

# Learning Time Delay Systems with Neural Ordinary Differential Equations

Xunbi A. Ji\* and Gábor Orosz\*,\*\*

\* *Department of Mechanical Engineering,*

\*\* *Department of Civil and Environmental Engineering,*

*University of Michigan, Ann Arbor, MI 48109 USA*

*(e-mail: [xunbij@umich.edu](mailto:xunbij@umich.edu), [orosz@umich.edu](mailto:orosz@umich.edu))*

**Abstract:** A novel way of using neural networks to learn the dynamics of time delay systems from sequential data is proposed. A neural network with trainable delays is used to approximate the right hand side of a delay differential equation. We relate the delay differential equation to an ordinary differential equation by discretizing the time history and train the corresponding neural ordinary differential equation (NODE) to learn the dynamics. An example on learning the dynamics of the Mackey-Glass equation using data from chaotic behavior is given. After learning both the nonlinearity and the time delay, we demonstrate that the bifurcation diagram of the neural network matches that of the original system.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

*Keywords:* neural ordinary differential equations, time delay systems, time delay neural networks, trainable time delays, Mackey-Glass equation.

## 1. INTRODUCTION

Artificial neural networks have been thriving in many fields over the past decade and their power in approximating and generalizing the underlying input-output relationships has been demonstrated for multiple examples. Multiple research studies related neural networks to dynamical systems; see, for instance, Kumpati et al. (1990); Pei et al. (2013); Brunton and Kutz (2019).

Knowledge from dynamical systems can help to design and train neural networks and many ideas in recurrent neural networks were inspired by dynamical systems principles. The deep equilibrium model (DEM) in Bai et al. (2019) represents the discrete-time map between hidden states. Since a globally stable autonomous dynamical system approaches a fix point, one can directly solve for the fix point via root-finding. Neural ordinary differential equation (NODE) in Chen et al. (2018) expresses the time derivative (i.e., the right hand side of the differential equation), while the evolution of the network for a given time period provides the map from the initial condition (input) and the terminal state (target). The depth of the network is related to the simulation time of the corresponding ordinary differential equation. Neural networks can be used to learn the dynamical systems. For example, Wong et al. (2021) provides an example of using discrete-time recurrent neural network to learn traffic dynamics while Li et al. (2005) uses a continuous-time recurrent neural network to express the right hand side of a differential equation. Also, see Turan and Jäschke (2021); Rahman et al. (2022); Koch et al. (2022) for some recent applications of NODEs on learning nonlinear dynamical systems.

In case of delayed dynamical systems, one shall include time delays in the neural networks. As a matter of fact, incorporating time delays in the neural networks may increase the capabilities of the networks. For example, in the

neural delay differential equation (NDDE) approach proposed by Zhu et al. (2021), adding the delayed state, the network is able to approximate trajectories that intersect in the original state space. One may increase the dimension of the system by defining extra states and construct an augmented ordinary differential equation (ANODE), as in Dupont et al. (2019), but this may not be able to capture the physical meaning of time delays in data-based systems.

Although time delays have been introduced to neural networks in both discrete time (Waibel et al. (1989); Zhang et al. (2022)) and continuous time (Marcus and Westervelt (1989); Stelzer et al. (2021)), these delays were fixed rather than treated as parameters to be learned. In this work, we consider a network with trainable time delays, a concept that was originally proposed in Ji et al. (2020). Different from NDDE, we use an ordinary differential equation to approximate the delayed dynamics, and thus, methods developed for training NODE can be directly applied. This construction is based on the idea of approximating the infinitesimal generator through discretizing the history as in Breda et al. (2014). For data-based models, our method bridges the gap between infinite-dimensional time delay systems and discrete-time maps obtained via full discretization.

The layout of the paper is as follows. In Section 2, we relate delay differential equations (DDEs) to ordinary differential equations (ODEs) via discretization of the history. In Section 3, we introduce trainable time delay neural networks and construct the NODE by discretizing the history. In Section 4 we illustrate the training algorithm with a specific loss function. We provide an example of learning the Mackey-Glass equation from simulation data in Section 5 and conclude the results in Section 6.

