Optimal Robust State Feedback Control Using Physics-Informed Neural Networks

Camilo E. Zambrano[1] and Eduardo Mojica-Nava[2]

***Abstract*— This paper presents a unified framework for optimal and robust state feedback control based on Physics-Informed Neural Networks. Traditional approaches to nonlinear optimal control, such as Hamilton–Jacobi–Bellman methods and dissipativity-based robustness analysis, suffer from the curse of dimensionality and rely on intractable analytical expressions for value functions or storage functions. To overcome these limitations, we formulate the robust optimal control problem as a physics-informed learning task in which the system dynamics, a Hamilton–Jacobi–Isaacs equation, and robustness constraints inspired by nonlinear $\mathcal{L}_2$-gain and passivity theory are embedded directly into the PINN loss structure. The resulting neural architecture simultaneously approximates the value function and its gradient, enabling the synthesis of a robust state feedback law without discretization or linearization. We provide conditions under which the proposed approach preserves key robustness properties of dissipative systems and demonstrate its effectiveness through numerical examples.**

I. INTRODUCTION

The design of optimal and robust feedback controllers for nonlinear systems remains a fundamental challenge in control engineering. Classical approaches such as Hamilton–Jacobi–Bellman (HJB) theory, dynamic programming, and nonlinear $\mathcal{L}_2$-gain and passivity methods provide rigorous stability and performance guarantees [1]–[3]. However, their practical application is often constrained by model complexity, the curse of dimensionality, and the difficulty of obtaining closed-form solutions to nonlinear partial differential equations (PDEs). As engineered systems grow more interconnected and uncertainty-rich, there is an increasing need for computational methods that exploit system physics while maintaining optimality and robustness [4], [5].

Physics-Informed Neural Networks (PINNs) have recently emerged as a powerful tool for solving forward and inverse problems governed by nonlinear differential equations [6]. By embedding the known system dynamics directly into the learning architecture, PINNs avoid the data dependence typical of black-box learning and maintain physical consistency. Their flexibility has motivated early explorations in control, including the use of neural techniques for solving HJB equations [7] and neural-network-based finite-difference least-squares solvers for Hamilton–Jacobi PDEs [5]. In this approach, PINN-based formulations are introduced that embed system dynamics and control inputs into the loss function to learn stabilizing feedback laws for nonlinear ODE systems, enabling long-horizon prediction and control without explicit discretization [8]–[10]. More recent developments address optimal and constrained control by integrating optimality conditions directly into the PINN architecture, allowing joint learning of state trajectories and optimal controls for both finite-dimensional and PDE systems [11]–[13]. Extensions to recurrent and ensemble PINN structures further improve predictive control performance and robustness in nonlinear process systems and high-dimensional dynamics [14]. Collectively, these approaches highlight the emerging role of PINNs as a flexible framework for model-based control that bridges data-driven learning and physics-consistent optimal control synthesis [15] and quantum control [16]. Yet, the integration of PINNs with robust optimal control theory remains only partially explored.

A fundamental body of work underpinning modern nonlinear robustness analysis comes from the development of the $\mathcal{L}_2$-gain and passivity framework [3]. This theory establishes powerful input–output characterizations for nonlinear systems using dissipativity concepts, enabling robustness certification through energy-based arguments rather than state–space linearization. The $\mathcal{L}_2$-gain approach provides a systematic method for bounding the worst-case amplification of disturbances, while passivity offers a complementary viewpoint ensuring stability under interconnections and model uncertainties. These tools have become central in robust nonlinear control, forming the basis of feedback design strategies that guarantee stability margins even under significant structural uncertainty. In the context of our work, the robustness constraints embedded in the PINN formulation are inspired by these dissipativity principles by integrating $\mathcal{L}_2$-gain–like terms into the learning loss, the resulting neural approximation of the value function inherits meaningful robustness guarantees.

In this work, we develop a unified framework for optimal robust state feedback control using PINNs. The key idea is to embed both the system dynamics and robustness constraints, motivated by nonlinear $\mathcal{L}_2$-gain and passivity theory, into a PINN-based approximation of the value function. By learning a solution to a robust Hamilton–Jacobi–Isaacs (HJI) equation, the PINN implicitly generates the gradient information required to synthesize an optimal feedback law without discretization or linearization. Compared to traditional PDE solvers, this representation leverages the expressiveness of deep networks while preserving the physics of the underlying control problem.

The main contributions of this paper are threefold. First, a PINN-based formulation of robust optimal control, combining physics-informed learning with robustness concepts from nonlinear $\mathcal{L}_2$-gain and passivity theory is presented.

[1,2]Both authors are with the Faculty of Electrical and Electronic Engineering, Universidad Nacional de Colombia, Bogotá, Colombia. czambrano@unal.edu.co, eamojican@unal.edu.co

PINNs are used to solve the Hamiltonian based on Lax-Friedrichs framework. Then, a neural approximation of the value function that enables direct extraction of a robust state feedback law is developed. Finally, a set of numerical experiments is presented to demonstrate the performance of the proposed method under structured uncertainty and unstable systems. In general, the results indicate that PINNs offer a promising alternative to traditional approaches for solving robust control problems, providing high approximation accuracy, scalability, and strong alignment with physics-based modeling.

