



Brief paper

Recursive nonlinear-system identification using latent variables[☆]Per Mattsson^{a,*}, Dave Zachariah^b, Petre Stoica^b^a Department of Electronics, Mathematics and Natural Sciences, University of Gävle, Gävle, Sweden^b Department of Information Technology, Uppsala University, Uppsala, Sweden

ARTICLE INFO

Article history:

Received 28 June 2016

Received in revised form 29 August 2017

Accepted 19 December 2017

Keywords:

Nonlinear systems

Multi-input/multi-output systems

System identification

ABSTRACT

In this paper we develop a method for learning nonlinear system models with multiple outputs and inputs. We begin by modeling the errors of a nominal predictor of the system using a latent variable framework. Then using the maximum likelihood principle we derive a criterion for learning the model. The resulting optimization problem is tackled using a majorization–minimization approach. Finally, we develop a convex majorization technique and show that it enables a recursive identification method. The method learns parsimonious predictive models and is tested on both synthetic and real nonlinear systems.

© 2018 Elsevier Ltd. All rights reserved.

1. Introduction

In this paper we consider the problem of learning a nonlinear dynamical system model with multiple outputs $\mathbf{y}(t)$ and multiple inputs $\mathbf{u}(t)$ (when they exist). Generally this identification problem can be tackled using different model structures, with the class of linear models being arguably the most well studied in engineering, statistics and econometrics (Barber, 2012; Bishop, 2006; Box, Jenkins, Reinsel, & Ljung, 2015; Ljung, 1998; Söderström & Stoica, 1988).

Linear models are often used even when the system is known to be nonlinear (Enqvist, 2005; Schoukens, Vaes, & Pintelon, 2016). However certain nonlinearities, such as saturations, cannot always be neglected. In such cases using block-oriented models is a popular approach to capture static nonlinearities (Giri & Bai, 2010). Recently, such models have been given semiparametric formulations and identified using machine learning methods, cf. Pillonetto (2013) and Pillonetto, Dinuzzo, Chen, De Nicolao, and Ljung (2014). To model nonlinear dynamics a common approach is to use NARMAX models (Billings, 2013; Sjöberg et al., 1995).

In this paper we are interested in recursive identification methods (Ljung & Söderström, 1983). In cases where the model structure is linear in the parameters, recursive least-squares can be applied. For certain models with nonlinear parameters, the

extended recursive least-squares has been used (Chen, 2004). Another popular approach is the recursive prediction error method which has been developed, e.g., for Wiener models, Hammerstein models, and polynomial state-space models (Mattsson & Wigren, 2016; Tayamon, Wigren, & Schoukens, 2012; Wigren, 1993).

Nonparametric models are often based on weighted sums of the observed data (Roll, Nazin, & Ljung, 2005). The weights vary for each predicted output and the number of weights increases with each observed datapoint. The weights are typically obtained in a batch manner; in Bai and Liu (2007) and Bijl, van Wingerden, Schön, and Verhaegen (2015) they are computed recursively but must be recomputed for each new prediction of the output.

For many nonlinear systems, however, linear models work well as an initial approximation. The strategy in Paduart et al. (2010) exploits this fact by first finding the best linear approximation using a frequency domain approach. Then, starting from this approximation, a nonlinear polynomial state-space model is fitted by solving a nonconvex problem. This two-step method cannot be readily implemented recursively and it requires input signals with appropriate frequency domain properties.

In this paper, we start from a nominal model structure. This class can be based on insights about the system, e.g. that linear model structures can approximate a system around an operating point. Given a record of past outputs, $\mathbf{y}(t)$ and inputs $\mathbf{u}(t)$, that is,

$$\mathcal{D}_t \triangleq \{ (\mathbf{y}(1), \mathbf{u}(1)), \dots, (\mathbf{y}(t), \mathbf{u}(t)) \},$$

a nominal model yields a predicted output $\mathbf{y}_0(t+1)$ which differs from the output $\mathbf{y}(t+1)$. The resulting prediction error is denoted $\boldsymbol{\varepsilon}(t+1)$ (Ljung, 1999). By characterizing the nominal prediction errors in a data-driven manner, we aim to develop a refined predictor model of the system. Thus we integrate classic and data-driven system modeling approaches in a natural way.

[☆] This work has been partly supported by the Swedish Research Council (VR) under contracts 621–2014–5874 and 2016–06079. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Brett Ninness under the direction of Editor Torsten Söderström.

* Corresponding author.

E-mail addresses: per.mattsson@hig.se (P. Mattsson), dave.zachariah@it.uu.se (D. Zachariah), ps@it.uu.se (P. Stoica).

The general model class and problem formulation are introduced in Section 2. Then in Section 3 we apply the principle of maximum likelihood to derive a statistically motivated learning criterion. In Section 4 this nonconvex criterion is minimized using a majorization–minimization approach that gives rise to a convex user-parameter free method. We derive a computationally efficient recursive algorithm for solving the convex problem, which can be applied to large data sets as well as online learning scenarios. In Section 5, we evaluate the proposed method using both synthetic and real data examples.

In a nutshell, the contribution of the paper is a modeling approach and identification method for nonlinear multiple input–multiple output systems that:

- explicitly separates modeling based on application-specific insights from general data-driven modeling,
- circumvents the choice of regularization parameters and initialization points,
- learns parsimonious predictor models,
- admits a computationally efficient implementation.

Notation: $\mathbf{E}_{i,j}$ denotes the ij th standard basis matrix. \otimes and \odot denote the Kronecker and Hadamard products, respectively. $\text{vec}(\cdot)$ is the vectorization operation. $\|\mathbf{x}\|_2$, $\|\mathbf{x}\|_1$ and $\|\mathbf{X}\|_{\mathbf{W}} = \sqrt{\text{tr}(\mathbf{X}^T \mathbf{W} \mathbf{X})}$, where $\mathbf{W} \succ \mathbf{0}$, denote ℓ_2 -, ℓ_1 - and weighted norms, respectively. The Moore–Penrose pseudoinverse of \mathbf{X} is denoted \mathbf{X}^\dagger .

Remark 1. An implementation of the proposed method is available at <https://github.com/magni84/lava>.

2. Problem formulation

Given \mathcal{D}_{t-1} , the n_y -dimensional output of a system can always be written as

$$\mathbf{y}(t) = \mathbf{y}_0(t) + \boldsymbol{\varepsilon}(t), \quad (1)$$

where $\mathbf{y}_0(t)$ is any one-step-ahead predictor based on a nominal model. Here we consider nominal models on the form

$$\mathbf{y}_0(t) = \Theta \boldsymbol{\varphi}(t), \quad (2)$$

where the $p \times 1$ vector $\boldsymbol{\varphi}(t)$ is a given function of \mathcal{D}_{t-1} and Θ denotes the unknown parameters.

Remark 2. A typical example of $\boldsymbol{\varphi}(t)$ is

$$\boldsymbol{\varphi}(t) = [\mathbf{y}^T(t-1) \cdots \mathbf{y}^T(t-n_a) \mathbf{u}^T(t-1) \cdots \mathbf{u}^T(t-n_b) \mathbf{1}]^T, \quad (3)$$

in which case the nominal predictor is linear in the data and therefore captures the linear system dynamics. Nonlinearities can be incorporated if such are known about the system, in which case $\boldsymbol{\varphi}(t)$ will be nonlinear in the data.

The popular ARX model structure, for instance, can be cast into the framework (1) and (2) by assuming that the nominal prediction error $\boldsymbol{\varepsilon}(t)$ is a white noise process (Ljung, 1998; Söderström & Stoica, 1988). For certain systems, (2) may accurately describe the dynamics of the system around its operation point and consequently the white noise assumption on $\boldsymbol{\varepsilon}(t)$ may be a reasonable approximation. However, this ceases to be the case even for systems with weak nonlinearities, cf. Enqvist (2005).

Next, we develop a data-driven model of the prediction errors $\boldsymbol{\varepsilon}(t)$ in (1), conditioned on the past data \mathcal{D}_{t-1} . Specifically, we assume the conditional model

$$\boldsymbol{\varepsilon}(t) | \mathcal{D}_{t-1} \sim \mathcal{N}(\mathbf{Z}\boldsymbol{\gamma}(t), \Sigma), \quad (4)$$

where \mathbf{Z} is an $n_y \times q$ matrix of unknown latent variables, Σ is an unknown covariance matrix, and the $q \times 1$ vector $\boldsymbol{\gamma}(t)$ is any given function of \mathcal{D}_{t-1} . This is a fairly general model structure that can capture correlated data-dependent nominal prediction errors.

