

# Semi-Implicit Neural Ordinary Differential Equations

Hong Zhang<sup>1</sup>, Ying Liu<sup>2</sup>, Romit Maulik<sup>1,3</sup>

<sup>1</sup>Argonne National Laboratory

<sup>2</sup>University of Iowa

<sup>3</sup>Pennsylvania State University

hongzhang@anl.gov, ying-liu-1@uiowa.edu, rmm7011@psu.edu

## Abstract

Classical neural ODEs trained with explicit methods are intrinsically limited by stability, crippling their efficiency and robustness for stiff learning problems that are common in graph learning and scientific machine learning. We present a semi-implicit neural ODE approach that exploits the partitionable structure of the underlying dynamics. Our technique leads to an implicit neural network with significant computational advantages over existing approaches because of enhanced stability and efficient linear solves during time integration. We show that our approach outperforms existing approaches on a variety of applications including graph classification and learning complex dynamical systems. We also demonstrate that our approach can train challenging neural ODEs where both explicit methods and fully implicit methods are intractable.

## 1 Introduction

Implicit neural networks (NNs) (Winston and Kolter 2020; Bai, Kolter, and Koltun 2019) have emerged as a new paradigm for designing infinite-depth networks. The traditional neural network consists of an *explicit* stack of layers. In contrast, an implicit neural network represents an implicit model that often involves a nonlinear system, such as neural ordinary differential equations (ODEs) (Chen et al. 2018; Dupont, Doucet, and Teh 2019; Massaroli et al. 2020; Kidger 2022) and equilibrium models (Winston and Kolter 2020; Bai, Kolter, and Koltun 2019). In order to solve the nonlinear system, an iterative procedure is often needed. For example, neural ODEs require a time integrator to obtain the solution at desired time points. Deep equilibrium (DEQ) models (Bai, Kolter, and Koltun 2019) solve an implicit system with fixed-point iteration. Differentiating through the iterative procedure with direct backpropagation can be intractable, however, due to the memory consumption that increases linearly. Key to the success of existing implicit NNs is using high-level adjoint methods for gradient calculation. The adjoint methods can reduce the memory footprint significantly for both equilibrium models (Winston and Kolter 2020) and ODE models (Chen et al. 2018; Zhuang et al. 2020; Zhang and Zhao 2022).

Despite the advancement in gradient calculation approaches, however, implicit NNs are still limited by the bottleneck created by computationally expensive algorithms for solving the *linear or nonlinear system*. New methods have been developed to circumvent the need to solve the original linear systems. The Jacobian-free backpropagation (JFB) method in (Fung et al. 2022) exploits a preconditioning technique. Proximal implicit neural ODEs (Baker et al. 2022) approximate the solution of the linear system with an optimization algorithm. However, all these methods are rooted in low-order iterative procedures, which eases differentiation but suffers from suboptimal convergence speed and instability. Traditional linear solvers, including those highly optimized for AI accelerators such as GPUs, have favorable convergence properties but are not widely used in machine learning due to training efficiency constraints.

For neural ODEs, an additional bottleneck comes from the *numerical stability*. Explicit solvers, especially those based on Runge-Kutta (RK) methods, have been the default choice for the training of neural ODEs. These methods are only *conditionally* stable, requiring small time steps to maintain stability. The step size can be extremely small for *stiff* ODEs. Stiffness may result from dramatically different time scales involved in the dynamics that the ODE models, or from the spatial discretization of partial differential equations (PDEs) on a grid. For example, the Graph Neural Diffusion (GRAND) model, recently proposed for building continuous-depth graph NNs (GNNs) (Xhonneux, Qu, and Tang 2020) using graph neural ODEs (Poli et al. 2019), results in stiff diffusion equations. With explicit methods, one has to use a step size as small as 0.005 to main stability, leading to 43,908 NN evaluations per epoch. See Sec. 5 for details on the experiments on GRAND and more examples.

This work makes the following contributions:

- We propose a *semi-implicit* neural ODE (SINODE), a new type of implicit neural network, based on a linear-nonlinear *partitioning* of the system. The partitioned system is integrated with implicit-explicit (IMEX) methods, providing significant improvement in computational efficiency and stability over traditional neural ODEs.
- We develop a *differentiable* IMEX ODE solver as the backbone of SINODE. The solver provides reverse-accurate gradients through a *discrete adjoint* approach, avoiding the cost of backpropagation through the ODE

solver; it also enables state-of-the-art linear solution techniques when solving the implicit functions for mini-batch training of SINODE.

- We demonstrate the efficacy and efficiency of SINODE with graph learning tasks and two time-series prediction problems that are inherently *nonlinear*, *stiff*, or *chaotic*. The results also show that off-the-shelf solvers can be adapted to machine learning models to provide *stability* and *efficiency*.

## 2 Preliminaries

**Adjoint and Automatic Differentiation** Reverse-mode automatic differentiation, as widely used in scientific computing and machine learning, is also an adjoint method implemented at a low level. The adjoint methods used for implicit networks are derived at a high level of abstraction. We will take the DEQ model and the neural ODE model as examples. In a continuous view, the forward pass of the DEQ model can be considered as a root-finding procedure for a nonlinear function, and the loss depends on the root:

$$z = \mathbf{f}(z, u; p), \quad \text{Loss } \ell = \Phi(z), \quad (1)$$

where  $z$  is the equilibrium state for input  $u$  and  $\mathbf{f}$  stands for the NN parameterized by  $p$ . The adjoint equation of the nonlinear problem (Bai, Kolter, and Koltun 2019) is a linear equation defined as

$$\left( \mathcal{I} - \frac{\partial \mathbf{f}(z, u; p)}{\partial z} \right)^T \boldsymbol{\lambda} = \left( \frac{d\ell}{dz} \right)^T, \quad (2)$$

where the derivative  $\frac{d\ell}{dz}$  is a row vector and the adjoint variable  $\boldsymbol{\lambda}$  is a column vector. The loss gradient (a column vector) can be calculated with

$$\nabla_p \ell = \left( \frac{\partial \mathbf{f}(z, u; p)}{\partial p} \right)^T \boldsymbol{\lambda}. \quad (3)$$

The adjoint model requires storing only  $z$  and  $u$  during training; therefore, the memory cost is independent of the number of iterations in the forward pass.

