Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Learning Non-Parametric Models with Guarantees: A Smooth Lipschitz Regression Approach ⋆

Emilio T. Maddalena * **Colin N. Jones** *

* *Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne 1015, Switzerland (e-mails: {emilio.maddalena,colin.jones}@epfl.ch).*

Abstract: We propose a non-parametric regression methodology that enforces the regressor to be fully consistent with the sample set and the ground-truth regularity assumptions. As opposed to the Nonlinear Set Membership technique, this constraint guarantees the attainment of everywhere differentiable surrogate models, which are more suitable to optimization-based controllers that heavily rely on gradient computations. The presented approach is named Smooth Lipschitz Regression (SLR) and provides error bounds on the prediction error at unseen points in the space. A numerical example is given to show the effectiveness of this method when compared to the other alternatives in a Model Predictive Control setting.

Keywords: Safe learning, error bounds, nonlinear set membership, non-parametric regression, model predictive control.

1. INTRODUCTION

For more than five decades the expression 'learning' has been utilized by the control community to describe architectures in which performance is improved over time as additional information on the environment is collected (Nikolic and Fu, 1966; Fu, 1970). Due to the increasing interest in machine learning methods in recent years, areas such as Reinforcement Learning (RL), Neuro-Dynamic Programming (NDP) and Iterative Learning Control (ILC) have become very active research topics (Kiumarsi et al., 2017; Rosolia and Borrelli, 2018). When exploring these data-driven approaches to dealing with dynamical systems, a natural question arises: should one directly seek to optimize his objective or build a model before proceeding to the control task? This turns out to be a long-standing debate (Hou and Wang, 2013; Formentin et al., 2014), but model-based solutions are more frequently adopted when the problem being tackled requires safety guarantees (Canale et al., 2014; Wabersich and Zeilinger, 2018). This is the direction taken herein.

Owing to their high representative power and success in diverse applications, Gaussian Processes (GPs) are a bayesian non-parametric technique that has being used to deal with a variety of analysis and control problems. In Ostafew et al. (2016) and Berkenkamp et al. (2016b), the authors explore GPs respectively for ground robot locomotion and tuning controller parameters for quadrotors; whereas the same formalism is used to estimate regions of attraction for nonlinear systems in Berkenkamp et al. (2016a). One critical ingredient used in many GP-based methodologies capable of guaranteeing a safe operation is the ability to bound the error between the GP mean function and the unknown true system dynamics. This bound is usually expressed as a scaled version of the posterior standard deviation (see Srinivas et al. (2012); Chowdhury and Gopalan (2017) for the derivations). The main challenge in these cases is efficiently estimating the various associated quantities such as the maximum information gain, especially in on-line learning.

An alternative approach to GPs which is also classified as non-parametric is the well known Nonlinear Set Membership (NSM) method presented in Milanese and Novara (2004) and refined later in a series of works. By assuming differentiability of the ground-truth, optimal upper and lower bounds on its unknown values can be obtained. The predictor that minimizes the worst-case error over the whole domain arises simply as the mean of such bounds. A similar theory was studied in Beliakov (2006), as well as in Calliess (2016); Limon et al. (2017) for Hölder-continuous functions. One limitation of these works is the non-differentiability of the yielded regressors, which are in general piecewise nonlinear and whose regions are defined by the intersection of two hyperbolic Voronoi diagrams. Although this might not constitue a problem in some domains, it certainly hinders their application in a real-time optimization-based controller context.

Contributions: In this work we propose a particular radial basis functions (RBF) regression methodology that overcomes the main limitation of the NSM and similar approaches, namely the regressor non-differentiability. The method is referred to as Smooth Lipschitz Regression (SLR), and can be readily used for learning non-parametric models of nonlinear dynamical systems from data while providing prediction error bounds. In contrast with the recent proposal of filtering interpolants a poste-

⋆ This work has received support from the Swiss National Science Foundation under the RISK project (Risk Aware Data-Driven Demand Response, grant number 200021 175627.

riori (Manzano et al., 2019), we directly obtain a nominal model with the desired regularity properties. A numerical example is provided to show several advantages of SLR over NSM in a Model Predictive Control problem, where the on-line computational times are significantly reduced.

Notation and basic definitions: $\mathbb{R}^n$ is the n-dimensional Euclidean space endowed with the usual metric. span(V) is the linear span of a collection of vectors V. A $n \times n$ diagonal matrix with entries $c_1, \ldots, c_n$ is written as $diag(c_1, \ldots, c_n)$. Prob(A) denotes the probability of event A in an appropriate probability space. A map $f : \mathcal{X} \to \mathcal{Y}$ between two metric spaces with metrics d_x and d_y is Lipschitz continuous if $\exists L : d_y(f(x_1), f(x_2)) \leq L\, d_x(x_1, x_2), \forall x_1, x_2 \in \mathcal{X}$. The smallest such constant L is known as the *best* Lipschitz constant. We shall focus on Lipschitz functions between Euclidean spaces. Given a function f, its gradient is a column vector denoted by ∇f. $\boldsymbol{\mathcal{C}}^0$ and $\boldsymbol{\mathcal{C}}^1$ are respectively the space of continuous and continuously differentiable functions from $\mathcal{X}$ to $\mathcal{Y}$.

