



HAL
open science

Controlled oscillation modeling using port-Hamiltonian neural networks

Maximino Linares, Guillaume Doras, Thomas Hélié, Axel Roebel

► To cite this version:

Maximino Linares, Guillaume Doras, Thomas Hélié, Axel Roebel. Controlled oscillation modeling using port-Hamiltonian neural networks. 2026. <hal-05515374>

HAL Id: hal-05515374

<https://hal.science/hal-05515374v1>

Preprint submitted on 17 Feb 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-SA 4.0 - Attribution - Non-commercial use - ShareAlike - International License

Controlled oscillation modeling using port-Hamiltonian neural networks

M. Linares^{a,*}, G. Doras^a, T. Hélie^a, A. Roebel^a

^aIRCAM, 1, place Igor-Stravinsky, 75004, Paris, France

Abstract

Learning dynamical systems through purely data-driven methods is challenging as they do not learn the underlying conservation laws that enable them to correctly generalize. Existing port-Hamiltonian neural network methods have recently been successfully applied for modeling mechanical systems. However, even though these methods are designed on power-balance principles, they usually do not consider power-preserving discretizations and often rely on Runge-Kutta numerical methods. In this work, we propose to use a second-order discrete gradient method embedded in the learning of dynamical systems with port-Hamiltonian neural networks. Numerical results are provided for three systems deliberately selected to span different ranges of dynamical behavior under control: a baseline harmonic oscillator with quadratic energy storage; a Duffing oscillator, with a non-quadratic Hamiltonian offering amplitude-dependent effects; and a self-sustained oscillator, which can stabilize in a controlled limit cycle through the incorporation of a nonlinear dissipation. We show how the use of this discrete gradient method outperforms the performance of a Runge-Kutta method of the same order. Experiments are also carried out to compare two theoretically equivalent port-Hamiltonian systems formulations and to analyze the impact of regularizing the Jacobian of port-Hamiltonian neural networks during training.

Keywords: physics-informed machine learning, port-Hamiltonian neural networks, discrete gradient, Jacobian regularization

1. Introduction

Purely data-driven methods for dynamical systems pose several challenges: the volume of useful data is generally limited, they produce accurate-but-wrong predictions, they are not capable of dealing with uncertainty and their predictions are not explainable nor interpretable [1]. At the same time, it is natural to leverage the prior knowledge obtained through centuries of scientific study in the form of inductive bias [2] when designing predictive models [3]. In this direction, a successful data-driven physical model is one whose inductive bias better captures the true dynamics and is able to predict a correct outcome for data not observed during training. These inductive bias are incorporated through soft and hard constraints [4]. Soft constraints add penalty terms to the training loss function, discouraging violations of physical laws. This approach is widely applicable, but the model must balance adherence to the constraints against fitting the observed data, without providing any formal guarantee. The seminal *Physics-Informed Neural Networks* (PINNs) [5] is an example of soft-constraining the model. In contrast, hard constraints ensure strict compliance with specified physical

laws by embedding them directly into the model's structure, independently of the available data. In this sense, hard constraints can be used to incorporate energy conservation laws, symmetry, numerical methods for PDEs or Koopman theory [6]. However, imposing hard constraints reduces the space of possible solutions and generally limits the model's expressiveness. As a result, hard constraints are difficult to apply in practice: incorrect assumptions about the physical system can lead to overly biased models with a poor generalization performance.

In this paper, we consider dynamical systems whose state \mathbf{x} is governed by the following ODE:

$$\frac{d\mathbf{x}}{dt} := \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) \quad (1)$$

and analyze how hard constraints based on energy conservation and power balance principles are incorporated into neural networks based on *physically consistent* port-Hamiltonian systems formulations. By *physically consistent*, we refer, in this work, to the combination of power-balanced state-space models with discrete gradient numerical methods, which preserve the system's energy during discretization. The central hypothesis is that enforcing this physical structure as a hard constraint improves interpretability and generalization with respect to a vanilla NeuralODE. To substantiate this claim, we conduct a systematic study on three controlled oscillatory systems of

*Corresponding author

Email addresses: maximino.linares@ircam.fr (M. Linares), guillaume.doras@ircam.fr (G. Doras), thomas.helie@ircam.fr (T. Hélie), axel.roebel@ircam.fr (A. Roebel)

increasing modeling complexity: a harmonic oscillator, a Duffing oscillator and a self-sustained oscillator. These systems are deliberately selected to include linear and nonlinear Hamiltonian dynamics as well as nonlinear dissipation mechanisms. The harmonic oscillator serves as the simplest baseline example with quadratic energy storage; the Duffing oscillator offers nonlinearities in the Hamiltonian, capturing amplitude-dependent effects; and the self-sustained oscillator incorporates a nonlinear dissipation which can stabilize the system in a controlled-limit cycle.

The main contributions of this paper are

- a comparison of two theoretically equivalent port-Hamiltonian systems (PHS) formulations: the semi-explicit PH-Differential-Algebraic-Equations (PH-DAE) and the input-state-output PHS with feedthrough; when they are implemented as port-Hamiltonian neural networks (PHNNs).
- a performance comparison between the Gonzalez discrete gradient method, which is a second-order energy-preserving numerical method, and a second-order explicit Runge-Kutta method when used to discretize the PHNN model during learning.
- an empirical study of the impact of regularizing the Jacobian of PHNN through two methods already applied to NeuralODEs and a new one tackling the stiffness of the learned ODE solutions.

The rest of this paper is organized as follows. Section 2 introduces the necessary preliminaries on dynamical systems, port-Hamiltonian systems, numerical methods, neural ordinary differential equations, and port-Hamiltonian neural networks. Section 3 presents the port-Hamiltonian formulations considered in this work, the enforced physical constraints, and the oscillatory examples used throughout the paper. Section 4 focuses on port-Hamiltonian neural networks, detailing how physical constraints are incorporated into the learning process, the comparison between continuous- and discrete-time models, and the Jacobian regularization in the port-Hamiltonian neural networks. Section 5 formulates the key research questions addressed by the experimental study. Section 6 reports and discusses the results of the experiments. Finally, Section 7 concludes the paper and outlines directions for future work. The code is publicly available: <https://github.com/mlinaresv/ControlledOscillationPHNNs> [will be released after paper acceptance].

2. Preliminaries

2.1. Dynamical systems

2.1.1. Dynamical system ODE

Consider a dynamical system governed by the following system of equations:

$$\begin{cases} \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) \\ \mathbf{y}(t) = \mathbf{h}(\mathbf{x}(t), \mathbf{u}(t)) \end{cases} \quad (2)$$

where $\mathbf{x}(t) \in \mathbb{R}^{n_x}$, $\mathbf{u}(t) \in \mathbb{R}^{n_u}$ and $\mathbf{y}(t) \in \mathbb{R}^{n_y}$ are referred to as the *state*, *input* and *output* of the system, respectively; and where $\mathbf{f} : \mathcal{D} \rightarrow \mathbb{R}^{n_x}$, $\mathbf{h} : \mathcal{D} \rightarrow \mathbb{R}^{n_y}$ are \mathcal{C}^1 functions with domain $\mathcal{D} \subseteq \mathbb{R}^{n_x} \times \mathbb{R}^{n_u}$. The system of equations (2) as a whole is referred to as *state-space model* [7].

In this work, we consider autonomous systems (i.e. where $\mathbf{u}(t)$ is constant for all t). Obtaining a solution for a given initial condition is often referred to as solving the *initial value problem* (IVP)

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) \quad \mathbf{x}(0) = \mathbf{x}_0. \quad (3)$$

In the following, we omit the explicit time dependence of $\mathbf{x}(t)$, $\mathbf{y}(t)$ and $\mathbf{u}(t)$ to simplify the notation, when there is no ambiguity, and we let $\mathbf{J}_{\mathbf{f}}(\mathbf{x})$ denote the *Jacobian matrix* of the function \mathbf{f} of (3).

2.1.2. Well-posed problems

The term *well-posed* was introduced to refer to problems where a) the solution exists, b) is unique, and c) depends continuously on the initial conditions and parameters [8]. A sufficient condition for an IVP to be well-posed is that \mathbf{f} is K -Lipschitz [9].

The *spectral norm* of a matrix \mathbf{A} is defined as:

$$\|\mathbf{A}\|_2 = \max_{\|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2 = \sigma_{max}(\mathbf{A}), \quad (4)$$

where $\sigma_{max}(\mathbf{A})$ is the largest *singular value* of \mathbf{A} [10]. It is a standard result that if the following condition holds:

$$\|\mathbf{J}_{\mathbf{f}}(\mathbf{x})\|_2 \leq K < \infty \quad \forall \mathbf{x} \in \mathcal{D}$$

then \mathbf{f} is K -Lipschitz [9] (see Appendix A).

In practice, controlling the well-posedness of the IVP reduces to enforcing an upper bound on the spectral norm of the Jacobian of \mathbf{f} .

2.1.3. Well-conditioned problems

The term *well-conditioned* refers to problems seen as a function \mathbf{f} where a small perturbation of \mathbf{x} yields only small changes in $\mathbf{f}(\mathbf{x})$ – the meaning of *small* depending

on the context [11]. This relationship can be characterized by the *relative condition number* of \mathbf{f} defined as:

$$\kappa_{\mathbf{f}}(\mathbf{x}) = \lim_{\delta \rightarrow 0} \sup_{\|\delta \mathbf{x}\| \leq \delta} \left(\frac{\|\mathbf{f}(\mathbf{x} + \delta \mathbf{x}) - \mathbf{f}(\mathbf{x})\|_2}{\|\mathbf{f}(\mathbf{x})\|_2} \bigg/ \frac{\|\delta \mathbf{x}\|_2}{\|\mathbf{x}\|_2} \right), \quad (5)$$

where $\delta \mathbf{x}$ and $\delta \mathbf{f}$ are infinitesimal. If $\mathbf{f} \in \mathcal{C}^1$, this rewrites as:

$$\kappa_{\mathbf{f}}(\mathbf{x}) = \frac{\|\mathbf{J}_{\mathbf{f}}(\mathbf{x})\|_2}{\|\mathbf{f}(\mathbf{x})\|_2 / \|\mathbf{x}\|_2}. \quad (6)$$

for $\|\mathbf{x}\|_2 \neq 0$ and its limit when $\|\mathbf{x}\|_2 \rightarrow 0$. If $\kappa_{\mathbf{f}}(\mathbf{x})$ is small (resp. large), the problem is said to be *well-conditioned* (resp. *ill-conditioned*).

For a linear system $\mathbf{f}(\mathbf{x}) = \mathbf{A}\mathbf{x}$, the Jacobian of \mathbf{f} becomes $\mathbf{J}_{\mathbf{f}}(\mathbf{x}) = \mathbf{A}$, $\forall \mathbf{x}$. It is a standard result that the condition number can then be bounded s.t.:

$$\kappa_{\mathbf{f}}(\mathbf{x}) \leq \|\mathbf{J}_{\mathbf{f}}(\mathbf{x})\|_2 \|\mathbf{J}_{\mathbf{f}}^{-1}(\mathbf{x})\|_2 = \frac{\sigma_{\max}(\mathbf{J}_{\mathbf{f}}(\mathbf{x}))}{\sigma_{\min}(\mathbf{J}_{\mathbf{f}}(\mathbf{x}))} \quad (7)$$

where $\sigma_{\max}(\mathbf{J}_{\mathbf{f}}(\mathbf{x}))$, $\sigma_{\min}(\mathbf{J}_{\mathbf{f}}(\mathbf{x}))$ are the maximum and minimum singular values of $\mathbf{J}_{\mathbf{f}}(\mathbf{x})$ [11]. In the following, the upper bound in (7) is denoted as $\kappa(\mathbf{J}_{\mathbf{f}}(\mathbf{x}))$.

In practice, controlling the well-conditioning of a linear problem \mathbf{f} reduces to enforcing an upper bound on its condition number.

2.1.4. Stiff ODE systems

The term *stiff* refers to an initial value problem for which certain numerical methods require prohibitively small step sizes to maintain stability. This behavior can be characterized by the *stiffness ratio* defined as:

$$\rho(\mathbf{J}_{\mathbf{f}}) = \frac{\max |\Re(\lambda(\mathbf{J}_{\mathbf{f}}))|}{\min |\Re(\lambda(\mathbf{J}_{\mathbf{f}}))|}, \quad (8)$$

where $\lambda(\mathbf{J}_{\mathbf{f}})$ are the *eigenvalues* of the Jacobian matrix and $\Re(\cdot) : \mathbb{C} \rightarrow \mathbb{R}$ is the real part operator [12]. Although it is rigorously only true for linear equations, a system with a large stiffness ratio is generally stiff [12].

In practice, controlling the stiffness of an ODE system reduces to enforcing an upper bound on its stiffness ratio.

2.2. Port-Hamiltonian systems

In the Hamiltonian formalism [13], the mechanical state of \mathcal{S} is represented by a vector $\mathbf{x} = [\mathbf{q}, \mathbf{p}]^{\top} \in \mathbb{R}^{n_x}$, where $\mathbf{q}, \mathbf{p} \in \mathbb{R}^{n_x/2}$ denote the generalized coordinates and conjugate momenta. The dynamics is governed through a scalar-valued energy function $H(\mathbf{x})$ known as *Hamiltonian*, and the time evolution of the system follows Hamilton's equations:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) = \mathbf{S} \nabla H(\mathbf{x}) = \begin{bmatrix} 0 & \mathbb{I}_n \\ -\mathbb{I}_n & 0 \end{bmatrix} \nabla H(\mathbf{x}), \quad (9)$$

where $\mathbf{S} \in \mathbb{R}^{n_x \times n_x}$ is the canonical symplectic matrix which imposes the energy conservation principle. This formalism is restricted to conservative closed systems and does not readily describes dissipation or external control, which are common to many real-world systems.

Port-Hamiltonian formalism [14–16] generalizes Hamiltonian mechanics to multi-physics open systems by explicitly modeling energy exchange with the environment through ports (inputs/outputs) and dissipation. In this work, the considered class of open systems is represented in this formalism as a network of:

- *energy storing components* with state \mathbf{x} and energy $E := H(\mathbf{x})$, with H positive definite and \mathcal{C}^1 -regular, so the stored power is $P_{\text{stored}} := \dot{E} = \nabla H(\mathbf{x})^{\top} \dot{\mathbf{x}}$,
- *dissipative components*, described by an effort law $\mathbf{z}(\mathbf{w})$ for a flow variable \mathbf{w} , where $P_{\text{diss}} := \mathbf{z}(\mathbf{w})^{\top} \mathbf{w}$; by convention, the dissipated power is counted positively, i.e. $P_{\text{diss}} \geq 0$, with $P_{\text{diss}} = 0$ for conservative components,
- *external components* are represented through ports by system inputs \mathbf{u} and outputs \mathbf{y} , with the convention that $P_{\text{ext}} := \mathbf{u}^{\top} \mathbf{y}$ is positive when received by these external components.

The coupling of internal flows \mathcal{F} and efforts \mathcal{E} governs the time evolution of the system, expressed as:

$$\underbrace{\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{w} \\ \mathbf{y} \end{bmatrix}}_{:=\mathcal{F}} = \mathbf{S} \underbrace{\begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{z}(\mathbf{w}) \\ \mathbf{u} \end{bmatrix}}_{:=\mathcal{E}}, \quad (10)$$

where $\mathbf{S} = -\mathbf{S}^{\top}$ is skew-symmetric so that

$$0 = \underbrace{P_{\text{stored}}}_{\nabla H(\mathbf{x})^{\top} \dot{\mathbf{x}}} + \underbrace{P_{\text{diss}}}_{\mathbf{z}(\mathbf{w})^{\top} \mathbf{w}} + \underbrace{P_{\text{ext}}}_{\mathbf{u}^{\top} \mathbf{y}} = \langle \mathcal{E} | \mathcal{F} \rangle \quad (11)$$

with $\langle \mathcal{E} | \mathcal{F} \rangle = \mathcal{E}^{\top} \mathcal{F}$. Indeed,

$$\mathcal{E}^{\top} \mathcal{F} \stackrel{(14)}{=} \mathcal{E}^{\top} \mathbf{S} \mathcal{E} \stackrel{(S=-S^T)}{=} 0 \quad (12)$$

The passivity of the system stems from the fact that $P_{\text{diss}} \geq 0$, which imposes that $P_{\text{stored}} = -P_{\text{diss}} - P_{\text{ext}} \leq -P_{\text{ext}}$. As a consequence, $P_{\text{stored}} \leq 0$ (non-increasing internal energy) when external sources are off. This intrinsically physics-consistent formulation is highly general and can be applied to a wide range of physical domains, including acoustics, fluid mechanics, quantum physics and others [17–21] (see [16] for formulations more general than (10)).

