Journal of Process Control 166 (2026) 103810

Contents lists available at ScienceDirect

Journal of Process Control

journal homepage: www.elsevier.com/locate/jprocont

Post-training inherited disk-margin certification for local-affine neural network controllers trained from PID teachers

Sebastián Madrigal [a,*], Orlando Arrieta [a,b], Ramon Vilanova [a]

[a] Departament de Telecomunicació i Enginyeria de Sistemes, Universitat Autònoma de Barcelona, 08193 Bellaterra, Barcelona, Spain
[b] Instituto de Investigaciones en Ingeniería, Facultad de Ingeniería, Universidad de Costa Rica, 11501-2060, San José, Costa Rica

ARTICLE INFO

Keywords:
Neural network control
Imitation learning
PID control
Disk margin
Integral quadratic constraints
Linear matrix inequalities

ABSTRACT

Neural-network (NN) controllers trained by imitation learning can reproduce the behavior of conventional PID controllers, but their deployment requires assessing whether the stability and robustness properties of the teacher are preserved after replacement. This paper proposes a post-training framework to certify robustness inheritance when a PID is replaced by a NN controller. The method uses the balanced disk margin of the PID-controlled loop as the reference robustness measure and evaluates the robustness level inherited by the local-affine realization of the trained NN around the nominal operating point. The certificate is formulated as an integral quadratic constraint (IQC) condition expressed through a linear matrix inequality (LMI), whose feasibility certifies a plant-side robustness bound for the local-affine NN controller closed-loop system. The framework is evaluated on the Van de Vusse continuous stirred-tank reactor (CSTR), an inverse response nonlinear benchmark. The results show that the proposed pipeline, combining imitation training, autonomous validation, local-affine LMI certification, and nonlinear simulations, identifies NN controllers that preserve autonomous viability and inherit a significant fraction of the PID teacher disk margin.

1. Introduction

Neural-network (NN) controllers have become an active research topic in feedback control because they can approximate nonlinear input–output mappings and reproduce control policies generated by model-based or expert controllers. In process control applications, this is particularly relevant when a reliable conventional controller is already available and can be used as a teacher, enabling a gradual transition from classical to learning-based control. Beyond replacing a single controller, the learned representation may support future adaptation or transfer to related control loops, plants, or operating points. Previous studies have therefore investigated transfer suitability and controller adaptation [1–3].

However, replacing or adapting a conventional controller by a learned one requires more than accurate supervised imitation or improved post-transfer performance. The autonomous NN closed-loop system may deviate from the teacher-controlled loop when operating under its own induced state distribution, and adaptation may improve local performance while degrading robustness properties inherited from the original controller [4,5]. Therefore, learned-controller viability should also be assessed from a stability and robustness perspective. This motivates post-training certificates capable of quantifying robustness inheritance, so that a learned controller can be evaluated not only by its closed-loop performance, but also by the robustness that it preserves.

Within this broader need for reliable learned feedback control, the stability analysis of NN controllers has received significant attention in recent years. One influential direction formulates NN-controlled systems through Integral Quadratic Constraints (IQC), enabling stability analysis of feedback loops containing static NN controllers [6]. Related developments have considered NN controllers for set-point tracking problems, where stability and offset-free behavior must be preserved despite the use of learned feedback policies [7,8]. Recurrent NN controllers have also been studied under closed-loop synthesis frameworks that aim to provide stability guarantees in partially observed settings [9,10]. Other works have extended quadratic-constraint ideas to event-triggered NN control of perturbed systems, showing how robust-control tools can be adapted to learned feedback laws [11]. Beyond quadratic-constraint formulations, several studies have focused on the structural and regional properties of NN-controlled systems. For piecewise-affine NN controllers, stability can be analyzed by exploiting the piecewise-affine nature of the closed-loop map and deriving conditions for reliable stabilization [12]. Regional stability conditions have also been proposed for recurrent NN-based control systems, emphasizing that guarantees for learned controllers are often valid only

* Corresponding author.
E-mail addresses: sebastian.madrigal@uab.cat (S. Madrigal), orlando.arrieta@ucr.ac.cr (O. Arrieta), ramon.vilanova@uab.cat (R. Vilanova).

https://doi.org/10.1016/j.jprocont.2026.103810
Received 2 June 2026; Received in revised form 10 July 2026; Accepted 27 July 2026
Available online 4 August 2026
0959-1524/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

within prescribed operating regions [13]. Robust stability under parametric uncertainty has been addressed through interval matrix descriptions, providing another route for certifying NN-controlled systems under model uncertainty [14]. Broader surveys and recent work further emphasize that stability and robustness guarantees are essential for deploying learned controllers in feedback operation [15,16]. This viewpoint is reflected in a recent review work related to stability and robustness guarantees for NN-in-the-loop systems [17].

A complementary line of research investigates how stability or dissipativity requirements can be incorporated during NN controller synthesis or training. For example, closed-loop dissipativity conditions have been used as a design objective when synthesizing NN controllers with formal guarantees [18,19]. Related work has also studied the interaction between stability and learning in networked dynamical systems, including applications involving physics-informed neural networks [20]. These approaches are valuable because they seek to shape the controller during the learning process. However, they differ from scenarios in which a NN controller has already been trained from an existing controller and must be evaluated after training.

Robustness margins provide a natural bridge between classical control analysis and NN controller certification. Disk margins characterize simultaneous gain and phase variations through a complex multiplicative perturbation, offering a unified alternative to treating gain and phase margins separately [21]. This is particularly relevant for PID-controlled loops, where margin-based robustness measures are commonly used to assess the reliability of a feedback design. Recent work has begun to connect disk-margin concepts with NN controller analysis by studying stability margins of NN-controlled systems through robustness-oriented formulations [22]. These results provide an important foundation for relating classical robustness guarantees to learned controllers. Nevertheless, much of the existing literature either analyzes the full nonlinear NN controller through sector or quadratic constraints, or aims to impose stability-related properties during controller synthesis.

This paper adopts a post-training certification perspective. Rather than synthesizing a NN controller subject to robustness constraints, the goal is to evaluate whether and in which way an already trained NN controller inherits the robustness margin of a PID teacher. The NN controller is trained by imitation learning from closed-loop data generated by the PID controller. This setting is practically relevant because it uses an existing stabilizing controller as the source of training data. At the same time, imitation learning may suffer from distribution-shift effects, since a controller trained on teacher-generated trajectories can behave differently when deployed autonomously under its own induced closed-loop state distribution [23–25]. Consequently, one-step imitation accuracy is not sufficient to justify replacement of the PID controller.

The proposed framework combines autonomous closed-loop validation with local robustness certification. After training, only NN controllers that remain viable without the PID teacher are retained, and each NN controller is then replaced by its frozen-slope local-affine realization around the nominal operating point. This produces an LTI controller $K_{N,0}(z)$ associated with a fixed activation pattern. The balanced disk-margin robustness channel of the resulting local-affine NN-controlled loop is then certified through a plant-side IQC/LMI condition based on the discrete-time bounded-real lemma [26]. The certified radius is finally compared with the balanced disk-margin radius of the PID teacher to quantify the inherited robustness fraction. The resulting certificate is local by construction. It applies to the local-affine realization of the trained NN controller around the selected operating point. This local formulation keeps the certificate directly connected to the classical disk-margin computation while providing a tractable post-training robustness measure. Nonlinear operation is subsequently employed as an independent validation step, unlike approximation based designs for uncertain nonlinear systems [27–29], the nonlinear process is not approximated by the NN, but rather is simulated directly from its differential equations, to assess whether the locally certified behavior remains meaningful under set-point changes and disturbances.

Fig. 1 in Section 2 situates the proposed certification method within a broader knowledge-transfer framework for control systems. In this context, the present work focuses on establishing whether a PID-trained NN controller can operate autonomously and preserve a certifiable fraction of the teacher's robustness margin before any reuse or adaptation is considered. Although the framework is instantiated here on the nonlinear Van de Vusse CSTR, its formulation is applicable to other industrial processes. Transfer-suitability assessment and adaptation to related target loops remain outside the scope of this paper; representative approaches are reported in [1–3].

Within this scope, the distinction with respect to existing IQC-based, piecewise-affine, and NN margin-analysis approaches is that the robustness reference is inherited from an existing controller. Rather than imposing a new robustness condition during NN synthesis, the proposed framework is applied after imitation training to quantify the fraction of the PID teacher's balanced disk margin preserved by its NN replacement. This control-theoretic certification stage provides a prerequisite for the reliable reuse or future transfer of the learned controller across related feedback loops. The main contributions are therefore summarized as follows:

- A teacher referenced robustness-inheritance metric is introduced to quantify the fraction of the PID teacher's balanced disk-margin preserved by an NN controller trained through imitation learning.
- A post-training local-affine certification method is developed by freezing the NN activation pattern around the certified operating point and verifying the inherited margin through a plant-side IQC/LMI condition.
- The complete workflow is evaluated on the nonlinear Van de Vusse CSTR benchmark, combining autonomous closed-loop screening, local-affine certification, and nonlinear servo-regulatory validation.

The rest of the paper is organized as follows. Section 2 formulates the robustness-inheritance problem, Section 3 presents the IQC/LMI certificate, and Section 4 details the CSTR implementation and NN training. Section 5 reports the validation and certification results, while Section 6 summarizes the main conclusions and future work.

2. Problem formulation

This work considers the post-training robustness certification of an NN controller designed by imitation learning from a PID teacher. The overall setting is illustrated in Fig. 1. During training, the PID controller K_T operates the nominal plant P and generates closed-loop data used to train an NN controller $K_N = f_\theta(\cdot)$ that approximates the teacher control action. During autonomous deployment, the PID controller is no longer present and the plant is controlled only by the trained NN. The objective is to determine whether the NN controller preserves a certifiable fraction of the balanced disk margin originally available for the PID teacher. Rather than certifying the complete nonlinear piecewise-affine NN globally, the proposed method considers the local-affine realization $K_{N,0}(z)$ obtained by freezing the NN activation slopes around a selected operating regressor $\xi^\star$. The resulting certificate is therefore local with respect to the activation pattern, while remaining directly comparable with the robustness margin of the PID teacher.

The Fig. 1 provides a two-level view of the proposed methodology. The upper part shows how the certification problem fits within a broader controller-knowledge-transfer workflow, whereas the lower part expands the specific sequence formalized in this section: closed-loop data generation with the PID teacher, imitation-based synthesis of K_N, autonomous deployment, construction of the local-affine model $K_{N,0}(z)$, and teacher-referenced disk-margin certification. Accordingly, the scheme serves as a roadmap for the formulation developed in the following subsections.

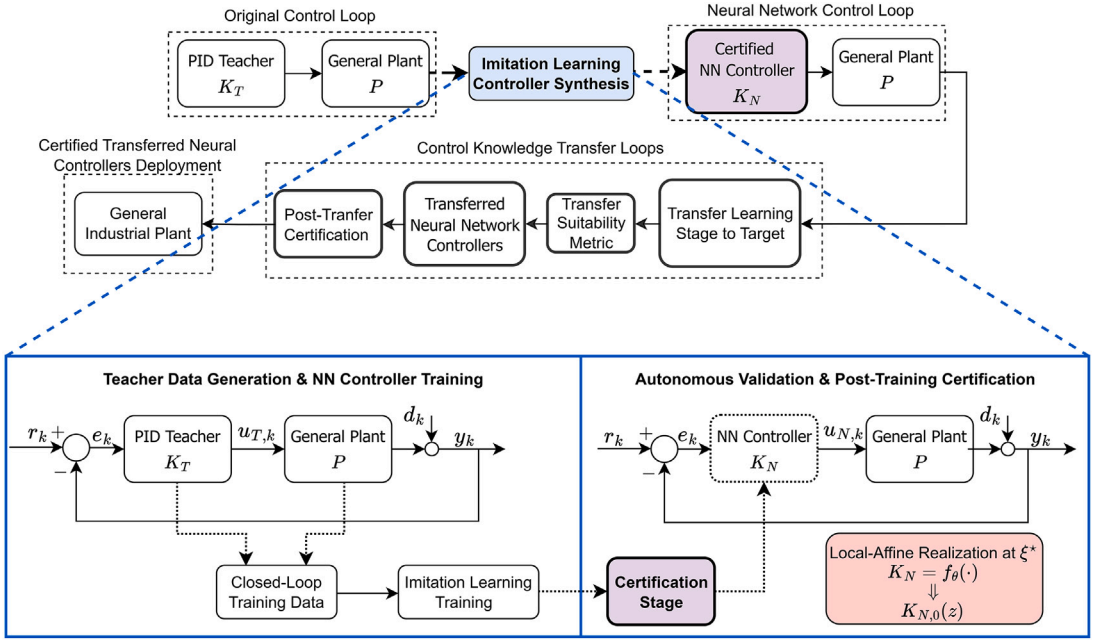

Fig. 1. General knowledge-transfer framework for control systems and scope of the proposed PID-to-NN certification procedure.

