

Operator-Learning-Inspired Modeling of Neural Ordinary Differential Equations

Woojin Cho^{1*}, Seunghyeon Cho^{1*}, Hyundong Jin¹, Jinsung Jeon¹,
Kookjin Lee², Sanghyun Hong³, Dongeun Lee⁴, Jonghyun Choi¹, Noseong Park¹

¹Yonsei University

²Arizona State University

³Oregon State University

⁴Texas A&M University-Commerce

{snowmoon, seunghyeoncho, tuzi04, jjsjjs0902}@yonsei.ac.kr,
kjlee8344@gmail.com, sanghyun.hong@oregonstate.edu, dongeun.lee@tamuc.edu, {jc, noseong}@yonsei.ac.kr

Abstract

Neural ordinary differential equations (NODEs), one of the most influential works of the differential equation-based deep learning, are to continuously generalize residual networks and opened a new field. They are currently utilized for various downstream tasks, e.g., image classification, time series classification, image generation, etc. Its key part is how to model the time-derivative of the hidden state, denoted $\frac{d\mathbf{h}(t)}{dt}$. People have habitually used conventional neural network architectures, e.g., fully-connected layers followed by non-linear activations. In this paper, however, we present a neural operator-based method to define the time-derivative term. Neural operators were initially proposed to model the differential operator of partial differential equations (PDEs). Since the time-derivative of NODEs can be understood as a special type of the differential operator, our proposed method, called branched Fourier neural operator (BFNO), makes sense. In our experiments with general downstream tasks, our method significantly outperforms existing methods.

Introduction

Recently, in deep-learning community, there have been growing attention in developing differential-equation(DE)-based neural network architectures. One of the most representative DE-based models is neural ordinary differential equations (NODEs) (Chen et al. 2018), which serve as a general purpose neural network architecture, powered by deep implicit layers, modeling a continuous-depth representation and achieving state-of-the-art (SOTA) performance in many different downstream tasks, e.g., time-series modeling (Chen et al. 2018; Song et al. 2020; Kidger et al. 2020; Morrill et al. 2021; Choi, Jeon, and Park 2021; Jhin et al. 2022; Kim et al. 2022; Jeon et al. 2022). Another fast growing area of deep learning is scientific machine learning (SML), where learning the deep-learning-based solution operator (also known as neural operators) of parametric partial differential equations has been of particular interest to researchers (Li et al. 2020b; Cao 2021; Gupta, Xiao, and Bogdan 2021; Rahman, Ross, and Aizzadenesheli 2022). Fourier neural operators (FNOs) (Li

et al. 2020b), a representative class of neural operators, have demonstrated their effectiveness in building surrogate models for solution operators of parametric partial differential equations (PDEs).

However, as opposed to NODEs, which have been explored in many different ML downstream tasks as generic neural network architectures, neural operators so far have been applied mostly to learning the solution operators of PDEs; the only exception we found is to use neural operators in computer vision task (Guibas et al. 2021). Thus, in this study, we aim to bring the expressive power of neural operators into the modeling NODEs resulting in neural-operator-based NODE and apply this novel combination to many different downstream tasks. We are motivated to conduct this study based on a couple of observations: (1) as ODEs can be interpreted as a particular type of PDEs, which involve derivatives in only one variable, the ODE function can be modeled as neural operators, and (2) as neural operators are effective in learning “parametric” PDE operators via a function-to-function mapping principle, the neural-operator-based NODE is expected to be equipped with stronger expressivity.

Our approach: NODEs model the dynamics of hidden state $\mathbf{h}(t)$ as a form of a system of ODEs from data such that

$$\frac{d\mathbf{h}(t)}{dt} = f(\mathbf{h}(t), t; \boldsymbol{\theta}_f), \quad (1)$$

where $\boldsymbol{\theta}_f$ denotes a set of learnable parameters. Taking the forward pass is equivalent to solving the initial value problem (IVP) in Eq. (2) with the initial value $\mathbf{h}(t_0)$:

$$\mathbf{h}(t_1) = \mathbf{h}(t_0) + \int_{t_0}^{t_1} f(\mathbf{h}(t), t; \boldsymbol{\theta}_f) dt, \quad (2)$$

and the last hidden state $\mathbf{h}(t_1)$ is typically used for downstream tasks (e.g., classification).

The ODE function (the right hand side, or RHS) in Eq. (1) can be considered as applying the differential operator to $\mathbf{h}(t)$ and it is natural to redesign RHS, f , as neural operators. From our empirical studies, we observed that naively adapting recent neural operator methods, such as Fourier neural operators (FNOs), to NODEs, however, results in degradation of performance. We speculate that this degradation is caused by a habitual choice of a rather simple architectural

*These authors contributed equally.

design, which uses conventional neural networks, (e.g., convolutional neural networks (CNNs) and fully-connected networks (FCNs)) to learn the differential operator. We, instead, propose a novel neural operator architecture called branched Fourier neural operators (BFNOs), which leads to learning of much more expressive neural operators that can be applied to complex ML downstream tasks.

For empirical evaluations, we test our method on various ML downstream tasks including image classification, time series classification, and image generation. For each task, we compare our method with SOTA NODE and non-NODE-based baselines. In particular, our main competitors are enhanced NODE models, such as heavy ball NODEs (HBNODEs). Our BFNO-based NODEs (BFNO-NODEs) outperforms all considered baselines. We summarize our contributions as follows:

- 1. Branched Fourier neural operator:** We design the ODE function of NODEs using our proposed BFNO method. In our design, we do not inherit previous architectures of FNO variants, i.e., applying a low-pass filter to the Fourier transform output, followed by a convolution operation, as it is rather restrictive to be used for general machine learning tasks. Instead, our BFNO uses dynamic global convolutional operations with multiple kernels. To our knowledge, we are the first using neural operators to design the ODE function of NODEs.
- 2. High performance in general ML tasks:** Through extensive experiments, we demonstrate that BFNO-NODEs significantly outperform considered baselines in various ML downstream tasks. Our experiments set include image classification, time series classification, and image generation. Our method improves on the performance of previous SOTA NODE and non-NODE methods by 20% in image classification, 1.7% in time series classification in terms of test accuracy, and 10% in image generation (in terms of the standard metrics for those tasks).

