

Learning-based Parameter-Adaptive Reference Governors

Chakrabarty, Ankush; Berntorp, Karl; Di Cairano, Stefano

TR2020-092 July 03, 2020

Abstract

Reference governors (RGs) provide an effective method for ensuring safety via constraint enforcement in closedloop control systems. When the parameters of the underlying systems are unknown, but constant or slowly-varying, robust formulations of RGs that consider only the worst-case effect may be overly conservative and exhibit poor performance. This paper proposes a parameter-adaptive reference governor (PARG) architecture that is capable of generating safe trajectories in spite of parameter uncertainties without being as conservative as robust RGs. The proposed approach leverages on-line data to inform algorithms for robust parameter estimation. Subsequently, confidence bounds around parameter estimates are fed to supervised machine learners for approximating robust constraint admissible sets leveraged by the PARG. While initially, due to the absence of on-line data, the PARG may be as conservative as a robust RG, as more data is gathered and the confidence bounds become tighter, such conservativeness reduces, as demonstrated in a simulation example.

American Control Conference (ACC)

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Learning-based Parameter-Adaptive Reference Governors

Ankush Chakrabarty[†], Karl Berntorp, Stefano Di Cairano

Abstract—Reference governors (RGs) provide an effective method for ensuring safety via constraint enforcement in closed-loop control systems. When the parameters of the underlying systems are unknown, but constant or slowly-varying, robust formulations of RGs that consider only the worst-case effect may be overly conservative and exhibit poor performance. This paper proposes a parameter-adaptive reference governor (PARG) architecture that is capable of generating safe trajectories in spite of parameter uncertainties without being as conservative as robust RGs. The proposed approach leverages on-line data to inform algorithms for robust parameter estimation. Subsequently, confidence bounds around parameter estimates are fed to supervised machine learners for approximating robust constraint admissible sets leveraged by the PARG. While initially, due to the absence of on-line data, the PARG may be as conservative as a robust RG, as more data is gathered and the confidence bounds become tighter, such conservativeness reduces, as demonstrated in a simulation example.

Index Terms—Reference governors, machine learning, uncertain systems, adaptive systems, particle filtering, invariant sets, robust control, learning for control.

I. INTRODUCTION

Due to their ability to enforce constraints without requiring a full re-design [1], [2], and with a relatively low computational burden, reference governors (RGs) have proven useful in multiple application domains, including vehicles, aerospace, manufacturing, and energy systems [3]–[6].

Despite being relatively common in real-world applications, there are relatively few designs for RGs when the parameters of the underlying systems are uncertain. To the best of our knowledge, the load governor approach proposed in [7] is the only parameter-adaptive reference governor (PARG) formulation in the literature. A major reason for the dearth of PARG frameworks is that the computation of robust constraint admissible sets under parameter uncertainty is extremely difficult due to complex geometries of these sets and inherent non-convexity, even for linear systems. Recently proposed sampling-driven machine learning approaches may provide computationally tractable and efficient frameworks for estimating these robust invariant sets on-line [8], [9] by offloading simulation and trajectory generation off-line.

In order to construct robust constraint admissible sets, one requires not only the point estimates of the parameters themselves, but also regions in the parameter space within which the true parameter lies with high probability. To this end, we propose the use of recursive parameter estimators that have demonstrated excellent performance in a wide

range of estimation problems [10], [11]. For nonlinear systems, these approaches are generally intractable, and particle filtering provides an effective alternative [12], [13], where the state space is estimated by predicting state trajectories (particles) and weighting them according to the likelihood of the measurements.

In this paper, we describe a PARG framework that is capable of enforcing constraints in parameter-uncertain closed-loop systems without modifying the control algorithm directly. As a specific realization of this PARG framework, we consider two components: (i) a recursive statistical parameter estimator based on Bayesian update laws for generating confidence intervals around a point estimate of the unknown parameter; and (ii) a support vector machine (SVM) algorithm that dynamically learns constraint admissible sets by combining off-line data based on sampling, and on-line data provided by the parameter estimator. An advantage of using interval-based estimates rather than point estimates of the parameter, is that the intervals can exhibit certain properties such as monotonicity that are crucial to ensure performance guarantees on the PARG. Another advantage of our proposed method is that learners with good approximation properties, such as with universal kernels [8], [9], can be employed to represent highly non-convex robust constraint admissible sets for black-box systems using simulations and systematic sampling; this is challenging via analytical methods.

