

Neural Identification of Feedback-Stabilized Nonlinear Systems

Mahrokh G. Boroujeni,^{1*} Laura Meroi,^{1*} Leonardo Massai,¹ Clara L. Galimberti,²
and Giancarlo Ferrari-Trecate¹

Abstract—Neural networks have demonstrated remarkable success in modeling nonlinear dynamical systems. However, identifying these systems from closed-loop experimental data remains a challenge due to the correlations induced by the feedback loop. Traditional nonlinear closed-loop system identification methods struggle with reliance on precise noise models, robustness to data variations, or computational feasibility. Additionally, it is essential to ensure that the identified model is stabilized by the same controller used during data collection, ensuring alignment with the true system’s closed-loop behavior. The dual Youla parameterization provides a promising solution for linear systems, offering statistical guarantees and closed-loop stability. However, extending this approach to nonlinear systems presents additional complexities. In this work, we propose a computationally tractable framework for identifying complex, potentially unstable systems while ensuring closed-loop stability using a complete parameterization of systems stabilized by a given controller. We establish asymptotic consistency in the linear case and validate our method through numerical comparisons, demonstrating superior accuracy over direct identification baselines and compatibility with the true system in stability properties.

I. INTRODUCTION

Neural networks have shown strong performance in learning nonlinear dynamical systems from input-output data, due to their flexibility and expressive power [1], [2]. However, in many practical situations, collecting data under open-loop conditions is impractical or unsafe, particularly for unstable systems. Even for stable systems, feedback control may be necessary for safety, efficiency, or operation near desired setpoints. For instance, drones are controlled to fly safely, robots are operated to achieve certain tasks, and biological systems are regulated by natural mechanisms. In many cases, disconnecting the controller to collect open-loop data is either infeasible or prohibitively costly.

In such scenarios, data is typically gathered while the system operates in closed-loop with a stabilizing controller [3], [4], [5]. This motivates closed-loop system identification, which aims to learn models from such data. Compared to the open-loop case, this setting introduces challenges, most notably the correlation between input signals and measurement noise caused by feedback [3], [4].

¹ Institute of Mechanical Engineering, EPFL, Switzerland. {mahrokh.ghoddousiboroujeni, laura.meroi, l.massai, giancarlo.ferraritrecate}@epfl.ch.

² Dalle Molle Institute for Artificial Intelligence (IDSIA), SUPSI, Switzerland. clara.galimberti@supsi.ch.

* The first two authors collaborated equally.

This work was supported as a part of NCCR Automation, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 51NF40.225155) and the NECON project (grant number 200021219431).

Classical methods for closed-loop identification are grouped into direct, joint input-output, and indirect approaches [6]. Direct methods estimate both the system and noise model independently of the feedback. They are consistent, i.e., can recover the true model with infinite data, if the noise model is accurate. However, real-world disturbances are often too complex to model precisely [3], [4].

Joint input-output methods first identify a system driven by white noise that generates the joint input-output signal, then simultaneously estimate the target system and controller [7]. While consistent even with imperfect noise models, these methods often suffer from theoretically unbounded variance [8] and high computational cost, especially for nonlinear controllers [6]. Indirect methods consider the closed-loop mapping and identify the system by exploiting knowledge of the controller [4]. These methods parameterize the model to reformulate the problem, eliminating noise effects at the plant input and enabling the use of standard open-loop techniques [3], [4]. This ensures consistency even when the noise model is misspecified, at least in the linear case.

One commonly used parameterization is the *dual Youla* parameterization, which describes all systems stabilized by a given controller [9]. This formulation is the dual of the Youla-Kučera parameterization, which characterizes all controllers that stabilize a given system [10]. A key advantage of this approach is that it inherently ensures the identified model is stabilized by the controller. This aligns with the stability of the true system under the given closed-loop configuration [4]. The dual Youla parameterization involves the coprime factorizations of both the controller and a nominal plant stabilized by it, along with Youla parameters that can be optimized to minimize the model’s discrepancy from the target system. For a linear plant and controller, the Youla parameters correspond to a stable rational transfer function, making optimization straightforward.

Given dual Youla’s success in identifying linear systems, prior work has explored extending this parameterization to cases with nonlinear systems but linear controllers (see [11] for a review). However, two key challenges arise:

- Coprime factorizations exist only for a subclass of nonlinear systems and are typically difficult to obtain even when they do exist [11].
- The parameterization relies on stable operators, the nonlinear counterparts of the Youla parameter, and their inverses which are challenging to optimize over. A similar challenge has been noted in the context of learning stabilizing nonlinear controllers using the Youla parameterization [12].

As a result, existing methods do not lend themselves to a computationally tractable framework.

More recently, [5] proposed a tractable framework for closed-loop identification of nonlinear systems. While effective in their experiments, this method has two key limitations: (i) the closed-loop system of the learned model and controller may become unstable, and (ii) it assumes a predefined parametric form for the system, limiting flexibility. This leaves closed-loop identification of expressive, nonlinear models a critical yet not sufficiently explored challenge.

Contributions. We propose a closed-loop system identification framework for learning complex, potentially unstable models using neural networks while ensuring closed-loop stability with a given controller. Our approach is based on a parameterization inspired by the dual Youla framework, adapted to settings where both the system and controller may be nonlinear. To this end, we first derive a complete characterization of all systems that can be stabilized by a given controller in terms of an operator in \mathcal{L}_p . Motivated by [12], we then focus on a parametric subset of these operators, resulting in a computationally tractable optimization framework. In the linear case, we establish the asymptotic consistency of our approach, aligning with existing results on the consistency of the dual Youla parameterization [3], [4]. Finally, numerical comparisons with direct baselines demonstrate the superiority of our method in capturing instability and achieving improved identification accuracy.

A. Notation

We represent sequences as matrices with infinitely many rows: $\mathbf{x} = [x_0, x_1, \dots]^\top$, where $x_t \in \mathbb{R}^{d_x} \forall t \in \mathbb{N}$. The set of such sequences is ℓ^{d_x} , and $\mathbf{x} \in \ell_p^{d_x} \subset \ell^{d_x}$ for $p \in \mathbb{N}$ if \mathbf{x} has bounded p -norm: $\|\mathbf{x}\|_p = (\sum_{t=0}^{\infty} |x_t|^p)^{1/p} < \infty$. We denote truncation as $x_{j:i} = [x_i, \dots, x_j]^\top$ for $j \geq i$ and $x_{j:i} = \emptyset$ otherwise.

An operator $\mathbf{Y} : \ell^{d_x} \rightarrow \ell^{d_y}$ is a mapping between two sequences and is *causal* if $\mathbf{Y}(\mathbf{x}) = (\Upsilon_0(x_0), \dots, \Upsilon_t(x_{t:0}), \dots)$, and *strictly causal* if $\mathbf{Y}(\mathbf{x}) = (\Upsilon_0(\emptyset), \dots, \Upsilon_t(x_{t-1:0}), \dots)$, where $\Upsilon_0(\emptyset)$ is a vector in \mathbb{R}^{d_y} . A causal operator \mathbf{Y} is:

- \mathcal{L}_p -stable ($\mathbf{Y} \in \mathcal{L}_p$) if $\mathbf{Y}(\mathbf{x}) \in \ell_p^{d_y}$ for all $\mathbf{x} \in \ell_p^{d_x}$,
- incrementally \mathcal{L}_p -stable if $\exists \gamma \in [0, \infty)$ such that $\forall \mathbf{x}_1, \mathbf{x}_2 \in \ell_p^{d_x}$, $\|\mathbf{Y}(\mathbf{x}_1) - \mathbf{Y}(\mathbf{x}_2)\|_p \leq \gamma \|\mathbf{x}_1 - \mathbf{x}_2\|_p$, where γ is the incremental finite gain (i.f.g.) of \mathbf{Y} .

