Decision-Focused Learning under Decision Dependent Uncertainty for Power Systems with Price-Responsive Demand

Petros Ellinas
Dept. of Wind and Energy Systems
Technical University of Denmark (DTU)
Lyngby, Denmark
petrel@dtu.dk

Vassilis Kekatos
Elmore Family School of ECE
Purdue University
West Lafayette, IN 47906, USA
kekatos@purdue.edu

Georgios Tsaousoglou
Dept. of Applied Math. & Computer Science
Technical University of Denmark (DTU)
Lyngby, Denmark
geots@dtu.dk

***Abstract*—Contemporary power systems are experiencing a growing integration of flexible resources on the consumer side. Different from flexible demand that submits specific bids to energy markets, price-responsive demand (PRD) adjusts its power consumption without notice, simply based on the resulting electricity prices as well as internal priorities and limitations. In this paper, we demonstrate how a high penetration of PRD (the behavior of which is invisible to the operator during the day-ahead unit commitment stage) results in systematic inefficiency costs and formulate the so-termed decision-focused learning problem of learning to provide a demand forecast which, once fed as an input to the operator's economic dispatch optimization problem, results in an efficient dispatch. Interestingly, the prescribed demand forecast affects the resulting prices, which in turn affect the actual demand realization, giving rise to a decision-dependent uncertainty. Motivated by the problem's hard-to-evaluate objective function, we solve it using Bayesian optimization. The empirical evaluations demonstrate significant savings in the effective real-time system cost, compared to the current practice of using the default demand forecast. Moreover, the method is shown to achieve a system cost that is fairly close to the one achieved by a system that fully integrates PRD into the day-ahead process; but without requiring any change in the operator's existing dispatch algorithm while avoiding all efforts necessary for the integration of flexible demand, which is a widely pursued field of ongoing research.**

***Index Terms*—decision-focused learning, value-oriented forecasting, decision-dependent uncertainty, Bayesian optimization, economic dispatch.**

I. Introduction

The primary role of a System Operator (SO) is to ensure that load demand is met in a safe and economically optimal way. This goal is typically pursued by SOs through two main operations, one taking place in day-ahead and one in real-time. The Day-Ahead (DA) process refers to solving the Unit Commitment and Economic Dispatch problem for the next day, while the Real-Time (RT) process refers to taking real-time balancing measures to ensure supply-demand balance.

This work was partially supported by Horizon 2020 ARV project under grant agreement no. 101036723 and by US NSF under Grant 2150596.

In the DA optimization problem, demand has been traditionally treated as a known parameter. The SO receives price-quantity bids from generators and balances supply with demand, which simultaneously establishes the DA electricity prices. However, with the surge of flexible demand, this traditional paradigm is challenged. Thereupon, there has been a large body of literature modeling and quantifying demand-side flexibility [1] and designing aggregation [2], [3], and market frameworks [4], [5] for integrating it into the operator's DA process as a flexible demand that comes on top of the base (inflexible) load forecast. Such developments necessarily pertain to loads with scheduling and communication capabilities.

On the flip side, there is a second type of flexible loads, termed here as Price-Responsive Demand (PRD), that simply shape their consumption by reacting to electricity prices, and without communicating their flexibility capabilities to an aggregator, which deems them invisible during the DA process. This can result in important system inefficiencies since PRD can render the base demand's forecast inaccurate, granted it constitutes an input parameter to the DA problem. Therefore, although the operator may treat the DA dispatch as optimal, it is in fact ill-informed and can trigger costly re-dispatch actions during the RT process. Moreover, in this work, we observe that PRD can systematically counteract the DA demand forecast: In times with a higher demand forecast, electricity prices will be higher, and thus, PRD is prompted to shift loads away from those times. In turn, this results in a reduction of the actual demand, thereby counteracting the forecast. Similarly, in times with a lower demand forecast, electricity prices will be lower, prompting PRD to shift loads into those times. This increases the actual demand, again counteracting the forecast.

This observation spurs this work's endeavor to engineer the demand input of the DA task, such that the ex-post system's cost after both the DA and RT processes, is efficient. The problem is to prescribe a forecast for next-day's demand profile, not with the goal of making an accurate prediction, but of favorably biasing the DA dispatch towards making decisions that will turn out to be efficient ex-post. Deliberately prescribing a forecast with the purpose of optimizing the resulting cost has been known in machine learning as *Decision-Focused Learning* (DFL) or *value-oriented forecasting* [6], [7].

The sparse applications of related techniques in power system operations are surveyed in [8]. Notable references include [9] and [10]. SOs are also known to actively bias the forecast empirically for similar purposes [11]. However, these references consider exogenous uncertainties (renewable generation), whose realization does not depend on the prescribed forecast. In stark contrast, in the presence of PRD, demand depends on prices, which in turn, depend on the prescribed DA forecast. This peculiarity brings the problem into the realm of *endogenous* or *decision-dependent* uncertainty [12].

Stochastic optimization under decision-dependent uncertainty has been recently considered in power system operations. Reference [13] introduces a robust model for energy and reserve dispatch across multiple stages considering deferrable loads and curtailable demand. Reference [14] uses decision-dependent uncertainty to model the effect of unit commitment decisions on forced outages, and takes it into account to optimize operational reliability. A chance-constrained program for optimal battery charging with battery's bounds modeled as uncertain and dependent on past service requests and incentives is presented in [15]. Finally, [16] conceptualizes a market for RES and load predictions, where the decision on the prediction purchases affects the uncertainty set in the subsequent robust generation dispatch problem. However, the aforementioned works do not deal with decision-focused learning problems.

To the best of our knowledge, the decision-focused learning problem of prescribing a demand forecast in systems with price-responsive demand (and, thereby, with decision-dependent uncertainty), has not been considered before. Accordingly, the first contribution of this work is to:

c1) Formulate the DFL problem (Section III) under PRD. To this end, our approach models (Section II): *i)* the DA task as a classic unit commitment and economic dispatch problem with a given demand forecast; *ii)* a PRD model as an optimization problem that performs load shifting based on the DA prices (resulting as the dual variables of the DA power-balance constraints); and *iii)* the RT task as an economic dispatch that determines balancing actions with given generator commitment decisions and known demand.

