Identification of linear models from quantized data: a midpoint-projection approach

Riccardo S. Risuleo, Giulio Bottegal, and Håkan Hjalmarsson *Fellow, IEEE*

***Abstract*— We consider the identification of linear models from quantized output data. We develop a variational approximation of the likelihood function which allows us to find variationally-optimal approximations of the maximum likelihood and maximum-a-posteriori estimates. We show that these estimates are obtained by projecting the midpoint in the quantization interval of each output measurement onto the column space of the input regression matrix. Interpreting the quantized output as a random variable, we derive its moments for generic noise distributions. For the case of Gaussian noise and Gaussian i.i.d. input, we give an analytical characterization of the bias which we use to build a bias-compensation scheme that leads to consistent estimates.**

***Index Terms*—**

I. Introduction

Quantization is the mapping of an analog—or continuous—variable into a discrete set of values. In the process, information contained in the original variable is lost: at any point in time, we do not know the value of the variable but only that it belongs to some *quantization interval*.

Quantization is necessary in all forms of computation and communication as computers, controllers, and network devices all operate with digital logic. While effective encoding can account for the loss of information due to noise on the digital signals, the loss of information due to quantization can only be mitigated with more refined (and expensive) analog-to-digital converters [1]. While the high-bit quantizers used in industry reduce the effect of quantization below the noise floor, there are applications where the quantization plays a significant role. For instance, in distributed sensor networks and control, designers are forced to use coarse quantizations to reduce bit rates and communication costs [2], [3]. This can have a strong impact on the accuracy of the identified models [4].

Quantization has been treated in depth in the area of digital signal processing where it has been analyzed in relation to roundoff errors in finite precision arithmetic [5]. More recently, the quantization of random variables has been interpreted as *area sampling* [6], [7] and the probability distribution of the quantized random variable can be related to the probability density function of the underlying continuous random variable. In particular, the probability of observing any one of the quantization levels is given by the probability that the continuous random variable falls within that quantization interval (see Figure 1). The relationship between the quantized and continuous densities can be analyzed using tools that remind of Nyquist's sampling theorem [7, Ch. 4] and the effect of quantization can be interpreted as an *additive uniform noise* plus an alias term.

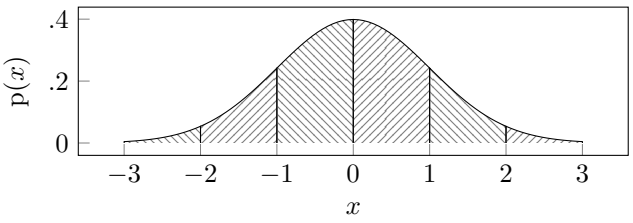

Fig. 1. The indicated areas represent the probability distribution of the outcomes of standard-normal random variable quantized with uniform quantization with unitary quantization interval (e.g., rounding to the nearest integer)

Identification of systems having quantized data is a challenging problem. The loss of information given by the presence of a quantization mechanism may cause standard prediction error-based algorithms to give unsatisfactory results. For this reason, the problem of quantization in system identification has been object of research for the last three decades; see [8] for an overview. The case of *binary quantization* has attracted particular attention. In the seminal work [9], the assumptions of periodic input and knowledge of the distribution of the disturbance are used to invert the output dynamics and estimate a finite impulse response (FIR) model consistently. The assumptions of having a known noise model and an FIR model are relaxed in the follow-up work [10]. In [11], the requirement of having a periodic input is replaced by the presence of a known dithering signal, which allows for consistent weighted least-squares estimation of an FIR model. Identification methods for FIR systems with binary quantizers have also been developed in an on-line setting, see [12], [13]. The main idea in these works is to update the parameter estimate when there is a mismatch between the

This work was supported in part by the Swedish Research Council under the project *System identification: Unleashing the algorithms* (contract 2015-05285) and under the research environment *NewLEADS - New Directions in Learning Dynamical Systems* (contract 2016-06079) and in part by the European Research Council (ERC) Horizon 2020 research and innovation programme under the *Advanced Research Grant SYSDYNET* (grant agreement No 694504).

R.S. Risuleo and H. Hjalmarsson are with the School of Electrical Engineering and Computer Science, KTH - Royal Institute of Technology, Stockholm, Sweden ({risuleo, hjalmars}@kth.se).

G. Bottegal is with the Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands (g.bottegal@tue.nl).

true and the predicted outputs. Experiment design techniques for binary quantizers are developed in [14], [15], where the problem is cast under a set membership system identification framework, and in [16], which describes a discrete-valued input design technique based on graph theory. More recently, the interest for system identification from quantized data has shifted toward multiple-level quantizers. A recursive algorithm for maximum likelihood identification of ARMA models is developed and studied in [17]; authors of [18] propose a two-step asymptotically efficient algorithm for systems with quantization blocks both on the input and the output side. In [19], the problem of FIR identification of a system with quantized data is formulated as a quadratic program where both the model parameters and the non-quantized output are estimated simultaneously. For the general case of *uniform quantizers*, estimators based on instrumental variables and quantiles are proposed in [20] and [21], respectively. Identification techniques for block-oriented nonlinear models with quantized output data are also discussed in the literature. In [22], an algorithm for Hammerstein systems is proposed by extending the methodology proposed in [9]. In [23], identification of Wiener systems with binary observations is tackled using Support Vector Machines (SVM); an analogous technique is also discussed in [24] for the standard linear case.

The focus of this paper is on *maximum likelihood* and Bayesian point estimators. Within the framework of quantized output data, methods relying on these estimators have been developed in recent years [25]–[29], using the Expectation-Maximization (EM) algorithm [30] as a pivotal technique to build iterative estimation schemes. Typically, the "expectation" step of the EM algorithm requires computing an integral which, in the framework under study, is difficult to evaluate—unless binary quantizers are considered, see [28]. To circumvent this issue, various strategies for numerical integration have been proposed, such as the scenario approach [26] or Markov Chain Monte Carlo methods [27], [29]. While these identification techniques generally provide accurate system estimates, numerical integration suffers from the curse of dimensionality, making the computational burden grow fast with the data sample size.

The computational complexity issue motivates the development of this work, where we also tackle the problem of identifying a linear model from quantized data. In particular, we consider a linear-in-the-parameters model with additive Gaussian white noise and we estimate the parameters using maximum-likelihood and maximum-a-posteriori approaches. Because the likelihood function is not available in closed form, we propose a variational approximation that allows us to obtain simple expressions for the approximate maximum likelihood and maximum-a-posteriori estimators. In particular, the analysis we propose shows that the variationally-optimal approximation of the maximum likelihood estimate of the system consists in a least-squares projection of the middle point of the active quantization intervals onto the subspace spanned by the column of the system regression matrix. Similarly, the variationally-optimal approximation of the maximum-a-posteriori estimate of the system consists in a penalized least-squares projection of the middle point of the active quantization intervals onto the same subspace. By relying on linear transformations of the data, the proposed approximate estimates are extremely simple to compute and can be used to initialize iterative methods for maximum likelihood/a-posteriori identification.

To derive the results, we develop a new variational lower bound on the normalization constant of the truncated Gaussian density. While many approaches for Gaussian integration exist in the literature [31]–[34], these usually aim at providing a good approximation for the Gaussian cumulative density function over the whole real line. In contrast, we are only interested in approximating the cumulative density function over the quantization intervals. In addition, the closed-form approximations available in the literature are intended for numerically accurate and stable evaluation and will not result in simple formulas for estimation.

As part of the contributions of this paper, we develop an analysis of the effect of the quantization on the output measurements in the special case of uniform quantizers. Interpreting quantization as a noise perturbation around the midpoint of the quantization interval, we provide expressions for its first and second moments. Using this analysis, we show that in general the approximate estimates proposed are asymptotically biased. However, for the case where the input is an i.i.d. Gaussian sequence we provide an analytic characterization of the bias, which we then use to construct a bias-compensation scheme that leads to consistent estimates of the impulse response.

A preliminary version of this work was presented in the conference contribution [35], where the variationally-optimal approximate identification was developed. In this paper, we extend the result by analyzing the asymptotic performance of the estimator and formulating a bias-compensation scheme that provides consistent estimates.

The rest of the paper is organized as follows. In Section II, we present the problem of system identification from quantized measurements. In Section III, we present some properties of the maximum-likelihood and maximum-a-posteriori estimators. In Section IV, we present a variational approximation of the truncated Gaussian distribution. In Section V, we present the variationally-optimal approximate estimators. In Section VI, we analyze the bias of the approximate maximum likelihood estimator for general noise distributions and we propose a consistent bias-compensated estimator for the case of Gaussian measurement noise and Gaussian input signal. In Section VIII, we present various simulation experiments to validate the results. In Section IX we draw some conclusions and present future research directions. In the Appendix, we present the proofs of the main results.

II. Problem formulation

We consider the problem of estimating linear systems from quantized output data. We consider a known discrete-time input signal u_k, with time index k, fed into a linear system G_0. The noise-corrupted output w_k of the linear system is then quantized by a quantizer $\mathcal{Q}$ (see Figure 2).

We suppose that the linear system is time-invariant and stable, so that the latent signal w_k can be represented by the noise-corrupted discrete convolution

$$w_k = U_k^T g_0 + \varepsilon_k, \tag{1}$$

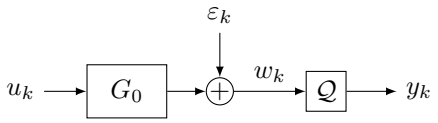

Fig. 2. The setup considered in this paper.

where g_0 is the $n \times 1$ vector of nonzero impulse response samples of G_0, U_k is the $n \times 1$ vector of input samples $[U_k]_i = u[k-i]$, and ε_k is a Gaussian white-noise sequence with variance σ^2.

Note that (1) may represent a finite impulse response model of order n or an output-error model (where the order n is chosen so that $g_k \approx 0$ for $k > n$).

We suppose that $\mathcal{Q}$ is an infinite quantizer with quantization levels L_n, so that

$$\mathcal{Q}(x) = \sum_{n=-\infty}^{\infty} n \, \mathbf{1}_{[L_n, L_{n+1})}(x),$$

where $\mathbf{1}_a(\cdot)$ is the indicator function of the interval a. In other words, the output of the quantizer is an integer n that tells us that the value x is in the nth quantization interval—that is, $\mathcal{Q}(x) = n$ if $x \in [L_n, L_{n+1})$ (see Figure 3).

The problem we consider is thus as follows.

Problem 1: Given N samples of the quantized output $y_k = \mathcal{Q}(w_k)$ of the linear system (1) subject to the input u_k, estimate the impulse response g_0.