For an autonomous neural ODE

$$\frac{du(t)}{dt} = \mathbf{f}(u(t); p), \quad t \in [0, T], \quad \text{Loss } \ell = \Phi(u(T)), \quad (4)$$

its continuous adjoint (Kidger, Chen, and Lyons 2020) is a new ODE integrated backward from time  $T$  to 0:

$$\frac{d\boldsymbol{\lambda}(t)}{dt} = - \left( \frac{\partial \mathbf{f}(u(t); p)}{\partial u(t)} \right)^T \boldsymbol{\lambda}(t). \quad (5)$$

The loss gradient is computed with

$$\nabla_p \ell = \int_0^T \left( \frac{\partial \mathbf{f}(u(t); p)}{\partial p} \right)^T \boldsymbol{\lambda}(t) dt. \quad (6)$$

For notational brevity, throughout this paper we omit  $p$  in all parameterized functions and drop the arguments  $(\cdot)$  for functions and variables such as  $\mathbf{f}$ ,  $u$ , and  $\boldsymbol{\lambda}$  when there is no ambiguity.

**Reverse-accurate Adjoint Approaches** In a discrete adjoint approach, the adjoint is derived based on the discretized version of the continuous equations (4). Without loss of generality, we assume (4) is discretized with a time-stepping algorithm, denoted by an operator  $\mathcal{N}$  that propagates the solution from one time step to another:

$$u_{n+1} = \mathcal{N}(u_n), \quad n = 0, 1, \dots, N. \quad (7)$$

The adjoint model of this time-stepping algorithm is

$$\begin{aligned} \boldsymbol{\lambda}_n &= \left( \frac{\partial \mathcal{N}}{\partial u}(u_n) \right)^T \boldsymbol{\lambda}_{n+1}, \\ \boldsymbol{\mu}_n &= \left( \frac{\partial \mathcal{N}}{\partial p}(u_n) \right)^T \boldsymbol{\lambda}_n + \boldsymbol{\mu}_{n+1}, \quad n = N-1, \dots, 0, \end{aligned} \quad (8)$$

with terminal condition  $\boldsymbol{\lambda}_N = \left( \frac{\partial \ell}{\partial u_N} \right)^T$ ,  $\boldsymbol{\mu}_N = \left( \frac{\partial \ell}{\partial p} \right)^T$ .

The adjoint variables  $\boldsymbol{\lambda}$  and  $\boldsymbol{\mu}$  denote the loss gradient with respect to the state  $u$  and the NN parameters  $p$ , respectively. For a practical implementation, (8) is typically derived in a case-by-case manner for different time-stepping algorithms (Zhang and Sandu 2014). The discrete adjoint model is conceptually equivalent to applying automatic differentiation to each iteration of the time-stepping procedure, therefore yielding the equally accurate gradient as obtained with direct backpropagation. Since the same trajectory (sequence of time steps) is used for the forward pass and the backward pass, one only needs to control step sizes during the forward pass when using adaptive time steps with the discrete adjoint approach.

## 3 The Partitioned Model

We consider an autonomous neural ODE framework with an additively partitioned right-hand side:

$$\frac{du}{dt} = \mathcal{G}(u) + \mathcal{H}(u), \quad (9)$$

where  $u(t) \in \mathbb{R}^d$  is the state,  $t$  is the time, and  $\mathcal{G} : \mathbb{R}^d \rightarrow \mathbb{R}^d$  and  $\mathcal{H} : \mathbb{R}^d \rightarrow \mathbb{R}^d$  are NNs, for example feed-forward NNs. Similar to conventional neural ODE, our framework takes an input  $u(t_0)$ , learns the representation of the dynamics, and predicts the output  $u(t_N)$  or a sequence of outputs  $u_{t_i}$ ,  $t = 1, \dots, N$  by solving (9) from the starting time  $t_0$  to the ending time  $t_N$  with a numerical ODE solver. While this framework is general and applicable for a wide range of applications, we are interested in a specific model with a nonlinear part and a linear part. For a purely linear neural ODE, which is rare in practice, we can use a partitioned model containing two linear parts. Without loss of generality, we assume  $\mathcal{H}$  is linear, and thus (9) can be rewritten to

$$\frac{du}{dt} = \underbrace{\mathcal{G}(u)}_{\text{Nonlinear}} + \underbrace{\mathcal{J}u}_{\text{Linear}}. \quad (10)$$

The reason for this choice will become clear later.

## 4 Training Algorithms

To train the partitioned NN, we adopt a semi-implicit approach to integrate (9) in a forward pass and compute the gradients with a discrete adjoint approach in a backward pass. The algorithm for each pass will be described below.

