
Hypersolvers: Toward Fast Continuous-Depth Models

Michael Poli*
KAIST, DiffEqML
poli_m@kaist.ac.kr

Stefano Massaroli*
The University of Tokyo, DiffEqML
massaroli@robot.t.u-tokyo.ac.jp

Atsushi Yamashita
The University of Tokyo
yamashita@robot.t.u-tokyo.ac.jp

Hajime Asama
The University of Tokyo
asama@robot.t.u-tokyo.ac.jp

Jinkyoo Park
KAIST
jinkyoo.park@kaist.ac.kr

Abstract

The infinite-depth paradigm pioneered by Neural ODEs has launched a renaissance in the search for novel dynamical system-inspired deep learning primitives; however, their utilization in problems of non-trivial size has often proved impossible due to poor computational scalability. This work paves the way for scalable Neural ODEs with *time-to-prediction* comparable to traditional discrete networks. We introduce hypersolvers, neural networks designed to solve ODEs with low overhead and theoretical guarantees on accuracy. The synergistic combination of hypersolvers and Neural ODEs allows for cheap inference and unlocks a new frontier for practical application of continuous-depth models. Experimental evaluations on standard benchmarks, such as sampling for *continuous normalizing flows*, reveal consistent pareto efficiency over classical numerical methods.

1 Introduction

The framework of *neural ordinary differential equations* (Neural ODEs) (Chen et al., 2018) has reinvigorated research in continuous deep learning (Zhang et al., 2014), offering new system-theoretic perspectives on neural network architecture design (Greydanus et al., 2019; Bai et al., 2019; Poli et al., 2019; Cranmer et al., 2020) and generative modeling (Grathwohl et al., 2018; Yang et al., 2019). Despite the successes, Neural ODEs have been met with skepticism, as these models are often slow in both training and inference due to heavy numerical solver overheads. These issues are further exacerbated by applications which require extremely accurate numerical solutions to the differential equations, such as physics-inspired neural networks (Raissi et al., 2019) and continuous normalizing flows (CNFs) (Chen et al., 2018). Common knowledge within the field is that these models appear too slow in their current form for meaningful large-scale or embedded applications. Several attempts have been made to either directly or indirectly address some of these limitations, such as redefining the forward pass as a root finding problem (Bai et al., 2019), introducing *ad hoc* regularization terms (Finlay et al., 2020; Massaroli et al., 2020a) and augmenting the state to reduce stiffness of the solutions (Dupont et al., 2019; Massaroli et al., 2020b).

Unfortunately, these approaches either give up on the Neural ODE formulation altogether, do not reduce computation overhead sufficiently or introduce additional memory requirements. Although there is no shortage of works utilizing Neural ODEs in forecasting or classification tasks (Yildiz et al., 2019;

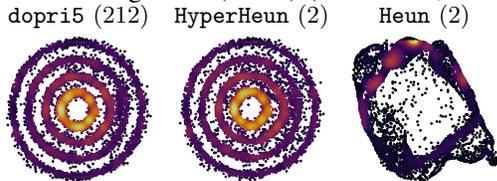


Figure 1: Hypersolvers for density estimation via *continuous normalizing flows*: dopri5 inference accuracy is achieved with 100x speedup.

*Equal contribution. Author order was decided by flipping a coin.

Jia & Benson, 2019; Kidger et al., 2020), current state-of-the-art is limited to offline applications with no constraints on inference time. In particular, high-potential application domains for Neural ODEs such as control and prediction often deal with tight requirements on inference speed and computation e.g robotics (Hester, 2013) that are not currently within reach. For example, a generic state-of-the-art convolutional Neural ODE takes at least an order of magnitude² longer to infer the label of a *single* MNIST image. This inefficiency results in inference passes far too slow for real-time applications.

Method	NFEs	Local Error
p -th order solver	$\mathcal{O}(pK)$	$\mathcal{O}(\epsilon^{p+1})$
adaptive-step solver	–	$\mathcal{O}(\tilde{\epsilon}^{p+1})$
Euler hypersolver	$\mathcal{O}(K)$	$\mathcal{O}(\delta\epsilon^2)$
p -th order hypersolver	$\mathcal{O}(pK)$	$\mathcal{O}(\delta\epsilon^{p+1})$

Figure 2: Asymptotic complexity comparison. Number of function evaluations (NFEs) needed to compute K solver’s steps. ϵ is the step size, $\tilde{\epsilon}$ is the max step size of adaptive solvers, $\delta \ll 1$ is correlated to the hypersolver training results.

hypersolvers; these auxiliary neural networks are trained to solve the initial value problem (IVP) emerging from the forward pass of continuous-depth models. Hypersolvers improve on the computation-correctness trade-off provided by traditional numerical solvers, enabling fast and arbitrarily accurate solutions during inference.

● **Pareto efficiency** The trade-off between solution accuracy and computation is one of the oldest and best-studied topics in the numerics literature (Butcher, 2016) and was mentioned in the seminal work (Chen et al., 2018) as a feature of continuous models. Traditional approaches shift additional compute resources into improved accuracy via higher-order adaptive-step methods (Prince & Dornand, 1981). For the most part, the computation-accuracy pareto front determined by traditional methods has been treated as optimal, allowing practitioners its traversal with different solver choices. We provide theoretical and practical results in support of the pareto efficiency of hypersolvers, measured with respect to both *number of function evaluations* (NFEs) as well as standard indicators of algorithmic complexity. Fig. 2 provides a comparison of hypersolvers and traditional methods.