Compared to optimization with decision-dependent uncertainty, the above problem bears additional complexities. The uncertainty realization (demand) depends on the dual rather than the primal variable of the economic dispatch, to which our decision (demand forecast) is only an input. Thus, evaluating the final system cost of any given DA demand forecast entails simulating the whole chain of the DA, PRD, and RT processes. This renders DFL not only non-convex but also plagued with an objective function that is computationally costly even to evaluate. These characteristics motivate using Bayesian optimization (BO) – a surrogate-model-based learning method for optimizing black-box, hard-to-evaluate functions, which gives rise to our second contribution:

c2) Solve DFL using BO (Section IV). The proposed solution is empirically evaluated on the standardized GMLC case study [17] and compared to two benchmarks (Section V): the *status quo* approach that uses standard demand forecast methods and does not consider PRD in the DA process, and the theoretically optimal approach of incorporating the (unknown in practice) PRD model into the DA process as if it was fully integrated.

II. System Model

Consider a power system where a set $\mathcal{G}$ of generators is managed across a set $\mathcal{T}$ of discrete intervals. The operational status of a generator $g \in \mathcal{G}$ at time $t \in \mathcal{T}$ is described by three binary variables $(n_{g,t}, r_{g,t}, e_{g,t})$, as in

$$n_{g,t}, r_{g,t}, e_{g,t} \in \{0,1\}, \quad g \in \mathcal{G}, t \in \mathcal{T}, \tag{1}$$

where $n_{g,t} = 1$ captures the case where the generator is ON, the start-up variable $r_{g,t} = 1$ encodes the case that the generator is turned from OFF to ON at time t, and the shut-down variable $e_{g,t}$ takes the value of 1 if the generator is turned from ON to OFF at time t. The logic connecting the three variables is enforced by

$$r_{g,t} - e_{g,t} = n_{g,t} - n_{g,t-1}, \quad g \in \mathcal{G}, t \in \mathcal{T}, \tag{2}$$

where $n_{g,0}$ denotes a predetermined initial state. Minimum up/down-time requirements are enforced as

$$n_{g,t} \geq \sum_{\tau=\max\{0,t-\overline{\mathrm{q}}_g+1\}}^{t} r_{g,\tau}, \qquad g \in \mathcal{G}, t \in \mathcal{T}. \tag{3a}$$

$$1 - n_{g,t} \geq \sum_{\tau=\max\{0,t-\underline{\mathrm{q}}_g+1\}}^{t} e_{g,\tau}, \quad g \in \mathcal{G}, t \in \mathcal{T}, \tag{3b}$$

so that generator g stays ON (OFF) for at least $\overline{\mathrm{q}}_g$ ($\underline{\mathrm{q}}_g$) times once turned ON (OFF). The continuous variable $p_{g,t}$ denotes the active power production of generator g at time t. Ramp-up and ramp-down constraints are imposed as

$$p_{g,t-1} - \mathrm{r}_g \leq p_{g,t} \leq p_{g,t-1} + \mathrm{r}_g, \quad g \in \mathcal{G}, t \in \mathcal{T}. \tag{4}$$

To capture piece-wise linear generation costs, the variable $p_{g,t}$ can be decomposed into multiple segments $\{p_{g,l,t}\}_{l \in \mathcal{L}_g}$, each one upper bounded by $\overline{\mathrm{P}}_{g,l}$, as in

$$p_{g,t} = \sum_{l \in \mathcal{L}_g} p_{g,l,t}, \qquad g \in \mathcal{G}, t \in \mathcal{T} \tag{5a}$$

$$0 \leq p_{g,l,t} \leq n_{g,t}\overline{\mathrm{P}}_{g,l}, \quad g \in \mathcal{G}, l \in \mathcal{L}_g, t \in \mathcal{T}. \tag{5b}$$

The total cost for generator g is expressed as the sum of the start-up cost $\mathrm{w}_g^{\mathtt{su}}$, fixed costs $\mathrm{w}_g^{\mathtt{fx}}$, and a piece-wise linear generation cost as

$$c_{g,t}^{\mathtt{DA}} = r_{g,t}\mathrm{w}_g^{\mathtt{su}} + n_{g,t}\mathrm{w}_g^{\mathtt{fx}} + \sum_{l \in \mathcal{L}_g} p_{g,l,t}\mathrm{h}_{g,l}\mathrm{f}_g \quad g \in \mathcal{G}, t \in \mathcal{T} \tag{6}$$

where $\mathrm{h}_{g,l}$ [Btu/MWh] is the incremental heat rate for level l (increasing for higher segments), and f_g [\$/Btu] is the generator's fuel price.

A. Day-Ahead Dispatch

The system operator (SO) is responsible for managing the generators so as to meet the system's electricity demand in an economically efficient way. The generators' binary variables and ramp constraints necessitate a day-ahead process where generators are committed to an operational state for each time slot of the day, and they are dispatched to a scheduled output. To ease the exposition, power flow constraints imposed by transmission lines have been ignored here, even though they can be readily included in the problem formulation and solution methodology. The supply-demand balance constraint for the day-ahead (DA) dispatch process is:

$$\sum_{g\in\mathcal{G}} p_{g,t} = D_t, \qquad t\in\mathcal{T} \quad : (\lambda_t) \tag{7}$$

where λ_t is the associated optimal dual variable, which instantiates the energy price at time t, and D_t is the expected net demand, i.e., a forecast of the system's electricity load minus renewable generation. Under these considerations, the DA unit-commitment and economic-dispatch problem reads

$$\begin{aligned} \min\ & \sum_{t\in T}\sum_{g\in\mathcal{G}} c^{\mathtt{DA}}_{g,t} && (\mathtt{DA})\\ \text{over}\ & \mathcal{V}^{\mathtt{DA}} := \left\{\left(n_{g,t}, r_{g,t}, e_{g,t}, p_{g,t}, \{p_{g,l,t}\}_{l\in\mathcal{L}_g}\right)\right\}_{g\in\mathcal{G},t\in\mathcal{T}}\\ \text{s.t.}\ & (1)-(7) \end{aligned}$$

Notice that the expected demand D_t is treated by the operator as a parameter of the DA optimization. Nonetheless, it is this quantity that we will be seeking to optimally learn in the DFL problem; the problem this work is after. The dual variables $\{\lambda_t\}_{t\in\mathcal{T}}$ of the balance constraints (7) instantiate the system's energy prices, based on which the PRD will adapt its consumption. This is modeled in the next subsection. Note that, technically, the dual variables are obtained by first solving the MILP problem ($\mathtt{DA}$), and then re-solving it as a linear program with the binary variables fixed.

