2017 IEEE 56th Annual Conference on Decision and Control (CDC)
December 12-15, 2017, Melbourne, Australia

Data-driven $\mathcal{H}_\infty$-norm estimation via expert advice

Gianmarco Rallo*, Simone Formentin*, Cristian R. Rojas**, Tom Oomen*** and Sergio M. Savaresi*

***Abstract*— $\mathcal{H}_\infty$-norm estimation is usually an important aspect of robust control design. The aim of this paper is to develop a data-driven estimation method exploiting iterative input design, without requiring parametric modeling. More specifically, the estimation problem is formulated as a *sequential game*, whose solution is derived within the *prediction with expert advice* framework. The proposed method is shown to be competitive with the state-of-the-art techniques.**

I. INTRODUCTION

In model-based control design, the quality of the controller highly depends on the accuracy of the model of the plant, in that any model mismatch may severely jeopardize the closed-loop performance. For this reason, in robust control, such a model mismatch is explicitly taken into account in the design phase [18]. Since in many applications the model of the plant under control is derived (via system identification [9]) from a set of input/output measurements, the issue of deriving *from data* also the model of the modeling error arises. This problem has been widely treated in the literature and it has led to the development of techniques such as *model error modeling* [12] or *stochastic embedding* [5].

In common robust control designs, it is not necessary to provide a full description of the model error, but a few model features are sufficient. For instance, when control design is based upon the Small Gain theorem, an estimate of the $\mathcal{H}_\infty$-norm of the model error is sufficient to guarantee the robust stability of the closed-loop system. In these cases, since the error dynamics is usually more complex than the one expressed by a low-order parametric model [17], one may focus the experimental efforts on the final objective of the identification, e.g., through a direct non-parametric estimation from data of the $\mathcal{H}_\infty$-norm of the model error.

In [16], two approaches are introduced to address the above issue: one method based on Discrete Fourier Transforms (DFT) and another one based on circulant matrices. In particular, the former method was shown to be very effective in data-driven controller design [3], [4]. Both the proposed solutions lead to one-shot accurate norm estimation, but suffer from limitations due to the finiteness of the dataset. Moreover, the method relying on circulant matrices is computationally very demanding even for simple systems.

* G. Rallo, S. Formentin and S.M. Savaresi are with the Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Via G. Ponzio 34/5, 20133 Milan, Italy. Email to: {gianmarco.rallo,simone.formentin}@polimi.it.

** C.R. Rojas is with the Department of Automatic Control, Electrical Engineering, KTH Royal Institute of Technology, S-100 44 Stockholm, Sweden. Email to: cristian.rojas@ee.kth.se.

*** T. Oomen is with the Department of Mechanical Engineering, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, Netherlands. Email to: t.a.e.oomen@tue.nl.

The procedure proposed in [7] and extended in [17] overcomes the limitations due to the finiteness of the dataset by playing with the choice of the input. Such a method could in fact be seen as an iterative input design procedure. The goal of the approach is to iteratively approximate the best input, that is a sinusoidal input of frequency corresponding to the peak of the (unknown) dynamical system. With such an excitation, the estimate of the norm of interest is the best possible for a given power constraint on the input signal. This method leads to global convergence, even for multivariable systems [10]. However, in its present formulation, measurement noise may lead to bias errors [14].

The main contribution of this paper is a new framework for estimating the $\mathcal{H}_\infty$-norm using *expert advice*, along the lines of [2]. In particular, it is shown that the iterative input design procedure can be seen as a sequential game, where – at each round – the player makes a decision (on where the peak is located) based exclusively on the advice of a certain number of "experts" (here, a pre-selected set of frequencies). The advice of each expert is to consider a sinusoid at a specific frequency as the optimal input and to employ that as the input at the next iteration. The goal of the player is to act like the best expert of the pool, by maximizing the profit derived from following its advice. At the end of each round, the performance of the experts is revealed to the player, who elaborates the strategy for the next round of the game.

Formally, input design is implemented by selecting the amplitude of the sine waves at each frequency, according to the reliability of each expert. At the end of each iteration, the performance of the experts will be assessed via a numerical score reflecting their credibility. As a final result, it will be numerically shown that the proposed procedure outperforms the approach in [14] in terms of accuracy, at the price of a higher convergence time.