The sets of causal and strictly causal \mathcal{L}_p -stable operators are denoted \mathcal{L}_p^C and \mathcal{L}_p^{SC} , respectively.

II. PLANTS STABILIZED BY A GIVEN CONTROLLER

A. Problem setup

We consider a discrete-time, time-varying dynamical system that may be nonlinear and unstable, described by:

$$\begin{cases} y_t = G_t(u_{t-1:0}) + v_t, & t \in \mathbb{N}, \\ y_0 = G_0(\emptyset) + v_0, \end{cases} \quad (1)$$

with input $u_t \in \mathbb{R}^{d_u}$, nominal initial output $G_0(\emptyset) \in \mathbb{R}^{d_y}$, and output $y_t \in \mathbb{R}^{d_y}$. The output is corrupted by independent noise v_t . The system is strictly causal, meaning its output at time t depends only on past inputs up to time $t-1$.




Fig. 1: *Left:* The closed-loop of system \mathbf{G} and controller \mathbf{K} , with input \mathbf{u} , noisy output \mathbf{y} , excitation signal \mathbf{r} , and noise \mathbf{v} . *Right:* In the ICI framework, the model $\hat{\mathbf{G}}$ is an interconnection of a copy of the controller \mathbf{K} and a trainable operator $\hat{\mathbf{S}} \in \mathcal{L}_p^{SC}$. The model input and output are denoted by $\hat{\mathbf{u}}$ and $\hat{\mathbf{y}}^\circ$, respectively.

The model (1) is an input-output relationship, representing all strictly causal dynamical systems. The next example shows how to derive it from a state-space model.

Example 1 (From state-space to input-output models):

Consider a noise-free, strictly causal, time-invariant system in the state-space form:

$$\begin{cases} x_{t+1} = F(x_t, u_t), & x_0 = \bar{x}, \\ y_t = H(x_t), \end{cases}$$

where $x_t \in \mathbb{R}^{d_x}$ is the state at time t and $\bar{x} \in \mathbb{R}^{d_x}$ is the initial state. By recursion, the functions G_t in (1) are defined as:

$$\begin{aligned} G_0(\emptyset) &= H(\bar{x}), \quad G_1(u_0) = H(F(\bar{x}, u_0)), \dots \\ G_t(u_{t-1:0}) &= H(\underbrace{F(\dots(F(\bar{x}, u_0), u_1), \dots, u_{t-1}))}_{t \text{ times}}), \dots \end{aligned}$$

These expressions show that G_t depends explicitly on the initial state \bar{x} for all t , which is coherent with the definition of input-output models in the literature [13].

Moreover, linear operators with zero initial conditions represent transfer functions, highlighting their universality in characterizing LTI systems [12].

The system (1) can be expressed in operator form as:

$$\mathbf{y} = \mathbf{G}(\mathbf{u}) + \mathbf{v}, \quad (2)$$

where \mathbf{G} is a strictly causal operator such that $\mathbf{G}(\mathbf{u}) = [G_0(\emptyset), G_1(u_0), \dots]^\top$, $\mathbf{u} = [u_0, u_1, \dots]^\top$, etc. We assume the system is controlled by a causal dynamic output feedback controller \mathbf{K} . In addition to the feedback loop, a persistently exciting (PE) signal $r_t \in \mathbb{R}^{d_u}$ is applied to the system. The resulting closed-loop system, illustrated in the left panel of Fig. 1, is:

$$\begin{cases} \mathbf{y} = \mathbf{G}(\mathbf{u}) + \mathbf{v}, \\ \mathbf{u} = \mathbf{r} + \mathbf{K}(\mathbf{y}), \end{cases} \quad (3)$$

where $\mathbf{K}(\mathbf{y}) = [K_0(y_0), K_1(y_{1:0}), \dots]^\top$ and $\mathbf{r} = [r_0, r_1, \dots]^\top$. We assume the closed-loop system formed by \mathbf{G} and \mathbf{K} is \mathcal{L}_p -stable, as defined below.

Definition 1 (Closed-loop \mathcal{L}_p -stability [14]): The closed-loop of \mathbf{G} and \mathbf{K} , as shown in the left panel of Fig. 1, is \mathcal{L}_p -stable if for every $\mathbf{r} \in \ell_p$ and $\mathbf{v} \in \ell_p$, it holds that $\mathbf{u} \in \ell_p$ and $\mathbf{y} \in \ell_p$. Equivalently, the mapping $(\mathbf{r}, \mathbf{v}) \mapsto (\mathbf{u}, \mathbf{y})$ is an \mathcal{L}_p operator.

Our goal is to learn a dynamical system $\hat{\mathbf{G}}$ that approximates \mathbf{G} , given knowledge of the controller \mathbf{K} . Crucially,

we allow $\hat{\mathbf{G}}$ to be unstable, enabling the identification of an unstable \mathbf{G} , while ensuring that the closed-loop system formed by $\hat{\mathbf{G}}$ and \mathbf{K} remains stable. This requirement maintains consistency with the fact that \mathbf{K} stabilizes \mathbf{G} .

B. Internal controller identification framework

Learning over the space of systems $\hat{\mathbf{G}}$ stabilized by controller \mathbf{K} is challenging, as this space is generally nonconvex. To address this, we parameterize all and only systems stabilized by \mathbf{K} using a dual formulation of the Internal Model Control (IMC) framework [15], which we call Internal Controller Identification (ICI). Here, $\hat{\mathbf{G}}$ is structured as an interconnection of controller \mathbf{K} and a trainable strictly causal operator $\hat{\mathbf{S}} \in \mathcal{L}_p^{SC}$. This transforms the problem of learning $\hat{\mathbf{G}}$ into optimizing $\hat{\mathbf{S}}$ within \mathcal{L}_p^{SC} operators, enabling unconstrained optimization [12], [14]. The corresponding model is illustrated in the right panel of Fig. 1, where the noise-free output $\hat{y}_t^\circ \in \mathbb{R}^{d_y}$ for an input $\hat{u}_t \in \mathbb{R}^{d_u}$ is given by:

$$\hat{y}^\circ = \hat{\mathbf{S}}(\hat{\mathbf{u}} - \mathbf{K}(\hat{y}^\circ)). \quad (4)$$

Remark 1 (Recursive implementation of $\hat{\mathbf{S}}$): Since $\hat{\mathbf{S}}$ is strictly causal, the output \hat{y}_t° depends on past outputs only up to time $t - 1$, i.e., $\hat{y}_t^\circ = \hat{S}_t(\hat{u}_{t-1:0} - K_{t-1:0}(\hat{y}_{t-1:0}^\circ))$. This structure enables a recursive computation of \hat{y}_t° , eliminating the need to solve the implicit nonlinear equation (4) at each time step. Additionally, if $\hat{\mathbf{S}}$ and \mathbf{K} are input-output operators associated with dynamical systems (see Example 1), their state-space representations can be used to implement (4) as a Markovian dynamical system, thereby avoiding the need to store the full history of $\hat{u}_{t-1:0}$ and $\hat{y}_{t-1:0}^\circ$.