II. MOTIVATION

A. Problem Statement

We consider the class of parametric discrete-time nonlinear systems

$$x_{t+1} = f(x_t, v_t) + \theta^\top g(x_t, v_t), \quad (1a)$$

$$y_t = h(x_t, v_t), \quad (1b)$$

where $t \in \mathbb{Z}_+$ denotes the time-index, $x_t \in \mathbb{X} \subset \mathbb{R}^{n_x}$ is the measured system state, $v_t \in \mathbb{V} \subset \mathbb{R}$ is the reference input, $y_t \in \mathbb{Y} \subset \mathbb{R}$ is the output that should track the reference, and f, g, h are nonlinearities that represent the model and output dynamics. The vector $\theta \in \Theta \subset \mathbb{R}^{n_\theta}$ models a set of system parameters. The output y_t must satisfy constraints described by the set $\mathbb{Y} \subset \mathbb{R}$ for each instant of time, that is

$$y_t \in \mathbb{Y} \text{ for every } t \geq 0. \quad (2)$$

Assumption 1. *The sets $\mathbb{X}, \mathbb{V}, \mathbb{Y}$ and Θ are compact and known to the designer. The sets \mathbb{X}, \mathbb{V} , and \mathbb{Y} contain the origin in their interiors. The set \mathbb{V} is convex.*

Note that Assumption 1 is mild. Although the classical reference governor literature does not require explicit

[†]Corresponding author: Ankush Chakrabarty, Phone: +1 (617) 758-6175.

All authors {chakrabarty, berntorp, dicairano}@merl.com are affiliated with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA.

boundedness of the state-space \mathbb{X} , we make this an explicit assumption because we will use sampling-based methods to characterize constraint admissible sets in this paper. In order to make these set approximations accurate, the domain over which the samples are drawn must be bounded, and known to the designer. Furthermore, the reference v is assumed to be scalar, which is the standard formulation for the reference governor. The results of this manuscript can be extended almost directly to vector-valued v , which is the case for vector reference governor and command governor; see for example, [6].

An implicit assumption made in the description above is that in the unconstrained setting, that is $\mathbb{Y} = \mathbb{R}$, the closed-loop system (1) exhibits good tracking performance. Thus, the closed-loop system (1) is asymptotically stable and for each $r \in \mathbb{V}$, when $v_t = r_t \equiv r$ for all $t \geq 0$, $y_t \rightarrow r$ as $t \rightarrow \infty$. The objective of a reference governor is to select v_t as close as possible to r_t while ensuring that (2) is enforced. In the literature [6], the commonly treated cases are when θ is known or when it is unknown and constantly varying within a given range. In this paper, we consider the case when the parameter vector $\theta \in \Theta \subset \mathbb{R}^{n_\theta}$ is unknown, but constant.

Our proposed PARG is given by the control law

$$\begin{aligned} v_t &= \bar{\mathcal{G}}(v_{t-1}, x_t, \hat{\Theta}_t, r_t) \\ &= v_{t-1} + \mathcal{G}(v_{t-1}, x_t, \hat{\Theta}_t, r_t)(r_t - v_{t-1}), \end{aligned} \quad (3)$$

where $\hat{\Theta}_t \subset \Theta$ is a bounded interval of parameter values, computed by a parameter estimator with the functional form

$$\hat{\Theta}_t = \mathcal{E}(v_{t-1}, x_t, \hat{\Theta}_{t-1}). \quad (4)$$

Our **objective** is to design \mathcal{G} and \mathcal{E} such that the closed-loop system (1), (3), and (4) satisfies constraints (2) in spite of parametric uncertainty and, when possible, tracks the desired reference r_t . A schematic diagram representation of the proposed PARG architecture is shown in Figure 1.

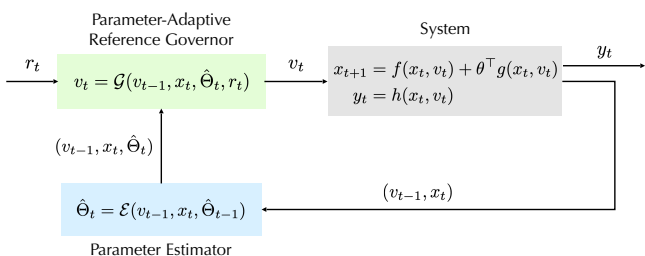


Fig. 1. Block diagram representation of the parameter-adaptive reference governor (PARG) added-on to the unconstrained system.

Remark 1. In many applications such as industrial motors [14], autonomous vehicles [15], and buildings [16], exact model parameters are not always known, although the set Θ is known from experience or archived literature. ■

B. Proposed Solution

Constructing PARGs for the uncertain system (1) poses two major difficulties. First, parameter estimators generally converge asymptotically to the true parameters, implying that

the current estimate of the parameter at any arbitrary finite time t is not necessarily correct. Thus, designing references with point estimates $\hat{\theta}_t \neq \theta_t$ does not guarantee constraint satisfaction for any finite t . Second, estimating parameter invariant sets on-line for employment in a reference governor using model-based analytical methods is computationally challenging, because the underlying constraint admissible invariant sets are typically non-convex with respect to the parameters and analytical methods do not always scale well into high-dimensional state-spaces.

