International Journal of Engineering Sciences Paradigms and Researches (Volume 47, Issue: Special Issue of January 2018)
ISSN (Online): 2319-6564 and Website: www.ijesonline.com

Controlling in power systems: Design and analytical recommendations for specific safety issues

[1]Mr. Sarada Prasad Sahoo, [2] Dr. Ajaya Kumar Swain
[1*] *Associater Professor, Dept. Of Electrical Engineering, NIT BBSR,*
Professor DEPT. of Electrical Engineering, NIT BBSR,
[1]sarada@thenalanda.com, ajaykumar@thenalanda.com

A B S T R A C T

Artificial intelligence (AI) and machine learning are advancing quickly, which has renewed interest in their potential applications in power systems for cutting-edge control methods that support the integration of higher levels of renewable generation and address rising levels of uncertainty and variability. The most important new safety hazards and issues that arise when depending on learning for control in electric grid operations are discussed in this study along with these new applications. We build on recent taxonomical work in AI safety and focus on four concrete safety problems. We draw on two case studies, one in frequency regulation and one in distribution system control, to exemplify these problems and show mitigating measures. We then provide general guidelines and literature to help people working on integrating learning capabilities for control purposes to make safety risks a central tenet of design

1. Introduction

Over the last decade, research on machine learning (ML) and artificial intelligence (AI) has been growing and maturing, leading to an extensive variety of efficient algorithms to learn parameters and functions from historical data and real-time measurements. Energy and power systems scholars are rapidly developing new learning-based strategies to control dynamics in various areas of power systems operations, at different timescales. Some reasons behind these new learning-based approaches are to develop more sophisticated levels of control, to account for unmodeled uncertainty or inherently random aspects of a system, to adapt to changing conditions, or to manage new aspects of smart grids. While these developments hold promise for operating networks more efficiently and under higher levels of renewable generation, they also introduce new vulnerabilities and safety risks [1], both technical and non-technical. These vulnerabilities are in part inherent to the nature of learning modalities, but also arise as new control methods are integrated in *legacy* infrastructure and practices. Traditional power systems tend to rely on relatively simple *fit-and-forget* control logics that historically did not need to be updated often, as systems were "overdimensioned" to sustain flow conditions for long term projections. The rapid and constant evolution of power systems, largely due to the energy transition, defies this traditional approach necessitating new control logics that leverage *flexibility* and try to get the most out of existing infrastructure to prevent expensive updates. For ML approaches to facilitate this transition, these have to be appropriately designed to *integrate* in existing and often aging infrastructure, in such a way that no new safety risks are introduced and operators can work with them effectively. This requires these techniques to also be adequately connected to the system operator's planning cycles to guarantee learned behaviors remain adequate in shifting environments. Understanding and mitigating new vulnerabilities that arise through the use of ML requires additional analysis and design thinking that is often overlooked in traditional power systems literature.

So far, ML approaches in power systems have mostly focused on their affordances and little work covers new safety problems. The work in [2] provides an overview on how reinforcement learning (RL) has been used in power systems, discussing general trends, common algorithms used, and specific power system applications. It does not discuss safety issues. The authors of [3] discuss stability in power systems with the use of reinforcement learning, explaining how central RL assumptions may break in power systems applications, making the use of these algorithms risky. In fact, RL approaches at large are mostly devoid of strong safety guarantees. Nevertheless, recent approaches outside the

power systems' community are using RL considering safety and stability guarantees to address these risks [4,5].

In this paper, we discuss general safety-critical issues that can arise with the use of ML and AI, even when their mathematical assumptions are satisfied. We examine concrete safety problems in the design of learning-based controllers for power systems and propose guidelines for adequately analyzing and designing for these, making three contributions. First, in Section 2, we review representative literature using machine learning techniques in power system control problems, and we discuss safety problems that arise in these contexts. We emphasize that there are plenty of other ML applications, such as security assessment and vulnerability, that are not directly explored in this paper but can also benefit from guidelines proposed in this work. In Section 3, we draw inspiration from [1] to categorize safety risks as expressions of typical AI failure modes, including negative (unintended) side effects, partial training data, safe exploration and distributional shifts. With these categories, we ground our analyses in Section 4, through two simulation results: controller design for frequency regulation in transmission systems [6], and data-driven decentralized control for (multi-agent) distributed energy resources in distribution networks [7]. These case studies provide sufficient breadth to discuss different power systems contexts and timescales and the impact of learning on critical notions of safety. In Section 5, we propose guidelines for how safety risks can be addressed along the research, design, training and implementation of the proposed controller in the system. Finally, Section 6 concludes our paper.

2. Machine learning and artificial intelligence for power system's control

In this study, we focus on the usage and risks of ML and generally AI for controlling power systems. There exist several applications for power systems that can benefit from using these techniques. In Table 1, we present a non-exhaustive but representative summary of ML applications in power systems. In the framework of control theory, ML and AI techniques can be used in different areas: system identification, state estimation, disturbance and target prediction, and control action. In this paper we narrow down our focus to the use of ML and AI for determining control actions in power systems operations. It is out of scope to explore the use of ML in the context of: forecasting in power systems (e.g. load, wind or solar capacity factors) [28,29,31,32], clustering of electricity demand profiles [33], state/parameter estimation [21,23], theft detection [26,27], etc. Following the problem areas in Table 1, we cover the references and discuss what safety concerns arise.

Frequency regulation

The work presented in [6,34] proposes a frequency regulation control scheme learned using least squares and lasso regression, respectively. In [6,34], power dynamics are modeled as time-variant due to the change in the inertia coefficients in the grid [11]. The training set in these studies is generated through solving optimization problems that find optimal actions for controlling frequency. A controller proportional to the states (frequency and angles) is assumed and learned from the training set. One of the challenges that this control design faces is robustness to safety-critical states (safe and persistent excitation/exploration), i.e. unusual high deviations in frequency. To address this, the training set includes optimal control actions when facing these less common but safety-critical initial states of high frequency deviations. If these stressed conditions are not included in the training set, the controller would perform well only under mild disturbances. Thus, it would not be able to steer frequency deviations back to zero at all times. Lastly, when designing a controller for a dynamical system, stability guarantees for the closed loop system are required for safety purposes. This work adds virtual inertia (controller proportional to the derivative of the frequency) to the controller in order to guarantee stability [6].

Table 1
Summary of applications of ML/AI on Power Systems.

ML/AI Application	Problem	References
Data-Driven Control Design	Voltage Regulation	[8–10]
	Frequency Regulation	[6,11]
	System Dispatch	[8,12,13]
	Demand Response	[14–19]
	Storage Management	[12,13,20]
System Identification	Fault restoration / identification	[21,22]
	Forecasting	
Parameter Estimation / Series	Network Observability	[23–25]
	Electricity Theft	[26,27]
	Unit Commitment	[28,29]

Voltage regulation

Relatively more studies propose data-driven controllers which can regulate voltages in distribution systems, e.g. [8–10,12]. Most of these employ offline optimal power flow calculations to derive the training dataset composed of the optimal DER setpoints. These are used to design local controllers using different ML models and features. The work in [8,9] apply multiple linear regression to a set of local features and derives the individual DER controllers for reactive power control. The method does not interfere with the operation of existing legacy equipment, such as load tap changers and capacitor banks. This method is further developed in [10], in which also active power control and voltage balancing are considered. Reference [12] also considers active power curtailment, energy storage and controllable loads, by using segmented regression and support vector machines. To account for uncertainties, the offline calculations are formulated using chance constraints. However, both references focused on expected operating behaviors without evaluating the system's behavior under changing conditions.