In the next section, we propose a solution to the problem based on the likelihood principle.

III. Maximum-likelihood estimation from quantized measurements

In this section, we analyze Problem 1 from a maximum-likelihood perspective. In particular we derive the likelihood function and we propose two identification approaches: the first based on maximum-likelihood, the second based on maximum-a-posteriori probability.

A. Maximum-likelihood estimation

To derive the likelihood function, note that the sequence of outputs y_k tells us the active quantization intervals at each time step; in particular we know that $w_k \in \big[L_{y_k}, L_{y_k+1}\big)$.

From the Gaussian assumption on the noise, we have that the latent variable w_k is a Gaussian random variable with mean $U_k^T g_0$ and variance σ^2. Hence, we can find the distribution of the (discrete) random variable y_k as follows

$$\begin{aligned}\mathbf{P}\big\{y_k = n \big| \, g_0\big\} &= \mathbf{P}\big\{w_k \in [L_n, L_{n+1})\big\} \\ &= \int_{L_n}^{L_{n+1}} N(z|U_k^T g_0, \sigma^2) \, \mathrm{d}z, \end{aligned} \tag{2}$$

where $N(x|m, s^2)$ denotes the probability density function of a normal random variable with mean m and variance s^2:

$$N(x|m, s^2) = \frac{1}{\sqrt{2\pi s^2}} \exp\left\{-\frac{1}{2s^2}(x-m)^2\right\}. \tag{3}$$

So, the probability of the event $\{y_k = n\}$ is equal to the normalization constant of a Gaussian density truncated between L_n and L_{n+1}. From the white-noise assumption on ε_k, we can write the likelihood function as

$$\mathrm{p}(\{y_k\}_{k=1}^N \,|\, g) = \prod_{k=1}^{N} \int_{L_{y_k}}^{L_{y_k+1}} N(z|U_k^T g, \sigma^2) \, \mathrm{d}z. \tag{4}$$

Once we have the likelihood function, we can find the maximum-likelihood estimate of the impulse response solving

$$g^{\mathrm{ML}} = \arg\max_g \prod_{k=1}^{N} \Phi\left(\tfrac{L_{y_k+1} - U_k^T g}{\sigma}\right) - \Phi\left(\tfrac{L_{y_k} - U_k^T g}{\sigma}\right) \tag{5}$$

where $\Phi(\cdot)$ is the standard-normal cumulative distribution function. Note that the cost function is nonconvex and possibly high-dimensional; so, to find the maximum likelihood estimate we have to resort to search methods, which may be difficult to initialize and converge to local solutions.

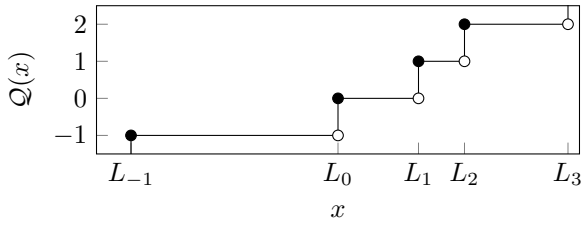

Fig. 3. Example of an infinite quantizer: the output is n when the input is between L_n and L_{n+1}.

One key result is that we can interpret the maximum likelihood estimate as the least-squares projection of a point in the interior of the active quantization intervals. To this end, define the $N \times n$ Toeplitz matrix U such that the kth row is the vector U_k^T; then, we have the following result:

Lemma 2: Let $\{y_k\}_{k=1}^N$ be the sequence of the quantized output of the system (1) to the input $\{u_k\}_{k=1}^N$. Then, there exists a vector $z^\star$, with $z_k^\star \in (L_{y_k}, L_{y_k+1})$, such that the maximum-likelihood solution (5) can be written as $g^{\mathrm{ML}} = \big(U^T U\big)^{-1} U^T z^\star$.

Proof: Consider the integral in (2). By the mean-value theorem, there exists an $z_k^\star \in (L_{y_k}, L_{y_k+1})$ such that

$$\int_{L_{y_k}}^{L_{y_k+1}} N(z|U_k^T g, \sigma^2) \, \mathrm{d}z = (L_{y_k+1} - L_{y_k}) N(z_k^\star|U_k^T g, \sigma^2).$$

Then, the logarithm of (4) can be written as

$$\log \mathrm{p}(\{y_k\}_{k=1}^N \,|g) = -\frac{1}{2\sigma^2} \sum_{k=1}^{N} \Big(z_k^\star - U_k^T g\Big)^2 + C,$$

where C is a constant independent of the impulse response g. The first-order optimality condition gives the result. ∎

While this result is not constructive—in the sense that it does not give us any way to find the point $z^\star$—it tells us that for sure there is a point inside the active quantization intervals that gives the maximum-likelihood estimate when projected onto the span of U. Hence, we may follow a *best worst-case* approach and take the midpoints of the active quantization intervals,

$$[\bar{z}]_k = \bar{z}_k = \frac{L_{y_k+1} + L_{y_k}}{2}, \tag{6}$$

to perform the estimation. For specific problems, this may return the maximum likelihood estimate, as shown in the following result.

Proposition 3: Under the assumptions of Lemma 2, suppose that there exists a $\hat{g}$ such that $\bar{z} = U\hat{g}$. Then, $\hat{g}$ is a local solution of (5).

Proof: See Appendix A

B. Bayesian estimation

In situations where data are scarce or noisy, the maximum likelihood estimate (5) may suffer from overfitting or high variance. In these situations, it may be beneficial to use regularization and kernel-based methods to obtain a favorable shift of the bias–variance tradeoff [36], [37]. From the Bayesian interpretation of kernel-based methods, we consider a Gaussian-process model for the impulse response with a kernel function $K(\cdot,\cdot)$. In other words, we suppose that the vector g of impulse response samples has a multivariate Gaussian distribution given by

$$\mathrm{p}(g) = N(g|0, K), \tag{7}$$

where the covariance matrix K is determined by the kernel as $[K]_{i,j} = K(i,j)$. Then, we can find the maximum a posteriori estimate of the impulse response as

$$g^{\mathrm{MAP}} = \arg\max_{g} \mathrm{p}(g) \prod_{k=1}^{N} \Phi\left(\tfrac{L_{y_k+1} - U_k^T g}{\sigma}\right) - \Phi\left(\tfrac{L_{y_k} - U_k^T g}{\sigma}\right). \tag{8}$$

Also in this case, the solution can be characterized as an appropriate penalized least-squares projection of an interior point of the active quantization interval:

Lemma 4: Under the assumptions of Lemma 2, there exists a vector $z^\star$, with $z_k^\star \in (L_{y_k}, L_{y_k+1})$, such that the maximum-a-posteriori solution (8) can be written as $g^{\mathrm{MAP}} = \left(U^T U + \sigma^2 K^{-1}\right)^{-1} U^T z^\star$.

Proof: Follows from the same arguments as Lemma 2.

As mentioned in the previous section, neither Lemma 2 nor Lemma 4 are constructive. There is, in other words, no way to find the point $z^\star$ except for numerical optimization—which is equivalent to solving (5) or (8) directly. However, we may intuitively consider the midpoint of the active quantization intervals $\bar{z}$ as a good approximation of $z^\star$. In the next section, we show that this choice is optimal in the sense that it minimizes a variational approximation of the likelihood function.

IV. Variational approximation of truncated Gaussian distributions

To simplify the notation, in this section we let $N(x)$ denote the Gaussian probability density function $N(x|m, s^2)$ defined in (3). Note that $N(x)$ is convex in $(x-m)^2$; hence, it admits a convex dual function which approximates it pointwise from below [38, Section 3.3], according to the following lemma.

Lemma 5: Let $N(x)$ be the probability density function of a Gaussian random variable with mean m and variance s^2. For any $x, \xi \in \mathbf{R}$,

$$N(x) \geq -\frac{N(\xi)}{2s^2}\left[(x-m)^2 - (\xi-m)^2 - 2s^2\right], \tag{9}$$

and, for any $x \in \mathbf{R}$, there exists a $\xi^\star(x)$ such that (9) holds with equality.

Proof: $N(x)$ is convex in $(x-m)^2$; hence, by convex duality, there exists a dual function $N^*(\lambda)$, given by

$$N^*(\lambda) = \max_{(x-m)^2} \left[\lambda (x-m)^2 - N(x)\right], \tag{10}$$

such that $N(x) = \max_\lambda \left[\lambda (x-m)^2 - N^*(\lambda)\right]$. From the first-order condition, we find that the maximum with respect to x in (10) is attained when

$$\lambda + \frac{N(x)}{2s^2} = 0.$$

Let $x = \xi$ be the point where this maximum is attained, then the dual function is given by

$$N^*\big(\lambda(\xi)\big) = \lambda(\xi)\,(\xi-m)^2 + 2s^2\lambda(\xi).$$

Plugging this expression into the convex duality, we have

$$N(x) = \max_{\xi}\Big[\lambda(\xi)(x-m)^2 - \lambda(\xi)(\xi-m)^2 - 2s^2\lambda(\xi)\Big].$$

from which, the result follows. ■

Now, we can use the result in Lemma 5 to find a lower bound on the normalization constant of the truncated Gaussian (2). To this end, we first approximate the Gaussian density over the integration interval, then we integrate the approximation. By maximizing with respect to the variational parameter, we find the best lower bound on the normalization constant. The result is summarized in the following lemma.

Lemma 6: Let $N(x)$ be the probability density function of a Gaussian random variable with mean m and variance s^2; then, for any $a, b \in \mathbf{R}$, $a < b$, we have that

$$\int_a^b N(x)\,\mathrm{d}x \geq I(a,b); \tag{11}$$

where

$$I(a,b) = \frac{b-a}{\sqrt{2\pi s^2}} \exp\left\{-\frac{(b-m)^3 - (a-m)^3}{6s^2\,(b-a)}\right\}. \tag{12}$$

Proof: See Appendix B.

As stated in Lemma 5, (9) holds with equality for some $\xi(x)$ for any $x \in \mathbf{R}$. This is not the case for (11): having integrated the dual function before maximizing with respect to ξ, we have invalidated the convex duality result. Hence, when we use $I(a,b)$ instead of the normalization constant of a truncated Gaussian density, we are making an approximation error. In the following result, we give a bound on this error.

Lemma 7: Let $N(x)$ be the probability density function of a Gaussian random variable with mean m and variance s^2 then

$$\left|\int_a^b N(x)\,\mathrm{d}x - I(a,b)\right| \leq \frac{\sqrt{2}}{\sqrt{\pi s^2}}\,(b-a)\,.$$

where $I(a,b)$ is given by (12).