Related Work

Continuous-depth neural networks: Continuous-depth neural networks allow greater modeling flexibility and have been investigated as a potential replacement for traditional deep feed-forward neural networks (Weinan 2017; Haber and Ruthotto 2017; Lu et al. 2018). NODEs (Chen et al. 2018) can model continuous-depth neural networks. By the continuous-depth property, NODEs describe the change of the hidden state $\mathbf{h}(t)$ over time. In particular, NODEs use the adjoint sensitivity method to address the main technical challenge of training continuous-depth networks. Many continuous-depth models that use the same computational formalism, such as neural controlled differential equations (Kidger et al. 2020), neural stochastic differential equations (Liu et al. 2019), and so on, have been proposed. In addition, NODE-based models with enhanced ODE function architectures have achieved significant improvements in various tasks in comparison with the original NODE design. Augmented NODEs (ANODEs (Dupont, Doucet, and Teh 2019)) suggested solving the homeomorphic limitation of ODEs with adding extra dimensions. Second-order NODEs (SONODEs (Norcliffe et al. 2020)) showed that they can expand the adjoint sensitivity

approach to the second-order ODEs efficiently. In order to further enhance the training and inference of NODEs, heavy ball NODEs (HBNODEs (Xia et al. 2021)) model the dynamics with the conventional momentum method. Nesterov NODEs (NesterovNODEs (Nguyen et al. 2022)) is proposed to use the Nesterov momentum expression with a fast convergence rate among optimizer methods. Adaptive moment estimation NODEs (AdamNODEs (Cho et al. 2022)) use an enhanced momentum-based method to define the ODE function. In this paper, we propose a novel continuous-depth NODE architecture with the infinite dimensional property inherited from characteristics of neural operators. Our enhanced NODE architecture, called BFNO-NODE, outperforms existing NODE enhancements.

Operator learning: Expressing and analyzing changes in phenomena over time and space as PDEs are an important issue in natural science and engineering. However, solving PDE problems is a mathematically difficult task. In the field of numerical analysis, these problems are approached using various numerical methods (Smith and Smith 1985; Reddy 2019). These methods perform well in many areas where PDE problems need to be analyzed, but there are fundamental challenges: i) it takes a lot of computational cost to solve the problem, ii) the higher its required accuracy, the longer its computation time, and iii) finally, it is difficult to analyze for various spatial resolutions with a single model.

To overcome these limitations, neural operators (Lu, Jin, and Karniadakis 2019; Kovachki et al. 2021) have recently been proposed in the field of deep learning. Neural operators aim to learn the (inverse) differential operator of PDEs. Furthermore, they also have the infinite dimensional property.

Recently, models with repetitive operator-based layers have been proposed in the field of neural operators (Li et al. 2020b,c; Wen et al. 2022). Typically, FNOs first define a kernel integral operator layer in the Fourier domain and repeatedly stack multiple such layers. In the field of computer vision, there was a study that improved the vision transformer (Dosovitskiy et al. 2020) using FNOs. Adaptive Fourier neural operator (AFNO) was proposed in (Guibas et al. 2021) by replacing the self-attention layers of the vision transformer with adaptive FNO layers. However, these operator architectures did not perform well when they were used in the ODE function in our preliminary experiments. Therefore, we propose our BFNO concept.

Preliminaries

Neural ordinary differential equations (NODEs): In order to determine $\mathbf{h}(t_1)$ from $\mathbf{h}(t_0)$, NODEs solve the integral problem in Eq. (2), where the ODE function $f(\mathbf{h}(t), t; \theta_f)$ is a neural network used to approximate $\frac{d\mathbf{h}(t)}{dt}$ (cf. Eq. (1)). Once f is trained, NODEs rely on numerical ODE solvers to solve the integral problem, such as the explicit Euler method, the Dormand-Prince (DOPRI (Dormand and Prince 1980)) method. ODE solvers discretize the time variable t and transform an integral into numerous additions. One distinguished characteristic of NODEs is that the gradient of loss w.r.t. θ_f , denoted $\nabla_{\theta_f} C = \frac{\partial C}{\partial \theta_f}$, where C is a task-dependent cost

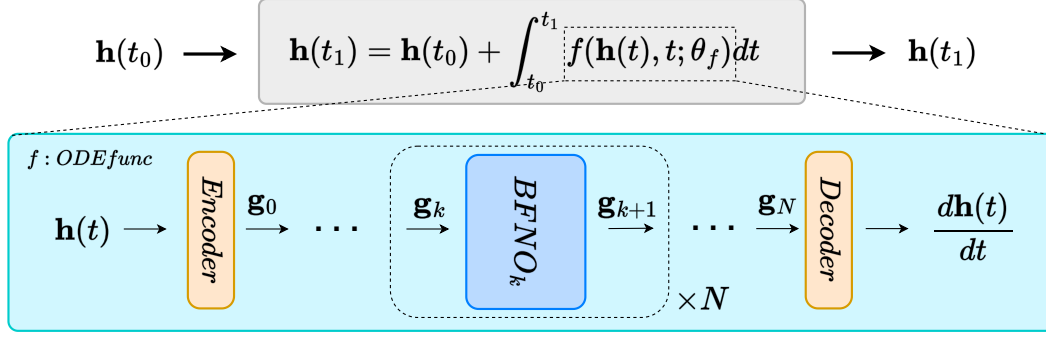


Figure 1: The overall architecture of BFNO-NODE. Our goal is to define the ODE function f based on our proposed neural operator.

function, can be calculated by a reverse-mode integration, which has $\mathcal{O}(1)$ space complexity. This gradient calculation method is known as the adjoint sensitivity method.