The authors of [7] compare the performance of the state-of-the-art data-driven controllers when the operating conditions differ from the training dataset of the controllers' design stage. The importance of the local features selection is highlighted concluding that voltage magnitudes comprise a significant local measurement that carries global information through the physics of the power flow.

Reference [35] uses deep reinforcement learning to learn a Volt-Var control policy trained to minimize operational costs while complying with the physical operational constraints. Contrary to optimization-based approaches, this method does not require accurate data for the topology and the network and its parameters. In order to statistically guarantee safety, this work uses a constrained policy optimization algorithm which guarantees satisfaction of the operational constraints in the form of expectation. However, since the learning of the controller is done offline using historical data, unseen real-time conditions impose the risk of violations. Furthermore, using realized historical data the method does not consider low frequency but safety-critical events. Lastly, since voltage violations are only taken into account as soft constraints in the objective function, there are no guarantees for the voltage to stay within acceptable bounds.

System dispatch and optimal power flow

The work in [36] proposes a multi-agent framework to restore power systems after a loss in generation. The agents problem is solved with a Q-learning algorithm to determine switching to energize or de-

T&G Planning	[30,31]	energize load in a network. The framework proposed uses the ad- vantages of centralized and decentralized architectures to achieve ac- curate decisions and fast responses when potential failures are detected.

Simulations show how the proposed framework performs better than

the traditional centralized and decentralized approaches. However, as the authors discuss, due to the usage of a reward function that only values the amount of load energized, physical constraints are not taken into account. This can result in violations of voltage, frequency, power flow constraints, etc. Therefore, safety in the system is not guaranteed.

While presented above as solutions to voltage regulation, the work in [7,9,10] is at heart a methodology to decentralize general OPF problems. The work in [10] considers training DERs to learn how to shape their nodal behaviors based on local information to collectively be independent from power flow coming from the substation. It also covers a case where voltages are balanced across three phases. The general ability to mimic OPF problems in a learning-based and decentralized fashion is quite powerful, but more study is needed to understand the exact breakdown scenarios of the method. The advantage is that all local controllers are open-loop and can be easily simulated and characterized for general safety analyses.

In [37] the authors propose a data-driven framework that uses limited information on forecast error distribution in order to calculate stochastic optimal power flow. The objective of the power flow controller is to minimize a function of operational costs and conditional value-at-risk of power systems network constraint violations. The control actions are power injections as well as power reserves to react to forecast errors from renewable energy sources (RES). The distribution errors are not known. The error information is only obtained through a finite training set. Using this, the authors propose a distributionally robust power flow optimization to determine power injections and reserve schedules that are robust to sampling errors from the dataset. This work is mindful of safety concerns by addressing through their method the issue of distributional shift in the data used for training.

Demand response

Demand response (DR) is also an important field in the literature that explores the use of ML/AI techniques in power systems. Different algorithms are proposed to provide several grid services. In [14,15] online convex optimization is used to track a setpoint with uncertain and flexible loads in demand response programs. Setpoint tracking has been studied in the past by posing it as a model predictive control (MPC) problem. The MPC formulation relies on precise load modeling and observations or state estimation. The main benefits of using ML/AI in this context are that load modeling, communication requirements, state estimation, and perfect information of the setpoint signal are not entirely necessary to compute scheduling decisions.

The work in [16] uses online learning (OL) in a multi-armed bandit framework to provide load shedding while learning load parameters. Similarly, [17] uses OL in a bandit framework to provide load shedding services while considering load constraints. Reward manipulation is a critical safety issue here, since both papers assume fixed curtailment parameters that can be potentially manipulated under strategic behavior. In [18], an OL algorithm is used to select thermostatically controlled loads (TCLs) to provide load shifting services to flatten the load. The model assumes that the TCLs are always available to schedule, this may not hold in scenarios in which the appliance is disconnected from the service.

In [19], RL and a deep neural network (DNN) are used to propose an incentive-based DR algorithm. A DNN is used to forecast electricity prices and load patterns, and a Q-learning algorithm is used by a service provider to compute incentive rates to consumers that promotes load reduction. The paper provides guidelines on how the proposed framework could be implemented in practice. However, no specifications are provided on how to set-up learning hyper-parameters in a practical scenario, and no discussion is included on the impacts of them in system performance.

Despite the many advantages of using ML/AI in DR settings, there are certain safety issues that arise from these papers. First, physical constraints of a network are not modeled. This can lead to violations of

power flow, voltage limits, and potential negative side effects. Second, as mentioned early, no methods are presented for the selection of some hyper-parameters used in the different ML algorithms. Finally, dispatch or reward manipulation and strategic behavior of participants can in- validate important assumptions that ensure theoretical results and ex- pected behavior. Thus, performance can vary widely depending on these effects, posing a challenge on the implementation of such algo-rithms in real power systems.

3. Review of safety risks in AI

The rapid adoption and ubiquitous experimentation with machine learning has led to concerns about safety [38,39]. Inspired by these taxonomies, we focus on safety risks of using ML in power systems. The above taxonomical work considers safety issues emerging in the standard setting of designing an intelligent *agent* within an *environment*. While it provides a good starting point, power systems issues tend to be of higher complexity than typically assumed in learning theory. First, many control design issues in power systems are more complex with multiple controlled nodes, requiring a multi-agent approach [9]. And second, the *environment* that agents need to learn about entail extensive existing *legacy* control systems, infrastructure and practices, having various physical, digital and social layers [40]. As such, for ML schemes to be effective and not defy any existing safety mechanisms, it is critical to take legacy seriously in making assumptions and design considera- tions, and frame learning problems as challenges of *integration* rather than "deployment" or "automation" [41].

Building on [38,39], we characterize the safety concerns that are most relevant for learning-based control in power systems. These problems are not necessarily mutually exclusive, but collectively describe the challenges that the authors believe deserve more attention in power systems and control research.

Avoiding negative side effects

An ML model optimizes its actions according to an objective func-tion, with possible constraints, that may not be able to capture all the behaviors to keep a system safe. What potential side effects can we expect, and can we account for these in formulating the learning problem?

Persistent excitation and safe exploration

ML models trained in an offline fashion rely on training data that *represents* the conditions in which the resulting models are used. What data is needed to ensure learned parameters result in safe behavior in practice? How do we make sure that scenarios that are safety-critical, but do not occur (often) in historical data are addressed effectively in training? Online ML approaches require randomized exploration to understand how to behave "optimally", which can lead a system into unsafe territory. What safety mechanisms are available to prevent unsafe exploration scenarios?

