Stable Compartmental Recurrent Neural Networks for Glucose–Insulin Dynamics Modeling in the Artificial Pancreas Framework

Stefano De Carli, Nicola Licini, Davide Previtali, Fabio Previdi, *Member, IEEE*, and Antonio Ferramosca, *Senior Member, IEEE*

***Abstract*— Models of glucose-insulin dynamics are essential in treating type 1 diabetes. Yet, the adoption of advanced methods such as Recurrent Neural Networks (RNNs), which provide superior predictive accuracy over traditional linear models, remains limited in safety-critical applications due to their black-box nature and inherent lack of theoretical guarantees. To overcome these limitations, recent advancements introduced the compartmental RNN for glucose prediction. This approach embeds the topology of the physiological system into the network, eliminating unphysiological cross-talk between insulin and meal effects, typical of purely black-box structures. However, these architectures still rely on standard RNNs, like gated recurrent unit networks, i.e., models that lack formal stability guarantees. In this work, we bridge this gap by introducing the stable compartmental RNN. Our proposal leverages chaos-free networks as internal compartmental units, ensuring that the resulting model inherently satisfies input-to-state stability by design, requiring no complex parametric constraints during network training. Numerical validation on the in silico UVA/Padova simulator confirms that the proposed architecture achieves high predictive accuracy while providing formal stability guarantees.**

***Index Terms*— Compartmental and positive systems, Grey-box modeling, stability of nonlinear systems.**

I. Introduction

TYPE 1 Diabetes Mellitus (T1DM) is characterized by the loss of endogenous insulin production, requiring lifelong exogenous insulin administration to regulate blood glucose levels [1]. The reliable estimation of glucose-insulin dynamics is a fundamental prerequisite for the safe and effective deployment of the Artificial Pancreas (AP), where control algorithms must anticipate the effects of exogenous insulin and carbohydrate intake while robustly handling intra- and inter-patient variability [2]. In model-based control frameworks, such as Model Predictive Control (MPC), the predictive model should not only be accurate, but also exhibit formal boundedness guarantees: a model that generates unbounded or diverging open-loop predictions can violate the Lyapunov decrease and recursive feasibility conditions on which closed-loop stability proofs are built [3].

Historically, linear compartmental models have been adopted due to their structural simplicity and intrinsic stability [4], but their limited expressivity does not capture the time-varying nature of human metabolism [5]. Recurrent Neural Networks (RNNs) [6, Ch. 15] have therefore emerged as powerful data-driven alternatives [7], [8], though purely black-box RNNs lack physiological interpretability and may exhibit unpredictable behaviors in closed-loop [9]. A significant step forward was the introduction of the Compartmental RNN (COMP-RNN) framework [10], which hard-codes physiological topology, decoupling the insulin and meal processing chains and eliminating unphysiological cross-talk. However, standard COMP-RNNs rely on standard gated architectures, such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks [6, Ch. 15], which may exhibit chaotic behavior with high sensitivity to initial conditions [11], providing no formal guarantee on how bounded inputs map to internal state trajectories. This is particularly consequential in the AP context: as established in robust predictive control, when the plant is subject to uncertainty, the stability and, possibly, the feasibility of the nominal MPC may be lost [3]; neural models with chaotic internal dynamics can amplify this risk by generating diverging open-loop predictions along the optimization horizon, leading to unphysiological glucose forecasts [9]. In a safety-critical context, such as the AP, conditions like prolonged exercise, extreme meal sizes, or high sensor noise may further amplify these modeling mismatches, leading to unphysiological predictions and potentially unsafe insulin dosing recommendations. While stability in RNNs can be promoted via ad hoc loss augmentation strategies [12], [13], these approaches typically introduce additional complexity and limit training efficiency. In contrast, a model that is

Received 17 March 2026; revised 17 May 2026; accepted 4 June 2026. Date of publication 17 June 2026; date of current version 23 June 2026. This work was supported by the National Plan for National Recovery and Resilience Plan (NRRP) Complementary Investments (Piano Nazionale Complementare (PNC), Established with the Decree-Law May 6, 2021, no. 59, Converted by Law no. 101 of 2021) in the Call for the Funding of Research Initiatives for Technologies and Innovative Trajectories in the Health and Care Sectors (Directorial Decree no. 931 of 06-06-2022) AdvaNced Technologies for Human-centrEd Medicine (ANTHEM) under Project PNC0000003. This article has not been presented at a conference. Recommended by Senior Editor S. Olaru. *(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; nicola.licini@unibg.it; davide.previtali@unibg.it; previdi@unibg.it; antonio.ferramosca@unibg.it).

Digital Object Identifier 10.1109/LCSYS.2026.3704428

intrinsically *Input-to-State Stability (ISS)* offers concrete and actionable guarantees. In particular, it ensures bounded state trajectories under bounded inputs and provides a structured framework for quantifying the effect of input disturbances on the system's evolution. These properties are especially valuable when integrating learned models within control pipelines, as they facilitate robustness analysis, support observer and estimator design, and contribute to the verification of multi-step predictions and safety properties [9].

In this letter, we bridge the gap between physiological interpretability and strict control-theoretic stability by introducing the *stable COMP-RNN*. We replace standard gated LSTM/GRU networks within the COMP-RNN framework with Chaos-Free Networks (CFNs) [11], an RNN designed to avoid chaotic dynamics, and prove that the resulting cascade interconnection satisfies ISS *by design*, requiring no changes in training. We validate our approach on the UVA/Padova T1DM simulator [14], confirming that the proposed stable COMP-RNNexhibits no statistically significant difference in predictive accuracy compared to the recent standard COMP-RNN [10], while providing formal stability guarantees.

The letter is structured as follows. Section II introduces the notation, the problem statement, and recalls the definition of ISS. Section III details the dynamics of the CFN architecture. Section IV defines the stable COMP-RNN along with a proof of its stability. Section V presents numerical experiments validating the theoretical findings on simulated patients from the UVA/Padova T1DM simulator [14]. Finally, Section VI draws conclusions and outlines future research.