It is worth to mention that the approach described in this paper has strong links to *adaptive input design* (see, e.g., [13]), while it differs from that field in two respects: firstly, the data-driven approach presented here is fully non-parametric, while adaptive input design schemes traditionally assume a parametric or semi-parametric model; secondly, adaptive schemes modify the input signal on every sample, while our approach works with independent "batches" of data, thus making this scheme generally easier to analyze.

It should also be remarked here that the expert advice framework has been rarely used in the control community, e.g., in the simple adaptive control setting [8] or within the framework of online convex optimization [11]. To the best of the authors' knowledge, this is the first time the approach of [2] is applied to an identification problem.

978-1-5090-2872-6/17/$31.00 ©2017 IEEE

The paper is organized as follows. The problem is formally stated in Section II. The estimation algorithm based on expert advice is described in detail in Section III. In Section IV, the approach is tested on a benchmark simulation example and compared with the state of the art techniques. The paper is ended by some concluding remarks.

II. PROBLEM STATEMENT

Consider the asymptotically stable single-input single-output (SISO) linear time-invariant (LTI) discrete-time system G given by:

$$y(t) = G(z)u(t) + e(t) = \bar{y}(t) + e(t), \tag{1}$$

where u denotes a quasi-stationary input [9], $\bar{y}$ is the noiseless output, e is a zero-mean stochastic noise process, independent of u, and y denotes the measured output. Throughout, t indicates the discrete time variable. The variable z denotes either the forward shift operator $zx(t) = x(t+1)$ in the time domain or the $\mathcal{Z}$ transform variable in the frequency domain, depending on the context.

The infinity norm of G equals the ℓ_2-induced norm:

$$\beta \coloneqq \|G\|_\infty = \sup_{u\in\ell_2, u\neq 0} \frac{\|Gu\|_2}{\|u\|_2} = \sup_{u\in\ell_2, u\neq 0} \frac{\|\bar{y}\|_2}{\|u\|_2}. \tag{2}$$

The latter term of (2) states that once the input u that maximizes $\|\bar{y}\|_2$ is found, then applying that input and computing $\|\bar{y}\|_2/\|u\|_2$ leads to an estimate of β. Such an input corresponds to a sinusoid whose frequency ω^o is such that $|G(e^{j\omega^o})|$ is maximum (ω^o will be called the *peak* frequency from now on).

In this work, an iterative approach to find the optimal input without requiring a model of the plant will be proposed, based on the learning approaches in [2]. At each iteration, the input signal is built as a multisine, where the frequencies of the harmonic components are

$$\omega_k = \frac{2\pi k}{NT_s}, \quad k \in \mathcal{K} = \left\{0, 1, \ldots, \frac{N}{2}\right\}, \tag{3}$$

T_s is the sample time and N indicates the number of samples of the signal, which is fixed throughout the duration of the procedure[1].

Then, the expression of the input, at the n^{th} iteration, reads

$$u^{(n)}(t) = \sum_{k=0}^{N/2} \left|U_k^{(n)}\right| \sin\left(\omega_k t + \varphi_k^{(n)}\right), \tag{4}$$

where the superscript index in parentheses reveals the iteration and $|U_k|$ (φ_k) indicates the amplitude (phase) of the harmonic with frequency ω_k. The shaping of the amplitudes will be the main objective of the procedure of this paper. Each phase φ_k, instead, is randomly sampled from the uniform probability distribution defined over the interval $[0, 2\pi]\, rad$ (the phases $\{\varphi_1, \ldots, \varphi_{N/2}\}$ are *independent and identically distributed (i.i.d.)* random variables). Alternatively, the phases can also be selected according to practical criteria, *e.g.*, the ones giving the smallest final input amplitude can be taken. This can be done, for instance, with nonlinear optimization schemes [6], or by synthesizing a binary input via a receding horizon approach [15].

[1]Eq. (3) coincides with the frequencies evaluated for the *Discrete Fourier Transform* (DFT) of a signal of length N and sample time T_s.

III. $\mathcal{H}_\infty$-NORM ESTIMATION WITH EXPERT ADVICE

In this section, the algorithm for estimating the $\mathcal{H}_\infty$-norm of a dynamical system from data is first explained and commented step by step. Then, the features of the approach are discussed and analyzed in detail.

