Chapter 1

Learning Stable Local Volt/Var Controllers in Distribution Grids

Zhenyi Yuan[a], Guido Cavraro[b], and Jorge Cortés[a]
[a] *University of California, San Diego, La Jolla, CA 92093, USA* , [b] *National Renewable Energy Laboratory, Golden, CO 80401, USA*

ABSTRACT
This chapter describes a framework to synthesize provably stable local Volt/Var controllers for distributed energy resources (DERs) in power distribution grids (DGs). The goal is to control the reactive power injections of DERs to improve the system performance as quantified by a generic optimal reactive power flow (ORPF) problem. To achieve this, we jointly design for each DER the control function, which prescribes the reactive power update rule, and the equilibrium function, which approximates the ORPF solutions from local measurements of voltages and powers. We provide conditions on the equilibrium functions and the control parameters ensuring the stability of the closed-loop system. In particular, we discuss the trade-offs between each set of conditions accounting for practical considerations, like fully exploiting the DERs' generation capabilities and reducing the optimality gap. These conditions are then translated into learning constraints on the neural networks' parameters that are enforced in the training phase. We validate our framework with numerical simulations on the IEEE 37-bus network and through a comparison with an optimized version of standard piece-wise linear control rules.

KEYWORDS
Distributed energy resources, Local Volt/Var control, Distribution grids, Machine learning, Closed-loop asymptotic stability

1.1 INTRODUCTION

Environmental and economic factors are promoting the deployment of DERs on DGs, leading to opportunities for improving power system performance and reducing greenhouse gas emissions. However, DERs' uncoordinated power injections pose challenges to system stability and operation, and may result in large voltage variations and frequency deviations. Exploiting the flexibility of their power electronic interface, DERs have the capability to provide ancillary services and in particular to perform Volt/Var control, i.e., to regulate the reactive power outputs to provide voltage regulation. This chapter provide an overview of our efforts [1–3] to realize this potential by synthesizing a framework that does not necessitate access to global information or tight requirements on communi-

cation among DERs. We combine smart decision making at the local level with machine learning techniques to propose local Volt/Var controllers incorporating optimality considerations and rigorous performance guarantees.

Literature Review

Volt/Var strategies aim to keep voltages within safe preassigned limits by having generators inject reactive power. The literature provides a variety of options for voltage regulation. In classic approaches, generators are controlled in an *open loop* fashion, precisely, reactive power outputs are computed by the system operator as solutions of *optimal power flow* (OPF) problems and then dispatched to the generators. OPF problems can be tackled using efficient solvers, e.g., [4–6]. The high penetration of DERs and the increased variability of DGs call for solving several OPF problems within a limited time frame, which poses significant challenges on online computational capabilities and the communication infrastructure. As a way of bypassing them, the introduction of learning techniques to obtain fast (approximate) OPF problems solutions has been motivated, e.g. [7,8], to mention a few. A graph neural network leveraging the connectivity of the power system is trained to infer AC-OPF solutions in [9]. In [10], a DNN is trained to fit not only OPF minimizers, but also their sensitivities with respect to the problem inputs. Once trained, the inference time for these approaches when presented with a new input is minimal. Nevertheless, in these approaches, (approximate) solutions of an OPF problem can be computed in general only if information from *all* the grid buses is available. Specifically, all the load demands and the generators' generation limits must be known. This requirement is most of the time prohibitive for actual DGs. Indeed, 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.

Closed-loop strategies instead use measurements to compensate missing grid information. Since future DGs will host a massive number of controllable devices, *decentralized* approaches look more suitable for practical applications. Decentralized strategies can be grouped into two main categories. The first is the class of *distributed* algorithms, where agents cooperate and exchange information with peers. Distributed algorithms can be designed to exactly solve a given OPF problem, see, e.g., [11]. This kind of schemes are often referred to as feedback-based optimization controllers and steer the network toward the OPF solutions by mean of the cyclical alternation of sensing, communication, and actuation [12–14]. Nevertheless, distributed strategies require a reliable real-time communication network satisfying strict requirements, e.g., every agent has to be able to communicate with its neighbors in the power grid [12], which may be difficult to satisfied in practice for DGs. The second is the class of *local* algorithms, where each agent makes decisions 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 [15], a number of options for local voltage regulation have been proposed [12,16–18]. Exploiting only limited information, local schemes feature intrinsic performance limitations due to their lack of coordination, e.g., they might fail to regulate voltages even if the overall generation resources are enough [19].

To enhance the performance of local schemes and to reduce the gap with distributed and/or optimal controllers, a recent research direction aims to design data-driven and learning-based control rules leveraging data from the grid. For instance, utilities have available an enormous amount of historical consumption and generation data retrieved from smart meters deployed in DGs. The work [20] proposes to learn the local controller by taking as an input both voltages and active power setpoints. The work [21] design piecewise linear control functions once set the number of break points, whereas in [22] a framework for tuning the parameters of standard piecewise linear local voltage regulators is proposed. The aforesaid works provide interesting insights on learning Volt/Var rules, but they do not assess the system stability and thus are not straightforwardly applicable in actual grids. For this reason, characterizing the system stability has become a priority. In the context of frequency control, [23,24] seeks to explicitly engineer the neural network structure to integrate the stability requirement. In the field of Volt/Var control field, reinforcement learning-based control schemes ensuring the system stability are proposed in [25] and [26]. Also, [27,28] optimize the local Volt/Var control schemes based on system forecasts and guaranteeing the grid stability.

What is Covered in this Chapter

This chapter describes a Volt/Var control framework that tackles the challenges described above. We provide an overview of our results in [1–3] combining control and machine learning techniques to design provably stable local Volt/Var controllers that steer the system toward efficient configurations. In our approach, efficient configurations correspond to equilibrium points that approximate the solutions of *optimal reactive power flow* (ORPF) problem, i.e., particular instances of the OPF problem in which the goal is to optimize the generator's reactive power injections. The control framework is composed of two main elements: the *equilibrium functions* and the *control functions*. The equilibrium function describes possible system equilibrium points and map local measurements into approximation of ORPF solutions. Equilibrium functions are learned from historical data. The control function represents an incremental algorithm in which the reactive power is updated as a convex combination of the previous power setpoint and the actual value of the equilibrium function.

The properties of the equilibrium function play a key role in determining the stability of the dynamics described by the control function. In this chapter, we discuss how to design equilibrium functions to improve the performance of

the local control schemes and guarantee the system stability. We look at two types of equilibrium functions depending on their arguments. First, we consider equilibrium functions that take as an input voltage measurements. Second, with the aim of reducing the optimality gap with respect to the centralized solution, we introduce equilibrium functions that depend not only on voltage magnitudes but also on reactive power injections, and can be written in a separable way. We provide various conditions ensuring the asymptotic stability of the control scheme. These conditions take the form of slope constraints on the equilibrium functions. One set of conditions requires coupling slope constraints on the functions of voltages and functions of reactive powers. Another set of conditions requires only slope constraints on the functions of reactive powers and decreasing functions of voltages. Given these conditions ensuring the system stability, our framework provides a way to train the equilibrium functions that meet them using historical data offline. The online implementation of the local control schemes is then guaranteed to steer the network towards the desired configuration described by the equilibrium functions, leading to significantly optimized performance.

Notation

The set of real and complex numbers are represented as $\mathbb{R}$ and $\mathbb{C}$, respectively. Matrices and column vectors are denoted with upper and lowercase boldface letters; calligraphic symbols are reserved for sets. Given a vector $\mathbf{a}$ (a diagonal matrix $\mathbf{A}$), its n-th (diagonal) entry is denoted by a_n (A_n). 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 $\mathfrak{R}(\cdot)$ and $\mathfrak{I}(\cdot)$ 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 matrix $\mathbf{A}$ with real eigenvalues, $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$ represent its largest and smallest eigenvalue, respectively. For any matrix $\mathbf{B}$, it holds that $\|\mathbf{B}\| = \sqrt{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}$.

1.2 GRID MODELING AND PROBLEM FORMULATION

Here, we describe the mathematical model of distribution grids and formulate the problem of interest.

