

## RESEARCH ARTICLE OPEN ACCESS

# Neural Ordinary Differential Equations for Model Order Reduction of Stiff Systems

Matteo Caldana<sup>1</sup>  | Jan S. Hesthaven<sup>2</sup>

<sup>1</sup>MOX, Dipartimento di Matematica, Politecnico di Milano, Milano, Italy | <sup>2</sup>Karlsruhe Institute of Technology, Karlsruhe, Germany

**Correspondence:** Matteo Caldana ([matteo.caldana@polimi.it](mailto:matteo.caldana@polimi.it))

**Received:** 25 November 2024 | **Revised:** 14 May 2025 | **Accepted:** 30 May 2025

**Funding:** This work was supported by Ministero dell'Università e della Ricerca (Grant Numbers: PRIN2020 n. 20204LN5N5).

**Keywords:** neural networks | ordinary differential equations | reduced-order model | Runge-Kutta | stiff equations

## ABSTRACT

Neural Ordinary Differential Equations (ODEs) represent a significant advancement at the intersection of machine learning and dynamical systems, offering a continuous-time analog to discrete neural networks. Despite their promise, deploying neural ODEs in practical applications often encounters the challenge of stiffness, a condition where rapid variations in some components of the solution demand prohibitively small time steps for explicit solvers. This work addresses the stiffness issue when employing neural ODEs for model order reduction by introducing a suitable reparametrization in time. The considered map is data-driven, and it is induced by the adaptive time-stepping of an implicit solver on a reference solution. We show that the map produces a non-stiff system that can be cheaply solved with an explicit time integration scheme. The original, stiff, time dynamic is recovered by means of a map learnt by a neural network that connects the state space to the time reparametrization. We validate our method through extensive experiments, demonstrating improvements in efficiency for the neural ODE inference while maintaining robustness and accuracy when compared to an implicit solver applied to the stiff system with the original right-hand side.

**MSC2020 Classification:** 65L99, 68T07

## 1 | Introduction

In computational science and engineering, the complexity and scale of numerical simulations have grown exponentially, driven by advances in hardware and algorithms. High-fidelity models—also called Full Order Models (FOMs)—can now capture intricate details of physical phenomena across diverse domains such as fluid dynamics, structural analysis, climate modeling, and materials science. However, the increasing complexity of these models often comes at a significant computational cost, rendering them impractical for many applications that require repeated model evaluations over a large number of parameter values. Indeed, the computational cost is driven up by the curse of dimensionality [1].

Reduced-Order Models (ROMs) have emerged as a powerful solution to this challenge, offering a balance between accuracy and computational efficiency. By compressing the essential features of FOMs into a more manageable form, usually called latent space, ROMs enable the efficient simulation and analysis of complex systems, especially those requiring real-time analysis [2], inverse problems [3–6], uncertainty quantification [7–9], shape optimization [10] or optimal control [11–13].

The field of ROMs has seen substantial growth, providing practitioners with a wide range of techniques to choose from. Each method comes with its own set of advantages and drawbacks, and the optimal choice of a ROM often depends on the specific requirements of the application. Indeed, the term ROMs

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *International Journal for Numerical Methods in Engineering* published by John Wiley & Sons Ltd.

refers to multiple approaches, which are rather different from one another.

One of the most well-established paradigms for model order reduction is proper orthogonal decomposition (POD) [14, 15], that is, a linear projection technique that provides a high compression rate and, for instance, achieves a particularly high level of accuracy in the case of diffusion processes.

If the governing equations of the high-fidelity model are explicitly employed in the ROM, they are called intrusive [16–20]. Usually, a greedy algorithm or POD is employed to project the FOM onto a lower-dimensional space. However, for nonlinear models, additional techniques like the (discrete) empirical interpolation method are required to handle the complexities, which can lead to an unfavorable tradeoff between accuracy and computational efficiency [21–23]. Moreover, in [24] it has been shown that for transport-dominated problems, a large latent dimension must be used to achieve reasonable accuracy when using POD, thus limiting their practical use.

In recent years, data-driven methods based on machine learning and deep learning have gained large traction due to advances in GPU-based hardware and developments in algorithms and software supporting the artificial intelligence ecosystem. This paradigm is based on the availability of a large quantity of high-fidelity data that can be used to train the model. This approach proves effective if the large amount of resources employed in the offline phase (generation of the solutions and training of the model) is repaid by a large enough number of (cheap) evaluations of the trained model (online phase). These models are called non-intrusive since only knowledge of the data is needed to build them [25, 26]. We also refer to these kinds of ROMs as surrogate models, given their ability to approximate the behavior of the original FOM from the data. The field is ever-growing, with a large number of techniques being developed and published. To the best of our knowledge, the most relevant examples of data-driven approaches are the following. Closure models are a hybrid between POD and data-driven approaches where a nonlinear term (a neural network) is added to the linear ROM (usually POD-Galerkin) to account for the effect of the unresolved dynamics [27]. In [28–32], densely connected or convolutional autoencoders have been used to extract the full latent variables. Other techniques have been proposed to enhance the accuracy. For instance, graph neural networks (GNNs) tackle the problem of complicated geometries discretized with unstructured meshes, offering a more natural representation of variables using a graph [33, 34]. Fourier neural operators are a very successful approach that instead exploits the frequency space to enrich the layers of the model [35–37]. Depending on the considered model, time evolution may be treated just as any parameter of the system or explicitly learned, examples are DeepONets [31, 38], dynamic mode decomposition [39, 40], sparse identification of reduced latent dynamics (SINDy) [41, 42], RNNs [43], LSTMs [32, 44] or Gaussian processes [45].

Another time evolution approach of particular interest is based on neural ODEs [46–49]. The core concept of this method is straightforward: it involves learning the right-hand side of a dynamical system with a dense neural network, which can be viewed as the continuous limit of a residual neural network. This

approach offers numerous advantages, including a more natural representation of continuous-time data and enhanced parameter and memory efficiency. The primary challenge with neural ODEs is that they are, in effect, ODEs that must be solved. This often leads to issues with stiffness, complicating their solution and potentially affecting their stability and accuracy. Stiffness is the condition where there are rapid changes in some components of the solution while other components change slowly. This phenomenon creates numerical challenges when solving the equations, as explicit numerical methods might become inefficient or unstable [50]. Thus, the necessity of employing implicit solvers significantly drives up the cost.

In the present work, we tackle the problem of building a ROM of a stiff system. In this case, the main bottleneck is not represented by the large number of variables of the system but by the stiffness of the system. Indeed, there are multiple types of stiffness, which are usually connected to features of the system, such as the presence of multiple time scales, the spectrum of the Jacobian of the system, or numerical instability, that make building a ROM a challenging task. In this case, employing fixed-time time-steppers is unfeasible since the rapidly varying scales require unreasonably small steps. On the other hand, treating time like an input of a dense neural network also proves to be problematic since neural networks are known to perform badly on high frequency [51]. Still, in literature, there can be found several examples of stiff problems solved with machine learning. A possible solution was first proposed in [52], where the authors employ the quasi-steady-state assumption (QSSA) to reduce the stiffness of the ODE systems and show that a physics-informed neural network (PINN) can then be successfully applied to the converted non-/mild-stiff systems. This is particularly interesting since the system can be solved with an explicit method. In [53], the authors show that computing a stabilized gradient and suitable scaling of the network outputs enables learning stiff neural ODEs. In [54], PINNs are used with extreme learning machine to solve stiff problems. In [55], stiff quantitative systems pharmacology models are accelerated with echo state networks. In [56], an autoencoder is used to produce a ROM of a chemical kinetic model for the simulation of a combustion system. Recently, in [57] DeepONets are used to learn the discretized solution of challenging stiff chemical kinetics.

Our strategy is to employ a neural ODE-based non-stiff surrogate for stiff dynamics. Namely, we aim to employ a neural ODE to learn a time-reparametrized stiff system, where the time map is suitably built to reduce the stiffness of the problem. The insight is to exploit the advantages of the neural ODEs we previously described while keeping the (online) cost of solving the ODE under control by using an explicit solver. Indeed, our main contribution is to employ the time reparametrization induced by the time-stepping of an implicit solver as a way to reduce the stiffness of the system. An advantage of this procedure is that it is purely data-driven, meaning that it is possible to avoid manually deriving explicit algebraic expressions for the QSS species as was done in [52]. On the other hand, a drawback of our approach is that it is only worthwhile for ROMs: the expensive phase of gathering data for building the time mapping must be offset by a cheap online phase. Our technique can also be interpreted as a way to embed adaptive (implicit) time-stepping directly into the neural ODE. Thus, we can avoid backpropagation through the implicit solver,

which incurs a cubic cost with respect to the number of neural network parameters [53], thereby greatly reducing the computational cost of training the neural ODE. Our reparametrization strategy is also connected with [58–60], where the authors leverage spatial mappings for adaptive mesh refining to improve computational efficiency and to enhance the representation of nonlinear dynamics in complex systems. Moreover, recent work [61] has demonstrated that a structure-preserving NODE with linear/nonlinear splitting can achieve long-term stability on stiff problems comparable to implicit methods. Once the non-stiff system is solved, the solution is mapped back to the stiff dynamics using a map learnt with a neural network. The neural network is trained once with the neural ODE and shows good generalization properties to parameters outside the training dataset. Particular care is taken in the definition of the map so that it has no dependence on time. Hence, the model maintains good generalization in time for periodic systems since there is only explicit dependence on the state space.

Our technique is designed especially for problems where stiffness can be traced back to the case of large eigenvalues with a negative real part. Indeed, another largely studied type of stiffness is high-frequency stiffness, which corresponds to the case of reasonably large eigenvalues but very high frequency [62]. In these cases, the system features a fast and a slow solution, where the fast solution must be solved accurately to track the slow one. While reparametrization approaches have been proposed [63], they have significant limitations since they need to resolve the fast frequency [64, 65]. Thus, this kind of stiffness remains a challenge for our approach and an area of ongoing research. Indeed, other authors have explored structure-preserving machine learning techniques, such as fast-slow neural networks, to bypass the direct resolution of fast transients in singularly perturbed systems [66].

We demonstrate the accuracy and speed of the proposed ROM on a set of five test problems, where explicit solvers usually fail, widely used in literature to benchmark stiff ODE solvers [67]. The results are compared with a state-of-the-art Runge-Kutta implicit solver of the Radau II A kind of the fifth order [50] applied to the stiff system with the original right-hand side.

The remainder of the paper is organized as follows. First, in Section 2 we introduce the mathematical settings and notation for the problems of interest. In particular, in Section 2.4 we review the neural ODEs approach, which is the foundation of this work. In Section 3 we introduce our methodology, discussing also the implementation details. In Section 4 we present a comprehensive set of numerical results. Finally, in Section 5 we draw some conclusions, discussing the strengths and limitations of the proposed approach.

## 2 | Problem Formulation

In this section, we first present the construction of ROMs in algebraic terms for approximating a parametric dynamical system. Then, we provide a concise overview of fundamental deep learning concepts and formulate the data-driven model order reduction problem within the framework of neural ODEs.

### 2.1 | Parametric Systems of Ordinary Differential Equations

We consider the following parameterized, finite-dimensional dynamical system, described by a set of first-order ODE:

$$\begin{cases} \dot{\mathbf{u}}(t; \boldsymbol{\mu}) = \mathbf{f}(t, \mathbf{u}; \boldsymbol{\mu}), & t \in (0, T] \\ \mathbf{u}(0; \boldsymbol{\mu}) = \mathbf{u}_0(\boldsymbol{\mu}) \end{cases} \quad (1)$$

where  $\boldsymbol{\mu}$  is a vector containing all the parameters of the system, which belongs to a compact set  $\Gamma \subset \mathbb{R}^{N_\mu}$ ,  $\mathbf{u} : [0, T] \times \Gamma \rightarrow \mathbb{R}^{N_u}$  is the parameterized solution,  $\mathbf{f} : [0, T] \times \mathbb{R}^{N_u} \times \Gamma \rightarrow \mathbb{R}^{N_u}$  is the right-hand side function (encoding the dynamical system),  $t \in [0, T]$  is the time variable,  $\mathbf{u}_0 : \Gamma \rightarrow \mathbb{R}^{N_u}$  is the initial condition, and  $\dot{\mathbf{u}}$  is the total derivative with respect to time  $t$ . Without loss of generality, we are considering that the initial time for the system is zero; indeed, there always exists a translation in time that brings us to this case.

This formulation encompasses not only ODEs but also problems stemming from the semi-discrete formulation of systems of partial differential equations (PDEs) through suitable methods (e.g., finite elements [68], spectral elements [69], discontinuous Galerkin [70, 71], etc.). In these cases, one may observe an increase in the computational cost of the simulation since the dimension of the state space  $N_u$  is usually large. Indeed, these methods usually rely on a fine mesh to discretize the spatial derivatives.

The meaning of the parameter  $\boldsymbol{\mu} \in \Gamma$  is broad; indeed, it may represent physical properties (e.g., material properties), initial conditions, or geometrical properties (e.g., the shape of the domain). In this work, we focus on cases where  $\boldsymbol{\mu}$  models physical properties.

### 2.2 | Solving Ordinary Differential Equations

There is a wide literature of methods to solve ODEs [50]. In this work, we concentrate on the well-known Runge-Kutta methods. Let  $r \in \mathbb{N}$  s.t.  $r \geq 1$ , let  $b_i, a_{ij}, (i, j = 1, \dots, r)$  be real numbers and let  $c_i$  be defined as

$$c_i = \sum_{j=1}^{i-1} a_{ij}, \quad \forall i = 1, \dots, r$$

Then, fixed a timestep  $\Delta t > 0$  and setting  $t^{(n)} = n\Delta t$ , we call the  $r$ -stage Runge-Kutta method the following discretization of Equation (1)

$$\begin{aligned} \mathbf{k}_i &= \mathbf{f}\left(t^{(n)} + c_i \Delta t, \mathbf{u}^{(n)} + \Delta t \sum_{j=1}^r a_{ij} \mathbf{k}_j\right), \quad i = 1, \dots, r \\ \mathbf{u}^{(n+1)} &= \mathbf{u}^{(n)} + \Delta t \sum_{i=1}^r b_i \mathbf{k}_i \end{aligned} \quad (2)$$

$\forall n \in \mathbb{N}$  s.t.  $t^{(n)} \leq T$ . Here,  $\mathbf{u}^{(n+1)} \approx \mathbf{u}(t^{(n)} + \Delta t)$ . If  $a_{ij} = 0$  for  $i \leq j$ , each stage of Equation (2) depends only on the previous ones, thus the method is explicit. If  $a_{ij} = 0$  for  $i < j$  and at least one

$a_{ii} = 0$ , the method is called diagonally implicit. In all the other cases, it is called implicit.

Of particular interest are embedded Runge-Kutta methods, which automatically adjust the step size  $\Delta t$  to achieve a prescribed tolerance of the local error. The embedded method exploits a second set of scalars  $b_i^*$ ,  $i = 1, \dots, r$  to compute another approximation of the solution of lower order  $u^*$ . Namely, we want to satisfy the component-wise inequality

$$|(u^{(n+1)})_i - (u^{*(n+1)})_i| \leq \delta_i,$$

$$\delta_i = \text{atol} + \text{rtol} \max(|(u^{(n+1)})_i|, |(u^{(n)})_i|), \quad i = 1, \dots, N_u$$

where  $\text{atol}$  and  $\text{rtol}$  are user-defined (positive) absolute and relative tolerances, respectively. Then, the error is estimated as

$$\varepsilon = \sqrt{\frac{1}{N_u} \sum_{i=1}^{N_u} \left( \frac{(u^{(n+1)})_i - (u^{*(n+1)})_i}{\delta_i} \right)^2}$$

and it is compared to one to find the next step size, namely

$$\Delta t^{(n+1)} = \Delta t^{(n)} \min(10, \max(0.1, \varepsilon^{-1/(q+1)}))$$

where  $q$  is usually chosen to be the order of the method, and 10 and 0.1 are the maximum and minimum change factors, respectively. The value of these factors can be tuned depending on the problem. If the normalized error is larger than one, the timestep is rejected and is computed again with the updated  $\Delta t$ .

