

Probabilistic interval predictor based on dissimilarity functions

A. D. Carnerero, D. R. Ramirez, and T. Alamo

Abstract—This work presents a new methodology to obtain probabilistic interval predictions of a dynamical system. The proposed strategy uses stored past system measurements to estimate the future evolution of the system. The method relies on the use of dissimilarity functions to estimate the conditional probability density function of the outputs. A family of empirical probability density functions, parameterized by means of two scalars, is introduced. It is shown that the proposed family encompasses the multivariable normal probability density function as a particular case. We show that the presented approach constitutes a generalization of classical estimation methods. A validation scheme is used to tune the two parameters on which the methodology relies. In order to prove the effectiveness of the presented methodology, some numerical examples and comparisons are provided.

Index Terms—Prediction intervals, system identification, nonlinear systems, uncertainty.

I. INTRODUCTION

Consider a discrete nonlinear system

$$y_k = f_0(x_k, w_k), \quad (1)$$

where $f_0(\cdot, \cdot)$ is not known, k is the discrete time instant, $y_k \in \mathcal{Y} \subseteq \mathbb{R}$ is the output of the system, w_k accounts for parametric uncertainty, noise, disturbances, etc. Also, vector $x_k \in X \subseteq \mathbb{R}^{n_x}$ represents the past inputs and outputs of the system, i.e., $x_k = [y_{k-1}, y_{k-2}, \dots, y_{k-n_y}, u_k, u_{k-1}, \dots, u_{k-n_u}]$ and $n_x = n_y + n_u + 1$. Note that nonlinear terms of past system inputs-outputs could be incorporated into vector x_k .

In this paper we focus on interval predictions. That is, given the regressor x_k , the objective is to compute an interval $I(x_k) = [y_k^-, y_k^+]$ such that we maximize the probability that y_k belongs to $I(x_k)$ while minimizing the interval width ($y_k^+ - y_k^-$). These two conflicting objectives can be reconciled if one minimizes the interval width with the constraint that $I(x_k)$ contains y_k with a pre-specified probability.

Interval predictions play a relevant role in the control of uncertain systems. Zonotopes and DC Programming are used to obtain interval state estimators in [1] and [2] respectively. Interval observers for linear time-varying systems have been proposed in [3] and [4]. Fault detection methods based on zonotopic bounds can be found in [5]. In [6], set theoretic approaches are also used in the context of fault detection. Set membership methods [7], [8] can also be used to obtain interval predictions. A mixed Bayesian/set-membership approach is proposed in [9].

There exists different methods in the literature that address the problem of obtaining interval predictions for system (1). For example, if the uncertain vector w_k is bounded and $f_0(\cdot, \cdot)$ satisfies some Lipschitz assumptions, one can resort to bounded error methods [10] that guarantee that y_k is always contained in $I(x_k)$. See, for example, [11] and [12]. Other bounded error strategies have been proposed in [13], [14], [15]. The statistical characterization of noise

This work was supported by the Agencia Estatal de Investigación (AEI)-Spain under Grant PID2019-106212RB-C41/AEI/10.13039/501100011033 and also by Junta de Andalucía and FEDER funds under grant P20_00546.

Departamento de Ingeniería de Sistemas y Automática. Universidad de Sevilla. Spain. acarnerero, danirr, talamo@us.es

and disturbances can be used to enhance the performance of interval estimation methods. See [16], [17], [18] and references therein. Also, probabilistic validation methods can be used to assess the performance of the interval predictors [19], [20], [21], [22].

Denote $F(\bar{y}|x_k)$ the cumulative distribution function of the associated output y conditioned to the regressor $x = x_k$. That is,

$$F(\bar{y}|x_k) = \text{Prob}\{y \leq \bar{y} : x = x_k\}.$$

Related with this probability is the notion of quantile [23], [24]. Given x_k , we say that \bar{y}_τ is the conditioned τ -quantile if

$$F(\bar{y}_\tau|x_k) = \text{Prob}\{y \leq \bar{y}_\tau : x = x_k\} = \tau.$$

The notion of quantile is closely related to the one of confidence intervals. The estimation of the conditioned quantiles is relevant in multiple applications (see [25] and [26]) and can be addressed using different methodologies. The most classical approach relies on the assumption that y_k and x_k are jointly normal. That is, the assumption that the (joint) probability density function of the (random) variables y and x is a multivariable normal probability density function. Under this assumption, the conditioned p.d.f. is a monovariate normal p.d.f. and the quantiles can be obtained in a simple and direct way [27]. Unfortunately, the methods based on normal distributions are very sensitive to the presence of outlier contamination. Moreover, in many long-tailed situations, the normal assumption is not well suited to characterize confidence intervals and one has to resort to non-Gaussian distributions. In these cases, generalizations of the Chebyshev inequality can be used to obtain probabilistic bounds [28], [29].

The computation of the conditioned quantiles can be also addressed by means of parametric regression techniques [24], [25]. If one assumes that there exists θ for which $y_k \approx \theta^\top x_k$, then parameter vector θ can be chosen as the one that minimizes a cost function of the error $\theta^\top x_k - y_k$. If one chooses a cost function that penalizes in an asymmetric way positive and negative errors then a quantile regressor is obtained. Given the training pairs (y_j, x_j) , $j = 1, \dots, N$ and $\tau \in (0, 1)$, the quantile regressor is defined in terms of the following optimization problem

$$\min_{\theta} \sum_{j=1}^N (1-\tau) \max\{0, \theta^\top x_j - y_j\} + \tau \max\{0, y_j - \theta^\top x_j\}.$$

This linear optimization problem penalizes the (training) errors $e_j = \theta^\top x_j - y_j$, $j = 1, \dots, N$ in an asymmetric way. The positive errors are weighted with coefficient $(1-\tau)$ and the negatives with coefficient τ . If $\tau \in (0, 1)$ is close to zero, then the positive errors will be highly penalized (in comparison with the negative ones). This means that every optimal solution θ_τ to the linear optimization problem will tend to make most of the errors negative. This implies that $\theta_\tau^\top x_k$ could be used as a probabilistic lower bound for y_k . In a similar way, a probabilistic upper bound could be obtained taking $\tau \in (0, 1)$ close to 1. Under rather mild assumptions, any minimizer θ_τ of the proposed optimization problem can be used to obtain an estimation of the τ quantile. That is, $\theta_\tau^\top x_k$ serves as an estimation of the τ quantile associated with y_k . See [24], [30] and [25] for further details.

One of the main limitations of quantile regression is that a large number of training samples N is required if one desires to obtain

probabilistic guarantees of the method when τ is chosen close to the extremes of the interval $(0, 1)$. This is due to the fact that estimating the probability of rare events requires a large number of samples. For example, the number of independent identically distributed samples required to obtain the $1 - \epsilon$ quantile of a monovariate random variable grows with $\frac{1}{\epsilon}$ (see [31], [20] and [21]).