Proof: We have that $N(x) \leq 1/\sqrt{2\pi s^2}$; hence

$$\int_a^b N(x)\,\mathrm{d}x \leq \frac{1}{\sqrt{2\pi s^2}}\,(b-a)\,.$$

In addition,

$$\exp\left\{-\frac{(b-m)^3-(a-m)^3}{6s^2\,(b-a)}\right\}\leq 1.$$

From the triangle inequality, the result follows. ■

While the upper bound in Lemma 7 is very conservative, it highlights two important (and intuitive) features of our approximation: the error is small if the interval $b-a$ is small or if the variance s^2 is large.

V. Variationally-optimal approximate identification

Using the variational lower bound $I(a,b)$, we can approximate (2) to find a closed-form expression for the approximations of the maximum likelihood estimate (5) and the maximum-a-posteriori estimate (8). To this end we note that, using Lemma 6, we can define an approximation (lower bound) on the likelihood function (4) according to

$$\mathrm{p}(\{y_k\}_{k=1}^N\,|g)\geq\prod_{k=1}^N I_k(g) \tag{13}$$

where

$$I_k(g)=\frac{L_{y_k+1}-L_{y_k}}{\sqrt{2\pi\sigma^2}}\exp\left\{-\frac{\left(L_{y_k+1}-U_k^Tg\right)^3-\left(L_{y_k}-U_k^Tg\right)^3}{6\sigma^2(L_{y_k+1}-L_{y_k})}\right\};$$

that is, $I_k(g)=I(L_{y_k},L_{y_k+1})$ where I is given by (12) when $m=U_kg$ and $s^2=\sigma^2$.

Using (13), we can define the approximate maximum likelihood estimator

$$\hat{g}^{\mathrm{LS}}=\arg\max_g\prod_{k=1}^N I_k(g), \tag{14}$$

This results in the intuitive approach of using the midpoints of the active quantization intervals to replace the data in a least-squares projection. We have the following result.

Theorem 8: The approximate maximum-likelihood estimator (14) is given by

$$\hat{g}^{\mathrm{LS}}=\left(U^TU\right)^{-1}U^T\bar{z}, \tag{15}$$

where $\bar{z}$ is given by (6).

Proof: See Appendix C.

We can also use (13) to define an approximation of the posterior distribution in (8):

$$\hat{g}^{\mathrm{MAP}}=\arg\max\mathrm{p}(g)\prod_{k=1}^N I_k(g). \tag{16}$$

Interestingly, also this variational approximation admits a closed form expression given by an appropriate penalized least-squares projection of the midpoints of the active quantization intervals:

Corollary 9: The approximate maximum-a-posteriori estimator (16) is given by

$$\hat{g}^{\mathrm{MAP}}=\left(U^TU+\sigma^2K^{-1}\right)^{-1}U^T\bar{z}, \tag{17}$$

where K is the prior covariance matrix in (7) and $\bar{z}$ is given by (6).

Proof: Follows from the same reasoning as Theorem 8.

Remark 10: While the approximate estimator (14) is independent of the noise variance, the maximum-a-posteriori estimator (16) depends on the noise variance σ^2. An estimate of σ^2 may be found by first estimating the impulse response using (14) and then using the sample variance formula

$$\left(\hat{\sigma}^{\mathrm{LS}}\right)^2=\frac{1}{N}\sum_{k=1}^N\left(\bar{z}_k-U_k\hat{g}^{\mathrm{LS}}\right)^2-\frac{1}{12}\left(L_{y_k+1}-L_{y_k}\right)^2, \tag{18}$$

where the second term is a first-order approximation of the quantization noise. Note that, while it is possible to obtain an estimate of the noise variance maximizing (13), this estimate is very biased and performs worse than (18).

Remark 11: When the covariance matrix of the prior distribution (7) depends on some unknown hyperparameters ρ—that is $\mathrm{p}(g)=N\left(g|0,K(\rho)\right)$—these are usually estimated using either cross validation or maximizing the marginal likelihood function

$$\mathrm{p}\left(\{y_k\}_{k=1}^N\right)=\int\mathrm{p}\left(\{y_k\}_{k=1}^N\,|g\right)\mathrm{p}(g)\,\mathrm{d}g, \tag{19}$$

(see [36], [37]). In the case at hand, integrating out g in (19) requires numeric integration or sampling. One alternative is to disregard the quantizer and consider the marginal distribution of the latent variable w_k; then, we can use the midpoint of the active quantization interval to approximate the latent variable in the marginal-likelihood inspired criterion

$$\hat{\rho}=\arg\min_\rho\bar{z}^T\Sigma_\rho^{-1}\bar{z}+\log\det\Sigma_\rho, \tag{20}$$

where $\Sigma_\rho=UK_\rho U^T+\sigma^2I$ is the marginal covariance matrix of w_k. Note that, while (16) is variationally optimal in the sense of Lemma 6, the criterion (20) is not motivated by variational arguments (in Section VIII-F, we compare the proposed approximate approach with the approach based on maximization of the marginal likelihood function).

VI. Bias compensated estimate

In this section, we propose an analysis of the error that results from using the approximate maximum-likelihood estimator (14) when the output is quantized with a *uniform quantizer*—that is, when $L_{n+1}-L_n=L$ for all n. In particular, we derive an expression for the bias of the midpoint $\bar{z}$ as an approximation of the output of the system Ug. Using this expression, we propose a bias-compensation scheme that gives consistent estimates.

First, note that $\bar{z}_k$ in (6) is a random variable (which depends on the latent random variable w_k) that can be written as

$$\bar{z}_k=L\mathcal{Q}(w_k)+L/2. \tag{21}$$

Hence, we can define the *quantization error*

$$\delta_k=\bar{z}_k-w_k. \tag{22}$$

Then, we have the following result.

Lemma 12: Let w_k be a random variable with characteristic function $\varphi_k(\omega)$, and let δ_k be the quantization error (22); then

$$\mathbf{E}[\delta_k] = \sum_{n=1}^{\infty} \frac{2}{\omega_n} \mathrm{Im}\{\varphi_k(\omega_n)\},$$
$$\mathbf{E}\left[\delta_k^2\right] = \frac{L^2}{12} + \sum_{n=1}^{\infty} \frac{4}{\omega_n^2} \mathrm{Re}\{\varphi_k(\omega_n)\},$$
$$\mathbf{E}[w_k \delta_k] = \left(m + i\varphi'(0)\right) \frac{L}{2} - \sum_{n \neq 0} \frac{\varphi'(\omega_n)}{\omega_n}.$$

where $\omega_n = 2\pi n/L$ and where $\varphi_k'(\omega)$ is the first derivative of $\varphi_k(\omega)$, assuming that the series converge.

Proof: See Appendix D.

Using the result of Lemma 12, we can find the expression of the bias in the Gaussian-noise case as follows.

Corollary 13: Let w_k be a Gaussian random variable with mean m and covariance σ^2, and let δ_k be the quantization error (22); then

$$\mathbf{E}[\delta_k] = \sum_{n=1}^{\infty} \frac{2}{\omega_n} \exp\left\{-\frac{\omega_n^2 \sigma^2}{2}\right\} \sin(\omega_n m),$$
$$\mathbf{E}\left[\delta_k^2\right] = \frac{L^2}{12} + \sum_{n=1}^{\infty} \frac{4}{\omega_n^2} \exp\left\{-\frac{\omega_n^2 \sigma^2}{2}\right\} \cos(\omega_n m),$$
$$\mathbf{E}\left[w_k \delta_k\right] = 2 \sum_{n=1}^{\infty} \exp\left\{-\frac{\omega_n^2 \sigma^2}{2}\right\} \left[\frac{m}{\omega_n} \sin(\omega_n m) + \sigma^2 \cos(\omega_n m)\right].$$

Corollary 13 highlights the same properties as Lemma 7 in that the error we commit in considering the midpoint as a proxy for the latent signal becomes smaller when either the quantization interval L becomes small or the noise variance becomes large.

Thanks to Corollary 13, we can characterize the bias in the variationally-optimal estimate in the case of Gaussian noise. If the input signal u_k is Gaussian, Bussgang's theorem guarantees that a linear regression performed on the data after the nonlinear transformation (21) will give a scaling of the true system [39]. If the input is white, we have an explicit expression of the scaling. To this end, assume that u_k is Gaussian white noise with known variance λ^2. Then, we use the quantization error (22) to say that

$$\bar{z}_k = U_k^T g_0 + \varepsilon_k + \delta_k; \tag{23}$$

and we find the asymptotic bias of the variationally-optimal approximate maximum-likelihood estimator.

Theorem 14: Let $\bar{z}_k$ be the vector of the midpoints of the active quantization intervals of the output of a linear system with uniform quantization, with step L, in response to a white Gaussian input with variance λ^2; then, as $N \to \infty$, the variationally-optimal estimator (15) converges to

$$\hat{g}^{\text{LS}} \to g_0 + g_0 B(g_0)$$

where

$$B(g_0) = 2 \sum_{n=1}^{\infty} \exp\left\{-\frac{\omega_n^2}{2}\left(\lambda^2 \|g_0\|^2 + \sigma^2\right)\right\}.$$

Proof: See Appendix E.

If we consider the variationally-optimal maximum-a-posteriori estimator with a fixed kernel matrix K, then, we have a similar asymptotic result:

Corollary 15: Under the conditions of Theorem 14, the variationally-optimal estimator (17) converges to $\hat{g}^{\text{MAP}} \to g_0 + g_0 B(g_0)$.

Proof: See Appendix E.

Remark 16: The function $B(g_0)$ is related to the *Jacobi theta function*

$$\vartheta(z, \tau) = \sum_{n=-\infty}^{+\infty} \exp\{-i\pi\tau n^2\} \exp\{2\pi i z n\}. \tag{24}$$

which arises in problems in number theory, elliptic functions, and quantum field theory. For more details, see [40, Ch. 21] and [41].

Thanks to Theorem 14, we can define a bias-compensated estimate $\hat{g}^{\text{BC}}$ as the solution to

$$\hat{g}^{\text{BC}} = \hat{g}^{\text{LS}} - \hat{g}^{\text{BC}} B(\hat{g}^{\text{BC}}). \tag{25}$$

Note that—because $B(\cdot)$ is a scalar function—any $\hat{g}^{\text{BC}}$ that solves this equation is proportional to the least squares estimate $\hat{g}^{\text{LS}}$ and can be found by setting $\hat{g}^{\text{BC}} = \alpha \hat{g}^{\text{LS}}$ and solving with respect to α the following scalar equation

$$\alpha + \alpha B(\alpha \hat{g}^{\text{LS}}) - 1 = 0. \tag{26}$$

From Theorem 14, as $N \to \infty$, the solution of (25) converges to the solution $g^\star$ of the estimation equation

$$g^\star + g^\star B(g^\star) = g_0 + g_0 B(g_0).$$

Note that any $g^\star$ that solves this equation is proportional to the true impulse response g_0; in addition, $g^\star = g_0$ is a trivial solution. The following results shows that $g^\star = g_0$ is the only solution and, hence, that the estimator $\hat{g}^{\text{BC}}$ is consistent.

