Learning for Constrained Optimization:
Identifying Optimal Active Constraint Sets

Sidhant Misra[1], Line Roald[2], and Yeesian Ng[3]

[1]sidhant@lanl.gov
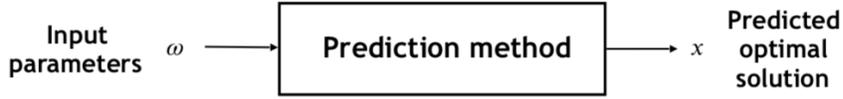[2]roald@wisc.edu
[3]yeesian@mit.edu

January 18, 2019

Abstract

In many engineered systems, optimization is used for decision making at time-scales ranging from real-time operation to long-term planning. This process often involves solving similar optimization problems over and over again with slightly modified input parameters, often under tight latency requirements. We consider the problem of using the information available through this repeated solution process to directly learn a model of the optimal solution as a function of the input parameters, thus reducing the need to solve computationally expensive large-scale parametric programs in real time. Our proposed method is based on learning relevant sets of active constraints, from which the optimal solution can be obtained efficiently. Using active sets as features preserves information about the physics of the system, enables interpretable models, accounts for relevant safety constraints, and is easy to represent and encode. However, the total number of active sets is also very large, as it grows exponentially with system size. The key contribution of this paper is a streaming algorithm that learns the relevant active sets from training samples consisting of the input parameters and the corresponding optimal solution, without any assumptions on the problem structure. The algorithm comes with theoretical performance guarantees, and is known to converge fast for problem instances with a small number of relevant active sets. It can thus be used to establish the practicability of the learning method. Through extensive experiments on the Optimal Power Flow problem, we observe that often only a few active sets are relevant in practice, suggesting that the active sets is the appropriate level of abstraction for a learning algorithm to target.
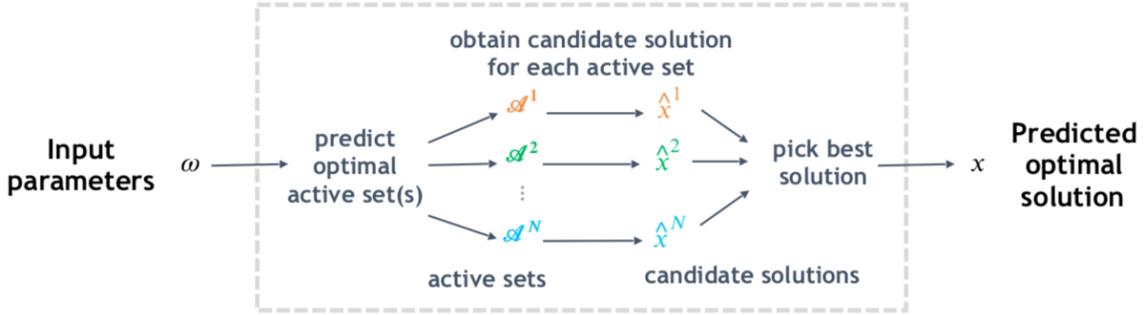
1 Introduction

The use of optimization methods to improve economic efficiency and enhance system performance is growing rapidly in engineering applications. Examples include the operation and planning of infrastructure like electric power grids, gas transmission systems, water and heating networks and transportation networks. Many of these systems operate in highly variable conditions, due to developments such as renewable energy integration in power systems and uncertain demand-side behaviour. To maintain safe and efficient operations, new optimized set-points are obtained by resolving the optimization problems as system conditions (and corresponding system parameters) change. The optimization problems exhibit two key features:

1. There exists a mathematical model of the physical system, including objective function and technical constraints. For optimization purposes, this model is assumed to be fully known except for those input parameters whose values are obtained in real-time, giving rise to parametric programming problems.

(a) Predicting optimal solution from input parameters.

(b) Predicting optimal solution using active sets as an intermediate step.

Figure 1: An illustration of our approach. Figure 1a depicts the general approach for predicting optimal solutions from input parameters. Figure 1b depicts our approach using active sets as an intermediate step.

2. Constraint satisfaction corresponds to enforcing operational safety limits and adhering to physical laws, and is of paramount importance.

As the optimization problems are solved over and over again in response to changing conditions, we build a rich history of input parameters and their corresponding optimal solutions. However, solving these parametric programs in real-time for realistically sized systems can be prohibitively difficult, owing to the combination of complex physical laws and tight latency requirements. This limits the frequency with which the system set-points can be updated.

The goal of this paper is to develop a framework that uses machine learning methods to enable these parametric programs to be solved more efficiently. An intuitive approach to apply machine learning would be to use data from existing solutions, and directly learn the optimal solution from the input data, as illustrated in Figure 1a. Previous literature on representing the mapping from parameters to optimal solutions in parametric programs often used reinforcement learning (RL), where the underlying physical constraints of the problem are assumed to be unknown. We refer the reader to Lillicrap et al. (2015), Mnih et al. (2015), Van Hasselt et al. (2016) for examples of this approach using neural networks. However, the neural networks mostly provide only an approximate model of the physical constraints, which may result in solutions that violate important system constraints. This has prompted a literature on "safety problems" (Amodei et al. 2016, García and Fernández 2015). Approaches to mitigate those safety concerns include modifying the exploration procedure by restricting it to a set of safe policies (Moldovan and Abbeel 2012), re-projecting the decisions made by the neural network into a known set of safe decisions (Achiam et al. 2017), or include an outer safety loop which overrides potentially unsafe decisions by the neural net (Fisac et al. 2017). Nonetheless, it remains difficult to enforce guarantees on behavioral constraints in the resulting policy (Moldovan and Abbeel 2012, Lu et al. 2017, Achiam et al. 2017) and verify that the resulting neural network policy provides a safe policy (Huang et al. 2017). Approaches that restrict or limit the neural net also may produce sub-optimal solutions.

A main drawback of these existing learning methods when applied in the context of engineered systems is hence their inability to enforce the constraints in the optimization accurately. One important reason for this drawback is that the methods are unable to leverage pre-existing knowledge about the mathematical form of the optimization problem. For best results, the

learning methods used should exploit any known structure of the optimization problem in consideration.

In this paper, we take a different approach to learning the optimal solution. Instead of directly learning the mapping from the input data to the optimal solution, we introduce the *active set* as an intermediate step in our learning procedure. To the best of our knowledge, this has never been explored before in the literature. The *optimal active set* in an optimization problem corresponds to the set of constraints that are satisfied with equality in the corresponding optimal solution. Learning the optimal active set arguably constitute the *right* level of abstraction for applying machine learning algorithm to optimization problems for several reasons:

(i) From an optimization perspective, the optimal active set in an optimization problem constitutes the minimal information required to recover the optimal solution. All optimization algorithms such as interior point methods, active set methods, and the simplex algorithm can see significant speed-up from the knowledge of the active set at optimality. In some cases, such as in linear programming (LP), the knowledge of the optimal active set can simplify the problem to the extent that it can be solved analytically, or in a single iteration.