B. Load Response to Price

This subsection presents a model of how PRD modifies consumption $\{d_t\}_{t\in\mathcal{T}}$ in response to the electricity prices $\{\lambda_t\}_{t\in\mathcal{T}}$ resulting from the ($\mathtt{DA}$) problem. Note that the model presented in this subsection is only to demonstrate the proposed methodology (to follow). The proposed method does not directly use the PRD model, but only the data resulting from it, and can work with *any* type of PRD. For simplicity, consumers are assumed to be subjected to the wholesale DA market prices directly. The consumers' response is modeled as the solution to the convex program:

$$\begin{aligned} \min\ & \sum_{t\in\mathcal{T}} \left(\lambda_t d_t + \beta s_t^2\right) && (8a)\\ \text{over}\ & \{d_t, s_t\}_{t\in\mathcal{T}} && (8b)\\ \text{s.t.}\ & \underline{\mathrm{d}}_t \le d_t \le \overline{\mathrm{d}}_t \qquad t\in\mathcal{T} && (8c)\\ & \sum_{t\in\mathcal{T}} d_t \ge \sum_{t\in\mathcal{T}} \mathrm{d}^0_t, && (8d)\\ & s_t = s_{t-1} + d_t - \mathrm{d}^0_t \qquad t\in\mathcal{T} && (8e) \end{aligned}$$

where $s_0 = 0$ and $\beta > 0$ is a given penalty factor. Parameter d^0_t denotes the spontaneous demand, that is the load's tentative demand at time t had there not been a pricing signal. The "state-of-energy" variable s_t models the load's capability to shift demand across time, e.g., by using batteries or the thermal storage capability of a building. Adjustments are confined within given per-instance limits $\{(\underline{\mathrm{d}}_t, \overline{\mathrm{d}}_t)\}_{t\in\mathcal{T}}$, which depend on consumers' flexibility. Adjustments are also constrained as a sum across time through constraint (8d). Constraint (8e) describes how demand can be time-shifted across the horizon, yet such shifting incurs an inconvenience to the owner of flexible loads modeled by the quadratic penalty βs_t^2 added to the objective function. Due to this quadratic penalty, it is easy to see that problem (8) is strictly convex over $\{d_t\}_{t\in\mathcal{T}}$, and hence, its minimizer is unique. The first term of the objective function penalizes the load variable d_t with the price acquired by the day-ahead commitment decisions. This will result in a load shift when λ_t is high. The second term of the objective function includes the flexibility constant β, which penalizes the load shift between timeslots. The demand sequence $\{d_t\}_{t\in\mathcal{T}}$ minimizing (8) will be denoted as $\{d_t(\boldsymbol{\lambda})\}_{t\in\mathcal{T}}$ where vector $\boldsymbol{\lambda}$ carries prices $\{\lambda_t\}_{t\in\mathcal{T}}$. This notation emphasizes that adjusted demands depend on prices across all times due to the capability of shifting flexible demand across time.

It is worth stressing that the response of demand to prices captured by (8) is actually unknown to the operator. This is because the flexible loads considered here do not communicate their flexibility via an aggregator or otherwise. Therefore, the SO does not take the demand-response model in (8) into account when solving the dispatch problem ($\mathtt{DA}$).

C. Real-Time Redispatch

To meet the actual real-time demand which typically deviates from the DA forecast, the dispatch decided by ($\mathtt{DA}$) undergoes a real-time adjustment decided by the so-termed real-time (RT) (re-)dispatch problem. To distinguish between the DA and the adjusted RT dispatch, we will henceforth denote RT decision variables with the superscript $(\cdot)^{\mathtt{RT}}$, while the values of DA variables are denoted with superscript $(\cdot)^{\mathtt{DA}}$. It is important to note that all DA variables are fixed and treated as parameters at the time the RT redispatch occurs.

The RT redispatch $p^{\mathtt{RT}}_{g,t}$ embodies the balancing action that comes on top of the generator's DA dispatch previously scheduled by the DA problem. Hence, the effective output $x^{\mathtt{RT}}_{g,t}$ of generator g at time t after solving the RT problem is

$$x^{\mathtt{RT}}_{g,t} = p^{\mathtt{RT}}_{g,t} + p^{\mathtt{DA}}_{g,t}. \tag{9}$$

As with DA, variable $x^{\mathtt{RT}}_{g,t}$ is decomposed into segments as

$$x^{\mathtt{RT}}_{g,t} = \sum_{l\in\mathcal{L}_g} x^{\mathtt{RT}}_{g,l,t} \tag{10a}$$

$$0 \le x^{\mathtt{RT}}_{g,l,t} \le n^{\mathtt{DA}}_{g,t}\cdot \overline{\mathrm{P}}_{g,l}, \qquad l\in\mathcal{L}_g, g\in\mathcal{G}, t\in\mathcal{T}. \tag{10b}$$

and a generator's final (effective) cost $c^{\mathtt{RT}}_{g,t}$ reads as

$$c^{\mathtt{RT}}_{g,t} = r^{\mathtt{DA}}_{g,t}\mathrm{w}^{\mathtt{su}}_g + n^{\mathtt{DA}}_{g,t}\mathrm{w}^{\mathtt{fx}}_g + \sum_{l\in\mathcal{L}_g} x^{\mathtt{RT}}_{g,l,t}\mathrm{h}_{g,l}\mathrm{f}_g, g\in\mathcal{G}, t\in\mathcal{T} \tag{11}$$

The balancing energy $p^{\text{RT}}_{g,t}$ for generator g is upper bounded by its spare capacity $\overline{\text{P}}_g - p^{\text{DA}}_{g,t}$, provided it is committed:

$$p^{\text{RT}}_{g,t} \leq n^{\text{DA}}_{g,t} \cdot (\overline{\text{P}}_g - p^{\text{DA}}_{g,t}), \qquad g \in \mathcal{G}, t \in \mathcal{T}. \tag{12}$$

Also, the balancing energy is lower bounded based on the scheduled dispatch and the generator's technical minimum as

$$p^{\text{RT}}_{g,t} \geq -n^{\text{DA}}_{g,t} \cdot (p^{\text{DA}}_{g,t} - \underline{\text{P}}_g), \qquad g \in \mathcal{G}, t \in \mathcal{T}. \tag{13}$$

Finally, the ramp constraints read as

$$x^{\text{RT}}_{g,t-1} - \text{r}_g \leq x^{\text{RT}}_{g,t} \leq x^{\text{RT}}_{g,t-1} + \text{r}_g, \quad g \in \mathcal{G}, t \in \mathcal{T}. \tag{14}$$

Given that flexible loads may have been adjusted in response to DA prices $\boldsymbol{\lambda}$, the power balance equation for the system during the RT redispatch becomes

$$d_t(\boldsymbol{\lambda}) - \sum_{g \in \mathcal{G}} x^{\text{RT}}_{g,t} = \ell^{\text{L}}_t - \ell^{\text{R}}_t, \qquad t \in \mathcal{T} \tag{15a}$$

$$\ell^{\text{L}}_t \geq 0 \quad \text{and} \quad \ell^{\text{R}}_t \geq 0, \qquad t \in \mathcal{T} \tag{15b}$$

where the introduced decision variables $\{(\ell^{\text{L}}_t, \ell^{\text{R}}_t)\}_{t \in \mathcal{T}}$ denote the lost load and renewable energy sources (RES), respectively. Recall that demand $d_t(\boldsymbol{\lambda})$ is the minimizer of (8), and is treated as a parameter when the operator solves the RT redispatch.