In order to address the first difficulty, we propose using parameter estimators that generate not only a point-estimate $\hat{\theta}_t$, but also an interval $\hat{\Theta}_t$ which contains the true θ with certainty (in deterministic estimators) or with high probability (in stochastic estimators). The advantage of using confidence intervals instead of point estimates is that they can be made to exhibit certain useful properties such as non-expansivity as more data becomes available. Unlike point estimates, which can be time-varying and unpredictable, confidence intervals can be designed to exhibit predictable dynamics, making them effective for constraint enforcement. Unfortunately, replacing a point-estimate with a confidence interval exacerbates the second difficulty as one now needs to estimate robust parameter invariant sets for the system (1) with varying confidence bounds. One of the contributions of this paper is to estimate these sets in a computationally efficient manner by exploiting supervised learning machines.

More formally, let $\mathcal{H} = \{(x, v) \in \mathbb{X} \times \mathbb{V} : h(x, v) \in \mathbb{Y}\}$ denote the set of state and reference inputs for which the output y satisfies the constraint (2). We present the following definition for parameter-robust constraint admissible sets.

Definition 1. The set $\mathcal{O}(\hat{\Theta}) \subset \mathcal{H}$ is a parameter-robust constraint admissible set for (1) if, for every initial condition $(x, v) \in \mathcal{O}(\hat{\Theta})$, when $x_0 = x$ and $v_t = v$ for all $t \geq 0$, $(x_t, v_t) \in \mathcal{H}$ for every $\theta \in \hat{\Theta}$ and for all $t > 0$. The set $\mathcal{O}(\hat{\Theta})$ is invariant.

In order to generate estimates of parameter-robust constraint admissible sets, we will adopt an off-line sampling-driven approach to collect data for learning the sets on-line as operational data becomes available.

An estimate of a parameter-robust constraint admissible set can subsequently be used to evaluate the control law (3) by solving for

$$\mathcal{G}(v_{t-1}, x_t, \hat{\Theta}_t, r_t) := \arg \min_{\gamma_t} (v_t - r_t)^2 \quad (5a)$$

$$\text{subject to: } (v_t, x_t) \in \mathcal{O}(\hat{\Theta}_t), \quad (5b)$$

$$v_t = v_{t-1} + \gamma_t(r_t - v_{t-1}), \quad (5c)$$

$$0 \leq \gamma_t \leq 1, \quad (5d)$$

$$v_t \in \mathbb{V}_\varepsilon(\hat{\Theta}_t) \quad (5e)$$

at each time instant t . Note that $\mathbb{V}_\varepsilon(\hat{\Theta}_t)$ denotes the set of references v such that a ball of radius $\varepsilon > 0$ centered at the corresponding steady state $x^{\text{ss}}(v, \theta)$ and v lies inside $\mathcal{O}(\hat{\Theta}_t)$,

$$\mathbb{V}_\varepsilon(\hat{\Theta}_t) \triangleq \left\{ v \in \mathbb{V} : \mathcal{B}_\varepsilon(x^{\text{ss}}(v, \theta), v) \subset \mathcal{O}(\hat{\Theta}_t), \forall \theta \in \hat{\Theta} \right\}.$$

III. DESIGN OF LEARNING-BASED PARG

In this section, we discuss the off-line data collection phase, the on-line parameter estimation phase, and the on-line learning phase of the proposed PARG.

A. Off-line data generation for on-line learning

As in [9], we will simulate trajectories of the closed loop system (1) off-line, from different initial states sampled from \mathbb{X} , reference inputs sampled from \mathbb{V} , and parameters within Θ . At the end of each off-line simulation, if an initial condition $x_i \in \mathbb{X}$ tracks a desired reference input $v_i \in \mathbb{V}$ without violating the constraint (2) at any time in the simulation, for a parameter θ_i sampled within Θ , then the combination (x_i, v_i) is labeled ‘+1’ to indicate it resides within the parameter-robust constraint admissible set $\mathcal{O}(\{\theta_i\})$. Contrarily, if the constraint is violated at any time point in the simulation, the feature (x_i, v_i) is labeled ‘-1’ to indicate it resides outside $\mathcal{O}(\theta_i)$. This sets up a binary classification problem which can be solved via supervised machine learning. This is stated more formally herein.

We extract N_x unique samples from \mathbb{X} and construct grids (not necessarily equidistantly spaced) on \mathbb{V} and Θ with N_v and N_θ nodes, respectively. Let x_i denote the i -th sampled state, v_j the j -th sampled reference input, and θ_k the k -th sampled parameter. For each (x_i, v_j, θ_k) , we simulate the model (1) forward in time over a finite horizon T_s with a constant reference v_j and parameter θ_k . The horizon T_s is chosen long enough that the tracking error is small (for example, $< 10^{-6}$) by the end of the simulation. For each simulation, we check whether $y_t \in \mathbb{Y}$ for every simulation time-point. We set the corresponding label of the sample x_i as follows:

$$\ell_i^{j,k} = \begin{cases} +1, & \text{if } y_t \in \mathbb{Y} \text{ for every } t \in \{0, 1, \dots, T_s\}, \\ -1, & \text{otherwise.} \end{cases} \quad (6)$$

At the end of this off-line data generation procedure, we have a fixed collection of initial $\{x_i\}_{i=1}^{N_x}$, and each initial condition x_i has a corresponding $N_v \times N_\theta$ matrix of labels

$$\ell_i = \begin{bmatrix} \ell_i^{1,1} & \dots & \ell_i^{1,N_\theta} \\ \vdots & \ddots & \vdots \\ \ell_i^{N_v,1} & \dots & \ell_i^{N_v,N_\theta} \end{bmatrix},$$

from which a labeled set will be generated on-line for robust invariant set estimation using supervised learning. Note that every element in ℓ_i is either +1 or -1 by (6).