2. PROBLEM DESCRIPTION

Let $\mathcal{X} \subset \mathbb{R}^n$ be a compact and convex set, and $\mathcal{Y} \subseteq \mathbb{R}$. The map $f : \mathcal{X} \to \mathcal{Y}$ is referred to as the *target function*, *regression function* or *ground-truth* and is unknown.

Assumption 1. The ground-truth f belongs to $\boldsymbol{\mathcal{C}}^1$ and $\max_{x \in \mathcal{X}} ||\nabla f(x)|| = L$, with L unknown.

Assumption 2. A collection of noise corrupted data points $D = \{(\mathsf{x}_i, \tilde{\mathsf{y}}_i) \,|\, \tilde{\mathsf{y}}_i = f(\mathsf{x}_i) + \delta_i, i = 1, \ldots, N\}$ is available.

Assumption 3. The following upper bounds are known:
- $\overline{L} \geq L$
- $\overline{\delta} \geq |\delta_i|, \forall i = 1, \ldots, N$

An additional property of the target function is formally derived based solely on Assumption 1.

Proposition 1. f is Lipschitz continuous with best constant L.

Proof: By using the multidimensional mean value theorem, $\forall x_1, x_2 \in \mathcal{X}$, $f(x_2) - f(x_1) = \nabla f(x_0)^T (x_2 - x_1)$, where $x_0 = \theta x_1 + (1 - \theta) x_2$ for some $\theta \in [0, 1]$. Then $|f(x_2) - f(x_1)| \leq ||\nabla f(x_0)||\, ||x_2 - x_1|| \leq L\, ||x_2 - x_1||$. L being the best constant follows from Rademacher's theorem and the convexity of $\mathcal{X}$ (Weaver, 2018, Corollary 1.42).

The set of all possible functions that are compatible with both our regularity assumption and the available data-set is called the 'feasible system set' in the NSM literature; nevertheless, at times the regressor is chosen from outside this same set (Cucker and Smale, 2002), causing the terminology to become inappropriate. For this reason, a more general term is preferred herein.

Definition 1. The function space $\mathbb{F}_D = \{f : \mathcal{X} \to \mathcal{Y} \,|\, f \in \boldsymbol{\mathcal{C}}^1,\, L_f \leq \overline{L},\, |f(\mathsf{x}_i) - \tilde{\mathsf{y}}_i| \leq \overline{\delta},\, \forall (\mathsf{x}_i, \tilde{\mathsf{y}}_i) \in D\}$, where L_f is a Lipschitz constant of f, is called the *consistent space*.

As shown in Milanese and Novara (2004); Beliakov (2006), any member $f \in \mathbb{F}_D$, and in particular the ground-truth itself, is *optimally* bounded by the functions

$$f_l(x) \leq f(x) \leq f_u(x),\ \forall x \in \mathcal{X} \tag{1}$$

where

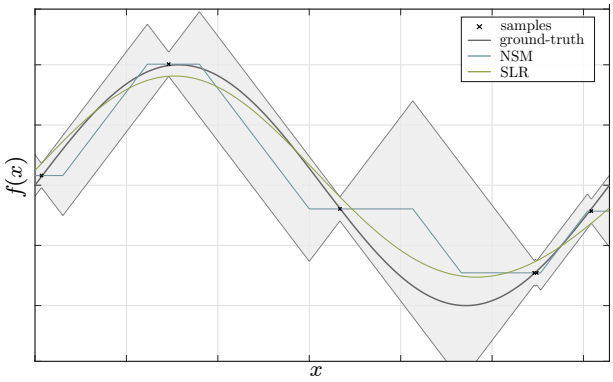

Fig. 1. A comparison between the NSM and SLR approaches when fitting noisy samples of a sinusoidal wave. The gray area enclosed by the upper and lower bounds define the space where the unknown ground-truth can lie.

$$f_u(x) = \min_{i=1,\ldots,N} (\tilde{\mathsf{y}}_i + \overline{\delta} + \overline{L}\, ||x - \mathsf{x}_i||) \tag{2a}$$

$$f_l(x) = \max_{i=1,\ldots,N} (\tilde{\mathsf{y}}_i - \overline{\delta} - \overline{L}\, ||x - \mathsf{x}_i||) \tag{2b}$$

are called respectively the ceiling and floor functions. Hence, $f_l(x)$ and $f_u(x)$ bound the space where the unknown ground-truth can be located. The more data-points are available, the tighter the bounds. An example is given in Fig. 1, where five noisy samples were collected from a sinusoidal wave.

