

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

Emilio T. Maddalena and Colin N. Jones

Abstract—We propose a non-parametric regression method that does not rely on the structure of the ground-truth, but only on its regularity properties. The methodology can be readily used for learning surrogate models of nonlinear dynamical systems from data, while providing bounds on the prediction error. In contrast with the well known Set Membership and Kinky Inference techniques that yield non-differentiable functions, the approach presented herein produces a smooth regressor. Consequently, it is more suitable to optimization-based controllers that heavily rely on gradient computations. A numerical example is provided to show the effectiveness of the method we call Smooth Lipschitz Interpolation (SLI) when compared to the aforementioned alternatives in a Model Predictive Control problem.

I. 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 [1]–[3]. 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 hot research topics [4]–[7]. 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 [8]–[10], but model-based solutions are more frequently adopted when the problem being tackled requires safety guarantees [11]–[17].

Owing to their high representative power and success in diverse applications, Gaussian Processes (GPs) are a bayesian non-parametric technique that has been exploited to deal with a variety of analysis and control problems. In [18] and [19], the authors explore GPs respectively for ground robot locomotion and tuning controller parameters for quadrotors; while the same formalism is used to estimate regions of attraction for nonlinear systems in [20]. 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 true system’s dynamics. This bound is usually expressed as a scaled version of the posterior standard deviation (see

[21], [22] for the derivations). The main challenge in these cases consists in efficiently estimating the various associated quantities such as the maximum information gain, especially in real-time applications with fast sampling frequencies.

An alternative approach to GPs which is also classified as non-parametric is the well known Nonlinear Set Membership (NSM) method presented in [23] and refined later in a series of works. By assuming Lipschitz continuity of the ground-truth, optimal upper and lower bounds on its unknown values can be obtained in closed-form. The predictor that minimizes the worst-case error over the whole domain arises simply as the mean of such bounds. The same theory was also studied in a different context by Beliakov [24], [25], and generalized to Hölder-continuous functions by Calliess, who coined the term Kinky Inference (KI) [26]. The later name emphasizes the fact that the NSM as well as the KI interpolants presents several kinks, which is clearly an obstacle for numerical optimization methods. Furthermore, they are defined on the intersection of two hyperbolic Voronoi diagrams, thus yielding piecewise-nonlinear models whose individual regions are in general not easily representable.

Contributions: In this work we propose a particular radial basis functions (RBF) interpolation methodology that overcomes the main limitation of both the NSM and KI approaches, namely non-differentiability of the regressor. The method is referred to as Smooth Lipschitz Interpolation (SLI), and can be readily used for learning non-parametric models of nonlinear dynamical systems from data while providing prediction error bounds. At the cost of possibly being more conservative, the SLI technique improves the computation times of optimization-based controllers. In contrast with the recent proposal of filtering the interpolant *a posteriori* [27], we directly obtain a function with the desired regularity properties. A numerical example is provided to show several advantages of SLI over the aforementioned alternatives in a Model Predictive Control problem.

Notation and basic definitions: \mathbb{R}^n is the n -dimensional Euclidean space endowed with the usual metric. Given a general set S , $cl(S)$ denotes its closure. $span(V)$ is the linear span of a collection of vectors V . A $n \times n$ diagonal matrix with entries c_1, \dots, c_n is written as $diag(c_1, \dots, c_n)$. $Prob(A)$ denotes the probability of event A in an appropriate probability space. \mathcal{C}^0 and \mathcal{C}^1 are respectively the space of continuous and continuously differentiable functions. A map $f : \mathcal{X} \rightarrow \mathcal{Y}$ 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

E. T. Maddalena and C. N. Jones are with Laboratoire d’Automatique, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne 1015, Switzerland. E-mails: {emilio.maddalena, colin.jones}@epfl.ch. 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).

on Lipschitz functions between Euclidean spaces. Given a function f , its gradient is a column vector denoted by ∇f .

II. PROBLEM DESCRIPTION

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

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

Assumption II.2. A collection of noise corrupted data points $D = \{(x_i, \tilde{y}_i) \mid \tilde{y}_i = f(x_i) + \delta_i, i = 1, \dots, N\}$ is available.

