

Learning Provably Stable Local Volt/Var Controllers for Efficient Network Operation

Zhenyi Yuan, *Student Member, IEEE*, Guido Cavraro, *Member, IEEE*,
Manish K. Singh, *Member, IEEE*, and Jorge Cortés, *Fellow, IEEE*

Abstract—This paper develops a data-driven framework to synthesize local Volt/Var control strategies for distributed energy resources (DERs) in power distribution grids (DGs). Aiming to improve DG operational efficiency, as quantified by a generic optimal reactive power flow (ORPF) problem, we propose a two-stage approach. The *first* stage involves learning the manifold of optimal operating points determined by an ORPF instance. To synthesize local Volt/Var controllers, the learning task is partitioned into learning local surrogates (one per DER) of the optimal manifold with voltage input and reactive power output. Since these surrogates characterize efficient DG operating points, in the *second* stage, we develop local control schemes that steer the DG to these operating points. We identify the conditions on the surrogates and control parameters to ensure that the locally acting controllers collectively converge, in a global asymptotic sense, to a DG operating point agreeing with the local surrogates. We use neural networks to model the surrogates and enforce the identified conditions in the training phase. AC power flow simulations on the IEEE 37-bus network empirically bolster the theoretical stability guarantees obtained under linearized power flow assumptions. The tests further highlight the optimality improvement compared to prevalent benchmark methods.

Index Terms—Distributed energy resources, global stability, local control, Volt/Var control.

I. INTRODUCTION

The deployment of a massive number of distributed energy resources (DERs) in power distribution grids (DGs) is dramatically changing the electric power grid. Primarily driven by sustainability and economic incentives, DERs present additional opportunities including reduction of the power generation cost and of greenhouse gas emissions. Nevertheless, DERs' uncoordinated power injections or sudden generation changes could pose challenges to system operations and stability, e.g., induce undesirable voltage deviations in distribution grids. To facilitate their integration in power grids, DERs are being provided with sensing and computation capabilities, hence becoming *smart agents*. DERs can exploit the flexibility of

their power electronic interface to perform, among other ancillary services, reactive power control. They can also take advantage of the widespread availability of data from DGs and the increased capabilities for storing and processing it to learn effective control policies. This paper aims to leverage learning in the synthesis of local Volt/Var controllers for voltage regulation incorporating optimality considerations and rigorous performance guarantees.

Literature Review: The main goal of Volt/Var control strategies is to keep voltages within safe preassigned limits by commanding DERs' reactive power injections. Classically, DER reactive power outputs are computed, in an *open-loop* fashion, by the system operator solving *optimal power flow* (OPF) problems. Efficient and advanced solvers for OPF problems are available, see e.g. [1]–[3]. However, high penetration of renewable generation and increased variability of DGs require solving numerous instances of OPF problems within a limited timeframe. Aiming at tackling this challenge, several learning-based approaches have been proposed to predict OPF solutions, see e.g. [4]–[9], to mention a few. Once trained, the inference time for these approaches when presented with a new input is minimal. Nevertheless, the solution of OPF problems requires information from *all* the buses. Specifically, power demands from loads, and generation limits from generators have to be precisely known. Such requirements are prohibitive for practical DGs since, in general, not all the buses are monitored in real time, individual loads are unlikely to announce their demand profiles in advance, and the availability of small size generators is hard to predict.

This has motivated the development of *closed-loop* strategies, which compensate for the lack of information with measurements retrieved from the field. Given the massive number of controllable devices envisioned to be hosted in future DGs, *decentralized* approaches are often advocated for practical applications. There are two notable classes of decentralized algorithms. The first consists of *distributed* algorithms in which agents cooperate and share information with peers. Distributed algorithms can achieve optimal performance, in the sense that they can be designed to exactly solve a given OPF instance, see e.g. [10]. Optimization-based feedback controllers that steer the network towards solutions to OPF problems based on the cyclical alternation of sensing, communication, and actuation have become recently popular [11]–[13]. Nevertheless, distributed strategies are suitable for systems endowed with a reliable real-time communication network meeting precise and strict requirements which are rarely satisfied in practice for DGs. For instance, in many works, each generator is required to share information with all its neighbors in the power network before every power output update [11]. The

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the NREL Laboratory Directed Research and Development Program. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes. This work was also partially supported by NSF Award ECCS-1947050.

Z. Yuan and J. Cortés are with the Department of Mechanical and Aerospace Engineering, UC San Diego. {z7yuan, cortes}@ucsd.edu. G. Cavraro is with the Power Systems Engineering Center, National Renewable Energy Laboratory. guido.cavraro@nrel.gov. M. K. Singh is with the Department of Electrical and Computer Engineering, University of Minnesota. msingh@umn.edu.

second class of decentralized algorithms consists of *local* approaches, in which each agent makes decision based only on information available locally. In local schemes, reactive power compensations are adjusted based merely on measurements taken locally. Even though pertinent standards allow DERs to provide reactive power compensations following static Volt/Var control rules, see IEEE standard 1547 [14], the literature has provided a variety of options for local voltage regulation [11], [15]–[17]. However, local schemes have intrinsic performance limitations, e.g., they might fail to regulate voltages even if the overall generation resources are enough [18].

Recent advances in data-driven and learning-based control seek to leverage data from the plant to learn optimal controllers. Though impressive results have been demonstrated, it has not yet been widely used in engineering practice due to the lack of closed-loop stability guarantees. Various results have been developed to tackle this problem, see e.g., a comprehensive review [19]. While most of these works achieve the stability of neural network control by using penalty functions to integrate stability requirement as *soft* constraints, recent work in power systems, primarily on frequency control, seek to explicitly engineer the neural network structure to integrate the stability requirement as *hard* constraints, see e.g. [20], [21]. In the context of Volt/Var control, the goal is to leverage learning techniques to enhance the performance of local control schemes and reduce its gap with distributed and/or optimal controllers, while retaining closed-loop system stability. Related works include [22] and [23], where reinforcement learning is used to learn stability-guaranteed local Volt/Var control schemes. However, the former enforces stringent derivative constraints on the policies to be searched, while the latter only guarantees the voltages converge to a region, instead of an equilibrium point, and both of them require the control policy to be continuously differentiable. Furthermore, neither of them takes into account the reactive power capacity limitations, which are critical when dealing with small-size generators. Other works provide interesting insights on learning Volt/Var rules, but do not assess the stability of the overall system and hence are not straightforwardly suitable for practical applications: [24], [25] leverage segmented linear regression techniques to learn local surrogates that predict OPF solutions, [26] proposes to learn local controller taking as an input both voltages and active power setpoints, and [27] proposes a framework for tuning the parameters of standard piecewise linear local voltage regulators.

Statement of Contributions: We propose a framework for designing local Volt/Var scheme whose goal is to not only regulate voltages but also act as local surrogates of *optimal reactive power flow* (ORPF) problem solvers. ORPF problems are particular instances of OPF problem in which the goal is to optimize the generator’s reactive power injections. We base our work on the distinction between the *control function* and the *equilibrium function*. The first represents the reactive power update rule, while the latter describes the possible system equilibrium points. In many works the equilibrium functions coincide with the control functions, but this is not the case for our framework. We advocate for a two-stage strategy. In the first stage, for each controllable node, the equilibrium

function providing ORPF solution surrogates is learned from historical data. Precisely, such a function receives as input the local voltage and gives as an output an approximation of the optimal reactive power set point. In the second stage, we devise a control function whose equilibrium points are exactly the ORPF approximated solutions provided by the equilibrium functions. The novelties of our paper with respect to the recent literature can be summarized as:

- The equilibrium functions are not forced to be (piecewise) linear and are not subject to slope limitations [11], [15], [16]. This relaxes several restrictive constraints on equilibrium functions, which leads to an enlarged search space of potential candidates of desired OPRF surrogates, thus reducing the optimality gap.
- The control rule is globally asymptotically stable, as opposed to algorithms that are locally asymptotically stable [17], [22] or whose stability is not analytically characterized. Our design provably steers the system to the desirable configuration described by equilibrium functions irrespective of the initial reactive power injection.
- We introduce the idea of pseudo data points to enhance the voltage regulation capability of learned controllers when voltages are not within desired limits.

With respect to its preliminary version [28], this work differs as follows: relaxing the differentiability requirement on the control rule; establishing global, rather than local, stability guarantees for the control scheme; providing a more general design of neural networks in the learning process; and the use of pseudo data points.

Outline: The paper is organized as follows. In Section II, we model a power distribution network and define the problem of interest. Section III introduces the control scheme, states its stability properties, and identifies conditions on the equilibrium functions needed for the system stability. Sections IV describes the learning process: precisely, it reports how we construct the dataset used for learning and how we parameterize the equilibrium function with a neural network to meet the required conditions by design. Numerical simulations in Section V validate the proposed approach and show significant improvements with respect to prevalent benchmark methods. Finally, Section VI concludes this work.

Notation: Throughout the paper, \mathbb{R} and \mathbb{C} denote the set of real and complex numbers, respectively. Upper and lower case boldface letters denote matrices and column vectors, respectively. Sets are represented by calligraphic symbols. Given a vector \mathbf{a} (a diagonal matrix \mathbf{A}), its n -th (diagonal) entry is denoted by a_n (A_n). $\mathbf{A} \succ (\succeq) \mathbf{0}$ denotes that matrix \mathbf{A} is positive (semi-) definite, and $\mathbf{A} \prec (\preceq) \mathbf{0}$ denotes that matrix \mathbf{A} is negative (semi-) definite. The symbol $(\cdot)^\top$ stands for transposition, and $\mathbf{1}$, $\mathbf{0}$, \mathbf{I} denote vectors of all ones and zeros and identity matrix with appropriate dimensions, respectively. Operators $\Re(\cdot)$ and $\Im(\cdot)$ extract the real and imaginary parts of a complex-valued argument, and act element-wise. With a slight abuse of notation, we use $|\cdot|$ to denote the absolute value for real-valued arguments, the magnitude for complex-valued arguments, and the cardinality when the argument is a set. $\|\cdot\|$ represents the Euclidean norm. Given a symmetric matrix \mathbf{A} , $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$ represent its largest and smallest eigenvalue, respectively. For any matrix \mathbf{B} , it holds that

$\|\mathbf{B}\| = \sqrt{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}$. The graph of a function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is the set of all points of the form $(x, \phi(x))$, whereas the range of ϕ is the set of its possible output values.