2.3. Numerical methods

For a given IVP (3), a *numerical method* approximates the solution without the need to analytically solve the ODE. A *discrete trajectory* \mathcal{T} is defined as the set $\{\mathbf{x}_n\}_{n \in \mathbb{N}}$

of consecutive states where $\mathbf{x}_n := \mathbf{x}(nh)$, $h = 1/sr$ is the time step and sr the sampling rate. Given the *initial state* $\mathbf{x}_0 = \mathbf{x}(0)$, the applied control \mathbf{u} and the ODE governing \mathbf{x} , a trajectory \mathcal{T} is obtained by such numerical method. If the governing equation is given by (10), the resulting \mathcal{T} is said to be a *discretization of a port-Hamiltonian system*. In this work, we use numerical methods that belong to two different categories:

2.3.1. Runge-Kutta (RK) methods

A well-known family of numerical methods is the *Runge-Kutta methods* [22]. Let $b_i, a_{i,j}$ ($i, j = 1, \dots, s$) be real numbers and let $c_i = \sum_{j=1}^s a_{i,j}$. An *s-stage* Runge-Kutta method is given by

$$\begin{cases} \mathbf{k}_i = \mathbf{f}\left(t_0 + hc_i, \mathbf{x}_n + h \sum_{j=1}^s a_{i,j} \mathbf{k}_j\right), & i = 1, \dots, s \\ \mathbf{x}_{n+1} = \mathbf{x}_n + h \sum_{i=1}^s b_i \mathbf{k}_i \end{cases} \quad (13)$$

where the weights $b_i, a_{i,j}$ are chosen to reach an accuracy order. In particular, the *explicit midpoint method* (RK2)

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \mathbf{f}\left(t_n + \frac{h}{2}, \mathbf{x}_n + \frac{h}{2} \mathbf{f}(t_n, \mathbf{x}_n)\right) \quad (14)$$

is a second-order explicit Runge-Kutta method.

2.3.2. Discrete gradient (DG) methods

Physical properties like energy-preservation are not in general respected when discretizing a port-Hamiltonian system using RK methods [23]. However, in Hamiltonian mechanics there is a rich theory of structure preserving integrators [22]. In particular, discrete gradient methods are a family of geometrical integrators that preserve the exact energy by construction [24]. A *discrete gradient* $\bar{\nabla}H : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_x}$ is an approximation of the gradient of a function $H : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$, satisfying the following two properties¹:

1. $\bar{\nabla}H(\mathbf{x}, \delta\mathbf{x})^T \delta\mathbf{x} = H(\mathbf{x} + \delta\mathbf{x}) - H(\mathbf{x})$
2. $\bar{\nabla}H(\mathbf{x}, \mathbf{0}) = \nabla H(\mathbf{x})$

where $\delta\mathbf{x} = \mathbf{x}' - \mathbf{x}$. In this work, we consider the *Gonzalez discrete gradient* method [25] (DG)

$$\begin{aligned} \bar{\nabla}_{mid}H(\mathbf{x}, \delta\mathbf{x}) &= \nabla H\left(\mathbf{x} + \frac{1}{2}\delta\mathbf{x}\right) \\ &+ \frac{H(\mathbf{x} + \delta\mathbf{x}) - H(\mathbf{x}) - \nabla H(\mathbf{x} + \frac{1}{2}\delta\mathbf{x})^T \delta\mathbf{x}}{\delta\mathbf{x}^T \delta\mathbf{x}} \delta\mathbf{x}. \end{aligned} \quad (15)$$

Note that this method is, in general, a second-order inverse-explicit integrator [26] that becomes linearly implicit when the Hamiltonian satisfies $H(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x}$.

¹In the literature, these two properties are usually written as

1. $\bar{\nabla}H(\mathbf{x}, \mathbf{x}')^T (\mathbf{x}' - \mathbf{x}) = H(\mathbf{x}') - H(\mathbf{x})$
2. $\bar{\nabla}H(\mathbf{x}, \mathbf{x}) = \nabla H(\mathbf{x})$.

2.4. Neural ODEs and Jacobian regularization

Neural Differential Ordinary Equations (NODEs) were introduced by Chen et al. [27] to define the evolution of a system's state using an ODE whose dynamics are modeled by a neural network:

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{f}_\theta(t, \mathbf{x}(t)) \quad (16)$$

where θ represents the parameters of the neural network. The forward pass of a NODE is defined as

$$\mathbf{x}(t_{n+1}) = \mathbf{x}(t_n) + \int_{t_n}^{t_{n+1}} \mathbf{f}_\theta(t, \mathbf{x}(t)) dt \quad (17)$$

where in practice the integration is approximated by any numerical method.

A well-documented challenge in the training of NODEs is the fact that the Jacobian of the learned dynamics \mathbf{f}_θ becomes poorly conditioned as the training progresses. Several studies have underlined this phenomenon. For instance, Dupont et al. [28] reported that when NODEs are overfitted on the MNIST dataset [29], the resulting flow may become so ill-conditioned that the numerical ODE solver is forced to take time steps smaller than machine precision, leading to numerical underflow and an increase in the number of function evaluations (NFE). Such pathological dynamics can also destabilize training and cause the loss to diverge. Motivated by these observations, Finlay et al. [30] emphasized the importance of explicitly constraining the learned vector field. They noted that, in the absence of such constraints, the learned dynamics may exhibit poor conditioning, which in turn degrades numerical integration and degrades training performance. To address this issue, they proposed regularizing the Jacobian of \mathbf{f}_θ using its Frobenius norm. In this direction, Josias et al. [31] proposed instead to regularize the Jacobian condition number, arguing that it reduced NFE without a significant loss in accuracy and controlling at the same time the Jacobian norm. Focusing on Jacobian regularization enables a connection to sensitivity analysis found in neural network literature, where also the spectral norm regularization has been employed [32] [33].

2.5. Port-Hamiltonian neural networks

In the context of learning Hamiltonian ODEs, Greydanus et al. [34] introduced *Hamiltonian Neural Network* (HNN), modifying the usual NODEs framework by parameterizing the Hamiltonian function of a given conservative physical system. Once the Hamiltonian is parameterized, (9) is leveraged in the loss function. This work led to many others trying to generalize the framework to learn more general physical systems. For example, Sosanya et al. [35] generalized the approach of HNN to non-conservative systems using Helmholtz's decomposition theorem to parameterize the dissipation potential. Several works integrating port-Hamiltonian systems theory into neural networks are

found in literature. Most of them implement the input-state-output representation² [16]

$$\begin{cases} \dot{\mathbf{x}} = (\mathbf{J}(\mathbf{x}) - \mathbf{R}(\mathbf{x}))\nabla_{\mathbf{x}}H(\mathbf{x}) + \mathbf{G}(\mathbf{x})\bar{\mathbf{u}} \\ \bar{\mathbf{y}} = \mathbf{G}^T(\mathbf{x})\nabla H(\mathbf{x}) \end{cases} \quad (18)$$

for designing the computation graph of the dynamics. In this context, Desai et al. [36] introduce the *Port-Hamiltonian Neural Network* (PHNN) to learn damped and controlled systems. Zhong et al. [37] use this formulation to model conservative systems with external control introducing the *Symplectic ODE-Net* (SymODEN). In another work, the same authors generalize the framework of SymODEN to learn also the dissipation [38]. Cherifi et al. [39] propose a framework based on input-state-output data to learn port-Hamiltonian systems described by (18). Roth et al. [40] introduce *Stable Port-Hamiltonian Neural Networks* (sPHNN) to learn dynamical systems with a single equilibrium under stability guarantees. A similar work to the one presented in this article, but based on pseudo-Hamiltonian formalism, is found in [4]. These approaches learn the dynamics through (18) and impose *a priori* knowledge, either on the expression of the Hamiltonian [37–41] or on the way dissipation affects the system [4, 36, 39, 40]. As in any other NODE framework, (18) is discretized, during or after the training, using a numerical method so that there is no preservation of the power balance and no guarantee of stability. Generally, this has been done either by high-order RK methods [34–36, 39, 40], but also via symplectic integrators [41–45], although the impact on the accuracy of one or the other approach has not been thoroughly studied. To the best of our knowledge, the use of discrete gradients methods has not yet been described in the PHNN literature.

3. Port Hamiltonian Models

3.1. Port-Hamiltonian formulations

Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{n_x}$ be the state of \mathcal{S} with an associated Hamiltonian $H : \mathcal{X} \rightarrow \mathbb{R}$ and $\mathbf{u}, \mathbf{y} \in \mathbb{R}^{n_u}$ be the input and output of the system, respectively. We consider three formulations in this work:

- (i) Semi-explicit PH-Differential-Algebraic-Equations (PH-DAE):

$$\underbrace{\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{w} \\ \mathbf{y} \end{bmatrix}}_{\text{flows } \mathcal{F}_i} = \mathbf{S} \underbrace{\begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{z}(\mathbf{w}) \\ \mathbf{u} \end{bmatrix}}_{\text{efforts } \mathcal{E}_i}, \quad (19)$$

where $\mathbf{S} = -\mathbf{S}^T \in \mathbb{R}^{(n_x+n_w+n_u) \times (n_x+n_w+n_u)}$ and $\mathbf{z}(\mathbf{w}) : \mathbb{R}^{n_w} \rightarrow \mathbb{R}^{n_w}$ is the resistive structure of \mathcal{S} .

The skew-symmetry of \mathbf{S} guarantees energy conservation and the resistive property of $\mathbf{z}(\mathbf{w})$ is given by $\mathbf{z}(\mathbf{w})^T \mathbf{w} \geq 0$, which guarantees passivity.

- (ii) Input-state-output PHS with feedthrough [46]:

$$\underbrace{\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{y} \end{bmatrix}}_{\text{flows } \mathcal{F}_{ii}} = (\mathbf{J} - \mathbf{R}(\mathbf{x}, \mathbf{u})) \underbrace{\begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{u} \end{bmatrix}}_{\text{efforts } \mathcal{E}_{ii}}, \quad (20)$$

where \mathbf{J} and $\mathbf{R}(\mathbf{x}, \mathbf{u}) \in \mathbb{R}^{(n_x+n_u) \times (n_x+n_u)}$ satisfy $\mathbf{J} = -\mathbf{J}^T$ and $\mathbf{R} = \mathbf{R}^T \succeq 0$. The skew-symmetry of \mathbf{J} accounts for conservative connections and the positive semi-definiteness of $\mathbf{R}(\mathbf{x}, \mathbf{u})$ guarantees passivity. This formulation can be retrieved from (i) in particular when $\mathbf{S}_{ww} = 0$ and $\mathbf{S}_{wx}, \mathbf{S}_{wu}$ do not depend on \mathbf{w} . It is also an extension of (18) for systems with direct feed-through.

- (iii) Skew-symmetric gradient PH-DAE: If the function $\mathbf{z}(\mathbf{w}) = \partial_{\mathbf{w}}Z(\mathbf{w})$ is derived from a potential Z (often referred to as the Rayleigh potential), the system (19) rewrites as the skew-gradient system [47]

$$\underbrace{\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{w} \\ \mathbf{y} \end{bmatrix}}_{\text{flows } \mathcal{F}_{iii}} = \mathbf{S} \nabla F \left(\underbrace{\begin{bmatrix} \mathbf{x} \\ \mathbf{w} \\ \mathbf{u} \end{bmatrix}}_{\text{efforts } \mathcal{E}_{iii}} \right), \quad (21)$$

where $\mathbf{S} = -\mathbf{S}^T \in \mathbb{R}^{(n_x+n_w+n_u) \times (n_x+n_w+n_u)}$ and $F = H(\mathbf{x}) + Z(\mathbf{w}) + \frac{\mathbf{u}^T \mathbf{u}}{2}$.

Only the formulations (i) and (ii) are considered in the experiments. The skew-symmetric gradient PH-DAE is presented as a generalization of the formulation used by Sosanya et al. [35]. The PHS formulations (i-iii) are passive (see Appendix B).

3.2. Physical constraints enforcement

The general class of physical systems considered in this work is such that the Hamiltonian H is \mathcal{C}^1 -regular positive definite. In addition, we consider the subclass of Hamiltonians of the form

$$H(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{Q}(\mathbf{x}) \mathbf{x}, \quad (22)$$

where $\mathbf{Q}(\mathbf{x})$ is a \mathcal{C}^1 -regular symmetric positive definite matrix function. As already mentioned, The PHS formulations (i)-(ii) satisfy two types of constraints by construction:

$$\begin{cases} \mathbf{S} \text{ (or } \mathbf{J}) \text{ skew-symmetric} & \text{(energy conservation)} \\ \mathbf{z}(\mathbf{w})^T \mathbf{w} \geq 0, \forall \mathbf{w} \text{ (or } \mathbf{R} \succeq 0) & \text{(passive laws)} \end{cases} \quad (23)$$

In the PH-DAE formulation, we consider the subclass of dissipation function of the form:

$$\mathbf{z}(\mathbf{w}) = \mathbf{\Gamma}(\mathbf{w})\mathbf{w} \quad (24)$$

²In the literature, this equation adopts the convention that $\bar{\mathbf{u}}^T \bar{\mathbf{y}} = \bar{P}_{ext}$ is positive when given to the system by the external components, so that $\bar{P}_{ext} = -P_{ext}$ and that $(\bar{\mathbf{u}}, \bar{\mathbf{y}}) = (\mathbf{u}, -\mathbf{y})$ or $(-\mathbf{u}, \mathbf{y})$, contrary to conventions used in (10) and all this paper

where $\Gamma(\mathbf{w}) \in \mathbb{R}^{n_w \times n_w}$ is \mathcal{C}^0 -regular positive semidefinite matrix function that admits a symmetric/skew-symmetric decomposition $\Gamma(\mathbf{w}) = \Gamma^{skew}(\mathbf{w}) + \Gamma^{sym}(\mathbf{w})$ with $\Gamma^{skew} = \frac{1}{2}(\Gamma - \Gamma^T)$ and $\Gamma^{sym} = \frac{1}{2}(\Gamma + \Gamma^T) \succeq 0$. The resistive property is satisfied as

$$\begin{aligned} \forall \mathbf{w}, \mathbf{z}(\mathbf{w})^T \mathbf{w} &= \mathbf{w}^T \Gamma(\mathbf{w})^T \mathbf{w} \\ &= \mathbf{w}^T \Gamma^{sym}(\mathbf{w})^T \mathbf{w} \stackrel{\Gamma^{sym} \succeq 0}{\geq} 0 \end{aligned} \quad (25)$$

Note that the classes of systems described by equations (22) to (25) cover a large spectrum of physical systems and remain fairly general.

3.3. Reparameterization of constraints

We now reparameterize the matrix constraints using the following properties:

- any symmetric positive semidefinite (resp. definite) matrix \mathbf{M}^{sym} can be written in the form $\mathbf{M}^{sym} = \mathbf{L}^T \mathbf{L}$ where \mathbf{L} is a lower triangular matrix with positive (resp. strictly positive) diagonal coefficients (Cholesky factorization [48]),
- any skew-symmetric matrix \mathbf{M}^{skew} can be written in the form $\mathbf{M}^{skew} = \mathbf{K} - \mathbf{K}^T$ where \mathbf{K} is a strictly lower triangular matrix.

Then, the physical constraints (23) are naturally satisfied considering the following reparameterizations:

$$\mathbf{Q}(\mathbf{x}) = \mathbf{L}_Q^T(\mathbf{x}) \mathbf{L}_Q(\mathbf{x}) \quad (26)$$

$$\Gamma(\mathbf{w}) = \mathbf{L}_\Gamma^T(\mathbf{w}) \mathbf{L}_\Gamma(\mathbf{w}) + \mathbf{K}_\Gamma(\mathbf{w}) - \mathbf{K}_\Gamma^T(\mathbf{w}) \quad \text{or} \quad (27)$$

$$\mathbf{R}(\mathbf{x}, \mathbf{u}) = \mathbf{L}_R^T(\mathbf{x}, \mathbf{u}) \mathbf{L}_R(\mathbf{x}, \mathbf{u}) \quad (28)$$

where $\mathbf{L}_{Q,R,\Gamma}$, resp. \mathbf{K}_Γ , are lower, resp. strictly lower, triangular matrices.

3.4. Oscillatory physical examples

Three oscillating systems that satisfy (22)-(24) are considered:

- the harmonic oscillator, which is quadratic in energy, with a linear dissipation;
- the Duffing oscillator, which is non-quadratic in energy, with a linear dissipation;
- a self-sustained oscillator, which is quadratic in energy, with a nonlinear dissipation.

This last physical system is of particular dynamical interest as it is designed to self-oscillate when combined with an adapted constant input, leading to a stable limit cycle [49]. Table 1 shows the different port-Hamiltonian structural elements for each of these systems.