Operator learning and discretization of function: Let S be a bounded, open set which is a subspace of \mathbb{R}^d , and let $P = P(S; \mathbb{R}^{d_p})$ and $Q = Q(S; \mathbb{R}^{d_q})$ be Banach spaces which define a function mapping from input $x \in S$ to output whose dimensions are d_p and d_q respectively. Then, a non-linear operator $\mathcal{L} : P \rightarrow Q$, where P and Q are the continuous function spaces, can be defined. By discretizing with finite numbers of observations, the function spaces can be represented in grid forms. At the end, neural operators learn a parameterized mapping $\mathcal{L}_\theta : P \rightarrow Q$ such that:

$$\mathcal{L}_\theta(P)(x) = Q(x), \quad \forall x \in S. \quad (3)$$

For the given two functions, $P(x)$ and $Q(x)$ with $x \in S$, the operator \mathcal{L}_θ can map these function spaces. As such, NODE’s time-derivative operator $\mathcal{D}_f = \frac{d(\cdot)}{dt}$ can be understood as a function mapping from $\mathbf{h}(t)$ to $\frac{d\mathbf{h}(t)}{dt}$, where t denotes the temporal coordinate. This differential operator can be regarded as a special case of the general operator representation. Therefore, we can design the ODE function f of NODEs as an operator-based network.

Kernel integral operators: As a way of expressing an operator, one can define the kernel integral operator \mathcal{K} as following:

$$\mathcal{K}(q)(x) := \int_S \kappa(x, y)q(y)dy, \quad \forall x \in S, \quad (4)$$

where κ is a continuous kernel function. If we use the Green’s kernel $\kappa(x, y) = \kappa(x - y)$, the RHS of Eq. (4) becomes the following global convolution operator:

$$\int_S \kappa(x - y)q(y)dy, \quad \forall x \in S. \quad (5)$$

Fourier neural operators (FNOs): For an efficient parameterization of the kernel κ , FNOs rely on the Fourier transform as a projection of a function onto the Fourier domain (Li et al. 2020a). Given an input function $\mathbf{g}(x)$, for

instance, FNOs use the following kernel integral operator \mathcal{K} parameterized by ψ (Li et al. 2020a):

$$(\mathcal{K}(\psi)\mathbf{g})(x) = \mathcal{F}^{-1}(\mathbf{R}_\psi \odot \mathcal{F}(\mathbf{g}))(x), \quad \forall x \in S, \quad (6)$$

where \mathcal{F} denotes the fast Fourier transform and \mathcal{F}^{-1} its inverse. \odot denotes the elementwise multiplication, and \mathbf{R}_ψ denotes a tensor representing a global convolutional kernel. Note that the global convolutional kernel \mathbf{R}_ψ is the only trainable parameter, i.e., $\psi \equiv \mathbf{R}_\psi$.

Proposed Method

In this section, we describe our motivations, followed by the overall workflow and detailed model architecture.

Motivations

The ODE function f in Eq. (1) is to learn the time-derivative operator of the hidden state $\mathbf{h}(t)$, in which case we can consider that f is a special operator suitable for defining NODEs. Therefore, our proposed NODE concept can be rewritten as follows:

$$\mathbf{h}(t_1) = \mathbf{h}(t_0) + \int_{t_0}^{t_1} \mathcal{D}_f(\mathbf{h}(t), t; \theta_f)dt, \quad (7)$$

where \mathcal{D}_f is the time-derivative operator. However, we found that FNOs (Eq. (6)) are not suitable for use in general machine learning tasks (see Section for detailed discussion). Therefore, we propose our own extended FNO method, called branched Fourier neural operator (BFNO).

Many NODE-based papers have habitually used conventional architectures, e.g., fully-connected layers, convolutional layers, and rectified linear units (ReLUs), to define the ODE function (Dupont, Doucet, and Teh 2019; Norcliffe et al. 2020; Xia et al. 2021; Nguyen et al. 2022; Cho et al. 2022). In comparison with them, our method provides a more rigorous way for defining the ODE function (or the time-derivative operator). Our main goal in this paper is to enhance the model accuracy by learning the ODE function more accurately.

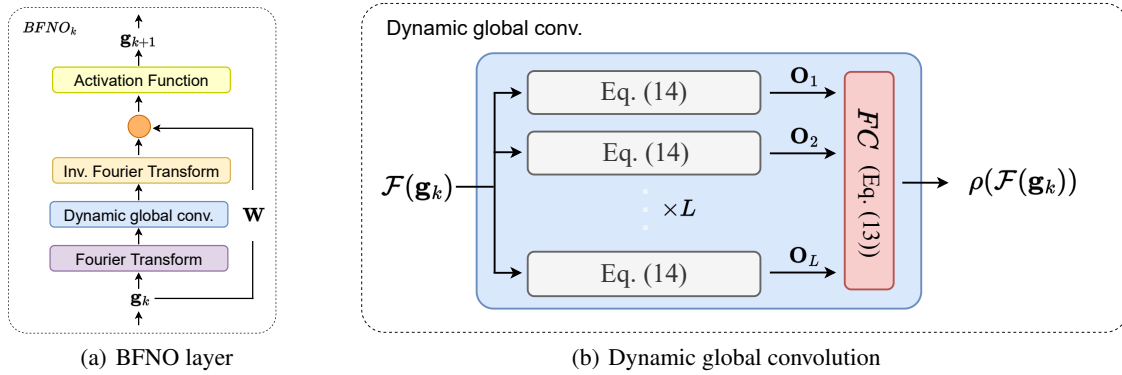


Figure 2: The detailed proposed architecture. We perform the Fourier transform, the dynamic global convolutional operation, and the inverse Fourier transform to process \mathbf{g}_k in conjunction with one additional transformation by the transformation matrix \mathbf{W} .

