) are computed using constraints (5r) and (5i). However, it is unlikely that there exist weights w such that the predictions actually satisfy the AC-OPF constraints, since the learning task is a highdimensional regression task. However, ignoring the constraints entirely leads to predictions that significantly violate the problem constraints as shown in [8], [26]. The approach from [9], [8] addresses this difficulty by using a Lagrangian dual method relying on constraint violations. The violation of a constraint $f(x) \ge 0$ is given by $\nu_c(x) = \max\{0, -f(x)\}$, while the violation of f(x) = 0 is $\nu_c(x) = |f(x)|$. Problem (10) can then be approximated by

$$\min_{\boldsymbol{w}} \quad \mathbb{L}(\boldsymbol{\lambda}, \boldsymbol{w}) = \mathbb{L}_0(\boldsymbol{\hat{v}}, \boldsymbol{\hat{\theta}}, \boldsymbol{\hat{p}^g}, \boldsymbol{\hat{q}^g}) + \boldsymbol{\lambda}^\top \boldsymbol{\bar{\nu}}$$
(11)
s.t. $(\boldsymbol{\hat{v}}, \boldsymbol{\hat{\theta}}, \boldsymbol{\hat{p}^g}, \boldsymbol{\hat{q}^g}) = \mathcal{O}[\boldsymbol{w}](\boldsymbol{p^d}, \boldsymbol{q^d})$

where $\boldsymbol{\lambda}^{\top} \boldsymbol{\bar{\nu}} = \sum_{c \in \mathbb{C}} \lambda_c \bar{\nu}_c(\boldsymbol{\hat{v}}, \boldsymbol{\hat{\theta}}, \boldsymbol{\hat{p}^g}, \boldsymbol{\hat{q}^g}), \lambda_c$ is the weight for the violation of constraint c, and $\bar{\nu}_c$ denotes the average violation of constraint c over all training instances. Again, the satisfaction of the constraints (5r), (5i) is guaranteed, since the power flows are computed indirectly from these constraints. For a fixed λ , $\mathbb{L}(\lambda, w)$ can be used as the loss function for training the neural network. Moreover, the constraint weights can be updated using a subgradient method inspired by Lagrangian dual methods that performs the following operations in iteration j.

$$\boldsymbol{w}^{j} = \operatorname*{arg\,min}_{\boldsymbol{w}} \mathbb{L}(\boldsymbol{\lambda}^{(j-1)}, \boldsymbol{w})$$
(12)
$$\boldsymbol{\lambda}^{j} = \boldsymbol{\lambda}^{(j-1)} + \rho \bar{\boldsymbol{\nu}}(\boldsymbol{w}^{j})$$

Learning the mapping \mathcal{O} is challenging for large-scale topologies. For instance, it takes 7 hours to train a network for a topology of 3500 buses [8]. This limits the potential applications of neural networks in large power systems which may be up to 50000 buses. Indeed, during operations, the topology of the system may change from day to day through line or bus switching, meaning that a different mapping needs to be learned. Similarly, the mapping \mathcal{O} depends on the commitment decisions in the day-ahead markets, again potentially changing the mapping to be learned.

The goal of this paper is to propose *a fast training procedure* to learn the mapping \mathcal{O} . Such a fast training procedure would have many advantages: the NN model could be trained after the day-ahead market clearing and/or in real time during operations when the network topology changes, and it could be tailored to the load profiles of specific times in the day (e.g., 2:00pm-4:00pm). These considerations are important, especially given the increasing share of renewable energy in the energy mix and the increasing prediction errors.

C. Exploiting Network Sparsity

One possible avenue to obtain a fast training procedure is to exploit the sparsity typically found in power system networks. Consider a partition $\{N^k\}_{k=1}^K$ of the buses, i.e.,

$$\bigcup_{k=1}^{K} \mathcal{N}^{k} = \mathcal{N}, \quad \mathcal{N}^{k} \cap \mathcal{N}^{k'} = \emptyset, k \neq k'$$

Denote the generators and loads of region k by \mathcal{G}^k and \mathcal{L}^k respectively and define

$$\mathcal{E}^{k} = \{(i, j) \in \mathcal{E} : i, j \in \mathcal{N}^{k}\}, k \in [K].$$
$$\mathcal{E}^{\leftrightarrow} = \mathcal{E} \setminus (\bigcup_{k=1}^{K} \mathcal{E}^{k}), \quad \mathcal{N}^{\leftrightarrow} = \{i : (i, j) \in \mathcal{E}^{\leftrightarrow} \lor (j, i) \in \mathcal{E}^{\leftrightarrow}\}$$

Here \mathcal{E}^k represents the lines within partition element k and $\mathcal{E}^{\leftrightarrow}$ the coupling lines that connect partition elements. In the French transmission system, the test case in this paper, $|\mathcal{N}| = 6705$, $|\mathcal{E}| = 8962$, and $|\mathcal{E}| \approx 1.3 \times |\mathcal{N}|$. Moreover, the system is organized in 12 geographical areas using 326 (3.6%) coupling lines and $\max_{k \in [K]} |\mathcal{N}^k| = 1156$ (17.2%) buses.