In a rather general setting, the regression problem being tackled is defined by the mathematical program below

$$\mathbb{P}1 : \min_{\hat{f} \in \mathbb{H}} \ c(\hat{f}) \tag{3a}$$

$$\text{s.t.} \ \ \hat{f} \in \mathbb{F}_D \tag{3b}$$

where $c : \mathbb{H} \to \mathbb{R}$ is a suitable real-valued functional, and $\mathbb{H}$ is the hypothesis space being used.

Remark 1. In the NSM framework, a closed-form solution to the problem exists and the obtained regressor $\hat{f}$ minimizes the cost $c(\hat{f}) = \sup_{f \in \mathbb{F}_D} ||f - \hat{f}||_p$ for different norms p. The optimal solution however lies not in $\mathbb{F}_D$, but rather on its boundary (Milanese and Novara, 2004), hence the non-differentiable points. In the SLR setting, we seek to enforce the membership $\hat{f} \in \mathbb{F}_D$ to obtain an everywhere differentiable function which is also fully consistent with the prior knowledge on the ground-truth.

Remark 2. Regularization terms of the form $\lambda ||\hat{f}||$ can be employed in the cost function to bias the solution towards smooth regressors. See for instance the discussion in Bhujwalla et al. (2016). Our approach, on the other hand, imposes constraints to *guarantee* (in a probabilistic sense) the desired model regularity.

3. SMOOTH LIPSCHITZ REGRESSION

3.1 Designing a consistent regressor

We restrict our attention to regressors represented as weighted sums of basis functions

$$\hat{f}(x) = \sum_{k=1}^{m} w_k\, \phi_k(x) \tag{4}$$

with smooth maps $\phi_k : \mathcal{X} \to \mathcal{Y}$ and scalars $w_k \in \mathbb{R}$, $\forall k = 1, \ldots, m$. The surrogate model $\hat{f}$ is therefore a member of $\mathbb{H} = \text{span}(\{\phi_1, \ldots, \phi_m\})$, which is a finite-dimensional Hilbert space when endowed with the standard operations and inner product. The number of elements is chosen to be $m = N+E$, where N is the number of available data-points and E is a design parameter. Next, we discuss how the condition $\hat{f} \in \mathbb{F}_D$ in $\mathbb{P}1$ can be approximately translated into computational terms.

From the very choice of $\{\phi_1, \ldots, \phi_m\}$, $\hat{f} \in \boldsymbol{\mathcal{C}}^1$ is guaranteed. Compatibility with the data-set D amounts to imposing $2N$ linear constraints on the weights, which are clearly always feasible. The only remaining condition is $L_{\hat{f}} \leq \overline{L}$. Since one usually knows L_{ϕ_k} for each basis function – or can numerically estimate it to any desired precision – the upper bound $\overline{L_{\hat{f}}} = \Sigma_{k=1}^{m} |w_k| \, L_{\phi_k}$ can be easily obtained for $L_{\hat{f}}$. Although enforcing $\overline{L_{\hat{f}}} \leq \overline{L}$ suffices to establish the desired condition, it turns out to be extremely conservative thus severely reducing the optimization problem feasible set. Instead, Proposition 1 links the maximum norm of the gradient and the best Lipschitz constant, hence we can exploit the fact that

$$L_{\hat{f}} \leq \overline{L} \tag{5}$$

$$\Leftrightarrow \max_{x \in \mathcal{X}} ||\nabla \hat{f}(x)|| \leq \overline{L} \tag{6}$$

$$\Leftrightarrow \forall x \in \mathcal{X}, ||\nabla \hat{f}(x)|| \leq \overline{L} \tag{7}$$

and use the last inequality. Finally, overcoming the uncountably infinite number of constraints is possible by employing random convex programs, i.e., the scenario approach theory (Calafiore and Campi, 2006).

Consider a uniform probability distribution over $\mathcal{X}$, and let $\{x_j\}_{j=1}^{S}$ be a set of 'scenarios' uniformly extracted from the domain with $S \geq \frac{2}{\epsilon} \ln(\frac{1}{\beta}) + 2m + \frac{2m}{\epsilon} \ln(\frac{2}{\epsilon})$, where both the confidence parameter $\beta \in (0, 1)$ and the level parameter $\epsilon \in (0, 1)$ are specified. The proposed approach to finding a smooth regressor then reads

$$\mathbb{P}2: \min_{w \in \mathbb{R}^m} \; c(w) \tag{8a}$$

$$\text{s.t.} \;\; w^T \phi(\mathsf{x}_i) \leq \mathsf{y}_i + \overline{\delta}, \, \forall i = 1, \ldots, N \tag{8b}$$

$$w^T \phi(\mathsf{x}_i) \geq \mathsf{y}_i - \overline{\delta}, \, \forall i = 1, \ldots, N \tag{8c}$$

$$\left|\left| \sum_{k=1}^{m} w_k \, \nabla \phi_k(x_j) \right|\right| \leq \overline{L}, \, \forall j = 1, \ldots, S \tag{8d}$$

where $w := [w_1 \; \ldots \; w_m]^T$, $\phi := [\phi_1 \; \ldots \; \phi_m]^T$, and $c(w)$ is an appropriate objective function. Let $w^\star := \text{argmin}\, \mathbb{P}2$ and $f^\star(x) := w^{\star T} \phi(x)$ denote respectively the optimal weight vector and optimal regressor function. The constraints (8d) are less restrictive than (3b) since $\hat{f}$ belongs to the consistent space $\mathbb{F}_D$ only if (8d) is imposed on the whole domain.