Robustness to distributional shift

The environment is subject to inevitable change, especially in many power systems where new appliances connect and disconnect. How do we make sure a machine learning model recognizes or is robust to such changes? And how do we ensure the model's own control actions do not cause detrimental distributional shift?

Safe Integration in legacy systems and practices

Most power systems rely on human operators to take control actions or intervene in contingency scenarios. How do we ensure a human operator can safely override or complement the actions of the learning-

based control scheme? And how do new learning-based controllers interact and complement existing legacy control mechanisms?

4. Case studies

In this section, we use two case studies to exemplify the four safety problems covered in Section 3. While all four problems are relevant in both case studies, space limitations constrain us to discuss each category for one of the case studies. We also focus our attention on supervised learning models that are trained in an offline fashion. This is motivated by the fact that such models, as compared to online/reinforcement learning models, have more structure and can be more readily designed to take safety considerations into account. Section 5 will provide some pointers for readers interested in safety concerns for online approaches.

Controller design for frequency regulation

Motivation

With the increasing penetration of non-synchronous variable RES in power grids, the system's inertia decreases and varies over time, affecting the efficacy of current control schemes to handle frequency regulation. Introducing time-varying inertia parameters complexifies the design of frequency controllers –now the controller has to perform well and be stable across all inertia parameter values. In order to prevent excessive tuning and analysis and produce a stabilizing control scheme, [6] proposes a learning approach to determine a *fixed* controller based on a dataset consisting of the input values produced by optimal linear-quadratic regulator (LQR) controllers designed across *different parameter settings*. The approach uses a least squares regression in which a linear feedback matrix is determined to best fit the LQR input values across all scenarios in the dataset. Mathematically, the approach learns a time-invariant controller of the form $u_t = K_L x_t$ where K_L is a constant matrix. The training dataset $(x^{(k)}, u^{(k)})$ represents optimal solutions for $k = 1, \ldots, K$ scenarios to the LQR problem,

$$\begin{aligned} \min_{\boldsymbol{x},\boldsymbol{u}} \quad & \sum_{t=0}^{T} x_t^\top Q x_t + u_t^\top R u_t , \\ \text{s.t.} \quad & x_0 = x^{(0)} , \\ & x_{t+1} = A_{q(t)} x_t + B_{q(t)} u_t, \quad t \in \{0, T-1\} , \end{aligned} \tag{1}$$

where Q is a positive semidefinite matrix, R is a positive definite matrix, T is the control time horizon, $x^{(0)}$ is the initial state, and the matrices $A_{q(t)}$ and $B_{q(t)}$ characterize the dynamical system with time-varying inertia coefficients represented by the hybrid mode $q(t)$ as described in [11]. The least squares problem reads

$$\min_{K_L} \sum_{k=1}^{K} \sum_{t=1}^{T} \| u_t^{(k)} - K_L x_t^{(k)} \|_2^2 . \tag{2}$$

Results show that the learned controller can be used to obtain a similar (satisfying) performance as the optimal LQR controllers in the different inertia modes.

4.1.2. Tackling safety risks

Following the taxonomy in Section 3, we identify the following risks in this case study.

Avoiding Negative Side Effects

When learning a controller for a dynamical system using least squares, we are purely optimizing the extent to which the controller $K_L x_t^{(k)}$ fits the optimal LQR controller values from (1). This exercise allows for some of the values to be poor fits, at the expense of improving others. As a result, the learned controller does not inherit the guarantee that an individual LQR controller has in terms of its closed loop stability. Resulting in a learned controller that can be unstable for instances, the learned controller may not be fast enough to compensate the rate of change of the frequency in the event of contingencies.

To prevent instability, this case study *complements* the learned controller with a term that depends on the derivative of the frequency, $K_V \dot{\omega}$, which represents a virtual inertia resource for the system, yielding

$$u = K_L (\theta^\top, \omega^\top)^\top + K_V \dot{\omega} . \tag{3}$$

The parameters of K_V are chosen using a heuristic based on a bisection method. K_V is assumed of the form $K_V = k_v I_{n \times n}$. Iterating over k_v, k_v is modified until the discretized closed loop system for the low inertia modes has all its eigenvalues inside the unit circle, making it stable as can be seen in Figure 1. As a result, adding virtual inertia guarantees stability for all the dynamical modes and learning can be safely used to design a fixed controller that is stable across all inertia scenarios.

The key insight of this mitigating measure, is that the ML procedure may not be able to capture all critical safety specifications explicitly, and additional analysis and control design may be needed to satisfy these.

Persistent Excitation and Safe Exploration

In order for the learned controller K_L to reflect the behavior of the LQR controllers it is relevant that the dataset sufficiently represents all the different scenarios, as well as safety-critical states. Put differently, for each scenario and LQR controller, we want to make sure its behavior is persistently excited and captured in the dataset, so that it can be integrated in the learned fixed controller.

In order to design a robust frequency controller, different safety-critical initial states x_0 of frequency deviations are simulated in the generation of the training set. These are randomly drawn from a normal distribution with zero mean and unitary variance. This allows the learned controller to be able to steer frequency back to its nominal value under more critical circumstances, as well as during typical deviations. If these safety-critical deviations would not be included in the training set, the learned controller would not be able to perform well in all scenarios. The controller would only be able to steer frequency deviation back to zero for mild cases of disturbances. Thus, it is imperative to include the worst case scenario in the training set.

Data-driven controllers for distributed energy resources

Motivation

Based on the available monitoring and communication infrastructure, DERs can be controlled via centralized, distributed or local approaches. Lately, data-driven control design methods have gained a lot of attention [7–10,12,13,42]. These methods are hybrid in the sense that the controllers are "trained" using offline *centralized* approaches, but the derived controls are *local*, and can be used when little or no communication infrastructure is available, thereby also allowing for less data sharing that may contribute to privacy needs.

In addition, these methods allow for principled and automated some of the time-varying inertia scenarios. This is more likely to happen in low inertia scenarios, such as when more solar is produced. In these

controller tuning, preventing excessive manual labor that can render distributed control methods economically unfeasible. Lastly, these approaches can help utilities to distribute infrastructure updates in time and space, thereby preventing more capital intensive updates [10].

Optimal Power Flow Formulation

The first step of the data-driven control design method is to compute optimal DER setpoints for different operating conditions through an Optimal Power Flow (OPF) procedure, under specific objectives, such as system losses minimization [12,42] or reference voltage tracking [8,10]. System safety and power quality considerations can be addressed by including appropriate constraints in the optimization problem.