● **Inference speed** By leveraging Hypersolved Neural ODEs, we obtain significant speedups on common benchmarks for continuous-depth models. In image classification tasks, inference is sped up by *at least* one order of magnitude. Additionally, the proposed approach is capable of solving *continuous normalizing flow* (CNF) (Chen et al., 2018; Grathwohl et al., 2018) sampling in few steps with little-to-no degradation of the sample quality as shown in Fig. 4.1. Moving beyond computational advantages at inference time, the proposed framework is compatible with *continual learning* (Parisi et al., 2019) or *adversarial learning* (Ganin et al., 2016) techniques where model and hypersolver are co-designed and jointly optimized. Sec. 6 provides an overview of this peculiar interplay.

2 Background: Continuous-Depth Models

We start by introducing necessary background on Neural ODE and numerical integration methods.

Neural ODEs We consider the following general Neural ODE formulation (Massaroli et al., 2020b)

$$\begin{cases} \dot{\mathbf{z}} = f_{\theta(s)}(s, \mathbf{x}, \mathbf{z}(s)) \\ \mathbf{z}(0) = h_x(\mathbf{x}) \\ \hat{\mathbf{y}}(s) = h_y(\mathbf{z}(s)) \end{cases} \quad s \in \mathcal{S} \quad (1)$$

with input $\mathbf{x} \in \mathbb{R}^{n_x}$, output $\hat{\mathbf{y}} \in \mathbb{R}^{n_y}$, *hidden* state $\mathbf{z} \in \mathbb{R}^{n_z}$ and \mathcal{S} is a compact subset of \mathbb{R} . Here $f_{\theta(s)}$ is a neural network, parametrized by $\theta(s)$ in some functional space. We equip the Neural ODE

²Compared with an equivalent-performance ResNet.

with input and output mappings $h_x : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_z}$, $h_y : \mathbb{R}^{n_z} \rightarrow \mathbb{R}^{n_y}$ which are kept linear as to avoid a collapse of the dynamics into a non-necessary map as discussed in (Massaroli et al., 2020b).

Solving the ODE Without any loss of generality, let $\mathcal{S} := [0, S]$ ($S \in \mathbb{R}^+$). The inference of Neural ODEs is carried out by solving the *initial value problem* (IVP) (1), i.e.

$$\hat{\mathbf{y}}(S) = h_y \left(h_x(\mathbf{x}) + \int_{\mathcal{S}} f_{\theta(\tau)}(\tau, \mathbf{x}, \mathbf{z}(\tau)) d\tau \right)$$

Due to the nonlinearities of $f_{\theta(s)}$, this solution cannot be defined in closed-form and, thus, a numerical solution should be obtained by iterating some predetermined ODE solver. Let us divide \mathcal{S} in K equally-spaced intervals $[s_k, s_{k+1}]$ such that for all $k \in \mathbb{N}_{<K}$ $s_{k+1} - s_k = S/K := \epsilon \in \mathbb{R}^+$. The numerical approximation of the IVP solution in \mathcal{S} can be computed by iterating

$$\begin{cases} \mathbf{z}_{k+1} = \mathbf{z}_k + \epsilon \psi(s_k, \mathbf{x}, \mathbf{z}_k) \\ \mathbf{z}_0 = h_x(\mathbf{x}) \\ \hat{\mathbf{y}}_k = h_y(\mathbf{z}_k) \end{cases} \quad k = 0, 1, \dots, K-1 \quad (2)$$

where ψ is a function performing the state update.

Numerical methods ODE solvers differ in how this map ψ is constructed³. In example, the Euler method is realized by setting $\psi(\mathbf{x}, s_k, \mathbf{z}_k) := f_{\theta(s_k)}(\mathbf{x}, s_k, \mathbf{z}_k)$. Note that, *higher-order* solvers compute $\psi(\mathbf{x}, s_k, \mathbf{z}_k)$ iteratively in p steps where p denotes the order of the solver. For example, in a p -th order Runge-Kutta (RK) (Runge, 1895) method ψ is computed as

$$\begin{aligned} \mathbf{r}_i &= f_{\theta(s_k)}(s_k + \mathbf{c}_i \epsilon, \mathbf{x}, \mathbf{z}_k + \tilde{\mathbf{z}}_k^i) \quad i = 1, \dots, p \\ \tilde{\mathbf{z}}_k^i &= \epsilon \sum_{j=1}^p \mathbf{a}_{ij} \mathbf{r}_j \quad i = 1, \dots, p \\ \psi &= \sum_{j=1}^p \mathbf{b}_j \mathbf{r}_j \end{aligned} \quad (3)$$

where $\mathbf{a} \in \mathbb{R}^{p \times p}$, $\mathbf{b} \in \mathbb{R}^p$, $\mathbf{c} \in \mathbb{R}^p$ fully characterize the method. Hence, the integration of a neural ODE in \mathcal{S} with a RK solver is $\mathcal{O}(pK)$ in memory efficiency and time complexity. On the other hand, *adaptive-step* solvers, e.g. the popular Dormand-Prince 5(4) (dopri5) have no explicit upper bounds in memory and time efficiency. This is especially critical as in many practical applications, a requirement for maximum memory consumption and/or inference time must be satisfied.

Common metrics In classic numerical analysis, two type of metrics are often defined, i.e. the *local truncation error* e_k

$$e_k := \|\mathbf{z}(s_{k+1}) - \mathbf{z}(s_k) - \epsilon \psi(s_k, \mathbf{x}, \mathbf{z}(s_k))\|_2,$$

representing the error accumulated in a single step, and the *global truncation error* \mathcal{E}_k is

$$\mathcal{E}_k = \|\mathbf{z}(s_k) - \mathbf{z}_k\|_2,$$

i.e. the error accumulated in the first k steps. Note that for a p -th order solver $e_k = \mathcal{O}(\epsilon^{p+1})$ and $\mathcal{E}_k = \mathcal{O}(\epsilon^p)$ (Butcher, 2016).

