

Regularized linear system identification using atomic, nuclear and kernel-based norms: The role of the stability constraint

Gianluigi Pillonetto, Tianshi Chen, Alessandro Chiuso, Giuseppe De Nicolao and Lennart Ljung

Linköping University Post Print



N.B.: When citing this work, cite the original article.

Original Publication:

Gianluigi Pillonetto, Tianshi Chen, Alessandro Chiuso, Giuseppe De Nicolao and Lennart Ljung, Regularized linear system identification using atomic, nuclear and kernel-based norms: The role of the stability constraint, 2016, Automatica, (69), 137-149.

<http://dx.doi.org/10.1016/j.automatica.2016.02.012>

Copyright: Elsevier

<http://www.elsevier.com/>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-130057>



Regularized linear system identification using atomic, nuclear and kernel-based norms: the role of the stability constraint

Gianluigi Pillonetto^a Tianshi Chen^b Alessandro Chiuso^a Giuseppe De Nicolao^c
Lennart Ljung^b

^a*Department of Information Engineering, University of Padova, Padova, Italy (e-mail: {giapi, chiuso}@dei.unipd.it)*

^b*Division of Automatic Control, Linköping University, Linköping, Sweden (e-mail: {tschen, ljung}@isy.liu.se)*

^c*Department of Computer Engineering and Systems Science, University of Pavia, Pavia, Italy (e-mail: giuseppe.denicolao@unipv.it)*

Abstract

Inspired by ideas taken from the machine learning literature, new regularization techniques have been recently introduced in linear system identification. In particular, all the adopted estimators solve a regularized least squares problem, differing in the nature of the penalty term assigned to the impulse response. Popular choices include atomic and nuclear norms (applied to Hankel matrices) as well as norms induced by the so called stable spline kernels. In this paper, a comparative study of estimators based on these different types of regularizers is reported. Our findings reveal that stable spline kernels outperform approaches based on atomic and nuclear norms since they suitably embed information on impulse response stability and smoothness. This point is illustrated using the Bayesian interpretation of regularization. We also design a new class of regularizers defined by “integral” versions of stable spline/TC kernels. Under quite realistic experimental conditions, the new estimators outperform classical prediction error methods also when the latter are equipped with an oracle for model order selection.

Key words: linear system identification; kernel-based regularization; atomic and nuclear norms; Hankel operator; Lasso; Bayesian interpretation of regularization; Gaussian processes; reproducing kernel Hilbert spaces

1 Introduction

Prediction error methods (PEM) are a classical tool for reconstructing the impulse response of a linear system starting from input-output measurements [32]. In the simplest scenario, one postulates a single model structure, e.g. a rational transfer function which depends on an unknown parameter vector g of known dimension. If the model contains the “true” system, and the noise source is Gaussian, estimation of g by PEM enjoys optimal asymptotic properties. In par-

ticular, this estimator can not be outperformed by any other unbiased estimator as the number of measurements goes to infinity.

However, in real applications, not only the number of measurements is always finite but also model complexity is typically unknown. This means that different model structures need to be introduced, e.g. rational transfer functions of different order. Each of them has to be fitted to data by PEM and then compared resorting e.g. to validation techniques (cross validation) or complexity criteria such as Akaike’s criterion (AIC) [3, 20]. Recent studies have however illustrated some limitations of this approach. For instance, when the data set size and/or the signal to noise ratio is not so large, it can return models with a non satisfactory prediction power on new data [41].

The above issues have motivated the development of an alternative route to system identification based on regularization techniques. The starting point is the use of high-order FIR models in combination with penalty terms (regularizers) on the impulse response g . In particular, an important class of estimators solves a convex optimization problem of

¹ This paper was not presented at any IFAC meeting. Corresponding author Gianluigi Pillonetto Ph. +390498277607. This research has been partially supported by the MIUR FIRB project RBFR12M3AC-Learning meets time: a new computational approach to learning in dynamic systems, by the Progetto di Ateneo CPDA147754/14-New statistical learning approach for multi-agents adaptive estimation and coverage control, by the Linnaeus Center CADICS, funded by the Swedish Research Council, and the ERC advanced grant LEARN, no 287381, funded by the European Research Council and by a research grant for junior researchers funded by Swedish Research Council (VR), under contract 2014-5894.

the form

$$\arg \min_g V(g) + \gamma J(g), \quad \gamma \geq 0.$$

Above, V is the so-called loss function that depends also on the input-output measurements and measures the adherence to experimental data. It is often given by the sum of squared residuals (the choice also adopted in this paper), leading to regularized least squares (ReLS). The term J is instead the regularizer which typically includes smoothness information on g . Finally, the positive scalar γ is the regularization parameter which has to suitably balance V and J . In practice, it is always unknown and has to be determined from data. An important advantage of ReLS over classical PEM is that the difficult model order determination can be replaced by estimation of few regularization variables. For instance, the stable spline estimator introduced in [40] depends only on two parameters: the regularization parameter γ and another one which enters J by determining how fast g is expected to decay to zero. Such variables can be determined e.g. by cross validation [20], C_p statistics [28, subsection 7.4] or marginal likelihood optimization [33, 34], an approach recently proved to be much effective [13, 36, 40]. In particular, in the spirit of Stein’s effect [30], a carefully tuned regularization can much reduce the variance of the estimates just introducing a small bias in the identification process. Hence, the mean squared error of the estimator can turn out inferior than that achieved by PEM [13, 40]. However, to obtain this, the choice of J is crucial since it has a major effect on the quality of impulse response reconstruction. It is thus of interest to investigate and compare the performance of different regularizers recently proposed in the system identification literature.

A recent regularization approach relies on the so called *nuclear norms*. The nuclear norm (or trace norm) of a matrix is the sum of its singular values. Being the convex envelope of the rank function on the spectral ball, it has been often used as a convex surrogate of the rank function [22]. In particular, since the McMillan degree (minimum realization order) of a discrete time time-invariant system coincides with the rank of the Hankel operator constructed from the impulse response coefficients, it is tempting to set J to the nuclear norm of the Hankel operator associated with the impulse response g . In this way, output data are described while encouraging a low McMillan degree, see [27, 31, 35, 50] and also [29, 49] for extensions of this idea for estimation e.g. of Box-Jenkins models.

Atomic norms have been also adopted as regularizers for system identification in the last years [10]. The function to be reconstructed is described as the sum of a (possibly infinite) number of basis functions, dubbed *atoms*. The penalty J is then given by the atomic norm defined, under some technical conditions, by the convex hull of the atomic set. For instance, the convex hull of one-sparse vectors of unit Euclidean norm leads to the ℓ_1 norm which enjoys important recovery properties [19], being also related to the popular LASSO procedure [52].

A motivation underlying the use of the atomic norm is that, in some sense, it represents the best convex penalty when the function is sum of few atoms. Successful applications in

signal processing and machine vision regard estimation of sparse vectors and low-rank matrices [2, 9, 18]. This technique has been recently introduced also in the system identification scenario in [48] exploiting low-order rational transfer functions as atomic set, see also [44, 45] for other approaches that use the ℓ_1 penalty.

Other recent system identification techniques exploit penalty terms induced by *kernel* functions [6, 47], which have to capture the expected features of the unknown function. In particular, the works [13, 39, 40] have proposed a class of kernels for system identification, including *stable spline (SS)/tuned-correlated (TC)*, which encodes information on smoothness and exponential stability of g , see also [12, 15] for other insights on the stable spline kernel structure and new constructions via orthonormal basis functions.

The goal of this paper is to compare the performance of ReLS equipped with atomic, nuclear or kernel-based norms via numerical studies. For this purpose, we will perform a Monte Carlo experiment where at every run the true system is a different rational transfer function randomly chosen by a MATLAB generator. The system input and the signal to noise ratio also varies from run to run, thus leading to a large variety of system identification problems. The results reveal that, in many cases, atomic and nuclear norms lead to unsatisfactory impulse response estimates. The drawbacks of these approaches are explained through the Bayesian interpretation of regularization, where different J are seen as different a priori probability density functions (pdf) assigned to g . This interpretation explains why the current implementations of atomic and nuclear norms fail in capturing important characteristics of stable dynamic systems. As far as the nuclear norm is concerned, when used in conjunction with a FIR model of order m , we shall see that it essentially corresponds to modeling the impulse response coefficients g_t as a non stationary white noise process with a variance roughly decaying as $1/t$ for $t < m/2$ and increasing as $1/(m-t)$ for $t > m/2$. This means that the prior underlying this approach does not embed any information on impulse response stability and smoothness, two key features to achieve good mean squared error properties [41]. As for ReLS equipped with the atomic norm described in [48], we will see that it suffers of the limitations of the LASSO procedure recently discussed in [4].

The results from the same Monte Carlo study will also show that estimators based on stable kernels outperform ReLS relying on atomic and Hankel nuclear norms. Furthermore, we will design a new class of regularizers induced by “integral” versions of stable spline kernels. This will lead to novel ReLS approaches for system identification that may outperform also the classical PEM equipped with an oracle for model complexity selection (the oracle selects the best model order exploiting information on the true system).

