

Chaos-Free Networks Are Stable Recurrent Neural Networks

Stefano De Carli^{ID}, Davide Previtali^{ID}, Mirko Mazzoleni^{ID}, *Associate Member, IEEE*,
and Fabio Previdi^{ID}, *Member, IEEE*

Abstract—Gated Recurrent Neural Networks (RNNs) are widely used for nonlinear system identification due to their high accuracy, although they often exhibit complex, chaotic dynamics that are difficult to analyze. This letter investigates the system-theoretic properties of the Chaos-Free Network (CFN), an architecture originally proposed to eliminate the chaotic behavior found in standard gated RNNs. First, we formally prove that the CFN satisfies Input-to-State Stability (ISS) by design. However, we demonstrate that the CFN architecture does not intrinsically guarantee Incremental ISS (δ ISS), as ensuring this property relies on specific parametric constraints. To address this, we introduce the Decoupled-Gate Network (DGN), a novel structural variant of the CFN that removes internal state connections in the gating mechanisms. Finally, we prove that the DGN unconditionally satisfies the δ ISS property, providing an incrementally stable architecture for identifying nonlinear dynamical systems without requiring complex network training modifications. Numerical results confirm that the DGN maintains the modeling capabilities of standard architectures while adhering to these rigorous stability guarantees.

Index Terms—Stability of nonlinear systems, nonlinear systems identification, neural networks.

I. INTRODUCTION

RECURRENT Neural Networks (RNNs) have become a standard tool for the identification of nonlinear dynamical systems [1]. While gated architectures, such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks [2, Ch. 15], demonstrate superior accuracy when it comes to modeling real systems [1], their deployment in safety-critical control frameworks is limited by their complexity. Specifically, standard gated RNNs can exhibit complex, irregular dynamics such as non-divergent but non-vanishing oscillations or chaotic behaviors, which are heavily influenced by the networks' initial conditions [3]. Further, the lack of intrinsic stability guarantees makes the derivation of theoretically sound RNN-based control strategies

challenging [4]. Consequently, the analysis of RNN stability has received significant attention in recent years [5], [6], [7], [8], particularly regarding *Input-to-State Stability* (ISS) and *Incremental Input-to-State Stability* (δ ISS), which are crucial for robust control frameworks and observer designs [4], [9], [10]. In parallel with these control-theoretic developments, the machine learning community has actively investigated the equilibrium properties of neural networks [11], [12] to ensure well-behaved forward propagation. In this context, the *Chaos-Free Network* (CFN) was proposed in the machine learning literature [3] to avoid the emergence of *strange* RNN state attractors, which are chaotic fractal structures typical of standard gated networks [3], [13, Ch. 9], ensuring instead that the CFN's state relaxes to a stable equilibrium when no input sequence is applied to it. While originally motivated by qualitative dynamical behavior, rigorous system-theoretic properties of the CFN remain unexplored.

In this work, we analyze the CFN through the lens of control theory. First, we formally prove that the absence of chaos in CFNs translates to the satisfaction of the ISS property by design. Second, we identify that the recurrent gating structure in CFNs does not yield an unconditional δ ISS guarantee by design. Indeed, promoting δ ISS requires enforcing complex parametric conditions during CFN training, analogous to those established for LSTM and GRU networks [5], [6], [8]. To address this, inspired by sequence modeling trends that favor simplified gating mechanisms for computational efficiency [14], we propose the *Decoupled-Gate Network* (DGN), a structural variant of the CFN that removes the state-dependent recurrence in the gating dynamics. While similar architectural designs have been recently explored in the machine learning community to enhance parallelization [14], we show that this structural choice provides an important system-theoretic benefit: it renders the DGN *unconditionally* δ ISS by design. Therefore, we establish a flexible neural network architecture for system identification that eliminates the need for constrained training procedures (e.g., stability penalties [4]) or complex post-training verification.

The remainder of this work is structured as follows. Section II introduces the notation and problem statement, and recalls the definitions of ISS and δ ISS. Section III details the dynamics of the CFN and introduces the proposed DGN architecture. Section IV provides the main theoretical contributions of the letter, deriving the stability proofs for both architectures. Section V presents numerical experiments validating the theoretical findings and comparing the modeling

Received 14 March 2026; revised 19 May 2026; accepted 8 June 2026. Date of publication 11 June 2026; date of current version 19 June 2026. This article has not been presented at a conference. Recommended by Senior Editor G. Giordano. (*Corresponding author: Stefano De Carli.*)

The authors are with the Department of Management, Information and Production Engineering, University of Bergamo, 24044 Dalmine, Bergamo, Italy (e-mail: stefano.decarli@unibg.it; davide.previtali@unibg.it; mirko.mazzoleni@unibg.it; previdi@unibg.it).

Digital Object Identifier 10.1109/LCSYS.2026.3702800




Fig. 1. Single-layer dynamics for the CFN and the DGN. The dashed paths represent the recurrent gate connections, which are absent in DGNs. Orange and purple represent the hidden state and input paths, respectively.

performance of the CFN and DGN against stable standard baselines. Finally, Section VI draws conclusions and outlines future research directions.

II. PRELIMINARIES AND PROBLEM STATEMENT

A. Notation and Preliminaries

Let \mathbb{R} , $\mathbb{R}_{\geq 0}$, $\mathbb{R}_{> 0}$, \mathbb{N} , \mathbb{Z}^+ denote the sets of real, non-negative real, positive real, natural numbers (including zero), and positive integers, respectively. Given $n, m \in \mathbb{Z}^+$, \mathbb{R}^n is the set of real column vectors of dimension n , while $\mathbb{R}^{n \times m}$ is the set of real matrices of dimension $n \times m$. Vectors are typeset in boldface lowercase, e.g., $\mathbf{v} = [v_1, \dots, v_n]^\top \in \mathbb{R}^n$. $\mathbf{0}$ and I denote the zero and identity matrices of appropriate dimensions. $|\mathcal{S}|$ is the cardinality of a set \mathcal{S} . The Hadamard (element-wise) product is indicated by \circ , and $\|\cdot\|_p$ represents the p -norm; unless otherwise specified, $\|\cdot\|$ denotes the infinity norm (i.e., $p = \infty$) for vectors and the induced infinity norm for matrices. We represent scalar discrete-time signals as s_k , where $k \in \mathbb{N}$ is the time step. Accordingly, an n -dimensional discrete-time signal is defined as $\mathbf{x}_k = [x_{1,k}, \dots, x_{n,k}]^\top \in \mathbb{R}^n$. The sequence of values attained by \mathbf{x}_k between $k = k_1$ and $k = k_2$ is denoted by $\{\mathbf{x}_k\}_{k=k_1}^{k_2}$. We use standard activation functions, i.e., the sigmoid $\sigma(x) = (1 + e^{-x})^{-1} \in (0, 1)$ and the hyperbolic tangent $\phi(x) = \tanh(x) \in (-1, 1)$ for $x \in \mathbb{R}$. For a scalar function $f: \mathbb{R} \rightarrow \mathbb{R}$, its bold counterpart $\mathbf{f}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ denotes the component-wise application. Finally, we employ standard definitions for comparison functions of classes \mathcal{K} , \mathcal{K}_∞ , and \mathcal{KL} as detailed in [15].

