

---

# Approximation Capabilities of Neural ODEs and Invertible Residual Networks

---

Han Zhang<sup>1</sup> Xi Gao<sup>1</sup> Jacob Unterman<sup>1</sup> Tom Arodz<sup>1</sup>

## Abstract

Neural ODEs and i-ResNet are recently proposed methods for enforcing invertibility of residual neural models. Having a generic technique for constructing invertible models can open new avenues for advances in learning systems, but so far the question of whether Neural ODEs and i-ResNets can model any continuous invertible function remained unresolved. Here, we show that both of these models are limited in their approximation capabilities. We then prove that any homeomorphism on a  $p$ -dimensional Euclidean space can be approximated by a Neural ODE operating on a  $2p$ -dimensional Euclidean space, and a similar result for i-ResNets. We conclude by showing that capping a Neural ODE or an i-ResNet with a single linear layer is sufficient to turn the model into a universal approximator for non-invertible continuous functions.

## 1. Introduction

A neural network block is a function  $F$  that maps an input vector  $x \in \mathcal{X} \subset \mathbb{R}^p$  to output vector  $F(x, \theta) \in \mathbb{R}^{p'}$ , and is parameterized by a weight vector  $\theta$ . We require that  $F$  is almost everywhere differentiable with respect to both of its arguments, allowing the use of gradient methods for tuning  $\theta$  based on training data and an optimization criterion, and for passing the gradient to preceding network blocks.

One type of neural building blocks that has received attention in recent years is a *residual block* (He et al., 2016), where  $F(x, \theta) = x + f(x, \theta)$ , with  $f$  being some differentiable, nonlinear, possibly multi-layer transformation. Input and output dimensionality of a residual block are the same,  $p$ , and such blocks are usually stacked in a sequence, a ResNet,  $x_{t+1} = x_t + f_t(x_t, \theta_t)$ . Often, the functional form of  $f_t$  is the same for all blocks  $t \in \{1, \dots, T\}$  in the

---

<sup>1</sup>Department of Computer Science, Virginia Commonwealth University, Richmond, Virginia, USA. Correspondence to: Tom Arodz <tarodz@vcu.edu>.

sequence. Then, we can represent the sequence through  $x_{t+1} - x_t = f_{\Theta}(x_t, t)$ , where  $\Theta$  consists of trainable parameters for all blocks in the sequence; the second argument,  $t$ , allows us to pick the proper subset of parameters,  $\theta_t$ . If we allow arbitrary  $f_{\Theta}$ , for example a neural network with input and output dimensionality  $p$  but with many hidden layers of dimensionality higher than  $p$ , a sequence of residual blocks can, in principle, model arbitrary mappings  $x \rightarrow \phi_T(x)$ , where we define  $\phi_T(x_0) = x_T$  to be the result of applying the sequence of  $T$  residual blocks to the initial input  $x_0$ . For example, a linear layer preceded by a deep sequence of residual blocks is a universal approximator for Lebesgue integrable functions  $\mathbb{R}^p \rightarrow \mathbb{R}$  (Lin & Jegelka, 2018).

Recently, models arising from residual blocks have gained attention as a means to construct invertible networks; that is, training a network results in a mapping  $x_0 \rightarrow x_T$  for which an inverse mapping  $x_T \rightarrow x_0$  exists. Ability to train a mapping that is guaranteed to be invertible has practical applications; for example, they give rise to normalizing flows (Deco & Brauer, 1995; Rezende & Mohamed, 2015), which allow for sampling from a complicated, multi-modal probability distribution by generating samples from a simple one, and transforming them through an invertible mapping. Thus, it is important to know whether an invertible model can be trained to approximate arbitrary invertible mappings, or if its approximation capabilities are limited.

### 1.1. Invertible Models

We focus our attention on two invertible models proposed recently: i-ResNet, a constrained ResNet, and Neural ODE, a continuous generalization of a ResNet.

**Invertible Residual Networks** While ResNets refer to arbitrary networks with any residual blocks  $x_{t+1} = x_t + f_{\Theta}(x_t, t)$ , that is, can have any residual mapping  $f_{\Theta}(x_t, t)$ , i-ResNets (Behrmann et al., 2019), and their improved variant, Residual Flows (Chen et al., 2019), are built from blocks in which  $f_{\Theta}$  is Lipschitz-continuous with constant lower than 1 as a function of  $x_t$  for fixed  $t$ , which we denote by  $\text{Lip}(f_{\Theta}) < 1$ . This constraint is sufficient (Behrmann et al., 2019) to guarantee invertibility of the residual network, that is, to make  $x_t \rightarrow x_{t+1}$  a one-to-one mapping.