Explicit Runge-Kutta methods are generally unsuitable for the solution of stiff equations because their region of absolute stability is small [50]. Indeed, in the present work, we will employ three different Runge-Kutta methods depending on the application. Namely, we use an implicit Runge-Kutta of the Radau II A kind of the fifth order [72] for solving the stiff problems, and we use a Runge-Kutta of the fourth order (in the adaptive and fixed step size variants) for non-stiff problems. We employ the following table for the Radau II A method

$\frac{2}{5} - \frac{\sqrt{6}}{10}$	$\frac{11}{45} - \frac{7\sqrt{6}}{360}$	$\frac{37}{225} - \frac{169\sqrt{6}}{1800}$	$-\frac{2}{225} + \frac{\sqrt{6}}{75}$
$\frac{2}{5} + \frac{\sqrt{6}}{10}$	$\frac{37}{225} + \frac{169\sqrt{6}}{1800}$	$\frac{11}{45} + \frac{7\sqrt{6}}{360}$	$-\frac{2}{225} - \frac{\sqrt{6}}{75}$
1	$\frac{4}{9} - \frac{\sqrt{6}}{36}$	$\frac{4}{9} + \frac{\sqrt{6}}{36}$	$\frac{1}{9}$
	$\frac{4}{9} - \frac{\sqrt{6}}{36}$	$\frac{4}{9} + \frac{\sqrt{6}}{36}$	$\frac{1}{9}$

Due to the super-convergence of the Radau II A method (classical order  $p = 2r - 1$ ) it is not possible to have an embedded method of order  $p - 1$  without extra cost. By taking a linear combination of  $\Delta t f(x_0, u_0)$  and the internal stage values  $k_i$  it is however possible to get an approximation of order  $r$ , we refer the interested reader to [72] for details. The following tableau is used for explicit Runge-Kutta of the fourth-order [73], also called the Dormand-Prince method or RK45

0							
1/5	1/5						
3/10	3/40	9/40					
4/5	44/45	-56/15	32/9				
8/9	19372/6561	-25360/2187	64448/6561	-212/729			
1	9017/3168	-355/33	46732/5247	49/176	-5103/18656		
1	35/384	0	500/1113	125/192	-2187/6784	11/84	
	35/384	0	500/1113	125/192	-2187/6784	11/84	0
	5179/57600	0	7571/16695	393/640	-92097/339200	187/2100	1/40

### 2.3 | Data-Driven Model Order Reduction

In this section, we discuss the concepts and notation that stand at the basis of data-driven model order reduction. We introduce the manifold of solutions of Equation (1)

$$\mathcal{M} = \{u(t; \mu) \mid t \in [0, T], \mu \in \Gamma\} \subset \mathbb{R}^{N_u} \quad (3)$$

which is a collection of the values of the solutions for each parameter  $\mu \in \Gamma$  and time  $t \in [0, T]$ . We suppose that it is possible to find a discretized approximation  $u(t^{(n)}; \mu)$  for some times  $t^{(n)} \in [0, T]$  up to arbitrary accuracy by integrating (1) with a suitable method. By solving the problem for  $N_\mu$  different choices of  $\mu$  we create a discretized manifold

$$\widetilde{\mathcal{M}} = \{u(t^{(n)}; \mu^{(j)}) \mid \mu^{(j)} \in \Gamma, 1 \leq j \leq N_\mu, t^{(n)} \in [0, T], 1 \leq n \leq N(\mu^{(j)})\} \quad (4)$$


which is a finite set of vectors that approximate the solution at a certain discretized point in time  $t^{(n)}$  and for a certain choice of parameter  $\mu^{(j)}$ . Let us stress that the time discretization  $\{0 \leq t_1 < \dots < t_{N(\mu^{(j)})} \leq T\}$  may change depending on the value of  $\mu^{(j)}$ , and need not be uniform. Indeed, it often happens that  $\mu^{(j)}$  changes the dynamics of the system, and thus a different time discretization is generated by the adaptive procedure. For the sake of simplicity, we omit the explicit dependence on  $\mu^{(j)}$  of the time discretization.

In practice, one might also obtain data from knowing the exact analytical solution of the problem for some specific values of  $\mu$  and  $t$  or from real-world measurements. Data-driven ROMs aim to leverage the information contained in the discretized solution manifold  $\widetilde{\mathcal{M}}$  to obtain in a fast and reliable way the solution  $u(t^{(n)}; \mu^{(j)})$  for  $\mu^{(j)}$  or  $t^{(n)}$  not in the discrete manifold used for training the model. Figure 1 shows a graphical representation of the solution manifold for the Van der Pol oscillator (cf. Section 4.1). The motivation behind ROMs is that classical time integration methods may become prohibitively expensive either due to the size of  $N_u$  or due to stiffness (cf. Section 3).

### 2.4 | Neural Odes

The model that we choose to use to build the ROM is neural ODEs [46]. Neural ODEs have emerged as powerful tools to model continuous-time dynamics using a differential equation. This approach provides several key benefits. Firstly, neural ODEs





**FIGURE 1** | Example of discrete solution manifold  $\tilde{\mathcal{M}}$  for the Van der Pol oscillator.

offer superior flexibility in modeling time-series data, as they can naturally handle irregularly sampled data and varying time steps. Secondly, they ensure smooth trajectories and better integration with physical laws and other domains where the underlying processes are continuous in nature. Furthermore, neural ODEs can lead to more efficient training and inference, as they require fewer parameters and can leverage advanced numerical solvers for differential equations. Neural ODEs are also more memory efficient compared to traditional recurrent neural networks, especially for long time-series data, because they avoid the need to store intermediate states of the network during backpropagation.

More precisely, a neural ODE is a dense fully connected feedforward (FNN) neural network  $\mathcal{NN}$  [74, 75], with parameters  $\theta$ , that evolves the dynamics as

$$\dot{\mathbf{u}}(t; \boldsymbol{\mu}) = \mathcal{NN}(t, \mathbf{u}(t); \boldsymbol{\mu}) \quad (5)$$


that is, the FNN is employed to approximate the right-hand side  $\mathbf{f}$  of the dynamical system (1), thus it takes as inputs the time  $t$ , the current state  $\mathbf{u}$ , and the parameters  $\boldsymbol{\mu}$ . In other words, this technique aims to predict  $\mathbf{u}$  by using  $\mathcal{NN}$  to learn the  $\mathbf{f}$  of Equation (1). One of the main drawbacks of this approach is that it is necessary to numerically integrate the system in the online phase. Given the recurrent nature of the architecture, adjoint solvers have been developed to enhance the memory efficiency during the training phase [46]; their cost, however, remains much larger when compared to FNN.

When used as an architecture for ROMs, the neural ODE is often paired with another neural network to reduce the dimension of the system from  $N_u$  to a small latent state space. This is usually achieved by means of autoencoders [28–32] or other reduction techniques [49]. However, if  $N_u$  is small, the advantages of using a small latent space representation may not be counterbalanced by the increased cost of evaluating the neural network. In this work, we propose a solution to this latter problem by introducing an appropriate change of variables that significantly reduces the stiffness of the problem. Indeed, we will show that our method makes it possible to use explicit solvers to integrate neural ODEs representing stiff systems.

We report in Figure 2 a representation of the neural ODE. The hat symbol  $\hat{\cdot}$  represents a suitable normalization of the data, see the next section for details. For the sake of simplicity, we will indicate in the captions the reference solution (the one that is comprised in the dataset) with the subscript “ref”, and the components of  $\mathbf{u}$  as predicted by the neural ODE with  $u_i$ .

#### 2.4.1 | Training a Neural ODE

There are two main approaches to train a neural ODE. The first one is the classic supervised approach. Assuming that the value of  $\mathbf{f}(t^{(n)}, \mathbf{u}(t^{(n)}; \boldsymbol{\mu}^{(j)}); \boldsymbol{\mu}^{(j)})$  is known for each sample in the discrete manifold  $\tilde{\mathcal{M}}$ , then it is possible to compute the MSE loss and minimize it through a gradient-based optimizer. However,



**FIGURE 2** | Architecture of the Neural ODE for model order reduction.

this kind of data is not always available. Instead, the second approach—the one originally proposed for the training of neural ODEs—minimizes the error with respect to the time-series data  $u(t^{(n)}; \mu^{(j)})$  for a fixed  $j$ . Namely, by using a differentiable ODE solver, we minimize the loss

$$\mathcal{L}(\theta) = \frac{1}{N_{\text{train}}} \sum_{j=1}^{N_{\text{train}}} \left\| u(t^{(N)}; \mu^{(j)}) - (u(t^{(0)}; \mu^{(j)})) + \int_{t^{(0)}}^{t^{(N)}} \mathcal{N}(\tau, u(\tau; \mu^{(j)}), \mu^{(j)}; \theta) d\tau \right\|_p \quad (6)$$

where  $\|\cdot\|_p$  indicates the discrete  $p$ -norm, the elements  $u(t^{(0)}; \mu^{(j)})$  and  $u(t^{(N)}; \mu^{(j)})$  belong to the solution manifold  $\tilde{\mathcal{M}}$ ,  $\mathcal{N}$  is the approximation of  $f$  made by a suitably trained neural network and the integral is usually computed with an explicit integration method like forward Euler or an explicit Runge-Kutta method.


The advantage of the supervised approach is that it is less computationally demanding since it does not need to compute the gradient through the ODE solver using automatic differentiation. On the other hand, in the supervised framework, the loss is not always indicative of the accuracy of the predictions on the time-series data  $u(t^{(n)}; \mu)$ ,  $n = 0, \dots, N$ . For this reason, we always employ as validation loss Equation (6).

Implementation details for Equation (6) are important when it is employed to train the neural network. We resample the discrete time-series to match the output of the time integrator, namely a Runge-Kutta method of the fourth order with a fixed timestep. To minimize the interpolation errors, we use a reference solution with a very fine time discretization. A uniform resampling allows for an efficient use of the partial results of the ODE solver for all times smaller than the final one. Moreover, the initial and final integration times are not  $(t^{(0)}, t^{(N)})$  but are batched by uniform extraction between  $[0, T]$  so that each time integration in the dataset is done with the same number of timesteps: this procedure enables vectorization of the computations, greatly reducing the computational cost. This optimization was originally introduced in `torchdiffeq` [46]. The number of timesteps used is


called the unroll length and is a hyperparameter that plays a key role in the optimization procedure. Finally, as commonly done in supervised learning, the computations are split into mini-batches to increase stochasticity. If not otherwise stated, when using Equation (6) for validation or testing, we employ as a time integrator an embedded Runge-Kutta method of order four with adaptive time-stepping. However, computations are significantly cheaper since we can turn off automatic differentiation.

In our experiments, we employ a two-phase training. First, a standard supervised approach is used to “initialize” the neural network since it is cheap and enables a fast reduction of the loss. The derivative data is estimated with a filtering approach from the trajectory data (cf. Section 4). Then, we switch to the differentiable ODE solver approach to fine-tune the ROM. Our experiments show that in this case, it is beneficial to use an increasingly larger unroll length to ensure the robustness and accuracy of the model. The two-phase training is not strictly needed, but we have found it significantly reduces the time and effort required for hyperparameter tuning compared to using only the trajectory data. Further details will be given in Section 4 for each test case. Employing a two-phase training approach not only reduces the computational cost but also increases the quality of the final model. Indeed, from our experiments, we have observed that with the supervised learning approach, the neural network is able to learn sharper gradients that may lead to unstable states. On the other hand, learning through the ODE solver produces stable results that might follow less precisely the reference solution. Figure 3 shows an example of two (underfitting) models. In Figure 4, we show the training and validation loss for a model trained in sequence with the supervised and ODE solver approaches.

Another important point of discussion is the explicit dependence on time  $t$  of the  $\mathcal{N}$ . In practice, we do not use  $t$  as input to the neural network  $\mathcal{N}$ . However, this does not necessarily limit us to autonomous systems of ODEs. Indeed, it is possible to use an input layer with memory, like an LSTM, to add temporal dependence. This approach is usually preferred to employing  $t$  as input to the neural network since it improves generalization properties in time.



**FIGURE 3** | Comparison of the supervised and ODE solver approaches for training the neural ODE from the same randomly initialized model for the Van der Pol oscillator (cf. Section 4.1). Blue and orange are the predictions of the model, and black is the reference solution. The supervised training is globally less accurate and stable, but better captures sharp gradients.



**FIGURE 4** | Example of the behavior of the loss, the validation loss, and the learning rate during the training of the ROM with the two-step approach: first the supervised approach and then the ODE solver approach.

To accelerate the training, it is key to apply suitable normalizations to the input  $\mu$ , the output  $f$ , and the time-series data  $u$ . The exact expressions for the normalizations depend on the test case considered and are detailed in Section 4. In some special cases, for example, if  $\mu$  is high-dimensional, further treatment of  $\mu$  may be needed to accelerate the training. The insight behind the choice of the normalization is that we aim to have data that has a magnitude close to one, does not have outliers, and has either a uniform distribution or a normal distribution with a standard deviation of magnitude close to one. In our experiments, strictly monotone normalizations achieved better results. We designed the normalizations by manual inspection of the histograms of the data.

### 3 | Addressing the Stiffness Issue

While many techniques have been developed to build ROMs of large systems stemming from semi-discretized PDEs, not many solutions are available for stiff ODEs. Indeed, it is a challenge to learn their complex dynamics, which often feature large gradients and widely different time scales. More rigorously, a commonly used formula to define the stiffness  $S$  index of a problem is

$$S = \frac{\Re(\lambda_{\max})}{\Re(\lambda_{\min})} T \quad (7)$$

where  $\lambda_i$  are the eigenvalues of the Jacobian of the system and  $\Re$  indicates the real part. However, it is known that this definition is not always helpful. For instance, it does not apply to the classic Robertson chemical kinetics problem since it has two zero eigenvalues [76]. A more pragmatic definition is given by Harier: “Stiff equations are problems for which explicit methods do not work” [50]. That is, explicit numerical time integrators such as the embedded Runge-Kutta method of the fourth order require timesteps so small that implicit methods become cheaper. Stiffness is a particularly important issue when working with neural ODEs. Indeed, it is a well-known fact that neural ODE training may fail due to stiffness. For instance, [53] shows issues of this kind and proposes techniques to improve the training. However, it is still necessary to use a costly implicit solver in the online phase, since the learnt dynamics is stiff.

Hence, to use the neural ODE as a ROM, we aim to reduce the stiffness of the system, allowing for the use of (cheap) explicit methods. To this end, we propose to change the dynamics of the problem by considering the change of variable in time induced by the adaptive time-stepping of the implicit solver used to solve the FOM. We remark that this change of variable is particularly effective in the case of stiffness characterized by a large  $S$ . However, in cases where stiffness is characterized by the presence of an

oscillating fast solution, reparametrization techniques show several limitations [64, 65]. Our approach is not an exception, and this is an area for further improvement.