2.1. Nominal plant and PID teacher

Consider a discrete-time LTI representation of the system:

$$\begin{aligned} x_{p,k+1} &= A_p x_{p,k} + B_p u_k \\ y_k &= C_p x_{p,k} \end{aligned} \tag{1}$$

where $x_{p,k} \in \mathbb{R}^{n_p}$ is the plant state, $u_k \in \mathbb{R}^{n_u}$ is the manipulated input, and $y_k \in \mathbb{R}^{n_y}$ is the measured output. In the implementation considered later, the system is SISO, although the formulation is written in state-space form to make the interconnection structure explicit. The control error is defined as $e_k = r_{T,k} - y_{T,k}$, where $r_{T,k}$ is the reference signal. The PID teacher is represented by the discrete-time state-space:

$$\begin{aligned} x_{T,k+1} &= A_T x_{T,k} + B_T e_k \\ u_{T,k} &= C_T x_{T,k} + D_T e_k \end{aligned} \tag{2}$$

where $x_{T,k} \in \mathbb{R}^{n_T}$ collects the internal states of the implemented PID controller, including the integral state and derivative-filter state. The input is the tracking error $e_k = r_{T,k} - y_{T,k}$, and the output $u_{T,k}$ is the PID teacher control action, $u_k = u_{T,k}$. The matrices A_T, B_T, C_T, and D_T therefore correspond to the chosen discrete-time PID implementation and are used only to express the teacher in a compact form for plant interconnection and comparison with the NN controller. Combining (1) and (2), the nominal PID-controlled closed-loop is written as:

$$\begin{aligned} \begin{bmatrix} x_{p,k+1} \\ x_{T,k+1} \end{bmatrix} &= \underbrace{\begin{bmatrix} A_p - B_p D_T C_p & B_p C_T \\ -B_T C_p & A_T \end{bmatrix}}_{A_{T,\mathrm{cl}}} \begin{bmatrix} x_{p,k} \\ x_{T,k} \end{bmatrix} + \underbrace{\begin{bmatrix} B_p D_T \\ B_T \end{bmatrix}}_{B_{T,\mathrm{cl}}} r_{T,k} \\ y_{T,k} &= \underbrace{\begin{bmatrix} C_p & 0 \end{bmatrix}}_{C_{T,\mathrm{cl}}} \begin{bmatrix} x_{p,k} \\ x_{T,k} \end{bmatrix} \end{aligned} \tag{3}$$

This nominal closed-loop model provides the reference behavior that the NN controller is trained to reproduce and also defines the loop for which the PID robustness margin is computed.

2.2. Disk margin as stability measure

The central stability and robustness measure used in this work is the disk margin. Unlike classical gain and phase margins treated separately, disk margins describe simultaneous gain and phase variations through a complex multiplicative perturbation factor. Let $L_T(z) = P(z)K_T(z)$ denote the nominal loop transfer function associated with the PID teacher. A perturbed loop is written as $L_f(z) = fL_T(z)$, where $f \in \mathbb{C}$ is a complex multiplicative uncertainty factor. The disk-margin framework defines the admissible perturbations as the disk set [21]:

$$f \in D(\alpha, \sigma) = \left\{ \frac{2 + (1-\sigma)\delta}{2 - (1+\sigma)\delta} \,:\, \delta \in \mathbb{C},\ |\delta| < \alpha \right\} \tag{4}$$

where $\alpha > 0$ determines the size of the uncertainty set and $\sigma \in \mathbb{R}$ is the skew parameter. The nominal case corresponds to $f = 1$. The real-axis intercepts of the disk are:

$$\gamma_{\max} = \frac{2 + \alpha(1-\sigma)}{2 - \alpha(1+\sigma)}, \qquad \gamma_{\min} = \frac{2 - \alpha(1-\sigma)}{2 + \alpha(1+\sigma)} \tag{5}$$

These quantities define the gain increase and gain decrease compatible with the same disk, while the tangent geometry of the disk determines the corresponding phase variation. Hence, $D(\alpha, \sigma)$ represents a coupled gain-phase uncertainty family rather than independent gain and phase perturbations. For a fixed skew σ, the largest stable disk size can be written as:

$$\alpha_{\max} = \left\| S_T + \frac{\sigma - 1}{2} \right\|_\infty^{-1} \tag{6}$$

where $S_T(z) = \left(1 + L_T(z)\right)^{-1}$ is the sensitivity function of the nominal PID-controlled loop. In this work, the balanced disk margin is employed, corresponding to $\sigma = 0$. In this case:

$$\alpha_T = \left\| \frac{S_T - T_T}{2} \right\|_\infty^{-1}, \qquad T_T(z) = L_T(z)\left(1 + L_T(z)\right)^{-1} \tag{7}$$

The quantity α_T is interpreted as the robustness level of the PID teacher under balanced simultaneous gain and phase uncertainty.

2.3. Imitation learning neural controller

The NN controller is trained to reproduce the PID control action using data generated by the teacher. Its autonomous control law is written as:

$$u_{N,k} = f_\theta(\xi_k) \tag{8}$$

where f_θ denotes the trained NN with parameters θ, and ξ_k is a causal regressor constructed from available closed-loop signals. In this work,

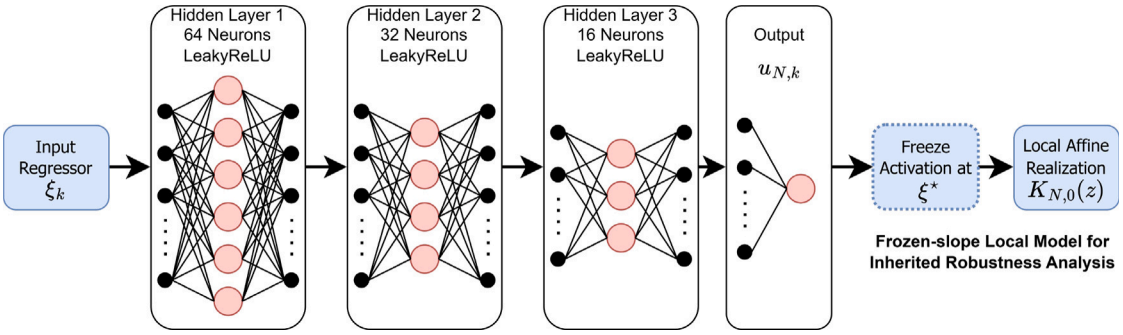

Fig. 2. Feedforward Multilayer Perceptron (MLP) neural-network architecture used for PID imitation.

the complete NN input regressor is defined as:

$$\xi_k = \begin{bmatrix} e_k & e_{k-1} & \cdots & e_{k-w_s+1} & i_{e,k} & u_{k-1} \end{bmatrix}^\top \tag{9}$$

where $e_k = r_k - y_k$ is the current control error, w_s is the error-window length, $i_{e,k}$ is the accumulated control error, and u_{k-1} is the previous control input. The integral state is updated according to $i_{e,k+1} = i_{e,k} + T_s e_k$, where T_s is the sampling period. The inclusion of u_{k-1} gives the feedforward NN controller a causal memory of the previous control action without introducing an algebraic loop. For the purpose of deriving a dynamic realization of the local-affine controller, the NN regressor is decomposed into the instantaneous control error and a regressor-state vector:

$$\xi_k = \begin{bmatrix} e_k \\ x_{R,k} \end{bmatrix} \tag{10}$$

where

$$x_{R,k} = \begin{bmatrix} e_{k-1} & \cdots & e_{k-w_s+1} & i_{e,k} & u_{k-1} \end{bmatrix}^\top \tag{11}$$

collects the memory terms of the regressor. Therefore, ξ_k is the complete static input vector of the NN, while $x_{R,k}$ is an auxiliary dynamic state containing the delayed errors, the accumulated error, and previous control input.

The dynamics of $x_{R,k}$ can be written compactly as:

$$x_{R,k+1} = A_R x_{R,k} + B_R e_k + E_R u_k \tag{12}$$

where A_R shifts the delayed-error states and propagates the integral state, B_R injects the current error e_k, and E_R updates the previous-control component using the current control action u_k. This representation separates the controller input e_k from the memory variables stored in $x_{R,k}$, which will be used later to convert the local-affine approximation of the NN into an equivalent dynamic LTI controller. The feedforward NN architecture used for the learned controllers is summarized in Fig. 2. Although the map $f_\theta(\cdot)$ is static, the autonomous NN controller is dynamic because its input regressor ξ_k contains delayed errors, the accumulated error, and the previous control input. During autonomous NN operation, the applied control action is $u_k = u_{N,k}$. Therefore, combining the regressor dynamics in (12) with the NN map gives the nonlinear state-space representation:

$$K_N : \begin{cases} x_{R,k+1} = A_R x_{R,k} + B_R e_k + E_R u_{N,k} \\ \xi_k = \begin{bmatrix} e_k \\ x_{R,k} \end{bmatrix} \\ u_{N,k} = f_\theta(\xi_k) \end{cases} \tag{13}$$

Equivalently, K_N is a nonlinear dynamic controller induced by the interconnection between the regressor-update dynamics and the feedforward NN. The local-affine realization derived next is obtained by freezing the activation pattern of this nonlinear controller around the certified operating regressor $\xi^\star$.

2.3.1. Local-affine realization

Although the deployed controller is the nonlinear NN $K_N = f_\theta(\cdot)$, the post-training robustness certificate developed in this work is applied to its local-affine realization around a certified operating point. This construction relies on the fact that feedforward networks with piecewise-linear activations define piecewise-affine input–output maps over regions determined by their activation patterns [30]. Since LeakyReLU is also piecewise linear, fixing the activation pattern around the certified operating point yields an affine representation of the trained NN restricted to that region. Let $\xi^\star$ denote the equilibrium regressor. In the nominal operating point used for certification, all signals are expressed as deviation variables, so that the equilibrium satisfies:

$$e_k^\star = 0, \qquad u_{N,k}^\star = 0 \tag{14}$$

The equilibrium regressor is chosen such that $f_\theta(\xi^\star) = 0$. Assume that no activation unit lies exactly at the kink of its activation function at $\xi^\star$. Then, within the corresponding fixed activation region, the NN is locally affine. In deviation variables:

$$\tilde{u}_{N,k} = K_{\xi,0}\tilde{\xi}_k \tag{15}$$

where

$$K_{\xi,0} = \left.\frac{\partial f_\theta}{\partial \xi}\right|_{\xi=\xi^\star} \tag{16}$$

is the frozen NN Jacobian evaluated at the operating point.