## 2. DISCRETIZATION OF THE HISTORY

Consider an autonomous time delay system

$$\dot{x}(t) = g(x_t), \quad (1)$$

where  $x \in \mathbb{R}^n$  and  $x_t \in \mathcal{X}$ . Here  $\mathcal{X} = C([-\tau_{\max}, 0], \mathbb{R}^n)$  is the space of continuous functions, i.e.,  $x_t(\vartheta) := x(t + \vartheta)$ ,  $\vartheta \in [-\tau_{\max}, 0]$ , and  $\tau_{\max} > 0$  is the maximum delay. The system (1) contains the functional  $g: \mathcal{X} \rightarrow \mathbb{R}^n$  and it can be re-written as an operator differential equation:

$$\dot{x}_t = \mathcal{G}(x_t), \quad (2)$$

where the operator  $\mathcal{G}: \mathcal{X} \rightarrow \mathcal{X}$  is defined as

$$\mathcal{G}(\varphi) = \begin{cases} g(\varphi), & \vartheta = 0, \\ \frac{d\varphi}{d\vartheta}, & \vartheta \in [-\tau_{\max}, 0), \end{cases} \quad (3)$$

see Stépán (1989); Breda et al. (2014).

We discretize in the variable  $\vartheta$  using a mesh of size  $M + 1$  and define the vector

$$X(t) = \begin{bmatrix} X_1(t) \\ X_2(t) \\ \vdots \\ X_{M+1}(t) \end{bmatrix} = \begin{bmatrix} x(t) \\ x(t-h) \\ \vdots \\ x(t-Mh) \end{bmatrix} \quad (4)$$

where  $h = \tau_{\max}/M$ . This way we can approximate (2) with the ordinary differential equation

$$\dot{X}(t) = G(X(t)), \quad (5)$$

where the function  $G: \mathbb{R}^{n(M+1)} \rightarrow \mathbb{R}^{n(M+1)}$  is given by

$$G(X) = \begin{bmatrix} \tilde{g}(X) \\ D_M X \end{bmatrix}. \quad (6)$$

Here  $\tilde{g}: \mathbb{R}^{n(M+1)} \rightarrow \mathbb{R}^n$  is the approximation of the functional  $g: \mathcal{X} \rightarrow \mathbb{R}^n$ . For example, considering

$$g(x_t) = f(x(t), x(t - (j + \alpha)h), x(t - \tau_{\max})) \quad (7)$$

with  $j \in \{0, 1, \dots, M - 1\}$ ,  $\alpha \in (0, 1)$ , and  $f: \mathbb{R}^{3n} \rightarrow \mathbb{R}^n$ , one may use linear interpolation to obtain

$$x(t - (j + \alpha)h) = (1 - \alpha)x(t - jh) + \alpha x(t - (j + 1)h), \quad (8)$$

yielding

$$\tilde{g}(X) = f(X_1, (1 - \alpha)X_{j+1} + \alpha X_{j+2}, X_{M+1}), \quad (9)$$

The matrix  $D_M \in \mathbb{R}^{(n-1)(M+1) \times n(M+1)}$  represents the numerical differentiation scheme used to approximate the time derivative of past states. For instance, choosing the forward Euler method to approximate the time derivative

$$\dot{x}(t - jh) = \frac{1}{h}(x(t - (j - 1)h) - x(t - jh)), \quad (10)$$

gives

$$\dot{X}_{j+1}(t) = \frac{1}{h}(X_j(t) - X_{j+1}(t)), \quad (11)$$

yielding

$$D_M = \frac{1}{h} \begin{bmatrix} I & -I & & & \\ & I & -I & & \\ & & \ddots & \ddots & \\ & & & I & -I \end{bmatrix}, \quad (12)$$

where  $I \in \mathbb{R}^{n \times n}$  is the  $n$ -dimensional unit matrix and only the nonzero blocks of  $D_M$  are spelled out.

