

# An L-BFGS-B approach for linear and nonlinear system identification under $\ell_1$ and group-Lasso regularization

Alberto Bemporad, *Fellow, IEEE*

**Abstract**—In this paper, we propose a very efficient numerical method based on the L-BFGS-B algorithm for identifying linear and nonlinear discrete-time state-space models, possibly under  $\ell_1$  and group-Lasso regularization for reducing model complexity. For the identification of linear models, we show that, compared to classical methods, the approach often provides better results, is much more general in terms of the loss and regularization terms used (such as penalties for enforcing system stability), and is also more stable from a numerical point of view. The proposed method not only enriches the existing set of linear system identification tools but can also be applied to identifying a very broad class of parametric nonlinear state-space models, including recurrent neural networks. We illustrate the approach on synthetic and experimental datasets and apply it to solve a challenging industrial robot benchmark for nonlinear multi-input/multi-output system identification. A Python implementation of the proposed identification method can be found in the package `jax-sysid`, accessible at <https://github.com/bemporad/jax-sysid>.

**Index Terms**—Nonlinear system identification, recurrent neural networks, linear system identification,  $\ell_1$ -regularization, group-Lasso regularization.

## I. INTRODUCTION

Model-based control design requires, as the name implies, a dynamical model of the controlled process, typically a linear discrete-time one. Learning dynamical models that explain the relation between input excitations and the corresponding observed output response of a physical system over time, a.k.a. “system identification” (SYSID), has been investigated since the 1950s [32], mostly for linear systems [19]. In particular, subspace identification methods like N4SID [25] and related methods [24] have been used with success in practice and are available in state-of-the-art software tools like the System Identification Toolbox for MATLAB [20] and the Python package SIPPY [4]. The massive progress in supervised learning methods for regression, particularly feedforward and recurrent neural networks (RNNs), has recently boosted research in *nonlinear* SYSID [21], [27]. In particular, RNNs are black-box models that can capture the system’s dynamics concisely in a deterministic nonlinear state-space form, and techniques have been proposed for identifying them [29], including those based on autoencoders [22].

In general, control-oriented modeling requires finding a balance between model simplicity and representativeness, as the complexity of the model-based controller ultimately depends on the complexity of the model used, in particular in model predictive control (MPC) design [10], [23]. Therefore, two conflicting objectives must be carefully balanced: maximize the quality of fit and minimize the complexity of the model (e.g., reduce the number of states and/or sparsify the matrices defining the model). For this reason, when learning small parametric nonlinear models for control purposes, quasi-Newton methods [18], due to their convergence to very high-quality solutions, and extended Kalman filters (EKF) [6], [28]

The author is with the IMT School for Advanced Studies, Piazza San Francesco 19, Lucca, Italy. Email: [alberto.bemporad@imtlucca.it](mailto:alberto.bemporad@imtlucca.it). This work has received support from the European Research Council (ERC), Advanced Research Grant COMPACT (Grant Agreement No. 101141351)

can be preferable to simpler (stochastic) gradient descent methods like Adam [17] and similar variants. Moreover, to induce sparsity patterns,  $\ell_1$ -regularization is often employed, although it leads to non-smooth objectives; optimization methods were specifically introduced to handle  $\ell_1$ -penalties [2] and, more generally, non-smooth terms [3], [8], [31], [34].

In this paper, we propose a novel approach to solve linear and nonlinear SYSID problems, possibly under  $\ell_1$  and group-Lasso regularization, based on the classical L-BFGS-B (Limited-memory Broyden–Fletcher–Goldfarb–Shanno with Box constraints) algorithm [15], an extension of the L-BFGS algorithm to handle bound constraints, whose software implementations are widely available. To handle such penalties, we consider the positive and negative parts of the parameters defining the model and provide two simple technical results that allow us to solve a regularized problem with well-defined gradients. To minimize the open-loop simulation error, we eliminate the hidden states by creating a condensed-form of the loss function, which is optimized with respect to the model parameters and initial states. By relying on the very efficient automatic differentiation capabilities of the JAX library [11] to compute gradients, we will first show that the approach can be effectively used to identify linear state-space models on both synthetic and experimental data, analyzing the effect of group-Lasso regularization for reducing the number of states or for selecting the most relevant input signals. We will also show that the approach is much more stable from a numerical viewpoint and often provides better results than classical linear subspace methods like N4SID, which, in addition, cannot handle non-smooth regularization terms, non-quadratic losses, and constraints on model parameters. We will also show how to identify open-loop stable linear models via additional non-quadratic regularization terms. Finally, we will apply the proposed method in a nonlinear SYSID setting to solve the challenging industrial robot benchmark for black-box nonlinear multi-input/multi-output SYSID proposed in [35], by identifying a recurrent neural network (RNN) model in residual form under  $\ell_1$ -regularization largely extending the preliminary results presented in [7].

## II. SYSTEM IDENTIFICATION PROBLEM

Given a sequence of input/output training data  $(u_0, y_0), \dots, (u_{N-1}, y_{N-1})$ ,  $u_k \in \mathbb{R}^{n_u}$ ,  $y_k \in \mathbb{R}^{n_y}$ , we want to identify a deterministic state-space model in the following form

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + f_x(x_k, u_k; \theta_x) \\ \hat{y}_k &= Cx_k + Du_k + f_y(x_k, u_k; \theta_y) \end{aligned} \quad (1)$$

where  $k$  denotes the sample instant,  $x_k \in \mathbb{R}^{n_x}$  is the vector of hidden states,  $A, B, C, D$  are matrices of appropriate dimensions,  $f_x : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ , and  $f_y : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_y}$  are nonlinear functions parametrized by  $\theta_x \in \mathbb{R}^{n_{\theta_x}}$  and  $\theta_y \in \mathbb{R}^{n_{\theta_y}}$ , respectively.

The training problem we want to solve is (cf. [8]):

$$\min_{z, x_1, \dots, x_{N-1}} r(z) + \frac{1}{N} \sum_{k=0}^{N-1} \ell(y_k, Cx_k + Du_k + f_y(x_k, u_k; \theta_y)) \quad (2a)$$

$$\begin{aligned} \text{s.t. } & x_{k+1} = Ax_k + Bu_k + f_x(x_k, u_k; \theta_x) \\ & k = 0, \dots, N-2 \\ & z \triangleq [x_0' \text{ col } A' \text{ col } B' \text{ col } C' \text{ col } D' \theta_x' \theta_y']' \end{aligned} \quad (2b)$$

where the optimization vector  $z \in \mathbb{R}^{n_z}$  collects the entries of the initial state  $x_0$  and of the model parameters  $A, B, C, D, \theta_x, \theta_y$ , with dimension  $n_z = n_x + (n_x + n_u)(n_x + n_y) + n_{\theta_x} + n_{\theta_y}$ ,  $\text{col}$  is the vectorization operator stacking matrix columns,  $\ell: \mathbb{R}^{n_y} \times \mathbb{R}^{n_y} \rightarrow \mathbb{R}$  is a differentiable loss function, and  $r: \mathbb{R}^{n_z} \rightarrow \mathbb{R}$  is a regularization term.

After eliminating  $x_k, k = 1, \dots, N-1$ , by replacing  $x_{k+1}$  as in (2b), Problem (2) can be rewritten as the unconstrained (and in general nonconvex) nonlinear programming (NLP) problem

$$\begin{aligned} \min_z & f(z) + r(z) \\ f(z) &= \frac{1}{N} \sum_{k=0}^{N-1} \ell(y_k, Cx_k + Du_k + f_y(x_k, u_k; \theta_y)). \end{aligned} \quad (3a)$$

We will sometimes also incorporate box constraints  $z_\ell \leq z \leq z_u$  in (3) to impose prescribed limits on the model parameters.

### A. Special cases

1) *Linear state-space models*: A special case of (2) is system identification of linear state-space models based on the minimization of the simulation error

$$\begin{aligned} \min_z & r(z) + \frac{1}{N} \sum_{k=0}^{N-1} \|y_k - Cx_k - Du_k\|_2^2 \\ \text{s.t. } & x_{k+1} = Ax_k + Bu_k, \quad k = 0, \dots, N-2 \\ & z \triangleq [x_0' \text{ col } A' \text{ col } B' \text{ col } C' \text{ col } D']'. \end{aligned} \quad (4a)$$

Specific model structures, such as  $y_k = [I \ 0]x_k$  or other sparsity patterns, can be imposed by simply removing components of  $z$ . Moreover, by using bound constraints on  $z$ , special properties of the model can be enforced; for example, *positive linear systems* can be identified by constraining the entries of  $A, B, C, D$  (and, possibly,  $x_0$ ) to be nonnegative.