1.2.1 Grid Modeling

A radial single-phase (or a balanced three-phase) DG having $N+1$ buses can be modeled by an undirected tree graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ rooted at the substation. Each node in the set $\mathcal{N} = \{0, 1, \ldots, N\}$ is associated with an electric bus; each edge in the set $\mathcal{E}$ is associated with a power line. The substation, labeled as node 0, is modeled as an ideal voltage source imposing the nominal voltage of 1 p.u. The

power grid state is fully described by the following variables:

- $u_n \in \mathbb{C}$ is the voltage phasor at bus $n \in \mathcal{N}$;
- $v_n := |u_n| \in \mathbb{R}$ is the voltage magnitude at bus $n \in \mathcal{N}$;
- $i_n \in \mathbb{C}$ is the injected current phasor at bus $n \in \mathcal{N}$;
- $s_n = p_n + iq_n \in \mathbb{C}$ is the nodal complex power injection at bus $n \in \mathcal{N}$; $p_n, q_n \in \mathbb{R}$ denote the active and the reactive power injections, respectively. We assume that if powers are injected into (absorbed from), then they have positive (negative) values, i.e., $p_n, q_n \geq 0$ ($p_n, q_n \leq 0$).

The vectors $\mathbf{u}, \mathbf{i}, \mathbf{s} \in \mathbb{C}^N$ collect the complex voltages, currents, and complex powers of buses $1, 2, \ldots, N$, and vectors $\mathbf{v}, \mathbf{p}, \mathbf{q} \in \mathbb{R}^N$ collect their voltage magnitudes, and active and reactive power injections. Let $z_e \in \mathbb{C}$ and $y_e = z_e^{-1} \in \mathbb{C}$ be the impedance and admittance of line $e = (m, n) \in \mathcal{E}$, respectively. The network bus admittance matrix $\mathbf{Y} \in \mathbb{C}^{(N+1)\times(N+1)}$ can be expressed as $\mathbf{Y} = \mathbf{Y}_L + \mathrm{diag}(\mathbf{y}_T)$, 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. $\mathbf{Y}$ is a symmetric matrix, whereas $\mathbf{Y}_L$ is a complex Laplacian matrix satisfying $\mathbf{Y}_L \mathbf{1} = \mathbf{0}$. We partition the bus admittance matrix to obtain

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

where $y_0 \in \mathbb{C}$, $\mathbf{y}_0 \in \mathbb{C}^N$, and $\tilde{\mathbf{Y}} \in \mathbb{C}^{N\times N}$. It can be proved that $\tilde{\mathbf{Y}}$ is invertible when the network is connected [29]. By defining $\tilde{\mathbf{Z}} := \tilde{\mathbf{Y}}^{-1}$ and $\hat{\mathbf{u}} := \tilde{\mathbf{Z}}\mathbf{y}_0$, the *power flow equations* read

$$\mathbf{u} = \tilde{\mathbf{Z}}\mathbf{i} + \hat{\mathbf{u}}, \tag{1.1a}$$

$$u_0 = 1, \tag{1.1b}$$

$$i_0 = \mathbf{1}^\top \mathbf{i}, \tag{1.1c}$$

$$u_n \bar{i}_n = p_n + jq_n, \tag{1.1d}$$

where $\bar{i}_n$ is the complex conjugate of i_n. The Kirchoff laws are captured by equation (1.1a); equations (1.1b) and (1.1c) come from the fact that the substation is the slack bus. Finally, equation (1.1d) provides the relation among voltages, currents, and powers. Even though voltages and powers are related by the nonlinear power flow equations (1.1), using a first-order Taylor expansion and defining $\tilde{\mathbf{R}} := \Re(\tilde{\mathbf{Z}})$ and $\tilde{\mathbf{X}} := \Im(\tilde{\mathbf{Z}}) \in \mathbb{R}^{N\times N}$, we can approximate [12] the power flow equations as

$$\mathbf{v} = \tilde{\mathbf{R}}\mathbf{p} + \tilde{\mathbf{X}}\mathbf{q} + \mathbf{1}. \tag{1.2}$$

Indeed, it is common to rely on linearizations of the power flow equations to study the control rule stability properties, e.g., see [13,14,17].

In this chapter, we assume that buses belonging to the subset $C \subseteq \mathcal{N}$, with $|C| = C$, are equipped with DERs. Every bus hosting a DER corresponds to a smart agent provided with some computational and sensing capabilities, i.e., it can measure its voltage magnitude. We will refer to buses in C as generators and to the remaining nodes in the set $\mathcal{L} = \mathcal{N} \setminus C$ as loads. For convenience, we partition reactive powers and voltage magnitudes by grouping together the nodes belonging to the load and generation sets,

$$\mathbf{q} = \begin{bmatrix} \mathbf{q}_C^\top & \mathbf{q}_\mathcal{L}^\top \end{bmatrix}^\top, \quad \mathbf{v} = \begin{bmatrix} \mathbf{v}_C^\top & \mathbf{v}_\mathcal{L}^\top \end{bmatrix}^\top.$$

The matrices $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{X}}$ can be partitioned as well, 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 $\mathbf{R}, \mathbf{X} \succ 0$, see [17]. Since the main goal is to present a framework for the design of Volt/Var schemes, we will assume that the variables $\mathbf{p}, \mathbf{q}_\mathcal{L}$ are not controlled. Using (1.2), we can express voltage magnitudes as functions exclusively of $\mathbf{q}_C$,

$$\mathbf{v}(\mathbf{q}_C) = \begin{bmatrix} \mathbf{X} \\ \mathbf{X}_\mathcal{L}^\top \end{bmatrix} \mathbf{q}_C + \hat{\mathbf{v}}, \tag{1.3}$$

where

$$\hat{\mathbf{v}} = \begin{bmatrix} \hat{\mathbf{v}}_C \\ \hat{\mathbf{v}}_\mathcal{L} \end{bmatrix} = \begin{bmatrix} \mathbf{X}_\mathcal{L} \\ \mathbf{X}_{\mathcal{L}\mathcal{L}} \end{bmatrix} \mathbf{q}_\mathcal{L} + \tilde{\mathbf{R}}\mathbf{p} + \mathbf{1}.$$

We point out that, although we employ the linearized power flow model in our stability analysis, the proposed algorithms can be implemented for the full nonlinear model. In fact, our simulations below use an exact AC power flow solver.

1.2.2 Problem Formulation

The deployment of DERs affects the DG operation, e.g., by introducing voltage quality issues. Indeed, sudden generation drops or excessive generation could have the voltages step outside desired operational limits. Since DERs are able to provide ancillary services, reactive power injection can be used to regulate the voltage profiles. Ideally, a system operator would like DER injections to be the solutions of an *optimal reactive power flow* (ORPF) problem, for which several

ORPF problem formulation exist. In this chapter, we consider

$$\mathbf{q}_C^\star(\mathbf{p}, \mathbf{q}_{\mathcal{L}}) := \arg\min_{\mathbf{q}_C} \; f(\mathbf{q}_C) \tag{1.4a}$$

$$\text{s.t.} \;\; (1.1a) - (1.1d)$$

$$\mathbf{v}_{\min} \le \mathbf{v} \le \mathbf{v}_{\max} \tag{1.4b}$$

$$\mathbf{q}_{\min} \le \mathbf{q}_C \le \mathbf{q}_{\max} \tag{1.4c}$$

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 the desired voltage lower and upper bounds on *all* the network buses; and $f : \mathbb{R}^C \to \mathbb{R}$ 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 (1.4b) via equation (1.1). Note that our framework can be adapted to OPRF problem instances different from (1.4). We assume that problem (1.4) admits a unique solution. When that is not the case, $\mathbf{q}_C^\star(\mathbf{p}, \mathbf{q}_{\mathcal{L}})$ can be chosen among the set of minimizers. For convenience, we denote the feasible set of reactive power injections as $\mathcal{Q} := \times_{n\in C}\mathcal{Q}_n$, with $\mathcal{Q}_n = \{q_n : q_{\min,n} \le q_n \le q_{\max,n}\}$.