The paper is so organized. Section 2 reports the problem statement while in Section 3 we briefly review the ReLS estimators based on the Hankel nuclear norm and the atomic norm. Section 4 gives insights about limitations of the tech-

niques described in Section 3 using the Bayesian interpretation of regularization. In Section 5 we briefly review stable spline kernels introduced in [38, 40], also introducing the concepts of stable spline atomic set. In Section 6, new integral versions of stable spline kernels are introduced and used to define novel ReLS estimators for linear system identification. All of the estimators are then tested via Monte Carlo studies in Section 7. Conclusions end the paper while some mathematical details are gathered in Appendix.

2 Problem statement

We use $u(t)$ and $y(t)$ to denote, respectively, the input and the noisy output of a SISO system observed at time instant t . For convenience of notation we shall also use the notation $y_i := y(t_i)$. The measurement model we consider is of the Output-Error (OE) type:

$$y_i = (g \otimes u)_i + e_i, \quad i = 1, \dots, N \quad (2.1)$$

where g is the system impulse response, $(g \otimes u)_i$ denotes the convolution in discrete time between g and u evaluated at t_i . Finally, e_i are independent Gaussian noises of constant variance σ^2 . The results could be straightforwardly extended to non stationary noises and/or ARMAX/BJ type structures. Our problem is to estimate the impulse response g assuming the input u is (deterministic and) known and having collected the measurements $Y = [y_1 \dots y_N]^T$.

Remark 2.1 *In the sequel, we will discuss regularized estimators based on IIR (infinite impulse response) or FIR (finite impulse response) models for g and equipped with a regularizer $J(g)$. In the IIR case, following the well known concept of BIBO stability, J is said to include the stability constraint if it embeds the constraint $\sum_{t=1}^{\infty} |g_t| < \infty$. FIR models are already structurally stable. However, also in this case we will often use expressions as lack of stability constraint to mean that J does not include the information that impulse response is expected to decay to zero as a function of the time lag. In particular, this concept will be formalized in Bayesian terms.*

3 ReLS based on Hankel nuclear/atomic norms

In this section we describe two ReLS approaches which use as regularizer J either the Hankel nuclear norm or a version of the atomic norm that can be approximated by the ℓ_1 norm.

3.1 Regularization via Hankel nuclear norm

We introduce a regularizer based on the Hankel nuclear norm. Recall that the nuclear norm of a matrix A is the sum of its singular values, i.e.

$$\|A\|_* = \sum_i \sigma_i(A).$$

The motivation underlying the use of this type of penalty for system identification stems from the fact that the minimum realization order (also known as McMillan degree) of a discrete time LTI system coincides with the rank of the Hankel operator $H(g)$ constructed from the impulse response coefficients g_t , i.e.

$$H(g) = \begin{pmatrix} g_1 & g_2 & g_3 & \cdots \\ g_2 & g_3 & g_4 & \cdots \\ g_3 & g_4 & g_5 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (3.1)$$

Then, as described e.g. in [27, 35], an estimator trading data fit with low McMillan degree can be obtained by solving

$$\arg \min_g \sum_{i=1}^N (y_i - (g \otimes u)_i)^2 + \gamma \|H(g)\|_* \quad (3.2)$$

3.2 Regularization via atomic norm

A different, but related regularization approach to linear system identification, called the atomic norm regularization approach, was recently suggested in [48]. It describes the model using “atoms” and defines a model complexity measure in terms of the “atomic norm”. This complexity measure is then used for regularization.

Let \mathcal{C} be the complex plane and $\mathcal{D} = \{w \in \mathcal{C}, |w| < 1\}$ and consider below the discrete-time case. A simplistic account of the idea is as follows: Given a set of “atoms”, $G_w(z)$, $w \in \mathcal{D}$ (which can be interpreted as basis functions for the model transfer function), we can construct a linear model via linear combinations of the atoms. For a concrete feeling of the concept, think of the atoms as normalized first order system with pole (possibly complex) denoted by $w \in \mathcal{D}$, so

$$G_w(z) = \frac{1 - |w|^2}{z - w}. \quad (3.3)$$

The transfer function of any linear system can be written as a finite or countable linear combination of these atoms:

$$G(z) = \sum_k a_k G_{w_k}(z), \quad w_k \in \mathcal{D}. \quad (3.4)$$

The atomic norm of this system is defined as

$$\|G(z)\|_{\mathcal{A}} = \inf \left\{ \sum_{w_k \in \mathcal{D}} |a_k| : (3.4) \text{ holds} \right\}. \quad (3.5)$$

It has been proved in [48] that the atomic norm well approximates the Hankel nuclear norm of $G(z)$, thus motivating its use in system identification.

For computational reasons, a finite number of atoms $G_{w_k}(z)$, $k = 1, \dots, p$ are selected. If $G_{w_k}(z)$ is selected and $w_k \in \mathcal{C}$, then $G_{w_k^*}(z)$ shall be also within the selected p atoms where w_k^* is the complex conjugate of w_k . Let the impulse response of $G_{w_k}(z)$ be ρ_{w_k} and denote

$$h_{w_k}(i) = (\rho_{w_k} \otimes u)_i$$

Then, recalling (3.4), the fit between the measured output and the model is

$$\sum_{i=1}^N \left(y_i - \sum_{k=1}^p a_k h_{w_k}(i) \right)^2$$

and, if a is the vector with k -th component a_k , the atomic norm regularized estimate becomes the LASSO like

$$\hat{a} = \underset{a}{\operatorname{argmin}} \sum_{i=1}^N \left(y_i - \sum_{k=1}^p a_k h_{w_k}(i) \right)^2 + \gamma \sum_{k=1}^p |a_k|. \quad (3.6)$$

4 Bayesian interpretation of regularization: the Hankel nuclear norm and the atomic norm case

Every ReLS estimator

$$\underset{g}{\operatorname{argmin}} \sum_{i=1}^N (y_i - (g \otimes u)_i)^2 + \gamma J(g) \quad (4.1)$$

can be given a simple interpretation in Bayesian terms. To simplify the exposition, assume that the measurements model (2.1) can be written as a FIR of order m . Then, abusing notation, we use g to denote the (column) vector containing the impulse response coefficients g_1, \dots, g_m . More specifically, assume that g is a random vector with pdf

$$p(g) \propto \exp \left(-\frac{J(g)}{2\lambda} \right) \quad (4.2)$$

where “ \propto ” stands for “proportional to” while λ is a positive scale factor. Assume also that Y is corrupted by additive zero mean white Gaussian noise, independent of g , of variance σ^2 , i.e.

$$p(Y|g) \propto \exp \left(-\frac{\sum_{i=1}^N (y_i - (g \otimes u)_i)^2}{2\sigma^2} \right).$$

By the Bayes rule, the posterior of g is the product of the likelihood and the prior, i.e. $p(g|Y) \propto p(Y|g)p(g)$, so that

$$-\log p(g|Y) = \frac{\sum_{i=1}^N (y_i - (g \otimes u)_i)^2}{2\sigma^2} + \frac{J(g)}{2\lambda} + \text{Const.}$$

Hence, setting $\gamma = \sigma^2/\lambda$, one can conclude that ReLS (4.1) can be always seen as a maximum a posteriori (MAP) estimator.

The estimator (4.1) obtained from this Bayesian perspective will perform well only provided the density $p(g)$, defined through (4.2), succeeds in capturing the essential features and possibly prior knowledge on the underlying dynamical system. In particular the prior should privilege stable and possibly smooth impulse responses. Now, it is of interest to elucidate the shape of the priors underlying the two ReLS approaches introduced in the previous section.

4.1 The Hankel nuclear norm case

According to the previous Bayesian interpretation of regularization, the nuclear norm regularizer

$$J(g) = \sum_i \sigma_i(H) \quad (4.3)$$

is associated with the prior

$$p_H(g) \propto \exp \left(-\frac{\sum_i \sigma_i(H)}{2\lambda} \right).$$

Now, we would like to understand how well p_H describes typical features of an impulse response.

By symmetry, it is easy to assess that g is zero-mean but then a simple closed form expression for the distribution of its components g_i is not available. In Appendix 9.1 we describe a Markov chain Monte Carlo (MCMC) scheme to efficiently sample from p_H [24], also exploiting a closed form approximation of p_H . Fig. 1 reports some results extracted from a chain of length 1e6 setting $2\lambda = 1$, $g \in \mathbb{R}^{99}$, $H \in \mathbb{R}^{50 \times 50}$. The left panel shows the standard deviations of the impulse response coefficients g_i (solid line) while the right panel displays information on the correlation between the components of g . The outcomes show that, under p_H , the impulse response coefficients g_i are (approximately) uncorrelated and with bathtub variance. This suggests that the prior associated with the Hankel nuclear norm provides a weakly informative guess when searching for stable and smooth impulse responses. This analysis also shows that even an apparently reasonable regularizer could be associated to a meaningless prior distribution.

Remark 4.1 *Lack of stability, manifesting itself in the fact that the impulse response prior variance does not decay to zero as a function of the lag index t , should not come as a surprise. In fact, for any fixed m and any given “stable” partial impulse response $\{g_t\}_{t=1}^m$, it is always possible to find an unstable impulse response $\{h_t\}_{t=1}^m$ with the same Hankel nuclear norm. A simple example goes as follows: consider the (finite) Hankel matrix $H_m(g)$ built with the first m impulse response coefficients of $g_t = ca^{t-1}b$ and the matrix $H_m(h)$*

built with $h_t = a^{m-1} c \left(\frac{1}{a}\right)^{t-1} b$. It is simple to check that $H_m(g)$ can be obtained from $H_m(h)$ by reversing the order of rows and columns and thus $\|H_m(h)\|_* = \|H_m(g)\|_*$.

