NEURAL DELAY DIFFERENTIAL EQUATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural Ordinary Differential Equations (NODEs), a framework of continuous-depth neural networks, have been widely applied, showing exceptional efficacy in coping with some representative datasets. Recently, an augmented framework has been developed for conquering some limitations emergent in application of the original framework. Here we propose a new class of continuous-depth neural networks with delay, named as Neural Delay Differential Equations (NDDEs), and, for computing the corresponding gradients, we use the adjoint sensitivity method to obtain the delayed dynamics of the adjoint. Since the differential equations with delays are usually seen as dynamical systems of infinite dimension possessing more fruitful dynamics, the NDDEs, compared to the NODEs, own a stronger capacity of nonlinear representations. We use several illustrative examples to demonstrate this outstanding capacity: (1) we successfully model the delayed dynamics where the trajectories in the lower-dimensional phase space could be mutually intersected and even chaotic, while the traditional NODEs without any argumentation are not directly applicable for such modeling, and (2) we achieve lower loss and higher accuracy not only for the data produced synthetically by complex models but also for the CIFAR10, a well-known image dataset. Our results on the NDDEs reveal that appropriately articulating the elements of dynamical systems into the network design is truly beneficial to promoting the network performance.

1 INTRODUCTION

A series of recent works have revealed a close connection between neural networks and dynamical systems (E, 2017; Li et al., 2017; Haber & Ruthotto, 2017; Chang et al., 2017; Li & Hao, 2018; Lu et al., 2018; E et al., 2019; Chang et al., 2019; Ruthotto & Haber, 2019; Zhang et al., 2019a; Pathak et al., 2018; Tang et al., 2020). On one hand, the deep neural networks can be used to solve the ordinary/partial differential equations that cannot be easily computed using the traditional algorithms. On the other hand, the elements of the dynamical systems can be useful for establishing novel and efficient frameworks of neural networks. Typical examples include the Neural Ordinary Differential Equations (NODEs), where the infinitesimal time of ordinary differential equations is regarded as the "depth" of a considered neural network (Chen et al., 2018).

Though the advantages of the NODEs were demonstrated through modeling continuous-time datasets and continuous normalizing flows with constant memory cost (Chen et al., 2018), the limited capability of representation for some functions were also studied (Dupont et al., 2019). Indeed, the NODEs cannot be directly used to describe the dynamical systems where the trajectories in the lower-dimensional phase space are mutually intersected. Also, the NODEs cannot model only a few variables from some physical or/and physiological systems where the effect of time delay is inevitably present. From a view point of dynamical systems theory, all these are attributed to the characteristic of finite-dimension for the NODEs.

In this article, we propose a novel framework of continuous-depth neural networks with delay, named as Neural Delay Differential Equations (NDDEs). We apply the adjoint sensitivity method to compute the corresponding gradients, where the obtained adjoint systems are also in a form of delay differential equations. The main virtues of the NDDEs include:

- feasible and computable algorithms for computing the gradients of the loss function based on the adjoint systems,

- representation capability of the vector fields which allow the intersection of the trajectories in the lower-dimensional phase space, and
- accurate reconstruction of the complex dynamical systems with effects of time delays based on the observed time-series data.

2 Related Works

NODEs Inspired by the residual neural networks (He et al., 2016) and the other analogous frameworks, the NODEs were established, which can be represented by multiple residual blocks as

$$\boldsymbol{h}_{t+1} = \boldsymbol{h}_t + \boldsymbol{f}(\boldsymbol{h}_t, \boldsymbol{w}_t),$$

where $\boldsymbol{h}_t$ is the hidden state of the t-th layer, $\boldsymbol{f}(\boldsymbol{h}_t; \boldsymbol{w}_t)$ is a differential function preserving the dimension of $\boldsymbol{h}_t$, and $\boldsymbol{w}_t$ is the parameter pending for learning. The evolution of $\boldsymbol{h}_t$ can be viewed as the special case of the following equation

$$\boldsymbol{h}_{t+\Delta t} = \boldsymbol{h}_t + \Delta t \cdot \boldsymbol{f}(\boldsymbol{h}_t, \boldsymbol{w}_t)$$

with $\Delta t = 1$. As suggested in (Chen et al., 2018), all the parameters $\boldsymbol{w}_t$ are unified into $\boldsymbol{w}$ for achieving parameter efficiency of the NODEs. This unified operation was also employed in the other neural networks, such as the recurrent neural networks (RNNs) (Rumelhart et al., 1986; Elman, 1990) and the ALBERT (Lan et al., 2019). Letting $\Delta t \to 0$ and using the unified parameter $\boldsymbol{w}$ instead of $\boldsymbol{w}_t$, we obtain the continuous evolution of the hidden state $\boldsymbol{h}_t$ as

$$\lim_{\Delta t \to 0} \frac{\boldsymbol{h}_{t+\Delta t} - \boldsymbol{h}_t}{\Delta t} = \frac{d\boldsymbol{h}_t}{dt} = \boldsymbol{f}(\boldsymbol{h}_t, t; \boldsymbol{w}),$$

which is in the form of ordinary differential equations. Actually, the NODEs can act as a feature extraction, mapping an input to a point in the feature space by computing the forward path of a NODE as:

$$\boldsymbol{h}(T) = \boldsymbol{h}(0) + \int_0^T \boldsymbol{f}(\boldsymbol{h}_t, t; \boldsymbol{w}) dt, \quad \boldsymbol{h}(0) = \text{input},$$

where $\boldsymbol{h}(0) = \text{input}$ is the original data point or its transformation, and T is the integration time (assuming that the system starts at $t = 0$). As emphasized in (Dupont et al., 2019), the flow of the NODEs cannot represent some functions omnipresently emergent in applications. Typical examples include the following two-valued function with one argument:

$$g_{\text{1-D}}(x) = \begin{cases} 1, & \text{if } x = -1, \\ -1, & \text{if } x = 1. \end{cases}$$

Our framework desires to conquer the representation limitation observed in applying the NODEs.