The rest of the paper is organized as follows. In Section II, the problem of the nonlinear state feedback controller is introduced. In Section III, we present the main numerical Hamiltonian based on Lax-Friedrichs framework. Then, a robust optimal control based on PINNs is developed in Section IV. Section V shows two numerical experiments, namely an inverted pendulum and a Kuramoto model, to illustrate the effectiveness of the proposed approach. Finally, some conclusions and future works are drawn in Section VI.

II. State Feedback Nonlinear H_∞ Control Problem

Many real-world systems operate under persistent disturbances, unmodeled effects, and uncertainties that cannot be ignored. For such systems, achieving robust performance is not only desirable but essential. The H_∞ framework provides a principled way to guarantee that the effect of unknown disturbances on the controlled outputs remains bounded. More precisely, it ensures that the system attenuates any admissible disturbance with an induced $\mathcal{L}_2-gain$ not exceeding a prescribed level γ. This makes H_∞ control especially attractive in applications where worst-case scenarios must be explicitly considered.

While the linear theory of H_∞ control is well developed and relies on algebraic Riccati equations, many engineering systems exhibit inherently nonlinear behavior, particularly when operating far from equilibrium or under large perturbations. In such cases, linearization-based control strategies may fail to capture the true dynamics, leading to performance degradation or even instability. Nonlinear H_∞ control seeks to overcome these limitations by directly incorporating the system's nonlinear structure into the robustness design. However, extending H_∞ methods to nonlinear settings introduces significant mathematical challenges. In the following,the basic results of nonlinear H_∞ optimal control are introduced, which are numerically solved in the next section using the proposed PINNs approach.

A. Basics of Nonlinear H_∞ Optimal Control

Consider the affine system with the full state available for feedback

$$\begin{aligned} \dot{x} = f(x,u,d) = a(x) + b(x)u + g(x)d \\ z = \begin{bmatrix} h(x) & u \end{bmatrix}^T \end{aligned} \tag{1}$$

with $u \in \mathbb{R}^m$, $d \in \mathbb{R}^r$, $x \in \chi \subset \mathbb{R}^n$ where χ is an open connected neighborhood of the origin, and $z \in \mathbb{R}^s$. Assume that $x = 0$ is an equilibrium of the unforced, undisturbed system, i.e., $a(0) = 0$ and $h(0) = 0$, and that the storage function satisfies $P(0) = 0$. Assume all data to be C^k (k-times continuously differentiable) with $k \geq 2$. Let $P(x)$ be a C^1 storage function $P \geq 0$ for the closed-loop system (1) with $\mathcal{L}_2\, gain$ supply rate $\frac{1}{2}\gamma^2||d||^2 - \frac{1}{2}||z||^2$ with $\gamma > 0$. On the other hand, from the theory of differential games, one may look at the H_∞ suboptimal problem as a two-player zero-sum differential game with cost criterion $\frac{1}{2}\gamma^2||d||^2 - \frac{1}{2}||z||^2$, where one player corresponds to the control input u, and the other player to the disturbance d. In the differential game formulation, let $p \in \mathbb{R}^n$ denote the costate variable associated with the state x, interpreted as the row vector of partial derivatives of the storage function, i.e., $p = P_x(x)$. The pre-Hamiltonian corresponding to this differential game is then given as

$$K_\gamma(x,p,d,u) := p^T[a(x) + b(x)u + g(x)d] - \frac{1}{2}\gamma^2||d||^2 + \frac{1}{2}||z||^2. \tag{2}$$

Computing the partial derivatives of $\frac{\partial K_\gamma}{\partial u} = 0$ and $\frac{\partial K_\gamma}{\partial d} = 0$, it is determined that $d^*(x,p)$ (worst-case disturbance) and $u^*(x,p)$ (optimal control) are

$$\begin{aligned} u^*(x,p) &= -b^T(x)p^T(x), \\ d^*(x,p) &= \frac{1}{\gamma^2}g^T(x)p^T(x), \end{aligned} \tag{3}$$

and they correspond to a saddle point of K_γ. Substituting (3) in (2) we obtain the optimal Hamiltonian

$$\begin{aligned} H_\gamma(x,p) =& p^T[a(x) + b(x)u^*(x,p) + g(x)d^*(x,p)] \\ &- \frac{1}{2}\gamma^2||d^*(x,p)||^2 + \frac{1}{2}||z(x,p)||^2. \end{aligned}$$

Then, if there exists $P(x)$ such that it is a solution to the Hamilton–Jacobi–Isaacs inequality, and $P_x(x)$ the row vector of partial derivatives $\left(\frac{\partial P}{\partial x_1}(x), ..., \frac{\partial P}{\partial x_n}(x)\right)$ satisfies

$$H_\gamma(x, P_x^T(x)) \leq 0 \tag{4}$$

then the state feedback $u = u^*(x, P_x^T(x))$ is such that the closed loop system (1) is stable and has $\mathcal{L}_2\, gain \leq \gamma$.