Note that when $\mathbf{Z} \equiv \mathbf{0}$, the prediction errors are temporally white and the nominal model (2) captures all relevant system dynamics. The latent variable is modeled as random here. Before data collection, we assume \mathbf{Z} to have mean $\mathbf{0}$ as we have no reason to depart from the nominal model assumptions until after observing data. Using a Gaussian distribution, we thus get

$$\text{vec}(\mathbf{Z}) \sim \mathcal{N}(\mathbf{0}, \mathbf{D}), \quad (5)$$

where \mathbf{D} is an unknown covariance matrix.

Our goal here is to identify a refined predictor of the form

$$\hat{\mathbf{y}}(t) = \underbrace{\hat{\Theta}\boldsymbol{\varphi}(t)}_{\hat{\mathbf{y}}_0(t)} + \underbrace{\hat{\mathbf{Z}}\boldsymbol{\gamma}(t)}_{\hat{\boldsymbol{\varepsilon}}(t)}, \quad (6)$$

from a data set \mathcal{D}_{t-1} , by maximizing the likelihood function. The first term is an estimate of the nominal predictor model while the second term tries to capture structure in the data that is not taken into account by the nominal model. Note that when $\hat{\mathbf{Z}}$ is sparse we obtain a parsimonious predictor model.

Remark 3. The model (1)–(4) implies that we can write the output in the equivalent form

$$\mathbf{y}(t) = \Theta \boldsymbol{\varphi}(t) + \mathbf{Z}\boldsymbol{\gamma}(t) + \mathbf{v}(t),$$

where $\mathbf{v}(t)$ is a white process with covariance Σ . In order to formulate a flexible data-driven error model (4), we overparametrize it using a high-dimensional $\boldsymbol{\gamma}(t)$. In this case, regularization of \mathbf{Z} is desirable, which is achieved by (5). Note that \mathbf{D} and Σ are both unknown. Estimating these covariance matrices corresponds to using a data-adaptive regularization, as we show in subsequent sections.

Remark 4. The nonlinear function $\boldsymbol{\gamma}(t)$ of \mathcal{D}_{t-1} can be seen as a basis expansion which is chosen to yield a flexible model structure of the errors. In the examples below we will use the Laplace operator basis functions (Solin & Särkkä, 2014). Other possible choices include the polynomial basis functions, Fourier basis functions, wavelets, etc. Ljung (1998), Sjöberg et al. (1995) and Van den Hof and Ninness (2005).

Remark 5. In (6), $\hat{\mathbf{y}}(t)$ is a one-step-ahead predictor. However, the framework can be readily applied to k -step-ahead prediction where $\boldsymbol{\varphi}(t)$ and $\boldsymbol{\gamma}(t)$ depend on $\mathbf{y}(1), \dots, \mathbf{y}(t-k)$.

3. Latent variable framework

Given a record of N data samples, \mathcal{D}_N , our goal is to estimate Θ and \mathbf{Z} to form the refined predictor (6). In Section 3.1, we employ the maximum likelihood approach based on the likelihood function $p(\mathbf{Y}|\Theta, \mathbf{D}, \Sigma)$, which requires the estimation of nuisance parameters \mathbf{D} and Σ . For notational simplicity, we write the parameters as $\Omega = \{\Theta, \mathbf{D}, \Sigma\}$ and in Section 3.2 we show how an estimator of \mathbf{Z} is obtained as a function of Ω and \mathcal{D}_N .

3.1. Parameter estimation

We write the output samples in matrix form as

$$\mathbf{Y} = [\mathbf{y}(1) \cdots \mathbf{y}(N)] \in \mathbb{R}^{n_y \times N}.$$

In order to obtain maximum likelihood estimates of Ω , we first derive the likelihood function by marginalizing out the latent variables from the data distribution:

$$p(\mathbf{Y}|\Omega) = \int p(\mathbf{Y}|\Omega, \mathbf{Z})p(\mathbf{Z})d\mathbf{Z}, \quad (7)$$

where the data distribution $p(\mathbf{Y}|\Omega, \mathbf{Z})$ and $p(\mathbf{Z})$ are given by (4) and (5), respectively.

To obtain a closed-form expression of (7), we begin by constructing the regression matrices

$$\Phi = [\varphi(1) \ \cdots \ \varphi(N)] \in \mathbb{R}^{p \times N},$$

$$\Gamma = [\gamma(1) \ \cdots \ \gamma(N)] \in \mathbb{R}^{q \times N}.$$

It is shown in Appendix A that (7) can be written as

$$p(\mathbf{Y}|\Omega) = \frac{1}{(2\pi)^{Nn_y}|\mathbf{R}|} \exp\left(-\frac{1}{2}\|\mathbf{y} - \mathbf{F}\theta\|_{\mathbf{R}^{-1}}^2\right), \quad (8)$$

where

$$\mathbf{y} = \text{vec}(\mathbf{Y}), \quad \theta = \text{vec}(\Theta), \quad \mathbf{z} = \text{vec}(\mathbf{Z}), \quad (9)$$

are vectorized variables, and

$$\mathbf{F} = \Phi^\top \otimes \mathbf{I}_{n_y}, \quad \mathbf{G} = \Gamma^\top \otimes \mathbf{I}_{n_y}, \quad (10)$$

$$\mathbf{R} \triangleq \mathbf{G}\mathbf{D}\mathbf{G}^\top + \mathbf{I}_N \otimes \Sigma. \quad (11)$$

Note that (8) is not a Gaussian distribution in general, since \mathbf{R} may be a function of \mathbf{Y} . It follows that maximizing (8) is equivalent to solving

$$\min_{\Omega} V(\Omega), \quad (12)$$

where

$$V(\Omega) \triangleq \|\mathbf{y} - \mathbf{F}\theta\|_{\mathbf{R}^{-1}}^2 + \ln|\mathbf{R}| \quad (13)$$

and $\mathbf{y} - \mathbf{F}\theta = \text{vec}(\mathbf{Y} - \Theta\Phi) = \text{vec}([\varepsilon(1) \ \cdots \ \varepsilon(N)])$ is nothing but the vector of nominal prediction errors.

3.2. Latent variable estimation

Next, we turn to the latent variable \mathbf{Z} which is used to model the nominal prediction error $\varepsilon(t)$ in (4). As we show in Appendix A, the conditional distribution of \mathbf{Z} is Gaussian and can be written as

$$p(\mathbf{Z}|\Omega, \mathbf{Y}) = \frac{1}{\sqrt{(2\pi)^{n_y q}|\Sigma_z|}} \exp\left(-\frac{1}{2}\|\mathbf{z} - \zeta\|_{\Sigma_z^{-1}}^2\right), \quad (14)$$

with conditional mean

$$\zeta = \mathbf{D}\mathbf{G}^\top \mathbf{R}^{-1}(\mathbf{y} - \mathbf{F}\theta), \quad (15)$$

and covariance matrix

$$\Sigma_z = (\mathbf{D}^{-1} + \mathbf{G}^\top(\mathbf{I}_N \otimes \Sigma^{-1})\mathbf{G})^{-1}. \quad (16)$$

An estimate $\hat{\mathbf{Z}}$ is then given by evaluating the conditional (vectorized) mean (15) at the optimal estimate $\hat{\Omega}$ obtained via (12).

4. Majorization–minimization approach

The quantities in the refined predictor (6) are readily obtained from the solution of (12). In general, $V(\Omega)$ may have local minima and (12) must be tackled using computationally efficient iterative methods to find the optimum. The obtained estimates will then depend on the choice of initial point Ω_0 . Such methods include the majorization–minimization approach (Hunter & Lange, 2004; Wu & Lange, 2010), which in turn include Expectation Maximization (Dempster, Laird, & Rubin, 1977) as a special case.

The majorization–minimization approach is based on finding a majorizing function $V'(\Omega|\tilde{\Omega})$ with the following properties:

$$V(\Omega) \leq V'(\Omega|\tilde{\Omega}), \quad \forall \Omega \quad (17)$$

with equality when $\Omega = \tilde{\Omega}$. The key is to find a majorizing function that is simpler to minimize than $V(\Omega)$. Let Ω_{k+1} denote the minimizer of $V'(\Omega|\Omega_k)$. Then

$$V(\Omega_{k+1}) \leq V'(\Omega_{k+1}|\Omega_k) \leq V'(\Omega_k|\Omega_k) = V(\Omega_k). \quad (18)$$

This property leads directly to an iterative scheme that decreases $V(\Omega)$ monotonically, starting from an initial estimate Ω_0 .