II. GRID MODELING AND PROBLEM FORMULATION

Consider a balanced three-phase power distribution network¹ with $N + 1$ buses represented by its single-phase equivalent and modeled as an undirected graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N} = \{0, 1, \dots, N\}$ are associated with the electrical buses and \mathcal{E} represents the set of the electrical lines between these buses. We label the substation node as 0, and assume that it behaves as an ideal voltage source imposing the nominal voltage of 1 p.u. Define the following quantities:

- $u_n \in \mathbb{C}$ is the voltage phasor at bus $n \in \mathcal{N}$;
- $v_n \in \mathbb{R}$ is the voltage magnitude at bus $n \in \mathcal{N}$;
- $i_n \in \mathbb{C}$ is the injected current phasor at bus $n \in \mathcal{N}$;
- $s_n = p_n + jq_n \in \mathbb{C}$ is the nodal complex power injection at bus $n \in \mathcal{N}$, where $p_n, q_n \in \mathbb{R}$ are the active and reactive powers, respectively. Powers take positive (negative) values, i.e., $p_n, q_n \geq 0$ ($p_n, q_n \leq 0$), when they are *injected into* (*absorbed from*) the grid.

We use vectors $\mathbf{u}, \mathbf{i}, \mathbf{s} \in \mathbb{C}^N$ to collect the complex voltages, currents, and complex powers of buses $1, 2, \dots, N$, and vectors $\mathbf{v}, \mathbf{p}, \mathbf{q} \in \mathbb{R}^N$ to collect their voltage magnitudes, active and reactive power injections. Denote by $z_e \in \mathbb{C}$ and by $y_e = z_e^{-1} \in \mathbb{C}$ respectively the impedance and the admittance of line $e = (m, n) \in \mathcal{E}$. The network bus admittance matrix $\mathbf{Y} \in \mathbb{C}^{(N+1) \times (N+1)}$ is a symmetric matrix that can be expressed as $\mathbf{Y} = \mathbf{Y}_L + \text{diag}(\mathbf{y}_T)$, where

$$(\mathbf{Y}_L)_{mn} = \begin{cases} -y_{(m,n)} & \text{if } (m, n) \in \mathcal{E}, m \neq n, \\ 0 & \text{if } (m, n) \notin \mathcal{E}, m \neq n, \\ \sum_{k \neq n} y_{(k,n)} & \text{if } m = n, \end{cases}$$

and the vector \mathbf{y}_T collects the shunt components of each bus. The matrix \mathbf{Y}_L is a complex Laplacian matrix, and hence satisfies $\mathbf{Y}_L \mathbf{1} = \mathbf{0}$. We partition the bus admittance matrix by separating the components associated with the substation and the ones associated with the other nodes, obtaining

$$\mathbf{Y} = \begin{bmatrix} y_0 & \mathbf{y}_0^\top \\ \mathbf{y}_0 & \tilde{\mathbf{Y}} \end{bmatrix},$$

with $y_0 \in \mathbb{C}$, $\mathbf{y}_0 \in \mathbb{C}^N$, and $\tilde{\mathbf{Y}} \in \mathbb{C}^{N \times N}$. If the network is connected, then $\tilde{\mathbf{Y}}$ is invertible [31]. Let $\tilde{\mathbf{Z}} := \tilde{\mathbf{Y}}^{-1}$, the power flow equations are

$$\mathbf{u} = \tilde{\mathbf{Z}} \mathbf{i} + \hat{\mathbf{u}}, \quad (1a)$$

$$u_0 = 1, \quad (1b)$$

$$i_0 = \mathbf{1}^\top \mathbf{i}, \quad (1c)$$

$$u_n \bar{v}_n = p_n + jq_n, \quad n \neq 0, \quad (1d)$$

$$v_n = |u_n| \quad (1e)$$

where \bar{v}_n denotes the complex conjugate of v_n and $\hat{\mathbf{u}} := \tilde{\mathbf{Z}} \mathbf{y}_0$. Eq. (1a) represents the Kirchoff equations and provides the

relation between voltages and currents. Eqs. (1b) and (1c) hold because the substation is modeled as the slack bus. Eq. (1d) comes from the fact that all the nodes, except the substation, are modeled to be constant power buses.

Assume a subset $\mathcal{C} \subseteq \mathcal{N}$ of buses host DERs, with $|\mathcal{C}| = C$. Every DER corresponds to a smart agent provided with some computational sensing capability, i.e., it measures its voltage magnitude. The remaining nodes constitute the set $\mathcal{L} = \mathcal{N} \setminus \mathcal{C}$ and are referred to as loads. For convenience, we partition reactive powers and voltage magnitudes by grouping together the nodes belonging to the load and generation sets

$$\mathbf{q} = [\mathbf{q}_{\mathcal{C}}^\top \quad \mathbf{q}_{\mathcal{L}}^\top]^\top, \quad \mathbf{v} = [\mathbf{v}_{\mathcal{C}}^\top \quad \mathbf{v}_{\mathcal{L}}^\top]^\top.$$

Motivated by practical considerations, We assume that the grid is not endowed with a communication network that can be used by agents, loads, and the system operator to share information in *real-time*. As a notable consequence, load demands are not known to the system operator or to agents in real time, preventing the use of open-loop strategies for computing the power outputs. The (averaged) demands may however be reported on a hourly/daily basis for slow timescale applications such as billing.

The massive deployment of DERs in DGs might induce voltage quality issues. For example, sudden generation drops could lead the voltages of a network with high penetration of renewables below desired operational limits and even close to collapse. Since DERs are able to provide ancillary services, reactive power compensation can be used to regulate voltage profiles. Ideally, one would like the DER reactive power setpoints to be the solution of an *optimal reactive power flow* (ORPF) problem of the form²

$$\mathbf{q}_{\mathcal{C}}^*(\mathbf{p}, \mathbf{q}_{\mathcal{L}}) := \arg \min_{\mathbf{q}_{\mathcal{C}}} f(\mathbf{q}_{\mathcal{C}}) \quad (2a)$$

$$\text{s.t. } (1a) - (1e)$$

$$\mathbf{v}_{\min} \leq \mathbf{v}(\mathbf{q}_{\mathcal{C}}) \leq \mathbf{v}_{\max} \quad (2b)$$

$$\mathbf{q}_{\min} \leq \mathbf{q}_{\mathcal{C}} \leq \mathbf{q}_{\max} \quad (2c)$$

where $\mathbf{q}_{\min}, \mathbf{q}_{\max} \in \mathbb{R}^C$ are the minimum and maximum DERs' reactive power injections, $\mathbf{v}_{\min}, \mathbf{v}_{\max} \in \mathbb{R}^N$ are desired voltage lower and upper bounds on *all* the network buses, and $f : \mathbb{R}^C \rightarrow \mathbb{R}$ is the cost function of interest. The minimizer depends on the uncontrolled variables \mathbf{p} and $\mathbf{q}_{\mathcal{L}}$, which appear implicitly in the constraint (2b) via equation (1e). Also, notice that for given (re)active loads and active generation, the voltage at node n and the vector of voltage magnitudes become functions exclusively of $\mathbf{q}_{\mathcal{C}}$, i.e.,

$$v_n = v_n(\mathbf{q}_{\mathcal{C}}), \quad \mathbf{v} = \mathbf{v}(\mathbf{q}_{\mathcal{C}}).$$

The complexity of solving (2) depends on the choice of the cost function $f(\cdot)$ and the non-convexity of (1). Tremendous advancements have been made to overcome the computational limitations via convex relaxations [1], linearized power-flow

¹While the analysis and algorithms developed in this work consider balanced grids, contemporary research has demonstrated applicability of related approaches to unbalanced multiphase networks [29], [30]. A rigorous and elaborate extension of our findings to encompass unbalanced grids constitutes our future research.

²More comprehensive ORPF problems could in principle be of interest in practical applications, e.g., considering line flows limitations as well. Although in this paper we focus on ORPF problems of the type (2), our approach can be readily applied to other ORPF formulations. Also, we restrict our attention to the case in which problem (2) admits a unique solution. When that is not the case $\mathbf{q}_{\mathcal{C}}^*(\mathbf{p}, \mathbf{q}_{\mathcal{L}})$ can be chosen among the set of minimizers.

equations [12], distributed optimization [10] and learning-based approaches [7]. However, solving (2) inevitably requires the knowledge of network-wide quantities ($\mathbf{p}, \mathbf{q}_{\mathcal{L}}$), centrally or via peer-to-peer communication. Since the needed supporting real-time communication infrastructure is not prevalent for most distribution systems, the optimal $\mathbf{q}_{\mathcal{C}}^*$ cannot be directly computed. Rather, inspired by the ongoing efforts towards designing local communication-free control rules for DERs and the recently reported success of neural-network-based surrogates for OPF, this work proposes a two-stage approach. In the first stage, termed the *learning stage*, historical data are used to learn functions that map voltages to (approximate) solutions of the ORPF problem (2). Specifically, for each agent $n \in \mathcal{C}$, we learn a function ϕ_n

$$\phi_n : \mathbb{R} \rightarrow \mathbb{R}, v_n \mapsto \phi_n(v_n)$$

that takes as an input the local voltage v_n , and provides as an output the approximated ORPF solution. Given any voltage v_n , $\phi_n(v_n)$ represents an approximation of the reactive power that the DER at node n would inject if its voltage is v_n and the network is operating at a solution of (2). The graph of ϕ_n , namely, points of the form $(v_n, \phi_n(v_n))$, consists then of ORPF solutions' surrogates and describes desirable network configurations. The second stage, termed the *control stage*, aims to design local control rules that steer the network to the aforesaid desirable configurations. That is, the controller equilibrium points are determined by the ϕ_n 's and, for this reason, the ϕ_n 's are called hereafter *equilibrium functions*. We introduce first in Section III the local control scheme and derive conditions ensuring system stability. We build on these conditions in Section IV to guide the learning of $\{\phi_n\}_{n \in \mathcal{C}}$ to promote the efficient operation of the DG.

III. PROVABLY STABLE LOCAL VOLT/VAR CONTROL

Here, we propose a local Volt/Var controller and analyze its stability properties. For each $n \in \mathcal{C}$, given the equilibrium function ϕ_n , consider the reactive power update

$$q_n(t+1) = q_n(t) + \epsilon(\phi_n(v_n(t)) - q_n(t)). \quad (3)$$

In (3), the new reactive power setpoint is chosen as a convex combination between the previous one and the equilibrium function evaluated at the current voltage. Rules like (3) are referred to as *incremental*, since the updated reactive power setpoint is obtained by adding to the previous one an increment weighted by the stepsize parameter $\epsilon \in [0, 1]$. An equilibrium point of (3), denoted $\mathbf{q}_{\mathcal{C}}^\sharp$, satisfies for $n \in \mathcal{C}$

$$q_n^\sharp = \phi_n(v_n) \quad (4a)$$

$$v_n^\sharp = v_n(\mathbf{q}_{\mathcal{C}}^\sharp). \quad (4b)$$

That is, $(\phi_n(v_n^\sharp), v_n^\sharp)$ belongs to the graph of ϕ_n and hence is a desirable configuration. It remains now to establish under what conditions the algorithm (3) converges to an equilibrium. The following convergence analysis assumes that uncontrolled variables, namely, \mathbf{p} and $\mathbf{q}_{\mathcal{L}}$, assume arbitrary, but fixed in time, values.