(ii) From a machine learning perspective, using the active sets as features allows us to reduce the dimension of the learning task. Instead of learning a complex multi-dimensional, continuous-valued mapping from the input parameters to the optimal solution, the problem is converted to the simpler task of learning the mapping between the input parameters and a finite number of active sets, allowing us to exploit the known mathematical model of the system. Thus, in contrast to the standard approach in reinforcement learning, we spend no effort on learning the mathematical model itself.

(iii) From an application perspective, active sets often correspond to meaningful operational patterns, allowing for interpretable models that can assist humans in decision-making. This point is critical in some engineering applications, where owing to the high stakes involved (such as avoiding black out in power grid), the final decision making is frequently performed by experienced system operators.

(iv) Our intuition tells us that human engineered systems often admits just a few modes of operation. This is borne out in the numerical experiments that we ran on the optimal power flow (OPF) problem across a wide variety of transmission networks, which translates into only a small number of active sets being relevant in practice. This offers significant advantage for a machine learning approach to learn these active sets from data.

Therefore, a key contribution of this paper is to show that learning the active set is an efficient and viable approach in many practical problems, with low number of relevant active sets. To this end, we propose a statistical learning approach which provides probabilistic guarantees for out-of-sample performance, and characterizes the complexity of the task in terms of the number of active sets required to produce an optimal solution with high probability. The input to the algorithm is a set of training samples consisting of the input parameters and the corresponding optimal solution, without any assumptions on the problem structure. The algorithm is designed to terminate when a sufficient number of optimal active sets have been discovered. The algorithm also allows for the user to set termination criteria that provides theoretical performance guarantees. In the worst case, the algorithm may require an exponential number of samples. However, we prove that it will terminate fast for so-called *low-complexity systems* where the number of optimal active sets is small. The algorithm thus simultaneously identifies the optimal active sets and establishes the practicability of the active-set-based learning method, which hinges on only a modest number of optimal active sets. As soon as the relevant active sets are identified, we can use them to predict the optimal solution, as outlined in Figure 1b.

A related, but distinct field of study is *explicit* Model Predictive Control (MPC), which was developed to improve solution times of MPC problems in an online setting with time-varying parameters. For parametric programs with a convex quadratic objective and linear constraints, the parameter space can be partitioned into a number of regions (corresponding to different optimal active sets), where the optimal solution in each region is an affine function of the input parameters. Explicit MPC approaches generally involve an offline procedure to identify the full

3

set of regions and corresponding optimal solutions (Borrelli et al. 2001, Bemporad et al. 2002a,b). They can then be stored in sophisticated data-structures to reduce memory storage requirements (Fuchs et al. 2010, Geyer et al. 2008, Jafargholi et al. 2014). There are also advances in fast search algorithms for the "point location problem" of identifying the corresponding active set for a given state point based on binary search trees (Tøndel et al. 2003, Johansen and Grancharova 2003, Bayat et al. 2012) and other data structures (Bayat et al. 2011, Herceg et al. 2013, Zhang et al. 2016, Zhang and Xiu 2018). Further, Karg and Lucia (2018) propose to design a deep neural net to specifically captures the piece-wise affine structure of the MPC solution manifold, and provides theoretical bounds on the number of neurons and layers to enable an exact representation of the solution.

An important difference to our approach is that explicit MPC searches through the full parameter space, and is hence limited to low dimensional systems, whereas our method uses information about the distribution over the input parameters to only investigate the practically relevant parts of the input space. Further, our approach is more general – explicit MPC is focused on solving convex quadratic programs, whereas the type of problems under consideration in this paper belong to a more general class of problems, sometimes including non-convex or integer constraints.

In the following we lay down the foundations for a statistical learning approach which provides probabilistic guarantees for out-of-sample performance, and characterizes the complexity of the task in terms of the number of active sets required to produce an optimal solution with high probability. Our main contributions in this paper are as follows:

(i) Establishing active sets as a natural and effective means to learn solutions to general parametric optimization problems. To the best of our knowledge, this has never been explored before in the literature.

(ii) Developing a streaming algorithm termed *DiscoverMass* with rigorous performance guarantees to discover the collection of relevant active sets, without any *a priori* assumptions on problem structure. The algorithm serves a dual purpose, as it both learns the relevant active sets and provides evidence of whether the suggested active-set-based learning method is applicable to a given problem instance.

(iii) Establishing the viability of the active-set approach on a practical problem. We employ *DiscoverMass* on a set of Optimal Power Flow benchmarks representative of optimization problems solved in operation of electric transmission grids to demonstrate that the number of relevant active sets is indeed small and learning them is highly efficient in terms of the number of samples required.

We emphasize that this paper *does not* provide a full end-to-end learning framework based on learning the active set. While we provide examples for how the discovered collection of active sets can be used to recover the optimal solution for a new realization of the input parameters using an ensemble policy or classification, the development of those frameworks will be considered in future work. In this paper, we lay the foundation for this future work by showing that methods based on learning the active sets are viable for practical problems.

The remainder of the paper is organized as follows. We first describe the main idea behind learning *active sets* in Section 2. We then develop a learning algorithm to identify a collection of active sets from training samples in Section 3 and establish theoretical performance guarantees. We demonstrate the effectiveness of our method using an application to power systems control in Section 5. Section 6 concludes the paper.

2 Active Set Learning for Optimization

The aim of this paper is to learn decision policies $x^*(\omega)$ which returns the optimal solution x^* for parametric programs of the form

$$x^*(\omega) \in \underset{x \in \mathbb{X}(\omega)}{\operatorname{argmin}} f(x, \omega) \tag{1}$$

where x denotes the decision variables, ω represent the uncertain input parameters, and

$$\mathbb{X}(\omega) = \{x \in \mathbb{R}^n : g_j(x, \omega) \leq 0 \quad \forall j \in [m]\} \tag{2}$$

defines the set of feasible solutions, and $f, g_1, \ldots, g_m$ are functions in x for all ω. We restrict ourselves to a set of parameters $\omega \in \Omega$, for which the optimization problem is actually feasible:

$$\Omega = \{\omega : \mathbb{X}(\omega) \neq \emptyset\}. \tag{3}$$

For any given solution $x \in \mathbb{X}(\omega)$, we define the *active set* corresponding to x as the set of constraints $\mathcal{A}$ that are binding at x. For a given parameter realization ω, we define the *optimal active set* $\mathcal{A}^*(\omega)$ as the active set at the optimal solution $x^*(\omega)$ [1]

If we know the optimal set $\mathcal{A}^*$ for a given $\omega \in \Omega$, we can solve the following optimization problem

$$x_{\mathcal{A}^*}^*(\omega) \in \underset{x \in \mathbb{X}_{\mathcal{A}^*}(\omega)}{\operatorname{argmin}} \quad f(x, \omega) \tag{4}$$

based on the reduced set of constraints $\mathcal{A}^*$

$$\mathbb{X}_{\mathcal{A}^*}(\omega) = \{x \in \mathbb{R}^n : \ g_j(x, \omega) \leq 0 \quad \forall j \in \mathcal{A}^*\} \tag{5}$$

The solution to (4) will be feasible and optimal for the original problem (1), but is easier to obtain since we typically have that $|\mathcal{A}^*| \ll m$. If we have a method to map the uncertain parameters ω to the corresponding optimal active set $\mathcal{A}^*$, we are thus able to obtain a the optimal solution efficiently.