To leverage the network sparsity, a natural first attempt would be to learn a mapping for each region, i.e.,

$$\mathcal{O}_0^k : \mathbb{R}^{2|\mathcal{L}^k|} \to \mathbb{R}^{2|\mathcal{N}^k|+2|\mathcal{G}^k|} \quad (k \in [K]).$$
(8)

The learning thus predicts the setpoints for generators in region k using only the loads of the same region. These learning tasks would be performed independently and in parallel. However, it is obvious that the loads \mathcal{L}^k are not sufficient to determine



Fig. 2: A Simple Network With $|\mathcal{N}| = 6$, $|\mathcal{E}| = 7$, $|\mathcal{G}| = 1$, and $|\mathcal{L}| = 4$.

the optimal setpoints for generators \mathcal{G}^k . In fact, \mathcal{O}_0^k is not even a function, since two inputs for the loads \mathcal{L}^k in the training set may be associated with different outputs due to loads in other parts of the network.

D. Capturing Flows on Coupling Lines

Consider the simplistic power system depicted in Figure 2. There are two areas, $\mathcal{N}^1 = \{1,2,3\}$ and $\mathcal{N}^2 = \{4,5,6\}$, which gives $\mathcal{N}^{\leftrightarrow} = \{3,4\}$, $\mathcal{E}^{\leftrightarrow} = \{(3,4),(4,3)\}$. The mapping \mathcal{O} views the setpoints for the generator at bus 2 (p_2^g, q_2^g) as a function of $(p_1^d, q_1^d, p_4^d, q_4^d, p_5^d, q_5^d, p_6^d, q_6^d)$. However, assume that flows $(p_{4,3}^f, q_{4,3}^f)$, along with the voltage magnitudes v_3, v_4 are fixed and respect constraints (4), (5r), (5i) associated with line (3, 4). In that case, the setpoints for the generator at bus 2 can be computed without the knowledge of $(p_4^d, q_4^d, p_5^d, q_5^d, p_6^d, q_6^d)$: the vector $(p_{4,3}^f, q_{4,3}^f, v_3)$ encodes all the information from area 2 needed to compute the generator setpoint. Hence, one may attempt to express (p_2^g, q_2^g) as a function of $(p_1^d, q_1^d, p_{4,3}^f, q_{4,3}^f, v_3)$ which decreases the input size from 8 to 5. The input size decreases by 3 in this example but the size reduction is significantly larger in actual systems.

With this in mind, the mappings in Equation 8 become

$$\mathcal{O}^{k}: \mathbb{R}^{2|\mathcal{L}^{k}| + |\mathcal{N}^{\rightarrow k}| + 2|\mathcal{E}^{\rightarrow k}|} \to \mathbb{R}^{|\mathcal{N}^{k} \setminus \mathcal{N}^{\rightarrow k}| + |\mathcal{E}^{k}| + 2|\mathcal{G}^{k}|} \tag{9}$$

where the coupling lines, buses of region k are defined as

$$\mathcal{E}^{\to k} = \{(i, j) \in \mathcal{E}^{\leftrightarrow} : i \in \mathcal{N}^k \ \lor j \in \mathcal{N}^k\}$$
$$\mathcal{N}^{\to k} = \mathcal{N}^k \cap \mathcal{N}^{\leftrightarrow}$$

 \mathcal{O}^k maps the loads in area k, the flows to area k, and the voltage of the coupling buses to the optimal generator setpoints in the area, i.e., the active and reactive outputs of the regional generators, the phase angle differences of the regional branches, and the voltage setpoints for the non-coupling buses of the region. For large transmission systems, the input/output dimensions of each mapping \mathcal{O}^k are significantly smaller that those of \mathcal{O} . The learning tasks can proceed in parallel and their complexity is reduced, since each mappings \mathcal{O}^k is an order of magnitude smaller in size than \mathcal{O} .

Unfortunately, this approach has a key limitation: each mapping \mathcal{O}^k can be learned from historical data but cannot be used for prediction since the coupling flows and voltages are not known at prediction time. Indeed, during training,

Algorithm 1: The First-Stage Coupling Training.