This paper presents a new methodology for the computation of interval predictions of a dynamical system. Dissimilarity functions are used to estimate the conditional probability density function of the outputs. The estimated probability density function is used to derive the interval prediction. It is shown that the standard linear regression is a particular case of the proposed methodology. The paper is organized as follows. In Section II a family of dissimilarity functions is proposed. In Section III the role of dissimilarity functions in linear regression is analyzed. The probabilistic interval predictors are presented in Section IV. The methodology is applied to some forecasting problems in Section V. The paper ends with a section of conclusions.

II. DISSIMILARITY FUNCTIONS

Given a data set

$$\mathcal{D} = \{z_i : i = 1, \dots, N\} \subset \mathbb{R}^n,$$

we are interested in determining if a given vector z can be considered to be similar to the other vectors of the data set \mathcal{D} . In a more precise way, we are looking for a function

$$J_d(\cdot, \cdot) : \mathbb{R}^n \times \mathcal{D} \rightarrow [0, \infty]$$

that measures the dissimilarity between a given point z and the data set \mathcal{D} . Large values of $J_d(z, \mathcal{D})$ represent a high degree of dissimilarity, while small values correspond to a high degree of similarity (i.e., a small degree of dissimilarity). Clearly, from a dissimilarity function $J_d(x, \mathcal{D})$ one can obtain a similarity function $J_s(x, \mathcal{D})$. For example, given $\sigma > 0$, $J_s(x, \mathcal{D}) = e^{-\sigma J_d(x, \mathcal{D})}$ is small when z is not similar to the points in \mathcal{D} and close to 1 when z is very similar to the elements of \mathcal{D} . Another possibility would be $J_s(x, \mathcal{D}) = (1 + \sigma J_d(z, \mathcal{D}))^{-1}$, where $\sigma > 0$.

There exists a wide class of operators that can serve as dissimilarity functions for the particular case in which \mathcal{D} is a singleton ($\mathcal{D} = \{z_{\mathcal{D}}\}$). For singleton \mathcal{D} , one popular choice is

$$J_d(z, z_{\mathcal{D}}) = \|z - z_{\mathcal{D}}\|,$$

where $\|\cdot\|$ is a given norm. One could also use the minimum distance to set \mathcal{D} . That is,

$$J_d(z, \mathcal{D}) = \min_{\hat{z} \in \mathcal{D}} \|z - \hat{z}\|. \quad (2)$$

Another possibility could be to consider as a dissimilarity function the mean value of the distances of z to each member of set \mathcal{D} . See chapter 2 of [32] and chapter 2 of [33] for a review of similarity and dissimilarity functions applied in the field of image registration and in the context of cluster analysis, respectively.

Dissimilarity and similarity functions can be used in the context of regression. Suppose that we have the pairs $\{x_i, y_i\}$, $i = 1, \dots, N$ and that we would like to estimate, given x , its corresponding output y . Given the similarity function $J_s(\cdot, \cdot)$, one possibility for the estimation \hat{y} of y is

$$\hat{y} = \sum_{i=1}^N \lambda_i y_i,$$

where the scalars λ_i are chosen in such a way that λ_i is small when the similarity function $J_s(x, x_i)$ is small. It is also reasonable to

normalize the sum of the scalars λ_i to the unity. That is, $\sum_{i=1}^N \lambda_i = 1$.

For example, one could choose

$$\lambda_i = \frac{J_s(x, x_i)}{\sum_{j=1}^N J_s(x, x_j)}, \quad i = 1, \dots, N.$$

Although this approach could be valid for some applications, more sophisticated approaches are required in many situations, as it is just a weighted average. We propose in this paper a convex optimization problem to obtain a measure of dissimilarity between a point z and a set \mathcal{D} . This is formally stated in the following definition.

Definition 1: Given $z \in \mathbb{R}^n$, a set of measurements $\mathcal{D} = \{z_1, \dots, z_N\} \subset \mathbb{R}^n$ and the scalar $\gamma \geq 0$, the dissimilarity function $J_\gamma(z, \mathcal{D})$ is defined as

$$\begin{aligned} J_\gamma(z, \mathcal{D}) &= \min_{\lambda_1, \dots, \lambda_N} \sum_{i=1}^N \lambda_i^2 + \gamma \sum_{i=1}^N |\lambda_i| \\ \text{s.t.} \quad z &= \sum_{i=1}^N \lambda_i z_i \\ 1 &= \sum_{i=1}^N \lambda_i. \end{aligned} \quad (3)$$

Remark 1: Note that non negative constant weights could be included into the cost function. That is, one could consider the cost function

$$\sum_{i=1}^N w_i \lambda_i^2 + \gamma \sum_{i=1}^N |\lambda_i|,$$

where the scalars w_i , $i = 1, \dots, N$ are used to weight the different elements in \mathcal{D} . These weights could be computed using a distance function between z and the singleton $\{z_i\}$ (for example, $w_i = \|z - z_i\|$) or any dissimilarity function.

This would be a way to incorporate local information into the analysis. This strategy could be useful when the considered system is non-linear. Although the results of the paper are stated for the particular case in which $w_i = 1$, $i = 1, \dots, N$, the generalization to the general case is not difficult.

Remark 2: We notice that the optimization problem (3) could be non-feasible. In order to rule out this possibility, we assume that the vectors that compose set \mathcal{D} span all the space.

Remark 3: Optimization problem (3) is similar to the one appearing in the context of direct weight optimization and kriging, where central predictions of a certain variable are obtained by means of the solution of an optimization problem [16], [15], [34], [35], [36].

It is important to remark that the proposed dissimilarity measure is invariant with respect to affine transformations. This is formally stated in the following property.

Property 1: Consider $z_{T,v}$ and $\mathcal{D}_{T,v}$ obtained from z and \mathcal{D} through the following affine transformation.

$$\begin{aligned} z_{T,v} &= Tz + v \\ \mathcal{D}_{T,v} &= \{Tz + v : z \in \mathcal{D}\}, \end{aligned}$$

where T is any non-singular matrix and v is any vector of adequate dimensions. Then

$$J_\gamma(z, \mathcal{D}) = J_\gamma(z_{T,v}, \mathcal{D}_{T,v}).$$

Proof: We first show that any feasible solution λ_i , $i = 1, \dots, N$ to the problem of computing $J_\gamma(z, \mathcal{D})$ is also a feasible solution for the computation of $J_\gamma(z_{T,v}, \mathcal{D}_{T,v})$. Suppose that $z = \sum_{i=1}^N \lambda_i z_i$ and

$\sum_{i=1}^N \lambda_i = 1$. Then

$$\begin{aligned} z_{T,v} &= Tz + v \\ &= T \left(\sum_{i=1}^N \lambda_i z_i \right) + \left(\sum_{i=1}^N \lambda_i \right) v \\ &= \sum_{i=1}^N \lambda_i (Tz_i + v). \end{aligned}$$

We notice that $Tz_i + v$, $i = 1, \dots, N$, are the elements of $\mathcal{D}_{T,v}$. Therefore λ_i , $i = 1, \dots, N$, is also a feasible solution for the problem that defines $J_\gamma(z_{T,v}, \mathcal{D}_{T,v})$. From this we infer that $J_\gamma(z_{T,v}, \mathcal{D}_{T,v}) \leq J_\gamma(z, \mathcal{D})$. On the other hand, since T is non-singular we can make a similar reasoning to show that any feasible solution for $J_\gamma(z_{T,v}, \mathcal{D}_{T,v})$ is a feasible solution for $J_\gamma(z, \mathcal{D})$. In this way we prove also that $J_\gamma(z, \mathcal{D}) \leq J_\gamma(z_{T,v}, \mathcal{D}_{T,v})$. Both inequalities prove the claimed equality. ■

This invariance property is very important because it guarantees that the analysis based on the proposed dissimilarity function is not affected by the choice of the coordinate system. We notice that many of the dissimilarity functions that can be found in the literature are not invariant. For example, any dissimilarity function based on the distance of z to the elements of \mathcal{D} , such as that of equation (2), will be dependent on the particular choice of coordinate system.