3 • Hypersolvers for Neural ODEs

Hypersolvers offers a computational framework for the interplay between Neural ODEs and their numerical solver. The core idea behind hypersolvers is to introduce an additional neural network g_ω to approximate the higher-order terms of a given solver, greatly increasing its accuracy while preserving the computational and memory efficiency. The simplest instance of Hypersolved Neural ODEs is based on Euler scheme:

$$\begin{cases} \mathbf{z}_{k+1} = \mathbf{z}_k + \epsilon f_{\theta(s_k)}(s_k, \mathbf{x}, \mathbf{z}_k) + \epsilon^2 g_\omega(\epsilon, s_k, \mathbf{x}, \mathbf{z}_k) \\ \mathbf{z}_0 = h_x(\mathbf{x}) \\ \hat{\mathbf{y}}_k = h_y(\mathbf{z}_k) \end{cases} \quad k = 0, 1, \dots, K-1 \quad (4)$$

³Numerical solvers which obey to (2) are called *explicit* solvers

where g_ω is a neural network approximating the second-order term of the Euler method. The derivation of the *Euler hypersolver* comes naturally from the following. Let $\mathbf{z}(s_k)$ be the true solution of (1) at $s_k \in \mathcal{S}$ and let $\epsilon > 0$ such that $s_k + \epsilon \in \mathcal{S}$. From the Taylor expansion of the solution around s_k , i.e.

$$\begin{aligned} \mathbf{z}(s_k + \epsilon) &= \mathbf{z}(s_k) + \epsilon \dot{\mathbf{z}}(s_k) + \frac{1}{2} \epsilon^2 \ddot{\mathbf{z}}(s_k) + \mathcal{O}(\epsilon^3) \\ &\approx \mathbf{z}(s_k) + \epsilon f_\theta(s_k, \mathbf{x}, \mathbf{z}(s_k)) \end{aligned}$$

we deduce that the classic Euler scheme corresponds to the first-order truncation of the above. The Euler hypersolver, instead, aims at approximating the second-order term, reducing the local truncation error of the overall scheme, while avoiding to compute and store further evaluations of $f_{\theta(s)}$, as required by higher order schemes, e.g. RK methods.

3.1 General formulation

A general formulation of Hypersolved Neural ODEs can be obtained extending (2). If we assume ψ to be the update step of a p -th order solver, then the general p -th order Hypersolved Neural ODE is defined as

$$\begin{aligned} \mathbf{z}_{k+1} &= \mathbf{z}_k + \overbrace{\epsilon \psi(s_k, \mathbf{x}, \mathbf{z}_k)}^{\text{solver step}} + \epsilon^{p+1} \overbrace{g_\omega(\epsilon, s_k, \mathbf{x}, \mathbf{z}_k)}^{\text{hypersolver net}} \quad k = 0, 1, \dots, K-1 \quad (5) \\ \mathbf{z}_0 &= h_x(\mathbf{x}) \\ \hat{\mathbf{y}}_k &= h_y(\mathbf{z}_k) \end{aligned}$$

Software implementation We implemented hypersolver variants of common low-order explicit ODE solvers, designed for compatibility with the TorchDyn (Poli et al., 2020) library⁴. The Appendix further includes a PyTorch (Paszke et al., 2017) module implementation.

3.2 Training hypersolvers

Assume to have available the *exact* solution of the Neural ODE evaluated at the mesh points s_k , practically obtained through an adaptive-step solver set up with low tolerances. With these solution checkpoints we construct the training set for the DE solver with tuples:

$$\{(s_k, \mathbf{z}(s_k))\}_{k \in \mathbb{N}_{\leq K}}$$

According to the introduced metrics e_k and \mathcal{E}_k , we introduce two types of loss functions aimed at improving each of the metrics.

Residual fitting We first start by defining the *residual* of the solver (2)

$$\mathcal{R}(s_k, \mathbf{z}(s_k), \mathbf{z}(s_{k+1})) = \frac{1}{\epsilon^{p+1}} [\mathbf{z}(s_{k+1}) - \mathbf{z}(s_k) - \epsilon \psi(s_k, \mathbf{x}, \mathbf{z}(s_k))] \quad (6)$$

which correspond to a scaled local truncation error without the neural correction term g_ω . Then, we can consider a loss measuring the discrepancy between the residual terms and the output of g_ω :

$$\ell = \frac{1}{K} \sum_{k=0}^{K-1} \|\mathcal{R}(s_k, \mathbf{z}(s_k), \mathbf{z}(s_{k+1})) - g_\theta(\epsilon, s_k, \mathbf{x}, \mathbf{z}(s_k))\|_2$$

If the hypersolver is trained to minimize ℓ_{local} , the following holds:

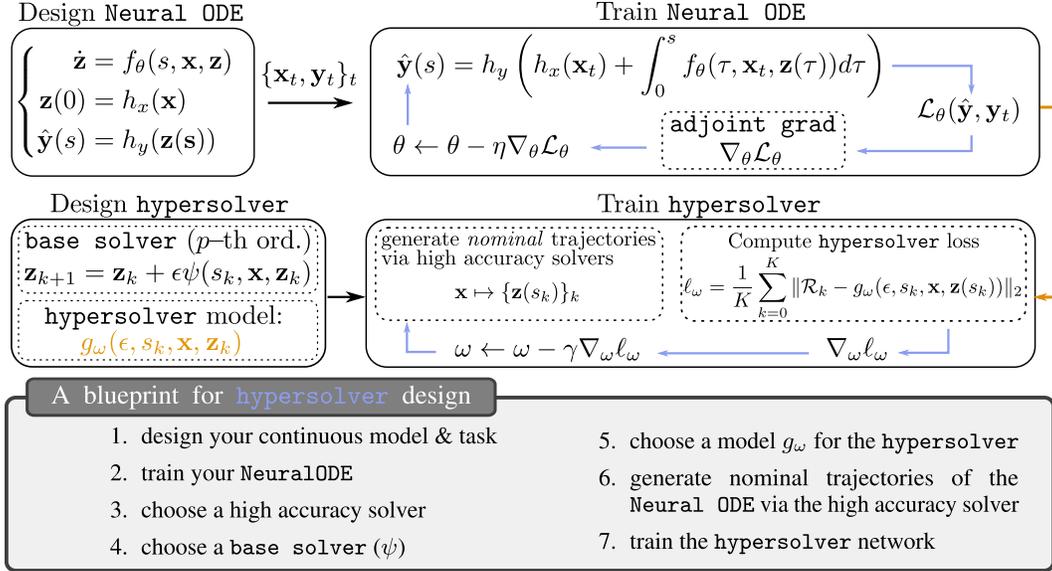
Theorem 1 (Hypersolver Local Truncation Error). *If g_ω is a $\mathcal{O}(\delta)$ approximator of \mathcal{R} , i.e.*