As customary in the literature of reactive power control for DGs, e.g., see [12], [13], [16], we rely on the following linearization of the power flow equations to study the stability

properties of (3). In general, voltage magnitudes are nonlinear functions of the nodal power injections. Define $\tilde{\mathbf{R}} := \mathfrak{R}(\tilde{\mathbf{Z}})$ and $\tilde{\mathbf{X}} := \mathfrak{S}(\tilde{\mathbf{Z}}) \in \mathbb{C}^{N \times N}$, using a first-order Taylor expansion, the power flow equation can be linearized to obtain [11]

$$\mathbf{v} = \tilde{\mathbf{R}}\mathbf{p} + \tilde{\mathbf{X}}\mathbf{q} + |\hat{\mathbf{u}}|. \quad (5)$$

Also, the matrices $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{X}}$ can be decomposed according to the former partition, yielding

$$\tilde{\mathbf{R}} = \begin{bmatrix} \mathbf{R} & \mathbf{R}_{\mathcal{L}} \\ \mathbf{R}_{\mathcal{L}}^\top & \mathbf{R}_{\mathcal{L}\mathcal{L}} \end{bmatrix}, \quad \tilde{\mathbf{X}} = \begin{bmatrix} \mathbf{X} & \mathbf{X}_{\mathcal{L}} \\ \mathbf{X}_{\mathcal{L}}^\top & \mathbf{X}_{\mathcal{L}\mathcal{L}} \end{bmatrix},$$

with $\mathbf{R}, \mathbf{X} \succ 0$ [16]. From (5), voltage magnitudes become functions exclusively of $\mathbf{q}_{\mathcal{C}}$:

$$\mathbf{v}(\mathbf{q}_{\mathcal{C}}) = \begin{bmatrix} \mathbf{X} \\ \mathbf{X}_{\mathcal{L}}^\top \end{bmatrix} \mathbf{q}_{\mathcal{C}} + \hat{\mathbf{v}}, \quad (6)$$

where

$$\hat{\mathbf{v}} := \begin{bmatrix} \hat{v}_{\mathcal{C}} \\ \hat{v}_{\mathcal{L}} \end{bmatrix} = \begin{bmatrix} \mathbf{X}_{\mathcal{L}} \\ \mathbf{X}_{\mathcal{L}\mathcal{L}} \end{bmatrix} \mathbf{q}_{\mathcal{L}} + \tilde{\mathbf{R}}\mathbf{p} + |\hat{\mathbf{u}}|. \quad (7)$$

Collecting the $\{\phi_n\}_{n \in \mathcal{C}}$ in the vector-valued function ϕ , and adopting the linearization (6), the system dynamics can be described in compact form as

$$\mathbf{q}_{\mathcal{C}}(t+1) = \mathbf{q}_{\mathcal{C}}(t) + \epsilon(\phi(\mathbf{v}_{\mathcal{C}}(t)) - \mathbf{q}_{\mathcal{C}}(t)), \quad (8a)$$

$$\mathbf{v}_{\mathcal{C}}(t+1) = \mathbf{X}\mathbf{q}_{\mathcal{C}}(t+1) + \hat{\mathbf{v}}_{\mathcal{C}}, \quad (8b)$$

where it is tacitly assumed that, at the timescale of the above iterates, the exogenous variables ($\mathbf{p}, \mathbf{q}_{\mathcal{L}}$) remain constant, resulting in a constant term $\hat{\mathbf{v}}_{\mathcal{C}}$ from (7). Let $(\mathbf{q}_{\mathcal{C}}^\sharp, \mathbf{v}_{\mathcal{C}}^\sharp)$ an equilibrium point of (8). By definition, it must satisfy

$$\mathbf{q}_{\mathcal{C}}^\sharp = \phi(\mathbf{v}_{\mathcal{C}}^\sharp), \quad (9a)$$

$$\mathbf{v}_{\mathcal{C}}^\sharp = \mathbf{X}\mathbf{q}_{\mathcal{C}}^\sharp + \hat{\mathbf{v}}_{\mathcal{C}}. \quad (9b)$$

In particular, it holds that $q_n^\sharp = \phi_n(v_n^\sharp)$ for each $n \in \mathcal{C}$, see (4), showing that the functions $\{\phi_n\}_{n \in \mathcal{C}}$ describe all the possible equilibrium points of (8a).

Define the reactive power injections feasible set $\mathcal{Q} := \times_{n \in \mathcal{C}} \mathcal{Q}_n$, with $\mathcal{Q}_n = \{q_n : q_{\min, n} \leq q_n \leq q_{\max, n}\}$. To ensure the stability of the proposed local algorithm, we require each ϕ_n to meet the following properties

- C1) ϕ_n is Lipschitz, i.e., there exists $L_n < \infty$ such that $|\phi_n(v) - \phi_n(v')| \leq L_n|v - v'|$, for all $v, v' \in \mathbb{R}$;
- C2) ϕ_n is non-increasing in v_n ;
- C3) $\text{range}(\phi_n) \subseteq \mathcal{Q}_n$, i.e., $\phi_n : \mathbb{R} \rightarrow \mathcal{Q}_n$.

The role of conditions C1) – C3) will be clear from the following formal results proved in the Appendix. The first one characterizes the equilibrium points of (8).

Proposition III.1. (*Feasibility of the reactive power update and uniqueness of the equilibrium*): Let $\{\phi_n\}_{n \in \mathcal{C}}$ satisfy the conditions C1) – C3), and assume $\mathbf{q}_{\mathcal{C}}(0) \in \mathcal{Q}$. The reactive power update (3) is feasible, i.e., $\mathbf{q}_{\mathcal{C}}(t) \in \mathcal{Q}$, $t \geq 0$. Moreover, the system (8) has an unique equilibrium point $(\mathbf{q}_{\mathcal{C}}^\sharp, \mathbf{v}_{\mathcal{C}}^\sharp)$.

The next result characterizes the stability properties of the equilibrium point identified in Proposition III.1.

Proposition III.2. (*Global asymptotic stability of the control rule (3)*): Let $\{\phi_n\}_{n \in \mathcal{C}}$ satisfy the conditions C1) – C3), and assume $\mathbf{q}_C(0) \in \mathcal{Q}$. Define $L = \max_{n \in \mathcal{C}} L_n$. If ϵ is such that

$$0 < \epsilon < \min \left\{ 1, \frac{2}{(\|\mathbf{X}\|L + 1)^2} \right\} \quad (10)$$

then the equilibrium of (8) is globally asymptotically stable.

Proposition III.2 indicates that, as long as $\{\phi_n\}_{n \in \mathcal{C}}$ meet the conditions C1) – C3), one can always find $\epsilon > 0$ so that $(\mathbf{q}_C, \mathbf{v}_C)$ converges to the unique equilibrium point $(\mathbf{q}_C^\#, \mathbf{v}_C^\#)$ under the reactive power update rule (3). Since \mathbf{q}_C and \mathbf{p} are fixed, the convergence of \mathbf{q}_C leads, cf. (6), also to the global asymptotic convergence of \mathbf{v} . Finally, we note that the proposed reactive power update rule (3) is a generalized version of the local control scheme proposed in [11] which only consider linear functions (instead of arbitrary nonlinear $\{\phi_n\}_{n \in \mathcal{C}}$ satisfying C1) – C3)).

Remark III.3. (*Global vs. local asymptotic stability of the equilibrium*): In our previous work [28], the equilibrium point of (8) under the reactive power update rule (3) was shown to be locally asymptotically stable if $0 < \epsilon < \min\{1, \frac{2}{\|\mathbf{X}\|L+1}\}$. The previous claim roughly implies that if $\mathbf{q}_C(0)$ is close enough to $\mathbf{q}_C^\#$, then it converges to $\mathbf{q}_C^\#$. Our result here extends the stability properties from local to global, at the cost of reducing the selection range of ϵ , as $\frac{2}{\|\mathbf{X}\|L+1} > \frac{2}{(\|\mathbf{X}\|L+1)^2}$. •

Remark III.4. (*Non-incremental vs. incremental reactive power update rules*): Many works in the literature consider local Volt/Var control schemes of the form [14]–[16]

$$q_n(t+1) = \varphi_n(v_n(t)). \quad (11)$$

Reactive power update rules like (11), where the new reactive power setpoints are determined based on the local voltage without explicitly exploiting a memory of past setpoints, are referred to as *non-incremental* [32]. The equilibrium points of (11) satisfies

$$\begin{aligned} q_n &= \varphi_n(v_n), \\ v_n &= v_n(\mathbf{q}_C), \end{aligned}$$

i.e., $\varphi_n(v_n)$ plays the double role of control function and equilibrium function. Thus, even the control rule

$$q_n(t+1) = \phi_n(v_n(t)). \quad (12)$$

looks appealing for our framework since its equilibria are determined by ϕ_n . Recall the proof of Proposition III.2, one can show that the algorithm (12) is globally asymptotically stable if the equilibrium functions meet not only C1) – C3) but also

$$\|\mathbf{X}\|L < \sqrt{2} - 1. \quad (13)$$

The former condition bounds the slope of the ϕ_n 's and appears often in the literature [11], [17], [28]. This means that, to ensure the stability of (11), we need to restrict the search space of potential candidates of $\{\phi_n\}_{n \in \mathcal{C}}$, thus risking a degradation in system performance in terms of the optimality gap at the equilibrium. The aforementioned restriction motivates adopting an incremental control like (3). •

IV. LEARNING EQUILIBRIUM FUNCTIONS FOR EFFICIENT NETWORK OPERATION

Having established the conditions on equilibrium functions for system stability, here we lay out a data-driven approach to synthesize the functions $\{\phi_n\}_{n \in \mathcal{C}}$ improving operational efficiency of the DG. Specifically, our goal is to learn local equilibrium functions $\{\phi_n\}_{n \in \mathcal{C}}$ under which the system equilibrium $\mathbf{q}_C^\#(\mathbf{p}, \mathbf{q}_C)$ is as close as possible to the ORPF problem solution $\mathbf{q}_C^*(\mathbf{p}, \mathbf{q}_C)$. The learning process consists of the following steps. First, given that the solution of (2) depends on $(\mathbf{p}, \mathbf{q}_C)$, we build a set $\{(\mathbf{p}^k, \mathbf{q}_C^k)\}_{k=1}^K$ of K load-generation scenarios. One can obtain the aforementioned scenarios via random sampling from assumed probability distributions, historical data, or from forecasted conditions for a look-ahead period. Second, we solve the ORPF problem (2) for these K scenarios to obtain a labeled dataset of corresponding minimizers $\mathcal{D} = \{(\mathbf{q}_{C,k}^*, \mathbf{v}_{C,k}^*)\}_{k=1}^K$, where the parametric dependencies are omitted for notational ease. Third, the entries for these minimizers are then separated for each $n \in \mathcal{C}$ to obtain datasets of the form $\mathcal{D}_n = \{(v_{n,k}^*, q_{n,k}^*)\}_{k=1}^K$, and each equilibrium function ϕ_n is then trained by solving

$$\begin{aligned} \min_{\phi_n} \quad & \sum_{k=1}^K |q_{n,k}^* - \phi_n(v_{n,k}^*)|^2 \\ \text{s.t.} \quad & \phi_n \text{ meets the conditions C1) – C3).} \end{aligned} \quad (14)$$

Typical approaches to solve (14) involves restricting the search space for function ϕ_n via convenient parameterization leading to approaches such as polynomial regression and neural network-based methods. Here we adopt the latter. Enforcing the properties C1) – C3) is in general not trivial and depends on many aspects, e.g., the number of considered layers and the used activation functions. In particular, designing monotone neural network may require additional considerations. Approaches to do so include, for instance, structure-based [33], gradient-constrained [34], and verification-based [35] methods. In the following, we provide a single hidden layer neural network design framework that achieves C1) – C3) and uses the ReLU activation function

$$\text{ReLU}(x) = \max(0, x).$$

Note however that, in principle, any continuous and monotonic activation function could be used in our framework, e.g., the Sigmoid or the Tanh³.