Solving (1.4) is challenging due to the non-convexity of (1.1). The literature proposes several methods for solving ORPF problems exploiting, e.g., convex relaxations [4], linearized power flow equations [13], distributed optimization [11], and learning-based approaches [30]. However, solving (1.4) inevitably requires the knowledge of the network-wide quantities $(\mathbf{p}, \mathbf{q}_{\mathcal{L}})$, centrally or via peer-to-peer communication. This means that, when a reliable real-time communication network is not present, the optimal $\mathbf{q}_C^\star$ cannot be directly computed.

Our proposed control framework is inspired by the ongoing efforts [12,16–18] on local Volt/Var control rules and the recently reported success [7–10] of neural-network-based surrogates for OPF. Historical data are used to learn functions that map local measurements to (approximate) solutions of the ORPF problem (1.4). Specifically, for each agent $n \in C$, we aim to learn a function

$$\gamma_n : \mathbb{R} \times \mathcal{Q}_n \to \mathcal{Q}_n, \;\; (v_n, q_n) \mapsto \gamma_n(v_n, q_n)$$

that takes as inputs the voltage v_n and reactive power q_n, and outputs the ORPF solution approximation. We consider then local controllers whose equilibrium points are determined by the functions $\{\gamma_n\}_{n\in C}$. For this reason, the $\{\gamma_n\}_{n\in C}$ are hereafter called *equilibrium functions*. Precisely, we propose the reactive power update

$$q_n(t+1) = (1-\epsilon)q_n(t) + \epsilon\gamma_n(v_n(t), q_n(t)), \quad n \in C \tag{1.5}$$

with $\epsilon \in [0, 1]$. The new reactive power setpoint is a convex combination between the previous one and the actual value of the equilibrium function. Assuming that $q_n(0) \in \mathcal{Q}_n, n \in C$, it trivially follows that $\mathbf{q}_C(t) \in \mathcal{Q}, t \ge 0$, i.e., the reactive power injection is always feasible. Let $\mathbf{q}_C^\sharp$ be an equilibrium point

of (1.5), i.e.,

$$q_n^\sharp = \gamma_n(v_n^\sharp, q_n^\sharp) \tag{1.6a}$$

$$v_n^\sharp = v_n(\mathbf{q}_C^\sharp). \tag{1.6b}$$

Notice that $(\gamma_n(v_n^\sharp, q_n^\sharp), v_n^\sharp)$ is a point of the graph of γ_n, i.e., equilibria of (1.5) are ORPF approximate solutions. It remains now to understand the conditions that ensure the convergence of (1.5).

The stability analysis we provide assumes that uncontrolled variables, i.e., $\mathbf{p}$ and $\mathbf{q}_{\mathcal{L}}$, take arbitrary but fixed values in time. This is motivated by the assumption that the control will act on a fast timescale in which $(\mathbf{p}, \mathbf{q}_{\mathcal{L}})$ can be considered constant. Also, our analysis employs the power flow approximation (1.2). That is, after collecting the functions $\{\gamma_n\}_{n\in C}$ in the vector-valued function $\boldsymbol{\gamma}$, we study the stability of

$$\mathbf{q}_C(t+1) = (1-\epsilon)\mathbf{q}_C(t) + \epsilon\boldsymbol{\gamma}(\mathbf{v}_C(t), \mathbf{q}_C(t)) \tag{1.7a}$$

$$\mathbf{v}_C(t+1) = \mathbf{X}\mathbf{q}_C(t+1) + \hat{\mathbf{v}}_C. \tag{1.7b}$$

We present the stability results for the case when the equilibrium function depends only on local voltage in Section 1.3, and generalize the treatment to the case when the equilibrium function depends not only on local voltage but also on local reactive power in Section 1.4. A learning framework for synthesizing local Volt/Var controllers with stability guarantees from data is given in Section 1.5.

1.3 EQUILIBRIUM FUNCTIONS DEPENDING ONLY ON VOLTAGE

In this section we consider the scenario where the equilibrium function only depends on local voltage, which is the case for most of existing works in the literature, see e.g. [1,2,18]. This means that

$$\gamma(v_n, q_n) = \phi_n(v_n),$$

for each $n \in C$. The function ϕ_n is Lipschitz with the Lipschitz constant $L_n < \infty$. The next result identifies properties of the functions $\{\phi_n\}_{n\in C}$ that ensure the closed-loop system stability.

Proposition 1.3.1. *(Uniqueness of the equilibrium and global asymptotic stability of* (1.7)*): Let* $\mathbf{q}_C(0) \in \mathcal{Q}$*, and let*

$$\{\phi_n\}_{n\in C} \text{ are non-increasing.} \tag{C1}$$

Then, (1.7) *has an unique equilibrium point* $(\mathbf{q}_C^\sharp, \mathbf{v}_C^\sharp)$ *which is also* globally asymptotically stable *if*

$$0 < \epsilon < \min\left\{1, \frac{2}{\|\mathbf{X}\|L+1}\right\}. \tag{1.8}$$

Proof. The uniqueness is proved in [2, Proof of Proposition III.1]. We prove here the global asymptotic stability to correct some imprecisions in the treatment of [2, Proof of Proposition III.2]. Consider the voltage evolution under (1.7),

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

Define a diagonal matrix $\mathbf{M}(t) \in \mathbb{R}^{C\times C}$ with the n-th diagonal entry being

$$M_n(t) = \begin{cases} \frac{|\phi_n(v_n(t)) - \phi(v'_n(t))|}{|v_n(t) - v'_n(t)|} & v_n \neq v'_n, \\ 0 & v_n = v'_n. \end{cases}$$

Then, for any $\mathbf{v}_C(0), \mathbf{v}'_C(0) \in \mathbb{R}^C$, it follows that

$$\begin{aligned}&\mathbf{v}_C(t+1) - \mathbf{v}'_C(t+1) = (1-\epsilon)(\mathbf{v}_C(t) - \mathbf{v}'_C(t)) + \epsilon\mathbf{X}\big(\boldsymbol{\phi}(\mathbf{v}_C(t)) - \boldsymbol{\phi}(\mathbf{v}'_C(t))\big) \\ &= (1-\epsilon)(\mathbf{v}_C(t) - \mathbf{v}'_C(t)) - \epsilon\mathbf{X}\mathrm{sign}(\mathbf{v}_C(\mathrm{t}) - \mathbf{v}'_C(\mathrm{t}))|\boldsymbol{\phi}(\mathbf{v}_C(\mathrm{t})) - \boldsymbol{\phi}(\mathbf{v}'_C(\mathrm{t}))| \\ &= (1-\epsilon)(\mathbf{v}_C(t) - \mathbf{v}'_C(t)) - \epsilon\mathbf{X}\mathbf{M}(t)\mathrm{sign}(\mathbf{v}_C(\mathrm{t}) - \mathbf{v}'_C(\mathrm{t}))|\mathbf{v}_C(\mathrm{t}) - \mathbf{v}'_C(\mathrm{t})| \\ &= ((1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}\mathbf{M}(t))(\mathbf{v}_C(t) - \mathbf{v}'_C(t)),\end{aligned}$$

where we have used in the second equality the fact that ϕ_n is non-increasing in v_n for each $n \in C$. Consequently,

$$\mathbf{v}_C(t+1) - \mathbf{v}'_C(t+1) = \underbrace{\left[\prod_{i=0}^{t}((1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}\mathbf{M}(i))\right]}_{:=\mathbf{g}_t}(\mathbf{v}_C(0) - \mathbf{v}'_C(0)). \quad (1.9)$$

We show that for t large enough, the operator $\mathbf{g}_t$ is a contraction. Notice that $(1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}\mathbf{M}(i) = \mathbf{X}^{\frac{1}{2}}((1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}^{\frac{1}{2}}\mathbf{M}(i)\mathbf{X}^{\frac{1}{2}})\mathbf{X}^{-\frac{1}{2}}$, therefore it holds that

$$\mathbf{g}_t = \mathbf{X}^{\frac{1}{2}}\left(\prod_{i=0}^{t}(1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}^{\frac{1}{2}}\mathbf{M}(i)\mathbf{X}^{\frac{1}{2}}\right)\mathbf{X}^{-\frac{1}{2}}.$$

