Available online at www.sciencedirect.com
SCIENCE DIRECT®

Computational Statistics & Data Analysis 47 (2004) 455–465

COMPUTATIONAL STATISTICS & DATA ANALYSIS

www.elsevier.com/locate/csda

On homogeneous least-squares problems and the inconsistency introduced by mis-constraining

Arie Yeredor[a,*], Bart De Moor[b]

[a] *Department of Electrical Engineering-Systems, Tel-Aviv University, P.O. Box 39040, Tel-Aviv 69978, Israel*

[b] *ESAT-SCD, Katholieke Universiteit Leuven, Kasteelpark Arenberg 10 Leuven B-3001, Belgium*

Received 1 December 2002; accepted 4 December 2003

Abstract

The term “homogeneous least-squares” refers to models of the form $\boldsymbol{Ya} \approx \mathbf{0}$, where $\boldsymbol{Y}$ is some data matrix, and $\boldsymbol{a}$ is an unknown parameter vector to be estimated. Such problems are encountered, e.g., when modeling auto-regressive (AR) processes. Naturally, in order to apply a least-squares (LS) solution to such models, the parameter vector $\boldsymbol{a}$ has to be somehow constrained in order to avoid the trivial solution $\boldsymbol{a} = \mathbf{0}$. Usually, the problem at hand leads to a “natural” constraint on $\boldsymbol{a}$. However, it will be shown that the use of some commonly applied constraints, such as a quadratic constraint, can lead to inconsistent estimates of $\boldsymbol{a}$. An explanation to this apparent discrepancy is provided, and the remedy is shown to lie with a necessary modification of the LS criterion, which is specified for the case of Gaussian model-errors. As a result, the modified LS minimization becomes a highly non-linear problem. For the case of quadratic constraints in the context of AR modeling, the resulting minimization involves the solution of an equation reminiscent of a “secular equation”. Numerically appealing solutions to this equation are discussed.

© 2003 Elsevier B.V. All rights reserved.

Keywords: Homogeneous least squares; Constraints; Inconsistency; Maximum likelihood

1. Introduction

In many estimation problems, e.g., in the context of the identification (parameter estimation) of linear systems (Söderström and Stoica, 1989), both the inputs and outputs

[*] Corresponding author. Tel.: +972-36405314; fax: +972-36407095.
E-mail addresses: arie@eng.tau.ac.il (A. Yeredor), bart.demoor@esat.kuleuven.ac.be (B. De Moor).

0167-9473/$ - see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/j.csda.2003.12.001

of a system are observed (possibly in the presence of additive noise), from which it is desired to estimate the system's parameters. The standard least-squares (LS) approach for this estimation problem is conceptually and computationally appealing. It consists of seeking the set of parameters with which the linear difference equations relating each output sample to past output and input samples are most closely satisfied (in the sense of a possibly weighted $L2$ norm of the errors vector). Unfortunately, however, the LS estimate is well-known to be biased and inconsistent whenever the past samples used in the regression equations involve noise. Therefore, for such problems LS is only used in cases of very high signal-to-noise ratio.

Extensive research has been addressed over the past two decades towards attempts to modify the LS estimate in such problems, so as to eliminate its bias and regain consistency (De Moor et al., 1994; Stoica and Söderström, 1982; Söderström and Stoica, 1983; Fernando and Nicholson, 1985; Zheng, 1988, 2002a,b; Van Pelt and Bernstein, 2001). Some of the well-known approaches, which have become nearly common practice in system identification, are, e.g., the instrumental variable (Stoica and Söderström, 1982; Söderström and Stoica, 1983) or the Koopmans–Levin (Fernando and Nicholson, 1985) methods.

There are, however, two exceptional cases (except for the trivial noiseless case) in which the LS estimate is generally unbiased and consistent:

- When the observed input is noiseless and the system has a finite impulse response (also termed a "zeros only" system). In this case the past samples involved in the regression equations are only the noiseless input samples. The LS model errors are exactly the output noise, so that for zero-mean noise the resulting LS estimate is unbiased. Moreover, if the output noise is (or can be uniquely transformed into) a sequence of independent, identically distributed random variables, the LS estimate is also consistent. Additionally, if the noise is Gaussian, then the properly weighted LS estimate coincides with the maximum-likelihood (ML) estimate.
- When the system identification problem is actually the identification of an autoregressive (AR) process, and the observed "output" (namely the process to be identified) is noiseless. The "input", or the process' "driving noise" in this case, is unobserved, which is equivalent to observations of zero, such that the "input observation noise" is actually (minus) the "driving noise". Thus, in such cases, the LS model errors are also the input noise, so that the same conditions (as mentioned above) for unbiased, consistent and ML-equivalent LS estimation prevail.

