

STATISTICAL LEARNING FOR PROBABILITY-CONSTRAINED STOCHASTIC OPTIMAL CONTROL

ALESSANDRO BALATA*, MICHAEL LUDKOVSKI†, ADITYA MAHESHWARI†, AND JAN PALCZEWSKI*

Abstract. We investigate Monte Carlo based algorithms for solving stochastic control problems with local probabilistic constraints. Our motivation comes from microgrid management, where the controller tries to optimally dispatch a diesel generator while maintaining low probability of blackouts at each step. The key question we investigate are empirical simulation procedures for learning the state-dependent admissible control set that is specified implicitly through a probability constraint on the system state. We propose a variety of relevant statistical tools including logistic regression, Gaussian process regression, quantile regression and support vector machines, which we then incorporate into an overall Regression Monte Carlo (RMC) framework for approximate dynamic programming. Our results indicate that using logistic or Gaussian process regression to estimate the admissibility probability outperforms the other options. Our algorithms offer an efficient and reliable extension of RMC to probability-constrained control. We illustrate our findings with two case studies for the microgrid problem.

Key words. Machine learning, stochastic optimal control, probabilistic constraints, regression Monte Carlo, microgrid control

AMS subject classifications. 93E20, 93E35, 49L20

1. Introduction. Stochastic control with probabilistic constraints is a natural relaxation of deterministic restrictions which tend to generate high costs forcing the avoidance of extreme events no matter their likelihood of occurrence. In contrast, with probabilistic constraints, constraint violation is tolerated up to a certain level offering a better trade-off between admissibility and cost. We refer to [16] for an overview of probability-constrained problems and list below some of our motivating settings and references:

1. Microgrid management: An electric power microgrid is a collection of intermittent renewable generator units, a conventional dispatchable diesel generator (or grid interconnection), and a battery energy storage system. The microgrid supplies electricity to a community in islanded mode, balancing fluctuating demand and supply. The operator achieves this by optimizing the use of the battery storage and the back-up dispatchable generator. Since perfect balancing is very expensive, it is common to allow for a small frequency of blackouts, i.e. occurrences where demand outstrips supply. Mixed-integer linear programming approaches to this problem through approximating with more conservative convex constraints appear in [23, 34].
2. Hydropower optimization: control of a hydropower dam with probabilistic constraints was discussed in [2]. Within this setup, the controller observes random inflows from precipitation, as well as fluctuating electricity prices. His objective is to control the downstream outflow from the dam to maximize profit from power sales, while ensuring a minimum dam capacity with high

*School of Mathematics, University of Leeds, Woodhouse Lane, Leeds LS2 9JT, United Kingdom (alessandro.balata@live.com; J.Palczewski@leeds.ac.uk).

†Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106 USA (ludkovski@pstat.ucsb.edu; maditya0310@gmail.com).

Funding: Michael Ludkovski and Aditya Maheshwari were supported in part by NSF grant DMS-1736439. Alessandro Balata was supported by the Natural Environment Research Council Doctoral training partnership Leeds-York

Corresponding Author: Aditya Maheshwari

probability. Other related works are [3, 32, 37].

3. Motion planning: finding the minimum-cost path for a robot from one location to another while avoiding colliding with objects that obstruct its path. Stochasticity in the environment implies that the robot motion is only partially controlled. Robust optimization that guarantees obstacle avoidance might be infeasible, making probabilistic constraints a viable alternative. Dynamic programming methods for unmanned aerial vehicles were introduced in [33] and the probabilistic-constrained motion of a robot was solved in [18].

Contribution. We consider state-dependent probabilistic constraints that are expressed through an expectation constraint at each system state. Our setup involves a continuous-state, continuous-time model with a discrete-step control, where the constraints are imposed step-by-step. Therefore, at each step and at each state, the controller must estimate which controls are admissible and then optimize over the latter. While this setting is simpler than global probabilistic constraints, it is much harder than unconstrained control, because a secondary numerical procedure is needed as part of dynamic programming to repeatedly compute the admissibility sets. The canonical setup involves finite-horizon control of a stochastic process described through a stochastic differential equation of Itô type. The solution paradigm involves the Bellman or Dynamic Programming equation (DPE), which works with discretized time-steps, but with a smooth spatial variable. In this context, we develop algorithms to solve stochastic optimal control problems with probabilistic constraints using regression Monte Carlo (RMC). To make this highly nontrivial extension to RMC, we investigate tools from machine learning (including support vector machines (SVM), Gaussian process (GP) regression, parametric density estimation, logistic regression and quantile regression) to statistically estimate the admissible set as a function of the system state. Our algorithm handles the two parts of the problem—the constraint estimation and the approximation of the conditional expectation—in parallel and with significantly lower simulation budget compared to a *naive* implementation.

After benchmarking the proposed approaches on two practical case-studies from energy battery management, our main finding is to recommend logistic regression and GP-smoothed probability estimation as the best procedures. These methods are stable, relatively fast and allow for a variety of further adjustments and speed-ups. In contrast, despite theoretical appeal, quantile regression and SVM are not well-suited for this task. On a higher level, our main take-away is that DPE-based stochastic control with probabilistic constraints (SCPC) is well within reach of cutting-edge RMC methods. Thus, it is now computationally feasible to tackle such problems, opening the door for new SCPC models and applications.

Related Models for Probabilistic Constraints. There is an extensive literature on one-period optimization with chance constraints and on global multi-period probabilistic constraints. For the one-period formulations, the most popular approach is to transform the problem into linear or non-linear programs over a set of scenarios [11, 12, 28, 31]. In particular, Monte Carlo scenarios as employed below are very common, but the typical setup involves a single optimization problem, while we face an infinite family of them indexed by the system state \mathbf{x} and the time-step n . Global probabilistic constraints in multi-period settings are tackled from multiple perspectives. The dynamic programming method [30] incorporates the constraint into the objective function via a Lagrange multiplier. The solution is then obtained by iteratively solving for the optimal control and the Lagrange multiplier. However, the solution is sub-optimal due to the duality gap. Mixed-integer linear programming [1] works by linearizing the

constraints and requires to discretize the state space. An alternative [6, 7] is to transform into a static problem which is however computationally feasible only under strict assumptions on the system dynamics and noise distribution. Yet another option is the stochastic viability approach [2, 15] that focuses on maximizing the probability of being admissible, which is defined both in terms of profit targets and satisfying constraints at every time step.

Compared to above, our model of a multi-period optimization with one-step probabilistic constraints applied at each time step is new in the literature. To our knowledge, the closest setup is studied recently in [19] to compute the hedging price of a portfolio whose risk is defined in terms of its future value with respect to a set of stochastic benchmarks. Besides a local probabilistic constraint, the authors also provide dynamic programming equations for multi-period constraints. However, their solution is driven by very specific loss functions and state processes. In contrast, we develop general purpose numerical schemes using statistical learning methods.

2. Problem formulation. We study numerical resolution of stochastic control problems on finite horizon $[0, T]$ with local implicit constraints, specifically we work with constraints defined through probabilistic conditions on the controlled state. Let $(\mathbf{X}(t))_{t \geq 0} \in \mathcal{X} \subset \mathbb{R}^d$ be a continuous time controlled Markov process adapted to a given filtration (\mathcal{F}_t) . The control is an (\mathcal{F}_t) -adapted process $(u(t))_{t \geq 0}$, taking values in $\mathcal{W} \subset \mathbb{R}$. We further assume that control decisions are made at discrete epochs $\{t_0, t_1, \dots, t_N = T\}$; between time-steps the value of $u(t)$ remains constant. Thus, the control process is piecewise-constant and càdlàg (right-continuous with left limits), and will be alternatively represented as $u(t) = \sum_{n=0}^{N-1} u_n \mathbb{1}_{[t_n, t_{n+1})}(t)$.

Our setup is essentially discrete-time as far as control is concerned, but we introduce the continuous-time system state because admissibility of actions depends on the trajectory of $\mathbf{X}(t)$ between control points. In our motivating examples, the dynamics of the system is described by a stochastic differential equation:

$$d\mathbf{X}(t) = b(t, \mathbf{X}(t), u(t))dt + \sigma(t, \mathbf{X}(t), u(t))dB(t),$$

where $(B(t))$ is a m -dimensional Brownian motion generating the filtration (\mathcal{F}_t) and $b : \mathbb{R}_+ \times \mathcal{X} \times \mathcal{W} \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}_+ \times \mathcal{X} \times \mathcal{W} \rightarrow \mathbb{R}^{d \times m}$ are measurable functions, such that a unique (weak) solution exists for admissible controls defined below and takes values in \mathcal{X} . For the convenience of notation, we will write \mathbf{X}_n for $\mathbf{X}(t_n)$, $n = 0, \dots, N$, and will not indicate explicitly the dependence on the control $(u(t))$ if the latter is clear from the context.

Admissibility is defined in feedback form via

$$(2.1) \quad \mathcal{U}_{n:N}(\mathbf{X}_n) = \left\{ (u_k)_{k=n}^N : P_k(\mathbf{X}_k, u_k) \in \mathcal{A}_k \ \forall k \in \{n, \dots, N-1\} \right\},$$

for deterministic functions $P_k : \mathcal{X} \times \mathcal{W} \rightarrow \mathbb{R}$ and given subsets $\mathcal{A}_k \subset \mathbb{R}$. For the rest of the article we assume $P_n(\mathbf{X}_n, u_n)$ and \mathcal{A}_n to be of the form

$$(2.2) \quad P_n(\mathbf{X}_n, u_n) \equiv p_n(\mathbf{X}_n, u_n) := \mathbb{P}\left(\mathcal{G}_n((\mathbf{X}(s))_{s \in [t_n, t_{n+1})}) > 0 \mid \mathbf{X}_n, u_n\right) \text{ and } \mathcal{A}_n := [0, p),$$

where \mathcal{G}_n is a functional defined on trajectories of $(\mathbf{X}(s))$ over the time interval $s \in [t_n, t_{n+1})$. In other words, we target the set of controls such that the conditional probability of the “failure” functional $\mathcal{G}_n(\cdot)$ of \mathbf{X} being greater than zero is bounded by a threshold p , i.e.

$$(2.3) \quad \mathcal{U}_n(\mathbf{X}_n) := \left\{ u \in \mathcal{W} : p_n(\mathbf{X}_n, u_n) < p \right\}.$$

The threshold p in equation (2.3) is interpreted as relaxing the strong constraint $\mathcal{G}_n(\cdot) \leq 0$ which may not be appropriate in a stochastic environment. Typical values of p would generally be small ($p \approx 0.05$). We assume in the following that at least one admissible control exists at any state and at any time, hence $\mathcal{U}_{n:N}$ is determined by $\mathcal{U}_k(\mathbf{x}) = \mathcal{U}_{k:k}(\mathbf{x})$, $k = n, \dots, N-1$, the non-empty sets of admissible controls satisfying the constraints at decision epochs t_k conditional on $\mathbf{X}_k = \mathbf{x}$.

The controller must optimize as well as evaluate the feasibility of proposed actions, with the performance criterion of the form

$$(2.4) \quad V_n(\mathbf{X}_n) = \inf_{(u_s)_{s=n}^N \in \mathcal{U}_{n:N}(\mathbf{X}_n)} \left\{ \mathbb{E} \left[\sum_{k=n}^{N-1} \int_{t_k}^{t_{k+1}} \pi_s(\mathbf{X}(s), u_k) ds + W(\mathbf{X}(t_N)) \middle| \mathbf{X}_n \right] \right\},$$

where $W(\cdot)$ represents the terminal penalty and $\pi_t(\cdot, \cdot)$ the running cost. We re-write (2.4) in terms of the corresponding dynamic programming equation at step n :

$$(2.5) \quad \begin{aligned} V_n(\mathbf{X}_n) &= \inf_{u \in \mathcal{U}_n(\mathbf{X}_n)} \left\{ \mathcal{C}_n(\mathbf{X}_n, u) \right\}, \\ \text{where } \mathcal{C}_n(\mathbf{X}_n, u) &= \mathbb{E} \left[\int_{t_n}^{t_{n+1}} \pi_s(\mathbf{X}(s), u) ds + V_{n+1}(\mathbf{X}(t_{n+1})) \middle| \mathbf{X}_n, u \right]. \end{aligned}$$

Above $\mathcal{C}_n(\mathbf{X}_n, u)$ is the continuation value, i.e. reward-to-go plus expectation of future rewards, from using the control u over $[t_n, t_{n+1}]$. Moreover, given the state \mathbf{X}_n , we say that $u^* \in \mathcal{U}_n(\mathbf{X}_n)$ is an optimal control if $V_n(\mathbf{X}_n) = \mathcal{C}_n(\mathbf{X}_n, u^*)$. Since the admissible set $\mathcal{U}_n(\mathbf{X}_n)$ is both time and state dependent, we need to estimate the continuation value $\mathcal{C}_n(\cdot, \cdot)$ and the admissible control set $\mathcal{U}_n(\cdot)$ at every time step. This is the major distinction from the standard scenario approach [28] in chance-constrained optimization where there is only a single problem to optimize over a fixed \mathcal{U} , but no further indexing by x and by n . The latter require a combination interpolation and optimization as part of the solution.

Alternative Formulation of Admissibility. We denote by $G_n(\mathbf{X}_n, u_n)$ as the regular conditional distribution [20] of the functional $\mathcal{G}_n(\cdot)$ given (\mathbf{X}_n, u_n) :

$$(2.6) \quad G_n(\mathbf{X}_n, u_n) := \mathcal{L} \left(\mathcal{G}_n((\mathbf{X}(s))_{s \in [t_n, t_{n+1}]}) \middle| \mathbf{X}_n, u_n \right),$$

where $\mathcal{L}(\cdot | \mathbf{X}_n, u_n)$ stands for a conditional law. When writing $\mathbb{P}(G_n(\mathbf{X}_n, u_n) > z)$ or $\mathbb{E}[g(G_n(\mathbf{X}_n, u_n))]$ we mean the probability or the expectation with respect to this conditional distribution.

We may rewrite equation (2.2) through the corresponding $(1-p)^{th}$ quantile $q(\mathbf{X}_n, u_n)$ of $G_n(\mathbf{X}_n, u_n)$:

$$(2.7) \quad q_n(\mathbf{X}_n, u_n) : (\mathbf{X}_n, u_n) \mapsto \arg \inf_z \left\{ \mathbb{P} \left(G_n(\mathbf{X}_n, u_n) > z \right) \leq p \right\}.$$

Then using

$$(2.8) \quad \mathcal{U}_n(\mathbf{X}_n) := \{u : p_n(\mathbf{X}_n, u) < p\} = \{u : q_n(\mathbf{X}_n, u) \leq 0\},$$

we can set $P'_n := q_n$ and $\tilde{\mathcal{A}} = (-\infty, 0]$ in (2.1). We will exploit this equivalence to propose quantile-based methods (Section 4) for the admissible set.

Remark 2.1. Assuming a one dimensional control $u_n \in \mathcal{W} \subset \mathbb{R}$, and the probability $p_n(\mathbf{X}_n, u_n)$ monotonically decreasing in u_n , estimating the admissible set $\mathcal{U}_n(\mathbf{X}_n)$ is equivalent to estimating the *minimum* admissible control

$$u_n^{\min}(\mathbf{X}_n) := \inf_{u \in \mathcal{W}} \left\{ u : p_n(\mathbf{X}_n, u) < p \right\}.$$

The corresponding admissible set is $\mathcal{U}_n(\mathbf{X}_n) = \{u \in \mathcal{W} : u \geq u_n^{\min}(\mathbf{X}_n)\}$.

Remark 2.2. A more general version are implicit constraints of the form

$$\{u \in \mathcal{W} : \mathbb{E}[g(G_n(\mathbf{X}_n, u))] \leq p\},$$

for a function $g : \mathbb{R} \rightarrow \mathbb{R}$. Even more abstractly, we can think of a generic implicit map $P_k(\cdot, \cdot)$ in (2.1) that defines $\mathcal{U}_{n:N}(\mathcal{X}_n)$, with the idea that inverting this map is numerically nontrivial (i.e. P_k is expensive to evaluate) and hence *a priori* it is not clear which controls satisfy constraints and which do not.

Remark 2.3. Equation (2.8) describes admissible controls u for a given state \mathbf{x} . The “dual” perspective is to consider the set of states $\mathcal{X}_n^a(u) \subset \mathcal{X}$ for which a given control u is admissible:

$$(2.9) \quad \mathcal{X}_n^a(u) := \{\mathbf{x} \in \mathcal{X} : p_n(\mathbf{x}, u) < p\}.$$

Often the cardinality of \mathcal{X} is infinite, while the control space \mathcal{W} is finite, so that enumerating (2.9) over $u \in \mathcal{W}$ is considerably easier than enumerating the uncountable family of sets $\mathbf{x} \mapsto \mathcal{U}_n(\mathbf{x})$ in equation (2.3). Furthermore, if $u \mapsto p_n(\mathbf{x}, u)$ is decreasing for all $\mathbf{x} \in \mathcal{X}$, then we obtain an ordering $\mathcal{X}_n^a(u_1) \subseteq \mathcal{X}_n^a(u_2)$ for $u_1 \leq u_2$. The latter nesting feature corresponds to ranking the controls in terms of their “riskiness” with respect to G_n , so that the safest control will have a very large $\mathcal{X}_n^a(u)$ (possibly all of \mathcal{X}), while the riskiest control will have a very small admissibility domain.