B. Obtaining confidence intervals from parameter estimators

The proposed approach leverages confidence intervals around parameter estimates to learn parameter-robust constraint admissible sets. An efficient way of determining such confidence intervals is by using Kalman filters [10], [11], [17], [18] (for the linear case) and adaptive particle filters [19] (for the nonlinear case). We do this by reformulating (1) in a probabilistic framework where θ is treated as an unknown disturbance with stochastic properties. In the current work, since x_t is known, we exploit the linearity of the system (1)

with respect to θ and use a Kalman filter for estimating θ and its confidence interval Θ_t . Note that the approach can be extended to the case when the state vector is not completely known and has to be estimated together with the parameter.

Specifically, we reformulate (1a) as

$$\theta_t = \theta_{t-1} + w_t, \quad (7a)$$

$$\bar{y}_t = g(x_{t-1}, v_{t-1})^\top \theta_t + e_t, \quad (7b)$$

where $\bar{y}_t = x_t^\top - f^\top(x_{t-1}, v_{t-1})$, that is, the dynamical system (1a) for x_t now plays the role of the measurement (output) equation in the Kalman filter.

The reason to address the parameter estimation problem in a Bayesian framework is that even if the state x_t is known, for instance, from measurements, such knowledge is typically imperfect due to inherent noise in the sensors measuring the state, even though we do not model the uncertainties explicitly in (1) for simplicity. Furthermore, a Bayesian framework provides a systematic approach to work with confidence intervals in recursive estimators. In a Bayesian context, we reformulate (7) as $\theta_{t+1} \sim p(\theta_t)$ and $\bar{y}_t \sim p(\theta_t)$. We address the parameter estimation problem by recursively estimating the posterior density function of the parameter θ_t , given by

$$p(\theta_t | \bar{y}_{0:t}). \quad (8)$$

using the measurement history $\bar{y}_{0:T} = \{\bar{y}_0, \dots, \bar{y}_T\}$.

The Bayesian updates for solving (8) can be summarized in the prediction and update equations

$$p(\theta_t | \bar{y}_{0:t-1}) = \int p(\theta_t | \theta_{t-1}) p(\theta_{t-1} | \bar{y}_{0:t-1}) d\theta_{t-1}, \quad (9a)$$

$$p(\theta_t | \bar{y}_{0:t}) = \frac{p(\bar{y}_t | \theta_t) p(\theta_t | \bar{y}_{0:t-1})}{p(\bar{y}_t | \bar{y}_{0:t-1})}, \quad (9b)$$

where $p(\bar{y}_t | \bar{y}_{0:t-1})$ is a normalization constant. If the process noise w_t and measurement noise e_t are Gaussian distributed then the Bayesian update recursions (9) result in the Kalman filter equations that estimate the parameter mean $\hat{\theta}_t$ and associated covariance P_t . Using the covariance, we estimate the confidence interval $\hat{\Theta}$ as

$$\hat{\Theta}_t^j = [\hat{\theta}_t^j - \beta P_t^{j,j}, \hat{\theta}_t^j + \beta P_t^{j,j}]$$

for each element j in the parameter vector θ_t and $\beta > 0$.

In order to provide theoretical guarantees on the PARG, we need to ensure that our confidence intervals do not expand with more available data, that is, $\hat{\Theta}_{t+1} \subseteq \hat{\Theta}_t$. While this is a natural consequence of applying Kalman filters to linear-in-parameter systems [20] such as (7), in general, exploration using nonlinear filters such as in particle filters could result in a violation of this condition. In such scenarios, we explicitly enforce contraction of confidence intervals. Specifically, if the filter computes an updated confidence interval $\hat{\Theta}_{t+1}$, we set

$$\hat{\Theta}_{t+1} := \begin{cases} \hat{\Theta}_t \cap \hat{\Theta}_{t+1} & \text{if } \hat{\Theta}_{t+1} \cap \hat{\Theta}_t \neq \emptyset \\ \hat{\Theta}_t, & \text{otherwise.} \end{cases} \quad (10)$$

This forces non-expansion of $\hat{\Theta}_t$ for all $t \geq 0$.