 $\begin{array}{c|c} \mathbf{1} & \lambda_c \leftarrow 0, \forall c \in \mathbb{C}^{\leftrightarrow} \\ \mathbf{2} \ \ \mathbf{for} \ \ i = 1, 2, ..., epochs_{\lambda} \ \mathbf{do} \\ \mathbf{3} & \quad \mathbf{for} \ \ j = 1, 2, ..., epochs_{w} \ \mathbf{do} \\ \mathbf{4} & \quad \left| \begin{array}{c} (\hat{\boldsymbol{v}}^{0}, \widehat{\boldsymbol{\Delta \theta}}^{0}) \leftarrow \mathcal{O}^{\leftrightarrow} [\boldsymbol{w}^{0}] (\boldsymbol{p}^{d}, \boldsymbol{q}^{d}) \\ \mathbb{L}^{\leftrightarrow} (\boldsymbol{\lambda}, \boldsymbol{w}^{0}) \leftarrow \mathbb{L}_{0} (\hat{\boldsymbol{v}}^{0}, \widehat{\boldsymbol{\Delta \theta}}^{0}) + \sum_{c \in \mathbb{C}^{\leftrightarrow}} \lambda_{c} \bar{\nu}_{c} (\hat{\boldsymbol{v}}^{0}, \widehat{\boldsymbol{\Delta \theta}}^{0}) \\ \mathbf{w}^{0} \leftarrow \boldsymbol{w}^{0} - \alpha \nabla_{\boldsymbol{w}^{0}} (\mathbb{L}^{\leftrightarrow} (\boldsymbol{\lambda}, \boldsymbol{w}^{0})) \\ \mathbf{s} & \quad \lambda_{c} \leftarrow \lambda_{c} + \rho \bar{\nu}_{c}, \quad \forall c \in \mathbb{C}^{\leftrightarrow} \\ \mathbf{9} \ \ \mathbf{end} \\ \mathbf{Result:} \ \ \mathrm{Weights} \ (\boldsymbol{w}^{0})^{*} \end{array}$

the learning task has access to the coupling values for each instance. However, this is not true at prediction time. The next section shows how to overcome this difficulty.

V. TWO-STAGE LEARNING OF AC-OPF

The fast training method for AC-OPF is a two-stage approach: the first stage is a NN that predicts the flow on the coupling lines and the second stage is a collection of NNs, each of which approximates a mapping \mathcal{O}^k .

A. Learning Coupling Voltages & Flows

The goal of the first stage is to learn the mapping

$$\mathcal{O}^{\leftrightarrow}: \mathbb{R}^{2|\mathcal{L}|} \to \mathbb{R}^{|\mathcal{N}^{\leftrightarrow}| + |\mathcal{E}^{\leftrightarrow}|}$$
(13)

from the loads to the voltages magnitude v^0 of the coupling buses and the phase angle difference $\Delta \theta^0$ of the coupling branches. Although the mapping considers all loads, it can be learned fast (e.g., under 30 minutes) even for large networks, because of the small number of coupling buses and lines. The coupling flows are then computed indirectly via constraints (5r) and (5i). Let $\mathbb{C}^{\leftrightarrow}$ denote the set of constraints (2) for $i \in \mathcal{N}^{\leftrightarrow}$, and (4) and (5r), (5i) for $(i, j) \in \mathcal{E}^{\leftrightarrow}$. The learning task uses a neural network $\mathcal{O}^{\leftrightarrow}[w^0]$ parameterized by weights w^0 and predicts the coupling voltages $(\hat{v}^0, \widehat{\Delta \theta}^0)$. The loss function is given by:

$$\mathbb{L}^{\leftrightarrow}(oldsymbol{\lambda},oldsymbol{w}^0) = \mathbb{L}_0(oldsymbol{\hat{v}}^0,\widehat{oldsymbol{\Delta}oldsymbol{ heta}}^0) + \sum_{c\in\mathbb{C}^{\leftrightarrow}}\lambda_car{
u}_c(oldsymbol{\hat{v}}^0,\widehat{oldsymbol{\Delta}oldsymbol{ heta}}^0)$$

The training follows equation (12) and the resulting optimal weights $(w^0)^*$ lead to the first-stage predictions

$$(\hat{\boldsymbol{v}}^0, \widehat{\boldsymbol{\Delta}} \widehat{\boldsymbol{ heta}}^0) = \mathcal{O}^{\leftrightarrow}[(\boldsymbol{w}^0)^*](\boldsymbol{p^d}, \boldsymbol{q^d})$$

and the resulting first-stage coupling flow predictions $(\hat{p}^f)^0, (\hat{q}^f)^0$. The first stage is summarized in Algorithm 1.

B. Training of Regional Systems

The training of the regional systems uses the first-stage predictions for the coupling flows and voltages. Note however that it could use the ground truth present in the instance data, but experimental results have shown that this degrades the overall prediction accuracy. The decoupling is illustrated in Figure 3, where the voltages of the coupling buses 3 and 4 and the incoming/outgoing flows for each region are fixed to the first-stage predictions.



Fig. 3: Illustration of the Decomposition.

Algorithm	2:	The	Second-Stage	Training	For	Sub-
Network k.						