Assumption II.3. The following upper bounds are known:

- $\bar{L} \geq L$
- $\bar{\delta} \geq |\delta_i|, \forall i = 1, \dots, N$

An additional property of the target function is derived based solely on Assumption II.1. In combination with the data-set, this will be employed to bound the prediction error.

Proposition II.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} [28, Corollary 1.42].

Next we define the set of all possible functions that are compatible with both our regularity assumption and the available data-set. This entity receives the name of 'feasible system set' in the NSM literature; nevertheless, at times the regressor is chosen from outside this same set [29], causing the terminology to become inappropriate. For this reason, a more general nomenclature is preferred herein.

Definition II.1. The function class $\mathcal{F}_D = \{f : \mathcal{X} \rightarrow \mathcal{Y} \mid f \in \mathcal{C}^1, L_f \leq \bar{L}, |f(x_i) - \tilde{y}_i| \leq \bar{\delta}, \forall (x_i, \tilde{y}_i) \in D\}$, where L_f is a Lipschitz constant of f , is called the *consistent space*.

As shown in [23]–[25], any member $f \in \mathcal{F}_D$, and in particular the ground-truth itself, is bounded by the functions

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

where

$$f_u(x) = \min_{i=1, \dots, N} (\tilde{y}_i + \bar{\delta} + \bar{L} \|x - x_i\|) \quad (2a)$$

$$f_l(x) = \max_{i=1, \dots, N} (\tilde{y}_i - \bar{\delta} - \bar{L} \|x - x_i\|) \quad (2b)$$

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.

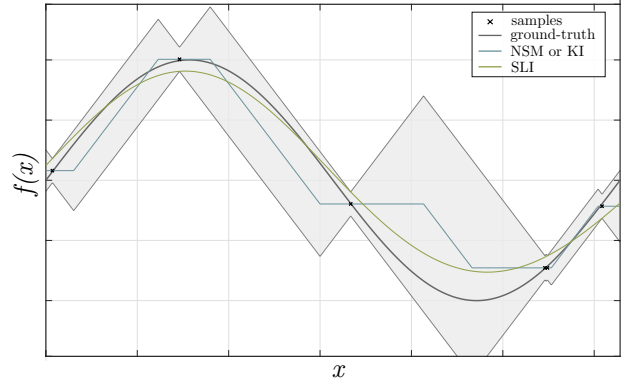


Fig. 1: A comparison between the NSM or KI and SLI approaches to interpolate samples of a sinusoidal wave. The gray area enclosed by the upper and lower bounds define the space where the unknown ground-truth can lie.

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

$$\text{IP1 : } \min_{\hat{f} \in \mathcal{H}} c(\hat{f}) \quad (3a)$$

$$\text{s.t. } \hat{f} \in \mathcal{F}_D \quad (3b)$$

where $c : \mathcal{H} \rightarrow \mathbb{R}$ is a suitable real-valued functional, and $\mathcal{H} \supseteq \mathcal{F}_D$ is a given Hilbert space.

Remark 1. The NSM approach becomes essentially a particular instance of IP1 with $c(\hat{f}) = \sup_{f \in \mathcal{F}_D} \|f - \hat{f}\|_p$, but only $\hat{f} \in \text{cl}(\mathcal{F}_D)$ is guaranteed since the regressor is not differentiable at its kinks. A closed-form solution to the optimization problem is obtained as $f^*(x) = 0.5(f_l(x) + f_u(x))$ for any L_p norm, which should match the one used in (1) and the Lipschitz constant. An illustration of this solution is depicted in Fig. 1.

III. PROPOSED SOLUTION

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

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

with *smooth* maps $\phi_k : \mathcal{X} \rightarrow \mathcal{Y}$ and constants $w_k \in \mathbb{R}, \forall k = 1, \dots, m$. The interpolant is therefore a member of the subspace $\mathcal{H} = \text{span}(\{\phi_1, \dots, \phi_m\})$ which is complete since it is finite-dimensional. 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. In the sequel, we discuss how the condition $\hat{f} \in \mathcal{F}_D$ in IP1 can be approximately translated into computational terms.