Remark 4.2 In [13], the problem of selecting a quadratic penalty for the reconstruction of stable exponentials was investigated under a deterministic framework. It has been shown that favorable mean squared error (MSE) properties can be obtained letting the diagonal (and off-diagonal) elements of the kernel (covariance) matrix decay exponentially to zero (as the entry number increases). In the Hankel case, it is shown in Appendix 9.1 that one approximately has $\text{Var}(g_t) \propto t^{-1}$ for the first impulse response coefficients and $\text{Var}(g_t) \propto (m-t+1)^{-1}$ for the last ones (recall that m is the order of the adopted FIR). So, the Hankel regularizer does not include exponential dynamics, a key point to well trade-off bias and variance. Note also from the particular profile of $\text{Var}(g_t)$ that the choice of m has an important influence on the structure of the Hankel regularizer.

4.2 The atomic norm case

It has been shown in [48] that, restricting to $\{g_t\}_{t \in \mathbb{Z}^+} \in \ell_2$, the atomic norm (3.6) is equivalent to the Hankel nuclear norm. For this equivalence to extend to the “finite” nuclear norm (4.3) the size of the Hankel matrix (3.1) (and thus the truncation index m) needs to grow to infinity. Note that the restriction $\{g_t\}_{t \in \mathbb{Z}^+} \in \ell_2$ is critical since, as we have seen, for any finite m the Hankel nuclear norm does not include a “stability” constraint. The situation is different with the atomic norm as any finite sum of (stable) atoms (3.4) will always result in a stable impulse response. Yet the ℓ_1 penalty on the coefficients in (3.6) (see also (4.4) below) may introduce severe bias. More insights are discussed below. Let $g_t = \sum_{j=1}^p a_j \rho_j(t)$. In view of (3.6), the atomic norm regularizer

$$J(g) = \sum_{j=1}^p |a_j| \quad (4.4)$$

amounts to a Bayesian prior

$$p_{AN}(g) \propto \exp\left(-\frac{\sum_{j=1}^p |a_j|}{2\lambda}\right)$$

so that the unknown parameters a_j are Laplace distributed independent random variables, all having the same variance. As in the case of LASSO, the solutions of (3.6) enjoy some sparsity properties meaning that, when γ increases, more and more elements of vector a are forced to zero. This may seem an appealing property but the results may fail to meet the expectations. In fact, the ℓ_1 penalty may introduce an excessive penalty on some large coefficients a_j to obtain sparsity, as documented in the Statistics literature [21, 56] and also recently demonstrated and discussed in [4]. In particular, assume that the true impulse response g is the sum of a finite number of atoms. For a large enough value of γ , all the a_j which do not contribute to g will go to zero

but the linear penalty in (3.6) will still yield to biased estimates of the other expansion coefficients, so that the resulting estimate is oversmoothed (biased). An example of this phenomenon is illustrated in Fig. 2 (left panel) which reports the results obtained via a Monte Carlo study where the impulse response g is fixed while different noise and input realizations are generated at every run, in the same way as described in Section 7. The estimates of g are obtained by (3.6) using 300 input-output samples, with γ chosen by an oracle. The latter has access to the true g and selects the regularization parameter which minimizes the mean squared error (the concept of oracle-based estimation is further detailed in subsection 7.2). Under a Bayesian viewpoint, the estimator suffers from the equal probability assigned to all the atomic functions so that the oracle can select few atoms only assuming that the prior variance (which is proportional to λ) is quite low. Obviously, this introduces a bias. It is worth asking if the adoption of different atoms and an unequal weighting on their expansion coefficients may be a remedy. We explore this issue in the next section.

5 Stable spline kernels and atoms

In the first part of this section we review the stable spline kernel, comparing its features with the prior induced by the Hankel nuclear norm discussed in section 4.1. We then compare the structure of the stable spline estimator with the atomic approach of section 4.2, also introducing the concept of stable spline atoms.

5.1 Stable spline kernels

According to its stochastic interpretation, we have seen that every ReLS can be seen as a Bayesian estimator. In particular, each quadratic penalty can be obtained modeling the impulse response as a particular (zero-mean) Gaussian process. This implies that fixing the covariance (kernel) is equivalent to fixing the quadratic regularizer.

In system identification the kernel should include information on smoothness and exponential stability of the impulse response g . One choice suggested in the literature is the so called first-order Stable Spline kernel [38], also called TC in [13]. Letting \mathbb{E} denote expectation, for $t = 1, 2, \dots$ and $s = 1, 2, \dots$ it is defined by

$$K(s, t) = \mathbb{E}[g_s g_t] \propto \alpha^{\max(s, t)}, \quad 0 \leq \alpha < 1, \quad (5.1)$$

while it is null elsewhere. The second-order version was proposed in [40]. It is given by

$$\frac{\alpha^{s+t} \alpha^{\max(s, t)}}{2} - \frac{\alpha^{3 \max(s, t)}}{6}, \quad 0 \leq \alpha < 1 \quad (5.2)$$

and leads to smoother impulse response realizations. Notice that both kernels are parametrized by α which is interpreted as an unknown hyperparameter.

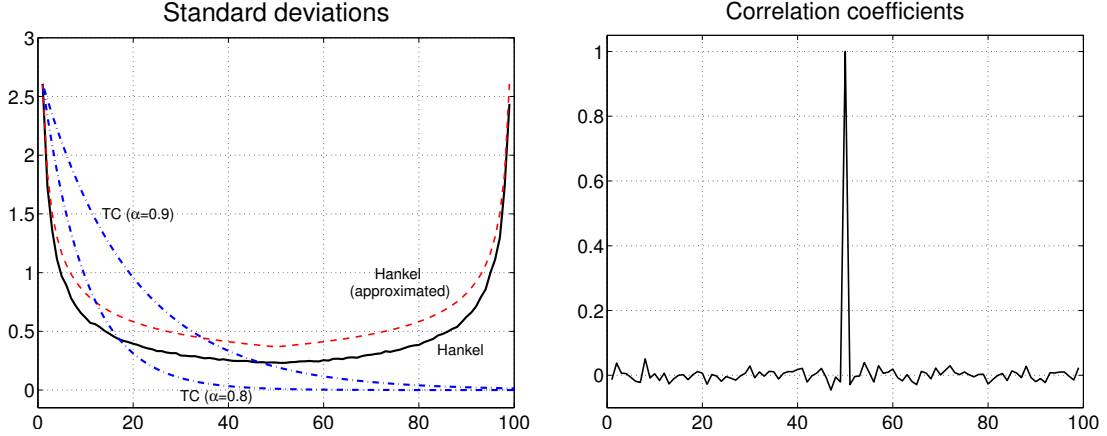


Fig. 1. **Prior induced by the Hankel Nuclear Norm:** the impulse response coefficients are contained in the vector $g \in \mathbb{R}^{99}$, modeled as a random vector with probability density function $p_H(g) \propto \exp(-\|H(g)\|_*)$. Here, $\|\cdot\|_*$ is the nuclear norm while $H(g)$ is the Hankel matrix (3.1) of size 50×50 . *Left:* standard deviations of the impulse response coefficients g_t reconstructed by MCMC (solid line) and approximated using the prior (9.1.9.2) (dashed line) derived in Appendix. The figure also displays the standard deviations of g_t when g is a Gaussian random vector with stable spline covariance (5.1) for two different values of α (dashdot lines). *Right:* 50-th row of the matrix containing the correlation coefficients returned by the MATLAB command `corrcoef(M)` where each row of the $1e6 \times 99$ matrix M contains one MCMC realization of g under the Hankel prior $p_H(g) \propto \exp(-\|H(g)\|_*)$.




Fig. 2. **Monte Carlo experiment with a fixed impulse response:** true impulse response (thick line) and mean \pm standard deviation (dashed line, computed after 100 runs) of the estimators At+Or (left) and iTS (implemented as described in subsection 7.2). Recall that, differently from iTS, At+Or is not implementable in practice since γ is tuned using the true impulse response.

The left panel of Fig. 1 reports the standard deviations of a random vector g whose covariance is the sampled version of (5.1) with α set to 0.8 and 0.9 respectively (dashdot lines). Differently from the bathtub prior induced by the Hankel nuclear norm (solid line), the stable spline kernel describes system dynamics which go exponentially to zero. Further, the hyperparameter α enhances model flexibility since it permits to tune the decay rate.

Also, while nonstationary white noise underlies the Hankel prior (see the right panel of Fig. 1), (5.1) introduces correlation among the impulse response coefficients, hence including information on impulse response smoothness. This is important to have good MSE properties as discussed in Remark 4.2. With the same remark in mind, note also that,

differently from the Hankel nuclear norm case, the stable spline prior shape is independent of the selected FIR order, thus making its effect on the estimation process more transparent. If m is increased e.g. to $2m$, the statistics of the first m impulse response coefficients remain the same.

5.2 Stable spline atoms and regularizer

Let g be a discrete-time Gaussian stochastic process with covariance proportional to the stable spline kernel (5.1). Then, the following proposition characterizes the architecture of the minimum variance estimator of g given the measurements (2.1).