However, it is difficult to analytically find the storage function $P(x)$ that solves (4). In the following sections, we are focused on how to approximate P using PINNs.

III. Numerical Hamiltonian Based On Lax-Friedrichs Scheme

The framework proposed in [4], centers on defining a residual functional based on a monotone and consistent numerical Hamiltonian constructed from forward and backward finite differences. For a discretization step $\delta > 0$, the gradient of a function $V(x)$ is approximated by the upwind and downwind finite differences:

$$D^+_{\delta,i}V(x) = \frac{V(x+\delta e_i) - V(x)}{\delta}$$
$$D^-_{\delta,i}V(x) = \frac{V(x) - V(x-\delta e_i)}{\delta}.$$

Collecting all directional differences, define the vectors

$$V^+ = (D^+_{\delta,1}V(x), \ldots, D^+_{\delta,n}V(x)) \in \mathbb{R}^n$$
$$V^- = (D^-_{\delta,1}V(x), \ldots, D^-_{\delta,n}V(x)) \in \mathbb{R}^n,$$

which serve as upwind and downwind approximations of the gradient $\nabla V(x)$.

The Lax-Friedrichs numerical Hamiltonian is then defined as

$$\widehat{H}_\alpha(x, V^+, V^-) = H\left(x, \frac{V^+ + V^-}{2}\right) - \alpha \sum_{i=1}^{n} \frac{V_i^+ - V_i^-}{2},$$

where $\alpha > 0$ introduces a numerical viscosity term ensuring stability and monotonicity. The parameter α must satisfy

$$\alpha \geq \max_{||p|| \leq L, x \in \Omega} ||\nabla_p H(x,p)||,$$

so that the operator remains monotone for any function with Lipschitz constant L. Let $\Omega \subset \mathbb{R}^n$ be a compact domain containing the origin, on which the storage function P is to be approximated. Let $\delta\mathbb{Z}^n := \{\delta k : k \in \mathbb{Z}^n\}$ denote the uniform Cartesian lattice with spacing $\delta > 0$. The discrete grid is then $\Omega_\delta := \delta\mathbb{Z}^n \cap \Omega$. Given a uniform grid $\Omega_\delta = \delta\mathbb{Z}^n \cap \Omega$, the discrete residual functional is $\mathcal{R}_{\alpha,\delta}(p) = \delta^n \sum_{x \in \Omega_\delta} \left[\widehat{H}_\alpha(x, D^+_\delta V(x), D^-_\delta V(x))\right]^2$.

A function p minimizes $\mathcal{R}_{\alpha,\delta}$ if and only if it satisfies $\widehat{H}_\alpha(x, D^+_\delta V(x), D^-_\delta V(x)) = 0 \forall x \in \Omega_\delta$, that is, p solves the discrete finite-difference equation defined by the numerical Hamiltonian. When $\widehat{H}_\alpha$ is consistent with H and monotone, every minimizer of this functional approximates the viscosity solution of the Hamiltonian-Jacobi equation $H(x, \nabla V) = 0$. This structure can be addressed with PINNs, the neural network $V_\theta(x)$ is trained by minimizing the expectation of the discrete residual. Because the Lax-Friedrichs operator contains a controlled amount of numerical diffusion through α, the optimization landscape is convexified, and then every stationary point corresponds to a discrete solution. Moreover, by the Barles-Souganidis convergence theorem proposed in [17], solutions of monotone and consistent schemes converge uniformly to the viscosity solution as $\delta \to 0^+$. Therefore, coupling a neural approximation V_θ with the Lax-Friedrichs residual yields a provably convergent PINN that inherits stability from the monotone scheme, ensures uniqueness through the viscosity formulation, and allows progressive refinement by annealing (α, δ) toward the nondiffusive limit.

To extend the framework from an equality constraint $H(x, \nabla V(x)) = 0$ to an inequality $H(x, \nabla V(x)) \leq 0$ as required in (4), the key idea is to reinterpret the PDE not as an exact residual equation, but as a variational inequality. In the inequality case, V must satisfy $H(x, \nabla V(x)) \leq 0$, which means that V is a viscosity subsolution of the Hamilton–Jacobi operator. At the continuous level, this ensures that V bounds from below the value function of the associated differential game. When H is replaced by the game Hamiltonian K_γ from (2), the corresponding Lax–Friedrichs numerical Hamiltonian is denoted $\widehat{K}_{\alpha,\delta}$, defined analogously as

$$\widehat{K}_{\alpha,\delta}(x, V^+, V^-) := K_\gamma\left(x, \frac{V^+ + V^-}{2}\right) - \alpha \sum_{i=1}^{n} \frac{V_i^+ - V_i^-}{2}.$$

With the Lax-Friedrichs scheme, the inequality constraint can be expressed directly as $\widehat{K}_{\alpha,\delta}(x, D^+_\delta V, D^-_\delta V) \leq 0$.