## 4.1 Forward Pass

In the forward pass, we solve the partitioned ODE (9) using implicit-explicit RK (IMEX-RK) methods. In particular, the right-hand side  $\mathcal{G}$  is treated explicitly while  $\mathcal{H}$  is treated implicitly. At each time step, the solution is advanced from  $t_n$  to  $t_{n+1}$  using

$$\begin{aligned} U^{(i)} &= u_n + \Delta t \sum_{j=1}^{i-1} a_{ij} \mathcal{G}^{(j)} + \Delta t \sum_{j=1}^i \tilde{a}_{ij} \mathcal{H}^{(j)}, \quad i = 1, \dots, s \\ u_{n+1} &= u_n + \Delta t \sum_{j=1}^s b_j \mathcal{G}^{(j)} + \Delta t \sum_{j=1}^s \tilde{b}_j \mathcal{H}^{(j)}, \end{aligned} \quad (11)$$

where  $U^{(i)}$  is the stage value,  $\Delta t$  is the step size,  $A = \{a_{ij}\}$  is strictly lower triangular,  $\tilde{A} = \{\tilde{a}_{ij}\}$  is lower triangular and can have zeros on the diagonal (these correspond to explicit stages), and  $\circ^{(i)}$  denotes stage index  $i$ .  $\mathcal{G}^{(j)}$  and  $\mathcal{H}^{(j)}$  are shorthand for  $\mathcal{G}(U^{(j)})$  and  $\mathcal{H}(U^{(j)})$ . The IMEX-RK scheme can be represented by the Butcher tableau consisting of  $(A = \{a_{ij}\}, b, c)$  for the explicit part and  $(\tilde{A} = \{\tilde{a}_{ij}\}, \tilde{b}, \tilde{c})$  for the implicit part. See the appendix for the coefficients of the IMEX methods used in this paper. The IMEX-RK methods are well known for their favorable stability properties for stiff systems (Ascher, Ruuth, and Spiteri 1997; Boscarino 2007; Boscarino and Russo 2009; Kennedy and Carpenter 2003).

**Linear Implicit System** At each stage, one needs to solve a nonlinear system for  $U^{(i)}$ , typically with an iterative method such as the Newton method. If  $\mathcal{H}$  is linear, however, (11) falls into a linear problem with respect to  $U^{(i)}$ :

$$(\mathcal{I} - \Delta t \tilde{a}_{ij} \mathcal{J}) U^{(i)} = u_n + \Delta t \sum_{j=1}^{i-1} (a_{ij} \mathcal{G}^{(j)} + \tilde{a}_{ij} \mathcal{J} U^{(j)}), \quad (12)$$

where  $\mathcal{J}$  is constant across the  $s$  stages and all the time steps. Thus, the linear-nonlinear partitioning can significantly increase computational efficiency by avoiding solving a nonlinear system. In practice, this partitioning is naturally applicable to many multiphysics problems. For example, in diffusion-reaction equations, diffusion can be represented by a linear term and treated implicitly in an IMEX setting. In neural-network-based models, layers such as convolution layers or fully connected layers without activation functions can be viewed as a linear map as well.

**Efficient Mini-batching** Consider a mini-batch  $\mathcal{B}_n = \{u_n^1, \dots, u_n^m\}$  consisting of  $m$  inputs at time  $t_n$ . The batch Jacobian in (12) can be represented in Kronecker product form:

$$\begin{aligned} [(\mathcal{I} - \Delta t \tilde{a}_{ij} \mathcal{J}) \otimes \mathcal{I}_m] U|_{\mathcal{B}}^{(i)} &= \mathcal{B}_n + \Delta t \sum_{j=1}^{i-1} a_{ij} \mathcal{G}_{\mathcal{B}}^{(j)} \\ &\quad + \Delta t \sum_{j=1}^{i-1} \tilde{a}_{ij} (\mathcal{J} \otimes \mathcal{I}_m) U_{\mathcal{B}}^{(j)}, \end{aligned} \quad (13)$$

where  $U|_{\mathcal{B}}$  is the corresponding stage values for  $\mathcal{B}_n$ . Since

the linear problem (13) is equivalent to

$$\begin{aligned} [\mathcal{I} - \Delta t \tilde{a}_{ij} \mathcal{J}] \text{reshape}(U|_{\mathcal{B}}^{(i)}, d, m) &= \text{reshape}(\mathcal{B}_n, d, m) \\ &\quad + \Delta t \sum_{j=1}^{i-1} a_{ij} \text{reshape}(\mathcal{G}_{\mathcal{B}}^{(j)}, d, m) \\ &\quad + \Delta t \sum_{j=1}^{i-1} \tilde{a}_{ij} \mathcal{J} \text{reshape}(U|_{\mathcal{B}}^{(j)}, d, m), \end{aligned} \quad (14)$$

we need to solve only a reduced linear system with  $m$  right-hand sides reshaped from the batched vectors. This problem can be solved conveniently by off-the-shelf linear solvers, including direct solvers and iterative solvers. With direct solvers, the factorization of the left-hand side  $\mathcal{I} - \Delta t \tilde{a}_{ij} \mathcal{J}$  can be reused across stages and time steps because  $\mathcal{J}$  is modified at the end of each training iteration when the NN parameters are updated by the optimizer. If  $\mathcal{J}$  is known a priori and not parameterized, the factorization can be reused even across mini-batches. Iterative solvers are typically chosen because of their excellent scalability and low memory requirement. If  $\mathcal{J}$  is large and sparse, using a matrix-free iterative solver may be advantageous.

## 4.2 Backward Pass

Unlike the vanilla neural ODE that employs a continuous adjoint approach, we use a discrete approach as an alternative to directly backpropagating through an ODE solver.