4. Port-Hamiltonian neural networks

4.1. Port-Hamiltonian neural network models

A port-Hamiltonian neural network (PHNN) parameterizes the right-hand side of (19)-(20), modeling the Hamiltonian and the dissipative terms with two distinct neural networks. In this work, we assume that the interconnection matrices \mathbf{J} and \mathbf{S} , are given a priori. The Hamiltonian function is parameterized identically for each formulation by a neural network $H_{\theta_H}(\mathbf{x})$, and its gradient is derived by auto-differentiation as proposed in the seminal HNN [34]. The implementation of the dissipative term depends on the formulation considered (i or ii), which yields two different PHNN architectures:

- PHNN-S models (19), where the dissipation function $\mathbf{z}(\mathbf{w}) : \mathbb{R}^{n_w} \rightarrow \mathbb{R}^{n_w}$ is parameterized by a neural network $\mathbf{z}_{\theta_z}(\mathbf{w})$, implementing the dynamic function:

$$\mathbf{f}_\theta(\mathbf{x}, \mathbf{u}) = \mathbf{f}_\theta^S(\mathbf{x}, \mathbf{u}) = \mathbf{S} \begin{bmatrix} \nabla H_{\theta_H}(\mathbf{x}) \\ \mathbf{z}_{\theta_z}(\mathbf{w}) \\ \mathbf{u} \end{bmatrix} \quad (29)$$

- PHNN-JR models (20), where the coefficients of the dissipation matrix $\mathbf{R}(\mathbf{x}, \mathbf{u})$ are parameterized by the outputs a neural network $\mathbf{R}_{\theta_R}(\mathbf{x}, \mathbf{u})$, implementing the dynamic function:

$$\mathbf{f}_\theta(\mathbf{x}, \mathbf{u}) = \mathbf{f}_\theta^{JR}(\mathbf{x}, \mathbf{u}) = (\mathbf{J} - \mathbf{R}_{\theta_R}(\mathbf{x}, \mathbf{u})) \begin{bmatrix} \nabla H_{\theta_H}(\mathbf{x}) \\ \mathbf{u} \end{bmatrix} \quad (30)$$

Figure 1 shows the architecture of the two different PHNN models according to each of the PHS formulations.

4.2. Physical constraints enforcement in neural networks

According to Section 3.2, the PHS model is passive if the Hamiltonian and the dissipative terms take the form of (26) and (27) or (28). The non-zero coefficients of a $n \times n$ lower (resp. strictly lower) triangular matrix can be parameterized by the $n(n+1)/2$ (resp. $n(n-1)/2$) outputs of a neural network \mathbf{L}_{θ_L} (resp. \mathbf{K}_{θ_K}), implementing the functions:

$$\begin{cases} H_{\theta_H}(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{L}_{\theta_{L_H}}^T(\mathbf{x}) \mathbf{L}_{\theta_{L_H}}(\mathbf{x}) \mathbf{x} \\ \mathbf{z}_{\theta_z}(\mathbf{w}) = \left(\mathbf{L}_{\theta_{L_z}}^T(\mathbf{w}) \mathbf{L}_{\theta_{L_z}}(\mathbf{w}) + \mathbf{K}_{\theta_{K_z}}(\mathbf{w}) - \mathbf{K}_{\theta_{K_z}}^T(\mathbf{w}) \right) \mathbf{w} \\ \mathbf{R}_{\theta_R}(\mathbf{x}, \mathbf{u}) = \mathbf{L}_{\theta_{L_R}}^T(\mathbf{x}, \mathbf{u}) \mathbf{L}_{\theta_{L_R}}(\mathbf{x}, \mathbf{u}) \end{cases} \quad (31)$$

This parametrization is widely used in the literature [37, 39, 40, 50, 51]. Table 2 shows the fixed and learned objects for the port-Hamiltonian networks based on formulations (i) and (ii).

System	Harmonic oscillator (linear)	Duffing oscillator (nonlinear)	Self-sustained oscillator (nonlinear)
$\mathbf{x} = \begin{bmatrix} q \\ p \end{bmatrix}$ \mathbf{u}	q : centered position (elongation) [m] p : momentum [Kg.m.s ⁻¹] $u = f$: force (applied or exterior) [N]		
$H(\mathbf{x})$	$\frac{p^2}{2m} + \frac{kq^2}{2}$	$\frac{p^2}{2m} + \frac{k_1 q^2}{2} + \frac{k_3 q^4}{4}$	$\frac{p^2}{2m} + \frac{kq^2}{2}$
\mathbf{J}	$\mathbf{J} = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}$		$\mathbf{J} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$
\mathbf{R}	$\mathbf{R} = \alpha \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ with $\alpha > 0$ [N/ms ⁻¹]		$\mathbf{R} = \Gamma(w) \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 1 \end{bmatrix}$ with $\Gamma(w) > 0$ [m/sN]
\mathbf{S}	$\mathbf{S} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -1 & 0 & -1 & -1 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$		$\mathbf{S} = \begin{bmatrix} 0 & -1 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ -1 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$
$\Gamma(w)$	α		$aw^2 + bw + c$
$z(w)$	$z(w) = \alpha w$		$z(w) = aw^3 + bw^2 + cw$
$Z(w)$	$Z(w) = \frac{\alpha w^2}{2}$		$Z(w) = \frac{\alpha w^4}{4} + \frac{bw^3}{3} + \frac{cw^2}{2}$
$\mathbf{J}_f(\mathbf{x})$	$\begin{bmatrix} 0 & \frac{1}{m} \\ -k & -\frac{\alpha}{m} \end{bmatrix}$	$\begin{bmatrix} 0 & \frac{1}{m} \\ -k_1 - 3k_3 q^2 & -\frac{\alpha}{m} \end{bmatrix}$	$\begin{bmatrix} -\Gamma'(w)kq - \Gamma(w)k & -\frac{1}{m} \\ k & 0 \end{bmatrix}$
$\ \mathbf{J}_f(\mathbf{x})\ _2$	$\sqrt{\lambda_{max}((\mathbf{J}_f(\mathbf{x}))^T \mathbf{J}_f(\mathbf{x}))}$ (see the exact details on the Appendix C)		
$\kappa(\mathbf{J}_f(\mathbf{x}))$	$\frac{\sigma_{max}(\mathbf{J}_f(\mathbf{x}))}{\sigma_{min}(\mathbf{J}_f(\mathbf{x}))} = \sqrt{\frac{\lambda_{max}((\mathbf{J}_f(\mathbf{x}))^T \mathbf{J}_f(\mathbf{x}))}{\lambda_{min}((\mathbf{J}_f(\mathbf{x}))^T \mathbf{J}_f(\mathbf{x}))}}$ (see the exact details on the Appendix C)		
$\rho(\mathbf{J}_f(\mathbf{x}))$	1, if the system is underdamped		

Table 1: Different port-Hamiltonian structural elements, Jacobian matrices and related quantities for the considered oscillators.



(a) Architecture of PHNN-S

(b) Architecture of PHNN-JR

Figure 1: Architecture of the two PHNN models considered in this work. White boxes with orange contour denote fixed algebraic operations whereas orange boxes indicate the trainable parameters.

4.3. Continuous- vs. discrete-time models

A *continuous model* $\mathbf{f}_\theta(\mathbf{x}_t, \mathbf{u}_t)$ represents the field $\dot{\mathbf{x}}_t$. During training, it learns an estimate of $\dot{\mathbf{x}}_t$ by minimizing the loss

$$\mathcal{L}_c = \|\dot{\mathbf{x}}_t - \mathbf{f}_\theta(\mathbf{x}_t, \mathbf{u}_t)\|_2^2 \quad (32)$$

During inference, it is integrated as the right-hand side of the ODE using a numerical method.

A *discrete model* $\mathbf{g}_{\theta,h}(\mathbf{x}_n, \mathbf{u}_n)$ represents the next state \mathbf{x}_{n+1} . During training, it learns an estimate of \mathbf{x}_{n+1} by

minimizing the loss

$$\mathcal{L}_d = \left\| \frac{\mathbf{x}_{n+1} - \mathbf{g}_{\theta,h}(\mathbf{x}_n, \mathbf{u}_n)}{h} \right\|_1 \quad (33)$$

via backpropagation through a differentiable numerical ODE solver using a discretization step h . We introduce (33) inspired by Zhu et al. [52], as it was shown to have theoretical guarantees for NeuralODEs using explicit Runge Kutta methods (see [52], Theorem 3.1). During inference, the discrete model is applied autoregressively to predict future states. Examples of discrete PHNN model learning framework can be found in literature [37, 38, 41, 53].

Formulation	(i) $\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{y} \end{bmatrix} = (\mathbf{J} - \mathbf{R}(\mathbf{x}, \mathbf{u})) \begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{u} \end{bmatrix}$	(ii) $\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{w} \\ \mathbf{y} \end{bmatrix} = \mathbf{S} \begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{z}(\mathbf{w}) \\ \mathbf{u} \end{bmatrix}$
Objects		
$H(\mathbf{x})$	$H_{\theta_H}(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{L}_{\theta_{L_H}}(\mathbf{x})^T \mathbf{L}_{\theta_{L_H}}(\mathbf{x}) \mathbf{x}$	
Dissipation	$\mathbf{R}_{\theta_R}(\mathbf{w}) = \mathbf{L}_{\theta_{L_R}}(\mathbf{x}, \mathbf{u})^T \mathbf{L}_{\theta_{L_R}}(\mathbf{x}, \mathbf{u})$	$\mathbf{z}_{\theta_z}(\mathbf{w}) = \mathbf{\Gamma}_{\theta_\Gamma}(\mathbf{w}) \mathbf{w}$ $\mathbf{\Gamma}_{\theta_\Gamma}(\mathbf{w}) = (\mathbf{L}_{\theta_{L_z}}^T(\mathbf{w}) \mathbf{L}_{\theta_{L_z}}(\mathbf{w})) + (\mathbf{K}_{\theta_{K_z}}(\mathbf{w}) - \mathbf{K}_{\theta_{K_z}}^T(\mathbf{w}))$
Interconnection matrix	$\mathbf{J} \in \mathbb{R}^{(n_x+n_u) \times (n_x+n_u)}$	$\mathbf{S} \in \mathbb{R}^{(n_x+n_w+n_u) \times (n_x+n_w+n_u)}$

Table 2: Fixed and learned objects for the different port-Hamiltonian networks based on formulations (i) and (ii).




Figure 2: Training and inference diagram for the continuous models \mathbf{f}_θ .




Figure 3: Training and inference diagram for the discrete models \mathbf{g}_θ .

As it can be seen on Figures 2 and 3, the neural network parameterizing the continuous- and the discrete-time models is the same. However, the ODE solver is used at different phases: during training and inference for discrete-time model, but only during inference for continuous-time model. In this work, we focus on discrete models, where the two different ODE solver schemes compared are shown in Table 3. This decision is based on the fact that in practice it is generally impossible to have access to state derivatives $\dot{\mathbf{x}}_t$. Following (33), training discrete models does not require state derivative measurements and relies only on state measurements. Nevertheless, their performance during inference is limited to using the same discretization step h as during training.

4.4. Jacobian regularization on PHNN

Let $\mathbf{J}_{\mathbf{f}_\theta}(\mathbf{x}, \mathbf{u})$ denote the Jacobian of a PHNN, where $\mathbf{f}_\theta(\mathbf{x}, \mathbf{u})$ corresponds to (29)-(30). Ideally, we would like

Accuracy order	Name	Characteristic	Power-balance
2	Explicit midpoint (RK2)	Explicit	×
2	Discrete gradient (DG)	Inverse-explicit	✓

Table 3: Accuracy order, name, characteristic and whether power-balance is respected for the chosen numerical schema.

$\mathbf{f}_\theta(\mathbf{x}, \mathbf{u})$ not only to be physically-consistent, which is already structurally guaranteed in the PHNNs, but also *numerically smooth*. This means that, among the different physically-consistent solutions learned by our models, those that are ill-posed, ill-conditioned or highly stiff should be penalized through some kind of Jacobian regularization, as proposed in the NeuralODE literature [28, 30, 31]. Following this approach, we soft-constrain the training loss \mathcal{L}_d to avoid non-desirable numerical behaviors in the PHNNs.

We experiment with regularizing the spectral norm $\|\mathbf{J}_{\mathbf{f}_\theta}\|_2$,

the condition number $\kappa(\mathbf{J}_{\mathbf{f}_\theta})$ or the stiffness ratio $\rho(\mathbf{J}_{\mathbf{f}_\theta})$ as defined in Section 2.1. Penalizing higher spectral norm values promotes solutions with a lower Lipschitz constant whereas doing so with condition number rewards well-conditioned solutions. The rationale for introducing $\rho(\mathbf{J}_{\mathbf{f}_\theta})$ as a new regularization term stems from the observation that ill-conditioned Jacobians, often associated with a high NFE in adaptive numerical methods, are characteristic of stiff ODE systems. Thus, penalizing large values of $\rho(\mathbf{J}_{\mathbf{f}_\theta})$ discourages the training procedure from converging toward stiff dynamics, thereby promoting indirectly better-conditioned ODE solutions. During the training, these regularization quantities are obtained after evaluating the Jacobian at the input data from a given batch, making the same assumption as in [31]: *that regularizing dynamics at the input data can lead to regularized dynamics across the entire solution space*. The following loss functions with Jacobian regularization terms are introduced:

$$\mathcal{L}_{CN} = \mathcal{L}_d + \lambda_{CN} \|\kappa(\mathbf{J}_{\mathbf{f}_\theta}(\mathbf{x}, \mathbf{u}))\|_2^2 \quad (34)$$

$$\mathcal{L}_{SN} = \mathcal{L}_d + \lambda_{SN} \|\mathbf{J}_{\mathbf{f}_\theta}(\mathbf{x}, \mathbf{u})\|_2 \quad (35)$$

$$\mathcal{L}_{SR} = \mathcal{L}_d + \lambda_{SR} \|\rho(\mathbf{J}_{\mathbf{f}_\theta}(\mathbf{x}, \mathbf{u})) - 1\|_2^2 \quad (36)$$

We set $\lambda_{SN} = 10^{-6}$, $\lambda_{CN} = 10^{-6}$ and $\lambda_{SR} = 10^{-4}$ in later experiments, and, as in [31], we add 10^{-6} to the denominator of $\kappa(\mathbf{J}_{\mathbf{f}_\theta})$ and $\rho(\mathbf{J}_{\mathbf{f}_\theta})$ to avoid underflow in the early stages of training.

5. Experiments

5.1. Our questions

Given a discrete trajectory \mathcal{T} and its initial state \mathbf{x}_0 , the objective of this work is to design PHNN models capable of accurately generating a trajectory $\tilde{\mathcal{T}} = \{\tilde{\mathbf{x}}_n\}_{n \in \mathbb{N}^+}$, with $\tilde{\mathbf{x}}_0 = \mathbf{x}_0$, that approximates \mathcal{T} as closely as possible. Each of the PHNN models presented in this work can be characterized by the:

1. PHS formulation: PHNN-S or PHNN-JR.
2. Type of model: Continuous or discrete.
3. Numerical method: RK2 or DG (see Table 3).
4. Number of trainable parameters: 800, 2k or 10k.
5. Number of training points: 25, 100 or 400.
6. Loss function: \mathcal{L}_d , \mathcal{L}_{SN} , \mathcal{L}_{CN} or \mathcal{L}_{SR} .

We set a *baseline* NODE [27] model and conduct several experiments to compare its performance to the proposed family of PHNNs models.

Study I: Impact of the number of training points.

The first study focuses on how the different PHNN architectures compare at inference stage with different numerical methods for the systems in Table 1. For this study, experiments are carried out for 25, 100 and 400 training points and the smallest number of training parameters is considered (≈ 800).

Study II: Impact of the number of trainable parameters. The second study uses the same combination of PHNN architectures and numerical methods but, opposite to Study I, the number of training points is fixed to 25 and the number of training parameters varies between *small* (≈ 800), *medium* ($\approx 4k$) and *large* ($\approx 20k$). Tables E.8-E.10 in Appendix D detail the design choices of the neural network components for each model. In each case, the criteria is to have a similar number of parameters for each formulation, which enables a fair comparison between them.

Study III: Impact of the Jacobian regularizations.

The third study focuses on how the different Jacobian regularizations in (34)-(36) influence the inference performance of the models. In this case, the same combination of PHNN architectures, numerical methods, training parameters and training points as in Study I is considered.

5.2. Implementation details

Table 4 shows the parameters considered for the generation of trajectories for each physical system as well as the training and inference parameters considered for the neural network experiments, whereas Figure 4 shows the distribution of the training and test points as well as two complete test trajectories for each oscillatory system.

Generation of synthetic data. The dataset \mathcal{T} consists of $N_{traj} = 12500$ trajectories. Each trajectory is generated synthetically according to the following criteria:

- i **Initial condition:** We fix E_{min}, E_{max} , and we sample an initial condition \mathbf{x}_0 such that

$$H(\mathbf{x}_0) \in \mathcal{I}_{E_0} = [E_{min}, E_{max}] \quad (37)$$

with $E_{min}, E_{max} \in \mathbb{R}^+$ (see Table 4 for numerical values).

- ii **Control:** For the harmonic and Duffing oscillators, external control is applied as a constant force \mathbf{u} so that

$$H(\mathbf{x}^*) \in \mathcal{I}_{E_{eq}} = [E_{eq}^{min}, E_{eq}^{max}] \quad (38)$$

where \mathbf{x}^* is the equilibrium point of the system. In control theory, this is usually referred to as *potential energy shaping* [54]. For the self-sustained oscillator, the external constant control \mathbf{u} is applied so that the system stabilizes in a limit cycle around \mathbf{x}^* (see Table 4 for numerical values).