Theorem 17: Under the assumptions of Theorem 14, the bias compensated estimator defined in (25) is consistent—that is, $\hat{g}^{\text{BC}} \to g_0$ as $N \to \infty$. In addition

$$N \mathbf{E}\left[\left(\hat{g}^{\text{BC}} - g_0\right)\left(\hat{g}^{\text{BC}} - g_0\right)^T\right] \to \Sigma \tag{27}$$

where $\Sigma = \frac{\sigma^2}{\lambda^2} I + \frac{L^2}{12\lambda^2} I + V - B(g_0)^2 g_0 g_0^T$ and

$$V = 4 \sum_{n=1}^{\infty} \left(\frac{1}{\omega_n^2} + \sigma^2\right)\left(\frac{1}{\lambda^2} I - g_0 g_0^T \omega_n^2\right) \exp\left\{-\frac{\omega_n^2}{2}\left(\lambda^2 \|g_0\|^2 + \sigma^2\right)\right\}.$$

Proof: See Appendix F.

The asymptotic variance (27) can be interpreted as composed of a term due to the noise, a term that is equivalent to a uniform white noise, and an *aliasing* term that comes from the correlation between the quantization error δ and the system output w (a similar result is presented in [4]).

For the problem of linear regression from quantized measurements, Gustafsson and Karlsson provide an expression for the Cramér-Rao lower bound [4, Theorem 6] for finite number of data which, using our notation, has the following asymptotic expression:

$$\Sigma \geq J(g_0)^{-1} \tag{28}$$

where $J(g)$ is the asymptotic Fisher information matrix,

$$J(g) := \sum_{n=-\infty}^{+\infty} \mathbf{E}\left[\frac{\left(N\left(\frac{L(n+1)-U^T g}{\sigma}\right) - N\left(\frac{Ln-U^T g}{\sigma}\right)\right)^2}{\Phi\left(\frac{L(n+1)-U^T g}{\sigma}\right) - \Phi\left(\frac{Ln-U^T g}{\sigma}\right)} UU^T \right],$$

where U is an zero-mean Gaussian vector with covariance matrix λI, and where $N(\cdot)$ is the density function and $\Phi(\cdot)$ is the distribution function of an standard normal random variable. A direct comparison of $J(g)^{-1}$ and the covariance Σ is intractable; hence, we show in some simulations (see Section VIII-B) that the asymptotic covariance matrix coincides with the Cramér-Rao lower bound. While this could indicate efficiency of the proposed bias-compensated method, a formal proof of this result is still beyond our reach.

In this section, we have derived explicit expressions for the case when both the input signal and the measurement noise are white and Gaussian. However, the result of Lemma 12 can be applied for general noise distributions. In particular, it can be used to derive the asymptotic bias of the approximate estimator for any distribution of the noise and the input signal. In the general case, we have the following result.

Theorem 18: Consider a linear system such as (1) with a wide-sense stationary ergodic input signal and white measurement noise with characteristic function $\varphi(\omega)$. Suppose that the output is quantized with uniform quantization of step L and let $\bar{z}$ be the vector of the midpoints of the active quantization intervals. Then, as $N \to \infty$, the estimator (15) converges to

$$\hat{g}^{\text{LS}} \to g_0 + R^{-1} \sum_{n=1}^{\infty} \frac{2}{\omega_n} \big(\zeta_n(g_0) \text{Re}\varphi(\omega_n) + \eta_n(g_0) \text{Im}\varphi(\omega_n) \big),$$

where $R = \mathbf{E}\big[UU^T\big]$, $\zeta_n(g_0) = \mathbf{E}\big[U \sin(\omega_n U^T g_0)\big]$, and $\eta_n(g_0) = \mathbf{E}\big[U \cos(\omega_n U^T g_0)\big]$, where $[U]_i = u_{k-i}$ for any k is an $n \times 1$ vector of samples of the input signal.

Proof: See Appendix H.

Theorem 18 can be used to define a bias-compensated estimate for general input and noise distributions. However, the statistical properties of this estimator depend strongly on distribution of the noise and of the input, and asymptotic unbiasedness or consistency may be difficult to verify.

VII. Iterative Algorithm

While it is possible to solve (25) directly using numerical methods, we can alternatively see (25) as an estimation equation with the following iterative version

$$\hat{g}^{(k+1)} = \hat{g}^{\text{LS}} - \hat{g}^{(k)} B(\hat{g}^{(k)}). \tag{29}$$

Under an appropriate *signal to quantization-noise variance* condition, the estimator $\hat{g}^{(k)}$ converges to $\hat{g}^{\text{BC}}$ and is, hence, consistent:

Proposition 19: Consider the assumptions of Theorem 14 and suppose, furthermore, that

$$\kappa = \frac{L^2 + 4\pi^2\sigma^2}{2\pi^2\sigma^2} \exp\left\{ -2\frac{\pi^2\sigma^2}{L^2} \right\} \leq 1. \tag{30}$$

Let $\hat{g}^{(k)}$ be defined by (29), then $\hat{g}^{(k)} \to \hat{g}^{\text{BC}}$ as $k \to \infty$, where $\hat{g}^{\text{BC}}$ solves (25). In addition,

$$\left\| \hat{g}^{\text{BC}} - \hat{g}^{(k+1)} \right\|_2 \leq \kappa \left\| \hat{g}^{\text{BC}} - \hat{g}^{(k)} \right\|_2. \tag{31}$$

Proof: See Appendix G.

Note that, the condition (30) is *sufficient* to ensure that (29) converges to $\hat{g}$. However, the iterative method (29) may converge even if $\kappa > 1$ and, in any case, the estimate (25) is consistent irrespective of the value of κ.

Remark 20: The condition (30) is very conservative; it allows for exponential convergence of the iteration (29) from any initial condition to the unique solution of (25) independently of the magnitude of the input signal and of the gain of the system. In the cases when (30) is not verified, the iteration (29) is guaranteed to converge to g_0 as $k \to \infty$ if

$$\exp\left\{ -2\frac{\pi^2\left(\lambda^2 \|g_0\|^2 + \sigma^2\right)}{L^2} \right\} \leq \frac{2\pi^2\left(\lambda^2 \|g_0\|^2 + \sigma^2\right)}{L^2 + 4\pi^2\left(\lambda^2 \|g_0\|^2 + \sigma^2\right)}.$$

Note that, while the condition (30) can be tested without the knowledge of the true system, this less conservative condition cannot be verified in practice. (See Section VIII-D for examples of the iterative solution and its convergence guarantees.)

VIII. Simulations

In this section, we present simulation experiments to validate the results presented in this paper. In the simulations, we consider FIR systems like (1) with infinite uniform quantizers and we compare the following estimators:

- $\hat{g}^{\text{ML}}$ The maximum-likelihood estimate found by solving (5). The optimization is done with local gradient-free optimization (Rowan's Subplex algorithm, [42]) stopped with a relative tolerance of 10^{-9} and initialized at the true parameters.
- $\hat{g}^{\text{LS}}$ The variationally-optimal approximation presented in Theorem 8. It estimates the system from a least-squares projection of the midpoint of the active quantization intervals.
- $\hat{g}^{\text{BC}}$ The bias-compensated estimate defined in (25). The solution is $\hat{g}^{\text{BC}} = \alpha \hat{g}^{\text{LS}}$ where α is computed solving (26) with a trust-region method with absolute tolerance of 10^{-9} started at $\alpha = 1$.
- $\hat{g}^{\text{OR}}$ The oracle estimator that uses the latent signal w to estimate the impulse response: $\hat{g}^{\text{OR}} = \left(U^T U\right)^{-1} U^T w$.

In the simulations, we compare the estimators with respect to the standard mean square error metric

$$\text{MSE}(g, g_0) = \|g - g_0\|^2.$$

A. Asymptotic performance

In the first simulation, we verify the asymptotic bias of the variationally-optimal estimate $\hat{g}^{\text{LS}}$ and the consistency of the bias compensated estimate $\hat{g}^{\text{BC}}$.

We consider the linear system

$$G_1(q) = 1 + 0.7q^{-1},$$

subject to a Gaussian noise input with variance $\lambda^2 = 1$. The output is subject to Gaussian white measurement noise with variance $\sigma^2 = 0.5$ and then collected with a quantizer with quantization step $L = 2.3$. Note that, the sufficient condition in Theorem 17 is verified with $\kappa \approx 0.3925$.

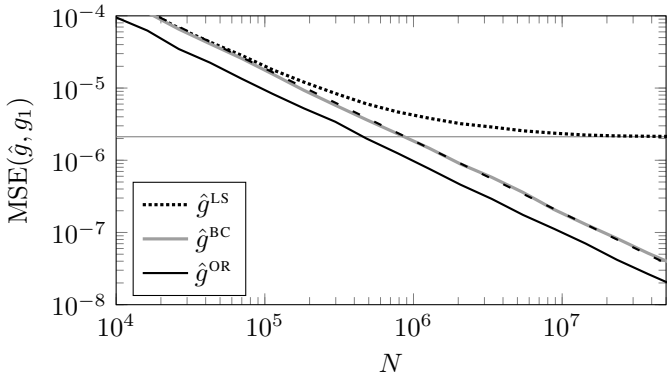

Fig. 4. Mean squared error of the estimates over 500 realizations of input signal and noise for different data sizes N. The approximate estimator (dotted line) converges to the theoretically predicted bias (horizontal line). The bias compensated estimate (gray line) follows the oracle estimate (black line) asymptotically. Note that the simulated error of $\hat{g}^{\mathrm{BC}}$ agrees with the error predicted by Theorem 17 (dashed line).

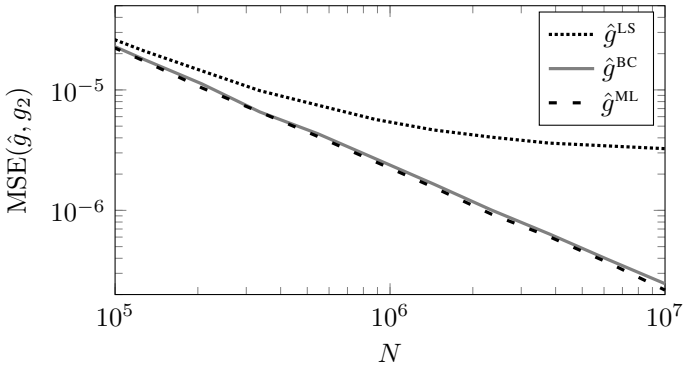

Fig. 5. Mean square error of the estimates over 500 realizations of input signal and noise for different data sizes N. The approximate estimator (dotted) converges to its asymptotic performance of $B(g_2)^2 \|g_2\|^2 \approx 3{\cdot}10^{-6}$ while the bias compensated estimate (gray) follows the maximum likelihood estimate (dashdotted).