At time t, the RT problem aims at minimizing the system's final effective cost, plus the cost $c^{\text{RT}}_{g,t}$ of lost load (weighted by the value of lost load V^{L}) and lost RES (weighted by the value of lost RES V^{R}):

$$\begin{aligned} W_t = \min \quad & \sum_{g \in \mathcal{G}} c^{\text{RT}}_{g,t} + \text{V}^{\text{L}} \cdot \ell^{\text{L}}_t + \text{V}^{\text{R}} \cdot \ell^{\text{R}}_t \qquad (\texttt{RT}_t) \\ \text{over} \quad & \mathcal{V}^{\text{RT}}_t := \left\{ \left\{ (p^{\text{RT}}_{g,t}, x^{\text{RT}}_{g,t}, \{x^{\text{RT}}_{g,l,t}\}_{l \in \mathcal{L}_g}) \right\}_{g \in \mathcal{G}}, \ell^{\text{L}}_t, \ell^{\text{R}}_t \right\} \\ \text{s.t.} \quad & (9) - (15). \end{aligned}$$

This problem is over the set $\mathcal{V}^{\text{RT}}_t$ of RT decision variables for time t. Since the SO is only informed about the actual demand $d_t(\boldsymbol{\lambda})$ in near real-time, the RT redispatch for t is solved in near real-time, i.e., independently per time slot t, with the dispatch $(x^{\text{RT}}_{g,t-1})_{g \in \mathcal{G}}$ taken as fixed (by the previously realized generators' output).

III. Decision-Focused Learning (DFL)

The effective system's cost $\sum_{t \in \mathcal{T}} W_t$ refers to the sum of the objective values of problems $(\texttt{RT}_t)$ for all t. This is the cost of the dispatch actually realized. However, this cost depends on the DA process in two ways:

1) Through the generators' DA commitment decisions $n^{\text{DA}}_{g,t}$ that constraint the RT dispatch and also directly affect the system's effective cost through (11); and
2) Through prices $\boldsymbol{\lambda}$ decided by the DA. This dependence is more subtle as prices affect the final flexible demand $d_t(\boldsymbol{\lambda})$ decided by (8), which then appears in constraint (15) of the RT problem.

In turn, the result of the DA problem depends on the forecasted demand D_t. Thus, parameter D_t affects the effective system cost through the chain of dependencies described. Moreover, the relation between D_t and the system's effective cost features an interesting intuition:

- During times t where D_t is relatively high compared to spontaneous demand, the DA problem will yield high prices, i.e., prices that will prompt (8) to shift loads away from those times. This will result in a reduction of demand; thereby counteracting the forecast.
- During times with a lower D_t, prices will be lower, prompting (8) to shift loads into those times, resulting in a demand increase; again counteracting the forecast.

This counteraction of the forecasted demand leads to *systematic imbalances*, left to be balanced by the RT problem. However, since the RT problem is constrained by DA decisions, a systematically ill-informed DA process causes severe inefficiencies. Namely, when the DA demand forecast is lower than the RT demand, the operator commits fewer generators than necessary, only to re-dispatch them later during the RT problem at increasingly costly higher output levels. Furthermore, if the DA-RT demand difference is so high that the committed generators cannot meet the demand, the system will be forced to resort to costly load curtailments. Likewise, when the DA demand forecast is higher than the RT demand, the system commits more generators than necessary, resulting in unnecessary fixed (e.g., start-up) costs. Additionally, if the DA-RT demand difference is too high, RES curtailments might also be necessary for RT to maintain the committed generators' dispatch above their technical minimum.

Given these insights about the effect of D_t, the cardinal question is: *how to select D_t in* (7)*, such that the system's effective cost resulting through the chain of three optimization problems* $(\texttt{DA})$*,* (8)*, and* $(\texttt{RT}_t)$ *is minimized.* This question is formalized as the *Decision-Focused Learning* (DFL) problem:

$$\begin{aligned} \min \quad & f(\mathbf{x}) := \sum_{t \in \mathcal{T}} W_t \left(\mathcal{V}^{\text{RT}}_t; \mathcal{V}^{\text{DA}}, d_t(\boldsymbol{\lambda}) \right) \qquad (\texttt{DFL}) \\ \text{over} \quad & \mathbf{x} := \{D_t\}_{t \in \mathcal{T}} \\ \text{s.t.} \quad & -\delta \leq \sum_{t \in \mathcal{T}} D_t - \bar{D} \leq \delta \\ & W_t \left(\mathcal{V}^{\text{RT}}_t; \mathcal{V}^{\text{DA}}, d_t(\boldsymbol{\lambda}) \right) : \text{optimal cost of } (\texttt{RT}_t), \quad t \in \mathcal{T} \\ & (\texttt{RT}_t) \text{ is parameterized by } \mathcal{V}^{\text{DA}} \text{ and } d_t(\boldsymbol{\lambda}), \quad t \in \mathcal{T} \\ & (d_t(\boldsymbol{\lambda}))_{t \in \mathcal{T}} \text{ is the minimizer of (8)}, \\ & \text{(8) is parameterized by } \boldsymbol{\lambda}, \\ & (\mathcal{V}^{\text{DA}}, \boldsymbol{\lambda}) \text{ are the minimizers of } (\texttt{DA}), \\ & (\texttt{DA}) \text{ is parameterized by } \{D_t\}_{t \in \mathcal{T}}. \end{aligned}$$

The first constraint confines the sum $\sum_{t \in \mathcal{T}} D_t$ within distance δ from a given value $\bar{D}$. The latter can be selected as the forecast for the sum of demands during DA dispatch, while distance δ can depend on the decision maker's confidence in the prediction. The dependencies between the different processes are illustrated in the upper three boxes of Fig. 1.