Proposition 2. If $c(w)$ is a convex function and an optimal solution $w^\star$ is found for $\mathbb{P}2$, then with probability no smaller than $1 - \beta$, it is also ϵ-level feasible for the robust problem obtained by imposing $\forall x \in \mathcal{X}$, $||\nabla \hat{f}(x)|| \leq \overline{L}$, i.e., $\text{Prob}(\, ||\nabla f^\star(x)|| > \overline{L} \,) < \epsilon$.

Proof: For $c(w)$ convex, $\mathbb{P}2$ is convex since (8b) and (8c) are linear, and (8d) is equivalent to a set of quadratic constraints. The proposition then follows directly from (Calafiore and Campi, 2006, Theorem 1 and Corollary 1).

Remark 3. Increasing the design parameter E provides greater flexibility for $\hat{f}(x)$ as the number of basis functions grows, but the number of constraints (8b) and (8c) remains unchanged. The regressor is then based on N samples and E 'free points'. As $m = N + E \to \infty$, a proper choice of basis functions $\phi(x)$ would ensure approximation of any continuous function arbitrarily well. As an example, Gaussian RBFs are dense in $\boldsymbol{\mathcal{C}}^0 \supset \mathbb{F}_D$ given compact domains. As a side effect, $m \to \infty$ would also increase the number of scenario constraints (8d) to infinity. Nevertheless, the feasible set of $\mathbb{P}2$ cannot become empty since at least the target function $f \in \mathbb{F}_D$ satisfies (8d) everywhere and thus belongs to the feasible set. Hence, by increasing the number of free points E, feasibility of $\mathbb{P}2$ is likely to be achieved.

Remark 4. Among the many possible objectives one could select, including a regularizer $||w||_1$ helps achieving sparser weight vectors w since it is an approximation of the nonconvex cardinality function. This is particularly useful when E is initially set to a large number aiming at finding a preliminary solution to $\mathbb{P}2$.

3.2 Probabilistic and deterministic guarantees

Next we examine the probabilistic property of the SLR regressor $f^\star(x)$ being completely contained inside the original ceiling and floor functions (2). Furthermore, a deterministic error bound is derived for $f^\star(x)$ irrespective of the later probability.

It is known that the effect of decreasing β on the sampling complexity S is minor and that in many applications this parameter can be set to 10^{-10} or even lower numbers (Campi et al., 2009). Therefore, Proposition 2 ensures with practical certainty that, if a solution is found, $||\nabla f^\star(x)||$ will be lower or equal to $\overline{L}$ on its whole domain except for a subset of size smaller than ϵ. One practical consequence of this fact is the possible violation of the bounds defined in (1). When $||\nabla f^\star(x')|| > \overline{L}$ for some $x' \in \mathcal{X}$ the optimal regressor may exceed the ceiling function or be inferior to the floor function. Nevertheless, it is also conceivable to have a function with higher gradients that still lies completely inside the aforementioned bounds. Consider for instance the target function $f(x_1, x_2) = 0.02\, x_1^2 + 2\, \cos(x_2)$ and a grid of 49 samples as the available data-set. The target function, two SLR regressors constructed with exponential RBFs, and the NSM interpolant are shown in Fig. 2(a). In Fig. 2(b) we present violation results obtained if the gradient constraints (8d) are completely neglected, i.e. $S = 0$ (left); and setting $\beta = 10^{-10}$, $\epsilon = 0.5$, producing $S = 1224$ scenarios (right). The plots illustrate that increasing the number of scenarios not only reduces the areas in which $||\nabla f^\star(x)||$ is greater than $\overline{L}$, but also ensures that the optimal regressor does not violate the ceiling and floor bounds. It is noteworthy that this result was achieved with a finite number of scenarios and a considerably high level parameter ϵ.

As shown below, deterministic error bounds on the prediction error can still be guaranteed due to the fact that the consistent space $\mathbb{F}_D$ is itself bounded.