$$\forall k \in \mathbb{N}_{\leq K} \quad \|\mathcal{R}(s_k, \mathbf{z}(s_k), \mathbf{z}(s_{k+1})) - g_\theta(\epsilon, s_k, \mathbf{x}, \mathbf{z}(s_k))\|_2 \leq \mathcal{O}(\delta),$$

then, the local truncation error e_k of the hypersolver is $\mathcal{O}(\delta \epsilon^{p+1})$.

The proof and further theoretical insights are reported in the Appendix.

⁴Supporting reproducibility code is at <https://github.com/DiffEqML/diffeqml-research/tree/master/hypersolver>



Trajectory fitting The second type of hypersolvers training aims at containing the global truncation error by minimizing the difference between the exact and approximated solutions in the whole depth domain \mathcal{S} , i.e.

$$L = \sum_{k=1}^K \|\mathbf{z}(s_k) - \mathbf{z}_k\|_2$$

It should be noted that *trajectory* and *residual fitting* can be combined into a single loss term, depending on the application.

4 Experimental Evaluation

The evaluation protocol is designed to measure hypersolver pareto efficiency, inference time speedups and generalizability across base solvers. We consider the following general benchmarks for Neural ODEs: standard image classification (Dupont et al., 2019; Massaroli et al., 2020b) and density estimation with continuous normalizing flows (CNFs) (Chen et al., 2018; Grathwohl et al., 2018).

4.1 Image Classification

We train standard convolutional Neural ODEs with input-layer augmentation (Massaroli et al., 2020b) on MNIST and CIFAR10 datasets. Following this initial optimization step, 2-layer convolutional *Euler hypersolvers*, HyperEuler, (4) are trained by residual fitting (6) on 10 epochs of the *training* dataset with solution mesh length set to $K = 10$. As ground-truth labels, we utilize the solutions obtained via `dopri5` with absolute and relative tolerances set to 10^{-4} on the same data. The objective of this first task is to show that hypersolvers retain their pareto efficiency when applied in high-dimensional data regimes. Additional details on hyperparameter choice and architectures are provided as supplementary material.

● **Pareto comparison** We analyze pareto efficiency of hypersolvers with respect to both ODE ODE solution accuracy and test task classification accuracy. It should be noted that residual fitting does not require task supervision; indeed, test data could be used for hypersolver training. Nonetheless, we decide to use only training data for residual fitting, in order to confirm hypersolver ability to generalize to unseen initial conditions of the Neural ODE.

Multiply-accumulate operations i.e MACs are used as a general algorithmic complexity measure. We opt for MACs instead of number of function evaluations (NFEs) of the Neural ODE vector field f_θ since the latter does not take into account computational overheads due to hypersolver network g_ω . It should be noted that for these specific architectures, single evaluations of f_θ and g_ω correspond to

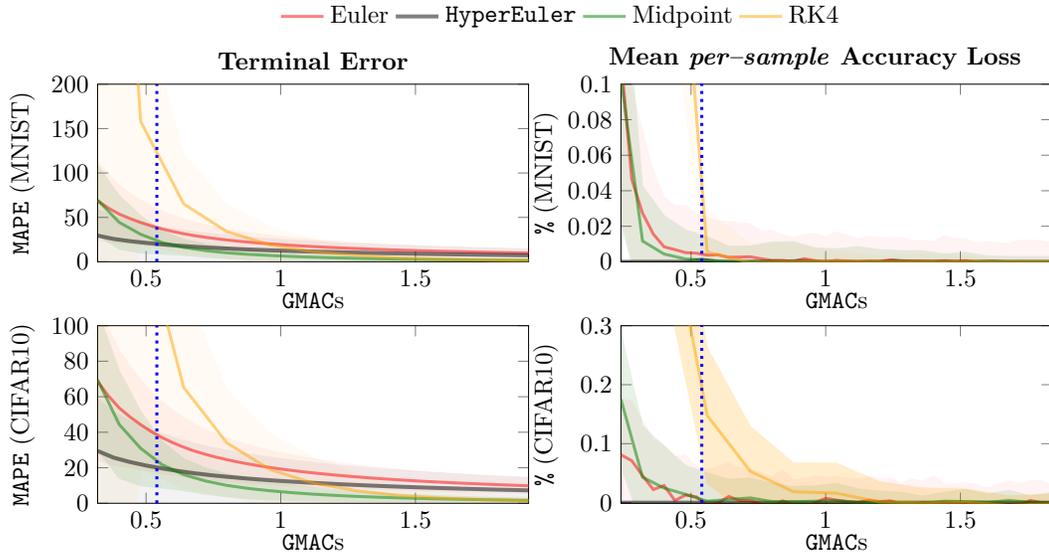


Figure 3: Test accuracy loss %-NFE and MAPE-GMAC Pareto fronts of different ODE solvers on MNIST and CIFAR10 test sets. HyperEuler shows higher Pareto efficiency for low function evaluations (NFEs) even over higher-order methods.