This implies that

$$\|\mathbf{g}_t\| \leq \|\mathbf{X}^{\frac{1}{2}}\|\|\mathbf{X}^{-\frac{1}{2}}\|\prod_{i=0}^{t}\|(1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}^{\frac{1}{2}}\mathbf{M}(i)\mathbf{X}^{\frac{1}{2}}\|.$$

Note that $\mathbf{X}^{\frac{1}{2}}\mathbf{M}(i)\mathbf{X}^{\frac{1}{2}} \geq 0$ is similar to $\mathbf{X}\mathbf{M}(i)$. Hence $\|(1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}^{\frac{1}{2}}\mathbf{M}(i)\mathbf{X}^{\frac{1}{2}}\| = \max\{|1-\epsilon-\epsilon\lambda_{\max}(\mathbf{X}\mathbf{M}(i))|, |1-\epsilon-\epsilon\lambda_{\min}(\mathbf{X}\mathbf{M}(i))|\}$. Therefore, $\|(1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}^{\frac{1}{2}}\mathbf{M}(i)\mathbf{X}^{\frac{1}{2}}\| < 1$ if and only if $1-\epsilon-\epsilon\lambda_{\max}(\mathbf{X}\mathbf{M}(i)) > -1$. Since $\lambda_{\max}(\mathbf{X}\mathbf{M}(i)) \leq \|\mathbf{X}\|L$, this directly follows from (1.8). Therefore, $\|(1-\epsilon)\mathbf{I} - \epsilon\mathbf{X}^{\frac{1}{2}}\mathbf{M}(i)\mathbf{X}^{\frac{1}{2}}\| < 1$ for any $0 \leq i \leq t$, and thus for sufficiently large t, we have that $\|\mathbf{g}_t\| < 1$, i.e., $\mathbf{g}_t$ is a contraction. Invoking the Banach's fixed-point theorem [31], we conclude that the equilibrium is globally asymptotically stable. □

This result means that, as long as the functions $\{\phi_n\}_{n\in\mathcal{C}}$ meet the conditions (C1), one can always find $\epsilon > 0$ so that $(\mathbf{q}_C, \mathbf{v}_C)$ converges to the unique equilibrium point $(\mathbf{q}_C^\sharp, \mathbf{v}_C^\sharp)$ under the reactive power update rule (1.5). Note that, in condition (1.8), $\|\mathbf{X}\|$ is fully determined by the DG and L can be computed once the functions $\{\phi_n\}_{n\in\mathcal{C}}$ have been selected. Because $\mathbf{q}_\mathcal{L}$ and $\mathbf{p}$ are fixed, the convergence of $\mathbf{q}_C$ leads also via (1.3) to the global asymptotic convergence of $\mathbf{v}$.

To conclude this section, we briefly compare (1.7) with the control schemes provided in the literature. Controllers like (1.5) are referred to as *incremental*, because the new reactive power is obtained by adding to the previous one an increment weighted by the stepsize parameter ϵ. Many works [15–17] propose local Volt/Var control schemes of the form

$$q_n(t+1) = \varphi_n(v_n(t)). \tag{1.10}$$

which can be referred to as *non-incremental* schemes [32]. Reactive powers are determined based on the local measurements and do not depend on the previous setpoints. The equilibrium points of (1.10) satisfy

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

i.e., $\varphi_n(v_n)$ plays the double role of the control function and the equilibrium function. Thus, even the control rule, which corresponds to $\epsilon = 1$ in (1.5),

$$q_n(t+1) = \phi_n(v_n(t)). \tag{1.11}$$

looks appealing for our framework because its equilibria are determined by ϕ_n. From the proof of Proposition 1.3.1, one can show that the algorithm (1.11) is globally asymptotically stable if the equilibrium functions meet not only (C1), but also

$$\|\mathbf{X}\|L < 1. \tag{1.12}$$

This condition bounds the slope of the functions $\{\phi_n\}_{n\in\mathcal{C}}$ and appears often in the literature [12,17,18,27]. This means that, to ensure the stability of non-incremental algorithms (1.10), 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 the adoption of an incremental control like (1.5).

1.4 EQUILIBRIUM FUNCTIONS WITH REACTIVE POWER AS AN ADDITIONAL ARGUMENT

In this section, we consider equilibrium functions that have as arguments both the local voltage and the reactive power. We focus in particular on separable

functions of the form

$$\gamma_n(q_n, v_n) = \phi_n(v_n) + \psi_n(q_n), \qquad n \in C \tag{1.13}$$

where ϕ_n and ψ_n are continuously differentiable functions depend solely on the local voltage and reactive power with Lipschitz constants $L_{\phi_n}, L_{\psi_n} < \infty$, respectively. For convenience, denote $L_{\boldsymbol{\phi}} := \max_{n\in C} L_{\phi_n}, L_{\boldsymbol{\psi}} := \max_{n\in C} L_{\psi_n}$. The next result provides conditions on the equilibrium functions that guarantee the uniqueness and asymptotic stability of the equilibrium.

Theorem 1.4.1. *(Uniqueness and global asymptotic stability of the equilibrium): Under equation (1.13), with ϕ_n and ψ_n continuously differentiable, the system (1.7) has an unique equilibrium point which is globally asymptotically stable if*

$$L_{\boldsymbol{\psi}} + L_{\boldsymbol{\phi}}\|\mathbf{X}\| < 1. \tag{C2}$$

We refer the reader to [3, Proof of Theorem III.1] for the proof. Note that, contrary to the classic solutions proposed in the literature, see e.g., [15,18], or to what we obtained in Section 1.3, condition (C2) does not constrain the equilibrium functions to be monotonic and allows them to have arbitrary shapes, as soon as their slopes "are not too steep". Indeed, it is trivial to see that the slopes of the functions $\{\phi_n, \psi_n\}_{n\in C}$ are bounded given that it must hold

$$L_{\boldsymbol{\phi}} < 1/\|\mathbf{X}\|, \quad L_{\boldsymbol{\psi}} < 1.$$

A consequence of this slope limitation is that the range of γ_n might be a strict subset of $\mathcal{Q}_n$, in particular when $q_{\max,n}$ is large. That is, the generation capability of agent n might not be fully exploited. This issue can be addressed by requiring the functions $\{\phi_n\}_{n\in C}$ to be *decreasing* and letting the functions $\{\psi_n\}_{n\in C}$ to have arbitrary shape but with limited slope. The next result, cf. [3, Theorem III.3], provides the condition for the stability in the aforesaid setup.

Theorem 1.4.2. *(Uniqueness and local asymptotic stability of the equilibrium): Under equation (1.13), with ϕ_n and ψ_n continuously differentiable, the system (1.7) has an unique equilibrium point if*

$$\{\phi_n\}_{n\in C} \text{ are decreasing and } L_{\boldsymbol{\psi}} < 1. \tag{C3}$$

Further, the equilibrium point is locally asymptotically stable if

$$\epsilon < \frac{2}{L_{\boldsymbol{\psi}} + L_{\boldsymbol{\phi}}\|\mathbf{X}\| + 1}. \tag{1.14}$$

The proof of this result can be found in [3, Proof of Theorem III.3]. We refer to condition (C2) as *coupled voltage-power slope constraint* (CVP-SC), whereas we refer to condition (C3) as *reactive power slope constraint* (RP-SC). We argue that the CVP-SC is suitable for DGs with many small-sized DERs since for

them the slope constraints do not lead to significant limits on their generation usage, and the more flexible shape of equilibrium functions enhances optimality. Instead, we the RP-SC is more suitable for the case of a DG with relatively big-size generators, since the lack of slope limitations on the functions $\{\phi_n\}_{n\in C}$ could help the generators make full use of their reactive power compensation capabilities and thus leads to better performance. We illustrate such trade-offs in the simulations through two different DG cases. Finally, we note that CVP-SC ensures the *global* asymptotic stability, whereas RP-SC ensures the *local* asymptotic stability, as the latter relies on the linearization of the operator $\mathbf{g}$ at equilibrium points. Moreover, CVP-SC allows arbitrary $\epsilon \in [0, 1]$, while RP-SC might be more restrictive on the selection of ϵ, cf. (1.14).