Optimal control As mentioned above, a closed connection between deep neural networks and dynamical systems have been emphasized in the literature and, correspondingly, theories, methods and tools of dynamical systems have been employed, e.g. the theory of optimal control (E, 2017; Li

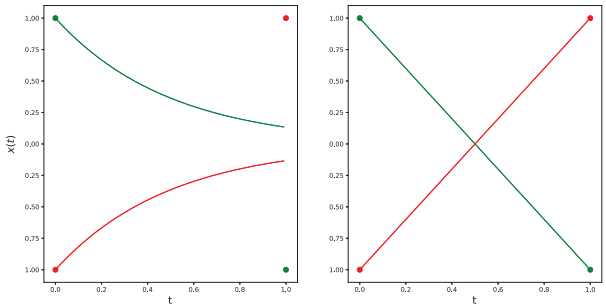

Figure 1: (Right) Two continuous trajectories generated by the DDEs are intersected, mapping -1 (resp., 1) to 1 (resp., -1), while (Left) the ODEs cannot represent such mapping.

et al., 2017; Haber & Ruthotto, 2017; Chang et al., 2017; Li & Hao, 2018; E et al., 2019; Chang et al., 2019; Ruthotto & Haber, 2019; Zhang et al., 2019a). Generally, we model a typical task using a deep neural network and then train the network parameters such that the given loss function can be reduced by some learning algorithm. In fact, training a network can be seen as solving an optimal control problem on difference or differential equations (E et al., 2019). The parameters act as a controller with a goal of finding an optimal control to minimize/maximize some objective function. Clearly, the framework of the NODEs can be formulated as a typical problem of optimal control on ODEs. Additionally, the framework of NODEs has been generalized to the other dynamical systems, such as the Partial Differential Equations (PDEs) (Han et al., 2018; Long et al., 2018; 2019; Ruthotto & Haber, 2019; Sun et al., 2020) and the Stochastic Differential Equations (SDEs) (Lu et al., 2018; Sun et al., 2018; Liu et al., 2019), where the theory of optimal control has been completely established. It is worthwhile to mention that the optimal control theory is tightly connected with and benefits from the method of the classical calculus of variations (Liberzon, 2011). We also will transform our framework into an optimal control problem, and finally solve it using the method of the calculus of variations.

3 THE FRAMEWORK OF NDDES

3.1 FORMULATION OF NDDES

In this section, we establish a framework of continuous-depth neural networks. To this end, we first introduce the concept of delay deferential equations (DDEs). The DDEs are always written in a form where the derivative of a given variable at time t is affected not only by the current state of this variable but also the states at some previous time instants or time durations (Erneux, 2009). Such kind of delayed dynamics play an important role in description of the complex phenomena emergent in many real-world systems, such as physical, chemical, ecological, and physiological systems. In this article, we consider a system of DDE with a single time delay:

$$\begin{cases} \frac{d\boldsymbol{h}_t}{dt} = f(\boldsymbol{h}_t, \boldsymbol{h}_{t-\tau}, t; \boldsymbol{w}), & t >= 0, \\ \boldsymbol{h}(t) = \phi(t), & t <= 0, \end{cases}$$

where the positive constant τ is the time delay and $\boldsymbol{h}(t) = \phi(t)$ is the *initial function* before the time $t = 0$. Clearly, in the initial problem of ODEs, we only need to initialize the state of the variable at $t = 0$ while we initialize the DDEs using a continuous function. Here, to highlight the difference between ODEs and DDEs, we provide a simple example:

$$\begin{cases} \frac{dx_t}{dt} = -2x_{t-\tau}, & t >= 0, \\ x(t) = x_0, & t <= 0. \end{cases}$$

where $\tau = 1$ (with time delay) or $\tau = 0$ (without time delay). As shown in Figure 1, the DDE flow can map -1 to 1 and 1 to -1; nevertheless, this cannot be made for the ODE whose trajectories are not intersected with each other in the t-x space in Figure 1.

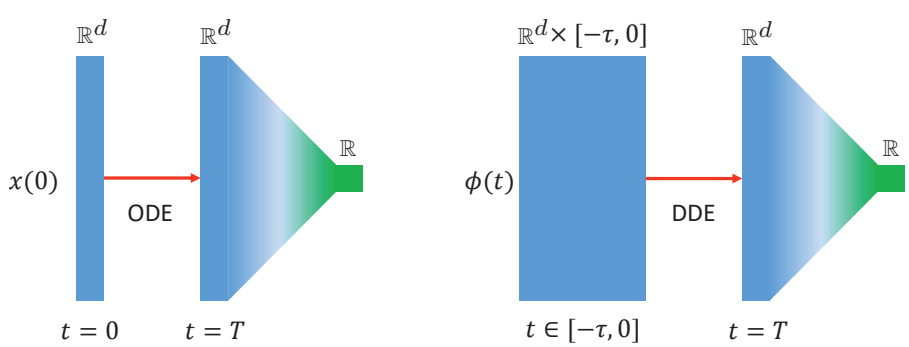

Figure 2: Sketchy diagrams of the NODEs and the NDDES, respectively, with the initial value $\boldsymbol{x}(0)$ and the initial function $\phi(t)$.

3.2 ADJOINT METHOD FOR NDDEs

Assume that the forward pass of DDEs is complete. Then, we need to compute the gradients in a reverse-mode differentiation by using the adjoint sensitivity method (Chen et al., 2018; Pontryagin et al., 1962). We consider an augmented variable, named as *adjoint* and defined as

$$\boldsymbol{\lambda}(t) = \frac{\partial L(\boldsymbol{h}(T))}{\partial \boldsymbol{h}(t)}, \tag{1}$$

where $L(\cdot)$ is the loss function pending for optimization. Notably, the resulted system for the adjoint is in a form of DDE as well.

Theorem 1 *(Adjoint method for NDDEs). Consider the loss function $L(\cdot)$. Then, the dynamics of adjoint can be written as*

$$\begin{cases} \dfrac{d\boldsymbol{\lambda}(t)}{dt} = -\boldsymbol{\lambda}(t)^\top \dfrac{\partial f(\boldsymbol{h}_t, \boldsymbol{h}_{t-\tau}, t; \boldsymbol{w})}{\partial \boldsymbol{h}_t} - \boldsymbol{\lambda}(t+\tau)^\top \dfrac{\partial f(\boldsymbol{h}_{t+\tau}, \boldsymbol{h}_t, t; \boldsymbol{w})}{\partial \boldsymbol{h}_t} \chi_{[0,T-\tau]}(t),\ t <= T \\ \boldsymbol{\lambda}(T) = \dfrac{\partial L(\boldsymbol{h}(T))}{\partial \boldsymbol{h}(T)}, \end{cases} \tag{2}$$

where $\chi_{[0,T-\tau]}(\cdot)$ is a typical characteristic function.