Consider a feedforward NN with L hidden layers. Let $z_{\ell,k}$ denote the output of hidden layer ℓ, and let $v_{\ell,k}$ denote its corresponding pre-activation vector. The NN is described recursively as:

$$z_{0,k} = \xi_k \tag{17}$$

$$v_{\ell,k} = W_\ell z_{\ell-1,k} + b_\ell, \quad z_{\ell,k} = \varphi_\ell(v_{\ell,k}), \quad \ell = 1, \ldots, L \tag{18}$$

and

$$u_{N,k} = W_o z_{L,k} + b_o \tag{19}$$

where W_ℓ and b_ℓ are the weight matrix and bias vector of hidden layer ℓ, W_o and b_o are the output-layer parameters, and $\varphi_\ell(\cdot)$ is the component-wise activation function of layer ℓ. In the implementation used in this work, the hidden layers employ LeakyReLU activations, a rectified activation family with nonzero negative-side slope [31,32]. For a LeakyReLU activation with negative slope a_ℓ, the activation is defined as:

$$\varphi_\ell(v) = \max(v, 0) + a_\ell \min(v, 0), \qquad 0 < a_\ell < 1 \tag{20}$$

The local-affine gain is obtained by freezing the activation slopes at the certified operating point. If $v_\ell^\star$ denotes the pre-activation vector of layer ℓ evaluated at $\xi^\star$, then

$$S_\ell = \operatorname{diag}\left(\varphi_\ell'(v_\ell^\star)\right), \qquad \ell = 1, \ldots, L \tag{21}$$

For the LeakyReLU activation, each diagonal entry of S_ℓ can be given by:

$$[S_\ell]_{ii} = \begin{cases} 1, & [v_\ell^\star]_i > 0 \\ a_\ell, & [v_\ell^\star]_i < 0 \end{cases} \tag{22}$$

Neurons satisfying $[v_\ell^\star]_i = 0$ lie exactly at the kink and are excluded from the local-affine differentiability assumption. Under this fixed activation pattern, the frozen Jacobian of the NN with respect to the physical regressor is:

$$K_{\xi,0} = W_o S_L W_L S_{L-1} W_{L-1} \cdots S_1 W_1^{\text{bar}} \tag{23}$$

where W_1^{bar} denotes the first-layer weight matrix after absorbing any input normalization into the affine map. If no input normalization is used, then $W_1^{\text{bar}} = W_1$. The local-affine gain is partitioned according to the regressor structure as:

$$K_{\xi,0} = \begin{bmatrix} k_0 & k_1 & \cdots & k_{w_s-1} & k_i & k_u \end{bmatrix} \tag{24}$$

Using the regressor state $x_{R,k}$ in (11), the local-affine NN controller can be written in deviation variables as:

$$\tilde{u}_{N,k} = k_0 \tilde{e}_k + K_R \tilde{x}_{R,k} \tag{25}$$

where

$$K_R = \begin{bmatrix} k_1 & \cdots & k_{w_s-1} & k_i & k_u \end{bmatrix} \tag{26}$$

Similarly, the regressor dynamics around the certified operating point are:

$$\tilde{x}_{R,k+1} = A_R \tilde{x}_{R,k} + B_R \tilde{e}_k + E_R \tilde{u}_{N,k} \tag{27}$$

Substituting (25) into (27) gives the LTI realization:

$$K_{N,0}(z) : \begin{cases} \tilde{x}_{R,k+1} = A_N \tilde{x}_{R,k} + B_N \tilde{e}_k \\ \tilde{u}_{N,k} = C_N \tilde{x}_{R,k} + D_N \tilde{e}_k \end{cases} \tag{28}$$

with $A_N = A_R + E_R K_R$, $B_N = B_R + E_R k_0$, $C_N = K_R$, and $D_N = k_0$. Here, $\tilde{x}_{R,k} = x_{R,k} - x_R^\star$, $\tilde{e}_k = e_k - e_k^\star$, and $\tilde{u}_{N,k} = u_{N,k} - u_{N,k}^\star$ denote deviations with respect to the certified operating point. The local-affine realization is associated with the activation region induced by the certified regressor $\xi^\star$. This region is not selected manually, once $\xi^\star$ is fixed from the nominal plant operating point, the NN activation pattern uniquely determines the corresponding activation cell. At this point, the deviation variables satisfy (14), the delayed errors are zero, and $i_e^\star$ can be chosen to satisfy the equilibrium consistency condition $f_\theta(\xi^\star) \approx 0$. The associated activation region is then:

$$\mathcal{R}(\xi^\star) = \left\{ \xi : v_{\ell,i}(\xi)\, v_{\ell,i}(\xi^\star) > 0,\ \forall \ell, i \right\} \tag{29}$$

where $v_{\ell,i}(\xi)$ is the pre-activation of neuron i in hidden layer ℓ. For all $\xi \in \mathcal{R}(\xi^\star)$, the LeakyReLU activation pattern, and hence the frozen-slope matrices S_ℓ, remain unchanged, so the NN is represented by the same local-affine realization $K_{N,0}(z)$. Points with $v_{\ell,i}(\xi^\star) = 0$ are excluded because the frozen-slope representation is not uniquely defined at activation kinks.

2.4. Local-affine inherited disk-margin

Let $D(\alpha_T, 0)$ denote the balanced disk-margin uncertainty set certified for the PID teacher. The objective of the proposed framework is to quantify how much of this teacher robustness is inherited by the local-affine realization $K_{N,0}(z)$ of the trained NN controller. The local-affine NN loop transfer function is $L_N(z) = P(z)K_{N,0}(z)$. The corresponding sensitivity and complementary sensitivity functions are:

$$S_N(z) = \left(1 + L_N(z)\right)^{-1}, \quad T_N(z) = L_N(z)\left(1 + L_N(z)\right)^{-1} \tag{30}$$

For the balanced disk-margin case $\sigma = 0$, the relevant robustness channel is:

$$G_N^{\text{DM0}}(z) = \frac{S_N(z) - T_N(z)}{2} \tag{31}$$

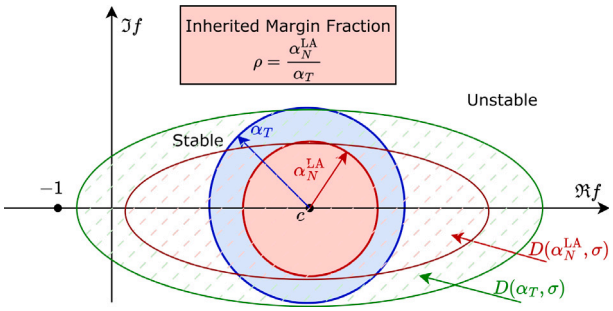

Fig. 3. Geometric interpretation of the inherited local-affine disk-margin concept. The outer disk $D(\alpha_T, 0)$ represents the balanced disk-margin robustness of the PID teacher, while the inner disk $D(\alpha_N^{\text{LA}}, 0)$ denotes the largest certified disk for the local-affine NN controller $K_{N,0}(z)$.

If the nominal local-affine closed-loop is internally stable, the inherited local-affine disk-margin radius is defined as:

$$\alpha_N^{\text{LA}} = \left\| G_N^{\text{DM0}} \right\|_\infty^{-1} = \left\| \frac{S_N - T_N}{2} \right\|_\infty^{-1} \tag{32}$$

The corresponding inherited local-affine disk-margin fraction is $\rho_{\text{inh}}^{\text{LA}} = \alpha_N^{\text{LA}} / \alpha_T$, which is local with respect to the NN activation pattern at the certified operating point. The geometric interpretation of this inherited margin is shown in Fig. 3. The outer disk represents the robustness margin $D(\alpha_T, 0)$ certified for the PID teacher, whereas the inner disk represents the largest disk $D(\alpha_N^{\text{LA}}, 0)$ certified for the local-affine NN-controlled loop. Therefore, the ratio $\rho_{\text{inh}}^{\text{LA}}$ measures the fraction of the teacher disk-margin radius that is preserved after replacing the PID controller by the local-affine realization of the trained NN controller.

Equivalently, the certified radius can be interpreted as the largest balanced disk radius that can be assigned to the local-affine NN-controlled loop:

$$\alpha_N^{\text{LA}} = \sup \left\{ \alpha \in [0, \alpha_T] : K_{N,0}(z) \text{ stabilizes } D(\alpha, 0) \right\} \tag{33}$$

Thus, $\rho_{\text{inh}}^{\text{LA}} \in [0, 1]$ quantifies the fraction of the PID teacher robustness level that remains certifiable after replacing the PID controller by the local-affine realization of the trained NN controller. The next section develops an equivalent plant-side IQC-LMI condition to certify α_N^{LA}.

3. Local-affine certification framework

The inherited local-affine disk-margin objective introduced in Section 2.4 can be certified through a plant-side IQC-LMI condition. Since the trained NN controller is replaced, for certification purposes, by its local-affine realization $K_{N,0}(z)$, the activation nonlinearities are not retained as uncertainty blocks in the final robust stability test. Instead, the only uncertainty channel considered in the certificate is the plant-side disk-margin perturbation associated with the candidate inherited radius. Thus, certification reduces to verifying whether the balanced disk-margin channel $G_N^{\text{DM0}}(z)$ in (31) satisfies:

$$\left\| G_N^{\text{DM0}} \right\|_\infty < \frac{1}{\alpha_N^\star} \tag{34}$$

for a candidate radius $\alpha_N^\star \in (0, \alpha_T]$. If this condition holds, then the local-affine NN-controlled loop robustly stabilizes all admissible norm-bounded perturbations associated with the balanced disk-margin channel of radius $\alpha_N^\star$. The corresponding plant-side robust feedback interconnection is shown in Fig. 4. In this representation, the LTI block $G_N^{\text{DM0}}(z)$ takes the perturbation output $w_{p,k}$ as input and produces the perturbation input $v_{p,k}$, while the uncertainty block Δ_p closes the feedback loop according to:

$$w_{p,k} = \Delta_p(v_{p,k}), \qquad \|\Delta_p\|_{2\to 2} \le \alpha_N^\star. \tag{35}$$

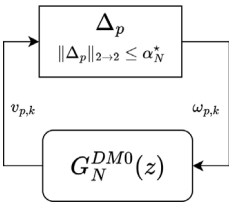

Fig. 4. Plant-side robust representation used for local-affine inherited disk-margin certification. The LTI block $G_N^{\mathrm{DM0}}(z)$ represents the balanced disk-margin robustness channel of the local-affine NN-controlled loop, while Δ_p denotes the plant-side norm-bounded perturbation satisfying $\|\Delta_p\|_{2\to 2} \le \alpha_N^\star$.

Here, $\|\Delta_p\|_{2\to 2}$ denotes the induced ℓ_2-gain of the uncertainty operator, meaning that Δ_p cannot amplify the input signal energy by more than the factor $\alpha_N^\star$. This notation is used because Δ_p is treated as a norm-bounded uncertainty operator, whereas $G_N^{\mathrm{DM0}}(z)$ is a stable LTI system whose induced ℓ_2-gain is given by its H_∞ norm. Thus, the certification problem is recast as a plant-side robust stability test for the feedback interconnection between $G_N^{\mathrm{DM0}}(z)$ and the norm-bounded operator Δ_p.

Let a minimal or numerically well-conditioned state-space realization of $G_N^{\mathrm{DM0}}(z)$ be given by:

$$G_N^{\mathrm{DM0}} : \begin{cases} x_{g,k+1} = A_g x_{g,k} + B_g w_k \\ v_k = C_g x_{g,k} + D_g w_k \end{cases} \tag{36}$$

where w_k and v_k are the input and output of the disk-margin robustness channel, respectively. In this form, the disk-radius uncertainty is represented as a norm-bounded operator, leading to the hard IQC:

$$\sum_{k=0}^{N} \begin{bmatrix} v_k \\ w_k \end{bmatrix}^\top \Pi_p(\alpha_N^\star) \begin{bmatrix} v_k \\ w_k \end{bmatrix} \ge 0, \qquad \forall N \ge 0 \tag{37}$$

with

$$\Pi_p(\alpha_N^\star) = \begin{bmatrix} (\alpha_N^\star)^2 I & 0 \\ 0 & -I \end{bmatrix} \tag{38}$$

This IQC states that the uncertainty output energy cannot exceed $(\alpha_N^\star)^2$ times the uncertainty input energy. Hence, the local-affine disk-margin certificate reduces to a plant-side small-gain/IQC condition applied to $G_N^{\mathrm{DM0}}(z)$. The next section derives a bounded-real LMI from (A_g, B_g, C_g, D_g) to certify whether $\|G_N^{\mathrm{DM0}}\|_\infty < 1/\alpha_N^\star$, thereby testing the feasibility of $\alpha_N^\star$.