In summary, (5,6) discretizes (2,3) in the history variable  $\vartheta$ . As illustrated in Fig. 1, we still maintain the continuous dynamics in time  $t$ . This enables us to propagate the state forward with high accuracy, which will be beneficial for learning. Note that, however, data samples will only be available with sampling time  $h$ .




Fig. 1. Illustration of discretization of a DDE in history.

## 3. TIME DELAY NEURAL NETWORKS

Time delay neural networks (TDNNs) were first proposed for speech recognition in Waibel et al. (1989). In particular, networks with input delays can be expressed by

$$\begin{aligned} z_1^t &= f(W_1 z_0^t + b_1), \\ z_l^t &= f_l(W_l z_{l-1}^t + b_l), \quad l \geq 2, \end{aligned} \quad (13)$$

where  $l = 1, 2, \dots, L$  denotes the layers of the neural network while  $t$  denotes time. The inputs for the network are the delayed values of the signal  $z_0^t$  which can be collected in the vector

$$\mathbf{z}_0^t = \begin{bmatrix} z_0^{t-\tau_1} \\ \vdots \\ z_0^{t-\tau_d} \end{bmatrix}, \quad (14)$$

with input delays  $0 \leq \tau_1 < \dots < \tau_d$ . The dimension of  $z_l$  depends on the number of neurons in each layer, and the network output  $z_L$  has the same dimension as  $z_0$ . The matrix  $W_l$  and the vector  $b_l$  are referred to as weight and bias respectively, while  $f_l$  denotes the element-wise nonlinear function in each layer. Building on our previous work in Ji et al. (2021), we use trainable time delay neural networks approximate the right hand side of the delay differential equations. We shape the network indirectly by propagating the states of the approximate system forward via numerical simulation (using time steps much smaller than  $h$ ) and comparing the result to data (which is available with sampling time  $h$ ). This provides an efficient way of learning the nonlinear dynamics and the delays.

In order to use the construction (13,14) for learning the dynamics (5,6), we define  $z_0^t = x(t)$ . We assume that the number of delays  $d$  in the system is known (or at least the upper limit of this number is known). Note that, however, the values of the delays  $0 \leq \tau_1 < \dots < \tau_d$  are not known. Then we define a matrix which allows us to approximate the delayed states  $z_0^{t-\tau_i}$  in (14) with the components  $X_j$  in (4), namely

$$\mathbf{z}_0^t = PX(t), \quad (15)$$

where  $P \in \mathbb{R}^{d \times n(M+1)}$  and since typically  $d \ll M$ ,  $P$  is a wide matrix. For the case of the linear interpolation given in example (7,8,9) we have

$$P = \begin{bmatrix} I & & & \\ & (1 - \alpha)I & \alpha I & \\ & & & \\ & & & I \end{bmatrix} \quad (16)$$

where the block  $(1 - \alpha)I$  is the  $j + 1$ -st column and the block  $\alpha I$  is the  $j + 2$ -nd column. Indeed, one may use higher-order interpolation methods which result in more nonzero blocks in  $P$ .

The output of the neural network (13) is set to be the derivative

$$\dot{x}(t) = z_L^t = \text{net}(PX(t)), \quad (17)$$

and thus the neural ordinary differential equation becomes

$$\dot{X}(t) = \widehat{G}(X(t)), \quad (18)$$

where the function  $\widehat{G}: \mathbb{R}^{n(M+1)} \rightarrow \mathbb{R}^{n(M+1)}$  is given by

$$\widehat{G}(X) = \begin{bmatrix} \text{net}(PX) \\ D_M X \end{bmatrix}, \quad (19)$$

cf. (5,6). We remark that to obtain  $D_M$ , instead of the forward Euler method we will use the central difference method which yields

$$D_M = \frac{1}{2h} \begin{bmatrix} I & O & -I \\ \ddots & \ddots & \ddots \\ & I & O & -I \\ 2I & & & -2I \end{bmatrix}, \quad (20)$$

where  $O \in \mathbb{R}^{n \times n}$  is the  $n$ -dimensional zero matrix while the other nonzero blocks are not spelled out; cf. (12). Note that the last row in  $D_M$  is still derived from forward Euler method.