Consider the case where we draw a new sample ω and do not know the optimal active set $\mathcal{A}^*$, but have a collection of possible candidates $\mathcal{A}_1, ..., \mathcal{A}_N$. Then, it is possible to obtain the optimal solution by solving the reduced problem (4) for each active set and check the resulting solutions for feasibility. Since each reduced problem (4) is a relaxation of (1), the solutions will either be optimal or infeasible for the original problem.

2.1 Learning the collection of relevant active sets

While the collection of all possible active sets is finite, it is also exponential in the problem size, making it a potentially complex learning task with high sample complexity. However, in many applications such as infrastructure systems that were engineered to accommodate particular modes of operation, only a few active sets are relevant in practice. The learning task at hand thus becomes to identify those important active sets.

In our approach, the *important* active sets are the active sets that have the highest probability of being optimal, under a given probability distribution over the uncertain input parameters ω. Our procedure to discover the active sets is illustrated in Figure 2. We first draw i.i.d. training samples ω_i from the probability distribution over ω. For each sample, we solve the optimization problem (1) to obtain the corresponding optimal solution and active set $\mathcal{A}_i$, which is now marked as an *observed* active set (in color in Figure 2). The collection of all active sets $\mathcal{A}_1, ..., \mathcal{A}_N$ that were observed during the first $\omega_i, ..., \omega_N$ samples form the set of observed active sets. The still *unobserved* active sets, marked in grey in Figure 2, are the active sets which were not optimal for any of our i.i.d. samples ω_i.

2.2 From active sets to the optimal solution

The knowledge of the relevant active sets can be used to more efficiently obtain the optimal solution. We describe two possible approaches.

Ensemble policy The ensemble policy obtains the optimal solution by solving the reduced problem (4) in parallel for each relevant active set $\mathcal{A}$, and then verifying the resulting solution for feasibility in (1). This method is described in detail in Ng et al. (2018), and illustrated in Figure 1b.

---

[1]When there are multiple optimal solutions, the degeneracy can be handled by using any consistent tie-breaking rule, such that the same realization ω will always be mapped to the same optimal active set.
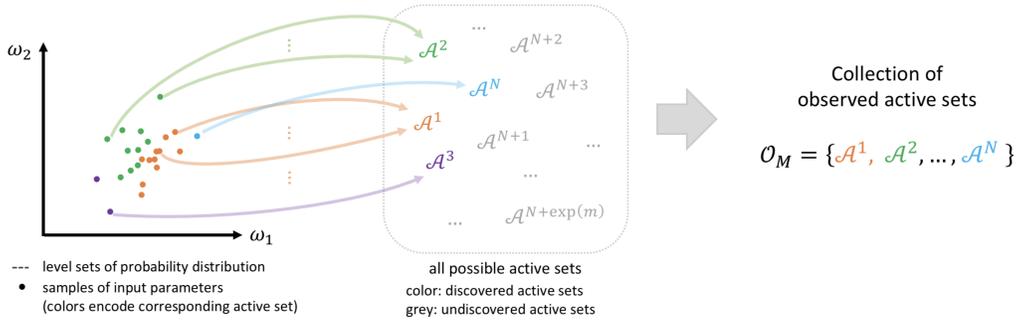
Figure 2: Learning procedure. Starting from samples (left), we discover the optimal active sets among the set of all possible active sets (middle). This provides a collection of important active sets, corresponding to those that are likely to be optimal (right).
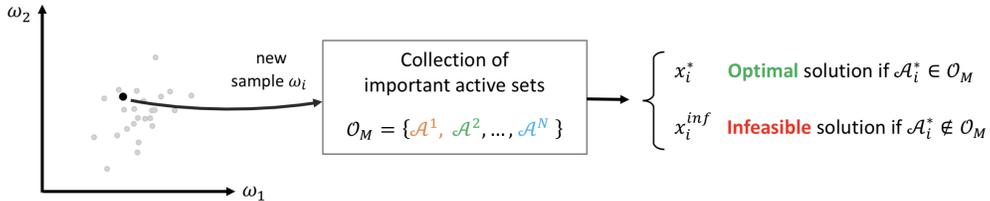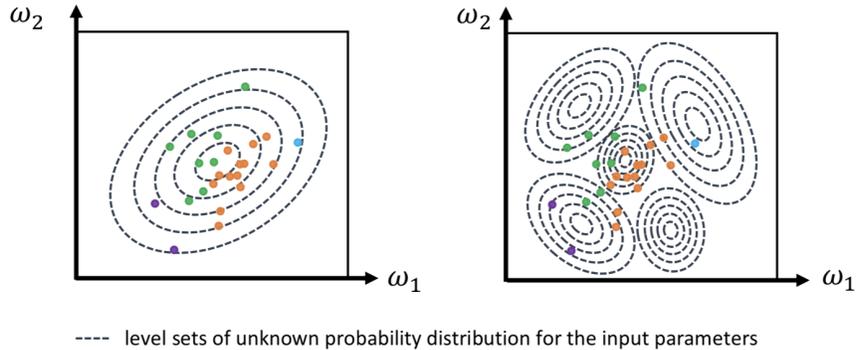

Figure 3: Prediction procedure. For a new realization ω (left), we evaluate a candidate solution for one or more observed actives sets (middle). If the optimal active set for ω corresponds to one of the observed actives sets, we recover the true optimal solution. If not, we will may end with an infeasible solution (right).
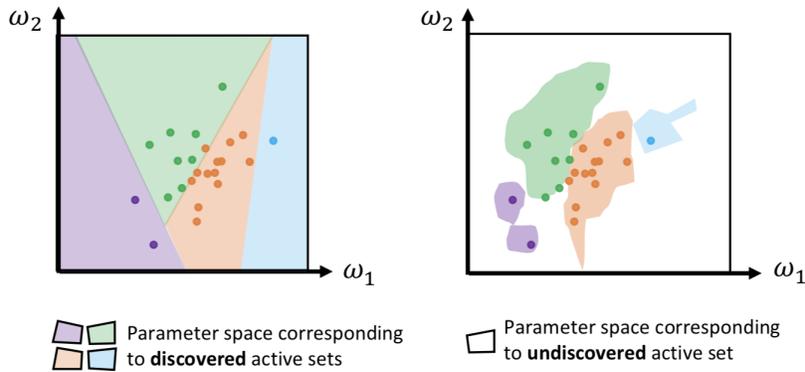
Classification This approach is based on training a classifier to learn the mapping from the parameter ω to the corresponding optimal active set. The number of classes in this classifier is same as the number of relevant active sets. Once the optimal active set is obtained, the optimal solution can be obtained by solving the reduced optimization problem in (4).