Overall Workflow

The overall model architecture of our proposed BFNO-NODE is shown in Figure 1. There are three major points in our design in contrast to the original NODE design: i) the neural operator-based ODE function, ii) the BFNO layer, and iii) the dynamic global convolution. We first sketch the overall mechanism before describing details as follows:

1. Given $\mathbf{h}(t)$ and t , we are interested in calculating $\frac{d\mathbf{h}(t)}{dt}$. Its first step is encoding $\mathbf{h}(t)$ and t into a representation \mathbf{g}_0 .
2. A series of N BFNO layers follows to calculate \mathbf{g}_N , which will be decoded into $\frac{d\mathbf{h}(t)}{dt}$. Therefore, the key in our design is how the BFNO layer works.
 - (a) In the k -th BFNO layer shown in Figure 2 (a), the first step is to apply the Fourier transform to its input \mathbf{g}_k , followed by a dynamic global convolution and the inverse Fourier transform.
 - (b) According to the convolution theorem (Bracewell and Bracewell 1986), the global convolution becomes an elementwise multiplication with a global filter in the Fourier domain (cf. Eq. (14)). Our dynamic global convolution i) applies L different filters, and ii) aggregates those L different outcomes with a fully-connected layer (cf. Eq. (13)) as shown in Figure 2 (b).
 - (c) There is another processing path for applying the linear transformation \mathbf{W} to \mathbf{g}_k .
 - (d) The previous two results are added, and a non-linear activation function finally outputs \mathbf{g}_{k+1} .

BFNO-NODE

Now, we describe our proposed method in detail.

Neural operator-based ODE functions: We propose the following ODE function with our newly designed BFNO layer:

$$\begin{aligned}
 \mathbf{g}_0 &= \text{Encoder}(\mathbf{h}(t), t), \\
 \mathbf{g}_{k+1} &= \text{BFNO}_k(\mathbf{g}_k), 0 \leq k \leq N-1, \\
 \frac{d\mathbf{h}(t)}{dt} &= \text{Decoder}(\mathbf{g}_N),
 \end{aligned} \tag{8}$$

where BFNO_k is the k -th BFNO layer. Both the time-augmented encoder (*Encoder*) and the decoder (*Decoder*) consist of fully-connected layers. The size of encoded vector \mathbf{g}_k , $0 \leq k \leq N$, is a hyperparameter. We use $\dim(\mathbf{g})$ to denote the size of \mathbf{g}_k . The output size of the decoder is the size of $\mathbf{h}(t)$. The size of $\mathbf{h}(t)$, denoted $\dim(\mathbf{h})$, is also a hyperparameter.

Branched Fourier neural operator (BFNO) layers: Figure 2 (a) shows the structure of the BFNO layer. The update process between the input value \mathbf{g}_k and the output value \mathbf{g}_{k+1} of the k -th BFNO layer is expressed as follows:

$$\mathbf{g}_{k+1}(x) = \text{BFNO}_k(\mathbf{g}_k(x)) \tag{9}$$

$$\Rightarrow \sigma((\mathcal{K}(\mathbf{h}, t; \psi)\mathbf{g}_k)(x) + \mathbf{W}\mathbf{g}_k(x)) \tag{10}$$

$$\Rightarrow \sigma(\mathcal{F}^{-1}(\rho(\mathcal{F}(\mathbf{g}_k)))(x) + \mathbf{W}\mathbf{g}_k(x)), \tag{11}$$

where the first term is a kernel integral operator parameterized by ψ , and the second term is a linear transformation parameterized by \mathbf{W} . σ is an activation function. ρ is the dynamic global convolutional operation for the input \mathbf{g}_k , which will be described shortly.

Under the regime of our proposed BFNO, \mathbf{g} is theoretically a spatial function. However, infinite-dimensional spatial functions cannot be processed by modern computers and therefore, all neural operator methods implement their discretized versions, which is also the case for our method. As a result, our actual implementations are as follows:

$$\mathbf{g}_{k+1} = \sigma(\mathcal{F}^{-1}(\rho(\mathcal{F}(\mathbf{g}_k))) + \mathbf{W}\mathbf{g}_k), \tag{12}$$

where \mathbf{g}_k (resp. \mathbf{g}_{k+1}) is an input (resp. output) vector. $\mathbf{W} \in \mathbb{R}^{\dim(\mathbf{g}) \times \dim(\mathbf{g})}$ is a linear transformation matrix, whose sizes are square matrix with the same number of rows and columns.

Dynamic global convolution ρ : We note that in our method, the dynamic global convolutional operation ρ is defined by the following method:

$$\rho(\mathcal{F}(\mathbf{g}_k)) = \text{FC}(\mathbf{O}_1, \mathbf{O}_2, \dots, \mathbf{O}_L), \tag{13}$$

$$\mathbf{O}_i = \mathbf{R}_i \odot \mathcal{F}(\mathbf{g}_k), \tag{14}$$

where FC means a fully-connected layer, and \mathbf{O}_i is the elementwise multiplication between the i -th global convolutional kernel \mathbf{R}_i and $\mathcal{F}(\mathbf{g}_k)$. Therefore, FC is to learn