Problem $(\texttt{DFL})$ is a multi-level, non-convex program over the vector $\mathbf{x}$ of demand forecasts across $t \in \mathcal{T}$. Simpler (namely, convex) versions of multi-level programs are oftentimes handled by substituting the inner problems by their

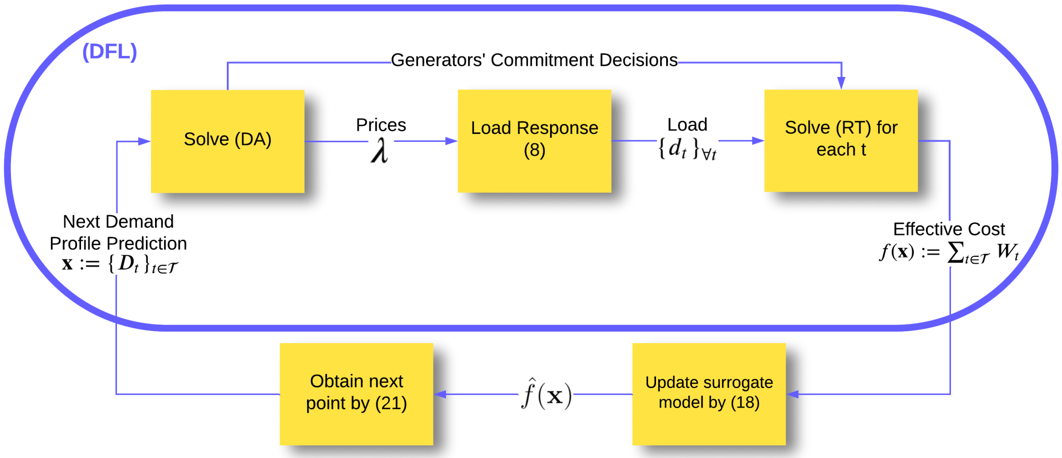

Fig. 1: Illustration of the DFL problem (circled in blue) and the proposed solution methodology.

Karush-Kuhn-Tucker (KKT) optimality conditions and introducing binary variables to model complementary slackness using the so-termed big-M trick. However, this approach is not applicable here, as the (DFL) problem is inherently more complex: the "lower-level" problem (8) depends, not on the primal, but on the dual variables of problem (DA) to which, in turn, the decision variable (i.e. the predicted demand) is only an input. To address this predicament, we resort to a zero-order optimization methodology described next.

IV. Solution Methodology

Although solving (DFL) is non-trivial, evaluating the objective value $f(\mathbf{x})$ at a feasible $\mathbf{x}$ can be accomplished at modest complexity. Bayesian optimization (BO) leverages function evaluations to build a surrogate model $\hat{f}(\mathbf{x})$ for $f(\mathbf{x})$, which is subsequently used for minimization. We next delineate how a surrogate model for (DFL)'s cost can be built and minimized.

The surrogate function $\hat{f}(\mathbf{x})$ is typically modeled as a zero-mean Gaussian process (GP) with covariance function $k(\mathbf{x},\mathbf{x}') := \mathbb{E}[f(\mathbf{x})f(\mathbf{x}')]$. A GP is a random process where any finite collection of its samples is a Gaussian random vector [18, Ch.1]. The DFL objective $f(\mathbf{x})$ can be interpreted as a random process over vector $\mathbf{x}$. If f has been evaluated already at some $\mathbf{x}_0$, then $f(\mathbf{x}_0)$ is known without uncertainty. For other $\mathbf{x}$'s, the values $f(\mathbf{x})$ could be unknown. The GP assumes that the prediction for $f(\mathbf{x})$ is correlated with $f(\mathbf{x}_0)$ depending on the distance between $\mathbf{x}$ and $\mathbf{x}_0$.

GP inference postulates a parameterized model for the covariance $k(\mathbf{x},\mathbf{x}')$ and learns the parameters of this model using the sampled function values. Suppose N samples of $\{\mathbf{x}_n, f(\mathbf{x}_n)\}_{n=1}^N$ are available. Let vector $\mathbf{f} := [f_1 \ \cdots \ f_N]^\top$ collect these samples of f. If $f(\mathbf{x})$ is modeled as a GP, then $\mathbf{f}$ is a Gaussian random vector. Without loss of generality, its mean value can be set to zero, while its covariance matrix $\boldsymbol{\Sigma}$ has entries $\boldsymbol{\Sigma}_{nm} = k(\mathbf{x}_n,\mathbf{x}_m)$ for all $n,m \in \{1,\ldots,N\}$, and $k(\mathbf{x}_n,\mathbf{x}_m)$ is a covariance function expressed in a parametric form such as the widely used Matérn kernel

$$k(\mathbf{x}_n,\mathbf{x}_m) = \alpha\left(1+\frac{\sqrt{5}r}{\beta}+\frac{5r^2}{3\beta^2}\right)e^{-\sqrt{5}r/\beta}+\gamma\delta_{nm} \quad (16)$$

where $r := \|\mathbf{x}_n - \mathbf{x}_m\|_2$ and δ_{nm} is the Kronecker delta function. The positive parameters $\{\alpha,\beta,\gamma\}$ are found using maximum likelihood estimation based on $\mathbf{f}$. Having a GP model on f enables us to make predictions on cost values for $\mathbf{x}$'s that have not been evaluated yet: Consider such an $\mathbf{x}$ and define the Gaussian vector

$$\begin{bmatrix}\mathbf{f}\\ f(\mathbf{x})\end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix}\mathbf{0}\\ 0\end{bmatrix}, \begin{bmatrix}\boldsymbol{\Sigma} & \mathbf{k}(\mathbf{x})\\ \mathbf{k}^\top(\mathbf{x}) & k(\mathbf{x},\mathbf{x})\end{bmatrix}\right). \quad (17)$$

Here the n-th entry of $\mathbf{k}(\mathbf{x})$ is $k(\mathbf{x}_n,\mathbf{x})$ for all n. Since $\mathbf{f}$ and $f(\mathbf{x})$ are jointly Gaussian, the conditional probability density function (PDF) of $f(\mathbf{x})$ given $\mathbf{f}$ is also Gaussian with

$$\hat{f}(\mathbf{x}) = \mu(\mathbf{x}) = \mathbf{k}^\top(\mathbf{x})\boldsymbol{\Sigma}^{-1}\mathbf{f} \quad (18a)$$

$$\sigma^2(\mathbf{x}) = k(\mathbf{x},\mathbf{x}) - \mathbf{k}^\top(\mathbf{x})\boldsymbol{\Sigma}^{-1}\mathbf{k}(\mathbf{x}). \quad (18b)$$

The mean can be used as an estimate of $f(\mathbf{x})$, while its variance captures the uncertainty of this estimate. Note that (17)–(18) hold for any $\mathbf{x}$. Therefore, function $\hat{f}(\mathbf{x})$ can be used as a surrogate of $f(\mathbf{x})$. Minimizing $\hat{f}(\mathbf{x})$ instead of $f(\mathbf{x})$ is easier as the former has an analytical form.