Both the above approaches are simpler and more efficient when the number of potential active sets is small. In this paper we utilize the *ensemble policy* in our experiments to test the performance of our method, and leave classification as a topic for future research.

We envision any learning approach based on learning active sets to rely on combining a new realization ω with one or more observed active sets, as illustrated in Figure 3. If the optimal active set corresponds to one of the observed active sets, then the optimal solution lies within the set of optimal solutions to the corresponding collection of relaxed problems (4). On the other hand, if the optimal active set of the new sample was previously unobserved, the solutions to the all relaxed problems (4) with $\mathcal{A}_i, ..., \mathcal{A}_N$ may be infeasible. Since this is an undesirable outcome, we would like to bound the probability of these outcomes. A main contribution of this paper is to provide a learning algorithm which guarantees that the probability of encountering a previously unobserved optimal active set falls below a pre-specified level. We develop a streaming algorithm which draws new samples and observes new optimal active sets until the desired performance level can be guaranteed. This helps us distinguish between cases where we have discovered active sets which cover a smaller or larger portion of the feasible space, as illustrated in Figure 4. The algorithm is formalized in Section 3.

---- level sets of unknown probability distribution for the input parameters

(a) The samples can come from different probability distributions.

Parameter space corresponding to **discovered** active sets

Parameter space corresponding to **undiscovered** active set

(b) Observed active sets can cover a larger or smaller part of the parameter space.

Figure 4: Our method makes no assumptions about the structure of the problem or the probability distribution we sample from, and it is hence very general. Figure 4a shows how the samples can arise from different probability distributions. Figure 4b shows how the active sets we discovered may cover a smaller or larger region of the parameter space. In general, we do not know whether what we have observed so far corresponds to a small region (left) or a large region (right), and how much of the probability mass is included in the discovered areas. Therefore, an essential part of our learning algorithm is to identify when we have discovered "enough" active sets, i.e. active sets that cover a sufficient portion of the probability mass in parameter space. Note that the active sets may have an arbitrary shape in general, as we make no assumptions on the problem structure.

3 Algorithm to learn relevant active sets

In this section, we describe a streaming algorithm that enables us to learn a collection of active sets whose combined probability mass exceeds a user-defined safety level. The algorithm draws random samples of the uncertain input parameters ω_i from the distribution $\mathbb{P}_\omega(.)$, solves the optimization problem (1) to obtain the corresponding active set $\mathcal{A}_i$ and uses a *stopping criterion* to ensure that sufficiently many active sets have been discovered. We provide details of the algorithm, and theoretical guarantees on its performance below.

Let $\mathcal{B}$ denote the finite set of all possible active sets, which is exponential in the number of constraints in the problem. We define *observed* and *unobserved* active sets in the following way.

Definition 1. *Let $\omega_1, \ldots, \omega_M$ be M i.i.d. samples drawn from the uncertainty distribution and $\mathcal{A}_1, \ldots, \mathcal{A}_M$ denote the optimal active sets corresponding to each ω_i. We call $\mathcal{O}_M = \cup_{i=1}^{M}\{\mathcal{A}_i\}$*

7

the set of observed active sets, and $\mathcal{U}_M = \mathcal{B} \setminus \mathcal{O}_M$ the set of unobserved active sets.

In Figure 4b, $\mathcal{O}_M$ and $\mathcal{U}_M$ correspond to the colored and white regions respectively. The streaming algorithm iteratively increases the number of samples M until the observed set $\mathcal{O}_M$ contains most of the important active sets. A crucial aspect of the algorithm is that knowledge of neither the shape and number of the active sets (colored regions) nor the number of samples M is required a-priori, but is decided on-the-fly. The algorithm therefore is guaranteed to converge for any system if enough samples are available, but can also be stopped at any time before termination with a lower safety guarantee.

In the following, we use the notion that the importance of an active set or a collection of active sets corresponds to the *mass* contained in it.

Definition 2. *The* mass *of an active set is defined as the probability that it is optimal for* (1) *under the distribution $\mathbb{P}_\omega(.)$ on ω. For $\mathcal{A} \in \mathcal{B}$,*

$$\pi(\mathcal{A}) = \mathbb{P}_\omega(\mathcal{A} = \mathcal{A}^*(\omega)) \tag{6}$$

For any subset $\mathcal{S} \subseteq \mathcal{B}$, we define the mass *of $\mathcal{S}$ as $\pi(\mathcal{S}) = \sum_{\mathcal{A} \in \mathcal{S}} \pi(\mathcal{A})$.*

After drawing M samples and observing a collection of active sets $\mathcal{O}_M$, we compute the rate of discovery. The rate of discovery quantifies the fraction of samples that correspond to observing an active set that was not observed among the first M samples, and is formally defined below.

Definition 3. *Let W be a positive integer denoting the window size. Let $\omega_1, \ldots, \omega_{M+W}$ be $M + W$ i.i.d. samples drawn from the uncertainty distribution and let $\mathcal{A}_i$ denote the optimal active set for ω_i. We denote by X_i the random variable that encodes whether a "new active set", i.e., an active set not observed within the first M samples, was observed in the $(M + i)^{th}$ sample, i.e.*

$$X_i = \begin{cases} 1, & \text{if } \mathcal{A}_{M+i} \notin \{\mathcal{A}_1\} \cup \ldots \cup \{\mathcal{A}_M\}, \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

Then the rate of discovery *over the window of size W is given by $\mathcal{R}_{M,W}$ and is defined as*

$$\mathcal{R}_{M,W} = \frac{1}{W} \sum_{i=1}^{W} X_i. \tag{8}$$

The rate of discovery $\mathcal{R}_{M,W}$ is directly related to the mass of the undiscovered set $\mathcal{U}_M$ of active sets. Intuitively, if we have already discovered the active sets that contain most of the mass, then the rate of discovery is expected to be low since most of the time we are likely to observe one of the active sets we have already observed so far. The following theorem makes this notion rigorous.

Theorem 1. *Let $\{\omega_i\}_{i=1}^{\infty}$ be a sequence of i.i.d. observations and define the unobserved set and rate of discovery as in Definition 1 and 3 respectively. Then the following statements hold. Let the probability mass in the unobserved set be denoted by $\pi(\mathcal{U}_M)$ and let the window size be such that*

$$W = W_M \geq c \max\{\log \underline{M}, \log M\}. \tag{9}$$

Then we can bound the probability mass in the undiscovered set by

$$\mathbb{P}\left(\pi(\mathcal{U}_M) - \mathcal{R}_{M,W} \leq \epsilon \quad \forall M \geq 1\right) > 1 - \frac{\gamma}{\gamma - 1} \frac{1}{(\underline{M} - 1)^{\gamma - 1}}, \quad \text{where } \gamma = \frac{c\epsilon^2}{2}. \tag{10}$$

The proof is deferred to Section 4. Theorem 1 forms the basis of the stopping criterion of our streaming algorithm. In the following subsection, we present this algorithm and provide rigorous guarantees on its performance based on the rate of discovery criterion.