1 $\lambda_c \leftarrow 0, \forall c \in \mathbb{C}^k$ 2 $(\hat{\boldsymbol{v}}^0, \widehat{\boldsymbol{\Delta}} \widehat{\boldsymbol{\theta}}^0) \leftarrow \mathcal{O}^{\leftrightarrow}[(\boldsymbol{w}^0)^*](\boldsymbol{p^d}, \boldsymbol{q^d})$ 3 Compute $(\hat{p}^f)^{0,k}, (\hat{q}^f)^{0,k}$ via (5r), (5i) 4 for $i = 1, 2, ..., epochs_{\lambda}$ do for $j = 1, 2, ..., epochs_w$ do 5 $(\hat{\boldsymbol{v}}^k, \widehat{\boldsymbol{\Delta}}\widehat{\boldsymbol{\theta}}^k, (\hat{\boldsymbol{p}}^{\boldsymbol{g}})^k, (\hat{\boldsymbol{q}}^{\boldsymbol{g}})^k) \leftarrow$ 6 $\mathcal{O}^{k}[\boldsymbol{w}^{k}]((\boldsymbol{p}^{d})^{k},(\boldsymbol{q}^{d})^{k},\hat{\boldsymbol{v}}^{0,k},(\hat{\boldsymbol{p}}^{f})^{0,k},(\hat{\boldsymbol{q}}^{f})^{0,k})$ $\mathbb{L}^k(\boldsymbol{\lambda}, \boldsymbol{w}^k) \leftarrow \mathbb{L}_0(\boldsymbol{\hat{v}}^k, \widehat{\boldsymbol{\Delta \theta}}^k, (\boldsymbol{\hat{p}^g})^k, (\boldsymbol{\hat{q}^g})^k) +$ 7 $\sum_{\substack{c \in \mathbb{C}^k \\ \boldsymbol{w}^k \leftarrow \boldsymbol{w}^k - \alpha \nabla_{\boldsymbol{w}^k} (\mathbb{L}^k(\boldsymbol{\lambda}, \boldsymbol{w}^g), \hat{\boldsymbol{q}^g}, \hat{\boldsymbol{v}^0}, (\hat{\boldsymbol{p}^f})^0, (\hat{\boldsymbol{q}^f})^0)^k)}$ 8 end 9 $\lambda_c \leftarrow \lambda_c + \rho \bar{\nu}_c, \quad \forall c \in \mathbb{C}^k$ 10 11 end **Result:** Weights $(\boldsymbol{w}^k)^*$

To learn mappings \mathcal{O}^k $(k \in [K])$, let \mathbb{C}^k denote the set of constraints associated with region k, i.e., constraint (2) for buses $i \in \mathcal{N}^k \setminus \mathcal{N}^{\to k}$, constraints (3r), (3i), (6r), and (6i) for buses $i \in \mathcal{N}^k$, and constraints (4), (5r), (5i) for branches $(i, j) \in \mathcal{E}^k$. In particular, the power balance constraint (6r), (6i) for region k becomes

$$\begin{aligned} p_i^g - (p_i^d - \sum_{(i,j)\in\mathcal{E}^{\rightarrow k}} (p_{ij}^f)^0) &= \sum_{(i,j)\in\mathcal{E}^k} p_{ij}^f \qquad i\in\mathcal{N}^k\\ q_i^g - (q_i^d - \sum_{(i,j)\in\mathcal{E}^{\rightarrow k}} (q_{ij}^f)^0) &= \sum_{(i,j)\in\mathcal{E}^k} q_{ij}^f \qquad i\in\mathcal{N}^k \end{aligned}$$

The learning task uses a collection $\{\mathcal{O}^k[\boldsymbol{w}^k]\}_{k \in [K]}$ of NNs and the loss function for each regional net is given by

$$\mathbb{L}^{k}(\boldsymbol{\lambda}, \boldsymbol{w}^{k}) = \mathbb{L}_{0}(\boldsymbol{\hat{v}}^{k}, \widehat{\boldsymbol{\Delta\theta}}^{k}, (\boldsymbol{\hat{p}}^{\boldsymbol{g}})^{k}, (\boldsymbol{\hat{q}}^{\boldsymbol{g}})^{k}) + \sum_{c \in \mathbb{C}^{k}} \lambda_{c} \bar{\nu}_{c}(\boldsymbol{\hat{v}}^{k}, \widehat{\boldsymbol{\Delta\theta}}^{k}, (\boldsymbol{\hat{p}}^{\boldsymbol{g}})^{k}, (\boldsymbol{\hat{q}}^{\boldsymbol{g}})^{k}, (\boldsymbol{\hat{q}}^{\boldsymbol{g}})^{k}, (\boldsymbol{\hat{q}}^{\boldsymbol{f}})^{0,k}, (\boldsymbol{\hat{q}}^{\boldsymbol{f}})^{0,k})$$

where $\hat{v}^{0,k}$ is the first-stage prediction for the voltage magnitude of the coupling buses of region k and $(\hat{p}^f)^{0,k}, (\hat{q}^f)^{0,k}$ the first-stage predictions for the incoming/outgoing flows of region k. The training, summarized in Algorithm 2, is performed using the approach in equation (12) and each region can be trained in parallel. Line 2 predicts the voltage setpoints for the coupling buses and the phase angle differences of the coupling lines. Line 3 computes the predicted coupling flows from these predictions. Line 6 computes the predictions for region k given the current NN parameters and constraint weights. Line 8 performs the back-propagation to update the weights and line 10 updates the constraint weights.