**Theorem 4.1.** *The gradient of loss  $\ell$  with respect to the NN parameters  $p$  for SINODE using the IMEX-RK methods in (11) can be calculated with the following discrete adjoint formula:*

$$\begin{aligned} (I - \Delta t \tilde{a}_{ii} \mathcal{H}_u^{T(i)}) \lambda_{n+1}^{(i)} &= \Delta t (b_i \mathcal{G}_u^{T(i)} + \tilde{b}_i \mathcal{H}_u^{T(i)}) \lambda_{n+1} \\ &\quad + \Delta t \mathcal{G}_u^{T(i)} \sum_{j=i+1}^s a_{ji} \lambda_{n+1}^{(j)} \\ &\quad + \Delta t \mathcal{H}_u^{T(i)} \sum_{j=i+1}^s \tilde{a}_{ji} \lambda_{n+1}^{(j)}, \\ &\quad i = s, \dots, 1, \\ \lambda_n &= \lambda_{n+1} + \sum_{j=1}^s \lambda_{n+1}^{(j)}, \\ \mu_{n+1}^{(i)} &= \Delta t (b_i \mathcal{G}_p^{T(i)} + \tilde{b}_i \mathcal{H}_p^{T(i)}) \lambda_{n+1} \\ &\quad + \Delta t \mathcal{G}_p^{T(i)} \sum_{j=i+1}^s a_{ji} \lambda_{n+1}^{(j)} \\ &\quad + \Delta t \mathcal{H}_p^{T(i)} \sum_{j=i}^s \tilde{a}_{ji} \lambda_{n+1}^{(j)}, \\ &\quad i = s, \dots, 1, \\ \mu_n &= \mu_{n+1} + \sum_{j=1}^s \mu_{n+1}^{(j)}, \end{aligned} \quad (15)$$

with terminal condition

$$\lambda_N = \frac{d\ell}{du}, \quad \mu_N = 0. \quad (16)$$

Here  $\lambda$  and  $\mu$  correspond to the partial derivatives of the loss with respect to the initial state and the NN parameters, respectively.

At the beginning of the backward pass, we initialize  $\lambda$  and  $\mu$  according to the terminal condition (16). Then they are propagated backward in time (the index  $n$  goes from  $N$  to 0) during the backward pass according to (15). At the end of the backward pass, we obtain the gradient

$$\nabla_p \ell = \mu_0^T.$$

If  $\tilde{a}_{ii} = 0$ , the first equation in (15) falls back to

$$\begin{aligned} \lambda_{n+1}^{(i)} &= \Delta t \left( b_i \mathcal{G}_u^{T(i)} + \tilde{b}_i \mathcal{H}_u^{T(i)} \right) \lambda_{n+1} \\ &+ \Delta t \mathcal{G}_u^{T(i)} \sum_{j=i+1}^s a_{ji} \lambda_{n+1}^{(j)} + \Delta t \mathcal{H}_u^{T(i)} \sum_{j=i+1}^s \tilde{a}_{ji} \lambda_{n+1}^{(j)}. \end{aligned} \quad (17)$$

For all  $\tilde{a}_{ii}$  that is nonzero, we need to solve a linear system that involves the transposed Jacobian. Nevertheless, the same strategy described in Sec. 4.1 can be applied to solving this transposed system. For efficiency, we compute the transposed Jacobian-vector products in (15) with automatic differentiation tools such as AutoGrad. Each time, the computational cost is one forward evaluation of the right-hand side  $\mathcal{G}$  or  $\mathcal{H}$  followed by a backward pass; and, more importantly, the memory cost depends only on the NN that approximates the ODE right-hand sides, not on the iterations needed by the linear solvers or the number of time steps. Note that the Jacobian depends on the ODE states  $u_n$  that can be stored during the forward pass. One can checkpoint selective states and recompute the missing states from the checkpoints during the backward pass. A checkpointing strategy (Zhang and Zhao 2022) can be used to balance memory consumption and recomputation.

The IMEX algorithms and their adjoints can be implemented in any major frameworks, such as PyTorch. We have implemented these algorithms as off-the-shelf solvers in PETSc (Balay et al. 2023) to leverage the state-of-the-art linear solvers (Mills et al. 2021) and the checkpointing techniques (Zhang and Constantinescu 2021, 2023) in PETSc. SINODE is implemented in the PNODE framework (Zhang and Zhao 2022) that integrates PyTorch and PETSc seamlessly.

## 5 Experiments

We validate and evaluate the performance of our methods when learning stiff ODE systems. Throughout this section we compare our methods with a variety of baseline methods including explicit and implicit methods. The neural network architectures we use follow the best settings identified in previous works (Linot et al. 2022; Chamberlain et al. 2021). The only modification necessary for using our method is to split the ODE right-hand side. Code is available at <https://github.com/caidao22/pnode>.

### 5.1 Graph Classification with GRAND

GRAND is an expansive and innovative class of GNNs that draws inspiration from the discretized diffusion PDE and

	NFEs		
	Cora	CoauthorCS	Photo
RK4	43908	7512	8572
Explicit Adams	10,995	1896	2169
Implicit Adams	1852	73	62
IMEX-RK2	<b>114</b>	<b>24</b>	<b>24</b>
IMEX-RK3	<b>228</b>	<b>48</b>	<b>48</b>
IMEX-RK4	<b>342</b>	<b>72</b>	<b>72</b>
IMEX-RK5	<b>456</b>	<b>96</b>	<b>96</b>
Crank–Nicolson	611	174	447

Table 1: NFEs per epoch for neural ODEs with explicit methods and implicit methods and SINODE with IMEX methods.

applies it to graphs. These architectures address prevailing challenges of current graph machine learning models such as depth, over-smoothing, bottlenecks, and graph rewiring. With a reinterpretation of the diffusion equation for graphs, the process can be depicted by an ODE system:

$$\frac{dx(t)}{dt} = (\mathbf{A}(x(t)) - \mathbf{I})x(t) \quad (18)$$

where  $\mathbf{A}(x(t)) = (a(x_i(t), x_j(t)))_{ij}$  is an attention matrix with the same structure as the adjacency matrix of the graph and  $\mathbf{I}$  is an identity matrix.