Proposition 5.1 *Let*

$$y_i = (g \otimes u)_i + e_i, \quad i = 1, \dots, N$$

where e_i are all mutually independent, Gaussian, with zero-mean and variance σ^2 . Assume also that g is a zero-mean Gaussian and causal process, independent of the noise, with covariance

$$\mathbb{E}[g_s g_t] = \lambda \alpha^{\max(s,t)} \quad (5.3)$$

with λ a positive scale factor. Then, one has

$$\alpha^{\max(s,t)} = \sum_{j=1}^{\infty} \zeta_j \rho_j(s) \rho_j(t) \quad (5.4)$$

where

$$\rho_j(t) = \sqrt{2} \sin\left(\frac{\alpha^t}{\sqrt{\zeta_j}}\right), \quad \zeta_j = \frac{1}{(j\pi - \pi/2)^2} \quad (5.5)$$

and

$$\mathbb{E}[g_t | Y] = \sum_{j=1}^{\infty} \hat{a}_j \rho_j(t) \quad (5.6)$$

where

$$\hat{a} = \arg \min_a \sum_{i=1}^N \left(y_i - \sum_{j=1}^{\infty} h_{ij} a_j \right)^2 + \gamma \sum_{j=1}^{\infty} \frac{a_j^2}{\zeta_j} \quad (5.7)$$

with

$$h_{ij} = (\rho_j \otimes u)_i, \quad \gamma = \frac{\sigma^2}{\lambda}.$$

Finally, the estimate admits also the closed-form expression

$$\mathbb{E}[g_t | Y] = \sum_{i=1}^N \hat{c}_i (K(\cdot, t) \otimes u(\cdot))_i \quad (5.8)$$

where \hat{c}_i are the components of the vector \hat{c}

$$\hat{c} := (A + \gamma I_N)^{-1} Y \quad (5.9)$$

with

$$A_{ij} = \sum_{t=1}^{\infty} u(j-t) \left(\sum_{k=1}^{\infty} u(i-k) K(t, k) \right) \quad (5.10)$$

■

The result above leads to the following comments:

- when the stable spline prior (5.1) is used, according to (5.6), the impulse response estimate is searched in the subspace spanned by the functions ρ_j given by (5.5). It is then natural to call

$$\mathcal{A}_\alpha = \left\{ \sin\left(\frac{\pi\alpha^t}{2}\right), \sin\left(\frac{3\pi\alpha^t}{2}\right), \sin\left(\frac{5\pi\alpha^t}{2}\right), \dots \right\}$$

the (first-order) *stable spline atomic set*. Such a set is parametrized by the positive scalar α which measures the distance from instability. Note also that the *stable spline atoms* ρ_j are ordered in such a way that their energy content at high frequencies increases as j augments and that their structure is much different w.r.t. that adopted in [10] and reported in (3.3);

- atomic norms are typically designed in such a way as to assign the same penalty to each expansion coefficient a_j . Instead, according to (5.7,5.5), given $g = \sum_{j=1}^{\infty} a_j \rho_j$, the (first-order) stable spline estimator uses the *stable spline regularizer*

$$J(g) = \sum_{j=1}^{\infty} \frac{a_j^2}{\zeta_j} \quad (5.11)$$

where ζ_j are weights decaying to zero. The expansion coefficients a_j are thus constrained to decay to zero at a rate which guarantees both impulse response regularity and system stability. Hence, penalizing high-frequency components, also by adjusting γ , the stable spline regularizer may privilege more parsimonious models¹, possibly leading to estimators which better trade bias and variance;

- the first step in the design of atomic-norm based estimators is the selection of the atomic set whose convex hull then defines the regularizer. In our stochastic framework, both of these steps are condensed in the choice of the covariance (kernel). Indeed, the kernel encodes both the atoms ρ_j and the regularizer, defined by ζ_j . This subsumes modeling and computational advantages: one needs neither to choose the number of atoms to be used (the kernel includes an infinite number of basis functions ρ_j so that no truncation affects the impulse response representation) nor store any basis function in memory. In fact, the estimate (5.8) can be computed using the kernel just inverting a matrix, as shown in (5.9). This feature is related to what is called the *kernel trick* in the machine learning literature [47].

6 Integral versions of stable spline kernels

When the stable spline kernel is used, we have seen from (5.5) that the atoms ρ_j depend on a single parameter α which establishes the decay rate of the eigenfunctions. Following also [11], a way to further enrich the hypothesis

¹ Note that, when using the approach proposed in [48], the atoms reported in (3.6) contain the term $1 - |w|^2$ which implicitly penalises basis functions close to the unit circle. However, no penalty is assigned to high-frequency impulse response components.

space, making it more flexible, is to exploit different values of α . However, differently from [11], our aim is the synthesis of a new kernel (in closed form) able to “contain” an infinite number of different decay rates, but which remains function only of a finite number of hyperparameters. In our Bayesian framework, we can obtain this by modeling the impulse response g as the sum of i.i.d. stochastic processes whose stable spline kernels differ in the value of α . In particular, let us introduce two hyperparameters: α_m , related to the fastest system pole, and α_M , connected with the slowest dynamics. Then, we can consider p values α_i satisfying $\alpha_m \leq \alpha_1 < \alpha_2 < \dots < \alpha_p \leq \alpha_M$ and equally spaced with step Δ_α , then building the kernel

$$\Delta_\alpha \sum_{i=1}^p \alpha_i^{\max(s,t)} \quad (6.1)$$

which induces the richer atomic set

$$\mathcal{A}_{[\alpha_m, \alpha_M]} = \left\{ \sin \left(\frac{\alpha_i^t}{\sqrt{\zeta_j}} \right), j = 1, 2, \dots \text{ and } i = 1, 2, \dots, p \right\} \quad (6.2)$$

where we still have $\zeta_j = \frac{1}{(j\pi - \pi/2)^2}$. Letting $p \rightarrow \infty$, so that $\Delta_\alpha \rightarrow 0$ in (6.1), the sum becomes the integral $\int_{\alpha_m}^{\alpha_M} \alpha^{\max(s,t)} d\alpha$ which leads to the new *first-order integral stable spline kernel* called iTC and reported in Table 1. The same procedure can be repeated starting from (5.2) obtaining the *second-order integral stable spline kernel*, called iSS in Table 1.

Finally, these two kernels iTC and iSS can be summed up, obtaining the kernel dubbed ITS in Table 1. This kernel thus synthesizes an infinite number of atoms that not only have different decay rates $\alpha \in [\alpha_m, \alpha_M]$ but also two different smoothness levels.

We can now introduce estimators based on the different regularization matrices P which are function of the hyperparameter vector η as reported in Table 1. The estimators have the same structure as documented in [40] and [13]. In particular, assuming a high order FIR, it is useful to rewrite the measurements model (2.1) as

$$Y = \Phi g + E \quad (6.3)$$

Here, as in section 4, g now denotes the m -dimensional vector whose components are the impulse response coefficients while the regression matrix Φ is defined by the input samples. Then, the noise variance σ^2 is estimated from the residuals obtained by fitting g via least squares. The hyperparameter vector is instead determined through marginal likelihood optimization [36], i.e.

$$\hat{\eta} = \arg \max_{\eta} Y^T \Sigma_\eta^{-1} Y + \log \det(\Sigma_\eta) \quad (6.4)$$

where $\Sigma_\eta = \Phi P \Phi^T + \sigma^2 I_N$. Finally, the impulse response

| | |
|------------|---|
| TC | $P_{kj}(\eta) = \lambda \alpha^{\max(k,j)};$ $\lambda \geq 0, 0 \leq \alpha < 1, \eta = [\lambda, \alpha]$ |
| SS | $P_{kj}(\eta) = \lambda \left(\frac{\alpha^{k+j+\max(k,j)}}{2} - \frac{\alpha^{3\max(k,j)}}{6} \right)$ $\lambda \geq 0, 0 \leq \alpha < 1, \eta = [\lambda, \alpha]$ |
| iTC | $P_{kj}(\eta) = \lambda \frac{\alpha_M^{\max(k,j)+1} - \alpha_m^{\max(k,j)+1}}{\max(k,j)+1}$ $\lambda \geq 0, 0 \leq \alpha_m \leq \alpha_M < 1; \eta = [\lambda, \alpha_m, \alpha_M]$ |
| iSS | $P_{kj}(\eta) = \frac{\lambda \alpha_M^{k+j+\max(k,j)+1}}{2(k+j+\max(k,j)+1)} - \frac{\lambda \alpha_M^{3\max(k,j)+1}}{18\max(k,j)+6}$ $- \frac{\lambda \alpha_m^{k+j+\max(k,j)+1}}{2(k+j+\max(k,j)+1)} + \frac{\lambda \alpha_m^{3\max(k,j)+1}}{18\max(k,j)+6}$ $\lambda \geq 0, 0 \leq \alpha_m \leq \alpha_M < 1; \eta = [\lambda, \alpha_m, \alpha_M]$ |
| ITS | $P_{kj}(\eta) = \lambda \frac{\alpha_M^{\max(k,j)+1} - \alpha_m^{\max(k,j)+1}}{\max(k,j)+1}$ $+ \frac{\lambda \alpha_M^{k+j+\max(k,j)+1}}{2(k+j+\max(k,j)+1)} - \frac{\lambda \alpha_M^{3\max(k,j)+1}}{18\max(k,j)+6}$ $- \frac{\lambda \alpha_m^{k+j+\max(k,j)+1}}{2(k+j+\max(k,j)+1)} + \frac{\lambda \alpha_m^{3\max(k,j)+1}}{18\max(k,j)+6}$ $\lambda \geq 0, 0 \leq \alpha_m \leq \alpha_M < 1; \eta = [\lambda, \alpha_m, \alpha_M]$ |

Table 1
List of (k, j) entries of different regularization matrices P built using Stable Kernels

estimate is the solution of

$$\arg \min_g \|Y - \Phi g\|^2 + g^T P^{-1} g$$

(note that $J(g)$ has become $g^T P^{-1} g$ above), which is given by

$$\hat{g} = P \Phi^T \Sigma_\eta^{-1} Y \quad (6.5)$$

with η set to its estimate $\hat{\eta}$.