The proposed optimization problem (3) is a strict convex optimization problem subject to convex constraints. This means that it has a unique solution [37]. From an optimization point of view, we notice that the numerical resolution can be addressed using a dual formulation. In the dual formulation for this particular optimization problem, the number of dual decision variables is equal to the number of equality constraints ($n + 1$) which is in many situations much smaller than the number of primal variables (N). On the other hand, the gradient of the objective function in the dual formulation can be obtained in a direct way because once the dual variables are fixed, the optimal values for the primal variables are obtained solving a separable optimization problem (which has an explicit solution). The numerical examples of this paper have been obtained using an accelerated gradient method in the dual variables. See [38], [39] and [40]. The alternating direction method of multipliers can also be used in this context [41].

As it is formally stated in the following property, the optimization problem has an explicit solution for the particular case $\gamma = 0$ (see Appendix A for proof).

Property 2: Suppose that $\mathcal{D} = \{z_1, z_2, \dots, z_N\}$, then $J_0(z, \mathcal{D})$ has the following explicit expression

$$J_0(z, \mathcal{D}) = N^{-1} + (z - \bar{z})^\top (ZZ^\top - N\bar{z}\bar{z}^\top)^{-1} (z - \bar{z}),$$

where $Z = [z_1 \ z_2 \ \dots \ z_N]$, $\bar{z} = N^{-1}Zu$ and $u \in \mathbb{R}^N$ is a vector with all its N components equal to 1.

The previous result shows that the dissimilarity function is a quadratic function on the argument z for the particular case $\gamma = 0$. For the more general case in which $\gamma > 0$ we can infer from the Karush-Kuhn-Tucker optimality conditions [37] that the dissimilarity function $J_\gamma(z, \mathcal{D})$ is a piecewise convex quadratic function with respect to z .

III. DISSIMILARITY FUNCTIONS AND REGRESSION

We show in this section how dissimilarity functions can be used in the context of regression. Imagine that the data set

$$\mathcal{D} = \{z_i = \begin{bmatrix} y_i \\ x_i \end{bmatrix} : i = 1, \dots, N\} \subset \mathcal{Y} \times \mathcal{X},$$

is available. Given x_k , and $\gamma \geq 0$, one could obtain an estimation \hat{y}_k for y_k minimizing the dissimilarity function of vector $\begin{bmatrix} y \\ x_k \end{bmatrix}$ with respect to the data set \mathcal{D} . That is,

$$\hat{y}_k = \arg \min_y J_\gamma \left(\begin{bmatrix} y \\ x_k \end{bmatrix}, \mathcal{D} \right).$$

Therefore, given x_k and $\gamma \geq 0$, the estimation \hat{y}_k could be obtained from the optimization problem

$$\begin{aligned} \min_{y, \lambda_1, \dots, \lambda_N} & \sum_{i=1}^N \lambda_i^2 + \gamma \sum_{i=1}^N |\lambda_i| \\ \text{s.t.} & \begin{bmatrix} y \\ x_k \end{bmatrix} = \sum_{i=1}^N \lambda_i \begin{bmatrix} y_i \\ x_i \end{bmatrix} \\ & 1 = \sum_{i=1}^N \lambda_i. \end{aligned} \quad (4)$$

Since the decision variable y appears only in the equality constraint (4), one could solve the optimization problem ignoring the equality constraint

$$y = \sum_{i=1}^N \lambda_i y_i,$$

and making the optimal value of y equal to

$$\hat{y}_k = y^* = \sum_{i=1}^N \lambda_i^* y_i, \quad (5)$$

where λ_i^* , $i = 1, \dots, N$ are the optimal values of the optimization problem

$$\begin{aligned} \min_{\lambda_1, \dots, \lambda_N} & \sum_{i=1}^N \lambda_i^2 + \gamma \sum_{i=1}^N |\lambda_i| \\ \text{s.t.} & x_k = \sum_{i=1}^N \lambda_i x_i \\ & 1 = \sum_{i=1}^N \lambda_i. \end{aligned}$$

Therefore, the estimation provided by the method is a linear combination of the outputs y_i . This is consistent with different results from the specialized literature in which it is shown that under certain assumptions, the optimal solution to an estimation problem is given by a linear combination of the observed outputs (see [16], and [42]). The central estimation provided by equation (5) is similar to other weighted methods like [15], [16] and [34].

As it is stated in the following property, the estimation obtained for the particular case $\gamma = 0$ matches the one given by the linear least squares method.

Property 3: Given a point x_k , the estimation

$$\hat{y}_k = \arg \min_y J_0 \left(\begin{bmatrix} y \\ x_k \end{bmatrix}, \mathcal{D} \right),$$

matches the estimation obtained by linear least squares using $\mathcal{D} = \{z_i = \begin{bmatrix} y_i \\ x_i \end{bmatrix} : i = 1, \dots, N\}$ as data set. The proof is provided in Appendix B.

Previous property shows that the estimation method proposed in this paper encompasses the least squares method for the particular case $\gamma = 0$. A family of optimal estimators is obtained if one considers γ as a tuning parameter. In the following sections, we

show not only how to tune the value of γ , but also how to use this methodology to obtain probabilistic interval estimations.

A. Empirical Probability Density Function

The dissimilarity of a given vector $z \in \mathcal{Z} \subseteq \mathbb{R}^n$ with respect to the elements of \mathcal{D} can be used to define an empirical probability density function. The next definition introduces the notion of empirical p.d.f.

Definition 2 (Empirical p.d.f.): Given a set of measurements $\mathcal{D} = \{z_1, \dots, z_N\} \subset \mathcal{Z}$, $\gamma \geq 0$ and $c \geq 0$, the empirical probability density function $\text{ep}_{\gamma,c}(z, \mathcal{D})$ is defined for every $z \in \mathcal{Z}$ as

$$\text{ep}_{\gamma,c}(z, \mathcal{D}) = \frac{\exp(-cJ_\gamma(z, \mathcal{D}))}{\int_{\mathcal{Z}} \exp(-cJ_\gamma(\hat{z}, \mathcal{D})) d\hat{z}}. \quad (6)$$

Note that expression (6) provides a family of probability density functions that are parameterized by constants γ and c . We notice that by construction,

$$\int_{\mathcal{Z}} \text{ep}_{\gamma,c}(\hat{z}, \mathcal{D}) d\hat{z} = 1.$$

We also notice that if \mathcal{Z} is a compact set, the choice $c = 0$ provides a uniform p.d.f. in \mathcal{Z} .