3.1. Discrete-time bounded real LMI

For a fixed candidate inherited radius $\alpha_N^\star$, define:

$$\gamma^\star = \frac{1}{\alpha_N^\star} \tag{39}$$

The condition $\left\|G_N^{\mathrm{DM0}}\right\|_\infty < \gamma^\star$ can be certified using the discrete-time bounded-real lemma [26]. For the realization (36), this amounts to finding a symmetric positive-definite matrix $X = X^\top \succ 0$ such that:

$$\begin{bmatrix} A_g^\top X A_g - X & A_g^\top X B_g & C_g^\top \\ B_g^\top X A_g & B_g^\top X B_g - (\gamma^\star)^2 I & D_g^\top \\ C_g & D_g & -I \end{bmatrix} \prec 0 \tag{40}$$

For a fixed $\gamma^\star$, (40) is an LMI in X. Equivalently, for a fixed candidate $\alpha_N^\star$, it is an LMI feasibility test with $\gamma^\star = 1/\alpha_N^\star$.

3.2. Certification theorem

Theorem Assumptions. The following result is stated for the nominal discrete-time LTI plant model $P(z)$, which is also the model used to compute the balanced disk margin of the PID teacher. The trained NN controller is not certified globally over all possible activation regions, instead, what is provided is a local result represented by the local-affine realization $K_{N,0}(z)$ obtained by freezing the activation pattern at the selected operating regressor $\xi^\star$. This requires that no neuron lies exactly at an activation kink at $\xi^\star$, so that the frozen-slope representation is well defined. Accordingly, the local-affine NN-controlled loop is defined as $L_N(z) = P(z)K_{N,0}(z)$, and its associated disk-margin channel $G_N^{\mathrm{DM0}}(z)$ is assumed to admit a stable state-space realization (A_g, B_g, C_g, D_g). The induced-gain argument is performed from zero initial condition, $x_{g,0} = 0$, for input sequences $w \in \ell_2$. Finally, the plant-side disk-margin uncertainty is represented in the normalized uncertainty channel by a causal norm-bounded operator Δ_p, satisfying $\|\Delta_p\|_{2\to 2} \le \alpha_N^\star$, with a well-posed feedback interconnection.

These assumptions are intended to define the precise scope of the proposed post-training certificate. The use of the nominal LTI plant is consistent with the computation of the PID teacher disk margin, which serves as the reference robustness level. The frozen-activation assumption is local by construction and only requires the selected operating regressor not to lie exactly on an activation boundary. The internal stability of the local-affine loop is verified as a necessary pre-check before solving the associated LMI, while the zero-initial-condition requirement is the standard setting for induced ℓ_2-gain certification. Finally, the norm-bounded operator Δ_p corresponds to the normalized plant-side disk-margin uncertainty channel used in the small-gain argument.

Theorem 1. *Under the preceding assumptions, let $G_N^{\mathrm{DM0}}(z)$ be the disk-margin channel defined in (31). For a candidate inherited disk radius $\alpha_N^\star \in (0, \alpha_T]$, define $\gamma^\star = 1/\alpha_N^\star$. If there exists $X = X^\top \succ 0$ such that the LMI (40) holds, then:*

$$\left\|G_N^{\mathrm{DM0}}\right\|_\infty < \frac{1}{\alpha_N^\star} \tag{41}$$

Consequently, the local-affine NN-controlled loop is robustly stable for all admissible norm-bounded operators in the normalized plant-side disk-margin channel:

$$\|\Delta_p\|_{2\to 2} \le \alpha_N^\star \tag{42}$$

Therefore, $\alpha_N^\star$ is a certified inherited local-affine disk-margin radius, and the corresponding certified inherited fraction is $\rho_{\mathrm{inh}}^\star = \frac{\alpha_N^\star}{\alpha_T}$.

Proof. Consider the quadratic storage function associated with the discrete-time bounded-real lemma:

$$V(x_{g,k}) = x_{g,k}^\top X x_{g,k}, \qquad X = X^\top \succ 0 \tag{43}$$

This function plays the role of a Lyapunov function for the dissipativity inequality used to certify an upper bound on the induced ℓ_2-gain of G_N^{DM0}. From the state-space realization (36), the one-step variation of V is:

$$V(x_{g,k+1}) - V(x_{g,k}) = \begin{bmatrix} x_{g,k} \\ w_k \end{bmatrix}^\top \begin{bmatrix} A_g^\top X A_g - X & A_g^\top X B_g \\ B_g^\top X A_g & B_g^\top X B_g \end{bmatrix} \begin{bmatrix} x_{g,k} \\ w_k \end{bmatrix} \tag{44}$$

The output equation gives:

$$v_k = C_g x_{g,k} + D_g w_k \tag{45}$$

Using (40), one obtains the dissipation inequality

$$V(x_{g,k+1}) - V(x_{g,k}) + v_k^\top v_k - (\gamma^\star)^2 w_k^\top w_k < 0 \tag{46}$$

Summing (46) from $k = 0$ to N gives:

$$V(x_{g,N+1}) - V(x_{g,0}) + \sum_{k=0}^{N} v_k^\top v_k - (\gamma^\star)^2 \sum_{k=0}^{N} w_k^\top w_k < 0. \tag{47}$$

Since the induced-gain argument is performed from zero initial condition, $x_{g,0} = 0$, it follows that $V(x_{g,0}) = 0$.

$$V(x_{g,N+1}) + \sum_{k=0}^{N} v_k^\top v_k < (\gamma^\star)^2 \sum_{k=0}^{N} w_k^\top w_k \tag{48}$$

Moreover, since $X \succ 0$, one has $V(x_{g,N+1}) \geq 0$. Therefore:

$$\sum_{k=0}^{N} v_k^{\top} v_k < (\gamma^{\star})^2 \sum_{k=0}^{N} w_k^{\top} w_k \tag{49}$$

Hence, the induced ℓ_2-gain from w to v is strictly smaller than $\gamma^{\star}$, which implies $\left\| G_N^{\mathrm{DM0}} \right\|_{\infty} < \gamma^{\star} = \frac{1}{\alpha_N^{\star}}$. Now consider the feedback interconnection with the plant uncertainty $w_k = \Delta_p(v_k)$, where $\|\Delta_p\|_{2\to 2} \leq \alpha_N^{\star}$, the small-gain condition gives:

$$\left\| G_N^{\mathrm{DM0}} \right\|_{\infty} \|\Delta_p\|_{2\to 2} < \frac{1}{\alpha_N^{\star}} \alpha_N^{\star} = 1 \tag{50}$$

Therefore, the feedback interconnection is robustly stable for all admissible plant uncertainties satisfying $\|\Delta_p\|_{2\to 2} \leq \alpha_N^{\star}$. Thus, the candidate radius $\alpha_N^{\star}$ is certified, and comparison with the PID teacher radius α_T gives the inherited fraction $\rho_{\mathrm{inh}}^{\star} = \alpha_N^{\star}/\alpha_T$. □

The certified local-affine inherited radius is then defined as the largest feasible candidate radius:

$$\alpha_N^{\mathrm{LMI}} = \sup \left\{ \alpha_N^{\star} \in [0, \alpha_T] : (40) \text{ is feasible } \gamma^{\star} = \frac{1}{\alpha_N^{\star}} \right\} \tag{51}$$

The corresponding LMI inherited fraction is $\rho_{\mathrm{inh}}^{\mathrm{LMI}} = \frac{\alpha_N^{\mathrm{LMI}}}{\alpha_T}$.

Remark 1 (*Local Nature of the Certificate*)**.** The robustness guarantee in Theorem 1 applies only to the frozen local-affine realization $K_{N,0}(z)$ around $\xi^{\star}$. It should therefore be interpreted as a local inherited disk-margin certificate, not as a global robustness guarantee for the full nonlinear NN controller.

Remark 2 (*Nominal Local Stability as a Necessary Precheck*)**.** The feasibility of (40) requires $G_N^{\mathrm{DM0}}(z)$ to be internally stable. Equivalently, the nominal local-affine NN-controlled closed-loop must be locally stable around the operating point. Therefore, before solving the LMI, the spectral radius of the local-affine closed-loop realization should be checked:

$$\rho(A_{N,\mathrm{cl},0}) < 1 \tag{52}$$

This precheck does not replace the disk-margin certificate. It only verifies that the local-affine NN controller is a viable candidate for certification.

3.3. Certification workflow algorithm

For a fixed candidate radius $\alpha_N^{\star}$, the condition (40) is an LMI in the decision variable X. Since $\gamma^{\star} = 1/\alpha_N^{\star}$, the inherited radius can be computed by bisection over the interval $[0, \alpha_T]$. In practice, the direct value $\alpha_N^{\mathrm{LA}} = \left\| G_N^{\mathrm{DM0}} \right\|_{\infty}^{-1}$ can also be used as an upper bound for the bisection.

This workflow provides a post-training stability and robustness certificate for the local-affine realization of the trained NN controller. The NN is not required to be trained with the LMI constraints. Instead, the LMI is solved after training, once the operating point and the activation pattern have been fixed.

Remark 3 (*Interpretation of Certification Failure and Retraining*)**.** Infeasibility of the LMI does not imply nonlinear instability, it only indicates that the selected local-affine realization $K_{N,0}(z)$ cannot be certified for the considered operating point, nominal model, and disk-margin channel. In this case, the NN controller is rejected as a certified candidate, and the certificate can be used as an a posteriori diagnostic to guide retraining, for instance by selecting checkpoints with smaller $\rho(A_{N,\mathrm{cl},0})$, improving local imitation accuracy, penalizing excessive control variation, or enforcing equilibrium consistency through $|f_{\theta}(\xi^{\star})|^2$. These refinements do not guarantee feasibility, but can increase the likelihood of obtaining a certifiable NN controller.

Algorithm 1 Local-affine inherited DM certification.

Require: Nominal plant $P(z)$, PID teacher $K_T(z)$, trained NN controller $K_N = f_{\theta}(\cdot)$, certified operating point $\xi^{\star}$ and tolerance ϵ_{α}.
Ensure: Certified local-affine inherited radius α_N^{LMI} and inherited fraction $\rho_{\mathrm{inh}}^{\mathrm{LMI}}$.
Compute the PID teacher balanced disk-margin radius α_T.
Freeze the NN activation slopes at $\xi^{\star}$ and construct the local-affine controller $K_{N,0}(z)$.
Form the local-affine loop transfer function $L_N(z) = P(z)K_{N,0}(z)$.
Compute $S_N(z) = (1 + L_N(z))^{-1}$, $T_N(z) = L_N(z)(1 + L_N(z))^{-1}$
Construct the balanced disk-margin channel $G_N^{\mathrm{DM0}}(z)$.
Check local-affine closed-loop stability. If unstable, set $\alpha_N^{\mathrm{LMI}} = 0$ and stop.
Obtain a state-space realization (A_g, B_g, C_g, D_g) of $G_N^{\mathrm{DM0}}(z)$.
Set $\alpha_{\min} = 0$ and $\alpha_{\max} = \alpha_T$.
while $\alpha_{\max} - \alpha_{\min} > \epsilon_{\alpha}$ **do**
  Set $\alpha_N^{\star} = (\alpha_{\min} + \alpha_{\max})/2$, $\gamma^{\star} = 1/\alpha_N^{\star}$.
  Solve the bounded-real LMI (40) for $X = X^{\top} \succ 0$.
  if (40) is feasible **then**
    Set $\alpha_{\min} = \alpha_N^{\star}$
  else
    Set $\alpha_{\max} = \alpha_N^{\star}$
  end if
end while
Set $\alpha_N^{\mathrm{LMI}} = \alpha_{\min}$.
Compute $\rho_{\mathrm{inh}}^{\mathrm{LMI}} = \alpha_N^{\mathrm{LMI}}/\alpha_T$.
return α_N^{LMI}, $\rho_{\mathrm{inh}}^{\mathrm{LMI}}$.

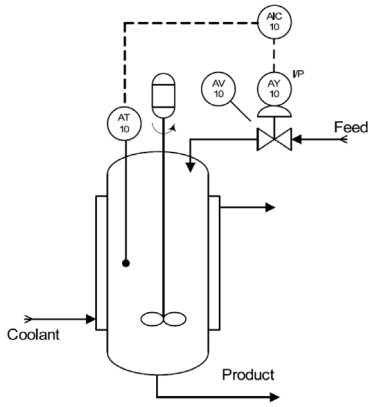

Fig. 5. Schematic representation of the Van de Vusse CSTR.