0.04 GMACs and 0.02 GMACs, respectively. HyperEuler is able to generalize to different step sizes not seen during training, which involved a 10 steps over an integration interval of 1s. Such residual training scheme over 9 residuals corresponds to a computational complexity for HyperEuler of 0.54 GMACs, highlighted in blue in Fig. 3. As shown in the Figure, HyperEuler enjoys Pareto optimality over alternative fixed-step methods. The hypersolver is able to generalize to different step sizes not seen during training, outperforming higher-order methods such as midpoint and RK4 at low NFEs. As expected, even though higher-order methods eventually surpass HyperEuler at higher NFEs as predicted by theoretical bounds, the hypersolver retains its Pareto optimality over Euler.

● **Wall-clock speedups** We measure wall-clock solution time speedups of various fixed-step methods over `dopri5` for image classification Neural ODEs. Here, absolute time refers to average time across batches of the MNIST test set required to solve the Neural ODE with different numerical schemes.

Each method performs the minimum number of steps to preserve total accuracy loss across the test set to less than 0.1%. As shown in Fig. 4, HyperEuler solves an MNIST Neural ODE roughly 8 times faster than `dopri5` and with comparable accuracy, achieving significant speedups even over its base method Euler. Indeed, Euler requires a larger number of steps due to its Pareto inefficiency compared to HyperEuler, leading to a slower overall solve. The measurements presented are collected on a single V100 GPU.

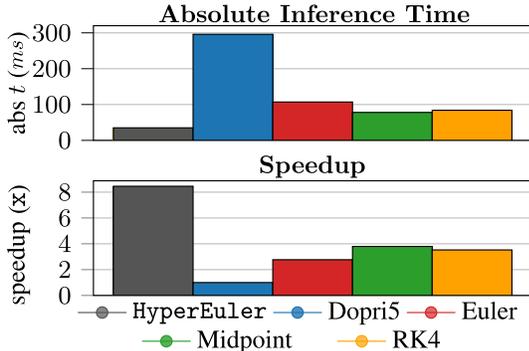


Figure 4: Absolute time (ms) speedup of fixed-step methods over `dopri5` (MNIST test set). HyperEuler solves the Neural ODE 8x faster than `dopri5` with the same accuracy.

General Butcher Tableau		Butcher Tableau of 2nd-order α family	
c	A	0	
	b	α	α
		$1 - \frac{1}{2\alpha}$	$\frac{1}{2\alpha}$

Figure 5: *Butcher Tableau* collecting coefficients of numerical methods see e.g (3). [left] general case. [Right] tableau of second-order α family. Note that $\alpha = 0.5$ recovers the *midpoint* method.

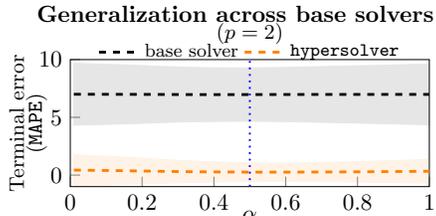


Figure 6: Neural ODE MAPE terminal solution error of HyperMidpoint and various members of the α -family.

• **Generalization across base solvers** We verify hypersolver capability to generalize across different base solvers of the same order. We consider the general family of second-order explicit methods parametrized by α (Süli & Mayers, 2003) as shown in Fig. 5. Employing a parametrizing family for second-order methods instead of specific instances such as midpoint or Heun allows for an analysis of gradual generalization performance as α is tuned away from its value corresponding to the chosen base solver. In particular we consider as midpoint, recovered by $\alpha = 0.5$, as the base solver for the corresponding hypersolver.

Fig. 6 shows average terminal MAPE solution error of MNIST Neural ODEs solved with both various α methods as well as a single HyperMidpoint. As with the previous experiments, the error is computed over dopri5 solutions, and averaged across test data batches. HyperMidpoint is then evaluated, without finetuning, by swapping its base solver with other members of the α family. The hypersolver generalizes to different base solvers, preserving its pareto efficiency over the entire α -family.

4.2 Lightweight Density Estimation

We consider sampling in the FFJORD (Grathwohl et al., 2018) variant of *continuous normalizing flows* (Chen et al., 2018) as an additional task to showcase hypersolver performance. We train CNFs closely following the setup of Grathwohl et al. (2018). Then, we optimize two-layer, second-order Heun hypersolvers, HyperHeun, with $K = 1$ residuals obtained against dopri5 with absolute tolerance 10^{-5} and relative tolerance 10^{-5} . The striking result highlighted in Fig. 7 is that with as little as two NFEs, Hypersolved CNFs provide samples that are as accurate as those obtained through the much more computationally expensive dopri5.

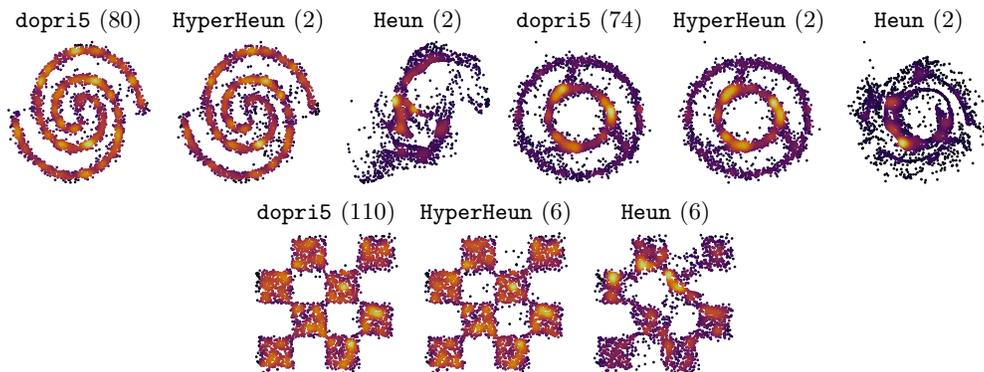


Figure 7: Reconstructed densities with continuous normalizing flows and Heun hypersolver HyperHeun. The inference accuracy of dopri5 is reached through the hypersolver with a significant speedup in terms of computation time and accuracy. Heun method fails to solve correctly the ODE with same NFEs of HyperHeun.