B. Problem Statement

We consider the problem of identifying a nonlinear discrete-time Multiple-Input Multiple-Output (MIMO) dynamical system from data. The objective is to learn the unknown map governing the evolution of the system output $\mathbf{y}_k \in \mathbb{R}^{n_y}$ based on the input $\mathbf{u}_k \in \mathbb{R}^{n_u}$, with $n_y, n_u \in \mathbb{Z}^+$. The system is characterized by an internal state $\mathbf{x}_k \in \mathbb{R}^{n_x}$, where $n_x \in \mathbb{Z}^+$. We assume that the input $\mathbf{u}_k \in \mathcal{U} \subset \mathbb{R}^{n_u}$ and the state $\mathbf{x}_k \in \mathcal{X} \subset \mathbb{R}^{n_x}$ belong to compact sets. The initial state $\mathbf{x}_0 \in \mathcal{X}$ is unknown and not subject to estimation. To identify the underlying dynamics, we use a dataset $\mathcal{D} = \{\mathcal{D}^{(v)}\}_{v=1}^V$ comprising $V \in \mathbb{Z}^+$ input-output sequences. Each sequence is defined as $\mathcal{D}^{(v)} := \{(\mathbf{u}_k^{(v)}, \mathbf{y}_k^{(v)})\}_{k=0}^{N^{(v)}-1}$, where $N^{(v)} \in \mathbb{Z}^+$ represents the length of the v -th sequence, with $v \in \{1, \dots, V\}$. Our goal is to identify the underlying dynamics using RNN architectures, focusing on models satisfying ISS and δ ISS. Stability is defined for a generic system $\mathbf{x}_{k+1} = \mathbf{w}(\mathbf{x}_k, \mathbf{u}_k)$, with $\mathbf{w}: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, assuming that \mathcal{X} is forward invariant with respect to \mathcal{U} , i.e., $\mathbf{w}(\mathbf{x}_k, \mathbf{u}_k) \in \mathcal{X}$, $\forall k \in \mathbb{N}$, for any $\mathbf{x}_k \in \mathcal{X}$ and $\mathbf{u}_k \in \mathcal{U}$.

Definition 1 (ISS [5], [16]): A system $\mathbf{x}_{k+1} = \mathbf{w}(\mathbf{x}_k, \mathbf{u}_k)$ is ISS in \mathcal{X} with respect to \mathcal{U} if there exist functions $\beta \in \mathcal{KL}$ and $\gamma_u, \gamma_b \in \mathcal{K}_\infty$ such that, for any $k \in \mathbb{N}$, any initial state $\mathbf{x}_0 \in \mathcal{X}$, any input sequence $\{\mathbf{u}_z \in \mathcal{U}\}_{z=0}^{k-1}$, and any bias vector $\mathbf{b} \in \mathbb{R}^{n_x}$, it holds that:

$$\|\mathbf{x}_k\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma_u \left(\max_{0 \leq z < k} \|\mathbf{u}_z\| \right) + \gamma_b (\|\mathbf{b}\|). \quad (1)$$

Definition 2 (δ ISS [5], [17]): A system $\mathbf{x}_{k+1} = \mathbf{w}(\mathbf{x}_k, \mathbf{u}_k)$ is δ ISS in \mathcal{X} with respect to \mathcal{U} if there exist functions $\beta_\delta \in \mathcal{KL}$ and $\gamma_{\delta u} \in \mathcal{K}_\infty$ such that, for any $k \in \mathbb{N}$, any pair of initial states $\mathbf{x}_0^a, \mathbf{x}_0^b \in \mathcal{X}$, and any pair of input sequences $\{\mathbf{u}_z^a \in \mathcal{U}\}_{z=0}^{k-1}$ and $\{\mathbf{u}_z^b \in \mathcal{U}\}_{z=0}^{k-1}$, the corresponding state trajectories \mathbf{x}_k^a and \mathbf{x}_k^b satisfy:

$$\|\mathbf{x}_k^a - \mathbf{x}_k^b\| \leq \beta_\delta (\|\mathbf{x}_0^a - \mathbf{x}_0^b\|, k) + \gamma_{\delta u} \left(\max_{0 \leq z < k} \|\mathbf{u}_z^a - \mathbf{u}_z^b\| \right). \quad (2)$$

We recall that δ ISS is a stronger property than ISS as the class of δ ISS systems constitutes a strict subset of the class of ISS systems [17, Proposition 1].

Remark 1: In this work, ISS and δ ISS are defined using the infinity norm for convenience, but the properties hold for any other vector norm [16].

III. CHAOS-FREE AND DECOUPLED-GATE NETWORKS

While standard LSTM and GRU networks are widely used and offer great modeling ability [1], [2, Ch. 15], they often exhibit chaotic behaviors [3] and lack theoretical stability guarantees by design. To address this, we analyze the chaos-free network architecture [3], which mitigates these issues, and propose a novel CFN variant, the decoupled-gate network, designed to guarantee unconditional δ ISS.

A. The Chaos-Free Network (CFN)

The CFN was introduced in [3] for general machine learning tasks. CFNs are designed to mitigate the chaotic behavior exhibited by standard gated RNNs (LSTM and GRU networks) while still managing the vanishing gradient problem. The network consists of $L \in \mathbb{Z}^+$ stacked layers. Each layer $l \in \mathcal{L} = \{1, \dots, L\}$ is a nonlinear discrete-time MIMO dynamical system in state-space form with hidden state $\mathbf{h}_k^{(l)} \in \mathbb{R}^{n_h^{(l)}}$. Here, (l) denotes the l -th layer, and $n_h^{(l)} \in \mathbb{Z}^+$ is its number of hidden units. The hidden state evolution of each l -th layer is governed by a forget gate $\mathbf{f}_k^{(l)} \in \mathbb{R}^{n_h^{(l)}}$, an input gate $\mathbf{i}_k^{(l)} \in \mathbb{R}^{n_h^{(l)}}$, and a candidate hidden state $\tilde{\mathbf{h}}_k^{(l)} \in \mathbb{R}^{n_h^{(l)}}$:

$$\mathbf{h}_{k+1}^{(l)} := \mathbf{f}_k^{(l)} \circ \phi(\mathbf{h}_k^{(l)}) + \mathbf{i}_k^{(l)} \circ \tilde{\mathbf{h}}_k^{(l)}, \quad (3a)$$

$$\mathbf{f}_k^{(l)} := \sigma(W_f^{(l)} \tilde{\mathbf{u}}_k^{(l)} + R_f^{(l)} \mathbf{h}_k^{(l)} + \mathbf{b}_f^{(l)}), \quad (3b)$$