2) *Training recurrent neural networks*: Recurrent neural networks (RNNs) are special cases of (1) in which  $f_x, f_y$  are multi-layer feedforward neural networks (FNNs) parameterized by weight/bias terms  $\theta_x \in \mathbb{R}^{n_{\theta_x}}$  and  $\theta_y \in \mathbb{R}^{n_{\theta_y}}$ , respectively, and  $A = 0, B = 0, C = 0, D = 0$ . In particular,  $f_x$  is a FNN with linear output function and  $L_x - 1$  layers parameterized by weight/bias terms  $\{A_i^x, b_i^x\}$  (whose components are collected in  $\theta_x$ ),  $i = 1, \dots, L_x$ , and activation functions  $f_i^x, i = 1, \dots, L_x - 1$ , and similarly  $f_y$  by weight/bias terms  $\{A_i^y, b_i^y\}, i = 1, \dots, L_y$  (forming  $\theta_y$ ) and activation functions  $f_i^y, i = 1, \dots, L_y - 1$ , followed by a possibly nonlinear output function  $f_{L_y}^y$  [6], [8]. We denote by  $n_1^x, \dots, n_{L_x-1}$  and  $n_1^y, \dots, n_{L_y-1}$  the number of neurons in the hidden layers of  $f_x$  and  $f_y$ , respectively.

3)  $\ell_2$  and  $\ell_1$  regularization: To prevent overfitting, both standard  $\ell_2$  and  $\ell_1$  regularization terms can be introduced. In Section IV we will use the elastic-net regularization

$$r(z) = \frac{1}{2} \left( \rho_\theta \|\Theta\|_2^2 + \rho_x \|x_0\|_2^2 \right) + \tau \|\Theta\|_1 \quad (5)$$

where  $\Theta$  collects all the model parameters related to the state-update and output function, i.e.,  $z = [x_0' \Theta']'$ ,  $\rho_\theta > 0, \rho_x > 0$ , and  $\tau \geq 0$ . While  $\ell_1$ -regularization promotes model sparsity, it also makes the objective function in (2) non-smooth. The quadratic  $\ell_2$ -penalty term is also beneficial from an optimization perspective, in that it regularizes the Hessian matrix of the objective function.

4) *Group-Lasso regularization*: To effectively reduce the order  $n_x$  of the identified model (1), the group-Lasso penalty [36]

$$r_g(z) = \tau_g \sum_{i=1}^{n_x} \|I_i z\|_2 \quad (6a)$$

can be included in (3a), where  $\tau_g \geq 0$  and  $I_i$  is the submatrix formed by collecting the rows of the identity matrix of order  $n_z$  corresponding to the entries of  $z$  related to the initial state  $x_{0i}$ , the  $i$ th column and  $i$ th row of  $A$ , the  $i$ th row of  $B$ , and the  $i$ th column of  $C$ . In the case of recurrent neural networks, the group also includes the  $i$ th columns of the weight matrices  $A_1^x$  and  $A_1^y$  of the first layer of the FNNs  $f_x, f_y$ , the  $i$ th row of the weight matrix  $A_{L_x}$ , and  $i$ th entry of the bias term  $b_{L_x}$  of the last layer of  $f_x$ , respectively [8].

Similarly, the group-Lasso penalty

$$r_g(z) = \tau_g \sum_{i=1}^{n_u} \|I_i z\|_2 \quad (6b)$$

can be used to select the input channels that are most relevant to describe the dynamics of the system, where now  $I_i$  is the submatrix formed by collecting the rows of the identity matrix of order  $n_z$  corresponding to the entries of  $z$  related to the  $i$ th column of  $B$  and  $i$ th column of  $D$ . This could be particularly useful, for example, when identifying Hammerstein models in which the input enters the model through a (possibly large) set of nonlinear basis functions. In the case of recurrent neural networks, the  $i$ th group also contains the column of the weight matrices  $A_1^x$  and  $A_1^y$  of the first layer of the FNNs  $f_x, f_y$  related to the  $i$ th input. Clearly, group-Lasso penalties can be combined with  $\ell_2$  and  $\ell_1$ -regularization.

We finally remark that *rather arbitrary* (possibly nonconvex) differentiable loss functions  $\ell$  can be handled by the approach proposed in this paper, as we use nonlinear programming to solve Problem (2).

### B. Handling multiple experiments

To handle multiple experiments  $\{u_0^j, y_0^j, \dots, u_{N_j-1}^j, y_{N_j-1}^j\}, j = 1, \dots, M$ , Problem (2) can be reformulated as in [8, Eq. (4)] by optimizing with respect to  $x_0^1, \dots, x_0^M, A, B, C, D, \theta_x, \theta_y$ , or just  $A, B, C, D, \theta_x, \theta_y$  with  $x_0^j = 0, \forall j = 1, \dots, M$ . Alternatively, as suggested in [5], [6], [22], we can introduce an initial-state encoder

$$x_0^j = f_e(v_0^j; \theta_e) \quad (7)$$

where  $v_0 \in \mathbb{R}^{n_v}$  is a measured vector available at time 0, such as a collection of  $n_a$  past outputs,  $n_b$  past inputs, and/or other measurements that are known to influence the initial state of the system, and  $f_e: \mathbb{R}^{n_v} \times \mathbb{R}^{n_{\theta_e}} \rightarrow \mathbb{R}^{n_x}$  is a FNN parameterized by a further optimization vector  $\theta_e \in \mathbb{R}^{n_{\theta_e}}$  to be learned jointly with  $A, B, C, D, \theta_x, \theta_y$ .

### C. Stability constraints

Although system instability is largely discouraged by minimizing the open-loop simulation error over the entire duration of the experiment(s), and the regularization terms also discourage the possible occurrence of unobservable modes, when learning linear models (4) the identified matrix  $A$  may have eigenvalues outside the unit circle, for example in the case of short experiments. Without loss of generality, as justified by next Lemma 1, asymptotic stability can be enforced by adding the following constraint

$$\|A\|_2^2 < 1 \quad (8)$$

in (3), where  $\|A\|_2$  denotes the spectral norm of  $A$ .

*Lemma 1*: Let  $\Sigma \triangleq (A, B, C, D)$  be an asymptotically stable discrete-time linear system. Then there exists a transformation  $T$  such

that the equivalent system  $\bar{\Sigma} \triangleq (\bar{A}, \bar{B}, \bar{C}, \bar{D})$  is such that  $\|\bar{A}\|_2 < 1$ , where  $\bar{A} = TAT^{-1}$ ,  $\bar{B} = TB$ ,  $\bar{C} = CT^{-1}$ ,  $\bar{D} = D$ .

*Proof:* See Appendix A. ■

Lemma 1 proves that the simple constraint (8) does not limit the expressivity of the linear model architecture. In our experiments, we relax (8) by including the additional penalty

$$r_A(z) = \rho_A \max\{\|A\|_2^2 - 1 + \epsilon_A, 0\}^2 \quad (9)$$

in  $r(z)$ , where  $\rho_A \gg 1$  and  $0 \leq \epsilon_A \ll 1$ . If the identified matrix  $A$  has eigenvalues outside the unit disk, the optimization can be repeated with a larger value of  $\rho_A$  and/or  $\epsilon_A$ .

### D. DC gain

Assuming a set of steady-state input/output pairs  $(y_j^{ss}, u_j^{ss})$  is also available,  $j = 1, \dots, S$ , we can improve the DC gain of the identified model by including the additional regularization term

$$r_s(z) = \frac{\rho_s}{S} \sum_{j=1}^S \|y_j^{ss} - Cx_j^{ss} - Du_j^{ss} - f_y(x_j^{ss}, u_j^{ss}; \theta_y)\|_2^2 \quad (10)$$

$$x_j^{ss} = Ax_j^{ss} + Bu_j^{ss} + f_x(x_j^{ss}, u_j^{ss}; \theta_x), \quad j = 1, \dots, S.$$

The steady-state state  $x_j^{ss}$  can be evaluated while optimizing  $z$  by an automatically differentiable solver, such as a linear system solver in the case of linear models (4), or, more generally, a nonlinear solver, e.g., the limited-memory Broyden solver [12].