Given the overparameterized error model (4), it is natural to initialize at points in the parameter space which correspond to the nominal predictor model structure (2). That is,

$$\Omega_0 = \{\Theta_0, \mathbf{0}, \Sigma_0\}, \quad \text{where } \Sigma_0 > \mathbf{0}, \quad (19)$$

at which $\hat{\mathbf{Z}} \equiv \mathbf{0}$.

4.1. Convex majorization

For a parsimonious parameterization and computationally advantageous formulation, we consider a diagonal structure of the covariance matrices in (4), i.e., we let

$$\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_{n_y}), \quad (20)$$

and we let $\mathbf{D}_i = \text{diag}(d_{i,1}, \dots, d_{i,q})$ denote the covariance matrix corresponding to the i th row of \mathbf{Z} , so that

$$\mathbf{D} = \sum_{i=1}^{n_y} \mathbf{D}_i \otimes \mathbf{E}_{i,i}. \quad (21)$$

We begin by majorizing (13) via linearization of the concave term $\ln|\mathbf{R}|$:

$$\ln|\mathbf{R}| \leq \ln|\tilde{\mathbf{R}}| - \text{tr}\{\tilde{\mathbf{R}}^{-1}(\mathbf{I}_N \otimes \tilde{\Sigma})\} - \text{tr}\{\mathbf{G}^\top \tilde{\mathbf{R}}^{-1} \mathbf{G}\tilde{\mathbf{D}}\} + \text{tr}\{\tilde{\mathbf{R}}^{-1}(\mathbf{I}_N \otimes \Sigma)\} + \text{tr}\{\mathbf{G}^\top \tilde{\mathbf{R}}^{-1} \mathbf{G}\mathbf{D}\}, \quad (22)$$

where $\tilde{\mathbf{D}}$ and $\tilde{\Sigma}$ are arbitrary diagonal covariance matrices and $\tilde{\mathbf{R}}$ is obtained by inserting $\tilde{\mathbf{D}}$ and $\tilde{\Sigma}$ into (11). The right-hand side of the inequality above is a majorizing tangent plane to $\ln|\mathbf{R}|$. The use of (22) yields a convex majorizing function of $V(\Omega)$ in (12):

$$V'(\Omega|\tilde{\Omega}) = \|\mathbf{y} - \mathbf{F}\theta\|_{\mathbf{R}^{-1}}^2 + \text{tr}\{\tilde{\mathbf{R}}^{-1}(\mathbf{I}_N \otimes \Sigma)\} + \text{tr}\{\mathbf{G}^\top \tilde{\mathbf{R}}^{-1} \mathbf{G}\mathbf{D}\} + \tilde{K}, \quad (23)$$

where $\tilde{K} = \ln|\tilde{\mathbf{R}}| - \text{tr}\{\tilde{\mathbf{R}}^{-1}(\mathbf{I}_N \otimes \tilde{\Sigma})\} - \text{tr}\{\mathbf{G}^\top \tilde{\mathbf{R}}^{-1} \mathbf{G}\tilde{\mathbf{D}}\}$ is a constant. To derive a computationally efficient algorithm for minimizing (23), the following theorem is useful:

Theorem 6. The majorizing function (23) can also be written as

$$V'(\Omega|\tilde{\Omega}) = \min_{\mathbf{Z}} V'(\Omega|\mathbf{Z}, \tilde{\Omega}) \quad (24)$$

where

$$V'(\Omega|\mathbf{Z}, \tilde{\Omega}) = \|\mathbf{Y} - \Theta\Phi - \mathbf{Z}\Gamma\|_{\Sigma^{-1}}^2 + \|\text{vec}(\mathbf{Z})\|_{\mathbf{D}^{-1}}^2 + \text{tr}\{\tilde{\mathbf{R}}^{-1}(\mathbf{I}_N \otimes \Sigma)\} + \text{tr}\{\mathbf{G}^\top \tilde{\mathbf{R}}^{-1} \mathbf{G}\mathbf{D}\} + \tilde{K}. \quad (25)$$

The minimizing \mathbf{Z} is given by the conditional mean (15).

Proof. The problem in (24) has a minimizing \mathbf{Z} which, after vectorization, equals ζ in (15). Inserting the minimizing \mathbf{Z} into (25) and using (A.3) yields (23).

Remark 7. The augmented form in (24), enables us to solve for the nuisance parameters \mathbf{D} and Σ in closed-form and also yields the conditional mean $\hat{\mathbf{Z}}$ as a by-product.

To prepare for the minimization of the function (25) we write the matrix quantities using variables that denote the i th rows of the following matrices:

$$\mathbf{Y} = \begin{bmatrix} \vdots \\ \mathbf{y}_i^\top \\ \vdots \end{bmatrix}, \quad \Theta = \begin{bmatrix} \vdots \\ \boldsymbol{\theta}_i^\top \\ \vdots \end{bmatrix}, \quad \mathbf{Z} = \begin{bmatrix} \vdots \\ \mathbf{z}_i^\top \\ \vdots \end{bmatrix}, \quad \Gamma = \begin{bmatrix} \vdots \\ \boldsymbol{\gamma}_i^\top \\ \vdots \end{bmatrix}.$$

Theorem 8. After concentrating out the nuisance parameters, the minimizing arguments Θ and \mathbf{Z} of the function (25) are obtained by solving the following convex problem:

$$\min_{\Theta, \mathbf{Z}} \sum_{i=1}^{n_y} (\|\mathbf{y}_i - \Phi^\top \boldsymbol{\theta}_i - \Gamma^\top \mathbf{z}_i\|_2 + \|\mathbf{w}_i \odot \mathbf{z}_i\|_1) \quad (26)$$

where

$$\mathbf{w}_i = [w_{i,1} \quad \cdots \quad w_{i,q}]^\top, \quad w_{i,j} = \sqrt{\frac{\boldsymbol{\gamma}_j^\top \tilde{\mathbf{R}}_i^{-1} \boldsymbol{\gamma}_j}{\text{tr}\{\tilde{\mathbf{R}}_i^{-1}\}}} \quad (27)$$

$$\tilde{\mathbf{R}}_i = \Gamma \tilde{\mathbf{D}}_i \Gamma^\top + \tilde{\sigma}_i \mathbf{I}_N. \quad (28)$$

The closed-form expression for the minimizing nuisance parameters (20) and (21) are given by

$$\hat{\sigma}_i = \frac{\|\mathbf{y}_i - \Phi^\top \boldsymbol{\theta}_i - \Gamma^\top \mathbf{z}_i\|_2}{\sqrt{\text{tr}\{\tilde{\mathbf{R}}_i^{-1}\}}}, \quad \hat{d}_{i,j} = \frac{|z_{i,j}|}{\sqrt{\boldsymbol{\gamma}_j^\top \tilde{\mathbf{R}}_i^{-1} \boldsymbol{\gamma}_j}}. \quad (29)$$

Proof. See Appendix B.

Remark 9. Problem (26) contains a data-adaptive regularizing term which typically leads to parsimonious estimates of \mathbf{Z} , cf. Stoica, Zachariah, and Li (2014).

Remark 10. Majorizing at a nominal predictor model, i.e. $\tilde{\Omega} = \{\hat{\Theta}, \mathbf{0}, \tilde{\Sigma}\}$ as discussed above, yields $\mathbf{R}_i = \tilde{\sigma}_i \mathbf{I}_N$ and the weights are given by

$$w_{i,j} = \frac{\|\boldsymbol{\gamma}_j\|_2}{\sqrt{N}}. \quad (30)$$

Then problem (26) and consequently the minimization of (24) becomes invariant with respect to $\tilde{\Sigma}$.

The iterative scheme (18) is executed by initializing at $\tilde{\Omega} := \Omega_0$ and solving (26). The procedure is then repeated by updating the majorization point using the new estimate $\tilde{\Omega} := \hat{\Omega}$. It follows that the estimates will converge to a local minima of (13). The following theorem establishes that the local minima found, and thus the resulting predictor (6), is invariant to Ω_0 in the form (19).

Theorem 11. All initial points Ω_0 in the form (19) result in the same sequence of minimizers $(\hat{\Theta}_k, \hat{\mathbf{Z}}_k)$ of (26), for all $k > 0$. Moreover, the sequence $(\hat{\mathbf{D}}_k, \hat{\Sigma}_k)$ converges to a unique point when $k \rightarrow \infty$.

Proof. See Appendix C.

Remark 12. As a result we may initialize the scheme (18) at $\Omega_0 = \{\mathbf{0}, \mathbf{0}, \mathbf{I}_{n_y}\}$. This obviates the need for carefully selecting an initialization point, which would be needed in e.g. the Expectation Maximization algorithm.

4.2. Recursive computation

We now show that the convex problem (26) can be solved recursively, for each new sample $\mathbf{y}(t)$ and $\mathbf{u}(t)$.