Each equilibrium function ϕ_n is then modeled using a single hidden layer neural network $\mathbf{N}(x)$ as

$$\begin{aligned} \phi_n(x) &= q_{\max,n} - \text{ReLU}(q_{\max,n} - \mathbf{N}(x)) \\ &\quad + \text{ReLU}(q_{\min,n} - \mathbf{N}(x)), \end{aligned} \quad (15)$$

where

$$\mathbf{N}(x) = \sum_{h=1}^H w_h \text{ReLU}(x - b_h) + \beta, \quad (16)$$

with w_h and b_h being the weight and bias of h -th neuron unit, respectively, β being an additional bias term applied in the output layer, and H is the number of neuron units in the

³Sigmoid(x) = $\frac{1}{1+e^{-x}}$, Tanh(x) = $\frac{e^x - e^{-x}}{e^x + e^{-x}}$.

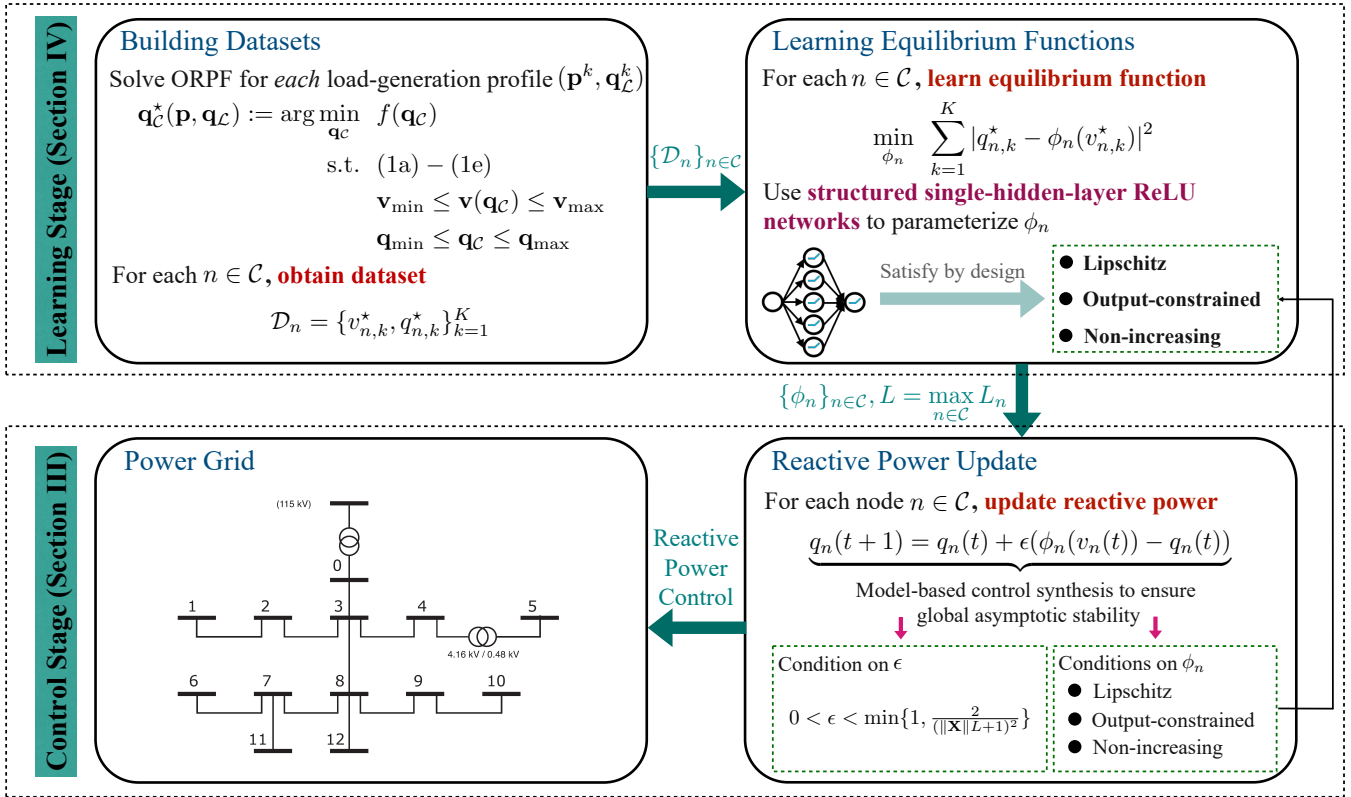


Fig. 1. Flowchart of the proposed data-driven control framework.

hidden layer. Note that conditions C1) and C3) for ϕ_n are automatically met because of the Lipschitzness of the ReLU function and the fact that the output of the neural network ϕ_n is constrained to the set \mathcal{Q}_n , cf. (15). Condition C2) is instead encoded by the next result.

Proposition IV.1. (Universal approximation of Lipschitz non-increasing function using ReLU activation function): Consider the single hidden layer neural network (16), and reorder the neuron units such that $b_1 \leq b_2 \leq \dots \leq b_H$. \mathbf{N} is non-increasing if and only if

$$\sum_{j=1}^J w_j \leq 0, \quad \forall J \in \{1, 2, \dots, H\}. \quad (17)$$

Furthermore, for any Lipschitz non-increasing function $g : \mathbb{R} \rightarrow \mathbb{R}$ and given any compact domain $\mathcal{X} \in \mathbb{R}$ and $\eta > 0$, there exist H , w_h , b_h , and β such that $|\mathbf{N}(x) - g(x)| \leq \eta$ for all $x \in \mathcal{X}$.

Its proof can be found in the Appendix, and note that monotonicity of \mathbf{N} remains after the projection (15). The above design poses weight constraints (17) and output constraints (15) to a single hidden layer neural network (16) so that it meets the conditions C1) – C3) by construction. Leveraging the universal approximation property, the optimization problem (14) is then equivalent to the parameterized formulation

$$\min_{\mathbf{w}, \mathbf{b}, \beta} \sum_{k=1}^K |q_{n,k}^* - \phi_n(v_{n,k}^*)|^2 \quad (18)$$

$$\text{s.t. (15), (16), (17)}$$

which can be solved to local optima using suitable renditions of (stochastic) gradient descent prevalent for neural network training. Also, exploiting the fact that the ReLU function is used as activation function, the Lipschitz constant L_n of each ϕ_n can be easily computed, see (17), as

$$L_n = \max_{J \in \{1, \dots, H\}} \left| \sum_{j=1}^J w_j \right|.$$

The two stages of the proposed framework are summarized in Algorithm 1. Fig. 1 provides a flowchart to illustrate the overall scheme.

Remark IV.2. (Enhancing the capability to regulate voltages when they are not within desired limits through pseudo data points): In the above exposition, the dataset points are solutions to the ORPF problem (2), which are subject to the constraint (2b). This is to say, for each $n \in \mathcal{C}$, the equilibrium function ϕ_n is trained only using data points such that $v_{\min, n} \leq v_{n,k} \leq v_{\max, n}$, i.e., not when the voltages exceed the limits. Nevertheless, in practical implementation, a DG might experience load-generation scenarios in which (2) is infeasible and the voltages do not meet the desired constraints. Engineering considerations suggest that in such cases the available reactive power capability should be maximally utilized to alleviate as much as possible the voltage violations. Namely, for each $n \in \mathcal{C}$, if $v_n < v_{\min, n}$ ($v_n > v_{\max, n}$), then $q_n = q_{\max, n}$ ($q_n = q_{\min, n}$), see [11], [15]. To ensure that the learned function ϕ_n meets this condition, we can add a certain

Algorithm 1 Local Volt/Var Control Framework

Learning Stage (Off-Line)

- Input:** Historical data in the form of K load-generation scenarios $\{(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k)\}_{k=1}^K$
- 1: Solve the ORPF problem (2) for each load-generation scenario, obtain the optimal reactive powers $\{\mathbf{q}_{\mathcal{C}}^{*,k}\}_{k=1}^K$
 - 2: Compute the optimal voltages $\{\mathbf{v}^{*,k}\}_{k=1}^K$ by solving the power flow equations (1) with the power injections $\{(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k, \mathbf{q}_{\mathcal{C}}^{*,k})\}_{k=1}^K$
 - 3: Build the dataset of optimal reactive powers and voltages $\mathcal{D} = \{(\mathbf{q}_{\mathcal{C}}^{*,k}, \mathbf{v}^{*,k})\}_{k=1}^K$
 - 4: Separate the entries of these minimizers to obtain, for each $n \in \mathcal{C}$, datasets of the form $\mathcal{D}_n = \{(v_n^{*,k}, q_n^{*,k})\}_{k=1}^K$
 - 5: Learn the equilibrium functions $\{\phi_n\}_{n \in \mathcal{C}}$ by solving (18)

Control Stage (On-Line)

- Input:** Learned equilibrium functions $\{\phi_n\}_{n \in \mathcal{C}}$, network parameter \mathbf{X} , constant $L = \max_{n \in \mathcal{C}} L_n$, stepsize parameter ϵ selected according to (10)

Each agent $n \in \mathcal{C}$, for $t \geq 0$, cyclically repeats:

- 1: Measuring its local voltage magnitude $v_n(t)$
 - 2: Updating the reactive power as per (3)
-

number of additional pseudo data points to the dataset, e.g., \bar{K} points of the form $\{(v_{n,k}, q_{\max,n})\}_{k=1}^{\bar{K}}$, and \bar{K} points of the form $\{(\bar{v}_{n,k}, q_{\min,n})\}_{k=1}^{\bar{K}}$, with $v_{n,k} \leq v_{\min,n}$, $\bar{v}_{n,k} \geq v_{\max,n}$. These points could be uniformly spaced or randomly sampled. Here we adopt the former method, illustrated in Fig. 2. •




Fig. 2. An illustration of the role of the pseudo data points for DER at node n . Blue and orange points respectively represent true and pseudo data points, while the dark red curves are instances of learned equilibrium functions. Adding pseudo data points helps the equilibrium functions reach maximum reactive power compensation capability when voltage exceeds the limits.