Further details about the sampling of the initial conditions and the control design can be found in Appendix E. As for the intrinsic parameters of each system, they are chosen to satisfy that the natural frequency f_0 is $1Hz$. For the harmonic and Duffing oscillator, a linear dissipation function $z(w) = cw$ is considered. In this case, c is chosen so that the damped harmonic oscillator has dissipated 99% of its energy in $D = 5T_0 = 5s$, where $T_0 = 1/f_0$

System	Harmonic Oscillator	Duffing Oscillator	Self-sustained oscillator
Intrinsic parameters	(m, k)	(m, k_1, k_3)	(m, k)
f_0 [Hz]		1 Hz	
m [kg]		0.16	
$1/k, 1/k_1$ [N/m]		0.16	
k_3 [N/m ³]	-	100k ₁	-
$z(w)$		0.9w [N]	1.3w ³ - 4w ² + 3w [m/s]
Constant control u [N]	Such that the equilibrium point energy lies in \mathcal{I}_{E_0}		Such that system reaches a limit cycle
$\mathcal{I}_{E_0}, \mathcal{I}_{E_{eq}}$ [J]			[0.1, 1]
sr_{gen} [Hz]			400f ₀
Generation numerical method	Gonzalez discrete gradient (15)		
α			0.31
β			5
D_{train} [s]			αT_0
D_{infer} [s]			βT_0
sr_{train}, sr_{infer} [Hz]			100f ₀
N_{train}			[25, 100, 400]
Batch size			64
N_{eval}			2500
N_{infer}			100

Table 4: Implementation hyperparameters for the different experiments.

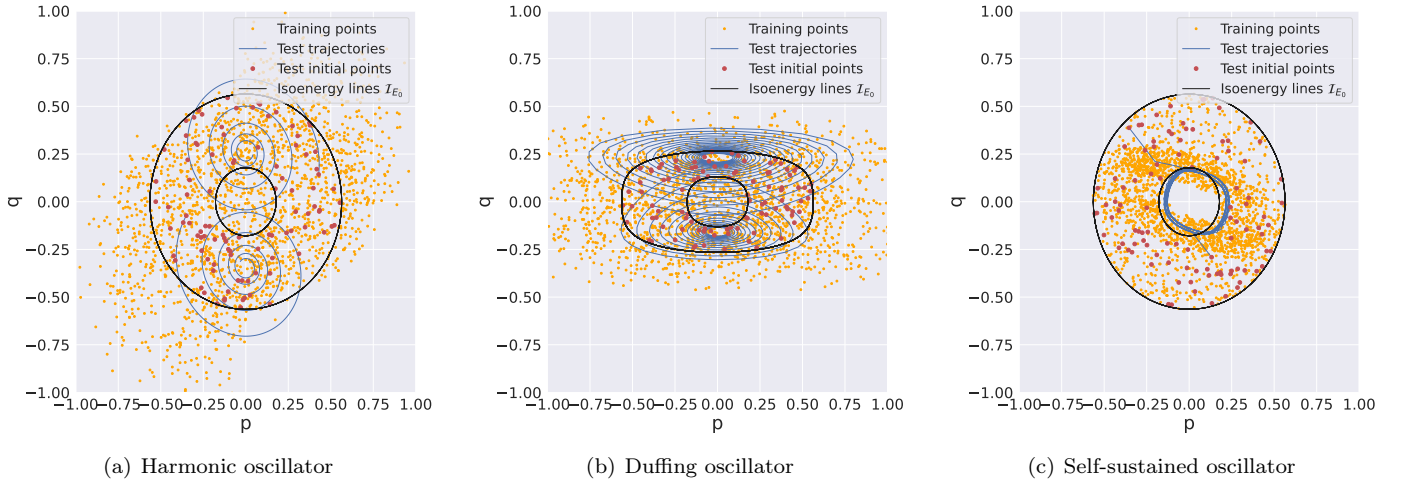


Figure 4: Training points, test initial points and two complete test trajectories for each of the three oscillatory systems. Note that in the case of the harmonic and Duffing oscillator, the applied control shifted the equilibrium point from $(p, q) = (0, 0)$ whereas in the case of the self-sustained oscillator, it stabilizes the trajectories in a limit cycle.

is the natural period. The same value of c is considered for the Duffing oscillator. As for the self-sustained oscillator, the considered nonlinear dissipation function is $z(w) = 1.3w^3 - 4w^2 + 3w$. In each physical system, the numerical scheme considered for the generation of the trajectories is the *Gonzalez discrete gradient* (15) for a duration D and a sampling rate $sr_{gen} = 400f_0$.

Training details. Let $D_{train} = \alpha T_0$ and $sr_{train} = \gamma f_0 < sr_{gen}$ be the training time and data sampling rate, respectively, where α and γ are hyperparameters. For each system, we consider α such that the harmonic os-

cillator has dissipated 25% of its initial energy in the absence of control and $\gamma = 100$. The set of training trajectories \mathcal{T}_{train} is built considering N_{train} trajectories from \mathcal{T} up until a duration T_{train} sub-sampled at sr_{train} . Once \mathcal{T}_{train} is created, one point $\xi_{train}^i = ((\mathbf{x}_n^i, \mathbf{u}_n^i), \mathbf{x}_{n+1}^i)$ is uniformly sampled from each trajectory $\tau_{train}^i \in \mathcal{T}_{train}$ (see Figure 5 for a graphical description). The set of points $\xi_{train} = \{\xi_{train}^i\}_{i=1, \dots, N_{train}}$ constitutes the training dataset. Note that training on a set of isolated points sampled at random from complete trajectories is closer to experimental conditions than training on a set of complete trajectories that might be more difficult to measure accu-




Figure 5: Schematic sampling and dataset construction procedure. A trajectory generated at sampling frequency sr_{gen} over a duration $D = D_{infer} = \beta T_0$ is shown as white dot markers. From this trajectory, a training point, highlighted with a green star marker, is uniformly sampled from a subset of samples, shown as black dot markers, obtained at frequency sr_{train} and restricted to $t \leq \alpha T_0$. The training horizon D_{train} and the inference horizon D_{infer} are indicated by arrows, with the vertical dashed line marking the end of the training interval.

rately in practice. It also probably adds complexity for the model, as isolated points are not obviously correlated. Experiments are carried out for $N_{train} \in [25, 100, 400]$. After each training epoch, models are validated on another dataset ξ_{eval} , which is constructed as ξ_{train} but for 2500 trajectories. For each of the datasets, experiments are carried out for 10 runs with different model initializations using the Adam optimizer [55] with a batch size 64 for $50k$ optimizer steps with a learning rate of 10^{-3} .

Inference details. The inference dataset \mathcal{T}_{infer} is constructed taking a subset of $N_{infer} = 100$ trajectories from \mathcal{T} for a duration $D_{infer} = \beta T_0$ and a sampling rate $sr_{infer} = sr_{train}$. Note that $\beta \gg \alpha$, i.e. the model has to generate in inference longer sequences that it has been trained for.

The performance of the model is then assessed generating autoregressively a trajectory $\tilde{\tau}$ from each of the initial conditions of the trajectories in \mathcal{T}_{infer} and comparing it with the reference trajectory τ . For each model initialization, we compute the Mean Squared Error (MSE) between the reference and predicted trajectories:

$$\mathcal{L}_{traj} = \frac{1}{N_{inf} \cdot T_{inf} \cdot f_{inf}} \sum_{i=1}^{N_{inf}} \|\tau_i - \tilde{\tau}_i\|^2 \quad (39)$$

6. Results and discussions

6.1. Impact of the number of training points

Figure 6 shows the inference error (see Table F.11 in the Appendix F for the detailed numerical values) for the

different models across the three values of N_{train} when learning the three controlled oscillatory systems with small model size. As expected, increasing the number of points decreases the error for all systems and architectures, except for when we train the PHNN-S model on the self-sustained oscillator where the performance seems to be fairly constant. The NODE architecture is consistently outperformed by one of the physically constrained models. Regarding the differences between the two numerical methods, the discrete gradient outperforms the RK2 in the low-data regime for the harmonic oscillator, and for every number of training points in the case of the nonlinear oscillators. For the Duffing oscillator, all performance errors show a high dispersion (as shown by large IQR) in the low-data regime, reflecting the model sensitivity to weight initialization when very few training points are available. Finally, whereas the PHNN-S model achieves the lowest performance error for the harmonic and Duffing oscillator, it is not the case for the self-sustained oscillator, where it is outperformed by the PHNN-JR and the NODE. Figure 7 compares the NODE with the best PHNN architecture for each oscillator when discretizing a trajectory starting on an initial condition \mathbf{x}_0 such that $H(\mathbf{x}_0) = 0.5J$.

6.2. Impact of the number of trainable parameters

Figure 8 shows the inference error (see Table F.12 in the Appendix F for the detailed numerical values) for the different models across the three regimes (small, medium, large) of trainable parameters when learning the three controlled oscillatory systems with $N_{train} = 25$. In this case, increasing the number of trainable parameters increases the error in general for all the systems and architectures. This can be explained by the fact that there are not enough

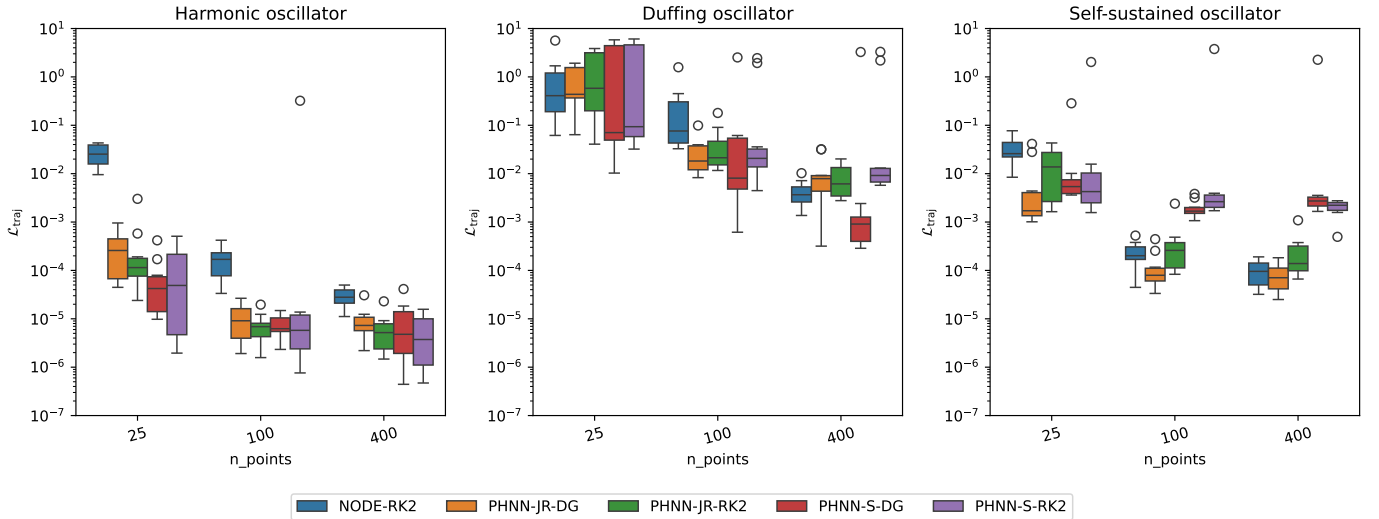


Figure 6: Boxplot of the inference errors for the three different oscillators and varying numbers of training points (small models). From left to right: harmonic, Duffing and self-sustained oscillators.

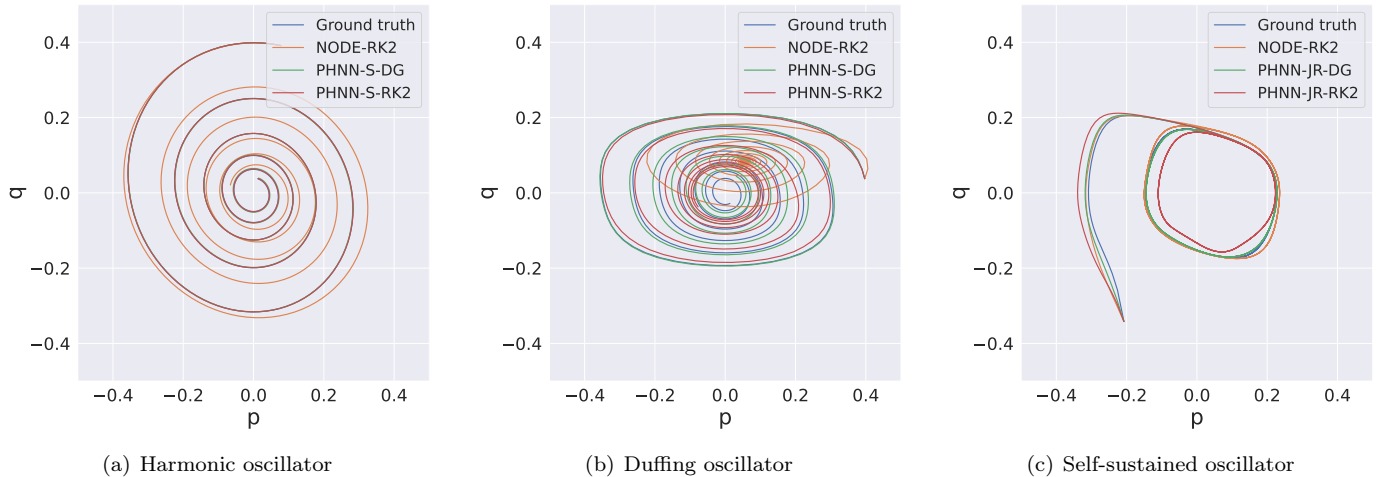


Figure 7: Comparison of the learned trajectory discretizations for each oscillatory system starting on \mathbf{x}_0 such that $H(\mathbf{x}_0) = 0.5J$. **(Left)** NODE and PHNN-S models trained on $N_{train} = 25$ from the harmonic oscillator. **(Center)** NODE and PHNN-S models trained on $N_{train} = 100$ from the Duffing oscillator. **(Right)** NODE and PHNN-JR models trained on $N_{train} = 25$ from the self-sustained oscillator. For the harmonic and Duffing oscillator, the trajectory is obtained with $\mathbf{u} = 0$; and for the self-sustained, with \mathbf{u} such that the system stabilizes in a limit cycle.

training samples to fit larger models and, thus, medium and large size models directly overfit and are not able to generalize correctly. The NODE is again consistently outperformed by one of the physically constrained models. Regarding the impact of using the different numerical methods, the discrete gradient only outperforms the RK2 when training small models. Interestingly, increasing the number of trainable parameters has a positive effect on the dispersion of the results for the Duffing oscillator, as we observe how the IQR is reduced with respect to the small size setting. Figure 9 shows how the trajectory discretization from the best PHNN architecture for each oscillator changes with respect to the number of trainable parameters starting on an initial condition \mathbf{x}_0 such that

$$H(\mathbf{x}_0) = 0.5J.$$

6.3. Impact of the Jacobian regularizations

Regularizing the Jacobian of the PHNN models is motivated by the potential correlation existing between the inference dispersion and a high condition number, spectral norm or stiffness ratio mean value at the end of the training. As we observe in Figure 10, the models with the highest mean condition number, spectral norm and stiffness ratio when modeling the Duffing oscillator with $N_{train} = 25$, i.e. PHNN-S-RK2 and PHNN-S-DG, are precisely the ones with a higher IQR value (see Figure 6 and Table F.11 in Appendix F). Furthermore, their stiffness ratio is higher than 1, meaning that these models are

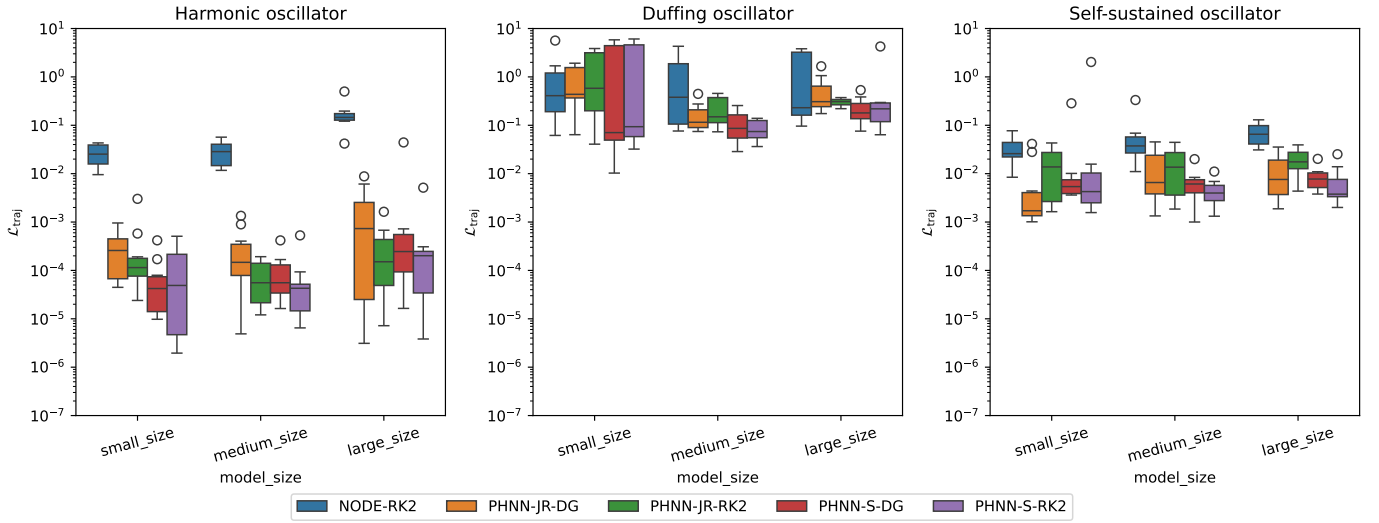


Figure 8: Boxplot of the inference errors for the three different oscillators and model sizes ($N_{train} = 25$). From left to right: harmonic, Duffing and self-sustained oscillators.