We introduce an ODE block where we apply the SINODE approach and split the right-hand side of (18) into a nonlinear part and a linear part:

$$\frac{dx(t)}{dt} = (A(x(t)) - I)x(t) = \underbrace{A(x(t))x(t)}_{\text{Nonlinear}} + \underbrace{(-Ix(t))}_{\text{Linear}}. \quad (19)$$

The nonlinear part is handled explicitly, and the linear part is handled implicitly with IMEX-RK methods. We use an iterative matrix-free solver in PETSc for solving the linear systems because of the simplicity of the linear part.

For assessment, we choose three benchmark datasets: Cora, CoauthorCS, and Photo. In addition to the baseline methods mentioned, we include the implicit Adams, which is claimed to perform well for this problem in (Chamberlain et al. 2021). Because of the stability constraints, we have to utilize a step size of 0.005 for explicit methods, while the implicit methods and the IMEX methods allow us to use a step size of 1 thanks to their superior stability properties.

Compared with the explicit methods and the implicit methods, our approach boasts advantages under two performance evaluation metrics. On the one hand, as depicted in Table 1, SINODE with IMEX-RK methods requires significantly fewer numbers of right-hand side function evaluations (NFEs) than explicit methods do. Each function evaluation corresponds strictly to one forward pass of the NN. Thus, NFE is usually a good metric that accounts for the dominant cost for training neural ODEs while being independent of the NN complexity. On the other hand, as illustrated in Fig. 1, SINODE with IMEX-RK methods takes much less

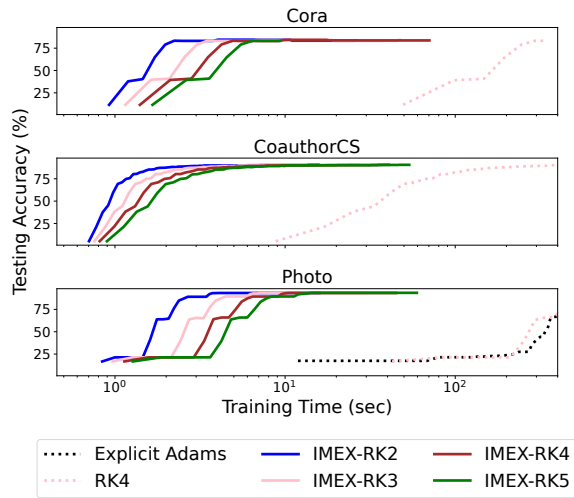


Figure 1: Testing accuracy versus training time for various time integrators using **explicit** and **semi-implicit** methods.

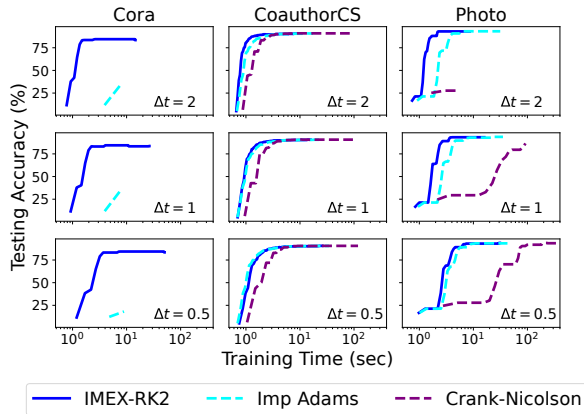


Figure 2: Testing accuracy vs. training time for various step sizes for **implicit** and **semi-implicit** methods. For Cora, Crank–Nicolson failed in the beginning of training, and implicit Adams failed in the middle of training.

training time than classical neural ODEs with explicit methods to achieve the same level of testing accuracy.

We have also performed a comparison between IMEX-RK2 and fully implicit methods for varying step sizes. Fig. 2 shows that IMEX-RK2 outperforms the implicit Adam and Crank–Nicolson significantly for Cora and Photo. Both implicit Adam and Crank–Nicolson cannot complete the training for Cora because the loss blows up.

## 5.2 Learning Dynamics for the Kuramoto–Sivashinsky (KS) Equation

The KS equation in one-dimensional space can be written as

$$\frac{\partial u}{\partial t} = -u \frac{\partial u}{\partial x} - \frac{\partial^2 u}{\partial x^2} - \frac{\partial^4 u}{\partial x^4}, \quad x \in [0, 22]. \quad (20)$$

It is a well-known nonlinear *chaotic* PDE for modeling complex spatiotemporal dynamics.

We train neural ODEs by approximating the nonlinear term  $u \frac{\partial u}{\partial x}$  with a fully connected multilayer perceptron (MLP) and approximating the remaining linear terms on the right-hand side of (20) with a fixed convolutional NN (CNN) layer. We use circular padding and a CNN filter of size 5 to mimic centered finite-difference operators for the antidiffusion and hyperdiffusion terms with periodic boundary conditions. The parameters of the filter are fixed to be  $[-1.0/\Delta x^4, 4.0/\Delta x^4 - 1.0/\Delta x^2, -6.0/\Delta x^4 + 2.0/\Delta x^2, 4.0/\Delta x^4 - 1.0/\Delta x^2, -1.0/\Delta x^4]$ , where the grid spacing  $\Delta x = L/N$ . A direct solver with reused LU factorizations is used to solve the linear systems associated with the CNN layer.




Figure 3: Ground truth (**Left**) vs. prediction (**Right**) using IMEX-RK3 with  $\Delta t = 0.2$ .

When applying SINODE, we treat the MLP part explicitly and the CNN part implicitly. A time step size 0.2 is used for the IMEX methods and the fully implicit method. Because of the stability constraints, the explicit methods do not work well until we decrease the step size to 0.001. Fig. 3 compares a true trajectory to predictions from SINODE. We can see that the predicted trajectories stay on the attractor for a long time and match the ground truth well.