## III. NON-SMOOTH NONLINEAR OPTIMIZATION

When both  $f$  and  $r$  are smooth functions in (3), the system identification problem can be solved by any general-purpose derivative-based NLP solver. The presence of  $\ell_1$  and group-Lasso regularization requires a little care. Nonlinear non-smooth optimization solvers exist that can deal with such penalties, see, e.g., the general purpose non-smooth NLP solvers [14], [16], [34], generalized Gauss Newton methods combined with ADMM [8], and methods specifically introduced for  $\ell_1$ -regularized problems such as the orphant-wise limited-memory quasi-Newton (OWL-QN) method [2]. *Stochastic* gradient descent methods based on mini-batches like Adam [17] would be quite inefficient to solve (2), due to the difficulty in splitting the learning problem in multiple sub-experiments and the very slow convergence of those methods compared to quasi-Newton methods, as we will show in Section IV (see also the comparisons reported in [8]).

When dealing with large datasets ( $Nn_y \gg 1$ ) and models ( $n_z \gg 1$ ), solution methods based on Gauss-Newton and Levenberg-Marquardt ideas, like the one suggested in [8], can be inefficient, due to the large size  $Nn_y \times n_z$  of the required *Jacobian* matrices of the residuals. In such situations, L-BFGS approaches [18] combined with efficient automatic differentiation methods can be more effective, as processing one epoch of data requires evaluating the *gradient*  $\nabla_z f$  of the loss function plus  $\nabla_z r$ .

We show now how the L-BFGS-B algorithm [15], an extension of L-BFGS to handle bound constraints, can be used to also handle  $\ell_1$  and group-Lasso regularization. To this end, we provide two simple technical results in the remainder of this section.

*Lemma 2:* Consider the  $\ell_1$ -regularized nonlinear programming problem

$$\min_x f(x) + \tau \|x\|_1 + r(x) \quad (11)$$

where  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ ,  $\tau > 0$ , and function  $r : \mathbb{R}^n \rightarrow \mathbb{R}$  is such that  $r(x) = \sum_{i=1}^n r_i(x_i)$ . In addition, let functions  $r_i : \mathbb{R} \rightarrow \mathbb{R}$  be convex and positive semidefinite ( $r_i(x_i) \geq 0$ ,  $\forall x_i \geq 0$ ,  $r_i(0) = 0$ ). Let  $g : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$  be defined by

$$g(y, z) = f(y - z) + \tau [1 \dots 1] \begin{bmatrix} y \\ z \end{bmatrix} + r(y) + r(-z) \quad (12)$$

and consider the following bound-constrained nonlinear program

$$\min_{y, z \geq 0} g(y, z). \quad (13)$$

Then any solution  $y^*, z^*$  of Problem (13) satisfies the complementarity condition  $y_i^* z_i^* = 0$ ,  $\forall i = 1, \dots, n$ , and  $x^* \triangleq y^* - z^*$  is a solution of Problem (11).

*Proof:* See Appendix B. ■

Note that, if functions  $r_i$  are symmetric, i.e.,  $r_i(x_i) = r_i(-x_i)$ ,  $\forall i = 1, \dots, n$ , then we can replace  $r(-z)$  with  $r(z)$  in (12). Hence, by applying Lemma 2 with  $r(x) = \rho \|x\|_2^2 = \rho (\sum_{i=1}^n x_i^2)$ , we get the following corollary:

*Corollary 1:* Consider the following nonlinear programming problem with elastic-net regularization

$$\min_x f(x) + \tau \|x\|_1 + \rho \|x\|_2^2 \quad (14a)$$

where  $\rho > 0$  and  $\tau > 0$ , and let

$$\begin{bmatrix} y^* \\ z^* \end{bmatrix} \in \arg \min_{y, z \geq 0} f(y - z) + \tau [1 \dots 1] \begin{bmatrix} y \\ z \end{bmatrix} + \rho \left\| \begin{bmatrix} y \\ z \end{bmatrix} \right\|_2^2. \quad (14b)$$

Then  $x^* \triangleq y^* - z^*$  is a solution of Problem (14a).

Note that (14b) includes the squared Euclidean norm  $\|y\|_2^2 + \|z\|_2^2$  instead of  $\|y - z\|_2^2$ , as it would result from a mere substitution of  $x \rightarrow y - z$  in both  $f$  and  $r$ . The regularization in (14b) provides a positive-definite, rather than only positive-semidefinite, term on  $[y' \ z']'$ , which in turn leads to better numerical properties of the problem, without changing its solution.

*Lemma 3:* Consider the  $\ell_1$ -regularized nonlinear program

$$\min_x f(x) + r(x) + \epsilon \|x\|_1 \quad (15)$$

where  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ , and let  $r : \mathbb{R}^n \rightarrow \mathbb{R}$  be convex and symmetric with respect to each variable, i.e., satisfying

$$r(x - 2x_i e_i) = r(x), \quad \forall x \in \mathbb{R}^n \quad (16a)$$

where  $e_i$  is the  $i$ th column of the identity matrix, and increasing on the positive orphant

$$r(x_0 + \gamma e_i) \geq r(x_0), \quad \forall x_0 \in \mathbb{R}^n, \quad x_0 \geq 0, \quad \gamma \geq 0 \quad (16b)$$

Let  $\epsilon > 0$  be an arbitrary (small) number, let  $g : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$  defined by

$$g(y, z) = f(y - z) + r(y + z) + \epsilon [1 \dots 1] \begin{bmatrix} y \\ z \end{bmatrix} \quad (17)$$

and consider the following bound-constrained nonlinear program

$$\min_{y, z \geq 0} g(y, z). \quad (18)$$

Then, any solution  $y^*, z^*$  of Problem (18) satisfies the complementarity condition  $y_i^* z_i^* = 0$ ,  $\forall i = 1, \dots, n$ , and  $x^* \triangleq y^* - z^*$  is a solution of Problem (15).

*Proof:* See Appendix C. ■

*Corollary 2:* Consider the following nonlinear programming problem with group-Lasso regularization

$$\min_x f(x) + \epsilon \|x\|_1 + \tau_g \sum_{i=1}^{n_g} \|I_i x\|_2 \quad (19a)$$

where  $\tau_g > 0$  and  $\epsilon > 0$  is an arbitrary small number. Let

$$\begin{bmatrix} y^* \\ z^* \end{bmatrix} \in \arg \min_{y, z \geq 0} f(y - z) + \epsilon [1 \dots 1] \begin{bmatrix} y \\ z \end{bmatrix} + \tau_g \sum_{i=1}^{n_g} \|I_i(y + z)\|_2 \quad (19b)$$

Then  $x^* \triangleq y^* - z^*$  is a solution of Problem (19a).

*Proof:* See Appendix D. ■

Corollary 2 enable us to also handle group-Lasso regularization terms by bound-constrained nonlinear optimization in which the objective function admits partial derivatives on the feasible domain. In particular, consider a basis vector  $e_j$  such that  $I_i e_j \neq 0$ . For  $r(y, z) = \|I_i(y + z)\|_2$ ,  $\lim_{\alpha \rightarrow 0^+} \frac{r(\alpha e_j, 0) - r(0, 0)}{\alpha} = \lim_{\alpha \rightarrow 0^+} \frac{r(0, \alpha e_j) - r(0, 0)}{\alpha} = \lim_{\alpha \rightarrow 0^+} \frac{|\alpha|}{\alpha} = 1$ , where the limit  $\alpha \rightarrow 0^+$  is taken since  $y, z \geq 0$ ; on the other hand, for  $r(x) = \|I_i x\|_2$ ,  $\lim_{\alpha \rightarrow 0} \frac{r(\alpha e_j) - r(0)}{\alpha} = \lim_{\alpha \rightarrow 0} \frac{|\alpha|}{\alpha}$  is not defined, where the limit  $\alpha \rightarrow 0$  is taken since there are no sign restrictions on the components of  $x$ .

Finally, it is easy to prove that Lemma 2 and Lemma 3 extend to the case in which only a subvector  $x_I$  of  $x$  enters the non-smooth regularization term,  $x_I \in \mathbb{R}^{n_I}$ ,  $n_I < n$ . In this case, we only need to replace  $x_I = y_I - z_I$ ,  $y_I, z_I \in \mathbb{R}^{n_I}$ ,  $y_I, z_I \geq 0$ .