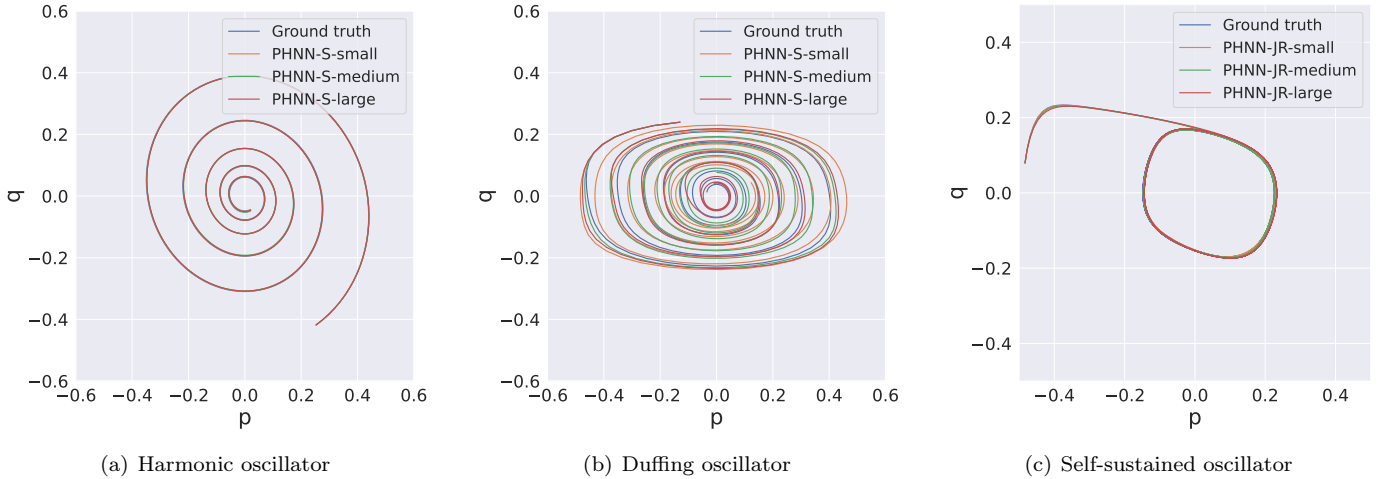


Figure 9: Comparison of the learned trajectory discretizations by the discrete gradient for each oscillatory system starting on \mathbf{x}_0 such that $H(\mathbf{x}_0) = 0.75J$. **(Left)** PHNN-S models trained on $N_{train} = 25$ from the harmonic oscillator. **(Center)** PHNN-S models trained on $N_{train} = 25$ from the Duffing oscillator. **(Right)** PHNN-JR models trained on $N_{train} = 25$ from the self-sustained oscillator. For the harmonic and Duffing oscillator, the trajectory is obtained with $\mathbf{u} = 0$; and for the self-sustained, with \mathbf{u} such that the system stabilizes in a limit cycle.

learning a representation that is stiffer than the original system from which we generate the data. Experiments are, thus, carried out to test whether controlling the numerical behavior through any of the proposed Jacobian regularizations lowers the inference error and/or dispersion.

Figure 11 shows the inference error (see Table F.13 in the Appendix F for the detailed numerical values) when learning the three oscillatory systems under the different Jacobian regularizations (34)-(36) with different small models and $N_{train} = 25$. BL refers to the baseline (i.e. no Jacobian regularization); CN, to the condition number regularization (34); SN, to the spectral norm regularization (35); and SR, to the stiffness ratio regularization (36). In

general, Jacobian regularization does not improve the results for those systems in which the model performances were already good (harmonic and self-sustained oscillator). However, CN-regularization considerably improves the results for the Duffing oscillator, where the dispersion was high. Interestingly, controlling the stiffness ratio towards the ground truth system value is not directly associated with an improvement in the accuracy nor the dispersion: the stiffness ratio values after SR-regularization are closer to 1 (see Figure 12) but the inference error or the dispersion is not always reduced.

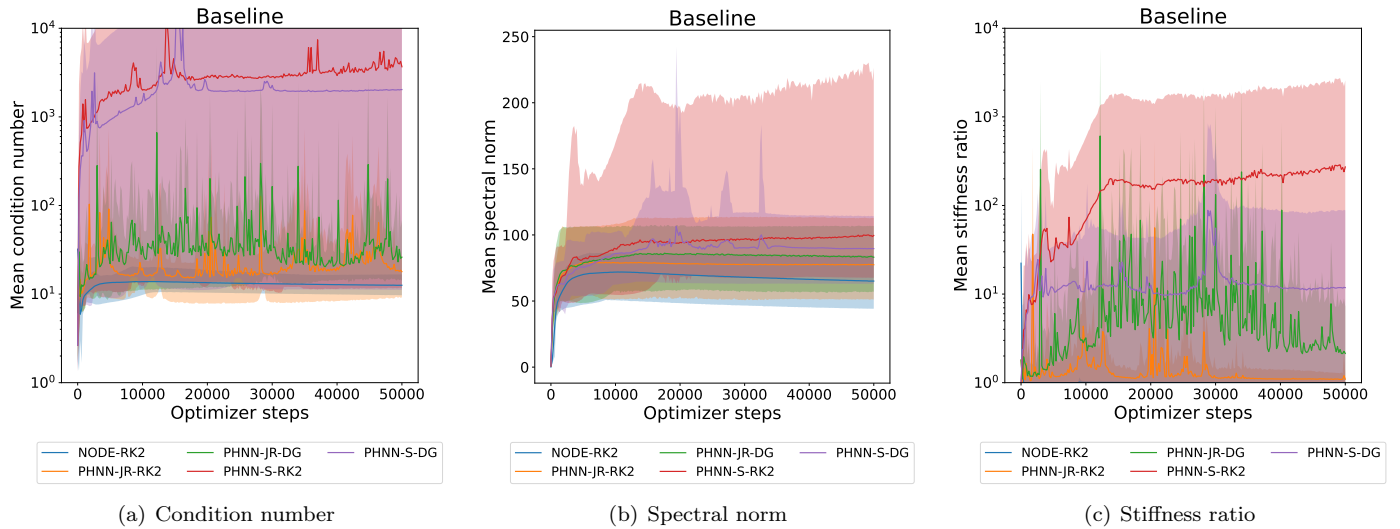


Figure 10: Mean condition number (**left**), spectral norm (**center**) and stiffness ratio (**right**) per optimizer step during the training for the different architectures when modeling the Duffing oscillator with $N_{train} = 25$ points and no Jacobian regularization (BL). For each optimizer step, the corresponding mean value (solid line) is obtained computing the average over the 10 model initializations, with the shaded area indicating the range between the minimum and maximum values.

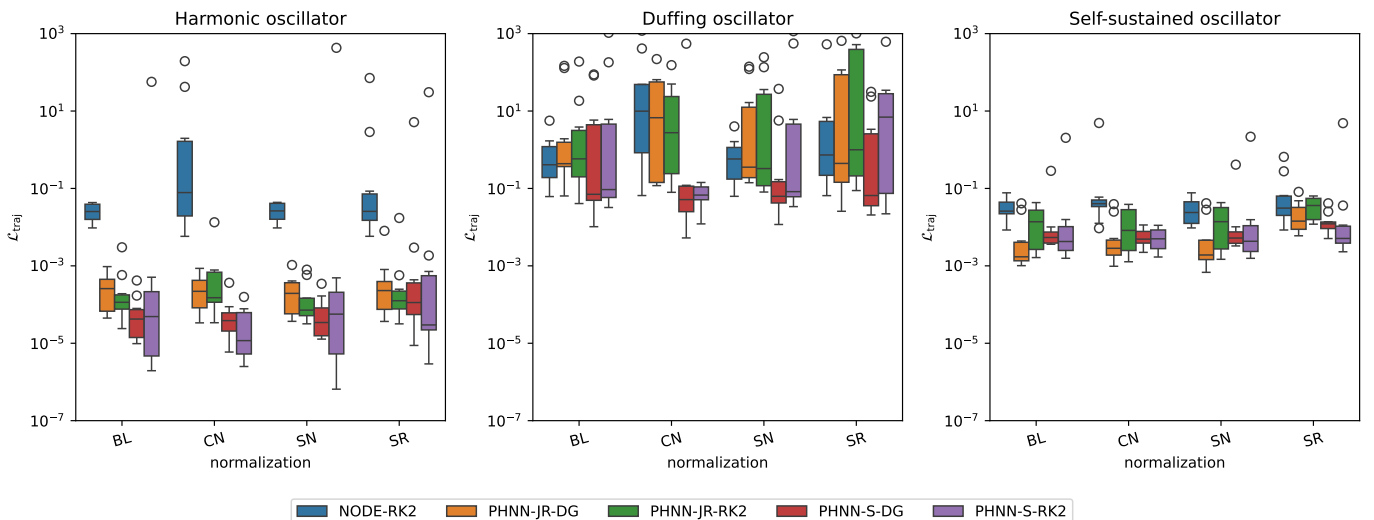


Figure 11: Boxplot of the inference errors for the three different oscillators and normalizations ($N_{train} = 25$ and small number of parameters). Notation: BL refers to the baseline (i.e. no Jacobian regularization); CN, to the condition number regularization (34); SN, to the spectral norm regularization (35); and SR, to the stiffness ratio regularization (36). From left to right: harmonic, Duffing and self-sustained oscillators.

6.4. Differences between PHNN-S and PHNN-JR

It is a recurrent conclusion in all the studies carried out in this work that the architecture that performs the best for the harmonic and Duffing oscillator is the PHNN-S, while for the self-sustained oscillator, is the PHNN-JR, independently of the numerical method used for discretization. We hypothesize that this could be due to whether the dissipation of the system is nonlinear or not, implying that PHNN-S has difficulty learning the system's dynamics in the presence of nonlinear dissipation, which is the case for the self-sustained oscillator.

7. Conclusion

Contributions. In this work, we have addressed the problem of designing *physically consistent* neural network models based on PHS formulations and energy-preserving numerical methods. The main contributions of this study include a comparison of two theoretically equivalent PHS formulations: the PH-DAE and the input-state-output PHS with feedthrough, when implemented as PHNNs; a performance comparison between a second-order energy-preserving numerical method and a Runge-Kutta method of the same order; and an empirical study of the impact of regularizing the Jacobian of PHNN through two methods already ap-

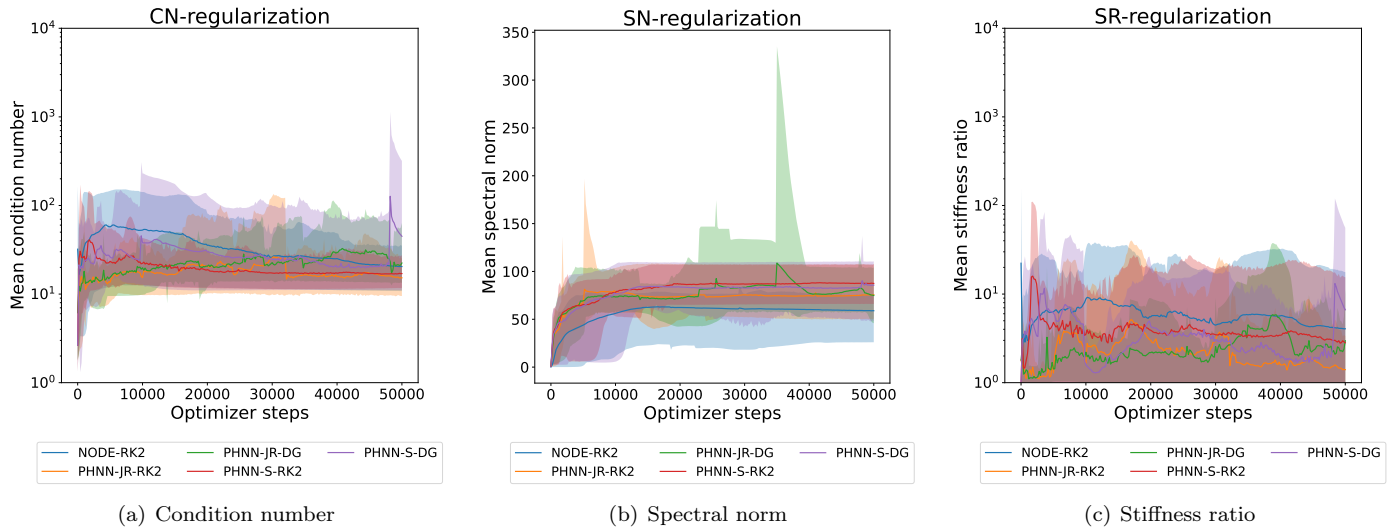


Figure 12: Impact on the mean condition number (**left**), spectral norm (**center**) and stiffness ratio (**right**) per optimizer step during the training for the different architectures when modeling the Duffing oscillator with $N_{train} = 25$ points and the proposed Jacobian regularizations. For each optimizer step, the corresponding mean value (solid line) is obtained computing the average over the 10 model initializations, with the shaded area indicating the range between the minimum and maximum values.

plied to NODEs and a new one tackling the stiffness of the learned ODE solutions. Our results demonstrate that the PHNN based on the discrete gradient method outperform those based on RK2, especially when the data comes from nonlinear systems and is learned by models with a small number of trainable parameters. The results also indicate that the best PHS formulation to be implemented as a neural network depends on the modeled system and where the nonlinearity operates. Interestingly, this dependence is also found in the results when Jacobian regularization is applied, where the harmonic and Duffing oscillator clearly benefited from these additional soft constraints while the self-sustained oscillator did not. Overall, the experiments on Jacobian regularization show how this technique could be used in PHNNs to improve the generalization when increasing the number of training points is not possible.

Limitations. While the proposed frameworks have shown promising performance, they are subject to certain limitations, including the fact the interconnection matrix (\mathbf{S} or \mathbf{J} , depending on the formulation) is always given a priori, and that we are only considering Hamiltonians of the form (22) and semi-explicit PHS formulations, which are not sufficiently general to describe any given physical system [16]. Numerically, our results are also limited by the fact that in our experiments we use second-order numerical methods, as the objective is to compare two classes of schemes (energy preserving versus non-preserving), rather than to achieve the best possible performance. As for the Jacobian regularizations, they are limited by the hypothesis that regularizing dynamics at the input data can lead to regularized dynamics across the entire solution space, which could not necessarily be the case when working with the nonlinear dynamics in-

duced by neural network-based models. Furthermore, the real impact of each Jacobian regularization is not fully explored in this study, as the value of λ_{CN} , λ_{SN} and λ_{SR} in (34)-(36) are chosen equally for each physical system and architecture, without an exhaustive grid search for each case.

Future works. Future work will focus on designing and testing learning frameworks capable of modelling PHS without knowing the interconnection matrix a priori and from which we only have access to its inputs and outputs. Of particular interest will be understanding the reasons why the performance is dependent on the PHS formulation and the modeled physical system, as well as detecting the factors that hinder the training of these models and how this relates to a well-conditioned and non-stiff learned Jacobian matrix. Whether the same conclusions hold when using a fourth-order discrete gradient and Runge-Kutta method remains to be explored. Finally, this framework could be used to model more complex ODE systems or even extended to tackle the learning of Hamiltonian PDEs, where power-preserving spatial discretizations should also be considered [56].

Overall, this paper provides a first comparison between the use of Runge-Kutta and discrete gradient methods in the PHNN framework as well as an extensive empirical study on the impact that different PHS formulations, number of training points, size of the models or Jacobian regularizations have on the modeling of a given physical system. The authors also hope that this type of work serves as a baseline for future methods and thus contribute to addressing the lack of homogeneity that characterizes similar work in the literature.

Acknowledgements. This project is co-funded by the European Union’s Horizon Europe research and innovation program Cofund SOUND.AI under the Marie Skłodowska-Curie Grant Agreement No 101081674. This project was also provided with computing HPC and storage resources by GENCI at IDRIS thanks to the grant 2024-102911 on the supercomputer Jean Zay’s V100 partition.