II. Preliminaries and Problem Statement

A. Notation and Preliminaries

Let $\mathbb{R}$, $\mathbb{R}_{\geq 0}$, $\mathbb{R}_{>0}$, $\mathbb{Z}^+$, and $\mathbb{N}$ denote the sets of real, non-negative real, positive real, positive integers, and natural numbers (including zero), respectively. $\mathbb{R}^n$ is the set of real column vectors of dimension n, while $\mathbb{R}^{n\times m}$ is the set of real matrices of dimension $n \times m$, for $n, m \in \mathbb{Z}^+$. $|\mathcal{S}|$ is the cardinality of a set $\mathcal{S}$. Vectors are typeset in boldface lowercase, e.g., $\boldsymbol{v} \in \mathbb{R}^n$. We denote by $\mathbf{0}$, $\mathbf{1}$, and I the zero, all-ones, and identity matrices of appropriate dimension, respectively. The Hadamard product is indicated by $\circ$, and $\|\cdot\|$ denotes the infinity norm. An n-dimensional discrete-time signal is represented as $\boldsymbol{x}_k \in \mathbb{R}^n$, where $k \in \mathbb{N}$ is the time step. The sequence of values attained by $\boldsymbol{x}_k$ between $k = k_1$ and $k = k_2$ is denoted by $\{\boldsymbol{x}_k\}_{k=k_1}^{k_2}$. We denote the sigmoid and hyperbolic tangent activation functions as $\sigma(x) = (1 + e^{-x})^{-1} \in (0, 1)$ and $\tanh(x) \in (-1, 1)$, respectively, for $x \in \mathbb{R}$. For a scalar function $f : \mathbb{R} \to \mathbb{R}$, its bold counterpart $\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^n$ denotes the component-wise application. We employ standard definitions for comparison functions of classes $\mathcal{KL}$ and $\mathcal{K}_\infty$ [15].

B. Problem Statement

We consider the identification of the discrete-time nonlinear dynamics of the T1DM-patient system from data. The objective is to learn the map governing the evolution of the system output $\boldsymbol{y}_k \in \mathbb{R}^{n_y}$ given the input $\boldsymbol{u}_k \in \mathbb{R}^{n_u}$, with $n_y, n_u \in \mathbb{Z}^+$. In the artificial pancreas framework, the inputs are $\boldsymbol{u}_k = [\iota_k, \mu_k]^\top \in \mathcal{U}$, $n_u = 2$, where ι_k [U] is the exogenous insulin delivery and μ_k [g] is the carbohydrate intake. The input set is $\mathcal{U} = \mathcal{U}^\iota \times \mathcal{U}^\mu \subset \mathbb{R}_{\geq 0}^{n_u}$, where $\mathcal{U}^\iota, \mathcal{U}^\mu \subset \mathbb{R}_{\geq 0}$ are compact sets of admissible insulin and carbohydrate values, respectively. The measured output $y_k = \gamma_k \in \mathbb{R}_{\geq 0}$ [mg/dL] is the Blood Glucose (BG) concentration, thus $n_y = 1$. The underlying dynamics are characterized by a state $\boldsymbol{x}_k \in \mathcal{X} \subset \mathbb{R}^{n_x}$, where $\mathcal{X}$ is a compact set and $n_x \in \mathbb{Z}^+$. To identify these 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 $\mathcal{D}^{(v)} := \{(\boldsymbol{u}_k^{(v)}, y_k^{(v)})\}_{k=0}^{N^{(v)}-1}$, where $N^{(v)} \in \mathbb{Z}^+$ is the length of the v-th sequence and $v \in \{1, \dots, V\}$. The objective is to identify the T1DM-patient system with physiologically-structured RNNs that inherently satisfy ISS. Stability is defined for a generic system $\boldsymbol{x}_{k+1} = \boldsymbol{w}(\boldsymbol{x}_k, \boldsymbol{u}_k)$, with $\boldsymbol{w} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$, assuming $\mathcal{X}$ is forward invariant with respect to $\mathcal{U}$.

Definition 1 (ISS [13]): A system $\boldsymbol{x}_{k+1} = \boldsymbol{w}(\boldsymbol{x}_k, \boldsymbol{u}_k)$ is ISS in $\mathcal{X}$ with respect to $\mathcal{U}$ if there exist functions $\beta \in \mathcal{KL}$ and $\chi_u, \chi_b \in \mathcal{K}_\infty$ such that, for any $k \in \mathbb{N}$, any initial state $\boldsymbol{x}_0 \in \mathcal{X}$, any input sequence $\{\boldsymbol{u}_z \in \mathcal{U}\}_{z=0}^{k-1}$, and any bias vector $\boldsymbol{b} \in \mathbb{R}^{n_x}$, it holds that:

$$\|\boldsymbol{x}_k\| \leq \beta(\|\boldsymbol{x}_0\|, k) + \chi_u\left(\max_{0\leq z<k} \|\boldsymbol{u}_z\|\right) + \chi_b(\|\boldsymbol{b}\|). \tag{1}$$

III. Chaos-Free Networks

While LSTM and GRU architectures are widely adopted for their superior modeling capabilities [6], [16], their dynamics often exhibit chaotic behavior [11] and lack theoretical stability guarantees by design [12], [13]. To address these limitations, the CFN was proposed [11]. This architecture is designed to mitigate chaotic dynamics while preserving the ability to handle the vanishing gradient problem. A single-layer CFN is structured as a discrete-time nonlinear dynamical system characterized by a selectable number of hidden units $n_\hbar \in \mathbb{Z}^+$, with a hidden state $\boldsymbol{h}_k \in \mathbb{R}^{n_\hbar}$ that evolves according to the following equations:

$$\boldsymbol{h}_{k+1} = \boldsymbol{f}_k \circ \tanh(\boldsymbol{h}_k) + \boldsymbol{i}_k \circ \tilde{\boldsymbol{h}}_k, \tag{2a}$$
$$\boldsymbol{f}_k = \boldsymbol{\sigma}\left(W_f \boldsymbol{u}_k + R_f \boldsymbol{h}_k + \boldsymbol{b}_f\right) \in \mathcal{G}, \tag{2b}$$
$$\boldsymbol{i}_k = \boldsymbol{\sigma}\left(W_i \boldsymbol{u}_k + R_i \boldsymbol{h}_k + \boldsymbol{b}_i\right) \in \mathcal{G}, \tag{2c}$$
$$\tilde{\boldsymbol{h}}_k = \tanh\left(W_{\tilde{h}} \boldsymbol{u}_k + \boldsymbol{b}_{\tilde{h}}\right) \in \mathcal{H}, \tag{2d}$$

where $\boldsymbol{f}_k, \boldsymbol{i}_k$ and $\tilde{\boldsymbol{h}}_k$ are respectively the forget gate, input gate and candidate hidden state. The activation functions strictly bound the internal variables: the gates evolve within $\mathcal{G} := (0, 1)^{n_\hbar} \subset \mathbb{R}^{n_\hbar}$, and the candidate hidden state within $\mathcal{H} := (-1, 1)^{n_\hbar} \subset \mathbb{R}^{n_\hbar}$. As demonstrated in [17, Proposition 1], such range of the activation functions guarantees that, for any initial condition, the hidden state $\boldsymbol{h}_k$ is strictly bounded and evolves within the compact forward invariant set $\mathcal{H}_{\text{inv}} := [-2, 2]^{n_\hbar} \subset \mathbb{R}^{n_\hbar}$. Going back to (2), $W_f, W_i, W_{\tilde{h}} \in \mathbb{R}^{n_\hbar \times n_u}$ are the input weights, $R_f, R_i \in \mathbb{R}^{n_\hbar \times n_\hbar}$ are the recurrent weights,

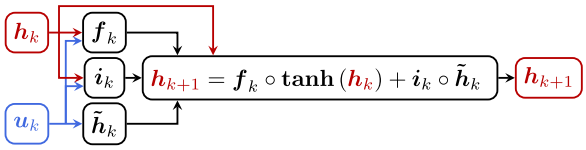

Fig. 1. Dynamics for the CFN layer in (2) at time k.

and $\boldsymbol{b}_f, \boldsymbol{b}_i, \boldsymbol{b}_{\tilde{h}} \in \mathbb{R}^{n_h}$ are the bias vectors. These single-layer CFN dynamics are illustrated in Fig. 1. The network output $\boldsymbol{y}_k \in \mathbb{R}^{n_y}$ is an affine mapping of $\boldsymbol{h}_{k+1}$:

$$\boldsymbol{y}_k = W_y \boldsymbol{h}_{k+1} + \boldsymbol{b}_y,$$

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

$$\theta := \{W_y, \boldsymbol{b}_y, W_f, W_i, W_{\tilde{h}}, R_f, R_i, \boldsymbol{b}_f, \boldsymbol{b}_i, \boldsymbol{b}_{\tilde{h}}\}. \tag{3}$$

Unlike traditional gated RNNs, the CFN's candidate hidden state $\tilde{\boldsymbol{h}}_k$ in (2d) only depends on the current input $\boldsymbol{u}_k$. However, the gating mechanisms retain full recurrent connectivity through the weights R_f and R_i. As demonstrated in [11, Lemma 2] and formalized in [17, Theorem 1], the CFN in (2) with parameters in (3) is ISS in $\mathcal{H}_{\text{inv}}$ with respect to a compact input set *by design*. This contrasts with the conditions for ISS of GRU and LSTM networks found in [12] and [13], which define constraints on the parameters of the respective networks and require custom training methodologies.

Remark 1: While CFNs can be extended to multi-layer architectures, this work focuses on a single-layer configuration. This choice is motivated by prior studies on glucose-insulin dynamics modeling [7], [8], [10], which suggest that a shallow architecture already captures the nonlinear dynamics while preventing overfitting.

IV. Stable Compartmental RNNs

In this Section, we present the stable COMP-RNN framework. This architecture builds upon the COMP-RNN introduced for the glucose-insulin system in [10]. While the original COMP-RNN embeds physiological topology to eliminate unphysiological cross-talk between meals and insulin, it relies on standard gated units (e.g., GRU and LSTM layers) that lack formal stability guarantees by design [12], [13]. We overcome this limitation by replacing the standard unstable recurrent units with CFNs, ensuring ISS of the architecture by design.

A. Reference Linear Compartmental Model

To provide a physiologically meaningful target for the network's internal states, we rely on a reference Linear Time-Invariant (LTI) compartmental model of the T1DM-patient system [4]. The model describes the metabolism through a 5-dimensional state vector $\boldsymbol{x}_k = [\gamma_k, \boldsymbol{x}_k^{\iota\top}, \boldsymbol{x}_k^{\mu\top}]^\top \in \mathbb{R}^5_{\geq 0}$. Here, γ_k is the BG concentration as specified in Section II-B, $\boldsymbol{x}_k^{\iota} \in \mathbb{R}^2_{\geq 0}$ [U/min] represents the insulin compartments, and $\boldsymbol{x}_k^{\mu} \in \mathbb{R}^2_{\geq 0}$ [g/min] models the carbohydrate absorption in the gut. The discrete-time dynamics are governed by:

$$\boldsymbol{x}_{k+1} = A_{\text{d}} \boldsymbol{x}_k + B_{\text{d}} \boldsymbol{u}_k + E_{\text{d}}, \tag{4}$$

where $\boldsymbol{u}_k$ is the exogenous input vector described in Section II-B, and the matrices $A_{\text{d}}, B_{\text{d}}, E_{\text{d}}$ are derived via forward Euler discretization of the continuous-time physiological parameters defined in [18].

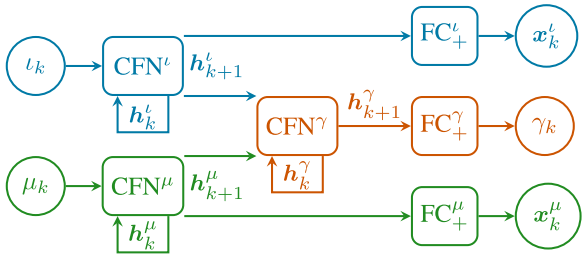

Fig. 2. Stable COMP-RNN architecture at time k.