Although the ordinary LS estimate in these cases is unbiased and consistent, a problematic aspect thereof lies in the possible formulation of such LS problem as constrained homogeneous least-squares (HLS) problems. More generally, HLS problems are problems in which an observed data matrix $\boldsymbol{Y}$ can be modeled approximately by

$$\boldsymbol{Y}\boldsymbol{a} \approx \mathbf{0}, \tag{1}$$

where $\boldsymbol{a}$ is an unknown parameters vector, to be estimated from the observations. The inequality often implies the presence of some "driving noise", which can also be regarded as "modeling errors" in terms of the deviation of $\boldsymbol{Y}\boldsymbol{a}$ from $\mathbf{0}$.

A straightforward LS approach would be to estimate $\boldsymbol{a}$ as the minimizer of the norm of $\boldsymbol{Ya}$. However, to avoid the trivial minimizer $\boldsymbol{a}=\boldsymbol{0}$, $\boldsymbol{a}$ has to be properly constrained. Usually in such problems, some "natural" constraint (or set of constraints) on $\boldsymbol{a}$ is dictated by the problem at hand (De Moor et al., 1994; Ninness, 1996; Van Pelt and Bernstein, 2001; Ysebaert et al., 2001). For example, the unbiased, consistent LS estimate for the two estimation problems mentioned above would be obtained by constraining the respective element of $\boldsymbol{a}$ to be 1.

However, it turns out that in general, the solution to the constrained minimization

$$\min_{\boldsymbol{a}} \boldsymbol{a}^{\mathrm{T}}\boldsymbol{Y}^{\mathrm{T}}\boldsymbol{Ya}, \qquad \text{s.t. } \boldsymbol{f}(\boldsymbol{a})=\boldsymbol{0}, \tag{2}$$

where $\boldsymbol{f}(\boldsymbol{a})=\boldsymbol{0}$ is the set of constraints, can often lead to a biased, inconsistent, and, in a sense, useless estimate of $\boldsymbol{a}$. We shall show that in order to obtain a consistent estimate (subject to a certain pre-specified constraint), the LS criterion $\boldsymbol{a}^{\mathrm{T}}\boldsymbol{Y}^{\mathrm{T}}\boldsymbol{Ya}$ has to be modified, taking into consideration the distribution of the model errors (or "driving noise"). We shall explicitly specify the modification for the case of Gaussian errors.

The paper is structured as follows. We begin with a simple example in the next section, illustrating the problematic aspects of mis-constraining. In Section 3 we present a possible remedy in the form of a modified LS criterion, based on equivalence to ML estimation, for the identification of a first-order AR process. In Section 4 we generalize the criterion to a general-order AR process. A possible computational approach to the minimization of the proposed modified LS criterion is presented in Section 5. Conclusions and summary appear in the closing section.

2. An example

We begin by considering an example in which we illustrate the problems induced by choosing a "wrong" constraint. The problem originates from the identification of an AR process as treated, e.g., in Lemmerling and De Moor (2001) (see also Söderström and Stoica (1989) or Yeredor (2000)).

Let y_n be a first-order AR (AR(1)) process satisfying the difference equation:

$$y_n = -a_1 y_{n-1} + e_n, \qquad n=1,2,\ldots,N, \tag{3}$$

where e_n is a white Gaussian noise process with zero mean and known variance σ_e^2. It is desired to estimate a_1 from the observations $y_1, y_2, \ldots y_N$. We further assume, for convenience, that y_0 is deterministically known to be zero.

We can formulate the model equations in matrix form as $\boldsymbol{Ya}=\boldsymbol{e}$, where

$$\boldsymbol{Y} \triangleq \begin{bmatrix} y_1 & 0 \\ y_2 & y_1 \\ \vdots & \vdots \\ y_N & y_{N-1} \end{bmatrix}, \quad \boldsymbol{a} \triangleq \begin{bmatrix} a_0 \\ a_1 \end{bmatrix}, \quad \boldsymbol{e} \triangleq \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_N \end{bmatrix}. \tag{4}$$

Any estimate $\hat{\boldsymbol{a}}$ of $\boldsymbol{a}$ in which $\hat{a}_0 = a_0 = 1$, leads to an estimate $\hat{\boldsymbol{e}}$ of $\boldsymbol{e}$, via $\hat{\boldsymbol{e}} = \boldsymbol{Y}\hat{\boldsymbol{a}}$. Our goal is then to choose $\hat{\boldsymbol{a}}$ such that the norm of $\hat{\boldsymbol{e}}$ is minimized, subject to the *linear* constraint $\hat{a}_0 = 1$:

$$\min_{\hat{\boldsymbol{a}}} \hat{\boldsymbol{a}}^{\mathrm{T}} \boldsymbol{Y}^{\mathrm{T}} \boldsymbol{Y} \hat{\boldsymbol{a}}, \qquad \text{s.t. } [1\,0] \cdot \hat{\boldsymbol{a}} = 1. \tag{5}$$