3.1 Streaming Algorithm for Learning Important Active Sets

We present the following algorithm called *DiscoverMass* which takes in as input parameters the quantities $\alpha, \epsilon, \delta > 0$. Here, α represents the limit on the undiscovered mass and requires that the algorithm returns a set of active sets of mass at least α. The quantity ϵ represents the difference between the undiscovered mass and the rate of discovery, and $1 - \delta$ represents the confidence. We can also choose the hyperparameter γ. We have explicitly left the dependence on all the parameters in the algorithm description so that interested readers can choose the values that may be suitable for their application. We specify the concrete set of chosen values along with a brief analysis of the effect of tuning the hyperparameters in the numerical experiment section.

Data: $\alpha, \epsilon, \delta, \gamma$

Compute $c = 2\gamma/\epsilon^2$ and $\underline{M} = 1 + \left(\frac{\gamma}{\delta(\gamma-1)} \right)^{\frac{1}{\gamma-1}}$;

Initialize $M = 1$ and $\mathcal{O} = \emptyset$;

repeat

    Calculate window size $W_M = c \max\{\log \underline{M}, \log M\}$;

    Draw $1 + W_M - W_{M-1}$ additional samples from $\mathbb{P}_\omega$ to obtain a total of $M + W_M$ samples ;

    Solve optimization problem (1) for each new sample;

    Add the newly observed active sets to $\mathcal{O}$;

    Compute $\mathcal{R}_{M,W_M}$ as given in (8);

    Update M = M+1;

until $\mathcal{R}_{M,W_M} < \alpha - \epsilon$;

return $\mathcal{O}, M, R_{M,W_M}$

Algorithm 1: DiscoverMass

DiscoverMass terminates with a set of observed active sets $\mathcal{O}$ that has at least $1 - \alpha$ of the mass with high probability. The following theorem guarantees that with the stopping criterion described above, *DiscoverMass* indeed succeeds in meeting the requirements imposed by the input parameters.

Theorem 2. *If the algorithm* DiscoverMass *is employed with input parameters* $\alpha, \epsilon, \delta, \gamma$, *then with probability at least* $\mathbf{1} - \delta$ *the algorithm will terminate having discovered active sets with mass at least* $\pi(\mathcal{O}) \geq 1 - \alpha$.

The proof of Theorem 2 can be found in Section 4. A key advantage of *DiscoverMass* supported by Theorem 2 is that it makes absolutely no *a priori* assumption regarding the underlying distribution of active sets in the system. This makes *DiscoverMass* a rigorous tool to quantify the unknown underlying distribution.

3.2 Low-complexity systems

Although Theorem 2 guarantees that the algorithm *DiscoverMass* will almost always succeed in finding the desired amount of mass, it gives no guarantees on the number of steps M the algorithm needs to discover this mass. In fact, the number of steps is highly dependent on the properties of the system and can vary significantly across systems. For example in a system where there are 2^N active sets all of which are equally probable, i.e. $\mathbb{P}_\omega(\mathcal{A}_i) = 2^{-N}$ for all $i \in [2^N]$, then it is clear that discovering $(1 - \alpha)$ fraction of the mass requires at least $M > (1 - \alpha)2^N$ samples. However, for some systems, which we refer to as low-complexity systems, it is possible to discover a large fraction of the mass quickly. Fortunately, many practical systems fall within this category. In Section 5, we have found that many practical systems can be described as *low-complexity* as it is possible to discover a large fraction of the mass quickly. We quantify this intuition with the following definition.

Definition 4. *We say that a system is a **low-complexity system** with parameters* α_0 *and* K_0 *if there are at most* K_0 *active sets that contain at least* $1 - \alpha_0$ *fraction of the mass. More precisely, there exists active sets* $\mathcal{A}_{i_1}, \ldots, \mathcal{A}_{i_{K_0}}$ *such that* $\sum_{j=1}^{K_0} \pi\left(\mathcal{A}_{i_j}\right) > 1 - \alpha_0$.

The next theorem guarantees that for *low-complexity systems* the algorithm terminates within a few iterations.

Theorem 3. *Suppose that we run the algorithm* DiscoverMass$(\alpha, \epsilon, \delta, \gamma)$ *on a low-complexity system with parameters* (α_0, K_0) *with* $\alpha > \alpha_0$. *Then with probability at least* $1 - \delta - \delta_0$ *the algorithm terminates in less iterations than* $M = \frac{1}{\alpha - \alpha_0}(K_0 \log 2 + \log 1/\delta_0)$.

The proof of Theorem 3 is found in the next section. From Theorem 3, we see that the sample complexity of discovering active sets with sufficient mass for low-complexity systems scales at most linearly with the number of underlying important active sets K_0.

4 Proofs

In this section, we provide the proofs of the three theorems stated in Section 3.

4.1 Proof of Theorem 1

We first prove part (a). For any M we can bound the probability that the rate of discovery is far from the mass of the unobserved set by

$$
\mathbb{P}(\pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon)
$$
$$
= \sum_u \mathbb{P}(\pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon \mid \mathcal{U}_M = u)\mathbb{P}(\mathcal{U}_M = u)
$$
$$
= \sum_u \mathbb{P}\left(\pi(u) - \frac{\sum_{i=1}^{W_M} X_i}{W_M} > \epsilon \mid \mathcal{U}_M = u\right)\mathbb{P}(\mathcal{U}_M = u), \tag{11}
$$

where the random variables X_i are defined in (8) and conditioned on $\mathcal{U}_M = u$ are i.i.d. Bernoulli random variables with mean $\pi(u)$. Using standard large deviation inequalities, we can bound the probability in (11) as

$$
\mathbb{P}(\pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon) < \sum_u e^{-W_M \epsilon^2/2}\mathbb{P}(\mathcal{U}_M = u)
$$
$$
= e^{-W_M \epsilon^2/2}. \tag{12}
$$

Therefore,

$$
\mathbb{P}(\exists\ M \geq 1 \quad \text{s.t.} \quad \pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon)
$$
$$
\leq \sum_{M=1}^{\infty} \mathbb{P}(\pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon)
$$
$$
= \sum_{M=1}^{\underline{M}} \mathbb{P}(\pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon)
$$
$$
+ \sum_{M=\underline{M}}^{\infty} \mathbb{P}(\pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon) \tag{13}
$$
$$
< \sum_{M=1}^{\underline{M}} e^{-c \log \underline{M} \epsilon^2/2} + \sum_{M=\underline{M}}^{\infty} e^{-c \log M \epsilon^2/2} \tag{14}
$$
$$
\leq \frac{1}{\underline{M}^{\frac{c\epsilon^2}{2}-1}} + \sum_{M=\underline{M}}^{\infty} \frac{1}{M^{\frac{c\epsilon^2}{2}}} \leq \frac{\gamma}{\gamma-1}\frac{1}{(\underline{M}-1)^{\gamma-1}}.
$$