We provide two ways to prove Theorem 1, which are, respectively, shown in Appendix. Using $\boldsymbol{h}(t)$ and $\boldsymbol{\lambda}(t)$, we compute the gradients with respect to the parameters $\boldsymbol{w}$ as:

$$\frac{dL}{d\boldsymbol{w}} = \int_T^0 -\boldsymbol{\lambda}(t)^\top \frac{\partial f(\boldsymbol{h}_t, \boldsymbol{h}_{t-\tau}, t; \boldsymbol{w})}{\partial \boldsymbol{w}} dt. \tag{3}$$

We solve the forward pass of $\boldsymbol{h}$ and backward pass for $\boldsymbol{h}$, $\boldsymbol{\lambda}$ and $\frac{dL}{d\boldsymbol{w}}$ by a piece-wise ODE solver, which is shown in Algorithm 1. For simplicity, we denote by $f(t)$ and $g(t)$ the vector filed of $\boldsymbol{h}$ and $\boldsymbol{\lambda}$, respectively. Moreover, in this paper, we only consider the initial function $\phi(t)$ as a constant function, i.e., $\phi(t) = \boldsymbol{h}_0$. Assume that $T = n \cdot \tau$ and denote $f_k(t) = f(k \cdot \tau + t)$, $g_k(t) = g(k \cdot \tau + t)$ and $\boldsymbol{\lambda}_k(t) = \boldsymbol{\lambda}(k \cdot \tau + t)$.

Algorithm 1 Piece-wise reverse-mode derivative of an DDE initial function problem

Input: dynamics parameters $\boldsymbol{w}$, time delay τ, start time 0, stop time $T = n \cdot \tau$, final state $\boldsymbol{h}(T)$, loss gradient $\partial L / \partial \boldsymbol{h}(T)$

$\frac{\partial L}{\partial \boldsymbol{w}} = 0_{|\boldsymbol{w}|}$

for i **in range**$(n-1, -1, -1)$:

  $s_0 = [\boldsymbol{h}(T), \boldsymbol{h}(T-\tau), ..., \boldsymbol{h}(\tau), \frac{\partial L}{\partial \boldsymbol{h}(T)}, ..., \frac{\partial L}{\partial \boldsymbol{h}((i+1)\cdot\tau)}, \frac{\partial L}{\partial \boldsymbol{w}}]$

  def aug_dynamics([$\boldsymbol{h}_{n-1}(t)$,...,$\boldsymbol{h}_0(t)$, $\boldsymbol{\lambda}_{n-1}(t)$,..,$\boldsymbol{\lambda}_i(t)$, .], t, $\boldsymbol{w}$):

    return $[f_{n-1}(t), f_{n-2}(t), ..., f_0(t), g_{n-1}(t), ..., g_i(t), -\boldsymbol{\lambda}_i(t)^\top \frac{\partial f_i(t)}{\partial \boldsymbol{w}}]$

  $[\frac{\partial L}{\partial \boldsymbol{h}(i\cdot\tau)}, \frac{\partial L}{\partial \boldsymbol{w}}]$=ODESolve($s_0$, aug_dynamics, τ, 0, $\boldsymbol{w}$)

return $\frac{\partial L}{\partial \boldsymbol{h}(0)}, \frac{\partial L}{\partial \boldsymbol{w}}$

In the traditional framework of the NODEs, we can calculate the gradients of the loss function and recompute the hidden states by solving another augmented ODEs in a reversal time duration. However, to achieve the reverse-mode of the NDDEs in Algorithm 1, we need to store the checkpoints of the forward hidden states $h(i \cdot \tau)$ for $i = 0, 1, ..., n$, which, together with the adjoint $\boldsymbol{\lambda}(t)$, can help us to recompute $h(t)$ backwards in every time periods. The main idea of the Algorithm 1 is to convert the DDEs as a piece-wise ODEs such that one can naturally employ the framework of the NODEs to solve it.

4 ILLUSTRATIVE EXPERIMENTS

4.1 EXPERIMENTS ON SYNTHETIC DATASETS

Here, we use some synthetic datasets produced by typical examples to compare the performance of the NODES and the NDDEs. In (Dupont et al., 2019), it is proved that the NODEs cannot represent

the function $g : \mathbb{R}^d \to \mathbb{R}$, defined by

$$g(\boldsymbol{x}) = \begin{cases} 1, & \text{if } \|\boldsymbol{x}\| \leq r_1, \\ -1, & \text{if } r_2 \leq \|\boldsymbol{x}\| \leq r_3, \end{cases}$$

where $0 < r_1 < r_2 < r_3$ and $\|\cdot\|$ is the Euclidean norm. The following proposition show that the NDDEs have stronger capability of representation.

Proposition 1 *The NDDEs can represent $g(\boldsymbol{x})$.*

To validate this proposition, we construct a special form of the NDDEs by

$$\begin{cases} \frac{d\boldsymbol{h}_i(t)}{dt} = \|\boldsymbol{h}_{t-\tau}\| - r, & t >= 0 \text{ and } i = 1, \\ \frac{d\boldsymbol{h}_i(t)}{dt} = 0, & t >= 0 \text{ and } i = 2, ..., d, \\ \boldsymbol{h}(t) = \boldsymbol{x}, & t <= 0. \end{cases} \tag{4}$$

where $r := (r_1 + r_2)/2$ is a constant and the final time point T is supposed to be equal to the time delay τ with some sufficient large value. Under such configurations, we can linearly separate the two clusters by some hyper plane. In Figure 3, we, respectively, show the original data of $g(\boldsymbol{x})$ for the dimension $d = 2$ and the transformed data by the DDEs (4) with the parameters $r_1 = 1$, $r_2 = 2$, $r_3 = 3$, $T = 10$ and $\tau = 10$. Clearly, the transformed data by the DDEs are linearly separable.

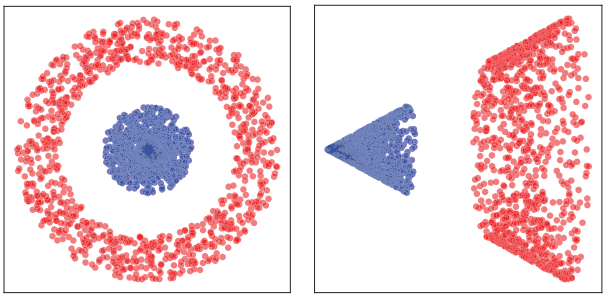

Figure 3: (Left) The data at the time $t = 0$ and (Right) the transformed data at a sufficient large final time T of the DDEs (4). Here, the transformed data are linearly separable.