VI. EXPERIMENTAL RESULTS

A. Experimental Setting

The datasets were generated by solving AC-OPF problems with the nonlinear solver IPOPT [27] using 2.5 GHz-i7 Intel Cores and 16GB of RAM. In total, 10^4 load profiles, which correspond to feasible AC-OPF problems, were generated for each test case. 80% of these instances were used for training and the remaining 20% for testing. The learning models were implemented using PyTorch [28] and trained using NVidia Tesla V100 GPUs with 16GB of memory. The training of each network utilizes mini-batches of size 120 and the learning rate α was set to be decreasing from 10^{-3} to 10^{-6} , while ρ was set to 10^{-3} .

The first-stage NN consists of two subnetworks with sizes $2|\mathcal{L}| \times |\mathcal{N}^{\leftrightarrow}| \times |\mathcal{N}^{\leftrightarrow}|$ and $2|\mathcal{L}| \times |\mathcal{E}^{\leftrightarrow}| \times |\mathcal{E}^{\leftrightarrow}|$ for the voltage magnitudes and phase angle differences respectively. For each region, each NN topology consists of 4 subnetworks, one for each predicted variable. The subnetworks have one hidden layer of size $2|\mathcal{L}^k| \times 3|\mathcal{L}^k| \times |\mathcal{G}^k|$. The subnetworks used to approximate the original mapping (equation (7)) are similar in structure and have size $2|\mathcal{L}| \times 3|\mathcal{L}| \times |\mathcal{G}|$.

B. Load Profiles

The power systems used as test cases (Table I) are parts of the actual French Transmission System. France is the French transmission system, France_EHV is the very high-voltage French system, and France_LYON is France_EHV with a detailed representation of the Lyon region. The French system is organized in 12 geographical regions. The dataset is generated by taking into account this geographical information. A load l in region k with nominal value $(p^d, q^d)^0$ is generated with the formula

$$(p^{d}, q^{d}) = (\alpha + \beta^{k} + \gamma^{l})((p^{d})^{0}, (q^{d})^{0})$$

where the following coefficients are randomly drawn from the following distributions

$$\alpha \sim \text{Uniform}[0.875, 0.975]$$

$$\beta^{k} \sim \text{Uniform}[-0.025, 0.025], \quad \forall k \in [K]$$

$$\gamma^{l} \sim \text{Uniform}[-0.0025, 0.0025], \quad \forall l \in \mathcal{L}$$

The term α captures the system-wide load level, while β^k is associated with differences in the loads between regions (e.g., due to potentially different weather conditions). The difference in coefficients may be up to 5% for two different regions. Finally, γ^l is the uncorrelated noise added to each individual load with a range of 0.5% of its nominal value.

The resulting dataset captures realistic load profiles: the uniform load perturbation, the load level differences between the regions, and the fixed active/reactive power ratio represent the typical behavior for aggregated demand in a largescale topology spanning several geographic regions. Randomly perturbing each individual load in an uncorrelated fashion would produce unrealistic load profiles: they would lead to an unnecessarily challenging learning task that would need to capture an exponential number of unrealistic behaviors of the power system. To highlight this point, Figure 4 depicts the



Fig. 4: Consumption for Three Regions in the French System over a 12-hour Interval.

actual consumption for three French regions over a 12 hour interval. Observe the strong correlation of the demand between the three regions. However, the correlation is not perfect and the ratios between the regional loads vary by small factors. The term β_k is used to account for this behavior. The resulting load profiles range from 0.85 of the nominal to the nominal load. This 15% difference is typical over a 12-hour interval as shown in Figure 4.

C. First-stage Predictions

This section presents the prediction errors of the first stage. The training time was limited to 30 minutes. Table II contains aggregate results for the active and reactive powers of the coupling branches, as well the voltage magnitudes, for all three test cases. The results are an average over all instances and coupling branches. The average error is close to 1 MW for the largest two test cases: France and France Lyon. Meanwhile, the 95-Quantile indicates that 95% of the predictions result in an error less that 5 MW. In the smaller France EHV, the errors are slightly higher reaching 3.5 MW on average. Given that the nominal load of the France system is 50,000 MW) and the nominal flow values 'are greater than 100 MW and up go to 1,000 MW), these results indicate that the prediction errors are small in percentage for all test cases. Table II also shows that the voltage magnitudes are predicted very accurately. Figure 5 contains detailed results on the active part of the flows of the coupling branches for the France test case, showing consistent results across all tested instances. The 95% quantile graph indicates that the prediction errors exceed 5 MW only for a very small percentage of the test cases and branches.