2.1. Regression Monte Carlo. In this article we focus on simulation-based techniques to solve (2.4). The overall framework is based on solving equation (2.5) through backward induction on $n = N - 1, N - 2, \dots$, replacing the true $V_n(\mathbf{x})$ with an estimate $\hat{V}_n(\mathbf{x})$. Since neither the conditional expectation, nor the admissibility constraint are generally available explicitly, those terms must also be replaced with their estimated counterparts. As a result, we work with the approximate Dynamic Programming recursion

$$(2.10) \quad \begin{aligned} \hat{V}_n(\mathbf{X}_n) &= \inf_{u_n \in \hat{\mathcal{U}}_n(\mathbf{X}_n)} \{\hat{\mathcal{C}}_n(\mathbf{X}_n, u_n)\}, \\ \text{where } \hat{\mathcal{C}}_n(\mathbf{X}_n, u_n) &:= \hat{\mathbb{E}} \left[\int_{t_n}^{t_{n+1}} \pi_s(\mathbf{X}(s), u_n) ds + \hat{V}_{n+1}(\mathbf{X}(t_{n+1})) \middle| \mathbf{X}_n, u_n \right]. \end{aligned}$$

Above, $\hat{\mathbb{E}}$ is the approximate projection operator and the set of admissible controls $\hat{\mathcal{U}}_n$ is also approximated via either $\hat{p}_n(\cdot, \cdot)$, i.e., $\hat{\mathcal{U}}_n(\mathbf{X}_n) := \{u : \hat{p}_n(\mathbf{X}_n, u) < p\}$, or $\hat{q}_n(\cdot, \cdot)$, i.e., $\hat{\mathcal{U}}_n(\mathbf{X}_n) = \{u : \hat{q}_n(\mathbf{X}_n, u) \leq 0\}$, see (2.8). The estimated optimal control $\hat{u}_n \in \hat{\mathcal{U}}_n(\mathbf{X}_n)$ satisfies $\hat{V}_n(\mathbf{X}_n) = \hat{\mathcal{C}}_n(\mathbf{X}_n, \hat{u}_n)$.

The key idea underlying our algorithm and defining the Regression Monte Carlo paradigm is that $\hat{\mathbb{E}}$ and $\hat{\mathcal{U}}$ are implemented through empirical regressions based on Monte Carlo simulations. In other words, we construct *random*, probabilistically defined approximations based on realized paths of \mathbf{X} . This philosophy allows to simultaneously handle the numerical integration (against the stochastic shocks in \mathbf{X}) and the numerical interpolation (defining $\hat{V}_n(x)$ for arbitrary x) necessary to solve (2.10).

To understand RMC, recall that specifying $\hat{\mathbb{E}}$ is equivalent to approximating the conditional expectation map $(\mathbf{x}, u) \mapsto \mathbb{E}[\psi((\mathbf{X}(s))_{s \in [t_n, t_{n+1}]}) | \mathbf{X}_n = \mathbf{x}, u_n = u] =: f(\mathbf{x}, u)$ where we specifically substitute

$$\psi((\mathbf{X}(s))_{s \in [t_n, t_{n+1}]}) = \int_{t_n}^{t_{n+1}} \pi_s(\mathbf{X}(s), u_n) ds + \hat{V}_{n+1}(\mathbf{X}(t_{n+1})).$$

To do so, we consider a dataset consisting of inputs $(\mathbf{x}_n^1, u_n^1), \dots, (\mathbf{x}_n^{M_c}, u_n^{M_c})$ and the corresponding pathwise realizations y^1, \dots, y^{M_c} with $y^j = \psi((\mathbf{x}(s))_{s \in [t_n, t_{n+1}]}^j)$, where $(\mathbf{x}(s))_{s \in [t_n, t_{n+1}]}^j$ is an independent draw from the distribution of the process $(\mathbf{X}(s))_{s \in [t_n, t_{n+1}]} | (\mathbf{x}_n^j, u_n^j)$. Then we use the training set $\{\mathbf{x}_n^j, u_n^j, y^j\}_{j=1}^{M_c}$ to compute \hat{f} , an estimator of f , via regression.

Similarly, estimating \mathcal{U}_n is equivalent to learning the conditional probability map $p_n(\mathbf{x}, u)$ (or the conditional quantile map $q_n(\mathbf{x}, u)$ in (2.3)) and then comparing to the threshold value p (zero, respectively). This statistical task, whose marriage with RMC is our central contribution, is discussed in Section 4.

The technique of using regressions for the approximation of the continuation value was developed in the celebrated works by [24] and [36] in the context of American option pricing and further enhanced in [10, 25]. This was extended for storage problems and controlled state process in [4, 8, 9, 13, 26]. Among the approaches for approximating f we mention [39] and [26] who exploit the structure of the problem to reduce the dimensionality of the regressions, [4, 5] who harness the distribution of process to reduce the variance of \hat{f} and [21, 26, 39] who use non-parametric regression methods for \hat{f} . Regression based approach has also been discussed in the context of stochastic dual dynamic programming for solving high dimensional storage problems in [38].

In contrast to the above well-developed literature, very little exists about estimating the set of admissible controls $\hat{\mathcal{U}}_n(\mathbf{X}_n)$, which requires approximating $p(\mathbf{x}, u)$ (or $q(\mathbf{x}, u)$) in equation (2.3). There are results about learning a single global admissibility set $\hat{\mathcal{U}}$ in a one-period setup, but those approaches do not transfer to our context of state- and time-dependent admissibility constraints. A naive approach is to estimate $\hat{\mathcal{U}}_n(\mathbf{X}_n)$ for every state realized during the backward induction through nested Monte Carlo. Namely for each pair (\mathbf{x}, u) encountered, we may estimate the probability of violating the constraint by simulating M_b samples from the conditional distribution $G_n(\mathbf{x}, u)$ as $\{g_n^b(\mathbf{x}, u)\}_{b=1}^{M_b}$. We then set $u \in \hat{\mathcal{U}}_n(\mathbf{x})$ if $\bar{p}_n(\mathbf{x}, u) < p$, where

$$(2.11) \quad \bar{p}_n(\mathbf{x}, u) := \sum_{b=1}^{M_b} \frac{\mathbb{1}_{g_n^b(\mathbf{x}, u) > 0}}{M_b}$$

is the empirical probability. Although simple to implement, this Nested Monte Carlo (NMC) method is computationally intractable for even the easiest problems. As an example, a typical RMC scheme employs $M_c \approx 100,000$ and assuming $M_b = 1000$ for inner simulations, which is necessary for good estimates of small probabilities $p \leq 0.1$, would require 10^8 simulation budget at every time-step to implement NMC. Note furthermore that NMC returns only the local estimates $\bar{p}(\mathbf{x}, u)$; no *functional* estimate of $\mathcal{U}_n(\mathbf{x})$ or $\mathcal{X}_n^a(u)$ is provided for an arbitrary \mathbf{x} or u , respectively. As a result, any out-of-sample evaluation (i.e. on a future sample path of \mathbf{X}) requires further inner simulations, making this implementation even more computationally prohibitive.

An important challenge in using $\hat{\mathcal{U}}$ is *verifying* admissibility. Since we are employing Monte Carlo samples to decide whether u is admissible at \mathbf{x} , this is a probabilistic statement and admissibility can never be guaranteed 100%. We may use statistical theories to quantify the accuracy of estimators of \mathcal{U} , for example, by applying Central Limit Theorem tools for the estimator $\bar{p}(\mathbf{x}, u)$ of the true $p(\mathbf{x}, u)$. In particular, we develop tools based on confidence intervals in order to make statements (with asymptotic guarantee) such as “ $u \in \mathcal{U}$ with 95% confidence” (equivalent to $p(\mathbf{x}, u) < p$ with 95% probability conditional on the data collected). Achieving reasonable confidence levels calls for “conservative” estimators of $\hat{\mathcal{U}}$. As we show, not doing so can make learning \mathcal{U} highly unreliable, frequently causing decisions that are inadmissible with respect the imposed probability constraint. Thus, the related construction of $\hat{\mathcal{U}}^{(\rho)}$ with specified confidence level ρ is a running theme in Section 4.

2.2. Motivation: controlling blackout probability in a microgrid. To make our presentation concrete, we illustrate the framework of (2.5) by formalizing the motivating application from microgrid

management. A microgrid comprises renewable and traditional generation sources, along with a medium of storage, designed and managed to provide electrical power to a community in a decentralized way. We consider a system composed of a dispatchable diesel generator, a renewable energy source and an electric battery storage. The microgrid topology is illustrated in the left panel of Figure 1 and is same as the example discussed in [26].

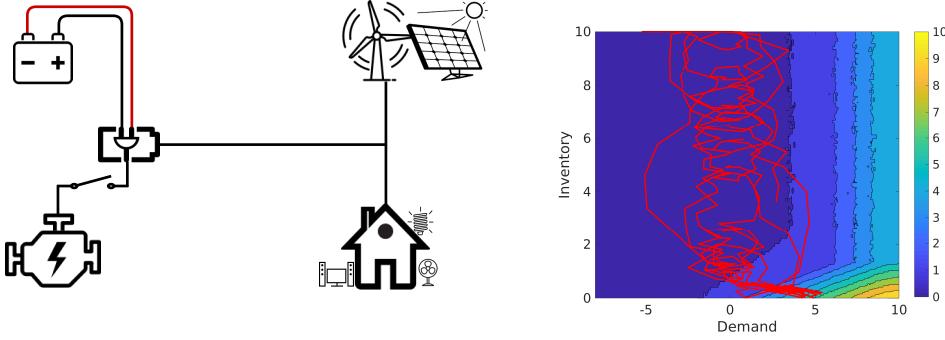


Fig. 1: Left panel: Microgrid topology: the load, the diesel generator, the battery and the renewables. Right: Contour plot for minimum admissible diesel output $(L, I, C) \mapsto u_n^{\min}(L, I, C)$ (see Remark 2.1). For $L < 0$, the constraint is not binding and $u_n^{\min}(L, I, C) = 0$. As demand increases, the constraint becomes more stringent, i.e. $u_n^{\min}(L, I, C)$ increases in L . Red curve represents a path of the controlled demand-inventory pair (L_n^u, I_n^u, C_n^u) following a myopic strategy choosing the minimum admissible control $u_n(L_n, I_n, C_n) = u_n^{\min}(L_n, I_n, C_n)$. The regime C can be visualised by observing when the red line crosses on the R.H.S. of the first contour line, indicating the the diesel generator should be turned on.

In this context, the state variables are $\mathbf{X}(t) = (L(t), I(t), C(t))$, where $L(t)$ is the net demand (demand net of renewable generation), $I(t) \in [0, I_{\max}]$ is the state of charge of the battery, referred to as “the inventory”, and $C(t) \in \{0, 1\}$ is the state of the diesel generator. $C(t) = 0$ refers to diesel being OFF and $C(t) = 1$ implies ON. The controller is in charge of the diesel through the control $u(t)$, which indicates the power output of the unit. We assume, for clarity of exposition, that the net demand $L(t)$ is an exogenous process, while $I(t)$ is controlled. We reiterate that the control decisions are made at discrete epochs $\{t_0, t_1, \dots, t_{N-1}\}$, however these decisions affect the state of the system continuously. The choice of $u(t_n) \equiv u_n$ at time t_n is based on minimizing the cost of running the microgrid, as well as controlling the probability of a blackout (i.e. failing to match the net demand) during $[t_n, t_{n+1})$. The blackout is described through the imbalance process $S(s) := L(s) - u_n - B(s)$, $\forall s \in [t_n, t_{n+1})$, representing the difference between the demand and supply, while the diesel output is held constant over the time step. The power output from the battery is a deterministic function of net demand, inventory and the control, $B(s) = \varphi(L(s), I(s), u_n)$ constrained by the physical limitations of the battery. $B(s) > 0$ implies supply of power from the battery and $B(s) < 0$ implies battery charging. The set of admissible controls is thus:

$$(2.12) \quad \mathcal{U}_n(L_n, I_n, C_n) := \left\{ u : \mathbb{P} \left(\sup_{s \in [t_n, t_{n+1})} S(s) > 0 \mid (L_n, I_n, C_n, u) \right) < p \right\}.$$

Thus in the context of microgrid, the conditional distribution G_n of equation (2.6) and the corresponding $p_n(L_n, I_n, C_n)$ are:

$$(2.13) \quad \begin{aligned} G_n(L_n, I_n, C_n, u_n) &= \mathcal{L} \left(\sup_{s \in [t_n, t_{n+1})} S(s) \mid (L_n, I_n, C_n, u_n) \right), \\ p_n(L_n, I_n, C_n, u_n) &= \mathbb{P}(G_n(L_n, I_n, C_n, u_n) > 0). \end{aligned}$$

Because p_n is not (in general) available analytically, the admissibility condition $p_n(L_n, I_n, C_n) < p$ is implicit. Recall that we denote by $\mathcal{W} = 0 \cup [\underline{u}, \bar{u}]$ the unconstrained control set. We assume that $u(t) = 0$ means that the diesel is OFF, while $u(t) > 0$ means that it is ON, and at output level $u(t)$. Thus, we define $C(s) = \mathbb{1}_{\{u_n > 0\}} \forall s \in (t_n, t_{n+1}]$ with the time interval left-open in order to allow for identification of switching on and off of the diesel generator at times t_n . Notice also that the process $C(t)$ does not satisfy the controlled diffusive dynamics, but this slight extension of the framework does not impact on the methods and results presented. We then look at the following formulation of the general problem:

$$(2.14) \quad \begin{aligned} V_n(L_n, I_n, C_n) &= \min_{\{u_k\}_{k=n}^{N-1}} \left\{ \mathbb{E} \left[\sum_{k=n}^{N-1} \left(\mathbb{1}_{\{C_k=0, u_k>0\}} \mathcal{K} + \rho(u_k) \Delta t_k \right) + W(L_N, I_N, C_N) \Big| (L_n, I_n, C_n) \right] \right\}, \\ \text{subject to} \quad &\mathbb{P} \left(\sup_{s \in [t_k, t_{k+1})} S(s) > 0 \Big| (L_k, I_k, C_k, u_k) \right) < p, \quad k = n, \dots, N-1, \end{aligned}$$

where $\Delta t_k = t_{k+1} - t_k$, $\rho(u_k)$ is the instantaneous cost of running the diesel generator with power output u_k and \mathcal{K} is the cost of switching it ON. We assume zero cost to turn the generator off. The DPE corresponding to (2.14) is the same as in (2.5) with the integral running cost $\int_{t_n}^{t_{n+1}} \pi_s(\mathbf{X}(s), u_n) ds$ replaced by

$$\mathbb{1}_{\{C_n=0, u_n>0\}} \mathcal{K} + \rho(u_n) \Delta t_n.$$

Remark 2.4. The admissible set $\mathcal{U} \subseteq \mathcal{W}$ for this problem has the special structure: if $u \in \mathcal{U}(\mathbf{x})$, then $\forall \mathcal{W} \ni \tilde{u} > u, \tilde{u} \in \mathcal{U}(\mathbf{x})$. Hence, we may represent $\mathcal{U}(\mathbf{x}) = [u_n^{\min}(\mathbf{x}), \bar{u}] \cap \mathcal{W}$ in terms of the minimal admissible diesel output $u_n^{\min}(\mathbf{x})$. Conversely, the admissibility domains for a fixed $u \in \mathcal{W}$ are nested: if $u_1 \leq u_2$ then $\mathcal{X}_n^a(u_1) \subseteq \mathcal{X}_n^a(u_2)$. This suggests to compute $\mathcal{X}_n^a(u)$ sequentially as u is increased and then invert to get $\mathcal{U}(\mathbf{x})$.

To visualize the minimum admissible control $u_n^{\min}(\mathbf{x})$, the right panel of Figure 1 presents the map $\mathbf{x} \rightarrow u_n^{\min}(\mathbf{x})$ under a constraint of $p = 0.01$ probability of blackout. We also present a path for $(L(t), I(t), C(t))_{t \geq 0}$ using a *myopic* strategy where the controller employs the minimum admissible control at each point, $u_n := u_n^{\min}(L_n, I_n, C_n) \forall n$. Notice how for the most part, $u_n^{\min}(\cdot) = 0$ so that $\mathcal{U}_n(\cdot) = \mathcal{W}$ and the blackout constraint is not binding. This is not surprising, as blackouts are only possible when $L(t) \gg 0$ is strongly positive and the battery is close to empty, $I(t) \simeq 0$. Thus, except for the lower-right corner, any control is admissible. As a result, only a small subset of the domain \mathcal{X} actually requires additional effort to estimate the admissible set $\mathcal{U}(\mathbf{x})$. In our experience this structure, where the constraint is not necessarily binding and where we mostly perform unconstrained optimization, is quite common.

3. Dynamic emulation algorithm. In this section we present our Dynamic emulation algorithm which provides approximation for the admissible set $\hat{\mathcal{U}}_n(\cdot)$ and the continuation value function $\hat{\mathcal{C}}_n(\cdot, \cdot)$. The crux of the algorithm are the following two steps, implemented in parallel at every time-step:

$$(3.1) \quad \begin{aligned} \text{Generate design} &\rightarrow \text{Generate 1-step paths \& admissibility statistic} \rightarrow \text{Estimate admissible set} \\ \text{Generate design} &\rightarrow \text{Generate 1-step paths \& pathwise profits} \rightarrow \text{Estimate continuation function} \end{aligned}$$

To estimate $\hat{\mathcal{C}}_n(\cdot, \cdot)$'s and $\hat{\mathcal{U}}_n(\cdot)$'s, we proceed iteratively backward in time starting with known terminal condition $W(\mathbf{X})$ and sequentially estimate $\hat{\mathcal{U}}_n$ and $\hat{\mathcal{C}}_n$ for $n = N-1, \dots, 0$. Assuming we have estimated $\hat{\mathcal{U}}_{n+1}, \dots, \hat{\mathcal{U}}_{N-1}$ and $\hat{\mathcal{C}}_{n+1}, \dots, \hat{\mathcal{C}}_{N-1}$, we first explain the estimation procedure for $\hat{\mathcal{U}}_n$ and $\hat{\mathcal{C}}_n$. This corresponds to a **fit** task. In the subsequent backward recursion at step $n-1$ we also need the **predict** task to actually *evaluate* $\hat{V}_n(\mathbf{X}_n)$ which requires evaluating $\hat{\mathcal{C}}_n(\cdot)$ at new (“out-of-sample”) inputs \mathbf{X}_n, u_n which of course do not coincide with the training inputs $(\mathbf{x}_n^1, u_n^1), \dots, (\mathbf{x}_n^{M_c}, u_n^{M_c})$.