4.2.1. Computing $\hat{\Theta}$

If we fix \mathbf{Z} and only solve for Θ , the solution is given by

$$\hat{\Theta} = \bar{\Theta} - \mathbf{Z}\mathbf{H}^\top \quad (31)$$

where

$$\bar{\Theta} = \mathbf{Y}\Phi^\dagger \quad \text{and} \quad \mathbf{H}^\top = \Gamma\Phi^\dagger.$$

Note that both $\bar{\Theta}$ and \mathbf{H} are independent of \mathbf{Z} , and that they can be computed for each sample t using a standard recursive least-squares (LS) algorithm:

$$\bar{\Theta}(t) = \bar{\Theta}(t-1) + (\mathbf{y}(t) - \bar{\Theta}(t-1)\boldsymbol{\varphi}(t))\boldsymbol{\varphi}^\top(t)\mathbf{P}(t) \quad (32)$$

$$\mathbf{H}(t) = \mathbf{H}(t-1) + \mathbf{P}(t)\boldsymbol{\varphi}(t)(\boldsymbol{\gamma}^\top(t) - \boldsymbol{\varphi}^\top(t)\mathbf{H}(t-1)) \quad (33)$$

$$\mathbf{P}(t) = \mathbf{P}(t-1) - \frac{\mathbf{P}(t-1)\boldsymbol{\varphi}(t)\boldsymbol{\varphi}^\top(t)\mathbf{P}(t-1)}{1 + \boldsymbol{\varphi}^\top(t)\mathbf{P}(t-1)\boldsymbol{\varphi}(t)}. \quad (34)$$

Remark 13. Natural initial values are $\bar{\Theta}(0) = \mathbf{0}$ and $\mathbf{H}(0) = \mathbf{0}$. The matrix Φ^\dagger equals $\Phi^\top(\Phi\Phi^\top)^{-1}$ when $t \geq p$ samples yield a full-rank matrix Φ . The matrix $\mathbf{P}(t)$ converges to $(\Phi\Phi^\top)^{-1}$. A common choice for the initial value of $\mathbf{P}(t)$ is $\mathbf{P}(0) = c\mathbf{I}$, where a larger constant $c > 0$ leads to a faster convergence of (34), cf. Söderström and Stoica (1988) and Stoica and Åhgren (2002).

4.2.2. Computing $\hat{\mathbf{Z}}$

Using (31), we concentrate out Θ from (26) to obtain

$$V'(\mathbf{Z}) = \sum_{i=1}^{n_y} (\|\boldsymbol{\xi}_i - (\Gamma^\top - \Phi^\top \mathbf{H}) \mathbf{z}_i\|_2 + \|\mathbf{w}_i \odot \mathbf{z}_i\|_1)$$

where

$$\boldsymbol{\xi}_i = \mathbf{y}_i - \Phi^\top \boldsymbol{\theta}_i.$$

In Appendix D it is shown how the minimum of $V'(\mathbf{Z})$ can be found via cyclic minimization with respect to the elements of \mathbf{Z} , similar to what has been done in Zachariah and Stoica (2015) in a simpler case. This iterative procedure is implemented using recursively computed quantities and produces an estimate $\hat{\mathbf{Z}}(t)$ at sample t .

4.2.3. Summary of the algorithm

The algorithm computes $\hat{\Theta}(t)$ and $\hat{\mathbf{Z}}(t)$ recursively by means of the following steps at each sample t :

- (i) Compute $\bar{\Theta}(t)$ and $\mathbf{H}(t)$, using (32)–(34).
- (ii) Compute $\hat{\mathbf{Z}}(t)$ via the cyclic minimization of $V'(\mathbf{Z})$. Cycle through all elements L times.
- (iii) Compute $\hat{\Theta}(t) = \bar{\Theta}(t) - \hat{\mathbf{Z}}(t)\mathbf{H}^\top(t)$

The estimates are initialized as $\hat{\Theta}(0) = \mathbf{0}$ and $\hat{\mathbf{Z}}(0) = \mathbf{0}$. In practice, small L works well since we cycle L times through all elements of \mathbf{Z} for each new data sample. The computational details are given in Algorithm 1 in Appendix D, which can be readily implemented e.g. in MATLAB.

5. Numerical experiments

In this section we evaluate the proposed method and compare it with two alternative identification methods.

5.1. Identification methods and experimental setup

The numerical experiments were conducted as follows. Three methods have been used: LS identification of affine ARX, NARX using wavelet networks (WAVE for short), and the latent variable method (LAVA) presented in this paper. From our numerical experiments we found that performing even only one iteration of the majorization–minimization algorithm produces good results, and doing so leads to a computationally efficient recursive implementation (which we denote LAVA-R for recursive).

For each method the function $\varphi(t)$ is taken as the linear regressor in (3). Then the dimension of $\varphi(t)$ equals $p = n_y n_a + n_u n_b + 1$. For affine ARX, the model is given by

$$\hat{\mathbf{y}}_{ARX}(t) = \bar{\Theta} \varphi(t),$$

where $\bar{\Theta}$ is estimated using recursive least squares (Söderström & Stoica, 1988). Note that in LAVA-R, $\bar{\Theta}$ is computed as a byproduct (32).

For the wavelet network, nlarx in the System Identification Toolbox for Matlab was used, with the number of nonlinear units automatically detected (Ljung, 2007).

For LAVA-R, the model is given by (6) and $\hat{\Theta}, \hat{\mathbf{Z}}$ are found by the minimization of (26) using $\hat{\Omega}_0 = \{\mathbf{0}, \mathbf{0}, \mathbf{I}_{n_y}\}$. The minimization is performed using the recursive algorithm in Section 4.2.3 with $L = 5$. The nonlinear function $\boldsymbol{\gamma}(t)$ of the data \mathcal{D}_{t-1} can be chosen to be a set of basis functions evaluated at $\varphi(t)$. Then $\mathbf{Z}\boldsymbol{\gamma}(t)$ can be seen as a truncated basis expansion of some nonlinear function. In the numerical examples, $\boldsymbol{\gamma}(t)$ uses the Laplace basis expansion due to its good approximation properties (Solin & Särkkä, 2014). Each element in the expansion is given by

$$\gamma_{k_1, \dots, k_{p-1}}(t) = \prod_{i=1}^{p-1} \frac{1}{\sqrt{\ell_i}} \sin\left(\frac{\pi k_i (\varphi_i(t) + \ell_i)}{2\ell_i}\right), \quad (35)$$

where ℓ_i are the boundaries of the inputs and outputs for each channel and $k_i = 1, \dots, M$ are the indices for each element of $\boldsymbol{\gamma}(t)$. Then the dimension of

$$\boldsymbol{\gamma}(t) = [\gamma_{1, \dots, 1}(t) \cdots \gamma_{p-1, \dots, p-1}(t)]^T$$

equals $q = M^{p-1}$, where M is a user parameter which determines the resolution of the basis.

Finally, an important part of the identification setup is the choice of input signal. For a nonlinear system it is important to excite the system both in frequency and in amplitude. For linear models a commonly used input signal is a pseudorandom binary sequence (PRBS), which is a signal that shifts between two levels in a certain fashion. One reason for using PRBS is that it has good correlation properties (Söderström & Stoica, 1988). Hence, PRBS excites the system well in frequency, but poorly in amplitude. A remedy to the poor amplitude excitation is to multiply each interval of constant signal level with a random factor that is uniformly distributed on some interval $[0, A]$, cf. Wigren (2006). Hence, if the PRBS takes the values -1 and 1 , then the resulting sequence will contain constant intervals with random amplitudes between $-A$ and A . We denote such a random amplitude sequence RS(A) where A is the maximum amplitude.

5.2. Performance metric

For the examples considered here the system does not necessarily belong to the model class, and thus there is no true parameter vector to compare with the estimated parameters. Hence, the different methods will instead be evaluated with respect to the simulated model output $\hat{\mathbf{y}}_s(t)$. For LAVA-R

$$\hat{\varphi}(t) = [\hat{\mathbf{y}}_s^T(t-1) \cdots \hat{\mathbf{y}}_s^T(t-n_a) \mathbf{u}^T(t-1) \cdots \mathbf{u}^T(t-n_b) \mathbf{1}]^T.$$

$$\hat{\mathbf{y}}_s(t) = \hat{\Theta} \hat{\varphi}(t) + \hat{\mathbf{Z}} \hat{\boldsymbol{\gamma}}(t)$$

and $\hat{\boldsymbol{\gamma}}(t)$ is computed as (35) with $\varphi(t)$ replaced by $\hat{\varphi}(t)$.