Remark 4: Recall that from property 2 we have

$$\begin{aligned} J_0(z, \mathcal{D}) &= N^{-1} + (z - \bar{z})^\top (ZZ^\top - N\bar{z}\bar{z}^\top)^{-1} (z - \bar{z}) \\ &= \frac{1}{N} + \frac{1}{N} (z - \bar{z})^\top \left(\frac{1}{N} ZZ^\top - \bar{z}\bar{z}^\top \right)^{-1} (z - \bar{z}), \end{aligned}$$

where $Z = [z_1 \ z_2 \ \dots \ z_N]$, $\bar{z} = N^{-1}Zu$ and u is a vector with all its components equal to 1. This means that if $c = \frac{N}{2}$ and $\gamma = 0$ then $\text{ep}_{\gamma,c}(z, \mathcal{D})$ is a multivariable normal distribution with mean \bar{z} and covariance matrix $\frac{1}{N}ZZ^\top - \bar{z}\bar{z}^\top$, which corresponds to the empirical covariance matrix of the data set \mathcal{D} .

As already commented, the proposed method provides a way to obtain a family of empirical probability density functions that encompasses the normal distributions and the uniform one. In order to obtain the parameters c and γ for a given data set \mathcal{D} generated by other distributions, one could use the maximum likelihood methodology. See, for example, [23]. We show in the following example how the proposed methodology can be used to estimate the probability density function that characterizes a given data set \mathcal{D} .

B. Clarifying example

A sample of 600 points in \mathbb{R} is obtained from a uniform probability function with support $[0, 1]$. One half of the available points are used as set \mathcal{D} . The other half is used as a test set. Then for every point in the test set, equation (6) is used to estimate the empirical probability density function associated with the considered pairs (c, γ) .

Figure 1 shows the empirical probability density functions estimated using different values of the parameters c and γ . In this case, $c = 1.5$ and $\gamma = 5$ is the pair that achieves the best fit for the distribution proposed in this example.

IV. INTERVAL ESTIMATION

This section presents the methodology to obtain, given $x_k \in \mathcal{X}$, an interval estimation of the corresponding output y_k . Given $x_k \in \mathcal{X}$, the data set

$$\mathcal{D} = \left\{ \begin{bmatrix} y_i \\ x_i \end{bmatrix} : i = 1, \dots, N \right\} \subset \mathcal{Y} \times \mathcal{X},$$




Fig. 1. Estimated probability distribution functions.

and the non negative scalars c and γ , the empirical **conditional** p.d.f. in $\mathcal{Y} \times \mathcal{X}$ is defined as

$$\text{ecp}_{\gamma,c}(y, x_k, \mathcal{D}) = \frac{\exp\left(-cJ_\gamma\left(\begin{bmatrix} y \\ x_k \end{bmatrix}, \mathcal{D}\right)\right)}{\int_{\mathcal{Y}} \exp\left(-cJ_\gamma\left(\begin{bmatrix} \hat{y} \\ x_k \end{bmatrix}, \mathcal{D}\right)\right) d\hat{y}}, \quad \forall y \in \mathcal{Y}.$$

The empirical conditional p.d.f. serves to model the probability of y given the occurrence of x_k . We now show how to use this notion to compute, given x_k , an interval estimation of y_k .

First, in order to simplify the numerical integration required to compute the interval estimations, we approximate set \mathcal{Y} with a set $\bar{\mathcal{Y}}$ of finite cardinality. That is, we consider the set

$$\bar{\mathcal{Y}} = \{\bar{y}_1, \dots, \bar{y}_M\},$$

where $\bar{y}_j < \bar{y}_{j+1}$, $j = 1, \dots, M-1$, and the extreme values of $\bar{\mathcal{Y}}$ are chosen to guarantee that y_k belongs to $[\bar{y}_1, \bar{y}_M]$ with high probability. A reasonable procedure to construct set $\bar{\mathcal{Y}}$ is to define $\bar{y}_1, \dots, \bar{y}_M$ as follows

$$\begin{aligned} \bar{y}_1 &= \min_{i=1, \dots, N} y_i \\ \bar{y}_M &= \max_{i=1, \dots, N} y_i \\ \bar{y}_j &= \bar{y}_1 + \left(\frac{\bar{y}_M - \bar{y}_1}{M-1} \right) (j-1), \quad j = 2, \dots, M-1. \end{aligned}$$

We notice that larger values of M provide better approximations of \mathcal{Y} at the expense of a larger computational burden. Given $x_k \in \mathcal{X}$, we define the *discrete* empirical conditional distribution, defined at each point of $\bar{\mathcal{Y}}$, as

$$\text{ecp}_{\gamma,c}(y, x_k, \mathcal{D}) = \frac{\exp\left(-cJ_\gamma\left(\begin{bmatrix} y \\ x_k \end{bmatrix}, \mathcal{D}\right)\right)}{\sum_{j=1}^M \exp\left(-cJ_\gamma\left(\begin{bmatrix} \bar{y}_j \\ x_k \end{bmatrix}, \mathcal{D}\right)\right)}. \quad (7)$$

By construction,

$$\sum_{\ell=1}^M \text{ecp}_{\gamma,c}(\bar{y}_\ell, x_k, \mathcal{D}) = 1.$$

This discrete empirical conditional p.d.f. defines a conditioned probability distribution on $\bar{\mathcal{Y}}$ (given x_k), that we denote $\text{Prob}_{\bar{\mathcal{Y}}|x_k}$. According to this discrete distribution, we have that \bar{y}_ℓ (i.e. the ℓ -th

element of $\bar{\mathcal{Y}}$) satisfies,

$$\text{Prob}_{\bar{\mathcal{Y}}|x_k} \{y \leq \bar{y}_\ell\} = \sum_{j=1}^{\ell} \overline{\text{ecp}}_{\gamma,c}(\bar{y}_j, x_k, \mathcal{D}), \quad (8)$$

$$\text{Prob}_{\bar{\mathcal{Y}}|x_k} \{y \geq \bar{y}_\ell\} = \sum_{j=\ell}^M \overline{\text{ecp}}_{\gamma,c}(\bar{y}_j, x_k, \mathcal{D}). \quad (9)$$

Given x_k , $\gamma \geq 0$ and $c \geq 0$, and $\tau \in (0, 1)$, we define the empirical upper conditioned τ -quantile, denoted by y_τ^+ , as the *smallest* element of $\bar{\mathcal{Y}}$ that satisfies

$$\text{Prob}_{\bar{\mathcal{Y}}|x_k} \{y \leq y_\tau^+\} \geq 1 - \tau.$$

In a similar way, the empirical lower conditional τ -quantile, denoted by y_τ^- , is defined as the *largest* element of $\bar{\mathcal{Y}}$ that satisfies

$$\text{Prob}_{\bar{\mathcal{Y}}|x_k} \{y \geq y_\tau^-\} \geq 1 - \tau.$$

Given x_k and $\tau \in (0, 1)$, the interval prediction for y_k is $[y_\tau^-, y_\tau^+]$. According to the discrete distribution $\text{Prob}_{\bar{\mathcal{Y}}|x_k}$, and the definition of y_τ^- and y_τ^+ we have

$$\text{Prob}_{\bar{\mathcal{Y}}|x_k} \{y \in [y_\tau^-, y_\tau^+]\} \geq 1 - 2\tau.$$

What precedes illustrates how to compute the interval prediction $[y_\tau^-, y_\tau^+]$, for a given x_k , τ and pair (γ, c) . See Algorithm 1 for a detailed description of the procedure.