From the very choice of $\{\phi_1, \dots, \phi_m\}$, $\hat{f} \in \mathcal{C}^1$ 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 $L_{\hat{f}} \leq \bar{L}$. Since one usually knows L_{ϕ_k} for each basis function – or can numerically estimate it to any desired

precision – the upper bound $\overline{L}_{\hat{f}} = \sum_{k=1}^m |w_k| L_{\phi_k}$ 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 II.1 links the maximum norm of the gradient and the best Lipschitz constant, hence we can exploit the fact that

$$L_{\hat{f}} \leq \overline{L} \quad (5)$$

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

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

and use the last inequality. Finally, overcoming the infinite number of constraints is possible by employing the statistical results of random convex programs, a theory also known as the scenario approach [30].

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 interpolant then reads

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

$$\text{s.t. } w^T \phi(x_i) \leq y_i + \bar{\delta}, \forall i = 1, \dots, N \quad (8b)$$

$$w^T \phi(x_i) \geq y_i - \bar{\delta}, \forall i = 1, \dots, N \quad (8c)$$

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

where $w := [w_1 \dots w_m]^T$, $\phi := [\phi_1 \dots \phi_m]^T$, and $c(w)$ is an appropriate objective function. Let $w^* := \arg\min \mathbb{P}2$ and $f^*(x) := w^{*T} \phi(x)$ 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 \mathcal{F}_D only if (8d) is imposed on the whole domain.

Proposition III.1. *If $c(w)$ is a convex function and an optimal solution w^* 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., $\text{Prob}(\|\nabla f^*(x)\| > \overline{L}) < \epsilon$.*

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 [30, Theorem 1 and Corollary 1].

Remark 2. Increasing the design parameter E provides greater flexibility for $f(x)$ 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 $m = N + E \rightarrow \infty$, 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 $\mathcal{C}^0 \supset \mathcal{F}_D$ given compact domains [31]. As a side effect, $m \rightarrow \infty$ 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 $f \in \mathcal{F}_D$ 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 3. 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 $\mathbb{P}2$.

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 [32]. Therefore, Proposition III.1 ensures with practical certainty that, if a solution is found, $\|\nabla f^*(x)\|$ 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 $\|\nabla f^*(x')\| > \overline{L}$ 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 $f(x_1, x_2) = 0.02x_1^2 + 2\cos(x_2)$ and a grid of 49 samples as the available data-set. The target function, two SLI 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 $\|\nabla f^*(x)\|$ 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 .

Imposing $f_l(x) \leq \hat{f}(x) \leq f_u(x), \forall x \in \mathcal{X}$ is necessary for a prediction $\hat{f}(x')$ at any unseen query point x' to be compatible with our knowledge. In other words, the target function would never attain a value that violates this condition. The SLI methodology yields a more regular regressor at the price of possibly predicting unattainable values in small subsets of the domain. Deterministic error bounds on the prediction error can still be guaranteed as shown next.

Proposition III.2. *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 $f(x)$. Then $\bar{e}_{\hat{f}}(x) \geq |e_{\hat{f}}(x)|$ holds for all $x \in \mathcal{X}$, where*

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

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, $|e_{\hat{f}}(x)| \leq \max\{|\hat{f}(x) - f_l(x)|, |\hat{f}(x) - f_u(x)|\}$.