We suppose the existence of a change of variable  $t_s = t_s(t)$  that makes the system less stiff. Namely, instead of learning Equation (1), we learn the following system

$$\begin{cases} \dot{\mathbf{u}}_s(t_s; \boldsymbol{\mu}) = \mathbf{f}_s(t_s, \mathbf{u}_s; \boldsymbol{\mu}), & t_s \in (0, 1] \\ \mathbf{u}_s(0; \boldsymbol{\mu}) = \mathbf{u}_{s,0}(\boldsymbol{\mu}) \end{cases} \quad (8)$$

where,  $\mathbf{u}_s(t_s; \boldsymbol{\mu}) = \mathbf{u}(t(t_s); \boldsymbol{\mu})$  and thus  $\mathbf{f}_s$  is

$$\begin{aligned} \mathbf{f}_s(t_s, \mathbf{u}_s; \boldsymbol{\mu}) &= \frac{d\mathbf{u}_s(t_s; \boldsymbol{\mu})}{dt_s} = \frac{d\mathbf{u}(t(t_s); \boldsymbol{\mu})}{dt} \frac{dt}{dt_s} \\ &= \mathbf{f}(t(t_s; \boldsymbol{\mu}), \mathbf{u}; \boldsymbol{\mu}) i(t_s; \boldsymbol{\mu}) \end{aligned}$$

We aim to build the time mapping  $t_s = t_s(t)$  in a data-driven manner, in particular we suppose that

$$t_s(t^{(n)}) = \frac{n}{N}$$

that is, the  $\Delta t$  chosen at the  $n$ -th step by the implicit time stepper is proportional to the derivative of the function that we employ as a change of variable in time. We remark that  $t_s$  is strictly increasing and thus invertible, with inverse  $t = t(t_s)$ . We guess this change of variables induces less stiff dynamics since the adaptive time-stepping spreads the sharp peaks and quickly adapts to the large timescales. The insight behind this approach is that using an implicit solver to train the neural ODE on Equation (1) has a cubic cost [53], instead, this change of variable allows the training via an explicit method, which is significantly cheaper. In other words, the adaptive time-stepping is embedded into the neural ODE, avoiding the backpropagation step through the implicit solver. While the proposed approach is closely related to classical embedded methods, it significantly reduces the cost of training a neural ODE.

While it could be possible to learn  $t(t_s; \boldsymbol{\mu})$  and  $i(t_s; \boldsymbol{\mu})$  by simple regression methods (such as monotonicity-preserving interpolation), this approach has two significant drawbacks for our use case: (i) it becomes unfeasible to use Equation (6) to train the neural ODE; (ii) these interpolation schemes do not generalize well for  $t > T$  and quickly become expensive as the dimension of  $\boldsymbol{\mu}$  increases.

Thus, in this framework, we have two neural networks. The first,  $\mathcal{NN}$  is used to approximate  $\mathbf{f}_s$ , which only requires pointwise estimates of these values. Further details for the training are given in Section 3.1. The second ingredient needed to accommodate for the time reparametrization is a way to compute  $t(t_s; \boldsymbol{\mu})$  in the online phase, to map the solution back to the original dynamics. Our experiments show that learning directly the map  $t = t(t_s; \boldsymbol{\mu})$  is a hard task for a neural network, probably due to the presence of abrupt changes in the function. Instead, we found that learning  $i(\mathbf{u}_s(t_s; \boldsymbol{\mu}))$  is a stable approach that produces more accurate results. We note that, provided  $t = t(t_s; \boldsymbol{\mu})$  is continuously differentiable (or just absolutely continuous, if we only require the derivative almost everywhere), it can

always be represented—and hence learned via the neural ODE approximation—as

$$t(t_s; \boldsymbol{\mu}) = \int_{t_s^{(0)}}^{t_s} i(\tau; \boldsymbol{\mu}) d\tau \approx \int_{t_s^{(0)}}^{t_s} \mathcal{NN}_t(\tau, \mathbf{u}_s(\tau; \boldsymbol{\mu}), \boldsymbol{\mu}) d\tau$$

by the fundamental theorem of calculus. In this case,  $\mathbf{u}_s(\tau; \boldsymbol{\mu})$  is just an extra input used to accelerate the training. However, to improve the generalization properties and keep the system autonomous, we prefer to remove the explicit dependence on  $t_s$ . This is possible only if the trajectories of the ODE are an injective immersion in  $\mathbb{R}^{N_u}$  or if the “velocity”  $i(t_s; \boldsymbol{\mu})$  is uniquely well defined at every point of the trajectory. Since it is non-trivial to prove that this second property holds, we have only empirically tested the hypothesis by analyzing the data. Indeed, in practice, this approach works even if the trajectories satisfy the previous properties almost everywhere. Hence, in all our test cases, the time mapping is computed with the following path integral

$$t(t_s; \boldsymbol{\mu}) \approx \int_{\mathbf{u}_s(t_s^{(0)})}^{\mathbf{u}_s(t_s)} \mathcal{NN}_t(\mathbf{u}_s; \boldsymbol{\mu}) d\mathbf{u}_s$$

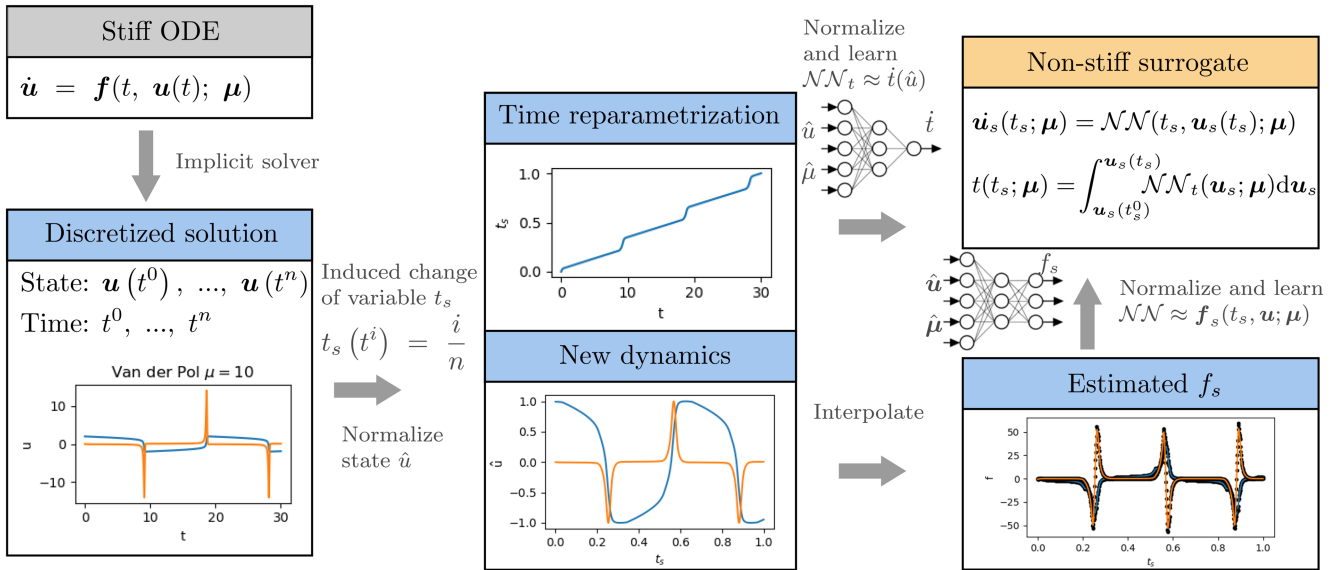
where the values of  $\mathbf{u}_s$  are the ones obtained by solving Equation (8) by time integration of the neural ODE  $\mathcal{NN}$ . Our experiments (see Section 4) demonstrate that the method generalizes well to parameters  $\boldsymbol{\mu}$  not present in the training set. The evaluation of the time dynamics is cheap since it can be done in one batch using vectorization. Indeed, in the first step, we compute  $\mathbf{u}_s(t_s; \boldsymbol{\mu})$ , then, we carry out the cumulative integration of  $i(\mathbf{u}_s(t_s; \boldsymbol{\mu}))$  with Simpson’s rule. In our experiments, this step requires only a few milliseconds and its cost is negligible with respect to the cost of the integration of the neural ODE. The main advantage of this approach is that it is cheap compared to autoencoders, especially for small problems since the only overhead is the computation of  $t = t(t_s; \boldsymbol{\mu})$ . However, it is also possible to combine the autoencoder with the described methodology: first, the autoencoder extracts a latent state representation, then, we could apply the time reparametrization in the latent space, reducing the stiffness of the latent dynamics.

Unlike traditional non-constant time-stepping, our data-driven time reparametrization fundamentally alters the underlying time dynamics by introducing a new, normalized time variable  $t_s$  based on the adaptive time-stepping behavior of the original system. While the results are in practice similar, this technique enables to predict the adaptive time-stepping for a new model parameter  $\boldsymbol{\mu}$  in a much cheaper way, without the necessity of evaluating the Jacobian or solving a linear system, like it is requested in an implicit ODE solver.

### 3.1 | Implementation Details

Applying this methodology to the training with the differentiable ODE solver does not require any change. Indeed, the values of the time-series  $\mathbf{u}_s(n/N; \boldsymbol{\mu}) = \mathbf{u}(t^{(n)}; \boldsymbol{\mu})$  are known, they only have a different time discretization, that we know pointwise. The main difference concerns the supervised learning approach. As explained before, this step allows us to “initialize” the neural network in a cheap way and requires us to estimate both





**FIGURE 5** | Time reparametrization workflow. Starting from the stiff ODE system (in gray), we compute the discrete solution and obtain the discrete manifold  $\tilde{\mathcal{M}}$  using an implicit method. We normalize the solution, compute the time reparametrization  $t_s(t)$ , and estimate  $f_s$  (in blue). Finally, we are able to train a non-stiff model that can be solved using an explicit method (in orange).

$\dot{u}_s(t_s; \mu) = f_s(t_s, u_s; \mu)$  and  $i$ , in order to train two neural network  $\mathcal{NN}$  and  $\mathcal{NN}_t$ , respectively. The approximation of the derivative is done using a Savitzky-Golay filter with low order (two or three) and small window size ( $< 10$ ). Figure 5 shows a scheme of the workflow. To accelerate the training also the values of the time derivative are normalized, in particular, we have found a logarithmic normalization to work particularly well.

The optimizer we employ is AdamW [77] with batch size 32 and a learning rate scheduler that geometrically reduces errors on plateaus of a factor of one-half with patience twenty epochs. The training lasts for a maximum of 400 epochs. The best model is chosen using early stopping based on the validation loss. The activation functions, the depth and width of the neural network, the initial learning rate, and the seed are hyperparameters that we optimize by means of random search on a problem-by-problem basis. In particular, the activation function is chosen among Gaussian error linear unit (GELU) [78], SiLU [79], Hard Swish [80], Leaky ReLU and ReLU, the depth between three and ten, the width between five and 100 and the initial learning rate between  $2 \cdot 10^{-3}$  and  $10^{-4}$ .

We employ a very similar approach to fine-tune the differentiable solver. The main difference is that the maximum number of epochs considered is 2000. The unroll length and the timestep size are tuned for each problem. As a rule of thumb, we employ increasingly larger unroll lengths with smaller and smaller learning rates. Namely, we start with covering about 5% of the time domain, and it increases up to 50%.

The neural network  $\mathcal{NN}_t$  approximating  $i(u_s(t_s; \mu))$  is also trained in a two-step process. First, it is trained in a supervised manner. Similar to what we do in the other supervised step, we employ as validation loss the accuracy of the predictions on the validation time-series when employing the ODE solver to

integrate both  $\mathcal{NN}_t$ , which approximates  $i(u_s(t_s; \mu))$ , and  $\mathcal{NN}$ , which approximates  $f_s(t_s, u_s; \mu)$ . Then, it is fine-tuned using as a target the time  $t$ . This is achieved by automatic differentiation through Simpson's integration rule. Namely, the loss is

$$\mathcal{L}(\theta) = \frac{1}{N_{\text{train}}} \sum_{j=1}^{N_{\text{train}}} \left\| t^{(N)} - \int_{u_s(t_s^{(0)})}^{u_s(t_s^{(N)})} \mathcal{NN}_t(u_s(\tau; \mu^{(j)}), \mu^{(j)}; \theta) d\tau \right\|_p$$

where  $t^{(0)}, t^{(N)}$  are sample times, and  $\mu^{(j)}$  are sample parameters used to build the discrete manifold  $\tilde{\mathcal{M}}$  and the integral is computed cumulatively on the discretization induced by the explicit Runge-Kutta method applied to the trained  $\mathcal{NN}$ . The time integration interval  $(t^{(0)}, t^{(N)})$  and the sequence of states  $u_s(\tau; \mu^{(j)})$  needed for the integration can be picked either from the solution manifold or the values obtained by applying the trained  $\mathcal{NN}$  to the training dataset. We tested  $p = 1, 2, 4$ , the best results were achieved with  $p = 1$ .

Finally, to achieve the best performances in the online phase (inference time), the model is optimized using the features of modern deep learning libraries, such as oneDNN graph fusion [81].

## 4 | Numerical Results

Our methodology is tested in five numerical experiments featuring stiff problems widely used in the literature to benchmark implicit ODE integrators [50, 67]. We do not consider dynamical systems derived from semi-discrete PDEs since it would require the use of autoencoders to extract the latent dynamics. Indeed, we aim to assess the behavior of our method as a stand-alone improvement to neural ODEs. Introducing other variables in the test cases, such as the error of the autoencoders, may be detrimental to this task. We are confident that, given the wide range of

test cases presented, our methodology could be extended to larger state spaces or semi-discrete PDEs without problems. Indeed, being able to outperform classical ODE solvers on small problems without the need to use a simplified latent space representation is a more difficult task.

Computations were carried out on a laptop with an AMD Ryzen 7 PRO 7840U. For the sake of reproducibility, we rely on the standard (CPU) implementations found in SciPy [82] of the Runge-Kutta methods for solving the original stiff system with the original right-hand side. Indeed, the evaluation of the original right-hand side is much cheaper than the evaluation of a neural ODE that approximates it. All the test cases feature systems that cost orders of magnitude more to solve with explicit solvers than with implicit ones. Moreover, to ensure a fair and rigorous comparison with the neural network, we executed both methods on a CPU. While leveraging a GPU would likely enable the neural network to significantly outperform the full-order method, our goal was to maintain a balanced evaluation framework. We remark that all the presented metrics, tables, and figures refer to the evaluation of  $\mathcal{NN}$  and  $\mathcal{NN}_t$  on a testing dataset, that is with values of the parameter  $\mu$  that are unseen by the neural network.

Before proceeding with the presentation of the test cases, let us introduce some metrics of interest that we will use to assess the performance of the ROMs.

- **time**: elapsed CPU time in seconds to solve the system up to the final time  $T$ .
- **# fev**: number of evaluations of the right-hand side (either  $f$  or  $\mathcal{NN}$ ).
- **# jev**: number of evaluations of the Jacobian of  $f$  (used for the implicit solver).
- **# lu**: number of LU decompositions (used for the implicit solver)
- **MSE <sub>$t_s$</sub>** : mean squared error of  $\mathcal{NN}$  in the  $t_s$  reparametrization.
- **MSE**: mean squared error in the  $t$  parametrization.
- **$L^2$** : relative integral error computed with Simpson's rule.
- **$d_{\text{peak}}$** : average distance of the peaks (any sample whose two direct neighbors have a smaller amplitude) between the reference solution and the low fidelity solution (used for periodic solutions).

Finally, we indicate with **tol** the tolerance of the solver. If not otherwise specified, it is intended that we set both the absolute (**atol**) and relative (**rtol**) tolerances of the embedded solver to this value.