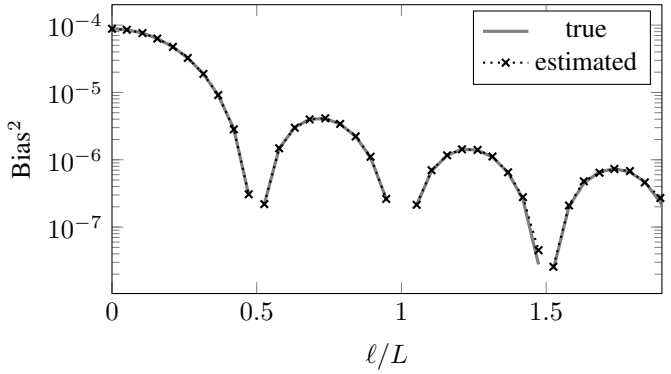

Fig. 6. Comparison of the asymptotic bias theoretically predicted by Theorem 18 and estimated by simulating g_1 for $N = 10^9$ samples and computing the bias of $\hat{g}^{\mathrm{ML}}$. Note that we have not simulated the system for multiples of $1/2$ because the estimated bias would be dominated by the variance error.

In Figure 4, we report the mean squared error of the estimates computed by the different methods for 500 different realizations of the input signal and the noise. As predicted by Theorem 14, the average error of the estimate $\hat{g}^{\mathrm{LS}}$ (dotted) converges to the value $B(g_1)^2\|g_1\|^2 \approx 2.116 \cdot 10^{-6}$ while the bias compensated estimate $\hat{g}^{\mathrm{BC}}$ (gray) follows the oracle estimator (black) which converges to g_1 asymptotically. In addition, we plot the asymptotic mean square error predicted by Theorem 17 (dashed) and we see that the empirical error of $\hat{g}^{\mathrm{BC}}$ agrees with the asymptotic expression of the mean square error—that is, $N\mathbf{E}\|\hat{g}^{\mathrm{BC}} - g_1\|^2 \to \mathrm{trace}\,\{\Sigma\} \approx 1.868$.

B. Comparison with Maximum-likelihood

In the second simulation, we compare the performance of the bias compensated estimate $\hat{g}^{\mathrm{BC}}$ with the performance of the maximum likelihood criterion $\hat{g}^{\mathrm{ML}}$. We consider the system,

$$G_2(q) = 1 + 0.3q^{-1} + 0.7q^{-2},$$

subject to a Gaussian white noise input with variance $\lambda^2 = 2.1$. The output is subject to a Gaussian white measurement noise with variance $\sigma^2 = 0.7$ and is collected after quantization with a step $L = 3.3$.

In Figure 5, we report the mean squared error of the estimates computed by the different methods for 500 different realizations of the input signal and the noise. From the simulation, It appears that the performance of the proposed bias-compensated estimate $\hat{g}^{\mathrm{BC}}$ (gray) compares very well with the maximum-likelihood estimator $\hat{g}^{\mathrm{ML}}$ (dashdotted). Also, it appears that, asymptotically, the variance of the bias-compensated estimator coincides with the variance of the maximum-likelihood estimator and reaches the Cramér-Rao lower bound as computed by (28)—that is, $N\mathbf{E}\|\hat{g}^{\mathrm{BC}} - g_2\|^2 \to \mathrm{trace}\,\{\Sigma\} \approx \mathrm{trace}\,\{J(g_2)^{-1}\} \approx 2.4717$.

This suggests that the proposed bias-compensated estimate is efficient (however, a formal analysis is still in preparation).

C. Asymptotic performance for uniform noise

In this section, we present a simulation experiment to validate the results of Theorem 18. In particular, we consider the asymptotic bias of the approximate estimator $\hat{g}^{\mathrm{LS}}$ when the measurement noise is white and uniformly distributed in $[-\ell, \ell]$. To this end, we consider the linear system $G_1(q)$, presented in Section VIII-A, subjected to a white and Gaussian input signal with variance $\lambda^2 = 1.1$ and we consider a uniform quantizer with step $L = 3.3$.

From Theorem 18, we have that $\hat{g}^{\mathrm{LS}}$ converges asymptotically to

$$\hat{g}^{\mathrm{LS}} \to g_1 + 2g_1 \sum_{n=1}^{\infty} \exp\left\{-\frac{\omega_n^2\lambda^2\|g_1\|^2}{2}\right\} \frac{\sin(\ell\omega_n)}{\ell\omega_n^2}.$$

Note, in particular, that if $\ell = mL/2$, for some positive integer m, then the method is asymptotically unbiased.

We compare the theoretically predicted bias with an empirical estimate of the bias obtained by estimating g_1 with $\hat{g}^{\mathrm{LS}}$ using $N = 10^9$ data points. In Figure 6, we present the average bias over 20 realizations of the input signal and the noise and we compare it with the bias predicted by Theorem 18. From the figure, we see that the theorem accurately predicts the asymptotic performance of the approximate method.

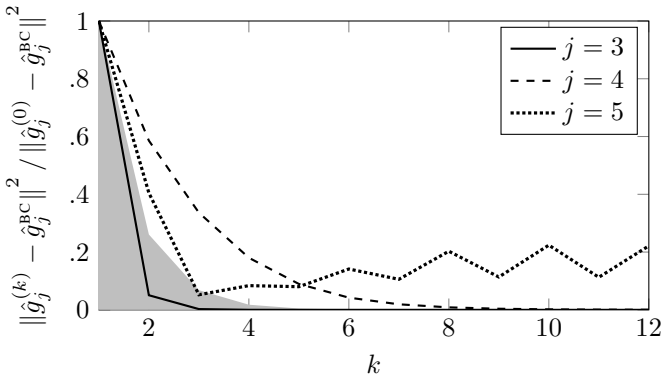

Fig. 7. Normalized distance between the iterative solution $\hat{\boldsymbol{g}}_j^{(k)}$ and the bias compensated estimate $\hat{\boldsymbol{g}}_j^{\mathrm{BC}}$. The first system ($j = 3$, solid) satisfies the hypothesis of Proposition 19, so the convergence is exponential with a rate guaranteed by (31) (shaded area). The second system ($j = 4$, dashed) satisfies the condition Remark 20 and converges even though $\kappa_2 > 1$. The third system ($j = 5$, dotted) does not have convergent iterations.

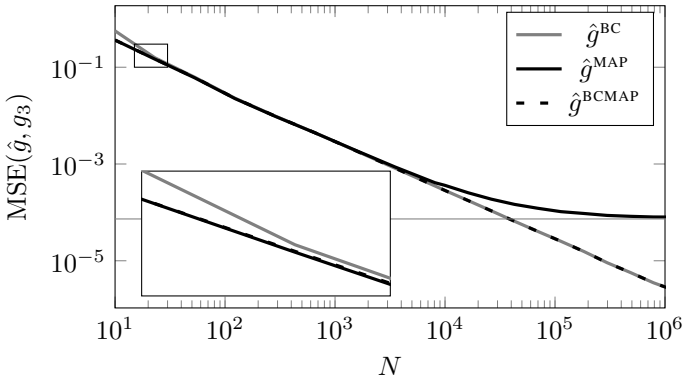

Fig. 8. Mean squared error of the estimates over 500 realizations of input signal and noise for different data sizes $\boldsymbol{N}$. The approximate Bayesian estimator (solid) converges to its asymptotic performance of $\boldsymbol{B(g_3)^2\,\|g_3\|^2 \approx 7.3 \cdot 10^{-5}}$ while the bias compensated estimates are asymptotically equivalent (dashed and gray). Note that, for short data lengths, the Bayesian estimates have lower mean-squared-error than the approximate maximum likelihood estimator.

D. Iterative solution

In this section, we present a simulation example where we validate the results of Proposition 19 and Remark 20. We consider the three systems,

$$g_3 = [0.4,\ 0.3,\ -0.7]^T \quad \lambda_3^2 = 1.1 \quad \sigma_3^2 = 0.6 \quad L_3 = 2.3,$$
$$g_4 = [0.4,\ 0.3,\ -0.4]^T \quad \lambda_4^2 = 1.0 \quad \sigma_4^2 = 0.2 \quad L_4 = 3.3,$$
$$g_5 = [0.4,\ 0.01]^T \quad \lambda_5^2 = 1.0 \quad \sigma_5^2 = 0.2 \quad L_5 = 6.3.$$

Only the first system satisfies the sufficient condition in Proposition 19 (with $\kappa_3 \approx 0.261$); the second system satisfies the extended condition in Remark 20; the third system does not satisfy any condition for convergence of the iterative estimator (29).

In Figure 7, we present the normalized distance of the iterative estimator to the fixed point $\hat{g}_j^{\mathrm{BC}}$ for $j = 3, 4, 5$, together with the theoretical bound given by (31). We see that the conditions in Proposition 19 and Remark 20 guarantee the convergence of the iteration to $\hat{g}^{\mathrm{BC}}$; in addition, we see that when the conditions are not fulfilled the iterative algorithm may diverge from the fixed point (all iterations start at $\hat{g}^{\mathrm{LS}}$).

E. Bias compensation and Bayesian estimation

In this section, we present a simulation where we compare the approximate maximum-likelihood and maximum-a-posteriori estimates to the bias-compensated estimates. We consider the system g_3 presented in the previous section and, in addition to the previously introduced $\hat{g}^{\mathrm{BC}}$, we cosider the following Bayesian estimators:

- $\hat{g}^{\mathrm{MAP}}$ the approximate maximum-a-posteriori estimator presented in Corollary 9. The parameters are modeled using a diagonal kernel, $K(\rho) = \rho I$. The hyperparameter ρ was estimated from the marginal-likelihood inspired criterion (20); and the noise variance is assumed to be known.
- $\hat{g}^{\mathrm{BCMAP}}$ A bias-compensated estimate based on Corollary 15. We define a bias-compensated maximum-a-posteriori estimate as $\hat{g}^{\mathrm{BCMAP}} = \alpha \hat{g}^{\mathrm{MAP}}$ where α is found solving $\alpha + \alpha B(\alpha \hat{g}^{\mathrm{MAP}}) - 1 = 0$, using a trust-region method with absolute tolerance of 10^{-9} started at $\alpha = 1$.