Formally, the OPF problem can be represented as

$$\min_{\mathbf{u}} c(\mathbf{x}, \mathbf{u}) \quad (4a)$$

$$\text{s.t.} \quad f(\mathbf{x}, \mathbf{u}, \mathbf{y}) = 0 \forall j, t \in (, 7, \quad (4b)$$

$$h_V(\mathbf{x}, \mathbf{u}, \mathbf{y}) \leq 0 \forall\ j, t \in (, 7, \tag{4c}$$

$$h_I(\mathbf{x}, \mathbf{u}, \mathbf{y}) \leq 0 \forall\ i, t \in \mathbb{Z}, 7, \tag{4d}$$

$$h_{\text{DER}}(\mathbf{x}, \mathbf{u}, \mathbf{y}) \leq 0 \forall\ j, t \in (, 7, \tag{4e}$$

$$g_{\text{DER}}(\mathbf{x}, \mathbf{u}, \mathbf{y}) = 0 \forall\ j, t \in (, 7. \tag{4f}$$

A distribution network with a set of nodes : =1, 2, ..., N_b (denoted by index j) and a set of lines : =1, 2, ..., N_{br} (denoted by index i) is considered. In order to account for the inter-temporal constraints of several DERs, it is necessary to solve the following multi-period AC OPF over the time horizon : =1, ..., N_{hor} (with each timestep denoted by index t). u represents the control vector, e.g. the DER active and reactive power setpoints, the position of the transformer taps, etc.; x corresponds to the state vector, i.e. the bus voltage magnitudes and angles (except for the slack bus, where the angle is set to 0 degrees and the magnitude is fixed); and y defines the constant parameters vector, comprising of the network topology, physical characteristics of the grid, and the thermal and voltage constraint limits. The operator optimizes the control vector u over the objective function given by (4a), where the function c(x, u) represents the various objectives. Equation (4b) corresponds to the power flow equations enforcing active and reactive power balances at each node. Equations (4c) - (4d) correspond to power quality constraints. Finally, (4e) - (4f) refer to DER models and constraints, e.g. technical and regulatory limitations on the DER operational power factor, constraints for the flexible loads depending on the nature of the load that can be shifted, setting a minimum and maximum per unit limit for the battery state of charge, a dynamic equation that updates the energy capacity at each time step based on the battery efficiency, etc. The second step uses the obtained optimal setpoints to design local DER controls for the real-time DN operation using ML techniques.

Data-driven Control Design The real-time response of the j^{th} inverter-based DER ($j \in [1, 2, ..., N_J]$) in terms of reactive power control $q_t^{(j)}$ is derived from the N_{OPF} optimal setpoints ($t \in [1, 2, ..., N_{OPF}]$) obtained in the offline calculations, and the final rules depend only on local features. The feature matrix $\mathbf{\Phi}^{(j)} \in \mathbf{R}^{N_{OPF} \times N_K}$ contains as columns the N_K features and as rows the N_{OPF} observations of the k^{th} input measurement $\boldsymbol{\phi}_k^{(j)} \in \mathbf{R}^{N_K}$, i.e. $\mathbf{\Phi}^{(j)} = [\boldsymbol{\phi}_1^{(j)}, \boldsymbol{\phi}_2^{(j)}, \cdots, \boldsymbol{\phi}_{N_K}^{(j)}]^T$.

As base features for the reactive power control the work follows [8] and uses the net active power demand $\phi_{1,t}^{(j)} = P_{g,j,t} - P_{l,j,t}$, the reactive power demand $\phi_{2,t}^{(j)} = Q_{l,j,t}$, and the maximum reactive power capability of the inverter $\phi_{3,t}^{(j)} = Q_{g,j,t}^{max}$. Combinations of these features are also considered, i.e. $\phi_{4,t}^{(j)} = \phi_{1,t}^{(j)} \cdot \phi_{2,t}^{(j)}$ and $\phi_{5,t}^{(j)} = (\phi_{1,t}^{(j)})^2$. Finally, the feature matrix is given by $\mathbf{\Phi}_t^{(j)} = [\phi_{1,t}^{(j)}, \phi_{2,t}^{(j)}, \phi_{3,t}^{(j)}, \phi_{4,t}^{(j)}, \phi_{5,t}^{(j)}]^T$. Using the least squares method, the local model for reactive power control is derived by solving

$$\min_{\alpha} \sum_{t \in N_{OPF}} (q_t^{(j)} - \tilde{q}_t^{(j)})^2, \tag{5a}$$

$$\tilde{q}_t^{(j)} = \alpha_0^{(j)} + \sum_{k \subset K} \alpha_k^{(j)} \cdot \mathbf{\Phi}_t^{(j)}, \tag{5b}$$

where $\alpha_k^{(j)}$ are the $k+1$ regression coefficients of the j^{th} unit for the $k \subset N_K$ features. A similar model for active power curtailment can be derived.

4.2.2. Tackling safety risks

According to the taxonomy of Section 3, we will focus on the following risks:

Robustness to Distributional Shift The data-driven DER controllers may suffer from two inherent forms of distributional shift.

Policy shift - the effect of control actions The historical data used to run OPF simulations to determine optimal setpoints for all DERs typically does not involve the control actions of the DERs yet. Once the

Table 2
Partitioning of nodal variables for all bus types.

	Exogenous		Endogenous
	Controllable - u	Uncontrollable - d_n	x_n^{end}
PQ generation	p_n, q_n		V_n, δ_n
PQ load		p_n, q_n	V_n, δ_n
PV generation	p_n, V_n		q_n, δ_n
slack bus	V_0	δ_0	p_0, q_0

the measured variables away from the distribution of the training data. This phenomena, called policy shift [43], can offset the intended effects of the controller and lead to unsafe behaviors such as constraint violation and instability.

We can address this challenge in a principled way by basic control theoretic analysis. We reformulate the state variables per bus as

$$x_n := \begin{bmatrix} V_n \\ \delta_n \\ p_n \\ q_n \end{bmatrix} \in \cdot^4. \tag{6}$$

We partition the state x_n into controllable inputs u_n, uncontrollable inputs or disturbances d_n and endogenous variables x_n^{end}. This partitioning is done per bus, based on the bus type, as suggested in [44] and summarized in Table 2.

In building a learning-based controller for a DER, the ML model can have as inputs a selection of the exogenous uncontrollable variables d_n and/or the endogenous state variables x_n^{end}, i.e. we are designing a local policy u_n: $=\pi_n(d_n, x_n^{end})$ for all buses n that have a controllable DER. Policy shift will only occur to the endogenous variables, as these are dynamically coupled with the input. To prevent policy shift, [10] only relies on exogenous variables, i.e. $u_n := \pi_n(d_n)$. In effect, this yields a *feedforward* controller which is typically used to give the desired response to command signals or objectives [45], in this case representing the OPF objectives and constraints. See [46, Sections 3.1 and 5.2] for more details on this control-theoretic perspective.

Natural shift - changes in the environment The second form of distributional shift is the inevitable occurrence of changes in the environment, which entail both unanticipated fault conditions as well as the (de-)installation of new electrical equipment. The data-driven schemes may perform very well in terms of mimicking the OPF setpoints seen in the training and test datasets, but what happens when the actual conditions deviate away from the ones seen in the training stage? Could data-driven schemes endanger the security of the system?