References

- [1] A. Cicirello, Physics-enhanced machine learning: a position paper for dynamical systems investigations, *Journal of Physics: Conference Series* 2909 (1) (2024) 012034. doi:10.1088/1742-6596/2909/1/012034.
- [2] J. Baxter, A model of inductive bias learning, *Journal of Artificial Intelligence Research* 12 (2000) 149–198. doi:10.1613/jair.731.
- [3] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, *Nature Reviews Physics* 3 (6) (2021) 422–440. doi:10.1038/s42254-021-00314-5.
- [4] S. Eidnes, A. J. Stasik, C. Sterud, E. Bøhn, S. Riemer-Sørensen, Pseudo-hamiltonian neural networks with state-dependent external forces, *Physica D: Nonlinear Phenomena* 446 (2023) 133673. doi:10.1016/j.physd.2023.133673.
- [5] M. Raissi, P. Perdikaris, 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* 378 (2019) 686–707. doi:10.1016/j.jcp.2018.10.045.
- [6] C. Meng, S. Griesemer, D. Cao, S. Seo, Y. Liu, When physics meets machine learning: A survey of physics-informed machine learning, *Machine Learning for Computational Science and Engineering* 1 (2025) 20. doi:10.1007/s44379-025-00016-0.
- [7] H. K. Khalil, *Nonlinear systems*, 3rd Edition, Prentice Hall, Upper Saddle River, N.J., 2002.
- [8] J. Hadamard, Sur les problèmes aux dérivées partielles et leur signification physique, *Princeton university bulletin* (1902) 49–52.
- [9] M. W. Hirsch, R. L. Devaney, S. Smale, *Differential equations, dynamical systems, and linear algebra*, 1st Edition, Vol. 60, Academic press, 1974.
- [10] G. Golub, C. Van Loan, *Matrix Computations*, 4th Edition, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, 2013.
- [11] L. N. Trefethen, D. Bau, III, *Numerical Linear Algebra*, 1st Edition, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997. doi:10.1137/1.9780898719574.
- [12] A. Iserles, *A First Course in the Numerical Analysis of Differential Equations*, 2nd Edition, A First Course in the Numerical Analysis of Differential Equations, Cambridge University Press, 2009.
- [13] J. R. Taylor, *Classical mechanics*, 1st Edition, Calif.:University Science Books, 2005.
- [14] B. Maschke, A. Van Der Schaft, P. Breedveld, An intrinsic hamiltonian formulation of network dynamics: non-standard poisson structures and gyrators, *Journal of the Franklin Institute* 329 (1992) 923–966. doi:10.1016/S0016-0032(92)90049-M.
- [15] V. Duindam, A. Macchelli, S. Stramigioli, H. Bruyninckx, *Modeling and control of complex physical systems: the port-Hamiltonian approach*, 1st Edition, Springer Science & Business Media, 2009.
- [16] A. Van Der Schaft, D. Jeltsema, et al., *Port-hamiltonian systems theory: An introductory overview*, *Foundations and Trends® in Systems and Control* 1 (2-3) (2014) 173–378.
- [17] A. Chaigne, J. Kergomard, *Acoustics of musical instruments*, 1st Edition, Springer, 2016.
- [18] S. Aoues, F. L. Cardoso-Ribeiro, D. Matignon, D. Alazard, Modeling and control of a rotating flexible spacecraft: A port-hamiltonian approach, *IEEE Transactions on Control Systems Technology* 27 (1) (2017) 355–362. doi:10.1109/TCST.2017.2771244.
- [19] T. Hélie, *Elementary tools on Port-Hamiltonian Systems with applications to audio/acoustics*, lecture (Mar. 2022). doi:hal-03986168.
- [20] F. L. Cardoso-Ribeiro, G. Haine, Y. Le Gorrec, D. Matignon, H. Ramirez, Port-hamiltonian formulations for the modeling, simulation and control of fluids, *Computers & Fluids* (2024) 106407doi:10.1016/j.compfluid.2024.106407.
- [21] D. Roze, T. Hélie, E. Rouhaud, Time-space formulation of a conservative string subject to finite transformations, *IFAC-PapersOnLine* 58 (6) (2024) 232–237. doi:10.1016/j.ifacol.2024.08.286.
- [22] E. Hairer, C. Lubich, G. Wanner, *Geometric numerical integration*, 2nd Edition, Vol. 31 of Springer Series in Computational Mathematics, Springer-Verlag, Berlin, 2006.
- [23] E. Celledoni, E. H. Høiseth, Energy-preserving and passivity-consistent numerical discretization of port-hamiltonian systems, arXiv preprint arXiv:1706.08621 (2017).

- [24] G. Quispel, G. S. Turner, Discrete gradient methods for solving odes numerically while preserving a first integral, *Journal of Physics A: Mathematical and General* 29 (13) (1996) L341. doi:10.1088/0305-4470/29/13/006.
- [25] O. Gonzalez, Time integration and discrete hamiltonian systems, *Journal of Nonlinear Science* 6 (1996) 449–467. doi:10.1007/BF02440162.
- [26] E. Celledoni, S. Eidnes, H. N. Myhr, Learning dynamical systems from noisy data with inverse-explicit integrators, *Physica D: Nonlinear Phenomena* 472 (2025) 134471. doi:10.1016/j.physd.2024.134471.
- [27] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations, in: *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, Curran Associates Inc., Red Hook, NY, USA, 2018, p. 6572–6583.
- [28] E. Dupont, A. Doucet, Y. W. Teh, Augmented neural odes, in: *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA, 2019, pp. 3140–3150.
- [29] L. Deng, The mnist database of handwritten digit images for machine learning research, *IEEE Signal Processing Magazine* 29 (6) (2012) 141–142. doi:10.1109/MSP.2012.2211477.
- [30] C. Finlay, J.-H. Jacobsen, L. Nurbekyan, A. M. Oberman, How to train your neural ode: the world of jacobian and kinetic regularization, in: *Proceedings of the 37th International Conference on Machine Learning, ICML’20*, JMLR.org, 2020, pp. 3154–3164.
- [31] S. Josias, W. Brink, Jacobian norm regularisation and conditioning in neural odes, in: *Artificial Intelligence Research*, Springer Nature Switzerland, Cham, 2022, pp. 31–45.
- [32] Y. Yoshida, T. Miyato, Spectral norm regularization for improving the generalizability of deep learning, arXiv preprint arXiv:1705.10941 (2017).
- [33] M. Takeru, K. Toshiki, K. Masanori, Y. Yuichi, Spectral normalization for generative adversarial networks, in: *International Conference on Learning Representations*, 2018, pp. 1–26.
- [34] S. Greydanus, M. Dzamba, J. Yosinski, Hamiltonian neural networks, *Advances in neural information processing systems* 32 (2019).
- [35] A. Sosanya, S. Greydanus, Dissipative hamiltonian neural networks: Learning dissipative and conservative dynamics separately, arXiv preprint arXiv:2201.10085 (2022).
- [36] S. A. Desai, M. Mattheakis, D. Sondak, P. Protopapas, S. J. Roberts, Port-hamiltonian neural networks for learning explicit time-dependent dynamical systems, *Physical Review E* 104 (3) (2021) 034312.
- [37] Y. D. Zhong, B. Dey, A. Chakraborty, Symplectic ode-net: Learning hamiltonian dynamics with control, arXiv preprint arXiv:1909.12077 (2019).
- [38] Y. D. Zhong, B. Dey, A. Chakraborty, Dissipative symoden: Encoding hamiltonian dynamics with dissipation and control into deep learning, arXiv preprint arXiv:2002.08860 (2020).
- [39] K. Cherifi, A. E. Messaoudi, H. Gernandt, M. Roschkowski, Nonlinear port-hamiltonian system identification from input-state-output data, arXiv preprint arXiv:2501.06118 (2025).
- [40] F. J. Roth, D. K. Klein, M. Kannapinn, J. Peters, O. Weeger, Stable port-hamiltonian neural networks, arXiv preprint arXiv:2502.02480 (2025).
- [41] Z. Chen, J. Zhang, M. Arjovsky, L. Bottou, Symplectic recurrent neural networks, arXiv preprint arXiv:1909.13334 (2019).
- [42] A. Zhu, P. Jin, Y. Tang, Deep hamiltonian networks based on symplectic integrators, arXiv preprint arXiv:2004.13830 (2020).
- [43] D. DiPietro, S. Xiong, B. Zhu, Sparse symplectically integrated neural networks, *Advances in Neural Information Processing Systems* 33 (2020) 6074–6085.
- [44] S. Xiong, Y. Tong, X. He, S. Yang, C. Yang, B. Zhu, Nonseparable symplectic neural networks, arXiv preprint arXiv:2010.12636 (2020).
- [45] H. Choudhary, C. Gupta, V. Kungurtsev, M. Leok, G. Korpas, Learning generalized hamiltonians using fully symplectic mappings, arXiv preprint arXiv:2409.11138 (2024).
- [46] A. Van der Schaft, *L2-gain and passivity techniques in nonlinear control*, 2nd Edition, Springer, 2000.
- [47] R. Muller, T. Hélie, Power-balanced modelling of circuits as skew gradient systems, in: *21 st International Conference on Digital Audio Effects (DAFx-18)*, 2018, pp. 1–8.
- [48] W. Press, S. Teukolsky, W. Vetterling, B. Flannery, *Numerical Recipes: The art of Scientific Computing*, Thrid Edition in C++, 3rd Edition, Cambridge University Press, 2007.
- [49] T. Hélie, M. Linares, G. Doras, Modèle passif minimal d’instrument musical auto-oscillant à configuration variable en temps, in: *CFA 2025 - 17e Congrès Français d’Acoustique*, Paris, France, 2025, pp. 1–21. doi:hal-05228704.

- [50] P. Schwerdtner, Port-hamiltonian system identification from noisy frequency response data, arXiv preprint arXiv:2106.11355 (2021).
- [51] P. Schwerdtner, T. Moser, V. Mehrmann, M. Voigt, Structure-preserving model order reduction for index one port-hamiltonian descriptor systems, arXiv preprint arXiv:2206.01608 (2022).
- [52] A. Zhu, P. Jin, B. Zhu, Y. Tang, On numerical integration in neural ordinary differential equations, in: International Conference on Machine Learning, PMLR, 2022, pp. 27527–27547.
- [53] C. Neary, U. Topcu, Compositional learning of dynamical system models using port-hamiltonian neural networks, in: Learning for Dynamics and Control Conference, PMLR, 2023, pp. 679–691.
- [54] R. Ortega, A. J. van der Schaft, I. Mareels, B. Maschke, Energy shaping control revisited, in: Advances in the control of nonlinear systems, Springer, 2007, pp. 277–307.
- [55] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).
- [56] V. Trenchant, H. Ramirez, Y. Le Gorrec, P. Kotyczka, Finite differences on staggered grids preserving the port-hamiltonian structure with application to an acoustic duct, Journal of Computational Physics 373 (2018) 673–697. doi:10.1016/j.jcp.2018.06.051.
- [57] S. H. Strogatz, Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering, 3rd Edition, CRC press, 2024.

Appendix A. Well-posedness

Theorem 1. *If $\|\mathbf{J}_f(\mathbf{x})\|_2 \leq K < \infty, \forall \mathbf{x} \in \mathcal{D}$, then \mathbf{f} is K -Lipschitz.*

Proof. Fix $\mathbf{x}, \mathbf{y} \in \mathcal{D}$. Define the segment between them like $\gamma(t) = \mathbf{x} + t(\mathbf{y} - \mathbf{x})$ for $t \in [0, 1]$. Let $\mathbf{g}(t) = \mathbf{f}(\gamma(t))$. By the chain rule, $\mathbf{g}'(t) = \mathbf{J}_f(\gamma(t))(\mathbf{y} - \mathbf{x})$. Using the fundamental theorem of calculus:

$$\mathbf{f}(\mathbf{y}) - \mathbf{f}(\mathbf{x}) = \int_0^1 \mathbf{J}_f(\gamma(t))(\mathbf{y} - \mathbf{x}) dt$$

Considering norms,

$$\|\mathbf{y} - \mathbf{x}\|_2 \leq \int_0^1 \|\mathbf{J}_f(\gamma(t))(\mathbf{y} - \mathbf{x})\|_2 dt$$

By operator norm inequality

$$\|\mathbf{J}_f(\gamma(t))(\mathbf{y} - \mathbf{x})\|_2 \leq \|\mathbf{J}_f(\gamma(t))\|_2 \|\mathbf{y} - \mathbf{x}\|_2$$

So

$$\|\mathbf{f}(\mathbf{y}) - \mathbf{f}(\mathbf{x})\|_2 \leq \int_0^1 \|\mathbf{J}_f(\gamma(t))\|_2 \|\mathbf{y} - \mathbf{x}\|_2 dt$$

From which, using $\|\mathbf{J}_f(\mathbf{x})\|_2 \leq K$,

$$\|\mathbf{f}(\mathbf{y}) - \mathbf{f}(\mathbf{x})\|_2 \leq K \|\mathbf{y} - \mathbf{x}\|_2$$

□

Following the proof of Theorem 1, it is easy to see that $K = \sup_{\mathbf{x} \in \mathcal{D}} \|\mathbf{J}_f(\mathbf{x})\|_2$ is a Lipschitz constant.

Appendix B. Passivity

The concept of *passivity* is related to the power-balance property and it is a powerful tool for the analysis of nonlinear open systems. In the following, we denote $\langle \mathbf{a} | \mathbf{b} \rangle = \mathbf{a}^T \mathbf{b}$.

Definition 1 (Passivity, [7, Def 6.3]). *A system of state $\mathbf{x}(t) \in \mathbb{R}^n$, input $\bar{\mathbf{u}}(t) \in \mathbb{R}^p$ and output $\bar{\mathbf{y}}(t) \in \mathbb{R}^p$ is said to be passive if there exists a continuously differentiable positive semidefinite function $V(\mathbf{x})$ (called the storage function³) such that for all the trajectories and inputs*

$$\langle \bar{\mathbf{u}}(t) | \bar{\mathbf{y}}(t) \rangle \geq \langle \nabla V(\mathbf{x}(t)) | \dot{\mathbf{x}}(t) \rangle, \quad \forall t \in \mathbb{R}^+ \quad (\text{B.1})$$

Moreover, it is said to be

- *lossless if*

$$\langle \bar{\mathbf{u}}(t) | \bar{\mathbf{y}}(t) \rangle = \langle \nabla V(\mathbf{x}(t)) | \dot{\mathbf{x}}(t) \rangle \quad (\text{B.2})$$

- *strictly passive if*

$$\langle \bar{\mathbf{u}}(t) | \bar{\mathbf{y}}(t) \rangle \geq \langle \nabla V(\mathbf{x}(t)) | \dot{\mathbf{x}}(t) \rangle + \psi(\mathbf{x}(t)) \quad (\text{B.3})$$

for some positive definite function $\psi(\mathbf{x})$.

Note that if $V(\mathbf{x})$ represents the system's energy and $\bar{P}_{ext} = \langle \bar{\mathbf{u}} | \bar{\mathbf{y}} \rangle$ the external power given to the system, the integral of (B.1) over any period of time $[0, t]$ reads

$$\int_0^t \langle \bar{\mathbf{u}}(s) | \bar{\mathbf{y}}(s) \rangle ds \geq V(\mathbf{x}(t)) - V(\mathbf{x}(0)) \quad (\text{B.4})$$

which interprets as "the system is passive if the energy given to the network through inputs and outputs over that period is greater or equal to the increase in the energy stored in the network over the same period". In the absence of external control $\mathbf{u} = 0$, the passivity condition (B.1) implies

$$V(\mathbf{x}(0)) \geq V(\mathbf{x}(t)) \quad \forall t \in \mathbb{R} \quad (\text{B.5})$$

Remark 1. If the system is governed by

$$\begin{cases} \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \bar{\mathbf{u}}) \\ \bar{\mathbf{y}} = \mathbf{h}(\mathbf{x}, \bar{\mathbf{u}}) \end{cases}, \quad (\text{B.6})$$

the criterion (B.1) can be written as

$$\langle \bar{\mathbf{u}} | \mathbf{h}(\mathbf{x}, \bar{\mathbf{u}}) \rangle \geq \nabla V(\mathbf{x})^T \mathbf{f}(\mathbf{x}, \bar{\mathbf{u}}), \quad \forall (\mathbf{x}, \bar{\mathbf{u}}) \in \mathbb{R}^n \times \mathbb{R}^p. \quad (\text{B.7})$$

³The storage function is a Lyapunov function if it is positive definite.

Appendix B.1. Passivity of PHS formulations (i)-(iii)

Proposition 1. *The semi-explicit port-Hamiltonian DAE (formulation (i)) governed by*

$$\underbrace{\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{w} \\ \mathbf{y} \end{bmatrix}}_{\mathcal{F}_i} = \mathbf{S} \underbrace{\begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{z}(\mathbf{w}) \\ \mathbf{u} \end{bmatrix}}_{\mathcal{E}_i}, \quad (\text{B.8})$$

with $\mathbf{S} = -\mathbf{S}^T$ and $\langle \mathbf{z}(\mathbf{w}) | \mathbf{w} \rangle \geq 0$, is passive in the sense of Definition 1 for storage function $V = H$ and input-output $(\bar{\mathbf{u}}, \bar{\mathbf{y}}) = (\mathbf{u}, -\mathbf{y})$.