4.2 Proof of Theorem 2

Let $\mathbb{E}$ be the event that the algorithm terminates after having discovered smaller than ϵ fraction of the mass. Then

$$\mathbb{P}\left(\mathbb{E}\right) \leq \mathbb{P}\left(\exists M \geq 1 \mid \mathcal{R}_{M,W} < \epsilon - \delta, \pi(\mathcal{U}_M) > \epsilon\right) \tag{15}$$

$$\leq \mathbb{P}\left(\exists M \geq 1 \mid \pi(\mathcal{U}_M) - \mathcal{R}_{M,W} > \epsilon\right) \tag{16}$$

$$\overset{(a)}{<} \frac{\gamma}{\gamma - 1} \frac{1}{(\underline{M} - 1)^{\gamma - 1}} \overset{(b)}{=} \delta, \tag{17}$$

where the last inequality (a) follows from Theorem 1, and (b) follows from the value of $\underline{M}$ in the initialization step in *DiscoverMass*.

4.3 Proof of Theorem 3

Let $I^* = \{\mathcal{A}_{i_1}, \ldots, \mathcal{A}_{i_{K_0}}\}$ denote the set of K_0 active sets defined in Definition 4 that contain at least $1 - \alpha_0$ of the mass. For any $M \geq 1$ recall that $\mathcal{U}_M$ denotes the set of unobserved active sets as in Definition 1. We define the sets $\mathcal{U}_{M_{I^*}} = \mathcal{U}_M \cap I^*$ and $\mathcal{U}_{M_{\bar{I}^*}} = \mathcal{U}_M \cap \bar{I}^*$ where $\bar{I}^*$ denotes the complement of the ensemble I^*, i.e., the set of active sets that are not a part of I^*. It follows that $\mathcal{U}_M = \mathcal{U}_{M_{I^*}} \cup \mathcal{U}_{M_{\bar{I}^*}}$ and $\pi(\mathcal{U}_M) = \pi(\mathcal{U}_{M_{I^*}}) + \pi(\mathcal{U}_{M_{\bar{I}^*}})$. We compute

$$\mathbb{P}\left(\pi(\mathcal{O}_\mathcal{M}) < 1 - \alpha\right) = \mathbb{P}(\pi(\mathcal{U}_M) > \alpha)$$

$$= \mathbb{P}(\pi(\mathcal{U}_{M_{I^*}}) + \pi(\mathcal{U}_{M_{\bar{I}^*}}) > \alpha)$$

$$\leq \mathbb{P}(\pi(\mathcal{U}_{M_{I^*}}) > \alpha - \alpha_0)$$

$$= \sum_{u \subseteq I^*} \mathbb{P}(\mathcal{U}_{M_{I^*}} = u) \, \mathbb{1}_{\{\pi(u) > \alpha - \alpha_0\}}$$

$$\leq \sum_{u \in I^*} \left(1 - (\alpha - \alpha_0)\right)^M$$

$$\leq 2^K e^{-(\alpha - \alpha_0)M} \overset{(a)}{<} \delta_0, \tag{18}$$

where (a) follows from the choice of M in the theorem. The proof follows by combining (18) with Theorem 1.

5 Numerical results

We test our learning framework on the DC Optimal Power Flow (OPF) problem, an optimization problem widely used in electricity market clearing as well as operation and planning of electric transmission grids. For real time operation, the reaction of generators to fluctuations in renewable energy must be calculated reliably within a short period of time, so it is an example with high practical relevance that is subject to stringent limits on computational time. A simple version of the DC OPF can be stated as follows

$$\rho(\omega) \in \underset{p \in \mathbb{R}^n}{\operatorname{argmin}} \ c^\top p \tag{19a}$$

$$\text{s.t. } e^\top (p - d + \omega) = 0 \tag{19b}$$

$$p^{\min} \leq p \leq p^{\max} \tag{19c}$$

$$f^{\min} \leq M(Hp - d + \omega) \leq f^{\max} \tag{19d}$$

Here, decisions $p \in \mathbb{R}^n$ are made on the active power generation at each generator. The aim is to provide power at minimum generation cost $c^\top p$ to a set of loads $d \in \mathbb{R}^v$ with random real-time variations ω. The solution is subject to a total power balance constraint (19b), as well as limits both on the minimum and maximum power generation $p^{\min}, p^{\max} \in \mathbb{R}^n$ at each generator (19c), and the minimum and maximum admissible flow $f^{\min}, f^{\max} \in \mathbb{R}^m$ on transmission lines (19d). The topology of the network is encoded in the matrix $H \in \mathbb{R}^{v \times n}$ mapping the power

	Normal distribution					**Uniform distribution**				
	K_M	M	W_M	$R_{M,W}$	$\mathbb{P}(p^*)$	K_M	M	W_M	$R_{M,W}$	$\mathbb{P}(p^*)$
Low-Complexity										
case3_lmbd	1	1	13'259	0.0	1.0	1	1	13'259	0.0	1.0
case5_pjm	1	1	13'259	0.0	1.0	1	1	13'259	0.0	1.0
case14_ieee	1	1	13'259	0.0	1.0	1	1	13'259	0.0	1.0
case30_ieee	1	1	13'259	0.0	1.0	1	1	13'259	0.0	1.0
case39_epri	2	2	13'259	0.0	1.0	2	2	13'259	0.0008	0.9998
case118_ieee	2	33	13'259	0.0	1.0	2	4	13'259	0.0019	0.9984
case57_ieee	2	2	13'259	0.0003	0.9997	3	46	13'259	0.0	1.0
case1888_rte	3	6	13'259	0.0	1.0	3	10	13'259	0.0	1.0
case1951_rte	5	47	13'259	0.0069	0.9943	11	63	13'259	0.0084	0.9901
case162_ieee_dtc	7	91	13'259	0.0054	0.9925	17	192	13'259	0.0085	0.9926
case24_ieee_rts	10	1456	18'209	0.0	1.0	11	64	13'259	0.0047	0.9941
High-Complexity										
case73_ieee_rts	19	1258	17'844	0.0087	0.9931	130	22'000	24'977	0.0136	-
case300_ieee	24	1257	17'842	0.0073	0.9919	293	9095	22'789	0.0099	0.9897
case200_pserc	174	4649	21'112	0.0099	0.9909	236	6741	22'040	0.0099	0.9901
case240_pserc	2993	22'000	24'997	0.0795	-	2993	22'000	24'997	0.0795	-

Table 1: Outcome of the learning algorithm for different systems for normal and uniformly distributed parameters, sorted by the number of discovered active sets. The number in the case name indicates the number of buses in the system. K_M: Number of active sets. M: Number of samples until termination. W_M: Window size. $R_{M,W}$: Rate of discovery. $\mathbb{P}(p^*)$: Probability of obtaining the optimal solution.

from each generator to their corresponding bus, and the matrix of power transfer distribution factors $M \in \mathbb{R}^{m \times v}$ (Christie et al. 2000). Here e is the vector of ones.