B. Model Architecture

The stable COMP-RNN partitions the architecture into three distinct sub-networks: CFN$^\iota$, CFN$^\mu$, and CFN$^\gamma$, representing the physiological compartments $\Phi = \{\iota, \mu, \gamma\}$. The state-space formulation of the stable COMP-RNN adopts a physiological cascade interconnection:

$$\boldsymbol{h}^\phi_{k+1} = \text{CFN}^\phi(\boldsymbol{h}^\phi_k, \phi_k), \quad \phi \in \{\iota, \mu\}, \tag{5a}$$

$$\boldsymbol{h}^\gamma_{k+1} = \text{CFN}^\gamma(\boldsymbol{h}^\gamma_k, \tilde{\boldsymbol{u}}^\gamma_k), \quad \tilde{\boldsymbol{u}}^\gamma_k = \left[\boldsymbol{h}^{\iota}_{k+1}{}^\top, \boldsymbol{h}^{\mu}_{k+1}{}^\top\right]^\top, \tag{5b}$$

$$\boldsymbol{x}^\phi_k = \text{FC}^\phi_+(\boldsymbol{h}^\phi_{k+1}), \quad \phi \in \{\iota, \mu\}, \tag{5c}$$

$$\gamma_k = \text{FC}^\gamma_+(\boldsymbol{h}^\gamma_{k+1}). \tag{5d}$$

Each sub-network CFN$^\phi$, $\phi \in \Phi$, follows the dynamics defined in (2), with a hidden state $\boldsymbol{h}^\phi_k \in \mathcal{H}^\phi_{\text{inv}} := [-2, 2]^{n^\phi_h}$, where $n^\phi_h \in \mathbb{Z}^+$ denotes the number of hidden units of compartment ϕ. The aggregate hidden state of the full architecture in (5) is defined as:

$$\boldsymbol{h}^{\text{full}}_k = \left[\boldsymbol{h}^{\iota\top}_k, \boldsymbol{h}^{\mu\top}_k, \boldsymbol{h}^{\gamma\top}_k\right]^\top \in \mathcal{H}^{\text{full}}_{\text{inv}} = \mathcal{H}^\iota_{\text{inv}} \times \mathcal{H}^\mu_{\text{inv}} \times \mathcal{H}^\gamma_{\text{inv}} \subset \mathbb{R}^{n^{\text{full}}_h}, \tag{6}$$

where $n^{\text{full}}_h = n^\iota_h + n^\mu_h + n^\gamma_h$. The insulin and meal compartments independently process their respective inputs as in (5a). Their internal states are concatenated to form the input $\tilde{\boldsymbol{u}}^\gamma_k$ for the glucose compartment in (5b), which computes the mixed metabolic response, as illustrated in Fig. 2. Unlike standard black-box RNNs, the stable COMP-RNN in (5) provides a fully interpretable augmented output comprising γ_k, $\boldsymbol{x}^\iota_k$, and $\boldsymbol{x}^\mu_k$. Outputs are enforced to be physiologically plausible via Fully Connected layers with non-negative activation functions (FC$_+$). Each FC$^\phi_+$ reads as:

$$\text{FC}^\phi_+(\boldsymbol{h}^\phi_k) = \mathbf{max}(\mathbf{0}, W^\phi_y \boldsymbol{h}^\phi_k + \boldsymbol{b}^\phi_y), \quad \phi \in \Phi, \tag{7}$$

where $W^\phi_y \in \mathbb{R}^{n^\phi_y \times n^\phi_h}$ and $\boldsymbol{b}^\phi_y \in \mathbb{R}^{n^\phi_y}$ are the output weight matrix and bias vector of compartment ϕ, and $n^\phi_y \in \mathbb{Z}^+$ is the output dimension ($n^\iota_y = n^\mu_y = 2$, $n^\gamma_y = 1$). The full set of tunable parameters of the stable COMP-RNN is:

$$\theta := \bigcup_{\phi \in \Phi} \left\{W^\phi_y, \boldsymbol{b}^\phi_y, W^\phi_f, W^\phi_i, W^\phi_{\tilde{h}}, R^\phi_f, R^\phi_i, \boldsymbol{b}^\phi_f, \boldsymbol{b}^\phi_i, \boldsymbol{b}^\phi_{\tilde{h}}\right\}, \tag{8}$$

where the per-compartment parameters follow (3).

C. Knowledge-Guided Training

To ensure that the virtual states $\boldsymbol{x}^\iota_k$ and $\boldsymbol{x}^\mu_k$ act as meaningful physiological quantities rather than unconstrained latent variables, the network parameters θ in (8) are optimized using a loss augmentation approach [10]. Specifically, the training

process minimizes an augmented loss function $\mathcal{L}(\theta;\mathcal{D}_{\text{tr}})$, evaluated over a training dataset $\mathcal{D}_{\text{tr}} \subset \mathcal{D}$, which incorporates prior physiological knowledge:

$$\mathcal{L}(\theta;\mathcal{D}_{\text{tr}}) = \alpha_{\text{D}}\mathcal{L}^{\text{D}}(\theta;\mathcal{D}_{\text{tr}}) + \alpha_{\text{B}}\mathcal{L}^{\text{B}}(\theta;\mathcal{D}_{\text{tr}}) + \alpha_{\text{A}}\mathcal{L}^{\text{A}}(\theta;\mathcal{D}_{\text{tr}}), \quad (9)$$

with $\alpha_{\text{D}},\alpha_{\text{B}},\alpha_{\text{A}} \in \mathbb{R}_{\geq 0}$ being weighting parameters. The term $\mathcal{L}^{\text{D}}(\theta;\mathcal{D}_{\text{tr}})$ is the Mean Squared Error (MSE) of the glucose prediction. The biological loss $\mathcal{L}^{\text{B}}(\theta;\mathcal{D}_{\text{tr}})$ penalizes discrepancies between the network predictions and the fundamental mass-conservation laws captured by the reference LTI model in (4). Finally, the auxiliary loss $\mathcal{L}^{\text{A}}(\theta;\mathcal{D}_{\text{tr}})$ explicitly guides the internal states $\boldsymbol{x}_k^{\text{i}}$ and $\boldsymbol{x}_k^{\mu}$ to emulate the compartmental transients of the LTI baseline. The exact mathematical formulations of these loss terms are reported in [10]. They are omitted here for brevity, as the core contribution of this work concerns the structural stability of the architecture rather than the training procedure. Crucially, because the internal compartments are built upon inherently stable CFN units, this training process *does not* require additional complex penalty terms to promote stability constraints, as usually done for standard approaches [12], [13].