Remark 5 (Conditioned median): Given the occurrence of x_k , a sensible estimation for y_k is the conditioned median $y_m(x_k, \gamma, c)$, that can be approximated by the center of the interval $[y_{0.5}^-, y_{0.5}^+]$.

Algorithm 1 Interval estimation $[y_\tau^-(x, \gamma, c), y_\tau^+(x, \gamma, c)]$.

Require: $x, \tau \in (0, 1), \gamma \geq 0, c \geq 0, \mathcal{D}, \bar{\mathcal{Y}} = \{\bar{y}_1, \dots, \bar{y}_M\}$.

Output: y_τ^-, y_τ^+ .

1: Obtain the dissimilarity function (see Definition 1) for each element of $\bar{\mathcal{Y}}$:

$$d_j = J_\gamma \left(\begin{bmatrix} \bar{y}_j \\ x \end{bmatrix}, \mathcal{D} \right), \quad j = 1, \dots, M.$$

2: Compute the conditioned probabilities (see equation 7):

$$p_j = \overline{\text{ecp}}_{\gamma,c}(\bar{y}_j, x, \mathcal{D}) = \frac{\exp(-cd_j)}{\sum_{\ell=1}^M \exp(-cd_\ell)}, \quad j = 1, \dots, M.$$

3: Compute the indexes ℓ_τ^+ and ℓ_τ^- corresponding to the lower and upper conditioned τ -quantiles (see (8) and (9)):

$$\ell_\tau^+ = \text{smallest integer } \ell \text{ satisfying } \sum_{j=1}^{\ell} p_j \geq 1 - \tau,$$

$$\ell_\tau^- = \text{largest integer } \ell \text{ satisfying } \sum_{j=\ell}^M p_j \geq 1 - \tau.$$

4: $y_\tau^- = \bar{y}_{\ell_\tau^-}$ and $y_\tau^+ = \bar{y}_{\ell_\tau^+}$.

The properties of the prediction intervals obtained using the procedure detailed above rely on the specific choice for γ and c , since they determine the underlining empirical distribution (see the example of section III-B). Given $\tau \in (0, 1)$, we now detail how to obtain a pair $(\gamma_\tau^*, c_\tau^*)$ such that sharp interval estimations are obtained, while meeting the probabilistic specifications (determined by τ) in a validation set

$$\mathcal{V} = \left\{ \begin{bmatrix} \tilde{y}_s \\ \tilde{x}_s \end{bmatrix} : s = 1, \dots, N_{\mathcal{V}} \right\} \subset \mathcal{Y} \times \mathcal{X}.$$

Let us now analyze the role of parameter $c \geq 0$ in the discrete empirical conditioned distribution given in equation (7). On the one hand, the choice $c = 0$ provides a flat distribution in which each element of $\bar{\mathcal{Y}}$ has a conditioned probability equal to $\frac{1}{M}$. On the other hand, large values of c provide narrow distributions centered around the point in $\bar{\mathcal{Y}}$ that minimizes, given x_k , the dissimilarity function $J_\gamma(\cdot, \mathcal{D})$. Consequently, for a fixed value of γ , larger values of c reduce the size of the obtained interval at the expense of increasing the fraction of outputs that are not contained in the interval estimations. Therefore, given γ , the corresponding value for c should be chosen as the largest value of c that guarantees in the validation set that the obtained intervals contain the outputs with the desired probability.

Algorithm 2 Optimal value of $c \geq 0$, for given $\gamma \geq 0$ and $\tau \in (0, 1)$

Require: $\tau \geq 0, \gamma \geq 0, c_{\max} > 0$ and $\epsilon > 0, \mathcal{D}, \bar{\mathcal{Y}}$ and the validation data set $\mathcal{V} = \left\{ \begin{bmatrix} \tilde{y}_s \\ \tilde{x}_s \end{bmatrix} : s = 1, \dots, N_{\mathcal{V}} \right\} \subset \mathcal{Y} \times \mathcal{X}$.

Output: c_γ .

1: $c_{\min} = 0$.

2: **while** $c_{\max} - c_{\min} \geq \epsilon$ **do**

3: $c = \frac{1}{2}(c_{\max} + c_{\min})$.

4: Compute, using Algorithm 1, the $N_{\mathcal{V}}$ intervals

$$I_s = [y_\tau^-(\tilde{x}_s, \gamma, c), y_\tau^+(\tilde{x}_s, \gamma, c)], \quad s = 1, \dots, N_{\mathcal{V}}.$$

5: Make n_{viol}^+ equal to the number of violations of the upper constraints

$$\tilde{y}_s \leq y_\tau^+(\tilde{x}_s, \gamma, c), \quad s = 1, \dots, N_{\mathcal{V}},$$

and n_{viol}^- equal to the number of violations of the lower constraints

$$\tilde{y}_s \geq y_\tau^-(\tilde{x}_s, \gamma, c), \quad s = 1, \dots, N_{\mathcal{V}}.$$

6: **if** $\frac{\max\{n_{\text{viol}}^+, n_{\text{viol}}^-\}}{N_{\mathcal{V}}} < \tau$ **then**

7: $c_{\min} = c$,

8: **else**

9: $c_{\max} = c$.

10: **end if**

11: **end while**

12: $c_\gamma = c_{\min}$.

From the discussion above, we have that the parameter c corresponding to a particular choice of $\gamma > 0$ (denoted c_γ) is determined by τ . As it is detailed in Algorithm 2, c_γ is chosen as the largest value of c (up to a given accuracy $\epsilon > 0$) that guarantees in the validation set that the obtained confidence intervals contain the outputs with the desired probability. That is, non smaller than $1 - 2\tau$.