We also train the NDDEs to represent the $g(\boldsymbol{x})$ for $d = 2$, whose evolutions during the training procedure are, respectively, shown in Figure 4. This figure also includes the evolutions of the NODEs. Notably, while the NODEs is struggled to break apart the annulus, the NDDEs easily separate them. The training losses and the flows of the NODEs and the NDDEs are depicted, respectively, in Figure 5. Particularly, the NDDEs achieve lower losses with the faster speed and directly separate the two clusters in the original 2-D space; however, the NODEs achieve it only by increasing the dimension of the data and separating them in a higher-dimensional space.

Additionally, NDDEs are suitable for fitting the time series with the delay effect in the original systems, which cannot be easily achieved by using the NODEs. To illustrate this, we use a model of 2-D DDEs, written as $\dot{\boldsymbol{x}} = \boldsymbol{A}\tanh(\boldsymbol{x}(t) + \boldsymbol{x}(t-\tau))$ with $\boldsymbol{x}(t) = \boldsymbol{x}_0$ for $t < 0$. Given the time series generated by the system, we use the NDDEs and the NODEs to fit it, respectively. Figure 6 shows that the NNDEs approach a much lower loss, compared to the NODEs. More interestingly, the NODEs prefer to fit the dimension of $\boldsymbol{x}_2$ and its loss always sustains at a larger value, e.g. 0.25 in Figure 6. The main reason is that two different trajectories generated by the autonomous ODEs cannot intersect with each other in the phase space.

We perform experiments on another two classical DDEs, i.e., the population dynamics and the Mackey-Glass system (Erneux, 2009). Specifically, the equation of the dimensionless population dynamics is $\dot{x} = rx(t)(1-x(t-1))$, where $x(t)$ is the ratio of the population to the carrying capacity of the population, r is the growth rate. The Mackey-Glass system is written as $\dot{x} = \beta\frac{x(t-\tau)}{1+x^n(t-\tau)} - \gamma x(t)$, where $x(t)$ is the number of the blood cells, β, n, τ, γ are the parameters of biological significance. The NODEs and the NDDEs are tested on these two dynamics. As shown in Figure 7, a very low training loss is achieved for the NDDEs while the the loss of the NODEs does not go down afterwards, always sustaining at some larger value. As for the predicting durations, the NDDEs thereby achieve a better performance than the NODEs.

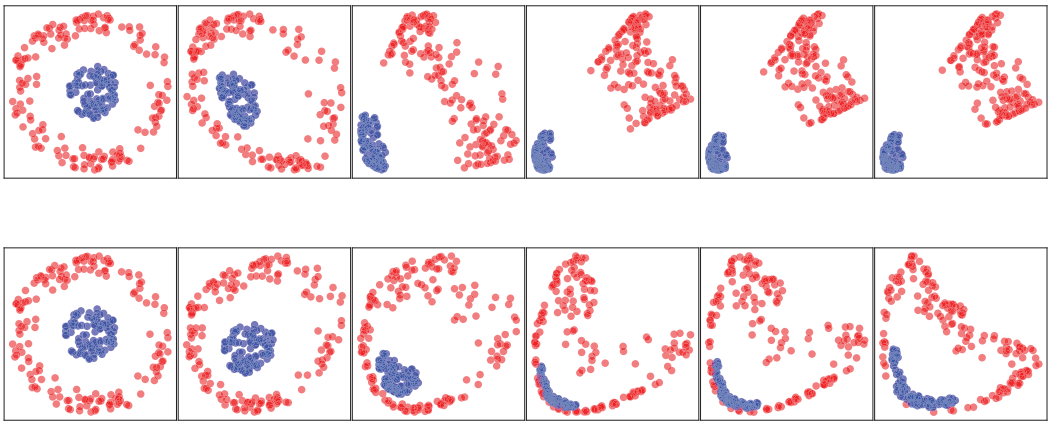

Figure 4: Evolutions of the NDDEs (top) and the NODEs (bottom) in the feature space during the training procedure.

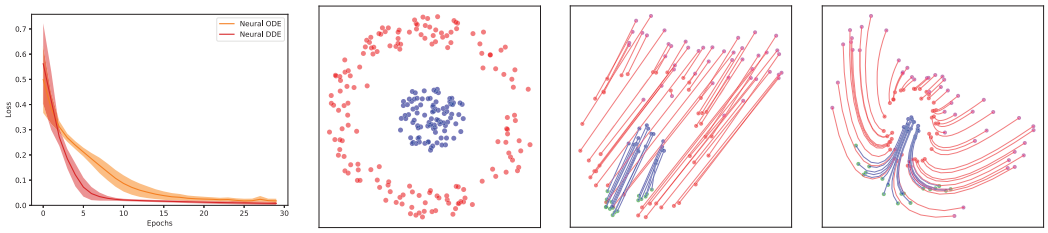

(a) (b) (c) (d)

Figure 5: Presented are the training losses (a) of the NODEs and the NDDEs on fitting the function $g(x)$ for $d = 2$. Also presented are the flows, from the data at the initial time point (b), of the NODEs (d) and the NDDEs (c) after training.

4.2 Experiments on image datasets

Here, we apply the NODEs and the NDDEs to the image datasets. For the NODEs, we model the vector filed $f(\boldsymbol{h}(t))$ as the convolutional architectures together with the same hyperparameter setups in (Dupont et al., 2019). The initial hidden state $\boldsymbol{h}(0) \in \mathbb{R}^{c\times h\times w}$ with respect to an image, where c, h, w are, respectively, the number of channels, height, and width of the image. For the NDDEs, we design the vector filed as $f(\mathrm{concat}(\boldsymbol{h}(t), \boldsymbol{h}(t-\tau)))$, mapping the space from $\mathbb{R}^{2c\times h\times w}$ to $\mathbb{R}^{c\times h\times w}$, where $\mathrm{concat}(\cdot,\cdot)$ executes the concatenation of two tensors on the channel dimension. Its initial function is designed as a constant function, i.e., $h(t)$ is the input/image for $t < 0$.