The following theorem establishes that a given stable controller, $\mathbf{K} \in \mathcal{L}_p$, can stabilize a system if and only if the system can be represented within the ICI framework.

Theorem 1: Let the controller $\mathbf{K} \in \mathcal{L}_p^C$ be an incrementally stable operator. Under the ICI framework, as shown in the right panel of Fig. 1 and described by (4), the following holds:

T.1 Let $\hat{\mathbf{S}}: \ell_p^{d_u} \rightarrow \ell_p^{d_y}$ be an \mathcal{L}_p^{SC} operator and consider the ICI model $\hat{\mathbf{G}}$ shown in the right plot of Fig. 1. Then, $\hat{\mathbf{G}}$ is stabilized by \mathbf{K} in the sense of Definition 1.

T.2 Let \mathbf{G} in Fig. 1 be a strictly causal system such that the closed-loop system is stable according to Definition 1. Then, there exists $\hat{\mathbf{S}} \in \mathcal{L}_p^{SC}$ such that $\mathbf{G} = \hat{\mathbf{G}}$, where $\hat{\mathbf{G}}$ is constructed as in the right panel of Fig. 1.

The proof can be found in the appendix of [16].

Since \mathbf{G} is stabilized by \mathbf{K} , by point T.2 of Theorem 1, it admits an ICI representation in terms of an operator $\hat{\mathbf{S}}$, which we refer to as the *true ICI* parameterization. The true operator $\hat{\mathbf{S}}$ serves as a benchmark for evaluating algorithms that learn $\hat{\mathbf{S}}$. Specifically, we assess a learning algorithm's effectiveness by its ability to recover $\hat{\mathbf{S}}$ in Sec. IV.

III. TRAINING PROCEDURE

In this section, we discuss learning $\hat{\mathbf{G}}$ from data to approximate the unknown system \mathbf{G} . The dataset is generated by applying $N \in \mathbb{N}$ excitation signals, each over a horizon $T \in \mathbb{N}$, to the closed-loop system in the left plot of Fig. 1,

and recording the resulting system inputs and noisy outputs, defined as:

$$\mathcal{D} = \{(r_{T-1:0}^n, u_{T-1:0}^n, y_{T-1:0}^n)\}_{n=0}^{N-1},$$

where $r_{T-1:0}^n$ represents the first T time steps of r in the experiment $n \in \{0, \dots, N-1\}$ (similarly for u and y).

Using this dataset, we explore three identification strategies: two direct and one indirect. These strategies aim to minimize the discrepancy between the measured and predicted outputs, $\{y_{T-1:0}^n\}_{n=0}^{N-1}$ and $\{\hat{y}_{T-1:0}^n\}_{n=0}^{N-1}$, based on the mean square error (MSE) cost function:

$$J(\mathcal{D}, \{\hat{y}_{T-1:0}^n\}_{n=0}^{N-1}) = \frac{1}{N} \sum_{n=0}^{N-1} \frac{1}{T} \sum_{t=0}^{T-1} \|y_t^n - \hat{y}_t^{\circ n}\|_2^2. \quad (5)$$

The key difference between the learning strategies lies in how the predicted output is obtained, as outlined below.

A. Direct identification

In direct identification, the measured inputs $\{u_{T-1:0}^n\}_{n=0}^{N-1}$ are applied to the model $\hat{\mathbf{G}}$, without considering the controller \mathbf{K} . This effectively treats the data as if collected in an open-loop setting. We use this approach to learn both a general $\hat{\mathbf{G}}$ (not necessarily in the ICI formulation) and $\hat{\mathbf{S}}$ within the ICI framework, leading to the following strategies.

Strategy 1: Direct identification of $\hat{\mathbf{G}}$, formulated as:

$$\begin{aligned} \hat{\mathbf{G}}^{dir-ID} &= \arg \min_{\hat{\mathbf{G}}} J(\mathcal{D}, \{\hat{y}_{T-1:0}^n\}_{n=0}^{N-1}), \\ &s.t. \hat{y}^{\circ n} = \hat{\mathbf{G}}(\mathbf{u}^n) \quad \forall n. \end{aligned}$$

The resulting model $\hat{\mathbf{G}}^{dir-ID}$ is not necessarily stabilized by \mathbf{K} , neglecting the knowledge that the true system \mathbf{G} is stabilized by \mathbf{K} . Nonetheless, we use this strategy as a baseline due to its widespread use. To ensure closed-loop stability, the next two strategies leverage the ICI framework.

Strategy 2: Direct identification of $\hat{\mathbf{S}}$ in the ICI framework, leading to:

$$\hat{\mathbf{S}}^{dir-ICI} = \arg \min_{\hat{\mathbf{S}} \in \mathcal{L}_p^{SC}} J(\mathcal{D}, \{\hat{y}_{T-1:0}^n\}_{n=0}^{N-1}),$$

$$s.t. \hat{y}^{\circ n} = \hat{\mathbf{S}}(\mathbf{u}^n - \mathbf{K}(\hat{y}^{\circ n})) \quad \forall n,$$

where the constraint follows from substituting $\hat{\mathbf{u}} = \mathbf{u}$ in (4). Recall that \hat{y}° can be computed recursively by Remark 1. The corresponding model, $\hat{\mathbf{G}}^{dir-ICI}$, is then defined as the closed-loop system of $\hat{\mathbf{S}}^{dir-ICI}$ and \mathbf{K} . Notably, $\hat{\mathbf{G}}^{dir-ICI}$ is stabilized by \mathbf{K} , as guaranteed by Theorem 1.

While direct methods lead to tractable optimization problems, they have two fundamental drawbacks.

D.1 Direct methods apply only when $\hat{\mathbf{G}}$ is stable [4] since the cost function J is evaluated using open-loop predictions of $\hat{\mathbf{G}}$, which may diverge if the model is unstable.

D.2 Direct methods can yield inconsistent estimates, meaning that the true \mathbf{G} is not recovered even with infinite data [6], [3], [17]. This inconsistency arises because the noise \mathbf{v} affects both the trainable operator's input, \mathbf{u} in **Strategy 1** and $\mathbf{u} - \mathbf{K}(\hat{y}^\circ)$ in **Strategy 2**, and the measured output \mathbf{y} , due to the feedback loop present during data collection, as illustrated in the left plot of Fig. 1. We further discuss this issue in Sec. IV.