D. Performance of the Learning Models

This section compares the model \mathbb{O} that directly approximates the mapping \mathcal{O} (Equation 7) with the proposed twostage approach \mathbb{D} . The results show that, with a time limit of 90 minutes, \mathbb{D} outperforms \mathbb{O} and is more scalable. In model \mathbb{D} , 30 minutes is allocated to the first stage, and 60 minutes to the second stage. The comparison is performed on the smaller systems, France_EHV and LYON, which represents

Benchmark	$ \mathcal{N} $	$ \mathcal{E} $	$ \mathcal{L} $	$ \mathcal{G} $	$ \mathcal{N}^k _{k=1}^K$	$ \mathcal{E}^{\leftrightarrow} $
France_EHV	1737	2350	1731	290	[338, 280, 233, 179, 143, 126, 124, 72, 67, 64, 57, 54]	148
France_Lyon	3411	4499	3273	771	[1158, 357, 294, 288, 264, 255, 231, 197, 184, 67, 62, 54]	219
France	6705	8962	6262	1708	[1156, 796, 748, 746, 627, 517, 497, 395, 325, 322, 298, 278]	326

TABLE I: The Power System Networks with Regional Information

		\hat{p}^f (MW)	v (P.U)		
Benchmark	Avg	95% Quantile	Avg	95% Quantile	
France_EHV	3.43	11.91	$25 \cdot 10^{-5}$	$76 \cdot 10^{-5}$	
France_Lyon	1.25	4.89	$27 \cdot 10^{-5}$	$82 \cdot 10^{-5}$	
France	0.99	4.11	$50 \cdot 10^{-5}$	$153 \cdot 10^{-5}$	

TABLE II: Absolute Errors for the Voltage Magnitude at the Coupling Buses and the Active Power Flow of the Coupling Branches.



Fig. 5: Prediction Errors (Average and 95% Quantile) over all Testing Instances for the Active Flow of the Coupling Branches for Testcase France. The Instances are Sorted in Increasing Order of Average Error.

the high-voltage French system and the high-voltage French with a detailed representation of the Lyon region. Experimental results on the full French system are only given for model \mathbb{D} , since the original model exceeds the capacity of the GPU memory. The comparison consists of three parts. The first part reports the accuracy for variables (\hat{v}, \hat{p}^g) (that are directly predicted) and the indirectly predicted \hat{p}^f . The second part considers constraint violations. The third part discusses how the predictions can be used to seed an optimization model that restores feasibility.

1) Prediction Accuracy: Figure 6 illustrate the convergence of the two models for the predicted variables \hat{v}, \hat{p}^g for a specific bus and generator from the France_LYON test case. The x-axis corresponds to test instances sorted by increasing system load. There is significant volatility in the ground truth values since instances that are close in the x-axis do not necessarily correspond to similar load vectors. Indeed, a similar overall system load may exhibit significant regional load differences. The results demonstrate that, for voltage magnitudes, model \mathbb{O} has significant errors. The same hold for active power. In constrast, model \mathbb{D} closely follows the ground truth for voltage magnitudes and exhibits minor errors for active power predictions. The difference between the two

	Mod	lel O	Model D		
Benchmark	Avg 95% Quantile		Avg	95% Quantile	
	20 10-5	105 10-5	00.10-5	at 10-5	
France_EHV	$39 \cdot 10^{-6}$	$125 \cdot 10^{-6}$	$22 \cdot 10^{-6}$	$61 \cdot 10^{-6}$	
France_Lyon	$45 \cdot 10^{-5}$	$127 \cdot 10^{-5}$	$22 \cdot 10^{-5}$	$78 \cdot 10^{-5}$	
France	-	-	$25 \cdot 10^{-5}$	$84 \cdot 10^{-5}$	

TABLE III: Prediction Errors (P.U.) for Voltage Magnitudes (\hat{v}) .

		Model O	Model \mathbb{D}		
Benchmark	Avg	Avg 95% Quantile		95% Quantile	
France_EHV	8.41	50.82	0.84	3.27	
France_Lyon	8.93	47.54	0.30	0.94	
France	-	-	0.19	0.70	

TABLE IV: Prediction Errors (MW) for Active Power (\hat{p}^g).

models is quite striking.

Tables III, IV, and V summarize the tprediction errors over all testcases, buses, generators, and lines, as well the 95%Quantile. The tables omit all power results for generators that are either off for all instances (due to potentially high cost) or constantly producing at their respective upper bounds (due to low cost). For voltage magnitudes, model \mathbb{D} divides the error in half compared to model \mathbb{O} . This difference is significant for the prediction of the power flows and constraint violations. Figure 7 demonstrates that model \mathbb{D} scales to the size of the France system and continues to produce highly accurate predictions. For active power, model delivers predictions whose errors are an order of magnitude smaller than those of model \mathbb{O} . The average errors are below 1 MW, which is small compared to the total system load ($\sim 50,000$ MW). Again, Figure 7 demonstrates that model \mathbb{D} nicely scales to the France system. The benefits of model \mathbb{D} are abundantly clear for the power flow predictions \hat{p}^f , which are indirectly predicted as a function of the predictions \hat{v} and $\hat{\theta}$. For France_LYON, the second largest test case, model O results in large errors (up to 50 MW). In contrast, model \mathbb{D} results in minor errors, with 95% of the predictions having an error of at most 1.04 MW in the largest benchmark. Compared to the overall system scale, these errors are small in percentage. Note that accurate predictions for power flows are critical for low violation degrees of the AC-OPF constraints.