| Method | Number of parameters | | | | Test Accuracy | | | |
|---------------|----------------------|----------|-----------|---------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| | MNIST | CIFAR-10 | CIFAR-100 | STL-10 | MNIST | CIFAR-10 | CIFAR-100 | STL-10 |
| NODE | 85,316 | 173,611 | 646,021 | 521,512 | 0.9531 ± 0.0042 | 0.5466 ± 0.0051 | 0.2178 ± 0.0024 | 0.3469 ± 0.0087 |
| ANODE | 85,462 | 172,452 | 645,121 | 520,180 | 0.9816 ± 0.0024 | 0.6025 ± 0.0032 | 0.2566 ± 0.0052 | 0.3711 ± 0.0089 |
| SONODE | 86,179 | 171,635 | 645,081 | 521,532 | 0.9824 ± 0.0013 | 0.6132 ± 0.0073 | 0.2591 ± 0.0058 | 0.3646 ± 0.0091 |
| HBNODE | 85,931 | 172,916 | 646,338 | 521,248 | 0.9814 ± 0.0011 | 0.5989 ± 0.0035 | 0.2549 ± 0.0057 | 0.3426 ± 0.0036 |
| GHBNODE | 85,931 | 172,916 | 646,338 | 521,248 | 0.9817 ± 0.0005 | 0.6085 ± 0.0050 | 0.2608 ± 0.0076 | 0.3648 ± 0.0078 |
| NesterovNODE | 85,930 | 172,915 | 647,869 | 523,351 | 0.9824 ± 0.0015 | 0.5996 ± 0.0033 | 0.2508 ± 0.0089 | 0.3595 ± 0.0051 |
| GNesterovNODE | 85,930 | 172,915 | 647,869 | 523,351 | 0.9807 ± 0.0013 | 0.6172 ± 0.0064 | 0.2568 ± 0.0060 | 0.3690 ± 0.0091 |
| AdamNODE | 85,832 | 171,904 | 644,974 | 521,192 | 0.9834 \pm 0.0004 | 0.6264 ± 0.0015 | 0.2405 ± 0.0049 | 0.3606 ± 0.0019 |
| BFNO-NODE | 85,591 | 171,333 | 644,547 | 514,291 | 0.9752 ± 0.0021 | 0.6289 \pm 0.0054 | 0.2890 \pm 0.0094 | 0.4455 \pm 0.0029 |

Table 1: The size of models and the test accuracy for MNIST, CIFAR-10, CIFAR-100, and STL-10 image classification. Our BFNO-based NODEs outperform other methods. Each baseline method has its own design guidance for the ODE function f as in our method.

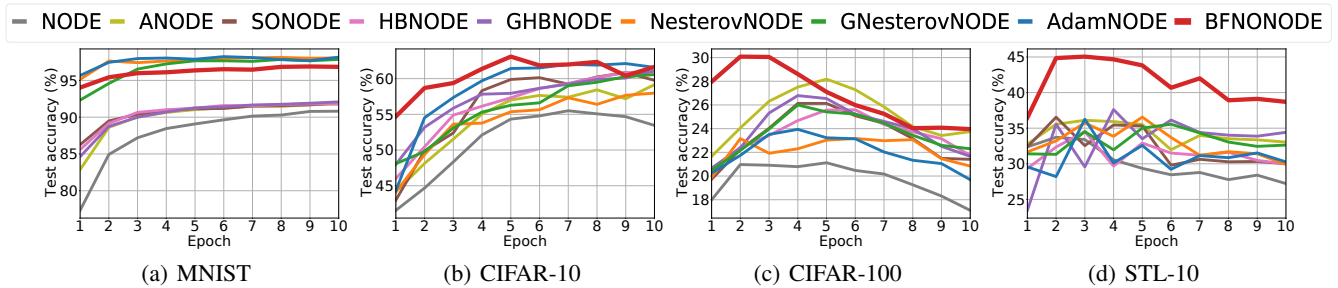


Figure 3: The proposed BFNO-NODE outperforms existing state-of-the-art methods for test accuracy — the results of train loss and NFE are in Appendix (Cho et al. 2023). We compare with other NODE-based models, such as vanilla NODE, ANODE, SONODE, HBNODE, GHBNODE, NesterovNODE, GNesterovNODE, and AdamNODE.

how to dynamically aggregate the L different global convolution outcomes, i.e., $\{\mathbf{O}_1, \mathbf{O}_2, \dots, \mathbf{O}_L\}$, where L is a hyperparameter. The learnable parameters (kernels) in \mathbf{R}_i , $1 \leq i \leq L$ and FC constitute our proposed dynamic global convolutional operation with respect to the input $\mathcal{F}(\mathbf{g}_k)$.

Experiments

In this section, we compare our model with existing baseline models for three general machine learning tasks: image classification, time series classification, and image generation. In comparison with the experiment set of HBNODE (Xia et al. 2021), our experiment set covers more general machine learning tasks and several more datasets. Our software and hardware environments are as follows: UBUNTU 18.04 LTS, PYTHON 3.6, TORCHDIFFEQ, PYTORCH 1.10.2, CUDA 11.4, i9 CPU, and NVIDIA RTX A6000.

Image Classification

Datasets: We test baselines and our model with the following four image classification benchmarks: MNIST (LeCun, Cortes, and Burges 2010), CIFAR-10 (Krizhevsky, Hinton et al. 2009), CIFAR-100 (Krizhevsky, Hinton et al. 2009), and STL-10 (Coates, Ng, and Lee 2011).

Experimental environments: In this experiment, NODE (Chen et al. 2018), ANODE (Dupont, Doucet, and

Teh 2019), SONODE (Norcliffe et al. 2020), HBNODE, GHBNODE (Xia et al. 2021), NesterovNODE, GNesterovNODE (Nguyen et al. 2022), and AdamNODE (Cho et al. 2022) are used as baselines, which cover a prominent set of enhancements for NODEs.

We use a learning rate of 0.001 and a batch size of 64. In all datasets, we use DOPRI with its default tolerance settings to solve the integral problem. For HBNODE and GHBNODE, the damping parameter γ is set to $\text{sigmoid}(\pi)$, where π is trainable and initialized to -3 . For all NODE-based baselines, we use three convolutional layers to define their ODE functions and for our method, we use two BFNO layers ($N = 2$) and two kernels ($L = 2$). For fair comparison, their model sizes (i.e., the number of parameters) are almost the same. (cf. Table 1).