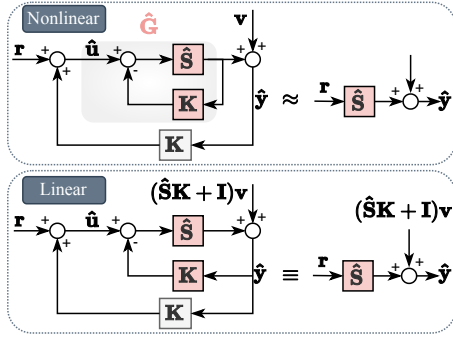


Fig. 2: *Top*: Closed-loop system of $\hat{\mathbf{G}}$ and \mathbf{K} (left) and its linearized approximation (right). *Bottom*: In the linear case, noise injection can be shifted to enable loop cancellation, yielding the exact simplified form on the right, consistent with (7).

B. Indirect identification

Indirect methods first identify the closed-loop system of $\hat{\mathbf{G}}$ and \mathbf{K} , using the excitation signal \mathbf{r} as input, as shown in the top-left plot of Fig. 2. They then leverage the known controller \mathbf{K} to extract a system model from the identified closed-loop system [6]. Since learning is based on the closed-loop system of $\hat{\mathbf{G}}$ and \mathbf{K} , these methods avoid direct open-loop simulation of $\hat{\mathbf{G}}$, thus addressing the problem **D.1**. Additionally, indirect methods assume a particular model parameterization in order to remove the noise contribution on the input of the trainable model, resolving issue **D.2** [3].

In the following, we develop an indirect strategy leveraging the ICI framework. The proposed scheme is inspired by indirect methods based on the dual Youla parameterization (see, e.g., [11]). As discussed in Sec. I, classical Youla-based methods are restricted to linear controllers, require intractable coprime factorizations for nonlinear systems, and involve operators that hinder computational optimization. By utilizing the ICI framework and recent techniques for parameterizing stable operators, we circumvent coprime factorizations and establish a computationally tractable framework.

We begin by computing the noisy output $\hat{\mathbf{y}}$ resulting from applying \mathbf{r} to the closed-loop system of $\hat{\mathbf{G}}$ and \mathbf{K} , as illustrated in the top-left plot of Fig. 2:

$$\begin{aligned}\hat{\mathbf{y}} &= \hat{\mathbf{S}}(\hat{\mathbf{u}} - \mathbf{K}(\hat{\mathbf{y}} - \mathbf{v})) + \mathbf{v} \\ &= \hat{\mathbf{S}}(\mathbf{r} + \mathbf{K}(\hat{\mathbf{y}}) - \mathbf{K}(\hat{\mathbf{y}} - \mathbf{v})) + \mathbf{v}.\end{aligned}\quad (6)$$

Next, following [11], we move the noise outside the operand of $\hat{\mathbf{S}}$ in (6) by linearizing $\hat{\mathbf{S}}$ and \mathbf{K} around \mathbf{r} and $\hat{\mathbf{y}}$, respectively. This results in:

$$\begin{aligned}\hat{\mathbf{y}} &\approx \hat{\mathbf{S}}(\mathbf{r}) + \hat{\mathbf{S}}_L(\mathbf{K}(\hat{\mathbf{y}}) - \mathbf{K}(\hat{\mathbf{y}} - \mathbf{v})) + \mathbf{v} \\ &\approx \underbrace{\hat{\mathbf{S}}(\mathbf{r})}_{\text{noise-independent signal}} + \underbrace{\hat{\mathbf{S}}_L(\mathbf{K}_L(\mathbf{v}))}_{\text{noise}} + \mathbf{v},\end{aligned}\quad (7)$$

where $\hat{\mathbf{S}}_L$ and \mathbf{K}_L denote the linearized versions of $\hat{\mathbf{S}}$ and \mathbf{K} , around \mathbf{r} and $\hat{\mathbf{y}}$, respectively. If $\hat{\mathbf{S}}$ and \mathbf{K} are linear, the approximation holds exactly, with $\hat{\mathbf{S}}_L = \hat{\mathbf{S}}$ and $\mathbf{K}_L = \mathbf{K}$.

Remark 2 (Validity of linear approximations): The above linearization of \mathbf{K} holds well at high output signal-to-noise

ratio (SNR), where $\hat{\mathbf{y}}$ has significantly higher power than \mathbf{v} .¹ Similarly, $\hat{\mathbf{S}}$ is well-approximated when \mathbf{r} has a much higher power than $\mathbf{K}(\hat{\mathbf{y}}) - \mathbf{K}(\hat{\mathbf{y}} - \mathbf{v})$. Since \mathbf{K} is incrementally stable, this condition simplifies to a high excitation SNR, where the power of \mathbf{r} is much greater than that of \mathbf{v} .

According to (7), the closed-loop response of the system formed by $\hat{\mathbf{G}}$ and \mathbf{K} to \mathbf{r} can be approximated by passing \mathbf{r} directly through $\hat{\mathbf{S}}$. Intuitively, the internal and external feedback loops nearly cancel out, allowing \mathbf{r} to propagate primarily through $\hat{\mathbf{S}}$ (see the top-right plot of Fig. 2). This approximation is exact in two cases. First, in the absence of noise, as is immediately evident from the top-left plot of Fig. 2. Second, when $\hat{\mathbf{S}}$ and \mathbf{K} are linear, rearranging the top-left plot in Fig. 2 by shifting the noise yields the bottom-left plot. As shown, the feedback loops cancel out, leading to the simplified diagram on the right, which aligns with (7).

In light of (7), we introduce the third strategy.

Strategy 3: Indirect identification of $\hat{\mathbf{S}}$ in the ICI framework via output linearization, formulated as:

$$\begin{aligned}\hat{\mathbf{S}}^{\text{indir-ICI}} &= \arg \min_{\hat{\mathbf{S}} \in \mathcal{L}_p^{\text{SC}}} J(\mathcal{D}, \{\mathcal{Y}_{T-1:0}^{\text{on}}\}_{n=0}^{N-1}), \\ &\text{s.t. } \hat{\mathbf{y}}^{\text{on}} = \hat{\mathbf{S}}(\mathbf{r}^n) \quad \forall n.\end{aligned}$$

The corresponding model $\hat{\mathbf{G}}^{\text{indir-ICI}}$ is then defined as the closed-loop system formed by $\hat{\mathbf{S}}^{\text{indir-ICI}}$ and \mathbf{K} .

Remark 3 (Open-loop equivalence): We observe from (7) that (i) $\hat{\mathbf{y}}$ can be approximated as $\hat{\mathbf{S}}$ applied to \mathbf{r} , which is independent of the noise, plus an additive noise term, and (ii) $\hat{\mathbf{S}}$ is a stable operator. As a result, **Strategy 3** is a *standard open-loop* problem, eliminating the instability (**D.1**) and correlation (**D.2**) issues of direct methods. This equivalence is well-known for classic indirect approaches, such as the dual Youla parameterization [4].

Following Remark 3, we can use any open-loop method, such as subspace identification or instrumental variables (see [6] for an overview), to obtain the system model. Here, we minimize the least squares cost in (5), aligning with the prediction error method [6]. Moreover, Remark 3 allows us to directly leverage the consistency properties of open-loop identification, as discussed in Sec. IV.