The surrogate model can be iteratively refined by evaluating f at more points. The additional points can be selected to reduce uncertainty $\sigma(\mathbf{x})$, but also focus on areas with smaller values of $\mu(\mathbf{x})$ as we are looking for the minimizer of $f(\mathbf{x})$. Let f_N^* denote the smallest cost value amongst these N evaluations, and $\mathbf{x}_N^*$ be the corresponding demands. We can improve upon f_N^* by sampling $f(\mathbf{x})$ at a new $\mathbf{x}$ featuring a potentially smaller cost. A meaningful criterion is the improvement in optimal cost captured thus far

$$\mathrm{I}(\mathbf{x}) := [f_N^* - f(\mathbf{x})]_+ \quad (19)$$

where $[z]_+ := \max\{z,0\}$. Ideally, we are looking for demands that maximize $\mathrm{I}(\mathbf{x})$. However, maximizing $\mathrm{I}(\mathbf{x})$ is as expensive as minimizing $f(\mathbf{x})$. Fortunately, thanks to the GP model, the DFL cost $f(\mathbf{x})$ in (19) can be replaced by $\hat{f}(\mathbf{x})$ and we thus maximize the expected value of the improvement

$$\text{EI}(\mathbf{x}) := \mathbb{E}\left[[f_N^* - \hat{f}(\mathbf{x})]_+\right] \tag{20}$$

where the expectation is with respect to the PDF in (18). Hence, the next point to sample is found as

$$\mathbf{x}_{N+1} = \arg\max_{\mathbf{x}\in\mathcal{X}} \ \text{EI}(\mathbf{x}). \tag{21}$$

Interestingly, function $\text{EI}(\mathbf{x})$ enjoys an analytical expression, which can be maximized using standard tools from continuous optimization Rather than the expected improvement (EI), other so-termed acquisition functions can be used to determine the demand vector to be evaluated next [19]. For the problem at hand, $\mathbf{x}$ is constrained to lie within a polytope, which allows for Bayesian optimization tools such as GPyOpt [20] to apply.

Having identified $\mathbf{x}_{N+1}$ from (21), we can evaluate $f(\mathbf{x}_{N+1})$ and either terminate the process or use the additional sample to update the surrogate model for f. A BO algorithm proceeds in three steps: *i)* A new $\mathbf{x}_{N+1}$ is identified by (21); *ii)* The cost is evaluated for $\mathbf{x}_{N+1}$, yielding $f(\mathbf{x}_{N+1})$, and the new sample is added to the training data; *iii)* The surrogate model is updated using the new datum per (18) with $\mathbf{k}(\mathbf{x})$ and $\mathbf{f}$ expanded by one entry. At each iteration, covariance parameters are updated using maximum likelihood estimation. The iterations terminate when a function evaluation budget is met. Upon completion, we select the sample $\mathbf{x}_n$ attaining the smallest cost $f(\mathbf{x}_n)$. Per Fig. 1, to sample a particular $\mathbf{x}$, problem (DA) is solved for demand vector $\mathbf{x}$ to compute DA prices, the load's reaction is then obtained via (8), and the corrective actions are taken by solving (RT_t) across all times.

V. Numerical Tests

A. Evaluation Setup

The proposed BO approach for solving (DFL) was numerically evaluated using the RTS-79 benchmark system along with generation and load parameters provided by [21]. This dataset includes a default day-ahead demand forecast $\widehat{\mathbf{x}}$ against which we benchmarked the proposed method. The proposed method was coded in Python using the libraries `GPyOpt` and `sklearn` [20], [22]. All optimization problems were solved using Gurobi. All tests were performed on an Intel i7 @ 2.2Ghz with 8GB RAM.

Unless stated otherwise, the flexibility parameter in (8a) was set to $\beta = 0.01$, the horizon length to $T = 6$ time slots, expected improvement was used as the acquisition function, results were averaged over 15 problem instances by simulating different days, the GP model was initially trained using $N = 100$ points randomly drawn from the feasible space, and it was subsequently refined by iteratively appending $M = 20$ additional function evaluations by optimizing the acquisition function online. The covariance function for the GP was selected as the Matérn-52 kernel [23], the parameters in (8c) were set as $\underline{\text{d}}_t = 0.8\text{d}_t^0$ and $\overline{\text{d}}_t = 1.2\text{d}_t^0$, and parameters V^L and V^R in (RT_t) were both set to 10000.

TABLE I: (DFL) ($T = 6$, $M = 20$) effective systems costs comparison between Proposed method, `Oracle` and `Default` benchmarks

Benchmark	Objective Value of Problem (DFL)
Default	790896.28
Oracle	601123.57
Proposed	658301.64

B. Performance, Benchmarks, and Comparison

The proposed model was compared against two alternatives:

1) `Default`: Using the default forecast $\widehat{\mathbf{x}}$ in (DA);
2) `Oracle`: Integrating price-responsive loads into (DA).

The second benchmark is only of theoretical interest as it is practically impossible to implement, since price-responsive loads respond in a spontaneous manner. The `Oracle` benchmark was simulated by solving

$$\min \sum_{t\in T}\sum_{g\in\mathcal{G}} c_{g,t} + \beta s_t^2 \tag{Oracle}$$
$$\text{over } \mathcal{V} := \left\{\{n_{g,t}, r_{g,t}, e_{g,t}, p_{g,t}, \{p_{g,l,t}\}_{l\in\mathcal{L}_g}\}_{g\in\mathcal{G}}, d_t, s_t\right\}_{t\in\mathcal{T}}$$
$$\text{s.t. } (1)-(7), (8c)-(8e)$$

Table I compares the three methods in terms of their effective system costs (i.e., the objective value of (DFL)). The reported costs empirically validate that the proposed method significantly outperforms the default forecast, while achieving a system cost of just 9.5% additional overhead compared to the theoretically optimal one. Interestingly, this suggests that instead of going through the enormous efforts required to integrate flexible demand into the DA process (cf. [2]–[5] and references therein), the proposed approach could achieve almost the same efficiency, by simply engineering a PRD-aware demand forecast.