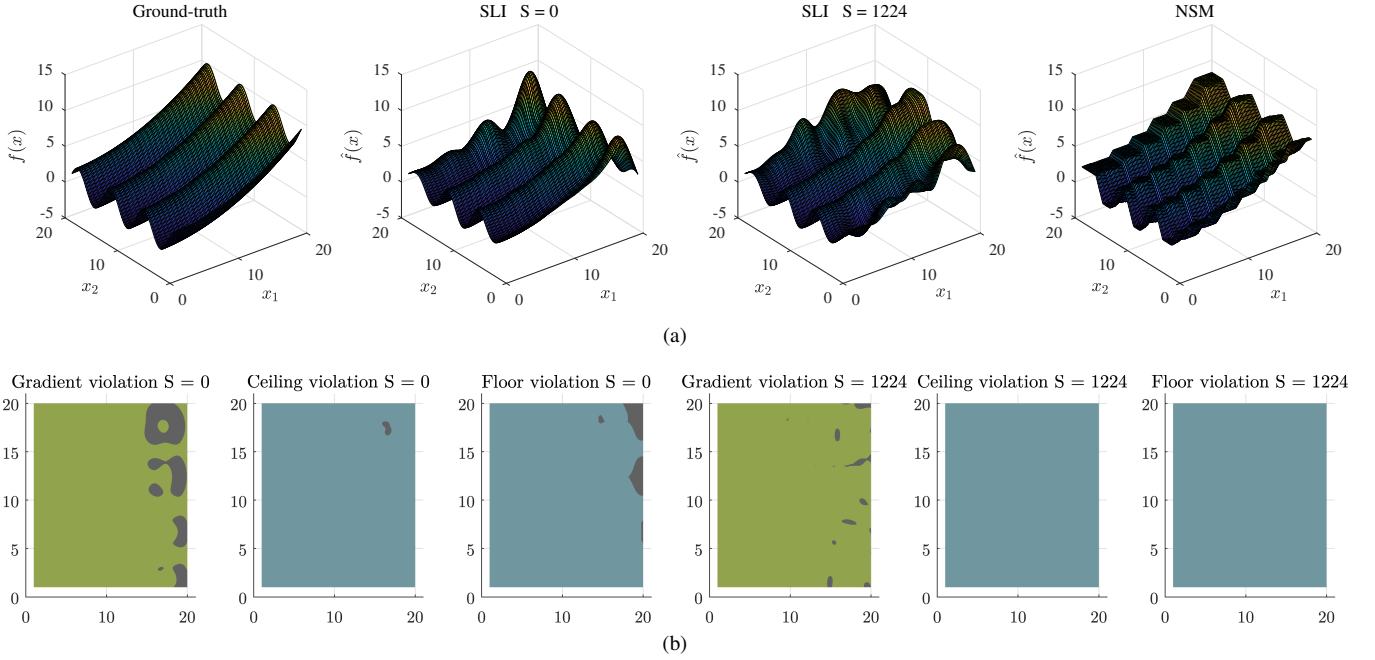


Fig. 2: (a) Target function, SLI and NSM regressors constructed based on a uniform grid of samples. (b) Effects of increasing the number of scenarios in SLI: violations areas are shown in gray for $S = 0$ (left) and $S = 1224$ (right).

IV. LEARNING DYNAMICAL SYSTEMS MODELS

The particular problem of performing regression for dynamical systems requires additional comments. Since usually some information is known about the plant to be controlled, the learning process should be carried out only for the unmodelled portion of the dynamics. Additive uncertainty is the most common model assumed in those cases [11], [33]. As an example, a mechanical system may be well described by a set of known linear differential equations, but learning may be used to capture nonlinear effects caused by static friction or backlash. Available information can be also incorporated into the data-set: if the equilibrium conditions are known a priori, the respective (x_i, y_i) pairs should be added to D . Finally, if the problem at hand is centered around a single output, autoregressive models are preferred since they allow for a tighter error bound between the real dynamics and the surrogate functions (see e.g. [15]).

From the control point of view, nonlinear MPC (NMPC) techniques have been proposed to tackle non-parametric models derived from data. Once an error bound $\bar{e}_{\hat{f}}(x)$ is established, its maximum worst-case value can be estimated offline. Next, if hard state constraints are present, a standard tightening procedure can be used along the prediction horizon to guarantee their satisfaction despite the mismatch between plant and model – assuming the set of feasible states does not eventually become empty [34]. Otherwise, if the model uncertainty is seen as a disturbance and it can be proven that the system reachable sets stay within the feasible space for all disturbance realizations, constraint satisfaction would also be verified [12]. When soft state constraints are used or only input variables are constrained, the true system

can be shown to be input-to-state stable with respect to the prediction errors under certain assumptions [15], [27].

V. NUMERICAL EXAMPLE

Consider the following discrete-time model of a unicycle

$$x_1^+ = x_1 + u_1 \cos(x_3) \quad (10a)$$

$$x_2^+ = x_2 + u_1 \sin(x_3) \quad (10b)$$

$$x_3^+ = x_3 + u_2 \quad (10c)$$

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 $x = [x_1 \ x_2 \ x_3]^T$, $u = [u_1 \ u_2]^T$; $f_1(x, u)$, $f_2(x, u)$ 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 II.1 and Lipschitz bounds $\bar{L}_1 = \bar{L}_2 = 1.8$ and $\bar{L}_3 = 3$ 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 $u \in \mathcal{U} := \{u \in \mathbb{R}^2 \mid [-0.5 \ -0.5]^T \leq u \leq [0.5 \ 0.5]^T\}$. 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 $\bar{\delta} = 0.03$.