Remark 2. If the state vector is available at every t , one can use a linear estimator to provide the confidence intervals, and, therefore, a more general approach using Bayesian recursions (9) is not needed. However, if the state is unavailable, the updates (9) can be employed to generate joint estimates of states and parameters via nonlinear recursive estimators. ■

C. Solving (5) by learning robust invariant sets

We will solve the problem (5b) efficiently (albeit approximately) using machine learning and gridding \mathbb{V} . Placing a grid on \mathbb{V} , along with the constraints (5c) and (5d), imply that the solution to (5) is contained within the sub-grid of \mathbb{V} defined by

$$\tilde{\mathbb{V}}_t := [\min\{r_{t,k}, v_{t,k}\}, \max\{r_{t,k}, v_{t,k}\}]. \quad (11)$$

Then, we can recast the problem (5) as a grid search,

$$v_t := \arg \min_{v \in \tilde{\mathbb{V}}_t} (v - r_t)^2 \quad (12a)$$

$$\text{subject to: } (v, x_t) \in \mathcal{O}(\hat{\Theta}_t), \quad (12b)$$

$$v \in \mathbb{V}_\varepsilon(\hat{\Theta}_t). \quad (12c)$$

Note that in order to solve the problem (12), we require an estimate of the set $\mathcal{O}(\hat{\Theta}_t)$, which is obtained by posing the estimation of this set as a supervised learning problem with concept drift.

In learning with concept drift, the features remain constant, but the labeled set changes with time [21]. In this work, the robust parameter invariant set changes with time because $\hat{\Theta}_t$ is time-varying. Therefore, a state x_i that was infeasible for $\hat{\Theta}_t$ may become feasible in the shrunken set $\hat{\Theta}_{t+1}$, even if v_t is fixed. This corresponds to a change in an element of the labeled set. If $v_t \neq v_{t+1}$, this could incur more drastic changes in $\hat{\Theta}_{t+1}$ and the labeled set.

We set up the learning problem as follows. At time instant t , consider a $\hat{\Theta}_t$ provided by the parameter estimator. Then, for each $v_j \in \tilde{\mathbb{V}}_t$ described in (11), and each $x_i \in \{x_i\}_{i=1}^{N_x}$ sampled off-line, we assign the label

$$z_{i,j}(\hat{\Theta}_t) = \min_{k \in \mathcal{I}_{i,j}(\hat{\Theta}_t)} \ell_i^{j,k}, \quad (13)$$

where

$$\mathcal{I}_{i,j}(\hat{\Theta}_t) := \{k : \theta_k \in \hat{\Theta}_t\}$$

is the index set of parameters contained in the current confidence interval $\hat{\Theta}_t$. Taking the minimum in (13) ensures that the estimated set is robust to *all* parameters within $\hat{\Theta}_t$. That is, if even one θ_i is infeasible for the particular v_j and x_i , then x_i does not belong to the robust parameter invariant set corresponding to $\hat{\Theta}_t$.

With the training data $D := \{(x_i, v_j), z_{i,j}\}$, we construct classifiers ψ_j , where $j = 1, \dots, |\tilde{\mathbb{V}}_t|$. For each v_j , a classifier is trained on features $\{x_i\}$ and their corresponding labels $\{z_{i,j}\}$. These classifiers need to represent inner approximations of the robust parameter invariant sets; to this end, one may select sub-level sets of the decision boundary $\psi_k = 0$ of the classifier until no infeasible sample is contained in the interior of the sub-level set [8]. Solving the problem (12)

then becomes identical to selecting the node v_j on the grid $\tilde{\mathbb{V}}_t$ that minimizes the cost (12a) while ensuring that $\psi_j(x_t) > 0$; that is, the current state is predicted by the j -th classifier to belong to the robust parameter invariant set induced by $\hat{\Theta}_t$.

As an exemplar classification algorithm, consider a 2-norm soft margin support vector machine (SVM) classifier trained on a dataset D by solving the optimization problem

$$(w_j^*, b_j^*, \xi_j^*) := \arg \min_{w, b, \xi} w^\top w + c \xi^\top \xi \quad (14a)$$

$$\text{subject to: } z_{i,j} (w^\top \varphi(x_i) + b) \geq 1 - \xi_i, \\ \forall i = 1, \dots, N_x.$$

Here, $c > 0$ is a regularization constant, w quantifies the margin of separation, b is a bias term, ξ are slack variables, and φ is a feature map into a reproducing kernel Hilbert space for a kernel function \mathcal{K} . The decision function of the SVM is given by

$$\psi_j(x) = \text{sign}((w_j^*)^\top \varphi(x) + b_j^*), \quad (14b)$$

where the inner product $(w_j^*)^\top \varphi(x)$ can be expressed efficiently by the kernel function \mathcal{K} . In the following section, we will provide probabilistic guarantees of the learning quality of this classification algorithm when used for estimating $\mathcal{O}(\hat{\Theta}_t)$.

IV. THEORETICAL GUARANTEES

A. Guarantees on the parameter estimator

We begin with the following lemma, which ensures that the true parameter lies in each confidence interval $\hat{\Theta}_t$ with high probability.