The performance can then be evaluated using the root mean squared error (RMSE) for each output channel i ,

$$\text{RMSE}_i = \sqrt{\frac{1}{T} \sum_{t=1}^T E[\|y_i(t) - \hat{y}_{s,i}(t)\|_2^2]}.$$

The expectations are computed using 100 Monte Carlo simulations on validation data.




Fig. 1. A block diagram of the system used in Example 5.3.

For the data set collected from a real system, it is not possible to evaluate the expectation in the RMSE formula. For such sets we use the fit of the data, i.e.,

$$\text{FIT}_i = 100 \left(1 - \frac{\|\mathbf{y}_i - \hat{\mathbf{y}}_{s,i}\|_2}{\|\mathbf{y}_i - \bar{\mathbf{y}}_i \mathbf{1}\|_2} \right),$$

where $\hat{\mathbf{y}}_{s,i}$ contains the simulated outputs for channel i , $\bar{\mathbf{y}}_i$ is the empirical mean of \mathbf{y}_i and $\mathbf{1}$ is a vector of ones. Hence, FIT compares the simulated output errors with those obtained using the empirical mean as the model output.

5.3. System with saturation

Consider the following state-space model,

$$x_1(t+1) = \text{sat}_2[0.9x_1(t) + 0.1u_1(t)] \quad (36)$$

$$x_2(t+1) = 0.08x_1(t) + 0.9x_2(t) + 0.6u_2(t) \quad (37)$$

$$\mathbf{y}(t) = \mathbf{x}(t) + \mathbf{e}(t) \quad (38)$$

where $\mathbf{x}(t) = [x_1(t) \ x_2(t)]^T$ and

$$\text{sat}_a(x) = \begin{cases} x & \text{if } |x| < a \\ \text{sign}(x)a & \text{if } |x| \geq a \end{cases}.$$

A block-diagram for the above system is shown in Fig. 1. The measurement noise $\mathbf{e}(t)$ was chosen as a white Gaussian process with covariance matrix $\sigma \mathbf{I}$ where $\sigma = 2.5 \cdot 10^{-3}$.

Data was collected from the system using an RS(A) input signal for several different amplitudes A . The identification was performed using $n_a = 1, n_b = 1, M = 4$, and $N = 1000$ data samples. This means that $p = 5$ and $q = 256$, and therefore there are 10 parameters in Θ and 512 in \mathbf{Z} .

Note that, for sufficiently low amplitudes $A, x_1(t)$ will be smaller than the saturation level $a = 2$ for all t , and thus the system will behave as a linear system. However, when A increases, the saturation will affect the system output more and more. The RMSE was computed for eight different amplitudes A , and the result is shown in Fig. 2. It can be seen that for small amplitudes, when the system is effectively linear, the ARX model gives a marginally better result than LAVA-R and WAVE. However, as the amplitude is increased, the nonlinear effects become more important, and LAVA-R outperforms both WAVE and ARX models.

5.4. Water tank

In this example a real cascade tank process is studied. It consists of two tanks mounted on top of each other, with free outlets. The top tank is fed with water by a pump. The input signal is the voltage applied to the pump, and the output consists of the water level in the two tanks. The setup is described in more detail in Wigren (2006). The data set consists of 2500 samples collected every five seconds. The first 1250 samples were used for identification, and the last 1250 samples for validation.




Fig. 2. The RMSE for Example (5.3) computed for different input amplitudes, using LAVA-R (solid), affine ARX (dashed) and WAVE (dash-dotted).




Fig. 3. The output in Example 5.4 (blue), plotted together with the output of the model identified by LAVA-R (red). The system was identified using the first 1250 data samples. The validation set consisted of the remaining 1250 samples. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
FIT for Example 5.4.

| | LAVA-R | WAVE | ARX |
|------------|--------|-------|-------|
| Upper tank | 91.6% | 79.2% | 84.9% |
| Lower tank | 90.8% | 76.9% | 78.6% |

The identification was performed using $n_a = 2$, $n_b = 2$ and $M = 3$. With two outputs, this results in 14 parameters in Θ and 1458 parameters in \mathbf{Z} . LAVA-R found a model with only 37 nonzero parameters in \mathbf{Z} , and the simulated output together with the measured output are shown in Fig. 3. The FIT values, computed on the validation data are shown in Table 1. It can be seen that an affine ARX model gives a good fit, but also that using LAVA-R the FIT measure can be improved significantly. In this example, WAVE did not perform very well.

5.5. Pick-and-place machine

In the final example, a real pick-and-place machine is studied. This machine is used to place electronic components on a circuit board, and is described in detail in Juloski, Heemels, and Ferrari-Trecate (2004). This system exhibits saturation, different modes,




Fig. 4. The output in Example 5.5 (blue), plotted together with the output of the model identified by LAVA-R (red). The system was identified using the first part of the data, while the validation set consisted of the remaining samples indicated after the dashed line. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2
FIT for Example 5.5.

| | LAVA-R | WAVE | ARX |
|-----|--------|-------|-------|
| FIT | 83.2% | 78.2% | 73.1% |

and other nonlinearities. The data used here are from a real physical process, and were also used in e.g. Bemporad, Garulli, Paoletti, and Vicino (2005), Juloski, Paoletti, and Roll (2006) and Ohlsson and Ljung (2013). The data set consists of a 15 s recording of the single input $u(t)$ and the vertical position of the mounting head $y(t)$. The data was sampled at 50 Hz, and the first 8 s ($N = 400$) were used for identification and the last 7 s for validation.

The identification was performed using $n_a = 2$, $n_b = 2$ and $M = 6$. For the SISO system considered here, this results in 5 parameters in Θ and 1296 parameters in \mathbf{Z} . LAVA-R found a model with 33 of the parameters in \mathbf{Z} being nonzero, the output of which is shown in Fig. 4.

The FIT values, computed on the validation data, for LAVA-R, WAVE and affine ARX are shown in Table 2. LAVA-R outperforms NARX using wavelet networks, and both are better than ARX.

6. Conclusion

We have developed a method for learning nonlinear systems with multiple outputs and inputs. We began by modeling the errors of a nominal predictor using a latent variable formulation. The nominal predictor could for instance be a linear approximation of the system but could also include known nonlinearities. A learning criterion was derived based on the principle of maximum likelihood, which obviates the tuning of regularization parameters. The criterion is then minimized using a majorization–minimization approach. Specifically, we derived a convex user-parameter free formulation, which led to a computationally efficient recursive algorithm that can be applied to large data sets as well as online learning problems.

The method introduced in this paper learns parsimonious predictor models and captures nonlinear system dynamics. This was illustrated via synthetic as well as real data examples. As shown in these examples a recursive implementation of the proposed method was capable of outperforming a batch method using a NARX model with a wavelet network.

Appendix A. Derivation of the distributions (8) and (14)

We start by computing $p(\mathbf{Y}|\Omega)$ given in (7). The function $p(\mathbf{Y}|\Omega, \mathbf{Z})$ can be found from (1)–(4) and the chain rule:

$$p(\mathbf{Y}|\Omega, \mathbf{Z}) = \prod_{t=1}^N p_e(\mathbf{y}(t) - \Theta\varphi(t) | \mathcal{D}_{t-1}, \mathbf{Z}), \quad (\text{A.1})$$

where we have neglected initial conditions (Söderström & Stoica, 1988). Since

$$p_e(\mathbf{y}(t) - \Theta\boldsymbol{\varphi}(t)|\mathcal{D}_{t-1}, \mathbf{Z}) \propto \exp\left(-\frac{1}{2}\|\mathbf{y}(t) - \Theta\boldsymbol{\varphi}(t) - \mathbf{Z}\boldsymbol{\gamma}(t)\|_{\Sigma}^2\right),$$

it follows that

$$p(\mathbf{Y}|\Omega, \mathbf{Z}) = \frac{1}{\sqrt{(2\pi)^{n_y N} |\Sigma|^N}} \exp\left(-\frac{1}{2}\|\mathbf{Y} - \Theta\Phi - \mathbf{Z}\Gamma\|_{\Sigma}^2\right). \quad (\text{A.2})$$

Using the vectorized variable in (9)–(10) we can see that

$$\text{vec}(\Theta\Phi) = \mathbf{F}\boldsymbol{\theta} \quad \text{and} \quad \text{vec}(\mathbf{Z}\Gamma) = \mathbf{G}\mathbf{z}$$

and thus,

$$\|\mathbf{Y} - \Theta\Phi - \mathbf{Z}\Gamma\|_{\Sigma}^2 = \|\mathbf{y} - \mathbf{F}\boldsymbol{\theta} - \mathbf{G}\mathbf{z}\|_{\mathbf{I}_N \otimes \Sigma}^2.$$

Next, we note that the following useful equality holds:

$$\begin{aligned} \|\mathbf{y} - \mathbf{F}\boldsymbol{\theta} - \mathbf{G}\mathbf{z}\|_{\mathbf{I}_N \otimes \Sigma}^2 + \|\mathbf{z}\|_{\mathbf{D}}^2 &= \\ \|\mathbf{y} - \mathbf{F}\boldsymbol{\theta}\|_{\mathbf{R}}^2 + \|\mathbf{z} - \boldsymbol{\zeta}\|_{\Sigma}^2 & \end{aligned} \quad (\text{A.3})$$

where \mathbf{R} is given by (11), $\boldsymbol{\zeta}$ by (15), and Σ_z by (16). To see that the equality holds, expand the norms on both sides of (A.3) and apply the matrix inversion lemma.