2) Feasibility: Table VI reports the constraint violations for the bounds on active power and voltage magnitude (constraints (2), (3r)). Model \mathbb{D} has minor violations for 99.9% of these constraints. Table VII report the violations of the active flow conservation constraints. Again, model \mathbb{D} has an



Fig. 6: Convergence of \mathbb{O} and \mathbb{D} Illustrated for a Bus and Generator.



(a) Voltage Magnitude

(b) Active power

Model \mathbb{O}

Avg 0.036

0.281

Max

0.193

0.987

Benchmark

France_EHV

France_Lyon

France

Model \mathbb{D}

Max

0.119

0.071

0.030

Avg

0.026

0.016

0.012

Fig. 7: Prediction Errors for the France System using Model \mathbb{D} .

		Model \mathbb{O}	Model \mathbb{D}		
Benchmark	Avg	95% Quantile	Avg	95% Quantile	
France_EHV	4.53	16.88	2.01	4.20	
France_Lyon	8.43	32.20	0.82	1.91	
France	-	-	0.45	1.04	

TABLE V: Prediction Errors (MW) for Active Power Flow (\hat{p}^f) .

	Mod	el 🛈	Model \mathbb{D}		
Benchmark	v	p^g	v	p^g	
France_EHV	99.90	98.46	99.74	99.94	
France_Lyon	99.72	99.24	99.91	99.99	
France	-	-	99.97	99.99	

TABLE VI: Percentage of AC-OPF bound Constraints with Violations under 1 MW (for \hat{p}^g) and under 10^{-4} P.U. (for \hat{v})

]	Model 0	Model \mathbb{D}		
Benchmark	Avg	Avg 95% Quantile		95% Quantile	
France_EHV	4.49	16.20	4.82	9.77	
France_Lyon	18.65	100.36	1.91	4.67	
France	-	-	1.05	2.39	

TABLE VII: Violation for the Active Power Balance Constraint (MW).

TABLE	VIII:	Objective	Value	Increase	of	the	Load-Flow
Solution	(in %) Compare	d to A	C-OPF OF	viec	tive	

average 1.05 MW violations in the France test case, which is insignificant compared to the scale of the system.

3) Load-Flow Analysis: This section shows how model \mathbb{D} can also be used for applications that require a highquality *feasible* solution to AC-OPF. It presents an optimization model, called a Load Flow (L-F) model, for finding the feasible AC-OPF solution that is closest to the prediction of model \mathbb{D} , i.e.,

min
$$||\boldsymbol{p}^{\boldsymbol{g}} - \hat{\boldsymbol{p}}^{\boldsymbol{g}}||_{2}^{2} + ||\boldsymbol{v} - \hat{\boldsymbol{v}}||_{2}^{2}$$
 (14)
s.t (2) - (6*i*).

Table VIII reports the objective increase in the load flow solution with models \mathbb{D} and \mathbb{O} compared to the AC-OPF solution, i.e.,

$$\frac{1}{N_{test}} \sum_{i=1}^{N_{test}} |1 - \frac{cost_{LF}}{cost_{AC}}| \times 100\%$$

	L-F O		L-F D		AC-OPF	
Benchmark	Avg	Max	Avg	Max	Avg	Max
France_EHV	7.14	11.72	6.70	10.83	7.10	10.51
France_Lyon	38.21	186.71	20.08	26.75	75.20	215.47
France	-	-	65.21	130.67	164.07	315.97

TABLE IX: Comparison of the Load Flow and AC-OPF Computation Times (in seconds).

where $cost_{LF}$ denotes the load flow cost and $cost_{AC}$ denotes the AC-OPF cost. The load flow based on model \mathbb{D} is within < 0.03% on average of the AC-OPF solution and one magnitude smaller compared to the solutions provided by \mathbb{O} on the France_Lyon testcase.

In terms of computational efficiency, model \mathbb{D} delivers a prediction in a few milliseconds, which makes it sufficient to compare the optimization results only. Table IX compares the execution times of the load flow and the AC-OPF optimizations. The results demonstrate that the load flow optimization is significantly faster compared to the AC-OPF optimization on the largest two benchmarks. *This indicates that a combination of machine learning and optimization is beneficial when a near-optimal AC-feasible solution to the OPF is desired.*

VII. CONCLUSION

This paper considered the design of a fast and scalable training for a machine-learning model that predicts AC-OPF solutions. It was motivated by the facts that (1) topology optimization and the stochasticity induced in renewable energy may lead to fundamentally different AC-OPF instances; and (2) existing machine-learning algorithms for AC-OPF require significant training time of and do not scale to the size of real transmission systems. The paper proposed a novel 2-stage approach that exploits a spatial network decomposition. The power network is viewed as a set of regions, the first stage learns to predict the flows and voltages on the buses and lines that are coupling the regions, and the second stage trains, in parallel, the machine-learning models for each region. Experimental results on the French transmission system (up to 6,700 buses and 9,000 lines) demonstrate the potential of the approach. Within a training time of 90 minutes, the approach predicts AC-OPF solutions with very high fidelity (e.g., an average error of 1 MW for an overall load of 50 GW) and minor constraint violations, producing significant improvements over the state-of-the-art. The results also show that the predictions can be used to seed a load flow optimization that returns a feasible solution within 0.03% of the AC-OPF objective, while reducing the running times by a factor close to 3. Future work will focus on generalizing the approach to security-constrained OPF, by studying how to merge the algorithm proposed in [3] to the AC setting and the proposed 2-stage approach.