A. The Algorithm

1. *Initialization.* Let $n = 1$. When no prior information on the dynamic system is available, select the same amplitude for all the harmonic components of the input. The value is determined according to the normalization

$$\sum_{k=0}^{N/2} \left|U_k^{(1)}\right|^2 = 1. \tag{5}$$

Consequently,

$$\left|U_k^{(1)}\right|^2 = \frac{1}{1 + N/2}, \forall k \in \mathcal{K} = \left\{0, 1, \ldots, \frac{N}{2}\right\}. \tag{6}$$

In the sequential game framework, this corresponds to assigning the same credibility (amplitude) to all the experts (frequencies ω_k) at the first instance of the game.
Notice that the first input design can be adjusted according to the previous knowledge of the dynamic system: the principle is to increase the excitation level (*i.e.* the values of $|U_k^{(1)}|$) in the region of the frequency domain where the peak is expected to be located, according to the available prior information.
Extract each phase $\varphi_k^{(1)} \in [0, 2\pi]\, rad$ (according to the uniform probability distribution). Then, build the initial input sequence $u^{(1)}$ as in (4), using the current values of the amplitudes and phases.

2. *Experiment.* Run a double-experiment, *i.e.* apply twice the same input $u^{(n)}$ to the system G, initially at rest (*i.e.*, zero initial condition). The two resulting output sequences $y_1^{(n)}$ and $y_2^{(n)}$ are obtained by collecting the measurements according to (1). The subscript index is used for distinguishing between the two experiments, which differ for the particular realizations of the additive noise e_1 and e_2 that corrupts the first and the second measurements, respectively.

3. *Experts assessment.* The idea behind the definition of the score values is to assign a higher reward to the experts ω_k's that are closer to the peak frequency ω^o.
The following estimate of the noiseless output DFT squared magnitude $|\bar{Y}_k^{(n)}|^2$ serves as a *basic* peak-proximity score $B_k^{(n)}$ of a given frequency ω_k:

$$\widehat{\left|\bar{Y}_k^{(n)}\right|^2} \coloneqq B_k^{(n)} = \Re\left(Y_{1,k}^{(n)} Y_{2,k}^{*(n)}\right). \tag{7}$$

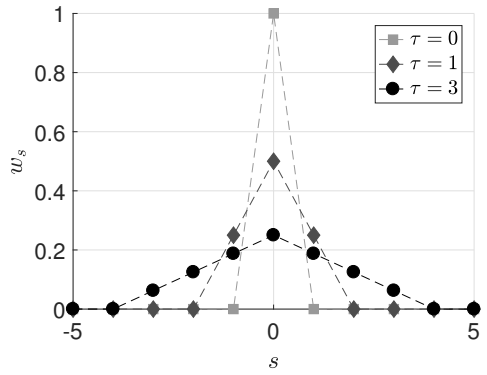

Fig. 1. Linear weight w_s for different values of the parameter τ.

Here $\Re(x)$ is the real part of the complex number x, while $Y_{1,k}^{(n)}$ and $Y_{2,k}^{(n)}$ are the DFTs of the outputs, whose values are defined for each frequency ω_k. The *star** in subscript indicates the conjugation operator. Notice that combining the results of the double experiment (in place of using $|Y_{1,k}^{(n)}|^2$ in (7)) allows to mitigate the output corruption caused by the measurement noise.

The final score S, *i.e.*, the one used for definitively evaluating the performance of each expert, exploits the defined basic scores B and it is built as follows:

$$S_k^{(n)} = \exp\left(\eta \sum_{s=-\tau}^{\tau} w_s B_{k+s}^{(n)}\right). \tag{8}$$

According to this definition, the score of the k-th expert not only depends on B_k but also on the basic scores achieved by the experts in its neighbourhood (from the $(k-\tau)$-th to the $(k+\tau)$-th). The motivation behind this choice is that S_k results more robust than B_k with respect to the noise. Moreover, this definition is in compliance with the considered application: it is, in fact, reasonable to assume that high values of B are reached also in the neighbourhood of the peak frequency, provided that the shape of the system frequency response is not infinitely sharp and the number of experts permits a sufficient sampling of the peak. The number of neighbours can be set through the design parameter τ.