Denoting $\hat{\boldsymbol{R}} \triangleq 1/N\, \boldsymbol{Y}^{\mathrm{T}}\boldsymbol{Y}$, we obtain the well-known LS solution

$$\hat{a}_1 = -\frac{\hat{\boldsymbol{R}}_{2,1}}{\hat{\boldsymbol{R}}_{2,2}}, \tag{6}$$

where $\hat{\boldsymbol{R}}_{i,j}$ denotes the (i,j)th element of $\hat{\boldsymbol{R}}$. If $|a_1| < 1$, then y_n is (asymptotically) stationary with autocorrelation satisfying

$$R[0] \triangleq E[y_n^2] = \frac{\sigma_e^2}{1 - a_1^2}, \qquad R[1] \triangleq E[y_n y_{n-1}] = -a_1 \cdot R[0]. \tag{7}$$

Moreover, we also have (asymptotically)

$$\hat{\boldsymbol{R}} \overset{N \to \infty}{\to} \begin{bmatrix} R[0] & R[1] \\ R[1] & R[0] \end{bmatrix}, \tag{8}$$

so that $\hat{a}_1 \to -R[1]/R[0] = a_1$ is a consistent estimator. Its consistency can be attributed to the fact that it is essentially the ML estimate, whose consistency is guaranteed in this problem setup.

Suppose now, that we want to use a quadratic constraint on $\hat{\boldsymbol{a}}$, and later “normalize” $\hat{a}_0$ to 1. Solving

$$\min_{\hat{\boldsymbol{a}}} \hat{\boldsymbol{a}}^{\mathrm{T}} \boldsymbol{Y}^{\mathrm{T}} \boldsymbol{Y} \hat{\boldsymbol{a}}, \qquad \text{s.t.} \quad \hat{\boldsymbol{a}}^{\mathrm{T}} \hat{\boldsymbol{a}} = 1, \tag{9}$$

reduces to an eigenvalue problem, and asymptotically, due to (8), we would eventually always get either $\hat{a}_1 = 1$ (if $R[1] < 0$), or $\hat{a}_1 = -1$ (if $R[1] > 0$), which is (almost) always inconsistent.

In the context of the original problem, the straightforward explanation is that the quadratic constraint is inappropriate (even when followed by normalization), and therefore there is no reason to expect a consistent estimator. We note in this context, that in Van Pelt and Bernstein (2001) it is shown, that it is generally possible to obtain a consistent estimate by using an alternative quadratic constraint of the form $\hat{\boldsymbol{a}}^{\mathrm{T}} \boldsymbol{N} \hat{\boldsymbol{a}} = 1$, where $\boldsymbol{N}$ is some symmetric (not necessarily positive-definite) matrix (that generally depends on the noise statistics). In fact, this is one of the proposed remedies for the LS estimator’s bias and inconsistency (as mentioned in the Introduction) in the general case. For our case it is evident, that using $N = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ is equivalent to the linear (sign-ambiguous) constraint $\hat{a}_0 = \pm 1$.

However, an interesting question is—what if the quadratic constraint were indeed part of the problem formulation—would such unreasonable estimates still be obtained?

In order to clarify this, consider an alternative formulation, in which we assume that the process y_n satisfies

$$a_0 y_n = -a_1 y_{n-1} + e_n, \quad n = 1, 2, \ldots, \tag{10}$$

(with the same characteristics of e_n as before), where it is now known that $a_0^2 + a_1^2 = 1$. It is desired to estimate a_0 and a_1 from $y_1, y_2, \ldots y_N$. Again we assume, for convenience, that y_0 is deterministically known to be zero.

Apparently, it is now legitimate to use the quadratically constrained minimization (9)—but then we would get the same highly inconsistent, nearly data-independent, totally unreasonable estimate.

Although this time the constraint is valid, the problem here lies with the objective function. As we shall show immediately, (9) is not the ML criterion, and therefore there is indeed no claim for consistency.

3. The correct (ML) criterion

As the matrix-vector product $\boldsymbol{Ya}$ is bilinear in the data $\boldsymbol{y}$ and the vector $\boldsymbol{a}$, we can also rewrite the model equations as $\boldsymbol{e} = \boldsymbol{Ya} = \boldsymbol{T}(\boldsymbol{a})\boldsymbol{y}$, where

$$\boldsymbol{T}(\boldsymbol{a}) \triangleq \begin{bmatrix} a_0 & & & \\ a_1 & a_0 & & \\ & \ddots & \ddots & \\ & & a_1 & a_0 \end{bmatrix}, \quad \boldsymbol{y} \triangleq \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}. \tag{11}$$

Note that the matrix $\boldsymbol{T}(\boldsymbol{a})$ defined above is square, due to the zero initial conditions ($y_0 = 0$) in (4). For higher-order AR processes this would generalize to assuming further zero initial conditions, namely $y_0 = y_{-1} = \cdots = y_{-p+1} = 0$ for an AR(p) process. However, often in practice the available data $\boldsymbol{y}$ are part of a stationary process, and then the “initial conditions” are not zeros, but must be treated as additional (random) unknowns. To avoid complications, it is common practice in these situations to ignore any equations involving “initial conditions” data, and then the respective first rows of $\boldsymbol{Y}$ in (4) are eliminated, resulting in a rectangular $\boldsymbol{T}(\boldsymbol{a})$. A “proper” way of incorporating non-zero initial conditions while maintaining $\boldsymbol{T}$ square can be found in Yeredor (2000).