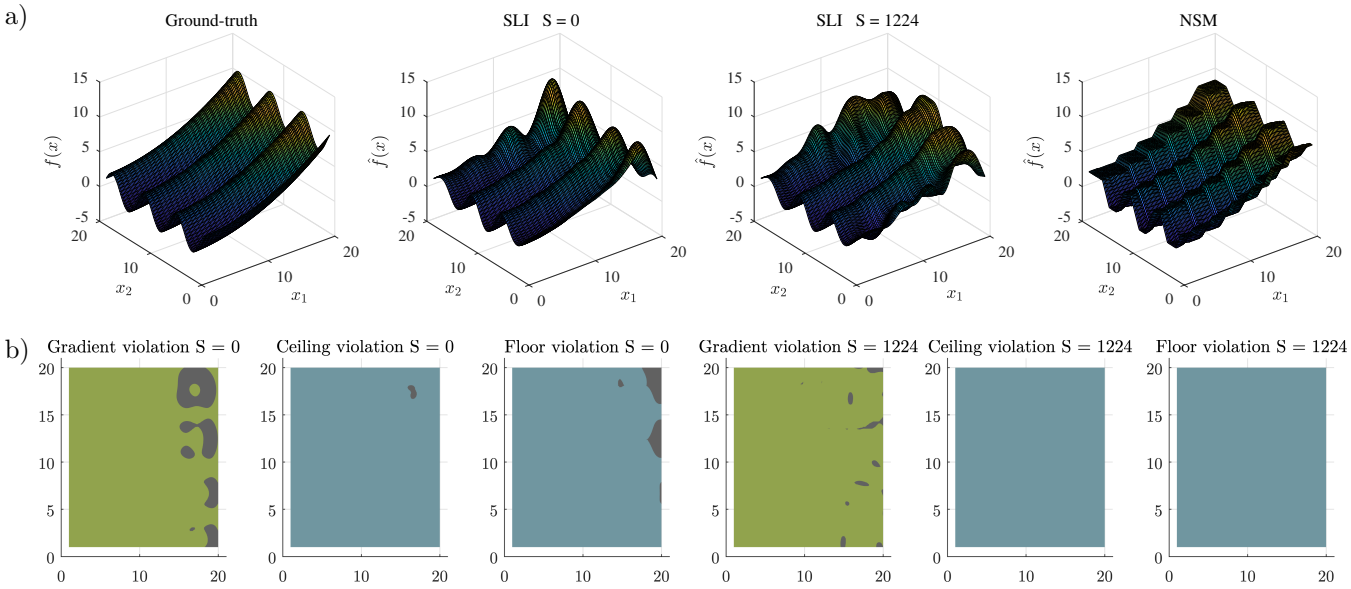

Fig. 2. (a) Target function, SLR and NSM regressors constructed based on a uniform grid of samples. (b) Effects of increasing the number of scenarios in SLR: violations areas are shown in gray for $S = 0$ (left) and $S = 1224$ (right).

Proposition 3. Let $e_{\hat{f}}(x) := \hat{f}(x) - f(x)$ denote the prediction error of a regressor $\hat{f}(x)$ with respect to the ground-truth $f(x) \in \mathbb{F}_D$. Then $\overline{e}_{\hat{f}}(x) \geq |e_{\hat{f}}(x)|$ holds for all $x \in \mathcal{X}$, where

$$\overline{e}_{\hat{f}}(x) := \max\{|\hat{f}(x) - f_l(x)|, |\hat{f}(x) - f_u(x)|\} \qquad (9)$$

Proof: In view of (1), $\forall x \in \mathcal{X}$, if $\hat{f}(x) \leq f(x)$, then $|e_{\hat{f}}(x)| \leq |\hat{f}(x) - f_u(x)|$, on the other hand, if $\hat{f}(x) \geq f(x)$, then $|e_{\hat{f}}(x)| \leq |\hat{f}(x) - f_l(x)|$. Therefore, $|e_{\hat{f}}(x)| \leq \max\{|\hat{f}(x) - f_l(x)|, |\hat{f}(x) - f_u(x)|\}$.

Taking the above proposition and the discussion on the number of scenarios into account, we note that while (9) could be used for robust control purposes, imposing an appropriate number of constraints L is still important to lowering the values attained by $\overline{e}_{f^\star}(x)$.

4. NUMERICAL EXAMPLE

Consider the following discrete-time model of a unicycle

$$x_1^+ = x_1 + u_1 \cos(x_3) \qquad (10a)$$
$$x_2^+ = x_2 + u_1 \sin(x_3) \qquad (10b)$$
$$x_3^+ = x_3 + u_2 \qquad (10c)$$

where x_1, x_2, $x_3 \in \mathbb{R}$ represent respectively the robot first and second coordinates in a fixed reference frame, and its orientation; u_1, $u_2 \in \mathbb{R}$ denote respectively the linear and angular velocity inputs. Let $x = \begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix}^T$, $u = \begin{bmatrix} u_1 & u_2 \end{bmatrix}^T$; $f_1(x,u)$, $f_2(x,u)$ and $f_3(x,u)$ be the difference equations associated respectively with x_1^+, x_2^+ and x_3^+. The three dynamics are to be learned separately given the validity of Assumption 1 and Lipschitz bounds $L_1 = L_2 = 1.8$ and $L_3 = 3$ for the target functions. The unicycle is free to move in the plane, i.e., there are no state constraints, whereas the control variables must satisfy $u \in \mathcal{U} := \{u \in \mathbb{R}^2 | \begin{bmatrix} -0.5 & -0.5 \end{bmatrix}^T \leq u \leq \begin{bmatrix} 0.5 & 0.5 \end{bmatrix}^T\}$. Random initial conditions were generated inside a box of unitary sides and random inputs were applied to the system, producing a total of $N = 300$ noise corrupted samples with $\overline{\delta} = 0.03$.