$$\mathbf{i}_k^{(l)} := \sigma(W_i^{(l)} \tilde{\mathbf{u}}_k^{(l)} + R_i^{(l)} \mathbf{h}_k^{(l)} + \mathbf{b}_i^{(l)}), \quad (3c)$$

$$\tilde{\mathbf{h}}_k^{(l)} := \phi(W_h^{(l)} \tilde{\mathbf{u}}_k^{(l)} + \mathbf{b}_h^{(l)}), \quad (3d)$$

$$\tilde{\mathbf{u}}_k^{(l)} := \begin{cases} \mathbf{u}_k & \text{if } l = 1, \quad n_u^{(l)} = n_u, \\ \mathbf{h}_{k+1}^{(l-1)} & \text{if } l \in \mathcal{L} \setminus \{1\}, \quad n_u^{(l)} = n_h^{(l-1)}, \end{cases} \quad (3e)$$

where $\tilde{\mathbf{u}}_k^{(l)} \in \mathbb{R}^{n_u^{(l)}}$ is the layer input, $W_f^{(l)}, W_i^{(l)}, W_h^{(l)} \in \mathbb{R}^{n_h^{(l)} \times n_u^{(l)}}$ are the input weights, $R_f^{(l)}, R_i^{(l)} \in \mathbb{R}^{n_h^{(l)} \times n_h^{(l)}}$ are the recurrent weights, and $\mathbf{b}_f^{(l)}, \mathbf{b}_i^{(l)}, \mathbf{b}_h^{(l)} \in \mathbb{R}^{n_h^{(l)}}$ are the bias vectors. These

single-layer CFN dynamics are shown in Fig. 1. A key aspect of CFNs is that, due to the ranges of the activation functions $\sigma(\cdot) \in (0, 1)$ and $\phi(\cdot) \in (-1, 1)$, every layer admits a compact forward invariant set.

Proposition 1 (Forward invariant set of $\mathbf{h}_k^{(l)}$): For any $l \in \mathcal{L}$, the set $\mathcal{H}_{\text{inv}}^{(l)} := [-2, 2]^{n_h^{(l)}}$ is a forward invariant compact set for the layer dynamics in (3). Thus, for any initial condition $\mathbf{h}_0^{(l)} \in \mathcal{H}_{\text{inv}}^{(l)}$, the hidden state satisfies $\mathbf{h}_k^{(l)} \in \mathcal{H}_{\text{inv}}^{(l)}$ for all $k \in \mathbb{N}$ and any input $\tilde{\mathbf{u}}_k^{(l)} \in \mathbb{R}^{n_u^{(l)}}$.

The proof is reported in Appendix A. In practice, $\mathbf{h}_0^{(l)}$ is typically initialized as the zero vector [18, Ch. 10], i.e., $\mathbf{h}_0^{(l)} = \mathbf{0} \in \mathcal{H}_{\text{inv}}^{(l)}, \forall l \in \mathcal{L}$. By Proposition 1, the complete network state $\mathbf{h}_k = [\mathbf{h}_k^{(1)\top}, \dots, \mathbf{h}_k^{(L)\top}]^\top$ evolves within $\mathcal{H}_{\text{inv}} = \times_{l=1}^L \mathcal{H}_{\text{inv}}^{(l)} \subset \mathbb{R}^{n_h}$, where $n_h = \sum_{l=1}^L n_h^{(l)}$ is the total number of hidden units. The network output \mathbf{y}_k follows from a fully connected mapping of the final layer hidden state:

$$\begin{cases} \mathbf{h}_{k+1}^{(l)} := \mathbf{f}_k^{(l)} \circ \phi(\mathbf{h}_k^{(l)}) + \mathbf{i}_k^{(l)} \circ \tilde{\mathbf{h}}_k^{(l)}, \forall l \in \mathcal{L}, \\ \mathbf{y}_k := W_y \mathbf{h}_{k+1}^{(L)} + \mathbf{b}_y, \end{cases} \quad (4)$$

where $W_y \in \mathbb{R}^{n_y \times n_h^{(L)}}$ and $\mathbf{b}_y \in \mathbb{R}^{n_y}$. Finally, the CFN relies on the following set of tunable parameters θ :

$$\begin{aligned} \theta &:= \{W_y, \mathbf{b}_y\} \cup \{\theta^{(l)} : l \in \mathcal{L}\}, \\ \theta^{(l)} &:= \{W_f^{(l)}, W_i^{(l)}, W_h^{(l)}, R_f^{(l)}, R_i^{(l)}, \mathbf{b}_f^{(l)}, \mathbf{b}_i^{(l)}, \mathbf{b}_h^{(l)}\}. \end{aligned} \quad (5)$$

A distinctive feature of CFNs, compared to common gated RNNs, is that the candidate state $\tilde{\mathbf{h}}_k^{(l)}, l \in \mathcal{L}$, depends *solely* on the current layer input $\tilde{\mathbf{u}}_k^{(l)}$ and not on the hidden state at the previous time step. However, the CFN retains full recurrent connectivity due to the presence of the weight matrices $R_f^{(l)}, R_i^{(l)}$. This coupling between the hidden state $\mathbf{h}_k^{(l)}$ and the gates introduces a state-dependent modulation that, while enabling rich dynamics, makes it challenging to provide a priori strong stability properties (such as δ ISS) independently of the specific parameter values, see Section IV.

B. The Proposed Decoupled-Gate Network (DGN)

To guarantee unconditional δ ISS without parametric constraints, we introduce the decoupled-gate network. Inspired by parallelizable RNN architectures [14], from which it inherits training efficiency benefits, the DGN structurally simplifies the CFN by eliminating the recurrent dependencies of the gates on the hidden state, setting $R_f^{(l)} = R_i^{(l)} = \mathbf{0}$ for all $l \in \mathcal{L}$. Consequently, the gates in (3b) and (3c) of a DGN layer depend exclusively on the layer input $\tilde{\mathbf{u}}_k^{(l)}$, removing state-dependent modulations that complicate stability analysis. Although this restricts single-layer dynamics, complex temporal processing is recovered through network depth ($L > 1$) [14]. Since deeper layers are driven by previous states ($\tilde{\mathbf{u}}_k^{(l)} = \mathbf{h}_{k+1}^{(l-1)}$), their gating mechanisms remain effectively state-dependent. This structural choice preserves model expressivity while rendering the DGN δ ISS *by design*, as proven in Section IV.

IV. STABILITY PROPERTIES

In this section, we analyze the stability of the CFN and DGN. All proofs are deferred to Appendix A. We adopt the following assumption on input boundedness.

Assumption 1 (Input boundedness): The inputs \mathbf{u}_k are bounded such that $\mathbf{u}_k \in \mathcal{U} := [-1, 1]^{n_u}$ for all $k \in \mathbb{N}$.