REFERENCES

 K. Lehmann, A. Grastien, and P. Van Hentenryck, "Ac-feasibility on tree networks is np-hard," *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 798–801, 2016.

- [2] A. Monticelli, M. Pereira, and S. Granville, "Security-constrained optimal power flow with post-contingency corrective rescheduling," *IEEE Transactions on Power Systems*, vol. 2, no. 1, pp. 175–180, 1987.
- [3] A. Velloso, P. V. Hentenryck, and E. S. Johnson, "An exact and scalable problem decomposition for security-constrained optimal power flow," 2019.
- [4] J. Wang, M. Shahidehpour, and Z. Li, "Security-constrained unit commitment with volatile wind power generation," *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1319–1327, 2008.
- [5] E. B. Fisher, R. P. O'Neill, and M. C. Ferris, "Optimal transmission switching," *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1346–1355, Aug 2008.
- [6] M. E. Baran and F. F. Wu, "Optimal capacitor placement on radial distribution systems," *IEEE Transactions on Power Delivery*, vol. 4, no. 1, pp. 725–734, Jan 1989.
- [7] Niharika, S. Verma, and V. Mukherjee, "Transmission expansion planning: A review," in *International Conference on Energy Efficient Tech*nologies for Sustainability, April 2016, pp. 350–355.
- [8] M. Chatzos, F. Fioretto, T. W. K. Mak, and P. V. Hentenryck, "Highfidelity machine learning approximations of large-scale optimal power flow," 2020.
- [9] F. Fioretto, T. W. Mak, and P. Van Hentenryck, "Predicting AC optimal power flows: Combining deep learning and lagrangian dual methods," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020, pp. 630–637.
- [10] L. Duchesne, E. Karangelos, and L. Wehenkel, "Machine learning of real-time power systems reliability management response," in 2017 IEEE Manchester PowerTech, 2017, pp. 1–6.
- [11] F. Hasan, A. Kargarian, and A. Mohammadi, "A survey on applications of machine learning for optimal power flow," in 2020 IEEE Texas Power and Energy Conference (TPEC), 2020, pp. 1–6.
- [12] S. Misra, L. Roald, and Y. Ng, "Learning for constrained optimization: Identifying optimal active constraint sets," 2019.
- [13] A. S. Xavier, F. Qiu, and S. Ahmed, "Learning to solve large-scale security-constrained unit commitment problems," 2019.
- [14] D. Deka and S. Misra, "Learning for DC-OPF: Classifying active sets using neural nets," in 2019 IEEE Milan PowerTech, June 2019.
- [15] F. Hasan, A. Kargarian, and J. Mohammadi, "Hybrid learning aided inactive constraints filtering algorithm to enhance ac opf solution time," 2020.
- [16] A. Robson, M. Jamei, C. Ududec, and L. Mones, "Learning an optimally reduced formulation of opf through meta-optimization," 2020.
- [17] K. Baker, "A learning-boosted quasi-newton method for ac optimal power flow," 2020.
- [18] —, "Learning warm-start points for ac optimal power flow," 2019.
- [19] L. Chen and J. E. Tate, "Hot-starting the ac power flow with convolutional neural networks," 2020.
- [20] X. Pan, T. Zhao, and M. Chen, "Deepopf: Deep neural network for dc optimal power flow," in 2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), 2019, pp. 1–6.
- [21] X. Pan, M. Chen, T. Zhao, and S. H. Low, "Deepopf: A feasibilityoptimized deep neural network approach for ac optimal power flow problems," 2020.
- [22] A. Zamzam and K. Baker, "Learning optimal solutions for extremely fast AC optimal power flow," *CoRR*, vol. abs/1910.01213, 2019. [Online]. Available: http://arxiv.org/abs/1910.01213
- [23] A. Venzke and S. Chatzivasileiadis, "Verification of neural network behaviour: Formal guarantees for power system applications," I E E E Transactions on Smart Grid, 2020.
- [24] A. Venzke, G. Qu, S. Low, and S. Chatzivasileiadis, "Learning optimal power flow: Worst-case guarantees for neural networks," 2020.
- [25] X. Pan, T. Zhao, M. Chen, and S. Zhang, "Deepopf: A deep neural network approach for security-constrained dc optimal power flow," 2020.
- [26] A. Velloso and P. V. Hentenryck, "Combining deep learning and optimization for security-constrained optimal power flow," 2020.
- [27] A. Wächter and L. T. Biegler, "On the implementation of an interiorpoint filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.
- [28] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.