5.1 Test set-up

We test our algorithm by running extensive simulations across a variety of networks (Li and Bo 2010, Lesieutre et al. 2011, Grigg et al. 1999, Birchfield et al. 2017, Price and Goodin 2011) from the IEEE PES PGLib-OPF v17.08 benchmark library (IEEE PES Task Force 2017). We report general results for 15 different test cases, varying in size from 3 to 1951 buses.

For each system, we consider two different distributions for the uncertain load deviations ω. First, we assume a *multivariate normal distribution*, where ω is defined as random vector of independent, zero mean variables with standard deviations $\sigma = 0.03d$ and zero correlation between loads. Second, we assume that each entry ω_i follows a *uniform distribution* with support $\omega_i \in [-3\sigma, +3\sigma] = [-0.09d, 0.09d]$. The uniform distribution is intentionally designed to spread the probability mass more evenly across a larger region, thus enabling us to compare cases with a different number of relevant active sets and different probability mass in each active set. Note that our method is able to handle any distribution, as long as a sufficient number of training samples is available.

The algorithm *DiscoverMass* has four input parameters, α, δ, ϵ, γ. We set the maximum mass of the undiscovered bases $\alpha = 0.05$ and the confidence level $\delta = 0.01$. The remaining parameters are set to $\epsilon = 0.04$ and $\gamma = 2$. Note that ϵ and γ are hyperparameters that can be tuned, leading to different minimum window size and $\underline{M}$. For details about the choice of parameters and their effects, we refer the reader to supplementary material. We run the algorithm until termination, or until $M = 22'000$. The procedure is implemented in Julia v0.6 (Bezanson et al. 2012), using JuMP v0.17 (Dunning et al. 2017), PowerModels.jl v0.5 (Coffrin et al. 2017) and OPFRecourse.jl (Ng et al. 2018).

5.2 Parameter choices in experiments

Varying the hyper-parameter γ results in different values for the window size W_M and minimum $\underline{M}$, while changing ϵ influences both the stopping criterion $R_{M,W}$ and the size of the window W. We note that there is a trade-off in the choice of both of those parameters. If γ is chosen such that $\underline{M}$ is low, the initial window size is smaller, but will rapidly start increasing as soon as $M > \underline{M}$. For ϵ a lower value implies a steep increase in the window size, but also reduces the requirements on the rate of discovery.

The difference between the probability mass of the undiscovered bases and the rate of discovery $R_{M,W}$ in the window W are set to $\epsilon = 0.04$, while we choose the hyperparameter $\gamma = 2$. Together, this implies the stopping criterion $R_{M,W} \leq 0.01$. The minimum window size $W_M \geq 13'259$ samples, and remains constant as long as $M \leq \underline{M} = 201$ samples. We run the algorithm until termination, or until $M = 22'000$. To assess how accurately we capture the probability of the undiscovered active sets, we run an out-of-sample test with 20'000 samples.

5.3 Numerical results

Table 1 lists the simulation results for the 15 different networks with sizes ranging from 3 to 1951 nodes. For each network, we report the number of samples to termination M, the number of active sets K_M discovered within the first M samples and the window size W_M required to guarantee a confidence of $\alpha < 0.05$. In addition, we list the rate of discovery $R_{M,W}$ at termination. All results are included for both the normal and uniform distributed parameters, and the system is sorted based on the number of active sets. Note that the choice of input parameters leads to a minimum window size $W_M \geq 13'259$ samples, and remains constant as long as $M \leq \underline{M} = 201$ samples.

Low-complexity systems We observe that for most systems, the learning algorithm performs well. Most systems have a relatively low number of relevant active sets (< 10), and terminate after less than 200 samples while the window size is still $W_M = W_{M,min} = 13'259$. System size is not a determining factor for the number of relevant active sets, with even the largest systems case1888_rte and case1951_rte exhibiting a low number of relevant active sets. This validates the intuition that many engineered systems (at least arising in the power systems benchmarks we've looked at) can be considered *low-complexity* as defined in Section 3.2. Comparing M and K_M in Table 1, we see a strong correlation between the number of active sets with majority of the mass and the samples till termination M. This is in agreement with the sample complexity derived in Theorem 3. The fact that most systems are low-complexity suggests that our proposed method for parameterizing and learning an optimal policy based on active sets might be effective in a wide variety of settings arising in engineered systems. We observe that the primary factor affecting the number of samples till termination is not the number of important active sets, but the existence of active sets with small but not insignificant probability mass (typically in the range 0.001-0.01). Due to their low probability mass, those active sets require a large number of samples for discovery, while the cumulative probability of the still undiscovered active sets is still high enough to exceed the stopping criterion $R_{M,W} \leq 0.01$.

High-complexity systems The outliers in this analysis are cases such as case24_ieee_rts, case73_ieee_rts, case200_pserc and case240_pserc. For case73_ieee_rts (uniform) and case240_pserc (normal and uniform), the algorithm does not terminate until we reach the upper limit of $M = 22'000$. For case240_pserc, the rate of discovery remains above $R_{M,W} \geq 0.07$ even after the discovery of 2993 active sets. We believe that this performance deficiency is due to atypical system characteristics. The two systems case200_pserc and case240_pserc, are the result of power grid models that have been subjected to network reduction, hence altering the properties of the systems relative to the other, more typical systems. Similarly, case24_ieee_rts and case73_ieee_rts are special in that they have a large number of generators relative to the number of nodes, and case73_ieee_rts consists of three identical areas. Due to the nature of their construction, these instances are useful adversarial systems for analyzing the performance of different learning algorithms.

Influence of parameter distribution We observe that the uniform distribution typically discovers a larger number of active sets compared to the normal distribution and requires more samples M before termination. With the probability mass spread more evenly across the probability space, there are a larger number of active sets with significant probability mass, and fewer active sets with high probability mass (for the multivariate normal, there is a concentration of mass in active sets closer to the forecasted value). However, for systems like case24_ieee_rts, the uniform distribution appears to encourage earlier termination due to a faster exploration rate.

Out of sample prediction performance The learning algorithm is compared against a full optimization approach where for each realization of ω the problem (1) is solved using standard optimization methods. This approach has higher computational requirements, but produces optimal solutions with probability one. Our learning based method has better computational efficiency but has a non-zero probability of failure, corresponding to returning infeasible solutions. To assess the performance of our learning algorithm we evaluated the predictions generated by the ensemble policy (where the solution is evaluated for all discovered active sets, then checked for feasibility) on $20'000$ test samples.

The results are shown in Table 1 which shows the probability of failure. We see that this value is consistently lower than the mass of the undiscovered set α provided as input to the *DiscoverMass* algorithm. Intuitively, this is to be expected, since Theorem 1 strives to provide upper bounds on the mass of the undiscovered set that are universal in nature and are guaranteed to work for all almost any continuous optimization problem and various kinds of uncertainty, without pre-specifying the number of samples M.