Proof. By skew-symmetry of \mathbf{S} , $\mathcal{E}^T \mathcal{F} = \mathcal{E}^T \mathbf{S} \mathcal{E} = 0$, so that the instantaneous power balance

$$\underbrace{\langle \nabla H(\mathbf{x}) | \dot{\mathbf{x}} \rangle}_{\text{stored power } P_S} + \underbrace{\langle \mathbf{z}(\mathbf{w}) | \mathbf{w} \rangle}_{\text{dissipated power } P_R \geq 0} + \underbrace{\langle \mathbf{u} | \mathbf{y} \rangle}_{\text{external power } P_P} = 0. \quad (\text{B.9})$$

After introducing $V = H$ and $(\bar{\mathbf{u}}, \bar{\mathbf{y}}) = (\mathbf{u}, -\mathbf{y})$, (B.9) implies

$$\langle \nabla V(\mathbf{x}) | \dot{\mathbf{x}} \rangle - \langle \bar{\mathbf{u}} | \bar{\mathbf{y}} \rangle = -\langle \mathbf{z}(\mathbf{w}) | \mathbf{w} \rangle \leq 0 \quad (\text{B.10})$$

from which we obtain the passivity condition ((B.1)). \square

Proposition 2. *The input-state-output port-Hamiltonian system with feed through:*

$$\underbrace{\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{y} \end{bmatrix}}_{\mathcal{F}_{ii}} = (\mathbf{J} - \mathbf{R}(\mathbf{x}, \mathbf{u})) \underbrace{\begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{u} \end{bmatrix}}_{\mathcal{E}_{ii}}, \quad (\text{B.11})$$

where $\mathbf{J} = -\mathbf{J}^T$ and $\mathbf{R} = \mathbf{R}^T \succeq 0$ is passive in the sense of Definition 1 for storage function $V = H$ and input-output $(\bar{\mathbf{u}}, \bar{\mathbf{y}}) = (\mathbf{u}, -\mathbf{y})$. Moreover, it is strictly passive if

$$\psi(\mathbf{x}) = \begin{bmatrix} \nabla H(\mathbf{x}) \\ 0 \end{bmatrix}^T \mathbf{R} \begin{bmatrix} \nabla H(\mathbf{x}) \\ 0 \end{bmatrix} > 0, \quad \forall \mathbf{x} \neq 0 \quad (\text{B.12})$$

Proof. The instantaneous power balance

$$\begin{aligned} \mathcal{E}^T \mathcal{F} &= \underbrace{\langle \nabla H(\mathbf{x}) | \dot{\mathbf{x}} \rangle}_{P_S} + \underbrace{\langle \mathbf{u} | \mathbf{y} \rangle}_{P_P} \\ &= - \underbrace{\left\langle \begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{u} \end{bmatrix} | \mathbf{R}(\mathbf{x}) | \begin{bmatrix} \nabla H(\mathbf{x}) \\ \mathbf{u} \end{bmatrix} \right\rangle}_{P_R \geq 0} < 0 \end{aligned} \quad (\text{B.13})$$

The passivity condition is obtained from (B.13) introducing $V = H$ and $(\bar{\mathbf{u}}, \bar{\mathbf{y}}) = (\mathbf{u}, -\mathbf{y})$. Note that, because $\mathbf{R} \succeq 0$, $-P_R < -\psi(\mathbf{x})$. Introducing this inequality into (B.13) gives

$$\langle \nabla V(\mathbf{x}) | \dot{\mathbf{x}} \rangle - \langle \bar{\mathbf{u}} | \bar{\mathbf{y}} \rangle = -P_R < -\psi(\mathbf{x}), \quad (\text{B.14})$$

from which the strict passivity condition ((B.3)) is obtained. \square

Proposition 3. *The skew-symmetric gradient PH-DAE:*

$$\underbrace{\begin{bmatrix} \dot{\mathbf{x}} \\ \mathbf{w} \\ \mathbf{y} \end{bmatrix}}_{\mathcal{F}_{iii}} = \mathbf{S} \underbrace{\nabla F \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{w} \\ \mathbf{u} \end{bmatrix} \right)}_{\mathcal{E}_{iii}}, \quad (\text{B.15})$$

where $\mathbf{S} = -\mathbf{S}^T$ and $F = H(\mathbf{x}) + Z(\mathbf{w}) + \frac{\mathbf{u}^T \mathbf{u}}{2}$, with $\nabla Z(\mathbf{w}) \mathbf{w} \geq 0$, is passive in the sense of Definition 1 for storage function $V = H$ and input-output $(\bar{\mathbf{u}}, \bar{\mathbf{y}}) = (\mathbf{u}, -\mathbf{y})$.

The proof is similar to Proposition 1 with $\mathbf{z}(\mathbf{w}) = \nabla Z(\mathbf{w})$.

Appendix C. Jacobian quantities for the physical systems

Tables C.5-C.7 show the Jacobian matrix $\mathbf{J}_f(\mathbf{x})$, the spectral norm $\|\mathbf{J}_f(\mathbf{x})\|_2$, the condition number $\kappa(\mathbf{J}_f(\mathbf{x}))$ and the stiffness ratio $\rho(\mathbf{J}_f(\mathbf{x}))$ for the different physical systems. The spectral norm is computed using the formula [10]

$$\|\mathbf{J}_f(\mathbf{x})\|_2 = \sqrt{\lambda_{\max}((\mathbf{J}_f(\mathbf{x}))^T \mathbf{J}_f(\mathbf{x}))} = \sigma_{\max}(\mathbf{J}_f(\mathbf{x})), \quad (\text{C.1})$$

where $\lambda_{\max}(\cdot), \sigma_{\max}(\cdot)$ denotes the maximum eigenvalue and singular value of a matrix. The condition number is computed using the formula [10]

$$\kappa(\mathbf{J}_f(\mathbf{x})) = \|\mathbf{J}_f(\mathbf{x})\|_2 \|\mathbf{J}_f(\mathbf{x})^{-1}\|_2 = \frac{\sigma_{\max}(\mathbf{J}_f(\mathbf{x}))}{\sigma_{\min}(\mathbf{J}_f(\mathbf{x}))} \quad (\text{C.2})$$

where $\sigma_{\max}(\cdot), \sigma_{\min}(\cdot)$ denote the maximum and minimum singular value of a matrix. The stiffness ratio is computed using the formula [12]

$$\rho(\mathbf{J}_f(\mathbf{x})) = \frac{\lambda_{\max}(\mathbf{J}_f(\mathbf{x}))}{\lambda_{\min}(\mathbf{J}_f(\mathbf{x}))} \quad (\text{C.3})$$

where $\lambda_{\max}(\cdot), \lambda_{\min}(\cdot)$ denote the maximum and minimum eigenvalue of a matrix.

Appendix D. Implementation details

Table E.8-E.10 show the design choices for each sub-network inside the PHNN that models the energy or the dissipation of the system.

Appendix E. Initial conditions and control design

Initial condition sampling: All the Hamiltonians in this work can be written in the form

$$H(q, p) = \frac{1}{2m} p^2 + \frac{1}{2} k_1 q^2 + \frac{\beta}{4} k_3 q^4 \quad (\text{E.1})$$

System	Harmonic oscillator (linear)
$\mathbf{J}_f(\mathbf{x})$	$\begin{bmatrix} 0 & \frac{1}{m} \\ -k & -\frac{\alpha}{m} \end{bmatrix}$
$\ \mathbf{J}_f(\mathbf{x})\ _2$	$\sqrt{\frac{1}{2} \left(k^2 + \frac{1+\alpha^2}{m^2} + \sqrt{\left(k^2 + \frac{1+\alpha^2}{m^2} \right)^2 - \frac{4k^2}{m^2}} \right)}$
$\kappa(\mathbf{J}_f(\mathbf{x}))$	$\sqrt{\frac{k^2 + \frac{1+\alpha^2}{m^2} + \sqrt{\left(k^2 - \frac{1+\alpha^2}{m^2} \right)^2 + \frac{4k^2\alpha^2}{m^2}}}{k^2 + \frac{1+\alpha^2}{m^2} - \sqrt{\left(k^2 - \frac{1+\alpha^2}{m^2} \right)^2 + \frac{4k^2\alpha^2}{m^2}}}}$
$\rho(\mathbf{J}_f(\mathbf{x}))$	1, if $\alpha^2 < 4km$ (underdamped)

Table C.5: Jacobian matrix, spectral norm, condition number and stiffness ratio for the harmonic oscillator.

System	Duffing oscillator (nonlinear)
$\mathbf{J}_f(\mathbf{x})$	$\begin{bmatrix} 0 & \frac{1}{m} \\ -k_1 - 3k_3q^2 & -\frac{\alpha}{m} \end{bmatrix}$
$\ \mathbf{J}_f(\mathbf{x})\ _2$	$\sqrt{\frac{1}{2} \left((k_1 + 3k_3q^2)^2 + \frac{1+\alpha^2}{m^2} + \sqrt{\left((k_1 + 3k_3q^2)^2 + \frac{1+\alpha^2}{m^2} \right)^2 - \frac{4(k_1+3k_3q^2)^2}{m^2}} \right)}$
$\kappa(\mathbf{J}_f(\mathbf{x}))$	$\sqrt{\frac{(k_1+3k_3q^2)^2 + \frac{1+\alpha^2}{m^2} + \sqrt{\left((k_1+3k_3q^2)^2 - \frac{1+\alpha^2}{m^2} \right)^2 + \frac{4\alpha^2(k_1+3k_3q^2)^2}{m^2}}}{(k_1+3k_3q^2)^2 + \frac{1+\alpha^2}{m^2} - \sqrt{\left((k_1+3k_3q^2)^2 - \frac{1+\alpha^2}{m^2} \right)^2 + \frac{4\alpha^2(k_1+3k_3q^2)^2}{m^2}}}}$
$\rho(\mathbf{J}_f(\mathbf{x}))$	1, if $\alpha^2 < 4m(k_1 + 3k_3q^2)$ (underdamped)

Table C.6: Jacobian matrix, spectral norm, condition number and stiffness ratio for the Duffing oscillator.

System	Self-sustained oscillator (nonlinear)
$\mathbf{J}_f(\mathbf{x})$	$\begin{bmatrix} -\Gamma'(w)kq - \Gamma(w)k & -\frac{1}{m} \\ k & 0 \end{bmatrix}$
$\ \mathbf{J}_f(\mathbf{x})\ _2$	$\sqrt{\frac{1}{2} \left(k^2(\Gamma'(w)q + \Gamma(w))^2 + k^2 + \frac{1}{m^2} + \sqrt{\left(k^2(\Gamma'(w)q + \Gamma(w))^2 + k^2 + \frac{1}{m^2} \right)^2 - \frac{4k^2}{m^2}} \right)}$
$\kappa(\mathbf{J}_f(\mathbf{x}))$	$\sqrt{\frac{k^2(\Gamma'(w)q + \Gamma(w))^2 + k^2 + \frac{1}{m^2} + \sqrt{\left(k^2(\Gamma'(w)q + \Gamma(w))^2 + k^2 - \frac{1}{m^2} \right)^2 + \frac{4k^2(\Gamma'(w)q + \Gamma(w))^2}{m^2}}}{k^2(\Gamma'(w)q + \Gamma(w))^2 + k^2 + \frac{1}{m^2} - \sqrt{\left(k^2(\Gamma'(w)q + \Gamma(w))^2 + k^2 - \frac{1}{m^2} \right)^2 + \frac{4k^2(\Gamma'(w)q + \Gamma(w))^2}{m^2}}}}$
$\rho(\mathbf{J}_f(\mathbf{x}))$	1, if $k^2(\Gamma'(w)q + \Gamma(w))^2 < \frac{4k}{m}$ (underdamped)

Table C.7: Jacobian matrix, spectral norm, condition number and stiffness ratio for the self-sustained oscillator.

For the harmonic (HO) and the self-sustained oscillator (SSO), the Hamiltonian satisfies $\beta = 0$, so that (E.1) particularizes to

$$H(q, p) = \frac{p^2}{2m} + \frac{1}{2}k_1q^2 \quad (\text{E.2})$$

We now want to sample an initial condition $\mathbf{x}_0 = (q_0, p_0)$ so that $H(\mathbf{x}_0) \in \mathcal{I}_{E_0} = [E_{min}, E_{max}]$. This problems is equivalent to sample (q, p) such that

$$E_0 = \frac{p^2}{2m} + \frac{1}{2}k_1q^2, \quad E_0 \in \mathcal{I}_{E_0} \quad (\text{E.3})$$

Note that (E.3) is the equation of an ellipse with semiaxes $a = \sqrt{\frac{2E_0}{k_1}}$ and $b = \sqrt{2E_0m}$. In parametric equations, this means that (q, p) can be written as:

$$\begin{cases} q = \sqrt{\frac{2E_0}{k_1}} \sin \theta \\ p = \sqrt{2mE_0} \cos \theta \end{cases} \quad E_0 \in \mathcal{I}_{E_0}, \theta \in [0, 2\pi) \quad (\text{E.4})$$

The initial condition sampling method (Algorithm 1) for the harmonic and self-sustained oscillator is based on (E.4).

Algorithm 1 Initial condition sampling for HO and SSO

Require: E_{\min}, E_{\max}

Ensure: (q, p)

- 1: Sample $\theta \sim \mathcal{U}(0, 2\pi)$
- 2: Sample $r \sim \mathcal{U}(E_{\min}, E_{\max})$
- 3: Compute $E_0 \leftarrow \sqrt{r}$
- 4: Compute

$$q \leftarrow \sqrt{\frac{2E_0}{k_1}} \sin \theta, \quad p \leftarrow \sqrt{2mE_0} \cos \theta$$

- 5: **return** (q, p)
-

For the Duffing oscillator (DO), the Hamiltonian satisfies $\beta > 0$ in (E.1). This Hamiltonian is not quadratic, which increases the complexity if we are to design a similar sampling technique to the one described before. In this case, the initial condition sampling is based on an acceptance-rejection method ((Algorithm 2)).

Algorithm 2 Initial condition sampling for DO

Require: $E_{\min}, E_{\max}, q_{\min}, q_{\max}, p_{\min}, p_{\max}$

Ensure: (q, p)

- 1: **repeat**
 - 2: Sample $q \sim \mathcal{U}(q_{\min}, q_{\max})$
 - 3: Sample $p \sim \mathcal{U}(p_{\min}, p_{\max})$
 - 4: Compute $H_0 \leftarrow H(q, p)$
 - 5: **until** $E_{\min} \leq H_0 \leq E_{\max}$ \triangleright Accept only if Hamiltonian in desired range
 - 6: **return** (q, p)
-

Control design: For the harmonic and Duffing oscillator, external control was applied as a constant force \mathbf{u} so that

$$H(\mathbf{x}^*) \in \mathcal{I}_{E_{eq}} = [E_{eq}^{\min}, E_{eq}^{\max}] \quad (\text{E.5})$$

The equilibrium ph-DAE formulation for these systems is

$$\begin{bmatrix} \dot{q} = 0 \\ \dot{p} = 0 \\ w^* \\ y^* \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -1 & 0 & -1 & -1 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \nabla_q H(q^*, p^*) = f(q^*) \\ \nabla_p H(q^*, p^*) = p^*/m \\ z(w^*) \\ u^* \end{bmatrix} \quad (\text{E.6})$$

which simplifies to

$$\begin{cases} 0 = p^*/m \\ 0 = -f(q^*) - z(w^*) - u^* \\ w^* = p^*/m \\ y^* = p^*/m \end{cases} \quad (\text{E.7})$$

From (E.7), $p^* = w^* = y^* = 0$, which also implies that $z(w^*) = cw^* = 0$. The equilibrium point of the system (q^*, p^*) satisfies

$$\begin{cases} H(q^*, p^*) = \frac{1}{2}k_1(q^*)^2 + \frac{\beta}{4}k_3(q^*)^4 = E_{eq} \in \mathcal{I}_{E_{eq}} \\ f(q^*) = k_1q^* + \beta k_3(q^*)^3 = -u^* \end{cases} \quad (\text{E.8})$$

The following algorithm is designed to find the control u^* that shifts the equilibrium point to the desired energy.