In order to evaluate in various aspects, we consider the following evaluation metrics: i) the convergence speed of train loss, and most importantly ii) test accuracy. In addition, the number of function evaluations (NFEs) is an effective metric to measure the complexity of the forward and backward computation for NODE-based models. We focus on the accuracy in the main paper and refer readers to Appendix (Cho et al. 2023) for additional analyses on the NFEs and train loss curves.

| Method | Accuracy |
|-----------------------------|-------------------------------------|
| RNN- Δ_t | 0.797 \pm 0.003 |
| RNN-Impute | 0.795 \pm 0.008 |
| RNN-D | 0.800 \pm 0.010 |
| GRU-D | 0.806 \pm 0.007 |
| RNN-VAE | 0.343 \pm 0.040 |
| ODE-RNN | 0.829 \pm 0.016 |
| GHBNODE-RNN | 0.838 \pm 0.017 |
| GNesterovNODE-RNN | 0.840 \pm 0.016 |
| Latent-ODE (RNN enc.) | 0.835 \pm 0.010 |
| Latent-ODE (ODE enc.) | 0.846 \pm 0.013 |
| Augmented-ODE | 0.811 \pm 0.031 |
| ACE-Latent-ODE | 0.859 \pm 0.011 |
| Latent-BFNO-NODE (ODE enc.) | 0.874 \pm 0.004 |

Table 2: Time series classification (HumanAct.)

Experimental results: Figure 3 shows the test accuracy curve in each dataset, and Table 1 summarizes the overall results. Our method, shown in red, had the fastest convergence speed, outperforming many baselines. For three datasets except for MNIST, our method shows far better test accuracy curves than those of other baselines. Other baselines’ train loss values were much higher than those of our method in MNIST. In particular, our method dominates the test accuracy in STL-10, achieving higher accuracy consistently. For MNIST, however, AdamNODE shows a higher accuracy than ours. Appendix (Cho et al. 2023) contains the train loss and NFE curves. In general, our method’s forward and backward NFES are the smallest (or close to the smallest), which is an outstanding achievement. Our method enhances not only the test accuracy but also the computational complexity. In Appendix (Cho et al. 2023), we also visualize the feature maps from various methods. While vanilla NODEs’ feature map shows some glitches, ours shows much better clustering outcomes.

Time Series Classification

Datasets: HumanActivity (Kaluža et al. 2010) and Physionet (Silva et al. 2010) benchmark datasets are used to train and evaluate models for time series classification. The HumanActivity dataset includes information from five people with four sensors at their left ankle, right ankle, belt, and chest while performing a variety of activities such as walking, falling, lying down, rising from a lying position, etc. Each activity of a person was recorded five times for the sake of reliability. In this experiment, we classify each person’s input into one of the seven activities.

PhysioNet consists of 8,000 time series samples and is used to forecast the mortality of intensive care unit (ICU) populations. The dataset had been compiled from 12,000 ICU stays. They documented up to 42 variables and removed brief stays of less than 48 hours. Additionally, the data recorded in this way have a timestamp of the elapsed time since admission to the ICU. In this task, the patient’s life or death is determined based on records.

| Method | AUROC |
|-----------------------------|-------------------------------------|
| RNN- Δ_t | 0.787 \pm 0.014 |
| RNN-Impute | 0.764 \pm 0.016 |
| RNN-D | 0.807 \pm 0.003 |
| GRU-D | 0.818 \pm 0.008 |
| RNN-VAE | 0.515 \pm 0.040 |
| ODE-RNN | 0.833 \pm 0.009 |
| Latent-ODE (RNN enc.) | 0.781 \pm 0.018 |
| Latent-ODE (ODE enc.) | 0.829 \pm 0.004 |
| Augmented-ODE (ODE enc.) | 0.851 \pm 0.002 |
| ACE-Latent-ODE (ODE enc.) | 0.853 \pm 0.003 |
| Latent-BFNO-NODE (ODE enc.) | 0.852 \pm 0.001 |

Table 3: Time series classification (PhysioNet)

Experimental environments: We consider a variety set of RNN-based models, ODE-RNN, and Latent-ODE, following the evaluation protocol in (Rubanova, Chen, and Duvenaud 2019). In addition, we build two more baselines by replacing the ODE function of ODE-RNN with various methods: i) GHBNODE-based design to enhance ODE-RNN, denoted by GHBNODE-RNN; ii) GNesterov-based design, denoted by GNesterovNODE-RNN. To enhance Latent-ODE, we adopt the attention mechanism for NODEs proposed in (Jhin et al. 2021) and call it ACE-Latent-ODE. Our model, Latent-BFNO-NODE, is implemented by replacing the ODE function of Latent-ODE (ODE enc.) with the BFNO-NODE-based design. In this process, the hyperparameter L is fixed to 1 for efficiency. We use accuracy for HumanActivity (since the dataset is balanced). AUROC is used for Physionet, considering its imbalanced nature.

Experimental results: In Table 2, we summarize the results for HumanActivity. In general, RNN-based models do not show good accuracy. RNN-VAE shows the worst accuracy. NODE-based models significantly outperform them. Our proposed Latent-BFNO-NODE shows the best accuracy, followed by ACE-Latent-ODE. Our method also marks a small standard deviation in accuracy, which shows the robust nature of our method.

Table 3 summarizes the results for PhysioNet. ACE-Latent-ODE shows good performance for this dataset, which shows the appropriateness of its attention mechanism. However, our Latent-BFNO-NODE has an AUROC score comparable to it with a much smaller standard deviation.

Image Generation

Datasets: For our image generation task, we use MNIST and CIFAR-10. These two datasets are the most widely used for conducting generative task experiments for NODE-based and invertible neural network-based models (Grathwohl et al. 2018).

Experimental environments: To verify the performance of the BFNO as ODE function, we consider the baselines in (Finlay et al. 2020), which cover a prominent set of NODE-based and invertible neural network-based models. The ODE

| Method | MNIST | CIFAR-10 |
|-------------------|-------------|-------------|
| REALNVP | 1.06 | 3.49 |
| I-RESNET | 1.05 | 3.45 |
| GLOW | 1.05 | 3.35 |
| FFJORD | 0.99 | 3.40 |
| FFJORD-RNODE | 0.97 | 3.38 |
| FFJORD-BFNO-RNODE | 0.88 | 3.33 |

Table 4: Negative log-likelihood (in bits/dim) on test images. Lower is better.