6   Conclusion

We present an approach to learn the mapping from a set of uncertain input parameters to the optimal solution for parametric programs. Our target application is engineered systems, where optimization problems used in operational decision making are typically solved repeatedly, but with varying input parameters. Our approach is based on learning the active sets of constraints at the optimal solution. By relying on active sets, we are able to decompose the learning task into a simpler classification task, and at the same time exploiting the knowledge of the exact mathematical model of the optimization problem, and enabling tight constraint enforcement. The viability of learning active sets and the performance of our approach is validated using a series of experiments based on benchmark optimization problem used in electric grid operation. The prediction step in this paper used the so-called ensemble policy. In future work we will explore classification and ranking algorithms to learn the mapping between the realization of the uncertain parameter onto the optimal active set. Further topics of exploration include application to long-term planning and multi-stage optimization problems.

References

Achiam J, Held D, Tamar A, Abbeel P (2017) Constrained policy optimization. *International Conference on Machine Learning*, 22–31.

Amodei D, Olah C, Steinhardt J, Christiano P, Schulman J, Mané D (2016) Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565* .

Bayat F, Johansen TA, Jalali AA (2011) Using hash tables to manage the time-storage complexity in a point location problem: Application to explicit model predictive control. *Automatica* 47(3):571–577.

Bayat F, Johansen TA, Jalali AA (2012) Flexible piecewise function evaluation methods based on truncated binary search trees and lattice representation in explicit mpc. *IEEE Transactions on Control Systems Technology* 20(3):632–640.

Bemporad A, Borrelli F, Morari M (2002a) Model predictive control based on linear Programming—The explicit solution. *IEEE Trans. Automat. Contr.* 47(12).

Bemporad A, Morari M, Dua V, Pistikopoulos EN (2002b) The explicit linear quadratic regulator for constrained systems. *Automatica* 38(1):3–20.

Bezanson J, Karpinski S, Shah VB, Edelman A (2012) Julia: A Fast Dynamic Language for Technical Computing. *CoRR* abs/1209.5145, URL `https://arxiv.org/abs/1209.5145`.

Birchfield AB, Xu T, Gegner KM, Shetye KS, Overbye TJ (2017) Grid structural characteristics as validation criteria for synthetic networks. *IEEE Trans. Power Systems* 32(4):3258–3265.

Borrelli L, Baotic T, Bemporad A, Morari T (2001) Efficient on-line computation of constrained optimal control. *Decision and Control, 2001. Proceedings of the 40th IEEE Conference on*, volume 2, 1187–1192 (IEEE).

Christie RD, Wollenberg BF, Wangensteen I (2000) Transmission management in the deregulated environment. *Proceedings of the IEEE* 88(2):170–195.

Coffrin C, Bent R, Sundar K, Ng Y, Lubin M (2017) Powermodels.jl: An open-source framework for exploring power flow formulations. URL `http://arxiv.org/abs/1711.01728`.

Dunning I, Huchette J, Lubin M (2017) JuMP: A modeling language for mathematical optimization. *SIAM Review* 59(2):295–320, URL `http://dx.doi.org/10.1137/15M1020575`.

Fisac JF, Akametalu AK, Zeilinger MN, Kaynama S, Gillula J, Tomlin CJ (2017) A general safety framework for learning-based control in uncertain robotic systems. *arXiv preprint arXiv:1705.01292* .

Fuchs AN, Jones C, Morari M (2010) Optimized decision trees for point location in polytopic data sets-application to explicit mpc. *American Control Conference (ACC), 2010*, 5507–5512 (IEEE).

Garcıa J, Fernández F (2015) A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* 16(1):1437–1480.

Geyer T, Torrisi FD, Morari M (2008) Optimal complexity reduction of polyhedral piecewise affine systems. *Automatica* 44(7):1728–1740.

Grigg C, Wong P, Albrecht P, Allan R, Bhavaraju M, Billinton R, Chen Q, Fong C, Haddad S, Kuruganty S, et al. (1999) The IEEE Reliability Test System-1996. A report prepared by the reliability test system task force of the application of probability methods subcommittee. *IEEE Trans. Power Systems* 14(3):1010–1020.

Herceg M, Mariethoz S, Morari M (2013) Evaluation of piecewise affine control law via graph traversal. *Control Conference (ECC), 2013 European*, 3083–3088 (IEEE).

Huang X, Kwiatkowska M, Wang S, Wu M (2017) Safety verification of deep neural networks. *International Conference on Computer Aided Verification*, 3–29 (Springer).

IEEE PES Task Force (2017) PGLib Optimal Power Flow Benchmarks. Published online at `https://github.com/power-grid-lib/pglib-opf`, accessed: October 4, 2017.

Jafargholi M, Peyrl H, Zanarini A, Herceg M, Mariéthoz S (2014) Accelerating space traversal methods for explicit model predictive control via space partitioning trees. *Control Conference (ECC), 2014 European*, 103–108 (IEEE).

Johansen TA, Grancharova A (2003) Approximate explicit constrained linear model predictive control via orthogonal search tree. *IEEE Transactions on Automatic Control* 48(5):810–815.

Karg B, Lucia S (2018) Efficient representation and approximation of model predictive control laws via deep learning. *arXiv preprint arXiv:1806.10644* .

Lesieutre BC, Molzahn DK, Borden AR, DeMarco CL (2011) Examining the limits of the application of semidefinite programming to power flow problems. *49th Allerton Conference*, 1492–1499 (IEEE).

Li F, Bo R (2010) Small test systems for power system economic studies. *IEEE PES General Meeting*, 1–4 (IEEE).

Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D (2015) Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* .

Lu T, Zinkevich M, Boutilier C, Roy B, Schuurmans D (2017) Safe exploration for identifying linear systems via robust optimization. *arXiv preprint arXiv:1711.11165* .

15

Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, et al. (2015) Human-level control through deep reinforcement learning. *Nature* 518(7540):529.

Moldovan TM, Abbeel P (2012) Safe exploration in markov decision processes. *Proceedings of the 29th International Coference on International Conference on Machine Learning*, 1451–1458 (Omnipress).

Ng YS, Misra S, Roald LA, Backhaus S (2018) Statistical learning for DC optimal power flow.

Price JE, Goodin J (2011) Reduced network modeling of WECC as a market design prototype. *IEEE PES General Meeting*, 1–6 (IEEE).

Tøndel P, Johansen TA, Bemporad A (2003) Evaluation of piecewise affine control via binary search tree. *Automatica* 39(5):945–950.

Van Hasselt H, Guez A, Silver D (2016) Deep reinforcement learning with double q-learning. *AAAI*, volume 16, 2094–2100.

Zhang J, Xiu X (2018) Kd tree based approach for point location problem in explicit model predictive control. *Journal of the Franklin Institute* .

Zhang J, Xiu X, Xie Z, Hu B (2016) Using a two-level structure to manage the point location problem in explicit model predictive control. *Asian Journal of Control* 18(3):1075–1086.