### A. Constraints on model parameters and initial states

Certain model structures require imposing bounds on model parameters and/or initial states, such as positive linear systems, as mentioned in Section II-A.1, and *input-convex neural networks*, that require nonnegative weight matrices [1]. General box constraints  $x_{\min} \leq x \leq x_{\max}$ , where  $x \in \mathbb{R}^n$  is the optimization vector, can be immediately enforced in L-BFGS-B as bound constraints. When we must split  $x = y - z$  to handle  $\ell_1$  and group-Lasso regularization, if  $x_{\max, i} > 0$  we can bound the corresponding positive part  $0 \leq y_i \leq x_{\max, i}$ , or, if  $x_{\max, i} < 0$ , constrain the negative part  $z_i \geq -x_{\max, i}$  and remove  $y_i$  from the optimization vector. Similarly, if  $x_{\min, i} < 0$  one can limit the corresponding negative part by  $0 \leq z_i \leq -x_{\min, i}$ , or, if  $x_{\min, i} > 0$ , constrain  $y_i \geq -x_{\min, i}$  and remove  $z_i$ .

General constraints  $h_I(z) \leq 0$ ,  $h_E(z) = 0$ ,  $h : \mathbb{R}^{n_z} \rightarrow \mathbb{R}^{n_{hI}}$ ,  $h : \mathbb{R}^{n_z} \rightarrow \mathbb{R}^{n_{hE}}$ , as shown in (9) for stability constraints, can be addressed in the learning problem (3a) via penalty functions, such as

$$\min_z f(z) + r(z) + \rho_h \left( \sum_{i=1}^{n_{hI}} \max\{h_{Ii}(z), 0\}^2 + \sum_{j=1}^{n_{hE}} h_{Ej}(z)^2 \right) \quad (20)$$

where  $\rho_h > 0$  is a (large) penalty parameter.

### B. Preventing numerical overflows

Solving (2) directly from an arbitrary initial condition  $z_0$  can lead to numerical instabilities during the initial optimization steps. During training, to prevent such an effect, when evaluating the loss  $f(z)$  we saturate the state vector  $x_{k+1} = \text{sat}(Ax_k + Bu_k + f_x(x_k, u_k; \theta_x), x_{\text{sat}})$ , where  $\text{sat}(x, x_{\text{sat}}) = \min\{\max\{x, -x_{\text{sat}}\}, x_{\text{sat}}\}$ , vector  $x_{\text{sat}} \in \mathbb{R}^n$ , and the min and max functions are applied componentwise. As minimizing the open-loop simulation error  $f(z)$  in (3b) and regularizing  $x_0$  and the model coefficients discourage the presence of unstable dynamics, in most cases the saturation term will not be active at the optimal solution if  $x_{\text{sat}}$  is chosen large enough. After identifying the model, it is easy to test on training data whether the saturation constraint is inactive, so to prove the redundancy of the saturation function.

At the price of additional numerical burden, to avoid introducing non-smooth terms in the problem formulation, the soft-saturation function  $\text{sat}_\gamma(x, x_{\text{sat}}) = x_{\text{sat}} + \frac{1}{\gamma} \log \frac{1+e^{-\gamma(x+x_{\text{sat}})}}{1+e^{-\gamma(x-x_{\text{sat}})}}$  could be used as an alternative to  $\text{sat}(x, x_{\text{sat}})$ , where the larger  $\gamma > 0$  the closer the function is to hard saturation.

### C. Initial state reconstruction for validation

To validate a given trained model on new test data  $\{\bar{u}(0), \bar{y}(0), \dots, \bar{u}(\bar{N}-1), \bar{y}(\bar{N}-1)\}$ , we need a proper initial condition  $\bar{x}_0$  for computing open-loop output predictions. Clearly, if an initial state encoder as described in (7) was also trained, we can simply set  $\bar{x}_0 = f_e(\bar{v}_0; \theta_e)$ . Otherwise, as we suggested in [6], we can solve the nonlinear optimization problem of dimension  $n_x$

$$\min_{\bar{x}_0} r_x(\bar{x}_0) + \frac{1}{\bar{N}} \sum_{k=0}^{\bar{N}-1} \ell(\bar{y}(k), \hat{y}(k)) \quad (21)$$

via, e.g., global optimizers such as Particle Swarm Optimization (PSO), where  $\hat{y}(k)$  are generated by iterating (1) from  $x(0) = \bar{x}_0$ .

In this paper, in alternative, we propose to run an EKF (forward in time) and Rauch-Tung-Striebel (RTS) smoothing [30, p. 268] (backward in time), based on the learned model, for  $N_e$  times. The approach is summarized in Appendix E for completeness. Clearly, for further refinement, the initial state obtained by the EKF/RTS pass can be used as the initial guess of a local optimizer (like L-BFGS) solving (21).

## IV. NUMERICAL RESULTS

We apply the nonlinear programming formulations described in Section III to identify linear state-space models and train recurrent neural networks on synthetic and real-world datasets. All experiments are run in Python 3.11 on an Apple M1 Max machine using JAX [11] for automatic differentiation. The L-BFGS-B solver [18] is used via the JAXopt interface [9] to solve bound-constrained nonlinear programs. When using group-Lasso penalties as in (19b), we constrain  $y, z \geq \epsilon$  to avoid possible numerical issues in JAX while computing partial derivatives when  $y = z = 0$ , with  $\epsilon = 10^{-16}$ . We apply standard scaling  $x \leftarrow (x - \bar{x})/\sigma_x$  to each input and output signal, where  $\bar{x}$  and  $\sigma_x$  are, respectively, the empirical mean and standard deviation computed on training data. When learning linear models, initial states are reconstructed using the EKF+RTS method described in Appendix E, running a single forward EKF and backward RTS pass ( $N_e = 1$ ). Unless stated differently, matrix  $A$  is initialized as  $0.5I$ , the remaining coefficients are initialized by drawing random numbers from the normal distribution with standard deviation 0.1. When running Adam as a gradient-descent method, due to the absence of line search, we store the model corresponding to the best loss found during the iterations.

For single-output systems, the quality of fit is measured by the classical  $R^2$ -score,  $R^2 = 100 \left( 1 - \frac{\sum_{k=1}^N (y_k - \hat{y}_k)^2}{\sum_{k=1}^N (y_k - \frac{1}{N} \sum_{i=1}^N y_k)^2} \right)$ , where  $\hat{y}_k$  is the output simulated in open-loop from  $x_0$ . In the case of multiple outputs, we consider the average  $R^2$ -score  $\bar{R}^2 \triangleq \frac{1}{n_y} \sum_{i=1}^{n_y} R_i^2$ , where  $R_i^2$  is the  $R^2$ -score measuring the quality of fit of the  $i$ th component of the output.

### A. Identification of linear models

1) *Cascaded-Tanks benchmark*: We consider the single-input/single-output Cascaded-Tanks benchmark problem described in [33]. The dataset consists of 1024 training and 1024 test standard-scaled samples. L-BFGS-B is run for a maximum of 1000 function evaluations, after being initialized by 1000 Adam iterations. Due to the lack of line search, which may reduce the step size considerably and trap the algorithm into a local minimum, we found that Adam is a good way to initially explore the optimization-vector space and provide a good initial guess to L-BFGS-B for refining the solution, without suffering from the slow convergence of gradient descent methods. We select the model with the best  $R^2$  on training data

$n_x$	$R^2$ (training)			$R^2$ (test)		
	lbfgs	sippy	M	lbfgs	sippy	M
1	87.43	56.24	87.06	83.22	52.38	83.18 (ssest)
2	94.07	28.97	93.81	92.16	23.70	92.17 (ssest)
3	94.07	74.09	93.63	92.16	68.74	91.56 (ssest)
4	94.07	48.34	92.34	92.16	45.50	90.33 (ssest)
5	94.07	90.70	93.40	92.16	89.51	80.22 (ssest)
6	94.07	94.00	93.99	92.17	92.32	88.49 (n4sid)
7	94.07	92.47	93.82	92.17	90.81	< 0 (ssest)
8	94.49	< 0	94.00	89.49	< 0	< 0 (n4sid)
9	94.07	< 0	< 0	92.17	< 0	< 0 (ssest)
10	94.08	93.39	< 0	92.17	92.35	< 0 (ssest)

**TABLE I:** Cascaded-tanks benchmark:  $R^2$ -scores on training and test data obtained by running Adam followed by L-BFGS from 5 different initial conditions and selecting the model with best  $R^2$ -score on training data (lbfgs), by Sippy [4] (sippy), and by the System Identification Toolbox for MATLAB, selecting the best model obtained by running both n4sid and ssest.

obtained out of 5 runs from different initial conditions. We also run the N4SID method implemented in Sippy [4] (sippy), and the n4sid method (with both focus on prediction and simulation) and ssest (prediction error method, with focus on simulation) of the System Identification Toolbox for MATLAB (M), all with stability enforcement enabled. Table I summarizes the obtained results. The average CPU time spent per run is about 2.4 s (lbfgs), 30 ms (sippy), 50 ms (n4sid with focus on prediction), 0.3 s (n4sid with focus on simulation), 0.5 s (ssest).