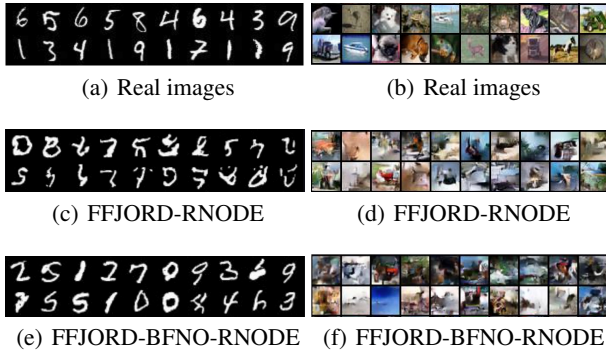


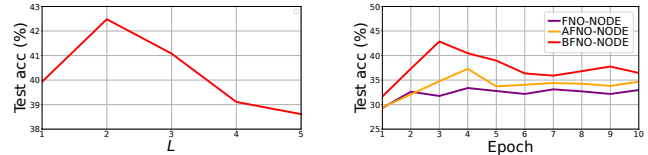
Figure 4: Real and produced samples from models for MNIST (left) and CIFAR-10 (right).

function $f(\mathbf{h}(t), t; \theta_f)$ in these baselines consists of four convolutional layers with 3×3 kernels and softplus activations. These layers contain 64 hidden units, and the time t is concatenated to the spatial input vector $\mathbf{h}(t)$ (as a side channel). The Gaussian Monte-Carlo trace estimator is used to calculate the log-probability of a generated sample via the change of variable theorem. Using the Adam optimizer with a learning rate of 0.001, we train on a single GPU with a batch size of 200 for 100 epochs. We use a model with the kinetic regularization proposed in (Finlay et al. 2020), and we propose FFJORD-BFNO-RNODE by replacing ODE function f of FFJORD-RNODE with BFNO structure. Our model consists of three BFNO layers and the softplus activations in between them. The hyperparameter L is fixed to 3. In addition, the transformation \mathbf{W} of BFNO is replaced with a convolutional operation with a 3×3 kernel since this task is generation.

Experimental results: In Table 4, it can be seen that the negative log-likelihood (NLL) of our model is better than those of all other baselines by large margins. In Figure 4, in addition, we show real images and fake images by our method and FFJORD-RNODE, which has the best performance among the baselines. In general, all methods show acceptable fake images.

Ablation Studies

We conduct two key ablation studies: 1) checking the model accuracy by changing the number of parallelized global convolutions (L) in BFNO, and 2) comparing BFNO with existing neural operator methods, such as FNO and AFNO.



(a) Ablation study on the number of parallelized global convolutions (L) (STL-10)

(b) Ablation study on the neural operator type (STL-10)

Figure 5: Ablation study on the two key design options of BFNO. Other figures are in Appendix (Cho et al. 2023).

Accuracy by the number of parallelized global convolutions (L) in BFNO: As shown in Figure 2 (b), a BFNO layer has multiple parallelized convolutions which are later merged by a fully-connected layer. The number of parallelized global convolutions in BFNO, denoted by L , is a key hyperparameter. Figure 5 (a) shows the model accuracy for various settings for L on image classification with STL-10. Too many or small convolutions lead to sub-optimal outcomes and $L = 2$ is the best configuration in this ablation study. We fathom that using too many convolutions result in overfitting, which degrades the model accuracy.

Comparison with FNO and AFNO: Figure 5 (b) compares our method with FNO and AFNO. As shown, our BFNO-NODE consistently outperforms them throughout the entire training epoch. FNO was developed to model PDE operators, and AFNO aims at improving the vision transformer by replacing its spatial mixer with a special neural operator. Since they were not designed for NODEs, however, they fail to show effectiveness in our experiments when used to define the ODE function of NODEs. Similar patterns are observed in other datasets, and the results are shown in Appendix (Cho et al. 2023).

Conclusions

Enhancing NODEs by adopting advanced ODE function architectures has been one active research trend for the past couple of years. However, existing approaches did not pay attention to the fact that the ODE function learns a special differential operator. In this work, we presented how neural operators can be used to define the ODE function and learn the differential operator. However, the naïve adoption of existing neural operators, such as FNO and AFNO, does not enhance NODEs. To this end, we designed a special neural operator architecture, called branched Fourier neural operators (BFNOs). Our dynamic global convolutional method with multiple parallelized global convolutions significantly improves the efficacy of various NODE-based models for three general machine learning tasks: image classification, time series classification, and image generation. Since the role of the ODE function is, in fact, applying a differential operator to the input $\mathbf{h}(t)$, our operator-based approach naturally shows the best fit in our experiments. Our ablation study with FNO and AFNO clearly shows that our BFNO design is one of the most important factors in our experiments.

Ethical Statement

PhysioNet contains much personal health information. However, it was released after removing 90% of observation to protect the privacy of the patients who have contributed their information. Therefore, our work does not have any related ethical concerns.

Acknowledgments

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded at the Korean government (MSIT) (No. 2020-0-01361, Artificial Intelligence Graduate School Program by Yonsei University, 10%), and (No.2022-0-00857, Development of AI/databased financial/economic digital twin platform, 90%).