#### 4.1 | Test Case 1: Van Der Pol Oscillator

The van der Pol oscillator is a second-order differential equation that exhibits limit cycle behavior, which makes it a valuable model for studying phenomena such as electrical circuits, cardiac rhythms, and neuronal firing patterns. The problem is governed

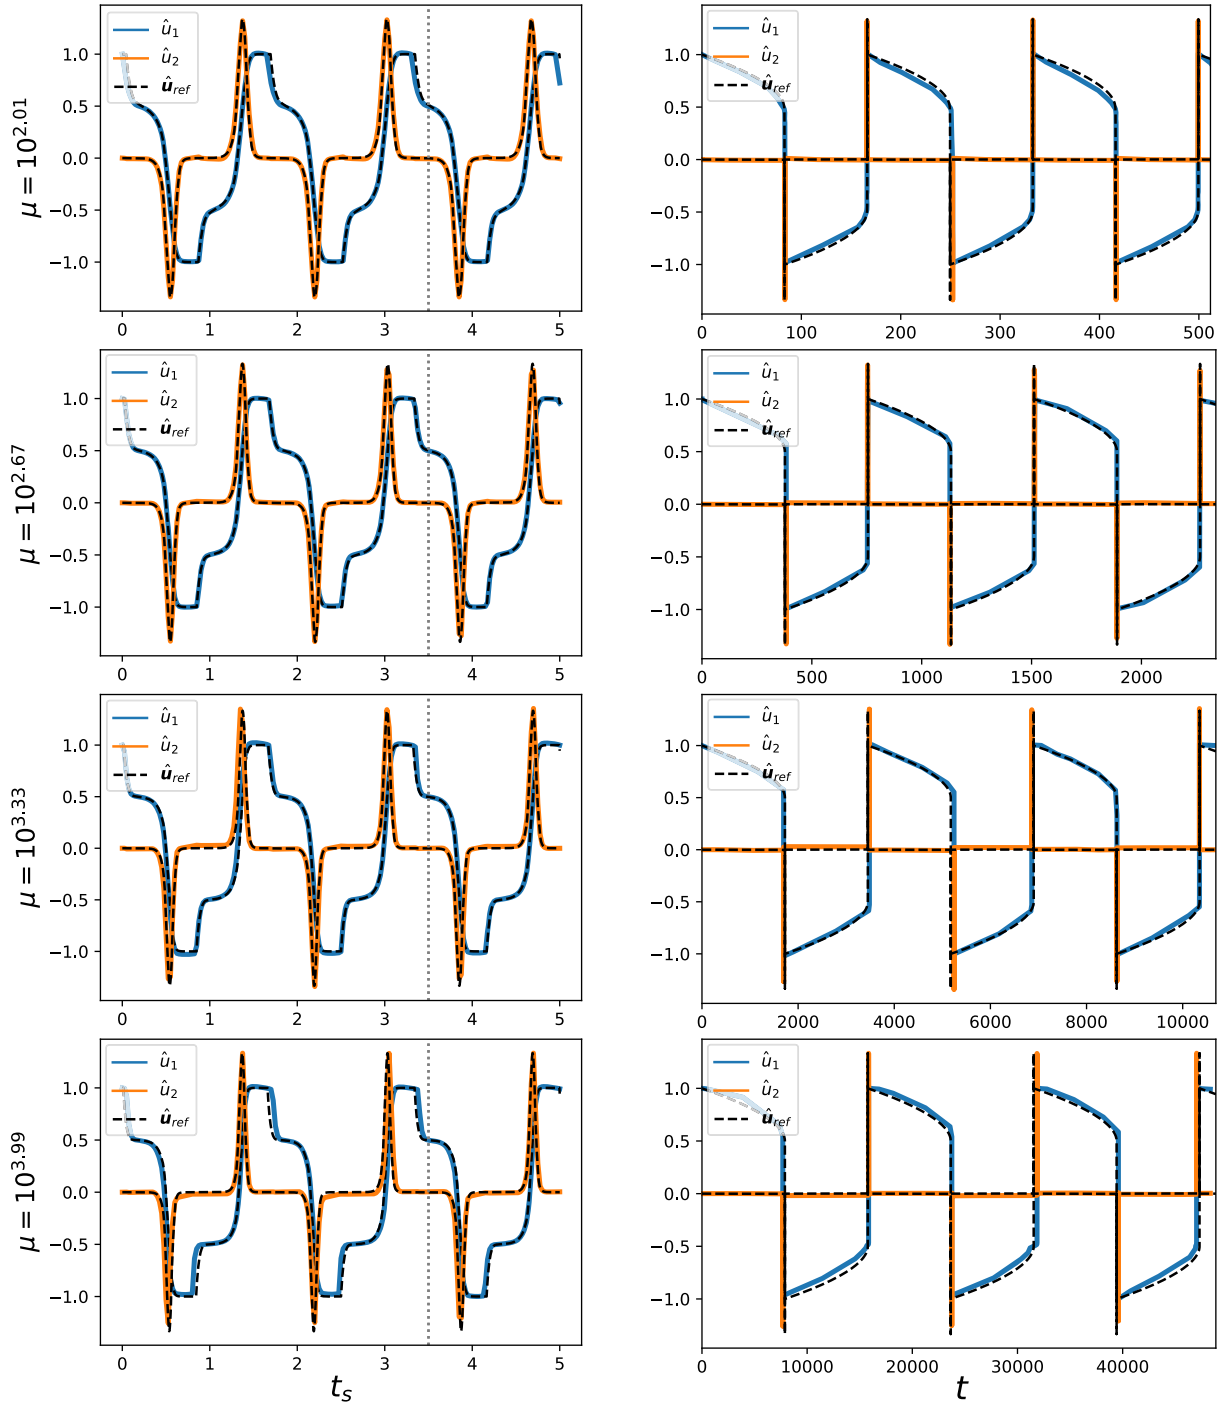
by a parameter  $\mu > 0$  and we consider its following formulation in the time interval  $t \in [0, 3.5\mu]$

$$\begin{cases} \dot{u}_1 = u_2, \\ \dot{u}_2 = \mu(1 - u_1^2)u_2 - u_1, \\ \mathbf{u}(0) = (2, 0) \end{cases} \quad (9)$$

This test case aims to assess the capabilities of our method on an autonomous periodic system. Indeed, the Van der Pol oscillator has two periodic solutions: a non-trivial periodic solution and an unstable zero solution. The parameter  $\mu > 0$  determines the significance of the nonlinear component of the equation. When  $\mu$  is large, the system becomes stiff. We set  $\Gamma = [10^2, 10^4]$  and consider the training dataset composed by a logarithmic discretization  $\Gamma$  with  $N_\mu = 51$  points. The validation dataset is the interval midpoints of discretized  $\Gamma$ . The reference solutions are computed using the Radau method, with absolute and relative tolerances set to  $10^{-12}$ .

The data for the supervised training is smoothed with a Savitzky-Golay filter of window seven and order two. Data is randomly subsampled to have about one thousand points for each period of the time-series. Fine-tuning is done by increasing the size of the unroll length starting from 20 up to 80 (exact values are a hyperparameter to be tuned) and time-stepping  $\Delta t = 1/40$ . The starting learning rate is reduced by a couple of orders of magnitude with respect to the supervised step.

Data is normalized using the following functions: parameter normalization:  $\hat{\mu} = \log_{10} \mu$ ; state normalization:  $\hat{\mathbf{u}} = (u_1/2, u_2/\mu)$ ; dynamics normalization:  $\hat{f} = (f_1/5, f_2/10)/\mu$ ; time normalization:  $\hat{T} = T/\mu$ . The model is tested for  $N_\mu^{\text{test}} = 4$  values of  $\mu \in [10^{2.01}, 10^{3.99}]$  spaced logarithmically (these values are chosen so that they are not present in the training nor the validation dataset). Moreover, we are also testing for time beyond what is present in the training and validation datasets, namely, we are integrating until a final time  $\hat{T} = 5$ . Figure 6 shows the results of applying the change of dynamics to this problem. On the left, the plot of the problem with the new dynamics shows smaller gradients. On the right, after the change of variable, we recover with good accuracy the original dynamics of the Van der Pol oscillator. In Table 1 we show a comparison of the computational cost. The number of right-hand side evaluations is one order of magnitude smaller for the explicit solver of the neural ODE w.r.t. the implicit solver of the FOM. This proves that the change of variables has significantly reduced the stiffness of the problem. Another interesting feature of our approach is that the ROM error is almost independent from  $\mu$ , on the other hand, the implicit solver exhibits a correlation between the two, suggesting that for larger values of  $\mu$  computational gains may be larger. In this case, the  $L^2$  error is computed only with respect to the first component of  $\mathbf{u}$ . Indeed, since  $u_2$  features sharp peaks, small variations in their position make the integral error large and not representative of the accuracy of the solution. For this reason, we have introduced the  $d_{\text{peak}}$  metric, which tries to measure the accuracy of the ROM in this regard. We stress that for this test case, the ROM is performing well also for unknown times, maintaining an accurate prediction of the period of the system. From the metrics in Table 1, it is also possible to notice that the main bottleneck in terms of accuracy for the ROM is the time mapping. Indeed, the model is very accurate in the  $t_s$  parametrization.



**FIGURE 6** | Test case 1: Van der Pol oscillator. On the left, in colors, the neural ODE predictions in  $t_s$ . On the right, in colors, the prediction is mapped to the original time  $t$ . In black is the reference solution. The vertical dotted line represents the final time in the training dataset.

In this case, we employ a fixed time-stepping in the online phase. This is a feasible approach only because the dynamics in the  $t_s$  parameterization exhibit similar behavior across all considered values of  $\mu$ .

#### 4.2 | Test Case 2: OREGO Problem

The OREGO problem consists of the following stiff system of 3 nonlinear ODEs defined in the time interval  $t \in [0, 1000]$

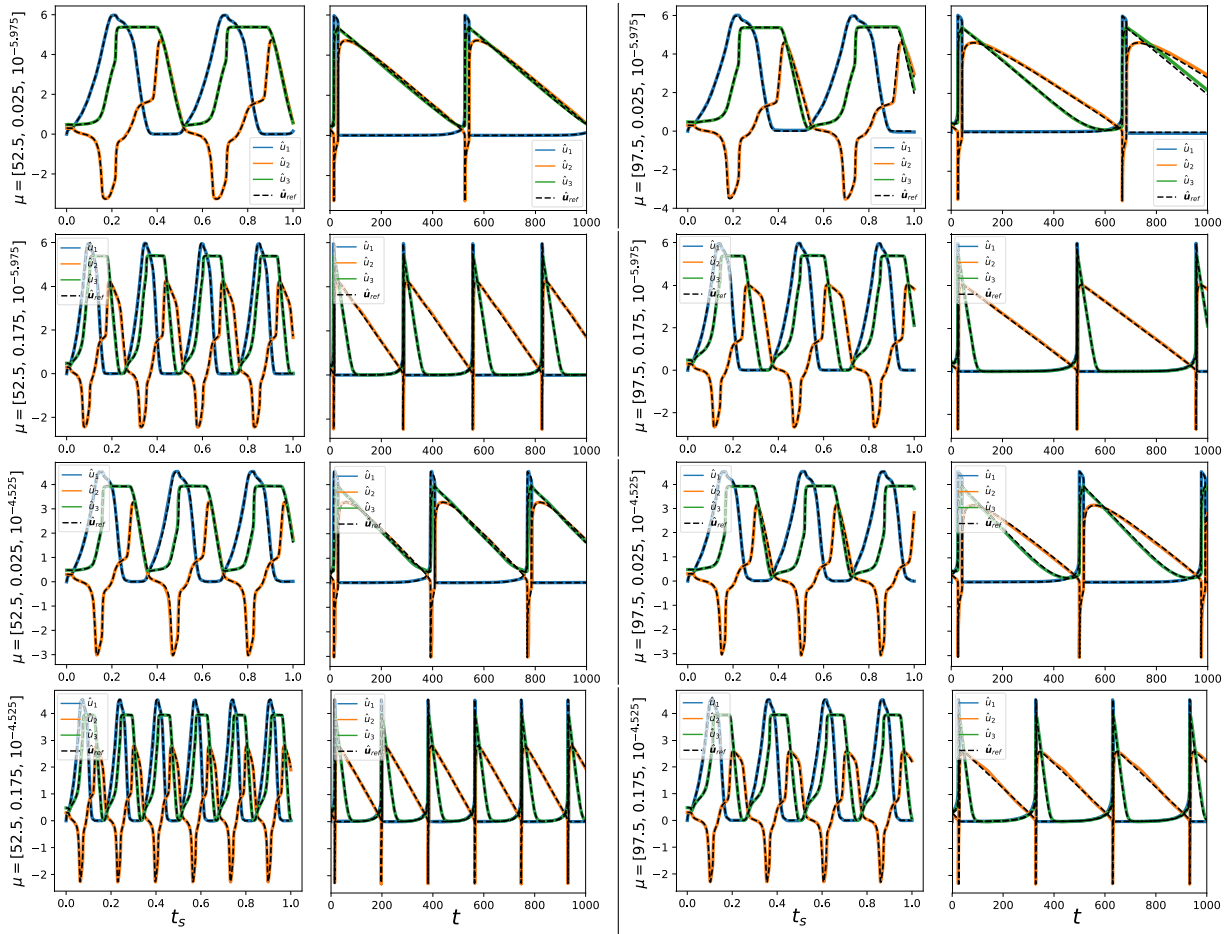
$$\begin{cases} \dot{u}_1 = \mu_1(u_2 - u_1 u_2 + u_1 - \mu_3 u_1^2), \\ \dot{u}_2 = \frac{1}{\mu_1}(-u_2 - u_1 u_2 + u_3), \\ \dot{u}_3 = \mu_2(u_1 - u_3), \\ \mathbf{u}(0) = (1, 2, 3) \end{cases} \quad (10)$$

Similar to the previous test case, the problem is an autonomous system that features a periodic solution. However, the problem originates from the description of a chemical reaction, thus it is key for the concentration  $\mathbf{u}$  to stay positive. Moreover, the values of  $\mathbf{u}$  span six orders of magnitude. The considered parameter

**TABLE 1** | Test case 1: Van der Pol.

$\mu$	Solver	tol	time [s]	# fev	# jev	# lu	$MSE_{t_i}$	$L^2$	$d_{peak}$
$10^{2.01}$	ROM	1/40	0.033	1000	0	0	4.13e-3	0.03520	0.662
	Radau	$10^{-2}$	0.046	2470	110	516	—	0.01447	0.811
		$10^{-1.5}$	0.041	2342	105	482	—	0.37252	14.59
$10^{2.67}$	ROM	1/40	0.033	1000	0	0	2.39e-3	0.02262	2.665
	Radau	$10^{-2}$	0.056	3316	133	688	—	0.00870	2.075
		$10^{-1.5}$	0.046	2591	118	576	—	0.53100	181.5
$10^{3.33}$	ROM	1/40	0.033	1000	0	0	5.94e-3	0.03098	19.27
	Radau	$10^{-2}$	0.072	3850	154	806	—	0.00742	8.083
		$10^{-1.5}$	0.052	2992	137	684	—	0.29987	655.2
$10^{3.99}$	ROM	1/40	0.033	1000	0	0	7.23e-3	0.04768	191.0
	Radau	$10^{-2}$	0.078	4418	156	938	—	0.04377	210.9
		$10^{-1.5}$	0.058	3337	148	768	—	0.17724	2822

Note: Comparison of computational cost and accuracy for the Radau solver and the neural ODE-based ROM on the test dataset. Refer to Section 4 for the definition of the metrics.