It is apparent from the table that, while our approach always provides reliable results, the other methods fail in some cases. We observed that the main reason for failure is due to numerical instabilities of the N4SID method (used by n4sid and, for initialization, by ssest).

2) *Group-Lasso regularization for model-order reduction:* To illustrate the effectiveness of group-Lasso regularization in reducing the model order, we generated 2000 training data from the following linear system

$$x_{k+1} = \begin{bmatrix} 0.96 & 0.26 & 0.04 & 0 & 0 & 0 \\ -0.26 & 0.70 & 0.26 & 0 & 0 & 0 \\ 0 & 0 & 0.93 & 0.32 & 0.07 & 0 \\ 0 & 0 & -0.32 & 0.61 & 0.32 & 0 \\ 0 & 0 & 0 & 0 & 0.90 & 0.38 \\ 0 & 0 & 0 & 0 & -0.38 & 0.52 \end{bmatrix} x_k + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0.07 & 0 \\ 0.32 & 0 \\ 0 & 0.10 \\ 0 & 0.38 \end{bmatrix} u_k + \xi_k$$


$$y_k = \begin{bmatrix} x_1 \\ x_3 \end{bmatrix} + \eta_k$$

where  $\xi_k$  and  $\eta_k$  are vectors whose entries are zero-mean Gaussian independent noise signals with standard deviation 0.01.


Figure 1 shows the  $R^2$ -score on training data and the corresponding model order obtained by running 1000 Adam iterations followed by at most 1000 L-BFGS-B function evaluations to handle the regularization term (6a) for different values of the group-Lasso penalty  $\tau_g$ . The remaining penalties are  $\rho_\theta = \rho_x = 10^{-3}$ , and  $\tau = \epsilon = 10^{-16}$ . For each value of  $\tau_g$ , the best model in terms of  $R^2$ -score on training data is selected out of 10 runs from different initial conditions. The figure clearly shows the effect of  $\tau_g$  in trading off the quality of fit and the model order. The CPU time is about 3.85 s per run.

3) *Group-Lasso regularization for input selection:* To illustrate the effectiveness of group-Lasso regularization in reducing the number of inputs of the system, we generate 10000 training data from a randomly generated stable linear system with 3 states, one output, 10 inputs, and zero-mean Gaussian process and measurement noise with standard deviation 0.01. The last 5 columns of the  $B$  matrix are divided by 1000 to make the last 5 inputs almost redundant.

Figure 2 shows the  $R^2$ -score on training data and the corresponding



**Fig. 1:**  $R^2$ -score on training data and resulting model order obtained with the group-Lasso penalty (6a) for different values of  $\tau_g$ .



**Fig. 2:**  $R^2$ -score on training data and resulting number of model inputs under the group-Lasso penalty (6b) for different values of  $\tau_g$ .

number of nonzero columns in the  $B$  matrix obtained by running 1000 Adam iterations followed by a maximum of 1000 L-BFGS-B function evaluations to handle the term (6b) with  $\rho_\theta = \rho_x = 10^{-8}$ ,  $\tau = \epsilon = 10^{-16}$ , and different values of the group-Lasso penalty  $\tau_g$ . For each value of  $\tau_g$ , the best model in terms of  $R^2$ -score on training data is selected out of 10 runs from different initial conditions. The figure clearly shows the effect of  $\tau_g$  in trading off the quality of fit and the number of inputs kept in the model. The CPU time is about 3.71 s per run.

### B. Stability constraint

To test the effectiveness of the approach to enforce stability via (8), we generated 1000 training and test data from the following slightly unstable linear system

$$x_{k+1} = \begin{bmatrix} 1.0001 & 0.5 & 0.5 \\ 0 & 0.9 & -0.2 \\ 0 & 0 & 0.7 \end{bmatrix} x_k + \begin{bmatrix} -0.4168 \\ -0.0563 \\ -2.1362 \end{bmatrix} u_k + \xi_k$$

$$y_k = [1.6403 \ -1.7934 \ -0.8417] x_k + \eta_k$$

where  $\xi_k \sim \mathcal{N}(0, 0.01^2 I)$ ,  $\eta \sim \mathcal{N}(0, 0.05^2)$ . The training problem is solved in 6.4 s under the additional penalty (9) with  $\rho_A = 10^3$  and  $\epsilon_A = 10^{-3}$ . The obtained  $R^2$ -score is 92.27 on training data and 91.41 on test data, the eigenvalues of the identified  $A$  matrix are 0.57871687, 0.99996862, and 0.92904764.

### C. Industrial robot benchmark

1) *Nonlinear system identification problem:* The dataset described in [35] contains input and output samples collected from KUKA KR300 R2500 ultra SE industrial robot movements, resampled at 10 Hz, where the input  $u \in \mathbb{R}^6$  collects motor torques (Nm) and the output  $y \in \mathbb{R}^6$  joint angles (deg). The dataset contains two experimental traces: a training dataset of  $N = 39988$  samples and a test dataset of  $N_t = 3636$  samples. Our aim is to obtain a discrete-time recurrent neural network (RNN) model in residual form (1), with

sample time  $T_s = 100$  ms, by minimizing the mean squared error (MSE) between the measured output  $y_k$  and the open-loop prediction  $\hat{y}_k$  obtained by simulating (1). Standard scaling is applied on both input and output signals.

The industrial robot identification benchmark introduces multiple challenges, as the system generating the data is highly nonlinear, multi-input multi-output, data are slightly over-sampled (i.e.,  $\|y_k - y_{k-1}\|$  is often very small), and the training dataset contains a large number of samples, that complicates solving the training problem. As a result, minimizing the open-loop simulation error on training data is rather challenging from a computational viewpoint. We select the model order  $n_x = 12$  and shallow neural networks  $f_x, f_y$  with, respectively, 36 and 24 neurons, and swish activation function  $\frac{x}{1+e^{-x}}$ .

We first identify matrices  $A, B, C$  and a corresponding initial state  $x_0$  on training data by solving problem (4) with  $\rho_\theta = \rho_x = 0.001$ ,  $\tau = 0$  in (5), which takes 9.12 s by running 1000 L-BFGS-B function evaluations, starting from the initial guess  $A = 0.99I$  and with all the entries of  $B$  and  $C$  normally distributed with zero mean and standard deviation equal to 0.1. The largest absolute value of the eigenvalues of matrix  $A$  is approximately 0.9776.

The average  $R^2$ -score on all outputs is  $\overline{R^2} = 48.2789$  on training data and  $R^2 = 43.8573$  on test data. For training data, the initial state  $x_0$  is taken from the optimal solution of the NLP problem (3), while for test data we reconstruct  $x_0$  by running EKF+RTS based on the obtained model for  $N_e = 10$  epochs.

For comparison, we also trained the same model structure using the N4SID method [25] implemented in the System Identification Toolbox for MATLAB R2023b [20] with focus on open-loop simulation. This took 36.21 s on the same machine and provided the lower-quality result  $R^2 = 39.2822$  on training data and  $R^2 = 32.0410$  on test data. The N4SID algorithm in `sippy` with default options did not succeed in providing meaningful results on the training dataset.

After fixing  $A, B, C$ , we train  $\theta_x, \theta_y, x_0$  by minimizing the MSE open-loop prediction loss under the elastic net regularization (5) to limit overfitting and possibly reduce the number of nonzero entries in  $\theta_x, \theta_y$ , with regularization coefficients  $\rho_\theta = 0.01$  and  $\rho_x = 0.001$  in (5), and different values of  $\tau$ . The total number of parameters to optimize is 1650, i.e.,  $\dim(\theta_x) + \dim(\theta_y) = 1590$  plus 12 components of the initial state  $x_0$ . This amounts to  $2 \cdot 1590 + 12 = 3192$  optimization variables when applying the method of Lemma 2. We achieved the best results by randomly sampling 100 initial conditions for the model parameters and picking up the best one in terms of loss  $f(z)$  as defined in (3b) before running Adam for 2000 iterations; then, we refined the solution by running a maximum of 2000 L-BFGS-B function evaluations. For a given value of  $\tau$ , solving the training problem took an average of 92 ms per iteration on a single core of the CPU. Solving directly the same  $\ell_1$ -regularized nonlinear programming problem as in (11) via Adam [17] (1590 variables) with constant learning rate  $\eta = 0.01$  took about 66 ms per iteration, where  $\eta$  is chosen to have a good tradeoff between the expected decrease and the variance of the function values generated by the optimizer.