Parameter $\gamma > 0$ can be obtained maximizing the likelihood ratio which, for a specific γ and corresponding c_γ , is defined as

$$L_\gamma = \sum_{s=1}^{N_{\mathcal{V}}} \log \left(\overline{\text{ecp}}_{\gamma,c_\gamma}(\tilde{y}_s, \tilde{x}_s, \mathcal{D}) \right)$$

Using $\bar{\mathcal{Y}} = \{\bar{y}_1, \dots, \bar{y}_M\}$, a numerical approximation to the optimal value of γ is given by

$$\gamma_\tau^* \approx \arg \max_{\gamma \in \Gamma} \sum_{s=1}^{N_{\mathcal{V}}} \log \left(\frac{\exp \left(-c_\gamma J_\gamma \left(\begin{bmatrix} \tilde{y}_s \\ \tilde{x}_s \end{bmatrix}, \mathcal{D} \right) \right)}{\sum_{j=1}^M \exp \left(-c_\gamma J_\gamma \left(\begin{bmatrix} \bar{y}_j \\ \tilde{x}_s \end{bmatrix}, \mathcal{D} \right) \right)} \right), \quad (10)$$

where Γ is a set containing all the possible values considered for γ .

Remark 6: Other criteria can be used to compute γ_τ^* . For example, γ_τ^* could be obtained by minimizing a cost function Q_γ that penalizes the average length of the intervals and/or the average prediction error with respect to the conditioned median introduced in Remark 5. We notice, however, that explicitly minimizing the size of the intervals may translate into an increased violation rate when the validation set has not a sufficiently large number of samples.

V. EXAMPLE

The Lorenz attractor is a system of ODEs known for having chaotic solutions with certain values of the parameters of the system. The equations that define the system are the following

$$\begin{aligned} \frac{do}{dt} &= \sigma(p - o) \\ \frac{dp}{dt} &= o(\rho - q) - p \\ \frac{dq}{dt} &= op - \beta q, \end{aligned} \quad (11)$$

where σ , ρ and β are real scalar parameters. In this example, these parameters take the values $\sigma = 10$, $\rho = 28$ and $\beta = 8/3$. Furthermore, in order to obtain the necessary data, the ODEs have been integrated numerically with a fixed time step of $T_s = 0.1s$ and initial conditions $o(0) = 1$, $p(0) = 1$ and $q(0) = 1$. Here, it is considered the task of forecasting the one-step ahead value of o , i.e., $y_k = o_k$, using the two previous values of o , that is, the regressor vector will be $x_k = [y_{k-1}, y_{k-2}]^\top$. To start with, 2500 data points are considered, normalized in the $[0, 1]$ range. Different sizes for the data set \mathcal{D} are considered in this example (200, 350 and 500 points). The validation set \mathcal{V} consists of 1000 data points and other 1000 data points are used as a test set, denoted by \mathcal{S} (note that \mathcal{D} , \mathcal{V} and \mathcal{S} are mutually disjoint sets). The set Γ is taken from $[0, 3]$ using a 0.1 sampling step. On the other hand, $\bar{\mathcal{Y}}$ is obtained from a grid of equally distant points in the interval $[-0.1893, 1.2298]$ sampled with a 1.4191×10^{-4} step.

Two different techniques will be considered as benchmarks. The first one is quantile regression [24], [25], a classical method for the estimation of conditioned quantiles. The second one is the set-membership method described in [7], [8]. This technique is a well-known method to generate interval bounds for a time series (usually produced from a dynamical system). For the sake of comparison, to guarantee that these bounds contain the output within a prescribed probability, we choose the parameters ϵ, γ of [7] such that the resulting empirical probability of containing a sample within the validation set \mathcal{V} is no smaller than $1 - 2\tau$.

The numerical results of the proposed approach and the two benchmark techniques are shown in table I for a $[o_{5\%}, o_{95\%}]$ interval, that is, $\tau = 0.05$, and in table II for $[o_{10\%}, o_{90\%}]$ ($\tau = 0.1$). The output of the test data should be contained in the first interval with a probability of 0.9 (0.8 for the second interval). The optimal value for γ has been chosen maximizing the likelihood function L_γ (see equation (10) and figure 3).

The empirical probability in the case of the quantile regression clearly does not meet the probabilistic specifications. In the case of the proposed approach and the set-membership method, the observed fraction of outputs that fall into the predicted intervals is much closer to the desired one. Note that, for all techniques, the obtained empirical probability can be below the desired probability. This could be solved relying on a probabilistic scaling scheme [22] or on probabilistic validation schemes [43], [44].

Regarding the interval width, the proposed approach clearly manages to obtain the smallest intervals for each data set. For the $[o_{5\%}, o_{95\%}]$ interval, the interval width is on average 24.35% smaller

than those provided by the set-membership method and 35.96% than the quantile regression. On the other hand, for the $[o_{10\%}, o_{90\%}]$ interval, the intervals with the proposed technique are 23.75% smaller than those with set-membership and 28.21% than those with the quantile regression. Taking into account the empirical probability values and the interval widths, we can conclude that the proposed approach obtains the best results.

Finally, in figure 2, we show the test set \mathcal{S} along with the computed intervals $[o_{5\%}, o_{95\%}]$ of the proposed approach. Note that the intervals are wider when there are trend changes in the output. Furthermore, figure 3 shows an example of the value of the maximum likelihood ratio L_γ as a function of γ (in this case for the data set of 200 points and interval $[o_{5\%}, o_{95\%}]$).




Fig. 2. Test set and computed intervals for Lorenz Attractor (interval $[o_{5\%}, o_{95\%}]$).




Fig. 3. Maximum likelihood ratio as a function of γ .

VI. CONCLUSIONS

This work presents a new approach to obtain an interval predictor to be used in nonlinear systems. The methodology relies on a parameterized family of dissimilarity functions, that are used to estimate the probability density function of the system output, conditioned to last inputs and outputs. A family of empirical probability density functions, parameterized by means of two parameters, is proposed. It is shown that the proposed family encompasses the multivariable normal probability density function as a particular case. The methodology allows us to provide probabilistic interval predictions of the output of the system. For a particular choice of the tuning parameters, the conditional probability density function of the output attains a maximum at the output estimated by least-squares regression. This shows that the proposed method constitutes a generalization of classical estimation methods. A validation scheme is used to tune the two parameters on which the methodology relies (c

TABLE I
RESULTS FOR THE LORENZ ATTRACTOR, INTERVAL $[o_{5\%}, o_{95\%}]$.

Data set length	Proposed approach		Quantile Regression		Set Membership	
	Empirical Probability	Interval Width	Empirical Probability	Interval Width	Empirical Probability	Interval Width
200	0.9140	2.0578	0.8290	3.0965	0.8960	2.9378
350	0.8990	1.9352	0.8260	3.0550	0.9100	2.4773
500	0.9070	2.0223	0.8410	3.2450	0.9120	2.5671

TABLE II
RESULTS FOR THE LORENZ ATTRACTOR, INTERVAL $[o_{10\%}, o_{90\%}]$.