Exponential radial basis functions

$$\phi_k(x) = \sigma^2 \exp\left(\frac{-||x - \mathsf{x}_k||^2}{2\,l^2}\right) \qquad (11)$$

$k = 1, \ldots, m$ were employed with $\sigma = 1$, $l = 2$, and $E = 20$ free points. The confidence and level parameters were chosen as $\beta = 10^{-5}$, $\epsilon = 0.3$, yielding $S = 4\,764$ scenarios. Next, in order to encourage sparsity the objective of $\mathbb{P}2$ was chosen to be $c(w) = ||w||_1$ and the problem was solved in MATLAB with the aid of MOSEK. The optimal solutions were found in approximately 10.96 seconds each on a 3.1 GHz Intel Core i7 machine. The regression performance of every obtained function was compared with the NSM alternative. A total of 100 000 points were uniformly sampled from the domain and used to calculate the root mean squared error (RMSE) and mean absolute error (MAE) indexes. The results are reported in Table 1, where it can be seen how SLR performed better in both negatively-oriented scores in all three cases: the cumulative errors are reduced by at least 81%. It can be however argued that the same E free points that guarantee additional flexibility to SLR may compromise its applicability later in an MPC controller since the regressor complexity is increased. The following simulations address this issue.

Table 1. NSM and SLR regression performance scores (negatively oriented)

	$\hat{f}_1(x,u)$		$\hat{f}_2(x,u)$		$\hat{f}_3(x,u)$	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
NSM	0.2008	0.1566	0.2114	0.1634	0.2410	0.1923
SLR	0.0282	0.0202	0.0344	0.0234	0.0435	0.0262

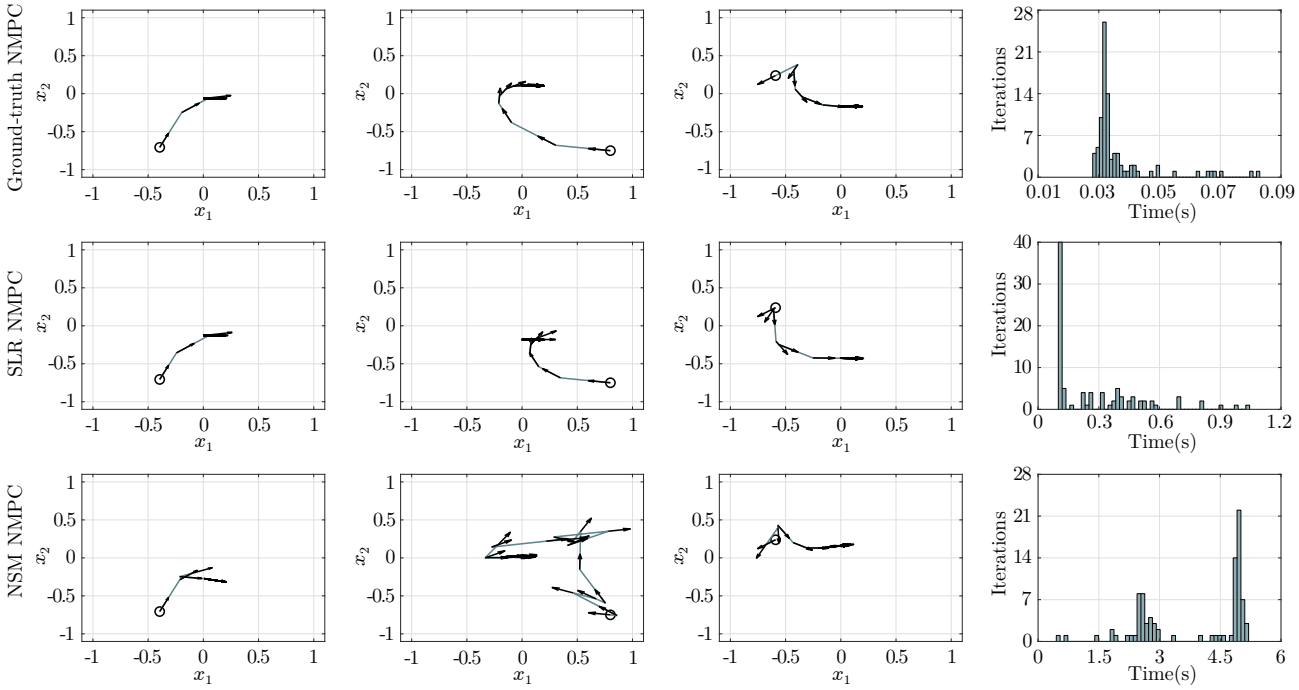

Fig. 3. Simulation results: Unicycle trajectories with NMPC designed with full knowledge of the dynamics (top row), SLR regressor (middle row) and NSM regressor (bottom row), where the same data-set was used for the last two techniques. In each of the first three columns, the system evolution is depicted starting from the same initial condition. The the computational times per NMPC iteration is presented in the rightmost column histograms.