3.1. Estimating the set of admissible controls. To estimate the set of admissible controls $\hat{\mathcal{U}}_n(\cdot)$ at time-step n , we choose design $\mathcal{D}_n^a := (\mathbf{x}_n^i, u_n^i, i = 1, \dots, M_a)$ and simulate trajectories of the state process $(\mathbf{X}(s))_{s \in [t_n, t_{n+1}]}^i$ starting from $\mathbf{X}^i(t_n) = \mathbf{x}_n^i$ and driven by control u_n^i . To evaluate the functional $\mathcal{G}((\mathbf{X}(s))_{s \in [t_n, t_{n+1}]}^i)$, we discretize the time interval $[t_n, t_{n+1}]$ into K finer sub-steps with $\Delta n_k := t_{n_{(k+1)}} - t_{n_k}$ and define the discrete trajectory $\mathbf{x}_n^i = \mathbf{x}_{n_0}^i, \mathbf{x}_{n_1}^i, \dots, \mathbf{x}_{n_{(K-1)}}^i, \mathbf{x}_{n_K}^i$. We then record

$$(3.2) \quad w_n^i := \mathbb{1}(\mathcal{G}((\mathbf{x}_{n_k}^i)_{k \in \{0, \dots, K-1\}}) > 0), \quad i = 1, \dots, M_a,$$

where, formally, we extend $(\mathbf{x}_{n_k}^i)_{k \in \{0, \dots, K-1\}}$ to a piecewise constant trajectory on $[t_n, t_{n+1}]$.

Analogous to standard RMC, we now select an approximation space \mathcal{H}_n^a to estimate the probability \hat{p}_n or the quantile \hat{q}_n , using the loss function \mathcal{L}_n^a and apply empirical projection:

$$(3.3) \quad \hat{p}_n := \arg \min_{f_n^a \in \mathcal{H}_n^a} \sum_{i=1}^{M_a} \mathcal{L}_n^a(f_n^a, w_n^i; \mathbf{x}_n^i, u_n^i).$$

See Section 4 for concrete examples of \mathcal{H}^a and \mathcal{L}^a . Note that the approximations \hat{p}_n and \hat{q}_n must be trained on joint state-control datasets $\{\mathbf{x}_n^i, u_n^i, w_n^i\}_{i=1}^{M_a}$ with w_n^i dependent on the method of choice and moreover yield random estimators (\hat{p}_n is a random variable).

Using the distribution of $\hat{p}_n(\mathbf{x}, u)$ we may obtain a more conservative estimator that provides better guarantees on the ultimate admissibility of (\mathbf{x}, u) . As a motivation, recall the NMC estimator $\bar{p}_n(\mathbf{x}, u)$ from (2.11); for reasonably large $M_b \gg 20$, the distribution of $\bar{p}_n(\mathbf{x}, u)$ is approximately Gaussian with mean $p_n(\mathbf{x}, u)$ and variance $\frac{p_n(\mathbf{x}, u)(1-p_n(\mathbf{x}, u))}{M_b}$. Defining

$$(3.4) \quad \hat{p}_n^{(\rho)}(\mathbf{x}, u) := \bar{p}_n(\mathbf{x}, u) + \xi_n^{(\rho)}(\mathbf{x}, u)$$

$$(3.5) \quad := \bar{p}_n(\mathbf{x}, u) + z_\rho \sqrt{\frac{\bar{p}_n(\mathbf{x}, u)(1-\bar{p}_n(\mathbf{x}, u))}{M_b}},$$

where z_ρ is the standard normal quantile at level ρ and $\xi_n^{(\rho)}(\mathbf{x}, u)$ represents a “safe” margin of error for \bar{p}_n at confidence level ρ . The corresponding approximate admissible set with confidence ρ is

$$(3.6) \quad \hat{\mathcal{U}}_n^{(\rho)}(\mathbf{x}) := \hat{\mathcal{U}}_n^{\xi^{(\rho)}}(\mathbf{x}) = \left\{ u : \hat{p}_n(\mathbf{x}, u) + \xi_n^{(\rho)}(\mathbf{x}, u) < p \right\}.$$

More generally, we set the admissible set for a site $\mathbf{x} \in \mathcal{X}$ to

$$(3.7) \quad \hat{\mathcal{U}}_n^\xi(\mathbf{x}) = \left\{ u : \hat{p}_n(\mathbf{x}, u) + \xi_n(\mathbf{x}, u) < p \right\},$$

where $\xi_n(\mathbf{x}, u)$ ensures “stronger” guarantee for the admissibility of u at \mathbf{x} . The margin of estimation error can also be fixed, $\xi_n(\mathbf{x}, u) = c \forall (\mathbf{x}, u) \in \mathcal{X} \times \mathcal{W}$, which can be applied when the sampling distribution of $\hat{p}_n(\mathbf{x}, u)$ is unknown. The corresponding admissible set

$$(3.8) \quad \hat{\mathcal{U}}_n^{\xi=c}(\mathbf{x}) = \left\{ u : \hat{p}_n(\mathbf{x}, u) + c < p \right\}.$$

is equivalent to estimating $\hat{\mathcal{U}}_n^{\xi=0}(\mathbf{x})$ with a shifted lower probability threshold $p - c$. To simplify notation, we use $\hat{\mathcal{U}}_n(\mathbf{x})$ to denote the unadjusted admissible set, $\hat{\mathcal{U}}_n(\mathbf{x}) := \hat{\mathcal{U}}_n^{\xi=0}(\mathbf{x})$ in the context of NMC. Analogously, we can adjust equations (3.2)-(3.3) based on learning the quantile $q_n(\mathbf{x}, u)$ to add a margin of error, $\hat{\mathcal{U}}_n^\xi(\mathbf{x}) = \{u : \hat{q}_n(\mathbf{x}, u) + \xi_n(\mathbf{x}, u) \leq 0\}$.

Remark 3.1. In the context of one-step static optimization, conservative estimates for the admissible set were explored in [12]. The idea is to consider u to be admissible iff $\bar{p}(\mathbf{x}, u) = 0$, i.e. there are no failures observed in all of the simulated samples. The choice of the sample size M_b is then determined using Chebyshev inequality that provides a theoretical guarantee on $\mathbb{P}(p(\mathbf{x}, u) > 0 | \bar{p}(\mathbf{x}, u; M_b) = 0)$. This approach is similar in our setup to a violation margin $c = p$ in (3.8). However, the guarantee on $p(\mathbf{x}, u)$ is only applicable locally at the design sites. We are not aware of any tools that would allow non-parametric guarantee for *any* $\mathbf{x} \in \mathcal{X}$. The regression based approach offers a model-based guarantee for admissibility of u at any \mathbf{x} by setting the parameter ρ in (3.6) or c in (3.8).

3.2. Estimating the continuation value. To estimate the continuation value $\mathcal{C}_n(\cdot, \cdot)$, we choose a simulation design $\mathcal{D}_n^c := (\mathbf{x}_n^j, u_n^j, j = 1 \dots, M_c)$ (which could be independent or equivalent to \mathcal{D}_n^a) and generate one-step paths for the state process $(\mathbf{X}(s))_{s \in [t_n, t_{n+1})}^j$ starting from $\mathbf{X}^j(t_n) = \mathbf{x}_n^j$ and driven by control u_n^j , comprising again finer sub-steps $\mathbf{x}_n^j = \mathbf{x}_{n_0}^j, \mathbf{x}_{n_1}^j, \dots, \mathbf{x}_{n_{K-1}}^j, \mathbf{x}_{n_K}^j$ (in principle the sub-steps could differ from the time discretization for $\hat{\mathcal{U}}_n$). Next, we compute the pathwise cost y_n^j :

$$(3.9) \quad y_n^j = \sum_{k=0}^{K-1} \pi_{n_k}(\mathbf{x}_{n_k}^j, u_n^j) \Delta n_k + v_{n+1}^j, \quad \text{where } v_{n+1}^j = \inf_{u \in \hat{\mathcal{U}}_{n+1}(\mathbf{x}_{n_K}^j)} \hat{\mathcal{C}}_{n+1}(\mathbf{x}_{n_K}^j, u), \quad j = 1 \dots M_c,$$

and we replace the time integral in (2.5) with a discrete sum over t_{n_k} 's. At the key step, we project $\{y_n^j\}_{j=1}^{M_c}$ onto an approximation space \mathcal{H}_n^c to evaluate the continuation value $\mathcal{C}_n(\cdot, \cdot)$:

$$(3.10) \quad \hat{\mathcal{C}}_n(\cdot, \cdot) := \arg \min_{f_n^c \in \mathcal{H}_n^c} \sum_{n=1}^{M_c} |f_n^c(\mathbf{x}_n^j, u_n^j) - y_n^j|^2.$$

The design sites $\{\mathbf{x}_n^j, u_n^j\}_{j=1}^{M_c}$ could be same or different from those used for learning the admissible sets in the previous subsection. Two standard approximation spaces \mathcal{H}_n^c used in this context are: global polynomial approximation and piecewise continuous approximation.

Remark 3.2. In the microgrid example of Section 2.2 the running cost over $[n, n+1)$ is known once the control u_n is chosen. Thus it can be taken outside the conditional expectation and the data to be regressed is simply $y^j = v_{n+1}^j$.

Global polynomial approximation:. This is a classical regression framework with polynomial bases $\phi_k(\cdot, \cdot)$ and $\hat{\mathcal{C}}_n^\alpha(\mathbf{x}, u) := \sum_k \alpha_k \phi_k(\mathbf{x}, u)$. The coefficients α are fitted via

$$(3.11) \quad \hat{\alpha} := \arg \min_{\alpha} \sum_{j=1}^{M_c} \left| \sum_k \alpha_k \phi_k(\mathbf{x}_n^j, u_n^j) - y_n^j \right|^2.$$

As an illustration, for the microgrid example of Section 2.2 we construct a quadratic polynomial approximation when diesel generator is ON, $u > 0$, using 10 bases $\{1, L, I, u, L^2, I^2, u^2, LI, Iu, LI\}$ and a separate quadratic approximation with the 6 basis functions $\{1, L, I, L^2, I^2, LI, LI\}$ when diesel generator is OFF, $u = 0$. Polynomial approximation is easy to implement but typically requires many degrees of freedom (lots of ϕ 's) to properly capture the shape of \mathcal{C} and can be empirically unstable, especially if there are sharp changes in the underlying function (see for example [21, 26]).

Piecewise continuous approximation: This is a state-of-art tool in low dimensions, $d \leq 3$. The main idea is to employ polynomial regression in a single dimension and extend to the other dimensions via linear interpolation. As an example, for the microgrid with diesel generator ON, we have three dimensions (L, I, u) . We discretize inventory I as $\{I^0, I^1, \dots, I^{M_I}\}$ and control u as $\{u^1, u^2, \dots, u^{M_u}\}$ and fit independent cubic polynomials in L for each pair (I^l, u^e) with $l \in \{0, 1, \dots, M_I\}$ and $e \in \{1, \dots, M_u\}$, i.e., $f_n^{l,e}(L) = \sum_k \alpha_k^{l,e} \phi_k(L)$. For any $(I, u) \in [I^l, I^{l+1}] \times [u^e, u^{e+1}]$ we then provide the interpolated approximation $\hat{C}_n(L, I, u)$ as

$$(3.12) \quad \hat{C}_n(L, I, u) = \frac{\begin{bmatrix} I^{l+1} - I & I - I^l \end{bmatrix} \begin{bmatrix} f_n^{l,e}(L) & f_n^{l,e+1}(L) \\ f_n^{l+1,e}(L) & f_n^{l+1,e+1}(L) \end{bmatrix} \begin{bmatrix} u^{e+1} - u \\ u - u^e \end{bmatrix}}{(u^{e+1} - u^e)(I^{l+1} - I^l)}.$$

Nonparametric approximation: Further alternatives for \mathcal{H}_n^c can be found in [26] who used Gaussian process regression and [21, 22] who used local polynomial regression. For semi-parametric approximation, [10] developed piecewise multivariate linear regression.

There are several possibilities for choosing the designs \mathcal{D}^a and \mathcal{D}^c , see [26] for a detailed discussion of different regression designs \mathcal{D}^c and their impact on the quality of the final solution.

3.3. Evaluation. We analyze the quality of the solution by computing three quantities on the out-of-sample dataset:

- estimate of the value function $V_0(\mathbf{x}_0)$ at $t = 0$ and state \mathbf{x}_0 ;
- empirical frequency of inadmissible decisions on the controlled trajectories $\mathbf{x}^{\hat{u}}$;
- statistical test for the realized number of constraint violations (blackouts for the microgrid).

Good solutions should minimize costs and not apply inadmissible controls. However, since we employ empirical estimators, \mathcal{U} is never known with certainty and we must handle the possibility that constraints are violated with probability more than p . In turn this leads to the trade-off between complying with (2.1) and optimizing costs. Similar treatment of constraints in the context of sample average approximation of probabilistic constrained optimization problems have been discussed in [27, 29]. Moreover, our framework implies that the whole algorithm is stochastic: multiple runs will lead to different results since both \hat{p}_n and \hat{C}_n are impacted by the random samples y_n^j and w_n^i .

Estimate of the value function: We evaluate the value function $\hat{V}_0(\mathbf{x}_0)$ at time $t_0 = 0$ and state \mathbf{x}_0 using M' out-of-sample paths $(\mathbf{x}_{0:N}^{\hat{u},m'})$, $m' = 1, \dots, M'$. Each trajectory $(\mathbf{x}_{0:N}^{\hat{u},m'})$ is generated by applying the estimated optimal control $\hat{u}_{0:N-1}$ based on the continuation value and admissible sets $(\hat{C}_n, \hat{\mathcal{U}}_n)_{n=0}^{N-1}$ leading to the realized pathwise cost

$$v_0(\mathbf{x}_{0:N}^{\hat{u},m'}) := \sum_{n=0}^{N-1} \sum_{k=0}^{K-1} \pi_{n_k}(\mathbf{x}_{n_k}^{\hat{u},m'}, \hat{u}_n^{m'}) \Delta n_k + W(\mathbf{x}_N^{\hat{u},m'}).$$

The resulting empirical Monte Carlo estimate is

$$(3.13) \quad \hat{V}_0(\mathbf{x}_0) \simeq \frac{1}{M'} \sum_{m'=1}^{M'} v_0(\mathbf{x}_{0:N}^{\hat{u},m'})$$

and represents an unbiased estimation of the value of the control policy and an asymptotic upper bound estimation of the value function, provided all controls used are admissible.

Empirical frequency of inadmissible decisions on the controlled trajectories: For the M' out-of-sample paths, we compare the estimated optimal control $\{\hat{u}_n(\mathbf{x}_n^{\hat{u},m'})\}_{n=m=1}^{N-1,M'}$ against the minimum admissible control $\{u_n^{\min}(\mathbf{x}_n^{\hat{u},m'})\}_{n=m=1}^{N-1,M'}$ assumed for a second to be known. Namely, for each path we compute the number of inadmissible decisions $w_0(\mathbf{x}_{0:N}^{\hat{u},m'})$ and the respective empirical frequency w_{freq} as:

$$(3.14) \quad w_0(\mathbf{x}_{0:N}^{\hat{u},m'}) := \sum_n \mathbb{1}_{\hat{u}_n(\mathbf{x}_n^{\hat{u},m'}) < u_n^{\min}(\mathbf{x}_n^{\hat{u},m'})} \quad \text{and} \quad w_{freq} := \frac{1}{N \cdot M'} \sum_{m'=1}^{M'} w_0(\mathbf{x}_{0:N}^{\hat{u},m'}),$$

respectively. We employ these metrics in Section 5, where a “gold standard” $\{u_n^{\min}(\mathbf{x}_n^{\hat{u},m'})\}_{n=m=1}^{N-1,M'}$ is obtained by brute force, utilizing a simulation budget 10^5 larger than for the actual methods we are comparing. Empirical gold standard is a common technique when analytical benchmark is unavailable, see e.g. [14]. A good estimation method should yield $w_{freq} \simeq 0$. However, if $w_{freq} > 0$, a controller can choose margin of error ξ_n (Equation (3.7)) to reduce w_{freq} at the expense of higher cost $v_0(\mathbf{x}_{0:N}^{\hat{u},m'})$.

Remark 3.3. In Equation (3.14) we exploit the structure of the admissible set for the microgrid control (cf. Remark 2.4). Generally, $w_0(\mathbf{x}_{0:N}^{\hat{u},m'}) := \sum_n \mathbb{1}_{\hat{u}_n(\mathbf{x}_n^{\hat{u},m'}) \notin \mathcal{U}_n(\mathbf{x}_n)}$, where $\mathcal{U}_n(\mathbf{x}_n)$ is either known in closed form or computed via empirical gold standard. The general setting, without any assumptions on the structure of $\mathcal{U}_n(\mathbf{x}_n)$, is computationally very expensive and beyond the scope of this work.

Statistical test: Next we propose statistical tests using the controlled trajectories to validate different methods for admissible set estimation. Such a test is essential to affirm the use of a numerical scheme for \mathcal{U}_n in the absence of a benchmark. As an example, in the context of microgrid we want to test the null hypothesis H_0 that the realized probability of blackouts is bounded to the required level against the alternative H_1 that their probability is too high. Let

$$(3.15) \quad B_n^{m'} = \mathbb{1}\left(\mathcal{G}(\mathbf{x}_{s \in [t_n, t_{n+1}]}^{\hat{u},m'}) > 0\right), \quad n = 0, \dots, N-1 \text{ and } m' = 1, \dots, M'.$$

Ignoring the correlation due to the temporal dependence in \mathbf{x}_n , we assume that $B_n^{m'} \sim \text{Bernoulli}(\tilde{p})$ are i.i.d. and \tilde{p} represents the true (unknown) probability of blackout. We want to test:

$$(3.16) \quad H_0 : \tilde{p} \leq p \quad \text{vs.} \quad H_1 : \tilde{p} > p.$$

A common approach to such composite null hypothesis is to replace H_0 with a more conservative hypothesis $\tilde{p} = p$ leading to the test statistic

$$(3.17) \quad \mathcal{T} := \frac{\sum_{m',n} (B_n^{m'} - p)}{\sqrt{M' \cdot N \cdot p \cdot (1-p)}} \sim \mathcal{N}(0, 1).$$

Hence, H_0 is rejected at a confidence level α if $\mathcal{T} > z_\alpha$ with $z_\alpha = \Phi^{-1}(\alpha)$, e.g. $z_\alpha = 1.65$ for $\alpha = 95\%$.

