



HAL
open science

Learning Reduced Nonlinear State-Space Models: an Output-Error Based Canonical Approach

Steeven Janny, Quentin Possamai, Laurent Bako, Madiha Nadri, Christian Wolf

► **To cite this version:**

Steeven Janny, Quentin Possamai, Laurent Bako, Madiha Nadri, Christian Wolf. Learning Reduced Nonlinear State-Space Models: an Output-Error Based Canonical Approach. 2022. hal-03672151v1

HAL Id: hal-03672151

<https://hal.science/hal-03672151v1>

Preprint submitted on 19 May 2022 (v1), last revised 17 Nov 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning Reduced Nonlinear State-Space Models: an Output-Error Based Canonical Approach

Steeven Janny¹, Quentin Possamai¹, Laurent Bako³, Madiha Nadri² and Christian Wolf⁴

Abstract—The identification of a nonlinear dynamic model is an open topic in control theory, especially from sparse input-output measurements. A fundamental challenge of this problem is that very few to zero prior knowledge is available on both the state and the nonlinear system model. To cope with this challenge, we investigate the effectiveness of deep learning in the modeling of dynamic systems with nonlinear behavior by advocating an approach which relies on three main ingredients: (i) we show that under some structural conditions on the to-be-identified model, the state can be expressed in function of a sequence of the past inputs and outputs; (ii) this relation which we call the state map can be modelled by resorting to the well-documented approximation power of deep neural networks; (iii) taking then advantage of existing learning schemes, a state-space model can be finally identified. After the formulation and analysis of the approach, we show its ability to identify three different nonlinear systems. The performances are evaluated in terms of open-loop prediction on test data generated in simulation as well as a real world data-set of unmanned aerial vehicle flight measurements.

Keywords: nonlinear system identification, state-space models, model reduction, deep learning, auto-encoding

I. INTRODUCTION

A large majority of deployed methods from control theory have as a prerequisite the provision of relatively precise dynamic model characterizing the temporal evolution of the state variables at stake. This model generally plays a central role, since the performance of the method is often directly related to the accuracy of the model [2], [4], [7], [8], [25]. Consequently, modeling and identification of a dynamic system is an essential preliminary step, since it will serve as foundation for additional processing, such as controller or observer design. However, this is not a trivial task: in the general case, the system is complex, non-linear, and involves physical phenomena that are often difficult or even impossible to model correctly without strongly impacting the required computation time. On the other hand, the identification of the parameters of a non-linear model is a non-convex problem, which can require tremendous hours of calibrations and experiments. Moreover, identification of a dynamic system often requires the intervention of domain experts and the ability to freely interact with the system.

¹Quentin Possamai and Steeven Janny are with INSA-Lyon, LIRIS UMR CNRS 5205. quentin.possamai@insa-lyon.fr, steeven.janny@insa-lyon.fr

²Madiha Nadri is with Université Claude Bernard Lyon 1, LAGEPP madiha.nadri-wolf@univ-lyon1.fr

³Laurent Bako is with Univ Lyon, Ecole Centrale de Lyon, INSA Lyon, Université Claude Bernard Lyon 1, CNRS, Ampère, UMR5005, 69130 Ecully, France, laurent.bako@ec-lyon.fr

⁴Christian Wolf is with Naver Labs Europe, Meylan, France, christian.wolf@naverlabs.com

The development of data-driven techniques for the identification of non-linear systems has provided a promising response to these issues and has received great interest over the last decade (see for example [12], [14], [17]). Specifically, neural networks propose to remove the burden of modeling by replacing it with the collection of massive datasets from the system of interest. Modeling methods based on Deep Learning constitute an alternative solution to painstaking physical modeling. The main idea is based on the use of an extremely versatile model, capable of approaching most dynamics with a certain degree of precision, that can be directly identified from pairs of input-output measurements in the case of dynamic systems, provided that these measurements gather enough information necessary to approximate the true dynamic. Nevertheless, the great flexibility of neural networks comes at the cost of a lack of mathematical structure making it difficult to perform theoretical analysis in terms of robustness, precision and stability. Moreover, learning complex, high-dimensional dynamical systems is not straightforward. The general formulation leads to latent dynamic models lacking of meaningful physical structure and requires large dimensional state spaces.

In this article, we propose a new identification structure for nonlinear state-space systems from a set of observation trajectories and associated inputs. We demonstrate the existence of a regressor inspired by finite impulse response models allowing to map a series of past observations to future outputs, and provide bounds derived from the prediction error during deployment. We then deduced a high-dimensional canonical state-space model discovered using an output-error based approach and propose to learn an auto-encoder projecting the dynamics into a smaller state-space. We evaluate our proposal on different systems in simulation and in the real world.

II. RELATED WORK

Data-driven dynamic models are widely studied in the community and get a lot of attention. In particular, [6], [9], [20] propose to find governing physical equations by performing a sparse regression from the data. At the junction between physical model and learning, [13], [15], [24], [26] use neural networks to model complementary phenomena not described by the initial physical model. In particular, [2], [18], [23] extend the dynamic model of a unmanned aerial vehicle with a neural network in charge of predicting aerodynamic disturbances, which are often very demanding and intractable for real time physical simulation.

Close to our work, a body of literature proposes to use deep learning for the identification of latent models, that is to say without direct physical meaning. This is notably the case for recent works around the Koopman operator [5], [11], [16], [19]. Another solution is to use an auto-encoder structure to model the latent dynamics of a system from past observations, following [3], [14]. Our proposal differs from this line of works in three main points: (1) we provide theoretical results and conditions for the existence of the dynamical system that we identify, (2) we propose to use a high-dimensional regressor structure without explicit state representation, which will be deduced from a dimensionality reduction operation and (3) we evaluate our approach on challenging and unstable systems.

Notation. \mathbb{R} and \mathbb{N} denote the sets of reals and natural numbers respectively. $\|\cdot\|$ refers to a generic norm on some appropriate space.