C. Parameterizing the operator $\hat{\mathbf{S}}$

To implement the ICI framework, we must parameterize the operator $\hat{\mathbf{S}} \in \mathcal{L}_p^{\text{SC}}$. This can be achieved using various families of parameterized $\mathcal{L}_p^{\text{SC}}$ operators, such as Recurrent Equilibrium Networks (RENs) [19], certain classes of State-Space Models (SSMs) [20], and neural network representations of Hamiltonian systems [21]. These operators all admit a state-space form, enabling recursive computation of the output of $\hat{\mathbf{G}}$ (see Remark 1). In our experiments, we follow [14], [22] and employ RENs, which define a broad class of \mathcal{L}_2 operators and embed arbitrarily deep neural networks [19].

For any chosen operator family, we denote the parameterized operator and the associated model by $\hat{\mathbf{S}}^\theta$ and $\hat{\mathbf{G}}^\theta$, with

¹The power of a signal \mathbf{y} is defined as $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \|\mathbf{y}_t\|_2^2$ [18].

learnable parameters $\theta \in \mathbb{R}^{d_\theta}$. Importantly, these families are constructed so that $\hat{\mathbf{S}}^\theta \in \mathcal{L}_p^{SC}$ for all $\theta \in \mathbb{R}^{d_\theta}$. Consequently, optimizing over $\hat{\mathbf{S}} \in \mathcal{L}_p^{SC}$ in **Strategy 2** and **Strategy 3** reduces to an unconstrained optimization over $\theta \in \mathbb{R}^{d_\theta}$, which can be efficiently solved using gradient descent methods.

IV. STATISTICAL ANALYSIS IN THE LINEAR CASE

The identification procedure is inherently stochastic, as it depends on the realizations of noise and excitation signals in the dataset. A central question is whether the learned models $\hat{\mathbf{G}}$ and $\hat{\mathbf{S}}$ asymptotically converge in probability to the true system \mathbf{G} or the true \mathbf{S} , defined in Sec. II-B, as the sample size grows. This notion of convergence is formalized through the concept of consistency:²

Definition 2 (Consistency [3]): An estimate $\hat{\mathbf{G}}$ (or $\hat{\mathbf{S}}$) obtained based on a dataset of N sequences, each of horizon T , is consistent if it converges to the true value \mathbf{G} (or \mathbf{S}) with probability 1 as the sample size grows indefinitely, i.e.,

$$\lim_{NT \rightarrow \infty} \hat{\mathbf{G}} = \mathbf{G}, \quad \text{or} \quad \lim_{NT \rightarrow \infty} \hat{\mathbf{S}} = \mathbf{S}.$$

Since $\hat{\mathbf{S}}$ and $\hat{\mathbf{G}}$ can be uniquely determined from one another, the consistency of $\hat{\mathbf{S}}$ is equivalent to that of $\hat{\mathbf{G}}$.

We discuss the consistency of the proposed strategies when \mathbf{K} , \mathbf{G} , and $\hat{\mathbf{G}}$ are linear. Throughout this section, we assume that the excitation signal is quasi-stationary and sufficiently rich to identify the plant dynamics [23], [4].

The consistency of direct methods (**Strategy 1** and **Strategy 2**) can only be analyzed for a stable $\hat{\mathbf{G}}$ (see **D.1**). Even in this case, it depends heavily on the noise model. Let the noise \mathbf{v} in (2) be expressed as $\mathbf{v} = \mathbf{H}(\mathbf{e})$, where \mathbf{e} is white noise and \mathbf{H} is a linear noise model. Direct methods typically estimate both the target system \mathbf{G} and the operator \mathbf{H} . It is well-established that a consistent estimate of \mathbf{G} can only be obtained if \mathbf{H} is either known or exactly learned [3], [4], [5], [24]. However, many real-world disturbances exhibit colored noise characteristics [25], which cannot be captured by a finite-order noise model \mathbf{H} . As a result, direct identification methods are often inconsistent in practice [24], [5].

Next, we discuss the consistency of **Strategy 3**. As noted in Remark 3, **Strategy 3** is equivalent to identifying $\hat{\mathbf{S}}$ in an open-loop setting. Consequently, consistency results from open-loop identification apply directly, as stated below.

Proposition 1 (Corollary 6 in [4]): Consider the stable closed-loop system given by (3), where the system \mathbf{G} and the controller \mathbf{K} are linear. Suppose the dataset \mathcal{D} consists of N sequences, each of horizon T , collected under quasi-stationary persistent excitation \mathbf{r} . Let $\hat{\mathbf{S}}^\theta$ be the estimate obtained by **Strategy 3** using \mathcal{D} , optimized over a parametric family of \mathcal{L}_p^{SC} operators. Then, both $\hat{\mathbf{S}}^\theta$ and the corresponding $\hat{\mathbf{G}}^\theta$ are consistent, provided that the true \mathbf{S} , as introduced in Sec. II-B, belongs to the chosen family.

The consistency of **Strategy 3** only requires the noise to be independent of the excitation signal, without imposing any additional assumptions on the noise.

²In some classic system identification literature, the term unbiasedness is also used to describe consistency (e.g., [4]). To align with machine learning terminology, we use consistency throughout.

Proposition 1 requires that the true \mathbf{S} belongs to the chosen parametric set. If this set covered the entire \mathcal{L}_p^{SC} space, consistency would be guaranteed. However, since most parametric families only span a subset of it, \mathbf{S} may fall outside, leading to inconsistency. This is a common trait in the analysis of learning algorithm consistency [4], which can be addressed by using a more flexible parameterization or increasing the dimensionality of the parameters, d_θ .

V. EXPERIMENTS

We consider two nonlinear systems: *i*) an unstable scalar system stabilized with nonlinear feedback, and *ii*) a stable planar robot, controlled to navigate to a target location. We compare our framework (**Strategy 3**) against the baselines (**Strategy 1** and **Strategy 2**), showing the effectiveness of our approach in both stable and unstable settings.³

Training datasets. For both experiments, the training dataset \mathcal{D} consists of $N = 40$ trajectories, each starting from a fixed initial condition, with a horizon of $T = 100$. The system input is given by $u_t = K_t(y_{t:0}) + r_t$, where the controller K_t will be specified for each experiment. The excitation signal r_t is independently sampled at each time step from a zero-mean Gaussian distribution, $r_t \sim \mathcal{N}(0, \sigma^2 I)$, ensuring sufficient exploration of the system dynamics [23]. The value of σ varies across experiments.

Evaluation criteria. We assess all strategies on a test set of 100 trajectories of length $T = 100$, using the same excitation and noise distributions as in training. Four criteria are used: *open-loop (OL) MSE*, *OL coefficient of determination (R^2)*, *closed-loop (CL) MSE*, and *CL R^2* . In the open-loop setup, the excitation signal is applied to both the true system and the learned model, and the MSE between their outputs is computed. Since MSE is scale-dependent, we also report the coefficient of determination (R^2), which provides a standardized measure of how well the model explains the variance in the dependent variable [26]: $R^2 = 1 - \frac{\sum_{n=1}^N \sum_{t=1}^T \|y_t^n - \hat{y}_t^n\|_2^2}{\sum_{n=1}^N \sum_{t=1}^T \|y_t^n - \bar{y}\|_2^2}$. It ranges from $-\infty$ to 1 (values closer to 1 indicate a better fit) and \bar{y} is the overall mean of y .