Remark IV.3. (*Deep neural network parameterization of non-increasing function via gradient penalization*): Though the single hidden layer neural network parameterization of (16) achieves universal approximation, it requires the “width” of the neural networks, i.e., H to be sufficiently large. Instead, in certain situations, deeper neural networks could achieve better approximations than the shallower ones even if they are much narrower [36]. On the other hand, the structure-based restrictions enforced on the neural networks to guarantee monotonicity may in some cases restrict expressibility and lead to unsatisfactory approximation results [35]. To overcome the aforementioned challenges, we describe here a deep neural network parameterization approach by incorporating




Fig. 3. The IEEE 37-bus feeder.

the monotonicity requirement as a penalization in the cost function of learning process. Suppose for each $n \in \mathcal{C}$, ϕ_n is parameterized by a deep neural network, and denote by $d\phi_n(v_n) \in \mathbb{R}$ the (sub)gradient of ϕ_n with respect to v_n . The cost function in (14) is then replaced by

$$\sum_{k=1}^K |q_{n,k}^* - \phi_n(v_{n,k}^*)|^2 + \gamma_n \max(0, d\phi_n(v_{n,k}^*)),$$

where $\gamma_n > 0$ is a tuning parameter. During implementation, one can gradually increase γ_n until ϕ_n is verified to be non-increasing. •

Remark IV.4. (*On the comparison with existing reinforcement learning-based approaches*): Recent literature has also investigated reinforcement learning approaches for learning stability guaranteed local Volt/Var controllers, e.g., [22], [23]. However, due to the lack of communication in the training phase, the cost function that the whole system seeks to minimize in such settings can only be separable, i.e., the summation of all the local cost functions at each node. Therefore, these approaches generally can not cope with cost functions like power losses which shows coupling among nodes. In contrast, since our approach only uses off-line collected data, any type of cost function could in principle be considered when solving the ORPF problem (2). •

V. NUMERICAL TESTS

We conduct case studies on the IEEE 37-bus feeder. We omit regulators, incorporate five solar generators, and convert it to its single-phase equivalent, see Fig. 3. The feeder has 25 buses with non-zero load, and the five solar generators are the DERs participating in reactive power compensation.

For our experiments, we use minute-based load and solar generation data retrieved from the Pecan Street dataset (June 1, 2018) [37]. The first 75 non-zero load buses from the dataset are aggregated every 3 loads and normalized to obtain 25 load profiles. Similarly, we obtain 5 solar generation profiles for the active power of DERs. The normalized load profiles for the 24-hour period are scaled so that the demand peak is 1.65 times the nominal load. We synthesize reactive loads by scaling active demand to match the power factors of the




Fig. 4. Minute-based data for the total (feeder-wise) solar power generation and active power demand.

IEEE 37-bus feeder. Fig. 4 shows the total demand and solar generation across the feeder.

To evaluate the effectiveness of our control framework, it is compared with the standard linear droop control from [11], [15] which dictates

$$q_n(t+1) = \varrho_n(v_n(t))$$

$$\varrho_n(v_n) := \begin{cases} q_{\max,n} & v_n(t) \leq v_{\min,n}, \\ q_{\min,n} & v_n(t) \geq v_{\max,n}, \\ -c_n(v_n(t) - v_{\min,n}) + q_{\max,n} & \text{otherwise,} \end{cases}$$

where $c_n = \frac{q_{\max,n} - q_{\min,n}}{v_{\max,n} - v_{\min,n}}$, and with the the optimized droop control design⁴ from [27]

$$q_n(t+1) = \rho_n(v_n(t))$$

$$\rho_n(v_n) := \begin{cases} q_{\max,n} & v_n \leq v_{\min,n}, \\ \frac{\bar{v}_{\min,n} - v_n}{\bar{v}_{\min,n} - v_{\min,n}} q_{\max,n} & v_{\min,n} < v_n < \bar{v}_{\min,n}, \\ 0 & \bar{v}_{\min,n} \leq v_n \leq \bar{v}_{\max,n}, \\ \frac{v_n - \bar{v}_{\max,n}}{\bar{v}_{\max,n} - \bar{v}_{\max,n}} q_{\min,n} & \bar{v}_{\max,n} < v_n < v_{\max,n}, \\ q_{\min,n} & v_n \geq v_{\max,n} \end{cases}$$

where $\bar{v}_{\min,n}$ and $\bar{v}_{\max,n}$ are parameters satisfying $v_{\min,n} < \bar{v}_{\min,n} \leq \bar{v}_{\max,n} < v_{\max,n}$ that are optimized given the dataset prescribing a day-ahead forecast.

A. The Learning Stage

We considered an optimization problem of the form (2) where the cost is set to

$$f(\mathbf{q}\mathbf{c}) = \alpha \underbrace{\|\mathbf{v}(\mathbf{q}\mathbf{c}) - \mathbf{1}\|}_{\textcircled{1}} + (1 - \alpha) \underbrace{(\mathbf{q}^\top \tilde{\mathbf{R}}\mathbf{q} + \mathbf{p}^\top \tilde{\mathbf{R}}\mathbf{p})}_{\textcircled{2}},$$

where $\textcircled{1}$ and $\textcircled{2}$ aims to minimize the voltage deviations and power losses [11], respectively, and the parameter α trades-off those two objectives. We assume the 5 DERs have uniform generation capabilities, precisely, $\mathbf{q}_{\max} = 0.4 \times \mathbf{1}$ MVAR and

⁴In [27], the voltage limits for maximum reactive power provision and absorption are selected as $v_{\min,n} = 0.9$ p.u. and $v_{\max,n} = 1.1$ p.u., respectively. To make sure the DERs reach maximum reactive power compensation capability when voltage exceeds the limits, as discussed in Remark IV.2, here they are set to $v_{\min,n} = 0.95$ p.u. and $v_{\max,n} = 1.05$

$\mathbf{q}_{\min} = -\mathbf{q}_{\max}$. The voltage limit vectors are set to $\mathbf{v}_{\max} = 1.05 \times \mathbf{1}$ p.u. and $\mathbf{v}_{\min} = 0.95 \times \mathbf{1}$ p.u.

The dataset for the learning process is built using the aforesaid power demands and generations obtained from the Pecan Street data, which are intended as day-ahead forecasts. We use the CVX toolbox [38] to solve the ORPF problem (2) with linearized power flow (5). Note however that, one can use any other power flow models to solve the ORPF problem. We add the pseudo data points to the obtained dataset as described in Remark IV.2 with $\underline{K} = \bar{K} = 700$, which results in a total of 2840 data points for each DER. We implement the neural network approach according to Proposition IV.1 using TensorFlow 2.7.0 and conduct the training process in Google Colab with a single TPU with 32 GB memory. The number of episodes and the number of neurons H are 2000 and 1000, respectively, and the neural networks are trained using the Adam optimizer [39] with the learning rate initialized at 0.01 and decays every 500 steps with a base of 0.5.

Fig. 5 plots the solutions to the ORPF problem (2) for all data profiles, the equilibrium function ϕ_{32} learned with and without pseudo points, the standard droop function ϱ_{32} , and the optimized droop function ρ_{32} for the DER at node 32 with $\alpha = \frac{1}{3}$. In contrast to the case in which no pseudo points are added in the learning process, the learned equilibrium function with pseudo points reaches maximum reactive power compensation capability when voltage exceeds the limits. We further summarize in Table I the average loss for the whole training dataset using the learned equilibrium functions and optimized/standard droop control functions, i.e.,

$$\frac{\sum_{k=1}^K \|\mathbf{q}_{\mathcal{C},k}^* - \square(\mathbf{v}_{\mathcal{C},k}^*)\|^2}{K \cdot C}, \quad (19)$$

where \square is ϕ for the data-based method, ρ for the optimized droop control, and ϱ for the standard droop control. The results illustrate the enhanced optimality of the learned equilibrium functions in approximating ORPF solutions compared to benchmarks.

TABLE I
AVERAGE LOSS VALUES FOR ALL DATA PROFILES

α	Learned equil. func.	Opt. droop func.	Std. droop func.
0	0.0090	0.0305	0.0438
1/3	0.0060	0.0171	0.0395
1/2	0.0071	0.0429	0.0662
2/3	0.0175	0.0724	0.0852
1	0.0377	0.0878	0.0965

We next illustrate the advantage of using incremental algorithm. Recall that to guarantee the convergence of the non-incremental algorithm, i.e., $\epsilon = 1$, one needs to further enforce an additional slope constraint (13) on the learned equilibrium functions, cf. Remark III.4. Fig. 6 shows that this additional slope constraint leads to larger approximation errors of the learned equilibrium functions in fitting the dataset (we do not consider the pseudo points during learning here for fairness), and thus degrades the optimality of system performance.

B. The Control Stage

We run the following simulations using the learned equilibrium functions for the case $\alpha = \frac{1}{3}$ with pseudo data points




Fig. 5. The orange solid and red dashed curves are respectively learned equilibrium functions with and without considering pseudo data points for the DER at node 32, while the blue dashdotted and green dotted curves are respectively optimized linear droop control function [27] and standard linear droop function [11], [15]. The comparison between orange solid and red dashed curves illustrates the role of pseudo data points in learning equilibrium functions. The former reaches the maximum reactive power compensation capability when voltage exceeds the limits, while the latter does not.




Fig. 6. Comparison of average training losses (MSE) for all DERs with and without the additional slope constraint (13) on the equilibrium functions as described in Remark III.4. The mean and standard deviations are evaluated based on 5 random seeds. This additional slope constraint leads to larger approximation error of the learned equilibrium functions in fitting the data points.

considered and assume that $\mathbf{q}_c(0) = \mathbf{0}$. Note that we use MATPOWER [40] to solve the power flow equation for all the tests below. We first verify the convergence properties of the proposed reactive power update rule (3) stated in Proposition III.2. Consider the scenario where load-generation profiles are fixed, Fig. 7 reports the evolution of the DERs' reactive power setpoints using load-generation profiles of the 695-th minute and consider 120 iterations of (3). For $\epsilon = 0.369$, the reactive power setpoint trajectories converge to their final values, cf. Fig. 7(a), while the case that $\epsilon = 1$ fails,

cf. Fig. 7(b). This is consistent with the sufficient condition

$$0 < \epsilon < \min \left\{ 1, \frac{2}{(\|\mathbf{X}\|L + 1)^2} \right\} = 0.3691$$

derived in Proposition III.2.