The training processes on MNIST and CIFAR10 are shown in Figure 8. For CIFAR10, the NDDEs have their training loss decreasing faster than the NODEs, achieving lower training and test loss. Also the test accuracy is much higher than the that of the NODEs (refer to Table 1). Naturally, the better performance of the NDDEs is attributed to the integration of the information not only on the hidden states at the current time t but at the previous time $t - \tau$. This kind of framework is akin to the key idea proposed in (Huang et al., 2017), where the information is processed on many hidden states. As for MNIST, the test accuracy of the NDDEs are competitive with the one of the NODEs (see Table 1); however, in the final stage of training, the loss of these NODEs is easy to diverge. Here, we run all the experiments for 5 times independently.

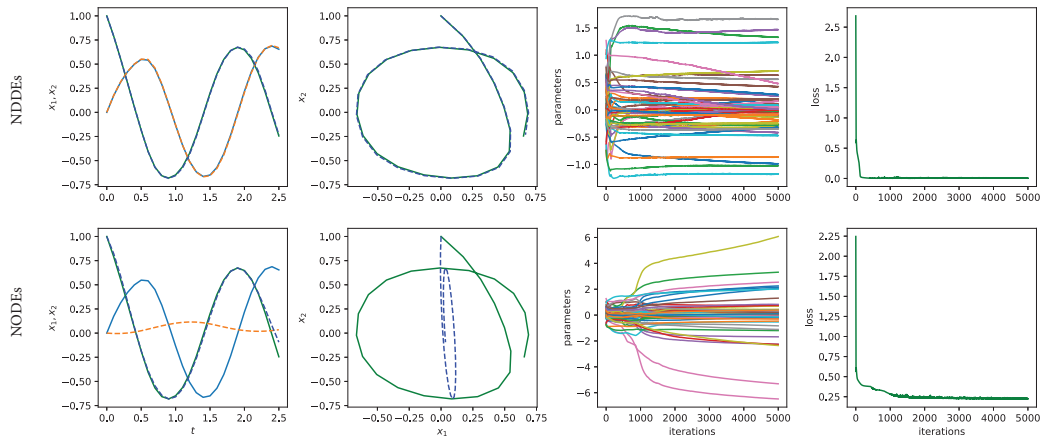

Figure 6: Comparison of the NDDEs (top) versus the NODEs (bottom) in the fitting a 2-D time series with the delay effect in the original system. From the left to the right, the true and the fitted time series, the true and the fitted trajectories in the phase spaces, the dynamics of the parameters in the neural networks, and the dynamics of the losses during the training processes.

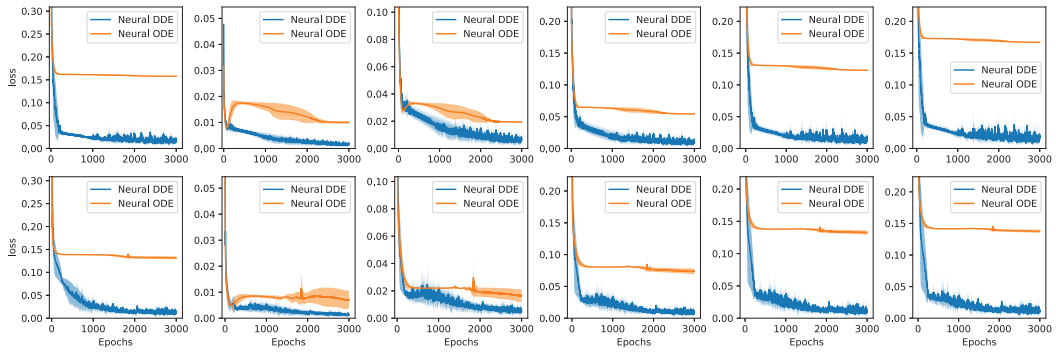

Figure 7: The training losses and the test losses of the population dynamics (top) and the Mackey-Glass system (bottom) by using the NDDEs and the NODEs. The figures from the left to the right present the results over the time intervals of the length 0.1τ, 0.5τ, τ, 1.5τ, 2τ. The delays for the two DDEs are designed as 1 and 2, respectively. The other parameters are set as $r = 1.8$, $\beta = 0.8$, $n = 10$, and $\gamma = 0.3$.

5 Discussion

In this section, we present the limitations of the NDDEs and further suggest several potential directions for future works. We add the delay effect to the NDDEs, which renders the model absolutely irreversible. Algorithm 1 thus requires a storage of the checkpoints of the hidden state $\boldsymbol{h}(t)$ at every instant of the multiple of τ. Actually, solving the DDEs is transformed to solving the ODEs of an increasingly high dimension with respect to the ratio of the final time T and the time delay τ. This

	NDDEs	NODEs
MNIST	$\mathbf{96.22\%} \pm \mathbf{0.55}$	$96.21\% \pm 0.66$
CIFAR10	$\mathbf{55.69\%} \pm \mathbf{0.39}$	$53.94\% \pm 0.65$

Table 1: The test accuracies with their standard deviations over the 5 realizations on the two image sets. Set the initial (final) time as 0 (1) for both neural networks, and the delay of the NDDEs as 1.

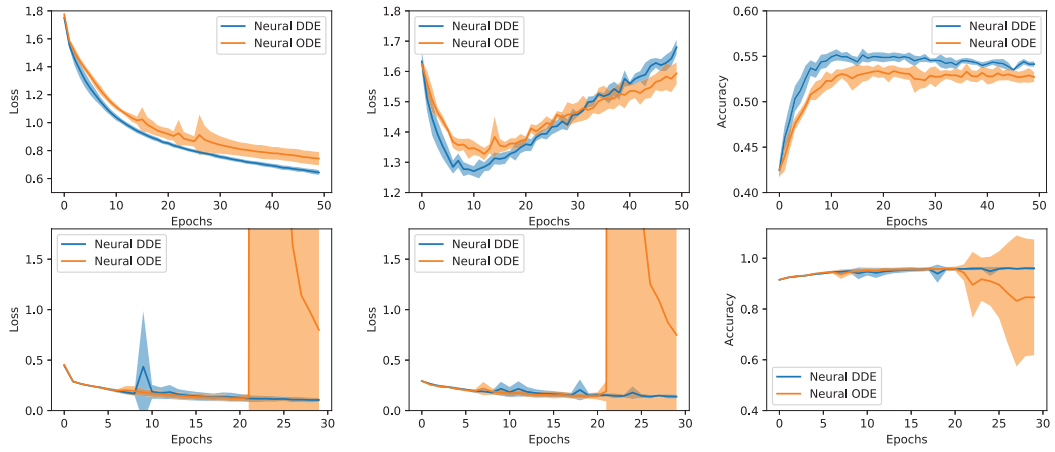

Figure 8: The training loss (left), the test loss (middle), and the accuracy (right) over the 5 realizations for the two image sets, i.e., CIFAR10 (top) and MNIST (bottom).