To enforce this condition within a PINN, we do not minimize the squared residual symmetrically, since that would penalize both positive and negative values of the discrete Hamiltonian. Instead, the loss function is made one-sided as

$$\mathcal{L}_{PDE}(\theta) = \mathbb{E}_{x \sim \omega}\left[\left(max\left\{0, \widehat{K}_{\alpha,\delta}\left(x, D^+_\delta \hat{V}, D^-_\delta \hat{V}\right)\right\}\right)^2\right].$$

This penalty acts only when the discrete Hamiltonian becomes positive. In the limit where the neural approximation V_θ achieves zero expected violation, the resulting function satisfies the discrete inequality $\widehat{H}_{\alpha,\delta} \leq 0$ everywhere, and by the monotonicity and consistency of the scheme, the limiting function V is a viscosity subsolution of $H(x, \nabla V(x)) \leq 0$.

Formally, the same Barles-Souganidis argument still applies. If the numerical Hamiltonian is monotone, stable, and consistent with the continuous operator, then any sequence of discrete subsolutions satisfying $\widehat{H}_{\alpha,\delta} \leq 0$ converges uniformly to the continuous viscosity subsolution. Thus, replacing the squared residual with a one-sided hinge loss preserves all the theoretical convergence properties of the equality case, while enforcing the correct variational inequality structure required for H_∞.

This operator inherits the monotonicity and consistency properties of the Lax–Friedrichs scheme.

IV. Robust Optimal Control Using Physics-Informed Neural Networks

To approximate the storage function P associated with the Hamilton–Jacobi–Isaacs inequality (4), a neural approximation $\hat{P}(x, \theta)$ is trained using the monotone Lax-Friedrichs residual described in the previous section, together with additional constraints ensuring the correct structural properties of P. Since the closed-loop equilibrium $x^* = 0$ must be a global minimum of the storage function and satisfy $P(0) = 0$, these features are incorporated as soft boundary conditions in the loss functional.

A convenient way to enforce them is to penalize both the value and the gradient of the network at the equilibrium point as $\mathcal{L}_{bc}(\theta) = \hat{P}(0, \theta)^2 + ||\nabla \hat{P}(0, \theta)||^2$.

The first term anchors the value of the function at the equilibrium, while the second suppresses any nonzero first-order variation at that point. Together, these conditions encourage the network to represent a nonnegative, locally minimal storage function. Non-negativity of the storage function, $P(x) \geq 0$, can be enforced by penalizing violations of this inequality and softly penalizing values near 0, obtaining $\mathcal{L}_{min}(\theta) = \mathbb{E}_{x\sim\omega}\left[\left(\textit{softplus}\left(-\hat{P}(x,\theta),\beta\right)\right)^2\right]$, where lower values of β imply a stronger penalization for values close to 0. This loss function drives the approximation toward the cone of positive-definite functions on Ω without interfering with regions where the constraint is already satisfied. These terms are combined with the Lax–Friedrichs residual loss enforcing the Hamilton–Jacobi–Isaacs inequality. The resulting objective function yields a neural approximation that respects the qualitative properties of storage functions and satisfies the viscosity subsolution condition required for robust optimal control.

We are now ready to state our main theoretical result in the following theorem.

Theorem 1: Let (1) be an affine system with inputs $u \in \mathbb{R}^m$ and disturbances $d \in \mathbb{R}^r$, with states $x \in \chi \subset \mathbb{R}^n$ and to-be-controlled outputs $z \in \mathbb{R}^s$, and $\gamma > 0$. Suppose there exists a storage function $P(x) \geq 0$ k-times continuously differentiable, with $k \geq 2$, such that there exists a solution for the Hamilton–Jacobi–Isaacs inequality (5) with optimal input and disturbance, and such that the closed-loop system has $\mathcal{L}_2$ *gain* $\leq \gamma$

$$u^*(x, P_x^T(x)) = -b^T(x)P_x^T(x)$$
$$d^*(x, P_x^T(x)) = \frac{1}{\gamma^2}g^T(x)P_x^T(x)$$

$$K_\gamma(x, P_x^T, d^*(x, P_x^T(x)), u^*(x, P_x^T(x))) \leq 0 \tag{5}$$

with $K_\gamma(x,p,d,u) := p^T[a(x)+b(x)u+g(x)d]-\frac{1}{2}\gamma^2||d||^2+\frac{1}{2}||z||^2$.