In Fig. 4 we show the training loss versus the training time for all the methods in order to evaluate their efficiency. As expected, SINODE using IMEX methods significantly outperforms the other methods. For example, SINODE with IMEX-RK2 is  $\sim 47$  times faster than the classical neural ODE with Dopri5 and  $\sim 33$  times faster than the neural ODE with Crank–Nicolson to decrease the training loss to  $10^{-3}$ . Table 2 compares the average NFEs per epoch for different solvers. It shows that SINODE using IMEX-RK methods can reduce the NFEs remarkably. The superior performance of SINODE is mainly attributed to superior stability properties and the computational advantages of semi-implicit methods. Note that the improvement in training efficiency becomes more significant as the number of time steps increases because of the increased opportunity to reuse the LU factorization for solving the linear systems. In our experiments we use only one time step during the training, which reflects the worst-case performance of SINODE.

Training with neural ODEs becomes more difficult when the grid resolution increases. After we change the grid size from 64 to 512, the performance of neural ODEs with explicit methods or fully implicit methods degrades dramatically while SINODE maintains high efficiency. We can still use the time step size 0.2 for the training of SINODE, but the explicit methods require a step size on the order of  $10^{-7}$




Figure 4: Training loss vs. training time for KS. **Left:** Grid size 64. **Right:** Grid size 512. Only SINODE with IMEX methods achieve feasible training completion for the case with grid size 512.

because of the severe stability constraint from the hyperdiffusion term. We are not able to make the fully implicit methods work with the step size 0.2 because the nonlinear solver diverges frequently during the training process. This is not uncommon for classical PDE solvers since the condition number of the system increases as the grid resolution increases. Decreasing the step size by 10 times stabilizes the nonlinear solvers. However, this leads to a per-epoch training time of more than one hour whereas SINODE typically takes less than one second per epoch.

An interesting observation is that the lower-order IMEX-RK methods lead to better training efficiency than higher-order IMEX-RK methods, despite their slightly slower convergence speed. This is because higher-order IMEX-RK methods require more stages, leading to more linear solves at each time step. For the coarse-grid case (64), the second-order IMEX-RK performs  $\sim 2.5$  times better than the fifth-order IMEX-RK. For the fine-grid case (512), its gain is larger than a factor of 5. We expect that the gain to increase as the grid resolution further increases.

### 5.3 Learning Dynamics for the Viscous Burgers Equation

We consider the one-dimensional viscous Burgers equation

$$\frac{\partial u}{\partial t} = -u \frac{\partial u}{\partial x} + \nu \frac{\partial^2 u}{\partial x^2}, \quad x \in [0, 1], \quad (21)$$

with viscosity  $\nu = 8 \cdot 10^{-4}$ , which makes the solution known as Burger turbulence.

The architecture of the neural ODEs approximates the nonlinear term  $-u \frac{\partial u}{\partial x}$  with a fully connected nonlinear NN, and the linear term  $\nu \frac{\partial^2 u}{\partial x^2}$  with a fixed linear layer derived from a finite difference method. Again we choose the direct solver to benefit from the reuse of LU factorizations. In the comparison of different methods, we choose a time step size of 0.05 as a conservative choice for the four IMEX-RK methods and the fully implicit method. Explicit methods cause the solution to explode until the time step size decreases to as small as 0.001 when the grid size is 512 and to 0.0005 when the grid size is 1024.

Fig. 5 shows that the predictions made by SINODE are in good agreement with the true solutions for an initial condition from the testing dataset.

	NFEs			
	KSE		Burgers	
	grid 64	grid 512	grid 512	grid 1024
Dopri5	12619	—	22819	45619
RK4	24000	—	15200	30400
IMEX-RK2	<b>60</b>	<b>60</b>	<b>152</b>	<b>152</b>
IMEX-RK3	<b>120</b>	<b>120</b>	<b>304</b>	<b>304</b>
IMEX-RK4	<b>180</b>	<b>180</b>	<b>456</b>	<b>456</b>
IMEX-RK5	<b>240</b>	<b>240</b>	<b>608</b>	<b>608</b>
Crank–Nicolson	3195	—	13072	11883

Table 2: NFEs for neural ODEs with explicit methods and implicit methods and SINODE with IMEX methods. The numbers are rounded to the closest integer.




Figure 5: Snapshots of VBE trajectories for grid size 512 using IMEX-RK3.

Fig. 6 show that SINODE significantly outperforms the classical neural ODE for the two different grid settings. IMEX-RK3 is consistently the best-performing method. For the case with grid size 512, IMEX-RK3 is  $\sim 6X$  faster than RK4 in terms of training time per epoch. The speedups increase to  $\sim 10X$  for the case with grid size 1024. The speedups over Crank–Nicolson are  $\sim 3.4X$  and  $\sim 2X$  for the two cases. The gain in time-to-accuracy appears to be larger than  $2X$  because the training converges slower with Crank–Nicolson than with IMEX-RK3. Table 2 lists NFEs for different ODE solvers.




Figure 6: Training loss vs. training time for VBE. **Left:** Grid size 512. **Right:** Grid size 1024. The explicit Adams does not converge due to instability and is thus not shown.