Furthermore, the basic scores achieved in the neighbourhood of ω_k are weighted by the function w_s, which is designed to de-emphasize the influence of the neighbours for increasing distance s from the considered k-th expert:

$$w_s = \frac{1 - \left|\frac{s}{\tau+1}\right|}{1 + 2\sum_{s=1}^{\tau} \frac{s}{1+\tau}}, \quad s = -\tau, \ldots, \tau. \tag{9}$$

Figure 1 shows the plots of the adopted piecewise linear weighting function for $\tau = \{0, 1, 3\}$ (null values of w_s are assumed for $|s| > \tau$). When $\tau = 0$, the score of each expert does not depend on its neighbours' performance.

The argument of the exponential function in (8) is scaled by the tuning parameter η.

It is worth to mention that the choice of the scores in (8) is closely related to the so-called *multiplicative weights update algorithm* [2], [1], one of the most popular algorithms for prediction with expert advice. However, unlike the standard multiplicative weights algorithm, here the scores have been modified to weigh the contributions of neighbouring frequencies, and, as will be discussed shortly, our algorithm does not choose one frequency, that is, a pure sinusoid, but a weighting function over the entire frequency range, that is, a multi-sine spectrum.

4. ω^o *and* β *estimation.* The estimated peak frequency for the current iteration corresponds to the expert that achieves the highest score

$$\hat{\omega}^{(n)} = \omega_{\hat{k}^{(n)}}, \tag{10}$$

where

$$\hat{k}^{(n)} = \arg\max_{k \in \mathcal{K}} S_k^{(n)}. \tag{11}$$

The estimate of β reads as follows:

$$\hat{\beta}^{(n)} = \frac{\sqrt{\sum_{t=1}^{N} y_1^{(n)}(t) y_2^{(n)}(t)}}{\|u^{(n)}\|_2}. \tag{12}$$

In this expression, the numerator is an estimate of $\|\bar{y}^{(n)}\|_2$. Notice that making use of the results of the double experiment produces a consistent estimate of the ℓ_2-norm of the noiseless output, *i.e.*, the expected value of the square error tends to zero for increasing N; otherwise, relying on a single experiment (*i.e.* considering $\|y_1^{(n)}\|_2$ in the numerator of (12)) generates a noise error term whose expected value does not converge to zero asymptotically.

5. *Input design for the following iteration.* Increase the amplitude of the sinusoids whose frequencies represent the best performing experts of the current iteration:

$$\left|U_k^{(n+1)}\right|^2 = \frac{\left|U_k^{(n)}\right|^2 S_k^{(n)}}{\sum_{s=0}^{\frac{N}{2}} \left|U_s^{(n)}\right|^2 S_s^{(n)}}. \tag{13}$$

Extract each phase $\varphi_k^{(n)} \in [0, 2\pi]\, rad$ (according to the uniform probability distribution). The next input $u^{(n+1)}$ is selected according to (4).

Notice that the normalization (5) is preserved, thus the energy of the input remains constant throughout the entire estimation procedure.

6. *End of each iteration.* Let $n \mapsto n+1$ and go to step 2.

B. Discussion

Definition 1 (Nominal best expert) The *nominal best expert* of the pool, ω_{k^*}, with respect to which our algorithm has to be compared, is defined as:

$$\omega_{k^*} : \left|G\left(e^{j\omega_{k^*}}\right)\right| \geq \left|G\left(e^{j\omega_k}\right)\right|, \forall k \in \mathcal{K}. \tag{14}$$

Therefore ω_{k^*} is the best approximation of the true peak frequency ω^o (given the fixed pool defined in (3)). The procedure succeeds if, after a certain number of rounds, the designed input converges to a sinusoid with frequency ω_{k^*}.

Remark 1 (Frequency quantization) The discrepancy between ω_{k^*} and ω^o depends on the quantization of the frequency axis according to (3), *i.e.*, on the position of the experts. It follows that N and T_s represent critical tuning knobs and must be carefully selected, if they are free. A thorough analysis of this point is left to future work.

At any given round n of the sequential game, the forecasting ability of each expert is quantified by its current score $S_k^{(n)}$, which represents the only measure of its performance. In the application at hand, the frequency scores are corrupted by measurement noise. Moreover, the amplitudes of the harmonic components of the next input are designed such that more excitation is provided to the frequencies that achieved higher scores, which therefore are most likely able to increment their performance. Hence the scores at iteration $n+1$ depend on the score achieved at iteration n. There is a practical way to verify if the advice of an expert prevails over the others. Assuming that the normalization (5) holds, it is sufficient to check if the amplitude of one harmonic tends to 1 for increasing n, *i.e.*,

$$\begin{cases} \hat{k}^{(n)} = \bar{k} \\ \left| U_{\bar{k}}^{(n)} \right| \geq 1 - \epsilon \\ \left| U_{\bar{k}}^{(n+1)} \right| \geq \left| U_{\bar{k}}^{(n)} \right| \end{cases}, \quad \forall n \geq n_{\mathrm{c}}. \tag{15}$$