Then, the numerical Hamiltonian $\widehat{K}_{\alpha,\delta}(\cdot)$ is monotone and consistent for some $\alpha > 0$ and some sufficiently small $\delta > 0$. And it can be approximated with a neural network $\hat{P}(x,\theta)$ using stochastic gradient descent and loss functions

$$\mathcal{L}_{PDE}(\theta) = \mathbb{E}_{x\sim\omega}\left[\left(\textit{max}\left\{0, \widehat{K}_{\alpha,\delta}\left(x, D_\delta^+\hat{P}, D_\delta^-\hat{P}\right)\right\}\right)^2\right]$$
$$\mathcal{L}_{bc}(\theta) = \hat{P}(0,\theta)^2 + ||\nabla\hat{P}(0,\theta)||^2$$
$$\mathcal{L}_{min}(\theta) = \mathbb{E}_{x\sim\omega}\left[\left(\textit{softplus}\left(-\hat{P}(x,\theta),\beta\right)\right)^2\right]$$

such that it is a local solution for the numerical Hamilton–Jacobi–Isaacs inequality (5).

Proof: Before proceeding with the technical steps, we outline the structure of the proof. We first characterize the saddle–point conditions of the Hamiltonian by computing the optimal input and worst–case disturbance that satisfy the first–order stationary conditions. Substituting these expressions into the Hamilton–Jacobi–Isaacs inequality allows us to recover the classical dissipation inequality that certifies the $\mathcal{L}_2$ gain bound of the closed-loop system. We then use this inequality, together with standard storage-function arguments, to establish asymptotic stability via LaSalle's Invariance Principle. Finally, we show that the discrete residual based on the monotone Lax–Friedrichs Hamiltonian satisfies the structural assumptions of the finite-difference scheme in [4], ensuring that the numerical approximation preserves the viscosity subsolution property required for robust stability.

Computing the partial derivatives $\frac{\partial K_\gamma}{\partial u} = 0$ and $\frac{\partial K_\gamma}{\partial d} = 0$, it results in $u^*(x,p) = -b^T(x)p^T(x)$, $d^*(x,p) = \frac{1}{\gamma^2}g^T(x)p^T(x)$. Thus, the critical point has saddle point dynamics, then

$$K_\gamma(x,p,d,u^*(x,p)) \leq K_\gamma(x,p,d^*(x,p),u^*(x,p)). \tag{6}$$

Let $P \geq 0$. Substituting $p = P_x^T(x)$ in (6) we obtain

$$P_x(x)\left[f(x,u^*(x,P_x^T(x)),d)\right] - \frac{1}{2}\gamma^2||d||^2 + \frac{1}{2}||z||^2$$
$$\leq K_\gamma(x, P_x^T(x), d^*(x,P_x^T(x)), u^*(x,P_x^T(x))).$$

Thus $K_\gamma(x, P_x^T(x), d^*(x,P_x^T(x)), u^*(x,P_x^T(x))) \leq 0$ for all d, then

$$P_x(x)\left[f(x,u^*(x,P_x^T(x)),d)\right] \leq \frac{1}{2}\gamma^2||d||^2 - \frac{1}{2}||z||^2 \tag{7}$$

showing that the closed-loop system has $\mathcal{L}_2$ *gain* $\leq \gamma$ with storage function P.

Owing that P is a storage function, by proposition 3.2.16 from [3] with the equation (7) evaluated in 0, fulfills LaSalle's Invariance principle in the closed-loop system, then it is asymptotically stable in $x = 0$.

Furthermore, the residual

$$\mathcal{R}_{\alpha,\delta}(p) = \delta^n \sum_{x\in\Omega_\delta}\left[\textit{max}\left\{0, \widehat{H}_\alpha(x, D_\delta^+V(x), D_\delta^-V(x))\right\}\right]^2 \tag{8}$$

holds the structure of Theorem 1 proposed in [4]. Considering the uniform Cartesian grid $\Omega_\delta := \delta\mathbb{Z}^n \cap \Omega$ for some $\delta = 1/N$ and $N \in \mathbb{N}$, if the critical point $v \in C(\bar{\Omega})$ of the residual (8) with a Lipschitz constant L and

$$2\alpha\sin^2\left(\frac{\pi}{2}\delta\right) > \max_{||p||\leq L, x\in\bar{\Omega}} ||\nabla_p H(x,p)||,$$

then $\widehat{H}_{\alpha,\delta}(x, D_\delta^+v(x), D_\delta^-v(x)) \leq 0$.

Thus the residual is still differentiable; the proof of the theorem remains for the inequality case. Finally, due to $\widehat{H}_{\alpha,\delta}$ is consistent with H, then the schema proposed is closed-loop stable with $\mathcal{L}_2$ *gain* $\leq \gamma$. ∎

Corollary 1.1: Let $\hat{P}(x)$ be the neural network approximation obtained from Theorem 1 with training loss $\mathcal{L}_{PDE}(\theta) = \epsilon_T$. Define the Hamiltonian residual $\epsilon_H = \sup_{x\in\Omega}\max\{0, K_\gamma(x, \hat{P}_x^T(x), \hat{d}, \hat{u})\}$.