III. PROBLEM STATEMENT AND PRELIMINARY RESULTS

A. Problem statement

We consider a nonlinear discrete-time system of the general form

$$\begin{cases} x_{t+1} = f^\circ(x_t, u_t) \\ y_t = h^\circ(x_t, u_t) + w_t, \end{cases} \quad (1)$$

with $x_t \in \mathcal{X} \subset \mathbb{R}^{n_x}$, $u_t \in \mathcal{U} \subset \mathbb{R}^{n_u}$, $y_t \in \mathcal{Y} \subset \mathbb{R}^{n_y}$ being the state, the input and the output of the system at discrete time $t \in \mathbb{N}$ respectively. $f^\circ : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ and $h^\circ : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_y}$ are some nonlinear vector-valued functions. As to $w_t \in \mathcal{W} \subset \mathbb{R}^{n_y}$, it represents measurement noise. We will make the following important assumptions:

- 1) The external signals u and w take values in compact sets \mathcal{U} and \mathcal{W} respectively.
- 2) The state-space \mathcal{X} is a known compact set containing the initial state x_0 .
- 3) $(\mathcal{X}, \mathcal{U}, \mathcal{W}, \mathcal{Y})$ and (f°, h°) satisfy the following invariance conditions:

$$\begin{aligned} \forall (x, u) \in \mathcal{X} \times \mathcal{U}, f^\circ(x, u) \in \mathcal{X} \\ \forall (x, u, w) \in \mathcal{X} \times \mathcal{U} \times \mathcal{W}, h^\circ(x, u) + w \in \mathcal{Y} \end{aligned}$$

- 4) f° (and h°) are uniformly Lipschitz continuous on $\mathcal{X} \times \mathcal{U} \subset \mathbb{R}^{n_x} \times \mathbb{R}^{n_u}$ with respect to \mathcal{U} , i.e., there exists a constant $\gamma_f > 0$ such that $\|f^\circ(x, u) - f^\circ(x', u)\| \leq \gamma_f \|x - x'\|$ for all $(x, x', u) \in \mathcal{X} \times \mathcal{X} \times \mathcal{U}$.

The assumptions 1–4 are required essentially to theoretically ensure the well-definedness of optimization problems that will be expressed later in the paper. Of course in the context of system identification, such types of assumptions are not intended to be checked prior to applying the method to be developed.

The problem of interest in this paper can be stated as follows: Given a finite number N of input-output data pairs $\{(u_t, y_t) : t = 1, \dots, N\}$ generated by a nonlinear system of the form (1), find an appropriate dimension n_x of a state-space representation along with estimates of the associated functions f° and h° . Here, the number n_y of outputs and the number n_u of inputs are known a priori. However the

dimension n_x of the state is a parameter of the model which needs to be estimated along with the maps (f°, h°) .

We develop a solution in three steps: first, a nonlinear regression model is derived from the data-generating system equations in (1). The underlying nonlinear map is then modelled by a deep neural network structure and trained with the available data following an output-error framework. Given this map, one can readily form an equivalent canonical state-space representation of (1) with, however, the drawback that its dimension may be high. Hence, the third and last step of the proposed procedure consists in model reduction¹, an objective which is achieved through the design of an appropriate encoder-decoder.

B. Preliminary results

An important challenge concerning the identification of the system (1) is the fact that the state x_t is not entirely measured. We therefore need to express it first as a function of the available past input-output measurements $\{(u_\tau, y_\tau) : \tau < t\}$. Indeed, if the noise w_t in (1) is assumed to be identically equal to zero, then under appropriate observability conditions on the system, there exists a time horizon ℓ and a map $\phi : \mathbb{R}^L \rightarrow \mathbb{R}^{n_x}$, with $L = \ell(n_u + n_y)$, such that the state x_t can be written as

$$x_t = \phi(z_t) \quad (2)$$

where

$$z_t = (u_{t-\ell}^\top \quad y_{t-\ell}^\top \quad \dots \quad u_{t-1}^\top \quad y_{t-1}^\top)^\top \quad (3)$$

is the so-called regressor vector. To show the existence of such a map ϕ , some observability conditions on the system to be identified (1) are needed. For this purpose let us start by introducing some notations. For a positive integer i , let $F_i : \mathbb{R}^{n_x} \times \mathbb{R}^{in_u} \rightarrow \mathbb{R}^{n_x}$ be the map defined recursively from the function f° in (1) as follows: for $x \in \mathbb{R}^{n_x}$ and $(u_1, \dots, u_i) \in \mathbb{R}^{in_u}$, $F_1(x, u_1) = f^\circ(x, u_1)$ and for all $i \geq 2$,

$$F_i(x, u_1, \dots, u_i) = f^\circ(F_{i-1}(x, u_1, \dots, u_{i-1}), u_i). \quad (4)$$

Before proceeding further, let us mention a useful property of the maps F_i .

Lemma 1: Under Assumption 3, if $f^\circ : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ is uniformly γ_f -Lipschitz on $\mathcal{X} \times \mathcal{U}$ with respect to \mathcal{U} , then the map F_i defined in (4) is uniformly γ_f^i -Lipschitz on $\mathcal{X} \times \mathcal{U}^i$ with respect to $\mathcal{U}^i \subset \mathbb{R}^{in_u}$.

Proof: The proof of this lemma is straightforward and is therefore omitted. ■

Now consider the function $\mathcal{O}_i : \mathbb{R}^{n_x} \times \mathbb{R}^{in_u} \rightarrow \mathbb{R}^{in_y}$ given by

$$\mathcal{O}_i(x, u_1, \dots, u_i) = \begin{pmatrix} h^\circ(x, u_1) \\ h^\circ(F_1(x, u_1), u_2) \\ \vdots \\ h^\circ(F_{i-1}(x, u_1, \dots, u_{i-1}), u_i) \end{pmatrix}. \quad (5)$$

¹By state-space model reduction, we mean the reduction of the state dimension. The process aims at finding another state-space model which is as close as possible to the primary one but with a compressed state.