The average  $R^2$ -score ( $\overline{R^2}$ ) as a function of the  $\ell_1$ -regularization parameter  $\tau$  is shown in Figure 3. For each value of  $\tau$ , we identified a model and its initial state starting from 30 different initial best-cost values of the model parameters and  $x_0 = 0$ , and the model with best  $R^2$  on test data is plotted for each  $\tau$ . It is apparent from the figure that the better the corresponding  $\overline{R^2}$  is on training data, the worse it is on test data; this denotes that the training dataset is not informative enough, as overfitting occurs.

The model obtained that leads to the best  $\overline{R^2}$  on test data corresponds to setting  $\tau=0.008$ . Table II shows the resulting  $R^2$ -




Fig. 3: Average  $R^2$ -score of open-loop predictions on training/test data and model sparsity, as a function of  $\tau$ . For each value of  $\tau$ , we show the best  $\overline{R^2}$ -score obtained on test data out of 30 runs and the corresponding  $\overline{R^2}$ -score on training data and model sparsity.




Fig. 4: Average  $R^2$ -score of open-loop predictions on test data from state estimates  $\hat{x}_{k|k}$  obtained by EKF ( $\tau=0.008$ ).

scores obtained by running such a model in open-loop simulation for each output, along with the scores obtained by running the linear model ( $A, B, C$ ). The latter coincide with the  $R^2$ -scores shown in Figure 3 for large values of  $\tau$ , which lead to  $\theta_x = 0, \theta_y = 0$ .

	$R^2$ (training) RNN model	$R^2$ (test) RNN model	$R^2$ (training) linear model	$R^2$ (test) linear model
$y_1$	86.0482	68.6886	63.0454	63.9364
$y_2$	77.4169	70.5481	53.1470	35.2374
$y_3$	72.3625	64.9590	63.7287	55.7936
$y_4$	75.7727	36.4175	29.9444	27.2043
$y_5$	65.1283	32.0540	35.4554	44.2490
$y_6$	86.1674	70.4031	44.3523	36.7233
average	77.1493	57.1784	48.2789	43.8573

TABLE II: Open-loop simulation:  $R^2$ -scores ( $\tau=0.008$ ).

To assess the benefits introduced by running L-BFGS-B, Table III compares the results obtained by running: (i) 2000 Adam iterations followed by 2000 L-BFGS-B function evaluations, (ii) 2000 OWL-QN function evaluations, (iii) only Adam for 6000 iterations, which is long enough to achieve a similar quality of fit, and (iv) 4000 L-BFGS-B iterations without first warm-starting with Adam (“No-Adam”). The table shows the best result obtained out of 10 runs in terms of the largest  $\overline{R^2}$  on test data. Model coefficients are considered zero if the absolute value is less than  $10^{-6}$ . It is apparent that L-BFGS-B or OWL-QN iterations, compared to pure Adam, better minimize the training loss and lead to much sparser models, providing similar results in terms of both model sparsity and training loss, with slightly different tradeoffs. We remark that, however, OWL-QN would not be capable of handling group-Lasso penalties (6). It also evident the benefit of warm-starting L-BFGS-B with Adam.

Finally, we run an extended Kalman filter based on model (1) (with  $\theta_x, \theta_y$  obtained with  $\tau=0.008$ ) and, for comparison, on the

linear model  $(A, B, C)$  to estimate the hidden states  $\hat{x}_{k|k}$ . In both cases, as routinely done in MPC practice [26], we computed EKF estimates on the model augmented by the output disturbance model  $q_{k+1} = q_k$  plus white noise, with  $q_k \in \mathbb{R}^6$ . The average  $R^2$ -score of the  $p$ -step ahead predictions  $\hat{y}_{k+p|k}$  are shown in Figure 4. This is a more relevant indicator of model quality for MPC purposes than the open-loop simulation error  $\hat{y}_{k|0} - y_k$  considered in Figure 3.

solver	adam iters	fcn evals	$\overline{R^2}$ training	$\overline{R^2}$ test	# zeros ( $\theta_x, \theta_y$ )	CPU time (s)
L-BFGS-B	2000	2000	77.1493	57.1784	556/1590	309.87
OWL-QN	2000	2000	74.7816	54.0531	736/1590	449.17
Adam	6000	0	71.0687	54.3636	1/1590	389.39
No-Adam	0	4000	66.8804	55.6671	302/1590	1361.08

**TABLE III:** Training results obtained by running L-BFGS-B or OWL-QN iterations after Adam iterations, just Adam, and without running Adam first ( $\tau=0.008$ ). The table shows the result out of 10 runs corresponding to the best  $R_2$ -score achieved on test data.

## V. CONCLUSIONS

The proposed approach based on bound-constrained nonlinear programming can solve a very broad class of nonlinear system identification problems, possibly under  $\ell_1$  and group-Lasso regularization. It is very general and can handle a wide range of smooth loss and regularization terms, bounds on model coefficients, and model structures. As shown in numerical experiments, the approach is also valid for the identification of linear models, as it seems more stable from a numerical point of view and often provides better results than classical linear subspace methods, possibly at the price of higher computations and the need to rerun the problem from different initial guesses due to local minima.

## REFERENCES

[1] B. Amos, L. Xu, and J.Z. Kolter. Input convex neural networks. In *Proc. 34th Int. Conf. on Machine Learning. Proceedings of Machine Learning Research*, volume 70, pages 146–155, Sydney, Australia, 2017.

[2] G. Andrew and J. Gao. Scalable training of  $\ell_1$ -regularized log-linear models. In *Proc. 24th Int. Conf. on Machine Learning*, pages 33–40, 2007.

[3] A.Y. Aravkin, R. Baraldi, and D. Orban. A Levenberg–Marquardt method for nonsmooth regularized least squares. *SIAM Journal on Scientific Computing*, 46(4):A2557–A2581, 2024.

[4] G. Armenise, M. Vaccari, R. Bacci Di Capaci, and G. Pannocchia. An open-source system identification package for multivariable processes. In *UKACC 12th Int. Conf. Control*, pages 152–157, Sheffield, UK, 2018.

[5] G. Beintema, R. Toth, and M. Schoukens. Nonlinear state-space identification using deep encoder networks. In *Proc. Machine Learning Research*, volume 144, pages 241–250, 2021.

[6] A. Bemporad. Recurrent neural network training with convex loss and regularization functions by extended Kalman filtering. *IEEE Transactions on Automatic Control*, 68(9):5661–5668, 2023.

[7] A. Bemporad. Training recurrent neural-network models on the industrial robot dataset under  $\ell_1$ -regularization. In *7th Workshop on Nonlinear System Identification Benchmarks*, Eindhoven, The Netherlands, April 2023.

[8] A. Bemporad. Training recurrent neural networks by sequential least squares and the alternating direction method of multipliers. *Automatica*, 156:111183, October 2023.

[9] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert. Efficient and modular implicit differentiation. *arXiv preprint arXiv:2105.15183*, 2021.

[10] F. Borrelli, A. Bemporad, and M. Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.

[11] J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

[12] C.G. Broyden. A class of methods for solving nonlinear simultaneous equations. *Mathematics of Computation*, 19(92):577–593, 1965.

[13] R.S. Bucy and P.D. Joseph. *Filtering for Stochastic Processing with Applications to Guidance*. Interscience Publishers, New York, 1968.

[14] J.V. Burke, F.E. Curtis, A.S. Lewis, M.L. Overton, and L.E.A. Simões. Gradient sampling methods for nonsmooth optimization. In A. Bagirov et al., editor, *Numerical Nonsmooth Optimization*, pages 201–225. 2020. <http://www.cs.nyu.edu/overton/software/hanso/>.

[15] R.H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.

[16] F.E. Curtis, T. Mitchell, and M.L. Overton. A BFGS-SQP method for nonsmooth, nonconvex, constrained optimization and its evaluation using relative minimization profiles. *Optimization Methods and Software*, 32(1):148–181, 2017. <http://www.timitchell.com/software/GRANSO/>.

[17] D.P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[18] D.C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.

[19] L. Ljung. *System Identification : Theory for the User*. Prentice Hall, 2 edition, 1999.

[20] L. Ljung. *System Identification Toolbox for MATLAB*. The Mathworks, Inc., 2001. <https://www.mathworks.com/help/ident>.