Then the closed-loop system with $\hat{u} = -b^T(x)\hat{P}_x^T$ satisfies

$$\int_{t_0}^{t_1} ||z||^2 dt \leq \gamma^2 \int_{t_0}^{t_1} ||d||^2 dt + 2\hat{P}(x_0) + 2\epsilon_H t_1 \tag{9}$$

where $\epsilon_H \leq \sqrt{\epsilon_T}$, with $\sqrt{\epsilon_T}$ the RMS value of the one-sided PDE loss. As $\epsilon_T \to 0$ and $\delta \to 0$ jointly, $\epsilon_H \to 0$ and the exact $\mathcal{L}_2$ *gain* bound of Theorem 1 is recovered.

Proof: By the saddle-point structure of K_γ, for any $d \in L_2$ and with $\hat{p} = (D_\delta^+ \hat{P} + D_\delta^- \hat{P})/2$: $K_\gamma(x, \hat{p}, d, \hat{u}) \leq K_\gamma(x, \hat{p}, \hat{d}, \hat{u}) \leq \epsilon_H(x)$. Since $\hat{P}_x f(x, \hat{u}, d) = K_\gamma(x, \hat{p}, d, \hat{u}) + (\gamma^2/2)||d||^2 - (1/2)||z||^2$, the chain rule gives

$$\frac{d\hat{P}}{dt} \leq \frac{\gamma^2}{2}||d||^2 - \frac{1}{2}||z||^2 + \epsilon_H. \qquad (10)$$

Integrating (10) over $[t_0, t_1]$ and using $\hat{P}(x(t_1)) \geq 0$ yields the states $\mathcal{L}_2$ *gain* bound. For the bound on ϵ_H: since $\mathcal{L}_{PDE}(\theta) = \mathbb{E}_{x \sim \omega}\left[(max\left\{0, \widehat{H}_\alpha\right\})^2\right] = \epsilon_T$, Cauchy-Schwarz gives $\mathbb{E}_{x \sim \omega}\left[max\left\{0, \widehat{H}_\alpha\right\}\right] \leq \sqrt{\epsilon_T}$. Because $\hat{H}_\alpha$ coincides with $K_\gamma(x, \hat{p}, \hat{d}, \hat{u})$ at the Lax-Friedrichs scheme ensure $\epsilon_H \to 0$ and the bound recovers the exact $\mathcal{L}_2$ *gain* $\leq \gamma$ of Theorem 1. ■

In the next section, we focus on showing the performance of the controller in nonlinear environments with uncertainty, presenting its capabilities to stabilize the systems and achieve the equilibrium point.

V. Simulations Results

In this section, we present two numerical examples to illustrate the effectiveness of the proposed method. First, an inverted pendulum with nonlinear dynamics is implemented. Then, a Kuramoto network model is simulated to illustrate the proposed control in a more complex nonlinear system.

A. Inverted Pendulum

Consider the inverted pendulum system in discrete time

$$\begin{aligned} x_1[k+1] &= x_1[k] + T_s x_2[k], \\ x_2[k+1] &= x_2[k] + T_s\left(-\frac{b}{ml^2}x_2[k] + \frac{g}{l}\sin(x_1[k])\right. \\ &\qquad \left. + \frac{1}{ml^2}u[k] + d[k]\right) \\ z &= \begin{bmatrix} x_1[k+1] & x_2[k+1] & u[k] \end{bmatrix}^T \end{aligned}$$

with $T_s = 0.01s$, k the discrete step, $l = 1\,\text{m}$, $m = 1\,\text{kg}$, $g = 9.8\,\text{m/s}^2$, and $b = 0.65\,\text{N·m·s/rad}$. For training the network, we choose $\gamma = 5$ and $\delta = 0.01$ for the Lax-Friedrichs approximation. The network consists of a multilayer perceptron with two layers of 80 neurons and Sigmoid activation functions, trained over 8,000 points and 50,000 epochs for 534 seconds. As a result, we obtained the losses in Fig. 1a and the smooth storage function $\hat{P}(x)$ in Fig. 2a.

The resulting Hamiltonian in Fig. 2b accomplishes the requirement to be $K_\gamma \leq 0$ ensuring that the closed-loop system has $\mathcal{L}_2$ *gain* $\leq \gamma$.

In order to evaluate the robustness of the controller, a Monte Carlo analysis was made over 150 simulation varying the parameters of the pendulum in a uniform

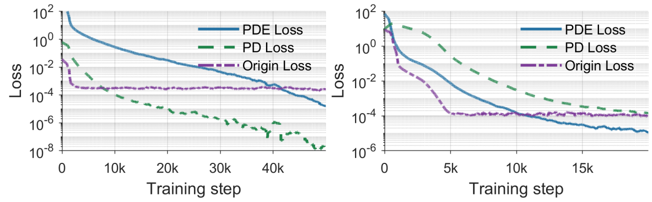

(a) Training losses functions for the pendulum system.
(b) Training losses functions for the Kuramoto oscillator system.