The sought-after distribution $p(\mathbf{Y}|\Omega)$ is given by (7). By using (A.3) it follows that

$$p(\mathbf{Y}|\Omega, \mathbf{Z})p(\mathbf{Z}) \propto \exp\left(-\frac{1}{2}\|\mathbf{y} - \mathbf{F}\boldsymbol{\theta}\|_{\mathbf{R}}^2\right) \exp\left(-\frac{1}{2}\|\mathbf{z} - \boldsymbol{\zeta}\|_{\Sigma}^2\right) \quad (\text{A.4})$$

with the normalization constant $((2\pi)^{n_y(N+q)}|\Sigma|^N|\mathbf{D}|)^{-1/2}$. Noting that

$$\int \exp\left(-\frac{1}{2}\|\mathbf{z} - \boldsymbol{\zeta}\|_{\Sigma}^2\right) d\mathbf{z} = \sqrt{(2\pi)^{n_y q} |\Sigma_z|}$$

it can be seen that

$$p(\mathbf{Y}|\Omega) = \frac{1}{\sqrt{(2\pi)^{N n_y} |\mathbf{R}|}} \exp\left(-\frac{1}{2}\|\mathbf{y} - \mathbf{F}\boldsymbol{\theta}\|_{\mathbf{R}}^2\right), \quad (\text{A.5})$$

which proves (8). To obtain an expression for $p(\mathbf{Z}|\Omega, \mathbf{Y})$ simply insert (A.4) and (A.5) into Bayes' rule to get (14).

Appendix B. Proof of Theorem 8

It follows from (21) that

$$\mathbf{R} = \sum_{i=1}^{n_y} (\mathbf{R}_i \otimes \mathbf{E}_{i,i})$$

where $\mathbf{R}_i = \Gamma^T \mathbf{D}_i \Gamma + \sigma_i \mathbf{I}_N$. Hence,

$$\mathbf{R}^{-1} = \sum_{i=1}^{n_y} (\mathbf{R}_i^{-1} \otimes \mathbf{E}_{i,i}).$$

Thus, we can rewrite (25) as (to within an additive constant):

$$V'(\Omega|\mathbf{Z}, \tilde{\Omega}) = \sum_{i=1}^{n_y} \left(\frac{1}{\sigma_i} \|\tilde{\mathbf{y}}_i\|_2^2 + \|\mathbf{z}_i\|_{\mathbf{D}_i}^2 + \sigma_i \text{tr}(\tilde{\mathbf{R}}_i^{-1}) + \text{tr}(\Gamma \tilde{\mathbf{R}}_i^{-1} \Gamma^T \mathbf{D}_i) \right) \quad (\text{B.1})$$

where $\tilde{\mathbf{y}}_i = \mathbf{y}_i - \Phi^T \boldsymbol{\theta}_i - \Gamma^T \mathbf{z}_i$.

We next derive analytical expressions for the Σ and \mathbf{D} that minimize $V'(\Omega|\mathbf{Z}, \tilde{\Omega})$. Note that

$$\frac{\partial}{\partial \sigma_i} V'(\Omega|\mathbf{Z}, \tilde{\Omega}) = -\frac{1}{\sigma_i^2} \|\tilde{\mathbf{y}}_i\|_2^2 + \text{tr}(\tilde{\mathbf{R}}_i^{-1}),$$

and setting the derivative to zero yields the estimate (29). In the same way it can be seen that the minimum of $d_{i,j}$ is attained at (29).

Inserting $\hat{\sigma}_i$ and $\hat{d}_{i,j}$ into (B.1), we see that we can find the minimizing Θ and \mathbf{Z} by minimizing

$$2 \sum_{i=1}^{n_y} \left(\sqrt{\text{tr}(\tilde{\mathbf{R}}_i^{-1})} \|\mathbf{y}_i - \Phi^T \boldsymbol{\theta}_i - \Gamma^T \mathbf{z}_i\|_2 + \sum_{j=1}^q |z_{i,j}| \sqrt{\tilde{\boldsymbol{\gamma}}_j^T \tilde{\mathbf{R}}_i^{-1} \tilde{\boldsymbol{\gamma}}_j} \right).$$

Since term i in the above sum is invariant with respect to $\boldsymbol{\theta}_k$ and \mathbf{z}_k for $k \neq i$, we can divide term i by $2\sqrt{\text{tr}(\tilde{\mathbf{R}}_i^{-1})}$, and see that minimizing the criterion above is equivalent to (26).

Appendix C. Proof of Theorem 11

Initializing (18) at $\bar{\Omega}_0 = \{\mathbf{0}, \mathbf{0}, \mathbf{I}_{n_y}\}$ and $\Omega_0 = \{\Theta_0, \mathbf{0}, \Sigma_0\}$ where $\Sigma_0 = \text{diag}(\sigma_1^{(0)}, \dots, \sigma_{n_y}^{(0)})$, produces two sequences denoted $\bar{\Omega}_k = \{\Theta_k, \bar{\mathbf{D}}_k, \bar{\Sigma}_k\}$ and $\Omega_k = \{\Theta_k, \mathbf{D}_k, \Sigma_k\}$ for $k > 0$, respectively. This results also in sequences $\bar{\mathbf{Z}}_k$ and \mathbf{Z}_k . The theorem states that:

$$\Theta_k = \bar{\Theta}_k \quad \text{and} \quad \mathbf{Z}_k = \bar{\mathbf{Z}}_k \quad (\text{C.1})$$

$$\mathbf{D}_k - \bar{\mathbf{D}}_k \rightarrow \mathbf{0} \quad \text{and} \quad \Sigma_k - \bar{\Sigma}_k \rightarrow \mathbf{0}. \quad (\text{C.2})$$

We now show the stronger result that the covariance matrices converge as

$$\mathbf{D}_i^{(k)} = c_i^{(k)} \bar{\mathbf{D}}_i^{(k)}, \quad \sigma_i^{(k)} = c_i^{(k)} \bar{\sigma}_i^{(k)}, \quad \forall k > 0, \quad (\text{C.3})$$

where $c_i^{(k)} = (\sigma_i^{(0)})^{\frac{1}{2^k}}$. Note that $c_i^{(k)} \rightarrow 1$ as $k \rightarrow \infty$. Hence (C.3) implies (C.2). We prove (C.3) and (C.1) by induction. That (C.3) and (C.1) hold for $k = 1$ follow directly from Theorem 8. Now assume that (C.3) holds for some $k \geq 1$. Let

$$\begin{aligned} \bar{\mathbf{R}}_i &= \Gamma \bar{\mathbf{D}}_i^{(k)} \Gamma^T + \bar{\sigma}_i^{(k)} \mathbf{I}_N, \\ \tilde{\mathbf{R}}_i &= \Gamma \mathbf{D}_i^{(k)} \Gamma^T + \sigma_i^{(k)} \mathbf{I}_N = c_i^{(k)} \bar{\mathbf{R}}_i, \end{aligned}$$

where the last equality follows by the assumption in (C.3). Therefore the weights used to estimate Θ_{k+1} and \mathbf{Z}_{k+1} are the same as those used to estimate $\bar{\Theta}_{k+1}$, $\bar{\mathbf{Z}}_{k+1}$:

$$w_{i,j} = \sqrt{\frac{\boldsymbol{\gamma}_j^T \tilde{\mathbf{R}}_i^{-1} \boldsymbol{\gamma}_j}{\text{tr}(\tilde{\mathbf{R}}_i^{-1})}} = \sqrt{\frac{\boldsymbol{\gamma}_j^T \bar{\mathbf{R}}_i^{-1} \boldsymbol{\gamma}_j}{\text{tr}(\bar{\mathbf{R}}_i^{-1})}},$$

so we can conclude that $\Theta_{k+1} = \bar{\Theta}_{k+1}$ and $\mathbf{Z}_{k+1} = \bar{\mathbf{Z}}_{k+1}$. The estimate \mathbf{D}_{k+1} is given by

$$d_{i,j}^{(k+1)} = \frac{|z_{i,j}|}{\sqrt{\boldsymbol{\gamma}_j^T \tilde{\mathbf{R}}_i^{-1} \boldsymbol{\gamma}_j}} = \frac{\sqrt{c_i^{(k)}} |z_{i,j}|}{\sqrt{\boldsymbol{\gamma}_j^T \bar{\mathbf{R}}_i^{-1} \boldsymbol{\gamma}_j}} = c_i^{(k+1)} \bar{d}_{i,j}^{(k+1)}$$

so $\mathbf{D}_i^{(k+1)} = c_i^{(k+1)} \bar{\mathbf{D}}_i^{(k+1)}$, and in the same way it can be seen that $\sigma_i^{(k+1)} = c_i^{(k+1)} \bar{\sigma}_i^{(k+1)}$. Hence by induction (C.3) and (C.1) are true for all $k > 0$ and Theorem 11 follows.