## 6 Related Work

**Implicit Networks** Whereas explicit networks consist of a sequence of layers, implicit networks define the output as a solution to a nonlinear system. For example, the well-known neural ODE model (Chen et al. 2018) solves an ODE problem in the forward pass, typically with explicit integration methods such as RK methods. The DEQ model (Winston and Kolter 2020; Bai, Kolter, and Koltun 2019) finds an equilibrium solution to a nonlinear system with a fixed-point iteration method in the forward pass. These models require solving a nonlinear system by embedding an iterative solver in the architecture. Therefore, common numerical issues associated with the solver, such as *convergence* and *stability*, may be more frequently encountered during the training process. For example, explicit methods can be restricted to *infinitely small* step sizes when neural ODEs are applied to stiff problems. Fully implicit methods (Zhang and Zhao 2022) are *absolutely stable*, but the nonlinear solvers may fail to converge for ill-conditioned systems. In addition, solving the nonlinear system poses a computational bottleneck for training implicit networks. These limitations have motivated a variety of acceleration techniques such as IMEXnet (Haber et al. 2019), equation scaling (Kim et al. 2021), JFB (Fung et al. 2022) and proximal implicit neural ODE (Baker et al. 2022), as well as this work.

**Gradient Calculation for Implicit Networks** Using backpropagation directly in the backward pass is often intractable for training implicit networks because they may require a large number of iterations and the memory to store the computation graphs increases with this number. In order to address this issue, adjoint methods (Chen et al. 2018; Winston and Kolter 2020; Zhuang et al. 2020) have been widely used to compute the gradient. When applied to ODEs, these methods typically require storing selective states for the Jacobian evaluation during a forward pass, known as checkpointing, and restoring missing states during a backward pass via recomputation (Chen et al. 2018; Zhuang et al. 2020) or interpolation (Daulbaev et al. 2020).

**Reverse-accurate Neural ODEs** Although adjoint methods provide a memory-efficient alternative to direct backpropagation, they do not always generate an accurate gradient. Various researchers have (Onken and Ruthotto 2020; Zhang and Zhao 2022) pointed out that the vanilla neural

ODE (Chen et al. 2018) adopts a continuous adjoint approach that is not reverse accurate. In particular, the gradient is not consistent with the forward pass, even if the same time integrator and the same step size are used for solving the continuous adjoint equation (5) as shown in (Zhang and Zhao 2022). Reverse-accurate variants of neural ODEs have been developed in (Gholaminejad, Keutzer, and Biros 2019; Zhang et al. 2019; Zhuang et al. 2020, 2021; Matsubara, Miyatake, and Yaguchi 2023). These variants rely on back-propagation, checkpointing and recomputation to alleviate the memory cost. Zhang and Zhao (2022) proposed a *discrete adjoint* approach that achieves reverse accuracy and can leverage optimal checkpointing schedules (Zhang and Constantinescu 2021, 2023) to balance the trade-off between recomputation and memory usage.

## 7 Conclusion

We propose a semi-implicit approach, named SINODE, for learning stiff neural ODEs. Our approach leverages a nonlinear-linear partitioning of the ODE right-hand side and a differentiable IMEX solver that provides reverse-accurate gradient computation. The excellent stability property of IMEX methods enables large time steps to be used for integrating the ODEs, which is especially critical for stiff problems where other methods are inefficient. In addition, our approach allows us to use off-the-shelf linear solvers in mini-batch training, opening the opportunity of adapting well-established techniques from the other fields. We demonstrate the advantages of SINODE when applied to graph learning tasks and time series modeling of two challenging dynamical systems. We show that SINODE typically allows us to use a time step that is two orders of magnitude larger than that allowed by explicit RK methods, leading to significant speedup over classic neural ODEs. We also show that SINODE succeeds for very stiff cases where existing methods fail due to instability.

**Limitations** SINODE should not be used for nonstiff problems where classic neural ODEs with explicit methods could be more efficient. When direct linear solvers are used, varying step sizes may degrade the performance of SINODE because the Jacobian cannot be reused across time steps. However, the performance is not affected much when matrix-free iterative linear solvers are used.

## Acknowledgments

This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program through the FASTMath Institute and DOE-FOA-2493 "Data-intensive scientific machine learning", under contract DE-AC02-06CH11357 at Argonne National Laboratory. We gratefully acknowledge the computing resources provided and operated by the Joint Laboratory for System Evaluation (JLSE) at Argonne National Laboratory.