Next, we test the proposed data-based control method in a scenario where load-generation profiles are time-varying. Specifically, we obtain load-generation profiles by randomly perturbing (5%) the consumption data used to learn the equilibrium functions. This can be interpreted as having the data from the dataset prescribing a day-ahead forecast, whereas their random perturbation act as the true realization of the load-generation scenarios. These loads and generations are minute-based and we consider 120 iterations of (3) per minute with $\epsilon = 0.369$. Fig. 8 compares the evolution of maximum/minimum voltages under the proposed data-based control method, optimized droop control method, ORPF solutions, and the case where no control action is taken. One can observe that, in contrast to the proposed data-based method, the optimized droop control method induces instability issues during 12:00 and 16:00 causing voltages to oscillate. We further compare in Fig. 9 the distances between the actual reactive power setpoints and ORPF solutions, i.e., $\|\mathbf{q}_c - \mathbf{q}_c^*\|$. Compared to the benchmarks, the reactive power setpoints during evolution under the proposed data-based method remains much closer to the ORPF solutions, thus leading to significantly improved optimality. To further illustrate the effectiveness and advantages of the proposed data-based control method, Table II summarizes the comparison results of the proposed data-based control method against the standard and optimized linear droop control methods, as well as the case where no control action is taken for different values of α . We quantify the performance by the average of distances of actual and optimal reactive power setpoints, i.e., average of $\|\mathbf{q}_c(t) - \mathbf{q}_c^*(t)\|$ for the entire day. It can be observed that the proposed data-based control method outperforms the benchmark methods in all cases.

TABLE II
AVERAGE DISTANCES BETWEEN ACTUAL REACTIVE POWER SETPOINTS AND ORPF SOLUTIONS FOR ENTIRE DAY

α	0	1/3	1/2	2/3	1
Data-based	0.1185	0.0985	0.0784	0.1115	0.1645
Opt. Droop	0.2786	0.2474	0.3728	0.4410	0.4854
Std. Droop	0.2886	0.3311	0.4076	0.4699	0.5047
No Control	0.3160	0.5081	0.6842	0.8029	0.8358

C. Robustness to Voltage Measurement Noise

Here, we test the robustness of the proposed data-based method against voltage measurement noise. Similar to Section V-B, we use the learned equilibrium functions to run simulations, but additionally add random perturbations to the measurement of local voltages. Table III summarizes the distances between actual reactive power setpoints and ORPF solutions under different voltage measurement noise levels. Specifically, 0.5% perturbation to the voltage measurement corresponds to a common level of precision among smart meters in the U.S. [41], while the case of 1% perturbation represents the biggest error allowed in power systems. Table III




Fig. 7. Evolution of reactive power setpoints under the proposed reactive power update rule (3) with (a) $\epsilon = 0.369$ and (b) $\epsilon = 1$, where we use the power data profiles of the 695-th minute and consider 120 iterations. This verifies the sufficient condition $0 < \epsilon < \min\{1, \frac{2}{(\|\mathbf{x}\|_{L+1})^2}\} = 0.3691$ in Proposition III.2 to ensure global asymptotic stability.

TABLE III
AVERAGE DISTANCES UNDER MEASUREMENT NOISE

Noise \ α	0	1/3	1/2	2/3	1
0.2%	0.1210	0.1017	0.0936	0.1291	0.1739
0.5%	0.1317	0.1176	0.1378	0.1862	0.2237
1.0%	0.1571	0.1553	0.2127	0.2762	0.3129

indicates that, even in the latter case, the proposed data-based method still significantly outperforms the benchmark methods without measurement noise.

D. Discussion

Our simulation results above validate the improved performance of the proposed data-based method compared to the linear droop control method for different control goals. In fact, apart from considering the minimization of voltage deviations and power losses, our framework allows the users to consider any other type of cost functions, depending on

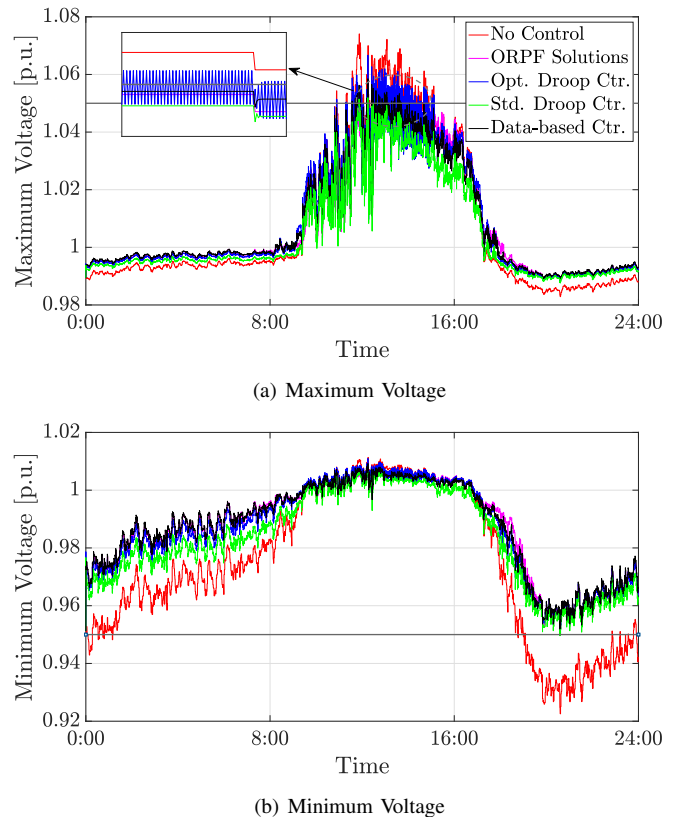


Fig. 8. Evolution of the maximum (top) and minimum (bottom) voltages of the IEEE 37-bus network under the proposed data-based, ORPF solutions, optimized linear droop, standard linear droop, and no control methods. For all minute-based data profiles, the ORPF problem is feasible and thus \mathbf{q}_C^* always exists. The optimized droop control induces voltage instability issues, causing voltages oscillations during 12:00 and 16:00, while the proposed data-based method guarantees the convergence of voltages for every minute-based data profile.

specific control goals, to learn purely local controllers that steer system operating points to approximated ORPF solutions. However, as one can observe in Table I, different cost functions may result in different optimality gaps between the proposed data-based method and the ORPF approach. Since the dataset we construct only maps the local voltage to the local optimal reactive power setpoint, it is possible that one fixed voltage corresponds to multiple optimal reactive power setpoints. On the other hand, it is also possible that the optimal solution pairs are not so close to the non-increasing shape as we require the equilibrium functions to be. We refer to these phenomena as *data inconsistency*. We note that different selections of cost function significantly influence the data inconsistency, and thus leads to very different optimality gaps. For example, as Table I suggests, the data becomes significantly more inconsistent when the minimization of voltage deviations takes a more important role in the cost function. As part of our follow-up work, we plan to include other available local information to alleviate the data inconsistency challenge, e.g., prevailing (re)active power injections as additional inputs of the equilibrium function. Another important observation is that, although the ORPF approach strictly guarantees that the voltages are within limits, our approach does not. For instance,

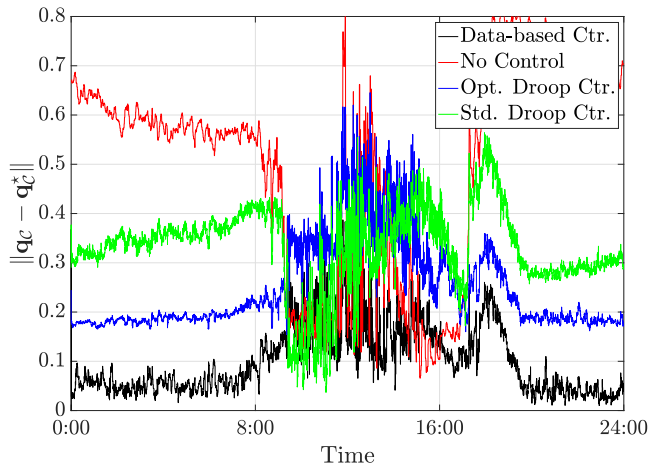


Fig. 9. Evolution of the distance between actual reactive power setpoints and ORPF solutions of the IEEE 37-bus network under the proposed data-based, optimized linear droop, standard linear droop, and no control methods.

in Fig. 8, the voltage nadir during evolution under the proposed data-based method slightly violates the voltage limits. The reason is that when α is relatively small, many of the optimal solutions given by the ORPF problem lie on the boundary of the voltages limits. Since the local surrogates only provide approximations of the optimal solutions, the actual converged voltages can easily go out of limits in such situations. On the other hand, as pointed out in [18], purely local control strategies generally have no guarantee on desired regulation, in the sense that the equilibrium $\mathbf{q}_C^\#$ of (8) could result in a $\mathbf{v}(\mathbf{q}_C^\#) \notin [\mathbf{v}_{\min}, \mathbf{v}_{\max}]$, even if there indeed exists \mathbf{q}_C such that $\mathbf{v}(\mathbf{q}_C) \in [\mathbf{v}_{\min}, \mathbf{v}_{\max}]$.

VI. CONCLUSIONS

We have presented a data-driven framework to design local Volt/Var controllers capable of steering a power distribution network towards efficient network configurations. Building on the idea of learning local surrogates that map local voltages to reactive power setpoints that approximate the ORPF solution, we have proposed a local control update scheme and identified conditions on surrogates and control parameters so that the reactive power point converges in a global asymptotic sense. By constructing a labeled dataset of ORPF solutions with different load and generation profiles, we have trained neural networks whose resulting parameterized functions meet the conditions on surrogates by design to fit the dataset. We have shown in AC power flow simulation tests that the proposed framework guarantees the voltage stability and significantly reduces operation cost compared to prevalent local control approaches. Future research directions include considering the regulation of legacy devices, enhancing data consistency by making use of other local information in building the dataset, reducing the optimality gap during the learning process, and extending the proposed framework to more general scenario where we take advantage of communication among neighboring agents.

APPENDIX A TECHNICAL LEMMAS

Lemma A.1. (*Bauer and Fike Theorem [42, Corollary 6.3.4]*): Let $\mathbf{A}, \mathbf{E} \in \mathbb{R}^{n \times n}$ with \mathbf{A} normal. If $\hat{\lambda}$ is an eigenvalue of $\mathbf{A} + \mathbf{E}$, then there exists an eigenvalue λ of \mathbf{A} such that $|\hat{\lambda} - \lambda| \leq \|\mathbf{E}\|$.