While OL criteria directly measure model accuracy, they do not apply to unstable systems, as their open-loop outputs diverge. Therefore, we also assess closed-loop performance. The same excitation signals and noise sequences are applied to both the true closed-loop system (comprising \mathbf{G} and controller \mathbf{K}) and the learned closed-loop system (formed by $\hat{\mathbf{G}}$ and \mathbf{K}), and the CL MSE and R^2 are computed. Since closed-loop outputs remain bounded when the controller stabilizes the system, CL criteria allow meaningful comparisons even for unstable plants. In the sequel, we repeat all experiments with 50 different random seeds and report 95% confidence intervals for each evaluation criterion.

Model details. We instantiate $\hat{\mathbf{G}}$ in **Strategy 1** and $\hat{\mathbf{S}}$ in **Strategy 2** and **Strategy 3** as RENs [19]. In all cases, the REN architecture comprises a state vector in \mathbb{R}^8 and 8 hidden

³The code is available on https://github.com/DecodEPFL/Nonlinear_system_identification.

layers, resulting in a total of 744 trainable parameters. This configuration provides sufficient modeling flexibility.

A. Scalar unstable system

We consider the following state-space scalar system:

$$\begin{cases} x_{t+1} = x_t^2 + 1 + u_t, & x_0 = \bar{x}, \\ y_t = x_t + v_t, \end{cases} \quad (8)$$

where x_t is the state and the excitation signal standard deviation is $\sigma = 0.5$. The noise v_t is sampled from a truncated normal distribution with a zero mean, standard deviation of 0.1, and bounded in the interval $(-0.25, 0.25)$.

The system in (8) is unstable, making closed-loop identification necessary. We use the controller $K_t(y_{t:0}) = -y_t^2 - 1 + 0.5y_t$, resulting in the closed-loop system:

$$\begin{cases} x_{t+1} = (0.5 - 2v_t)x_t + 0.5v_t - v_t^2 + r_t, & x_0 = \bar{x}, \\ y_t = x_t + v_t. \end{cases}$$

Since $v_t \in (-0.25, 0.25)$, the output converges on average for any initial condition and does not grow infinitely.

Results. Since the system is unstable, an unstable model $\hat{\mathbf{G}}$ must be used, making the direct strategies (**Strategy 1** and **Strategy 2**) inapplicable (see **D.1**). For **Strategy 3**, we compare the OL and CL responses against the true system in Fig. 3. In all plots, solid lines represent the average output over noise and excitation realizations, while shaded regions indicate the confidence bounds. The OL plots highlight **Strategy 3**'s ability to capture unstable dynamics, as the model's outputs diverge significantly within just 5 time steps. The OL confidence bounds are not visible due to the large output scale. As a result of these extreme values, OL criteria are not computable. The CL plots show a high similarity between the CL trajectories of the true and identified models, with a low CL MSE of 0.0034 ± 0.0006 . We also observe that the output converges to zero. This experiment showcases the effectiveness of our proposed approach in learning an unstable system.

B. Point-mass robot

We follow the experimental setup of [14], modeling a planar robot as a point mass $m \in \mathbb{R}_+$. The state is the robot's position $x_{1,t} \in \mathbb{R}^2$ and velocity $x_{2,t} \in \mathbb{R}^2$, while the output is a noisy position measurement. The input u_t is a planar force driving the motion according to the state-space equations:

$$\begin{cases} \begin{bmatrix} x_{1,t+1} \\ x_{2,t+1} \end{bmatrix} = \begin{bmatrix} x_{1,t} \\ x_{2,t} \end{bmatrix} + T_s \begin{bmatrix} x_{2,t} \\ \frac{1}{m}(-C(x_{2,t}) + u_t) \end{bmatrix}, \\ y_t = x_{1,t} + v_t, \end{cases}$$

where $T_s > 0$ is the sampling time and $v_t \sim \mathcal{N}(0, 0.1)$. Moreover, the function $C: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ models nonlinear drag forces, such as air or water resistance. We set $C(x_2) = b_1 x_2 - b_2 \|x_2\|_2 x_2$ for some $0 < b_2 < b_1$.

This system is stable due to the presence of drag forces. However, we assume the robot is controlled by a proportional controller that drives it toward a predefined target position $y^* = 0 \in \mathbb{R}^2$, potentially for economic reasons. The




Fig. 3: Comparison of system and model outputs in the scalar unstable experiment. *Top*: true system output; *Bottom*: output of the model learned by **Strategy 3**; *Left*: open-loop; *Right*: closed-loop. Solid lines show the average output across multiple realizations of noise and excitation signals and shaded regions indicate 95% confidence bounds. The open-loop confidence bounds are not visible due to their much smaller scale relative to the output. The learned model mirrors the true system's stability properties: it is unstable in open-loop but stable in closed-loop. Closed-loop responses closely match, with a low MSE of 0.0034 ± 0.0006 .

control law is given by $K_t(y_{t:0}) = \text{diag}(\kappa_1, \kappa_2)(y^* - y_t)$, with $\kappa_1, \kappa_2 \in \mathbb{R}_+$. Specific parameter values for the system and the controller ($m, b_1, b_2, \kappa_1, \kappa_2$) are provided in the codebase.

We analyze the impact of the excitation signal's standard deviation, σ . First, we consider a small $\sigma = 10$. The training trajectories are shown in Fig. 4, where the robot starts with an initial velocity of $[10, 0]^\top$. With low excitation, the trajectories remain smooth, allowing the robot to reach its target with minimal deviation. Next, we apply a high $\sigma = 50$. While larger excitation improves identification (as shown in the results section), Fig. 4 illustrates that it also introduces undesired deviations during the robot's movement.

Results. Since the system is stable, we can compare all strategies without encountering the instability issue noted in **D.1**. The results are summarized in Tab. I. For $\sigma = 10$, the excitation signal has less influence on the system input \mathbf{u} than the feedback controller, resulting in a less informative \mathbf{u} . As direct methods (**Strategy 1** and **Strategy 2**) rely on \mathbf{u} , they perform poorly under such conditions. In contrast, **Strategy 3** remains effective, as it identifies the system based on \mathbf{r} instead of \mathbf{u} . It can be seen from Tab. I that **Strategy 3** achieves lower OL and CL MSE and higher R^2 scores.

For $\sigma = 50$, all strategies are expected to improve due to increased excitation, which is confirmed by the increase in OL and CL R^2 . Note that MSE is not suitable for assessing improvement in this case, as it is sensitive to the scale of the output signal, which increases significantly between experiments. While stronger excitation benefits all methods, **Strategy 3** consistently outperforms the baselines across all evaluation criteria, further demonstrating its effectiveness.

VI. CONCLUSIONS AND FUTURE WORK

This paper presents a novel computationally tractable framework for closed-loop identification of nonlinear dynamical systems using a complete parameterization of sys-

TABLE I: Evaluation of different strategies on the point-mass robot experiment, considering two excitation signal levels: $\sigma = 10$ and $\sigma = 50$. Each experiment is repeated 50 times with different random seeds, and results are shown with 95% confidence intervals. The proposed **Strategy 3** consistently outperforms baselines across all settings and metrics. Colored cells indicate the best results.