Here, n_{c} indicates the iteration when the algorithm has converged according to a user-defined threshold $0 < \epsilon \ll 1$. In this case, the expert

$$\bar{\omega} = \omega_{\bar{k}} \tag{16}$$

is definitively considered as the best of the pool and it represents the final estimate of the peak frequency.

Remark 2 (Counteracting noise) If (15) holds, then the input tends to a sinusoid with unitary amplitude and frequency $\bar{\omega}$ only for $n \to \infty$. However, once the convergence has been assessed – provided that $\epsilon \ll 1$ – it is possible to assume that

$$u_t^{(n)} \approx \sin(\bar{\omega} t), \quad \forall n \geq n_{\mathrm{c}}. \tag{17}$$

Then, from the n_{c}-th round, the estimate $\hat{\beta}^{(n)}$ is more sensitive to the particular measurement noise realizations (which affect the numerator in (12)) than to the quality of the input. For this reason, a cumulative moving average may be considered to smoothen the estimate for $n > n_{\mathrm{c}}$:

$$\hat{\beta}_{\mathrm{m}}^{(n)} = \begin{cases} 0, & \text{if } n < n_{\mathrm{c}} \\ \frac{1}{n - n_{\mathrm{c}} + 1} \sum_{i=n_{\mathrm{c}}}^{n} \hat{\beta}^{(i)}, & \text{if } n \geq n_{\mathrm{c}}. \end{cases} \tag{18}$$

Notice also that, when measurement noise is present, it has a doubly detrimental effect on the procedure. In fact, besides deteriorating the estimate of β, it affects the scores S_k and thus the credibility of the experts. In the next section, we will show the overall effect of noise on a benchmark simulation example.

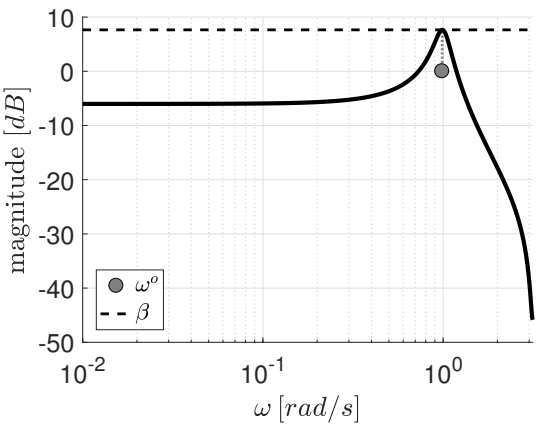

Fig. 2. Bode diagram of the magnitude of G.

IV. SIMULATION STUDY

The simulation example introduced in [14] is used for assessing the performance of the proposed estimator. Consider the discrete time (sample time $T_s = 1\,s$) LTI system

$$G(z) = \frac{0.2155 z^{-1} + 0.2012 z^{-2}}{1 - 0.9854 z^{-1} + 0.8187 z^{-2}}. \tag{19}$$

The $\mathcal{H}_\infty$ norm of this system is $\beta = 2.4114$. Such value has been computed by evaluating $|G(e^{j\omega})|$ at 5000 frequencies logarithmically distributed in $[0, \pi/T_s]\ rad/s$; the peak frequency is $\omega^o = 0.9879\,rad/s$. The Bode magnitude diagram of G is depicted in Figure 2.

Now, let us apply the proposed estimation procedure, with

$$\eta = 10, \quad \tau = 3, \quad \epsilon = 10^{-4}. \tag{20}$$

A Monte-Carlo simulation campaign of $R = 1000$ runs has been performed: for each run, $n_{\mathrm{t}} = 200$ iterations are executed. The length of the experiments is $N = 500$ samples. The measurement noise is added to the noiseless outputs with a variance of σ^2 chosen to maintain a constant signal-to-noise ratio $snr = \mathrm{var}(\bar{y}^{(n)})/\sigma^2 = 5$ in each experiment.