Remark 3.4. The above test assumes independence and identical distribution of $B_n^{m'}$'s. In the context of the microgrid example, neither of the two assumptions are valid; $B_n^{m'}$ has a different distribution because the state of the system affects the probability of a blackout, thus \tilde{p} varies with n, m' . Furthermore, $B_n^{m'}$ are not independent as they are derived from a single, sequentially controlled trajectory.

Remark 3.5. In the microgrid setup, the blackout constraint is frequently not binding (the net demand is negative half of the time). Therefore, \mathcal{T} as defined in equation (3.17) is most likely negative leading to accept the H_0 even when the method fails to choose the admissible control when the constraint

is binding. We fix this by evaluating the sum only when the constraint is binding, i.e.

$$(3.18) \quad \tilde{\mathcal{T}} := \frac{\sum_{m',n} (B_n^{m'} - p) \mathbb{1}_{\hat{u}_n^{\min}(\mathbf{x}_n^{\hat{u},m'}) > 0}}{\sqrt{p \cdot (1-p) \cdot M' \cdot N \cdot w_{bind}}} \quad \text{where} \quad w_{bind} = \frac{\sum_{m',n} \mathbb{1}_{u_n^{\min}(\mathbf{x}_n^{\hat{u},m'}) > 0}}{M' \cdot N},$$

where $\hat{u}_n^{\min}(\mathbf{x}_n^{\hat{u},m'})$ represents model estimate of the true minimum admissible control $u_n^{\min}(\mathbf{x}_n^{\hat{u},m'})$.

To wrap up this section, Algorithm 3.1 (dubbed Dynamic Emulation due to similarities with a related algorithm for unconstrained stochastic control from [26]) summarizes the overall sequence of steps. Lines 1-2 contain the parameters to the algorithm, Lines 3-8 (and 14-19) yield the stochastic simulator which generates designs and corresponding one-step paths. Line 10 (and again Line 20) computes pathwise one-step costs. Line 12 is the admissible set estimation. Line 13 is the estimation of the continuation value.

Algorithm 3.1 carries several advantages. First and foremost it is very general, and does not make any restrictions on the distribution $G_n(\mathbf{X}_n, u)$ defining \mathcal{U}_n or the form of the payoffs $\pi(\mathbf{x}, u)$. Hence it can be generically applied across a wide spectrum of SCPC problems. Second, the same template (in particular based on having two independent sub-modules) accommodates a slew of techniques for learning \mathcal{C} and \mathcal{U} bringing plug-and-play functionality, such as straightforward switching from probability to quantile estimation. Third, it allows for computational savings through parallelizing the estimation of \mathcal{U} and \mathcal{C} , or by re-using the same design and simulations $\mathcal{D}_n^a \equiv \mathcal{D}_n^c$ for the computation of the two sub-modules.

Remark 3.6. The challenge of RMC methods is that the errors recursively propagate backward. As a result, poor estimation at one step can affect the overall quality of the solution. In our algorithm, the errors at every step occur due to:

- Approximation architecture \mathcal{H}_n^a for $\hat{\mathcal{U}}_n$ \Rightarrow Projection error in admissible control set estimation;
- Approximation architecture \mathcal{H}_n^c for $\hat{\mathcal{C}}_n$ \Rightarrow Projection error in estimating continuation value;
- Designs \mathcal{D}_n^a and \mathcal{D}_n^c \Rightarrow Finite-sample Monte Carlo errors (difference between empirical estimates and theoretical projection-based ones)
- Discretization of the time interval $[t_n, t_{n+1}]$ using Δn_k \Rightarrow Integration error in approximating the integral $\int_{t_n}^{t_{n+1}} \pi_s(\mathbf{X}(s), u) ds$ and the admissible set \mathcal{U}_n .
- Numerical approximation of the solution of the controlled dynamics of $\mathbf{X}(t)$.
- Optimization errors in maximizing for \hat{u} over $\hat{\mathcal{U}}$, especially when the control set \mathcal{W} is continuous.

4. Admissible set estimation. In this section we propose two different approaches to estimate the admissible set of controls \mathcal{U}_n in equation (2.3):

- **Probability estimation:** Given a state $\mathbf{X}_n = \mathbf{x}$ and $u \in \mathcal{W}$, we estimate, via simulation, the probability of violating the constraint

$$\hat{p}_n(\mathbf{x}, u) \simeq \mathbb{P}(G_n(\mathbf{x}, u) > 0).$$

It follows that $u \in \hat{\mathcal{U}}_n(\mathbf{x}) \Leftrightarrow \hat{p}_n(\mathbf{x}, u) < p$. Particularly, to compute $\hat{p}_n(\mathbf{x}, u)$ we consider Gaussian process smoothing of empirical probabilities, logistic regression and parametric density fitting.

- **Quantile estimation:** We approximate the quantile $q_n(\mathbf{x}, u)$ of $G_n(\mathbf{x}, u)$ via empirical ranking, support vector machines and quantile regression methods. The admissible sets $\mathcal{U}_n(\mathbf{x})$ and $\mathcal{X}_n^a(u)$

Algorithm 3.1 Dynamic Emulation Algorithm

1: N (time steps), M_c (simulation budget for conditional expectation),
2: M_a (simulation budget for admissible set estimation)
3: Generate designs:
4: $\mathcal{D}_{N-1}^a := (\mathbf{x}_{N-1}^{\mathcal{D}_{N-1}^a}, u_{N-1}^{\mathcal{D}_{N-1}^a})$ of size M_a for estimating $\hat{\mathcal{U}}$.
5: $\mathcal{D}_{N-1}^c := (\mathbf{x}_{N-1}^{\mathcal{D}_{N-1}^c}, u_{N-1}^{\mathcal{D}_{N-1}^c})$ of size M_c for estimating $\hat{\mathcal{C}}$.
6: Generate one-step paths:
7: $\mathbf{x}_{N-1}^{i, \mathcal{D}_{N-1}^a} \mapsto \mathbf{x}_N^{i, \mathcal{D}_N^a}$ using $u_{N-1}^{\mathcal{D}_{N-1}^a}$ for $i = 1, \dots, M_a$
8: $\mathbf{x}_{N-1}^{j, \mathcal{D}_{N-1}^c} \mapsto \mathbf{x}_N^{j, \mathcal{D}_N^c}$ using $u_{N-1}^{\mathcal{D}_{N-1}^c}$ for $j = 1, \dots, M_c$
9: Terminal condition:
10: $y_{N-1}^j \leftarrow \sum_{k=0}^{K-1} \pi_{(N-1)_k}(\mathbf{x}_{(N-1)_k}^{j, \mathcal{D}_{N-1}^c}, u_{(N-1)_k}^{j, \mathcal{D}_{N-1}^c}) \Delta n_k + W(\mathbf{x}_N^{j, \mathcal{D}_N^c})$ for $j = 1, \dots, M_c$
11: **for** $n = N-1, \dots, 1$ **do**
12: Estimate $\hat{\mathcal{U}}_n(\cdot)$ using methods in Section 4 and paths $\mathbf{x}_n^{i, \mathcal{D}_n^a} \mapsto \mathbf{x}_{n+1}^{i, \mathcal{D}_n^a}$
13: $\hat{\mathcal{C}}_n(\cdot, \cdot) \leftarrow \arg \min_{f_n \in \mathcal{H}_n^c} \sum_{j=1}^{M_c} |f_n(\mathbf{x}_n^{j, \mathcal{D}_n^c}, u_n^{j, \mathcal{D}_n^c}) - y_n^j|^2$
14: Generate designs:
15: $\mathcal{D}_{n-1}^a := (\mathbf{x}_{n-1}^{\mathcal{D}_{n-1}^a}, u_{n-1}^{\mathcal{D}_{n-1}^a})$ of size M_a for estimating $\hat{\mathcal{U}}$.
16: $\mathcal{D}_{n-1}^c := (\mathbf{x}_{n-1}^{\mathcal{D}_{n-1}^c}, u_{n-1}^{\mathcal{D}_{n-1}^c})$ of size M_c for estimating $\hat{\mathcal{C}}$.
17: Generate one-step paths:
18: $\mathbf{x}_{n-1}^{i, \mathcal{D}_{n-1}^a} \mapsto \mathbf{x}_n^{i, \mathcal{D}_n^a}$ using $u_{n-1}^{\mathcal{D}_{n-1}^a}$ for $i = 1, \dots, M_a$
19: $\mathbf{x}_{n-1}^{j, \mathcal{D}_{n-1}^c} \mapsto \mathbf{x}_n^{j, \mathcal{D}_n^c}$ using $u_{n-1}^{\mathcal{D}_{n-1}^c}$ for $j = 1, \dots, M_c$
20: $y_{n-1}^j \leftarrow \sum_{k=0}^{K-1} \pi_{(n-1)_k}(\mathbf{x}_{(n-1)_k}^{j, \mathcal{D}_{n-1}^c}, u_{(n-1)_k}^{j, \mathcal{D}_{n-1}^c}) \Delta n_k + \max_{u \in \hat{\mathcal{U}}_n(\mathbf{x}_n^{j, \mathcal{D}_n^c})} \{\hat{\mathcal{C}}(n, \mathbf{x}_n^{j, \mathcal{D}_n^c}, u)\} \forall j$
21: **end for**
22: **return** $\{\hat{\mathcal{C}}_n(\cdot, \cdot), \hat{\mathcal{U}}_n(\cdot)\}_{n=1}^{N-1}$

are then defined as:

$$\hat{\mathcal{U}}_n(\mathbf{x}) := \{u : \hat{q}_n(\mathbf{x}, u) \leq 0\} \text{ and } \hat{\mathcal{X}}_n^a(u) := \{\mathbf{x} : \hat{q}_n(\mathbf{x}, u) \leq 0\}.$$

To implement all of the above techniques we use Monte Carlo simulation, specifying first the simulation design and then sampling (independently across draws) the G 's or Y 's to be used as training data. We work in a flexible framework where samples of $G_n(\mathbf{x}, u)$ are generated in batches of M_b simulations from each design site $\{\mathbf{x}^i, u^i\}_{i=1}^{M_a}$. The case of $M_b = 1$ corresponds to a classical regression approach, while large $M_b \gg 1$ can be interpreted as nested Monte Carlo averaging along M_b inner samples.

Remark 4.1. In section 3.2, we parameterized the elements of the approximation space \mathcal{H}_n^c for estimation of the continuation value function $\hat{\mathcal{C}}(\cdot, \cdot)$ via vectors $\boldsymbol{\alpha}$ i.e. $f_n^c(\mathbf{x}, u) \equiv f_n^c(\mathbf{x}, u; \boldsymbol{\alpha})$, cf. (3.10) and (3.11). To distinguish, in the following sections we use $\boldsymbol{\beta}$ to parameterize the approximators in \mathcal{H}_n^a for estimating the admissible set: $f_n^a(\mathbf{x}, u) \equiv f_n^a(\mathbf{x}, u; \boldsymbol{\beta})$ in eq. (3.3). The meaning and dimension of $\boldsymbol{\beta}$ will vary from method to method.

4.1. Probability estimation.

4.1.1. Interpolated nested Monte Carlo (INMC). Recall the NMC method from Section 2.1 where we select M_a design sites of state-action pairs and simulate multiple paths from each site to *locally*

assess the probability of $G_n(\mathbf{x}, u) > 0$ (in what follows, we suppress in the notation the dependence on n). Specifically, for each design site (\mathbf{x}^i, u^i) , $i = 1, \dots, M_a$, we simulate M_b batched samples from the distribution $G(\mathbf{x}^i, u^i)$ as $\{g^b(\mathbf{x}^i, u^i)\}_{b=1}^{M_b}$. The unbiased point estimator of $p(\mathbf{x}^i, u^i)$ is:

$$(4.1) \quad \bar{p}(\mathbf{x}^i, u^i) := \sum_{b=1}^{M_b} \frac{\mathbb{1}_{g^b(\mathbf{x}^i, u^i) > 0}}{M_b}.$$

Since (4.1) only yields M_b local estimates $\bar{p}(\mathbf{x}^i, u^i)$, for Algorithm 3.1 we have to extend them to an arbitrary state-action $(\mathbf{x}, u) \mapsto \hat{p}_{\text{INMC}}(\mathbf{x}, u)$. This is achieved by interpolating $\bar{p}(\mathbf{x}^i, u^i)$'s, e.g. linearly. The admissible set with confidence level ρ becomes:

$$\hat{\mathcal{U}}_{\text{INMC}}^{(\rho)}(\mathbf{x}) := \left\{ u : \hat{p}_{\text{INMC}}(\mathbf{x}, u) \leq p - z_\rho \sqrt{\frac{\hat{p}_{\text{INMC}}(\mathbf{x}, u)(1 - \hat{p}_{\text{INMC}}(\mathbf{x}, u))}{M_b}} \right\}.$$

However, especially for M_b small, interpolation performs poorly because the underlying point estimates $\bar{p}(\mathbf{x}^i, u^i)$ are *noisy*. Therefore, smoothing should be applied via a statistical *regression* model. Regression borrows information cross-sectionally to mitigate the estimation noise, reducing the variance of \bar{p} .

4.1.2. Gaussian process regression (GPR). GPR is a flexible non-parametric regression method that views the map $(\mathbf{x}, u) \rightarrow p(\mathbf{x}, u)$ as a realization of a Gaussian random field so that any finite collection of $\{p(\mathbf{x}, u)\}, (\mathbf{x}, u) \in \mathcal{X} \times \mathcal{W}$ is multivariate Gaussian. For any n design sites $\{(\mathbf{x}^i, u^i)\}_{i=1}^n$, GPR posits that

$$p(\mathbf{x}^1, u^1), \dots, p(\mathbf{x}^n, u^n) \sim \mathcal{N}(\vec{m}_n, \mathbf{K}_n)$$

with mean vector $\vec{m}_n := [m(\mathbf{x}^1, u^1; \boldsymbol{\beta}), \dots, m(\mathbf{x}^n, u^n; \boldsymbol{\beta})]$ and $n \times n$ covariance matrix \mathbf{K}_n comprised of $\kappa(\mathbf{x}^i, u^i, \mathbf{x}^{i'}, u^{i'}; \boldsymbol{\beta})$, for $1 \leq i, i' \leq n$. The vector $\boldsymbol{\beta}$ represents all the hyperparameters for this model.

Given the training dataset $\{(\mathbf{x}^i, u^i), \bar{p}^i\}_{i=1}^{M_a}$ (where \bar{p}^i is a shorthand for $\bar{p}(\mathbf{x}^i, u^i)$), GPR infers the posterior of $p(\cdot, \cdot)$ by assuming an observation model of the form $\bar{p}(\mathbf{x}, u) = p(\mathbf{x}, u) + \epsilon$ with a Gaussian noise term $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Conditioning equations for multivariate normal vectors imply that the posterior predictive distribution $p(\mathbf{x}, u) | \{(\mathbf{x}^i, u^i), \bar{p}^i\}_{i=1}^{M_a}$ at any arbitrary site (\mathbf{x}, u) is also Gaussian with the posterior mean $\hat{p}_{\text{GPR}}(\mathbf{x}, u)$ that is the proposed estimator of $p(\mathbf{x}, u)$:

$$(4.2) \quad \begin{aligned} \hat{p}_{\text{GPR}}(\mathbf{x}, u) &:= m(\mathbf{x}, u) + K^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} (\vec{p} - \vec{m}) = \mathbb{E}[p(\mathbf{x}, u) | \vec{x}, \vec{u}, \vec{p}] \\ \text{where } \vec{x} &= [\mathbf{x}^1, \dots, \mathbf{x}^{M_a}]^T, \quad \vec{u} = [u^1, \dots, u^{M_a}]^T, \quad \vec{p} = [\bar{p}^1, \dots, \bar{p}^{M_a}]^T, \\ K^T &= [\kappa(\mathbf{x}, u, \mathbf{x}^1, u^1; \boldsymbol{\beta}), \dots, \kappa(\mathbf{x}, u, \mathbf{x}^{M_a}, u^{M_a}; \boldsymbol{\beta})], \\ (4.3) \quad \vec{m} &= [m(\mathbf{x}^1, u^1; \boldsymbol{\beta}), \dots, m(\mathbf{x}^{M_a}, u^{M_a}; \boldsymbol{\beta})], \end{aligned}$$

and \mathbf{K} is $M_a \times M_a$ covariance matrix described through the kernel function $\kappa(\cdot, \cdot; \boldsymbol{\beta})$.