For notational simplicity, let us pose $\bar{u}_{1|i} = (u_1, \dots, u_i)$ so that $\mathcal{O}_i(x, u_1, \dots, u_i)$ in the previous equality can be replaced by $\mathcal{O}_i(x, \bar{u}_{1|i})$.

Definition 1: The system (1) is said to be finite-time observable over a time horizon $r \in \mathbb{N}$ if for each $\bar{u} \in \mathcal{U}^r$, the function $\mathcal{O}_r(\cdot, \bar{u})$, with \mathcal{O}_r defined as in (5), is injective.

Note that if the observability property in Definition 1 holds for some $r \in \mathbb{N}$ then it holds as well for any $i \geq r$.

Proposition 1 (Existence of the map ϕ): If the nonlinear system (1) (considered under the assumption that $w \equiv 0$) is finite-time observable in the sense of Definition 1, then there exist $\ell \in \mathbb{N}$ and a (nonlinear) map $\phi : \mathbb{R}^L \rightarrow \mathbb{R}^{n_x}$ such that (2) holds for all time $t \geq \ell$, any initial state in \mathcal{X} and any input signal taking values in \mathcal{U} .

Proof: For discrete time indices (i, j) with $i \leq j$, let $\bar{y}_{i|j} = (y_i^\top \ \dots \ y_j^\top)^\top$ be a vector of outputs from time i to time j . Likewise define $\bar{u}_{i|j} = (u_i^\top \ \dots \ u_j^\top)^\top$.

By iterating the system equations, it is easy to see that

$$\bar{y}_{t-\ell|t-1} = \mathcal{O}_\ell(x_{t-\ell}, \bar{u}_{t-\ell|t-1}). \quad (6)$$

By the finite-time observability assumption of the system, $\mathcal{O}_\ell(\cdot, \bar{u})$ admits an inverse for any given $\bar{u} \in \mathcal{U}^\ell$. Denote with $\mathcal{O}_\ell^*(\cdot, \bar{u}) : \mathbb{R}^{\ell n_y} \rightarrow \mathbb{R}^{n_x}$ the inverse map of $\mathcal{O}_\ell(\cdot, \bar{u})$ which is such that $\mathcal{O}_\ell^*(\mathcal{O}_\ell(x, \bar{u}), \bar{u}) = x$.

It hence follows from (6) that

$$x_{t-\ell} = \mathcal{O}_\ell^*(\bar{y}_{t-\ell|t-1}, \bar{u}_{t-\ell|t-1}) \quad (7)$$

which, by recursively applying the first equation of (1), gives

$$x_t = F_\ell(\mathcal{O}_\ell^*(\bar{y}_{t-\ell|t-1}, \bar{u}_{t-\ell|t-1}), \bar{u}_{t-\ell|t-1}) \triangleq \phi(z_t). \quad (8)$$

Consider now the more realistic scenario where the (unknown) measurement noise sequence $\{w_t\}$ is nonzero. Then Eq. (6) becomes

$$\bar{y}_{t-\ell|t-1} = \mathcal{O}_\ell(x_{t-\ell}, \bar{u}_{t-\ell|t-1}) + \bar{w}_{t-\ell|t-1}. \quad (9)$$

As a consequence, the state can no longer be obtained exactly by Eq. (7) or (8) since $\bar{y}_{t-\ell|t-1}$ does not lie in the range of $\mathcal{O}_\ell(\cdot, \bar{u}_{t-\ell|t-1})$. Let in this case the state $x_{t-\ell}$ and x_t be estimated by

$$\hat{x}_{t-\ell} \in \arg \min_{x \in \mathcal{X}} \|\bar{y}_{t-\ell|t-1} - \mathcal{O}_\ell(x, \bar{u}_{t-\ell|t-1})\| \quad (10)$$

$$\hat{x}_t = F_\ell(\hat{x}_{t-\ell}, \bar{u}_{t-\ell|t-1}), \quad (11)$$

for some norm $\|\cdot\|$ on $\mathbb{R}^{\ell n_y}$. The optimization problem in (10) is well-defined since, by Assumptions 1-4, the function $x \mapsto \|\bar{y}_{t-\ell|t-1} - \mathcal{O}_\ell(x, \bar{u}_{t-\ell|t-1})\|$ is defined on a compact set \mathcal{X} and is continuous. These, by the extreme value theorem, are sufficient conditions for the existence of a minimum and for the existence of the minimizer $\hat{x}_{t-\ell}$ as defined above.

In contrast, the estimates $\hat{x}_{t-\ell}$ and \hat{x}_t need not be uniquely defined in a general setting. Uniqueness would require some more strict conditions on the system under consideration.

Here, we will be content with a set-valued version $\hat{\phi}$ of ϕ in the noisy estimation scenario. Hence let $\hat{\phi}$ be defined by

$$\hat{\phi}(z_t) = \{F_\ell(\hat{x}_{t-\ell}, \bar{u}_{t-\ell|t-1}) : \hat{x}_{t-\ell} \text{ as in (10)}\}.$$

A question we ask now is how far the noisy estimate (11) lies from the true state x_t . To study this, a stronger notion of observability is introduced as follows.

Definition 2: The system (1) is called *finite-time uniformly observable* over a time horizon $\ell \in \mathbb{N}$ if there exists a constant $\alpha_\ell > 0$ such that for each $\bar{u} \in \mathcal{U}^\ell$,

$$\|\mathcal{O}_\ell(x, \bar{u}) - \mathcal{O}_\ell(x', \bar{u})\| \geq \alpha_\ell \|x - x'\| \quad (12)$$

for all $(x, x') \in \mathbb{R}^{n_x} \times \mathbb{R}^{n_x}$. Here $\|\cdot\|$ denotes a generic norm defined on appropriate spaces.

Based on this property, it is possible to bound the error between the noisy estimate (11) and the true state.

Proposition 2: Under Assumptions 1-4, if the system (1) is *finite-time uniformly observable* over a time horizon $\ell \in \mathbb{N}$ in the sense of Definition 2, then

$$\|\hat{x}_t - x_t\| \leq 2\gamma_f^\ell \alpha_\ell^{-1} \|\bar{w}_{t-\ell|t-1}\| \quad (13)$$

where γ_f is the Lipschitz constant of f° (See Assumption 4) and α_ℓ is the constant appearing in (12).