## References

- Ascher, U. M.; Ruuth, S. J.; and Spiteri, R. J. 1997. Implicit-explicit Runge–Kutta methods for time-dependent partial differential equations. *Applied Numerical Mathematics*, 25(2-3): 151–167.
- Bai, S.; Kolter, J. Z.; and Koltun, V. 2019. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, volume 32.
- Baker, J.; Xia, H.; Wang, Y.; Cherkaev, E.; Narayan, A.; Chen, L.; Xin, J.; Bertozzi, A. L.; Osher, S. J.; and Wang, B. 2022. Proximal Implicit ODE Solvers for Accelerating Learning Neural ODEs. arXiv:2204.08621.
- Balay, S.; Abhyankar, S.; Adams, M. F.; Benson, S.; Brown, J.; Brune, P.; Buschelman, K.; Constantinescu, E.; Dalcin, L.; Dener, A.; Eijkhout, V.; Faibussowitsch, J.; Gropp, W. D.; Hapla, V.; Isaac, T.; Jolivet, P.; Karpeev, D.; Kaushik, D.; Knepley, M. G.; Kong, F.; Kruger, S.; May, D. A.; McInnes, L. C.; Mills, R. T.; Mitchell, L.; Munson, T.; Roman, J. E.; Rupp, K.; Sanan, P.; Sarich, J.; Smith, B. F.; Zampini, S.; Zhang, H.; Zhang, H.; and Zhang, J. 2023. PETSc/TAO Users Manual. Technical Report ANL-21/39 – Revision 3.19, Argonne National Laboratory.
- Boscarino, S. 2007. Error analysis of IMEX Runge–Kutta methods derived from differential-algebraic systems. *SIAM Journal on Numerical Analysis*, 45(4): 1600–1621.
- Boscarino, S.; and Russo, G. 2009. On a class of uniformly accurate IMEX Runge–Kutta schemes and applications to hyperbolic systems with relaxation. *SIAM Journal on Scientific Computing*, 31(3): 1926–1945.
- Chamberlain, B.; Rowbottom, J.; Gorinova, M. I.; Bronstein, M.; Webb, S.; and Rossi, E. 2021. Grand: Graph neural diffusion. In *International Conference on Machine Learning*, 1407–1418. PMLR.
- Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D.; Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. 2018. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems*.
- Daulbaev, T.; Katrutsa, A.; Markeeva, L.; Gusak, J.; Cichocki, A.; and Oseledets, I. 2020. Interpolation Technique to Speed Up Gradients Propagation in Neural ODEs. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 16689–16700. Curran Associates, Inc.
- Dupont, E.; Doucet, A.; and Teh, Y. W. 2019. Augmented Neural ODEs. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Fung, S. W.; Heaton, H.; Li, Q.; McKenzie, D.; Osher, S.; and Yin, W. 2022. JFB: Jacobian-Free Backpropagation for Implicit Networks. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence, AAAI 2022*, volume 36.
- Gholaminejad, A.; Keutzer, K.; and Biros, G. 2019. AN-ODE: Unconditionally Accurate Memory-Efficient Gradients for Neural ODEs. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 730–736. International Joint Conferences on Artificial Intelligence Organization.
- Haber, E.; Lensink, K.; Treister, E.; and Ruthotto, L. 2019. IMEXnet A Forward Stable Deep Neural Network. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 2525–2534. PMLR.
- Kennedy, C. A.; and Carpenter, M. H. 2003. Additive Runge–Kutta schemes for convection–diffusion–reaction equations. *Applied Numerical Mathematics*, 44(1-2): 139–181.
- Kidger, P. 2022. On Neural Differential Equations. arXiv preprint arXiv:2202.02435.
- Kidger, P.; Chen, R. T. Q.; and Lyons, T. 2020. "Hey, that's not an ODE": Faster ODE Adjoint via Seminorms. In *International Conference on Machine Learning*.
- Kim, S.; Ji, W.; Deng, S.; Ma, Y.; and Rackauckas, C. 2021. Stiff neural ordinary differential equations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(9).
- Linot, A. J.; Burby, J. W.; Tang, Q.; Balaprakash, P.; Graham, M. D.; and Maulik, R. 2022. Stabilized Neural Ordinary Differential Equations for Long-Time Forecasting of Dynamical Systems. *Journal of Computational Physics*, 474.
- Massaroli, S.; Poli, M.; Park, J.; Yamashita, A.; and Asama, H. 2020. Dissecting Neural ODEs. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 3952–3963. Curran Associates, Inc.
- Matsubara, T.; Miyatake, Y.; and Yaguchi, T. 2023. The Symplectic Adjoint Method: Memory-Efficient Backpropagation of Neural-Network-Based Differential Equations. *IEEE Transactions on Neural Networks and Learning Systems*, 1–13.
- Mills, R. T.; Adams, M. F.; Balay, S.; Brown, J.; Dener, A.; Knepley, M.; Kruger, S. E.; Morgan, H.; Munson, T.; Rupp, K.; Smith, B. F.; Zampini, S.; Zhang, H.; and Zhang, J. 2021. Toward performance-portable PETSc for GPU-based exascale systems. *Parallel Computing*, 108: 102831.
- Onken, D.; and Ruthotto, L. 2020. Discretize-Optimize vs. Optimize-Discretize for Time-Series Regression and Continuous Normalizing Flows. arXiv preprint arXiv:2005.13420.

- Poli, M.; Massaroli, S.; Park, J.; Yamashita, A.; Asama, H.; and Park, J. 2019. Graph Neural Ordinary Differential Equations. *arXiv preprint arXiv:1911.07532*.
- Winston, E.; and Kolter, J. Z. 2020. Monotone operator equilibrium networks. In *Advances in Neural Information Processing Systems*, volume 2020-December.
- Xhonneux, L.-P.; Qu, M.; and Tang, J. 2020. Continuous Graph Neural Networks. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 10432–10441. PMLR.
- Zhang, H.; and Constantinescu, E. M. 2021. Revolve-based adjoint checkpointing for multistage time integration. In *International Conference on Computational Science*. (online), in main track.
- Zhang, H.; and Constantinescu, E. M. 2023. Optimal checkpointing for adjoint multistage time-stepping schemes. *Journal of Computational Science*, 66: 101913.
- Zhang, H.; and Sandu, A. 2014. FATODE: a library for forward, adjoint, and tangent linear integration of ODEs. *SIAM Journal on Scientific Computing*, 36(5): C504–C523.
- Zhang, H.; and Zhao, W. 2022. A Memory-Efficient Neural Ordinary Differential Equation Framework Based on High-Level Adjoint Differentiation. *IEEE Transactions on Artificial Intelligence*, 1–11.
- Zhang, T.; Yao, Z.; Gholami, A.; Gonzalez, J. E.; Keutzer, K.; Mahoney, M. W.; and Biros, G. 2019. ANODEV2: A Coupled Neural ODE Framework. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Zhuang, J.; Dvornik, N.; Li, X.; Tatikonda, S.; Papademetris, X.; and Duncan, J. 2020. Adaptive Checkpoint Adjoint Method for Gradient Estimation in Neural ODE. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, 11639–11649.
- Zhuang, J.; Dvornik, N. C.; Tatikond, S.; and Duncan, J. S. 2021. MALI: A memory efficient and reverse accurate integrator for Neural ODEs. In *International Conference on Learning Representations*.