1.5 LEARNING EQUILIBRIUM FUNCTIONS FROM DATA

Having established the conditions on equilibrium functions for system stability, in this section we lay out a data-driven approach to synthesize the equilibrium functions, characterizing the equilibrium points of (1.7).

For the case that equilibrium functions depend only on voltages, cf. Section 1.3, $\gamma_n(v_n, q_n) = \phi_n(v_n)$ for each $n \in C$, the data set used in the training can be obtained following the steps in Box I.

Box I

1. Build a set $\{(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k)\}_{k=1}^K$ of K load-generation scenarios. The scenarios can be obtained from historical data retrieved by smart meters, via random sampling from given a priori probability distributions, or from load predictions.
2. Solve[1], for each scenario, the ORPF problem (1.4) to obtain $\mathbf{q}_C^{\star,k}(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k)$ and associated $\mathbf{v}_C^{\star,k}(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k, \mathbf{q}_C^{\star,k})$.
3. Build the labeled data set of minimizers $\mathcal{D} = \{(\mathbf{q}_C^{\star,k}, \mathbf{v}_C^{\star,k})\}_{k=1}^K$. Obtain data sets $\mathcal{D}_n = \{(v_n^{\star,k}, q_n^{\star,k})\}_{k=1}^K$ from $\mathcal{D}$ by separating the entries associated with each $n \in C$.

Then, for each $n \in C$, ϕ_n is learned using the data in $\mathcal{D}_n = \{(v_n^{\star,k}, q_n^{\star,k})\}_{k=1}^K$

1. Note that here, we tacitly assumed that, for any scenario $(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k)$, the solution $\mathbf{q}_{C,k}^{\star}$ exists or, in other words, that the ORPF problem is feasible. However, a DG might be under load-generation scenarios for which (1.4) does not have a solution and the voltages constraints are not satisfied. In such cases, engineering sense suggests to use the entire reactive power generation capability to alleviate the voltage violations as much as possible, i.e., to set $q_n = q_{\max,n}$ ($q_n = q_{\min,n}$) when $v_n < v_{\min,n}$ ($v_n > v_{\max,n}$), see [15]. To obtain the functions $\{\phi_n\}_{n\in C}$ that comply with the former observation, a certain number of additional data points, referred to as *pseudo points*, can be added to the data set, e.g., pairs of the form $\{(\underline{v}_{n,k}, q_{\max,n})\}_{k=1}^{\underline{K}}$, $\underline{v}_{n,k} \leq v_{\min,n}$ and $\{(\overline{v}_{n,k}, q_{\min,n})\}_{k=1}^{\overline{K}}$, $\overline{v}_{n,k} \geq v_{\max,n}$. These points could be uniformly spaced or randomly sampled.

by solving a problem of the form

$$\min_{\phi_n} \sum_{k=1}^{K} |q_n^{\star,k} - \phi_n(v_n^{\star,k})|^2 \quad \text{s.t. } \phi_n \text{ meets the condition (C1).} \tag{1.15}$$

For the case that equilibrium functions depend on both voltages and reactive powers, cf. Section 1.4, $\gamma_n(v_n, q_n) = \phi_n(v_n) + \psi_n(q_n)$ for each $n \in C$, the data set is instead obtained following the steps in Box II.

Box II

1. Build a set $\{(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k, \mathbf{q}_C^k)\}_{k=1}^K$ of K load-generation scenarios. The scenarios can be obtained similarly to what we described in Box II.
2. Solve, for each scenario, the power flow equation (1.1) and the ORPF problem (1.4) to obtain $\mathbf{v}_C^k(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k, \mathbf{q}_C^k)$ and $\mathbf{q}_C^{\star,k}(\mathbf{p}^k, \mathbf{q}_{\mathcal{L}}^k)$, respectively.
3. Build the labeled data set of minimizers $\mathcal{D} = \{(\mathbf{v}_C^k, \mathbf{q}_C^k, \mathbf{q}_C^{\star,k})\}_{k=1}^K$. Obtain data sets $\mathcal{D}_n = \{(v_n^k, q_n^k, q_n^{\star,k})\}_{k=1}^K$ from $\mathcal{D}$ by separating the entries associated with each $n \in C$.

Then, for each $n \in C$, γ_n is learned using the data in $\mathcal{D}_n = \{(v_n^k, q_n^k, q_n^{\star,k})\}_{k=1}^K$ by solving a problem of the form

$$\min_{\gamma_n} \sum_{k=1}^{K} |q_n^{\star,k} - \gamma_n(v_n^k, q_n^k)|^2 \quad \text{s.t. } \gamma_n \text{ meets the conditions (C2) or (C3).} \tag{1.16}$$

Typical approaches to solve (1.15) or (1.16) involve restricting the search space for the function ϕ_n and ψ_n via a convenient parameterization, e.g., polynomial regression or neural network approximation methods. In the following, we provide a single-hidden-layer neural network design framework that achieves (C1) – (C3) and uses the Rectified Linear Unit (ReLU) activation function

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

Consider the single-hidden-layer neural network of the form

$$\mathsf{N}(x) = \sum_{h=1}^{H} w_h \text{ReLU}(x - b_h) + \beta, \tag{1.17}$$

with w_h and b_h being the weight and bias of the h-th neuron unit, respectively; β being an additional bias term applied in the output layer; and H being the number of neuron units in the hidden layer. The neural network (1.17) is a composite of

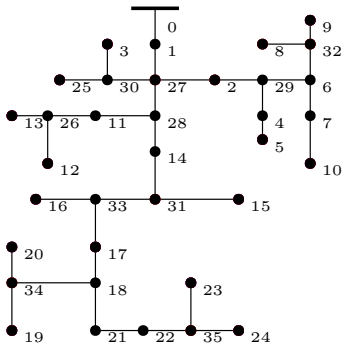

FIGURE 1.1 The modified IEEE 37-bus feeder.

H linear segments. Reorder the neuron units such that $b_1 \leq b_2 \leq \cdots \leq b_H$, an important property is that the slope of J-th segment of N can be expressed as [2]

$$\sum_{j=1}^{J} w_j, \quad J \in \{1, 2, \ldots, H\}. \tag{1.18}$$

Leveraging this property, and by parameterizing ϕ_n and ψ_n using the neural network N, one can easily design the conditions on the weights to ensure (C1) – (C3). Indeed, by making sure $\sum_{j=1}^{J} w_j < (\leq) 0$ for all $J \in \{1, 2, \ldots, H\}$, one can guarantee the learned functions to be decreasing (non-increasing), and, by posing constraints to $\max_{J \in \{1,..,H\}} \left| \sum_{j=1}^{J} w_j \right|$, one can also bound the Lipschitz constant of the learned functions. This helps to convert the optimization problems (1.15) and (1.16) to convenient parameterized formulations. Then they can be solved using suitable renditions of (stochastic) gradient descent prevalent for neural network training, e.g., the Adam algorithm [33].

1.6 CASE STUDY

In this section, we test the proposed framework on the modified single-phase equivalent of the IEEE 37-bus system, where we have also omitted its regulators, shown in Fig. 1.1. The data sets for the learning process are built using the same minute-based load and solar generation data profiles used in [2], which are intended as day-ahead forecasts. The voltage limits are set to $v_{\min,n} = 0.95$ p.u. and $v_{\min,n} = 1.05$ p.u. for every $n \in \mathcal{N}$.