Thus, if $\boldsymbol{e}$ is a zero-mean Gaussian vector with covariance $\sigma_e^2\boldsymbol{I}$, then $\boldsymbol{y} = \boldsymbol{T}^{-1}(\boldsymbol{a})\boldsymbol{e}$ is also a zero-mean Gaussian vector, but its covariance is $\sigma_e^2\boldsymbol{T}^{-1}(\boldsymbol{a})\boldsymbol{T}^{-\mathrm{T}}(\boldsymbol{a})$. Therefore its distribution is given by

$$f(\boldsymbol{y};\boldsymbol{a}) = \frac{1}{|2\pi\sigma_e^2\boldsymbol{T}^{-1}(\boldsymbol{a})\boldsymbol{T}^{-\mathrm{T}}(\boldsymbol{a})|^{1/2}}\,\mathrm{e}^{-\frac{1}{2\sigma_e^2}\boldsymbol{y}^{\mathrm{T}}\boldsymbol{T}^{\mathrm{T}}(\boldsymbol{a})\boldsymbol{T}(\boldsymbol{a})\boldsymbol{y}}, \tag{12}$$

the logarithm of which is given by

$$\begin{aligned} L(\boldsymbol{y};\boldsymbol{a}) &= \log f(\boldsymbol{y};\boldsymbol{a}) \\ &= c + \log|\boldsymbol{T}(\boldsymbol{a})| - \frac{1}{2\sigma_e^2}\boldsymbol{y}^{\mathrm{T}}\boldsymbol{T}^{\mathrm{T}}(\boldsymbol{a})\boldsymbol{T}(\boldsymbol{a})\boldsymbol{y} \\ &= c + \log|\boldsymbol{T}(\boldsymbol{a})| - \frac{1}{2\sigma_e^2}\boldsymbol{a}^{\mathrm{T}}\boldsymbol{Y}^{\mathrm{T}}\boldsymbol{Y}\boldsymbol{a}, \end{aligned} \tag{13}$$

where c is an irrelevant constant and $|\cdot|$ denotes the determinant. Evidently, in an AR problem, it is easy to observe from (11), that $|\boldsymbol{T}(\boldsymbol{a})| = a_0^N$, so that the maximization of the likelihood $L(\boldsymbol{y};\boldsymbol{a})$ is equivalent to the following minimization problem:

$$\min_{\hat{\boldsymbol{a}}}\{\hat{\boldsymbol{a}}^{\mathrm{T}}\boldsymbol{Y}^{\mathrm{T}}\boldsymbol{Y}\hat{\boldsymbol{a}} - N\sigma_e^2\log\hat{a}_0^2\}, \quad \text{s.t. } \hat{\boldsymbol{a}}^{\mathrm{T}}\hat{\boldsymbol{a}} = 1, \tag{14}$$

which would yield the consistent ML estimate of $\boldsymbol{a}$, in contrast to the “wrong” minimization problem of (9).

In the appendix we verify the consistency of the resulting estimate by deriving the closed-form solution to this simple, two-dimensional problem.

4. Generalization and discussion

Straightforward generalization to the general-order AR(q) ($q \geqslant 1$) process with general constraints $\boldsymbol{f}(\boldsymbol{a}) = \boldsymbol{0}$, maintains the same objective function:

$$\boxed{\min_{\hat{\boldsymbol{a}}}\{\hat{\boldsymbol{a}}^{\mathrm{T}}\boldsymbol{Y}^{\mathrm{T}}\boldsymbol{Y}\hat{\boldsymbol{a}} - N\sigma_e^2\log\hat{a}_0^2\} \quad \text{s.t. } \boldsymbol{f}(\hat{\boldsymbol{a}}) = \boldsymbol{0}.}$$

The general constraints can be any (linear or nonlinear) constraints that are justified by the model, such as (but not limited to) linear constraints reflecting known coefficients (most commonly $\hat{a}_0 = 1$) or known poles. Evidently, if (and only if) one or more of the constraints impose $\hat{a}_0 = 1$, then the second term of the objective function is zeroed out, and we obtain the classical objective function of (9). It is interesting to note, however, that no artificial constraints are actually needed (unless the available a-priori information would dictate so), because the “trivial” solution $\hat{\boldsymbol{a}} = \boldsymbol{0}$ is no longer a minimizer (due to the log term).