7 Simulated data: Monte Carlo studies

7.1 Data sets and performance index

We compare different estimators for discrete-time system identification. To this aim, we resort to two Monte Carlo studies whose implementation details are the same as described in Section 7.2 of [41]. Here, we just recall that each Monte Carlo consists of 1000 runs. At each run, g is defined by a different rational transfer function, given by the ratio of two polynomials of the same order (varying from 1 to 30), randomly obtained by a Matlab generator. The system, initially at rest, is fed with an input obtained by filtering a

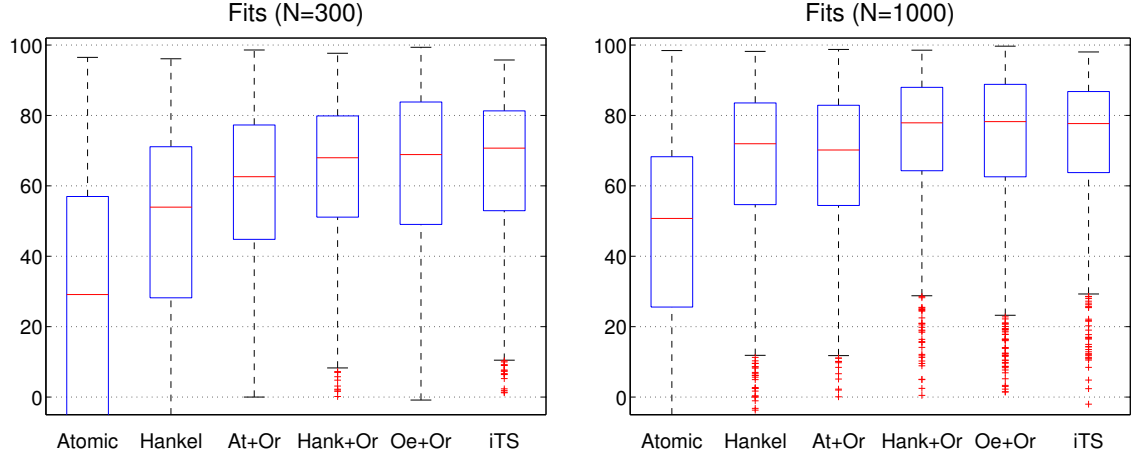


Fig. 3. **Monte Carlo experiment:** boxplots of the fits using 300 (left panel) or 1000 (right panel) input-output samples. Recall that At+Or, Hank+Or and Oe+Or are not implementable in practice since they exploit the knowledge of the true impulse response to tune model complexity.

| | <i>Atomic</i> | <i>Hankel</i> | <i>At+Or</i> | <i>Hank+Or</i> | <i>Oe+Or</i> | <i>iTS</i> |
|------------|---------------|---------------|--------------|----------------|--------------|------------|
| $N = 300$ | 4.5 | 31.9 | 59.5 | 63.7 | 64.8 | 65.1 |
| $N = 1000$ | 39.2 | 63.3 | 67.1 | 73.7 | 72.6 | 72.9 |

Table 2

Monte Carlo experiment: average fit achieved by PEM equipped with oracle (Oe+Or), ReLS based on Atomic norms (Atomic and At+Or), Hankel nuclear norms (Hankel and Hank+Or) and the new estimator based on the stable kernel iTS, using 300 or 1000 input-output samples. Recall that At+Or, Hank+Or and Oe+Or are not implementable in practice since they exploit the knowledge of the true impulse response to tune model complexity.




Fig. 4. **Monte Carlo experiment:** boxplots of the fits returned by ReLS equipped with different stable kernels using 300 (left panel) or 1000 (right panel) input-output samples. All the estimators are implementable in practice.

| | <i>iTS</i> | <i>iTC</i> | <i>iSS</i> | <i>TC</i> | <i>SS</i> |
|------------|------------|------------|------------|-----------|-----------|
| $N = 300$ | 65.1 | 64.8 | 64.4 | 63.2 | 62.3 |
| $N = 1000$ | 72.9 | 72.7 | 71.6 | 70.6 | 69.2 |

Table 3

Monte Carlo experiment: average fit returned by ReLS equipped with different stable kernels using 300 or 1000 input-output samples. All the estimators are implementable in practice.




Fig. 5. **Monte Carlo experiment using a different system generator:** boxplots of the fits achieved by Oe+Or and iTS using 300 (left panel) or 1000 (right panel) input-output samples.

zero mean unit variance white Gaussian noise through a 2nd order rational transfer function which also varies from run to run. Output data are corrupted by a white Gaussian noise, with the SNR (ratio between the variance of the noiseless output and the noise) randomly drawn at every run in the interval $[1, 10]$. The first and the second Monte Carlo study differ in the number N of available input-output measurement, equal to 300 or 1000, respectively. Given an estimator \hat{g} , its performance index is evaluated computing the fit measures

$$\mathcal{F}_j = 100 \times \left(1 - \frac{\|g^j - \hat{g}^j\|}{\|g^j\|} \right), \quad j = 1, \dots, 1000 \quad (7.1)$$

where g^j and \hat{g}^j are the true and the estimated impulse response at the j -th run.

7.2 Estimators compared via the Monte Carlo study

Oe+Or All the impulse response estimators introduced so far depend on an unknown parameter vector, denoted by η , which controls model complexity. For instance, in all the regularized techniques η contains at least the regularization parameter γ . When using PEM, η instead represents the order of different model structures. For instance, consider the use of rational transfer functions for discrete-time system identification. Then, η is the degree of the polynomials $B(z)$ and $A(z)$ composing the transfer function given, in the z -transfer domain, by

$$G(z) = \frac{B(z)}{A(z)} \quad (7.2)$$

Hereby, \hat{g} is said to be an oracle-based estimator if, having access to the true impulse response g , it determines model complexity by maximizing the fit measure (7.1) w.r.t. η . Note that such an impulse response estimator is never implementable in practice since the true impulse response g is not available. This identification procedure is however useful since it provides a performance reference. In particular,

$Oe + Or$ denotes the following PEM procedure. First, the Matlab function `Oe.m` is used to fit the model structures (7.2) to data for $\eta = 1, \dots, 30$ (the information that system was initially at rest is given to the estimator). Then, among the 30 impulse response estimates obtained, $Oe + Or$ returns that maximizing (7.1).

Hankel, Hank+Or We have implemented two variants of the estimator (3.2) based on the Hankel nuclear norm adopting a FIR model of order $m = 99$ with the size of the Hankel matrix $H(g)$ equal to 50×50 . In the first version, dubbed Hankel, the regularization parameter is estimated via cross validation. Data are divided into an identification and validation data set of equal size. For every value of γ in the grid defined by the MATLAB command `logspace(-5, 4, 50)`, the solution (3.2) is obtained from the identification data using the software CVX [25,26]. The regularization parameter $\hat{\gamma}$ is the one leading to the best prediction on the validation set according to a quadratic fit criterion. Finally, the impulse response estimate is achieved by solving (3.2) using $\gamma = \hat{\gamma}$ and all the available data. The second version of the estimator exploits an oracle which, at every run, selects the value of γ in the set `logspace(-5, 4, 50)` which maximizes the fit (7.1).

Atomic, At+Or Two types of atomic estimators (3.6) have been implemented. In both the variants, the atomic set is defined by the poles $w_k = \alpha e^{\sqrt{-1}\beta}$ and their complex conjugate w_k^* where, using a MATLAB notation, α and β take values on the two grids `[0.02:0.02:0.98 0.99 0.999]` and `[0 pi/50:pi/50:pi]`, respectively. Optimization (3.6) is then performed constraining each couple of expansions coefficients a_k related to complex conjugate poles to be equal and real.² The first estimator, dubbed Atomic, obtains at each Monte Carlo run the value of the regularization parameter via 10-fold cross validation. It has been implemented exploiting the MATLAB software package `glmnet` [23], providing the information on system initial conditions

² We have also implemented the estimator (3.6) equipped with a different atomic set given by discrete-time Laguerre basis functions [54]. Results (not shown) are similar to those described in the sequel obtained adopting (3.3).

and input delay. We have also used the MATLAB commands `options.lambda=logspace(-5,4,100)` to define the grid where the regularization parameter is searched and `options.intr=0` to specify that the relation between g and the system output is linear (and not affine). The second estimator, dubbed At+Or, is implemented in the same way, except that the regularization parameter is selected by the oracle.