The nominal best expert among a pool of size 251 (corresponding to the DFT frequencies for a sample size $N = 500$) turns out to be $\omega_{k^*} = 0.9927\,rad/s$ (where the index $k^* = 80$).

The instantaneous estimate $\hat{\beta}^{(n),[r]}$ (where the superscript $[r]$ indicates the run) for the first run ($r = 1$) is depicted with a grey line in Figure 3. The convergence (defined according to (15)) has been reached after $n_{\mathrm{c}}^{[r]} = 13$ iterations. The cumulative moving average of the estimate, computed as in (18), is also shown (black solid line). The black dotted line is obtained by computing for each n the mean value of $\hat{\beta}_{\mathrm{m}}^{(n)[r]}$ over the total number of runs R.

For 576 runs $\bar{\omega} = \omega_{k^*}$ while for 423 runs the procedure converges to $\bar{\omega} = 0.9802\,rad/s$, where in this case $\bar{k} =$

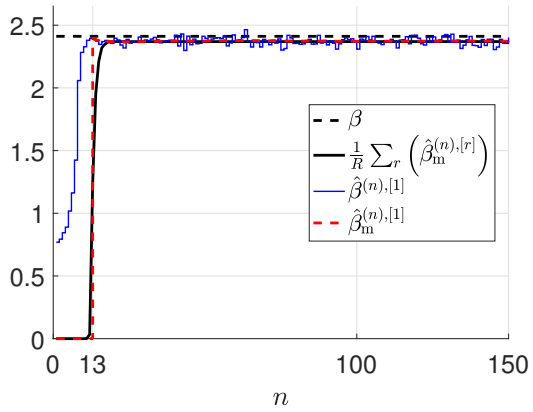

Fig. 3. $\mathcal{H}_\infty$-norm estimate trend (run $r = 1$ out of $R = 1000$).

79. Notice that, according to the values set for N and T_s, ω^o is located between ω_{79} and $\omega_{80} = \omega_{k^*}$. For one run, $\bar{\omega} = 1.0053\,rad/s$ ($\bar{k} = 81$). Therefore, the peak frequency estimate (*i.e.*, the best expert) is rather robust in this particular noisy environment characterized by a low *snr* $= 5$.

However, the noise has also impact on the instantaneous β estimation and although the cumulative moving average helps to reduce this corruption, the final value of the estimate still exhibits a small variability considering a finite number of iterations n_{t}.

Five indexes are here introduced for quantitatively evaluating the performance:

$$J_1 = \frac{1}{R}\sum_{r=1}^{R} |\beta - \hat{\beta}_{\mathrm{m}}^{(n_{\mathrm{t}})[r]}|, \tag{21}$$

$$J_2 = \frac{1}{R-1}\sum_{r=1}^{R} (|\beta - \hat{\beta}_{\mathrm{m}}^{(n_{\mathrm{t}})[r]}| - J_1)^2, \tag{22}$$

$$J_3 = \frac{1}{R}\sum_{r=1}^{R} n_{\mathrm{c}}^{[r]}, \tag{23}$$

$$J_4 = \frac{1}{R}\sum_{r=1}^{R} |\omega^o - \omega_{\bar{k}^{[r]}}|, \tag{24}$$

$$J_5 = \sum_{r=1}^{R} \left(\bar{k} \neq k^-\right) \wedge \left(\bar{k} \neq k^+\right), \tag{25}$$

where $\wedge$ is the logical operator AND, while

$$k^-, k^+ \in \mathcal{K} : \omega_{k^-} < \omega^o < \omega_{k^+},\ k^+ - k^- = 1. \tag{26}$$

The first index is the mean absolute error evaluated at the last iteration and it measures the overall quality of the β estimate; J_2 (the variance of the absolute error at n_{t}) allows us to take into account the variability of the estimate; J_3 indicates how many iterations, on average, are needed for the procedure to converge; J_4 is the mean value of the absolute peak frequency estimation error. The last index is defined assuming that the choice of N and T_s permits a proper sampling of the frequency response peak (*cf.* Remark 1), as in this case, where $k^+ = k^* = 80$.

snr [unit]	J_1 $[10^{-2}]$	J_2 $[10^{-4}]$	J_3 [iterations]	J_4 $[rad/s]$	J_5
1	4.64	1.25	13.54	$8.3 \cdot 10^{-3}$	75
5	4.2	0.17	13.76	$7.8 \cdot 10^{-3}$	1
20	4.18	0.12	13.94	$7.7 \cdot 10^{-3}$	0
100	4.16	0.1	14.06	$7.7 \cdot 10^{-3}$	0
∞	4.16	0.1	14.11	$7.7 \cdot 10^{-3}$	0

TABLE I

PERFORMANCE INDEXES FOR DIFFERENT *snr*.