For the more general case of an HLS (not necessarily AR) problem, the entire matrix $\boldsymbol{T}(\boldsymbol{a})$ has to be incorporated into the LS criterion. If, in addition, the model errors are assumed to have a general (known) covariance structure $\boldsymbol{C}_e$, then the resulting minimization assumes the form

$$\boxed{\min_{\hat{\boldsymbol{a}}}\{\hat{\boldsymbol{a}}^{\mathrm{T}}\boldsymbol{Y}^{\mathrm{T}}\boldsymbol{C}_e^{-1}\boldsymbol{Y}\hat{\boldsymbol{a}} - 2\log|\boldsymbol{T}(\boldsymbol{a})|\} \quad \text{s.t. } \boldsymbol{f}(\hat{\boldsymbol{a}}) = \boldsymbol{0}.}$$

A potentially weak point of this approach, is that prior knowledge of the noise variance σ_e^2 (or the noise covariance $\boldsymbol{C}_e$) is required. In some cases the noise statistics are indeed known a-priori, e.g., through knowledge of a physical model or of technical specifications, such as a receiver's input noise-figure. In other cases it is sometimes possible to estimate the noise statistics "off-line" when the signal of interest is muted. When such means are not available, it may still be possible to employ an iterative strategy in which, following an intelligent initial guess, the noise level is re-estimated from the implied residual $\hat{e}_n$, and the parameter estimates are refined accordingly in each iteration; however, such a strategy should be applied with caution, so as to avoid a misleading feedback of error between iterations.

5. Minimization of the modified LS criterion

In this section, we discuss the minimization of the modified LS criterion (14) for the general-order AR(p) model with a quadratic constraint. Given the estimated (symmetric, positive definite) $M \times M$ correlation matrix $\hat{\boldsymbol{R}}$ and the model error variance σ_e^2, we wish to minimize (with respect to (w.r.t.) $\boldsymbol{a}$)

$$\min_{\boldsymbol{a}}\{\boldsymbol{a}^{\mathrm{T}}\hat{\boldsymbol{R}}\boldsymbol{a} - \sigma_e^2 \log(a_0^2)\} \quad \text{s.t. } \boldsymbol{a}^{\mathrm{T}}\boldsymbol{a} = 1, \tag{15}$$

where a_0 denotes the first element of $\boldsymbol{a}$.

We form the Lagrangian,

$$L(\boldsymbol{a},\mu) = \boldsymbol{a}^{\mathrm{T}}\hat{\boldsymbol{R}}\boldsymbol{a} - \sigma_e^2 \log(a_0^2) - \mu(\boldsymbol{a}^{\mathrm{T}}\boldsymbol{a} - 1), \tag{16}$$

differentiating w.r.t. $\boldsymbol{a}$ (taking advantage of the symmetry of $\hat{\boldsymbol{R}}$), and equating zero, we obtain

$$(\hat{\boldsymbol{R}} - \mu\boldsymbol{I})\boldsymbol{a} = \frac{\sigma_e^2}{a_0} \cdot \boldsymbol{i}_1, \tag{17}$$

where $\boldsymbol{I}$ denotes the $M \times M$ identity matrix and $\boldsymbol{i}_1$ denotes its first column. Naturally, differentiation of $L(\boldsymbol{a},\mu)$ w.r.t. μ further yields the constraint $\boldsymbol{a}^{\mathrm{T}}\boldsymbol{a} = 1$.

Using the eigenvalue decomposition $\hat{\boldsymbol{R}} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{\mathrm{T}}$ (with $\boldsymbol{U}$ a unitary matrix and $\boldsymbol{\Lambda}$ diagonal) and defining $\tilde{\boldsymbol{a}} \triangleq a_0 \cdot \boldsymbol{a}$, we may rewrite (17) as

$$\boldsymbol{U}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})\boldsymbol{U}^{\mathrm{T}}\tilde{\boldsymbol{a}} = \sigma_e^2\boldsymbol{i}_1, \tag{18}$$

hence

$$\tilde{\boldsymbol{a}} = \sigma_e^2\boldsymbol{U}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})^{-1}\boldsymbol{U}^{\mathrm{T}}\boldsymbol{i}_1, \tag{19}$$

or, defining $\boldsymbol{v} \triangleq \boldsymbol{U}^{\mathrm{T}}\boldsymbol{i}_1$,

$$\tilde{\boldsymbol{a}} = \sigma_e^2\boldsymbol{U}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})^{-1}\boldsymbol{v}. \tag{20}$$