This is a customary assumption for neural networks, satisfied by input data normalization. Furthermore, considering the layer input $\tilde{\mathbf{u}}_k^{(l)}$ in (3e), it follows from Assumption 1 and Proposition 1 that $\tilde{\mathbf{u}}_k^{(l)} \in \tilde{\mathcal{U}}^{(l)} := [-2, 2]^{n_u^{(l)}}$ for all $l \in \mathcal{L}$. Indeed, $\tilde{\mathbf{u}}_k^{(1)} = \mathbf{u}_k \in \mathcal{U} \subset \tilde{\mathcal{U}}^{(1)}$, while for $l \in \mathcal{L} \setminus \{1\}$, $\tilde{\mathbf{u}}_k^{(l)} = \mathbf{h}_{k+1}^{(l-1)} \in \mathcal{H}_{\text{inv}}^{(l-1)} = \tilde{\mathcal{U}}^{(l)}$.

A. Stability of CFN Architectures

Based on the above assumptions, we derive sufficient parametric conditions for ISS and δ ISS of CFN layers.

Theorem 1 (ISS of a CFN layer): Under Assumption 1, the l -th layer, $l \in \mathcal{L}$, of the CFN in (3) is ISS in $\mathcal{H}_{\text{inv}}^{(l)}$ with respect to $\tilde{\mathcal{U}}^{(l)}$.

Theorem 2 (δ ISS of a CFN layer): Under Assumption 1, the l -th layer, $l \in \mathcal{L}$, of the CFN in (3) is δ ISS in $\mathcal{H}_{\text{inv}}^{(l)}$ with respect to $\tilde{\mathcal{U}}^{(l)}$ if:

$$\rho^{(l)} := \bar{\sigma}_f^{(l)} + \frac{1}{4} \|\mathbf{R}_f^{(l)}\| + \frac{1}{4} \|\mathbf{R}_i^{(l)}\| \bar{\phi}_h^{(l)} < 1, \quad (6)$$

where:

$$\bar{\sigma}_f^{(l)} := \sigma \left(\left\| \begin{bmatrix} 2W_f^{(l)} & 2R_f^{(l)} & \mathbf{b}_f^{(l)} \end{bmatrix} \right\| \right), \quad (7a)$$

$$\bar{\phi}_h^{(l)} := \phi \left(\left\| \begin{bmatrix} 2W_h^{(l)} & \mathbf{b}_h^{(l)} \end{bmatrix} \right\| \right). \quad (7b)$$

We now generalize the results to deep networks ($L > 1$).

Theorem 3 (ISS of a CFN): Under Assumption 1, the L -layer CFN in (4) is ISS in \mathcal{H}_{inv} with respect to \mathcal{U} .

Theorem 4 (δ ISS of a CFN): Under Assumption 1, the L -layer CFN in (4) is δ ISS in \mathcal{H}_{inv} with respect to \mathcal{U} if Condition (6) holds for every layer $l \in \mathcal{L}$.

B. Extension to DGN Architectures

The DGN architecture, being a structural simplification of the CFN, inherits similar but stronger stability properties.

Theorem 5 (δ ISS of a DGN): Under Assumption 1, the L -layer DGN, defined by (4) with $R_f^{(l)} = R_i^{(l)} = \mathbf{0}$ for all $l \in \mathcal{L}$, is δ ISS in \mathcal{H}_{inv} with respect to \mathcal{U} .

V. NUMERICAL RESULTS

In this section, we evaluate the modeling capabilities of the CFN and the proposed DGN architectures for system identification tasks, a domain *unexplored* in prior works, since CFNs were validated only for word-level language modeling [3]. We assess their performance on two standard benchmarks for RNN stability: the pH Neutralization Reactor [5] and the Four-Tank system [6].

A. Training Configuration

Recall the dataset \mathcal{D} introduced in Section II-B. For each benchmark, we partition the set of indices $\{1, \dots, V\}$ into three disjoint subsets \mathcal{I}_{tr} (training), \mathcal{I}_{val} (validation), and \mathcal{I}_{tst} (test), defining the corresponding datasets \mathcal{D}_{tr} , \mathcal{D}_{val} , and \mathcal{D}_{tst} . To satisfy Assumption 1, inputs and outputs are linearly rescaled within the range $[-1, 1]$ using parameters computed exclusively from \mathcal{D}_{tr} . We follow the standard training procedure for RNNs [18, Ch. 10], also shared with baseline comparisons

[5], [6]. The parameters θ in (5) are estimated by minimizing the Mean Squared Error (MSE) on \mathcal{D}_{tr} :

$$J(\theta) := \frac{1}{|\mathcal{I}_{\text{tr}}|} \sum_{v \in \mathcal{I}_{\text{tr}}} \frac{1}{N^{(v)} - N_w} \sum_{k=N_w}^{N^{(v)}-1} \left\| y_k^{(v)} - \hat{y}_k^{(v)}(\theta) \right\|_2^2, \quad (8)$$

where $\hat{y}_k(\theta)$ is the output predicted by the network with parameters θ for the sequence $v \in \mathcal{I}_{\text{tr}}$, given the input $\mathbf{u}_k^{(v)}$ at time k . We apply a washout period of $N_w = 25$ steps to the training and validation sequences, ensuring that the prediction error during the initial transient is excluded from the loss calculation [18, Ch. 10]. The models are trained for 2000 epochs using the Adam optimizer with a base learning rate of 0.001, decaying by 0.9 every 200 epochs to facilitate convergence. The training data is processed in mini-batches of size $b \in \mathbb{Z}^+$, which is adjusted according to the specific dataset dimensions. To prevent overfitting, we apply a dropout rate of 0.05. We adopt an early stopping strategy, selecting the parameters θ^* that yield the minimal MSE on \mathcal{D}_{val} , which is evaluated every epoch. Crucially, for the DGN architecture, *no additional constraints are imposed during training*, as the model is δ ISS by design (Theorem 5). In contrast, the stable baselines (LSTM and GRU) used for comparison require minimizing a loss function augmented with a stability penalty term to promote δ ISS compliance [5], [6]. Lastly, we train the CFN *without* auxiliary stability penalties. Notably, the resulting CFN models consistently violated the sufficient δ ISS condition in (6), highlighting that δ ISS compliance for this architecture remains parameter-dependent rather than structurally guaranteed a priori.

B. PH Reactor

We first utilize the pH neutralization benchmark [5] to compare the CFN and the DGN against the δ ISS-LSTM network baseline established in the literature [5], which was trained using a stability-promoting augmented loss. The system describes the nonlinear control of the pH level (y) via an alkaline flow (u) under acid disturbances, representing a single-input single-output system. \mathcal{D} contains $V = 20$ sequences generated by a Multilevel Pseudo-Random Signal (MPRS) with a sampling time of 10 s. \mathcal{D} is partitioned into $|\mathcal{D}_{\text{tr}}| = 15$, $|\mathcal{D}_{\text{val}}| = 4$ and $|\mathcal{D}_{\text{tst}}| = 1$ sequences, following the original experimental setup in [5].