Proof: It follows from the definition (10) of $\hat{x}_{t-\ell}$ that

$$\begin{aligned} \|\bar{y}_{t-\ell|t-1} - \mathcal{O}_\ell(\hat{x}_{t-\ell}, \bar{u}_{t-\ell|t-1})\| \\ \leq \|\bar{y}_{t-\ell|t-1} - \mathcal{O}_\ell(x, \bar{u}_{t-\ell|t-1})\|, \end{aligned}$$

for all $x \in \mathbb{R}^{n_x}$. In particular this inequality holds for $x = x_{t-\ell}$. By then invoking (9) we get

$$\begin{aligned} \|\mathcal{O}_\ell(x_{t-\ell}, \bar{u}_{t-\ell|t-1}) - \mathcal{O}_\ell(\hat{x}_{t-\ell}, \bar{u}_{t-\ell|t-1}) + \bar{w}_{t-\ell|t-1}\| \\ \leq \|\bar{w}_{t-\ell|t-1}\|. \end{aligned}$$

By the triangle inequality property of norms, it follows that

$$\begin{aligned} \|\mathcal{O}_\ell(x_{t-\ell}, \bar{u}_{t-\ell|t-1}) - \mathcal{O}_\ell(\hat{x}_{t-\ell}, \bar{u}_{t-\ell|t-1}) + \bar{w}_{t-\ell|t-1}\| \\ \geq \|\mathcal{O}_\ell(x_{t-\ell}, \bar{u}_{t-\ell|t-1}) - \mathcal{O}_\ell(\hat{x}_{t-\ell}, \bar{u}_{t-\ell|t-1})\| \\ - \|\bar{w}_{t-\ell|t-1}\| \end{aligned}$$

Combining with the previous inequality yields

$$\begin{aligned} \alpha_\ell \|\hat{x}_{t-\ell} - x_{t-\ell}\| \\ \leq \|\mathcal{O}_\ell(x_{t-\ell}, \bar{u}_{t-\ell|t-1}) - \mathcal{O}_\ell(\hat{x}_{t-\ell}, \bar{u}_{t-\ell|t-1})\| \\ \leq 2 \|\bar{w}_{t-\ell|t-1}\|. \end{aligned}$$

Here, the first inequality is a consequence of the assumption of uniform finite-time observability. As a consequence, $\|\hat{x}_{t-\ell} - x_{t-\ell}\| \leq 2\alpha_\ell^{-1} \|\bar{w}_{t-\ell|t-1}\|$. The result follows now by applying (11), the uniform Lipschitz assumption on f° stated in 4 and Lemma 1. ■

It can be observed from the expression of the error bound (13) that the more strongly the system is observable (that is, the larger the constant α_ℓ), the more robust the estimate \hat{x}_t . Indeed α_ℓ , when it exists, can be defined as

$$\inf_{\substack{\bar{u}, x, x' \in \mathcal{U}^\ell \times \mathcal{X} \times \mathcal{X} \\ x \neq x'}} \frac{\|\mathcal{O}_\ell(x, \bar{u}) - \mathcal{O}_\ell(x', \bar{u})\|}{\|x - x'\|}.$$




Fig. 1: (a) Block diagram of our canonical reduced state space representation as defined in (18) (b) Training: the dynamic is modeled by \hat{H} acting as a regressor from a short history of previous observations to future. The encoder-decoder model is used to reduce the size of the state z_t to x_t . We train each network by minimizing the prediction error as well as the reconstruction error.

IV. MODELING AND LEARNING

A. Nonlinear regression model

A starting point of our identification method for system (1) is to solve a nonlinear regression problem. To formulate this problem, note by Proposition 2 that the true state of the system can be written as $x_t = \hat{x}_t + \delta_t$ with $\|\delta_t\| \leq \gamma_f^\ell \alpha_\ell^{-1} \|\bar{w}_{t-\ell|t-1}\|$. Consider now plugging the state estimate (11) into the output equation of (1), which gives

$$\begin{aligned} y_t &= h^\circ(\hat{x}_t + \delta_t, u_t) + w_t \\ &= h^\circ(\hat{x}_t, u_t) + \xi_t, \end{aligned}$$

with ξ_t being an error component entirely due to the noise. It is indeed equal zero whenever $w \equiv 0$. It can be shown that ξ_t can be written as $\xi_t = w_t + \tilde{\delta}_t$ with $\|\tilde{\delta}_t\| \leq \gamma_h \gamma_f^\ell \alpha_\ell^{-1} \|\bar{w}_{t-\ell|t-1}\|$, where γ_h is the Lipschitz constant of the measurement function h° of system (1). Since \hat{x}_t is a function of z_t we end up with

$$y_t = H^\circ(z_t, u_t) + \xi_t \quad (14)$$

for some nonlinear function H° .

Remark 1: In the absence of noise, the exact expression of H° is $H^\circ(z_t, u_t) = h^\circ(F_\ell(\mathcal{O}_\ell^*(z_t), \eta(z_t)), u_t)$ with $\eta(z_t) = \bar{u}_{t-\ell|t-1}$.

The first step of the identification method is to construct a high dimensional state-space representation whose state is

the vector z_t defined in (3). More precisely, consider

$$\begin{cases} z_{t+1} = \bar{A}z_t + \bar{B}u_t + \bar{S}H^\circ(z_t, u_t) + \bar{S}\xi_t \\ y_t = H^\circ(z_t, u_t) + \xi_t, \end{cases} \quad (15)$$

where $\bar{A} = A \otimes I_{n_u+n_y}$, $\bar{B} = e_{n_x-1} \otimes I_{n_u}$, $\bar{S} = e_{n_x} \otimes I_{n_y}$, $e_i \in \mathbb{R}^{n_x}$ being the canonical basis vector which has 1 in its i -th entry and zero everywhere else, \otimes referring to the Kronecker product and $A \in \mathbb{R}^{n_x \times n_x}$ given by the canonical form

$$A = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & \ddots & 1 & 0 \\ 0 & \cdots & \cdots & 0 & 1 \\ 0 & \cdots & \cdots & 0 & 0 \end{pmatrix}.$$

From (14) it can be seen that (15) constitutes a state-space representation for system (1) since both models have the same input-output behavior for $t \geq \ell$.