When training the NODE (18,19), the weights  $W_l$  and biases  $b_l$  in (13) as well as the delays  $\tau_i$  are updated at every iteration. The update in  $\tau_i$  iteratively changes  $P$  defined in (15). The loss function and the training methods are introduced in the next section.

#### 4. TRAINING THE NETWORKS WITH DELAYS

After constructing of the NODE (18,19), one can simulate the resulting differential equations with given initial condition  $X(0)$ . The corresponding solution is denoted by  $\widehat{X}(t)$  and we have  $\hat{x}(t) = \widehat{X}_1(t)$ . During the learning process this is compared to the data, which is obtained as the solution  $x(t)$  of (1). We assume that the data is available with the sampling time  $h$  and we compare the solutions along the time horizon  $T = Nh, N \in \mathbb{N}$ . We construct the loss function using the one-norm:

$$L = \sum_{j=1}^N \|\hat{x}(jh) - x(jh)\|_1, \quad (21)$$

which depends on the initial history  $X(0)$  and the trainable parameters  $\theta = \{W_l, b_l, \tau_i\}, l = 1, \dots, L, i = 1, \dots, d$ .

Learning algorithms require the gradient

$$g = \frac{\partial L}{\partial \theta}, \quad (22)$$

to iterate the parameters so that (21) decreases. The gradient of the loss with respect to delay  $\tau_i$  is given by

$$\frac{\partial L}{\partial \tau_i} = \frac{\partial L}{\partial x(-\tau_i)} \frac{\partial x(-\tau_i)}{\partial \tau_i} = \frac{\partial L}{\partial x(-\tau_i)} (-\dot{x}(-\tau_i)), \quad (23)$$

where  $\frac{\partial L}{\partial x(-\tau_i)}$  is the gradient with respect to the initial delayed state and  $\dot{x}(-\tau_i)$  is the time derivative of  $x$  at  $-\tau_i$ . The latter one is approximated from data again using central differences while the former one can be calculated via back propagation or through the adjoint variable method as in Chen et al. (2018) after the forward simulation pass is established. The gradient of the loss with respect to other parameters  $\frac{\partial L}{\partial W_l}$  and  $\frac{\partial L}{\partial b_l}$  can be also obtained via back propagation.

With the availability of the gradient information, one may update the parameters using adaptive moment estimation (Adam, Kingma and Ba (2015)). The formula for updating each individual parameter is given by

$$\theta_{k+1} = \theta_k - \frac{\eta}{\sqrt{\hat{v}_k + \epsilon}} \hat{m}_k, \quad (24)$$

with

$$\begin{aligned} \hat{m}_k &= \frac{m_k}{1 - \beta_1^k}, \\ \hat{v}_k &= \frac{v_k}{1 - \beta_2^k}, \end{aligned} \quad (25)$$

where  $m_k$  and  $v_k$  are the first moment estimate and second moment estimate, respectively. The updates for  $m_k$  and  $v_k$  are given by

$$\begin{aligned} m_k &= \beta_1 m_{k-1} + (1 - \beta_1) g_k, \\ v_k &= \beta_2 v_{k-1} + (1 - \beta_2) g_k^2, \end{aligned} \quad (26)$$

using the gradient information  $g_k$  at iteration  $k$ ; cf. (22). The parameters  $\eta, \epsilon, \beta_1, \beta_2$  in (24,25,26) can be tuned by the user to achieve better performance; see Kingma and Ba (2015).