We train single-layer CFN and DGN models with $n_h^{(1)} = 7$ hidden units, matching the δ ISS-LSTM model in [5], and use a mini-batch size of $b = 5$. Performance is evaluated for a generic sequence $v \in \mathcal{I}_{\text{tst}}$ using the $\text{Fit}^{(v)}$ metric:

$$\text{Fit}^{(v)}(\theta^*) := 100 \left(1 - \frac{\sqrt{\sum_{k=0}^{N^{(v)}-1} \left(y_k^{(v)} - \hat{y}_k^{(v)}(\theta^*) \right)^2}}{\sqrt{\sum_{k=0}^{N^{(v)}-1} \left(y_k^{(v)} - \bar{y}^{(v)} \right)^2}} \right)^2,$$

where $\bar{y}^{(v)}$ is the sample mean of y along sequence v . The results, summarized in terms of Fit on \mathcal{D}_{tst} , demonstrate the competitive performance of the proposed architectures, as visually confirmed by the time-domain comparison in Fig. 2. The CFN achieves a Fit of 96.1%, closely matching the 96.5% of the δ ISS-LSTM benchmark in [5]. Notably, the DGN reaches 95.6%, practically matching the CFN

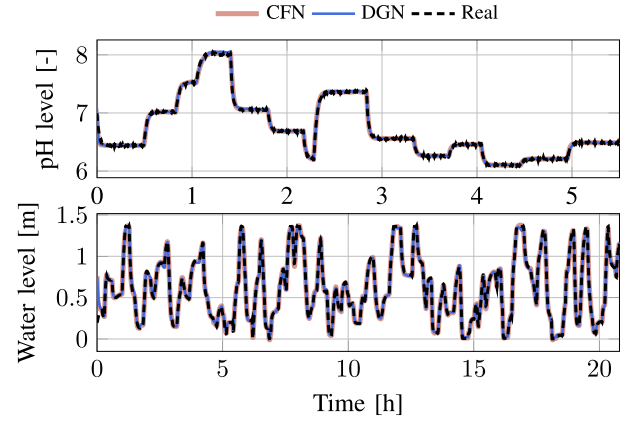


Fig. 2. Test predictions for pH Reactor (top) and Four-Tank (bottom). The plots compare the ground truth (dashed black line) against the CFN (thick orange line) and DGN (thin blue line) predictions. For the Four-Tank, median performance realizations are shown.

performance. Therefore, the DGN maintains high accuracy while providing rigorous unconditional stability, bypassing the complex stability-promoting augmented loss functions required by standard gated RNNs [5], [6].

C. Four-Tank System

We also evaluate the architectures on the Four-Tank benchmark, already used for δ ISS-GRU network validation in [6]. The task is to model the water level of the lower tank (y) given two input pump flow rates (\mathbf{u}), constituting a multiple-input single-output problem. A 25,000-step MPRS trajectory with a sampling time of 15 s is partitioned into 10 blocks. To assess the robustness of the results, we repeat the training and validation phases for all 45 possible combinations of these blocks, selecting 8 blocks for training and 2 for validation in each run. Windowing these blocks into subsequences of length 500 (with 50% overlap) practically yields $|\mathcal{D}_{\text{tr}}| = 77$ and $|\mathcal{D}_{\text{val}}| = 18$ subsequences. Then, an independent test trajectory ($|\mathcal{D}_{\text{tst}}| = 1$) of 5000 steps is generated. To assess performance under realistic conditions, we add white Gaussian noise to the output measurements with a signal-to-noise ratio of 100 (i.e., 20 dB). We employ a 3-layer architecture with $n_h^{(l)} = 7$ units per layer for both the CFN and DGN to match the δ ISS-GRU model in [6], and use a mini-batch size of $b = 25$. The models demonstrate remarkable performance on \mathcal{D}_{tst} for this task as well (see Fig. 2, where the realization corresponding to the median Fit performance is shown). The aggregate results, reported as mean \pm standard deviation, confirm the high robustness of the proposed models. Specifically, the CFN attains an average Fit of $94.03\% \pm 0.21\%$, while the DGN follows with $92.55\% \pm 0.19\%$; notably, if the first N_w transient steps are excluded from the Fit calculation, we obtain $96.22\% \pm 0.19\%$ for the CFN and $95.50\% \pm 0.22\%$ for the DGN. We compare these against the δ ISS-GRU baseline reported in [6], obtained with a stability-promoting augmented loss strategy. Although this baseline achieved a Fit of 97.0%, it is crucial to note that it was trained and tested with different data from the same system.

In conclusion, while structural constraints in DGNs slightly reduce expressivity, this is a favorable trade-off for obtaining unconditional δ ISS *by design*, eliminating the need for complex augmented losses during training [5], [6], [8].

VI. CONCLUSION

In this work, we investigated the stability properties of the CFN and the proposed DGN, a novel architecture designed to ensure unconditional δ ISS in nonlinear system identification. We proved that the CFN inherently satisfies ISS by design, while the DGN unconditionally satisfies δ ISS without the restrictive training constraints required by standard gated RNNs used in previous works [5], [6], [8]. This contrasts with standard gated RNNs (e.g., LSTM and GRU networks), where δ ISS must be enforced by augmenting the loss function with penalty terms during the network training process [5], [6], [8]. Numerical results confirm that the DGN maintains predictive accuracies comparable to the CFN and stable gated RNN baselines while offering intrinsically strong structural stability guarantees. Future research will focus on exploiting the intrinsic δ ISS property of the DGN to design robust state observers and integrating these models into safety-critical control frameworks, such as model predictive control [9], [10].

APPENDIX A PROOFS

PROOF OF PROPOSITION 1. Consider the hidden state $\mathbf{h}_k^{(l)} = [h_{1,k}^{(l)}, \dots, h_{n_h^{(l)},k}^{(l)}]^\top$ for $l \in \mathcal{L}$. For any component $j \in \mathcal{J}^{(l)} := \{1, \dots, n_h^{(l)}\}$, the update (3a) satisfies $|h_{j,k+1}^{(l)}| \leq |f_{j,k}^{(l)}| |\phi(h_{j,k}^{(l)})| + |i_{j,k}^{(l)}| |\tilde{h}_{j,k}^{(l)}|$ by the triangle inequality. Since $\sigma(\cdot) \in (0, 1)$ and $\phi(\cdot) \in (-1, 1)$, it follows that $|h_{j,k+1}^{(l)}| \leq 2$. Thus, $\forall k \in \mathbb{N}$, if $\mathbf{h}_k^{(l)} \in \mathcal{H}_{\text{inv}}^{(l)} := [-2, 2]^{n_h^{(l)}}$, then $\mathbf{h}_{k+1}^{(l)} \in \mathcal{H}_{\text{inv}}^{(l)}$ for any $\tilde{\mathbf{u}}_k^{(l)} \in \mathbb{R}^{n_u^{(l)}}$, making $\mathcal{H}_{\text{inv}}^{(l)}$ a forward invariant set. ■