	Strategy	OL MSE	CL MSE	OL R ²	CL R ²
$\sigma = 10$	Strategy 1	17.6847 \pm 0.0992	0.4800 \pm 0.0090	0.8368 \pm 0.0007	0.9902 \pm 0.0002
	Strategy 2	11.5468 \pm 0.0674	0.3916 \pm 0.0086	0.8934 \pm 0.0005	0.9920 \pm 0.0002
	Strategy 3	6.7351 \pm 0.0586	0.2398 \pm 0.0049	0.9378 \pm 0.0004	0.9951 \pm 0.0001
$\sigma = 50$	Strategy 1	16.2409 \pm 0.4997	6.2671 \pm 0.1598	0.8842 \pm 0.0029	0.9122 \pm 0.0019
	Strategy 2	3.8891 \pm 0.1109	1.8150 \pm 0.0258	0.9723 \pm 0.0006	0.9746 \pm 0.0003
	Strategy 3	2.6998 \pm 0.0567	1.3535 \pm 0.0209	0.9807 \pm 0.0003	0.9810 \pm 0.0002




Fig. 4: Training trajectories in the robot experiment. *Left*: excitation standard deviation $\sigma = 10$; *Right*: $\sigma = 50$. The robot has an initial velocity of $[10, 0]^T$ and operates in a loop with a controller that drives it to the origin. As expected, the robot deviates less from its intended trajectory under lower excitation, and more significantly when the excitation is higher. Colors indicate time progress.

tems stabilized by a given controller. The proposed method addresses key challenges, including instability risks and computational complexity. We theoretically show asymptotic consistency in the linear setting, aligning with established results. Numerical experiments demonstrate that our method outperforms existing work by accurately capturing system instabilities and achieving superior prediction performance.

Future work includes conducting a non-asymptotic statistical analysis in the spirit of [27], applying the identified model to adaptive control, and designing an iterative procedure that integrates system identification and control.

REFERENCES

- [1] J. Sjöberg, H. Hjalmarsson, and L. Ljung, “Neural networks in system identification,” *IFAC Proceedings Volumes*, vol. 27, no. 8, pp. 359–382, 1994.
- [2] G. Pillonetto, A. Aravkin, D. Gedon, L. Ljung, A. H. Ribeiro, and T. B. Schön, “Deep networks for system identification: A survey,” *Automatica*, vol. 171, p. 111907, 2025.
- [3] P. Van den Hof, “Closed-loop issues in system identification,” *Annual Reviews in Control*, vol. 22, pp. 173–186, 1998.
- [4] U. Forssell and L. Ljung, “Closed-loop identification revisited,” *Automatica*, vol. 35, no. 7, pp. 1215–1241, 1999.
- [5] I. Maruta and T. Sugie, “A simple framework for identifying dynamical systems in closed-loop,” *IEEE Access*, vol. 9, pp. 31 441–31 453, 2021.
- [6] T. Söderström and P. Stoica, *System identification*. Prentice-Hall, Inc., 1988.
- [7] R. Chen, A. Srivastava, M. Yin, and R. S. Smith, “Closed-loop identification of stabilized models using dual input–output parameterization,” *European Journal of Control*, vol. 80, p. 101089, 2024.
- [8] M. Abdalmoaty and R. S. Smith, “Small noise analysis of non-parametric closed-loop identification,” in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 2024, pp. 7988–7993.
- [9] F. Hansen, G. Franklin, and R. Kosut, “Closed-loop identification via the fractional representation: Experiment design,” in *1989 American Control Conference*, 1989, pp. 1422–1427.

- [10] D. Youla, H. Jabr, and J. Bongiorno, “Modern Wiener-Hopf design of optimal controllers—part ii: The multivariable case,” *IEEE Transactions on Automatic Control*, vol. 21, no. 3, pp. 319–338, 1976.
- [11] B. D. Anderson, “From Youla–Kucera to identification, adaptive and nonlinear control,” *Automatica*, vol. 34, no. 12, pp. 1485–1506, 1998.
- [12] C. L. Galimberti, L. Furieri, and G. Ferrari-Trecate, “Parametrizations of all stable closed-loop responses: From theory to neural network control design,” *Annual Reviews in Control*, vol. 60, p. 101012, 2025.
- [13] A. Paice and A. van der Schaft, “The class of stabilizing nonlinear plant controller pairs,” *IEEE Transactions on Automatic Control*, vol. 41, no. 5, pp. 634–645, 1996.
- [14] L. Furieri, C. L. Galimberti, and G. Ferrari-Trecate, “Learning to boost the performance of stable nonlinear systems,” *IEEE Open Journal of Control Systems*, vol. 3, pp. 342–357, 2024.
- [15] C. Economou, M. Morari, and B. Palsson, “Internal model control: Extension to nonlinear system,” *Industrial and Engineering Chemistry Process Design and Development*, vol. 25, no. 2, pp. 403–411, 1986.
- [16] M. G. Boroujeni, L. Meroi, L. Massai, C. L. Galimberti, and G. Ferrari-Trecate, “Neural identification of feedback-stabilized nonlinear systems,” 2025. [Online]. Available: <https://arxiv.org/abs/2503.22601>
- [17] G. C. Goodwin and J. S. Welsh, “Bias issues in closed loop identification with application to adaptive control,” *Communications in Information and Systems*, vol. 2, no. 4, pp. 349–370, 2002.
- [18] A. V. Oppenheim, A. S. Willsky, and I. T. Young, *Signals and systems*. Prentice-Hall, 1983.
- [19] M. Revay, R. Wang, and I. R. Manchester, “Recurrent equilibrium networks: Flexible dynamic models with guaranteed stability and robustness,” *IEEE Transactions on Automatic Control*, vol. 69, no. 5, pp. 2855–2870, 2023.
- [20] A. Gu, K. Goel, and C. Ré, “Efficiently modeling long sequences with structured state spaces,” in *International Conference on Learning Representations*, 2022.
- [21] M. Zakwan and G. Ferrari-Trecate, “Neural port-Hamiltonian models for nonlinear distributed control: An unconstrained parametrization approach,” *arXiv preprint arXiv:2411.10096*, 2024.
- [22] M. G. Boroujeni, C. L. Galimberti, A. Krause, and G. Ferrari-Trecate, “A PAC-Bayesian framework for optimal control with stability guarantees,” in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, 2024, pp. 8237–8244.
- [23] L. Ljung, *System Identification: Theory for the User*. Prentice Hall PTR, 1999.
- [24] L. Ljung and U. Forssell, “An alternative motivation for the indirect approach to closed-loop identification,” *IEEE Transactions on Automatic Control*, vol. 44, no. 11, pp. 2206–2209, 1999.
- [25] P. Bak, C. Tang, and K. Wiesenfeld, “Self-organized criticality: An explanation of the $1/f$ noise,” *Physical Review Letters*, vol. 59, no. 4, p. 381, 1987.
- [26] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: With Applications in R*. Springer, 2021.
- [27] D. Jones and M. A. Dahleh, “Closed loop system identification with known feedback: A non asymptotic viewpoint,” in *2022 American Control Conference (ACC)*, 2022, pp. 4094–4100.
- [28] D. Ho, “A system level approach to discrete-time nonlinear systems,” in *2020 American Control Conference (ACC)*, 2020, pp. 1625–1630.
- [29] C. A. Desoer and R.-W. Liu, “Global parametrization of feedback systems with nonlinear plants,” *Systems & Control Letters*, vol. 1, no. 4, pp. 249–251, 1982.