The loss function described in (21) is based on one simulation with a given initial history. When using multiple trajectories (from multiple initial histories) for one update, the sum of simulation losses for all the trajectories in the batch is considered. Then the gradient is calculated by taking the average of the gradients in the batch. We remark that when updating the delay parameters, we limit the delay value to be between 0 and  $\tau_{\max}$  using  $\tau_i = \max(\min(\tau_i, \tau_{\max}), 0)$ . Since the delay values are explicitly learned, one may in fact use the first row of the NODE (18,19) as a neural delay differential equation (NDDE) whose form is similar to (1).

#### 5. MACKEY-GLASS EQUATION

In this section, we examine the ability of the tools developed above while learning the dynamics of the Mackey-Glass equation (Mackey and Glass (1977)), a scalar autonomous time delay system of the form

$$\dot{x}(t) = \frac{\beta x(t - \tau)}{1 + (x(t - \tau))^\delta} - \gamma x(t), \quad (27)$$

with  $\beta = 4, \gamma = 2, \delta = 9.65$ . The system exhibits different qualitative behaviors as the time delay  $\tau$  is varied. For small values of  $\tau$ , there exists a globally stable equilibrium. This equilibrium loses stability as  $\tau$  is increased via a supercritical Hopf bifurcation, giving rise to stable limit cycle oscillations. With  $\tau$  increased further, the limit cycle loses stability via a supercritical period doubling bifurcation, and a period doubling cascade eventually leads to chaotic behavior; see the bifurcation diagram in Fig. 7(a). We set  $\tau = 1$ , which yields chaotic behavior, and generate 100 trajectories using constant initial histories  $x(t) \equiv c, t \in [-\tau_{\max}, 0]$  where  $c = 0.5 + i/99, i = 0, \dots, 99$  and  $\tau_{\max} = 1.5$ . The simulations are done in MATLAB using dde23.

According to (5,6) one may also generate data using the approximate ODE

$$\dot{X}(t) = \begin{bmatrix} \frac{\beta X_r(t)}{1 + (X_r(t))^\delta} - \gamma X_1(t) \\ D_M X(t) \end{bmatrix}. \quad (28)$$




Fig. 2. Data obtained for training and testing. One of the 100 trajectories is shown with initial history  $c = 0.5$ .

This may serve as the baseline for the performance of the NODE explained below for which we use  $h = 0.05$  yielding  $M = \tau_{\max}/h = 30$  and  $r = \tau/h + 1 = 21$  and matrix (20).

When constructing the NODE (18,19), we use two-hidden-layer trainable time delay neural networks with 5 neurons in each hidden layer and choose  $\tanh(\cdot)$  as the nonlinear activation function. Thus, following (17), we can write

$$\dot{X}(t) = \begin{bmatrix} W_3 \tanh(W_2 \tanh(W_1 P X(t) + b_1) + b_2) \\ D_M X(t) \end{bmatrix}. \quad (29)$$

where  $W_1 \in \mathbb{R}^{5 \times 2}$ ,  $W_2 \in \mathbb{R}^{5 \times 5}$ ,  $W_3 \in \mathbb{R}^{1 \times 5}$ ,  $b_1, b_2 \in \mathbb{R}^{5 \times 1}$  and the  $\tanh$  is applied element-wise. The matrix  $P$  is chosen to be the same as in (16). We focus on analyzing the case with knowledge of the number of delays. Recall that  $\tau_1 = 0$  and only the second delay  $\tau_2 > 0$  is learned. Both (28) and (29) are simulated using `ode45` in MATLAB. Finally, we extract the NDDE

$$\dot{x}(t) = W_3 \tanh\left(W_2 \tanh\left(W_1 \begin{bmatrix} x(t) \\ x(t - \tau_2) \end{bmatrix} + b_1\right) + b_2\right), \quad (30)$$

utilizing the first row of (29) and simulate the corresponding equations using `dde23`.