Exponential radial basis functions

$$\phi_k(x) = \sigma^2 \exp\left(\frac{-\|x - x_k\|^2}{2l^2}\right) \quad (11)$$

$k = 1, \dots, m$ 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 $S = 4764$ scenarios. Next, in

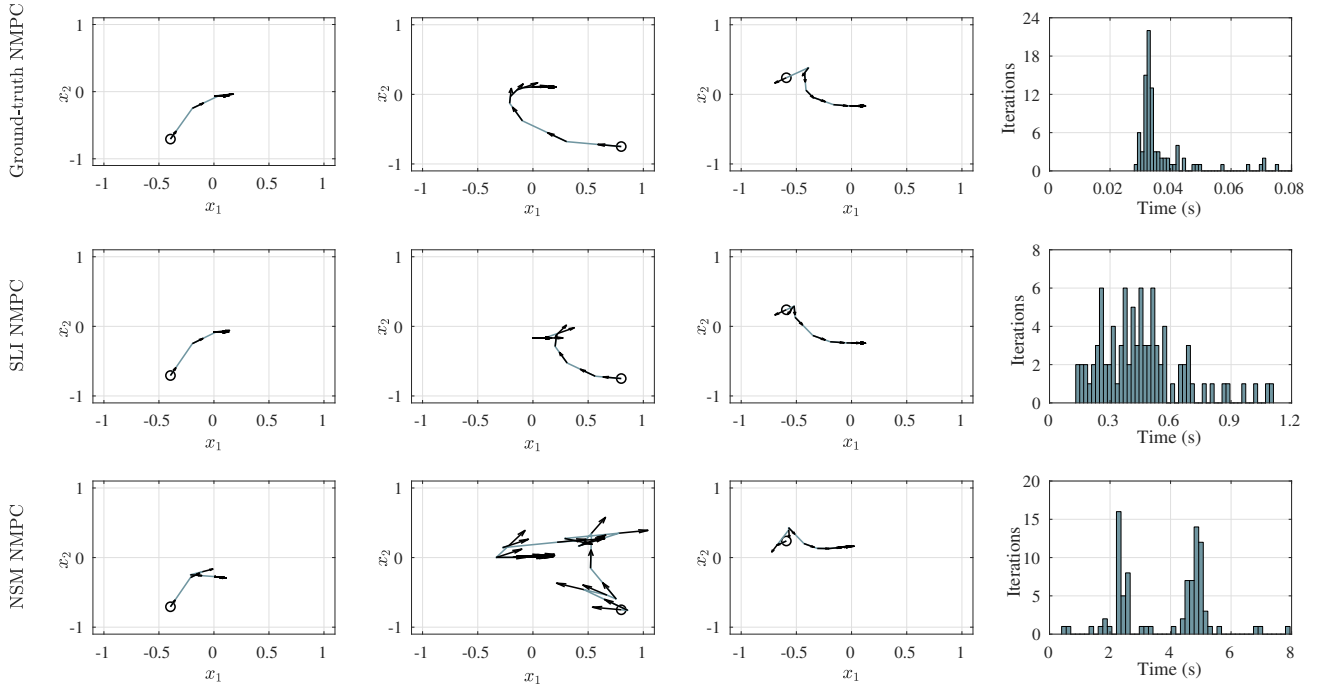


Fig. 3: Simulation results: Unicycle trajectories with NMPC designed with full knowledge of the dynamics (top row), SLI regressor (middle row) and NSM regressor (bottom row), where the same data-set was used for the last two techniques. In each of the first three columns the system evolution is depicted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

order to encourage sparsity the objective of IP2 was chosen to be $c(w) = \|w\|_1$ and the problem was solved in MATLAB with the aid of MOSEK. The optimal interpolants were found in approximately 46 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 I, where it can be seen how SLI performed better in both negatively-oriented scores for all 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 SLI may compromise its applicability later in an MPC setting since the regressor complexity is increased. The following simulations address this issue.