**FIGURE 7** | Test case 2: OREGO. On the left, in colors, the neural ODE predictions in  $t_s$ . On the right, in colors, the prediction is mapped to the original time  $t$ . In black is the reference solution. The state scale is logarithmic.

space is  $\Gamma = [50, 100] \times [0.002, 0.02] \times [10^{-6}, 10^{-4}]$ , which contains the value  $\mu = (77.27, 0.161, 8.375 \cdot 10^{-5})$  usually employed in literature for this test case. The training dataset is built by subsampling a discretization of  $\Gamma$  with a uniform grid of 6, 19, and 21 points in each direction, respectively. The validation dataset is based on a discretization of  $\Gamma$  in the midpoints of the training discretization. Reference solutions are computed with a tolerance of  $10^{-10}$ . We apply the following normalizations: parameter

normalization:  $\hat{\mu} = (\mu_1/77.27, \mu_2/0.161, \log_{10} \mu_3/5)$ ; state normalization  $\hat{u} = \log_{10} u$ ; dynamics normalization:  $\hat{f} = f$ . We test the performance of the ROM for eight parameters

$$\mu^{\text{test}} \in \{52.5, 97.5\} \times \{0.025, 0.175\} \times \{10^{-5.975}, 10^{-4.525}\}$$

Figure 7 shows that the ROM is able to accurately follow the reference solution and track its period. A quantitative comparison with the Radau solver is reported in Table 2. We remark that a



**TABLE 2** | Test case 2: OREGO.

$\mu$	Solver	tol	time [s]	# fev	# jev	# lu	$MSE_{t_s}$	$L^2$	$d_{\text{peak}}$
(52.5, 0.025, $10^{-5.975}$ )	ROM	$2 \cdot 10^{-4}$	0.037	1130	0	0	1.20e-3	1.19e-2	1.668
	Radau	$10^{-3}$	0.065	3777	167	656	—	1.31e-5	3.032
		$10^{-1.5}$	0.059	3213	165	664	—	<b>0.4691</b>	6.092
(97.5, 0.025, $10^{-5.975}$ )	ROM	$2 \cdot 10^{-4}$	0.034	1046	0	0	4.52e-3	7.75e-2	5.041
	Radau	$10^{-3}$	0.067	3445	159	600	—	2.12e-5	4.632
		$10^{-1.5}$	0.054	2815	141	556	—	<b>0.0099</b>	6.574
(52.5, 0.175, $10^{-5.975}$ )	ROM	$2 \cdot 10^{-4}$	0.065	1970	0	0	2.64e-3	5.67e-2	1.810
	Radau	$10^{-3}$	0.111	6590	310	1100	—	8.47e-5	1.115
		$10^{-1.5}$	0.095	5342	277	1068	—	0.0079	1.432
(97.5, 0.175, $10^{-5.975}$ )	ROM	$2 \cdot 10^{-4}$	0.057	1694	0	0	1.41e-3	3.08e-2	1.134
	Radau	$10^{-3}$	0.074	4649	212	744	—	1.43e-5	0.467
		$10^{-1.5}$	0.068	3654	197	712	—	0.0163	3.831
(52.5, 0.025, $10^{-4.525}$ )	ROM	$2 \cdot 10^{-4}$	0.048	1454	0	0	3.45e-3	5.14e-2	5.976
	Radau	$10^{-3}$	0.063	3895	183	666	—	3.73e-5	2.885
		$10^{-1.5}$	0.033	1814	92	364	—	<b>3.8838</b>	464.3
(97.5, 0.025, $10^{-4.525}$ )	ROM	$2 \cdot 10^{-4}$	0.049	1466	0	0	1.85e-3	6.27e-2	2.355
	Radau	$10^{-3}$	0.062	3809	172	658	—	8.83e-5	7.598
		$10^{-1.5}$	0.037	2021	100	392	—	<b>2.6084</b>	257.8
(97.5, 0.175, $10^{-4.525}$ )	ROM	$2 \cdot 10^{-4}$	0.088	2690	0	0	2.17e-3	1.17e-2	1.366
	Radau	$10^{-3}$	0.117	7148	319	1202	—	1.33e-5	1.100
		$10^{-1.5}$	0.098	5433	292	1054	—	0.0039	1.888
(52.5, 0.175, $10^{-4.525}$ )	ROM	$2 \cdot 10^{-4}$	0.059	1770	0	0	2.57e-3	2.28e-2	1.920
	Radau	$10^{-3}$	0.077	4781	217	814	—	2.35e-5	1.102
		$10^{-1.5}$	0.066	3635	196	712	—	0.0101	0.992

Note: Comparison of computational cost and accuracy for the Radau solver and the neural ODE-based ROM on the test dataset. Refer to Section 4 for the definition of the metrics. In bold, the error occurs when the solution is not positive.

tolerance of  $10^{-2}$  for the Radau solver did not produce a positive solution for all the problems in the test dataset. When this happens, the error is reported in bold in the table. The number of right-hand side evaluations of the explicit solver is less than half the one of the implicit one, showing that our method has made the system non-stiff. Also, the computational cost is consistently smaller than the one needed by the Radau solver, however, the latter has better accuracy. On the other hand, the distance of the peaks is comparable, proving that the ROM accurately captures the period of the system. In Figure 8 we present a comparison of the stiffness ratio for the original system and the one with simplified dynamics. On average, the time reparametrization reduces the stiffness by a couple of orders of magnitude, showing the effectiveness of our technique.

#### 4.2.1 | Parametric Initial Conditions

To prove that the parameter  $\mu$  is not limited to parametrizing the physical properties of the system, in this section we deal with the following variation of Equation (10)

$$\begin{cases} \dot{u}_1 = 77.27(u_2 - u_1 u_2 + u_1 - 8.375 \cdot 10^{-6} u_1^2), \\ \dot{u}_2 = 77.27^{-1}(-u_2 - u_1 u_2 + u_3), \\ \dot{u}_3 = 0.161(u_1 - u_3), \\ \mathbf{u}(0) = (\mu_1, \mu_2, \mu_3) \end{cases} \quad (11)$$

where  $\mu$  parameterizes the initial condition  $\mathbf{u}_0$ . The training samples for the  $\mu$  are in

$$\Gamma = \{(1, 2, 3), (5, 10, 15), (10, 20, 30)\}$$

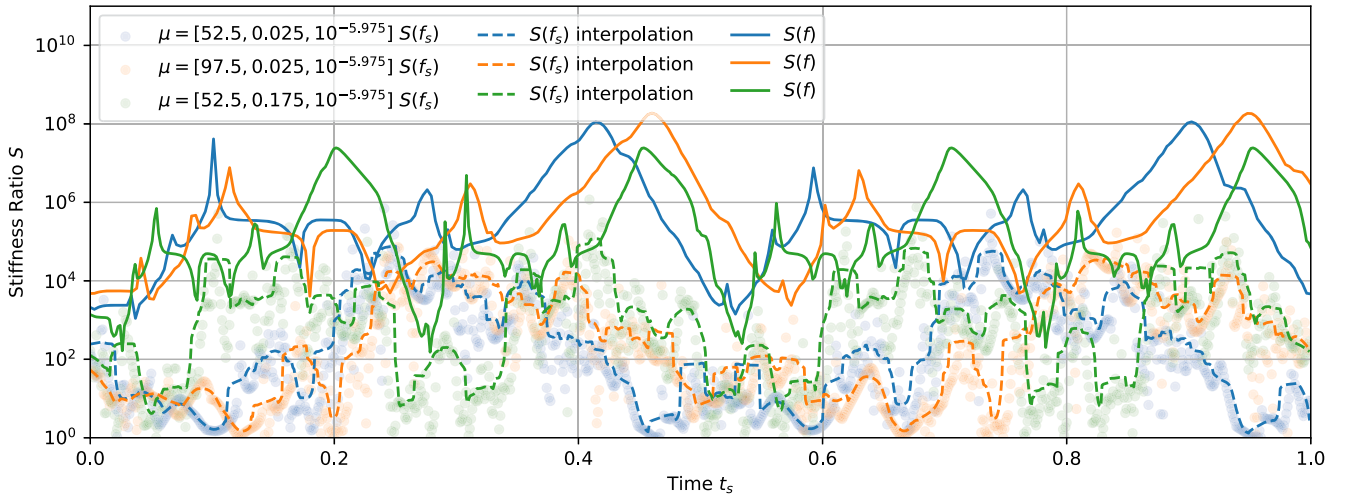
the normalization of the state space and the right-hand size are the same as before. The parameter  $\mu$  is not normalized. The problem is tested on  $\mu^{\text{test}} \in \{(2.5, 5, 7.5), (7.5, 15, 22.5)\}$ . The results are reported in Figure 9. Our model is able to capture the period of the system and the different number of peaks present in the considered time window. The  $L^2$ -error is 0.0408 and 0.0256 in the first and second case, respectively.

#### 4.3 | Test Case 3: ROBER Problem

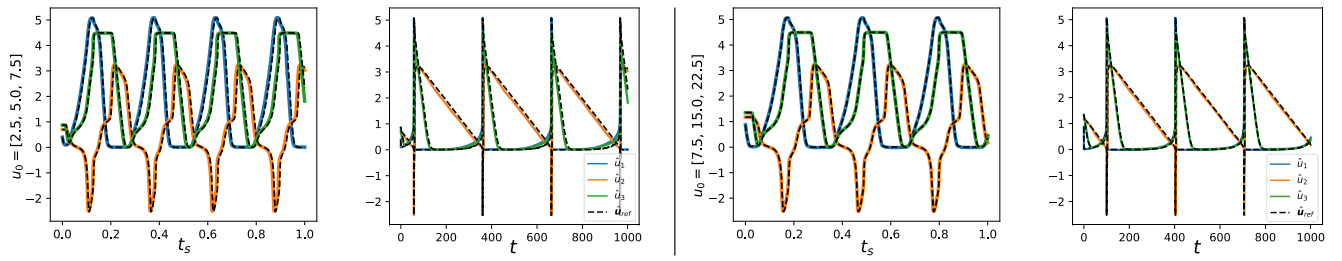
The ROBER problem is a stiff system of 3 nonlinear ODEs defined as follows:

$$\begin{cases} \dot{u}_1 = -\mu_1 u_1 + \mu_2 u_2 u_3, \\ \dot{u}_2 = \mu_1 u_1 - \mu_2 u_2 u_3 - \mu_3 u_2^2, \\ \dot{u}_3 = \mu_3 u_2^2, \\ \mathbf{u}(0) = (1, 0, 0) \end{cases} \quad (12)$$

The system describes the kinetics of an autocatalytic chemical reaction and is considered one of the most popular benchmarks for stiff solvers. When used in this context, the time integration window used is usually large, namely  $t \in [0, 10^{11}]$ , and the



**FIGURE 8** | Comparison of the stiffness ratio  $S$  of OREGO system in the original and simplified dynamics.



**FIGURE 9** | Test case 2: OREGO, parametric initial conditions. On the left, in colors, the neural ODE predictions in  $t_s$ . On the right, in colors, the prediction is mapped to the original time  $t$ . In black is the reference solution. The state scale is logarithmic.

parameters are  $\mu = [0.04, 10^4, 3 \cdot 10^7]$ . Indeed, codes might fail if  $u_2$  accidentally becomes negative, since it then tends to  $-\infty$ , causing overflow. This feature makes the ROBER problem ideal for testing the ability of our ROM to be accurate across different time scales.

We set  $\Gamma = [0.005, 0.05] \times [10^3, 10^5] \times [3 \cdot 10^7]$  and consider training dataset composed by a logarithmic discretization  $\Gamma$  with 16 uniform points and 31 logarithmically spaced points in the first two directions, respectively. The validation dataset is the interval midpoints of the discretized  $\Gamma$ . The reference solutions are computed using the Radau method, with absolute and relative tolerances set to  $\text{rtol} = 10^{-10}$ ,  $\text{atol} = 10^{-14}$ .

Data is normalized using the following functions:  $\hat{\mu} = (-\log_{10} \mu_1, \log_{10} \mu_2/4, \log_{10} \mu_3/7)$ ; state normalization  $\hat{u} = (u_1, u_2/10^{(\hat{\mu}_1 - \hat{\mu}_3)/2}, u_3)$ ; dynamics normalization:  $\hat{f} = f$ . The model is tested for four values of the parameters

$$\mu^{\text{test}} \in \{0.006, 0.49\} \times \{10^{3.025}, 10^{4.975}\}$$

Figure 10 shows that the ROM is able to accurately follow the reference solution across the different time scales. A quantitative comparison with the Radau solver is reported in Table 3. We remark that a tolerance of  $\text{rtol} = 10^{-4}$ ,  $\text{atol} = 10^{-7}$  for the Radau solver is among the smallest pair of tolerances that produce a positive solution for all the problems in the test dataset. The number of right-hand side evaluations of the explicit solver

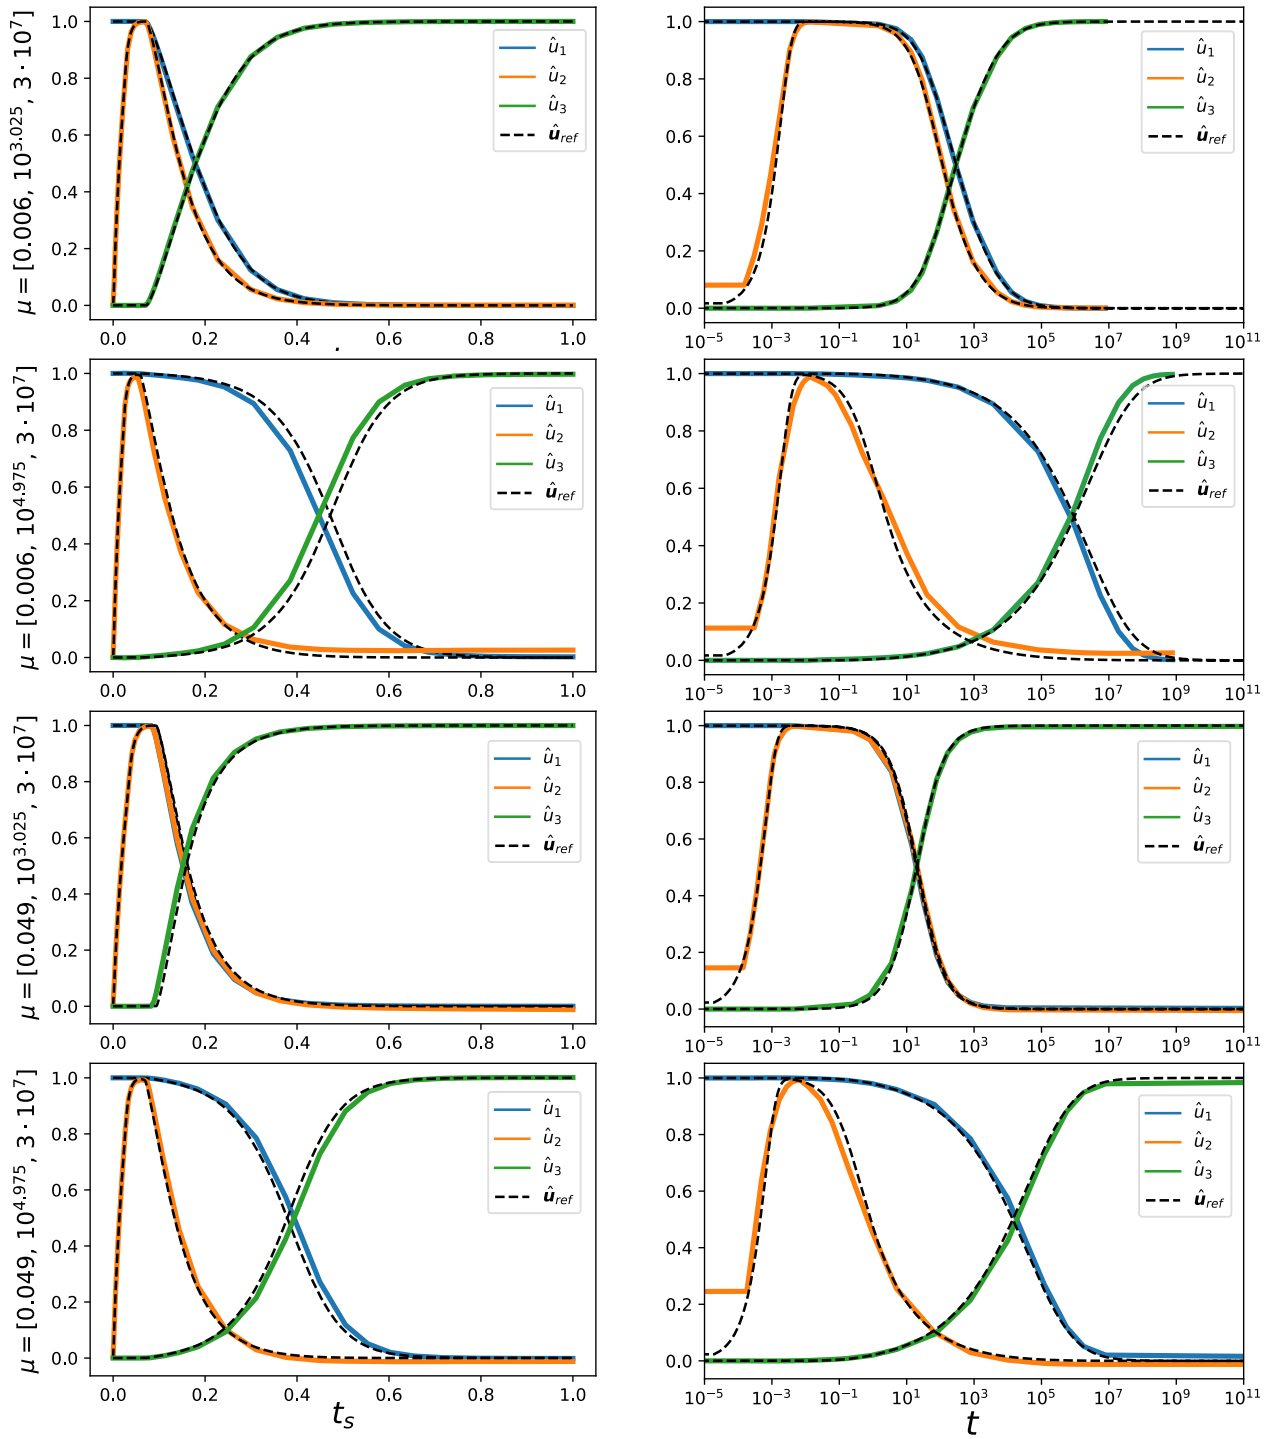
is almost one order of magnitude smaller than the implicit one, showing that our method has made the system non-stiff. Moreover, the computational cost is consistently smaller than the one needed by the Radau solver. However, the latter achieves a far smaller error. We highlight that the ROM sometimes fails to reach the correct final time, namely, the map to  $t$  sometimes stops at a time much smaller than  $10^{11}$ . Nevertheless, the ROM still captures the interesting part of the reaction, having issues only when the solution reaches a plateau. Similarly, we observe that the model is not accurate for  $t < 10^{-3}$  as it fails to follow the inflection point at  $t \sim 10^{-4}$  (cf. the right column of Figure 10). We show that these problems can be mitigated by reducing the tolerance of the RK45 solver to  $\text{tol} = 2 \cdot 10^{-5}$ , as shown in Figure 11.