[21] L. Ljung, C. Andersson, K. Tiels, and T.B. Schön. Deep learning and system identification. *IFAC-PapersOnLine*, 53(2):1175–1181, 2020.

[22] D. Masti and A. Bemporad. Learning nonlinear state-space models using autoencoders. *Automatica*, 129:109666, 2021.

[23] D.Q. Mayne, J.B. Rawlings, and M.M. Diehl. *Model Predictive Control: Theory and Design*. Nob Hill Publishing, LCC, Madison,WI, 2 edition, 2018.

[24] P. Van Overschee and B. De Moor. A unifying theorem for three subspace system identification algorithms. *Automatica*, 31(12):1853–1864, 1995.

[25] P. Van Overschee and B.L.R. De Moor. N4SID: Subspace algorithms for the identification of combined deterministic-stochastic systems. *Automatica*, 30(1):75–93, 1994.

[26] G. Pannocchia, M. Gabbicini, and A. Artoni. Offset-free MPC explained: novelties, subtleties, and applications. *IFAC-PapersOnLine*, 48(23):342–351, 2015.

[27] G. Pillonetto, A. Aravkin, D. Gedon, L. Ljung, A.H. Ribeiro, and T.B. Schön. Deep networks for system identification: a survey. *Automatica*, 171:111907, 2025.

[28] G.V. Puskorius and L.A. Feldkamp. Neurocontrol of nonlinear dynamical systems with Kalman filter trained recurrent networks. *IEEE Transactions on Neural Networks*, 5(2):279–297, 1994.

[29] H. Salehinejad, S. Sankar, J. Barfett, E. Colak, and S. Valaee. Recent advances in recurrent neural networks. 2017. <https://arxiv.org/abs/1801.01078>.

[30] S. Särkkä and L. Svensson. *Bayesian filtering and smoothing*, volume 17. Cambridge University Press, 2023.

[31] M. Schmidt, D. Kim, and S. Sra. Projected Newton-type methods in machine learning. In S. Sra, S. Nowozin, and S.J. Wright, editors, *Optimization for Machine Learning*, pages 305–329. MIT Press, 2012.

[32] J. Schoukens and L. Ljung. Nonlinear system identification: A user-oriented road map. *IEEE Control Systems Magazine*, 39(6):28–99, 2019.

[33] M. Schoukens, P. Mattson, T. Wigren, and J.-P. Noël. Cascaded tanks benchmark combining soft and hard nonlinearities. Technical report, Eindhoven University of Technology, 2016.

[34] L. Stella, A. Themelis, and P. Patrinos. Forward-backward quasi-Newton methods for nonsmooth optimization problems. *Computational Optimization and Applications*, 67(3):443–487, 2017. <https://github.com/kul-optec/ForBES>.

[35] J. Weigand, J. Götz, J. Ulmen, and M. Ruskowski. Dataset and baseline for an industrial robot identification benchmark. 2022. <https://www.nonlinearbenchmark.org/benchmarks/industrial-robot>.

[36] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.

## APPENDIX

### A. Proof of Lemma 1

Since  $\Sigma$  is asymptotically stable, there exist symmetric and positive definite matrices  $P, Q \in \mathbb{R}^{n_x \times n_x}$  such that the Lyapunov equation  $P - A^T P A = Q$  is satisfied. Let  $T^T T = P$  be a Cholesky decomposition of  $P$ . Then  $T^{-T}(T^T T - A^T T^T A)T^{-1} = \bar{Q}$ , where  $\bar{Q} \triangleq T^{-T} Q T^{-1} \succ 0$ . Hence,  $I - \bar{A}' \bar{A} = \bar{Q}$ , and therefore  $x'x - x' \bar{A}' \bar{A} x = x' \bar{Q} x \geq \alpha \|x\|_2^2, \forall x \neq 0$ , where  $\alpha$  is the smallest eigenvalue of  $\bar{Q}$ ,  $\alpha > 0$ , or equivalently  $1 - \|\bar{A}\|_2^2 / \|x\|_2^2 \leq \alpha$  and hence  $\frac{\|\bar{A}x\|_2}{\|x\|_2} \leq \sqrt{1 - \alpha}$ . Therefore,  $\|\bar{A}\|_2 = \sup_{\|x\|_2 \neq 0} \frac{\|\bar{A}x\|_2}{\|x\|_2} \leq \sqrt{1 - \alpha} < 1$ . ■

### B. Proof of Lemma 2

By contradiction, assume that  $y_i^* z_i^* \neq 0$  for some index  $i$ ,  $1 \leq i \leq n$ . Since  $y_i^*, z_i^* \geq 0$ , this implies that  $y_i^* z_i^* > 0$ . Let  $\alpha \triangleq \min\{y_i^*, z_i^*\}$  and set  $\bar{y} \triangleq y^* - \alpha e_i, \bar{z} \triangleq z^* - \alpha e_i$ , where  $e_i$  is the  $i$ th column of the identity matrix and clearly  $\alpha > 0$ . By setting  $\beta_y \triangleq \frac{\alpha}{y_i^*}, \beta_z \triangleq \frac{\alpha}{z_i^*}$ , we can rewrite  $\bar{y} = (1 - \beta_y)y^* + \beta_y(y^* - y_i^* e_i), \bar{z} = (1 - \beta_z)z^* + \beta_z(z^* - z_i^* e_i)$ , where clearly  $\beta_y, \beta_z \in (0, 1]$ . Since  $r_i$  are convex functions,  $r$  is also convex and satisfies Jensen's inequality:  $r(\bar{y}) = r((1 - \beta_y)y^* + \beta_y(y^* - y_i^* e_i)) \leq (1 - \beta_y)r(y^*) + \beta_y r(y^* - y_i^* e_i)$ . Since  $r$  is separable,  $r_i(0) = 0$ , and  $r_i(y_i^*) \geq 0$ , we also get  $r(\bar{y}) \leq r(y^*) - \beta_y(r(y^*) - r(y^* - y_i^* e_i)) = r(y^*) - \beta_y(r_i(y_i^*) - r_i(0)) \leq r(y^*)$ . Similarly,  $r(-\bar{z}) = r((1 - \beta_z)(-z^*) + \beta_z(-z^* + z_i^* e_i)) \leq (1 - \beta_z)r(-z^*) + \beta_z r(-z^* + z_i^* e_i) = r(-z^*) - \beta_z(r_i(-z_i^*) - r_i(0)) \leq r(-z^*)$ . Then, since  $\alpha > 0$  and  $\tau > 0$ , we obtain  $g(\bar{y}, \bar{z}) = f(y^* - \bar{z}^*) + \tau[1 \dots 1] \begin{bmatrix} y^* \\ z^* \end{bmatrix} - 2n\tau\alpha + r(\bar{y}) + r(-\bar{z}) < f(y^* - \bar{z}^*) + \tau[1 \dots 1] \begin{bmatrix} y^* \\ z^* \end{bmatrix} + r(y^*) + r(-z^*) = g(y^*, z^*)$ . This contradicts  $(y^*, z^*) \in \arg \min (13)$  and therefore the initial assumption  $y_i^* z_i^* \neq 0$ .

Now let  $x^* \triangleq y^* - z^*$  and assume by contradiction that  $x^* \notin \arg \min (11)$ . Let  $\bar{x} \in \arg \min (11)$  and set  $\bar{y} \triangleq \max\{\bar{x}, 0\}, \bar{z} \triangleq \max\{-\bar{x}, 0\}$ . Then, since  $\bar{y}_i \bar{z}_i = 0$  and  $r_i(0) = 0, \forall i = 1, \dots, n$ , we get  $g(\bar{y}, \bar{z}) = f(\bar{x}) + \tau\|\bar{x}\|_1 + \sum_{i=1}^n r_i(\bar{y}_i) + r_i(-\bar{z}_i) = f(\bar{x}) + \tau\|\bar{x}\|_1 + \sum_{i=1}^n r_i(\bar{x}_i) < f(x^*) + \tau\|x^*\|_1 + \sum_{i=1}^n r_i(x_i^*) = g(y^*, z^*)$ . This contradicts  $(y^*, z^*) \in \arg \min (13)$  and hence the assumption  $x^* \notin \arg \min (11)$ . ■