The performance achieved for *snr* $= 5$ is summarized in the grey row of Table I, which also shows the results obtained for different levels of *snr* (considering same values for N, η, τ and R). The noiseless case is indicated with *snr* $= \infty$.

It is possible to observe that the precision of the estimate increases for lower noise energy levels. Notice that in the noiseless case the procedure converges to k^* for 609 runs and to $k^- = 79$ for 391 runs: this variability – that consequently affects the β estimate – is due to the random selection of the phases at each iteration.

For the considered application, the experts ω_k are equally distributed in the frequency domain and their number depending on N strongly affects the final performance. Let us fix *snr* $= 5$ (η, τ and R remain the same) and assess the performance for various N. If N increases, there are more chances for the nominal best expert ω_{k^*} to be close to ω^o. Moreover the DFT operator, used for evaluating the score, could provide a more accurate result. However, when N is high, there are two main drawbacks: i) more iterations are needed for determining the best expert (*i.e.*, to converge) since the procedure has to take into account a larger set of ω_k's; ii) when several experts reside in the neighbourhood of the peak frequency, it is more difficult to establish who is performing better since their scores are noisy and thus it is harder to converge to the nominal best expert.

For what concerns the first comment, the trade-off curve between J_1 and J_3 is shown in Figure 4 in correspondence to 5 different choices of N. The improvement of the estimation performance (*i.e.*, achievement of lower J_1) is becoming increasingly tenuous for higher N: this is due to the second drawback discussed in the previous paragraph.

A. Comparison with the power iterations method

The performance of the proposed method is here compared with the results achieved by the power iterations approach. In particular we refer to the power iterations method described in [14], which does not need a double experiment for each iteration as in the first version proposed in [17]. The reason behind this comparison (which is not exactly fair) is that the best performance (in terms of accuracy and experimental cost) of the approach described in this paper is obtained by

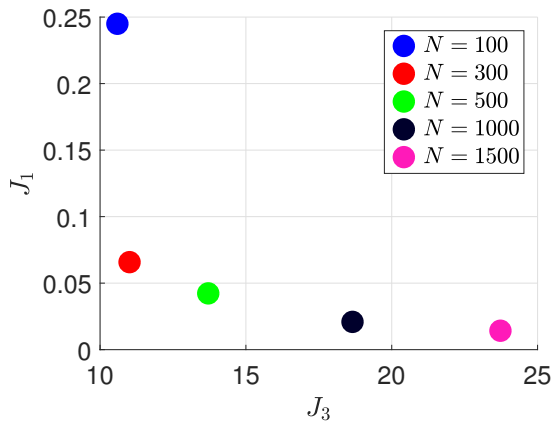

Fig. 4. Trade-off curves for different values of N.

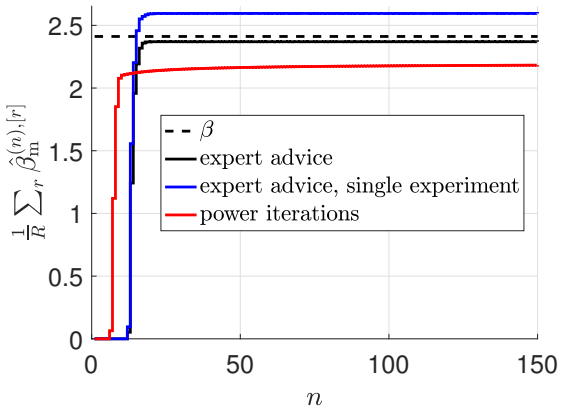

Fig. 5. Data-driven $\mathcal{H}_\infty$-norm estimate: a comparison.

exploiting the double experiment only when the procedure has converged (according to (15)).