Here two strategies may be viable. The first is to see if some of these changes can be anticipated and simulated to be added to the training set a priori. The second is to see if feedback can be added to the controller policies. With feedforward control alone, we do miss out on the ability of feedback, i.e. using endogenous state variables x_n^{end} as inputs to the learning-based controller, which can help to improve robustness and effective disturbance attenuation [45]. To include feedback from en- learning-based controllers are implemented the control actions will impact the dynamical state variables thereby shifting the distribution of

dogenous state variables requires careful analysis of stability and convergence [47]. More recent work in [7] combines feedforward and feedback, thereby allowing some policy shift in return for controllers that are more robust against natural distributional shifts.

Here, we showcase the effect of this strategy for a scenario where a new PV unit is installed at node 11 of the Cigré benchmark radial residential LV grid presented in [48]. Without retraining the local controllers, this causes overvoltages at node 19. Fig. 2 shows the evolution of the voltage magnitude difference for node 19, phase C, due to the additional PV unit at node 11. We investigate the difference of the following cases:

Method 0: The PV units are operating according to the existing gridcode [49].

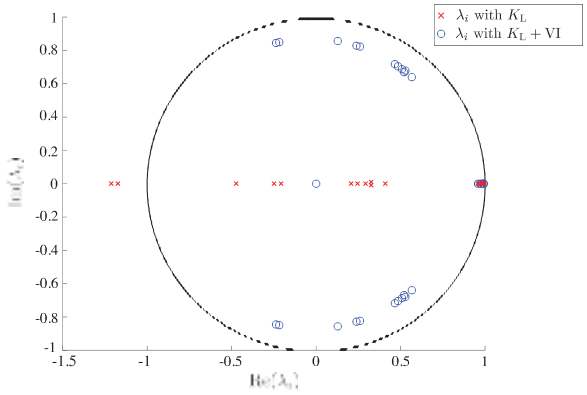

Fig. 1. Eigenvalue placement for the closed loop system in mode with lowest inertia using the learned controller K_L (crosses) and adding virtual inertia control K_L + VI (circles).

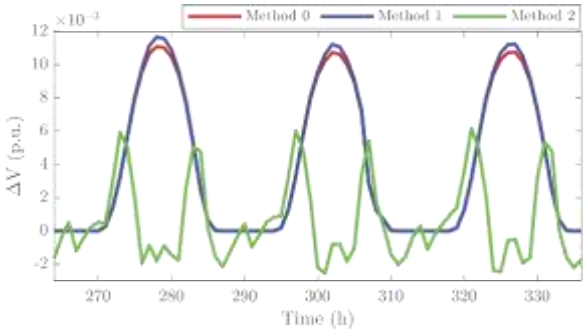

Fig. 2. Comparison of the voltage magnitude difference at node 19, phase C, due to the installation of a new PV unit at node 11.

- Method 1: The PV units are controlled based on the data-driven schemes derived before the addition of the new PV according to the method summarized in Section 4.2.1.2.
- Method 2: The same data-driven scheme is used before the addition of the new PV, but this time, the local voltage is used as an additional feature to derive the final scheme.

We observe that the data-driven scheme without the voltage magnitude feature (Method 1) results in a marginally larger impact on voltages than the current fit-and-forget approach (Method 0). In this case the difference is marginal, but highlights the difficulties in anticipating the response of data-driven controllers under distributional shifts. In order to consider such unsafe cases, one can either consider them in the offline OPF conditions, or design controllers to react to the local voltages as in Method 2.

Safe Integration in Legacy Systems and Practices

The behavior of the controllers depends on legacy control equipment and the decisions made by these in an automated fashion or by human operators, for instance for a substation's transformer equipped with on-load-tap-changers or shunt elements. In most practical contexts, these practices cannot be part of the learning-based controller design, and instead the learned controllers have to be able to collaborate and interact with these in a safe and effective manner. In this context, the operation of tap-changers does shift the distribution of voltage measurements, and hence controllers that rely on voltage are affected by it. In earlier work, these impacts of legacy equipment were an extra reason to opt for a feedforward policy that only relies on exogenous disturbance variables [8], as discussed above. As a result, the controllers are merely adjusting the nodal load profiles, without reacting to the voltage, thereby they are fully decoupled from other controllers, including other local learning-based controllers and legacy equipment. With this approach, an operator can safely change its tap

settings or even switches without a measurable response from the learned controllers, making this approach more easy to integrate in existing grid operation practices.

Avoiding negative side effects

The initial results that OPF results could be reconstructed and mimicked in a fully decentralized fashion under broad circumstances came as a surprise to control theorists, as it is generally assumed and analyzed that communication is needed to implement voltage regulation schemes that converge and satisfy voltage constraints [50]. The OPF formulation is of crucial importance in the design stage of data-driven controllers. Neglecting constraints may lead to undesired behavior and controllers that are not sensitized to safety boundaries. Put positively, the offline centralized OPF computations to build a training set that includes the effect of constraints makes it possible to sensitive learning-based controllers, even when these are all acting based on local information. This makes the use of supervised learning approaches much more valuable than online/reinforcement learning approaches, for which the integration of constraints in policy learning is still an understudied area [51].

To illustrate the importance of constraints for safety problems, consider the maximum acceptable voltage unbalance factor (VUF) [52], which, if exceeded in practice, endangers the proper functioning of specific loads. This can happen when using a single-phase representation of the distribution grid in the design stage.

Table 3 shows the maximum VUF at selected nodes for a 30-day operation using data-driven controllers that have (VUF^{max}) and have not (VUF^{max}) considered the VUF constraint in the training stage as an explicit constraint. The maximum acceptable value set in the OPF formulation is 2% and higher values could harm the equipment. We observe that node 12 shows lower values since it is electrically closer to the transformer when the three-phase voltage is regulated. On the contrary, nodes further away show larger and unacceptable values. Although both approaches experience VUF violations, only the case that considers the VUF constraint satisfies the grid code requirement to maintain the value below 2% for at least of 95% of each week. For node 19, the node with the maximum VUF values, the constraint was satisfied 96.2% of the month compared to 87.1% for the case that does not consider it.

5. Guidelines to address safety risks

Drawing together a combination of the best practices from the case studies in Section 4 and relevant literature, we present guidelines to address the most relevant safety risks. These guidelines are divided into considerations done in problem formulation and design, in empirical tests and tools and in sociotechnical aspects of power systems operations.

Explicit considerations in problem formulation and control design

Offline versus online learning

While most applications of machine learning for control or decision-making deploy some form of offline or supervised learning, in recent years the machine learning and broader artificial intelligence community has spent much energy on forms of *online or reinforcement learning* in which controller parameters are learned while operating a system.

Table 3
Maximum voltage unbalance factor using data-driven controllers over a period of a month.

	Node				
	12	16	17	18	19
^{max}	1.64	3.02	2.65	3.14	3.31
VUF^{max}	1.58	2.30	2.1	2.37	2.53

While reinforcement learning control schemes can yield sophisticated nonlinear control behaviors, these generally lack the structure and ability to shape or constrain dynamical behavior for safety (see below for some nascent work). In addition, online methods may take a long time to learn such behaviors. Without proper safety mechanisms, there are zero guarantees for the controlled system not to fail, crash or otherwise break down. As such, while potentially less expressive, offline methods may form a more welcome alternative in safety-critical situations that require guarantees and a priori safety analysis.