The mean function is often assumed to be constant $m(\mathbf{x}, u; \boldsymbol{\beta}) = \beta_0$ or described using a linear model $m(\mathbf{x}, u; \boldsymbol{\beta}) = \sum_{k=1}^K \beta_k \phi(\mathbf{x}^k, u^k)$ with $\phi(\cdot, \cdot)$ representing a polynomial basis. A popular choice for the kernel $\kappa(\cdot, \cdot, \cdot)$ is squared exponential (see equation (4.4)) with $\{\{\beta_{\text{len}, k}\}_{k=1}^d, \beta_{\text{len}, u}\}$ termed the lengthscales and σ_p^2 the process variance of $p(\cdot, \cdot)$:

$$(4.4) \quad \kappa(\mathbf{x}^i, u^i, \mathbf{x}^{i'}, u^{i'}) = \sigma_p^2 \exp \left(- \sum_{k=1}^d \frac{(x^{i,k} - x^{i',k})^2}{\beta_{\text{len}, k}} - \frac{(u^i - u^{i'})^2}{\beta_{\text{len}, u}} \right).$$

The hyperparameters $\boldsymbol{\beta} := (\{\beta_k\}_{k=1}^K, \{\beta_{\text{len},k}\}_{k=1}^d, \beta_{\text{len},u}, \sigma_p^2, \sigma_e^2)$ represent different attributes of $\hat{p}_{\text{GPR}}(\cdot, \cdot)$. The parameters $\{\beta_k\}_{k=1}^K$ determine the trend of $\hat{p}(\cdot, \cdot)$ over the input domain. The lengthscales $\{\beta_{\text{len},k}\}_{k=1}^d$ and $\beta_{\text{len},u}$ determine spatial smoothness or how quickly the function changes. Small lengthscales make \hat{p} to be “bumpy”, while large lengthscales make \hat{p} close to linear. The process variance σ_p^2 determines the amplitude of fluctuations in \hat{p}_{GPR} and σ_e^2 represents the sampling noise variance. These hyperparameters are estimated by maximizing the log-likelihood function based on the dataset $\{(x^i, u^i), \bar{p}^i\}_{i=1}^{M_a}$. Besides squared exponential kernel (Equation 4.4), other popular kernels include Matérn-3/2 and Matérn-5/2 [35].

A conservative estimate $\hat{p}_{\text{GPR}}^{(\rho)}(\mathbf{x}, u)$ at confidence level ρ is obtained by explicitly incorporating the (estimated) standard error of $\bar{p}(\mathbf{x}^i, u^i)$ into the GPR smoothing. Namely, we adjust the training dataset to $\{(x^i, u^i), \bar{p}_\rho^i\}_{i=1}^{M_a}$, where $\bar{p}_\rho^i := \bar{p}(\mathbf{x}^i, u^i) + z_\rho \sqrt{\frac{\bar{p}(\mathbf{x}^i, u^i)(1 - \bar{p}(\mathbf{x}^i, u^i))}{M_b}}$. The resulting $\hat{p}_{\text{GPR}}^{(\rho)}(\mathbf{x}, u)$ is the counterpart of (4.2) using $\{(x^i, u^i), \bar{p}_\rho^i\}_{i=1}^{M_a}$.

In Figure 2b we present the dataset $\{L^i, I^i, 0, \bar{p}^i\}_{i=1}^{M_a}$ (background colormap) for the microgrid case study. The thick red line indicates the contour $\{\hat{p}_{\text{GPR}} = 5\%\}$, dividing the state space \mathcal{X} for $u = 0$ into admissible $\mathcal{X}^a(0)$ (left of red line) and inadmissible region $(\mathcal{X}^a(0))^c$ (right of red line).

4.1.3. Logistic regression (LR). In the previous section, we created local batches to estimate $p(\mathbf{x}^i, u^i)$ pointwise and then regressed these estimates to build a global approximator. A classical alternative is to learn the probability of $G(x, u) > 0$ using a logistic regression model. This setup uses a single sample $g(\mathbf{x}^i, u^i)$ from $G(\mathbf{x}^i, u^i)$ from each design site (\mathbf{x}^i, u^i) and transforms it to a binary response $y^i = \mathbb{1}_{g(\mathbf{x}^i, u^i) > 0}$. The probability $\hat{p}(\mathbf{x}, u)$ is then modeled as a generalized linear model with a logit link function

$$(4.5) \quad \mathbb{P}(Y = 1 | \mathbf{x}, u) = \frac{1}{1 + e^{-\boldsymbol{\beta}^T \phi(\mathbf{x}, u)}} =: \hat{p}_{\text{LR}}(\mathbf{x}, u; \boldsymbol{\beta}).$$

The basis functions $\phi(\mathbf{x}, u)$ could be polynomials, e.g. quadratic or cubic in coordinates of (\mathbf{x}, u) . The regression coefficients $\boldsymbol{\beta}$ are fitted using the dataset $\{\mathbf{x}^i, u^i, \mathbf{y}^i\}_{i=1}^{M_s}$, as the solution to

$$(4.6) \quad \arg \max_{\boldsymbol{\beta}} \sum_{i=1}^{M_s} \{y^i \log p_{\text{LR}}(\mathbf{x}^i, u^i; \boldsymbol{\beta}) + (1 - y^i) \log(1 - p_{\text{LR}}(\mathbf{x}^i, u^i; \boldsymbol{\beta}))\}.$$

We may again create a more conservative estimate $\hat{\mathcal{U}}_{\text{LR}}^{(\rho)}(\mathbf{x})$ of $\hat{\mathcal{U}}_{\text{LR}}(\mathbf{x})$ at confidence level ρ by utilizing the standard error for \hat{p}_{LR} using the Delta method [40]:

$$\hat{\mathcal{U}}_{\text{LR}}^{(\rho)}(\mathbf{x}) := \left\{ u : \hat{p}_{\text{LR}}(\mathbf{x}, u; \boldsymbol{\beta}) \leq p - z_\rho \sqrt{\hat{p}_{\text{LR}}(\mathbf{x}, u)(1 - \hat{p}_{\text{LR}}(\mathbf{x}, u))\phi^T \text{Var}(\boldsymbol{\beta})\phi} \right\}.$$

In Figure 2a, we present the original realizations $y^i \in \{0, 1\}$ (in blue) for a design in the input subspace $(L, I, u = 0)$ of the microgrid case study. The figure indicates the resulting logistic regression fit $\hat{p}_{\text{LR}}(L, I, 0)$ at levels 1%, 5% and 10% (i.e. contour lines of $\hat{p}_{\text{LR}}(\hat{\boldsymbol{\beta}}) \in \{0.01, 0.05, 0.1\}$). The admissibility set for $u = 0$, $\mathcal{X}_n^a(0)$ is the region to the left of the thick red contour.

Remark 4.2. Similar to Section 4.1, we can simulate batched samples from each design site for the logistic regression, leading to “binomial” observation likelihood instead of (4.6).

Remark 4.3. A non-parametric variant of equation (4.5) is kernel logistic regression, where the basis functions are $\phi_j(\mathbf{x}, u) = \kappa(\mathbf{x}, u, \mathbf{x}^j, u^j)$ for a kernel function κ centered at (\mathbf{x}^j, u^j) . One common choice is radial basis functions (RBF) where $\kappa(\mathbf{x}, u, \mathbf{x}^j, u^j) = \exp(-\gamma_1 \|\mathbf{x} - \mathbf{x}^j\|_2^2 - \gamma_2 \|u - u^j\|_2^2)$. RBF can be interpreted as the squared-exponential kernel for a logistic Gaussian Process model, with a fixed bandwidth parameter γ_i . In contrast, in GPR the bandwidths are estimated through MLE.

4.1.4. Parametric density fitting (PF). This approach fits the distribution $G(\mathbf{x}, u)$, and then analytically infers the probability $\mathbb{P}(G(\mathbf{x}, u) > 0)$ from the corresponding cumulative distribution function. This is done by proposing a parametric family $\{f(\cdot; \Theta)\}$ of densities, fitting the underlying parameters Θ based on an empirical sample from G and then evaluating the resulting analytical probability $\bar{p}_{PF}(\mathbf{x}, u) := \int_0^\infty f_G(z|\hat{\Theta}(\mathbf{x}, u))dz$. PF yields a “universal” solution across a range of constraint levels p .

At a design site (\mathbf{x}, u) , the probability $p(\mathbf{x}, u)$ is estimated in a two-step procedure: first estimated locally over a design $\mathcal{D}^a = \{\mathbf{x}^i, u^i\}$ and then regressed over the full input domain $\mathcal{X} \times \mathcal{W}$. For the first step, we apply nested Monte Carlo to generate a collection of realized $\{g^b(\mathbf{x}^i, u^i)\}_{b=1}^{M_b}$ that is used to construct a parametric density via the maximum likelihood estimate:

$$(4.7) \quad \hat{\Theta}^i := \arg \max_{\Theta} \sum_{b=1}^{M_b} \log f_G(g^b(\mathbf{x}_i, u_i)|\Theta).$$

In the second step, we evaluate $\tilde{p}_{PF}(\mathbf{x}^i, u^i) := \int_0^\infty f_G(z|\hat{\Theta}(\mathbf{x}^i, u^i))$ and extend it to the full domain $\mathcal{X} \times \mathcal{W}$ based on the computed $\{\mathbf{x}^i, u^i, \tilde{p}_{PF}(\mathbf{x}^i, u^i)\}_{i=1}^{M_a}$ using \mathcal{L}_2 projection:

$$(4.8) \quad \hat{p}_{PF} = \arg \min_{\hat{p} \in \mathcal{M}_T} \|\hat{p}(\mathbf{x}^i, u^i) - \tilde{p}_{PF}(\mathbf{x}^i, u^i)\|^2,$$

where \mathcal{M}_T is an approximation space chosen for regression. The admissible set $\mathcal{U}(\mathbf{x})$ is estimated as:

$$\hat{\mathcal{U}}_{PF}(\mathbf{x}) := \{u : \hat{p}_{PF}(\mathbf{x}, u) \leq p\}.$$

A transformation of the distribution $G(\mathbf{x}, u)$ might be important for above distribution fitting. For example, in the context of microgrid, in Section 2.2, $G = \mathcal{L}(\sup_{s \in [t_n, t_{n+1})} S(s))$ has a point mass at 0 and thus, any continuous distribution will lead to poor statistical estimation. Using a transformation that preserves the probability of the target event,

$$(4.9) \quad \mathbb{P}\left(\sup_{s \in [t_n, t_{n+1})} S(s) > 0 | \mathcal{F}_n\right) = \mathbb{P}\left(\sup_{s \in [t_n, t_{n+1})} [L(s) - u_n - \frac{I(s)}{\delta s} \wedge B_{\max}] > 0 | \mathcal{F}_n\right),$$

we work with $G'(L_n, I_n, u_n) := \mathcal{L}(\sup_{s \in [t_n, t_{n+1})} [L(s) - u_n - \frac{I(s)}{\delta s} \wedge B_{\max}])$. In Figure 2c we present the empirical and estimated probability $z \mapsto \mathbb{P}(G'(L_n, I_n, u_n) > z)$ when $L_n = 5.5, I_n = 1.48$ and $u_n \in \{0, 1\}$ for the microgrid example. We model the distribution G' using a truncated normal distribution, $\mathbb{P}(G' \leq g) = \Phi(\frac{g - \theta_2}{\theta_3}) \mathbb{1}_{g \geq \theta_1}$, with parameters $\Theta = (\theta_1, \theta_2, \theta_3)$ representing the location of censoring, the mean and the standard deviation respectively. At $L_n = 5.5, I_n = 1.48, u_n = 1.0$ and inner simulation budget $M_b = 100$, the estimated parameters $(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3) = (-1.5, -1.12, 0.53)$ result in probability $\tilde{p}_{PF}(5.5, 1.48, 1.0) = 0.016$. The corresponding probability after \mathcal{L}_2 projection (equation (4.8)) is $\hat{p}_{PF}(5.5, 1.48, 1.0) = 0.017$. Thus at $p = 0.05$, the control $u = 1.0 \in \hat{\mathcal{U}}_n$ is admissible. However, at $u_n = 0$, $(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3) = (-0.5, -0.12, 0.55)$, $\tilde{p}_{PF}(5.5, 1.48, 0.0) = 0.414$ and $\hat{p}_{PF}(5.5, 1.48, 0.0) = 0.429$, thus the control $u = 0 \notin \hat{\mathcal{U}}_n$ is inadmissible.

4.2. Quantile estimation. In this section we consider methods for modeling and estimating $q(\mathbf{x}^i, u^i)$, the $(1-p)$ -th quantile of the distribution $G(\mathbf{x}^i, u^i)$. Admissibility corresponds to the quantile being negative.

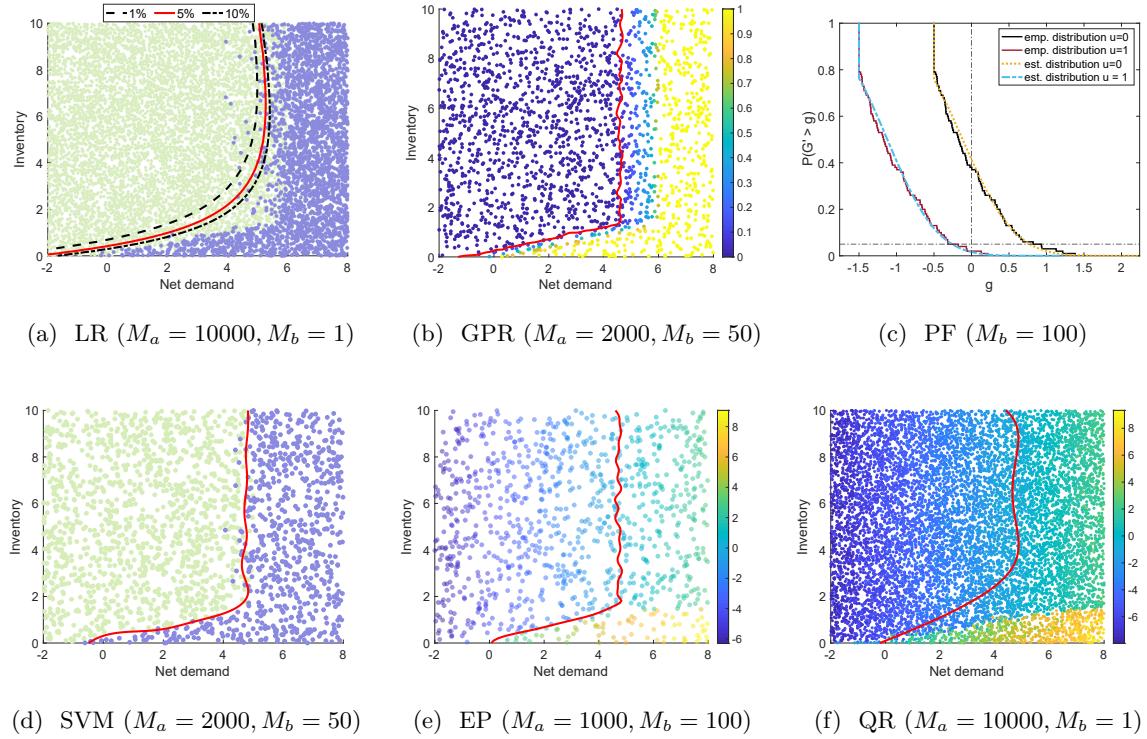


Fig. 2: Training data and fitted models for the methods of Section 4 at $u = 0$. Top row: probability estimation schemes, bottom row: quantile estimation schemes. *Top/left panel:* Training set $\{L^i, I^i, y^i\}_{i=1}^{M_a}$ for the LR model, color-coded according to the value of $y^i \in \{0, 1\}$, along with the estimated contours for $\hat{p}_{LR}(L, I)$ at levels $\{1\%, 5\%, 10\%\}$. *Top/center:* Training set $\{L^i, I^i, \bar{p}^i\}_{i=1}^{M_a}$ color-coded according to \bar{p}^i for GPR along with the contour $\{\hat{p}_{GPR}(L, I) = 5\%\}$. *Top/right:* parametric density fitting at $L_0 = 5.5, I_0 = 1.48$ and $u \in \{0, 1\}$. We show the empirical and fitted inverse cdf $\mathbb{P}(G' > g)$ based on a truncated Gaussian distribution. *Bottom/left:* Training set $\{L^i, I^i, y^i\}_{i=1}^{M_a}$ for SVM (color-coded according to $y^i \in \{-1, 1\}$) and the decision boundary in red. *Bottom/center:* Training set $\{L^i, I^i, \bar{q}^i\}_{i=1}^{M_a}$ color-coded according to \bar{q}^i for EP and the contour $\{\hat{q} = 0\}$. *Bottom/right:* Training set $\{L^i, I^i, g^i\}_{i=1}^{M_a}$ color-coded according to g^i for QR along with the contour $\{\hat{q}_{QR}(L, I) = 0\}$. All models share the same ground truth, so the red contours are identical up to model-specific estimation errors.

4.2.1. Empirical percentiles (EP). As before, we start by choosing M_a design sites of state-action pairs and generate batched samples $\{g^b(\mathbf{x}^i, u^i)\}_{b=1}^{M_b}$ from each design site (\mathbf{x}^i, u^i) . The empirical estimate of $q(\mathbf{x}^i, u^i)$ is simply the $(1-p)^{th}$ percentile of the realized $\{g^b\}_{b=1}^{M_b}$ (which requires $M_b > p^{-1}$):

$$\bar{q}(\mathbf{x}^i, u^i) = \text{percentile} \left(\{g^b\}_{b=1}^{M_b}, 100(1-p)\% \right).$$

Similar to previous methods, we extend to arbitrary $(\mathbf{x}, u) \mapsto \hat{q}(\mathbf{x}, u)$ using regression on the dataset $\{\mathbf{x}^i, u^i, \bar{q}(\mathbf{x}^i, u^i)\}_{i=1}^{M_a}$ and an approximation space \mathcal{M}_q . The set of admissible controls for \mathbf{x} is: $\hat{\mathcal{U}}_{EP}(\mathbf{x}) := \{u : \hat{q}(\mathbf{x}, u) \leq 0\}$. In Figure 2e we show the estimated $\hat{q}(\cdot, \cdot, 0)$ indicated via the background colormap. The thick red line indicates the zero-contour $\hat{q} = 0$, so that the admissibility set for $u = 0$, $\mathcal{X}_n^a(0)$, is the region to the left of the contour.

Remark 4.4. This approach is similar to the INMC approach discussed in Section 4.1, however, here we model the quantile rather than the probability of exceeding zero. Furthermore, we can use the regression standard error of $\hat{q}(\cdot, \cdot)$ to construct a more conservative estimate of the admissible set $\mathcal{U}_{EP}(\mathbf{x})$.