We now need to address the constraint $\boldsymbol{a}^{\mathrm{T}}\boldsymbol{a} = 1$. Noting that

$$\|\tilde{\boldsymbol{a}}\|_2^2 = a_0^2 \cdot \|\boldsymbol{a}\|_2^2, \tag{21}$$

we conclude that the constraint is satisfied if and only if $\|\tilde{\boldsymbol{a}}\|_2^2 = a_0^2$. From (20) we have

$$\|\tilde{\boldsymbol{a}}\|_2^2 = \sigma_e^4\boldsymbol{v}^{\mathrm{T}}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})^{-2}\boldsymbol{v}. \tag{22}$$

On the other hand, a_0^2 is simply the first element of $\tilde{\boldsymbol{a}}$, given by

$$a_0^2 = \boldsymbol{i}_1^{\mathrm{T}}\tilde{\boldsymbol{a}} = \sigma_e^2 \boldsymbol{v}^{\mathrm{T}}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})^{-1}\boldsymbol{v}. \tag{23}$$

Our constraint can thus be expressed as

$$\sigma_e^2 \boldsymbol{v}^{\mathrm{T}}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})^{-2}\boldsymbol{v} = \boldsymbol{v}^{\mathrm{T}}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})^{-1}\boldsymbol{v}, \tag{24}$$

or

$$\sigma_e^2 \sum_{m=1}^{M} \frac{v_m^2}{(\lambda_m - \mu)^2} = \sum_{m=1}^{M} \frac{v_m^2}{\lambda_m - \mu}, \tag{25}$$

where v_m and λ_m denote the mth elements of $\boldsymbol{v}$ and the (m,m)th element of $\boldsymbol{\Lambda}$, respectively. This expression leads to a polynomial (of degree $2M-1$ at most) in μ,

$$\sum_{m=1}^{M} \left\{ v_m^2(\mu + \sigma_e^2 - \lambda_m) \prod_{n \neq m} (\mu - \lambda_n)^2 \right\} = 0, \tag{26}$$

whose rooting would yield at most $2M-1$ possible real-valued solutions in μ. Each of the candidate (real-valued) solution can be plugged into (20), yielding a candidate $\tilde{\boldsymbol{a}}$. Dividing each element of $\tilde{\boldsymbol{a}}$ by $\sqrt{\tilde{a}_0}$ would yield a candidate unit-norm solution for $\boldsymbol{a}$. Of the resulting $2M-1$ (at most) solutions, the one that yields the smallest objective-function value $\boldsymbol{a}^{\mathrm{T}}\hat{\boldsymbol{R}}\boldsymbol{a} - \sigma_e^2 \log(a_0^2)$ is to be chosen as the global minimizer.

To avoid the need for general polynomial rooting, one may observe the resemblance of the left-hand side (LHS) or right-hand side (RHS) of (25) to the form known as a "secular equation" (see, e.g. (Gu and Eisenstat, 1995a,b)). Then, with vertical asymptotes at the locations of the eigenvalues λ_m (a typical situation is illustrated in Fig. 1), the graph for the LHS (as a function of μ) would resemble "parabolic" curves between these asymptotes, while the graph for the RHS would resemble monotonic "cubic power" curves between the asymptotes (each extending from $-\infty$ at the left asymptote to $+\infty$ at the right asymptote). The solutions in that case are particularly easy to compute numerically, because they interlace with the eigenvalues, and can be found by simple bisections; additionally, denoting the smallest and largest eigenvalues by $\lambda_{\min}$ and $\lambda_{\max}$ (respectively), there are no solutions larger than $\lambda_{\max}$ (since for all $\mu > \lambda_{\max}$, the LHS is positive and the RHS is negative), and there is always at least one real-valued solution, smaller than $\lambda_{\min}$ (since when $\mu \to \lambda_{\min}$ from below, the LHS is larger than the RHS, whereas when $\mu \to -\infty$ the LHS is smaller than the RHS- and both are continuous in $(-\infty, \lambda_{\min})$, so they must intersect).

Additionally, although in general any solution of (25) is merely a stationary point of the Lagrangian, and can therefore be either a minimum, a maximum or a saddle point, it can be observed that μ values below $\lambda_{\min}$ are guaranteed to be associated with (at least local) minima. To show that, we examine the second derivative matrix (Hessian) of the Lagrangian (16) w.r.t. $\boldsymbol{a}$:

$$\boldsymbol{H} \triangleq \frac{\partial^2 L(\boldsymbol{a}, \mu)}{\partial \boldsymbol{a}^2} = \hat{\boldsymbol{R}} - \mu\boldsymbol{I} + \frac{\sigma_e^2}{a_0^2}\boldsymbol{I}_{11} = \boldsymbol{U}(\boldsymbol{\Lambda} - \mu\boldsymbol{I})\boldsymbol{U}^{\mathrm{T}} + \frac{\sigma_e^2}{a_0^2}\boldsymbol{I}_{11}, \tag{27}$$