C. Performance Analysis

This subsection presents a more thorough analysis of the benefits of the proposed approach by presenting the cost savings achieved (expressed as a percentage of the `Default` benchmark's costs) for different configurations of relevant parameters. We first tested the cost improvement for different horizon lengths T and different numbers M of function evaluations (samples). The results of Fig. 2 indicate important savings even under conditions of 24-time slot day-ahead scheduling and limited function evaluations. The savings exhibit a diminishing trend with increasing T. This does not necessarily reflect an intrinsic property of the (DFL) problem per se. It can rather be attributed to the expanding search space for Bayesian optimization that could hamper the method's efficacy. On the other hand, further training of the surrogate model (higher M) yields higher system savings.

We also tested the effect of using different acquisition functions for the BO. The selection of points to sample the cost function of interest is guided by different acquisition functions, which aim at balancing exploration and exploitation depending on the uncertainty of the GP model. For example Lower

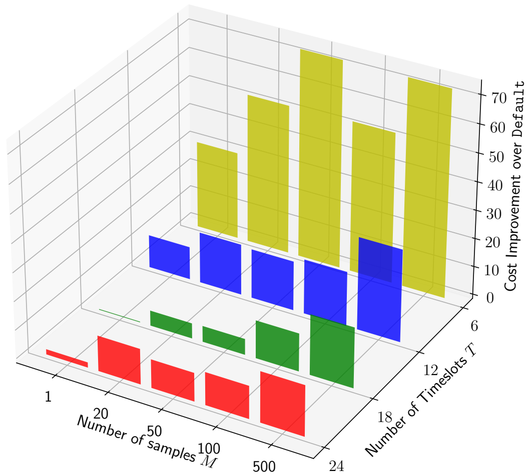

Fig. 2: Relative improvement in cost savings achieved by the Decision-Focused Learning in (DFL) over the `Default` benchmark. Problem (DFL) was solved using Bayesian optimization for varying horizon lengths T and cost function evaluations M.

TABLE II: Cost improvement and running time of BO for (DFL) ($T = 6$, $M = 20$) with different acquisition functions

Acquisition function	Cost improvement (%)	Comp. Time (s)
EI	41.9	210.3
MPI	22.8	203.1
LCB	25.6	215.6
EI MCMC	55.1	216.2
LCB MCMC	25.4	220.7
MPI MCMC	52.3	218.9

Confidence Bound (LCB) tries to strike a balance between the mean and model uncertainty; the Maximum Probability of Improvement (MPI) maximizes the chances of improvement over the present best; and the Expected Improvement (EI) seeks to maximise the expected improvement over the current best point. All three acquisition functions can be paired with Markov Chain Monte Carlo (MCMC) sampling to enable more effective and reliable exploration in high-dimensional spaces. Therefore, we have tested EI, MPI, LCB, and their MCMC variants. The results of the average improvement are shown in Table II together with the corresponding computational time. It is worth noting, that the computational time shown includes both the online evaluations $M = 20$ and the initial $N = 100$ evaluations. EI MCMC achieved the best performance.

D. Sensitivity Analysis: System's flexibility

This subsection presents a sensitivity analysis of the cost savings as a function of the system's flexibility levels. The latter is controlled via parameter β of the PRD model in (8a) with lower values of β modeling more flexible demand. The box plots illustrated in Fig. 3 show the relative cost improvement for varying β, where the boxes indicate the 25-75% range

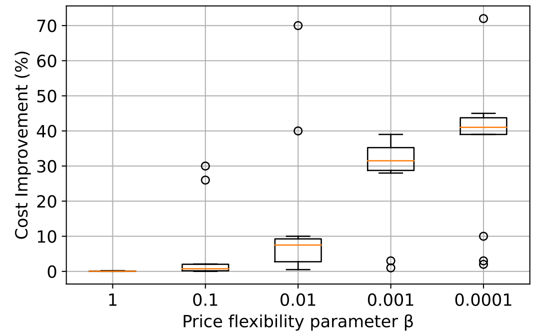

Fig. 3: Cost improvement for varying levels of price-responsiveness of loads controlled by parameter β of (8).

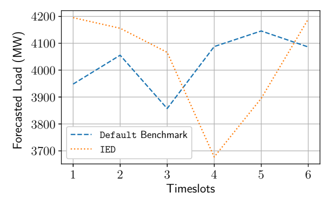

(a) Forecast day-ahead load

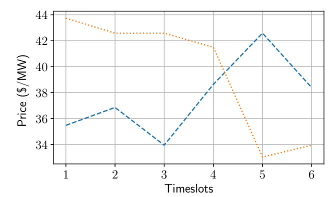

(b) Prices

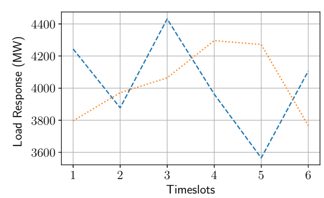

(c) Load response

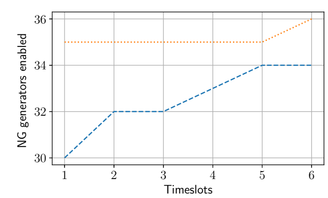

(d) Unit commitment

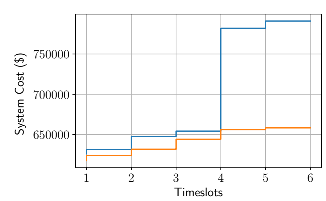

(e) Cumulative system cost

Fig. 4: Comparison between the proposed approach and the `Default` benchmark.

(for different operational days), the bars indicate the max/min values excluding outliers, and the circles indicate the outliers. As expected, the savings achieved by the Decision-Focused Learning increase as demand flexibility increases.

E. Interpretation of System's Behavior

To interpret why and how (DFL) solved via BO outperforms the `Default` benchmark, we contrasted their respective results in Fig. 4. In particular, Fig. 4a shows the predicted demand $\mathbf{x}$ found by (DFL) alongside the default forecast $\widehat{\mathbf{x}}$, and Fig. 4b shows the resulting prices upon solving (DA) using those demands. Figure 4c depicts the actual resulting demand $d_t(\boldsymbol{\lambda})$ for both cases and Fig. 4d the generators' commitment decisions. Finally, Fig. 4e presents the cumulative system cost across time.