A popular alternative to adjusting \bar{q} 's via regression standard errors is to replace the empirical percentile with the empirical conditional tail expectation (CTE):

$$\overline{\text{CTE}}(\mathbf{x}^i, u^i) := \frac{\sum_{b=1}^{M_b} g^b \mathbb{1}_{g^b \geq \bar{q}(\mathbf{x}^i, u^i)}}{\sum_{b=1}^{M_b} \mathbb{1}_{g^b \geq \bar{q}(\mathbf{x}^i, u^i)}},$$

and then running a regression on the training set $(\mathbf{x}^i, u^i, \overline{\text{CTE}}(\mathbf{x}^i, u^i))$ to obtain the fitted CTE surface $\widehat{\text{CTE}}(\mathbf{x}, u)$ and finally $\hat{\mathcal{U}}_{\text{CTE}}(\mathbf{x}) := \{u : \widehat{\text{CTE}}(\mathbf{x}, u) \leq 0\}$. This idea is similar to regularizing the Value-at-Risk estimation with the Conditional VaR.

4.2.2. Support Vector Machines (SVM). For a fixed control u , finding the admissible set $\mathcal{X}_n^a(u)$ in (2.9) can be interpreted as classifying each input \mathbf{x} as being in $\mathcal{X}_n^a(u)$ or not. Therefore, we consider the use of classification techniques, specifically support vector machines (SVM). This approach does not estimate the $(1 - p)$ -quantile $q(\mathbf{x}, u)$, but rather its 0-level set with respect to (\mathbf{x}, u) . The starting point is to use the nested Monte Carlo simulations to compute $\bar{p}(\mathbf{x}^i, u^i)$ with much smaller batch size M_b compared to Section 4.1. Next, we construct a binary classification objective with a training dataset $\{\mathbf{x}^i, u^i, y^i\}_{i=1}^{M_a}$ where the ± 1 -labels are

$$(4.10) \quad y^i := \begin{cases} 1, & \text{if } \bar{p}(\mathbf{x}^i, u^i) < p; \\ -1, & \text{otherwise.} \end{cases}$$

The boundary separating the two classes is evaluated by solving the optimization problem:

$$(4.11) \quad \min_{\boldsymbol{\beta} \in \mathbb{R}^K} \left\{ \sum_{i=1}^{M_a} \left(1 - y^i [\boldsymbol{\beta}^T \phi(\mathbf{x}^i, u^i) + \beta_0] \right)_+ + \frac{C}{2 \cdot M_a} \|\boldsymbol{\beta}\|^2 \right\},$$

where $\phi(\mathbf{x}, u) = [\phi_1(\mathbf{x}, u), \phi_2(\mathbf{x}, u), \dots, \phi_K(\mathbf{x}, u)]^T$ are the K basis functions and C is the penalty parameter. We estimate the set of admissible controls corresponding to \mathbf{x} as:

$$\hat{\mathcal{U}}_{\text{SVM}}(\mathbf{x}) := \{u : \hat{\boldsymbol{\beta}}^T \phi(\mathbf{x}, u) + \hat{\beta}_0 \geq 0\}.$$

Figure 2d displays the estimated $\hat{\mathcal{X}}_n^a(u)$ and the corresponding dataset $(L^i, I^i, 0, y^i)$ ($u = 0$ is fixed). The region where $u = 0$ is admissible is to the left of the (thick red) decision boundary.

Remark 4.5. A conservative estimate $\hat{\mathcal{U}}_{\text{SVM}}^{(\rho)}$ is obtained by biasing the decision boundary to the left by re-labeling the training points in (4.10) via:

$$(4.12) \quad y^i = \begin{cases} 1, & \text{if } \bar{p}(\mathbf{x}^i, u^i) + z_\rho \sqrt{\frac{\bar{p}(\mathbf{x}^i, u^i)(1 - \bar{p}(\mathbf{x}^i, u^i))}{M_b}} < p \\ -1, & \text{otherwise.} \end{cases}$$

4.2.3. Quantile Regression (QR). QR directly constructs a parametric model for $q(\mathbf{x}, u)$:

$$\hat{q}(\mathbf{x}, u; \boldsymbol{\beta}) := \sum_k \beta_k \phi_k(\mathbf{x}, u).$$

To estimate the coefficients $\boldsymbol{\beta} \in \mathbb{R}^K$, we use the dataset $\{\mathbf{x}^i, u^i, g^i\}_{i=1}^{M_a}$ (where g^i is a sample from the distribution $G(\mathbf{x}^i, u^i)$) to maximize the negative log likelihood:

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^K} \left\{ \sum_{i=1}^{M_a} \mathcal{L}^{(p)} \left(g^i - \sum_{k=1}^K \beta_k \phi_k(\mathbf{x}^i, u^i) \right) \right\},$$

where $\mathcal{L}^{(p)}(y) = y(p - 1_{\{y < 0\}}) = py_+ + (1 - p)y_-$. As for the parametric density fitting, a transformation of $G(\mathbf{x}, u)$ might be beneficial when applying quantile regression. Figure 2f presents the dataset $\{L^i, I^i, 0, g^i\}_{i=1}^{M_a}$ (background colormap) and the estimated admissible set $\hat{\mathcal{X}}^a(0)$ which is the region to the left of the contour $\{\hat{q}_{QR}(L, I) = 0\}$ (thick red line).

Relying on the Delta method again to compute the variance of the estimated quantile $\hat{q}(\mathbf{x}, u; \hat{\beta})$ as $\phi(\mathbf{x}, u)'Var(\hat{\beta})\phi(\mathbf{x}, u)$, the admissible set at \mathbf{x} at confidence level ρ is:

$$\hat{\mathcal{U}}_{QR}^{(\rho)}(\mathbf{x}) := \left\{ u : \hat{q}(\mathbf{x}, u; \hat{\beta}) + z_\rho \sqrt{\phi(\mathbf{x}, u)'Var(\hat{\beta})\phi(\mathbf{x}, u)} \leq 0 \right\}.$$

5. Case Studies. Recall the problem introduced in Section 2.2 where we control the operation of a diesel generator in order to match demand at minimal cost while maintaining the probability of blackout between each decision epoch below a given threshold p . In this section, we discuss two variants of such microgrid control. In the first example, we assume a time-homogeneous net-demand process which reduces the problem of estimating admissible set to a pre-processing step. In the second example, we use time-dependent net demand process calibrated to data obtained from a microgrid in Huatacondo, Chile. Time-inhomogeneity requires to estimate the admissible set at every step. The microgrid features a perfectly efficient battery, so that the respective power output at t_{n_k} (recall t_{n_k} is a generic time instance on the finely discretized time grid) is given by:

$$B_{n_k} = \max \left(\min \left(L_{n_k} - u_n, \frac{I_{n_k}}{\Delta n_k} \right), -\frac{I_{\max} - I_{n_k}}{\Delta n_k} \right).$$

Table 1 lists other microgrid parameters, i.e. capacity of the battery I_{\max} , maximum charging rate B_{\min} , maximum discharging rate B_{\max} and diesel switching cost \mathcal{K} .

Table 1: Parameters for the Microgrid example.

$I_{\max} = 10$ (kWh), $B_{\min} = -6$, $B_{\max} = 6$ (kW), $\mathcal{K} = 5$
$T = 48$ (hours), $\Delta t = 0.25$ (hours)

5.1. Implementation details. Numerical Gold Standard: In the absence of analytic benchmark, we use a high-budget gold standard to compare the output from the models discussed in Section 4. For each fixed time-step t_n we discretize the domain $\mathcal{X} = (L, I)$ into 10,000 design sites over a grid of 100×100 . For each design site (L^i, I^j) , $i, j \in \{1, \dots, 100\}$ and $u^k \in 0 \cup \{1 = u_1, \dots, u_{101} = 10\}$, we evaluate $\hat{p}(L^i, I^j, u^k)$ using (4.1) with batch size $M_b = 10,000$. Thus, the total simulation budget is $100 \times 100 \times 102 \times 10000 \approx 10^{10}$. We then evaluate the local minimal admissible control

$$u_n^{\min}(L^i, I^j) = \min \{u : \hat{p}(L^i, I^j, u) < p\}.$$

To evaluate $u_n^{\min}(L, I)$ at new sites we employ linear interpolation on the dataset $\{L^i, I^j, u_n^{\min}(L^i, I^j)\}_{i,j=1}^{100}$.

To estimate the continuation function, we use the piecewise continuous approximation of Section 3.2 with $M_I = 15$, $M_u = 15$ and $M_L = 1500$ sites in L . The design \mathcal{D}^c is constructed as the Cartesian product $\{L^1, L^2, \dots, L^{M_L}\} \times \{I^0, I^1, \dots, I^{M_I}\} \times \{u^0, u^1, \dots, u^{M_u}\}$, where $L^i, i = 1, \dots, M_L$ are sampled uniformly from the interval $[L_{\min}, L_{\max}] = [-8, 8]$. The inventory $\{I^0, I^1, \dots, I^{M_I}\}$ and control $\{u^1, \dots, u^{M_u}\}$ discretizations are equispaced in $[0, I_{\max}]$ and $[\underline{u}, \bar{u}]$ respectively.

Low budget policies: We approximate the continuation value function \mathcal{C} using a piecewise continuous approximation with degree-3 in L combined with interpolation in other dimensions. The design \mathcal{D}^c is the same as for the numerical gold standard with discretization levels $M_L = 1000, M_I = 10, M_u = 10$.

We approximate the admissible set \mathcal{U} using the methods described in Section 4 and compare the performance of each method by using a fixed set of $M' = 20,000$ out-of-sample simulations. To address the discontinuity in $\mathcal{W} = 0 \cup [u, \bar{u}]$, we implement two separate statistical models to learn $\mathcal{U}_n(\cdot)$. As an example, with logistic regression of Section 4.1.3 we estimate two sets of parameters in equation (4.5): the first one uses one-step paths generated with $u = 0$ and a two-dimensional regression of $y^{i,(1)}$ against (L^i, I^i) . The second one uses design sites in the three-dimensional space (L, I, u) where $u \in [1, 10]$ and a 3-D regression of $y^{i,(2)}$ against (L^i, I^i, u^i) . For both regressions we choose a Sobol sequence space-filling experimental design: $\mathcal{D}^a = (L^i, I^i) \in [-2, 8] \times [0, 10]$ when $u = 0$, and $\mathcal{D}^a = (L^i, I^i, u^i) \in [-2, 8] \times [0, 10] \times [1, 10]$ otherwise. The control space $[1, 10]$ is discretized into 51 levels.

Additional parameters used for each method are specified in Table 2. We found that Matern-3/2 kernels work better than (4.4) for smoothing $\bar{p}(L, I, u)$ (GPR) and $\hat{p}(L, I, u)$ (PF) because the respective input-output maps feature steep transitions as a function of (L, i, u) . It is known that “rougher” kernels are better suited for such learning tasks compared to the C^∞ -smooth squared exponential kernel (4.4) by allowing the fitted \hat{p} to have more “wiggle room”. On the other hand, in the context of EP and CTE the input observations of $\hat{q}(L, I, u)$ and $\overline{\text{CTE}}(L, I, u)$ are quite smooth in (L, i, u) and both GP kernel families perform equally well. The algorithms are implemented in `python 2.7`. We used “`GaussianProcessRegressor`” and “`SVM.SVC`” functions from `sklearn` library for GPR and SVM respectively. For LR and QR we used “`Logit`” and “`quantile_regression`” functions from `statsmodels` library.

Table 2: Parameters for the estimation of the admissible sets for each method. We use total simulation budget of 10^5 for all models except the Gold Standard.

Method	Budget ($M_a \times M_b$)	Further parameters
Gaussian Process (GPR)	2000×50	Matern-3/2 kernel
Logistic Regression (LR)	$10^5 \times 1$	Degree-2 polynomials
Parametric Density Fitting (PF)	2000×50	Truncated Gaussian, Matern-3/2 kernel
Empirical Percentile (EP)	1000×100	Squared exponential kernel
Conditional Tail Expectation (CTE)	1000×100	Squared exponential kernel
Quantile Regression (QR)	$10^5 \times 1$	Degree-4 polynomials
Support Vector Machine (SVM)	2000×50	$C = 1$, RBF kernel
Gold Standard (GS)	$10^6 \times 10^4$	budget = 10^{10}

5.2. Example 1: Microgrid with Stationary Net Demand. In this subsection, we assume time-homogeneous Ornstein-Uhlenbeck dynamics of the net demand process

$$(5.1) \quad dL(t) = -\lambda L(t)dt + \sigma dB(t) \implies L(t) = L(0)e^{-\lambda t} + \sigma \int_0^t e^{-\lambda(t-s)}dB(s),$$

where $(B(t))$ is a standard Brownian motion. This scenario reduces the complexity of learning the probability constraints since we need to estimate the admissible set $\mathcal{U}_0(\cdot)$ only once as a pre-processing

step before starting the approximate dynamic programming scheme for the continuation values. The simplified setting offers a good testbed to evaluate the performance of different admissible set estimation methods of Section 4; we show that the relative performance remains similar as we extend to more realistic dynamics in Section 5.3. For this example, we assume the mean reversion parameter $\lambda = 0.5$ and volatility $\sigma = 2$.

Figure 3a plots the resulting costs $\hat{V}_0(0, 5)$ versus the frequency of inadmissible decisions w_{freq} for different methods of Section 4. We show the results both for $p = 0.05$ (dark blue), and $p = 0.01$ (light grey) and benchmark both cases against the numerical gold standard. Since the probabilistic constraints form the crux of the problem, we require schemes to maintain $\hat{u} \in \mathcal{U}_n$ as much as possible, i.e., $w_{freq} \approx 0$. At $p = 5\%$, we observe 0.09%, 0.54% and 1.36% frequency of inadmissible decisions with logistic regression (LR), Gaussian process regression (GPR) and parametric density fitting (PF), respectively. Such accuracy might be deemed acceptable. However w_{freq} is much worse (as high as 8.4% with EP) for the other methods. While all the methods are a priori consistent, admissible set estimation via probability-based methods clearly seems to outperform quantile-based ones. Our experiments suggest that at low simulation budget, estimators of $p(\mathbf{x}, u)$ have significantly lower bias compared to estimators of $q(\mathbf{x}, u)$, thus partially explaining the difference. For a more stringent threshold $p = 1\%$, we find the cost of all the methods to increase, without significant difference in the frequency of inadmissible decisions w_{freq} . Indeed, Figure 3 illustrates the trade-off between lower costs and lower w_{freq} (i.e. more conservative estimate of the constraints).

Table 3 expands Figure 3 by also reporting the corresponding $\tilde{\mathcal{T}}$ statistic, the average inadmissibility margin w_{avm} and realized frequency of violations (i.e. blackouts) w_{rlzd} defined as:

$$(5.2) \quad w_{avm} := \frac{\sum_{n, m'} |\hat{u}_n(\mathbf{x}_n^{\hat{u}, m'}) - u_n^{\min}(\mathbf{x}_n^{\hat{u}, m'})| \mathbb{1}_{\hat{u}_n(\mathbf{x}_n^{\hat{u}, m'}) - u_n^{\min}(\mathbf{x}_n^{\hat{u}, m'}) < 0}}{\sum_{n, m'} \mathbb{1}_{\hat{u}_n(\mathbf{x}_n^{\hat{u}, m'}) - u_n^{\min}(\mathbf{x}_n^{\hat{u}, m'}) < 0}};$$

$$(5.3) \quad w_{rlzd} := \frac{1}{N \cdot M'} \sum_{n, m'} \mathbb{1}_{\sup_{s \in [t_n, t_{n+1}]} S^{m'}(s) > 0}.$$

We find the realized frequency of violations w_{rlzd} to be lowest for LR, GPR and PF. The average inadmissibility margin w_{avm} is also lowest for GPR and PF (the large value of w_{avm} for LR is attained in very small region as evident from $w_{freq} \approx 0$). The $\tilde{\mathcal{T}}$ statistic is negative for LR, GPR and PF and positive for the rest, meaning that all other methods fail to statistically respect the probability constraints when binding. Due to small frequency of inadmissible decisions w_{freq} , cost $\hat{V}_0(0, 5)$ similar to the numerical gold standard and negative test statistic $\tilde{\mathcal{T}}$, we recommend LR, GP and PF methods for the problem at hand.

Next, we test the sensitivity of the cost in terms of the probability threshold p (employing logistic regression $\hat{\mathcal{U}}_{LR}$) in Figure 3b. Increasing p decreases V as the set of admissible controls \mathcal{U} monotonically increases in p . For example, any admissible control at $p = 1\%$ threshold is also feasible for $p > 1\%$, thus the respective cost at 1% will be at least as much as at, say, 10% threshold.

As previously discussed, the constraint is binding for only approximately 10% of time-steps. In fact, that probability varies across the methods since the estimate of $\hat{\mathcal{U}}$ affects the choice of \hat{u}_n and ultimately the *distribution* of \hat{X}_n . Intuitively, the realized system states are driven by the estimates of the probabilistic constraints. Typically, more conservative estimates of \mathcal{U} will push $\hat{X}_{0:N}$ away from the “risky” regions. This is also confirmed in Figure 3 where as $p \rightarrow 1$, $w_{rlzd} \rightarrow 20\% = w_{bind}$ while in Table 3 $w_{bind} \simeq 10\%$.