Data set length	Proposed approach		Quantile Regression		Set Membership	
	Empirical Probability	Interval Width	Empirical Probability	Interval Width	Empirical Probability	Interval Width
200	0.8060	1.6053	0.7450	2.2776	0.8160	2.4248
350	0.8060	1.6164	0.7270	2.0607	0.7900	1.9797
500	0.8100	1.6195	0.7630	2.4371	0.8100	2.0021

and γ). The method has been applied to generate interval predictions, which have been compared favourably with the ones obtained by means of quantile regression and set-membership methods.

APPENDIX A

Proof of Property 2: Denote $\lambda = [\lambda_1 \lambda_2 \dots \lambda_N]^\top$. We solve the optimization problem using a dual formulation where $\mu \in \mathbb{R}^{n+1}$ denotes the multipliers associated with the equality constraint

$$z = \sum_{i=1}^N \lambda_i z_i = Z\lambda,$$

and ν is the multiplier corresponding to the equality

$$1 = \sum_{i=1}^N \lambda_i = u^\top \lambda.$$

the Lagrange function is

$$\mathcal{L}(\lambda, \mu, \nu) = \lambda^\top \lambda + \mu^\top (Z\lambda - z) + \nu(u^\top \lambda - 1).$$

Denote λ^* , μ^* and ν^* the optimal values for the primal and dual variables. From $\frac{\partial \mathcal{L}(\lambda^*, \mu^*, \nu^*)}{\partial \lambda} = 0$ we obtain that the optimal vector λ^* is given by

$$\lambda^* = -\frac{1}{2}(Z^\top \mu^* + u\nu^*). \quad (12)$$

Since $u^\top \lambda^* = 1$, $Zu = N\bar{z}$ and $u^\top u = N$ we can premultiply both terms of last equality by u^\top to obtain

$$\begin{aligned} 1 &= -\frac{1}{2}(u^\top Z^\top \mu^* + N\nu^*) \\ &= -\frac{N}{2}(\bar{z}^\top \mu^* + \nu^*). \end{aligned}$$

Therefore,

$$\nu^* = -\frac{2}{N} - \bar{z}^\top \mu^*.$$

Substituting the expression for ν^* in (12) yields,

$$\begin{aligned} \lambda^* &= -\frac{1}{2} \left(Z^\top \mu^* - u \left(\frac{2}{N} + \bar{z}^\top \mu^* \right) \right) \\ &= \frac{u}{N} - \frac{1}{2} (Z^\top - u\bar{z}^\top) \mu^*. \end{aligned} \quad (13)$$

Premultiplying by Z we obtain

$$Z\lambda^* = \bar{z} - \frac{1}{2} (ZZ^\top - N\bar{z}\bar{z}^\top) \mu^*. \quad (14)$$

From the equality constraint $Z\lambda^* = z$ and (14) we have

$$\mu^* = -2(ZZ^\top - N\bar{z}\bar{z}^\top)^{-1}(z - \bar{z}).$$

Substituting μ^* in equation (13) we infer

$$\lambda^* = \frac{u}{N} + (Z^\top - u\bar{z}^\top)(ZZ^\top - N\bar{z}\bar{z}^\top)^{-1}(z - \bar{z}).$$

Finally, taking into account that

$$u^\top (Z^\top - u\bar{z}^\top) = (N\bar{z}^\top - N\bar{z}^\top) = 0$$

we obtain

$$(Z^\top - u\bar{z}^\top)^\top (Z^\top - u\bar{z}^\top) = ZZ^\top - N\bar{z}\bar{z}^\top.$$

From last equality and the expression for λ^* we conclude

$$\begin{aligned} J_0(z, \mathcal{D}) &= (\lambda^*)^\top \lambda^* \\ &= \frac{1}{N} + (z - \bar{z})^\top (ZZ^\top - N\bar{z}\bar{z}^\top)^{-1}(z - \bar{z}). \quad \blacksquare \end{aligned}$$

APPENDIX B

Proof of Property 3: From equation (5) we have that the optimal value for the estimation is $\hat{y}_k = y^* = \sum_{i=1}^N \lambda_i^* y_i$, where for the particular case $\gamma = 0$, λ_i^* , $i = 1, \dots, N$, are the optimal values of the optimization problem

$$\begin{aligned} \min_{\lambda_1, \dots, \lambda_N} \quad & \sum_{i=1}^N \lambda_i^2 \\ \text{s.t.} \quad & x_k = \sum_{i=1}^N \lambda_i x_i \\ & 1 = \sum_{i=1}^N \lambda_i. \end{aligned}$$

Defining

$$R = \begin{bmatrix} x_1 & x_2 & \dots & x_N \\ 1 & 1 & \dots & 1 \end{bmatrix}, \quad r_k = \begin{bmatrix} x_k \\ 1 \end{bmatrix},$$

we have that the equality constraints can be rewritten as

$$R\lambda = r_k, \quad (15)$$

where $\lambda = [\lambda_1 \lambda_2 \dots \lambda_N]^\top$. From the Karush-Kuhn-Tucker optimality conditions we infer that the optimal solution is given by (see subsection 10.1.1 in [37])

$$\begin{bmatrix} I & R^\top \\ R & 0 \end{bmatrix} \begin{bmatrix} \lambda^* \\ \varphi^* \end{bmatrix} = \begin{bmatrix} 0 \\ r_k \end{bmatrix},$$

where φ^* corresponds to the optimal dual decision variables corresponding to the equality constraint (15), (see [37]). The previous equation can be rewritten as

$$\begin{aligned}\lambda^* &= -R^\top \varphi^* \\ R\lambda^* &= r_k.\end{aligned}$$

From here we obtain $-RR^\top \varphi^* = r_k$ which implies $\varphi^* = -(RR^\top)^{-1}r_k$. We finally obtain

$$\lambda^* = R^\top (RR^\top)^{-1}r_k.$$

Therefore,

$$\begin{aligned}\hat{y}_k &= Y^\top \lambda^* \\ &= Y^\top R^\top (RR^\top)^{-1}r_k \\ &= r_k^\top (RR^\top)^{-1}RY,\end{aligned}$$

where $Y = [y_1 \ y_2 \ \dots \ y_N]^\top$. We notice that this corresponds to the least squares estimation obtained when we consider as regressors the vectors $[x_j^\top \ 1]^\top$, $j = 1, \dots, N$ (see [45], [23]). ■