### C. Proof of Lemma 3

Similarly to the proof of Lemma 2, assume by contradiction that  $y_i^* z_i^* \neq 0$  for some index  $i$ ,  $1 \leq i \leq n$ . Since  $y_i^*, z_i^* \geq 0$ , this implies  $y_i^* z_i^* > 0$ . Let  $\alpha \triangleq \min\{y_i^*, z_i^*\}$  and set  $\bar{y} \triangleq y^* - \alpha e_i, \bar{z} \triangleq z^* - \alpha e_i$ , where clearly  $\alpha > 0$ . By setting  $\beta \triangleq \min\left\{\frac{\alpha}{y_i^*}, \frac{\alpha}{z_i^*}\right\}$ , we can rewrite  $\bar{y} = (1 - \beta)y^* + \beta(y^* - y_i^* e_i), \bar{z} = (1 - \beta)z^* + \beta(z^* - z_i^* e_i)$ , where clearly  $\beta \in (0, 1]$ . Then,  $r(\bar{y} + \bar{z}) = r((1 - \beta)(y^* + z^*) + \beta(y^* + z^* - (y_i^* + z_i^*) e_i)) \leq (1 - \beta)r(y^* + z^*) + \beta r(y^* + z^* - (y_i^* + z_i^*) e_i) \leq r(y^* + z^*)$ , where the first inequality follows from Jensen's inequality due to the convexity of  $r$  and the second inequality from (16b). Then, since  $\alpha > 0$ , we obtain  $g(\bar{y}, \bar{z}) = f(y^* - \bar{z}^*) + r(\bar{y} + \bar{z}) + \epsilon[1 \dots 1] \begin{bmatrix} y^* \\ z^* \end{bmatrix} - 2n\epsilon\alpha < f(y^* - \bar{z}^*) + r(y^* + z^*) + \epsilon[1 \dots 1] \begin{bmatrix} y^* \\ z^* \end{bmatrix} = g(y^*, z^*)$ . This contradicts  $(y^*, z^*) \in \arg \min (18)$  and therefore the initial assumption  $y_i^* z_i^* \neq 0$ .

Now let  $x^* \triangleq y^* - z^*$  and assume by contradiction that  $x^* \notin \arg \min (11)$ . Let  $\bar{x} \in \arg \min (11)$  and set  $\bar{y} \triangleq \max\{\bar{x}, 0\}, \bar{z} \triangleq \max\{-\bar{x}, 0\}$ . Then, since  $\bar{y}_i \bar{z}_i = 0, \forall i = 1, \dots, n$ , we get  $g(\bar{y}, \bar{z}) = f(\bar{y} - \bar{z}) + r(\bar{y} + \bar{z}) + \epsilon[1 \dots 1] \begin{bmatrix} \bar{y} \\ \bar{z} \end{bmatrix} = f(\bar{x}) + r(\bar{x}) + \epsilon\|\bar{x}\|_1 < f(x^*) + r(x^*) + \epsilon\|x^*\|_1 = f(y^* - z^*) + r(y^* - z^*) + \epsilon[1 \dots 1] \begin{bmatrix} y^* \\ z^* \end{bmatrix} = f(y^* - z^*) + r(y^* + z^*) + \epsilon[1 \dots 1] \begin{bmatrix} y^* \\ z^* \end{bmatrix}$ , where the equivalences  $r(\bar{y} + \bar{z}) = r(\bar{y} - \bar{z}) = r(\bar{x})$  and  $r(y^* + z^*) = r(y^* - z^*) =$

$r(x^*)$  follow by (16a) and the complementarity conditions  $\bar{y}_i \bar{z}_i = 0, y_i^* z_i^* = 0$ , respectively. This contradicts  $(y^*, z^*) \in \arg \min (13)$  and hence the assumption  $x^* \notin \arg \min (11)$ . ■

### D. Proof of Corollary 2

By the triangle inequality, we have that  $\|I_i(\alpha x_1 + (1 - \alpha)x_2)\|_2 \leq \|\alpha I_i x_1\|_2 + \|(1 - \alpha)I_i x_2\|_2 = \alpha \|I_i x_1\|_2 + (1 - \alpha) \|I_i x_2\|_2$ , for all  $0 \leq \alpha \leq 1$  and  $i = 1, \dots, n_g$ . Hence,  $g(x) = \sum_{i=1}^{n_x} \|I_i x\|_2$  satisfies Jensen's inequality and is therefore convex. Moreover, since  $\|I_i x\|_2^2 = \sum_{j \in J_i} x_j^2$ , where  $J_i$  is the set of indices of the components of  $x$  corresponding to the  $i$ th group,  $g(x)$  is clearly symmetric with respect to each variable  $x_j$ . Finally, for all vectors  $x_0 \geq 0$  and  $\gamma \geq 0$ , we clearly have that  $x_{0j} + \gamma \geq x_{0j}$ , and therefore  $\|I_i(x_0 + \gamma e_j)\|_2 \geq \|I_i x_0\|_2$ , proving that  $g$  is also increasing on the positive orthant. Hence, the corollary follows by applying Lemma 3. ■

### E. EKF and RTS smoother for initial state reconstruction

Algorithm 1 reports the procedure used to estimate the initial state  $x_0$  of a generic nonlinear parametric model

$$\begin{aligned} x_{k+1} &= f(x_k, u_k; \theta) \\ \hat{y}_k &= g(x_k, u_k; \theta) \end{aligned} \quad (22)$$

for a given input/output dataset  $(u_0, y_0), \dots, (u_{N-1}, y_{N-1})$  based on multiple runs of EKF with Joseph stabilized covariance updates [13, Eq. (4.26)] and RTS smoothing, where  $\theta$  is the learned vector of parameters.

The procedure requires storing  $\{x_{k|k}, x_{k+1|k}, P_{k|k}, P_{k+1|k}\}_{k=0}^{N-1}$  and, in the case of nonlinear models, also the Jacobian matrices  $A_k \triangleq \frac{\partial f(x_{k|k}, u_k; \theta)}{\partial x}$  to avoid recomputing them twice. In the reported examples we used  $x_0^s = 0, P_0^s = \frac{1}{\rho_x} I$ , which is equivalent to the regularization term  $\frac{\rho_x}{2} \|x_0\|_2^2$  (cf. Eq. (12) in [6]), and we set  $Q = 10^{-8} I, R = I$ .

#### Algorithm 1 Initial state reconstruction by EKF and RTS smoothing

**Input:** Dataset  $(u_k, y_k), k = 0, \dots, N - 1$ ; number  $N_e \geq 1$  of epochs; initial state  $x_0^s \in \mathbb{R}^{n_x}$  and covariance matrix  $P_0^s \in \mathbb{R}^{n_x \times n_x}$ ; output noise covariance matrix  $R \in \mathbb{R}^{n_y \times n_y}$  and process noise covariance matrix  $Q \in \mathbb{R}^{n_x \times n_x}$ .

1. **For**  $e = 1, \dots, N_e$  **do:**
2.  $x_{0|-1} \leftarrow x_0^s, P_{0|-1} \leftarrow P_0^s;$
- 2.1. **For**  $k = 0, \dots, N - 1$  **do:** [forward EKF]
  - 2.1.1.  $C_k \leftarrow \frac{\partial g}{\partial x}(x_{k|k-1}, u_k; \theta);$
  - 2.1.2.  $M_k \leftarrow P_{k|k-1} C_k' (R + C_k P_{k|k-1} C_k')^{-1};$
  - 2.1.3.  $e_k \leftarrow y_k - g(x_{k|k-1}, u_k; \theta);$
  - 2.1.4.  $x_{k|k} \leftarrow x_{k|k-1} + M_k e_k;$
  - 2.1.5.  $P_{k|k} \leftarrow (I - M_k C_k) P_{k|k-1} (I - M_k C_k)' + M_k R M_k';$
  - 2.1.6.  $A_k \leftarrow \frac{\partial f}{\partial x}(x_{k|k}, u_k; \theta);$
  - 2.1.7.  $P_{k+1|k} \leftarrow A_k P_{k|k} A_k' + Q;$
  - 2.1.8.  $x_{k+1|k} \leftarrow f(x_{k|k}, u_k; \theta);$
- 2.2.  $x_N^s \leftarrow x_{N|N-1}; P_N^s \leftarrow P_{N|N-1};$
- 2.3. **For**  $k = N - 1, \dots, 0$  **do:** [backward RTS smoother]
  - 2.3.1.  $G_k \leftarrow P_{k|k} A_k' P_{k+1|k}^{-1};$
  - 2.3.2.  $x_k^s \leftarrow x_{k|k} + G_k (x_{k+1}^s - x_{k+1|k});$
  - 2.3.3.  $P_k^s \leftarrow P_{k|k} + G_k (P_{k+1}^s - P_{k+1|k}) G_k';$

3. **End.**

**Output:** Initial state estimate  $x_0^s$ .