Fig. 1: Training loss convergence for the proposed PINN framework. Fig. 1a Inverted pendulum system over 50,000 steps, with final PDE, positive-definiteness, and origin losses reaching the order of 10^{-4}, 10^{-8}, and 10^{-4}, respectively. Fig. 1b Kuramoto oscillator network over 20,000 steps, with corresponding final losses of order 10^{-5}, 10^{-4}, and 10^{-4}. Convergence of all three loss components confirms that the learned storage function $\hat{P}(x)$ satisfies the HJI subsolution condition, positive-definiteness, and the boundary condition $\hat{P}(0) = 0$.

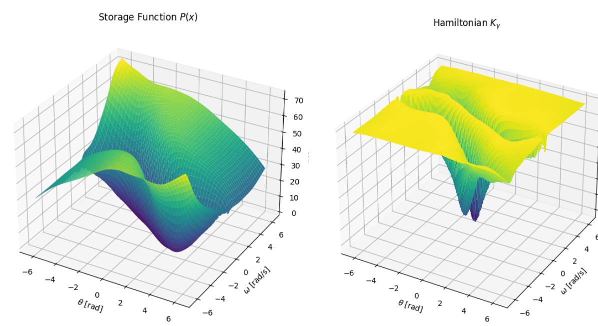

(a) Storage function $\hat{P}(x)$.
(b) Hamiltonian K_γ.

Fig. 2: Fig. 2a shows the storage function that satisfies the Hamiltonian for the inverted pendulum system. This function fulfills the requirements to be a storage function, such as $\hat{P}(0) = 0$ and positive over the domain. Fig. 2b shows the surface of the resulting Hamiltonian for the corresponding storage function $\hat{P}(x)$ and the inverted pendulum system.

distribution within the intervals $b \in [0, 2]\,\text{N·m·s/rad}$, $l \in [0.5, 1.5]\,\text{m}$, $m \in [0.5, 1.5]\,\text{kg}$, $g \in [8, 11]\,\text{m/s}^2$. The results in Fig. 3a show that the controller can handle the uncertainty and reach the inverted position (0 rad).

B. Kuramoto Oscillator Network

Consider the Kuramoto Oscillator networked system

$$\begin{aligned} \phi_i[k+1] = \phi_i[k] + T_s(w_i + \sum_{j=1}^{N} K_i a_{ij} \sin(\phi_j[k] - \phi_i[k]) \\ + u[k] + d[k]) \end{aligned}$$

where N denotes the number of nodes, K_i is the coupling gain associated with node i, a_{ij} represents the adjacency between nodes i and j, and $\phi[k]$ is the phase of node i at time step k. The terms $u[k]$ and $d[k]$ correspond to the control input and an external disturbance, respectively. For training,

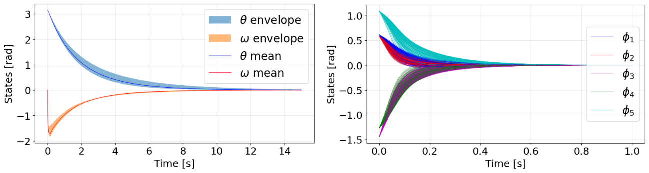

(a) Pendulum states. (b) Kuramoto oscillator states.

Fig. 3: Fig. 3a Shows the pendulum states with parameter variations in the domain $b \in [0, 2]$ N·m·s/rad, $l \in [0.5, 1.5]$ m, $m \in [0.5, 1.5]$ kg, $g \in [8, 11]$ m/s^2. The shaded areas correspond to the minimum and maximum values of each state. Fig. 3b shows the Kuramoto oscillator states over the Monte Carlo analysis, varying the coupling factor and network topology. As shown in the figure, each state converged to the reference point $\phi = 0$.

we used a fully connected 5 nodes Kuramoto oscillator with coupling $K = 1$, $\gamma = 5$, and sampling time $T_s = 0.01$. The neural network employed has 5 inputs and 1 output, 2 hidden layers with Sigmoid activation and 100 neurons in each one. Trained with a learning rate of 0.001 over $20,000$ data points and $20,000$ epochs in 176 seconds. As a result, we got the loss functions in Fig. 1b.

In order to test the performance of the controller, we implement a Monte Carlo analysis varying the coupling factor K with a uniform distribution in $[0.1, 3]$. Furthermore, for each simulation, we use a random network topology, including disconnected nodes, ring, hierarchical, mesh, and star topologies.

As a result of 100 simulations with the same initial conditions, we can observe in Fig. 3b that every node of every topology reaches the equilibrium point $\phi = 0$, in less than 1 second, even if the network has nodes disconnected. Although the controller has information about every node, the dynamics of each one evolve differently for every topology.