Exploiting the laws of physics (hybrid modeling)

If information is available about the physics of a system, it can help to exploit it by training a ML model based on simulations of such physics. The work in [8,10,12,42] exploits this to incorporate important safety constraints, that the resulting ML models end up respecting with very high probability, despite the fact that the models are implemented in a fully decentralized fashion. In RL, model-based methods allow to encode more structure into models [53], and recent work explores learning constraints [51].

Variable partitioning

Machine learning models used to inform or act as controllers, will use variables as input features. The choice of what variables to use has profound effects. Variables can be partitioned in two groups, those representing dynamic behavior *endogenous* to the system itself or, more often, of variables that act as *exogenous* disturbance variables. Using endogenous state variables comes with the issue of *policy shift* [43,54], which means that either the training data is skewed by control actions taken during collection, or the training data was collected without control actions and is skewed by the introduction of the learning-based controllers itself. In situations where data can be gathered through controlled perturbation experiments in such a way that all relevant dynamic modes can be persistently excited, the learning problem can be rephrased as a system identification [55] or adaptive control problem. For learning exogenous behavior, like weather variables or other variables that are not dynamically coupled in the system, applying machine learning is more straight-forward. Instead of persistently exciting a system, for these variables we want to make sure we include data for all possible scenarios these variable may be in.

Safe learning

In situations where online methods are needed, one may consider combining learning with a safety controller that serves as a backup in situations where the exploration process becomes risky. Recent work has considered constructing control methods that allow for *safe learning*. These methods construct an unsafe set of behaviors that the system cannot enter into, and construct a controller that switches between a learning mode and a safety control mode once the boundary of the unsafe set is hit [5,56]. Because the method is learning about the environment, the unsafe set can be updated when new information is gathered, either to further restrict the allowable behavior or to increase the range and make the system more agile.

Stability guarantees

As discussed in Section 4, using learning to develop a feedback policy may lead to an unstable system, which we addressed by doing a posteriori checks on the resulting policy to ensure it is stable. Recently, researchers developed a way to integrate explicit stability guarantees into reinforcement learning procedures [4].

Formal verification

Verifying formal properties of learning-based systems is generally considered an extremely challenging problem, due to the general inability to appropriately model the environment, to formally specify a requirement, to the challenge of modeling systems that learn (online methods), the lack of effective computational engines and lack of

methods to integrate formal verification into the design process [57]. That said, the need to build more secure and safer ML systems does motivate ongoing work in this area [58].

Empirical tests and tools

Worst case and adversarial testing

The changes that happen in power systems may form challenges to the settings of any control scheme. It is of high importance to fail test a learning-based controller for a variety of scenarios, both expected and worst case, to understand its robustness against changes and unexpected phenomena. In addition, machine learning models are vulnerable to being strategically gamed or hacked, especially as the model's input dimension increases [59]. Further, stakeholders may have an interest in the outcomes of a model and behave strategically [60]. Lastly, reinforcement learning algorithm themselves may cut corners, receiving higher rewards through unintended and potentially harmful behavior [61]. It is crucial that control schemes that rely on ML models are tested in an adversarial sense to understand what damages may occur. Such information can inform a practitioner to decide whether the approach is robust enough and what extra measures are needed to mitigate adversarial and worst-case situations.

Data generation (historical versus simulated)

In addition to historical data, changes in the environment that can be anticipated can be simulated and augmented in a training dataset. This is especially relevant in control schemes for distribution systems where new energy resources are connected at an increasing rate. By training a ML model with an appropriate proxy of the future changes, performance and safety issues may be alleviated. That said, it does not form a formal guarantee and requires analysis to understand the value and limitations for a particular application.

Calculation of risk probabilities and mutual information

Most ML models produce point predictions, that do not communicate how certain the model is. New advances in probabilistic ML try to develop models that also account for uncertainty in both model paramaters and predictions [62]. These are crucial for scientific analysis and may also inform the safety analysis for learning-based control. Related, a ML model can often also be interpreted as a reconstruction of a dataset or optimization problem. In [9], the mutual information is estimated between the outputs of an ML model and the results of an OPF problem that were used to train the ML model. This way one can determine how well the ML model reconstructs the OPF problem, providing information on whether the model should be improved with extra input features or communication between buses in the network.

Sociotechnical design considerations

Engaging with legacy systems and practices (before you design)

The guidelines so far are all broadly applicable but do not form a "silver bullet" to get your ML based control design right. In practice, every control challenge is unique and situated in a complex environment with legacy infastructure, systems and practices. The most can be learned and attained by engaging closely with this legacy to understand what kinds of assumptions can be made and how the learning problem can benefit from existing structure and domain knowledge. In the engaged work the authors have done, some general lessons have emerged that are worth mentioning. First, the increasing integration of planning and operation [42]. This is both motivated by the higher frequency of changes in the grid, but further realized through learning based controllers which can be a tool not just in operations but also in simulating behaviors in the grid over longer time horizons. Second, the importance of interpretability. Having a functioning model that has fewer input variables or parameters is generally preferred by engineers who end up having to work with the models and debug and analyse them for safety

problems. Third, the interaction with older control equipment is often inevitable. While researchers may want to champion a drastically new approach that replaces existing software or hardware, the economic reality of most utilities is that building on existing infrastructure is a given. This fundamentally shapes what is possible and, while forming a constraint on the design space, may lead to new theoretical advances [9]. Lastly, building learning-based systems includes adopting and making lots of assumptions about values that can lead to undesirable biases and impact how people interact with a system. Anticipating such value conflicts early on in the design process can help shape the possible outcomes and prevent issues down the line [63].

6. Conclusion

In this work, demand response, ideal power flow, frequency regulation, and voltage regulation were all evaluated in relation to the use of machine learning to the design of power system controllers. Relevant safety issues have been found and discussed in recent research. For two case studies that also included mitigation strategies and architectural components, these challenges were made concrete. A series of recommendations for including safety as a key component of learning-based controller design was offered based on lessons learned. These recommendations cover problem formulation, empirical testing, and interacting with the domain and its legacy infrastructure, systems, and practises.These guidelines aim to stimulate a more rigorous and thoughtful consideration of safety pro- blems in the expanding use of machine learning in power systems and other safety-critical contexts.

References

[1] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, D. Mané, Concrete Problems in AI Safety, (2016). [Online]. Available: http://arXiv.org/abs/1606.06565

[2] M. Glavic, R. Fonteneau, D. Ernst, Reinforcement learning for electric power system decision and control: Past considerations and perspectives, IFAC-PapersOnLine 50 (1) (2017) 6918–6927, https://doi.org/10.1016/j.ifacol.2017.08.1217. 20th IFAC World Congress

[3] D. Ernst, M. Glavic, L. Wehenkel, Power systems stability control: reinforcement learning framework, IEEE Trans. Power Syst. 19 (1) (2004) 427–435, https://doi.org/10.1109/TPWRS.2003.821457.