When generating data via (27) we drop the transients, that is, we only consider the trajectories in the time domain  $t \in [10, 20]$  as indicated in Fig. 2. Then we sample the data with time step  $h = 0.05$  and use the first 141 samples ( $t \in [10, 17]$ ) for training and the next 60 samples ( $t \in [17, 20]$ ) for testing. When calculating the loss (21), we set the simulation horizon to be  $N = 10$  steps (i.e.,  $T = Nh = 0.5$ ). From the 141 samples in each trajectory, we take segments of 31 samples as initial histories and obtain predictions for next 10 steps. Thus, 101 input-output pairs (i.e.,  $[x(t - \tau_{\max}), \dots, x(t)] \rightarrow [x(t + h), \dots, x(t + T)]$ ,  $t = 11.50, 11.55, \dots, 17$ ) are generated. This way, we obtain 10100 pairs from the 100 trajectories for training.

We train the network for 2000 iterations and use the batch size 1000, such that 10 out of the 101 pairs from each trajectory are used for each update. The delay  $\tau_2$  is initialized uniformly between  $[0, \tau_{\max}]$ ,  $b_1$  and  $b_2$  are initialized as zeros, while  $W_1, W_2$  and  $W_3$  are initialized using Glorot initialization (Glorot and Bengio (2010)). In the Adam updating rule (24,25,26) the learning rate  $\eta$  is set to 0.01 while  $\beta_1, \beta_2, \epsilon$  are set to 0.9, 0.999,  $10^{-8}$ , respectively; see Kingma and Ba (2015). From different parameter initializations, the networks may converge to different local minima. We plot the delay learning path and training loss of the networks with three different learned delays as function of the iteration number in Fig. 3. For




Fig. 3. Delay values (top) and training loss (bottom) for three different parameter initializations. The final learned delays are  $\tau_2 \approx 0, 0.09, 1$ , respectively. Note that the network parameters may jump over a local minimum and converge to another one later due to the constant learning rate, which may cause the temporary increase in training loss.




Fig. 4. Simulation of three trained NODEs with different learned delays for training (top) and testing (bottom).

the two paths leading to small delay values, the training loss stops decreasing very early, while for the third path the loss keeps decreasing with the iterations. This may be used to detect undesired local minima and to filter out the bad initialization at the early stages of training.

The predictions of three trained networks are shown in Fig. 4. The simulation of the network with correctly learned delay (blue solid curve) fits the data (red and black dashed curves) very well and makes good short-term predictions about the future. The NODEs with incorrect delays (magenta and orange curves) perform poorly both for training and for testing.

We also study the effect of the simulation horizon on the training process. We trained 100 networks with different initial parameters using 3-step, 10-step and 40-step hori-




Fig. 5. Histograms of the learned delay at iteration 2000 when using different prediction horizons as indicated.




Fig. 6. Comparing the nonlinearity of the Mackey-Glass DDE to that of the learned NDDE (extracted from the NODE for  $\tau_2 = 1.00$  in Fig. 4). (a), (b): top view of the nonlinear surfaces; (c): top view of the error surface, with training data indicated by gray dots; (d): limit cycle predicted by the NDDE (green curve) compared with that of DDE (dashed red curve) for  $\tau_1 = 0, \tau_2 = 0.8$ .

zons and plot the distribution of the learned delays at iteration 2000 in Fig. 5. We observe a trade-off between the performance and training time. Increasing the horizon can reduce the chance of being trapped at small delay values, but the training process is slower for longer horizons.

To illustrate on how well the learned network performs, we compare the nonlinearity obtained through NDDE (30) with that of the original DDE (27). The coloring in Fig. 6(a) shows  $\dot{x}(t)$  (i.e., the right hand side of the DDE) as a function of  $x(t)$  and  $x(t - \tau)$ . The black curve indicates  $\dot{x} = 0$  and its intersection with the 45 degree red line gives the equilibria of the system, namely,  $x(t) \equiv 0$  and  $x(t) \equiv 1$ . Panel (b) shows the network output  $\hat{x}(t)$  (i.e., the right hand side of the NDDE) as a function of  $x(t)$  and  $x(t - \tau)$ , while panel (c) depicts the difference between the nonlinearities of the DDE and the NDDE. The gray trajectory indicates the chaotic attractor of the DDE which covers a significant portion of the plane.