PROOF OF THEOREM 1. Consider a single layer $l \in \mathcal{L}$ and omit the superscript (l) for ease of discussion. For any $j \in \mathcal{J}$, we can upper bound the forget gate component $f_{j,k}$ from (3b) by exploiting the monotonic properties of the sigmoid function:

$$\begin{aligned} |f_{j,k}| &\leq \max_{\tilde{\mathbf{u}}_k \in \tilde{\mathcal{U}}, \mathbf{h}_k \in \mathcal{H}_{\text{inv}}} \|\sigma(W_f \tilde{\mathbf{u}}_k + R_f \mathbf{h}_k + \mathbf{b}_f)\| \\ &\leq \sigma(\| [2W_f \ 2R_f \ \mathbf{b}_f] \|) := \bar{\sigma}_f \in (0, 1). \end{aligned} \quad (9)$$

Thus, $\|\mathbf{f}_k\| \leq \bar{\sigma}_f$ for all $k \in \mathbb{N}$. Taking the norm of the state update in (3a) and applying the triangle inequality yields:

$$\|\mathbf{h}_{k+1}\| \leq \|\mathbf{f}_k\| \|\phi(\mathbf{h}_k)\| + \|\mathbf{i}_k\| \|\tilde{\mathbf{h}}_k\|.$$

Using the 1-Lipschitz continuity of the hyperbolic tangent function and the bounds $\|\mathbf{f}_k\| \leq \bar{\sigma}_f$, $\|\mathbf{i}_k\| \leq 1$, we obtain:

$$\|\mathbf{h}_{k+1}\| \leq \bar{\sigma}_f \|\mathbf{h}_k\| + \|\tilde{\mathbf{h}}_k\|. \quad (10)$$

From (3d), the 1-Lipschitz continuity of the hyperbolic tangent function implies that the candidate hidden state is bounded as $\|\tilde{\mathbf{h}}_k\| \leq \|W_{\tilde{h}}\| \|\tilde{\mathbf{u}}_k\| + \|\mathbf{b}_{\tilde{h}}\|$. Substituting this back into (10), we get the recursive inequality:

$$\|\mathbf{h}_{k+1}\| \leq \bar{\sigma}_f \|\mathbf{h}_k\| + \|W_{\tilde{h}}\| \|\tilde{\mathbf{u}}_k\| + \|\mathbf{b}_{\tilde{h}}\|.$$

We propagate this inequality back to the initial state:

$$\begin{aligned} \|\mathbf{h}_k\| &\leq \bar{\sigma}_f^k \|\mathbf{h}_0\| + \sum_{z=0}^{k-1} \bar{\sigma}_f^{k-1-z} (\|W_{\tilde{h}}\| \|\tilde{\mathbf{u}}_z\| + \|\mathbf{b}_{\tilde{h}}\|) \\ &\leq \bar{\sigma}_f^k \|\mathbf{h}_0\| + \left(\|W_{\tilde{h}}\| \max_{0 \leq z < k} \|\tilde{\mathbf{u}}_z\| + \|\mathbf{b}_{\tilde{h}}\| \right) \sum_{z=0}^{k-1} \bar{\sigma}_f^{k-1-z}. \end{aligned}$$

Since $\bar{\sigma}_f \in (0, 1)$, the geometric series can be bounded as $\sum_{z=0}^{k-1} \bar{\sigma}_f^{k-1-z} \leq \frac{1}{1-\bar{\sigma}_f}$. We can now deduce the \mathcal{KL} and the two \mathcal{K}_∞ functions for the ISS condition in (1):

$$\begin{aligned} \beta(\|\mathbf{h}_0\|, k) &:= \bar{\sigma}_f^k \|\mathbf{h}_0\|, \\ \gamma_u \left(\max_{0 \leq z < k} \|\tilde{\mathbf{u}}_z\| \right) &:= \frac{\|W_{\tilde{h}}\|}{1-\bar{\sigma}_f} \max_{0 \leq z < k} \|\tilde{\mathbf{u}}_z\|, \\ \gamma_b(\|\mathbf{b}_{\tilde{h}}\|) &:= \frac{1}{1-\bar{\sigma}_f} \|\mathbf{b}_{\tilde{h}}\|. \end{aligned}$$

Thus, the l -th CFN layer, $l \in \mathcal{L}$, is ISS in $\mathcal{H}_{\text{inv}}^{(l)}$ with respect to $\tilde{\mathcal{U}}^{(l)}$ due to the existence of these functions. ■

PROOF OF THEOREM 2. Consider a single layer $l \in \mathcal{L}$ and omit the superscript (l) for ease of discussion. For any $j \in \mathcal{J}$, we first derive the bound of the candidate hidden state component $\tilde{h}_{j,k}$ from (3d). As with $\bar{\sigma}_f$ in (9), since the hyperbolic tangent function is monotonically increasing:

$$\begin{aligned} |\tilde{h}_{j,k}| &\leq \max_{\tilde{\mathbf{u}}_k \in \tilde{\mathcal{U}}} \|\phi(W_{\tilde{h}} \tilde{\mathbf{u}}_k + \mathbf{b}_{\tilde{h}})\| \\ &\leq \phi(\|[2W_{\tilde{h}} \ \mathbf{b}_{\tilde{h}}]\|) := \bar{\phi}_{\tilde{h}} \in (0, 1). \end{aligned} \quad (11)$$

Now, consider two state trajectories $\mathbf{h}_k^a, \mathbf{h}_k^b$ starting from $\mathbf{h}_0^a, \mathbf{h}_0^b \in \mathcal{H}_{\text{inv}}$ and driven by the input sequences $\{\tilde{\mathbf{u}}_z^a \in \tilde{\mathcal{U}}\}_{z=0}^{k-1}$ and $\{\tilde{\mathbf{u}}_z^b \in \tilde{\mathcal{U}}\}_{z=0}^{k-1}$. Let $\xi_k^\Delta := \xi_k^a - \xi_k^b$ denote the incremental dynamics for any variable ξ_k . We analyze the component-wise evolution of the state difference for any $j \in \mathcal{J}$, based on the state update equation (3a):

$$h_{j,k+1}^\Delta = (f_{j,k}^a \phi(h_{j,k}^a) + i_{j,k}^a \tilde{h}_{j,k}^a) - (f_{j,k}^b \phi(h_{j,k}^b) + i_{j,k}^b \tilde{h}_{j,k}^b).$$