4. Implementation

The proposed post-training certification framework is evaluated on the nonlinear Van de Vusse CSTR, an inverse-response benchmark with operating-point-dependent dynamics also considered in [33]. The implementation follows three stages: PID teacher design from a reduced nominal model, imitation-based NN controller training on nonlinear CSTR trajectories, and post-training certification through local-affine realizations followed by nonlinear validation.

4.1. Nonlinear CSTR benchmark

The considered process is the Van de Vusse reactor shown in Fig. 5. The controlled variable is the concentration of species B, denoted by $C_B(t)$, and the manipulated variable is the inlet flow rate $Q(t)$. The reaction scheme is:

$$\begin{aligned} A &\xrightarrow{k_1} B \xrightarrow{k_2} C \\ 2\,A &\xrightarrow{k_3} D \end{aligned} \tag{53}$$

where B is the desired intermediate product. Assuming constant reactor volume, the concentration dynamics are:

$$\begin{aligned} \frac{dC_A(t)}{dt} &= -k_1 C_A(t) - k_3 C_A^2(t) + \frac{Q(t)}{V}\left[C_{Ai}(t) - C_A(t)\right] \\ \frac{dC_B(t)}{dt} &= k_1 C_A(t) - k_2 C_B(t) - \frac{Q(t)}{V} C_B(t) \end{aligned} \tag{54}$$

The nominal operating point and kinetic parameters are reported in [33]: $k_1 = \frac{5}{6}$ min^{-1}, $k_2 = \frac{5}{3}$ min^{-1}, $k_3 = \frac{1}{6}$ L/gmol min, with $V = 70$ L, $C_{Ai0} = 10$ gmol/L, $Q_0 = 40$ L/min and steady-state concentrations $C_{A0} = 3$ gmol/L, $C_{B0} = 1.117$ gmol/L. The nonlinear model is integrated with a sampling period of $T_s = 0.12$ s $= 0.002$ min, which is also the sampling time used in the discrete-time controller and certification implementation.

4.2. Nominal model and PID teacher

Although the neural controllers are trained and validated on the nonlinear CSTR, the PID teacher is designed from the nominal inverse-response model used in the original benchmark. This nominal model is obtained from the average representation of two local models identified around the operating point by applying positive and negative 10% perturbations to the manipulated variable $Q(t)$, following the procedure reported in [33]. The resulting reduced-order inverse-response model is:

$$P(s) = \frac{K(-bTs + 1)}{(Ts + 1)(aTs + 1)} \tag{55}$$

with $K = 0.00840$, $T = 0.557$, $a = 0.553$, $b = 0.679$. This reduced model is used only for the design and robustness-margin evaluation of the PID teacher and for the nominal plant representation in the certification step. In contrast, the training data for the neural controllers are generated by applying the PID teacher directly to the nonlinear CSTR model in (54). The teacher is implemented as a classical PID controller with derivative filtering:

$$K_T(s) = K_p\left(1 + \frac{1}{T_i s} + \frac{T_d s}{T_f s + 1}\right), \qquad T_f = \frac{T_d}{N} \tag{56}$$

The PID teacher parameters are $K_p = 104.4853$, $T_i = 0.7745$, $T_d = 0.1731$, $N = 10$, obtained by minimizing tracking error subject to a balanced disk-margin constraint with target $\alpha_T = 0.875$. The controller output is used as the deviation variable $u_{T,k} = Q_{T,k} - Q_0$, with the applied flow rate saturated in $Q \in [0.5Q_0, 1.5Q_0]$.

4.3. Teacher-data generation

The neural controllers are trained from closed-loop trajectories generated by the PID teacher on the nonlinear CSTR. The training reference consists of random ramp-hold profiles for $C_{B,r}$, constrained to remain inside the steady-state C_B range induced by $Q_0(1 \pm 10\%)$. This provides teacher data around the nominal operating point while still exciting the inverse-response dynamics. An independent validation trajectory is generated with the same procedure but using a different random seed. The training and validation trajectories have durations of 240 min and 60 min, respectively, and are sampled with $T_s = 0.002$ min, yielding $N_{\mathrm{tr}} = 120000$ and $N_{\mathrm{val}} = 30000$ samples. At each sample, the NN input regressor is defined as in (9), with $e_k = C_{B,r,k} - C_{B,k}$, window size $w_s = 10$, accumulated error $i_{e,k}$, and previous applied control deviation u_{k-1}. Hence, the NN input dimension is $n_\xi = 12$, and the target output is the PID teacher action $u_{T,k}$. The resulting teacher datasets are:

$$\mathcal{D}_{T,\mathrm{tr}} = \left\{(\xi_k, u_{T,k})\right\}_{k=1}^{N_{\mathrm{tr}}}, \qquad \mathcal{D}_{T,\mathrm{val}} = \left\{(\xi_k, u_{T,k})\right\}_{k=1}^{N_{\mathrm{val}}} \tag{57}$$

In addition, a local auxiliary dataset is generated around the nominal operating point using small-amplitude reference variations of 0.5% in C_B, together with anchor samples close to the equilibrium regressor $\xi^\star$. The local trajectory contains 30000 samples and is augmented with 1000 equilibrium anchors, giving $N_{\mathrm{loc}} = 31000$ samples. This dataset is used to improve the NN behavior near the operating point considered for local certification.

4.4. Neural controller training

All learned controllers use the NN architecture and regressor defined in Section 2.3. The four variants considered in the case study differ only in the training protocol: standard batch learning (BL), smoothed batch learning ($\mathrm{BL_S}$), enhanced scheduled sampling (ESS), and enhanced scheduled sampling with fine tuning ($\mathrm{ESS_{FT}}$). In all cases, the NN predicts the manipulated-variable deviation $u_{N,k} = Q_{N,k} - Q_0$, and the input regressor is standardized using the mean and standard deviation of the teacher-training data. The complete training and validation workflow is repeated over six independent training seeds, and seed-dependent results are reported as mean ± standard deviation.[1]

The BL controller is trained from the PID teacher dataset by minimizing the one-step imitation loss:

$$J_{\mathrm{BL}} = \frac{1}{N_{\mathrm{tr}}}\sum_{k=1}^{N_{\mathrm{tr}}}\left(f_\theta(\xi_k) - u_{T,k}\right)^2 \tag{58}$$

This objective evaluates the NN only on teacher generated regressors. $\mathrm{BL_S}$ uses the same supervised imitation principle, but adds a temporal smoothness penalty on consecutive NN control predictions. This variant is included to assess whether regularizing the learned control action improves autonomous closed-loop behavior without changing the imitation setting. The ESS controller is trained using a closed-loop scheduled-sampling procedure presented by [24]. During training, the nonlinear plant is simulated with the mixed applied action:

$$u_k^{\mathrm{app}} = p_{\mathrm{NN}} u_{N,k} + \left(1 - p_{\mathrm{NN}}\right) u_{T,k} \tag{59}$$

where p_{NN} is gradually increased. Hence, early training is close to teacher-forced imitation, while later stages expose the NN to states induced by its own previous actions. $\mathrm{ESS_{FT}}$ NN controller is obtained by fine tuning the selected ESS model at a high value of p_{NN}, so that the network is further adapted to near-autonomous operation. Small control perturbations are also injected during this phase to improve the local robustness of the learned policy. For all variants, the training objective also includes a local regularization term:

$$J_{\mathrm{loc}} = \frac{1}{N_{\mathrm{loc}}}\sum_{k=1}^{N_{\mathrm{loc}}}\left(f_\theta(\xi_k^{\mathrm{loc}}) - u_k^{T,\mathrm{loc}}\right)^2 + \lambda_{\mathrm{eq}}\left|f_\theta(\xi^\star)\right|^2 \tag{60}$$

The first term promotes teacher-consistent behavior around the nominal operating point, while the equilibrium penalty enforces a small NN output at the certified regressor $\xi^\star$, where the manipulated-variable deviation should be zero. Model selection combines validation imitation accuracy with a local stability precheck based on the spectral radius $\varrho(A_{N,\mathrm{cl},0})$ of the frozen-slope local-affine closed loop. A checkpoint is eligible only if $\varrho(A_{N,\mathrm{cl},0}) < 1$. This condition is used only as a necessary local-stability precheck during training, then the inherited disk-margin certificate is computed using the procedure described in the next subsection.

4.5. LMI-based disk-margin certification

After training, each NN controller is converted into a local-affine realization around the nominal operating point. The certified regressor is written as $\xi^\star = \begin{bmatrix}0 & \cdots & 0 & i_e^\star & 0\end{bmatrix}^\top$, where $i_e^\star$ is selected so that $f_\theta(\xi^\star) \approx 0$. The LeakyReLU slopes are then frozen at their values at $\xi^\star$, yielding the local-affine NN controller $K_{N,0}(z)$, which is interconnected with the discretized nominal plant in (55). The inherited disk-margin certificate is computed for this frozen-slope realization, rather than by treating the activation nonlinearities as uncertainty blocks. The analysis quantifies the fraction of the PID teacher margin α_T preserved by the local-affine NN-controlled loop, using both a direct disk-margin computation α_N^{LA} and a bounded-real LMI certificate α_N^{LMI} on G_N^{DM0}.

[1] Training routines for the four NN controller variants are available at https://github.com/sebastian-mb-UAB/LocalInheritedDMCertification4NNControl.

The LMI certificate is implemented in YALMIP using MATLAB LMI Lab. The margin α_N^{LMI} is obtained by bisection with tolerance 10^{-6} and at most 100 iterations, using a 0.999 safety factor on the upper endpoint to avoid testing exactly at the boundary. Feasibility is imposed with $X \succeq 10^{-15} I$, and each solution is checked a posteriori from the symmetry corrected eigenvalues of X and of the bounded-real LMI matrix with tolerance 10^{-6}. Before solving, the disk-margin channel is minimally realized and scaled or balanced when possible, the LMI is tested only after the local-affine closed-loop satisfies the Schur precheck $\varrho(A_{N,\mathrm{cl},0}) < 1$. Thus, the near coincidence between α_N^{LA} and α_N^{LMI} serves as a numerical consistency check, while the LMI formulation provides the methodological basis for extending the certificate to richer uncertainty descriptions and more complex NN controller structures.

4.6. Nonlinear validation workflow

Finally, the trained NN controllers are tested autonomously on the nonlinear CSTR under two validation scenarios:

- In Case 1, the set-point is changed to $C_{B,r} = 1.2$ at $t = 3.0$ min, followed by a negative flow-rate disturbance $\Delta Q = -0.10 Q_0$ at $t = 12.5$ min, and a positive inlet-concentration disturbance $\Delta C_{Ai} = +0.05 C_{Ai0}$ at $t = 20.0$ min.
- In Case 2, the set-point is changed to $C_{B,r} = 1.0$, followed by $\Delta Q = +0.10 Q_0$ and $\Delta C_{Ai} = -0.05 C_{Ai0}$ at the same times.

For a generic interval $\mathcal{I}$, the tracking and control-variation indices are computed as:

$$J_e(\mathcal{I}) = \int_{\mathcal{I}} |e(t)|\, dt, \qquad \mathrm{TV}_u(\mathcal{I}) = \sum_{k \in \mathcal{I}} |Q_k - Q_{k-1}| \tag{61}$$

Additional closed-loop indicators are also reported: the saturation fraction S_F, the peak tracking error $\max_k |e_k|$, the maximum manipulated-variable excursion $\max_k |Q_k - Q_0|$, and the maximum normalized operating-region displacement $\max_k d_{\mathrm{op}}(k)$. These metrics allow the nonlinear validation to assess whether the certified local behavior is consistent with the observed closed-loop performance under representative servo and regulatory conditions.