D. Theoretical Stability Analysis

We now demonstrate that embedding stable CFN units into the compartmental structural topology yields a global network architecture that is unconditionally ISS *by design*.

Theorem 1 (ISS of the stable COMP-RNN): Consider the stable COMP-RNN architecture in (5), with state $\boldsymbol{h}_k^{\text{full}} \in \mathcal{H}_{\text{inv}}^{\text{full}}$ defined in (6) and external inputs $\boldsymbol{u}_k = [\iota_k, \mu_k]^\top \in \mathcal{U}$. The stable COMP-RNN is ISS in $\mathcal{H}_{\text{inv}}^{\text{full}}$ with respect to $\mathcal{U}$.

The proof is reported in Appendix. The ISS property of the stable COMP-RNN has direct implications for its use within closed-loop control frameworks, notably MPC. By Definition 1, bounded insulin and carbohydrate inputs guarantee bounded internal states for any initial condition, regardless of prediction horizon length. This is a prerequisite for certifying the closed-loop stability of model-based controllers such as MPC, where recursive feasibility and stability proofs typically assume a bounded prediction model [3].

V. In Silico Results

In this Section, we evaluate the modeling capabilities of the proposed stable COMP-RNN. The primary objective is to show that the use of CFNs can provide unconditional structural stability without degrading the predictive accuracy with respect to standard COMP-RNNs [10].

A. Experimental Setup

We evaluate the models on the 10-patient adult cohort of the commercial UVA/Padova T1DM simulator [14], the FDA-recognized benchmark for replacing preclinical animal trials in insulin therapy research. To do so, we apply uniform input conditions to all patients and sample data at 5-minute intervals. The generated dataset $\mathcal{D}$ is partitioned into mutually exclusive training $\mathcal{D}_{\text{tr}}$, validation $\mathcal{D}_{\text{val}}$, and test $\mathcal{D}_{\text{tst}}$ sets. All datasets are generated according to standard clinical meal and bolus protocols. In particular, $\mathcal{D}_{\text{tst}}$ reproduces the original 3-day scenario reported in [19, Table 3], whereas $\mathcal{D}_{\text{tr}}$ extends the patterns

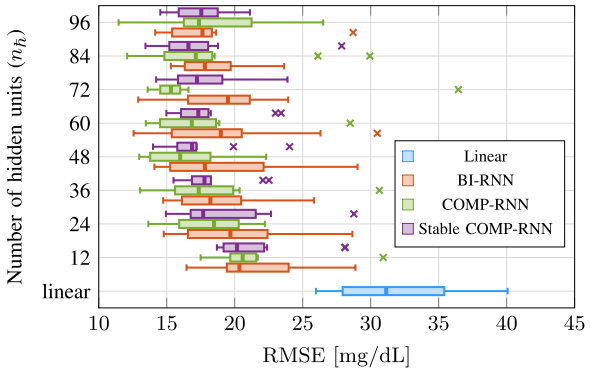

Fig. 3. Distribution of the RMSE evaluated on test data across the 10 virtual patients for the baseline linear model, BI-RNN, standard COMP-RNN, and the proposed stable COMP-RNN.

of [19, Table 2] over a 30-day horizon. The validation set $\mathcal{D}_{\text{val}}$ is instead obtained from a single 14-day simulation with three structured meals per day. To improve clinical realism, all datasets include intra-patient day-to-day variability by randomly perturbing meal timings ($\pm$20 minutes), carbohydrate amounts ($\pm$20%), and meal durations ($\pm$10 minutes). Insulin delivery follows a standard basal–bolus therapy regimen, while common patient-related errors are simulated through bolus miscalculations (over-/under-estimation) and administration delays ranging from 5 to 30 minutes. For training, the 30-day trajectory of $\mathcal{D}_{\text{tr}}$ is segmented into 113 overlapping 48-hour windows with a stride of 6 hours, enabling truncated backpropagation through time during optimization [6, Ch. 15]. It is crucial to note that models operate as plant simulators: predictions are performed over long horizons relying solely on exogenous inputs *without* autoregressive feedback (see the model in (5)).

B. Training Configuration and Evaluation Metrics

Inputs and outputs are normalized in $[-1, 1]$ prior to training to satisfy [17, Assumption 1]. We evaluate the models configuring the total number of hidden units from 12 to 96 in increments of 12. We distribute these hidden units uniformly across the three physiological compartments (e.g., a stable COMP-RNN with $n_h^{\text{full}} = 96$ has $n_h^{\phi} = 32$ for each $\phi \in \Phi$). To ensure a fair and direct comparison with standard COMP-RNNs, the training procedure and all the specific network hyperparameters follow the ones established in [10]. Specifically, the network parameters θ in (8) are optimized by minimizing the augmented loss function $\mathcal{L}(\theta;\mathcal{D}_{\text{tr}})$ introduced in Section IV-C. We set the loss weights to $\alpha_{\text{D}} = 0.5$, $\alpha_{\text{B}} = 0.25$, and $\alpha_{\text{A}} = 0.25$. Training is performed using the Adam optimizer [6, Ch. 8] with 2 mini-batches and an initial learning rate of 0.02, which is decayed by a factor of 0.8 every 500 epochs. To prevent overfitting, we apply a dropout rate of 0.1. The optimization runs for 1250 epochs, employing an early stopping strategy: we evaluate the performance on $\mathcal{D}_{\text{val}}$ every 15 iterations and save the parameter configuration θ^* that achieves the minimum MSE on $\mathcal{D}_{\text{val}}$. Performance is evaluated on the test sequence $\mathcal{D}^{(v)} \in \mathcal{D}_{\text{tst}}$ using the Root Mean Squared Error (RMSE) [mg/dL].