TC,SS,iTC,iSS,iTS These are the estimators (6.5) equipped with the stable kernels described in Section 6, with the dimension of g set to $m = 100$ and the hyperparameter vector η estimated via marginal likelihood optimization (6.4). Note that all of these estimators are implementable in practice

7.3 Results

Boxplots of the fits achieved by Atomic, At+Or, Hankel, Hank+Or, Oe+Or and iTS in the two Monte Carlo studies are displayed in the left and right panel of Fig. 3. Table 2 also shows the average fits.

It is apparent that the fits of Hankel are significantly smaller than those returned by Oe+Or. The more so in the first experiment where the average fit of Hankel is 31.9 while that of Oe+Or is 64.8. Instead, the performance of Hank+Or (γ is chosen by the oracle) is virtually identical to that of Oe+Or. As for Atomic, its performance is not satisfactory, inferior than that of Hankel. Only using the oracle-based estimator At+Or, the performance becomes comparable with that of Oe+Or. Instead, the integral prior iTS largely outperforms the other regularized system identification approaches proposed in the literature. Remarkably, its performance is very close or also superior to that of the oracle-based approaches. For instance, in the two experiments iTS provides average fits equal to 65.1 and 72.9 whereas Oe+Or return 64.8 and 72.6.

The beneficial effect of the unequal weighting on the expansion coefficients is evident also reconsidering Fig. 2: iTS (right panel, results with γ estimated via marginal likelihood) outperforms At+Or (left panel, results with γ estimated by the oracle).

Finally, Fig. 4 and Table 3 permit to compare the performance of all the ReLS approaches equipped with the stable kernels. As a matter of fact, all the estimators perform well, revealing the importance of informing the estimation process of system stability. Notice also that iTS provides the best results, pointing out benefits of integral versions of stable kernels.

7.4 Use of a different system generator

The random generator adopted so far defines challenging systems, with poles often located at high frequencies. However, a visual inspection reveals that the average number of significant Hankel singular values is rather small, being around 10. A consequence is that the mean order selected by Oe+Or is around 5. We have thus found of interest to repeat the experiment adopting a different “higher-order” generator. The rational transfer function order is now fixed

to 30 and the poles are selected iterating the following procedure at every Monte Carlo run: With equal probability a real or a couple of complex conjugate poles is added to the denominator until its order reaches 30. In the case of a real pole, it is randomly drawn from a uniform distribution on $[-0.95, 0.95]$, while the absolute values and phases of the complex conjugate pairs are independent random variables uniform on $[0, 0.95]$ and $[0, \pi]$, respectively. The zeros are then selected in the same way except that their absolute values are drawn in the interval $[0, 2]$.

Boxplots of the fit values are reported in Fig. 5, restricting the comparison to Oe+Or and iTS. Remarkably, the advantage of iTS over Oe+Or increases: its average fit is 65.3 vs 59.3 when $N = 300$ and 76.7 vs 73.3 when $N = 1000$. This can be explained considering that, on average, Oe+Or now selects models of order 13 and this may further undermine optimal asymptotic properties of PEM under correct order specification. In addition, such an estimator has now to hinge on higher-dimensional non convex problems, possibly more prone to local minima. Under these circumstances, ReLS equipped with empirical Bayes can be especially useful, see also [36, 37] for insights on marginal likelihood effectiveness in controlling model complexity.

8 Real data: temperature prediction

To test the algorithms on real data we have also considered thermodynamic modeling of buildings. We placed sensors in two rooms of a small two-floor residential building of about 80 m^2 and 200 m^3 ; the sensors have been placed only on one floor (approximately 40 m^2) and their location is approximately shown in Fig. 6 (top panel). The larger room is the living room while the smaller is the kitchen. The experimental data was collected through a wireless sensor network made of 8 *Tmote-Sky* nodes produced by Moteiv Inc, each of them is provided with a temperature sensor. The building was inhabited during the measurement period, which lasted for 8 days starting from February 24th, 2011; samples were taken every 5 minutes. The heating system was controlled by a thermostat; the reference temperature was manually set every day depending upon occupancy and other needs. The location of the 8 sensors was as follows:

- Node #1 (label 137 in Fig. 6) was above a cabinet (2.5 meters high).
- Node #2 (label 111) was above a sideboard, about 1.8 meters high, close to thermoconvector.
- Node #3 (label 139) was above a cabinet (2.5 meters high).
- Node #4 (label 140) was placed on a bookshelf (1.5 meters high).
- Node #5 (label 141) was placed outside.
- Node #6 (label 153) was placed above the stove (2 meters high).
- Node #7 (label 156) was placed in the middle of the room, hanging from the ceiling (about 2 meters high).
- Node #8 (label 160) was placed above one radiator and was meant to provide a proxy of water temperature in

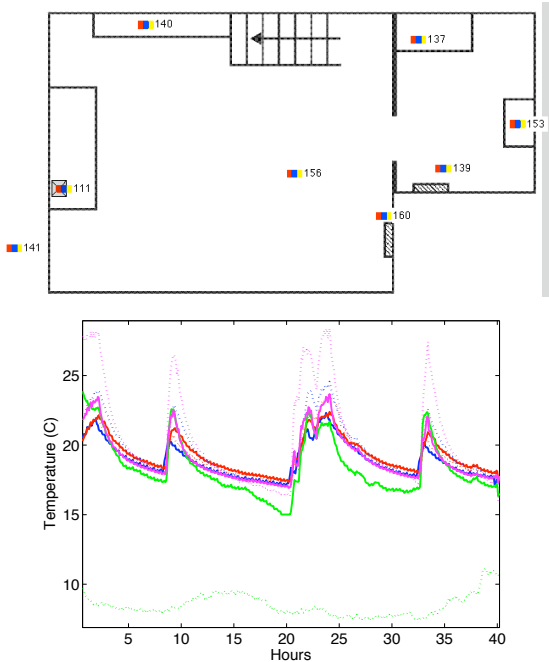


Fig. 6. Nodes location (top) and measured temperatures during the first 40 hours (bottom).

the heating system.

This gives a total of 8 temperature profiles. A preliminary inspection of the measured signals, reported in Fig. 6 (bottom panel) reveals the high level of collinearity which is well-known to complicate the estimation process in System Identification [7, 32, 51].

We only consider Multiple Input-Single Output (MISO) models, with the temperature from the first node as output (y_i) and the other 7 temperatures as inputs (u_i^j , $j = 1, \dots, 7$). We leave identification of a full Multiple Input-Multiple Output (MIMO) model for future investigation. We split the available data into 2 parts; the first $N_{id} = 1000$ temperature samples are the identification data while the last $N_{test} = 1500$ are used for test purposes. The notation y^{test} identifies the test data. Note that $N_{id} = 1000$, with 5 minute sampling times, corresponds to around $\simeq 80$ hours; this is a rather small time interval and, as such, models based on these data cannot capture seasonal variations. Consequently, in our experiments we assume a “stationary” environment and normalize the data so as to have zero mean and unit variance before identification is performed.

We envision that model predictive based methodologies, (see [8] and the recent papers [55], [43], [17]), may be effective for these applications and, as such, we evaluate our new estimators based on their ability to predict future data. The predictive power of the model is measured for

k -step-ahead prediction on *test* data, as the fit:

$$\mathcal{F}^k := 100 \times \left(1 - \frac{\sqrt{\sum_{i=k}^{N_{test}} (y_i^{test} - \hat{y}_{i|i-k})^2}}{\sqrt{\sum_{i=k}^{N_{test}} (y_i^{test})^2}} \right). \quad (8.1)$$

Identification has been performed using ARMAX+Or, Hankel and iTS. More specifically, ARMAX+Or exploits ARMAX models formed by polynomials of the same order. It has access to the test set and selects that model order which maximizes $\sum_{k=1}^{100} \mathcal{F}^k$. In particular, it turns out that setting the order to 7 provides the best average fit on an horizon of length around 8 hours.

As for the two regularized estimators, consider ARX models of the form $y_i = (g^1 \otimes y)_i + \sum_{j=1}^7 (g^{j+1} \otimes u^j)_i + e_i$, where the $\{g^j\}$ are the 8 unknown one-step ahead predictor impulse responses. Then, both Hankel and iTS assume the form

$$\arg \min_{g^j} \sum_{i=1}^N \left(y_i - (g^1 \otimes y)_i - \sum_{j=1}^7 (g^{j+1} \otimes u^j)_i \right)^2 + \gamma \sum_{j=1}^8 J(g^j),$$

differing only in the adopted J . Both these estimators are then implemented in the same way described in the previous section.

The three estimators are compared using as performance indexes \mathcal{F}^k defined in (8.1). The results are reported in Fig. 7 (left panel): similarly to what happened using simulated data, the performance of ARMAX+Or and iTS is similar and superior than that of Hankel. Sample trajectories of one-hour-ahead test data prediction obtained by iTS and Hankel are also visible in Fig. 7 (right panel).

9 Conclusions

The results of this paper highlight the importance of Bayesian interpretation of ReLS for system identification. In particular, the Bayesian framework offers transparent guidelines for selecting regularizers that capture crucial system features. Another use regards the assessment of existing regularizers: for instance, the drawbacks of the Hankel nuclear norm have been linked to the adoption of a Bayesian prior which models the impulse response as a nonstationary white noise.