We can observe that, for the default case, when the load spiked in time slot 2, the resulting price spiked too, resulting in a reduction of the PRD in time slot 2. On the other hand, in time slot 3, where the prediction declined sharply, the prices had a corresponding reduction. Therefore the real-time demand rose sharply. These observations offer an empirical validation of our intuition of how the presence of PRD causes the forecast to systematically counteract itself.

Interestingly, by Fig. 4d, the (`DFL`) method achieved its cost savings by committing more generators than the `Default`, and not less. The explanation is that the proposed method opted for committing more, but smaller units (with lower minimum capacity but higher costs) to handle the deviations caused by PRD, immunizing the system against variabilities. On the contrary, the `Default` benchmark committed fewer generators and ended up in increased effective costs in real-time. This is especially manifested at timeslot 4, where the difference between the forecast and actual load resulted in a severe effective cost, caused by unavoidable RES curtailments due to the reduced ability of the few available generators to ramp down quickly enough.

VI. Conclusions

Price-responsive demand (PRD) causes a systematic inefficiency as it reacts to prices without notice and, thus, its flexible behavior is invisible during the day-ahead unit commitment stage. Spurred by this inefficiency, this work has put forth the so-called *Decision-Focused Learning* (DFL) problem under decision-dependent uncertainty, which refers to prescribing a day-ahead demand forecast, that results in ex-post efficient dispatch decisions and prices, also anticipating the PRD response to the resulting prices. The proposed data-driven BO-based solution was compared favorably to the standard method of using the default forecast, demonstrating significant savings in the system's effective costs. The proposed approach is compatible with and can be directly incorporated into existing operations as it only changes the operator's demand forecast without altering the dispatch mechanism or the market procedure. Importantly, the results indicate that, despite the Operator being agnostic to the demand's response to price, a near-optimal dispatch can be obtained by simply biasing the demand forecast, all while completely bypassing the much-discussed "market-integration" of flexible demand.

References

[1] P. Pinson, H. Madsen *et al.*, "Benefits and challenges of electrical demand response: A critical review," *Renewable and Sustainable Energy Reviews*, vol. 39, pp. 686–699, 2014.

[2] S. Taheri, V. Kekatos, S. Veeramachaneni, and B. Zhang, "Data-driven modeling of aggregate flexibility under uncertain and non-convex device models," *IEEE Trans. Smart Grid*, vol. 13, no. 6, pp. 4572–4582, 2022.

[3] G. Tsaousoglou, I. Sartzetakis, P. Makris, N. Efthymiopoulos, E. Varvarigos, and N. G. Paterakis, "Flexibility aggregation of temporally coupled resources in real-time balancing markets using machine learning," *IEEE Trans. on Industrial Informatics*, vol. 18, no. 7, pp. 4342–4351, 2021.

[4] G. Tsaousoglou, J. S. Giraldo, and N. G. Paterakis, "Market mechanisms for local electricity markets: A review of models, solution concepts and algorithmic techniques," *Renewable and Sustainable Energy Reviews*, vol. 156, p. 111890, 2022.

[5] G. Tsaousoglou, R. Junker, M. Banaei, S. S. Tohidi, and H. Madsen, "Integrating distributed flexibility into TSO-DSO coordinated electricity markets," *IEEE Trans. on Energy Markets, Policy, and Regulation*, 2023, (early access).

[6] J. Mandi, J. Kotary, S. Berden, M. Mulamba, V. Bucarey, T. Guns, and F. Fioretto, "Decision-focused learning: Foundations, state of the art, benchmark and future opportunities," 2023.

[7] B. Wilder, B. Dilkina, and M. Tambe, "Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization," 2018.

[8] H. Zhang and R. Li, "Decision-oriented learning for future power system decision-making under uncertainty," 2024.

[9] J. Morales, M. Muñoz, and S. Pineda, "Prescribing net demand for two-stage electricity generation scheduling," *Operations Research Perspectives*, vol. 10, p. 100268, 2023. [Online]. Available: http://dx.doi.org/10.1016/j.orp.2023.100268

[10] T. Carriere and G. Kariniotakis, "An integrated approach for value-oriented energy forecasting and data-driven decision-making application to renewable energy trading," *IEEE Trans. Smart Grid*, vol. 10, no. 6, pp. 6933–6944, 2019.

[11] California Independent System Operator, "2019 annual report on market issues and performance," Department of Market Monitoring, Tech. Rep., 2020. [Online]. Available: http://www.caiso.com/Documents/2019AnnualReportonMarketIssuesandPerformance.pdf

[12] K. Wood and E. Dall'Anese, "Stochastic saddle point problems with decision-dependent distributions," *SIAM Journal on Optimization*, vol. 33, no. 3, pp. 1943–1967, 2023.

[13] K.-M. Aigner, J.-P. Clarner, F. Liers, and A. Martin, "Robust approximation of chance constrained dc optimal power flow under decision-dependent uncertainty," *European Journal of Operational Research*, vol. 301, no. 1, pp. 318–333, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221721009000

[14] B. Hu, C. Pan, C. Shao, K. Xie, T. Niu, C. Li, and L. Peng, "Decision-dependent uncertainty modeling in power system operational reliability evaluations," *IEEE Trans. Power Syst.*, vol. 36, no. 6, pp. 5708–5721, 2021.

[15] N. Qi, P. Pinson, M. R. Almassalkhi, L. Cheng, and Y. Zhuang, "Chance-constrained generic energy storage operations under decision-dependent uncertainty," *IEEE Trans. Sustain. Energy*, 2023.

[16] R. Xie, P. Pinson, Y. Xu, and Y. Chen, "Robust generation dispatch with purchase of renewable power and load predictions," *IEEE Trans. Sustain. Energy*, 2024.

[17] "GitHub - GridMod/RTS-GMLC: Reliability Test System - Grid Modernization Lab Consortium — github.com," https://github.com/GridMod/RTS-GMLC, [Accessed 26-Jun-2023].

[18] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006.

[19] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," *Advances in neural information processing systems*, vol. 25, 2012.

[20] "Gpyopt: A Bayesian optimization framework in Python," http://github.com/SheffieldML/GPyOpt, 2016.

[21] "GitHub - GridMod/RTS-GMLC: Reliability Test System - Grid Modernization Lab Consortium — github.com," https://github.com/GridMod/RTS-GMLC, [Accessed 22-08-2023].

[22] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[23] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 11 2005. [Online]. Available: https://doi.org/10.7551/mitpress/3206.001.0001