Lemma 1. *If the estimator (7) uses the update (10), then*

$$\Pr[\theta \in \hat{\Theta}_t, \forall t \geq 0] \geq \pi_\theta. \quad (15)$$

B. Guarantees on the learning algorithm

The parameter estimator is not the only statistical method used in this paper. For certain classes of learners used to approximate $\mathcal{O}(\hat{\Theta}_t)$, one can also provide probabilistic guarantees of approximation quality.

We require the following definition from [22] before we can present the lemma.

Definition 2. *A continuous kernel \mathcal{K} is universal in \mathbb{X} if the space of all functions induced by \mathcal{K} is dense in the space of all continuous functions defined on \mathbb{X} .*

Remark 3. The Gaussian radial basis function kernel used in nonlinear SVM is an example of a commonly used universal kernel. ■

Suppose the SVM learner (14) used for estimating the parameter-robust constraint admissible set comprises a kernel \mathcal{K} that is universal in the feature space \mathbb{X} . Let Φ denote the set of decision functions induced by the kernel \mathcal{K} . Let $\phi_j \in \Phi$ be the decision function generated by solving (14) with N_x data samples for each $j = 1, \dots, |\tilde{\mathbb{V}}_t|$. We define

the empirical risk of ϕ_j on the feature set \mathbb{X} and the labeled set \mathbb{Z} by

$$\mathcal{R}(\phi_j, N_x, c) := \int_{\mathbb{X} \times \mathbb{Z}} \mathbf{1}_{\{\phi_j(x) \neq z\}} \mu(df, dz),$$

where $\mu(\cdot, \cdot)$ is a probability measure on $\mathbb{X} \times \mathbb{Z}$.

The following lemma states that as long as the two sets of features corresponding to opposite labels are a strictly positive distance from each other and the true decision boundary is not pathological, they can be separated with a given margin at zero empirical risk by solving (14), given a sufficiently large number of data points, as long as the SVM learner employs a universal kernel.

Lemma 2. *For any $t \geq 0$, suppose that the training set be partitioned into two compact sets \mathbb{D}^+ and \mathbb{D}^- containing only positive and only negative labels, respectively, and suppose the true separating boundary has bounded curvature and finite perimeter. If the distance between these two sets is strictly positive, then the universal kernel \mathcal{K} separates \mathbb{D}^+ and \mathbb{D}^- with margin $\gamma > 0$. Furthermore, for any regularizer $c > 0$ of the SVM, there exists a decision function ψ_j of the form (14b) and scalar $\pi_\gamma \in (0, 1)$ such that*

$$\Pr[\mathcal{R}(\psi_j, N_x, c) = 0] \geq \pi_\gamma \quad (16)$$

for sufficiently large N_x and every $j = 1, \dots, |\tilde{\mathbb{V}}_t|$.

As argued in [8], one can make the learner more conservative to ensure that (for a given finite dataset) no infeasible point is labeled as feasible (for safety). In terms of the true learning problem (e.g. as the dataset grows to infinity), one can shrink the decision boundary in such a way that the probability of labeling an infeasible point as feasible is arbitrarily small at the expense of asymmetrically labeling feasible points as infeasible.

Remark 4. A lower bound on N_x for a given π_γ can be computed using the value of c , γ , and the covering number of the space \mathbb{X} . ■

C. Guarantees on the PARG

The following lemma ensures that the parameter-robust constraint admissible sets do not contract for a fixed reference input, as long as the confidence intervals are non-expansive.

Lemma 3. *The update (10) implies that $\mathcal{O}(\hat{\Theta}_t) \supseteq \mathcal{O}(\hat{\Theta}_{t-1})$ and $\mathcal{O}(\hat{\Theta}_t) \subseteq \mathcal{H}$ for all $t \geq 0$.*

Lemmas 1–3 enable the following guarantees on the constraint satisfaction performance of the PARG-in-the-loop system.

Theorem 1. *Suppose Assumption 1 hold. Let π_θ and π_γ be as defined in (15) and (16), respectively. Let $t_0 \geq 0$ denote a time instant at which $(x_{t_0}, v) \in \mathcal{O}(\hat{\Theta}_{t_0})$ for some $v \in \mathbb{V}_\varepsilon(\hat{\Theta}_{t_0})$. Then, and for any $r(t) \in \mathbb{V}$, the closed-loop system (1), (5), (10) satisfies the constraints (2) for all $t \geq t_0$, with probability at least $\pi_\theta \pi_\gamma$.*

Theorem 2. *Let the conditions of Theorem 1 hold. Let $r(t) = r$ for all $t \geq 0$, and let there be a finite time \hat{t} such that $r \in$*

$\mathbb{V}_\varepsilon(\hat{\Theta}_{\hat{t}})$. Then, there exists a finite $\bar{t} \geq \hat{t}$ such that $v(\bar{t}) = r$, with probability at least $\pi_\theta \pi_\gamma$.