The metrics w_{freq} (3.14), w_{bind} (3.18), w_{rlzd} (5.3) are closely linked. As the inadmissible decisions can occur only when the constraint is binding, $u^{\min} > 0$, we expect $w_{freq} \leq w_{bind}$ and $w_{freq} \approx w_{bind}$ for a method with a bias in overestimating the admissible set (e.g. $\mathcal{X}^{a,EP}(u) \supset \mathcal{X}^{a,GS}(u) \forall u \in \mathcal{W}$). The realized violations (blackouts) w_{rlzd} can be represented as a sum of three:

$$w_{rlzd} = p_1 w_{freq} + p_2 (w_{bind} - w_{freq}) + p_3 (1 - w_{bind}), \quad p_1 + p_2 + p_3 = 1,$$

where the weights p_1, p_2, p_3 depend on the distribution of the controlled trajectories. The first term represents the instances when the constraint is binding but the controller chooses an inadmissible control (i.e. mis-estimates $\hat{\mathcal{U}}$). The second term represents instances when the constraint is binding and correctly estimated, but due to random shocks violations take place (with a conditional frequency below the specified $p = 0.05$). The last term represents instances when the constraint is not binding but some violations still occur with the intrinsic conditional frequency strictly less than p . Note that due to $w_{bind} \ll 1$, most of the violations are of the latter type, i.e. take place when $u^* = 0$ and the conditional violation probability is below p . We illustrate these scenarios in Figure 3c using the LR model. Thus, the first term counts the instances when violations occur at the same time as controller makes an inadmissible decision (circle encircling triangle), the second term counts the triangles when $I \approx 0$, and the third term the triangles in the grey region where the constraint is not binding (violations when $u^{\min} = 0$).

Although we observed poor performance of quantile based methods, asymptotically (with respect to the simulation budget) we expect them to perform similar to the probability based methods. As an example, in Appendix A Table 6, we present the performance of SVM for thresholds $p = 5\%$ and $p = 1\%$ with increasing budget. For $p = 5\%$ and by increasing the simulation budget from 10^5 to 10^8 , we find the frequency of inadmissible decisions w_{freq} to drop from 5.93% to 1.5%, average inadmissibility amount w_{avm} from 0.78 kW to 0.27 kW, frequency of realized blackouts w_{rlzd} from 2.80% to 0.30% and the test statistic which rejected the method at 10^5 simulation budget ($\mathcal{T} \gg 0$) suggests to accept it ($\mathcal{T} \ll 0$) at 10^8 simulation budget. We observe similar behavior at $p = 1\%$. The main challenge with quantile-based methods is the underlying bias in learning $q(\mathbf{x}, u)$. This bias is known to converge to zero slowly, necessitating a relatively large M_b . A glimpse of this can be observed in Appendix A where SVM does not perform adequately all the way up to total budget of 10^8 . Increasing M_b with M_a fixed offers only a limited improvement that tapers off quickly, as was observed previously in [26]. Keeping a constant budget, a high M_b forces a low M_a which offsets these performance gains as the regression is not able to properly explore the space.

There are many techniques to improve \bar{q} , for example variance reduction tools. We have experimented with variance reduction for the SVM method by using antithetic variables. Although the results were marginally better, they were not meaningfully different from those reported in Appendix A. A completely different way to improve EP and SVM could be to judiciously choose the simulation design \mathcal{D}^a . Paper [26] explores various approaches to build simulation designs for estimating the continuation value function and documents their strong impact on performance. We anticipate that a similar strategy for \mathcal{D}^a may improve convergence.

Conservative estimators for \mathcal{U} . Algorithms for SCPC are expected to respect the probabilistic constraints, so that it is critical to minimize the occurrence of inadmissible decisions. As discussed in Section 4, one way to raise the statistical guarantee for admissibility of $\hat{\mathcal{U}}$ is by adding a margin of error $\xi(\mathbf{x}, u)$. This yields a more conservative (i.e. smaller) $\hat{\mathcal{U}}$ and hence lowers w_{freq} . In Table 4 we examine three scenarios for $\xi(\mathbf{x}, u)$ sorted from least to most conservative (in all cases we maintain $p = 5\%$):

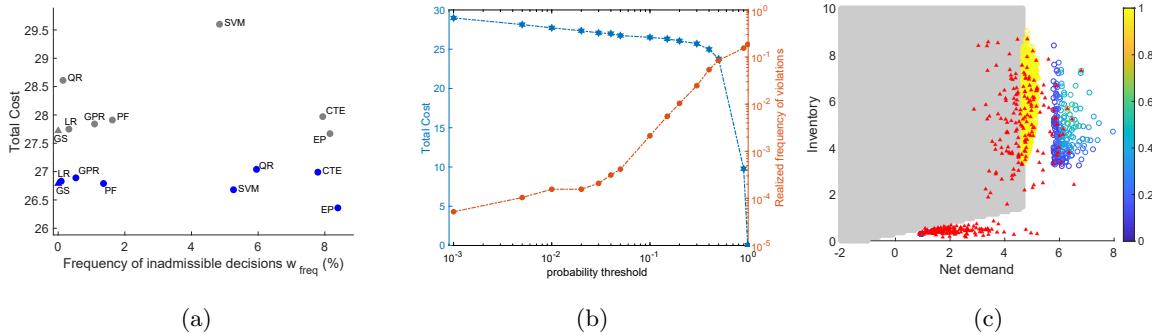


Fig. 3: *Left panel:* Trade-off between cost $\hat{V}_0(0, 5)$ and frequency of inadmissible decisions w_{freq} for the stationary model. Dark blue points correspond to $p = 5\%$ probability constraint threshold and light grey ones to $p = 1\%$. *Center:* Total cost $\hat{V}_0(0, 5)$ (left axis, blue stars) and realized frequency of violations w_{rlzd} (right axis, red circles) as functions of p employing the LR model. *Right:* Locations (L, I) of realized violations $\sup_{s \in [t_n, t_{n+1})} S^{m'}(s) > 0$ (red triangles), inadmissible decisions $\hat{u}(n, \mathbf{x}_n^{\hat{u}, m'}) - u_n^{\min}(\mathbf{x}_n^{\hat{u}, m'}) < 0$ (circles with color representing the inadmissibility margin) on 5000 out-of-sample simulations using LR model. The constraint is binding in the white region and is not binding in the grey region.

Table 3: Cost of running the microgrid $\hat{V}_0(0, 5)$, frequency of inadmissible decisions w_{freq} , average inadmissibility margin w_{avm} , realized frequency of violations (i.e. blackouts) w_{rlzd} , test statistic $\tilde{\mathcal{T}}$ and frequency of binding constraint w_{bind} for the example in Section 5.2.

Method	$\hat{V}_0(0, 5)$ (\$)	w_{freq} (%)	w_{avm} (kW)	w_{rlzd} (%)	$\tilde{\mathcal{T}}$	w_{bind} (%)
GS	26.79	0.00	0.00	0.37	-	-
LR	26.83	0.09	0.82	0.03	-125	8.69
GPR	26.89	0.53	0.16	0.11	-98	8.10
PF	26.79	1.36	0.27	0.21	-69	8.51
SVM	26.68	5.26	0.55	1.83	388	9.67
QR	27.04	5.95	0.33	0.98	145	9.49
CTE	26.99	7.79	0.43	1.63	320	9.93
EP	26.36	8.39	0.49	1.98	403	10.45

- Scenario 1: unadjusted $\xi = 0\%$ (same as Table 3);
- Scenario 2: $\xi^{(\rho)}(\mathbf{x}, u)$ at 95% confidence level, $z_\rho = 1.96$;
- Scenario 3: fixed $\xi = 4\%$, which is equivalent to lowering the violation threshold to $p - \xi = 1\%$.

Table 4 confirms the intuition that the frequency of inadmissible decisions w_{freq} should be decreasing from scenario 1 to 3. This is illustrated in Figure 4 that shows how the minimum admissible control is affected by $\xi(\mathbf{x}, u)$. Although adding a margin of error does lower w_{freq} , this mechanism does not really alter the relative performance of the different methods. Thus, for all three scenarios, we find SVM, CTE and EP to give unreliable estimates of \mathcal{U} (since $\tilde{\mathcal{T}} \gg 0$). An exception is QR which yields high w_{freq} for $\xi = 0$ but does become acceptable ($\tilde{\mathcal{T}} < 0$) in scenario 3. In contrast, LR, GPR and PF perform

well throughout. Table 7 in Appendix B provides additional results as we vary the confidence level to $\rho = 90\%, 99\%$ and 99.95% , with the same general conclusions. (Observe that a w_{freq} very close to zero likely implies that $\hat{\mathcal{U}} \subset \mathcal{U}$ is strictly smaller and the controller is so conservative as to rule out some admissible actions.)

We generally expect the ultimate cost $\hat{V}_0(0, 5)$ to increase as $\hat{\mathcal{U}}$ becomes more conservative, see the estimated \hat{V} 's across each row of Table 4. The increase in costs arises due to two factors: when the diesel generator is started sooner (due to $u = 0$ becoming inadmissible as ξ is raised), and the higher level of \hat{u} once the diesel is ON. This can be seen in Figure 4 where in Scenarios 2 and 3 the controller switches the generator at a lower net demand and once the diesel is running picks a higher power output ($\hat{u}^{\min}(\cdot, I; p = 5\%, \xi) - \hat{u}^{\min}(\cdot, I; p = 5\%, \xi = 0) > 0$). We stress that the link between $\hat{\mathcal{U}}$ and \hat{V} is complicated by the fact that as $\hat{\mathcal{U}}$ changes, so does the distribution of the controlled paths. So for example in Table 4 the cost for QR falls in Scenario 2, although it remains within two Monte Carlo standard errors.

Table 4: Impact of margin of error ξ on the estimated cost of running the microgrid $\hat{V}_0(0, 5)$, frequency of inadmissible decisions w_{freq} , and test statistic $\tilde{\mathcal{T}}$ from (3.17). The probabilistic constraint is $p = 5\%$.

Method	$\xi = 0\%$			$\xi^{(0.95)}(\mathbf{x}, u)$			$\xi = 4\%$		
	$\hat{V}_0(0, 5)$	w_{freq}	$\tilde{\mathcal{T}}$	$\hat{V}_0(0, 5)$	w_{freq}	$\tilde{\mathcal{T}}$	$\hat{V}_0(0, 5)$	w_{freq}	$\tilde{\mathcal{T}}$
GS	26.79	0.00	-	-	-	-	-	-	-
LR	26.83	0.09	-125	26.95	0.08	-124	27.86	0.04	-112
GPR	26.89	0.53	-98	28.00	0.01	-110	28.12	0.00	-107
PF	26.79	1.36	-69	-	-	-	27.91	0.44	-96
SVM	26.68	5.26	388	29.65	3.41	225	29.60	3.41	225
QR	27.04	5.95	145	26.89	5.17	72	28.61	0.00	-117
CTE	26.99	7.79	320	27.36	7.52	274	28.44	6.83	248
ER	26.36	8.39	403	26.97	7.78	225	28.13	7.08	283

Take-aways. Our experiments demonstrate the following: (i) To accurately estimate admissible sets of the form (2.3) we recommend to use LR, GPR or PF which all model the underlying probability of violations $p(\mathbf{x}, u)$. Although asymptotically equivalent, the approach of quantile estimation leads to poor estimates $\hat{\mathcal{U}}_n$ for practical budgets. (ii) Frequency of inadmissible decisions can be partly controlled by using conservative estimates $\hat{\mathcal{U}}_n^\xi$ at the expense of higher costs. However, even a conservative $\hat{\mathcal{U}}^\xi$ fails to make quantile-based methods acceptable, except for QR. (iii) For a new application, our suggested approach is to first evaluate the test statistic $\tilde{\mathcal{T}}$ at $\xi = 0\%$ using one of the recommended methods. Depending on how close is $\tilde{\mathcal{T}}$ to zero, one can then adjust $\hat{\mathcal{U}}$ via ξ to improve the statistical guarantees on the frequency of inadmissible decisions w_{freq} .

5.3. Example 2: Microgrid with seasonal demand. Unlike the previous example, where we assumed time-homogeneous net demand, in practice there is seasonality: during the day renewable generation is high and net demand is often negative; during morning/evening demand exceeds supply making $L(t) > 0$. To incorporate this seasonality we use time-dependent Ornstein Uhlenbeck process (see [17]

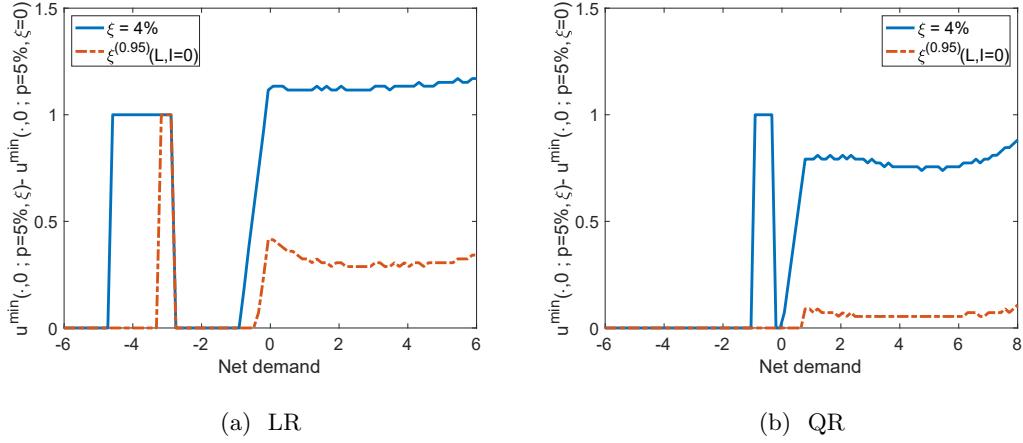


Fig. 4: Impact of the margin of error $\xi(\cdot, \cdot)$ on minimum admissible control \hat{u}^{\min} . We plot the difference between minimum admissible control for scenario 2 ($\hat{u}^{\min}(\cdot, I; \xi^{(0.95)}(L, I))$) and scenario 3 ($\hat{u}^{\min}(\cdot, I; \xi = 4\%)$) with respect to scenario 1 ($\hat{u}^{\min}(\cdot, I; \xi = 0\%)$) using LR (left panel) and QR (right panel) models.

for a similar microgrid control problem):

$$(5.4) \quad dL(t) = \left[\frac{\partial \mu}{\partial t}(t) + \lambda(\mu(t) - L(t)) \right] dt + \sigma(t)dB(t).$$

Here, λ represents the speed of mean reversion towards the seasonal mean $\mu(t)$, while $\sigma(t)$ represents the time-varying volatility. Using Itô's lemma and integration by parts one can prove that

$$L(t) = \mu(t) + e^{-\lambda t} (L(0) - \mu(0)) + \int_0^t e^{-\lambda(t-s)} \sigma(s) dB(s).$$

Thus,

$$\mathbb{E}[L(t)] = \mu(t) + e^{-\lambda t} (L(0) - \mu(0)).$$

We calibrate $\mu(t)$ and $\sigma(t)$ in (5.4) using iterative methodology described in [17] and the data from a solar-powered microgrid in Huatacondo, Chile¹. Specifically, we compute the mean and variance of the residual demand over 24 hours at 15-minute intervals using data from Spring 2014, i.e. compute $\{\mu_1, \mu_2, \dots, \mu_{96}\}$ and $\{\sigma_1, \sigma_2, \dots, \sigma_{96}\}$. The estimated $\mu(t)$ can be seen in the left panel of Figure 5 that plots the empirical average of $L(t)$. As expected, during the day, i.e., $t \in [12, 20]$ (noon-8:00 pm), the expected net-demand is negative ($\mu(t) < 0$) while it is positive ($\mu(t) > 0$) in the morning and during the night. The volatility $\sigma(t)$ is higher during the day due to the intermittent and unpredictable nature of solar irradiance. The mean reversion parameter was estimated to be $\lambda = 0.3416$.

To visualize the interplay of the net demand, inventory and optimal control, the left panel of Figure 5 presents the average trajectories of the three processes over 48 hours. During the morning hours when the demand $L(t)$ is high and the battery is empty, the controller uses the diesel generator. During the day when the renewable output is high and $L(t)$ is negative, the controller switches off the diesel and the battery charges itself. However, the non-trivial region is when the average net-demand changes

¹<https://microgrid-symposiums.org/microgrid-examples-and-demonstrations/huatacondo-microgrid/>

sign, either from positive to negative around noon or negative to positive in the evening. During the former time interval, the optimal control process is in $\{0, 1\}$ (recall that minimum diesel output is 1). Similarly, during the evening when the net demand becomes positive (as the renewable output declines), the controller quickly ramps up the diesel to match $L(t) \gg 0$. The right panel of Figure 5 repeats the average control and inventory curves, also showing their 2-standard deviation bands (in terms of the out-of-sample trajectories of $\hat{L}_{0:T}^{\hat{u}}$). As expected, the time periods around ramp-up or ramp-down of the diesel generator is when \hat{u}_n experiences the greatest path-dependency and dispersion and differs most from the demand curve.

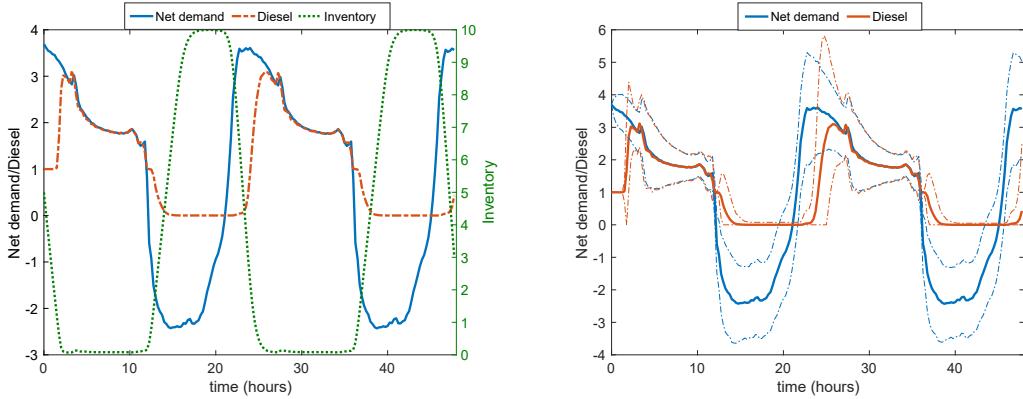


Fig. 5: Model parameters, average trajectory of the state variables, control and their variance. Left panel: Average values of net demand $\frac{1}{M'} \sum_{m'=1}^{M'} L_n^{\hat{u}, m'}$, inventory $\frac{1}{M'} \sum_{m'=1}^{M'} I_n^{\hat{u}, m'}$ and optimal control (diesel) $\frac{1}{M'} \sum_{m'=1}^{M'} \hat{u}_n^{m'}$ processes using the gold standard strategy. Right panel: 95% confidence bands for net demand $L_n^{\hat{u}}$ and realized optimal diesel control \hat{u}_n . Net demand and diesel output is measured in kW and Inventory in kWh.