Similar drawbacks affect also other recent approaches that model the impulse response as the superimposition of a large set of atoms, employing an atomic norm as regularizer. In the literature, it has been argued that any reasonable penalty function should be constant on the set of atoms [10]. Actually, this can result in penalty functions with poor capability of controlling model complexity, thus leading to estimators with large variance. Indeed, if the number of atoms to be included tends to infinity, all of them being e.g. of unit ℓ_1 norm, any regularizer including smoothness and system stability can not assign the same penalty to the

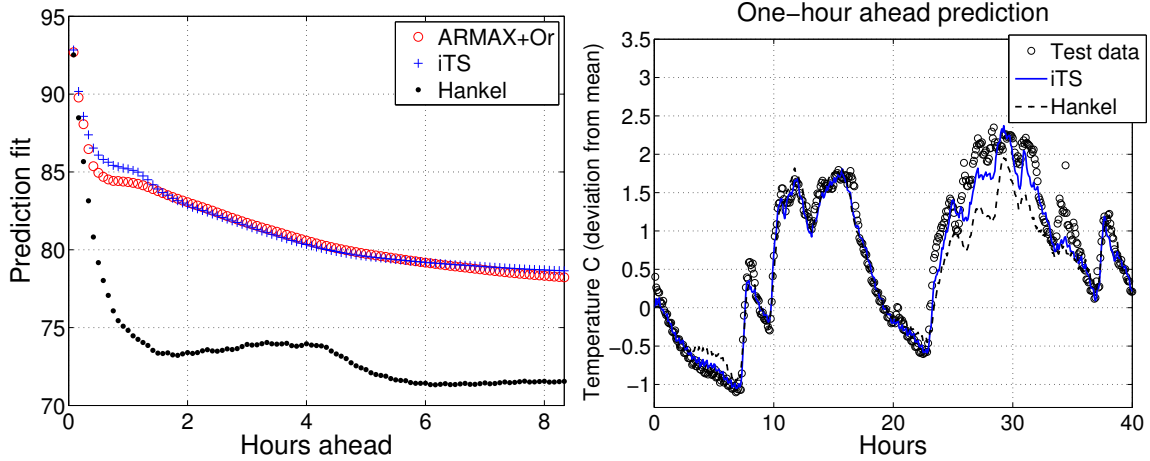


Fig. 7. Test set prediction fit \mathcal{F}^k as defined in (8.1) (left) and 1-hour ahead test set prediction (right). Recall that ARMAX+Or is not implementable in practice since it exploits the knowledge of the test set to tune model complexity.

atoms. Including smoothness and stability constraints in the estimation process, e.g. using stable spline kernels, instead leads to a kind of regularizer whose weights are non uniform. Advantages are confirmed by the simulation studies: in the Monte Carlo experiments the family of regularizers induced by stable kernels, including the three novel covariances iTC, iSS and iTS, performs systematically better than other nuclear and atomic techniques.

In conclusion, we also stress that, even if drawbacks of the Hankel norm have been here illustrated, this does not mean that such regularizer can not be useful for system identification. For example, the MIMO case has not been well investigated yet and there can be also cases where large magnitude corruptions or un-modelled dynamics can be well described by Hankel or weighted Hankel norms [35, 46]. An interesting perspective is also the design of estimators that combine stable kernels and atomic norms. Preliminary work on this can be found in [14, 42].

Appendix

9.1 Hankel nuclear norm prior: approximation and MCMC reconstruction

Our aim is to design an MCMC scheme able to reconstruct in sampled form the prior

$$p_H(g) \propto \exp\left(-\frac{\sum_i \sigma_i(H)}{2\lambda}\right)$$

associated to the Hankel nuclear norm regularizer. The key point is the definition of a proposal density leading to an efficient Metropolis-Hastings update, e.g. see [24]. For these purposes, it is useful to introduce the novel regularizer

$$J(g) = \sum_i \sigma_i^2(H).$$

According to the Bayesian interpretation of regularization, the associated prior is

$$\begin{aligned} \tilde{p}_H(g) &\propto \exp\left(-\frac{\sum_i \sigma_i^2(H)}{2\lambda}\right) \\ &= \exp\left(-\frac{\text{tr}(HH^T)}{2\lambda}\right). \end{aligned} \quad (9.1)$$

where $\text{tr}(HH^T)$ is the trace of HH^T . The last equality, together with simple calculations, leads to the following result.

Proposition 9.1 *Let $g \in \mathbb{R}^m$ and $H \in \mathbb{R}^{p \times p}$, where $m = 2p - 1$. If g is a random vector with pdf $\tilde{p}_H(g)$, then all the g_t are independent and Gaussian. In particular, one has*

$$g_t \sim \begin{cases} \mathcal{N}\left(0, \frac{\lambda}{t}\right) & \text{if } 1 \leq t \leq \frac{m+1}{2} \\ \mathcal{N}\left(0, \frac{\lambda}{m-t+1}\right) & \text{if } \frac{m+1}{2} < t \leq m \end{cases} \quad (9.2)$$

■

Thus, \tilde{p}_H describes the impulse response coefficients as white noise whose variance first decreases until $t = \frac{m+1}{2}$, and then increases, a stochastic process whose realizations are hardly similar to those of a stable system. Note also that the choice of the dimension m of g has an important influence on the prior shape as the minimum value is reached for $t = \frac{m+1}{2}$.

Coming back to the original Hankel prior, we have exploited the prior \tilde{p}_H to generate a Markov chain converging to $p_H(g) \propto \exp\left(-\frac{\sum_i \sigma_i(H)}{2\lambda}\right)$ setting $2\lambda = 1, g \in \mathbb{R}^{99}, H \in \mathbb{R}^{50 \times 50}$. In particular, results displayed in Fig. 1 and discussed in subsection 4.1, have been obtained generating a chain of length 1e6 by a random walk Metropolis scheme. More specifically, when the state chain is g^k , the pro-

posed sample is generated as $h^{k+1} = g^k + \alpha s^k$, where all the s^k are i.i.d. random vectors drawn from \tilde{p}_H while the scale factor α is set to 0.3. Then, with probability $\min\left(1, \frac{p_H(h^{k+1})}{p_H(g^k)}\right)$ the Markov chain state g^{k+1} is set to h^{k+1} , otherwise $g^{k+1} = g^k$. The left panel of Fig. 1 also shows the standard deviations of the impulse response coefficients g_i under the approximated prior \tilde{p}_H characterized by (9.1) (dashed line, scaled so that the variances of g_1 under \tilde{p}_H and p_H are equal). The similarity between \tilde{p}_H and p_H confirms that the bad performance of the nuclear norm regularizer is due to a prior which shares the same flaws pointed out by (9.2).

9.2 Proof of Proposition 5.1

We start discussing the functional nature of the problem (5.7), then obtaining its Bayesian interpretation. Just for a while, it is useful to reason in continuous-time and introduce the following Sobolev space [1] of functions $h : [0, 1] \rightarrow \mathbb{R}$

$$\mathcal{S} = \left\{ h : h(0) = 0, h \text{ abs. cont.}, \int_0^1 \dot{h}^2(t) dt < \infty \right\}$$

with (squared) norm $\|h\|_{\mathcal{S}}^2 = \int_0^1 \dot{h}^2(t) dt$. It is well known that this is a RKHS with reproducing kernel which coincides with the covariance of the Brownian motion and is given, for $t, s \geq 0$ by

$$S(s, t) = \min(s, t) = 2 \sum_{j=1}^{\infty} \zeta_j \sin\left(\frac{t}{\sqrt{\zeta_j}}\right) \sin\left(\frac{s}{\sqrt{\zeta_j}}\right) \quad (9.3)$$

where $\zeta_j = \frac{1}{(j\pi - \pi/2)^2}$. Combining (9.3) and Theorem 4 on pag. 37 of [16], \mathcal{S} can be also expressed as

$$\mathcal{S} = \left\{ h \mid h(t) = \sum_{j=1}^{\infty} h_j \sqrt{2} \sin\left(\frac{t}{\sqrt{\zeta_j}}\right) \quad t \in [0, 1], \sum_{j=1}^{\infty} \frac{h_j^2}{\zeta_j} < \infty. \right\} \quad (9.4)$$

Now, note that $\min(\alpha^t, \alpha^s) = \alpha^{\max(t, s)}$. Then, still using (9.3) we obtain

$$\alpha^{\max(s, t)} = 2 \sum_{j=1}^{\infty} \zeta_j \sin\left(\frac{\alpha^t}{\sqrt{\zeta_j}}\right) \sin\left(\frac{\alpha^s}{\sqrt{\zeta_j}}\right), \quad (9.5)$$

which coincides with (5.4) when t and s are restricted to the set of natural numbers \mathbb{N} . Hence, the results contained on pag. 351 of [5] on the RKHSs derived by sampling kernels allow us to conclude that the RKHS induced by the stable spline kernel with domain on $\mathbb{N} \times \mathbb{N}$ is

$$\mathcal{H} = \left\{ g \mid g(t) = \sum_{j=1}^{\infty} g_j \sqrt{2} \sin\left(\frac{\alpha^t}{\sqrt{\zeta_j}}\right) \quad t \in \mathbb{N}, \sum_{j=1}^{\infty} \frac{g_j^2}{\zeta_j} < \infty. \right\} \quad (9.6)$$

$$\text{with } \|g\|^2 = \min_{\{g_j\}} \sum_{j=1}^{\infty} \frac{g_j^2}{\zeta_j} \text{ s.t. } g = \sum_{j=1}^{\infty} g_j \sqrt{2} \sin\left(\frac{t}{\sqrt{\zeta_j}}\right).$$

In view of the RKHS connection, the solution of (5.7) can be now obtained by the representer theorem for system identification. More specifically, Theorem 3 on pag. 671 of [41] and the connection between Bayes estimation of Gaussian processes reported in Sections 1.4 and 1.5 of [53] lead to (5.8) and this completes the proof.