definitely indicates a high computational cost. To further apply and improve the framework of the NDDEs, a few potential directions for future works are suggested, including:

Applications to more real-world datasets. In the real-world systems such as physical, chemical, biological, and ecological systems, the delay effects are inevitably omnipresent, truly affecting the dynamics of the produced time-series (Bocharov & Rihan, 2000; Kajiwara et al., 2012). The NDDEs are undoubtedly suitable for realizing model-free and accurate prediction (Quaglino et al., 2019). Additionally, since the NODEs have been generalized to the areas of Computer Vision and Natural Language Processing He et al. (2019); Yang et al. (2019); Liu et al. (2020), the framework of the NDDEs probably can be applied to the analogous areas, where the delay effects should be ubiquitous in those streaming data.

Extension of the NDDEs. A single constant time-delay in the NDDEs can be further generalized to the case of multiple or/and distributed time delays (Shampine & Thompson, 2001). As such, the model is likely to have much stronger capability to extract the feature, because the model leverages the information at different time points to make the decision in time. Besides, we can model the initial function as an ODE, driving the NDDEs to have a stronger nonlinear representation. All these extensions could be potentially suitable for some complex tasks.

Time-dependent controllers. From a viewpoint of control theory, the parameters in the NODEs/NDDEs could be regarded as time-independent controllers, viz. constant controllers. A natural generalization way is to model the parameters as time-dependent controllers. In fact, such controllers were proposed in (Zhang et al., 2019b), where the parameters $\boldsymbol{w}(t)$ were treated as another learnable ODE $\dot{\boldsymbol{w}} = q(\boldsymbol{w}(t), \boldsymbol{p})$, $q(\cdot, \cdot)$ is a different neural network, and the parameters $\boldsymbol{p}$ and the initial state $\boldsymbol{w}(0)$ are pending for optimization. Also, the idea of using a neural network to generate the other one was initially conceived in some earlier works including the study of the hypernetworks (Ha et al., 2016).

6 Conclusion

In this article, we establish the framework of NDDEs, whose vector fields are determined mainly by the states at the previous time. We employ the adjoint sensitivity method to compute the gradients of the loss function. The obtained adjoint dynamics backward follow another DDEs coupled with the forward hidden states. We show that the NDDEs can represent some typical functions that cannot be represented by the original framework of NODEs. We also demonstrate the exceptional efficacy of the proposed framework by using the synthetic data or real-world datasets. All these reveal that integrating the elements of dynamical systems to the architecture of neural networks could be potentially beneficial to the promotion of network performance.

AUTHOR CONTRIBUTIONS

To be added after the review.

ACKNOWLEDGMENTS

To be added after the review.

REFERENCES

Gennadii A Bocharov and Fathalla A Rihan. Numerical modelling in biosciences using delay differential equations. *Journal of Computational and Applied Mathematics*, 125(1-2):183–199, 2000.

Bo Chang, Lili Meng, Eldad Haber, Frederick Tung, and David Begert. Multi-level residual networks from dynamical systems view. *arXiv preprint arXiv:1710.10348*, 2017.

Bo Chang, Minmin Chen, Eldad Haber, and Ed H Chi. Antisymmetricrnn: A dynamical system view on recurrent neural networks. *arXiv preprint arXiv:1902.09689*, 2019.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in neural information processing systems*, pp. 6571–6583, 2018.

Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. In *Advances in Neural Information Processing Systems*, pp. 3140–3150, 2019.

Weinan E. A proposal on machine learning via dynamical systems. *Communications in Mathematics and Statistics*, 5(1):1–11, 2017.

Weinan E, Jiequn Han, and Qianxiao Li. A mean-field optimal control formulation of deep learning. *Research in the Mathematical Sciences*, 6(1):10, 2019.

Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.

Thomas Erneux. *Applied delay differential equations*, volume 3. Springer Science & Business Media, 2009.

David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.

Eldad Haber and Lars Ruthotto. Stable architectures for deep neural networks. *Inverse Problems*, 34(1):014004, 2017.

Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Xiangyu He, Zitao Mo, Peisong Wang, Yang Liu, Mingyuan Yang, and Jian Cheng. Ode-inspired network design for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1732–1741, 2019.

Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

Tsuyoshi Kajiwara, Toru Sasaki, and Yasuhiro Takeuchi. Construction of lyapunov functionals for delay differential equations in virology and epidemiology. *Nonlinear Analysis: Real World Applications*, 13(4):1802–1826, 2012.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.

Qianxiao Li and Shuji Hao. An optimal control approach to deep learning and applications to discrete-weight neural networks. *arXiv preprint arXiv:1803.01299*, 2018.

Qianxiao Li, Long Chen, Cheng Tai, and E Weinan. Maximum principle based algorithms for deep learning. *The Journal of Machine Learning Research*, 18(1):5998–6026, 2017.

Daniel Liberzon. *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, 2011.

Xuanqing Liu, Tesi Xiao, Si Si, Qin Cao, Sanjiv Kumar, and Cho-Jui Hsieh. Neural sde: Stabilizing neural ode networks with stochastic noise. *arXiv preprint arXiv:1906.02355*, 2019.

Xuanqing Liu, Hsiang-Fu Yu, Inderjit Dhillon, and Cho-Jui Hsieh. Learning to encode position for transformer with continuous dynamical model. *arXiv preprint arXiv:2003.09229*, 2020.

Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. Pde-net: Learning pdes from data. In *International Conference on Machine Learning*, pp. 3208–3216, 2018.

Zichao Long, Yiping Lu, and Bin Dong. Pde-net 2.0: Learning pdes from data with a numeric-symbolic hybrid deep network. *Journal of Computational Physics*, 399:108925, 2019.

Yiping Lu, Aoxiao Zhong, Quanzheng Li, and Bin Dong. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In *International Conference on Machine Learning*, pp. 3276–3285. PMLR, 2018.

Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical Review Letters*, 120(2), 2018.

LS Pontryagin, VG Boltyanskij, RV Gamkrelidze, and EF Mishchenko. The mathematical theory of optimal processes. 1962.

Alessio Quaglino, Marco Gallieri, Jonathan Masci, and Jan Koutník. Snode: Spectral discretization of neural odes for system identification. *arXiv preprint arXiv:1906.07038*, 2019.

David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

Lars Ruthotto and Eldad Haber. Deep neural networks motivated by partial differential equations. *Journal of Mathematical Imaging and Vision*, pp. 1–13, 2019.

Lawrence F Shampine and Skip Thompson. Solving ddes in matlab. *Applied Numerical Mathematics*, 37(4):441–458, 2001.

Qi Sun, Yunzhe Tao, and Qiang Du. Stochastic training of residual networks: a differential equation viewpoint. *arXiv preprint arXiv:1812.00174*, 2018.

Yifan Sun, Linan Zhang, and Hayden Schaeffer. Neupde: Neural network based ordinary and partial differential equations for modeling time-dependent data. In *Mathematical and Scientific Machine Learning*, pp. 352–372. PMLR, 2020.

Yang Tang, Jürgen Kurths, Wei Lin, Edward Ott, and Ljupco Kocarev. Introduction to focus issue: When machine learning meets complex systems: Networks, chaos, and nonlinear dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(6):063151, 2020.

Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4541–4550, 2019.

Dinghuai Zhang, Tianyuan Zhang, Yiping Lu, Zhanxing Zhu, and Bin Dong. You only propagate once: Painless adversarial training using maximal principle. *arXiv preprint arXiv:1905.00877*, 2 (3), 2019a.

Tianjun Zhang, Zhewei Yao, Amir Gholami, Joseph E Gonzalez, Kurt Keutzer, Michael W Mahoney, and George Biros. Anodev2: A coupled neural ode framework. In *Advances in Neural Information Processing Systems*, pp. 5151–5161, 2019b.

A Appendix

A.1 The First Proof of Theorem 1

Here, we present a direct proof of the adjoint method for the NDDEs. For the neat of the proof, the following notations are slightly different from those in the main text.

Let $x(t)$ obey the DDE written as

$$\begin{cases} \dot{x}(t) = f(x(t), y(t), \theta(t)),\ y(t) = x(t-\tau),\ t \in [0, T], \\ x(t) = x_0,\ t \in [-\tau, 0]. \end{cases} \tag{5}$$

whose adjoint state is defined as

$$\lambda(t) := \frac{\partial L}{\partial x(t)}, \tag{6}$$

where L is the loss function, that is, $L := L(X(T))$. After discretizing the above DDE, we have

$$\begin{aligned} x(t+\Delta t) &= x(t) + \Delta t \cdot f(x(t), y(t), \theta(t)), \\ &= x(t) + \Delta t \cdot f(x(t), x(t-\tau), \theta(t)), \\ x(t+\tau+\Delta t) &= x(t+\tau) + \Delta t \cdot f(x(t+\tau), y(t+\tau), \theta(t+\tau)), \\ &= x(t+\tau) + \Delta t \cdot f(x(t+\tau), x(t), \theta(t+\tau)). \end{aligned} \tag{7}$$

According to the definition of $\lambda(t)$ and applying the chain rule, we have

$$\begin{aligned} \lambda(t) &= \frac{\partial L}{\partial x(t+\Delta t)} \frac{\partial x(t+\Delta t)}{\partial x(t)} + \frac{\partial L}{\partial x(t+\tau+\Delta t)} \frac{\partial x(t+\tau+\Delta t)}{\partial x(t)} \cdot \chi_{[0,T-\tau]}(t) \\ &= \lambda(t+\Delta t) \frac{\partial x(t+\Delta t)}{\partial x(t)} + \lambda(t+\tau+\Delta t) \frac{\partial x(t+\tau+\Delta t)}{\partial x(t)} \cdot \chi_{[0,T-\tau]}(t) \\ &= \lambda(t+\Delta t)(I + \Delta t \cdot f_x(x(t), y(t), \theta(t))) \\ &\quad + \lambda(t+\tau+\Delta t)\Delta t \cdot f_y(x(t+\tau), y(t+\tau), \theta(t+\tau))) \cdot \chi_{[0,T-\tau]}(t) \\ &= \lambda(t+\Delta t) + \Delta t \cdot \left[\lambda(t+\Delta t) \cdot f_x(t) + \lambda(t+\tau+\Delta t) \cdot f_y(t+\tau) \cdot \chi_{[0,T-\tau]}(t)\right] \end{aligned} \tag{8}$$

which implies,

$$\begin{aligned} \dot{\lambda}(t) &= \lim_{\Delta t \to 0} \frac{\lambda(t+\Delta t) - \lambda(t)}{\Delta t} \\ &= -\lambda(t) \cdot f_x(t) - \lambda(t+\tau) \cdot f_y(t+\tau) \cdot \chi_{[0,T-\tau]}(t), \end{aligned} \tag{9}$$

where $\chi_{[0,T-\tau]}(\cdot)$ is a characteristic function. For the parameter $\theta(t)$, the result can be analogously derived as

$$\begin{aligned} \frac{\partial L}{\partial \theta(t)} &= \frac{\partial L}{\partial x(t+\Delta t)} \frac{\partial x(t+\Delta t)}{\partial \theta(t)} \\ &= \Delta t \cdot \lambda(t+\Delta t) \cdot f_\theta(t) \end{aligned} \tag{10}$$

In this article, $\theta(t)$ is considered to be a constant variable, i.e., $\theta(t) \equiv \theta$, which yields:

$$\begin{aligned} \frac{\partial L}{\partial \theta} &= \lim_{\Delta t \to 0} \sum \Delta t \cdot \lambda(t+\Delta t) \cdot f_\theta(t) \\ &= \int_0^T \lambda(t) \cdot f_\theta(t) dt \\ &= \int_T^0 -\lambda(t) \cdot f_\theta(t) dt. \end{aligned} \tag{11}$$