Given a finite set of input-output data points $\{(u_t, y_t)\}_{t=1}^{T+\ell}$, an estimate of the function H° can be obtained in a certain nonlinear model class \mathcal{H} as

$$\hat{H} \in \arg \min_{H \in \mathcal{H}} J(H),$$

where $J(H)$ is a regression loss given as

$$J(H) = \frac{1}{T} \sum_{t=\ell+1}^{T+\ell} \alpha_t \|y_t - H(\hat{z}_t, u_t)\|^2 \quad (16)$$

$$\text{s.t. } \hat{z}_{t+1} = \bar{A}\hat{z}_t + \bar{B}u_t + \bar{S}H(\hat{z}_t, u_t), \quad \hat{z}_{\ell+1} = z_{\ell+1}, \quad (17)$$

where α_t is a weighting coefficient such as $\alpha_t = 1$ except for $\alpha_{\ell+1} = 10$. Regression starts after a burn-in phase of ℓ steps (ie. the window size), which are needed to construct a full state-representation.

The choice of the model class \mathcal{H} is fundamental for several reasons: it must be sufficiently large to represent a good approximation of H° , and it should not be too large in order to ensure learnability and generalization to unseen conditions [22]. In other words, this class needs to be complex enough to capture the behavior of the unknown nonlinear function H° while being easily identifiable from rather limited set of measurements. We rely on the well-documented approximation power of deep neural networks and propose to select \mathcal{H} to be the class of multi-layer perceptrons (MLP), whose capacity shall be limited. In particular, the number of hidden layers and number of neurons of each layer is considered as a hyperparameter optimized over a validation set independent of training and evaluation splits, as classically done in machine learning.

B. Model reduction

The last step of the proposed method consists in model reduction. In effect, the model described in (15), although structurally simple, may suffer from a high dimensional state vector z_t . This may be a concern for some applications. We therefore propose a second deep learning structure allowing nonlinear state-space model reduction by encoding the state

variable z_t into a low dimensional state variable $\bar{x}_t \in \mathbb{R}^{\bar{n}}$ for some user-defined dimension $\bar{n} \in \mathbb{N}$. Formally, we train an auto-encoder $(\mathcal{E}, \mathcal{D})$ such that $\bar{x}_t = \mathcal{E}(z_t)$ and $z_t = \mathcal{D}(\bar{x}_t)$. By applying these maps to Eq. (15) and neglecting the noise terms, we get an approximate representation of the initial system (1) as follows:

$$\begin{cases} \bar{x}_{t+1} = \mathcal{E}\left(\bar{A}\mathcal{D}(\bar{x}_t) + \bar{B}u_t + \bar{S}\hat{H}(\mathcal{D}(\bar{x}_t), u_t)\right) \\ \bar{y}_t = \hat{H}(\mathcal{D}(\bar{x}_t), u_t). \end{cases} \quad (18)$$

The state space equation is summarized in figure 1a. The parameters of the encoder \mathcal{E} and decoder \mathcal{D} are trained with the following reconstruction loss from data samples $\{z_t\}$ collected from the training set (see also Fig. 1b).

$$(\mathcal{E}, \mathcal{D}) = \arg \min_{\mathcal{E}', \mathcal{D}'} \|z_t - \mathcal{D}'(\mathcal{E}'(z_t))\| \quad (19)$$

Note that the dimension \bar{n} of the compressed state \bar{x}_t in the estimated model (18) is potentially different from the true state dimension n_x . Another observation is that by going from (15) to (18), one reduces the dimension of the state vector but at the cost of introducing some structural complexity. Hence the computational price associated with simulating a model such as (18) may still be high depending on the complexity of the auto-encoder $(\mathcal{E}, \mathcal{D})$.

From a formal point of view, this reduced state has several advantages. The constructed state-space z_t is not part of the update equation (18) anymore. We can also experimentally show (see section V), that this method can discover state representations of smaller size with a method which is generic in nature and can be applied to a broad class of problems.

V. EXPERIMENTAL RESULTS

We illustrate and evaluate the proposed nonlinear dynamical model identification approach on the estimation and prediction of the state of a system with unknown dynamics. To demonstrate the practical feasibility of our model, we propose to study its behavior in three different scenarios. First, we demonstrate the capabilities of our regressor function H° for output prediction on simulated systems.² We also study the influence of key parameters, namely the length ℓ of the time window and the impact of the state reduction.

A. Dynamical systems and benchmarks

We use two simulated and one real system to validate our contributions.

Tank — we test the proposed method on the cascade tank system introduced in [21]. This system relates the water level in two connected tanks without consideration of overflow. It has the form (1) of a discrete-time state-space model with f° and h° implicitly instantiated as follows

$$\begin{cases} x_{1,t+1} = x_{1,t} - k_1\sqrt{x_{1,t}} + k_2u_t \\ x_{2,t+1} = x_{2,t} + k_3\sqrt{x_{1,t}} - k_4\sqrt{x_{2,t}} \\ y_t = x_{2,t} + w_t, \end{cases} \quad (20)$$

²Code and datasets will be made publicly available upon acceptance.

with $x_t = (x_{1,t} \ x_{2,t}) \in \mathbb{R}^2$ being the state and $k_i, i = 1, \dots, 4$ known parameters.