Algorithm 3 Control design for the HO and DO

Require: $E_{eq}^{\min}, E_{eq}^{\max}$

Ensure: Control input u^*

- 1: Sample $E_{eq} \sim \mathcal{U}(E_{eq}^{\min}, E_{eq}^{\max})$
 - 2: Find q^* such that $H(q^*, 0) = E_{eq}$
 - 3: Compute control input $u^* \leftarrow -k_1q^* - \beta k_3(q^*)^3$
 - 4: **return** u^*
-

For the case of the self-sustained oscillator, we want to apply a constant control \mathbf{u} such that the system stabilizes in a limit cycle around \mathbf{x}^* . The equilibrium ph-DAE formulation for this system is

$$\begin{bmatrix} \dot{q} = 0 \\ \dot{p} = 0 \\ w^* \\ y^* \end{bmatrix} = \begin{bmatrix} 0 & -1 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ -1 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \nabla_q H(q^*, p^*) = kq^* \\ \nabla_p H(q^*, p^*) = p^*/m \\ z(w^*) \\ u^* \end{bmatrix} \quad (\text{E.9})$$

which simplifies to

$$\begin{cases} 0 = -p^*/m + z(w^*) \\ 0 = kq^* \\ w^* = -kq^* - u^* \\ y^* = z(w^*) \end{cases} \quad (\text{E.10})$$

We note that the system has an equilibrium point in $(q^*, p^*) = (0, m \cdot z(w^*))$. In order to determine whether closed orbits arise in this system, we use the Poincaré-Bendixson theorem [57]. As there is only one fixed point, and the phase space is of dimension 2, the existence of a close orbit (or, limit cycle) is guaranteed if the fixed point is unstable. The fixed point is unstable if $\Delta > 0$ and $\tau < 0$ (see [57]), where Δ and τ are the determinant and the trace of the Jacobian matrix \mathbf{J}_f of $f = (\dot{q}, \dot{p})$ evaluated at the fixed point. From (E.10)

$$\mathbf{J}_f(q^*, p^*) = \begin{bmatrix} z'(w^*) \cdot k & -\frac{1}{m} \\ k & 0 \end{bmatrix} \quad (\text{E.11})$$

Model	Components	Network type	Input dimension	Layers	Hidden units	Output dimension	Number of parameters	
							(per component)	(per model)
NODE	Black-box	MLP	3	3	100	2	21,4k	21,4k
PHNN-JR	$\mathbf{L}_{\theta_{L_H}}(\mathbf{x})$	MLP	2	2	100	3	10,7k	21,8k
	$\mathbf{L}_{\theta_{L_R}}(\mathbf{x}, \mathbf{u})$	MLP	3	2	100	3	11,1k	
PHNN-S	$\mathbf{L}_{\theta_{L_H}}(\mathbf{x})$	MLP	2	2	100	3	10,7k	21,1k
	$\mathbf{L}_{\theta_{L_z}}(\mathbf{w})$	MLP	1	2	100	1	10,4k	
	$\mathbf{K}_{\theta_{L_z}}(\mathbf{w})$	MLP	1	2	100	0	0	

Table E.8: Specificities of the implemented port-Hamiltonian neural networks with large number of parameters.

Model	Components	Network type	Input dimension	Layers	Hidden units	Output dimension	Number of parameters	
							(per component)	(per model)
NODE	Black-box	MLP	3	2	60	2	4,3k	4,3k
PHNN-JR	$\mathbf{L}_{\theta_{L_H}}(\mathbf{x})$	MLP	2	2	42	3	2,1k	4,3k
	$\mathbf{L}_{\theta_{L_R}}(\mathbf{x}, \mathbf{u})$	MLP	3	2	42	3	2,2k	
PHNN-S	$\mathbf{L}_{\theta_{L_H}}(\mathbf{x})$	MLP	2	2	42	3	2,1k	4,1k
	$\mathbf{L}_{\theta_{L_z}}(\mathbf{w})$	MLP	1	2	42	1	2k	
	$\mathbf{K}_{\theta_{L_z}}(\mathbf{w})$	MLP	1	2	42	0	0	

Table E.9: Specificities of the implemented port-Hamiltonian neural networks with medium size of parameters.

Model	Components	Network type	Input dimension	Layers	Hidden units	Output dimension	Number of parameters	
							(per component)	(per model)
NODE	Black-box	MLP	3	2	24	2	848	848
PHNN-JR	$\mathbf{L}_{\theta_{L_H}}(\mathbf{x})$	MLP	2	2	16	3	371	809
	$\mathbf{L}_{\theta_{L_R}}(\mathbf{x}, \mathbf{u})$	MLP	3	2	16	3	438	
PHNN-S	$\mathbf{L}_{\theta_{L_H}}(\mathbf{x})$	MLP	2	2	16	3	371	852
	$\mathbf{L}_{\theta_{L_z}}(\mathbf{w})$	MLP	1	2	20	1	481	
	$\mathbf{K}_{\theta_{L_z}}(\mathbf{w})$	MLP	1	2	20	0	0	

Table E.10: Specificities of the implemented port-Hamiltonian neural networks with low size of parameters.

The Jacobian matrix \mathbf{J}_f has determinant $\Delta = km > 0$ and trace $\tau = z'(w^*) \cdot k$. In order to generate self-oscillations, the external force u must be chosen so that $z'(w^*) = z'(-kq^* - u) < 0$. Taking into account that in the equilibrium $q^* = 0$ and $w^* = -u^*$, the set of external controls that sets the system in a self-oscillation is

$$u^* = -w^* \text{ where } w^* \in \{w \in \mathbb{R} \mid z'(w) < 0\} \quad (\text{E.12})$$

The above information gives the following algorithm

Algorithm 4 Control design for the SSO

Require: $z(w)$

Ensure: Control input u^*

- 1: Find $I_w = \{w \in \mathbb{R} \mid z'(w) < 0\}$.
 - 2: Sample $w \sim \mathcal{U}(I_w)$
 - 3: Compute control input $u^* = -w$
 - 4: **return** u^*
-

Appendix F. Table of results

Tables F.11-F.13 show the exact inference errors for each of the studies presented in Section 5. In each table

cell, the median value and IQR for 10 model initializations are reported.

Harmonic	25	100	400
NODE-RK2	2.53e-02 [2.33e-02]	1.69e-04 [1.54e-04]	2.79e-05 [1.84e-05]
PHNN-JR-DG	2.58e-04 [3.82e-04]	9.10e-06 [1.22e-05]	7.29e-06 [5.07e-06]
PHNN-JR-RK2	1.14e-04 [1.02e-04]	6.88e-06 [3.81e-06]	5.18e-06 [5.50e-06]
PHNN-S-DG	4.22e-05 [5.99e-05]	6.24e-06 [4.97e-06]	4.78e-06 [1.21e-05]
PHNN-S-RK2	4.89e-05 [2.11e-04]	5.77e-06 [9.55e-06]	3.73e-06 [8.89e-06]
Duffing	25	100	400
NODE-RK2	4.09e-01 [1.01e+00]	7.62e-02 [2.62e-01]	3.67e-03 [2.73e-03]
PHNN-JR-DG	4.36e-01 [1.19e+00]	1.83e-02 [2.55e-02]	7.89e-03 [4.76e-03]
PHNN-JR-RK2	5.82e-01 [2.96e+00]	2.13e-02 [3.15e-02]	6.15e-03 [9.92e-03]
PHNN-S-DG	7.06e-02 [4.37e+00]	8.08e-03 [4.93e-02]	9.08e-04 [8.69e-04]
PHNN-S-RK2	9.35e-02 [4.54e+00]	2.07e-02 [1.85e-02]	9.17e-03 [6.13e-03]
Self-sustained	25	100	400
NODE-RK2	2.59e-02 [2.21e-02]	2.01e-04 [1.38e-04]	9.56e-05 [9.18e-05]
PHNN-JR-DG	1.71e-03 [2.71e-03]	7.93e-05 [5.03e-05]	7.06e-05 [7.01e-05]
PHNN-JR-RK2	1.38e-02 [2.48e-02]	2.59e-04 [2.63e-04]	1.39e-04 [2.20e-04]
PHNN-S-DG	5.40e-03 [3.60e-03]	1.68e-03 [4.83e-04]	2.72e-03 [1.12e-03]
PHNN-S-RK2	4.26e-03 [7.81e-03]	2.64e-03 [1.60e-03]	2.22e-03 [7.85e-04]

Table F.11: Median inference error and IQR (in brackets) for different numbers of training points when learning the three oscillatory systems with small models. In bold, the best result for each combination of system and number of training points.

Harmonic	Small size	Medium size	Large size
NODE-RK2	2.53e-02 [2.33e-02]	2.85e-02 [2.59e-02]	1.45e-01 [4.67e-02]
PHNN-JR-DG	2.58e-04 [3.82e-04]	1.47e-04 [2.68e-04]	7.33e-04 [2.52e-03]
PHNN-JR-RK2	1.14e-04 [1.02e-04]	5.56e-05 [1.19e-04]	1.51e-04 [3.87e-04]
PHNN-S-DG	4.22e-05 [5.99e-05]	5.56e-05 [9.58e-05]	2.45e-04 [4.59e-04]
PHNN-S-RK2	4.89e-05 [2.11e-04]	4.28e-05 [3.73e-05]	2.01e-04 [2.13e-04]
Duffing	Small size	Medium size	Large size
NODE-RK2	4.09e-01 [1.01e+00]	3.80e-01 [1.77e+00]	2.31e-01 [3.08e+00]
PHNN-JR-DG	4.36e-01 [1.19e+00]	1.15e-01 [1.19e-01]	3.08e-01 [4.01e-01]
PHNN-JR-RK2	5.82e-01 [2.96e+00]	1.50e-01 [2.60e-01]	3.08e-01 [7.61e-02]
PHNN-S-DG	7.06e-02 [4.37e+00]	8.63e-02 [1.10e-01]	1.80e-01 [1.46e-01]
PHNN-S-RK2	9.35e-02 [4.54e+00]	7.42e-02 [7.00e-02]	2.19e-01 [1.70e-01]
Self-sustained	Small size	Medium size	Large size
NODE-RK2	2.59e-02 [2.21e-02]	3.74e-02 [3.06e-02]	6.53e-02 [5.74e-02]
PHNN-JR-DG	1.71e-03 [2.71e-03]	6.56e-03 [2.02e-02]	7.59e-03 [1.54e-02]
PHNN-JR-RK2	1.38e-02 [2.48e-02]	1.36e-02 [2.37e-02]	1.75e-02 [1.50e-02]
PHNN-S-DG	5.40e-03 [3.60e-03]	6.12e-03 [3.48e-03]	7.74e-03 [5.18e-03]
PHNN-S-RK2	4.26e-03 [7.81e-03]	3.98e-03 [2.95e-03]	3.77e-03 [4.27e-03]

Table F.12: Median inference error and IQR (in brackets) for different model sizes when learning the three oscillatory systems with $N_{train} = 25$. In bold, the best result for each combination of system and model size.

Harmonic	BL	CN	SN	SR
NODE-RK2	2.53e-02 [2.33e-02]	7.82e-02 [1.63e+00]	2.63e-02 [2.53e-02]	2.56e-02 [5.70e-02]
PHNN-JR-DG	2.58e-04 [3.82e-04]	2.20e-04 [3.40e-04]	1.94e-04 [3.09e-04]	2.30e-04 [3.19e-04]
PHNN-JR-RK2	1.14e-04 [1.02e-04]	1.50e-04 [5.70e-04]	7.18e-05 [9.39e-05]	1.26e-04 [1.46e-04]
PHNN-S-DG	4.22e-05 [5.99e-05]	3.84e-05 [4.05e-05]	3.44e-05 [6.62e-05]	1.13e-04 [3.07e-04]
PHNN-S-RK2	4.89e-05 [2.11e-04]	1.17e-05 [5.67e-05]	5.66e-05 [2.02e-04]	2.97e-05 [5.35e-04]
Duffing	BL	CN	SN	SR
NODE-RK2	4.09e-01 [1.01e+00]	9.93e+00 [4.77e+01]	5.77e-01 [9.72e-01]	7.33e-01 [5.18e+00]
PHNN-JR-DG	4.36e-01 [1.19e+00]	6.72e+00 [5.63e+01]	3.55e-01 [1.24e+01]	4.43e-01 [8.66e+01]
PHNN-JR-RK2	5.82e-01 [2.96e+00]	2.76e+00 [2.35e+01]	3.26e-01 [2.69e+01]	9.97e-01 [3.91e+02]
PHNN-S-DG	7.06e-02 [4.37e+00]	5.15e-02 [8.92e-02]	6.26e-02 [1.07e-01]	6.57e-02 [2.54e+00]
PHNN-S-RK2	9.35e-02 [4.54e+00]	6.74e-02 [5.83e-02]	8.27e-02 [4.53e+00]	6.97e+00 [2.80e+01]
Self-sustained	BL	CN	SN	SR
NODE-RK2	2.59e-02 [2.21e-02]	4.04e-02 [1.63e-02]	2.39e-02 [3.27e-02]	3.10e-02 [4.33e-02]
PHNN-JR-DG	1.71e-03 [2.71e-03]	2.84e-03 [2.74e-03]	1.91e-03 [3.14e-03]	1.43e-02 [2.36e-02]
PHNN-JR-RK2	1.38e-02 [2.48e-02]	8.24e-03 [2.58e-02]	1.38e-02 [2.93e-02]	3.62e-02 [3.97e-02]
PHNN-S-DG	5.40e-03 [3.60e-03]	4.90e-03 [3.89e-03]	5.23e-03 [3.72e-03]	1.22e-02 [4.37e-03]
PHNN-S-RK2	4.26e-03 [7.81e-03]	5.00e-03 [5.67e-03]	4.31e-03 [8.56e-03]	5.11e-03 [6.72e-03]

Table F.13: Median inference error and IQR (in brackets) after applying the different Jacobian regularizations when learning the three oscillatory systems with $N_{train} = 25$ and small models. The best result for each combination of system and Jacobian regularization is shown in bold. Notation: BL refers to the baseline (i.e. no Jacobian regularization); CN, to the condition number regularization (34); SN, to the spectral norm regularization (35); and SR, to the stiffness ratio regularization (36)

Figure captions

Figure 1: Architecture of the two PHNN models considered in this work. White boxes with orange contour denote fixed algebraic operations whereas orange boxes indicate the trainable parameters.

Figure 2: Training and inference diagram for the continuous models f_θ .

Figure 3: Training and inference diagram for the discrete models g_θ .

Figure 4: Training points, test initial points and two complete test trajectories for each of the three oscillatory systems. Note that in the case of the harmonic and Duffing oscillator, the applied control shifted the equilibrium point from $(p, q) = (0, 0)$ whereas in the case of the self-sustained oscillator, it stabilizes the trajectories in a limit cycle.

Figure 5: Schematic sampling and dataset construction procedure. A trajectory generated at sampling frequency sr_{gen} over a duration $D = D_{infer} = \beta T_0$ is shown as white dot markers. From this trajectory, a training point, highlighted with a green star marker, is uniformly sampled from a subset of samples, shown as black dot markers, obtained at frequency sr_{train} and restricted to $t \leq \alpha T_0$. The training horizon D_{train} and the inference horizon D_{infer} are indicated by arrows, with the vertical dashed line marking the end of the training interval.

Figure 6: Boxplot of the inference errors for the three different oscillators and varying numbers of training points (small models). From left to right: harmonic, Duffing and self-sustained oscillators.

Figure 7: Comparison of the learned trajectory discretizations for each oscillatory system starting on \mathbf{x}_0 such that $H(\mathbf{x}_0) = 0.5J$. **(Left)** NODE and PHNN-S models trained on $N_{train} = 25$ from the harmonic oscillator. **(Center)** NODE and PHNN-S models trained on $N_{train} = 100$ from the Duffing oscillator. **(Right)** NODE and PHNN-JR models trained on $N_{train} = 25$ from the self-sustained oscillator. For the harmonic and Duffing oscillator, the trajectory is obtained with $\mathbf{u} = 0$; and for the self-sustained, with \mathbf{u} such that the system stabilizes in a limit cycle.

Figure 8: Boxplot of the inference errors for the three different oscillators and model sizes ($N_{train} = 25$). From left to right: harmonic, Duffing and self-sustained oscillators.

Figure 9: Comparison of the learned trajectory discretizations by the discrete gradient for each oscillatory system starting on \mathbf{x}_0 such that $H(\mathbf{x}_0) = 0.75J$. **(Left)**

PHNN-S models trained on $N_{train} = 25$ from the harmonic oscillator. **(Center)** PHNN-S models trained on $N_{train} = 25$ from the Duffing oscillator. **(Right)** PHNN-JR models trained on $N_{train} = 25$ from the self-sustained oscillator. For the harmonic and Duffing oscillator, the trajectory is obtained with $\mathbf{u} = 0$; and for the self-sustained, with \mathbf{u} such that the system stabilizes in a limit cycle.

Figure 10: Mean condition number **(left)**, spectral norm **(center)** and stiffness ratio **(right)** per optimizer step during the training for the different architectures when modeling the Duffing oscillator with $N_{train} = 25$ points and no Jacobian regularization (BL). For each optimizer step, the corresponding mean value (solid line) is obtained computing the average over the 10 model initializations, with the shaded area indicating the range between the minimum and maximum values.

Figure 11: Boxplot of the inference errors for the three different oscillators and normalizations ($N_{train} = 25$ and small number of parameters). Notation: BL refers to the baseline (i.e. no Jacobian regularization); CN, to the condition number regularization (34); SN, to the spectral norm regularization (35); and SR, to the stiffness ratio regularization (36). From left to right: harmonic, Duffing and self-sustained oscillators.

Figure 12: Impact on the mean condition number **(left)**, spectral norm **(center)** and stiffness ratio **(right)** per optimizer step during the training for the different architectures when modeling the Duffing oscillator with $N_{train} = 25$ points and the proposed Jacobian regularizations. For each optimizer step, the corresponding mean value (solid line) is obtained computing the average over the 10 model initializations, with the shaded area indicating the range between the minimum and maximum values.