[4] F. Berkenkamp, M. Turchetta, A. Schoellig, A. Krause, Safe model-based reinforcement learning with stability guarantees, Advances in neural information processing systems, (2017), pp. 908–918.

[5] J.F. Fisac, A.K. Akametalu, M.N. Zeilinger, S. Kaynama, J. Gillula, C.J. Tomlin, A general safety framework for learning-based control in uncertain robotic systems, IEEE Trans Automat Contr 64 (7) (2019) 2737–2752, https://doi.org/10.1109/TAC.2018.2876389.

[6] P. Hidalgo-Gonzalez, R. Henriquez-Auba, D.S. Callaway, C.J. Tomlin, Frequency regulation using data-driven controllers in power grids with variable inertia due to renewable energy, IEEE PES General Meeting, IEEE, 2019.

[7] S. Karagiannopoulos, R. Dobbe, P. Aristidou, D. Callaway, G. Hug, Data-driven Control Design Schemes in Active Distribution Grids: Capabilities and Challenges, IEEE PES PowerTech, (2019).

[8] O. Sondermeijer, R. Dobbe, D. Arnold, C. Tomlin, T. Keviczky, Regression-based inverter control for decentralized optimal power flow and voltage regulation, arXiv version of paper presented at IEEE PES General Meeting 2016 arXiv:1902.08594 (2019).

[9] R. Dobbe, D. Fridovich-Keil, C. Tomlin, Fully Decentralized Policies for Multi-Agent Systems: An Information Theoretic Approach, Conference on Neural Information Processing Systems, (2017). Long Beach, CA, USA

[10] R. Dobbe, O. Sondermeijer, D. Fridovich-Keil, D. Arnold, D. Callaway, C. Tomlin, Towards distributed energy services: decentralizing optimal power flow with machine learning, IEEE Trans. Smart Grid (2019) 1, https://doi.org/10.1109/TSG.2019.2935711.

[11] P. Hidalgo-Gonzalez, D.S. Callaway, R. Dobbe, R. Henriquez-Auba, C.J. Tomlin, Frequency regulation in hybrid power dynamics with variable and low inertia due to renewable energy, 2018 IEEE Conference on Decision and Control (CDC), (2018), pp. 1592–1597.

[12] S. Karagiannopoulos, P. Aristidou, 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 (2019), https://doi.org/10.1109/TSG.2019.2905348. doi: 10.1109/TSG.2019.2905348

[13] F. Bellizio, S. Karagiannopoulos, P. Aristidou, G. Hug, Optimized local control schemes for active distribution grids using machine learning techniques, IEEE PES General Meeting, (2018).

[14] A. Lesage-Landry, J.A. Taylor, Setpoint tracking with partially observed loads, IEEE Trans. Power Syst. 33 (5) (2018) 5615–5627.

[15] A. Lesage-Landry, J.A. Taylor, Online convex optimization for demand response, Proc. Bulk Power Syst. Dynamics Control Symp. (2017), pp. 1–8.

[16] D. Kalathil, R. Rajagopal, Online learning for demand response, 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), IEEE, 2015, pp. 218–222.

[17] R. Henríquez, A. Lesage-Landry, J.A. Taylor, D. Olivares, M. Negrete-Pincetic, Managing load contract restrictions with online learning, 2017 IEEE Global Conference on Signal and Information Processing (GlobalSIP), IEEE, 2017, pp. 1035–1039.

[18] A. Lesage-Landry, J.A. Taylor, Learning to shift thermostatically controlled loads, Proceedings of the 50th Hawaii International Conference on System Sciences, (2017).

[19] R. Lu, S.H. Hong, Incentive-based demand response for smart grid with reinforcement learning and deep neural network, Appl. Energy 236 (2019) 937–949.

[20] R. Xiong, J. Cao, Q. Yu, Reinforcement learning-based real-time power management for hybrid energy storage system in the plug-in hybrid electric vehicle, Appl. Energy 211 (2018) 538–548.

[21] O. Ardakanian, Y. Yuan, V. Wong, R. Dobbe, S. Low, A. von Meier, C.J. Tomlin, On identification of distribution grids, IEEE Trans. Control Network Syst. (2019).

[22] S. Brahma, R. Kavasseri, H. Cao, N.R. Chaudhuri, T. Alexopoulos, Y. Cui, Real-time identification of dynamic events in power systems using pmu data, and potential applications models, promises, and challenges, IEEE Trans. Power Delivery 32 (1) (2016) 294–301.

[23] Y. Liao, Y. Weng, G. Liu, R. Rajagopal, Urban mv and lv distribution grid topology estimation via group lasso, IEEE Trans. Power Syst. 34 (1) (2018) 12–27.

[24] D. Deka, S. Backhaus, M. Chertkov, Estimating distribution grid topologies: a graphical learning based approach, 2016 Power Systems Computation Conference (PSCC), IEEE, 2016, pp. 1–7.

[25] D. Deka, S. Backhaus, M. Chertkov, Structure learning in power distribution networks, IEEE Trans. Control Network Syst. 5 (3) (2017) 1061–1074.

[26] Y. Gao, B. Foggo, N. Yu, A physically inspired data-driven model for electricity theft detection with smart meter data, IEEE Trans. Ind. Inf. (2019) 1, https://doi.org/10.1109/TII.2019.2898171.

[27] P. Jokar, N. Arianpoo, V.C. Leung, Electricity theft detection in ami using customers consumption patterns, IEEE Trans. Smart Grid 7 (1) (2015) 216–226.

[28] C. Duan, L. Jiang, W. Fang, J. Liu, Data-driven affinely adjustable distributionally robust unit commitment, IEEE Trans. Power Syst. 33 (2) (2017) 1385–1398.

[29] C. Zhao, Y. Guan, Data-driven stochastic unit commitment for integrating wind generation, IEEE Trans. Power Syst. 31 (4) (2015) 2587–2596.

[30] A. Bagheri, J. Wang, C. Zhao, Data-driven stochastic transmission expansion planning, IEEE Trans. Power Syst. 32 (5) (2016) 3461–3470.

[31] M. Sun, J. Cremer, G. Strbac, A novel data-driven scenario generation framework for transmission expansion planning with high renewable energy penetration, Appl. Energy 228 (2018) 546–555.

[32] Y. Chen, Y. Tan, B. Zhang, Exploiting vulnerabilities of load forecasting through adversarial attacks, Proc. of the Tenth ACM International Conference on Future Energy Systems, (2019), pp. 1–11.

[33] I.P. Panapakidis, T.A. Papadopoulos, G.C. Christoforidis, G.K. Papagiannis, Pattern recognition algorithms for electricity load curve analysis of buildings, Energy Build. 73 (2014) 137–145, https://doi.org/10.1016/j.enbuild.2014.01.002.