For this simulation campaign we consider $N = 500$, $snr = 5$, $n_{\mathrm{t}} = 150$ and $R = 1000$. The design parameters η and τ are selected as in (20). The mean trends of the β estimate are shown in Figure 5. From the comparison, it comes out that the power iterations method of [14] is faster (it requires on average 10 iterations to complete its transient). However, the proposed algorithm with the double-experiment achieves the best results in terms of estimation error. Then we tested the strategy described in this paper when a single experiment is performed for each iteration. The estimation performance is worse in this case (see Figure 5) but $J_3 = 13.59$ (which is close to the value assessed for the double experiment version) and $J_5 = 0$. In fact, $\bar{k} = k^*$ for 575 runs while $\bar{k} = 79$ for 425 runs. For this reason, we can conclude that the degradation of the performance is mainly due to a higher error in the β estimation phase than to the best expert evaluation, thus the best strategy can rely on a single experiment in the first stage of the procedure.

V. CONCLUSIONS

In this paper, an iterative data-driven strategy for the $\mathcal{H}_\infty$-norm estimation is proposed within the prediction with expert advice framework.

A thorough simulation campaign on a benchmark example in several environmental conditions (noise level, length of the experiment) and for different choices of the design parameters of the algorithm has been carried out, also comparing this method with state of the art techniques. The results show that the present framework allows one to obtain more accurate estimates at the price of a slower convergence.

Ongoing work is devoted to gain further understanding of the underlying mechanisms from a theoretical perspective.

REFERENCES

[1] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.

[2] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.

[3] S. Formentin and A. Karimi. A data-driven approach to mixed-sensitivity control with application to an active suspension system. *IEEE Transactions on Industrial Informatics*, 9(4):2293–2300, 2013.

[4] 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*, 28(10):882–897, 2014.

[5] G.C. Goodwin, J.H. Braslavsky, and M.M. Seron. Non-stationary stochastic embedding for transfer function estimation. *Automatica*, 38(1):47–62, 2002.

[6] P. Guillaume, J. Schoukens, R. Pintelon, and I. Kollar. Crest-factor minimization using nonlinear Chebyshev approximation methods. *IEEE Transactions on Instrumentation and Measurement*, 40(6):982–989, 1991.

[7] H. Hjalmarsson. From experiment design to closed-loop control. *Automatica*, 41(3):393–438, 2005.

[8] P.R. Kumar. An adaptive controller inspired by recent results on learning from experts. In K.J. Åström, G.C. Goodwin, and P.R. Kumar, editors, *Adaptive Control, Filtering, and Signal Processing*, pages 199–204. Springer, 1995.

[9] L. Ljung. *System Identification: Theory for the User, 2nd Ed.* Prentice Hall, 1999.

[10] T. Oomen, R. van der Maas, C.R. Rojas, and H. Hjalmarsson. Iterative data-driven $\mathcal{H}_\infty$ norm estimation of multivariable systems with application to robust active vibration isolation. *IEEE Transactions on Control Systems Technology*, 22(6):2247–2260, 2014.

[11] M. Raginsky, A. Rakhlin, and S. Yüksel. Online convex programming and regularization in adaptive control. In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, pages 1957–1962, 2010.

[12] W. Reinelt, A. Garulli, and L. Ljung. Comparing different approaches to model error modeling in robust identification. *Automatica*, 38(5):787–803, 2002.

[13] C.R. Rojas, H. Hjalmarsson, L. Gerencsér, and J. Mårtensson. An adaptive method for consistent estimation of real-valued non-minimum phase zeros in stable LTI systems. *Automatica*, 47(7):1388–1398, 2011.

[14] C.R. Rojas, T. Oomen, H. Hjalmarsson, and B. Wahlberg. Analyzing iterations in identification with application to nonparametric $\mathcal{H}_\infty$-norm estimation. *Automatica*, 48(11):2776–2790, 2012.

[15] C.R. Rojas, J.S. Welsh, and G.C. Goodwin. A receding horizon algorithm to generate binary signals with a prescribed autocovariance. In *Proceedings of the 2007 American Control Conference (ACC)*, pages 122–127, New York, 2007.

[16] K. van Heusden, A. Karimi, and D. Bonvin. Data-driven estimation of the infinity norm of a dynamical system. In *46th IEEE Conference on Decision and Control*, pages 4889–4894. IEEE, 2007.

[17] B. Wahlberg, M. Barenthin Syberg, and H. Hjalmarsson. Non-parametric methods for $\mathcal{L}_2$-gain estimation using iterative experiments. *Automatica*, 46(8):1376–1381, 2010.

[18] K. Zhou and J.C. Doyle. *Essentials of robust control*. Prentice Hall, 1998.