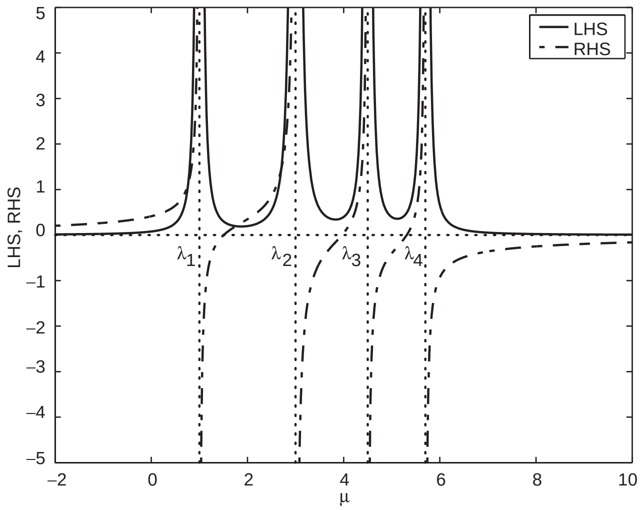

Fig. 1. A typical pattern of the LHS vs. the RHS of (25).

where $\boldsymbol{I}_{11}$ denotes an $M \times M$ matrix with all-zeros entries, except for its $(1,1)$ element, which is 1. Indeed, for $\mu < \lambda_{\min}$, the first term is positive-definite, hence (since $\boldsymbol{I}_{11}$ is positive-semidefinite) the Hessian is positive-definite and the associated solution is guaranteed to be a minimum.

Although the converse is not guaranteed in general,[1] we conjecture that the solution associated with the *smallest* μ (which is always below $\lambda_{\min}$) is always the global minimizer of criterion (15). This conjecture has been supported by extensive experimentation, yet we were unable to provide a rigorous proof. It may be interesting to note, in this context, that when $\mu = 0$, the solution to (17) coincides with the maximum entropy (ME) estimate of $\boldsymbol{a}$, and for different values of μ the associated solutions of (17) can be regarded as "modified ME" estimates of $\boldsymbol{a}$. Thus, the smallest (absolute) value of μ implies "the least modification" of the ME estimate.

6. Conclusion

We have presented and explained the observation, that the constrained HLS approach can sometimes yield useless estimates. The reason is that the LS criterion in these cases has to be supplemented with an additional term in order to yield consistent estimates in a statistical framework. Fortunately, in several common applications this term is automatically zeroed-out by the constraint; however, when using constraints that do not guarantee zero value to this term, one has to take precaution not to exclude the term from the minimization. Thus, from an algebraic point of view, monic constraints

[1] For $\mu > \lambda_{\min}$ the Hessian may still be positive-definite.

may often be the most 'manageable', albeit theoretically not the only ones possible (even not from the point of view of ML).

We specified this additional term for the case of Gaussian "driving noise" (or "model errors"), and discussed possible numerical approaches for the resulting minimization. Note, however, that for the problems illustrated, the consistent estimators of the parameters are based on consistent estimates of second-order statistics of the data, so for ergodic processes with finite second-order moments, the consistency is maintained regardless of the distribution, which may be non-Gaussian.

Acknowledgements

The research of Dr. De Moor is supported by the Research Council KUL: GOA-Mefisto 666, several Ph.D./postdoc and fellow Grants; Flemish Government:—FWO: Ph.D/postdoc grants, projects, G.0240.99 (multilinear algebra), G.0407.02 (support vector machines), G.0197.02 (power islands), G.0141.03 (Identification and cryptography), G.0491.03 (control for intensive care glycemia), G.0120.03 (QIT), research communities (ICCoS, ANMMM);—AWI: Bil. Int. Collaboration Hungary/ Poland;—IWT: PhD Grants, Soft4s (softsensors), Belgian Federal Government: DWTC (IUAP IV-02 (1996–2001) and IUAP V-22 (2002–2006)), PODO-II (CP/40: TMS and Sustainibility); EU: CAGE; ERNSI; Eureka 2063-IMPACT; Eureka 2419-FliTE; Contract Research/agreements: Data4s, Electrabel, Elia, LMS, IPCOS, VIB.

Appendix A. Closed-form solution for the AR(1) example

We shall show that the solution of the "correct" minimization problem (14) yields a consistent estimate. Dividing by N, we obtain the equivalent problem

$$\min_{\hat{\boldsymbol{a}}}\{\hat{\boldsymbol{a}}^{\mathrm{T}}\hat{\boldsymbol{R}}\hat{\boldsymbol{a}} - \sigma_e^2 \log \hat{a}_0^2\} \quad \text{s.t. } \hat{\boldsymbol{a}}^{\mathrm{T}}\hat{\boldsymbol{a}} = 1. \tag{A.1}$$