Appendix D. Derivation of the proposed recursive algorithm

In order to minimize $V'(\mathbf{Z})$ we use a cyclic algorithm. That is, we minimize with respect to one component at a time. We follow an

approach similar to that in Zachariah and Stoica (2015), with the main difference being that here we consider arbitrary nonnegative weights \mathbf{w}_i .

Note that minimization of $V'(\mathbf{Z})$ with respect to $z_{i,j}$ is equivalent to minimizing

$$V'(z_{i,j}) = \|\tilde{\xi}_{i,j} - \mathbf{c}_j z_{i,j}\|_2 + w_{i,j} |z_{i,j}|$$

where

$$\tilde{\xi}_{i,j} = \xi_i - \sum_{k \neq j} \mathbf{c}_k z_{i,k}, \quad \mathbf{c}_j = [\Gamma^\top - \Phi^\top \mathbf{H}].$$

As in Zachariah and Stoica (2015) it can be shown that the sign of the optimal $\hat{z}_{i,j}$ is given by $\text{sign}(\hat{z}_{i,j}) = \text{sign}(\mathbf{c}_j^\top \tilde{\xi}_{i,j})$. Hence we only have to find the absolute value $r_{i,j} = |z_{i,j}|$. Let

$$\alpha_{i,j} = \|\tilde{\xi}_{i,j}\|_2^2, \quad \beta_j = \|\mathbf{c}_j\|_2^2, \quad g_{i,j} = \mathbf{c}_j^\top \tilde{\xi}_{i,j}.$$

It is then straightforward to verify that the minimization of $V'(z_{i,j})$ is equivalent to minimizing

$$V'(r_{i,j}) = (\alpha_{i,j} + \beta_j r_{i,j}^2 - 2g_{i,j} r_{i,j})^{1/2} + w_{i,j} r_{i,j},$$

over all $r_{i,j} \geq 0$, and then setting $\hat{z}_{i,j} = \text{sign}(g_{i,j}) \hat{r}_{i,j}$. From the Cauchy–Schwarz inequality it follows that

$$\alpha_{i,j} \beta_j \geq g_{i,j}^2.$$

Using this inequality it was shown in Zachariah and Stoica (2015) that $V'(r_{i,j})$ is a convex function. The derivative of $V'(r_{i,j})$ is given by (dropping the subindices for now),

$$\frac{dV'}{dr} = \frac{\beta r - |g|}{(\beta r^2 - 2|g|r + \alpha)^{1/2}} + w. \quad (\text{D.1})$$

Since $V'(r)$ is convex it follows that it is minimized by $r = 0$ if and only if $dV'(0)/dr \geq 0$, i.e., if and only if

$$\alpha w^2 \geq g^2. \quad (\text{D.2})$$

Next we study the case when $g^2 > \alpha w^2$. It then follows from (D.1) that the stationary points of $V'(r)$ satisfy

$$(\beta r - |g|) = -w(\beta r^2 - 2|g|r + \alpha)^{1/2}. \quad (\text{D.3})$$

Solving this equation for r we get the stationary point

$$\hat{r} = \frac{|g|}{\beta} - \frac{w}{\beta \sqrt{\beta - w^2}} \sqrt{\alpha \beta - g^2}.$$

Hence we can conclude that the minimizer of $V'(z_{i,j})$ is given by

$$\hat{z}_{i,j} = \begin{cases} \text{sign}(g_{i,j}) \hat{r}_{i,j} & \text{if } \alpha_{i,j} w_{i,j}^2 < g_{i,j}^2 \\ 0 & \text{otherwise} \end{cases}.$$

Next we show how to obtain this estimate using only recursively computed quantities. Let

$$\mathbf{T} = (\Gamma^\top - \Phi^\top \mathbf{H})^\top (\Gamma^\top - \Phi^\top \mathbf{H}) \quad (\text{D.4})$$

$$\kappa_i = \|\xi_i\|_2^2 \quad (\text{D.5})$$

$$\rho_i = (\Gamma^\top - \Phi^\top \mathbf{H})^\top \xi_i \quad (\text{D.6})$$

$$\eta_i = \|\xi_i - (\Gamma^\top - \Phi^\top \mathbf{H}) \mathbf{z}_i\|_2^2 \quad (\text{D.7})$$

$$\zeta_i = (\Gamma^\top - \Phi^\top \mathbf{H}) (\xi_i - (\Gamma^\top - \Phi^\top \mathbf{H}) \mathbf{z}_i). \quad (\text{D.8})$$

Then it is straightforward to show that

$$\alpha_{i,j} = \eta_i + \beta_j z_{i,j}^2 + 2\zeta_{i,j} z_{i,j}$$

$$\beta_j = \mathbf{T}_{j,j}$$

$$g_{i,j} = \zeta_{i,j} + \beta_j z_{i,j}.$$

Also define $\Psi^{\mathbf{a},\mathbf{b}}(t)$ recursively, for any two vector-valued signals $\mathbf{a}(t), \mathbf{b}(t)$, as

$$\Psi^{\mathbf{a},\mathbf{b}}(0) = \mathbf{0} \quad (\text{D.9})$$

$$\Psi^{\mathbf{a},\mathbf{b}}(t+1) = \Psi^{\mathbf{a},\mathbf{b}}(t) + \mathbf{a}(t)\mathbf{b}^\top(t). \quad (\text{D.10})$$

Note that $\Psi^{\mathbf{a},\mathbf{b}}(t) = (\Psi^{\mathbf{b},\mathbf{a}}(t))^\top$. It can be verified that all quantities (D.4)–(D.8), and thus $\hat{z}_{i,j}$, can be computed from $\Psi^{\cdot,\cdot}(N)$.

The full algorithm for updating the needed quantities, including the update of $\hat{\Theta}$ and \mathbf{H} , is summarized in Algorithm 1. Note that the iterations of the outer for-loop can be executed in parallel.

Algorithm 1 : Recursive solution to (26)