Comparing Table 5, which lists the estimated cost $\hat{V}_0(\mu(0), 5)$ along with related statistics, with Figure 3 indicates that incorporating seasonal net-demand process does not change the relative order of performance between the methods. The cost goes up as the diesel generator has to be used throughout the mornings and the evenings to match demand.

As in the previous example, the performance of LR, GPR and PF almost matches the gold standard despite significantly lower simulation budget. In this setting the constraint is binding approximately 45% of the time (except for GPR and PF where it is 30% and 25% of the time). Frequency of inadmissible decisions w_{freq} is 0.03% for LR, 1.17% for GPR, and 0.02% for PF. In contrast w_{freq} is 43% for QR, 22% for EP, 43% for SVM and 22% for CTE, implying that all these schemes are highly unreliable for learning $\hat{\mathcal{U}}$. The average inadmissibility margin w_{avm} is also significantly lower for GPR (0.14 kW) and PF (0.26 kW) compared to the rest of the methods. Here again we observe larger inadmissibility margin and very low frequency of inadmissible decisions for logistic regression. Similar behavior is also evident for the test statistic $\tilde{\mathcal{T}}$ and realized frequency of violations w_{rlzd} .

To illustrate the typical behavior over a trajectory, Figure 6 plots the average control $Ave(\hat{u}_n) := \frac{1}{M'} \sum_{m=1}^{M'} \hat{u}_n(\mathbf{x}_n^{\hat{u}, m'})$ corresponding to different methods and the average minimum admissible control $Ave(u_n^{\min}) := \frac{1}{M'} \sum_{m'=1}^{M'} u_n^{\min}(\mathbf{x}_n^{\hat{u}, m'})$ computed using the gold standard. Notice that the latter is dependent upon the controlled trajectories $\mathbf{x}_n^{\hat{u}}$, resulting in different $Ave(u_n^{\min})$ across methods. We expect $Ave(\hat{u}_n)$ above $Ave(u_n^{\min})$ if a given method does not violate the constraint most of the time. This is true

Table 5: Cost of running the microgrid $\hat{V}_0(\mu(0), 5)$, frequency of inadmissible decisions w_{freq} , average inadmissibility margin w_{avm} , realized frequency of violations w_{rlzd} and frequency of the constraint being binding w_{bind} for the case study in Section 5.3.

Method	$\hat{V}_0(\mu(0), 5)$	w_{freq} (%)	w_{avm} (kW)	w_{rlzd} (%)	$\tilde{\mathcal{T}}$	w_{bind} (%)
GS	53.38	0	0	0.30	-	-
LR	53.78	0.03	0.79	0.01	-301	45.2
GPR	54.04	1.17	0.14	0.19	-220	31.0
PF	54.55	0.02	0.26	0.01	-226	25.7
SVM	40.52	43.37	0.91	43.37	5,306	46.4
QR	52.56	42.87	0.28	38.41	4,772	46.3
CTE	53.02	21.62	0.21	10.43	1,079	46.0
EP	52.82	21.91	0.23	11.57	1,227	46.1

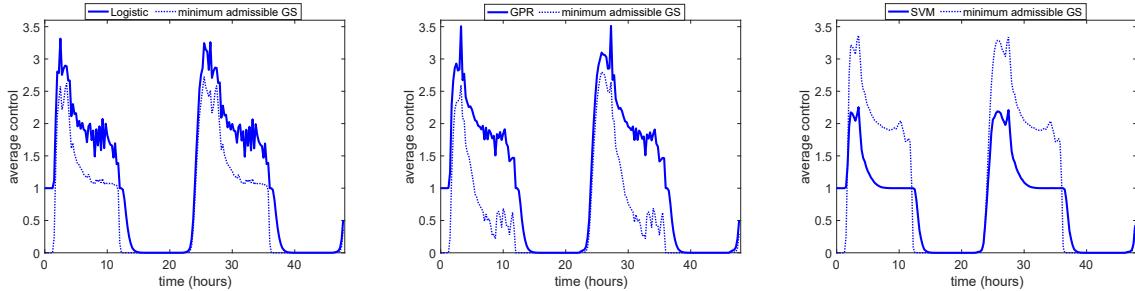


Fig. 6: Average control $Ave(\hat{u}_n)$ for LR, GPR and SVM and the average minimum admissible control $Ave(u_n^{\min})$ using Gold Standard across forward controlled trajectories.

for LR and GPR, but SVM quite obviously fails, as the dashed line in the leftmost panel of Figure 6 is significantly higher than the solid line at numerous time steps. Furthermore, the conservative nature of GPR is reflected in the large difference between the average minimum admissible control and the average optimal control. This is also evident through $w_{bind} \approx 30\%$ for GPR compared to approximately 45% for the rest of the methods.

6. Conclusion. We developed a statistical learning framework to solve stochastic optimal control problems with local probabilistic constraints. The key objective of our algorithm is to efficiently estimate the set of admissible controls $\mathcal{U}(\cdot)$ and the continuation value function $\mathcal{C}(\cdot, \cdot)$ covering a general formulation of the state process dynamics and rewards. Since SCPC problems require estimating the admissible set repeatedly during the backward induction, we use regression based functional representation of $\mathbf{x} \mapsto \mathcal{U}(\mathbf{x})$. This perspective also provides a natural way of uncertainty quantification for admissibility, in particular offering conservative estimates that bring statistical guarantees regarding $\hat{\mathcal{U}}$. At the same time, our dynamic emulation algorithm allows parallel computation of \mathcal{U} and \mathcal{C} for additional computational efficiency.

Thanks to the plug-and-play functionality of the dynamic emulation algorithm, it was straightforward

to test a large variety of schemes for learning \mathcal{U} . Our numerical results suggest that estimating probabilistic constraints via logistic regression, Gaussian process smoothing and parametric density fitting is more accurate than estimating the corresponding quantile (empirical ranking, SVM or quantile regression). A future line of research would be to additionally parametrize (e.g. using another GP model) the optimal control map $\mathbf{x} \mapsto \hat{u}_n(\mathbf{x})$ [14] which would speed-up the algorithm in the context of continuous action spaces. Another direction would be to extend the one-dimensional control framework described in this paper to a multi-dimensional setting. In reference to the microgrid example, multi-dimensional control will allow us to control the diesel output and the demand response or to control output from multiple dispatchable generators.

Appendix A. Effect of Simulation Budget ($M_b \times M_a$).

Table 6: Impact of simulation budget on performance of SVM for the case study in Section 5.2 and probability thresholds $p = 5\%$ and $p = 1\%$. The reported values are averages over 10 runs of each scheme. The total simulation budget is divided into batch size M_b and number of design sites M_a . For total budget 10^5 : $(M_b, M_a) = (100, 1000)$; for 10^6 : $(M_b, M_a) = (500, 2000)$; for 10^7 : $(M_b, M_a) = (2000, 5000)$; for 10^8 : $(M_b, M_a) = (10000, 10000)$.

p	Budget	$\hat{V}_0(0, 5)$ (\$)	w_{freq} (%)	w_{avm} (kW)	w_{rlzd} (%)	$\tilde{\mathcal{T}}$	w_{bind} (%)
5%	10^5	26.38	5.93	0.78	2.80	665	9.73
	10^6	26.55	5.28	0.55	1.84	386	9.77
	10^7	26.68	4.96	0.53	1.64	330	9.75
	10^8	26.79	1.50	0.27	0.30	-51	9.22
1%	10^5	28.32	6.63	0.93	2.43	1,460	9.87
	10^6	28.26	5.17	0.66	1.09	631	9.56
	10^7	28.52	0.55	0.24	0.03	-39	8.78
	10^8	28.41	0.15	0.22	0.01	-51	8.82

Appendix B. Effect of Adaptive Margin of Error Level ρ .

Table 7: Impact of conservative $\mathcal{U}^{(\rho)}$ estimators for the case study in Section 5.2. The probabilistic constraint is set at $p = 5\%$.

Method	$\rho = 90\%$			$\rho = 99\%$			$\rho = 99.95\%$		
	$\hat{V}_0(0, 5)$	w_{freq}	w_{rlzd}	$\hat{V}_0(0, 5)$	w_{freq}	w_{rlzd}	$\hat{V}_0(0, 5)$	w_{freq}	w_{rlzd}
LR	26.74	0.090	0.034	26.87	0.085	0.032	27.04	0.085	0.026
GPR	27.34	0.012	0.055	28.06	0.007	0.037	28.06	0.005	0.029
SVM	27.35	4.975	1.732	29.20	3.481	1.117	29.72	3.395	1.088
QR	27.20	5.373	0.793	27.18	4.880	0.676	27.04	4.409	0.591
CTE	27.93	7.158	1.153	28.31	6.766	0.888	28.61	6.163	0.714
EP	26.78	7.990	1.497	27.17	7.629	1.183	27.96	7.102	0.956

REFERENCES

- [1] S. AHMED AND A. SHAPIRO, *Solving chance-constrained stochastic programs via sampling and integer programming*, in INFORMS TutORials in Operations Research, Z.-L. Chen and S. Raghavan, eds., INFORMS, 2014, pp. 261–269.
- [2] J.-C. ALAIS, P. CARPENTIER, AND M. DE LARA, *Multi-usage hydropower single dam management: chance-constrained optimization and stochastic viability*, Energy Systems, 8 (2017), pp. 7–30.
- [3] L. ANDRIEU, R. HENRION, AND W. RÖMISCH, *A model for dynamic chance constraints in hydro power reservoir management*, European Journal of Operational Research, 207 (2010), pp. 579–589.
- [4] A. BALATA AND J. PALCZEWSKI, *Regress-Later Monte Carlo for Optimal Inventory Control with applications in energy*, arXiv:1703.06461, (2017).
- [5] A. BALATA AND J. PALCZEWSKI, *Regress-Later Monte Carlo for optimal control of Markov processes*, arXiv:1703.09705, (2018).
- [6] L. BLACKMORE, M. ONO, A. BEKTASSOV, AND B. C. WILLIAMS, *A probabilistic particle-control approximation of chance-constrained stochastic predictive control*, IEEE Transactions on Robotics, 26 (2010), pp. 502–517.
- [7] L. BLACKMORE, M. ONO, AND B. C. WILLIAMS, *Chance-constrained optimal path planning with obstacles*, IEEE Transactions on Robotics, 27 (2011), pp. 1080–1094.
- [8] A. BOOGERT AND C. DE JONG, *Gas storage valuation using a Monte Carlo method*, The Journal of Derivatives, 15 (2008), pp. 81–98.
- [9] A. BOOGERT AND C. DE JONG, *Gas storage valuation using a multi-factor price process*, Journal of Energy Markets, 4 (2011), pp. 29–52.
- [10] B. BOUCHARD AND X. WARIN, *Monte Carlo valuation of American options: Facts and new algorithms to improve existing methods*, in Numerical Methods in Finance: Bordeaux, June 2010, R. A. Carmona, P. Del Moral, P. Hu, and N. Oudjane, eds., Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 215–255.
- [11] G. C. CALAFIORE AND M. C. CAMPI, *The scenario approach to robust control design*, IEEE Transactions on Automatic Control, 51 (2006), pp. 742–753.
- [12] M. C. CAMPI AND G. C. CALAFIORE, *Notes on the scenario design approach*, IEEE Transactions on Automatic Control, 54 (2009), pp. 382–385.
- [13] R. CARMONA AND M. LUDKOVSKI, *Valuation of energy storage: an optimal switching approach*, Quantitative Finance, 10 (2010), pp. 359–374.
- [14] M. P. DEISENROTH, C. E. RASMUSSEN, AND J. PETERS, *Gaussian process dynamic programming*, Neurocomputing, 72 (2009), pp. 1508 – 1524.
- [15] L. DOYEN AND M. D. LARA, *Stochastic viability and dynamic programming*, Systems & Control Letters, 59 (2010), pp. 629 – 634.
- [16] A. GELETU, M. KLÖPPEL, H. ZHANG, AND P. LI, *Advances and applications of chance-constrained approaches to systems optimisation under uncertainty*, International Journal of Systems Science, 44 (2013), pp. 1209–1232.
- [17] B. HEYMANN, J. F. BONNANS, F. SILVA, AND G. JIMENEZ, *A stochastic continuous time model for microgrid energy management*, in 2016 European Control Conference (ECC), June 2016, pp. 2084–2089.
- [18] L. JANSON, E. SCHMERLING, AND M. PAVONE, *Monte Carlo motion planning for robot trajectory optimization under uncertainty*, in Robotics Research: Volume 2, A. Bicchi and W. Burgard, eds., Springer International Publishing, Cham, 2017, pp. 343–361.
- [19] Y. JIAO, O. KLOPFENSTEIN, AND P. TANKOV, *Hedging under multiple risk constraints*, Finance and Stochastics, 21 (2017), pp. 361–396.
- [20] I. KARATZAS AND S. E. SHREVE, *Brownian Motion*, Springer New York, New York, NY, 1998.
- [21] N. LANGRENÉ, T. TARNOPOLSKAYA, W. CHEN, Z. ZHU, AND M. COOKSEY, *New regression Monte Carlo methods for high-dimensional real options problems in minerals industry*, in 21st International Congress on Modelling and Simulation, 2015.
- [22] N. LANGRENÉ AND X. WARIN, *Fast and stable multivariate kernel density estimation by fast sum updating*, Journal of Computational and Graphical Statistics, 28 (2019), pp. 596–608.
- [23] C. LIU, X. WANG, Y. ZOU, H. ZHANG, AND W. ZHANG, *A probabilistic chance-constrained day-ahead scheduling model for grid-connected microgrid*, in 2017 North American Power Symposium (NAPS), 2017, pp. 1–6.
- [24] F. A. LONGSTAFF AND E. S. SCHWARTZ, *Valuing American options by simulation: A simple least-squares approach*, The Review of Financial Studies, 14 (2001), pp. 113–147.
- [25] M. LUDKOVSKI, *Kriging metamodels and experimental design for Bermudan option pricing*, Journal of Computational

Finance, 22 (2018), pp. 37–77.

[26] M. LUDKOVSKI AND A. MAHESHWARI, *Simulation methods for stochastic storage problems: A statistical learning perspective*, Energy Systems, 11 (2020), pp. 377–415.

[27] J. LUEDTKE AND S. AHMED, *A sample approximation approach for optimization with probabilistic constraints*, SIAM Journal on Optimization, 19 (2008), pp. 674–699.

[28] A. NEMIROVSKI AND A. SHAPIRO, *Scenario approximations of chance constraints*, in Probabilistic and Randomized Methods for Design under Uncertainty, G. Calafiore and F. Dabbene, eds., Springer London, London, 2006, pp. 3–47.

[29] A. NEMIROVSKI AND A. SHAPIRO, *Convex approximations of chance constrained programs*, SIAM Journal on Optimization, 17 (2007), pp. 969–996.

[30] M. ONO, M. PAVONE, Y. KUWATA, AND J. BALARAM, *Chance-constrained dynamic programming with application to risk-aware robotic space exploration*, Autonomous Robots, 39 (2015), pp. 555–571.

[31] A. PEÑA-ORDIERES, J. R. LUEDTKE, AND A. WÄCHTER, *Solving chance-constrained problems via a smooth sample-based nonlinear approximation*, arXiv:1905.07377, (2019).

[32] A. PRÉKOPA AND T. SZÁNTAI, *Flood control reservoir system design using stochastic programming*, in Mathematical Programming in Use, M. L. Balinski and C. Lemarechal, eds., Springer Berlin Heidelberg, Berlin, Heidelberg, 1978, pp. 138–151.

[33] S. A. P. QUINTERO, M. LUDKOVSKI, AND J. P. HESPAÑHA, *Stochastic optimal coordination of small UAVs for target tracking using regression-based dynamic programming*, Journal of Intelligent & Robotic Systems, 82 (2016), pp. 135–162.

[34] A. RAVICHANDRAN, S. SIROUSPOUR, P. MALYSZ, AND A. EMADI, *A chance-constraints-based control strategy for microgrids with energy storage and integrated electric vehicles*, IEEE Transactions on Smart Grid, 9 (2018), pp. 346–359.

[35] O. ROUSTANT, D. GINSBOURGER, AND Y. DEVILLE, *DiceKriging, DiceOptim: Two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization*, Journal of Statistical Software, 51 (2012), pp. 1–55.

[36] J. N. TSITSIKLIS AND B. VAN ROY, *Regression methods for pricing complex American-style options*, IEEE Transactions on Neural Networks, 12 (2001), pp. 694–703.

[37] W. VAN ACKOOIJ, R. HENRION, A. MÖLLER, AND R. ZORGATI, *Joint chance constrained programming for hydro reservoir management*, Optimization and Engineering, 15 (2014), pp. 509–531.

[38] W. VAN-ACKOOIJ AND X. WARIN, *On conditional cuts for Stochastic Dual Dynamic Programming*, ArXiv e-prints, (2017).

[39] X. WARIN, *Gas storage hedging*, in Numerical Methods in Finance: Bordeaux, June 2010, R. A. Carmona, P. Del Moral, P. Hu, and N. Oudjane, eds., Springer, Berlin, Heidelberg, 2012, pp. 421–445.

[40] J. XU AND J. S. LONG, *Confidence intervals for predicted outcomes in regression models for categorical outcomes*, Stata Journal, 5 (2005), pp. 537–559.