[34] P. Hidalgo-Gonzalez, R. Henriquez-Auba, D.S. Callaway, C.J. Tomlin, Frequency regulation using sparse learned controllers in power grids with variable inertia dueto renewable energy, (To appear) 2019 IEEE Conference on Decision and Control (CDC), (2019), pp. 1–7.

[35] W. Wang, N. Yu, J. Shi, Y. Gao, Volt-var control in power distribution systems with deep reinforcement learning, (To appear) 2019 IEEE SmartGridComm, (2019),pp. 1–7.

[36] D. Ye, M. Zhang, D. Sutanto, A hybrid multiagent framework with q-learning for power grid systems restoration, IEEE Trans. Power Syst. 26 (4) (2011) 2434–2441, https://doi.org/10.1109/TPWRS.2011.2157180.

[37] Y. Guo, K. Baker, E. Dall'Anese, Z. Hu, T. Summers, Stochastic optimal power flow based on data-driven distributionally robust optimization, 2018 Annual American Control Conference (ACC), (2018), pp. 3840–3846, https://doi.org/10.23919/ACC. 2018.8431542.

[38] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, D. Mané, Concrete problems in AI safety, arXiv preprint arXiv:1606.06565(2016).

[39] J. Leike, M. Martic, V. Krakovna, P.A. Ortega, T. Everitt, A. Lefrancq, L. Orseau, S. Legg, Ai safety gridworlds, arXiv preprint arXiv:1711.09883(2017).

[40] P.N. Edwards, S.J. Jackson, G.C. Bowker, C.P. Knobel, UnderstandingInfrastructure: Dynamics, Tensions, and Design, (2007).

[41] A. Mateescu, M.C. Elish, AI In context: the labor of integrating new technologies,Data Soc. Rep. (2019).

[42] S. Karagiannopoulos, P. Aristidou, G. Hug, Hybrid approach for planning and op-erating active distribution grids, IET Gener. Trans. Distrib. (2017) 685–695, https://doi.org/10.1049/iet-gtd.2016.0642.

[43] P. Schulam, S. Saria, Reliable Decision Support using Counterfactual Models, in:

I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems 30, Curran Associates, Inc., 2017, pp. 1697–1708. [Online]. Available: http://papers.nips.cc/paper/6767-reliable-decision-support-using-counterfactual-models.pdf

[44] A. Hauswirth, A. Zanardi, S. Bolognani, F. Drfler, G. Hug, Online optimization in closed loop on the power flow manifold, PowerTech, 2017 IEEE Manchester, IEEE, 2017.

[45] K.J. strm, R.M. Murray, Feedback systems: An Introduction for Scientists and Engineers, Princeton University Press, 2010.

[46] R.I.J. Dobbe, An Integrative Approach to Data-Driven Monitoring and Control of Electric Distribution Networks, UC Berkeley, 2018 Ph.D. thesis. [Online]. Available: https://escholarship.org/uc/item/1bw112wx

[47] M. Farivar, L. Chen, S. Low, Equilibrium and dynamics of local voltage control in distribution systems, 2013 IEEE 52nd Annual Conference on Decision and Control (CDC), (2013), pp. 4329–4334, https://doi.org/10.1109/CDC.2013.6760555.

[48] K. Strunz, E. Abbasi, C. Abbey, C. Andrieu, F. Gao, T. Gaunt, A. Gole, N. Hatziargyriou, R. Iravani, Benchmark systems for network integration of renewable and distributed energy resources, CIGRE, Task Force C6.04 (273) (2014) 4–6.

[49] VDE-AR-N 4105, Generators connected to the LV distribution network - technical requirements for the connection to and parallel operation with low-voltage distribution networks, Technical Report, FNN, 2011.

[50] G. Cavraro, S. Bolognani, R. Carli, S. Zampieri, The value of communication in the voltage regulation problem, Decision and Control (CDC), 2016 IEEE 55th Conference on, IEEE, 2016, pp. 5781–5786.

[51] J. Achiam, D. Held, A. Tamar, P. Abbeel, Constrained Policy Optimization, Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17, JMLR.org, 2017, pp. 22–31. Event-place: Sydney, NSW, Australia [Online]. Available: http://dl.acm.org/citation.cfm?id=3305381.3305384

[52] P. Pillay, M. Manyage, Definitions of Voltage Unbalance, IEEE Power Engineering Review 21 (5) (2001) 49–51, https://doi.org/10.1109/MPER.2001.4311362. [Online]. Available: http://ieeexplore.ieee.org/document/4311362/

[53] C. Xie, S. Patil, T. Moldovan, S. Levine, P. Abbeel, Model-based reinforcement learning with parametrized physical models and optimism-driven exploration, 2016 IEEE International Conference on Robotics and Automation (ICRA), (2016), pp. 504–511, https://doi.org/10.1109/ICRA.2016.7487172.

[54] S. Saria, A. Subbaswamy, Tutorial: safe and reliable machine learning, arXiv:1904.07204 [cs] (2019). ArXiv: 1904.07204.

[55] L. Ljung, System identification, Wiley Encyclop. Electr. Electron. Eng. (1999) 1–9.

[56] A.K. Akametalu, J.F. Fisac, J.H. Gillula, S. Kaynama, M.N. Zeilinger, C.J. Tomlin, Reachability-based safe learning with Gaussian processes, 53rd IEEE Conference on Decision and Control, (2014), pp. 1424–1431, https://doi.org/10.1109/CDC.2014.7039601.

[57] S.A. Seshia, D. Sadigh, S.S. Sastry, Towards verified artificial intelligence, arXiv preprint arXiv:1606.08514(2016).

[58] S.A. Seshia, A. Desai, T. Dreossi, D.J. Fremont, S. Ghosh, E. Kim, S. Shivakumar, M. Vazquez-Chanlatte, X. Yue, Formal specification for deep neural networks, International Symposium on Automated Technology for Verification and Analysis, Springer, 2018, pp. 20–34.

[59] N. Carlini, D. Wagner, Towards Evaluating the Robustness of Neural Networks, 2017 IEEE Symposium on Security and Privacy (SP), (2017), pp. 39–57, https://doi.org/10.1109/SP.2017.49.

[60] S. Milli, J. Miller, A.D. Dragan, M. Hardt, The Social Cost of Strategic Classification, Conference on Fairness, Accountability and Transparency in Sociotechnical Systems, ACM, Atlanta, 2019. ArXiv: 1808.08460 [Online]. Available: http://arxiv.org/abs/1808.08460

[61] D. Hadfield-Menell, A. Dragan, P. Abbeel, S. Russell, The Off-Switch Game, Workshops at the Thirty-First AAAI Conference on Artificial Intelligence, (2017).

[62] Z. Ghahramani, Probabilistic machine learning and artificial intelligence, Nature 521 (7553) (2015) 452–459, https://doi.org/10.1038/nature14541.

[63] R. Dobbe, S. Dean, T. Gilbert, N. Kohli, A Broader View on Bias in Automated Decision-Making: Reflecting on Epistemology and Dynamics, Workshop on Fairness, Accountability and Transparency in Machine Learning, Stockholm, (2018). [Online]. Available: https://arxiv.org/abs/1807.00553