1.6.1 Equilibrium Functions Depending Only On Voltage

Here we validate the results obtained in Section 1.3. The grid hosts 5 DERs are placed at nodes $C_1 = \{27, 31, 32, 34, 35\}$ with uniform generation capabilities, i.e., $q_{\max,n} = 0.4$ MVAR, and $q_{\min,n} = -q_{\max,n}$ for every $n \in C_1$. We benchmark

the proposed strategy against the optimized droop control design from [22]

$$q_n(t+1) = \rho_n(v_n) := \begin{cases} q_{\max,n} & v_n \le v_{\min,n}, \\ \frac{\bar{v}_{\min,n} - v_n}{\bar{v}_{\min,n} - v_{\min,n}} q_{\max,n} & v_{\min,n} < v_n < \bar{v}_{\min,n}, \\ 0 & \bar{v}_{\min,n} \le v_n \le \bar{v}_{\max,n}, \\ \frac{v_n - \bar{v}_{\max,n}}{\bar{v}_{\max,n} - \bar{v}_{\max,n}} q_{\min,n} & \bar{v}_{\max,n} < v_n < v_{\max,n}, \\ q_{\min,n} & v_n \ge v_{\max,n} \end{cases}$$

where the parameters $\bar{v}_{\min,n}$ and $\bar{v}_{\max,n}$, with $v_{\min,n} < \bar{v}_{\min,n} \le \bar{v}_{\max,n} < v_{\max,n}$, are optimized given day-ahead forecasts.

We consider the cost in (1.4) to be

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

where ① and ② stand for the voltage deviations and power losses [12], respectively, and α is a trade-off parameter for those two objectives. The ORPF problem (1.4) is solved with linearized power flow (1.2) using the CVX toolbox [34]. However, we note that any other power flow models can be used to solve the ORPF problem.

We add pseudo data points to the obtained data set with $\underline{K} = \overline{K} = 700$, which results in a total of 2840 data points for each DER. We implement the neural network approach according to Section 1.5 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 [33] with the learning rate initialized at 0.01, decaying every 500 steps with a base of 0.5.

The equilibrium functions are computed by solving (1.15). Fig. 1.2 plots the equilibrium function ϕ_{32} learned with and without pseudo points, and the optimized droop function ρ_{32} for the DER at node 32 with $\alpha = \frac{1}{3}$. Compared 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, with little change to the learned curve when voltage is with in the limits. We further summarize in Table 1.1 the average loss for the whole training data set using the learned equilibrium functions and optimized droop control functions, i.e.,

$$\frac{\sum_{k=1}^{K} \|\mathbf{q}^\star_{C,k} - \square(\mathbf{v}^\star_{C,k})\|^2}{K}, \tag{1.19}$$

where $\square$ is $\boldsymbol{\phi}$ for the data-based method and $\boldsymbol{\rho}$ for the optimized droop control. The results illustrate the enhanced optimality of the learned equilibrium functions in approximating ORPF solutions compared to the benchmark method.

TABLE 1.1 Average loss values for all data profiles

α	0	1/3	1/2	2/3	1
Learned equilibrium function	0.045	0.030	0.036	0.088	0.189
Optimal droop function	0.175	0.086	0.215	0.362	0.439

We also illustrate the advantage of using the incremental algorithm in Fig. 1.2. Recall that to guarantee the convergence of the non-incremental algorithm, i.e., $\epsilon = 1$, one needs to further enforce an additional slope constraint (1.12) on the learned equilibrium functions. Fig. 1.2 shows that this additional slope constraint leads to worse approximation performance of the learned equilibrium functions in fitting the data set (we do not consider the pseudo points during learning here for fairness), and thus degrades the optimality of system performance.

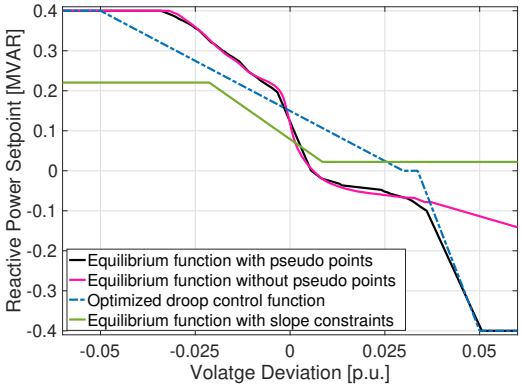

FIGURE 1.2 The learned equilibrium functions under different settings and the optimized droop function.

Having learned the equilibrium functions, next we run simulations for the case $\alpha = \frac{1}{3}$, with the pseudo data points considered, and assume that $\mathbf{q}_C(0) = \mathbf{0}$. We use MATPOWER [35] to solve the power flow equation. First, we verify the convergence properties of the proposed reactive power update rule (1.5) stated in Proposition 1.3.1. Consider the scenario where the load-generation profiles are fixed. Fig. 1.3 reports the evolution of the DERs' reactive power setpoints using load-generation profiles of the 695-th minute and considers 20 iterations of (1.5). For $\epsilon = 0.369$, the reactive power setpoint trajectories converge to their final values, cf. Fig. 1.3(a), whereas for $\epsilon = 1$ it fails, cf. Fig. 1.3(b). This is consistent with the sufficient condition

$$0 < \epsilon < \min\left\{1, \frac{2}{\|\mathbf{X}\|L + 1}\right\} = 0.859$$

derived in Proposition 1.3.1.

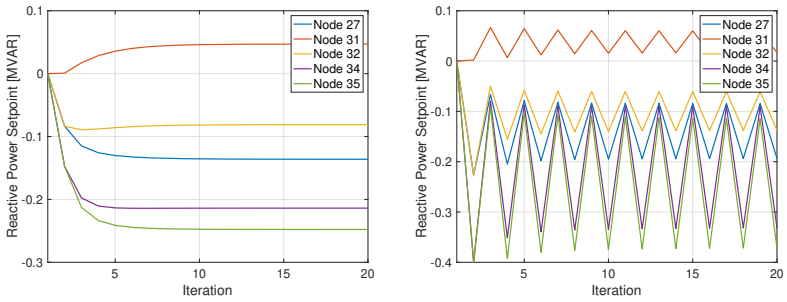

(a) $\epsilon = 0.369$ (b) $\epsilon = 1$

FIGURE 1.3 Evolution of reactive power setpoints under the proposed reactive power update rule (1.5) with (a) $\epsilon = 0.369$ and (b) $\epsilon = 1$, where we use the power data profiles of the 695-th minute and consider 20 iterations. This verifies the sufficient condition $0 < \epsilon < \min\{1, \frac{2}{\|\mathbf{X}\|L+1}\} = 0.859$ in Proposition 1.3.1 to ensure global asymptotic stability.

Next, we test the proposed data-based control method in a scenario where the load-generation profiles are time-varying. We obtain the testing 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 data set prescribing a day-ahead forecast, whereas their random perturbation acts as the true realization of the load-generation scenarios. These loads and generations are minute-based and we consider 120 iterations of (1.5) per minute with $\epsilon = 0.369$. Fig. 1.4 compares the evolution of the maximum/minimum voltages under the proposed data-based control method, the optimized droop control method, the ORPF solutions, and the case where no control action is taken. One can observe that, the optimized droop control method induces instability issues within 12:00 and 16:00 causing voltages to oscillate. To further illustrate the effectiveness and advantages of the proposed data-based control method, Table 1.2 summarizes the comparison results of the proposed data-based control method against the 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 the distances of the actual and optimal reactive power setpoints, i.e., average of $\|\mathbf{q}_C(t) - \mathbf{q}_C^\star(t)\|$ for the entire day. It can be observed that the proposed data-based control method outperforms the benchmark method in all cases.

TABLE 1.2 Average distances between actual reactive power setpoints and ORPF solutions for entire day

α	0	1/3	1/2	2/3	1
Data-based Control	0.119	0.099	0.078	0.112	0.165
Optimal Droop Control	0.279	0.247	0.373	0.441	0.485
No Control	0.316	0.508	0.684	0.803	0.836

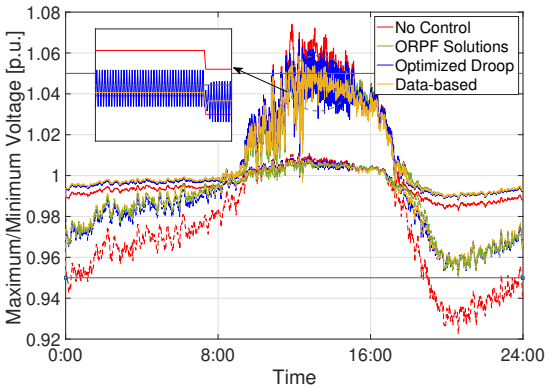

FIGURE 1.4 Evolution of the maximum (solid line) and minimum (dashed line) voltages of the IEEE 37-bus network under the proposed data-based, ORPF solutions, optimized linear droop, and no control methods. For all minute-based data profiles, the ORPF problem is feasible and thus $\mathbf{q}_C^\star$ 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.