In Figure 8, we report the mean squared error of the estimates computed by the different methods for 500 different realizations of the input signal and the noise. As predicted by Corollary 15, the average error of the estimate $\hat{g}^{\mathrm{MAP}}$ (solid black) converges to the same limit as $\hat{g}^{\mathrm{ML}}$—that is, the value $B(g_3)^2\|g_3\|^2 \approx 7.319 \cdot 10^{-5}$. The bias compensated maximum-a-posteriori estimate $\hat{g}^{\mathrm{BCMAP}}$ (dashed black) follows asymptotically the bias compensated estimate $\hat{g}^{\mathrm{BC}}$ (solid gray) whose asymptotic covariance converges to the one predicted by the Cramér-Rao lower bound (28):

$$N\|\hat{g}^{\mathrm{BC}} - g_3\|^2 \approx\!\!\rightarrow \mathrm{trace}\,\{\Sigma\} \approx J(g_3)^{-1} \approx 2.8402.$$

Note that, for shorter data records, the bias introduced by the kernel improves the performance of the approximate Bayesian estimators compared to the approximate maximum-likelihood estimator.

F. Bayesian estimation of high-order FIR models

In this section, we present some simulations regarding the variationally-optimal Bayesian estimator (8). We consider the impulse responses of the first 500 systems in the S1D1 dataset in [36]; for simplicity, we truncate all the impulse responses after the first 70 samples. We subject the systems to a Gaussian white-noise input with unitary variance and we collect the output with Gaussian additive white noise with variance $1/10$th of the variance of the corresponding noiseless output. In the simulations, we consider 4 different uniform quantizers with steps $L^2 = 12\sigma^2,\ 48\sigma^2,\ 120\sigma^2$ —which, in a first-order approximation of (23), correspond to a relative importance of the quantization over the measurement noise of 1, 4, and 10. In addition, we consider the case with no quantization (that is $L \to 0$) and the ceiling-type quantizer (that is $L = 1$), considered in [29].

In the simulation, we compare the estimators

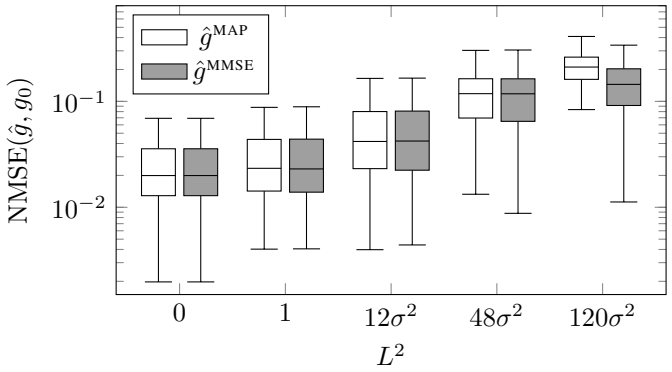

Fig. 9. Boxplots of the normalized mean square errors of the 500 estimated impulse responses, for different levels of quantization.

$\hat{g}^{\text{MAP}}$ The approximate maximum-a-posteriori estimator presented in Corollary 9. The impulse responses are modeled with the *stable-spline kernel* [37],

$$\big[K(\rho)\big]_{i,j} = \rho_1\, \rho_2^{\max(i,j)}. \tag{32}$$

The hyperparameters and the noise variance are estimated from the marginal-likelihood inspired criterion (20) using expectation maximization. The hyperparameters were initialized arbitrarily as $\rho_1 = 1$, $\rho_2 = 0.6$; the noise variance σ^2 was initialized according to (18).

$\hat{g}^{\text{MMSE}}$ The Bayesian minimum mean-square-error estimator presented in [29]. The impulse response is modeled with the stable-spline kernel (32). The hyperparameters are estimated from the marginal-likelihood function using Monte-Carlo expectation-maximization with Gibbs sampling. The iterations were initialized at the hyperparameter and noise variance values found by $\hat{g}^{\text{MAP}}$. At each iteration 400 samples were used; during the first iteration, 100 samples were discarded for burnin. The posterior means were computed with Monte-Carlo using 1000 samples after 100 samples of burnin.

Remark 21: Note that, as $L \to 0$, the compared methods coincide with the standard Bayesian kernel-based approach [36], [37] of estimating the hyperparameters from the marginal-likelihood function of the data and estimating the impulse response from the mean of the posterior density (which is, when $L \to 0$, Gaussian).

The results of the simulation are presented in Figure 9. In the figure, we show the boxplots of the normalized mean square error, $\text{NMSE}(g, g_0) = \text{MSE}(g, g_0)/\|g_0\|^2$, where g is an estimate of the true impulse response g_0.

From the simulation, we see that the approximate method $\hat{g}^{\text{MAP}}$ has a fit that is comparable to the performance of $\hat{g}^{\text{MMSE}}$ for a wide range of quantization levels with a decrease in performance for coarser quantizations.

IX. Conclusions

In this paper, we have considered the identification of finite impulse-response models from quantized data. First, we have proposed a variational approximation of the likelihood function which leads to closed-form approximations for the maximum likelihood and maximum-a-posteriori estimates. Then, we have analyzed the quantization error for different noise distributions and we have derived an expression of the asymptotic bias of the approximate maximum likelihood estimator for the case of Gaussian measurement noise and Gaussian input signal. Compensating for the bias, we have found a consistent estimator. We have validated the results in different simulation examples.

In the analysis, we have derived the asymptotic covariance of the consistent estimator. While numeric simulations seem to indicate that the asymptotic covariance is matches the covariance of the maximum-likelihood estimator (which would indicate that the proposed method is efficient), a formal proof of this result is complicated by the lack of a closed-form expression for the Cramér-Rao lower bound in the case of white Gaussian inputs. Research in this direction is ongoing and we hope to be able to present our results in future publications.

X. Acknowledgment

The authors thank Mikhail Skopenkov at the Institute for Information Transmission Problems of the Russian Academy of Sciences and Michael Somos at Department of Mathematics and Statistics of Georgetown University for suggesting the use of Jacobi's identities.

Appendix

A. Proof of Proposition 3

The first-order condition for the maximum-likelihood criterion (5), is given by

$$\sum_{k=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} \left(\mathrm{e}^{-\frac{b_k^2}{2\sigma^2}} - \mathrm{e}^{-\frac{a_k^2}{2\sigma^2}} \right) \prod_{j \neq k} \Big\{ \Phi(b_k) - \Phi(a_k) \Big\} \frac{U_k^T}{\sigma} = 0,$$

where $b_k = L_{y_k+1} - U_k^T g$ and $a_k = L_{y_k} - U_k^T g$.

Let now $\Delta_k = (L_{y_k+1} - L_{y_k})/2$. If $g = \hat{g}$, then $b_k = L_{y_k+1} - \bar{z}_k = \Delta_k$ and $a_k = L_{y_k} - \bar{z}_k = -\Delta_k$ and the first-order condition is verified.

Consider now the second-order condition in $g = \hat{g}$, we get that the Hessian is given by

$$H = -\sum_{k=1}^{N} \frac{\Delta_k}{\sqrt{2\pi\sigma^2}} \mathrm{e}^{-\frac{\Delta_k^2}{2\sigma^2}} \prod_{j \neq k} \Big\{ \Phi(\Delta_k) - \Phi(-\Delta_k) \Big\} \frac{U_k U_k^T}{\sigma^2};$$

because the terms in the sum are all positive, the Hessian is negative definite, and $g = \hat{g}$ is a local solution of (5). ■

B. Proof of Lemma 6

From Lemma 5, we have that

$$\int_a^b N(x)\,\mathrm{d}x \geq -\frac{N(\xi)}{2s^2} \int_a^b \Big[(x-m)^2 - (\xi-m)^2 - 2s^2 \Big]\,\mathrm{d}x$$
$$= -\frac{N(\xi)}{2s^2} \left\{ \frac{(b-m)^3 - (a-m)^3}{3} - (\xi-m)^2\,(b-a) - 2s^2\,(b-a) \right\}.$$

This expressions holds for all $\xi \in \mathbf{R}$. Let $\xi^\star$ be the maximizer of the right-hand site of the inequality. From the first order-optimality condition, we find that

$$\frac{N'(\xi^\star)}{2s^2}\left\{\frac{(b-m)^3-(a-m)^3}{3}-(\xi^*-m)^2(b-a)-2s^2(b-a)\right\} - 2\frac{N(\xi^\star)}{2s^2}(\xi-m)(b-a)=0.$$

where $N'(\xi)$ is the first derivative of $N(\xi)$. Note that $N'(\xi) = -N(\xi)(\xi-m)/s^2$; hence, the optimality condition becomes

$$(\xi^*-m)^2=\frac{(b-m)^3-(a-m)^3}{3(b-a)}.$$

Plugging this back into the inequality, we find

$$\int_a^b N(x)\,\mathrm{d}x \geq N(\xi^*)(b-a),$$

from which, we have the result. ■

C. Proof of Theorem 8

Consider the logarithm of the cost function in (14); from the first-order condition

$$\sum_{k=1}^N \frac{\left(L_{y_k+1}^2-L_{y_k}^2\right)-2\left(L_{y_k+1}-L_{y_k}\right)U_k^T\hat{g}^{\mathrm{LS}}}{2\sigma^2\left(L_{y_k+1}-L_{y_k}\right)}U_k^T=0,$$

$$\sum_{k=1}^N \frac{L_{y_k+1}+L_{y_k}}{2}U_k^T=\sum_{k=1}^N\left(U_k^T\hat{g}^{\mathrm{LS}}\right)U_k^T,$$

from which, $\bar{z}^TU = g^{\mathrm{LS}\,T}U^TU$. ■

D. Proof of Lemma 12

Assuming that all series converge, we use the Poisson summation formula (see [43], Appendix 10A), to find that

$$\mathbf{E}\{\delta\}-L/2=\int_{-\infty}^{+\infty}\left(L\mathcal{Q}(w)-w\right)\mathrm{p}(w)\,\mathrm{d}w$$

$$=-\sum_{n=-\infty}^{\infty}\int_0^L x\,\mathrm{p}(x+nL)\,\mathrm{d}x=-\sum_{n=-\infty}^{\infty}\varphi(\omega_n)\int_0^L\frac{x}{L}\mathrm{e}^{i\omega_n x}\,\mathrm{d}x,$$

where $\omega_n = 2\pi n/L$. Now for $n \neq 0$, we have that $\int_0^L x\mathrm{e}^{i\omega_n x}/L\,\mathrm{d}x = i/\omega_n$. So

$$\mathbf{E}\{\delta\}=\sum_{n\neq0}-\frac{i}{\omega_n}\varphi(\omega_n)=\sum_{m=1}^{\infty}\frac{2}{\omega_m}\mathrm{Im}\varphi(\omega_m);$$

from which, we have the result.