The comparison includes the BI-RNN [8], a single-layer unstructured GRU network trained with the same augmented loss in (9), using all n_h^{full} hidden units without compartmental

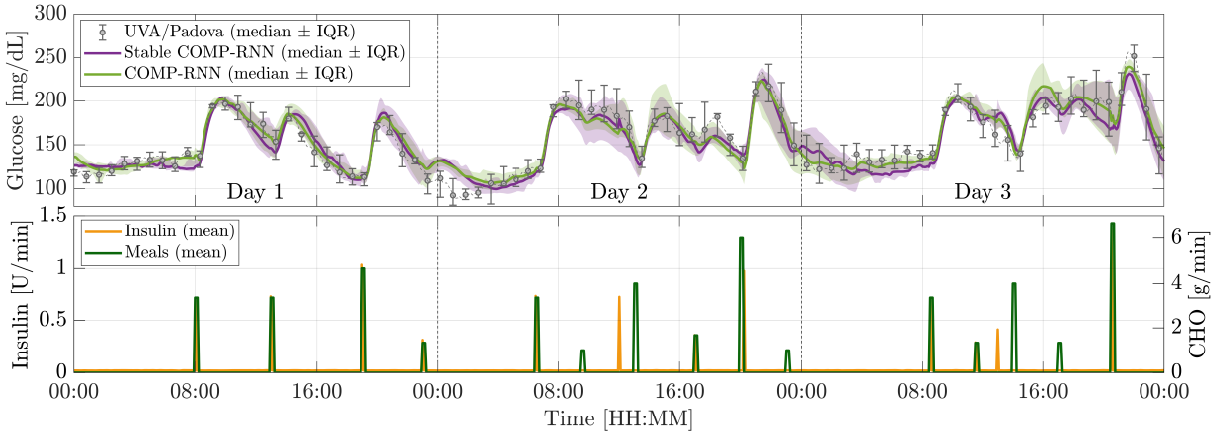

Fig. 4. Performance on $\mathcal{D}_{\text{tst}}$: COMP-RNN and stable COMP-RNN ($n_h^{\text{full}} = 48$) show high predictive accuracy for tracking BG on all patients.

partitioning. The standard COMP-RNN [10] shares the exact compartmental topology of the proposed model but relies on standard GRU layers, which lack formal stability guarantees and can exhibit chaotic behavior [11].

Crucially, because the stable COMP-RNN is intrinsically ISS by design (Theorem 1), its training process does not require stability-promoting penalty terms unlike standard stable RNN approaches [12], [13], significantly simplifying the optimization landscape.

C. Results and Discussion

We compare the stable COMP-RNN against the standard COMP-RNN [10] and the BI-RNN [8], introduced in Section V-B.

The predictive performance of the models across the 10-patient cohort is depicted in Fig. 3. As expected, the addition of physiological structure allows both compartmental architectures to consistently outperform the unstructured BI-RNN and the baseline linear model. Further, the box plots in Fig. 3 confirm that the compartmental structure also narrows the inter-patient error distribution. This confirms that physiological structure is the primary driver of predictive accuracy in this context, independently of the stability properties of the constituent RNNs. Most importantly, the stable COMP-RNN exhibits a predictive accuracy that strictly matches its unstable counterpart, with minor fluctuations in RMSE that fall well within the standard deviation bounds. This preservation of performance is *non-trivial*: replacing GRU networks with CFNs introduces a fundamentally different dynamic structure, which could, in principle, reduce model expressivity and degrade predictive performance. The fact that no such degradation is observed demonstrates that the physiological compartmental structure of the COMP-RNN framework is sufficiently rich to absorb the ISS constraint without loss of representational capacity. In other words, the ISS guarantee is achieved without structural compromises: it is a consequence of architectural design, not of a trade-off against accuracy. To evaluate this comparison, we conducted a two-tailed Wilcoxon Signed-Rank test on the paired RMSE values, which failed to reject the null hypothesis of zero median difference between the stable and standard COMP-RNN across all hidden unit configurations, confirming no statistically significant difference in predictive accuracy. Crucially, this result has a direct implication for closed-loop development: since the stable COMP-RNN is provably ISS (Theorem 1) and exhibits predictive performance with no statistically significant degradation relative to the standard COMP-RNN, adopting the proposed architecture provides additional theoretical guarantees without sacrificing predictive accuracy, which is particularly desirable in safety-critical settings [9]. Specifically, the ISS property ensures a "fading memory" behavior [11] that effectively bounds the impact of state perturbations which, in a closed-loop control scenario, naturally arise from observer estimation errors. Unlike standard architectures that can exhibit chaotic divergence [12], [13], this structural guarantee ensures that open-loop predictions remain bounded and reliable even in worst-case scenarios, providing the necessary robustness for safe closed-loop control [3].

The results further confirm that the proposed architecture successfully integrates two independent design benefits: ISS by design via CFNs and physiological consistency via compartmental decoupling, eliminating unphysiological cross-talk [8]. The performance reaches an optimal plateau at $n_h^{\text{full}} = 48$. Beyond this point, larger architectures do not yield significant predictive gains, suggesting that 48 units represent the ideal trade-off between model expressivity and parametric efficiency. The tracking capability of this optimal configuration ($n_h^{\text{full}} = 48$) is illustrated in Fig. 4, where both the COMP-RNN and the stable COMP-RNN accurately follow the BG transients of all patients, confirming their high fidelity across various post-prandial scenarios.