Three initial conditions were given to the unicycle, and NMPC controllers were designed with the goal of steering it to the origin with zero orientation. The cost function was set to $J = \sum_{t=0}^{N_{hor}} (x^T(t+1)Qx(t+1)) + (u^T(t)Ru(t))$,

with parameters $N_{hor} = 5$, $Q = \text{diag}(10^3, 10^3, 1)$, $R = \text{diag}(10^{-2}, 10^{-2})$, and hard constraints were imposed on the inputs $u(t) \in \mathcal{U}$, $t = 0, \dots, N_{hor} - 1$. First, full knowledge of (10a)–(10c) was assumed and the ground-truth functions were employed as dynamic constraints in the NMPC formulation. Next, the SLI and NSM regressors constructed with the same data-set were considered. All optimizations were carried out using `fmincon` under the same settings, except for the SLI case where analytical gradients of the constraints were given to the solver. It is important to highlight that this certainly constitutes the main advantage of the proposed technique over NSM, where gradients cannot be obtained analytically. The resulting trajectories are shown in Fig. 3 for 30 steps, with arrows indicating the unicycle orientation at each time instant. Clearly, using the SLI 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.

TABLE I: Regression performance scores comparison

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

TABLE II: NMPC computational times per iteration

	Mean (s)	Std. (s)	Worst case (s)
Ground-truth	0.0373	0.0106	0.0806
SLI	0.4585	0.2093	1.1085
NSM	3.8118	1.4684	7.9096

Computational times per controller iteration are reported in the histograms of Fig. 3 and additional measures are given in Table II ('Std.' is used for standard deviation). Despite the 20 additional free points, optimization with the SLI regressor was significantly faster than with the NSM counterpart, reducing the mean iteration time by 88% and the standard deviation by 85%. 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 SLI case.

VI. CONCLUSIONS AND FUTURE WORKS

The SLI method was proposed to interpolate data from the same class of functions studied in the NSM literature, which is also very similar to the one considered in the KI works. By exploiting suitable basis functions, the methodology introduced in this paper overcomes the non-differentiability problem present in both NSM and KI. Deterministic bounds were given for the SLI prediction error. The scenario approach was adopted to reduce the norm of the regressor gradient in random domain points. An example was given to illustrate that a finite number of scenarios can already lead to non-violation of the original ceiling and floor functions. Finally, the potential of the SLI method was verified in a NMPC problem, where not only the computational times per iteration were significantly lower compared to the NSM alternative, but also the resulting system trajectories were closer to the case of completely known dynamics. Future investigations will include a detailed comparison between SLI and NSM-based predictive controllers when handling systems with hard state constraints.