A. Proof of Theorem 1

We first state a useful corollary from [28].

Corollary 1: Consider a causal operator Υ such that $\Upsilon^\circ := \Upsilon - \mathbf{I}$ is strictly causal, where \mathbf{I} is the identity operator. Then, the causal operator Υ^{-1} exists, and the sequence $\mathbf{b} = \Upsilon^{-1}\mathbf{a}$ satisfies, for $t = 0, 1, \dots$:

$$b_t = a_t - \Upsilon_t^\circ(b_{t-1:0}).$$

We start by proving **T.1**, meaning that for any $\mathbf{v} \in \ell_p^{d_y}$ and $\mathbf{r} \in \ell_p^{d_u}$, one has $\hat{\mathbf{y}} \in \ell_p^{d_y}$ and $\hat{\mathbf{u}} \in \ell_p^{d_u}$. We denote the input to the operator $\hat{\mathbf{S}}$ by $\hat{\boldsymbol{\omega}}$. Then, from the right panel of Fig. 1,

$$\hat{\boldsymbol{\omega}} = \hat{\mathbf{u}} - \mathbf{K}(\hat{\mathbf{y}}^\circ) = \mathbf{r} + \mathbf{K}(\hat{\mathbf{S}}(\hat{\boldsymbol{\omega}}) + \mathbf{v}) - \mathbf{K}(\hat{\mathbf{S}}(\hat{\boldsymbol{\omega}})).$$

We now upper bound its norm:

$$\begin{aligned} \|\hat{\boldsymbol{\omega}}\|_p &= \|\mathbf{r} + \mathbf{K}(\hat{\mathbf{S}}(\hat{\boldsymbol{\omega}}) + \mathbf{v}) - \mathbf{K}(\hat{\mathbf{S}}(\hat{\boldsymbol{\omega}}))\|_p \\ &\leq \|\mathbf{r}\|_p + \|\mathbf{K}(\hat{\mathbf{S}}(\hat{\boldsymbol{\omega}}) + \mathbf{v}) - \mathbf{K}(\hat{\mathbf{S}}(\hat{\boldsymbol{\omega}}))\|_p \\ &\leq \|\mathbf{r}\|_p + \gamma_K \|\hat{\mathbf{S}}(\hat{\boldsymbol{\omega}}) + \mathbf{v} - \hat{\mathbf{S}}(\hat{\boldsymbol{\omega}})\|_p = \|\mathbf{r}\|_p + \gamma_K \|\mathbf{v}\|_p, \end{aligned}$$

where the second line is due to the triangle inequality and $\gamma_K > 0$ is the i.f.g. of \mathbf{K} . Since $\mathbf{v} \in \ell_p^{d_y}$ and $\mathbf{r} \in \ell_p^{d_u}$, we have $\hat{\boldsymbol{\omega}} \in \ell_p^{d_u}$. Additionally, since $\hat{\mathbf{S}} \in \mathcal{L}_p$, $\hat{\mathbf{y}} = \hat{\mathbf{S}}(\hat{\boldsymbol{\omega}}) + \mathbf{v} \in \ell_p^{d_y}$. Consequently, since $\mathbf{K} \in \mathcal{L}_p$, $\hat{\mathbf{u}} = \mathbf{K}(\hat{\mathbf{y}}) + \mathbf{r} \in \ell_p^{d_u}$. Therefore, the mapping from (\mathbf{r}, \mathbf{v}) to $(\hat{\mathbf{u}}, \hat{\mathbf{y}})$ is an \mathcal{L}_p operator.

Next, we prove **T.2**. Given a strictly causal system $\hat{\mathbf{G}}$ stabilized by \mathbf{K} , define $\hat{\mathbf{S}} = \hat{\mathbf{G}}(\mathbf{I} - \mathbf{K}\hat{\mathbf{G}})^{-1}$. Note that the

inverse of $(\mathbf{I} - \mathbf{K}\hat{\mathbf{G}})$ exists due to Corollary 1. Additionally, set the initial condition of $\hat{\mathbf{S}}$ such that:

$$\hat{\mathbf{S}}_0(\emptyset - \mathbf{K}_0(\emptyset)) = \hat{\mathbf{G}}_0(\emptyset). \quad (9)$$

In [29], it has been shown that $\hat{\mathbf{S}} \in \mathcal{L}_p^{SC}$.

We apply the same input $\hat{\mathbf{u}}$ to both $\hat{\mathbf{G}}$ and the closed-loop system formed by $\hat{\mathbf{S}}$ and \mathbf{K} , as illustrated in the right panel of Fig. 1, and denote their respective outputs by $\hat{\mathbf{y}}^{\circ des}$ and $\hat{\mathbf{y}}^\circ$. Our goal is to prove that $\hat{\mathbf{y}}^{\circ des} = \hat{\mathbf{y}}^\circ$, or equivalently, $\hat{y}_t^{\circ des} = \hat{y}_t^\circ$ for all $t \geq 0$. At $t = 0$, this is ensured by (9).

First, we compute the outputs. For $\hat{\mathbf{G}}$, we have $\hat{\mathbf{y}}^{\circ des} = \hat{\mathbf{G}}(\hat{\mathbf{u}})$. From (4), we obtain that:

$$\hat{\mathbf{y}}^\circ = \hat{\mathbf{S}}(\hat{\mathbf{u}} - \mathbf{K}(\hat{\mathbf{y}}^\circ)) = \hat{\mathbf{G}} \underbrace{(\mathbf{I} - \mathbf{K}\hat{\mathbf{G}})^{-1}(\hat{\mathbf{u}} - \mathbf{K}(\hat{\mathbf{y}}^\circ))}_{\mathbf{b}} = \hat{\mathbf{G}}\mathbf{b}.$$

Converting the above into time equations gives:

$$\hat{y}_t^{\circ des} = \hat{G}_t(\hat{u}_{t-1:0}) \quad , \quad \hat{y}_t^\circ = \hat{G}_t(b_{t-1:0}), \quad (10)$$

where we use the strong causality of $\hat{\mathbf{G}}$ to truncate the operand at time $t - 1$. Additionally, b_t can be computed using the recursive algorithm of Corollary 1 as:

$$b_t = \hat{u}_t - K_t(\hat{y}_{t:0}^\circ) + K_t(\hat{G}_{t:0}(b_{t-1:0})) \quad \forall t \geq 0. \quad (11)$$

From (10), one has $\hat{y}_t^{\circ des} = \hat{y}_t^\circ$ for all $t \geq 0$ if $\hat{u}_t = b_t$ for all $t \geq 0$. By (11), this holds if $K_t(\hat{y}_{t:0}^\circ) = K_t(\hat{G}_{t:0}(b_{t-1:0}))$ for all $t \geq 0$. Since $\hat{y}_t^\circ = \hat{G}_{t:0}(b_{t-1:0})$ from (10), the condition is satisfied, thus completing the proof.