Three initial conditions were given to the unicycle, and NMPC controllers were designed with cost function $J = \Sigma_{t=0}^{N_{hor}}(x^T(t+1)Q\,x(t+1))+(u^T(t)R\,u(t))$, with parameters $N_{hor} = 5$, $Q = diag(10^3, 10^3, 1)$, $R = diag(10^{-2}, 10^{-2})$, and hard constraints were imposed on the inputs $u(t) \in \mathcal{U}$, $t = 0, \ldots, N_{hor} - 1$. First, full knowledge of (10a)–(10c) was assumed and the ground-truth functions were employed as dynamic constraints in the NMPC formulation. Next, the SLR and NSM regressors constructed with the same data-set were considered. All optimizations were carried out using an interior-point method under *exactly the same settings*. The resulting trajectories are shown in Fig. 3 for 30 steps, with arrows indicating the unicycle orientation at each time instant. Clearly, using the SLR regressor instead of the NSM one led to a system evolution more similar to the full-knowledge case. Even though all techniques conducted the unicycle to approximately the same final conditions, the movement under the NSM-based NMPC was more abrupt in specific parts of the space.

Computational times per controller iteration are reported in the histograms of Fig. 3 and additional measures are given in Table 2 ('Std.' is used for standard deviation). Despite the 20 additional free points, optimization with the SLR regressor was significantly faster than with the NSM counterpart, reducing the mean iteration time by 93% with lower standard deviation. This is due to the piecewise non-linear non-differentiable nature of the equality constraints established by the NSM model, whereas a single smooth differentiable nonlinear equality constraint is enforced per time step in the SLR case. In practice, an even further speed up could be achieved in the SLR case by dropping the basis functions whose weights were negligible, thus simplifying the regressor.

Table 2. NMPC computational times

	Mean (s)	Std. (s)	Worst case (s)
Ground-truth	0.0368	0.0112	0.0830
SLR	0.2856	0.2269	1.0421
NSM	3.8635	1.2794	5.1685

Even though the NMPC problem formulated in this section has no state constraints, the error bounds of the obtained regressors were analyzed. A set of 10^6 samples were uniformly sampled from the joint state-control space and used to evaluate the SLR bound through (9) and the NSM bound – simply $0.5\,|f_u(x,u) - f_l(x,u)|$. The histograms of normalized obtained values for each one of the three dynamics are shown in Fig. 4. The SLR technique resulted in larger error bounds with respect to NSM in all cases, with an increase in the empirical mean of 19%, 20% and 13% for the first, second and third dynamics, respectively. This behavior was expected given that the NSM regressor minimizes (9) by being equidistant to the ceiling and floor functions (see e.g. Fig. 1). Moreover, the two uppermost histograms are similar to each other since one target function is a translated version of the other; whereas the last histogram presented the largest mean and maximal error bound, which are mainly due to the more conservative Lipschitz constant $\overline{L}_3$.

5. CONCLUSION

The SLR methodology was proposed to construct surrogate models from datapoints of everywhere differentiable unknown functions defined on bounded convex domains. In contrast with the NSM approach, the presented math-

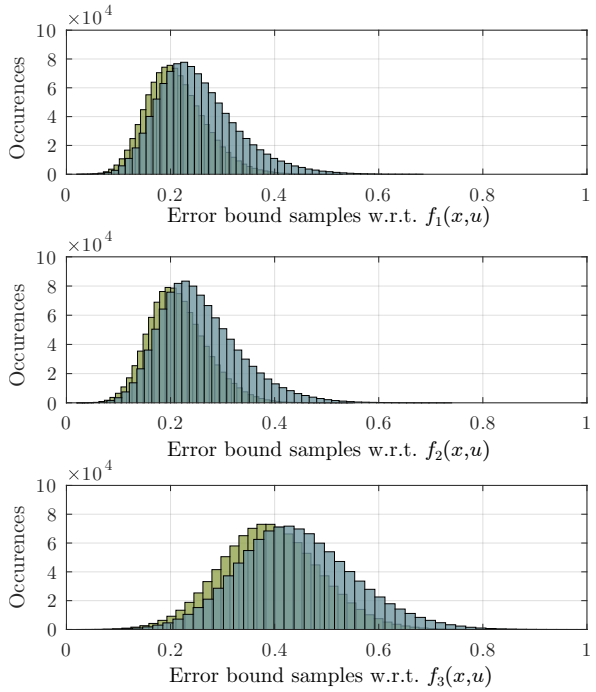

Fig. 4. Histograms showing the results of sampling the NSM (green) and SLR (blue) error bounds.

ematical program enforces the regressor to be fully consistent with the prior knowledge about the ground-truth. The obtained semi-infinite optimization problem was tackled with the scenario approach, yielding smooth regressors that were shown to be suitable to optimization-based controllers mainly due to their regularity properties.