5. Results and discussion

5.1. Autonomous closed-loop validation

Autonomous closed-loop validation is used as a deployment-oriented screening step for the trained NN controllers. Unlike a one-step imitation error evaluated on teacher-generated data, this test connects each NN controller directly to the plant, so that its own previous actions determine the subsequent closed-loop signals. The PID teacher is only used as a reference baseline. This setting is therefore able to reveal whether the learned control law remains reliable under its induced state distribution, which is a necessary condition before any local robustness certificate can be meaningfully interpreted. The trajectory in Fig. 6 shows a representative response for one of the training seeds, while Table 1 summarizes the corresponding multi-seed behavior. The representative trajectory highlights a clear difference between pure batch imitation and the controllers trained with additional closed-loop regularization or exposure mechanisms. The standard BL controller follows the teacher only initially, but then accumulates deviations as its own actions drive the plant away from the distribution on which it was trained. This behavior is consistent with the distribution-shift limitation of supervised imitation learning. In contrast, $\mathrm{BL_S}$, ESS, and $\mathrm{ESS_{FT}}$ preserve the qualitative closed-loop behavior of the PID teacher throughout the validation trajectory, indicating that these controllers have learned a control action that remains effective when recursively deployed.

The multi-seed metrics in Table 1 confirm that this is not an isolated rollout effect. The BL controller exhibits both larger mean error and much larger seed-to-seed variability, with IAE $= 1.035 \pm 2.01$ and $\max|e| = 0.048 \pm 0.03$. The standard deviation being comparable to, or larger than, the mean indicates that the autonomous behavior of BL is strongly initialization-dependent. By contrast, $\mathrm{BL_S}$, ESS, and $\mathrm{ESS_{FT}}$ achieve IAE values close to the PID baseline and maximum tracking errors essentially at the same level as the teacher, with very small dispersion across seeds. Among them, $\mathrm{ESS_{FT}}$ gives the lowest mean IAE and manipulated-variable total variation, while ESS provides similarly consistent tracking. These results motivate excluding BL from the inherited-margin certification stage and retaining the NN controllers with reliable autonomous closed-loop behavior.

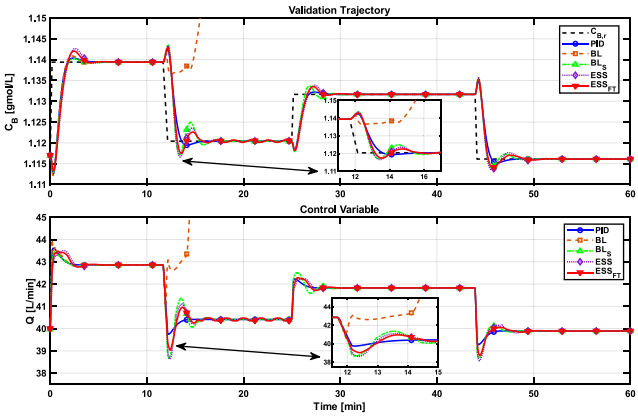

Fig. 6. Representative autonomous closed-loop validation trajectory obtained with the 6th training seed. On top the reference and output responses of the PID teacher and the NN controllers, in the lower panel the manipulated variable.

Table 1
Autonomous closed-loop metrics on the validation trajectory.

Controller	IAE	TV_u	$\max\|e\|$
PID	**0.073**	**13.518**	**0.026**
BL	1.035 ± 2.01	29.706 ± 8.87	0.048 ± 0.03
$\mathrm{BL_S}$	0.082 ± 0.003	19.568 ± 0.98	0.026 ± 0.0001
ESS	0.080 ± 0.0004	18.350 ± 0.39	$0.026 \pm 8.99 \times 10^{-5}$
$\mathrm{ESS_{FT}}$	0.078 ± 0.0005	17.964 ± 0.47	$0.026 \pm 7.30 \times 10^{-5}$

5.2. Inherited disk margin certification

After the autonomous screening step, the inherited disk-margin certificate is computed for the $\mathrm{BL_S}$, ESS, and $\mathrm{ESS_{FT}}$ based controllers. This restriction is deliberate: the certificate is intended to quantify robustness inheritance for NN controllers that have already demonstrated autonomous closed-loop viability. The standard BL controller is therefore excluded from the certification summary and kept only as a naive imitation baseline, since its closed-loop behavior is not suitable in autonomous operation.

Table 2 reports the local-affine certification results for the selected NN controllers. The spectral radius $\varrho(A_{N,\mathrm{cl},0})$ is first used as a stabilizing pre-check for the frozen-slope NN-controlled linearization. All evaluated NN controllers satisfy $\varrho(A_{N,\mathrm{cl},0}) < 1$, confirming that the corresponding local-affine closed loops are Schur stable around the certified operating regressor $\xi^\star$. This is an essential prerequisite for interpreting the disk-margin calculation, since the inherited robustness margin is meaningful only for a locally stable closed-loop realization.

The certified margins show that the NN controllers preserve a substantial portion of the PID teacher robustness level. With $\alpha_T = 0.875$, the $\mathrm{BL_S}$ based NN controller achieves the largest mean inherited fraction, with $\rho_{\mathrm{inh}}^{\mathrm{LMI}} = 77.29\% \pm 14.82\%$. The ESS and $\mathrm{ESS_{FT}}$ based NN controllers inherit $66.94\% \pm 8.76\%$ and $68.50\% \pm 12.05\%$, respectively. Thus, all controllers recover more than two thirds of the teacher disk-margin radius

Table 2
Local-affine inherited disk-margin certification for the NN controllers retained after autonomous closed-loop validation. Values are reported as mean ± standard deviation over the training seeds.

Controller	$\rho(A_{N,\mathrm{cl},0})$	α_N^{LA}	$\rho_{\mathrm{inh}}^{\mathrm{LA}}(\%)$	α_N^{LMI}	$\rho_{\mathrm{inh}}^{\mathrm{LMI}}(\%)$
$\mathrm{BL_S}$	0.9987 ± 0.00027	0.6770 ± 0.1298	77.37 ± 14.84	0.6763 ± 0.1297	77.29 ± 14.82
ESS	0.9985 ± 0.00049	0.5879 ± 0.0749	67.19 ± 8.56	0.5857 ± 0.0767	66.94 ± 8.76
$\mathrm{ESS_{FT}}$	0.9984 ± 0.00040	0.6000 ± 0.1056	68.57 ± 12.07	0.5994 ± 0.1055	68.50 ± 12.05

Reference teacher margin: $\alpha_T = 0.875$.

on average, while $\mathrm{BL_S}$ reaches the highest mean certified margin but also shows the largest dispersion. The NN controllers trained with the ESS approach provides the most repeatable inherited-margin estimate, and $\mathrm{ESS_{FT}}$ offers a slightly higher mean margin than ESS with moderate additional variability.

A key observation is the close agreement between the direct local-affine disk-margin calculation and the LMI-certified value. Across all NN controllers, α_N^{LA} and α_N^{LMI} are nearly identical at the reported precision. This agreement verifies that the proposed bounded-real LMI formulation certifies the same robustness level obtained from the direct disk-margin computation, while expressing the result in a form suitable for post-training robustness certification. This is important beyond the present numerical comparison, because the LMI formulation provides a systematic certificate that can be integrated into more general robustness-analysis pipelines, including future extensions with additional uncertainty descriptions, operating regions, or controller-structure constraints. Therefore, the LMI result is not only a consistency check, but also the certification mechanism that makes the proposed inherited-margin analysis extensible. In this sense, Table 2 provides the main quantitative evidence of the proposed certification step: NN controllers trained by imitation-learning can be locally linearized through their frozen activation pattern and assigned a certified inherited disk-margin fraction relative to the PID teacher.

The scope of this results is the certified operating neighborhood associated with $\xi^\star$ and the frozen-slope realization $K_{N,0}(z)$. Therefore, the margin should be interpreted as a local robustness-inheritance certificate, not as a standalone performance ranking over the full nonlinear trajectory. This distinction is relevant for the later nonlinear validation, the certification quantifies how much PID robustness is inherited by the local NN realization, while the subsequent simulations evaluate how the corresponding trained controllers behave on the nonlinear CSTR benchmark.

5.3. Nonlinear CSTR validation

The nonlinear validation step evaluates the PID teacher and the NN controllers on the original Van de Vusse CSTR under the two servo-regulatory scenarios defined in Section 4. These simulations are independent from the local-affine certificate and are used to assess whether the controllers evaluated after autonomous validation remain practically viable when deployed on the nonlinear inverse-response plant. The trajectories in Figs. 7 and 8 provide representative responses for one training seed, whereas Tables 3 and 4 summarize the corresponding multi-seed performance.

Case 1 represents the milder nonlinear scenario. As shown in Fig. 7, the NN controllers reproduce the qualitative behavior of the PID teacher and remain well behaved throughout the set-point and disturbance changes. The quantitative results in Table 3 show that their tracking indices remain of the same order as the PID baseline. Among the NN controllers, the ones based on ESS and $\mathrm{ESS_{FT}}$ provide the lowest integrated tracking errors, with $\mathrm{ESS_{FT}}$ giving the best values for J_{er}, J_{eQ}, and J_{eCAi}. $\mathrm{BL_S}$ is also viable and has relatively low manipulated-variable variation, especially in the regulatory intervals, but its integrated tracking errors are slightly larger than those of the ESS-based controllers.

The additional closed-loop indicators give a more complete view of the nonlinear response. In Case 1, the NN controllers exhibit zero saturation indicator (S_F), while the standard BL controller shows a nonzero

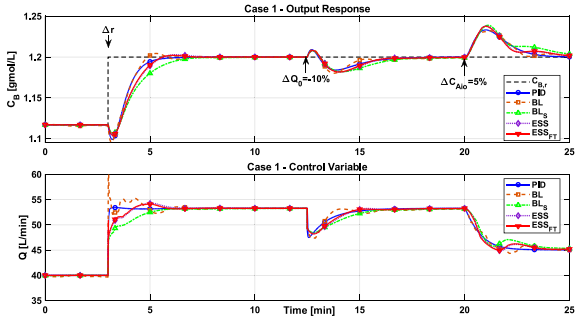

Fig. 7. Representative nonlinear CSTR validation for Case 1 obtained with the 6th training seed. On top the output response and the lower panel shows the manipulated variable.

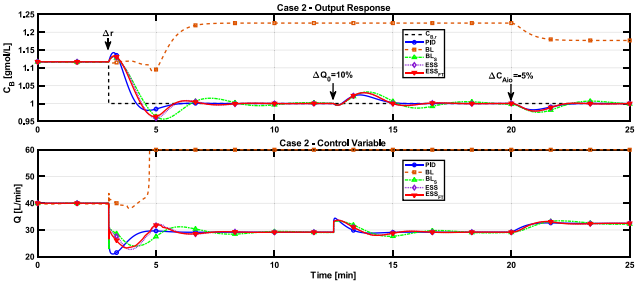

Fig. 8. Representative nonlinear CSTR validation for Case 2 obtained with the 6th training seed. On top the output response and the lower panel shows the manipulated variable.

Table 3
Nonlinear CSTR performance metrics closed-loop indicators for Case 1. Values are reported as mean ± standard deviation over six training seeds.

Metric	PID	K_N^{BL}	$K_N^{\mathrm{BL_S}}$	K_N^{ESS}	$K_N^{\mathrm{ESS_{FT}}}$
J_{er}	**0.098**	0.127 ± 0.027	0.130 ± 0.007	0.116 ± 0.002	0.114 ± 0.002
TV_{ur}	**5.12**	35.2 ± 8.3	17.7 ± 2.5	24.7 ± 4.6	25.6 ± 5.1
J_{eQ}	**0.036**	0.050 ± 0.012	0.048 ± 0.002	0.042 ± 0.001	0.041 ± 0.002
TV_{uQ}	**7.51**	9.7 ± 6.1	5.03 ± 0.18	5.45 ± 0.22	5.53 ± 0.21
J_{eCAi}	**0.061**	0.084 ± 0.033	0.085 ± 0.004	0.072 ± 0.002	0.070 ± 0.002
TV_{uCAi}	**8.21**	18.4 ± 14.5	9.3 ± 2.6	10.5 ± 1.8	10.7 ± 1.8
S_F $[10^{-3}]$	**0**	2.9 ± 5.1	0	0	0
$\max\lvert e\rvert$	**0.101**	0.095 ± 0.009	0.095 ± 0.001	0.095 ± 0.0003	0.096 ± 0.0003
$\max\lvert Q-Q_0\rvert$	**13.42**	19.1 ± 2.0	13.31 ± 0.06	14.1 ± 1.1	14.4 ± 1.6
$\max_k d_{\mathrm{op}}(k)$	**0.202**	0.217 ± 0.009	0.216 ± 0.003	0.213 ± 0.002	0.213 ± 0.002

values and the largest manipulated-variable excursion, $\max|Q-Q_0| = 19.1 \pm 2.0$. This agrees with the representative trajectory, where BL requires a more aggressive control action. The peak tracking errors are close across all NN controllers and comparable to the PID teacher, but the operating-region excursion $\max_k d_{\mathrm{op}}(k)$ is slightly larger for the BL and $\mathrm{BL_S}$ based controllers than for the ones trained with ESS and $\mathrm{ESS_{FT}}$. Thus, Case 1 shows that the NN controllers can reproduce the PID nonlinear behavior, while already revealing that pure batch imitation leads to a less regular control action.