References

- Bracewell, R. N.; and Bracewell, R. N. 1986. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York.
- Cao, S. 2021. Choose a transformer: Fourier or galerkin. *Advances in Neural Information Processing Systems*, 34: 24924–24940.
- Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. *Advances in neural information processing systems*, 31.
- Cho, S.; Hong, S.; Lee, K.; and Park, N. 2022. AdamNODEs: When Neural ODE Meets Adaptive Moment Estimation. *arXiv preprint arXiv:2207.06066*.
- Cho, W.; Cho, S.; Jin, H.; Jeon, J.; Lee, K.; Hong, S.; Lee, D.; Choi, J.; and Park, N. 2023. Operator-learning-inspired Modeling of Neural Ordinary Differential Equations. *arXiv:2312.10274*.
- Choi, J.; Jeon, J.; and Park, N. 2021. LT-OCF: learnable-time ODE-based collaborative filtering. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 251–260.
- Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 215–223. JMLR Workshop and Conference Proceedings.
- Dormand, J.; and Prince, P. 1980. A family of embedded Runge-Kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1): 19 – 26.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Dupont, E.; Doucet, A.; and Teh, Y. W. 2019. Augmented neural odes. *Advances in Neural Information Processing Systems*, 32.
- Finlay, C.; Jacobsen, J.-H.; Nurbekyan, L.; and Oberman, A. M. 2020. How to train your neural ODE: the world of Jacobian and kinetic regularization. In *ICML*.
- Grathwohl, W.; Chen, R. T.; Bettencourt, J.; Sutskever, I.; and Duvenaud, D. 2018. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*.
- Guibas, J.; Mardani, M.; Li, Z.; Tao, A.; Anandkumar, A.; and Catanzaro, B. 2021. Adaptive fourier neural operators: Efficient token mixers for transformers. *arXiv preprint arXiv:2111.13587*.
- Gupta, G.; Xiao, X.; and Bogdan, P. 2021. Multiwavelet-based operator learning for differential equations. *Advances in Neural Information Processing Systems*, 34: 24048–24062.
- Haber, E.; and Ruthotto, L. 2017. Stable architectures for deep neural networks. *Inverse problems*, 34(1): 014004.
- Jeon, J.; Kim, J.; Song, H.; Cho, S.; and Park, N. 2022. GT-GAN: General Purpose Time Series Synthesis with Generative Adversarial Networks. *arXiv preprint arXiv:2210.02040*.
- Jhin, S. Y.; Jo, M.; Kong, T.; Jeon, J.; and Park, N. 2021. Ace-node: Attentive co-evolving neural ordinary differential equations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 736–745.
- Jhin, S. Y.; Lee, J.; Jo, M.; Kook, S.; Jeon, J.; Hyeong, J.; Kim, J.; and Park, N. 2022. EXIT: Extrapolation and Interpolation-based Neural Controlled Differential Equations for Time-series Classification and Forecasting. In *Proceedings of the ACM Web Conference 2022*, 3102–3112.
- Kaluža, B.; Mirchevska, V.; Dovgan, E.; Luštrek, M.; and Gams, M. 2010. An agent-based approach to care in independent living. In *International joint conference on ambient intelligence*, 177–186. Springer.
- Kidger, P.; Morrill, J.; Foster, J.; and Lyons, T. 2020. Neural controlled differential equations for irregular time series. *Advances in Neural Information Processing Systems*, 33: 6696–6707.
- Kim, J.; Lee, C.; Shin, Y.; Park, S.; Kim, M.; Park, N.; and Cho, J. 2022. SOS: Score-based Oversampling for Tabular Data. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 762–772.
- Kovachki, N.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2021. Neural operator: Learning maps between function spaces. *arXiv preprint arXiv:2108.08481*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- LeCun, Y.; Cortes, C.; and Burges, C. 2010. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2.
- Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2020a. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*.
- Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2020b. Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*.
- Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Stuart, A.; Bhattacharya, K.; and Anandkumar, A. 2020c. Multipole

graph neural operator for parametric partial differential equations. *Advances in Neural Information Processing Systems*, 33: 6755–6766.

Liu, X.; Xiao, T.; Si, S.; Cao, Q.; Kumar, S.; and Hsieh, C.-J. 2019. Neural sde: Stabilizing neural ode networks with stochastic noise. *arXiv preprint arXiv:1906.02355*.

Lu, L.; Jin, P.; and Karniadakis, G. E. 2019. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*.

Lu, Y.; Zhong, A.; Li, Q.; and Dong, B. 2018. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In *International Conference on Machine Learning*, 3276–3285.

Morrill, J.; Salvi, C.; Kidger, P.; and Foster, J. 2021. Neural rough differential equations for long time series. In *International Conference on Machine Learning*, 7829–7838. PMLR.

Nguyen, N.; Nguyen, T. M.; Huyen, V. T. K.; Osher, S.; and Vo, T. 2022. Improving Neural Ordinary Differential Equations with Nesterov’s Accelerated Gradient Method. In *Advances in Neural Information Processing Systems*.

Norcliffe, A.; Bodnar, C.; Day, B.; Simidjievski, N.; and Liò, P. 2020. On second order behaviour in augmented neural odes. *Advances in Neural Information Processing Systems*, 33: 5911–5921.

Rahman, M. A.; Ross, Z. E.; and Azizzadenesheli, K. 2022. U-no: U-shaped neural operators. *arXiv preprint arXiv:2204.11127*.

Reddy, J. N. 2019. *Introduction to the finite element method*. McGraw-Hill Education.

Rubanov, Y.; Chen, R. T.; and Duvenaud, D. K. 2019. Latent ordinary differential equations for irregularly-sampled time series. *Advances in neural information processing systems*, 32.

Silva, I.; Moody, G.; Scott, D. J.; Celi, L. A.; and Mark, R. G. 2010. Predicting In-Hospital Mortality of ICU Patients: The PhysioNet/Computing in Cardiology Challenge 2012. *Comput Cardiol*, 39: 245–248.

Smith, G. D.; and Smith, G. D. 1985. *Numerical solution of partial differential equations: finite difference methods*. Oxford university press.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.

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

Wen, G.; Li, Z.; Azizzadenesheli, K.; Anandkumar, A.; and Benson, S. M. 2022. U-FNO—An enhanced Fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163: 104180.

Xia, H.; Suliafu, V.; Ji, H.; Nguyen, T.; Bertozzi, A.; Osher, S.; and Wang, B. 2021. Heavy ball neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 34.