VI. Conclusion

In this work, we advance the compartmental modeling of glucose-insulin dynamics by introducing the stable COMP-RNN framework. By replacing standard gated units with CFNs, we bridge the gap between high-fidelity data-driven modeling and formal control-theoretic guarantees in the AP application. Our theoretical analysis demonstrates that the resulting architecture satisfies ISS by design, while in silico validation confirms that these stability properties are achieved without degrading predictive performance. Specifically, the numerical results indicate that the stable COMP-RNN exhibits no statistically significant difference in predictive accuracy compared to its unstable counterpart across all tested configurations, with an optimal trade-off between expressivity

and parametric efficiency reached at $n_{\hbar}^{\text{full}} = 48$ hidden units. Furthermore, the proposed model strictly maintains the physiological decoupling of insulin and meal pathways, eliminating the unphysiological cross-talk often observed in unstructured architectures [8]. Crucially, the intrinsic stability of the CFN units allows the training process to avoid complex stability-promoting regularization terms or strict parametric conditions, significantly simplifying the optimization landscape compared to traditional approaches [12], [13]. Being limited to an in silico 10-patient cohort, the validation in this letter is still partial but provides a foundation and calls for future clinical validation. Future work will integrate this approach within MPC frameworks, testing on real-world datasets, and including further metabolic factors such as physical activity [20].

Appendix
Proof of Theorem 1

From [17, Theorem 1], each CFN satisfies, for any $\phi \in \Phi$:

$$\left\|\boldsymbol{h}_k^{\phi}\right\| \leq \bar{\sigma}_f^{\phi}\left\|\boldsymbol{h}_{k-1}^{\phi}\right\| + \left\|W_{\tilde{h}}^{\phi}\right\|\left\|\boldsymbol{u}_{k-1}^{\phi}\right\| + \left\|\boldsymbol{b}_{\tilde{h}}^{\phi}\right\|, \tag{10}$$

where $\bar{\sigma}_f^{\phi} = \sigma\left(\left\|\left[2W_f^{\phi}\ 2R_f^{\phi}\ \boldsymbol{b}_f^{\phi}\right]\right\|\right) \in (0,1)$, and $\boldsymbol{u}_k^{\phi}$ denotes the input to compartment ϕ (i.e., ι_k, μ_k, or $\tilde{\boldsymbol{u}}_k^{\gamma}$). For the glucose compartment, since $\tilde{\boldsymbol{u}}_{k-1}^{\gamma} = [\boldsymbol{h}_k^{\iota\top}, \boldsymbol{h}_k^{\mu\top}]^{\top}$ by (5b), it holds that:

$$\left\|\tilde{\boldsymbol{u}}_{k-1}^{\gamma}\right\| = \max\left(\left\|\boldsymbol{h}_k^{\iota}\right\|, \left\|\boldsymbol{h}_k^{\mu}\right\|\right) \leq \left\|\boldsymbol{h}_k^{\iota}\right\| + \left\|\boldsymbol{h}_k^{\mu}\right\|.$$

Substituting this into (10) for $\phi = \gamma$, using (10) for $\phi \in \{\iota, \mu\}$, and denoting $\boldsymbol{p}_k := \left[\left\|\boldsymbol{h}_k^{\iota}\right\|, \left\|\boldsymbol{h}_k^{\mu}\right\|, \left\|\boldsymbol{h}_k^{\gamma}\right\|\right]^{\top}$, $\boldsymbol{q}_k := \left[|\iota_k|, |\mu_k|\right]^{\top}$, and $\boldsymbol{r} := \left[\left\|\boldsymbol{b}_{\tilde{h}}^{\iota}\right\|, \left\|\boldsymbol{b}_{\tilde{h}}^{\mu}\right\|, \left\|\boldsymbol{b}_{\tilde{h}}^{\gamma}\right\|\right]^{\top}$, the three inequalities can be written in the following matrix form:

$$\boldsymbol{p}_k \leq A\boldsymbol{p}_{k-1} + B\boldsymbol{q}_{k-1} + C\boldsymbol{r}, \tag{11}$$

where, denoting $\alpha^{\phi} := \left\|W_{\tilde{h}}^{\phi}\right\|$ and $\zeta^{\phi} := \bar{\sigma}_f^{\phi}$, the matrices $A \in \mathbb{R}_{\geq 0}^{3\times 3}$, $B \in \mathbb{R}_{\geq 0}^{3\times 2}$, $C \in \mathbb{R}_{\geq 0}^{3\times 3}$ are:

$$A = \begin{bmatrix} \zeta^{\iota} & 0 & 0 \\ 0 & \zeta^{\mu} & 0 \\ \alpha^{\gamma}\zeta^{\iota}\alpha^{\gamma}\zeta^{\mu}\zeta^{\gamma} & & \end{bmatrix}, B = \begin{bmatrix} \alpha^{\iota} & 0 \\ 0 & \alpha^{\mu} \\ \alpha^{\gamma}\alpha^{\iota}\alpha^{\gamma}\alpha^{\mu} & \end{bmatrix}, C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \alpha^{\gamma}\alpha^{\gamma}1 & & \end{bmatrix}.$$

The matrix A is lower triangular with eigenvalues $\bar{\sigma}_f^{\iota}$, $\bar{\sigma}_f^{\mu}$, $\bar{\sigma}_f^{\gamma}$, all strictly within the unit circle, making A Schur stable. Propagating (11) back to the initial state yields:

$$\boldsymbol{p}_k \leq A^k\boldsymbol{p}_0 + \max_{0\leq z<k}\|\boldsymbol{u}_z\|\sum_{z=0}^{k-1}A^zB\mathbf{1} + \sum_{z=0}^{k-1}A^zC\boldsymbol{r},$$