V. NUMERICAL EXAMPLE

We illustrate our proposed method on a second-order nonlinear electromagnetically actuated mass-spring damper system studied in [1]. The closed-loop system without the reference governor is given by the forward Euler discretization of

$$\begin{aligned} \dot{x}_1 &= x_2, \\ \dot{x}_2 &= -\frac{c + c_d}{m} x_2 + \theta \left(\frac{1}{m} v - \frac{1}{m} x_1 \right), \\ y &= x_1, \end{aligned}$$

with sampling time $\tau = 1$ ms. The parameter $\theta = 38.94$ represents the unknown spring constant; other parameter values are $c_d = 4.00$, $c = 0.66$, and $m = 1.54$. The set \mathbb{Y} is described by the constraints $x_1 \leq 8 \times 10^{-3}$ and $0 \leq u(x, v) \leq 0.3$, where

$$u(x, v) = \frac{1}{\alpha} (\theta v - c_d x_2) (d_0 - x_1)^\nu$$

is the legacy tracking controller whose structure and parameters θ , c_d , $\alpha = 4.5 \times 10^{-5}$, $d_0 = 1.02 \times 10^{-2}$, and $\nu = 1.99$, cannot be altered. At design time, we know that $\theta \in [10, 90]$; thus, we initialize $\hat{\Theta}_0 := [0, 90]$. We also know that the set of reference inputs $\mathbb{V} := [0.5 \times 10^{-3}, 7.5 \times 10^{-3}]$, and $\mathbb{X} := [-8 \times 10^{-3}, 8 \times 10^{-3}] \times [-4 \times 10^{-2}, 4 \times 10^{-2}]$.

For off-line data generation, we randomly select $N_x = 2000$ low-discrepancy samples on \mathbb{X} , and uniformly partition \mathbb{Y} and \mathbb{V} into 50 and 80 sub-intervals, respectively. For each x_i , we simulate the closed-loop dynamics forward in time for $T_s = 5$ seconds, and check whether the constraints were violated at any time point.

To perform parameter estimation, we use the recursive filter described in Sec. III-B. We model the measurement covariance to be diagonal according to $R_f = \text{diag}([10^{-4}, 10^{-8}])$ and set $Q_f = 10^{-8}I$. The 99% confidence interval $\hat{\Theta}_t$ generated by the parameter estimator is used for on-line updating of the robust parameter invariant sets. Upon updating the label set based on $\hat{\Theta}_t$ as described in (13), we use support vector machine (SVM) bi-classifiers to perform invariant set estimation. The SVM kernel is composed of radial basis functions, and default hyperparameters in MATLAB's Statistics and Machine Learning Toolbox are assigned. In order to promote the generation of strictly feasible sets (where no point labeled '-1' is contained in the set) we use an asymmetric cost function with cost matrix: $\text{antidiag}([1 \ 0.5])$. Since the computation times for updating these sets are generally larger than the sampling time, we adopt a heuristic for skipping updates if $\|\hat{\Theta}_t - \hat{\Theta}_{t-1}\| \leq 0.001$ and $\|r_t - r_{t-1}\| \leq 10^{-5}$.

In Figure 2, we compare the performance of the learning-based PARG to a non-adaptive RG which assumes a parameter value of $\tilde{\theta} = 45$, which is the point estimate $\hat{\theta}$ after 0.1 s (which means 100 data points, since $\tau = 0.001$). The output of the parameter estimator is shown in subplot




Fig. 2. [A] Mean parameter estimate, the true parameter value, and the 99% confidence interval. [B] Measured output of closed-loop system with conventional RG with incorrect parameter estimate. [C] Measured output of closed-loop system with learning-based PARG. No constraints are violated. PE = parameter estimator, RG = reference governor without adaptation, PARG = parameter-adaptive reference governor.

[A] (dotted line) along with the true parameter value (green continuous line). The point estimate $\hat{\theta}$ converges to a small neighborhood around θ within 1 s, and the 99% confidence intervals (blue continuous lines) start contracting to a tight set around θ around 20 s. Note that the contractions of $\hat{\Theta}_t$ occur when the desired reference r_t jumps and v_t varies. This happens because the v_t transient dynamics excite the closed-loop system and parameter estimation is abetted by satisfaction of weak persistence of excitation conditions. In subplot [B] and [C], we illustrate the benefit of the learning-based PARG. In subplot [B], we see that the non-adaptive RG cannot satisfy constraints at all time $t \geq 0$ because the constraint admissible set is generated based on an incorrect estimate of θ . Conversely, as evident from subplot [C], the PARG, which uses parameter-robust constraint admissible set, does not violate constraints anywhere.

VI. CONCLUSIONS

In this paper, we developed an adaptation mechanism for reference governors that can handle constraint satisfaction in systems with parametric uncertainties. We demonstrated that machine learning and confidence interval estimation is effective for approximating robust invariant sets in a computationally tractable manner, thereby enabling the PARG to run on-line.

ACKNOWLEDGMENTS

The authors would like to thank Uroš V. Kalabić and Matt Brand at MERL for helpful comments.