5 Related Work

Neural network solvers There is a long line of research leveraging the universal approximation capabilities of neural networks for solving differential equations. A recurrent theme of the existing work (Lagaris et al., 1997, 1998; Li-ying et al., 2007; Li & Li, 2013; Mall & Chakraverty, 2013; Raissi et al., 2018; Qin et al., 2019) is direct utilization of noiseless analytical solutions and evaluations in

low dimensional settings. Application specific attempts (Xing & McCue, 2010; Breen et al., 2019; Fang et al., 2020) provide empirical evidence in support of the earlier work, though the approximation task is still cast as a gradient–matching regression problem on noiseless labels. Deep neural network base solvers have also been used in the distributed parameters setting for PDEs (Han et al., 2018; Magill et al., 2018; Weinan & Yu, 2018; Raissi, 2018; Piscopo et al., 2019; Both et al., 2019; Khoo & Ying, 2019; Winovich et al., 2019; Raissi et al., 2019). Techniques to use neural networks for fast simulation of physical systems have been explored in (Grzeszczuk et al., 1998; James & Fatahalian, 2003; Sanchez-Gonzalez et al., 2020). More recent advances involving symbolic regressions include (Winovich et al., 2019; Regazzoni et al., 2019; Long et al., 2019).

The hypersolver approach is different in several key aspects. To the best knowledge of the authors, this represents the first example where neural network solvers show both *consistent* and *significant* pareto efficiency improvements over traditional solvers in high–dimensional settings. The performance advantages are demonstrated in the absence of analytic solutions and are supported by theoretical guarantees, ultimately yielding large inference speedups of practical relevance for Neural ODEs.

Multi–stage residual networks After seminal research (Sonoda & Murata, 2017; Lu et al., 2017; Chang et al., 2017; Hauser & Ray, 2017; Chen et al., 2018) uncovered and strengthened the connection between ResNets and ODE discretizations, a variety of architecture and objective specific adjustments have been made to the vanilla formulation. The above allow, for example, to accommodate irregular observations in sequence data (Demeester, 2019) or inherit beneficial properties from the corresponding numerical methods (Zhu et al., 2018). Although these approaches share some structural similarities with the Hypersolved formulation (4), the objective is drastically different. Indeed, such models are optimized for task–specific metrics without concern about preserving ODE properties, or developing a synergistic connection between model and solver.

6 Discussion

Hypersolvers can be leveraged beyond the inference step of continuous–depth models. Here, we provide avenues of further development of the framework.

Hypersolver overhead The source of the computational (and memory) overheads caused by the use of hypersolver is indeed represented by the evaluation of g_ω at each solver step. Nonetheless, this overhead (e.g. in terms of multiply–accumulate operations, MACs) decreases as the solver order increases. In fact, in a p th order solver where f_θ should be evaluated p times, g_ω is evaluated only once. Let MAC_f , MAC_g be indicators of algorithmic complexity of f_θ and g_ω , respectively. We have that the *relative overhead* (in terms of MACs) O_r is

$$O_r = \frac{p\text{MAC}_f + \text{MAC}_g}{p\text{MAC}_f} = 1 + \frac{1}{p} \frac{\text{MAC}_g}{\text{MAC}_f}$$

and $O_r \rightarrow 1$ for $p \rightarrow \infty$. Thus, the experiments on pareto efficiency and wall–clock speedup using HyperEuler showcased in Sec. 4.1 should be regarded as worst–case scenario, i.e. the most expensive computational–wise.

Even in the worst-case scenario, hypersolvers remain pareto efficient over traditional methods

Beyond fixed–step explicit hypersolvers In this work, we focus on developing hypersolvers as enhancements to fixed–step explicit methods for Neural ODEs. Although this approach is already effective during inference, hypersolvers are not constrained to this setting. Indeed, the proposed framework can be used to systematically blend learning models and numerical solvers beyond the fixed–step, explicit case. In principle, we could employ hypersolvers into *predictor–corrector* scheme where we may learn higher–order terms of either the (explicit) predictor or the (implicit) corrector, effectively reducing the overall truncation error. Similarly, adaptive stepping might be achieved by augmenting, in example, the *Dormand–Prince* (dopri5) scheme. dopri5 uses six NFEs to calculate fourth- and fifth-order Runge–Kutta solutions and obtain the error estimate for step adaptation. Here, we could substitute RK5 with an HyperRK4 and/or train a NN to perform the adaptation given the error estimate.

Hypersolver are not limited to fixed-step explicit base solvers.

Accelerating Neural ODE training Speeding up continuous-depth model training with hypersolvers involves additional challenges. In particular, it is necessary to ensure that the hypersolver network remains a $O(\delta)$ approximator of residuals \mathcal{R} across training iterations. A theoretical toolkit to tackle such a task may be offered by *continual learning* (Parisi et al., 2019).

Consider the problem of approximating the solution of a Neural ODE at training iteration $t + 1$ having optimized the hypersolver on flows generated by the model $f_{\theta_t(s)}(\mathbf{x}_t, s, \mathbf{z}(s))$ at the previous training step t . This setting involves a certifiably smooth transition between tasks that is directly controlled by the learning rate η , leading to the following result

Proposition 1 (Vector field training sensitivity). *Let the model parameters θ_t be updated according to the gradient-based optimizer step $\theta_{t+1} = \theta_t + \eta\Gamma(\nabla_{\theta}\mathcal{L}_t)$, $\eta > 0$ to minimize a loss function \mathcal{L}_t and let f_{θ_t} be Lipschitz w.r.t. θ . Then,*

$$\forall \mathbf{z} \in \mathbb{R}^{n_z}, \quad \|\Delta f_{\theta_t}(s, \mathbf{x}, \mathbf{z})\|_2 \leq \eta L_{\theta} \|\Gamma(\nabla_{\theta}\mathcal{L})\|_2$$

being L_{θ} the Lipschitz constant.