Our simulation results above validate the improved performance of the proposed data-based method compared to the optimized 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 specific control goals, to learn purely local controllers that steer system operating points to approximated ORPF solutions. An important observation is that, although the ORPF approach strictly guarantees that the voltages are within limits, our approach does not. For instance, in Fig. 1.4, 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 [19], purely local control strategies generally have no guarantee on desired regulation, in the sense that the equilibrium $\mathbf{q}_C^\sharp$ of (1.7) could result in $\mathbf{v}(\mathbf{q}_C^\sharp) \notin [\mathbf{v}_{\min}, \mathbf{v}_{\max}]$, even if there indeed exists $\mathbf{q}_C$ such that $\mathbf{v}(\mathbf{q}_C) \in [\mathbf{v}_{\min}, \mathbf{v}_{\max}]$.

1.6.2 Equilibrium Functions with Reactive Power as an Additional Argument

Here we validate the approach developed in Section 1.4. We benchmark the resulting control rules against the method in Section 1.3, where the equilibrium functions solely depend on local voltages.

We set the cost in (1.4) to be

$$f(\mathbf{q}_C) = \mathbf{q}^\top \tilde{\mathbf{R}} \mathbf{q} + \mathbf{p}^\top \tilde{\mathbf{R}} \mathbf{p},$$

meaning purely minimizing power losses. We consider two scenarios. In `Case-1`, there are 5 generators, placed at buses $C_1 = \{27, 31, 32, 34, 35\}$, with generation capability $\mathbf{q}_{\max} = 0.4 \times \mathbf{1}$ MVAR; in `Case-2`, there are 10 generators, placed at buses $C_2 = C_1 \cup \{6, 18, 28, 29, 33\}$, with generation capability $\mathbf{q}_{\max} = 0.2 \times \mathbf{1}$ MVAR. In both the cases, $\mathbf{q}_{\min} = -\mathbf{q}_{\max}$. Notice that in `Case-2`, we have a bigger number of smaller generators than in `Case-1`. For our experiments, we use the same minute-based load and solar generation data in Section 1.6.1. Moreover, we randomly generate five reactive power injection $\mathbf{q}_C$ from $[\mathbf{q}_{\min}, \mathbf{q}_{\max}]$ for each minute-based data, resulting in a total of $K = 1440 \times 5 = 7200$ load-generation profiles. We use the CVX toolbox [34] to solve the power flow equation (1.2) as well as the ORPF problem (1.4) for all load-generation profiles. The training setup is the same as in Section 1.6.1.

The equilibrium functions are computed by solving (1.16). Consistent with our discussion in Section 1.4, we use RP-SC in `Case-1` since it is more suitable for the case with relatively big size generators, and use CVP-SC in `Case-2` due to its advantages for the case with relatively small size generators. In all cases, we use the method in Section 1.3 as the baseline. Table 1.3 illustrates the learning performance for the two test cases. In both, the algorithms developed under CVP-SC in Theorem 1.4.1 and RP-SC in Theorem 1.4.2, when compared with the one in Section 1.3, achieve lower training loss, computed as $\frac{1}{K} \sum_{n \in C} \sum_{k=1}^{K} |q_n^{\star,k} - \gamma_n(v_n^k, q_n^k)|^2$, cf. Table 1.3. This shows that the inclusion of reactive power as argument of the equilibrium function helps increase the prediction accuracy. Figs. 1.5 and 1.6 plot the learned functions ϕ_{35} and ψ_{35} for `Case-1` and `Case-2`, respectively. In `Case-1`, although CVP-SC allows the function ϕ_n to have arbitrary shape, the learned function ϕ_{35} in Fig. 1.5 is decreasing. Thus, the more restrictive slope limitations of CVP-SC make its performance worse than RP-SC in approximating the ORPF solutions, which explains CVP-SC yielding greater training loss than RP-SC for `Case-1` in Table 1.3. Instead, in `Case-2`, although CVP-SC has more restrictive slope limitations on the functions $\{\phi_n\}_{n \in C}$, it does not affect much the performance since the generation capability is relatively small. Instead, the monotonicity requirement for the functions $\{\phi_n\}_{n \in C}$ in RP-SC degrades the prediction accuracy as one can see that the learned function ϕ_{35} in Fig. 1.6 for CVP-SC is not always decreasing. This is the reason that CVP-SC works better for `Case-2`, as Table 1.3 suggests. These observations are consistent with our discussion in Section 1.4.

We test the control performance of the proposed methods under CVP-SC and RP-SC. Although our stability analysis is done for the linearized power flow model, here we employ MATPOWER [35] to solve the AC power flow to

1. The method in Section 1.3 is a special case of RP-SC with $\psi_n = 0$ for all $n \in C$.
2. The method in Section 1.3 is a special case of RP-SC with $\psi_n = 0$ for all $n \in C$.

TABLE 1.3 Average training loss and optimality gap distance between actual reactive power setpoints and ORPF solutions[2]

	Training loss		Average distance	
	Baseline	Improvement	Baseline	Improvement
`Case-1`	3.15×10^{-2}	39.7%	0.1969	13.09%
`Case-2`	7.9×10^{-3}	14.0%	0.1235	39.41%

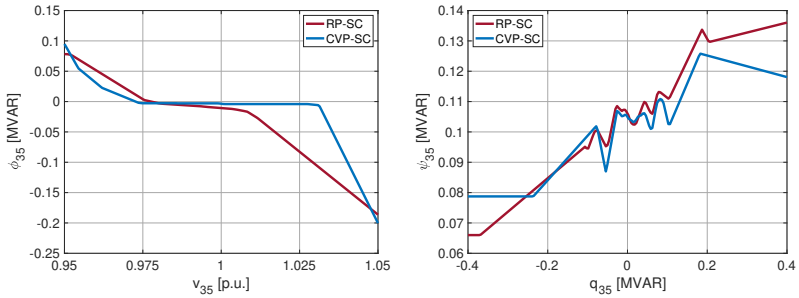

(a) ϕ_{35} (b) ψ_{35}

FIGURE 1.5 Leaned functions (a) ϕ_{35} and (b) ψ_{35} of node 35 for `Case-1` under CVP-SC and RP-SC.

run the simulations. We test the proposed methods in a scenario where load-generation profiles are time-varying. Specifically, we randomly perturb (5%) the consumption data between 15:00 and 17:00 to obtain the testing load-generation profiles. We set $\epsilon = 0.1$ and consider 120 iterations of (1.5) using the controllers developed under CVP-SC and RP-SC. Table 1.3 summarizes the average distances between the actual reactive power setpoints and the ORPF solutions. The performance displayed here by CVP-SC and RP-SC illustrates their respective advantages in different DG cases and their significant improvement compared to the baseline. Notably, we observe that the performance improvement achieved by CVP-SC in `Case-2` is greater than that by RP-SC in `Case-1`. This is because in `Case-2` CVP-SC enjoys the performance improvement from the inclusion of reactive power as argument of the equilibrium function as well as the more flexible shape of equilibrium function. Instead, in `Case-1`, the performance improvement achieved by RP-SC is only due to the inclusion of reactive power as argument of the equilibrium function.