Adding and subtracting the cross-terms $f_{j,k}^a \phi(h_{j,k}^b)$ and $i_{j,k}^a \tilde{h}_{j,k}^b$, we can directly group the terms as follows:

$$\begin{aligned} h_{j,k+1}^\Delta &= \underbrace{f_{j,k}^a (\phi(h_{j,k}^a) - \phi(h_{j,k}^b))}_{\tau_1} + \underbrace{(f_{j,k}^a - f_{j,k}^b) \phi(h_{j,k}^b)}_{\tau_2} \\ &\quad + \underbrace{i_{j,k}^a (\tilde{h}_{j,k}^a - \tilde{h}_{j,k}^b)}_{\tau_3} + \underbrace{(i_{j,k}^a - i_{j,k}^b) \tilde{h}_{j,k}^b}_{\tau_4}. \end{aligned} \quad (12)$$

We bound each term individually leveraging $|i_{j,k}^a| \leq 1$, $|\phi(h_{j,k}^b)| \leq 1$, $|f_{j,k}^a| \leq \bar{\sigma}_f$ from (9), and $|\tilde{h}_{j,k}^b| \leq \bar{\phi}_{\tilde{h}}$ from (11). Applying the 1-Lipschitz continuity property of the hyperbolic tangent function and the $\frac{1}{4}$ -Lipschitz continuity property of the sigmoid function to the gates, we bound the absolute values of the incremental components:

$$\begin{aligned} |\tau_1| &\leq \bar{\sigma}_f \|\mathbf{h}_k^\Delta\|, \quad |\tau_3| \leq \|W_{\tilde{h}}\| \|\tilde{\mathbf{u}}_k^\Delta\|, \\ |\tau_2| &\leq \frac{1}{4} (\|W_f\| \|\tilde{\mathbf{u}}_k^\Delta\| + \|R_f\| \|\mathbf{h}_k^\Delta\|), \\ |\tau_4| &\leq \frac{1}{4} (\|W_{i\tilde{h}}\| \|\tilde{\mathbf{u}}_k^\Delta\| + \|R_{i\tilde{h}}\| \|\mathbf{h}_k^\Delta\|) \bar{\phi}_{\tilde{h}}. \end{aligned}$$

Substituting these bounds into (12) and taking the maximum with respect to $j \in \mathcal{J}$, we obtain:

$$\|\mathbf{h}_{k+1}^\Delta\| \leq \rho \|\mathbf{h}_k^\Delta\| + \Gamma_{\tilde{u}} \|\tilde{\mathbf{u}}_k^\Delta\|, \quad (13)$$

where $\rho := \bar{\sigma}_f + \frac{1}{4} \|R_f\| + \frac{1}{4} \|R_{i\tilde{h}}\| \bar{\phi}_{\tilde{h}}$ matches the left-hand side of (6), and $\Gamma_{\tilde{u}} := \|W_{\tilde{h}}\| + \frac{1}{4} \|W_f\| + \frac{1}{4} \|W_{i\tilde{h}}\| \bar{\phi}_{\tilde{h}}$. Assuming

Condition (6) holds ($\rho < 1$) and unwinding the recursion back to the initial state difference, we obtain:

$$\begin{aligned} \|\mathbf{h}_k^\Delta\| &\leq \rho^k \|\mathbf{h}_0^\Delta\| + \Gamma_{\tilde{u}} \sum_{z=0}^{k-1} \rho^{k-1-z} \|\tilde{\mathbf{u}}_z^\Delta\| \\ &\leq \rho^k \|\mathbf{h}_0^\Delta\| + \frac{\Gamma_{\tilde{u}}}{1-\rho} \max_{0 \leq z < k} \|\tilde{\mathbf{u}}_z^\Delta\|, \end{aligned}$$

where the input sequence is bounded by its maximum and the geometric series is bounded as in the proof of Theorem 1. We can now deduce the \mathcal{KL} and the \mathcal{K}_∞ functions for the δ ISS condition in (2), i.e.,

$$\begin{aligned} \beta_\delta(\|\mathbf{h}_0^a - \mathbf{h}_0^b\|, k) &:= \rho^k \|\mathbf{h}_0^a - \mathbf{h}_0^b\|, \\ \gamma_{\delta u} \left(\max_{0 \leq z < k} \|\tilde{\mathbf{u}}_z^a - \tilde{\mathbf{u}}_z^b\| \right) &:= \frac{\Gamma_{\tilde{u}}}{1-\rho} \max_{0 \leq z < k} \|\tilde{\mathbf{u}}_z^a - \tilde{\mathbf{u}}_z^b\|. \end{aligned}$$

Thus, the l -th CFN layer, $l \in \mathcal{L}$, is δ ISS in $\mathcal{H}_{\text{inv}}^{(l)}$ with respect to $\tilde{\mathcal{U}}^{(l)}$ if Condition (6) holds. ■

PROOF OF THEOREM 3. Every CFN layer is unconditionally ISS (Theorem 1). The entire network is a cascade of ISS subsystems. As proven in [16], the cascade of ISS systems is ISS. Thus, the network is ISS. ■

PROOF OF THEOREM 4. Consider two state trajectories $\mathbf{h}_k^a, \mathbf{h}_k^b$ for the L -layer CFN driven by the input sequences $\{\mathbf{u}_z^a \in \mathcal{U}\}_{z=0}^{k-1}$ and $\{\mathbf{u}_z^b \in \mathcal{U}\}_{z=0}^{k-1}$, starting from initial states $\mathbf{h}_0^a, \mathbf{h}_0^b \in \mathcal{H}_{\text{inv}}$. As in the proof of Theorem 2, the superscript Δ denotes the incremental dynamics. Let $\boldsymbol{\eta}_k^\Delta \in \mathbb{R}_{\geq 0}^L$ be the vector collecting the state difference norms, i.e., $\boldsymbol{\eta}_k^\Delta := \left[\|\mathbf{h}_k^{(1),\Delta}\|, \dots, \|\mathbf{h}_k^{(L),\Delta}\| \right]^\top$. Recall the inter-layer connections $\|\tilde{\mathbf{u}}_k^{(l),\Delta}\| = \|\mathbf{h}_{k+1}^{(l-1),\Delta}\|$ for $l \in \mathcal{L} \setminus \{1\}$ in (3e), while $\|\tilde{\mathbf{u}}_k^{(1),\Delta}\| = \|\mathbf{u}_k^\Delta\|$. Recursively unwrapping (13) through the layers [6] yields the overall vector dynamics (\leq is applied element-wise):

$$\boldsymbol{\eta}_k^\Delta \leq A_\delta \boldsymbol{\eta}_{k-1}^\Delta + B_{\delta u} \|\mathbf{u}_{k-1}^\Delta\|, \quad (14)$$

where $A_\delta \in \mathbb{R}_{\geq 0}^{L \times L}$ is lower triangular with diagonal elements $\rho^{(l)}$, defined in (6), and $B_{\delta u} \in \mathbb{R}_{\geq 0}^L$. If Condition (6) holds $\forall l \in \mathcal{L}$, A_δ is Schur stable as its eigenvalues (its diagonal elements) satisfy $|\rho^{(l)}| < 1$. Propagating (14) backwards to $k = 0$ yields:

$$\boldsymbol{\eta}_k^\Delta \leq A_\delta^k \boldsymbol{\eta}_0^\Delta + \sum_{z=0}^{k-1} A_\delta^{k-1-z} B_{\delta u} \max_{0 \leq z < k} \|\mathbf{u}_z^\Delta\|.$$