2D Drone — we introduce a model of a 2-dimensional drone, i.e. an unmanned aerial vehicle, which moves in a 2D plane. The drone is equipped with two propellers and its dynamic is modeled by:

$$\begin{cases} \ddot{p}_x = -\frac{k_T}{m}(\Omega_1^2 + \Omega_2^2)\sin(\theta) - \frac{\gamma}{m}(\Omega_1 + \Omega_2)\dot{p}_x \\ \ddot{p}_z = \frac{k_T}{m}(\Omega_1^2 + \Omega_2^2)\cos(\theta) - \frac{\gamma}{m}(\Omega_1 + \Omega_2)\dot{p}_z - g \\ \ddot{\theta} = \frac{k_T L}{J}(\Omega_2^2 - \Omega_1^2) \\ y = (p_x \ p_z \ \theta)^T, \end{cases} \quad (21)$$

where (p_x, p_z) is the position, k_T the thrust constant, Ω_i the rotational speed of the i^{th} propeller, L the length of the UAV, m its mass, J its inertia and γ its friction coefficient. The main interest of such a system is its naturally unstable dynamics, which complicates the identification process. The system has been discretized.

3D Drone — We also evaluate on recordings of the Blackbird UAV flight dataset [1], which consists of 10 hours of aggressive quadrotor flights, measured with an accurate motion capture device. We use this real world data to demonstrate that our observer discovers a state representation containing the same information as in the physical state without any supervision.

We processed the raw data gathered from the on-board inertial measurement unit (IMU) and propeller rotation speeds as observation and command signals. The regressor is trained to simulate the IMU measurements, i.e. acceleration and angular speed of the drone expressed in the local frame.

Noise has been added to the observations for the two simulated systems, tank and 2D Drone. More details about the dataset generation is provided in the appendix.

B. Baseline methods

To experimentally compare our model to competing approaches from the literature, we introduce a neural baseline in the form of gated recurrent units (GRU) [10], the state-of-the-art variant of recurrent neural networks. This is a pure data-driven technique from the machine learning field, where the learned state representation directly corresponds to the hidden state vector of the GRU. For a fair comparison, we limit the size of the hidden vector to fit the corresponding size of z_t . We refer to this model as *classic GRU*. Its update equations are given by

$$\begin{aligned} h_{t+1} &= \text{GRU}([y_t; u_t], h_t), \quad h_0 = 0 \\ y_t &= \text{MLP}(h_t), \end{aligned} \quad (22)$$

where $\text{GRU}(\cdot)$ is a shorthand notation corresponding to the classical update equations of GRUs [10]. For simplicity, and as usually done, gates have been omitted from the notation.

The baseline is evaluated in a setting which is comparable to the proposed model. In particular, the model has access to the same window of input/output pairs $[y_{t+k}, u_{t+k}]_{k=1..l}$ during the initial burn-in phase. However, these values do not explicitly make up the state, as in our model. This data-driven baseline is sufficiently general to be able to learn the




Fig. 2: Visual example of the output prediction made by \hat{H} (**Ours (MLP)**) for different value of ℓ on the three datasets.

same state representation in theory, but there is no guarantee that training will lead to this solution.

We also experiment with the model introduced in [14] which consists of an auto-encoder with a learned latent dynamics that operates on the reduced state representation. This model has been evaluated on the same tank system, yet, with a different data collection technique. Train and test trajectory in the Tank dataset as proposed in [14] are generated from PRBS-like signals, which is a classical approach for system identification. Our version of Tank dataset is much more challenging: observations are collected from closed-loop simulations with targets generated in a procedural manner and PID control. In our dataset, we took care to explore a wide range of possible states with sparse measurements in the train set to prevent over-fitting on a specific command design.

C. Extension: a hybrid state-space model

We also introduce an extension of our model, which combines the advantages of both methodologies. It uses our proposed state representation z_t , but implements the mapping H° by a GRU in place of the MLP proposed in section IV. Formally, the GRU updates a zero-initialized hidden vector using the previous observations and command. This vector is then decoded by a MLP to the desired observation. Equation (15) is then used for forward prediction. We refer to this model as *Ours (GRU)*, it is given as follows:

$$\begin{aligned}
 z_{t+1} &= \bar{A}z_t + \bar{B}u_t + \bar{S} \text{MLP}(h_t) \\
 h_{i+1} &= \text{GRU}([y_i; u_i], h_i), \quad h_{t-l} = 0 \\
 y_t &= \text{MLP}(h_t).
 \end{aligned} \quad (23)$$

D. Output prediction and parameter analysis

Output forecasting – the identified dynamic model can be evaluated by performing open-loop forward prediction from initial conditions and the set of inputs applied to the real system. The model then forecasts outputs, which may be compared to actual ground truth measurements. We assessed the first stage of our method using this task, i.e. the resolution of the regression problem. Table I reports the mean squared error on 100 step roll-out predictions for each baseline and different window sizes $\ell \in \{5, 10, 15, 20, 25, 30\}$. Our method shows excellent prediction error even for low window sizes, and consistently outperforms the closest competing method from the literature, Masti et al. [14], by a large margin. We conjecture two key arguments to justify this difference : (1) the structure proposed by [14] suffers from complex interaction between the auto-encoder and the latent dynamics that penalizes learning, and (2) the model design process over-fitted on the simpler dataset used in the original paper.

Machine learning baseline – is competitive with our contribution. However, its structure forces to observe only one couple (y_t, u_t) at a time. Relevant information needs to be stored in its memory, the vectorial hidden state, and this storage process is fully learned by gradient descent, a difficult process. In principle, these models can learn a state representation which is similar or even identical to our designed state-map, but there is no guarantee that this representation emerges. Our state map model can therefore be seen as a form of useful inductive bias for recurrent neural models.