where, for all $z \in \{0,\ldots,k-1\}$, since $\mathcal{U} = \mathcal{U}^{\iota}\times\mathcal{U}^{\mu}$, we used $|\iota_z| \leq \|\boldsymbol{u}_z\|$ and $|\mu_z| \leq \|\boldsymbol{u}_z\|$, hence $[|\iota_z|, |\mu_z|]^{\top} \leq \mathbf{1}\|\boldsymbol{u}_z\|$, and then $\|\boldsymbol{u}_z\| \leq \max_{0\leq j<k}\left\|\boldsymbol{u}_j\right\|$. Since A, B, C are non-negative and A is Schur stable, the finite sums are upper-bounded by the respective infinite series, where $\sum_{z=0}^{\infty}A^z = (I-A)^{-1}$ [21, Ch. 7]. Denoting $\Gamma_B := (I-A)^{-1}B$ and $\Gamma_C := (I-A)^{-1}C$, applying the norm, and using $\left\|\boldsymbol{h}_k^{\text{full}}\right\| = \left\|[\left\|\boldsymbol{h}_k^{\iota}\right\|, \left\|\boldsymbol{h}_k^{\mu}\right\|, \left\|\boldsymbol{h}_k^{\gamma}\right\|]^{\top}\right\|$:

$$\left\|\boldsymbol{h}_k^{\text{full}}\right\| \leq \left\|A^k\right\|\left\|\boldsymbol{h}_0^{\text{full}}\right\| + \|\Gamma_B\|\max_{0\leq z<k}\|\boldsymbol{u}_z\| + \|\Gamma_C\|\left\|\boldsymbol{b}_{\tilde{h}}^{\text{full}}\right\|, \tag{12}$$

where $\boldsymbol{b}_{\tilde{h}}^{\text{full}} := \left[\boldsymbol{b}_{\tilde{h}}^{\iota\top}, \boldsymbol{b}_{\tilde{h}}^{\mu\top}, \boldsymbol{b}_{\tilde{h}}^{\gamma\top}\right]^{\top}$. Since A is Schur stable, leveraging its Jordan form [21, Ch. 7], there exists a $\omega \geq 1$ such that $\left\|A^k\right\| \leq \omega\sum_{z=0}^{m-1}\binom{k}{z}\left(\max_{\phi\in\Phi}\bar{\sigma}_f^{\phi}\right)^{k-z} \to 0$ as $k \to \infty$, where $m \in \mathbb{Z}^+$ is the size of the largest Jordan block of A. Therefore, the three terms in (12) identify $\beta^{\text{full}} \in \mathcal{KL}$ and $\chi_u^{\text{full}}, \chi_b^{\text{full}} \in \mathcal{K}_{\infty}$, matching Definition 1. ■

References

[1] A. Katsarou et al., "Type 1 diabetes mellitus," *Nat. Rev. Dis. Primers*, vol. 3, no. 1, p. 17016, 2017.

[2] N. Licini, M. A. Santos, F. Previdi, and A. Ferramosca, "A robust tube-based MPC approach for safer insulin dosing," *IFAC J. Syst. Control*, vol. 35, Mar. 2026, Art. no. 100393.

[3] D. Limón et al., *Input-to-State Stability: A Unifying Framework for Robust Model Predictive Control*. Berlin, Germany: Springer, 2009, pp. 1–26.

[4] Y. Ruan, M. E. Wilinska, H. Thabit, and R. Hovorka, "Modeling day-to-day variability of glucose–insulin regulation over 12-week home use of closed-loop insulin delivery," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 6, pp. 1412–1419, Jun. 2017.

[5] E. J. Mansell, P. D. Docherty, and J. G. Chase, "Shedding light on grey noise in diabetes modelling," *Biomed. Signal Process. Control*, vol. 31, pp. 16–30, Jan. 2017.

[6] K. P. Murphy, *Probabilistic Machine Learning*. Cambridge, MA, USA: MIT Press, 2022.

[7] F. Iacono, L. Magni, and C. Toffanin, "Personalized LSTM-based alarm systems for hypoglycemia and hyperglycemia prevention," *Biomed. Signal Process. Control*, vol. 86, Sep. 2023, Art. no. 105167.

[8] S. De Carli, N. Licini, D. Previtali, F. Previdi, and A. Ferramosca, "Integrating biological-informed recurrent neural networks for glucose-insulin dynamics modeling," *IFAC-PapersOnLine*, vol. 59, no. 2, pp. 91–96, 2025.

[9] 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, Apr. 2022.

[10] S. De Carli, N. Licini, D. Previtali, F. Previdi, and A. Ferramosca, "Knowledge-guided recurrent neural networks for glucose-insulin dynamics modeling," *IFAC J. Syst. Control*, vol. 2026, May 2026, Art. no. 100406.

[11] T. Laurent and J. H. V. Brecht, "A recurrent neural network without chaos," in *Proc. ICLR*, 2016, pp. 1–8.

[12] 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.

[13] 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. Eur. Control Conf. (ECC)*, Jun. 2025, pp. 911–916.

[14] C. D. Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli, "The UVA/PADOVA type 1 diabetes simulator: New features," *J. Diabetes Sci. Technol.*, vol. 8, no. 1, pp. 26–34, Jan. 2014.

[15] C. M. Kellett, "A compendium of comparison function results," *Math. Control, Signals, Syst.*, vol. 26, no. 3, pp. 339–374, Sep. 2014.

[16] G. Pillonetto et al., "Deep networks for system identification: A survey," *Automatica*, vol. 171, Mar. 2025, Art. no. 111907.

[17] S. De Carli, D. Previtali, M. Mazzoleni, and F. Previdi, "Chaos-free networks are stable recurrent neural networks," 2026, *arXiv:2603.14106*.

[18] P. Abuin, P. S. Rivadeneira, A. Ferramosca, and A. H. González, "Artificial pancreas under stable pulsatile MPC: Improving the closed-loop performance," *J. Process Control*, vol. 92, pp. 246–260, Feb. 2020.

[19] M. Messori et al., "Model individualization for artificial pancreas," *Comput. Methods. Programs Biomed.*, vol. 171, pp. 133–140, May 2019.

[20] N. Licini, B. Sonzogni, P. Abuin, F. Previdi, A. H. González, and A. Ferramosca, "Artificial pancreas under stable pulsatile model predictive control: Including the physical activity effect," *IFAC J. Syst. Control*, vol. 36, Jun. 2026, Art. no. 100414.

[21] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*. Philadelphia, PA, USA: SIAM, 2023.

Open Access funding provided by 'Universita Degli Studi di Bergamo' within the CRUI CARE Agreement