Case 2 is more discriminating, Fig. 8 shows that the standard BL controller loses practical closed-loop reliability: the manipulated variable reaches saturation and the output departs substantially from the

Table 4
Nonlinear CSTR performance metrics closed-loop indicators for Case 2. Values are reported as mean ± standard deviation over six training seeds.

Metric	PID	K_N^{BL}	$K_N^{\mathrm{BL_S}}$	K_N^{ESS}	$K_N^{\mathrm{ESS_{FT}}}$
J_{er}	**0.119**	0.845 ± 0.937	0.229 ± 0.107	0.155 ± 0.009	0.153 ± 0.009
TV_{ur}	**15.94**	55.0 ± 38.4	41.3 ± 21.2	38.5 ± 4.9	37.4 ± 3.1
J_{eQ}	**0.033**	0.634 ± 0.787	0.106 ± 0.106	0.058 ± 0.011	0.058 ± 0.010
TV_{uQ}	**7.21**	14.9 ± 15.0	12.4 ± 11.4	9.38 ± 1.14	9.35 ± 1.20
J_{eCAi}	**0.025**	0.401 ± 0.476	0.070 ± 0.064	0.032 ± 0.009	0.032 ± 0.008
TV_{uCAi}	**3.67**	4.4 ± 4.0	8.4 ± 8.1	5.57 ± 0.38	5.39 ± 0.38
S_F [10^{-3}]	**0**	318.6 ± 439.2	0	0.10 ± 0.14	0.14 ± 0.20
$\max \lvert e \rvert$	**0.143**	0.186 ± 0.074	0.131 ± 0.002	0.134 ± 0.0001	0.134 ± 0.0001
$\max \lvert Q - Q_0 \rvert$	**19.00**	19.7 ± 0.7	16.6 ± 1.4	18.0 ± 1.8	18.2 ± 1.8
$\max_k d_{\mathrm{op}}(k)$	**0.233**	0.284 ± 0.100	0.266 ± 0.036	0.249 ± 0.007	0.247 ± 0.008

PID response. This behavior is reflected in Table 4, where BL presents the largest integrated errors and very large seed-to-seed dispersion, for example $J_{er} = 0.845 \pm 0.937$ and $J_{eQ} = 0.634 \pm 0.787$. The saturation indicator also increases sharply, $S_F = 318.6 \pm 439.2$ in units of 10^{-3}, confirming that the poor tracking behavior is associated with loss of effective control in the nonlinear simulations.

Among the NN controllers, $\mathrm{BL_S}$ remains viable in Case 2 and avoids saturation, while also achieving the smallest NN peak tracking error, $\max|e| = 0.131 \pm 0.002$. However, its integrated error indices exhibit larger seed-to-seed dispersion than those obtained with the scheduled-sampling-based controllers. The ESS-based designs show a more uniform response in this more demanding scenario. Their integrated errors are close to each other, their standard deviations remain low, and their saturation indicators are practically negligible. In particular, $\mathrm{ESS_{FT}}$ gives the lowest values for J_{er} and J_{eCAi}, whereas ESS achieves nearly identical tracking with comparable operating-region excursions. These results suggest that exposing the NN controller to its own closed-loop decisions during training improves nonlinear consistency when the validation trajectory moves farther from the nominal operating condition.

Overall, the nonlinear validation shows that the evaluated NN controllers are not merely accurate one-step imitators of the PID teacher, they remain deployable on the nonlinear CSTR under representative servo-regulatory operation. The PID teacher still provides the reference behavior and is not outperformed in every metric. However, $\mathrm{BL_S}$, ESS, and $\mathrm{ESS_{FT}}$ based controllers preserve meaningful closed-loop behavior, whereas the standard BL baseline becomes unreliable in the more demanding case. This supports the workflow by combining autonomous validation, local robustness certification, and nonlinear validation of the retained NN controllers.

5.3.1. Activation-region consistency

Since the proposed certificate is local to the activation pattern at $\xi^\star$, Table 5 reports activation-region consistency along the nonlinear CSTR validation trajectories. The Mean Neuron Agreement (MNA) measures how much of the certified activation pattern is preserved, the Near-Kink Neuron-Sample (NKNS) metric quantifies how often pre-activations approach the LeakyReLU switching boundary, and $\max_k \|\xi_k - \xi^\star\|$ reports the largest regressor displacement from the certified operating point.

The MNA values, around 60%–66%, indicate that a substantial fraction of the hidden-neuron activation pattern associated with the certified regressor $\xi^\star$ is preserved along the nonlinear trajectories. Thus, the NN controllers retain a partial local-affine consistency with the certified operating region, even when the closed-loop is evaluated under set-point and disturbance changes. This suggests that the certified local model remains dynamically relevant along the tested trajectories, but only as a local approximation around $\xi^\star$, not as an exact description of the full NN behavior over the entire validation path.

However, NKNS is extremely small in all cases, indicating that the trajectories rarely pass close to activation boundaries; hence, the activation-pattern changes are not dominated by numerical ambiguity near the kink. The regressor-distance metric further shows that Case 2 induces larger excursions from $\xi^\star$, particularly for ESS and $\mathrm{ESS_{FT}}$, which is consistent with the more demanding nonlinear validation scenario. Overall, these results support the intended interpretation of the certificate as a local inherited-margin guarantee around $\xi^\star$, complemented by nonlinear closed-loop validation.

5.4. Parametric-uncertainty validation

To assess robustness against plant mismatch, a Monte Carlo parametric-uncertainty validation was performed around the nominal CSTR operating point. The PID teacher and the trained NN controllers were kept fixed, and $N_{\mathrm{MC}} = 100$ nonlinear plants were generated by independently perturbing k_1, k_2, k_3, V, and C_{Ai0} within a uniform ±10% box, while keeping Q_0 fixed as the nominal manipulated-variable bias. For each sampled plant, the two nonlinear validation cases; PID teacher and NN controllers using the 6th trained seed; were repeated. Table 6 summarizes the sensitivity of each controller to parametric variations by comparing perturbed trajectories with the nominal nonlinear trajectory of the same controller and case. The metrics $\Delta C_{B,\max}^{95}$ and $\Delta C_{B,\max}^{\mathrm{wc}}$ report the 95th-percentile and worst-case output deviations, while $\Delta Q_{\max}^{95}$ gives the 95th-percentile manipulated-variable deviation. The percentage probabilities P_{sat} and P_{NR} denote, respectively, the percentage of samples with actuator saturation and the percentage of samples that fail to recover within the prescribed tolerance after a set-point or disturbance change. Deviation and saturation metrics are normalized with respect to the PID baseline: values below, equal to, and above one indicate lower, equal, and higher sensitivity than the PID, respectively. The non-recovery probability P_{NR} is reported in absolute percentage values because normalization is not meaningful in Case 2, where $P_{\mathrm{NR}}^{\mathrm{PID}} = 0\%$.

The results show that the naive controller K_N^{BL} is considerably more sensitive to parametric uncertainty than the PID teacher and the retained NN controllers. It increases output deviations and saturation in Case 1, and shows severe degradation in the second case, with more than twice the PID output deviation and $P_{\mathrm{NR}} = 96\%$. In contrast, $K_N^{\mathrm{BL_S}}$, K_N^{ESS}, and $K_N^{\mathrm{ESS_{FT}}}$ remain closer to the PID baseline, with near PID sensitivity in Case 1 and moderate increases in Case 2. Moreover, K_N^{ESS} and $K_N^{\mathrm{ESS_{FT}}}$ achieve $P_{\mathrm{NR}} = 0\%$ in Case 2, matching the PID baseline. These results support the workflow as a screening and certification procedure for selected NN controllers that preserve closed-loop behavior and PID inherited robustness trends under plant-parameter variations.

5.5. Certification and nonlinear behavior analysis

The preceding results show that the proposed workflow provides a consistent sequence for screening, certifying, and validating NN controllers. The autonomous closed-loop validation first distinguishes controllers that remain reliable when deployed under their own actions from controllers that only reproduce the teacher in a supervised setting. Under this criterion, $K_N^{\mathrm{BL_S}}$, K_N^{ESS}, and $K_N^{\mathrm{ESS_{FT}}}$ are retained for certification, whereas K_N^{BL} is kept only as a naive imitation baseline due to its larger tracking errors and seed-to-seed variability. For the NN controllers, the local-affine certification results show Schur-stable frozen-slope closed loops and substantial inherited disk-margin fractions with respect to the PID teacher. The LMI-certified inherited margins are 77.29% ± 14.82% for $K_N^{\mathrm{BL_S}}$, 66.94% ± 8.76% for K_N^{ESS}, and 68.50% ± 12.05% for $K_N^{\mathrm{ESS_{FT}}}$. The close agreement between the direct local-affine margins and the LMI-certified margins confirms the numerical consistency of the bounded-real certification procedure. Thus, the certificate provides a quantitative measure of local robustness inheritance around the certified operating regressor $\xi^\star$.

The nonlinear CSTR results complement this certificate by evaluating the NN controllers in the original nonlinear plant. In Case 1, $K_N^{\mathrm{BL_S}}$, K_N^{ESS}, and $K_N^{\mathrm{ESS_{FT}}}$ remain close to the PID teacher in tracking performance, with the main differences appearing in the manipulated-variable variation. In Case 2, the difference between local inherited

Table 5
Activation-region consistency metrics of the NN controllers during nonlinear CSTR validation.

Metric	Case 1			Case 2		
	$K_N^{\mathrm{BL_S}}$	K_N^{ESS}	$K_N^{\mathrm{ESS_{FT}}}$	$K_N^{\mathrm{BL_S}}$	K_N^{ESS}	$K_N^{\mathrm{ESS_{FT}}}$
MNA [%]	64.71 ± 6.28	60.33 ± 6.28	60.03 ± 5.77	66.05 ± 4.58	63.14 ± 4.91	63.19 ± 4.47
NKNS [10^{-4}%]	0.00 ± 0.00	1.14 ± 0.64	2.71 ± 2.59	0.57 ± 0.93	0.29 ± 0.39	2.43 ± 2.84
$\max_k \|\xi_k - \xi^\star\|$	17.29 ± 0.08	17.30 ± 0.05	17.35 ± 0.08	17.18 ± 1.92	19.64 ± 4.08	21.11 ± 5.75

Table 6
Parametric-uncertainty validation inside a ±10% plant-parameter box around the nominal CSTR operating point.