By parameterizing the constraint as $\hat{a}_0=\cos(\hat{\theta})$ and $\hat{a}_1=\sin(\hat{\theta})$ for some single parameter $\hat{\theta}$, we obtain the following equivalent *unconstrained* minimization:

$$\min_{\hat{\theta}}\{\cos^2(\hat{\theta})\hat{\boldsymbol{R}}_{1,1} + 2\cos(\hat{\theta})\sin(\hat{\theta})\hat{\boldsymbol{R}}_{1,2} + \sin^2(\hat{\theta})\hat{\boldsymbol{R}}_{2,2} - \sigma_e^2 \log(\cos^2(\hat{\theta}))\}. \tag{A.2}$$

As already mentioned, assuming that $|a_1/a_0| < 1$, y_n is (asymptotically) a stationary process, and $\hat{\boldsymbol{R}}$ tends (as $N \to \infty$) to the true $\boldsymbol{R}$ (as in (8)). Thus, with $\hat{\boldsymbol{R}}_{1,1} = \hat{\boldsymbol{R}}_{2,2} = R[0]$, and $\hat{\boldsymbol{R}}_{1,2} = R[1]$, this minimization problem reduces to

$$\min_{\hat{\theta}}\{\sin(2\hat{\theta})R[1] - \sigma_e^2 \log(\cos^2(\hat{\theta}))\}. \tag{A.3}$$

(Note that without the second term, a minimum is always obtained either at $\hat{\theta} = \pi/4$ or at $\hat{\theta} = 3\pi/4$, depending only on the sign of $R[1]$, as already observed for the "wrong" minimization problem earlier). Differentiating and equating to zero we obtain that the minimizing $\hat{\theta}$ should satisfy

$$\frac{\tan(\hat{\theta})}{\cos(2\hat{\theta})} = -\frac{R[1]}{\sigma_e^2}. \tag{A.4}$$

Indeed, for the stationary process

$$y_n = -\frac{\sin(\theta)}{\cos(\theta)}\, y_{n-1} + \frac{1}{\cos(\theta)}\, e_n, \tag{A.5}$$

we have

$$R[0] = \frac{\sigma_e^2}{\cos^2(\theta)}\,\frac{1}{1-\tan^2(\theta)} = \frac{\sigma_e^2}{\cos(2\theta)} \tag{A.6}$$

and

$$R[1] = -\tan(\theta)R[0] = -\sigma_e^2\,\frac{\tan(\theta)}{\cos(2\theta)}, \tag{A.7}$$

from which the consistency of $\hat{\theta}$ is evident in view of (A.4).

References

De Moor, B., Gevers, M., Goodwin, G., 1994. L2-Overbiased, L2-underbiased and L2-unbiased estimation of transfer functions. Automatica 30 (5), 893–898.

Fernando, K.V., Nicholson, H., 1985. Identification of linear systems with input and output noise: the Koopmans–Levin method. IEE Proc. D 132, 30–36.

Gu, M., Eisenstat, S.C., 1995a. A Divide-and-conquer algorithm for the bidiagonal SVD. SIAM J. Matrix Anal. Appl. 16, 79–92.

Gu, M., Eisenstat, S.C., 1995b. A Divide-and-conquer algorithm for the symmetric tridiagonal eigenvalue problem. SIAM J. Matrix Anal. Appl. 16, 172–191.

Lemmerling, P., De Moor, B., 2001. Misfit versus latency. Automatica 37, 2057–2067.

Ninness, B., 1996. Integral constraints on the accuracy of least-squares estimation. Automatica 32 (3), 391–397.

Söderström, T., Stoica, P., 1983. Instrumental Variable Methods for System Identification. Springer, New York.

Söderström, T., Stoica, P., 1989. System Identification. Prentice-Hall, Englewood Cliffs, NJ.

Stoica, P., Söderström, T., 1982. Bias correction in least-squares identification. Internal J. Control 35, 449–457.

Van Pelt, T.H., Bernstein, D.S., 2001. Quadratically constrained least-squares identification. Proceedings of the American Control Conference, Arlington, VA, USA, 25–27 June 2001, pp. 3684–3689.

Yeredor, A., 2000. The joint MAP-ML criterion and its relation to ML and to extended least-squares. IEEE Trans. Signal Process. 48 (12), 3484–3492.

Ysebaert, G., Van Acker, K., Moonen, M., De Moor, B., 2001. Constraints in channel shortening equalizer design for DMT-based systems. Internal Report 01–27, ESAT-SISTA, K.U. Leuven, Leuven, Belgium, 2001.

Zheng, W.-X., 1988. Consistent estimation of parameters of stochastic feedback systems in the presence of correlated disturbances. Adv. Modelling Simulation 14, 15–26.

Zheng, W.-X., 2002a. Noisy input–output system identification—the Koopmans–Levin method revisited. Proceedings of the Conference on Decision and Control, Las Vegas, NV, December 2002, pp. 636–637.

Zheng, W.-X., 2002b. A bias correction method for identification of linear dynamic errors-in-variables models. IEEE Trans. Automat. Control 47, 1142–1147.