For moderate window sizes, our model benefits from the immediate availability of all the components of z_t in its state. For very large window sizes or complex dynamical

| Window size ℓ | Tank ($\times 10^{-4}$) | | | | 2D Drone ($\times 10^{-2}$) | | | | 3D Drone ($\times 10^{-2}$) | | | |
|-----------------------|---------------------------|-------------------|--------------|--------------|-------------------------------|-------------------|------------|-------------|-------------------------------|-------------------|-------------|------------|
| | Classic GRU [†] | Masti et al. [14] | Ours (GRU) | Ours (MLP) | Classic GRU [†] | Masti et al. [14] | Ours (GRU) | Ours (MLP) | Classic GRU [†] | Masti et al. [14] | Ours (GRU) | Ours (MLP) |
| 5 | 163 | 1030 | 138 | 7.14 | 62.8 | 60.5 | 106 | 31.4 | 24.8 | 14.6 | 6.44 | 15.2 |
| 10 | 41.7 | 1070 | 5.60 | 0.930 | 82.7 | 58.6 | 68.9 | 9.95 | 23.7 | 14.5 | 5.32 | 14.2 |
| 15 | 4.57 | 957 | 3.06 | 0.960 | 61.9 | 58.2 | 35.2 | 7.52 | 23.4 | 13.0 | 5.07 | 13.6 |
| 20 | 4.04 | 914 | 1.07 | 0.761 | 78.4 | 55.3 | 19.3 | 8.06 | 22.7 | 12.6 | 4.68 | 13.5 |
| 25 | 0.600 | 915 | 0.481 | 0.606 | 80.3 | 53.6 | 23.0 | 5.17 | 21.3 | 12.1 | 4.83 | 13.6 |
| 30 | 1.73 | 917 | 0.193 | 0.448 | 104 | 51.3 | 25.0 | 3.13 | 19.2 | 12.6 | 4.61 | 13.1 |

[†] The size of the hidden state of each GRU model is adapted to the window size s.t. fits the size of the equivalent regressor model.

TABLE I: Quantitative evaluation: we report MSE error over 100-step rollouts by the learned regression model and compare with baselines, for different windows sizes ℓ . Our model consistently outperforms all baselines.

systems (such as 3D Drone), the GRU extension (*Ours (GRU)*) outperforms the MLP regressor. In this situation, the GRU takes advantage of its incrementally updated memory, and manages to manipulate the large dimension of z_t by processing it piecewise, whereas the MLP must manipulate the entire vector. Figure 2 shows samples of predicted trajectory using the MLP regressor approach for each dataset.

E. Model reduction

The reduction step is performed downstream of the regression model training. Nevertheless, the difficulty of the reduction task is directly related to the initial size of the state representation z_t , that is, to the size of the window ℓ . In order to accurately evaluate the compression capabilities of our approach, we trained several auto-encoders for each value of $\ell \in \{5, 10, 15, 20, 25, 30\}$ corresponding to different rates of compression increasing by steps of 15%.

Figure 4 shows the compression capabilities of our encoder-decoder structure for different window sizes ℓ . The results are consistent on the three datasets. The compression rate is more sensitive on small input dimensions, and conversely, a larger dimension can be reduced extensively with negligible loss of accuracy. Indeed, increasing the number of inputs arguably leads to an increase in the redundancies exploitable by the encoder to reduce the dimension of the state space and reconstruct it with limited deviation with respect to the initial vector.

Yet such reduction introduces noise to the state representation that the regressor will have to cope with. We thus evaluate the impact of state-space reduction on the output forecasting capabilities of our model, and summarize the results in figure 3. Our reduction method manages to reduce the dimension of the state in a consistent way up to 60% for the two datasets in simulation without sensible variation of the prediction error. The error bars reflect the double dependence of our approach both on the performance of the regression model H° but also on the quality of the encoding-decoding. We compare favorably to the baseline in [14].

VI. CONCLUSION

In this work, we take advantage of the power of high-capacity deep neural networks to design a new methodology for estimating nonlinear dynamical systems from a set of input/output data pairs. We show that the state can be expressed as a state map computed as a function of past inputs

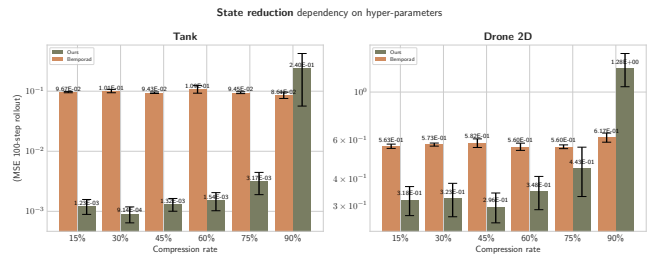


Fig. 3: We studied the impact of state compression for multiple configurations of final latent state dimension and temporal window and aggregate the results by this two parameters on the synthetic datasets. Specifically, we measure the MSE on observation prediction error for 100 step in the future (*Ours (MLP)*).




Fig. 4: Heatmap of MSE for the encoder-decoder structure depending on both the window size (which relates to the initial state dimension $\dim z$) and the compression rate.

and outputs. We learn a mapping from this representation to model outputs from training data using deep networks and show that this approach is competitive. We tackled the problem of reducing the state space, showing that this way a state of similar size than the original problem can be obtained through machine learning with an auto-encoding solution. The proposed approach can be used to reduce the order of a given nonlinear model, such as infinite-dimensional discrete systems. The methodology was validated using three numerical examples and using a data-set from real experiments from the literature.

REFERENCES

- [1] Amado Antonini, Winter Guerra, Varun Murali, Thomas Sayre-McCord, and Sertac Karaman. The Blackbird Dataset: A large-scale dataset for UAV perception in aggressive flight. *The International Journal of Robotics Research*, 2018.
- [2] Leonard Bauersfeld, Elia Kaufmann, Philipp Foehn, Sihao Sun, and Davide Scaramuzza. NeuroBEM: Hybrid Aerodynamic Quadrotor Model. *Robotics: Science and Systems*, 2021.