REFERENCES

- [1] T. Alamo, J. Bravo, and E. Camacho, "Guaranteed state estimation by zonotopes," *Automatica*, vol. 41, no. 6, pp. 1035–1043, 2005.
- [2] T. Alamo, J. Bravo, M. Redondo, and E. Camacho, "A set-membership state estimation algorithm based on DC programming," *Automatica*, vol. 44-1, pp. 216–224, 2008.
- [3] R. Thabet, T. Raïssi, C. Combastel, D. Efimov, and A. Zolghadri, "An effective method to interval observer design for time-varying systems," *Automatica*, vol. 50, no. 10, pp. 2677 – 2684, 2014.
- [4] S. Chebotarev, D. Efimov, T. Raïssi, and A. Zolghadri, "Interval observers for continuous-time LPV systems with l_1/l_2 performance," *Automatica*, vol. 58, pp. 82 – 89, 2015.
- [5] S. A. Raka and C. Combastel, "Fault detection based on robust adaptive thresholds: A dynamic interval approach," *Annual Reviews in Control*, vol. 37, no. 1, pp. 119 – 128, 2013.
- [6] F. Xu, V. Puig, C. Ocampo-Martinez, F. Stoican, and S. Oлару, "Actuator-fault detection and isolation based on set-theoretic approaches," *Journal of Process Control*, vol. 24, no. 6, pp. 947 – 956, 2014.
- [7] M. Milanese and C. Novara, "Set membership identification of nonlinear systems," *Automatica*, vol. 40, no. 6, pp. 957–975, 2004.
- [8] —, "Unified set membership theory for identification, prediction and filtering of nonlinear systems," *Automatica*, vol. 47, no. 10, pp. 2141–2151, 2011.
- [9] R. Fernandez-Canti, J. Blesa, S. Tornil-Sin, and V. Puig, "Fault detection and isolation for a wind turbine benchmark using a mixed Bayesian/Set-membership approach," *Annual Reviews in Control*, vol. 40, pp. 59 – 69, 2015.
- [10] M. Milanese, J. Norton, H. Piet-Lahanier, and E. Walter, *Bounding Approaches to System Identification*. Plenum Press, New York, 1996.
- [11] M. Milanese and C. Novara, "Set membership prediction of nonlinear time series," *IEEE Transactions on Automatic Control*, vol. 50, no. 11, pp. 1655–1669, 2005.
- [12] J. M. Manzano, D. Limon, D. M. de la Peña, and J. P. Calliess, "Robust learning-based MPC for nonlinear constrained systems," *Automatica*, vol. 117, p. 108948, 2020.
- [13] E. Bai, Y. Ye, and R. Tempo, "Bounded error parameter estimation: A sequential analytic center approach," *IEEE Transactions on Automatic control*, vol. 44, no. 6, pp. 1107–1117, 1999.
- [14] L. Jaulin, "Interval constraint propagation with application to bounded-error estimation," *Automatica*, vol. 36, pp. 1547–1552, 2000.
- [15] J. Bravo, T. Alamo, M. Vasallo, and M. Gegundez, "A general framework for predictors based on bounding techniques and local approximation," *IEEE Transactions on Automatic Control*, vol. 62, no. 7, pp. 3430–3435, 2017.
- [16] J. Roll, A. Nazin, and L. Ljung, "Nonlinear system identification via direct weight optimization," *Automatica*, vol. 41, no. 3, pp. 475–490, 2005.
- [17] J. Bravo, T. Alamo, M. Gegundez, and D. Marin, "Combined stochastic and deterministic interval predictor for time-varying systems," in *23th Mediterranean Conference on Control and Automation (MED)*. IEEE, 2015, pp. 833–839.
- [18] C. Combastel, "Zonotopes and Kalman observers: Gain optimality under distinct uncertainty paradigms and robust convergence," *Automatica*, vol. 55, pp. 265 – 273, 2015.
- [19] B. Efron and R. Tibshirani, "Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy," *Statistical science*, pp. 54–75, 1986.
- [20] T. Alamo, R. Tempo, A. Luque, and D. Ramirez, "Randomized methods for design of uncertain systems: Sample complexity and sequential algorithms," *Automatica*, vol. 52, pp. 160–172, 2015.
- [21] T. Alamo, J. Manzano, and E. Camacho, "Robust design through probabilistic maximization," in *Uncertainty in Complex Networked Systems*. Springer, 2018, pp. 247–274.
- [22] V. Mirasierra, M. Mammarella, F. Dabbene, and T. Alamo, "Prediction error quantification through probabilistic scaling," *IEEE Control Systems Letters*, vol. 6, pp. 1118–1123, 2022.
- [23] K. Murphy, *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- [24] R. Koenker and G. Bassett, "Regression quantiles," *Econometrica: Journal of the Econometric Society*, pp. 33–50, 1978.
- [25] C. Davino, M. Furno, and D. Vistocco, *Quantile Regression. Theory and Applications*. Wiley, 2014.
- [26] G. Bassett and H. Chen, "Portfolio style: Return-based attribution using quantile regression," in *Economic Applications of Quantile Regression*. Springer, 2002, pp. 293–305.
- [27] A. Papoulis and S. Pillai, *Probability, Random Variables and Stochastic Processes*. Mc Graw Hill, 2002.
- [28] J. Navarro, "A very simple proof of the multivariate Chebyshev's inequality," *Communications in Statistics-Theory and Methods*, vol. 45, no. 12, pp. 3458–3463, 2016.
- [29] B. Stellato, B. Van Parys, and P. J. Goulart, "Multivariate Chebyshev inequality with estimated mean and variance," *The American Statistician*, vol. 71, no. 2, pp. 123–127, 2017.
- [30] S. Portnoy and R. Koenker, "The Gaussian hare and the Laplacian tortoise: computability of squared-error versus absolute-error estimators," *Statistical Science*, vol. 12, no. 4, pp. 279–300, 1997.
- [31] R. Tempo, E. Bai, and F. Dabbene, "Probabilistic robustness analysis: explicit bounds for the minimum number of samples," *Systems & Control Letters*, vol. 30, pp. 237–242, 1997.
- [32] A. Goshtasby, *Image Registration. Principles, Tools and Methods*. Springer, 2012.
- [33] S. T. Wierchoń and M. Kłopotek, *Modern algorithms of cluster analysis*. Springer, 2018.
- [34] J. R. Salvador, D. R. Ramirez, T. Alamo, and D. M. de la Peña, "Offset free data driven control: application to a process control trainer," *IET Control Theory & Applications*, vol. 13, no. 18, pp. 3096–3106, 2019.
- [35] N. Cressie, "Kriging nonstationary data," *Journal of the American Statistical Association*, vol. 81, no. 395, pp. 625–634, 1986.
- [36] J. R. Salvador, D. M. de la Peña, T. Alamo, and A. Bemporad, "Data-based predictive control via direct weight optimization," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 356–361, 2018.
- [37] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [38] A. Beck, *First-order methods in optimization*. SIAM, 2017.
- [39] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [40] Y. Nesterov, *Lectures on convex optimization*. Springer, 2018, vol. 137.
- [41] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010.
- [42] M. Milanese and G. Belforte, "Estimation theory and uncertainty intervals evaluation in presence of unknown but bounded error: linear families of models and estimators," *IEEE Transactions on Automatic Control*, vol. 27, pp. 408–414, 1982.
- [43] A. D. Carnerero, D. R. Ramirez, T. Alamo, and D. Limon, "Probabilistically certified management of data centers using predictive control," *IEEE Transactions on Automation Science and Engineering*, 2021.
- [44] B. Karg, T. Alamo, and S. Lucia, "Probabilistic performance validation of deep learning-based robust NMPC controllers," *arXiv preprint arXiv:1910.13906*, 2019.
- [45] L. Ljung, *System Identification*. Prentice Hall, Upper Saddle River, NJ, 1999.