REFERENCES

- [1] E. Gilbert and I. Kolmanovsky, "Nonlinear tracking control in the presence of state and control constraints: a generalized reference governor," *Automatica*, vol. 38, no. 12, pp. 2063–2073, 2002.
- [2] E. Gilbert and K. Tan, "Linear systems with state and control constraints: the theory and applications of maximal output admissible sets," *IEEE Trans Automatic Control*, vol. 36, no. 9, pp. 1008–1020, 1991.
- [3] U. Kalabić, R. Gupta, S. Di Cairano, A. Bloch, and I. Kolmanovsky, "Constrained spacecraft attitude control on SO(3) using reference governors and nonlinear model predictive control," in *Proc. of the American Control Conference (ACC)*. IEEE, 2014, pp. 5586–5593.
- [4] S. Di Cairano, U. V. Kalabić, and I. V. Kolmanovsky, "Reference governor for network control systems subject to variable time-delay," *Automatica*, vol. 62, pp. 77–86, 2015.
- [5] W. Lucia, G. Franzè, and M. Szaier, "A hybrid command governor scheme for rotary wings unmanned aerial vehicles," *IEEE Transactions on Control Systems Technology*, 2018.
- [6] E. Garone, S. Di Cairano, and I. Kolmanovsky, "Reference and command governors for systems with constraints: A survey on theory and applications," *Automatica*, vol. 75, pp. 306–328, 2017.
- [7] J. Sun and I. Kolmanovsky, "Load governor for fuel cell oxygen starvation protection: A robust nonlinear reference governor approach," *Proc. of the American Control Conference*, vol. 1, no. 6, pp. 828–833, 2004.
- [8] A. Chakrabarty, V. Dinh, M. J. Corless, A. E. Rundell, S. H. Žak, and G. T. Buzzard, "Support vector machine informed explicit nonlinear model predictive control using low-discrepancy sequences," *IEEE Transactions on Automatic Control*, vol. 62, no. 1, pp. 135–148, 2016.
- [9] A. Chakrabarty, A. Raghunathan, S. Di Cairano, and C. Danielson, "Data-driven estimation of backward reachable and invariant sets for unmodeled systems via active learning," in *Proc. of the IEEE Conf. on Decision and Control (CDC)*. IEEE, 2018, pp. 372–377.
- [10] S. Särkkä and A. Nummenmaa, "Recursive noise adaptive Kalman filtering by variational Bayesian approximations," *IEEE Transactions on Automatic Control*, vol. 54, no. 3, pp. 596–600, 2009.
- [11] Y. Huang, Y. Zhang, Z. Wu, N. Li, and J. Chambers, "A novel adaptive kalman filter with inaccurate process and measurement noise covariance matrices," *IEEE Transactions on Automatic Control*, vol. 63, no. 2, pp. 594–601, 2017.
- [12] A. Doucet and A. M. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," in *Handbook of Nonlinear Filtering*, D. Crisan and B. Rozovsky, Eds. Oxford University Press, 2009.
- [13] K. Berntorp and S. Di Cairano, "Process-noise adaptive particle filtering with dependent process and measurement noise," in *Int. Conf. Decision and Control*, Las Vegas, NV, Dec. 2016.
- [14] J. Chen, J. Huang, and Y. Sun, "Resistances and speed estimation in sensorless induction motor drives using a model with known regressors," *IEEE Transactions on Industrial Electronics*, vol. 66, no. 4, pp. 2659–2667, 2018.
- [15] K. Berntorp and S. Di Cairano, "Tire-stiffness and vehicle-state estimation based on noise-adaptive particle filtering," *IEEE Transactions on Control Systems Technology*, vol. 27, no. 3, pp. 1100–1114, 2018.
- [16] P. Radecki and B. Hency, "Online building thermal parameter estimation via unscented Kalman filtering," in *Proc. of the 2012 American Control Conference (ACC)*, 2012, pp. 3056–3062.
- [17] R. Mehra, "On the identification of variances and adaptive Kalman filtering," *IEEE Transactions on Automatic Control*, vol. 15, no. 2, pp. 175–184, 1970.
- [18] X. Rong Li and Y. Bar-Shalom, "A recursive multiple model approach to noise identification," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 30, no. 3, pp. 671–684, 1994.
- [19] E. Özkan, V. Šmídl, S. Saha, C. Lundquist, and F. Gustafsson, "Marginalized adaptive particle filtering for nonlinear models with unknown time-varying noise parameters," *Automatica*, vol. 49, no. 6, pp. 1566–1575, 2013.
- [20] F. Gustafsson, *Statistical Sensor Fusion*. Lund, Sweden: Utbildningshuset/Studentlitteratur, 2010.
- [21] G. Widmer and M. Kubat, "Learning in the presence of concept drift and hidden contexts," *Machine Learning*, vol. 23, no. 1, pp. 69–101, 1996.
- [22] I. Steinwart, "On the influence of the kernel on the consistency of support vector machines," *Journal of Machine Learning Research*, vol. 2, no. Nov, pp. 67–93, 2001.