Since the time delay and the nonlinearity are learned separately, we can vary the time delay while keeping the net-

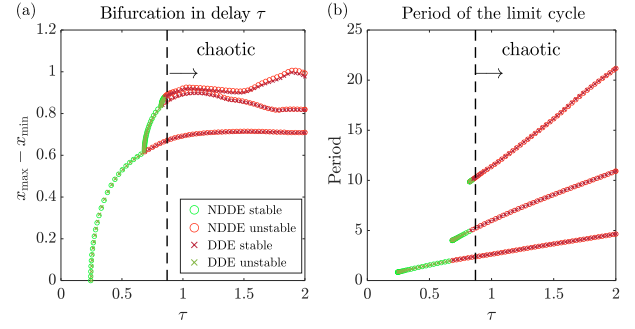


Fig. 7. Bifurcations diagram of the learned NDDE and the Mackey-Glass DDE. The equilibrium  $x(t) \equiv 1$  loses its stability via Hopf bifurcation at  $\tau = 0.24$  and the arising stable periodic orbit undergoes a period-doubling cascade (with the first two bifurcations located at  $\tau = 0.61$  and  $\tau = 0.84$ ). The chaotic behavior starts at  $\tau = 0.87$ .




Fig. 8. Delay learning paths with 3 delays (top) and 4 delays (bottom). Both networks are trained with 30-step simulation horizon.

work with same weights and biases and analyze the behavior of the corresponding NDDE. Fig. 6(d) depicts the trajectory for  $\tau_2 = 0.8$  where the system approaches a stable period two limit cycle. (Recall that the training was performed for  $\tau_2 = 1$  where the system approached a chaotic attractor.) The constant initial condition  $x(t) = 0.2$  is used which lies outside the training dataset. The neural network gives a good prediction for the limit cycle, since the nonlinearity of the learned NDDE is similar to the of the DDE. This demonstrates the advantage of learning the delays in addition to learning the weights and biases. With the explicitly learned time delays, the network generalizes well and provides insights to the system dynamics.

We exhibit the bifurcation diagram of the trained NDDE and compare it to that of the original DDE in Fig. 7. One may observe the formation of limit cycles and their stability changes as the delay increases from 0 to 2. The network is able to predict appearance of Hopf and period doubling bifurcations and give good predictions for the amplitude, period and stability of the limit cycles without requiring any retraining.

Finally, the method still works if we do not have knowledge about the zero delay, or even about the number of the delays in the system. When (moderately) overestimating the number of delays, we can still learn the dynamics. Fig. 8 shows the delay learning paths of two different network configurations with 3 delays and with 4 delays, respectively. Observe that the delays converge to values around 0 and 1. The nice generalization ability of the networks in this example does depend on the system and training data. Thanks to the chaotic behavior, the training data is rich enough to recover the nonlinearity locally for this autonomous single-delay scalar system. However, similar to other data-driven methods, the performance of this method is limited by the quality of the data. For autonomous systems with multiple states and/or multiple delays, obtaining good training data with rich dynamics might be challenging in many applications.

## 6. CONCLUSION

A method for learning the dynamics of time delay systems was presented. The key idea is to transform the dynamics described by a delay differential equation to an ordinary differential equation through discretizing the history, while retaining continuous time evolution of the systems to enable high accuracy predictions. The dynamics was learned by constructing a neural ordinary differential equation, whose right hand side, represented by a neural network with certain structure, approximated the right hand side of the ODE. A neural delay differential equation can be extracted from the obtained NODE, and it generalized very well when either the parameters or the initial conditions were changed. The developed methods were applied to the Mackey-Glass equation where both the nonlinearity and the time delay were learned with high accuracy. Future research includes applying the developed method to real engineering systems with and without first principle models.