To summarize, we get the gradients with respect to $x(0)$ and θ in a form of augmented DDEs. These DDEs are backward in time and written in an integral form of

$$\begin{aligned} x(0) &= x(T) + \int_T^0 f(x, y, \theta)dt, \\ \frac{\partial L}{\partial x(0)} &= \frac{\partial L}{\partial x(T)} + \int_T^0 -\lambda(t) \cdot f_x(t) - \lambda(t+\tau) \cdot f_y(t+\tau) \cdot \chi_{[0,T-\tau]}(t)dt, \\ \frac{\partial L}{\partial \theta} &= \int_T^0 -\lambda(t) \cdot f_\theta(t)dt. \end{aligned} \tag{12}$$

A.2 The Second Proof of Theorem 1

Here, we mainly employ the Lagrangian multiplier method and the calculus of variation to derive the adjoint method for the NDDEs. First, we define the Lagrangian by

$$\mathcal{L} := L(X(T)) + \int_0^T \lambda(t)(\dot{x} - f(x(t), y(t), \theta))dt. \tag{13}$$

We thereby need to find the so-called Karush-Kuhn-Tucker (KKT) conditions for the Lagrangian, which is necessary for finding the optimal solution of θ. To obtain the KKT conditions, the calculus of variation is applied by taking variations with respect to $\lambda(t)$ and $x(t)$.

For $\lambda(t)$, let $\hat{\lambda}(t)$ be a continuous and differentiable function with a scalar ϵ. We add the perturbation $\hat{\lambda}(t)$ to the Lagrangian $\mathcal{L}$, which results in a new Lagrangian, denoted by

$$\mathcal{L}(\epsilon) := L(x(T)) + \int_0^T (\lambda(t) + \epsilon\hat{\lambda}(t))(\dot{x} - f(x(t), y(t), \theta))dt. \tag{14}$$

In order to obey the optimal condition for $\lambda(t)$, we require the following equation

$$\frac{d\mathcal{L}(\epsilon)}{d\epsilon} = \int_0^T \hat{\lambda}(t)(\dot{x} - f(x(t), y(t), \theta))dt \tag{15}$$

to be zero. Due to the arbitrariness of $\hat{\lambda}(t)$, we obtain

$$\dot{x} - f(x(t), y(t), \theta) = 0, \quad \forall t \in [0, T], \tag{16}$$

which is exactly the DDE forward in time.

Analogously, we can take variation with respect to $x(t)$ and let $\hat{x}(t)$ be a continuous and differentiable function with a scalar ϵ. Here, $\hat{x}(t) = 0$ for $t \in [-\tau, 0]$. We also denote by $\hat{y}(t) := \hat{x}(t-\tau)$ for $t \in [0, T]$. The new Lagrangian under the perturbation of $\hat{x}(t)$ becomes

$$\mathcal{L}(\epsilon) := L(x(T) + \epsilon\hat{x}(T)) + \int_0^T \lambda(t) \left(\frac{dx(t) + \epsilon\hat{x}(t)}{dt} - f(x(t) + \epsilon\hat{x}(t), y(t) + \epsilon\hat{y}(t), \theta) \right) dt. \tag{17}$$

We then compute the $\frac{d\mathcal{L}(\epsilon)}{d\epsilon}$, which gives

$$\begin{aligned} \frac{d\mathcal{L}(\epsilon)}{d\epsilon}|_{\epsilon=0} &= \frac{\partial L}{\partial x(T)}\hat{x}(T) + \int_0^T \lambda(t) \left(\frac{d\hat{x}(t)}{dt} - f_x(x(t), y(t), \theta)\hat{x}(t) - f_y(x(t), y(t), \theta)\hat{y}(t) \right) dt \\ &= \frac{\partial L}{\partial x(T)}\hat{x}(T) + \lambda(t)\hat{x}(t)|_0^T + \int_0^T -\hat{x}(t)\frac{d\lambda(t)}{dt} \qquad \text{integration by parts} \\ &\quad + \int_0^T -\lambda(t)f_x(x(t), y(t), \theta)\hat{x}(t) - \lambda(t)f_y(x(t), y(t), \theta)\hat{y}(t)dt \\ &= \left(\frac{\partial L}{\partial x(T)} + \lambda(T) \right) \hat{x}(T) + \int_0^T -\hat{x}(t)\frac{d\lambda(t)}{dt} - \lambda(t)f_x(x(t), y(t), \theta)\hat{x}(t) \\ &\quad - \int_0^T \lambda(t)f_y(x(t), y(t), \theta)\hat{x}(t-\tau)dt \end{aligned}$$

$$
\begin{aligned}
&= \left(\frac{\partial L}{\partial x(T)} + \lambda(T)\right)\hat{x}(T) + \int_0^T -\hat{x}(t)\frac{d\lambda(t)}{dt} - \lambda(t) f_x(x(t), y(t), \theta)\hat{x}(t) \\
&\quad - \int_{-\tau}^{T-\tau} \lambda(t'+\tau) f_y(x(t'+\tau), y(t'+\tau), \theta)\hat{x}(t')dt' \\
&= \left(\frac{\partial L}{\partial x(T)} + \lambda(T)\right)\hat{x}(T) + \int_0^T -\hat{x}(t)\frac{d\lambda(t)}{dt} - \lambda(t) f_x(x(t), y(t), \theta)\hat{x}(t) \\
&\quad - \int_{0}^{T-\tau} \lambda(t') f_y(x(t'+\tau), y(t'+\tau), \theta)\hat{x}(t')dt' \qquad \text{no variation on the interval } [-\tau, 0] \\
&= \left(\frac{\partial L}{\partial x(T)} + \lambda(T)\right)\hat{x}(T) \\
&\quad + \int_0^T \hat{x}(t)\left(-\frac{d\lambda(t)}{dt} - \lambda(t) f_x(x(t), y(t), \theta) - \lambda(t+\tau) f_y(x(t+\tau), y(t+\tau), \theta)\chi_{[0,T-\tau]}(t)\right) dt
\end{aligned} \tag{18}
$$

Notice that $\frac{d\mathcal{L}(\epsilon)}{d\epsilon}|_{\epsilon=0} = 0$ is satisfied for all continuous differentiable $\hat{x}(t)$ at the optimal $x(t)$. Thus, we have

$$
\frac{d\lambda(t)}{dt} = -\lambda(t) f_x(x(t), y(t), \theta) - \lambda(t+\tau) f_y(x(t+\tau), y(t+\tau), \theta)\chi_{[0,T-\tau]}(t), \quad \lambda(T) = \frac{\partial L}{\partial x(T)}. \tag{19}
$$

Therefore, the adjoint state follows a DDE as well.