#### 4.4 | Test Case 4: E5 Problem

The E5 problem consists of the following stiff system of 4 nonlinear ODEs:

$$\begin{cases} \dot{u}_1 = -\mu_1 u_1 + \mu_2 u_1 u_3, \\ \dot{u}_2 = \mu_1 u_1 - \mu_3 \mu_4 u_2 u_3, \\ \dot{u}_3 = \mu_1 u_1 - \mu_2 u_1 u_3 - \mu_3 \mu_4 u_2 u_3 + \mu_3 u_4, \\ \dot{u}_4 = \mu_2 u_1 u_3 - \mu_3 u_4, \\ u(0) = (1.76 \cdot 10^{-3}, 0, 0, 0) \end{cases} \quad (13)$$

The problem describes a chemical pyrolysis where  $u$  represents the concentration of the reactants. The largely different rates of



**FIGURE 10** | Test case 3: ROBER. On the left, in colors, the neural ODE predictions in  $t_s$ . On the right, in colors, the prediction is mapped to the original time  $t$ . In black is the reference solution.

reaction that occur in the same system are the cause of stiffness. As a test problem, it is usually integrated in a large time window, namely, we choose the interval  $t \in [0, 10^{11}]$ . This means the problem features widely different scales in time and space. The problem is particularly challenging. Indeed, the formulation (13) is affected by the cancellation of digits and it is recommended to use the relation  $\dot{u}_3 = \dot{u}_2 - \dot{u}_4$  to compute  $f(t, u)$ . Moreover, it has been shown that a very small tolerance on the absolute scalar


error must be used to obtain a reliable reference solution. Indeed, they are computed using the Radau method, with absolute and relative tolerances set to  $\text{atol} = 10^{-24}$ ,  $\text{rtol} = 10^{-10}$ .

We set  $\Gamma = [5 \cdot 10^{-10}, 5 \cdot 10^{-9}] \times [10^7, 10^8] \times \{1.13 \cdot 10^3\} \times \{10^6\}$  which contains the parameter  $\mu = (7.89 \cdot 10^{-10}, 1.1 \cdot 10^7, 1.13 \cdot 10^3, 10^6)$ , which is the usual choice for the benchmark. The training dataset is built by a discretization  $\Gamma$  with 11 logarithmically

**TABLE 3** | Test case 3: ROBER.

$\mu$	Solver	rtol (atol)	time [s]	# fev	# jev	# lu	MSE <sub>t<sub>s</sub></sub>	MSE
(0.006, 10 <sup>3.025</sup> , 3 · 10 <sup>7</sup> )	ROM	2 · 10 <sup>-4</sup>	0.008	230	0	0	1.21e-5	2.40e-4
	Radau	10 <sup>-4</sup> (10 <sup>-7</sup> )	0.019	927	44	210	—	6.72e-6
		10 <sup>-3</sup> (10 <sup>-6</sup> )	0.051	3055	132	798	—	9.58e11
(0.006, 10 <sup>4.975</sup> , 3 · 10 <sup>7</sup> )	ROM	2 · 10 <sup>-4</sup>	0.009	272	0	0	1.19e-3	1.11e-3
	Radau	10 <sup>-4</sup> (10 <sup>-7</sup> )	0.021	1100	52	240	—	5.74e-6
		10 <sup>-3</sup> (10 <sup>-6</sup> )	0.016	790	53	194	—	2.65e-5
(0.049, 10 <sup>3.025</sup> , 3 · 10 <sup>7</sup> )	ROM	2 · 10 <sup>-4</sup>	0.007	206	0	0	4.34e-4	4.97e-5
	Radau	10 <sup>-4</sup> (10 <sup>-7</sup> )	0.021	1081	46	256	—	4.49e-6
		10 <sup>-3</sup> (10 <sup>-6</sup> )	0.030	1787	73	406	—	1.9626
(0.049, 10 <sup>4.975</sup> , 3 · 10 <sup>7</sup> )	ROM	2 · 10 <sup>-4</sup>	0.009	260	0	0	4.46e-4	9.80e-4
	Radau	10 <sup>-4</sup> (10 <sup>-7</sup> )	0.024	1121	55	248	—	3.35e-6
		10 <sup>-3</sup> (10 <sup>-6</sup> )	0.014	795	50	198	—	1.50e-5

Note: Comparison of computational cost and accuracy for the Radau solver and the neural ODE-based ROM on the test dataset. Refer to Section 4 for the definition of the metrics.



**FIGURE 11** | Test case 3: ROBER. The prediction of the model with an integration tolerance of  $\tau_{\text{tol}} = 2 \cdot 10^{-5}$  mapped to the original time  $t$ . The smaller tolerance results in more accurate predictions, see Figure 10.

spaced points in the first two directions. The validation dataset is the interval midpoints of the discretized  $\Gamma$ . Data is normalized using the following functions:  $\hat{\mu} = \log_{10} \mu$ ; state normalization  $\hat{u} = \log_{10} u/10$ ; dynamics normalization:

$$\hat{f} = \begin{cases} f_i & \text{if } |f_i| < 2, \\ \text{sgn}(f_i)(\log(|f_i| - 1) + 2) & \text{otherwise} \end{cases}$$

The insight behind this normalization is simple: the data features very long tails due to the presence of large gradients. Thus, we apply a logarithmic transformation to these large values. Suitable constants are added to ensure continuity with the linear part of the transformation.

The model is tested for four values of the parameters

$$\mu^{\text{test}} \in \{5 \cdot 10^{-9.975}, 5 \cdot 10^{-9.025}\} \times \{10^{7.025}, 10^{7.975}\} \\ \times \{1.13 \cdot 10^3\} \times \{10^6\}$$

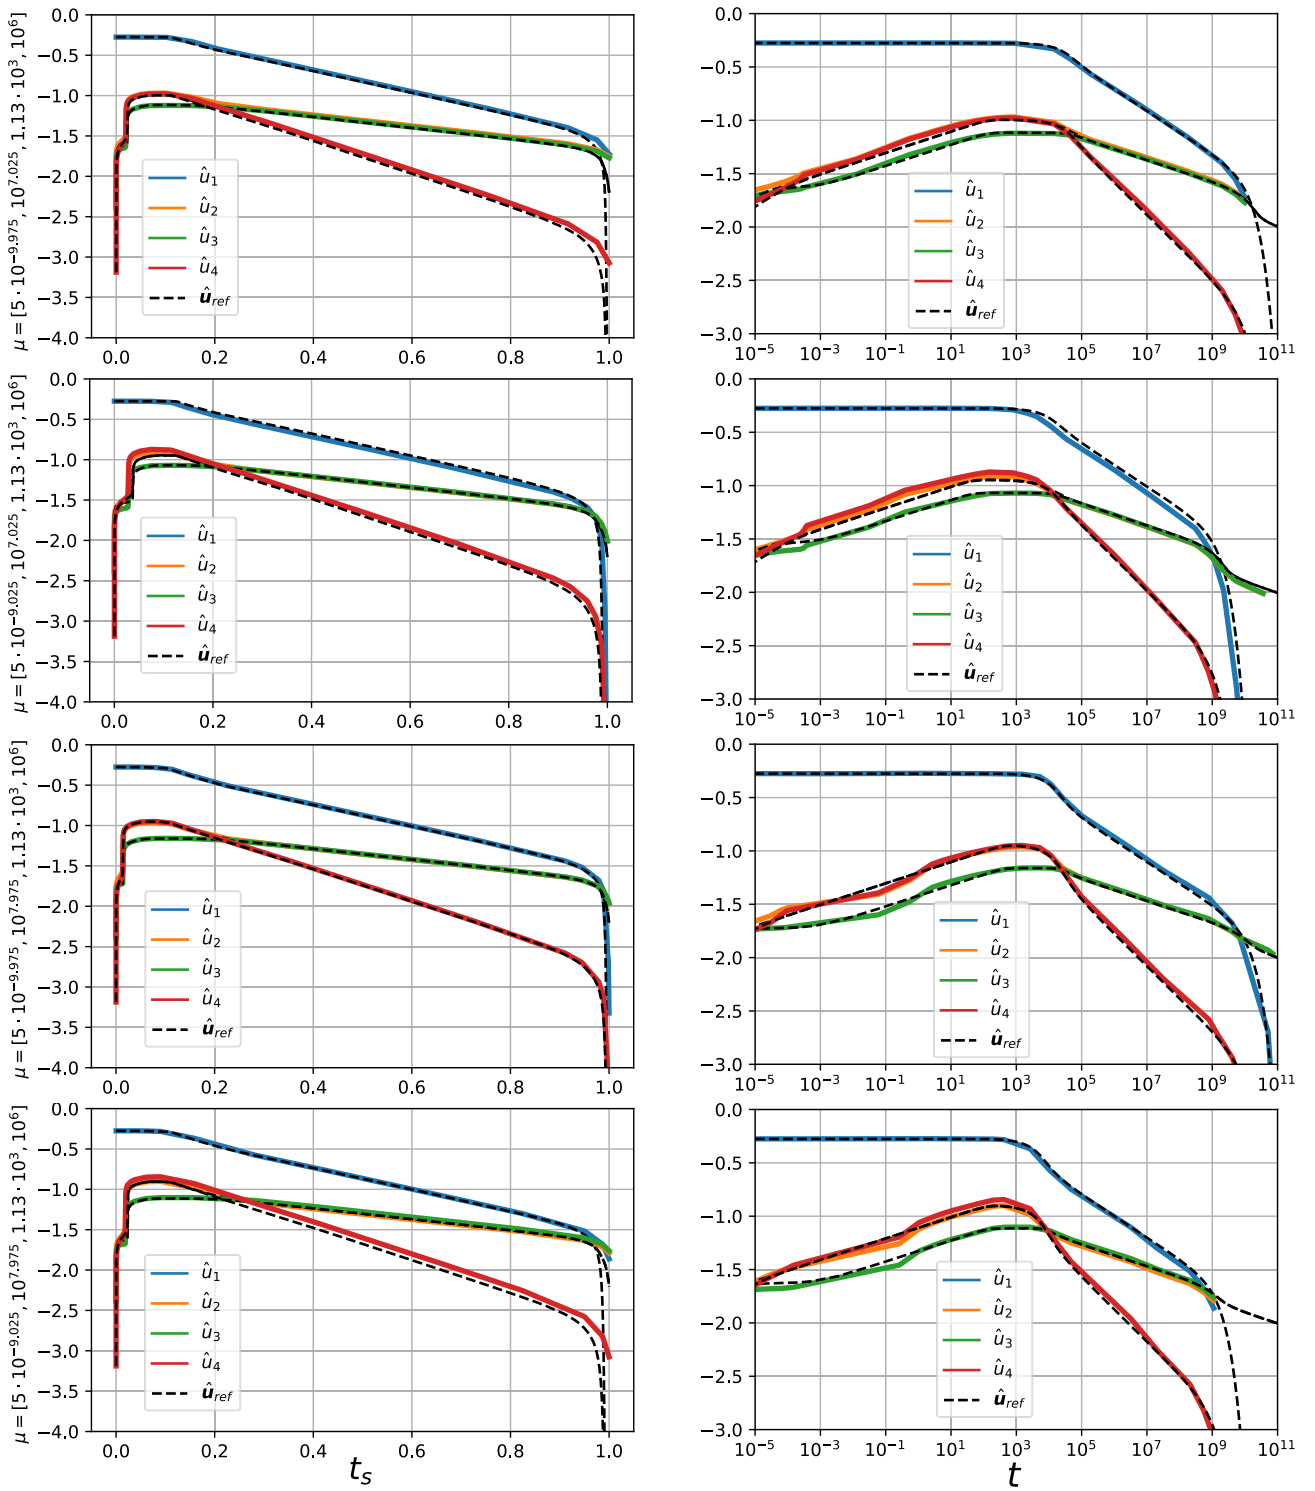
Figure 12 shows that the ROM is able to accurately follow the reference solution across the different time and space scales. A

quantitative comparison with the Radau solver is reported in Table 4. We remark that a tolerance of  $\text{rtol} = 10^{-6}$ ,  $\text{atol} = 10^{-20}$  for the Radau solver is among the smallest pair of tolerances that produce a positive and stable solution for all the problems in the test dataset. The number of right-hand side evaluations of the explicit solver is more than one order of magnitude smaller than the implicit one, showing that our method has indeed made the system non-stiff. The work precision tradeoff is particularly favorable for our ROM. Indeed, it has comparable accuracy at a much lower computational cost. The main limitation of our model is that it sometimes fails to reach the correct final time, namely, the map to  $t$  sometimes stops at a time much smaller than  $10^{11}$ . However, the ROM still captures the interesting part of the reaction, having issues only when the concentration plummets to values close to zero ( $< 10^{-20}$ ).

#### 4.5 | Test Case 5: POLLU Problem

To show the effectiveness of the proposed method on large state spaces, we test our methodology on the POLLU problem: as a





**FIGURE 12** | Test case 4: E5. On the left, in colors, the neural ODE predictions in  $t_s$ . On the right, in colors, the prediction is mapped to the original time  $t$ . In black is the reference solution. The state scale is logarithmic.