VI. Conclusions and Future Work

In this work, we presented a framework for computing robust optimal state-feedback controllers for nonlinear systems by combining Physics-Informed Neural Networks with a monotone Lax-Friedrichs numerical Hamiltonian. By recasting the nonlinear H_∞ state-feedback synthesis problem as the search for a viscosity subsolution of the Hamilton–Jacobi–Isaacs inequality, we showed that the proposed architecture can approximate valid storage functions that guarantee stability and an $\mathcal{L}_2$ gain bound. The use of a numerical Hamiltonian with controlled artificial viscosity provides a convexified optimization landscape and ensures monotonicity and consistency, enabling PINNs to inherit the convergence properties of classical finite-difference schemes. The method successfully produced smooth storage functions and stabilizing policies for two nonlinear benchmark systems: the inverted pendulum and the Kuramoto oscillator network. By performing Monte Carlo simulations, it is demonstrated that the resulting controllers exhibit strong robustness properties, maintaining stability and convergence even under wide variations in physical parameters, model uncertainties, and network topologies.

Several aspects remain for future exploration. While the Lax-Friedrichs Hamiltonian provides theoretical guarantees, its computational cost scales poorly with high dimension. Future research may explore different strategies to expand the approach to higher-dimensional nonlinear systems and attempt to avoid online differentiation. Furthermore, the present framework assumes full-state feedback, a natural challenge is to expand this approach to output feedback H_∞. Finally, sim-to-real experiments will be conducted in order to identify implementation limitations and determine the performance over real systems.

References

[1] J. C. Willems, "Dissipative dynamical systems," *European Journal of Control*, vol. 13, no. 2-3, pp. 134–151, 2007.

[2] H. K. Khalil and P. V. Kokotović, *Nonlinear Systems*. Springer, 2002.

[3] A. van der Schaft, *L2-Gain and Passivity Techniques in Nonlinear Control*. Springer International Publishing, 2017.

[4] C. Esteve-Yagüe, R. Tsai, and A. Massucco, "Finite-difference least square methods for solving hamilton-jacobi equations using neural networks," *Journal of Computational Physics*, vol. 524, p. 113721, 2025.

[5] Z. Kang and J. Wilkening, "Finite-difference least-square methods for solving hamilton–jacobi equations using neural networks," *Journal of Scientific Computing*, vol. 92, no. 3, pp. 1–32, 2022.

[6] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.

[7] K. Nishi, S. Sonoda, M. Ikeda, and N. Murata, "Finite-difference least-square approaches for solving hamilton–jacobi–bellman equations by neural networks," *Neural Networks*, vol. 142, pp. 328–342, 2021.

[8] Y. Zheng, C. Hu, X. Wang, and Z. Wu, "Physics-informed recurrent neural network modeling for predictive control of nonlinear processes," *Journal of Process Control*, vol. 128, p. 103005, 2023.

[9] T. X. Nghiem, J. Drgoňa, C. Jones, Z. Nagy, R. Schwan, B. Dey, A. Chakrabarty, S. Di Cairano, J. A. Paulson, A. Carron, M. N. Zeilinger, W. Shaw Cortez, and D. L. Vrabie, "Physics-informed machine learning for modeling and control of dynamical systems," in *2023 American Control Conference (ACC)*, 2023, pp. 3735–3750.

[10] J. Drgoňa, T. X. Nghiem, T. Beckers, M. Fazlyab, E. Mallada, C. Jones, D. Vrabie, S. L. Brunton, and R. Findeisen, "Safe physics-informed machine learning for dynamics and control," in *2025 American Control Conference (ACC)*. IEEE, 2025, pp. 591–606.

[11] M.-C. Lai, Y. Song, X. Yuan, H. Yue, and T. Zeng, "The hard-constraint pinns for interface optimal control problems," *SIAM Journal on Scientific Computing*, vol. 47, no. 3, pp. C601–C629, 2025.

[12] C. J. García-Cervera, M. Kessler, and F. Periago, "Control of partial differential equations via physics-informed neural networks," *Journal of Optimization Theory and Applications*, vol. 196, pp. 391–414, 2023.

[13] J. Barry-Straume, A. Sarshar, A. A. Popov, and A. Sandu, "Physics-informed neural networks for pde-constrained optimization and control," *Communications in Applied Mathematics and Computation*, 2025.

[14] E. A. Antonelo, E. Camponogara, L. O. Seman, J. P. Jordanou, E. R. de Souza, and J. F. Hübner, "Physics-informed neural nets for control of dynamical systems," *Neurocomputing*, vol. 579, p. 127419, 2024.

[15] L. K. Miyatake, E. Camponogara, E. A. Antonelo, and A. Pavlov, "Physics-informed neural networks for control of single-phase flow systems governed by partial differential equations," *Engineering Applications of Artificial Intelligence*, vol. 160, p. 111869, 2025.

[16] A. Norambuena, M. Mattheakis, F. J. González, and R. Coto, "Physics-informed neural networks for quantum control," *Physical Review Letters*, vol. 132, no. 1, p. 010801, 2024.

[17] G. Barles and P. Souganidis, "Convergence of approximation schemes for fully nonlinear second order equations," in *29th IEEE Conference on Decision and Control*, 1990, pp. 2347–2349 vol.4.