Lemma A.2. (*Positive semidefiniteness of $\mathbf{X}\mathbf{M}$ and upper bound of $\|\mathbf{X}\mathbf{M}\|$*): The matrix $\mathbf{X}\mathbf{M}$ is positive semidefinite. Moreover, it holds

$$\|\mathbf{X}\mathbf{M}\| \leq \|\mathbf{X}\|L.$$

Proof. Let $(\lambda_i, \boldsymbol{\xi}_i)$ be a left eigenpair for $\mathbf{X}\mathbf{M}$. Then, $(\lambda_i, \boldsymbol{\xi}_i \mathbf{X}^{\frac{1}{2}})$ is a left eigenpair for the symmetric matrix $\mathbf{X}^{\frac{1}{2}} \mathbf{M} \mathbf{X}^{\frac{1}{2}} \succeq 0$. Indeed,

$$\boldsymbol{\xi}_i \mathbf{X}^{\frac{1}{2}} \mathbf{X}^{\frac{1}{2}} \mathbf{M} \mathbf{X}^{\frac{1}{2}} = \boldsymbol{\xi}_i \mathbf{X} \mathbf{M} \mathbf{X}^{\frac{1}{2}} = \lambda_i \boldsymbol{\xi}_i \mathbf{X}^{\frac{1}{2}}.$$

Therefore, $\mathbf{X}\mathbf{M}$ is positive semidefinite as well. Also,

$$\|\mathbf{X}\mathbf{M}\| \leq \|\mathbf{X}\| \|\mathbf{M}\| \leq \|\mathbf{X}\|L,$$

which completes the proof. \square

APPENDIX B PROOFS OF PROPOSITIONS

Proof of Proposition III.1. We first show by induction the feasibility of the reactive power update (3). The initial power injection $q_n(0)$ belongs to \mathcal{Q}_n by hypothesis. Assume now that $q_n(t) \in \mathcal{Q}_n$ and C3) holds. Then $q_n(t+1)$ is the convex combination of two elements of \mathcal{Q}_n . We next show that (8) has a unique equilibrium point. From (9), the equilibrium exists if $\mathbf{q}_C = \mathbf{h}(\mathbf{q}_C)$ has a solution, where $\mathbf{h} : \mathcal{Q} \rightarrow \mathcal{Q}$ is a continuous vector function with $\mathbf{h}(\mathbf{q}_C) = \phi(\mathbf{X}\mathbf{q}_C + \hat{\mathbf{v}}_C)$. Since \mathcal{Q} is convex and compact, according to Brouwer's Fixed Point Theorem [43, Corollary 6.6], such solution exists. Finally, to show uniqueness, we reason by contradiction. Assume both $(\mathbf{q}_C^\#, \mathbf{v}_C^\#)$ and $(\mathbf{q}_C^\#, \mathbf{v}_C^\#)$ are equilibrium points for (8) with $\mathbf{q}_C^\# \neq \mathbf{q}_C^\#$. From (9a),

$$\mathbf{q}_C^\# - \mathbf{q}_C^\# = \phi(\mathbf{v}_C^\#) - \phi(\mathbf{v}_C^\#) = \mathbf{D}(\mathbf{v}_C^\# - \mathbf{v}_C^\#), \quad (20)$$

where $\mathbf{D} \in \mathbb{R}^{C \times C}$ is a diagonal matrix with

$$D_n = \begin{cases} \frac{\phi_n(v_n^\#) - \phi_n(v_n^\#)}{v_n^\# - v_n^\#} & v_n^\# \neq v_n^\#, \\ 0 & v_n^\# = v_n^\#. \end{cases}$$

From C2), ϕ_n is non-increasing in v_n for all $n \in \mathcal{C}$. Hence, $D_n \leq 0, \forall n \in \mathcal{C}$, and $\mathbf{D} \preceq 0$. On the other hand, (9b) yields

$$\mathbf{q}_C^\# - \mathbf{q}_C^\# = \mathbf{X}^{-1}(\mathbf{v}_C^\# - \mathbf{v}_C^\#).$$

Then, it follows that

$$(\mathbf{X}^{-1} - \mathbf{D})(\mathbf{v}_C^\# - \mathbf{v}_C^\#) = \mathbf{0}.$$

Since $\mathbf{X} \succ 0$, it holds that $\mathbf{X}^{-1} \succ 0$, $\mathbf{X}^{-1} - \mathbf{D} \succ 0$. As a consequence, $\mathbf{v}_C^\# - \mathbf{v}_C^\# = \mathbf{0}$ and $\mathbf{q}_C^\# - \mathbf{q}_C^\# = \mathbf{0}$, cf. (9a), which is a contradiction. This completes the proof. \square

Proof of Proposition III.2. Consider the voltage evolution under (8),

$$\mathbf{v}_C(t+1) = \mathbf{X}\mathbf{q}_C(t+1) + \hat{\mathbf{v}}_C$$

$$\begin{aligned}
&= (1 - \epsilon)\mathbf{X}\mathbf{q}_C(t) + \epsilon\mathbf{X}\phi(\mathbf{v}_C(t)) + (1 - \epsilon)\hat{\mathbf{v}}_C + \epsilon\hat{\mathbf{v}}_C \\
&= (1 - \epsilon)\mathbf{v}_C(t) + \epsilon(\mathbf{X}\phi(\mathbf{v}_C(t)) + \hat{\mathbf{v}}_C) := \mathbf{g}(\mathbf{v}_C(t)).
\end{aligned}$$

We show that, for small enough values of ϵ , the operator $\mathbf{g} : \mathbb{R}^C \rightarrow \mathbb{R}^C$ is a contraction, i.e.,

$$\frac{\|\mathbf{g}(\mathbf{v}_C) - \mathbf{g}(\mathbf{v}'_C)\|}{\|\mathbf{v}_C - \mathbf{v}'_C\|} < 1, \quad (21)$$

for any $\mathbf{v}_C, \mathbf{v}'_C \in \mathbb{R}^C$. Indeed, define a diagonal matrix $\mathbf{M} \in \mathbb{R}^{C \times C}$ with the n -th diagonal entry being

$$M_n = \begin{cases} \frac{|\phi_n(v_n) - \phi_n(v'_n)|}{|v_n - v'_n|} & v_n \neq v'_n, \\ 0 & v_n = v'_n. \end{cases}$$

Then, it follows that

$$\begin{aligned}
&\|\mathbf{g}(\mathbf{v}_C) - \mathbf{g}(\mathbf{v}'_C)\| \\
&= \|(1 - \epsilon)(\mathbf{v}_C - \mathbf{v}'_C) + \epsilon\mathbf{X}(\phi(\mathbf{v}_C) - \phi(\mathbf{v}'_C))\| \\
&= \|(1 - \epsilon)\text{sign}(\mathbf{v}_C - \mathbf{v}'_C)|\mathbf{v}_C - \mathbf{v}'_C| \\
&\quad - \epsilon\mathbf{X}\text{sign}(\mathbf{v}_C - \mathbf{v}'_C)|\phi(\mathbf{v}_C) - \phi(\mathbf{v}'_C)|\| \\
&= \|(1 - \epsilon)|\mathbf{v}_C - \mathbf{v}'_C| - \epsilon\mathbf{X}|\phi(\mathbf{v}_C) - \phi(\mathbf{v}'_C)|\| \\
&\leq \|(1 - \epsilon)\mathbf{I} - \epsilon\mathbf{X}\mathbf{M}\| \|\mathbf{v}_C - \mathbf{v}'_C\|,
\end{aligned}$$

where we have used in the second equality the fact that ϕ_n is non-increasing in v_n for each $n \in \mathcal{C}$, and thus $\text{sign}(\phi(\mathbf{v}_C) - \phi(\mathbf{v}'_C)) = -\text{sign}(\mathbf{v}_C - \mathbf{v}'_C)$.

To prove (21), it is sufficient to show that there always exists ϵ such that $\|(1 - \epsilon)\mathbf{I} - \epsilon\mathbf{X}\mathbf{M}\| < 1$, which is equivalent to proving that $\lambda_{\max}(\mathbf{\Gamma}) < 1$, where $\mathbf{\Gamma} \succeq 0$ is

$$\begin{aligned}
\mathbf{\Gamma} &\triangleq [(1 - \epsilon)\mathbf{I} - \epsilon\mathbf{X}\mathbf{M}]^\top [(1 - \epsilon)\mathbf{I} - \epsilon\mathbf{X}\mathbf{M}] \\
&= (1 - \epsilon)^2\mathbf{I} - \epsilon(1 - \epsilon)(\mathbf{M}\mathbf{X}^\top + \mathbf{X}\mathbf{M}) + \epsilon^2\mathbf{M}\mathbf{X}^\top\mathbf{X}\mathbf{M}.
\end{aligned}$$

Rewrite $\mathbf{\Gamma}$ as

$$\begin{aligned}
\mathbf{\Gamma} &= \underbrace{(1 - 2\epsilon)\mathbf{I} - \epsilon(\mathbf{M}\mathbf{X}^\top + \mathbf{X}\mathbf{M})}_{:=\mathbf{A}} \\
&\quad + \epsilon^2 \underbrace{(\mathbf{I} + \mathbf{X}\mathbf{M} + \mathbf{M}\mathbf{X}^\top + \mathbf{M}\mathbf{X}^\top\mathbf{X}\mathbf{M})}_{:=\mathbf{E}}.
\end{aligned}$$

Note that \mathbf{A} is symmetric. According to Lemma A.1, provided in the Appendix, it holds that $0 \leq \lambda_{\max}(\mathbf{\Gamma}) \leq \lambda_{\max}(\mathbf{A}) + \epsilon^2\|\mathbf{E}\|$. Now it remains to show that there always exists ϵ such that

$$\lambda_{\max}(\mathbf{A}) + \epsilon^2\|\mathbf{E}\| < 1.$$

Let $\underline{\lambda} = \lambda_{\min}(\mathbf{M}\mathbf{X}^\top + \mathbf{X}\mathbf{M})$. According to Lemma A.2, provided in the Appendix, $\mathbf{M}\mathbf{X}^\top, \mathbf{X}\mathbf{M} \succeq 0$, and therefore $\underline{\lambda} \geq 0$. The above condition is then equivalent to

$$1 - \epsilon(2 + \underline{\lambda}) + \epsilon^2\|\mathbf{E}\| < 1,$$

which yields

$$0 < \epsilon < \frac{(2 + \underline{\lambda})}{\|\mathbf{E}\|}.$$

Finally, according to Lemma A.2, it holds

$$\begin{aligned}
\|\mathbf{E}\| &\leq 1 + 2\|\mathbf{X}\mathbf{M}\| + \|\mathbf{X}\mathbf{M}\|^2 \\
&\leq 1 + 2\|\mathbf{X}\|L + \|\mathbf{X}\|^2L^2 = (\|\mathbf{X}\|L + 1)^2,
\end{aligned}$$

and hence

$$\frac{2 + \underline{\lambda}}{\|\mathbf{E}\|} \geq \frac{2}{\|\mathbf{E}\|} \geq \frac{2}{(\|\mathbf{X}\|L + 1)^2}.$$

Considering that we require $\epsilon \in [0, 1]$, (10) follows, concluding the proof. \square

Proof of Proposition IV.1. To show the sufficiency of (17), notice that, for any $x \in \mathbb{R}$, if $x < b_1$, then $N(x) = \beta$; and for $x \geq b_1$, $N(x)$ is divided into H segments with the slope of the J -th segment being $\sum_{j=1}^J w_j$ for $J \in \{1, 2, \dots, H\}$. Given (17), it follows that N is non-increasing. To show the necessity of (17), suppose N is non-increasing. Then, there exists a segment of the equilibrium function, say the J -th, $J \in \{1, 2, \dots, H\}$ which is increasing. Hence, $\sum_{j=1}^J w_j > 0$ which is a contradiction.