REFERENCES

- [1] Z. Nikolic and K. Fu, "An algorithm for learning without external supervision and its application to learning control systems," *IEEE Transactions on Automatic Control*, vol. 11, no. 3, pp. 414–422, 1966.
- [2] K.-S. Fu, "Learning control systems—review and outlook," *IEEE Transactions on Automatic Control*, vol. 15, no. 2, pp. 210–221, 1970.
- [3] D. H. Nguyen and B. Widrow, "Neural networks for self-learning control systems," *IEEE Control Systems Magazine*, vol. 10, no. 3, pp. 18–23, 1990.
- [4] B. Kiumarsi, F. L. Lewis, and Z.-P. Jiang, " H_∞ control of linear discrete-time systems: Off-policy reinforcement learning," *Automatica*, vol. 78, pp. 144–152, 2017.
- [5] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-dynamic programming*. Athena Scientific Belmont, MA, 1996.
- [6] H.-S. Ahn, Y. Chen, and K. L. Moore, "Iterative learning control: Brief survey and categorization," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 6, pp. 1099–1121, 2007.
- [7] U. Rosolia and F. Borrelli, "Learning model predictive control for iterative tasks. a data-driven control framework," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1883–1896, 2018.
- [8] Z.-S. Hou and Z. Wang, "From model-based control to data-driven control: Survey, classification and perspective," *Information Sciences*, vol. 235, pp. 3–35, 2013.
- [9] S. Formentin, K. Van Heusden, and A. Karimi, "A comparison of model-based and data-driven controller tuning," *International Journal of Adaptive Control and Signal Processing*, vol. 28, no. 10, pp. 882–897, 2014.
- [10] S. Tu and B. Recht, "The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint," *arXiv preprint arXiv:1812.03565*, 2018.
- [11] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013.
- [12] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, "Learning-based model predictive control for safe exploration," in *Proc. of the Conference on Decision and Control (CDC)*, 2018.
- [13] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," in *Proc. of the Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 908–918.
- [14] A. K. Akametalu, S. Kaynama, J. F. Fisac, M. N. Zeilinger, J. H. Gillula, and C. J. Tomlin, "Reachability-based safe learning with gaussian processes," in *Proc. of the Conference on Decision and Control (CDC)*, 2014, pp. 1424–1431.
- [15] M. Canale, L. Fagiano, and M. Signorile, "Nonlinear model predictive control from data: a set membership approach," *International Journal of Robust and Nonlinear Control*, vol. 24, no. 1, pp. 123–139, 2014.
- [16] F. Berkenkamp and A. P. Schoellig, "Safe and robust learning control with gaussian processes," in *Proc. of the European Control Conference (ECC)*, 2015, pp. 2496–2501.
- [17] K. P. Wabersich and M. N. Zeilinger, "Linear model predictive safety certification for learning-based control," in *Proc. Conference on Decision and Control (CDC)*, 2018, pp. 7130–7135.
- [18] C. J. Ostafew, A. P. Schoellig, T. D. Barfoot, and J. Collier, "Learning-based nonlinear model predictive control to improve vision-based mobile robot path tracking," *Journal of Field Robotics*, vol. 33, no. 1, pp. 133–152, 2016.
- [19] F. Berkenkamp, A. P. Schoellig, and A. Krause, "Safe controller optimization for quadrotors with gaussian processes," in *Proc. of the International Conference on Robotics and Automation (ICRA)*, 2016, pp. 491–496.
- [20] F. Berkenkamp, R. Moriconi, A. P. Schoellig, and A. Krause, "Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes," in *Proc. of the Conference on Decision and Control (CDC)*, 2016, pp. 4661–4666.
- [21] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.
- [22] S. R. Chowdhury and A. Gopalan, "On kernelized multi-armed bandits," in *Proc. of the International Conference on Machine Learning (ICML)*, 2017, pp. 844–853.
- [23] M. Milanese and C. Novara, "Set membership identification of non-linear systems," *Automatica*, vol. 40, no. 6, pp. 957–975, 2004.
- [24] G. Beliakov, "Interpolation of lipschitz functions," *Journal of Computational and Applied Mathematics*, vol. 196, no. 1, pp. 20–44, 2006.
- [25] —, "Smoothing lipschitz functions," *Optimisation Methods and Software*, vol. 22, no. 6, pp. 901–916, 2007.
- [26] J.-P. Calliess, "Lazily adapted constant kinky inference for non-parametric regression and model-reference adaptive control," *arXiv preprint arXiv:1701.00178*, 2016.
- [27] J. M. Manzano, D. Limon, D. M. de la Peña, and J. P. Calliess, "Output feedback mpc based on smoothed projected kinky inference," *IET Control Theory & Applications*, 2019.
- [28] N. Weaver, *Lipschitz algebras (second edition)*. World Scientific, 2018.
- [29] F. Cucker and S. Smale, "On the mathematical foundations of learning," *Bulletin of the American mathematical society*, vol. 39, no. 1, pp. 1–49, 2002.
- [30] G. C. Calafiore and M. C. Campi, "The scenario approach to robust control design," *IEEE Transactions on Automatic Control*, vol. 51, no. 5, pp. 742–753, 2006.
- [31] C. A. Micchelli, Y. Xu, and H. Zhang, "Universal kernels," *Journal of Machine Learning Research*, vol. 7, no. Dec, pp. 2651–2667, 2006.
- [32] M. C. Campi, S. Garatti, and M. Prandini, "The scenario approach for systems and control design," *Annual Reviews in Control*, vol. 33, no. 2, pp. 149–157, 2009.
- [33] J. Umlauf, L. Pöhler, and S. Hirche, "An uncertainty-based control lyapunov approach for control-affine systems modeled by gaussian process," *IEEE Control Systems Letters*, vol. 2, no. 3, pp. 483–488, 2018.
- [34] J. Manzano, D. Limon, D. M. de la Peña, and J. Calliess, "Robust data-based model predictive control for nonlinear constrained systems," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 505–510, 2018.