1.7 CONCLUSIONS

We have presented a data-driven framework to design local Volt/Var controllers to steer a DG towards efficient configurations described by the equilibrium func-

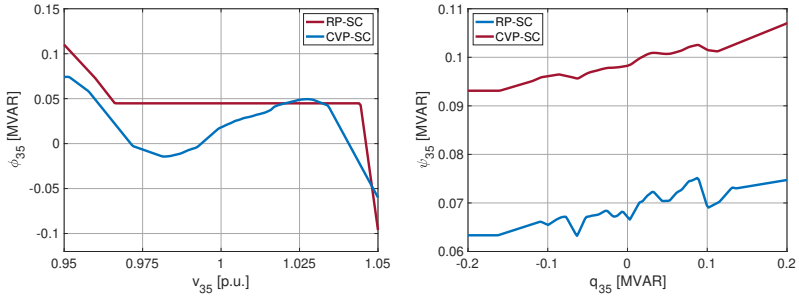

(a) ϕ_{35} (b) ψ_{35}

FIGURE 1.6 Leaned functions (a) ϕ_{35} and (b) ψ_{35} of node 35 for `Case-2` under CVP-SC and RP-SC.

tions. Equilibrium functions are learned using labeled data sets built using historical data and solving ORPF problems. The control function represents an incremental rule that updates, for every agent, the reactive power injection as a convex combination between the previous setpoint and the present equilibrium function value. Conditions on the equilibrium functions and on the control parameters ensuring the stability of the overall systems are formally derived and discussed. We then have validated our approach via AC power flow simulations, which show significant improvements compared to prevalent local control approaches. Future research directions include considering the regulation of legacy devices and unbalanced DGs, enhancing data consistency by making use of other local information in building the data set, reducing the optimality gap during the learning process, and extending the proposed framework to a more general scenario where we take advantage of communication among neighboring agents.

ACKNOWLEDGEMENT

The authors would like thank Dr. Manish K. Singh from the University of Minnesota, Twin Cities for fruitful discussions. 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.

[1] G. Cavraro, Z. Yuan, M. K. Singh, J. Cortés, Learning local Volt/Var controllers towards efficient network operation with stability guarantees, in: IEEE Conf. on Decision and Control, Cancun, Mexico, 2022, pp. 5056–5061.
[2] Z. Yuan, G. Cavraro, M. K. Singh, J. Cortés, Learning provably stable local Volt/Var controllers for efficient network operation, IEEE Transactions on Power SystemsTo appear.
[3] Z. Yuan, G. Cavraro, J. Cortés, Constraints on OPF surrogates for learning stable local Volt/Var controllers, IEEE Control Systems LettersTo appear.
[4] S. H. Low, Convex relaxation of optimal power flow - Part I: Formulations and equivalence, IEEE Transactions on Control of Network Systems 1 (1) (2014) 15–27.
[5] B. Cui, X. A. Sun, A new voltage stability-constrained optimal power-flow model: Sufficient condition, SOCP representation, and relaxation, IEEE Transactions on Power Systems 33 (5) (2018) 5092–5102.
[6] L. Gan, N. Li, U. Topcu, S. H. Low, Optimal power flow in tree networks, in: IEEE Conf. on Decision and Control, Firenze, Italy, 2013, pp. 2313–2318.
[7] X. Pan, T. Zhao, M. Chen, S. Zhang, DeepOPF: A deep neural network approach for security-constrained DC optimal power flow, IEEE Transactions on Power Systems 36 (3) (2021) 1725–1735.
[8] F. Fioretto, T. W. K. Mak, P. Van Hentenryck, Predicting AC optimal power flows: Combining deep learning and lagrangian dual methods, in: AAAI Conference on Artificial Intelligence, New York, USA, 2020, pp. 630–637.
[9] D. Owerko, F. Gama, A. Ribeiro, Optimal power flow using graph neural networks, in: IEEE Int. Conf. on Acoustics, Speech and Signal Processing, Barcelona, Spain, 2020, pp. 5930–5934.
[10] M. K. Singh, V. Kekatos, G. B. Giannakis, Learning to solve the AC-OPF using sensitivity-informed deep neural networks, IEEE Transactions on Power Systems 37 (4) (2022) 2833–2846.
[11] E. Dall'Anese, H. Zhu, G. B. Giannakis, Distributed optimal power flow for smart microgrids, IEEE Transactions on Smart Grid 4 (3) (2013) 1464–1475.
[12] G. Cavraro, R. Carli, Local and distributed voltage control algorithms in distribution networks, IEEE Transactions on Power Systems 33 (2) (2017) 1420–1430.
[13] E. Dall'Anese, A. Simonetto, Optimal power flow pursuit, IEEE Transactions on Smart Grid 9 (2) (2018) 942–952.
[14] G. Qu, N. Li, Optimal distributed feedback voltage control under limited reactive power, IEEE Transactions on Power Systems 35 (1) (2020) 315–331.
[15] 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) (2018) 1–138doi:10.1109/IEEESTD.2018.8332112.
[16] K. Turitsyn, P. Sulc, S. Backhaus, M. Chertkov, Options for control of reactive power by distributed photovoltaic generators, Proceedings of the IEEE 99 (6) (2011) 1063–1073.
[17] H. Zhu, H. J. Liu, Fast local voltage control under limited reactive power: Optimality and stability analysis, IEEE Transactions on Power Systems 31 (5) (2015) 3794–3803.
[18] X. Zhou, M. Farivar, Z. Liu, L. Chen, S. H. Low, Reverse and forward engineering of local voltage control in distribution networks, IEEE Transactions on Automatic Control 66 (3) (2021) 1116–1128.
[19] S. Bolognani, R. Carli, G. Cavraro, S. Zampieri, On the need for communication for voltage regulation of power distribution grids, IEEE Transactions on Control of Network Systems 6 (3) (2019) 1111–1123.
[20] X. Sun, J. Qiu, J. Zhao, Optimal local Volt/Var control for photovoltaic inverters in active distribution networks, IEEE Transactions on Power Systems 36 (6) (2021) 5756–5766.
[21] S. Karagiannopoulos, P. Aristidou, G. Hug, Data-driven local control design for active distri-

bution grids using off-line optimal power flow and machine learning techniques, IEEE Trans. Smart Grid 10 (6) (2019) 6461–6471. doi:10.1109/TSG.2019.2905348.

[22] H. Ji, C. Wang, P. Li, J. Zhao, G. Song, F. Ding, J. Wu, A centralized-based method to determine the local voltage control strategies of distributed generator operation in active distribution networks, Applied Energy 228 (2018) 2024–2036.

[23] W. Cui, Y. Jiang, B. Zhang, Reinforcement learning for optimal primary frequency control: A Lyapunov approach, IEEE Transactions on Power SystemsTo appear.

[24] Z. Yuan, C. Zhao, J. Cortés, Reinforcement learning for distributed transient frequency control with stability and safety guarantees, Systems & Control LettersSubmitted.

[25] W. Cui, J. Li, B. Zhang, Decentralized safe reinforcement learning for inverter-based voltage control, Electric Power Systems Research 211 (2022) 108609.

[26] Y. Shi, G. Qu, S. H. Low, A. Anandkumar, A. Wierman, Stability constrained reinforcement learning for real-time voltage control, in: American Control Conference, Atlanta, GA, 2022, pp. 2715–2721.

[27] C. Zhang, Y. Xu, Y. Wang, Z. Y. Dong, R. Zhang, Three-stage hierarchically-coordinated Voltage/Var control based on PV inverters considering distribution network voltage stability, IEEE Transactions on Sustainable Energy 13 (2) (2021) 868–881.

[28] S. Gupta, S. Chatzivasileiadis, V. Kekatos, Deep learning for optimal Volt/VAR control using distributed energy resources, arXiv preprint arXiv:2211.09557.

[29] A. M. Kettner, M. Paolone, On the properties of the compound nodal admittance matrix of polyphase power systems, IEEE Transactions on Power Systems 34 (1) (2019) 444–453.

[30] M. K. Singh, S. Gupta, V. Kekatos, G. Cavraro, 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, 2020.

[31] A. Granas, J. Dugundji, Fixed Point Theory, Vol. 14, Springer, 2003.

[32] M. Farivar, X. Zhou, 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, 2015, pp. 732–737.

[33] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: International Conference on Learning Representations, San Diego, CA, USA, 2015.

[34] M. Grant, S. Boyd, CVX: Matlab software for disciplined convex programming, version 2.1, available at http://cvxr.com/cvx (Mar. 2014).

[35] R. D. Zimmerman, C. E. Murillo-Sánchez, R. J. Thomas, Matpower: Steady-state operations, planning and analysis tools for power systems research education, IEEE Transactions on Power Systems 26 (1) (2011) 12–19.