- 1: Input: $\mathbf{y}(t), \varphi(t), \gamma(t), \check{\Theta}$ and $\check{\mathbf{Z}}$
 - 2: Update $\mathbf{P}(t), \hat{\Theta}(t)$ and $\mathbf{H}(t)$ according to (32)–(34).
 - 3: Update $\Psi^{\varphi,\varphi}(t), \Psi^{\gamma,\gamma}(t), \Psi^{\mathbf{y},\mathbf{y}}(t), \Psi^{\varphi,\gamma}(t), \Psi^{\varphi,\mathbf{y}}(t)$ and $\Psi^{\gamma,\mathbf{y}}(t)$ according to (D.10).
 - 4: $\mathbf{T} = \Psi^{\gamma,\gamma} - \Psi^{\gamma,\varphi} \mathbf{H} - \mathbf{H}^\top \Psi^{\varphi,\gamma} + \mathbf{H}^\top \Psi^{\varphi,\varphi} \mathbf{H}$.
 - 5: **for** $i = 1, \dots, n_y$ **do**
 - 6: $\kappa = \Psi^{\mathbf{y},\mathbf{y}} + \check{\theta}_i^\top \Psi^{\varphi,\varphi} \check{\theta}_i - 2\check{\theta}_i^\top [\Psi^{\varphi,\mathbf{y}}]_i$.
 - 7: $\rho = [\Psi^{\gamma,\mathbf{y}}]_i - \Psi^{\gamma,\varphi} \check{\theta}_i - \mathbf{H}^\top [\Psi^{\varphi,\gamma}]_i + \mathbf{H}^\top \Psi^{\varphi,\varphi} \check{\theta}_i$.
 - 8: $\eta = \kappa - 2\rho^\top \check{\mathbf{z}}_i + \check{\mathbf{z}}_i^\top \mathbf{T} \check{\mathbf{z}}_i$.
 - 9: $\zeta = \rho - \mathbf{T} \check{\mathbf{z}}_i$
 - 10: **repeat**
 - 11: **for** $j = 1, \dots, q$ **do**
 - 12: $\alpha = \eta + \mathbf{T}_{j,j} \check{z}_{i,j}^2 + 2\zeta_j \check{z}_{i,j}$.
 - 13: $g = \zeta_j + \mathbf{T}_{j,j} \check{z}_{i,j}$.
 - 14: $\beta = \mathbf{T}_{j,j}$.
 - 15: $\hat{r} = \frac{|g|}{\beta} - \frac{w_{i,j}}{\beta \sqrt{\beta - w_{i,j}^2}} \sqrt{\alpha \beta - g^2}$.
 - 16: $\hat{z}_{i,j} = \begin{cases} \text{sign}(g) \hat{r} & \text{if } \alpha w_{i,j}^2 < g^2 \\ 0 & \text{otherwise} \end{cases}$
 - 17: $\eta := \eta + \mathbf{T}_{j,j} (\check{z}_{i,j} - \hat{z}_{i,j}) + 2(\zeta_j - \hat{z}_{i,j}) \zeta_j$.
 - 18: $\zeta := \zeta + [\mathbf{T}]_j (\check{z}_{i,j} - \hat{z}_{i,j})$.
 - 19: $\check{z}_{i,j} := \hat{z}_{i,j}$.
 - 20: **end for**
 - 21: **until** number of iterations equals L .
 - 22: **end for**
 - 23: $\hat{\mathbf{Z}} = \hat{\mathbf{Z}}$.
 - 24: $\hat{\Theta} = \hat{\Theta} - \hat{\mathbf{Z}} \mathbf{H}$.
 - 25: Output: $\hat{\Theta}, \hat{\mathbf{Z}}$.
-

References

- Bai, E. W., & Liu, Y. (2007). Recursive direct weight optimization in nonlinear system identification: A minimal probability approach. *IEEE Transactions on Automatic Control*, 52(7), 1218–1231.
- Barber, D. (2012). *Bayesian reasoning and machine learning*. Cambridge University Press.
- Bemporad, A., Garulli, A., Paoletti, S., & Vicino, A. (2005). A bounded-error approach to piecewise affine system identification. *IEEE Transactions on Automatic Control*, 50(10), 1567–1580.
- Bijl, H., van Wingerden, J. W., Schön, T., & Verhaegen, M. (2015). Online sparse gaussian process regression using fitc and pitc approximations. *IFAC-PapersOnLine*, 48(28), 703–708.
- Billings, S. (2013). *Nonlinear system identification: NARMAX methods in the time, frequency, and spatio-temporal domains*. Wiley.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
- Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control*. John Wiley & Sons.
- Chen, H. (2004). Extended recursive least squares algorithm for nonlinear stochastic systems. In *American control conference, 2004. Proceedings of the 2004, Vol. 5* (pp. 4758–4763). IEEE.
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1–38.
- Enqvist, M. (2005). *Linear models of nonlinear systems* (Ph.D. thesis), Linköping, Sweden: Linköping University.

- Giri, F., & Bai, E. W. (2010). *Block-oriented nonlinear system identification*, Vol. 1. Springer.
- Hunter, D. R., & Lange, K. (2004). A tutorial on MM algorithms. *The American Statistician*, 58(1), 30–37.
- Juloski, A., Heemels, W., & Ferrari-Trecate, G. (2004). Data-based hybrid modelling of the component placement process in pick-and-place machines. *Control Engineering Practice*, 12(10), 1241–1252.
- Juloski, A. L., Paoletti, S., & Roll, J. (2006). Recent techniques for the identification of piecewise affine and hybrid systems. In L. Menini, L. Zaccarian, & C. Abdallah (Eds.), *Current trends in nonlinear systems and control* (pp. 79–99). Boston: Birkhäuser.
- Ljung, L. (1998). *System identification: Theory for the user*. Pearson Education.
- Ljung, L. (1999). Model validation and model error modeling. In *The Åström symposium on control* (pp. 15–42). Studentlitteratur.
- Ljung, L. (2007). *System identification toolbox for use with MATLAB*. The MathWorks, Inc..
- Ljung, L., & Söderström, T. (1983). *Theory and practice of recursive identification*. Cambridge, MA: MIT Press.
- Mattsson, P., & Wigren, T. (2016). Convergence analysis for recursive Hammerstein identification. *Automatica*, 71, 179–186.
- Ohlsson, H., & Ljung, L. (2013). Identification of switched linear regression models using sum-of-norms regularization. *Automatica*, 49(4), 1045–1050.
- Paduart, J., Lauwers, L., Swevers, J., Smolders, K., Schoukens, Johan, & Pintelon, Rik (2010). Identification of nonlinear systems using polynomial nonlinear state space models. *Automatica*, 46(4), 647–656.
- Pillonetto, G. (2013). Consistent identification of Wiener systems: A machine learning viewpoint. *Automatica*, 49(9), 2704–2712.
- Pillonetto, G., Dinuzzo, F., Chen, T., De Nicolao, G., & Ljung, Lennart (2014). Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, 50(3), 657–682.
- Roll, J., Nazin, A., & Ljung, L. (2005). Nonlinear system identification via direct weight optimization. *Automatica*, 41(3), 475–490.
- Schoukens, J., Vaes, M., & Pintelon, R. (2016). Linear system identification in a nonlinear setting: Nonparametric analysis of the nonlinear distortions and their impact on the best linear approximation. *IEEE Control Systems*, 36(3), 38–69.
- Sjöberg, J., Zhang, Q., Ljung, L., Benveniste, A., Delyon, Bernard, Glorennec, Pierre-Yves, et al. (1995). Nonlinear black-box modeling in system identification: a unified overview. *Automatica*, 31(12), 1691–1724.
- Söderström, T., & Stoica, P. (1988). *System identification*. Prentice-Hall, Inc..
- Solin, A., & Särkkä, S. (2014). Hilbert space methods for reduced-rank Gaussian process regression. arXiv preprint arXiv:1401.5508.
- Stoica, P., & Åhgren, P. (2002). Exact initialization of the recursive least-squares algorithm. *International Journal of Adaptive Control and Signal Processing*, 16(3), 219–230.
- Stoica, P., Zachariah, D., & Li, J. (2014). Weighted SPICE: A unifying approach for hyperparameter-free sparse estimation. *Digital Signal Processing*, 33, 1–12.
- Tayamon, S., Wigren, T., & Schoukens, J. (2012). Convergence analysis and experiments using an RPEM based on nonlinear ODEs and midpoint integration. In *2012 IEEE 51st annual conference on decision and control* (pp. 2858–2865). IEEE.
- Van den Hof, P., & Ninness, B. (2005). System identification with generalized orthonormal basis functions. In *Modelling and identification with rational orthogonal basis functions* (pp. 61–102). Springer.
- Wigren, T. (1993). Recursive prediction error identification using the nonlinear Wiener model. *Automatica*, 29(4), 1011–1025.
- Wigren, T. (2006). Recursive prediction error identification and scaling of non-linear state space models using a restricted black box parameterization. *Automatica*, 42(1), 159–168.
- Wu, T. T., & Lange, K. (2010). The mm alternative to em. *Statistical Science*, 25(4), 492–505.
- Zachariah, D., & Stoica, P. (2015). Online hyperparameter-free sparse estimation method. *IEEE Transactions on Signal Processing*, 63(13), 3348–3359.

Per Mattsson is a researcher at University of Gävle in Sweden. He received the M.Sc. degree in engineering physics from Uppsala University, Sweden, in 2010. He received the Ph.D. degree in automatic control and system identification from Uppsala University in 2016. His research interests include system identification, machine learning and automatic control.

Dave Zachariah is an associate senior lecturer at Uppsala University in Sweden. He received the M.S. degree in electrical engineering from Royal Institute of Technology (KTH), Stockholm, Sweden, in 2007. During 2007–2008 he worked as a research engineer at Global IP Solutions, Stockholm. He received the Tech. Lic. and Ph.D. degrees in signal processing from KTH in 2011 and 2013, respectively. His research interests include statistical signal processing, machine learning, estimation theory, and localization.

Petre Stoica is a researcher and educator in the field of signal processing and its applications to radar/sonar, communications and biomedicine. He is a professor of signal and system modeling at Uppsala University in Sweden, and a member of the Royal Swedish Academy of Engineering Sciences, the Romanian Academy (honorary), the European Academy of Sciences, and the Royal Society of Sciences in Uppsala.