Regarding the second moment, using the same arguments we used in the first step,

$$\mathbf{E}\left\{\delta^2\right\}=\int_{-\infty}^{+\infty}\left(w-L\mathcal{Q}(w)-L/2\right)^2\mathrm{p}(w)\,\mathrm{d}w$$

$$=\sum_{n=-\infty}^{\infty}\int_0^L\frac{\left(x-L/2\right)^2}{L}\varphi(\omega_n)\,\mathrm{e}^{i\omega_n x}\,\mathrm{d}x=\frac{L^2}{12}+\sum_{n\neq0}\frac{2}{\omega_n^2}\varphi(\omega_n),$$

For the covariance, with similar arguments, we have

$$\mathbf{E}\{w\delta\}=\int_{-\infty}^{+\infty}w\left(L\mathcal{Q}(w)+L/2-w\right)\mathrm{p}(w)\,\mathrm{d}w$$

$$=-\int_0^L\sum_{n=-\infty}^{\infty}(x+nL)x\mathrm{p}(x+nL)\,\mathrm{d}x+\frac{L}{2}m$$

Then, from the Poisson summation formula applied to $nL\mathrm{p}(x+nL)$, we have that $\sum_{n=-\infty}^{+\infty}nL\mathrm{p}(x+nL) = \sum_{n=-\infty}^{+\infty}-\frac{i}{L}\varphi(\omega_n) - \frac{x}{L}\varphi'(\omega_n)$ So $\mathbf{E}\{w\delta\} = \sum_{n=-\infty}^{\infty}\frac{i}{L}\varphi'(\omega_n)\int_0^L x\mathrm{e}^{i\omega_n x}\,\mathrm{d}x + \frac{L}{2}m$ from which the result follows. ■

E. Proof of Theorem 14 and Corollary 15

We consider the approximate estimator $\hat{g}^{\mathrm{LS}}$. Then, by (23),

$$\hat{g}^{\mathrm{LS}}-g_0=\left[\frac{1}{N}\sum_{k=1}^N U_kU_k^T\right]^{-1}\sum_{k=1}^N\frac{U_k(\varepsilon_k+\delta_k)}{N}.$$

Because u_k is a Gaussian white noise sequence with variance λ^2, independent (by assumption) of ε_k, when $N \to \infty$,

$$\frac{1}{N}\sum_{k=1}^N U_kU_k^T\to\lambda^2I,\qquad\frac{1}{N}\sum_{k=1}^N U_k\varepsilon_k\to0.$$

Similarly,

$$\left[\frac{1}{N}\sum_{k=1}^N U_k^T\delta\right]_i\to\mathbf{E}[u_{k-i}\delta_k]=\mathbf{E}\left[u_{k-i}\mathbf{E}\left[\delta_k\middle|\{u_t\}_{t=1}^N\right]\right]$$

$$=\sum_{n=1}^{\infty}\frac{L}{\pi n}\exp\left\{-\frac{2\sigma^2\pi^2n^2}{L^2}\right\}\mathbf{E}\left[u_{k-i}\sin\left(\frac{2\pi n}{L}U_k^Tg_0\right)\right]$$

from which, we have the result.

Consider now the approximate estimator $\hat{g}^{\mathrm{MAP}}$. Because u_k is Gaussian white noise we have that $\left(U^TU+\sigma^2K^{-1}\right)/N \to \lambda I$, and, asymptotically, $\hat{g}^{\mathrm{MAP}} \approx \hat{g}^{\mathrm{ML}}$. ■

F. Proof of Theorem 17

First, consider the case $2\pi\lambda^2\|\hat{g}^{\mathrm{LS}}\|^2 = L^2$ and $\sigma = 0$. Then, any solution to (25) is given by $\hat{g} = \alpha\hat{g}^{\mathrm{LS}}$ where α solves

$$\alpha f(\alpha)=\alpha\left(1+2\sum_{n=1}^{\infty}\exp\{-\pi\alpha^2n^2\}\right)=1\qquad(33)$$

Note that $f(\alpha) > 0$; hence, from (33), we have that $\alpha > 0$ and that $f(\alpha)$ is a monotone decreasing function of α. Note that $f(\alpha) = \vartheta(0; i\alpha^2)$ where ϑ is the Jacobi theta function (24). By Jacobi's identities (also known as Jacobi's formulae, see [40], Section 21.51),

$$\vartheta\left(\frac{z}{\tau};-\frac{1}{\tau}\right)=\sqrt{-i\tau}\exp\left\{\frac{\pi}{\tau}iz^2\right\}\vartheta(z;\tau).$$

Hence, in our case, $\alpha f(\alpha) = f(1/\alpha)$. Because $f(\alpha)$ is monotone decreasing, $f(1/\alpha)$ is monotone increasing in α. Hence, $\alpha f(\alpha)$ is monotone increasing and thus (33) has only one solution and (25) has only one solution.

Let now $N \to \infty$, then $\hat{g}^{\mathrm{LS}} \to g_0 + g_0B(g_0)$ by Theorem 14 and, by continuity, $\hat{g} \to \alpha^\star g_0$ where α is solution to

$\alpha^\star f(\alpha^\star) = f(1.)$. By the same argument as before, this equation has a unique solution given by $\alpha^\star = 1$; hence $\hat{g} \to g_0$ as $N \to \infty$ and we have the proof.

In the general case, note that $\hat{g} = \alpha\hat{g}^{\text{LS}}$ because $B(\cdot)$ is a scalar function; note also that

$$B(\alpha\hat{g}^{\text{LS}}) = f\left(\pi\sqrt{\alpha^2\lambda^2\|\hat{g}^{\text{LS}}\|^2 + \sigma^2}\right);$$

Applying the same reasoning to $B(\cdot)$ gives the result.

Regarding the asymptotic covariance matrix of the estimates, using the results of Theorem 14, we have that asymptotically

$$\begin{aligned}&(\hat{g}-g_0)(\hat{g}-g_0)^T \to (\hat{g}^{\text{LS}}-g_0)(\hat{g}^{\text{LS}}-g_0)^T - B(g_0)^2 g_0 g_0^T \\ &\to \frac{1}{N\lambda^4}\left[\mathbf{E}\left[U_k U_k^T \varepsilon_k^2\right] + \mathbf{E}\left[U_k U_k^T \delta_k^2\right] + 2\mathbf{E}\left[U_k U_k^T \varepsilon_k \delta_k\right]\right]\end{aligned}$$

Using the results of Corollary 13, we have that

$$\begin{aligned}&\left[\mathbf{E}\left[U_k U_k^T \delta_k^2\right]\right]_{i,j} = \mathbf{E}\left[u_{k-i}u_{k-j}\mathbf{E}\left[\delta_k^2 \middle| \{u_t\}_{t=1}^N\right]\right] \\ &= \frac{L^2\lambda^2}{12} + \sum_{n=1}^{+\infty}\frac{4}{\omega_n^2}\exp\left\{-\frac{\omega_n^2\sigma^2}{2}\right\}\mathbf{E}\left[u_{k-i}u_{k-j}\cos\left(\omega_n U_k^T g_0\right)\right],\end{aligned}$$

and that

$$\begin{aligned}&\left[\mathbf{E}\left[u_{k-i}u_{k-j}\varepsilon_k\delta_k\right]_{i,j}\right] = \mathbf{E}\left[u_{k-i}u_{k-j}\left(w - U_k^T g_0\right)\delta_k\right]_{i,j} \\ &= 2\sigma^2\sum_{n=1}^{+\infty}\exp\left\{-\frac{\omega_n^2\sigma^2}{2}\right\}\mathbf{E}\left[u_{k-i}u_{k-j}\cos\left(\omega_n U_k^T g_0\right)\right];\end{aligned}$$

from which, we have the result. ■

G. Proof of Proposition 19

Let $F(g) = \hat{g}^{\text{LS}} - gB(g)$; then $\hat{g}$ is a fixed point of the mapping $F(\cdot)$. Note that, for any g_1, g_2,

$$\left\|F(g_1) - F(g_2)\right\|_2 = \left\|g_2 B(g_2) - g_1 B(g_1)\right\|_2$$

Then, note that

$$\begin{aligned}&\sum_{n=1}^{\infty}\exp\left\{-2\frac{\pi^2 n^2}{L^2}\left(\lambda^2\|g_i\|^2 + \sigma^2\right)\right\} \\ &\le \exp\left\{-2\frac{\pi^2}{L^2}\left(\lambda^2\|g_i\|^2+\sigma^2\right)\right\} + \int_1^{\infty}\exp\left\{-2\frac{\pi^2}{L^2}\left(\lambda^2\|g_i\|^2+\sigma^2\right)x^2\right\}\mathrm{d}x \\ &\le \exp\left\{-2\frac{\pi^2}{L^2}\left(\lambda^2\|g_i\|^2 + \sigma^2\right)\right\}\left(1 + \frac{L^2}{4\pi^2\left(\lambda^2\|g_i\|^2 + \sigma^2\right)}\right) \\ &\le \exp\left\{-2\frac{\pi^2\sigma^2}{L^2}\right\}\left(1 + \frac{L^2}{4\pi^2\sigma^2}\right) = \kappa,\end{aligned}$$

where κ is a constant independent of g. Then,

$$\left\|F(g_1) - F(g_2)\right\|_2 \le 2\kappa\|g_2 - g_1\|_2 .$$

Under condition (25), $2\kappa < 1$ so $F(\cdot)$ is a contractive mapping and thus has a unique fixed point $\hat{g}$ and the iteration $\hat{g}^{(k+1)} = F(\hat{g}^{(k)})$ is such that $\hat{g}^{(k)} \to \hat{g}$ if $k \to \infty$ with the indicated rate of convergence. ■

H. Proof of Theorem 18

We consider the approximate estimator $\hat{g}^{\text{LS}}$. Then, by (23),

$$\hat{g}^{\text{LS}} - g_0 = \left[\frac{1}{N}\sum_{k=1}^{N}U_k U_k^T\right]^{-1}\sum_{k=1}^{N}\frac{U_k\left(\varepsilon_k + \delta_k\right)}{N}.$$

Because u_k is wide-sense stationary, ergodic, and (by assumption) independent of ε_k, when $N \to \infty$,

$$\frac{1}{N}\sum_{k=1}^{N}U_k U_k^T \to R, \qquad \frac{1}{N}\sum_{k=1}^{N}U_k\varepsilon_k \to 0.$$

where $[R]_{i,j} = \mathbf{E}\left[u_i u_j\right]$. Similarly,

$$\begin{aligned}&\left[\frac{1}{N}\sum_{k=1}^{N}U_k^T\delta\right]_i \to \mathbf{E}[u_{k-i}\delta_k] = \mathbf{E}\left[u_{k-i}\mathbf{E}\left[\delta_k \middle| \{u_t\}_{t=1}^N\right]\right] \\ &= \sum_{n=1}^{\infty}\frac{2}{\omega_n}\mathbf{E}\left[u_{k-i}\mathrm{Im}\left\{\varphi(\omega_n)\exp\left\{i\omega_n U^T g_0\right\}\right\}\right].\end{aligned}$$

from which, we have the result. ■

REFERENCES

[1] R. Isermann, *Digital control systems*. Springer Science & Business Media, 2013.