Metric	Case 1					Case 2				
	PID	K_N^{BL}	$K_N^{\mathrm{BL_S}}$	K_N^{ESS}	$K_N^{\mathrm{ESS_{FT}}}$	PID	K_N^{BL}	$K_N^{\mathrm{BL_S}}$	K_N^{ESS}	$K_N^{\mathrm{ESS_{FT}}}$
Normalized metrics with respect to the PID baseline										
$\Delta C_{B,\max}^{95}$	**0.1112**	1.53	1.02	1.02	1.02	**0.0952**	2.42	1.11	1.13	1.14
$\Delta C_{B,\max}^{\mathrm{wc}}$	**0.1456**	1.29	0.99	0.99	0.99	**0.1049**	2.54	1.13	1.12	1.12
$\Delta Q_{\max}^{95}$	**20.003**	1.29	1.02	1.09	1.08	**23.421**	0.96	1.06	0.94	0.96
P_{sat}	**39.0**	2.15	1.00	1.00	1.00	**52.0**	1.92	0.65	0.88	0.69
Absolute metric										
P_{NR} [%]	**38.0**	63.0	49.0	38.0	38.0	**0.0**	96.0	5.0	0.0	0.0

margin and nonlinear consistency becomes clearer. Although $K_N^{\mathrm{BL_S}}$ has the largest mean inherited margin, its nonlinear response is more dispersed than that of the ESS-based controllers. Thus, while it remains Schur stable and certified locally, it is less consistent over the more demanding nonlinear trajectory. By contrast, K_N^{ESS} and $K_N^{\mathrm{ESS_{FT}}}$ provide the most uniform nonlinear behavior among the NN controllers. This distinction is important because the inherited disk-margin certificate measures local robustness preservation near $\xi^\star$, whereas the nonlinear simulations assess finite-amplitude closed-loop behavior. Together, they show that the workflow rejects the fragile K_N^{BL} baseline, certifies meaningful local robustness inheritance for the NN controllers, and identifies K_N^{ESS} and $K_N^{\mathrm{ESS_{FT}}}$ as the most consistent nonlinear replacements.

Taken together, these results clarify the practical role of the proposed framework. The objective is not to show that the NN controllers outperform the PID teacher in every metric, but to determine whether a NN controller trained by imitation-learning can replace the teacher while preserving reliable autonomous behavior and a quantifiable fraction of its robustness margin. The evidence suggests that reliable PID-to-NN replacement is possible when training promotes closed-loop consistency. In this setting, the certificate acts as a post-training robustness filter, complemented by nonlinear and uncertainty validations.

6. Conclusions and future work

This paper proposed a post-training certification framework for assessing the replacement of a PID teacher by an imitation-trained NN controller. The framework combines autonomous closed-loop deployment, local-affine inherited disk-margin certification, activation-region analysis, nonlinear validation, and parametric-uncertainty assessment to distinguish supervised imitation accuracy from deployable closed-loop behavior. The Van de Vusse CSTR study shows that naive imitation can lead to fragile autonomous operation, whereas viable NN controllers can preserve a meaningful fraction of the PID teacher robustness margin and maintain practical nonlinear closed-loop behavior. The results also highlight the complementary roles of the analyses: the certificate provides a local robustness guarantee, the activation-region metrics assess the relevance of the local representation, and the nonlinear and uncertainty tests evaluate practical behavior under operation and plant mismatch. Future work will extend the certificate toward regional and less conservative guarantees, develop training objectives that promote certification, and consider richer uncertainty descriptions, more complex NN controller structures, and fault scenarios.

CRediT authorship contribution statement

Sebastián Madrigal: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Orlando Arrieta:** Writing – review & editing, Validation, Supervision, Formal analysis. **Ramon Vilanova:** Writing – review & editing, Validation, Supervision, Resources, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work has received support from the Spanish Government under MICINN projects PID2024-156522OB-C33 and PCI2023-145977-2 co-funded by European Union.

Data availability

Data will be made available on request.

References

[1] I. Pisa, A. Morell, J.L. Vicario, R. Vilanova, Transfer learning in wastewater treatment plants control: Measuring the transfer suitability, J. Process Control 124 (2023) 36–53, http://dx.doi.org/10.1016/j.jprocont.2023.02.006.

[2] P. Comas, A. Morell, R. Vilanova, J.L. Vicario, Predicting transfer learning suitability in ANN-based control of FOPDT industrial processes, J. Process Control 152 (2025) 103476, http://dx.doi.org/10.1016/j.jprocont.2025.103476.

[3] I. Pisa, A. Morell, J.L. Vicario, R. Vilanova, Transfer learning approach for the design of basic control loops in wastewater treatment plants, in: 2021 26th IEEE International Conference on Emerging Technologies and Factory Automation, ETFA, 2021, pp. 1–8, http://dx.doi.org/10.1109/ETFA45728.2021.9613360.

[4] M. Fazlyab, A. Robey, H. Hassani, M. Morari, G. Pappas, Efficient and accurate estimation of Lipschitz constants for deep neural networks, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, R. Garnett (Eds.), in: Advances in Neural Information Processing Systems, vol. 32, Curran Associates, Inc., 2019.

[5] P. Pauli, A. Koch, J. Berberich, P. Kohler, F. Allgöwer, Training robust neural networks using Lipschitz bounds, IEEE Control. Syst. Lett. 6 (2022) 121–126, http://dx.doi.org/10.1109/LCSYS.2021.3050444.

[6] H. Yin, P. Seiler, M. Arcak, Stability analysis using quadratic constraints for systems with neural network controllers, IEEE Trans. Autom. Control 67 (4) (2021) 1980–1987, http://dx.doi.org/10.1109/tac.2021.3069388.
[7] P. Pauli, J. Köhler, J. Berberich, A. Koch, F. Allgöwer, Offset-free setpoint tracking using neural network controllers, Learn. Dyn. Control. (2021) 992–1003.
[8] H. Ai, Y. Zheng, Z. Xu, L. Jiang, A. Zhu, X. Fu, Bounded barrier Lyapunov function-based neural network control of space robot with output tracking error constraints and input saturation, ISA Trans. (2026) http://dx.doi.org/10.1016/j.isatra.2026.04.012.
[9] F. Gu, H. Yin, L.E. Ghaoui, M. Arcak, P. Seiler, M. Jin, Recurrent neural network controllers synthesis with stability guarantees for partially observed systems, Proc. AAAI Conf. Artif. Intell. 36 (5) (2022) 5385–5394, http://dx.doi.org/10.1609/aaai.v36i5.20476.
[10] F. Xu, H. He, M. Song, X. Xu, Iterative neural network adaptive robust control of a maglev planar motor with uncertainty compensation ability, ISA Trans. 140 (2023) 331–341, http://dx.doi.org/10.1016/j.isatra.2023.05.010.
[11] C. De Souza, A. Girard, S. Tarbouriech, Event-triggered neural network control using quadratic constraints for perturbed systems, Automatica 157 (2023) 111237, http://dx.doi.org/10.1016/j.automatica.2023.111237.
[12] F. Fabiani, P.J. Goulart, Reliably-stabilizing piecewise-affine neural network controllers, IEEE Trans. Autom. Control 68 (9) (2022) 5201–5215, http://dx.doi.org/10.1109/tac.2022.3216978.
[13] A. La Bella, M. Farina, W. D'Amico, L. Zaccarian, Regional stability conditions for recurrent neural network-based control systems, Automatica 174 (2025) 112127, http://dx.doi.org/10.1016/j.automatica.2025.112127.
[14] Y. Zhang, X. Xu, Robust stability of neural network control systems with interval matrix uncertainties, Automatica 177 (2025) 112289, http://dx.doi.org/10.1016/j.automatica.2025.112289.
[15] G. Norris, G.J.J. Ducard, C. Onder, Neural networks for control: A tutorial and survey of stability-analysis methods, properties, and discussions, in: 2021 International Conference on Electrical, Computer, Communications and Mechatronics Engineering, ICECCME, 2021, pp. 1–6, http://dx.doi.org/10.1109/iceccme52200.2021.9590912.
[16] C. Liu, L. Wang, Z. Chu, H. Su, Neural network-based adaptive decentralized safe control for interconnected nonlinear systems with time delays, ISA Trans. 166 (2025) 168–178, http://dx.doi.org/10.1016/j.isatra.2025.07.016.
[17] I.R. Manchester, R. Wang, N.H. Barbara, Neural networks in the loop: Learning with stability and robustness guarantees, Annu. Rev. Control. Robot. Auton. Syst. 9 (1) (2026) 181–214, http://dx.doi.org/10.1146/annurev-control-022624-032459.
[18] N. Junnarkar, M. Arcak, P. Seiler, Synthesizing neural network controllers with closed-loop dissipativity guarantees, 2024, http://dx.doi.org/10.48550/arxiv.2404.07373, arXiv (Cornell University).
[19] H. Zheng, Y. Guo, Z. Wang, Z. Wang, Z. Hao, Robust tracking control of uncertain autoloaders by implicit Lyapunov method and scleronomic Lagrangian mechanics-informed neural network, ISA Trans. 167 (Pt A) (2025) 969–983, http://dx.doi.org/10.1016/j.isatra.2025.08.025.
[20] L.B. De Giuli, A. La Bella, M. Farina, R. Scattolini, Stability and learning of networked systems with application to physics-informed neural networks, Automatica 189 (2026) 113005, http://dx.doi.org/10.1016/j.automatica.2026.113005.
[21] P. Seiler, A. Packard, P. Gahinet, An introduction to disk margins [lecture notes], IEEE Control. Syst. 40 (5) (2020) 78–95, http://dx.doi.org/10.1109/mcs.2020.3005277.
[22] N. Junnarkar, M. Arcak, P. Seiler, Stability margins of neural network controllers*, in: 2025 American Control Conference, ACC, 2025, pp. 1355–1360, http://dx.doi.org/10.23919/acc63710.2025.11107746.
[23] S. Ross, G.J. Gordon, J.A. Bagnell, A reduction of imitation learning and structured prediction to no-regret online learning, 2010, http://dx.doi.org/10.48550/arxiv.1011.0686, arXiv (Cornell University).
[24] P. Comas, J.L. Vicario, A. Morell, R. Vilanova, Enhanced scheduled sampling training framework for ANN-based PID control of a continuous stirred-tank reactor, in: 2025 IEEE 30th International Conference on Emerging Technologies and Factory Automation, ETFA, 2025, pp. 1–4, http://dx.doi.org/10.1109/ETFA65518.2025.11205624.
[25] X. Ge, S. Putri, F. Moazeni, J. Khazaei, Enhanced imitation learning of robust nonlinear model predictive control via temporal convolutional neural network for DC shipboard microgrid, ISA Trans. (2026) http://dx.doi.org/10.1016/j.isatra.2026.04.001.
[26] S. Boyd, L.E. Ghaoui, E. Feron, V. Balakrishnan, Linear Matrix Inequalities in System and Control Theory, Society for Industrial and Applied Mathematics, 1994, http://dx.doi.org/10.1137/1.9781611970777.
[27] A. Bounemeur, M. Chemachema, General fuzzy adaptive fault-tolerant control based on Nussbaum-type function with additive and multiplicative sensor and state-dependent actuator faults, Fuzzy Sets and Systems 468 (2023) 108616, http://dx.doi.org/10.1016/j.fss.2023.108616.
[28] A. Bounemeur, M. Chemachema, Finite-time output-feedback fault tolerant adaptive fuzzy control framework for a class of MIMO saturated nonlinear systems, Int. J. Syst. Sci. 56 (4) (2024) 733–752, http://dx.doi.org/10.1080/00207721.2024.2409853.
[29] B. Abdelhamid, C. Mohamed, Robust fuzzy adaptive fault-tolerant control for a class of second-order nonlinear systems, Internat. J. Adapt. Control Signal Process. 39 (1) (2024) 15–30, http://dx.doi.org/10.1002/acs.3916.
[30] B. Sattelberg, R. Cavalieri, M. Kirby, C. Peterson, R. Beveridge, Locally linear attributes of ReLU neural networks, Front. Artif. Intell. 6 (2023) 1255192, http://dx.doi.org/10.3389/frai.2023.1255192.
[31] A.L. Maas, A.Y. Hannun, A.Y. Ng, Rectifier nonlinearities improve neural network acoustic models, in: ICML Workshop on Deep Learning for Audio, Speech, and Language Processing, Atlanta, GA, USA, 2013.
[32] B. Xu, N. Wang, T. Chen, M. Li, Empirical evaluation of rectified activations in convolutional network, 2015, http://dx.doi.org/10.48550/arxiv.1505.00853, arXiv (Cornell University).
[33] V.M. Alfaro, R. Vilanova, Robust tuning of 2dof five-parameter PID controllers for inverse response controlled processes, J. Process Control 23 (4) (2013) 453–462, http://dx.doi.org/10.1016/j.jprocont.2013.01.005.