We next show the universal approximation property. Given a compact domain \mathcal{X} of the form $\underline{x} \leq x \leq \bar{x}$, consider an equispaced partition of \mathcal{X} into H intervals, with the length of each interval being $s = \frac{\bar{x} - \underline{x}}{H}$. Consider the function

$$N(x) = \sum_{h=1}^H w_h \text{ReLU}(x - b_h) + \beta,$$

whose parameters are defined as: $\beta = g(\underline{x})$, $b_h = \underline{x} + (h - 1)s$, $w_1 = \frac{g(\underline{x} + s) - g(\underline{x})}{s}$, $w_2 = \frac{g(\underline{x} + 2s) - g(\underline{x} + s)}{s} - w_1, \dots, w_h = \frac{g(\underline{x} + hs) - g(\underline{x} + (h-1)s)}{s} - \sum_{k=1}^{h-1} w_k$. Note that it holds that

$$\sum_{k=1}^h w_k = \frac{g(\underline{x} + hs) - g(\underline{x} + (h-1)s)}{s} \leq 0$$

because $g(x)$ is nonincreasing, for $h \in \{1, 2, \dots, H\}$. Hence, inequality (17) is satisfied. It can be observed that

$$|N(x) - g(x)| \leq L_g s$$

where L_g is the Lipschitz constant of g . Hence, by setting $H \geq \frac{L_g(\bar{x} - \underline{x})}{\eta} = O(1/\eta)$, it holds that $|N(x) - g(x)| \leq L_g s \leq \eta$ for all $x \in \mathcal{X}$. This completes the proof. \square

REFERENCES

- [1] S. H. Low, "Convex relaxation of optimal power flow - Part I: Formulations and equivalence," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 15–27, 2014.
- [2] B. Cui and X. A. Sun, "A new voltage stability-constrained optimal power-flow model: Sufficient condition, SOCP representation, and relaxation," *IEEE Transactions on Power Systems*, vol. 33, no. 5, pp. 5092–5102, 2018.
- [3] L. Gan, N. Li, U. Topcu, and S. H. Low, "Optimal power flow in tree networks," in *IEEE Conf. on Decision and Control*, (Firenze, Italy), pp. 2313–2318, Dec. 2013.
- [4] X. Pan, T. Zhao, M. Chen, and S. Zhang, "DeepOPF: A deep neural network approach for security-constrained DC optimal power flow," *IEEE Transactions on Power Systems*, vol. 36, no. 3, pp. 1725–1735, 2021.
- [5] A. S. Zamzam and K. Baker, "Learning optimal solutions for extremely fast AC optimal power flow," in *IEEE Int. Conf. on Communications, Control, and Computing Technologies for Smart Grids*, (Tempe, AZ, USA), Nov. 2020.
- [6] D. Owerko, F. Gama, and A. Ribeiro, "Optimal power flow using graph neural networks," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, (Barcelona, Spain), pp. 5930–5934, 2020.
- [7] M. K. Singh, S. Gupta, V. Kekatos, G. Cavraro, and A. Bernstein, "Learning to optimize power distribution grids using sensitivity-informed deep neural networks," in *IEEE Int. Conf. on Communications, Control, and Computing Technologies for Smart Grids*, (Tempe, AZ, USA), Nov. 2020.

- [8] M. K. Singh, V. Kekatos, and G. B. Giannakis, "Learning to solve the AC-OPF using sensitivity-informed deep neural networks," *IEEE Transactions on Power Systems*, vol. 37, no. 4, pp. 2833–2846, 2022.
- [9] F. Fioretto, T. W. K. Mak, and P. Van Hentenryck, "Predicting AC optimal power flows: Combining deep learning and lagrangian dual methods," in *AAAI Conference on Artificial Intelligence*, (New York, USA), pp. 630–637, Feb. 2020.
- [10] E. Dall'Anese, H. Zhu, and G. B. Giannakis, "Distributed optimal power flow for smart microgrids," *IEEE Transactions on Smart Grid*, vol. 4, no. 3, pp. 1464–1475, 2013.
- [11] G. Cavraro and R. Carli, "Local and distributed voltage control algorithms in distribution networks," *IEEE Transactions on Power Systems*, vol. 33, no. 2, pp. 1420–1430, 2017.
- [12] E. Dall'Anese and A. Simonetto, "Optimal power flow pursuit," *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 942–952, 2018.
- [13] G. Qu and N. Li, "Optimal distributed feedback voltage control under limited reactive power," *IEEE Transactions on Power Systems*, vol. 35, no. 1, pp. 315–331, 2020.
- [14] "IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces," *IEEE Std 1547-2018 (Revision of IEEE Std 1547-2003)*, pp. 1–138, 2018.
- [15] K. Turitsyn, P. Sulc, S. Backhaus, and M. Chertkov, "Options for control of reactive power by distributed photovoltaic generators," *Proceedings of the IEEE*, vol. 99, no. 6, pp. 1063–1073, 2011.
- [16] H. Zhu and H. J. Liu, "Fast local voltage control under limited reactive power: Optimality and stability analysis," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3794–3803, 2015.
- [17] X. Zhou, M. Farivar, Z. Liu, L. Chen, and S. H. Low, "Reverse and forward engineering of local voltage control in distribution networks," *IEEE Transactions on Automatic Control*, vol. 66, no. 3, pp. 1116–1128, 2021.
- [18] S. Bolognani, R. Carli, G. Cavraro, and S. Zampieri, "On the need for communication for voltage regulation of power distribution grids," *IEEE Transactions on Control of Network Systems*, vol. 6, no. 3, pp. 1111–1123, 2019.
- [19] C. Dawson, S. Gao, and C. Fan, "Safe control with learned certificates: A survey of neural Lyapunov, barrier, and contraction methods," *arXiv preprint arXiv:2202.11762*, 2022.
- [20] W. Cui, Y. Jiang, and B. Zhang, "Reinforcement learning for optimal primary frequency control: A Lyapunov approach," *IEEE Transactions on Power Systems*, 2022. To appear.
- [21] Z. Yuan, C. Zhao, and J. Cortés, "Reinforcement learning for distributed transient frequency control with stability and safety guarantees," *Systems & Control Letters*, 2022. Submitted.
- [22] W. Cui, J. Li, and B. Zhang, "Decentralized safe reinforcement learning for inverter-based voltage control," *Electric Power Systems Research*, vol. 211, p. 108609, 2022.
- [23] Y. Shi, G. Qu, S. H. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for real-time voltage control," in *American Control Conference*, (Atlanta, GA), pp. 2715–2721, June 2022.
- [24] S. Karagiannopoulos, P. Aristidou, and G. Hug, "Data-driven local control design for active distribution grids using off-line optimal power flow and machine learning techniques," *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6461–6471, 2019.
- [25] P. S. Torre and P. Hidalgo-Gonzalez, "Decentralized optimal power flow for time-varying network topologies using machine learning," *Electric Power Systems Research*, vol. 212, p. 108575, 2022.
- [26] X. Sun, J. Qiu, and J. Zhao, "Optimal local Volt/Var control for photovoltaic inverters in active distribution networks," *IEEE Transactions on Power Systems*, vol. 36, no. 6, pp. 5756–5766, 2021.
- [27] H. Ji, C. Wang, P. Li, J. Zhao, G. Song, F. Ding, and J. Wu, "A centralized-based method to determine the local voltage control strategies of distributed generator operation in active distribution networks," *Applied Energy*, vol. 228, pp. 2024–2036, 2018.
- [28] G. Cavraro, Z. Yuan, M. K. Singh, and J. Cortés, "Learning local Volt/Var controllers towards efficient network operation with stability guarantees," in *IEEE Conf. on Decision and Control*, (Cancun, Mexico), Dec. 2022. To appear.
- [29] S. Gupta, S. Chatzivasileiadis, and V. Kekatos, "Deep learning for optimal Volt/VAR control using distributed energy resources," *arXiv preprint arXiv:2211.09557*, 2022.
- [30] J. Feng, Y. Shi, G. Qu, S. H. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for real-time voltage control in distribution systems," *arXiv preprint arXiv:2209.07669*, 2022.
- [31] A. M. Kettner and M. Paolone, "On the properties of the compound nodal admittance matrix of polyphase power systems," *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 444–453, 2019.
- [32] M. Farivar, X. Zhou, and L. Chen, "Local voltage control in distribution systems: An incremental control algorithm," in *IEEE Int. Conf. on Communications, Control, and Computing Technologies for Smart Grids*, (Miami, FL, USA), pp. 732–737, Nov. 2015.
- [33] H. Daniels and M. Velikova, "Monotone and partially monotone neural networks," *IEEE Transactions on Neural Networks*, vol. 21, no. 6, pp. 906–917, 2010.
- [34] A. Wehenkel and G. Louppe, "Unconstrained monotonic neural networks," in *Conference on Neural Information Processing Systems*, vol. 32, (Vancouver, Canada), pp. 1543–1553, Dec. 2019.
- [35] X. Liu, X. Han, N. Zhang, and Q. Liu, "Certified monotonic neural networks," in *Conference on Neural Information Processing Systems*, vol. 33, (Vancouver, Canada), pp. 15427–15438, Dec. 2020.
- [36] I. Safran and O. Shamir, "Depth-width tradeoffs in approximating natural functions with neural networks," in *International Conference on Machine Learning*, (Sydney, Australia), pp. 2979–2987, Aug. 2017.
- [37] "Dataport: The world's largest energy data resource," *Pecan Street Inc*, 2015. Available at <https://dataport.pecanstreet.org/>.
- [38] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," Mar. 2014. Available at <http://cvxr.com/cvx>.
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations*, (San Diego, CA, USA), May 2015.
- [40] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, "Matpower: Steady-state operations, planning and analysis tools for power systems research education," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 12–19, 2011.
- [41] EEI-AEIC-UTC, "Smart meters and smart meter systems: A metering industry perspective," *Washington, DC, USA, Edison Elect. Inst., White Paper*, 2011.
- [42] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 2012.
- [43] K. C. Border, *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge, UK: Cambridge University Press, 1985.