## REFERENCES

- Bai, S., Kolter, J.Z., and Koltun, V. (2019). Deep equilibrium models. *Advances in Neural Information Processing Systems*, 32.
- Breda, D., Maset, S., and Vermiglio, R. (2014). *Stability of linear delay differential equations: A numerical approach with MATLAB*. Springer.
- Brunton, S.L. and Kutz, J.N. (2019). *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press.
- Chen, R.T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. (2018). Neural ordinary differential equations. *Advances in Neural Information Processing Systems*.
- Dupont, E., Doucet, A., and Teh, Y.W. (2019). Augmented neural ODEs. *Advances in Neural Information Processing Systems*, 32.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, 249–256. JMLR Workshop and Conference Proceedings.
- Ji, X.A., Molnár, T.G., Avedisov, S.S., and Orosz, G. (2020). Feed-forward neural networks with trainable delay. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120, 127–136. PMLR.
- Ji, X.A., Molnár, T.G., Avedisov, S.S., and Orosz, G. (2021). Learning the dynamics of time delay systems with trainable delays. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144, 930–942. PMLR.
- Kingma, D.P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun (eds.), *Proceedings of the 3rd International Conference on Learning Representations*.
- Koch, J., Maxner, T., Amatya, V., Ranjbari, A., and Dowling, C. (2022). Physics-informed machine learning of parameterized fundamental diagrams. *arXiv preprint arXiv:2208.00880*.
- Kumpati, S.N., Kannan, P., et al. (1990). Identification and control of dynamical systems using neural networks. *IEEE Transactions on Neural Networks*, 1(1), 4–27.
- Li, X.D., Ho, J.K., and Chow, T.W. (2005). Approximation of dynamical time-variant systems by continuous-time recurrent neural networks. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 52(10), 656–660.
- Mackey, M.C. and Glass, L. (1977). Oscillation and chaos in physiological control systems. *Science*, 197(4300), 287–289.
- Marcus, C. and Westervelt, R. (1989). Stability of analog neural networks with delay. *Physical Review A*, 39(1), 347.
- Pei, J.S., Mai, E.C., Wright, J.P., and Masri, S.F. (2013). Mapping some basic functions and operations to multi-layer feedforward neural networks for modeling nonlinear dynamical systems and beyond. *Nonlinear Dynamics*, 71(1-2), 371–399.
- Rahman, A., Drgoña, J., Tuor, A., and Strube, J. (2022). Neural ordinary differential equations for nonlinear system identification. In *2022 American Control Conference*, 3979–3984.
- Stelzer, F., Röhm, A., Vicente, R., Fischer, I., and Yanchuk, S. (2021). Deep neural networks using a single neuron: folded-in-time architecture using feedback-modulated delay loops. *Nature Communications*, 12(1), 1–10.
- Stépán, G. (1989). *Retarded Dynamical Systems: Stability and Characteristic Functions*. Longman.
- Turan, E.M. and Jäschke, J. (2021). Multiple shooting for training neural differential equations on time series. *IEEE Control Systems Letters*, 6, 1897–1902.
- Waibel, A., Hanazawa, T., Hinton, G., Shikano, K., and Lang, K.J. (1989). Phoneme recognition using time-delay neural networks. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(3), 328–339.
- Wong, S., Jiang, L., Walters, R., Molnár, T.G., Orosz, G., and Yu, R. (2021). Traffic forecasting using vehicle-to-vehicle communication. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144, 917–929. PMLR.
- Zhang, X., Xu, X., and Zhu, Y. (2022). An improved time delay neural network model for predicting dynamic heat and mass transfer characteristics of a packed liquid desiccant dehumidifier. *International Journal of Thermal Sciences*, 177, 107548.
- Zhu, Q., Guo, Y., and Lin, W. (2021). Neural delay differential equations. In *International Conference on Learning Representations*.