Since A_δ and $B_{\delta u}$ are non-negative, the finite sum is upper-bounded by the infinite series. Further, due to Schur stability, we have $\sum_{z=0}^{\infty} A_\delta^z = (I - A_\delta)^{-1}$ [19, Ch. 7]. Then, $\sum_{z=0}^{k-1} A_\delta^{k-1-z} B_{\delta u} \max_{0 \leq z < k} \|\mathbf{u}_z^\Delta\| \leq (I - A_\delta)^{-1} B_{\delta u} \max_{0 \leq z < k} \|\mathbf{u}_z^\Delta\|$. Applying the infinity norm to both sides and noting $\|\boldsymbol{\eta}_k^\Delta\| = \|\mathbf{h}_k^\Delta\|$ results in:

$$\|\mathbf{h}_k^\Delta\| \leq \|A_\delta^k\| \|\mathbf{h}_0^\Delta\| + \|(I - A_\delta)^{-1} B_{\delta u}\| \max_{0 \leq z < k} \|\mathbf{u}_z^\Delta\|.$$

Next, we leverage the Jordan form of A_δ to derive an upper bound for $\|A_\delta^k\|$ [19, Ch. 7]:

$$\|A_\delta^k\| \leq \omega \sum_{z=0}^{m-1} \binom{k}{z} \left(\max_{l \in \mathcal{L}} \rho^{(l)} \right)^{k-z}, \quad (15)$$

for some $\omega \geq 1$, where $m \in \mathbb{Z}^+$, $m \leq L$, is the size of the largest Jordan block of A_δ . Conventionally, we set $\binom{k}{z} = 0$ for

any $z > k$. Since $\rho^{(l)} \in (0, 1)$, the bound in (15) tends to 0 as $k \rightarrow \infty$. Thus, the δ ISS condition in (2) holds with:

$$\beta_\delta(\|\mathbf{h}_0^a - \mathbf{h}_0^b\|, k) := \omega \sum_{z=0}^{m-1} \binom{k}{z} \left(\max_{l \in \mathcal{L}} \rho^{(l)} \right)^{k-z} \|\mathbf{h}_0^a - \mathbf{h}_0^b\|,$$

$$\gamma_{\delta u} \left(\max_{0 \leq z < k} \|\mathbf{u}_z^a - \mathbf{u}_z^b\| \right) := \|(I - A_\delta)^{-1} B_{\delta u}\| \max_{0 \leq z < k} \|\mathbf{u}_z^a - \mathbf{u}_z^b\|.$$

Hence, the L -layer CFN is δ ISS in \mathcal{H}_{inv} with respect to \mathcal{U} if Condition (6) holds for every layer $l \in \mathcal{L}$. ■

PROOF OF THEOREM 5. The DGN is a special case of the CFN where $R_f^{(l)} = R_i^{(l)} = \mathbf{0}, \forall l \in \mathcal{L}$. Substituting these zero matrices into the δ ISS sufficient condition in (6), the left-hand side simplifies to $\rho^{(l)} = \bar{\sigma}_f^{(l)}$. Since $\bar{\sigma}_f^{(l)} \in (0, 1)$ as in (9), the condition $\rho^{(l)} < 1$ is satisfied. Since this holds $\forall l \in \mathcal{L}$, the L -layer DGN is δ ISS by Theorem 4. ■

REFERENCES

- [1] 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, Jan. 2025, Art. no. 111907.
- [2] K. P. Murphy, *Probabilistic Machine Learning*. Cambridge, MA, USA: MIT Press, 2022.
- [3] T. Laurent and J. von Brecht, "A recurrent neural network without chaos," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2017. [Online]. Available: <https://openreview.net/forum?id=S1dlzvc1g>
- [4] F. Bonassi, M. Farina, J. Xie, and R. Scattolini, "On recurrent neural networks for learning-based control: Recent results and ideas for future developments," *J. Process Control*, vol. 114, pp. 92–104, Jun. 2022.
- [5] E. Terzi, F. Bonassi, M. Farina, and R. Scattolini, "Learning model predictive control with long short-term memory networks," *Int. J. Robust Nonlinear Control*, vol. 31, no. 18, pp. 8877–8896, 2021.
- [6] F. Bonassi, M. Farina, and R. Scattolini, "On the stability properties of gated recurrent units neural networks," *Syst. Control Lett.*, vol. 157, Nov. 2021, Art. no. 105049.
- [7] W. D'Amico, A. La Bella, and M. Farina, "An incremental input-to-state stability condition for a class of recurrent neural networks," *IEEE Trans. Autom. Control*, vol. 69, no. 4, pp. 2221–2236, Apr. 2024.
- [8] S. De Carli, D. Previtali, L. Pitturelli, M. Mazzoleni, A. Ferramosca, and F. Previdi, "Infinity-norm-based input-to-state-stable long short-term memory networks: A thermal systems perspective," in *Proc. ECC*, 2025, pp. 911–916.
- [9] I. Schimperna and L. Magni, "Robust offset-free constrained model predictive control with long short-term memory networks," *IEEE Trans. Autom. Control*, vol. 69, no. 12, pp. 8172–8187, Dec. 2024.
- [10] I. Schimperna and L. Magni, "Robust constrained nonlinear model predictive control with gated recurrent unit model," *Automatica*, vol. 161, Mar. 2024, Art. no. 111472.
- [11] J. Miller and M. Hardt, "Stable recurrent models," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2019. [Online]. Available: <https://openreview.net/forum?id=Hygxb2CqKm>
- [12] S. Bai, J. Z. Kolter, and V. Koltun, "Deep equilibrium models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., Curran Associates, Inc., 2019. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2019/file/01386bd6d8e091c2ab4c7c7de644d37b-Paper.pdf
- [13] S. Strogatz, *Nonlinear Dynamics and Chaos*. Boca Raton, FL, USA: CRC Press, 2019.
- [14] L. Feng, F. Tung, M. O. Ahmed, Y. Bengio, and H. Hajimirsadeghi, "Were RNNs all we needed?," 2024, *arXiv:2410.01201*.
- [15] C. M. Kellett, "A compendium of comparison function results," *Math. Control, Signals, Syst.*, vol. 26, no. 3, pp. 339–374, Sep. 2014.
- [16] Z.-P. Jiang and Y. Wang, "Input-to-state stability for discrete-time nonlinear systems," *Automatica*, vol. 37, no. 6, pp. 857–869, Jun. 2001.
- [17] F. Bayer, M. Burger, and F. Allgower, "Discrete-time incremental ISS: A framework for robust NMPC," in *Proc. ECC*, 2013, pp. 2068–2073.
- [18] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [19] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*. Philadelphia, PA, USA: SIAM, 2023.