References

- [1] R.A. Adams and J. Fournier. *Sobolev Spaces*. Academic Press, 2003.
- [2] S. Aja-Fernandez, R. Garcia, D. Tao, and X. Li. Tensors in image processing and computer vision. In Springer, editor, *Advances in Pattern Recognition*. 2009.
- [3] H. Akaike. A new look at the statistical model identification. *IEEE Trans. on Automatic Control*, AC-19:716–723, 1974.
- [4] A. Aravkin, J. Burke, A. Chiuso, and G. Pillonetto. Convex vs non-convex estimators for regression and sparse estimation: the mean squared error properties of ARD and GLasso. *Journal of Machine Learning Research*, 15:217–252, 2014.
- [5] N. Aronszajn. Theory of reproducing kernels. *Trans. of the American Mathematical Society*, 68:337–404, 1950.
- [6] L. Bottou, O. Chapelle, D. DeCoste, and J. Weston, editors. *Large Scale Kernel Machines*. MIT Press, Cambridge, MA, USA, 2007.
- [7] G. Box, G.M. Jenkins, and G. Reinsel. *Time Series Analysis: Forecasting & Control*. 3rd edition.
- [8] E.F. Camacho and C. Bordons. *Model Predictive Control*. Advanced Textbooks in Control and Signal Processing. Springer Verlag, 2004.
- [9] E.J. Cands and B. Recht. Exact matrix completion via convex optimization. *Found. Comp. Math.*, 9:717–772, 2009.
- [10] V. Chandrasekaran, B. Recht, P.A. Parrilo, and A.S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.
- [11] T. Chen, M. S. Andersen, L. Ljung, A. Chiuso, and G. Pillonetto. System identification via sparse multiple kernel-based regularization using sequential convex optimization techniques. *IEEE Transactions on Automatic Control*, 59(11):2933–2945, 2014.
- [12] T. Chen and L. Ljung. Regularized system identification using orthonormal basis functions. In *Proceedings of the 14th European Control Conference*, Linz, Austria.
- [13] T. Chen, H. Ohlsson, and L. Ljung. On the estimation of transfer functions, regularizations and Gaussian processes - revisited. *Automatica*, 48(8):1525–1535, 2012.
- [14] A. Chiuso, T. Chen, L. Ljung, and G. Pillonetto. Regularization strategies for nonparametric system identification. In *Proceedings of the 52nd Annual Conference on Decision and Control (CDC)*, 2013.
- [15] A. Chiuso, T. Chen, L. Ljung, and G. Pillonetto. On the design of multiple kernels for nonparametric linear system identification. In *Proceedings of the 53rd Annual Conference on Decision and Control (CDC)*, 2014.
- [16] F. Cucker and S. Smale. On the mathematical foundations of learning. *Bulletin of the American mathematical society*, 39:1–49, 2001.

- [17] H. Dong, X. Yan, F. Chao, and Y. Li. Predictive control model for radiant heating system based on neural network. In *2008 International Conference on Computer Science and Software Engineering*, pages 5106 – 5111, 2008.
- [18] D.L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52:1289–1306, 2006.
- [19] D.L. Donoho. For most large undetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Commun. Pure Appl. Math.*, 59:797–829, 2006.
- [20] B. Efron. The estimation of prediction error: Covariance penalties and cross-validation. *Journal of the American Statistical Association*, 99(14):619–632(14), 2004.
- [21] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, december 2001.
- [22] M. Fazel, H. Hindi, and S.P. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *American Control Conference, 2001. Proceedings of the 2001*, volume 6, pages 4734–4739 vol.6, 2001.
- [23] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33:1–22, 2010.
- [24] W.R. Gilks, S. Richardson, and D.J. Spiegelhalter. *Markov chain Monte Carlo in Practice*. London: Chapman and Hall, 1996.
- [25] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- [26] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [27] C. Grossmann, C.N. Jones, and M. Morari. System identification via nuclear norm regularization for simulated moving bed processes from incomplete data sets. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC)*, pages 4692–4697, 2009.
- [28] T. J. Hastie, R. J. Tibshirani, and J. Friedman. *The Elements of Statistical Learning. Data Mining, Inference and Prediction*. Springer, Canada, 2001.
- [29] H. Hjalmarsson, J. Welsh, and C.R. Rojas. Identification of box-jenkins models using structured ARX models and nuclear norm relaxation. In *16th IFAC Symposium on System Identification*, pages 322–327. IFAC, 2012.
- [30] W. James and C. Stein. Estimation with quadratic loss. In *Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability, Vol. I*, pages 361–379. University of California Press, 1961.
- [31] Z. Liu and L. Vandenberghe. Interior-point method for nuclear norm approximation with application to system identification. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1235–1256, 2009.
- [32] L. Ljung. *System Identification - Theory for the User*. Prentice-Hall, Upper Saddle River, N.J., 2nd edition, 1999.
- [33] D.J.C. MacKay. Bayesian interpolation. *Neural Computation*, 4:415–447, 1992.
- [34] J. S. Maritz and T. Lwin. *Empirical Bayes Method*. Chapman and Hall, 1989.
- [35] K. Mohan and M. Fazel. Reweighted nuclear norm minimization with application to system identification. In *American Control Conference (ACC)*, pages 2953–2959, 2010.
- [36] G. Pillonetto and A. Chiuso. Tuning complexity in kernel-based linear system identification: the robustness of the marginal likelihood estimator. In *Proceedings of the 13th European Control Conference (ECC), Strasbourg*, 2014.
- [37] G. Pillonetto and A. Chiuso. Tuning complexity in regularized kernel-based regression and linear system identification: the robustness of the marginal likelihood estimator. *Automatica*, 58:106–117, 2015.
- [38] G. Pillonetto, A. Chiuso, and G. De Nicolao. Regularized estimation of sums of exponentials in spaces generated by stable spline kernels. In *Proceedings of the IEEE American Cont. Conf., Baltimore, USA*, 2010.
- [39] G. Pillonetto, A. Chiuso, and G. De Nicolao. Prediction error identification of linear systems: a nonparametric Gaussian regression approach. *Automatica*, 47(2):291–305, 2011.
- [40] G. Pillonetto and G. De Nicolao. A new kernel-based approach for linear system identification. *Automatica*, 46(1):81–93, 2010.
- [41] 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*, 50(3):657–682, 2014.
- [42] G. Prando, A. Chiuso, and G. Pillonetto. Bayesian and regularization approaches to multivariable linear system identification: the role of rank penalties. In *IEEE Conference on Decision and Control (CDC 2014)*, 2014.
- [43] S. Prvara, J. Siroky, L. Ferkl, and J. Cigler. Predicting hourly building energy use: the great energy predictor shootout: overview and discussion of results. *Energy and Buildings*, 43:45–48, 2011.
- [44] C.R. Rojas, R. Toth, and H. Hjalmarsson. Sparse estimation of polynomial and rational dynamical models. *IEEE Transactions on Automatic Control*, 59(11):2962–2977, 2014.
- [45] C.R. Rojas, B. Wahlberg, and H. Hjalmarsson. A sparse estimation technique for general model structures. In *Proceedings of the European Control Conference (ECC13)*, 2013.
- [46] D. Sadigh, H. Ohlsson, S.S. Sastry, and S.A. Seshia. Robust subspace system identification via weighted nuclear norm optimization. In *Proceedings of the 19th World Congress of the International Federation of Automatic Control (IFAC)*.
- [47] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. (Adaptive Computation and Machine Learning). MIT Press, 2001.
- [48] P. Shah, B.N. Bhaskar, G. Tang, and B. Recht. Linear system identification via atomic norm regularization. In *Proceedings of the 51st Annual Conference on Decision and Control (CDC)*, pages 6265–6270, 2012.
- [49] M. Signoretto and J.A.K. Suykens. Convex estimation of cointegrated VAR models by a nuclear norm penalty. In *Proceedings of the 16th IFAC Symposium on System Identification (Sysid)*, 2012.
- [50] R.S. Smith. Frequency domain subspace identification using nuclear norm minimization and Hankel matrix realizations. *IEEE Transactions on Automatic Control*, 59(11):2886–2896, 2014.
- [51] T. Söderström and P. Stoica. *System Identification*. Prentice-Hall, 1989.
- [52] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1994.
- [53] G. Wahba. *Spline models for observational data*. SIAM, Philadelphia, 1990.
- [54] B. Wahlberg. System identification using Laguerre models. *IEEE Transactions on Automatic Control*, 36(5):551–562, 1991.
- [55] M. Yudong, F. Borrelli, B. Hancey, B. Coffey, S. S. Bengea, and P. Haves. Model predictive control for the operation of building cooling systems. In *American Control Conference*, pages 5106 – 5111, 2010.
- [56] H. Zou. The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476):1418–1429, 2006.