Regarding its limitations, the deterministic error bounds associated with the SLR were shown to be more conservative when compared to the NSM case. By increasing the number of gradient constraints, this effect can be alleviated at the price of having to solve a more complex regression problem. Future investigations could address the issue of finding more suitable representations for the error bounds so as to enable efficient robust control formulations.

REFERENCES

Beliakov, G. (2006). Interpolation of lipschitz functions. *Journal of Computational and Applied Mathematics*, 196(1), 20–44.

Berkenkamp, F., Moriconi, R., Schoellig, A.P., and Krause, A. (2016a). Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes. In *Proc. Conference on Decision and Control (CDC)*, 4661–4666.

Berkenkamp, F., Schoellig, A.P., and Krause, A. (2016b). Safe controller optimization for quadrotors with gaussian processes. In *Proc. International Conference on Robotics and Automation (ICRA)*, 491–496.

Bhujwalla, Y., Laurain, V., and Gilson, M. (2016). The impact of smoothness on model class selection in nonlinear system identification: An application of derivatives in the RKHS. In *Proc. American Control Conference (ACC)*, 1808–1813.

Calafiore, G.C. and Campi, M.C. (2006). The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5), 742–753.

Calliess, J.P. (2016). Lazily adapted constant kinky inference for nonparametric regression and model-reference adaptive control. *arXiv preprint arXiv:1701.00178*.

Campi, M.C., Garatti, S., and Prandini, M. (2009). The scenario approach for systems and control design. *Annual Reviews in Control*, 33(2), 149–157.

Canale, M., Fagiano, L., and Signorile, M. (2014). Nonlinear model predictive control from data: A set membership approach. *International Journal of Robust and Nonlinear Control*, 24(1), 123–139.

Chowdhury, S.R. and Gopalan, A. (2017). On kernelized multi-armed bandits. In *Proc. International Conference on Machine Learning (ICML)*, 844–853.

Cucker, F. and Smale, S. (2002). On the mathematical foundations of learning. *Bulletin of the American Mathematical society*, 39(1), 1–49.

Formentin, S., Van Heusden, K., and Karimi, A. (2014). A comparison of model-based and data-driven controller tuning. *International Journal of Adaptive Control and Signal Processing*, 28(10), 882–897.

Fu, K.S. (1970). Learning control systems–review and outlook. *IEEE Transactions on Automatic Control*, 15(2), 210–221.

Hou, Z.S. and Wang, Z. (2013). From model-based control to data-driven control: Survey, classification and perspective. *Information Sciences*, 235, 3–35.

Kiumarsi, B., Lewis, F.L., and Jiang, Z.P. (2017). H∞ control of linear discrete-time systems: Off-policy reinforcement learning. *Automatica*, 78, 144–152.

Limon, D., Calliess, J., and Maciejowski, J.M. (2017). Learning-based nonlinear model predictive control. *IFAC-PapersOnLine*, 50(1), 7769–7776.

Manzano, J.M., Limón, D., de la Peña, D.M., and Calliess, J. (2019). Output feedback MPC based on smoothed projected kinky inference. *IET Control Theory & Applications*, 13(6), 795–805.

Milanese, M. and Novara, C. (2004). Set membership identification of nonlinear systems. *Automatica*, 40(6), 957–975.

Nikolic, Z. and Fu, K. (1966). An algorithm for learning without external supervision and its application to learning control systems. *IEEE Transactions on Automatic Control*, 11(3), 414–422.

Ostafew, C.J., Schoellig, A.P., Barfoot, T.D., and Collier, J. (2016). Learning-based nonlinear model predictive control to improve vision-based mobile robot path tracking. *Journal of Field Robotics*, 33(1), 133–152.

Rosolia, U. and Borrelli, F. (2018). Learning model predictive control for iterative tasks. A data-driven control framework. *IEEE Transactions on Automatic Control*, 63(7), 1883–1896.

Srinivas, N., Krause, A., Kakade, S.M., and Seeger, M.W. (2012). Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5), 3250–3265.

Wabersich, K.P. and Zeilinger, M.N. (2018). Linear model predictive safety certification for learning-based control. In *Proc. Conference on Decision and Control (CDC)*, 7130–7135.

Weaver, N. (2018). *Lipschitz algebras (second edition)*. World Scientific.