By leveraging the above result, or pretraining the hypersolver on a sufficiently large collection of dynamics, it might be possible to construct a training procedure for Neural ODEs which maximizes hypersolver reuse across training iterations. Similar to other application areas such as language processing (Howard & Ruder, 2018; Devlin et al., 2018), we envision pretraining techniques to play a fundamental part in the search for easy-to-train continuous-depth models.

Maximizing hypersolver reuse represents an important objective for faster Neural ODE training.

Model-solver joint optimization Hypersolver and Neural ODE training can be carried out *jointly* during optimization for the main task. Beyond numerical accuracy metrics, other task specific losses can be considered for hypersolvers. In the standard setting, numerical solvers act as adversaries preserving the ODE solution accuracy at the cost of expressivity. Taking this analogy further, we propose adversarial optimization in the form $\min_{\omega} \max_{\theta} \sum_{k=0}^K \|\mathbf{z}_k - \bar{\mathbf{z}}_k\|_2$ where $\bar{\mathbf{z}}_k$ is the solution at mesh point k given by an adaptive step solver. When used either during hypersolver pretraining or as a regularization term for the main task, the above gives rise to emerging behaviors in the dynamics $f_{\theta(s)}$ which exploit solver weaknesses. We observe, as briefly discussed in the Appendix, that direct adversarial training teaches $f_{\theta(s)}$ to leverage *stiffness* (Shampine, 2018) of the differential equation to increase the hypersolver solution error.

Adversarial training may be used to enhance hypersolver resilience to challenging dynamics.

7 Conclusion

Computational overheads represent a great obstacle for the utilization of continuous-depth models in large scale or real-time applications. This work develops the novel hypersolver framework, designed to alleviate performance limitations by leveraging the key model-solver interplay of continuous-depth architectures. Hypersolvers, neural networks trained to solve Neural ODEs accurately and with low overhead, improve solution accuracy at a negligible computational cost, ultimately improving pareto efficiency of traditional methods. Indeed, the synergistic combinations of Hypersolvers and Neural ODEs enjoy large speedups during inference steps of standard benchmarks of continuous-depth models, allowing in example accurate sampling from *continuous normalizing flows* (CNFs) in as little as 2 *number of function evaluations* (NFEs). Finally, we discuss how the hypersolver paradigm can be extended to enhance Neural ODE training through continual learning, pretraining or joint optimization of model and hypersolver.

Broader Impact

Major application areas for continuous deep learning architectures so far have been generative modeling (Grathwohl et al., 2018) and forecasting, particularly in the context of patient medical data (Jia & Benson, 2019). While these models have an intrinsic interpretability advantages over discrete counterparts, it is important that future iterations preserve these properties in the search for greater scalability. Early adoption of the hypersolver paradigm would speed up widespread utilization of Neural ODEs in these domains, ultimately leading to positive impact in healthcare applications.

Accurate forecasting is at the foundation of system identification and control, two additional application areas set to be greatly impacted by continuous models. Unfortunately, theoretical guarantees of robustness in the worst-case scenario are challenging to construct for data-driven approaches. As these approaches are refined, they are also likely to negatively impact the employment market by accelerating job automation in critical areas.

Acknowledgment

We thank Patrick Kidger for helpful discussions. This work was supported by the Korea Agency for Infrastructure Technology Advancement (KAIA) grant, funded by the Ministry of Land, Infrastructure and Transport under Grant 19PIYR-B153277-01.

References

- Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, pp. 688–699, 2019.
- Gert-Jan Both, Subham Choudhury, Pierre Sens, and Remy Kusters. Deepmod: Deep learning for model discovery in noisy data. *arXiv preprint arXiv:1904.09406*, 2019.
- Philip G Breen, Christopher N Foley, Tjarda Boekholt, and Simon Portegies Zwart. Newton vs the machine: solving the chaotic three-body problem using deep neural networks. *arXiv preprint arXiv:1910.07291*, 2019.
- John Charles Butcher. *Numerical methods for ordinary differential equations*. John Wiley & Sons, 2016.
- JR Cash. Efficient numerical methods for the solution of stiff initial-value problems and differential algebraic equations. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 459(2032):797–815, 2003.
- 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.
- Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in neural information processing systems*, pp. 6571–6583, 2018.
- Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. *arXiv preprint arXiv:2003.04630*, 2020.
- Thomas Demeester. System identification with time-aware neural sequence models. *arXiv preprint arXiv:1911.09431*, 2019.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- John R Dormand and Peter J Prince. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.
- Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented neural odes. In *Advances in Neural Information Processing Systems*, pp. 3134–3144, 2019.
- Jie Fang, Chenglian Liu, TE Simos, and I Th Famelis. Neural network solution of single-delay differential equations. *Mediterranean Journal of Mathematics*, 17(1):30, 2020.
- Chris Finlay, Jörn-Henrik Jacobsen, Levon Nurbekyan, and Adam M Oberman. How to train your neural ode. *arXiv preprint arXiv:2002.02798*, 2020.

- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, pp. 15353–15363, 2019.
- Radek Grzeszczuk, Demetri Terzopoulos, and Geoffrey Hinton. Neuroanimator: Fast neural network emulation and control of physics-based models. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pp. 9–20, 1998.
- Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018. ISSN 0027-8424.
- Michael Hauser and Asok Ray. Principles of riemannian geometry in neural networks. In *Advances in neural information processing systems*, pp. 2807–2816, 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- Todd Hester. *TEXPLORE: Temporal Difference Reinforcement Learning for Robots and Time-Constrained Domains*. Springer, 2013.
- Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*, 2018.
- Doug L James and Kayvon Fatahalian. Precomputing interactive dynamic deformable scenes. *ACM Transactions on Graphics (TOG)*, 22(3):879–887, 2003.
- Junteng Jia and Austin R Benson. Neural jump stochastic differential equations. In *Advances in Neural Information Processing Systems*, pp. 9843–9854, 2019.
- Yuehaw Khoo and Lexing Ying. Switchnet: a neural network model for forward and inverse scattering problems. *SIAM Journal on Scientific Computing*, 41(5):A3182–A3201, 2019.
- Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *arXiv preprint arXiv:2005.08926*, 2020.
- Ivan Kobyzev, Simon Prince, and Marcus A Brubaker. Normalizing flows: Introduction and ideas. *arXiv preprint arXiv:1908.09257*, 2019.
- Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural network methods in quantum mechanics. *Computer Physics Communications*, 104(1-3):1–14, 1997.
- Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5):987–1000, 1998.
- Shuai Li and Yangming Li. Nonlinearly activated neural network for solving time-varying complex sylvester equation. *IEEE transactions on cybernetics*, 44(8):1397–1407, 2013.
- Xu Li-ying, Wen Hui, and Zeng Zhe-zhao. The algorithm of neural networks on the initial value problems in ordinary differential equations. In *2007 2nd IEEE Conference on Industrial Electronics and Applications*, pp. 813–816. IEEE, 2007.
- 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.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Yiping Lu, Aoxiao Zhong, Quanzheng Li, and Bin Dong. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. *arXiv preprint arXiv:1710.10121*, 2017.

- Martin Magill, Faisal Qureshi, and Hendrick de Haan. Neural networks trained to solve differential equations learn general representations. In *Advances in Neural Information Processing Systems*, pp. 4071–4081, 2018.
- Susmita Mall and Snehashish Chakraverty. Comparison of artificial neural network architecture in solving ordinary differential equations. *Advances in Artificial Neural Systems*, 2013, 2013.
- Stefano Massaroli, Michael Poli, Michelangelo Bin, Jinkyoo Park, Atsushi Yamashita, and Hajime Asama. Stable neural flows. *arXiv preprint arXiv:2003.08063*, 2020a.
- Stefano Massaroli, Michael Poli, Jinkyoo Park, Atsushi Yamashita, and Hajime Asama. Dissecting neural odes. *arXiv preprint arXiv:2002.08071*, 2020b.
- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Maria Laura Piscopo, Michael Spannowsky, and Philip Waite. Solving differential equations with neural networks: applied to the calculation of cosmological phase transitions. *arXiv preprint arXiv:1902.05563*, 2019.
- Michael Poli, Stefano Massaroli, Junyoung Park, Atsushi Yamashita, Hajime Asama, and Jinkyoo Park. Graph neural ordinary differential equations. *arXiv preprint arXiv:1911.07532*, 2019.
- Michael Poli, Stefano Massaroli, Atsushi Yamashita, Hajime Asama, and Jinkyoo Park. Torchdyn: A neural differential equations library. *arXiv preprint arXiv:2009.09346*, 2020.
- Peter J Prince and John R Dormand. High order embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics*, 7(1):67–75, 1981.
- Tong Qin, Kailiang Wu, and Dongbin Xiu. Data driven governing equations approximation using deep neural networks. *Journal of Computational Physics*, 395:620–635, 2019.
- Maziar Raissi. Deep hidden physics models: Deep learning of nonlinear partial differential equations. *The Journal of Machine Learning Research*, 19(1):932–955, 2018.
- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Multistep neural networks for data-driven discovery of nonlinear dynamical systems. *arXiv preprint arXiv:1801.01236*, 2018.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- Francesco Regazzoni, Luca Dede, and Alfio Quarteroni. Machine learning for fast and reliable solution of time-dependent differential equations. *Journal of Computational Physics*, 397:108852, 2019.
- Carl Runge. Über die numerische auflösung von differentialgleichungen. *Mathematische Annalen*, 46(2):167–178, 1895.
- Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter W Battaglia. Learning to simulate complex physics with graph networks. *arXiv preprint arXiv:2002.09405*, 2020.
- Lawrence F Shampine. *Numerical solution of ordinary differential equations*. Routledge, 2018.
- Lawrence F Shampine and Charles William Gear. A user’s view of solving stiff ordinary differential equations. *SIAM review*, 21(1):1–17, 1979.
- Sho Sonoda and Noboru Murata. Double continuum limit of deep neural networks. In *ICML Workshop Principled Approaches to Deep Learning*, 2017.
- Endre Süli and David F Mayers. *An introduction to numerical analysis*. Cambridge university press, 2003.
- E Weinan and Bing Yu. The deep ritz method: a deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
- Nick Winovich, Karthik Ramani, and Guang Lin. Convpde-ug: Convolutional neural networks with quantified uncertainty for heterogeneous elliptic partial differential equations on varied domains. *Journal of Computational Physics*, 394:263–279, 2019.

- Zhiliang Xing and Leigh McCue. Modeling ship equations of roll motion using neural networks. *Naval Engineers Journal*, 122(3):49–60, 2010.
- 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.
- Çağatay Yıldız, Markus Heinonen, and Harri Lähdesmäki. ODE2VAE: Deep generative second order odes with bayesian neural networks. *arXiv preprint arXiv:1905.10994*, 2019.
- Huaguang Zhang, Zhanshan Wang, and Derong Liu. A comprehensive review of stability analysis of continuous-time recurrent neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 25(7):1229–1262, 2014.
- Mai Zhu, Bo Chang, and Chong Fu. Convolutional neural networks combined with runge-kutta methods. *arXiv preprint arXiv:1802.08831*, 2018.