[2] R. Min, M. Bhardwaj, S.-H. Cho, E. Shih, A. Sinha, A. Wang, and A. Chandrakasan, "Low-power wireless sensor networks," in *Int. Conf. VLSI Des.* IEEE, 2001, pp. 205–210.

[3] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, "Feedback control under data rate constraints: An overview," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 108–137, 2007.

[4] F. Gustafsson and R. Karlsson, "Statistical results for system identification based on quantized observations," *Automatica*, vol. 45, no. 12, pp. 2794–2801, 2009.

[5] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing*. Pearson, 1975.

[6] B. Widrow, I. Kollar, and M.-C. Liu, "Statistical theory of quantization," *IEEE Trans. Instrum. Meas.*, vol. 45, no. 2, pp. 353–361, 1996.

[7] B. Widrow and I. Kollár, *Quantization noise*, 2008.

[8] L. Y. Wang, G. G. Yin, J.-F. Zhang, and Y. Zhao, *System identification with quantized observations*. Springer, 2010.

[9] L. Y. Wang, J.-F. Zhang, and G. G. Yin, "System identification using binary sensors," *IEEE Trans. Autom. Control*, vol. 48, no. 11, pp. 1892–1907, 2003.

[10] L. Y. Wang, G. G. Yin, and J.-F. Zhang, "Joint identification of plant rational models and noise distribution functions using binary-valued observations," *Automatica*, vol. 42, no. 4, pp. 535–547, 2006.

[11] E. Colinet and J. Juillard, "A weighted least-squares approach to parameter estimation problems based on binary measurements," *IEEE Trans. Autom. Control*, vol. 55, no. 1, pp. 148–152, 2010.

[12] K. Jafari, J. Juillard, and M. Roger, "Convergence analysis of an online approach to parameter estimation problems based on binary observations," *Automatica*, vol. 48, no. 11, pp. 2837–2842, 2012.

[13] J. Guo and Y. Zhao, "Recursive projection algorithm on FIR system identification with binary-valued observations," *Automatica*, vol. 49, no. 11, pp. 3396–3401, 2013.

[14] M. Casini, A. Garulli, and A. Vicino, "Input design in worst-case system identification using binary sensors," *Automatic Control, IEEE Transactions on*, vol. 56, no. 5, pp. 1186–1191, 2011.

[15] ——, "Input design in worst-case system identification with quantized measurements," *Automatica*, vol. 48, no. 12, pp. 2997–3007, 2012.

[16] B. Godoy, P. Valenzuela, C. Rojas, J. Aguero, and B. Ninness, "A novel input design approach for systems with quantized output data," in *Control Conference (ECC), 2014 European*. IEEE, 2014, pp. 1049–1054.

[17] D. Marelli, K. You, and M. Fu, "Identification of ARMA models using intermittent and quantized output observations," *Automatica*, vol. 49, no. 2, pp. 360–369, 2013.

[18] J. Guo, L. Wang, G. Yin, Y. Zhao, and J.-F. Zhang, "Asymptotically efficient identification of FIR systems with quantized observations and general quantized inputs," *Automatica*, vol. 57, pp. 113–122, 2015.

[19] J. Wang and Q. Zhang, "Identification of FIR systems based on quantized output measurements: a quadratic programming-based method," *IEEE Transactions on Automatic Control*, vol. 60, no. 5, pp. 1439–1444, 2015.
[20] H. Suzuki and T. Sugie, "System identification based on quantized I/O data corrupted with noises and its performance improvement," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2006, pp. 3684–3689.
[21] A. Moschitta, J. Schoukens, and P. Carbone, "Parametric system identification using quantized data," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 8, pp. 2312–2322, 2015.
[22] Y. Zhao, J.-F. Zhang, L. Y. Wang, and G. G. Yin, "Identification of Hammerstein systems with quantized observations," *SIAM Journal on Control and Optimization*, vol. 48, no. 7, pp. 4352–4376, 2010.
[23] G. Li and C. Wen, "Identification of Wiener systems with clipped observations," *IEEE Transactions on Signal Processing*, vol. 60, no. 7, pp. 3845–3852, 2012.
[24] A. Goudjil, M. Pouliquen, E. Pigeon, O. Gehan, and M. M'Saad, "Identification of systems using binary sensors via support vector machines," in *Proc. IEEE Conf. Decis. Control*. IEEE, dec 2015.
[25] A. Chiuso, "A note on estimation using quantized data," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 12 486–12 491, 2008.
[26] B. I. Godoy, G. C. Goodwin, J. C. Agüero, D. Marelli, and T. Wigren, "On identification of FIR systems having quantized output data," *Automatica*, vol. 47, no. 9, pp. 1905–1915, 2011.
[27] B. I. Godoy, J. C. Agüero, R. Carvajal, G. C. Goodwin, and J. I. Yuz, "Identification of sparse FIR systems using a general quantisation scheme," *Int. J. Control*, vol. 87, no. 4, pp. 874–886, 2014.
[28] T. Chen, Y. Zhao, and L. Ljung, "Impulse response estimation with binary measurements: A regularized FIR model approach," *IFAC Proc. Vol.*, vol. 45, no. 16, pp. 113–118, 2012.
[29] G. Bottegal, H. Hjalmarsson, and G. Pillonetto, "A new kernel-based approach to system identification with quantized output data," *Automatica*, vol. 85, pp. 145–152, 2017.
[30] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. R. Stat. Soc. Ser. B (Methodol.)*, pp. 1–38, 1977.
[31] G. Pólya *et al.*, "Remarks on computing the probability integral in one and two dimensions," in *Proc. Berkeley Symp. Math. Stat. Prob.* University of California Press, 1949, pp. 63–78.
[32] R. G. Hart, "A formula for the approximation of definite integrals of the normal distribution function," *Mathematical Tables and Other Aids to Computation*, vol. 11, no. 60, p. 265, 1957.
[33] ——, "A close approximation related to the error function," *Math. Comput.*, vol. 20, no. 96, pp. 600–602, 1966.
[34] R. Yerukala and N. K. Boiroju, "Approximations to standard normal distribution function," *Int. J. Sci. Eng. Res.*, vol. 6, no. 4, pp. 515–518, 2015.
[35] R. S. Risuleo, G. Bottegal, and H. Hjalmarsson, "Approximate maximum-likelihood identification of linear systems from quantized measurements," in *Proc. IFAC Symp. System Identification (SYSID)*, 2018, (to appear).
[36] T. Chen, H. Ohlsson, and L. Ljung, "On the estimation of transfer functions, regularizations and Gaussian processes—Revisited," *Automatica*, vol. 48, no. 8, pp. 1525–1535, 2012.
[37] G. Pillonetto, F. Dinuzzo, T. Chen, G. De Nicolao, and L. Ljung, "Kernel methods in system identification, machine learning and function estimation: A survey," *Automatica*, vol. 50, no. 3, pp. 657–682, 2014.
[38] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
[39] M. Enqvist and L. Ljung, "Linear approximations of nonlinear FIR systems for separable input processes," *Automatica*, vol. 41, no. 3, pp. 459–473, 2005.
[40] E. T. Wittaker and G. N. Watson, *A course of modern analysis*. Cambridge University Press, 1927.
[41] W. P. Reinhardt and P. L. Walker, "Theta functions," in *NIST Handbook of Mathematical Functions*, F. W. J. Olver, D. M. Lozier, R. F. Boisvert, and C. W. Clark, Eds. Cambridge University Press, 2010.
[42] T. H. Rowan, "Functional stability analysis of numerical algorithms," Ph.D. dissertation, 1990.
[43] A. S. Papoulis and U. Pillai, *Probability, Random Variables, and Stochastic Processes*, 2nd ed. McGraw-Hill Education Ltd, 2002.

Riccardo Sven Risuleo received the M.S. degree summa cum laude in Automatic Control from "La Sapienza" university of Rome, Italy, in 2013. He received the Licentiate degree and the Ph.D. degree in Electrical Engineering in 2016 and 2018, respectively, from the KTH Royal Institute of Technology in Stockholm, Sweden. He is currently a postdoctoral researcher at KTH and Uppsala University where he carries out research on system identification and machine learning.

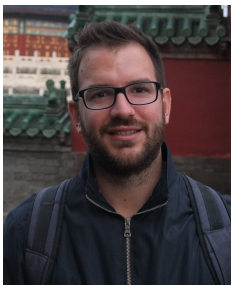

Giulio Bottegal received the M.S. degree in Automation Engineering (cum laude) in 2009 and the Ph.D. degree in Information Engineering in 2013 from the University of Padova. From 2013 to 2018 he held postdoctoral positions at the School of Electrical Engineering, KTH Royal Institute of Technology (Sweden), at ESAT, KU Leuven (Belgium), and at the Department of Electrical Engineering, TU Eindhoven (The Netherlands). He held visiting research positions at the Australian National University and at the University of Melbourne. Currently, he is a Senior Design Engineer at ASML B.V., working on the development of machine learning algorithms for semiconductor metrology.

Håkan Hjalmarsson was born in 1962. He received the M.S. degree in Electrical Engineering in 1988, and the Licentiate degree and the Ph.D. degree in Automatic Control in 1990 and 1993, respectively, all from Linköping University, Sweden. He has held visiting research positions at California Institute of Technology, Louvain University and at the University of Newcastle, Australia. He has served as an Associate Editor for Automatica (1996–2001), and IEEE Transactions on Automatic Control (2005–2007) and been Guest Editor for European Journal of Control and Control Engineering Practice. He is Professor at the School of Electrical Engineering, KTH, Stockholm, Sweden. He is an IEEE Fellow and past Chair of the IFAC Coordinating Committee CC1 Systems and Signals. In 2001 he received the KTH award for outstanding contribution to undergraduate education. His research interests include system identification, signal processing, control and estimation in communication networks and automated tuning of controllers.