- [3] Gerben Beintema, Roland Toth, and Maarten Schoukens. Nonlinear state-space identification using deep encoder networks. In *Conference on Learning for Dynamics and Control*, 2021.
- [4] Alberto Bemporad. Model predictive control design: New trends and tools. In *Conference on Decision and Control*, 2006.
- [5] Steven L. Brunton Bethany Lusch, J. Nathan Kutz. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 2018.
- [6] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *National Academy of Sciences*, 2016.
- [7] Lucian Buşoniu, Tim de Bruin, Domagoj Tolić, Jens Kober, and Ivana Palunko. Reinforcement learning for control: Performance, stability, and deep approximators. *Annual Reviews in Control*, 2018.
- [8] Chien-Chern Cheah, Chao Liu, and Jean-Jacques E Slotine. Adaptive jacobian tracking control of robots with uncertainties in kinematic, dynamic and actuator models. *Transactions on Automatic Control*, 2006.
- [9] Zhao Chen, Yang Liu, and Hao Sun. Physics-informed learning of governing equations from scarce data. *Nature Communications*, 2021.
- [10] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Conference on Empirical Methods in Natural Language Processing*, 2014.
- [11] S. Janny, V. Andrieu, M. Nadri, and C. Wolf. Deep kkl: Data-driven output prediction for non-linear systems. In *Conference on Decision and Control*, 2021.
- [12] Lennart Ljung, Carl Andersson, Koen Tiels, and Thomas B. Schön. Deep learning and system identification. *International Federation of Automatic Control*, 2020.
- [13] Yun Long, Xueyuan She, and S. Mukhopadhyay. Hybridnet: Integrating model-based and data-driven learning to predict evolution of dynamical systems. In *Conference on Robot Learning*, 2018.
- [14] Daniele Masti and Alberto Bemporad. Learning nonlinear state–space models using auto encoders. *Automatica*, 2021.
- [15] Viraj Mehta, Ian Char, Willie Neiswanger, Youngseog Chung, Andrew Nelson, Mark Boyer, Egemen Kolenen, and Jeff Schneider. Neural Dynamical Systems: Balancing Structure and Flexibility in Physical Prediction. In *APS Division of Plasma Physics Meeting Abstracts*, 2021.
- [16] J. Peralez, F. Galuppo, P. Dufour, C. Wolf, and M. Nadri. Data-driven multimodel control waste for heat recovery system on a heavy duty truck engine. In *Conference on Decision and Control*, 2020.
- [17] Gianluigi Pillonetto, Francesco Dinuzzo, Tianshi Chen, Giuseppe De Nicolao, and Lennart Ljung. Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, 2014.
- [18] Quentin Possamaï, Steeven Janny, Madiha Nadri, Laurent Bako, and Christian Wolf. Learning to estimate uav created turbulence from scene structure observed by onboard cameras. In *ArXiv pre-print*, 2022.
- [19] Clarence W. Rowley, Igor Mezic, Shervin Bagheri, Philipp Schlatter, and Dan S. Henningson. Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics*, 2009.
- [20] Subham Sahoo, Christoph H. Lampert, and Georg Martius. Learning equations for extrapolation and control. In *International Conference on Machine Learning*, 2018.
- [21] M. Schoukens and J.P. Noël. Three benchmarks addressing open challenges in nonlinear system identification. *International Federation of Automatic Control*, 2017.
- [22] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press, 2014.
- [23] Guanya Shi, Xichen Shi, Michael O’Connell, Rose Yu, Kamyar Azizzadenezeheli, Animashree Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural Lander: Stable Drone Landing Control using Learned Dynamics. *International Conference on Robotics and Automation*, 2019.
- [24] Qi Wang, Feng Li, Yi Tang, and Yan Xu. Integrating model-driven and data-driven methods for power system frequency stability assessment and control. *Transactions on Power Systems*.
- [25] Alexander Weinmann. *Uncertain models and robust control*. Springer Science & Business Media, 2012.

A. Dataset details

Tank – dataset is generated by uniform sampling of five waypoints lying in $[0, 5]$ evenly distributed on time on a

- [26] Yuan Yin, Vincent Le Guen, Jérémie Dona, Ibrahim Aayed, Emmanuel de Bézenac, Nicolas Thome, and Patrick Gallinari. Augmenting physical models with deep networks for complex dynamics forecasting. In *Conference on Learning Representations*, 2021.

200 steps reference constructed by cubic spline interpolation between the waypoints. This reference is then tracked with a PID controller. Our dataset contains 60 trajectories for the train set, and 20 for both the validation and test set. For simulation, we use $k_1 = 0.5, k_2 = 0.4, k_3 = 0.2$ and $k_1 = 0.3$.

3D drone – dataset built on the BlackBird dataset [1]. We extracted IMU measurements and commands from raw data and apply pre-processing as follows : temporal synchronization of both signals, noise filtering have using Butterworth filters, and sampling rate reduction to 50Hz. To create train/valid/test splits, we sampled 20 flights to form the validation and 10 for test split. The remaining 146 flights were used for the training set. Each flight have been sliced in 200-steps chunks to facilitate training.

2D drone – dataset is generated by uniform sampling of 5 to 10 2D waypoints lying in $[-2, 2]$ evenly distributed on time on a 600 steps reference constructed by cubic spline interpolation between the waypoints. This reference is then tracked with an model predictive control approach. Our dataset contains 500 flights for training, and 20 flights for validation and test sets. For simulation, we choose $k_T = 4 \times 10^{-4}, \gamma = 10^{-9}, L = 0.15, m = 1$ and $J = 2.7 \times 10^{-3}$. The system is simulated with Euler integration scheme at 30Hz.

B. Models details

Classic GRU – is a 2 layer gated recurrent unit. The hidden vector size is chosen such that the cumulated dimension of the two hidden vectors matches the one of the corresponding state z_t , formally $n_h = \frac{1}{2}(n_u + n_y)$. The hidden vector is then decoded by a multilayer perceptron with one hidden unit of size n_h .

Ours (MLP) – uses a MLP to model H° with 3 hidden units of size 256 for simulated dataset and 2 layer with 2048 units for the 3D Drone. The encoder-decoder is modeled with 2 MLP with 2 layers of 512 units.

Ours (GRU) – model H° with a GRU with three layers, and hidden size of 128. The encoder-decoder is identical as Ours (MLP).

Each model is implemented in Pytorch and trained with Adam optimizer, with learning rate of 10^{-4} . We trained the regressor (both MLP and GRU) for 10,000 epochs, and the encoder for 3,000 for the simulated datasets and respectively 300 epochs for the 3D drone dataset.