**TABLE 4** | Test case 4: E5.

$\mu$	Solver	rtol (atol)	time [s]	# fev	# jev	# lu	MSE <sub>t<sub>s</sub></sub>	MSE
$(5 \cdot 10^{-9.975}, 10^{7.025}, 1.13e3, 10^6)$	ROM	$2 \cdot 10^{-4}$	0.013	308	0	0	4.12e-2	5.06e-3
	Radau	$10^{-6}(10^{-20})$	0.064	4119	93	416	—	8.36e-4
$(5 \cdot 10^{-9.025}, 10^{7.025}, 1.13e3, 10^6)$	ROM	$2 \cdot 10^{-4}$	0.015	344	0	0	4.97e-2	4.00e-3
	Radau	$10^{-6}(10^{-20})$	0.065	4156	102	426	—	1.84e-2
$(5 \cdot 10^{-9.975}, 10^{7.975}, 1.13e3, 10^6)$	ROM	$2 \cdot 10^{-4}$	0.013	302	0	0	1.56e-2	4.80e-3
	Radau	$10^{-6}(10^{-20})$	0.063	4003	103	430	—	9.87e-4
$(5 \cdot 10^{-9.025}, 10^{7.975}, 1.13e3, 10^6)$	ROM	$2 \cdot 10^{-4}$	0.011	260	0	0	6.48e-2	5.53e-3
	Radau	$10^{-6}(10^{-20})$	0.068	4055	112	438	—	1.88e-2

Note: Comparison of computational cost and accuracy for the Radau solver and the neural ODE-based ROM on the test dataset. Refer to Section 4 for the definition of the metrics.

stiff system of 20 non-linear ODEs. The system is the chemical reaction part of the air pollution model developed at the Dutch National Institute of Public Health and Environmental Protection. The problem features 25 reactions ( $\mu \in \mathbb{R}^{25}$ ) and 20 reacting compounds ( $u \in \mathbb{R}^{20}$ ). Namely, we have that the system is defined by

$$f(u; \mu) = \begin{bmatrix} -r_1 - r_{10} - r_{14} - r_{23} - r_{24} + r_2 + r_3 + r_9 + r_{11} + r_{12} + r_{22} + r_{25} \\ -r_2 - r_3 - r_9 - r_{12} + r_1 + r_{21} \\ -r_{15} + r_1 + r_{17} + r_{19} + r_{22} \\ -r_2 - r_{16} - r_{17} - r_{23} + r_{15} \\ -r_3 + 2r_4 + r_6 + r_7 + r_{13} + r_{20} \\ -r_6 - r_8 - r_{14} - r_{20} + r_3 + 2r_{18} \\ -r_4 - r_5 - r_6 + r_{13} \\ r_4 + r_5 + r_6 + r_7 \\ -r_7 - r_8 \\ -r_{12} + r_7 + r_9 \\ -r_9 - r_{10} + r_8 + r_{11} \\ r_9 \\ -r_{11} + r_{10} \\ -r_{13} + r_{12} \\ r_{14} \\ -r_{18} - r_{19} + r_{16} \\ -r_{20} \\ r_{20} \\ -r_{21} - r_{22} - r_{24} + r_{23} + r_{25} \\ -r_{25} + r_{24} \end{bmatrix}$$

where the reactions  $r$  are connected to  $u$  and  $\mu^{(0)}$  by rates reported in Table 5. The initial condition is

$$u_0 = (0, 0.2, 0, 0.04, 0, 0, 0.1, 0.3, 0.01, 0, 0, 0, 0, 0, 0, 0, 0.007, 0, 0, 0)$$

It is usually considered that the time interval  $t \in [0, 60]$  is representative of the behavior of the reactants sufficiently. The reference solution is obtained by using a tolerance of  $10^{-10}$ . Full details about the model can be found in [83]. The parameter space  $\Gamma$  is built by changing the three components of  $\mu$  that most impact the system. Namely, a preliminary sensibility analysis has shown that the components 4, 6, and 14 of  $\mu$  are the most relevant, that

is, they produced the largest variation in the solution  $u$  when perturbed. Hence, we define the parameter space as

$$\Gamma = \bigtimes_{i \notin \{4,6,14\}} \{\mu_i^{(0)}\} \times \bigtimes_{i \in \{4,6,14\}} \left[ \frac{1}{2}\mu_i^{(0)}, 2\mu_i^{(0)} \right]$$

where  $\mu^{(0)}$  is the reference parameter value used for the system, reported in Table 5. The intervals are discretized with 16 uniform points in each coordinate direction. The validation dataset is made by the midpoints of the training datasets.

The model is tested for two parameters  $\mu^{\text{test}(1)}, \mu^{\text{test}(2)}$  which are equal to  $\mu^{(0)}$  apart from the three most relevant components, namely  $\mu_i^{\text{test}(1)} = 0.525\mu_i^{(0)}$ ,  $\mu_i^{\text{test}(2)} = 1.975\mu_i^{(0)}$ ,  $i = 4, 6, 14$ . Results are shown in Figure 13. The ROM is able to follow the reference solution across all the reactions. A quantitative comparison with the Radau solver is reported in Table 6. The number of right-hand side evaluations of the explicit solver is similar to the implicit one. On the other hand, each evaluation of the right-hand side, each evaluation of the Jacobian, and each LU decomposition is more costly than in the previous test cases since  $N_u$  is larger. However, the cost of evaluating the neural network is similar to the previous test cases since the dimensions of the networks are comparable. Indeed, the cost of the ROM is still smaller than the cost of the Radau solver while keeping comparable accuracy. These results are of particular significance because they prove that as the state space gets larger, our methodology becomes more cost-effective.

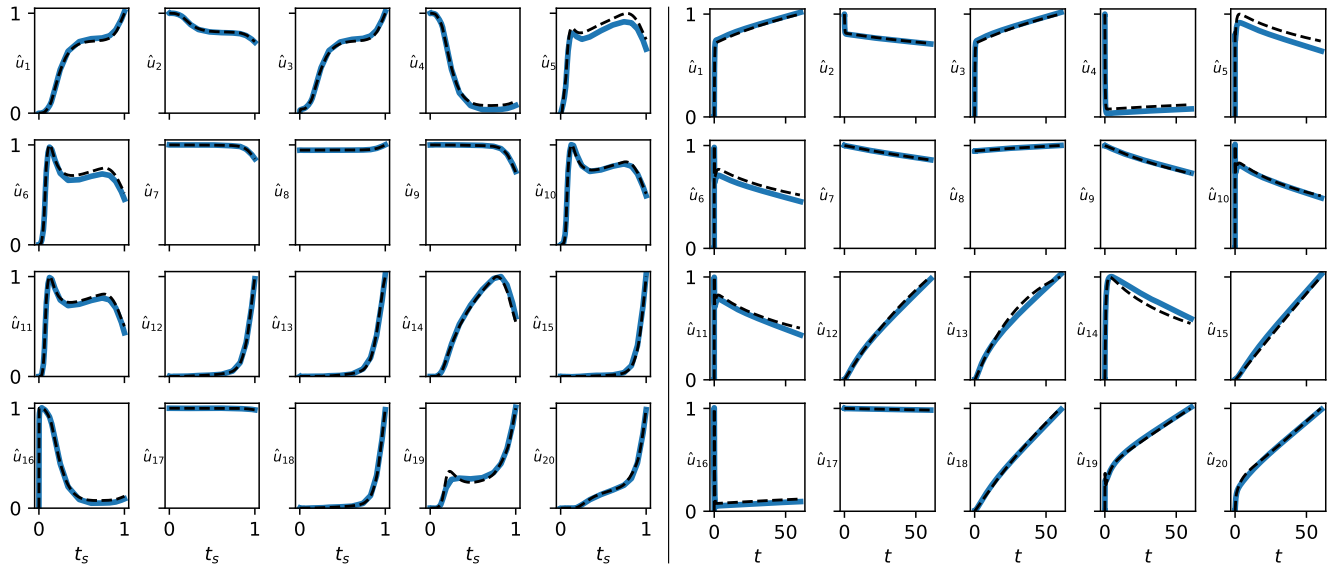
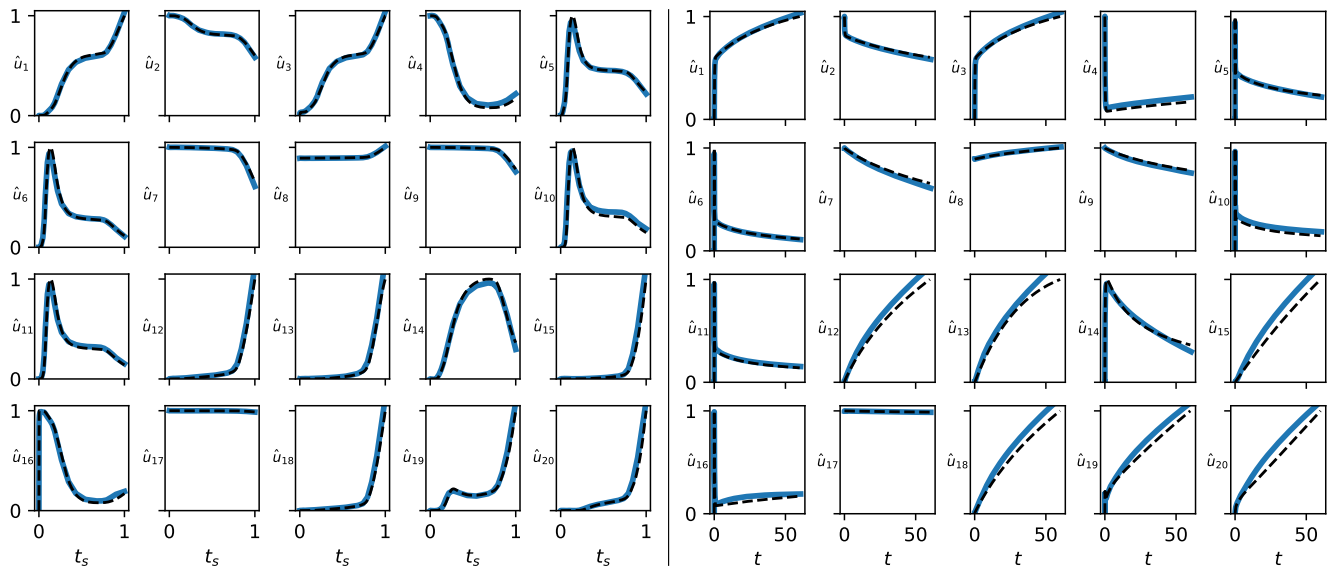
## 5 | Conclusions

In this work, we have developed a novel methodology to tackle some kinds of stiffness in neural ODEs. The approach hinges on a suitable time reparametrization of the system that significantly reduces the stiffness, making it possible to efficiently apply explicit solvers to the neural ODEs. The construction of the time reparametrization is general and completely data-driven, making it ideal for an application to ROMs.

By leveraging the intrinsic capabilities of neural ODEs to handle continuous-time data, we capture the dynamics of stiff systems effectively. The accuracy, robustness, and efficiency of the methodology were tested in five famous benchmark test cases from the literature. Namely, when applied to periodic systems,

**TABLE 5** | Reference reaction rates for the POLLU model that define  $\mu^{(0)}$ .

Reaction	$r$	$\mu^{(0)}$	Reaction	$r$	$\mu^{(0)}$	Reaction	$r$	$\mu^{(0)}$
1	$\mu_1 u_1$	0.350e0	10	$\mu_{10} u_{11} u_1$	0.900e4	19	$\mu_{19} u_{16}$	0.444e12
2	$\mu_2 u_2 u_4$	0.266e2	11	$\mu_{11} u_{13}$	0.220e-1	20	$\mu_{20} u_{17} u_6$	0.124e4
3	$\mu_3 u_5 u_2$	0.120e5	12	$\mu_{12} u_{10} u_2$	0.120e5	21	$\mu_{21} u_{19}$	0.210e1
4	$\mu_4 u_7$	0.860e-3	13	$\mu_{13} u_{14}$	0.188e1	22	$\mu_{22} u_{19}$	0.578e1
5	$\mu_5 u_7$	0.820e-3	14	$\mu_{14} u_1 u_6$	0.163e5	23	$\mu_{23} u_1 u_4$	0.474e-1
6	$\mu_6 u_7 u_6$	0.150e5	15	$\mu_{15} u_3$	0.480e7	24	$\mu_{24} u_{19} u_1$	0.178e4
7	$\mu_7 u_9$	0.130e-3	16	$\mu_{16} u_4$	0.350e-3	25	$\mu_{25} u_{20}$	0.312e1
8	$\mu_8 u_9 u_6$	0.240e5	17	$\mu_{17} u_4$	0.175e-1			
9	$\mu_9 u_{11} u_2$	0.165e5	18	$\mu_{18} u_{16}$	0.100e9			

Test case 5.1,  $\mu = \mu^{\text{test}(1)}$ Test case 5.2,  $\mu = \mu^{\text{test}(2)}$ **FIGURE 13** | Test case 5: POLLU. On the left, the neural ODE predictions in  $t_s$ . On the right, the prediction is mapped to the original time  $t$ . In blue the neural ODE prediction of  $\hat{u}$ , in black the reference solution.

**TABLE 6** | Test case 5: POLLU.

$\mu$	Solver	tol	time [s]	# fev	# jev	# lu	$MSE_{t_s}$	MSE
$\mu_i = 0.525\mu_i^{(0)}, i = 4, 6, 14$	ROM	$10^{-4}$	0.006	194	0	0	6.89e-4	9.80e-4
	Radau	$10^{-4}$	0.011	182	10	50	—	3.81e-3
		$10^{-3}$	0.011	149	10	40	—	0.2077
$\mu_i = 1.975\mu_i^{(0)}, i = 4, 6, 14$	ROM	$10^{-4}$	0.007	212	0	0	1.16e-3	1.69e-3
	Radau	$10^{-4}$	0.017	268	14	70	—	9.03e-4
		$10^{-3}$	0.014	203	12	54	—	4.63e-3

Note: Comparison of computational cost and accuracy for the Radau solver and the neural ODE-based ROM on the test dataset. Refer to Section 4 for the definition of the metrics.

the model generalizes to data well beyond the time training interval. We attribute this characteristic primarily to the autonomous nature of the neural ODE and the time reparametrization model. When applied to chemical equations, our model always maintained the positivity of the concentration, even when applied to problems suffering from numerical instability. Moreover, the work/precision tradeoff of our models often proved comparable or favorable when confronted with a state-of-the-art Runge-Kutta implicit solver of the Radau II A kind. In a few cases where this was not observed, our model provided a computationally cheaper alternative for approximate solutions while preserving concentration positivity, which is not possible with the Radau solver for small tolerances.

In conclusion, our work highlights the promising potential of neural ODEs in creating efficient and accurate ROMs for stiff ODEs. Stiffness characterized by reasonably large eigenvalues but very high frequency remains a challenge for our approach, and it is an active topic of research. Future research is directed toward improving the accuracy of the time mapping that, as of right now, stands as the major accuracy bottleneck of the ROM. We also aim to test this framework to handle the semi-discrete formulation of PDEs, and other higher-dimensional systems. Moreover, techniques like model distillation, network pruning, and better hyperparameter tuning could further reduce the computational cost of our approach.

### Author Contributions

**Matteo Caldana:** conceptualization, methodology, software, validation, formal analysis, investigation, data curation, visualization, writing – original draft. **Jan S. Hesthaven:** conceptualization, methodology, resources, writing – review and editing, supervision, project administration.

### Acknowledgments

Matteo Caldana has been partially funded by PRIN2020 n. 20204LN5N5 funded by the Italian Ministry of Universities and Research (MUR). Matteo Caldana is a member of INdAM-GNCS. Open access publishing facilitated by Politecnico di Milano, as part of the Wiley - CRUI-CARE agreement.

### Conflicts of Interest

The authors declare no conflicts of interest.

### Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

### References

1. P. Benner, S. Gugercin, and K. Willcox, “A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems,” *SIAM Review* 57, no. 4 (2015): 483–531.
2. M. Peirlinck, F. S. Costabal, J. Yao, et al., “Precision Medicine in Human Heart Modeling: Perspectives, Challenges, and Opportunities,” *Biomechanics and Modeling in Mechanobiology* 20 (2021): 803–831.
3. T. Cui, Y. M. Marzouk, and K. E. Willcox, “Data-Driven Model Reduction for the Bayesian Solution of Inverse Problems,” *International Journal for Numerical Methods in Engineering* 102, no. 5 (2015): 966–990.
4. M. Frangos, Y. Marzouk, K. Willcox, and B. van Bloemen Waanders, “Surrogate and Reduced-Order Modeling: A Comparison of Approaches for Large-Scale Statistical Inverse Problems,” In *Large-Scale Inverse Problems and Quantification of Uncertainty* (Wiley, 2010), 123–149.
5. T. Lähivaara, A. Malehm, A. Pasanen, L. Kärkkäinen, J. M. Huttunen, and J. S. Hesthaven, “Estimation of Groundwater Storage From Seismic Data Using Deep Learning,” *Geophysical Prospecting* 67, no. 8 (2019): 2115–2126.
6. Y. Maday, A. T. Patera, J. D. Penn, and M. Yano, “A Parameterized-Background Data-Weak Approach to Variational Data Assimilation: Formulation, Analysis, and Application to Acoustics,” *International Journal for Numerical Methods in Engineering* 102, no. 5 (2015): 933–965.
7. T. Bui-Thanh, K. Willcox, and O. Ghattas, “Parametric Reduced-Order Models for Probabilistic Analysis of Unsteady Aerodynamic Applications,” *AIAA Journal* 46, no. 10 (2008): 2520–2529.
8. P. Chen, A. Quarteroni, and G. Rozza, “Reduced Basis Methods for Uncertainty Quantification,” *SIAM/ASA Journal on Uncertainty Quantification* 5, no. 1 (2017): 813–869.
9. B. Sudret and A. Der Kiureghian, “Stochastic Finite Element Methods and Reliability: A State-of-the-Art Report,” 2000 Department of Civil and Environmental Engineering, University of California.
10. A. Manzoni, A. Quarteroni, and G. Rozza, “Shape Optimization for Viscous Flows by Reduced Basis Methods and Free-Form Deformation,” *International Journal for Numerical Methods in Fluids* 70, no. 5 (2012): 646–670.
11. M. Bergmann and L. Cordier, “Optimal Control of the Cylinder Wake in the Laminar Regime by Trust-Region Methods and POD Reduced-Order Models,” *Journal of Computational Physics* 227, no. 16 (2008): 7813–7840.



12. F. Negri, G. Rozza, A. Manzoni, and A. Quarteroni, "Reduced Basis Method for Parametrized Elliptic Optimal Control Problems," *SIAM Journal on Scientific Computing* 35, no. 5 (2013): A2316–A2340.
13. S. S. Ravindran, "A Reduced-Order Approach for Optimal Control of Fluids Using Proper Orthogonal Decomposition," *International Journal for Numerical Methods in Fluids* 34, no. 5 (2000): 425–448.
14. P. Benner, M. Ohlberger, A. Cohen, and K. Willcox, *Model Reduction and Approximation: Theory and Algorithms* (SIAM, 2017).
15. J. S. Hesthaven, G. Rozza, B. Stamm, et al., *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, vol. 590 (Springer, 2016).
16. A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems* (SIAM, 2005).
17. P. Benner, V. Mehrmann, and D. C. Sorensen, *Dimension Reduction of Large-Scale Systems*, vol. 45 (Springer, 2005).
18. T. Bui-Thanh, K. Willcox, and O. Ghattas, "Model Reduction for Large-Scale Systems With High-Dimensional Parametric Input Space," *SIAM Journal on Scientific Computing* 30, no. 6 (2008): 3270–3288.
19. J. S. Hesthaven, C. Pagliantini, and G. Rozza, "Reduced Basis Methods for Time-Dependent Problems," *Acta Numerica* 31 (2022): 265–345.
20. C. Prud'Homme, D. V. Rovas, K. Veroy, et al., "Reliable Real-Time Solution of Parametrized Partial Differential Equations: Reduced-Basis Output Bound Methods," *Journal of Fluids Engineering* 124, no. 1 (2002): 70–80.
21. M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera, "An 'Empirical Interpolation' Method: Application to Efficient Reduced-Basis Discretization of Partial Differential Equations," *Comptes Rendus Mathématique* 339, no. 9 (2004): 667–672.
22. C. Canuto, T. Tonn, and K. Urban, "A Posteriori Error Analysis of the Reduced Basis Method for Nonaffine Parametrized Nonlinear PDEs," *SIAM Journal on Numerical Analysis* 47, no. 3 (2009): 2001–2022.
23. S. Chaturantabut and D. C. Sorensen, "Nonlinear Model Reduction via Discrete Empirical Interpolation," *SIAM Journal on Scientific Computing* 32, no. 5 (2010): 2737–2764.
24. K. Lee and K. T. Carlberg, "Model Reduction of Dynamical Systems on Nonlinear Manifolds Using Deep Convolutional Autoencoders," *Journal of Computational Physics* 404 (2020): 108973.
25. J. S. Hesthaven and S. Ubbiali, "Non-Intrusive Reduced Order Modeling of Nonlinear Problems Using Neural Networks," *Journal of Computational Physics* 363 (2018): 55–78.
26. Q. Wang, J. S. Hesthaven, and D. Ray, "Non-Intrusive Reduced Order Modeling of Unsteady Flows Using Artificial Neural Networks With Application to a Combustion Problem," *Journal of Computational Physics* 384 (2019): 289–307.
27. Q. Wang, N. Ripamonti, and J. S. Hesthaven, "Recurrent Neural Network Closure of Parametric POD-Galerkin Reduced-Order Models Based on the Mori-Zwanzig Formalism," *Journal of Computational Physics* 410 (2020): 109402.
28. S. L. Brunton, B. R. Noack, and P. Koumoutsakos, "Machine Learning for Fluid Mechanics," *Annual Review of Fluid Mechanics* 52 (2020): 477–508.
29. S. Fresca, L. Dede', and A. Manzoni, "A Comprehensive Deep Learning-Based Approach to Reduced Order Modeling of Nonlinear Time-Dependent Parametrized PDEs," *Journal of Scientific Computing* 87 (2021): 1–36.
30. R. Maulik, B. Lusch, and P. Balaprakash, "Reduced-Order Modeling of Advection-Dominated Systems With Recurrent Neural Networks and Convolutional Autoencoders," *Physics of Fluids* 33, no. 3 (2021): 037106.
31. V. Oommen, K. Shukla, S. Goswami, R. Dingreville, and G. E. Karniadakis, "Learning Two-Phase Microstructure Evolution Using Neural Operators and Autoencoder Architectures," *NPJ Computational Materials* 8, no. 1 (2022): 190.
32. P. R. Vlachas, G. Arampatzis, C. Uhler, and P. Koumoutsakos, "Multiscale Simulations of Complex Systems by Learning Their Effective Dynamics," *Nature Machine Intelligence* 4, no. 4 (2022): 359–366.
33. L. Pegolotti, M. R. Pfaller, N. L. Rubio, et al., "Learning Reduced-Order Models for Cardiovascular Simulations With Graph Neural Networks," *Computers in Biology and Medicine* 168 (2024): 107676.
34. F. Pichi, B. Moya, and J. S. Hesthaven, "A Graph Convolutional Autoencoder Approach to Model Order Reduction for Parametrized PDEs," *Journal of Computational Physics* 501 (2024): 112762.
35. Z. Li, D. Z. Huang, B. Liu, and A. Anandkumar, "Fourier Neural Operator With Learned Deformations for Pdes on General Geometries," *Journal of Machine Learning Research* 24, no. 388 (2023): 1–26.
36. Z. Li, N. Kovachki, K. Azizzadenesheli, et al., "Fourier Neural Operator for Parametric Partial Differential Equations," (2020), *arXiv Preprint arXiv:2010.08895*.
37. G. Wen, Z. Li, K. Azizzadenesheli, A. Anandkumar, and S. M. Benson, "U-FNO – An Enhanced Fourier Neural Operator-Based Deep-Learning Model for Multiphase Flow," *Advances in Water Resources* 163 (2022): 104180.
38. L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis, "Learning Nonlinear Operators via DeepONet Based on the Universal Approximation Theorem of Operators," *Nature Machine Intelligence* 3, no. 3 (2021): 218–229.
39. S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control* (Cambridge University Press, 2022).
40. P. J. Schmid, "Dynamic Mode Decomposition of Numerical and Experimental Data," *Journal of Fluid Mechanics* 656 (2010): 5–28.
41. S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering Governing Equations From Data by Sparse Identification of Nonlinear Dynamical Systems," *Proceedings of the National Academy of Sciences* 113, no. 15 (2016): 3932–3937.
42. K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, "Data-Driven Discovery of Coordinates and Governing Equations," *Proceedings of the National Academy of Sciences* 116, no. 45 (2019): 22445–22451.
43. Y. Liu, J. N. Kutz, and S. L. Brunton, "Hierarchical Deep Learning of Multiscale Differential Equation Time-Steppers," *Philosophical Transactions of the Royal Society A* 380, no. 2229 (2022): 20210200.
44. P. Conti, M. Guo, A. Manzoni, and J. S. Hesthaven, "Multi-Fidelity Surrogate Modeling Using Long Short-Term Memory Networks," *Computer Methods in Applied Mechanics and Engineering* 404 (2023): 115811.
45. M. Guo and J. S. Hesthaven, "Data-Driven Reduced Order Modeling for Time-Dependent Problems," *Computer Methods in Applied Mechanics and Engineering* 345 (2019): 75–99.
46. R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural Ordinary Differential Equations," *Advances in Neural Information Processing Systems* 31 (2018), [http://papers.nips.cc/paper\\_files/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html](http://papers.nips.cc/paper_files/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html).
47. G. D. Portwood, P. P. Mitra, M. D. Ribeiro, et al., "Turbulence Forecasting via Neural ODE," (2019), *arXiv:1911.05180 [physics]*.
48. F. Regazzoni, L. Dede, and A. Quarteroni, "Machine Learning for Fast and Reliable Solution of Time-Dependent Differential Equations," *Journal of Computational Physics* 397 (2019): 108852.
49. F. Regazzoni, S. Pagani, M. Salvador, L. Dede, and A. Quarteroni, "Learning the Intrinsic Dynamics of Spatio-Temporal Processes Through

- Latent Dynamics Networks,” *Nature Communications* 15, no. 1 (2024): 1834–2024.
50. G. Wanner and E. Hairer, *Solving Ordinary Differential Equations II*, vol. 375 (Springer Berlin Heidelberg New York, 1996).
51. S. Wang, X. Yu, and P. Perdikaris, “When and Why PINNs Fail to Train: A Neural Tangent Kernel Perspective,” *Journal of Computational Physics* 449 (2022): 110768.
52. W. Ji, W. Qiu, Z. Shi, S. Pan, and S. Deng, “Stiff-Pinn: Physics-Informed Neural Network for Stiff Chemical Kinetics,” *Journal of Physical Chemistry A* 125, no. 36 (2021): 8098–8106.
53. S. Kim, W. Ji, S. Deng, Y. Ma, and C. Rackauckas, “Stiff Neural Ordinary Differential Equations,” *Chaos (Woodbury, N.Y.)* 31, no. 9 (2021): 093122.
54. M. De Florio, E. Schiassi, and R. Furfaro, “Physics-Informed Neural Networks and Functional Interpolation for Stiff Chemical Kinetics,” *Chaos (Woodbury, N.Y.)* 32, no. 6 (2022): 063107.
55. R. Anantharaman, A. Abdelrehim, A. Jain, et al., “Stably Accelerating Stiff Quantitative Systems Pharmacology Models: Continuous-Time Echo State Networks as Implicit Machine Learning,” *IFAC-PapersOnLine* 55, no. 23 (2022): 1–6, 9th IFAC Conference on Foundations of Systems Biology in Engineering FOSBE 2022.
56. P. Zhang and R. Sankaran, “Autoencoder Neural Network for Chemically Reacting Systems,” *Journal of Machine Learning for Modeling and Computing* 3, no. 4 (2022): 1–28.
57. S. Goswami, A. D. Jagtap, H. Babaee, B. T. Susi, and G. E. Karniadakis, “Learning Stiff Chemical Kinetics Using Extended Deep Neural Operators,” *Computer Methods in Applied Mechanics and Engineering* 419 (2024): 116674.
58. N. Barral, T. Taddei, and I. Tifouti, “Registration-Based Model Reduction of Parameterized Pdes With Spatio-Parameter Adaptivity,” *Journal of Computational Physics* 499 (2024): 112727.
59. M. Khamlich, F. Pichi, M. Girfoglio, A. Quaini, and G. Rozza, “Optimal Transport-Based Displacement Interpolation With Data Augmentation for Reduced Order Modeling of Nonlinear Dynamical Systems,” (2024), arXiv Preprint arXiv:2411.08750.
60. R. L. Van Heyningen, N. C. Nguyen, P. Blonigan, and J. Peraire, “Adaptive Model Reduction of High-Order Solutions of Compressible Flows via Optimal Transport,” *International Journal of Computational Fluid Dynamics* 37, no. 6 (2023): 541–563.
61. A. A. Loya, D. A. Serino, and Q. Tang, “Structure-Preserving Neural Ordinary Differential Equations for Stiff Systems (2025),” arXiv Preprint arXiv:2503.01775.
62. L. R. Petzold, L. O. Jay, and J. Yen, “Numerical Solution of Highly Oscillatory Ordinary Differential Equations,” *Acta Numerica* 6 (1997): 437–483.
63. W. Huang and B. Leimkuhler, “The Adaptive Verlet Method,” *SIAM Journal on Scientific Computing* 18, no. 1 (1997): 239–256.
64. U. M. Ascher and S. Reich, “On Some Difficulties in Integrating Highly Oscillatory Hamiltonian Systems,” in *Computational Molecular Dynamics: Challenges, Methods, Ideas: Proceedings of the 2nd International Symposium on Algorithms for Macromolecular Modelling, Berlin, May 21–24, 1997* (Springer, 1999), 281–296.
65. A. Richardson and J. Finn, “Symplectic Integrators With Adaptive Time Steps,” *Plasma Physics and Controlled Fusion* 54, no. 1 (2011): 014004.
66. D. A. Serino, A. A. Loya, J. W. Burby, I. G. Kevrekidis, and Q. Tang, “Intelligent Attractors for Singularly Perturbed Dynamical Systems,” (2024), arXiv Preprint arXiv:2402.15839.
67. F. Mazzia, F. Iavernaro, and C. Magherini, “Test Set for Initial Value Problem Solvers,” (2003), Department of Mathematics, University of Bari.
68. O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu, *The Finite Element Method: Its Basis and Fundamentals* (Elsevier, 2005).
69. C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang, *Spectral Methods: Evolution to Complex Geometries and Applications to Fluid Dynamics* (Springer Science & Business Media, 2007).
70. B. Cockburn, G. E. Karniadakis, and C.-W. Shu, *Discontinuous Galerkin Methods: Theory, Computation and Applications*, vol. 11 (Springer Science & Business Media, 2012).
71. J. S. Hesthaven and T. Warburton, *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications* (Springer Science & Business Media, 2007).
72. E. Hairer and G. Wanner, “Stiff Differential Equations Solved by Radau Methods,” *Journal of Computational and Applied Mathematics* 111, no. 1 (1999): 93–111.
73. J. R. Dormand and P. J. Prince, “A Family of Embedded Runge-Kutta Formulae,” *Journal of Computational and Applied Mathematics* 6, no. 1 (1980): 19–26.
74. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016), <http://www.deeplearningbook.org>.
75. A. Zhang, Z. C. Lipton, M. Li, and A. J. Smola, *Dive Into Deep Learning* (Cambridge University Press, 2023).
76. L. F. Shampine and S. Thompson, “Stiff Systems,” *Scholarpedia* 2, no. 3 (2007): 2855.
77. I. Loshchilov and F. Hutter, “Decoupled Weight Decay Regularization,” (2017), arXiv Preprint arXiv:1711.05101.
78. D. Hendrycks and K. Gimpel, “Gaussian Error Linear Units (GELUs),” (2016), arXiv Preprint arXiv:1606.08415.
79. S. Elfving, E. Uchibe, and K. Doya, “Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning,” *Neural Networks* 107 (2018): 3–11.
80. A. Howard, M. Sandler, G. Chu, et al., “Searching for MobileNetV3,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision* (IEEE Computer Society, 2019), 1314–1324.
81. A. Paszke, S. Gross, S. Chintala, et al., “Automatic Differentiation in Pytorch,” (2017), NIPS 2017 Workshop on Autodiff.
82. P. Virtanen, R. Gommers, T. E. Oliphant, et al., “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods* 17 (2020): 261–272.
83. J. G. Verwer, “Gauss–Seidel Iteration for Stiff ODEs From Chemical Kinetics,” *SIAM Journal on Scientific Computing* 15, no. 5 (1994): 1243–1250.