Given the constraint on the Lipschitz constant, an invert-

ible mapping  $x \rightarrow 2x$  cannot be performed by a single i-ResNet layer. But a stack of two layers, each of the form  $x \rightarrow x + (\sqrt{2} - 1)x$  and thus Lipschitz-continuous with constant lower than 1, yields the desired mapping. A single i-ResNet layer  $x_{t+1} = (I + f_{\Theta})(x_t, t)$ , where  $I$  is the identity mapping, is  $\text{Lip}(I + f_{\Theta}) = k < 2$ , and a composition of  $T$  such layers has Lipschitz constant of at most  $K = k^T$ . Thus, for any finite  $K$ , it might be possible to approximate any invertible mapping  $h$  with  $\text{Lip}(h) \leq K$  by a series of i-ResNet layers, with the number of layers depending on  $K$ . However, the question whether the possibility signaled above is true, and i-ResNet have universal approximation capability within the class of invertible continuous mappings, has not been considered thus far.

**Neural Ordinary Differential Equations** Neural ODEs (ODE-Nets) (Chen et al., 2018) are a recently proposed class of differentiable neural network building blocks. ODE-Nets were formulated by observing that processing an initial input vector  $x_0$  through a sequence of residual blocks can be seen as evolution of  $x_t$  in time  $t \in \{1, \dots, T\}$ . Then, a residual block (eq. 1) is a discretization of a continuous-time system of ordinary differential equations (eq. 2)

$$x_{t+1} - x_t = f_{\Theta}(x_t, t), \quad (1)$$

$$\frac{dx_t}{dt} = \lim_{\delta_t \rightarrow 0} \frac{x_{t+\delta_t} - x_t}{\delta_t} = f_{\Theta}(x_t, t). \quad (2)$$

The transformation  $\phi_T : \mathcal{X} \rightarrow \mathcal{X}$  taking  $x_0$  into  $x_T$  realized by an ODE-Net for some chosen, fixed time  $T \in \mathbb{R}$  is not specified directly through a functional relationship  $x \rightarrow f(x)$  for some neural network  $f$ , but indirectly, through the solutions to the initial value problem (IVP) of the ODE

$$x_T = \phi_T(x_0) = x_0 + \int_0^T f_{\Theta}(x_t, t) dt \quad (3)$$

involving some underlying neural network  $f_{\Theta}(x_t, t)$  with trainable parameters  $\Theta$ . By a  $p$ -ODE-Net we denote an ODE-Net that takes a  $p$ -dimensional sample vector on input, and produces a  $p$ -dimensional vector on output. The underlying network  $f_{\Theta}$  must match those dimensions on its input and output, but in principle can have arbitrary internal architecture, including multiple layers of much higher dimensionality.

By the properties of ODEs, ODE-Nets are always invertible, we can just reverse the limits of integration, or alternatively integrate  $-f_{\Theta}(x_t, t)$ . The adjoint sensitivity method (Pontryagin et al., 1962) based on reverse-time integration of an expanded ODE allows for finding gradients of the IVP solutions  $\phi_T(x_0)$  with respect to parameters  $\Theta$  and the initial values  $x_0$ . This allows training ODE-Nets using gradient descent, as well as combining them with other neural network blocks. Since their introduction, ODE-Nets have seen

improved implementations (Rackauckas et al., 2019) and enhancements in training and stability (Gholami et al., 2019; Zhang et al., 2019).

Unlike an unconstrained residual block, a Neural ODE on its own does not have universal approximation capability. Consider a continuous, differentiable, invertible function  $f(x) = -x$  on  $\mathcal{X} = \mathbb{R}$ . There is no ODE defined on  $\mathbb{R}$  that would result in  $x_T = \phi_T(x_0) = -x_0$  (Dupont et al., 2019). Informally, in ODEs, paths  $(x_t, t)$  between the initial value  $(x_0, 0)$  and final value  $(x_T, T)$  have to be continuous and cannot intersect in  $\mathcal{X} \times \mathbb{R}$  for two different initial values, and paths corresponding to  $x \rightarrow -x$  and  $0 \rightarrow 0$  would need to intersect. By contrast, in an unconstrained residual block sequence, a discrete dynamical system on  $\mathcal{X}$ , we do not have continuous paths, only points at unit-time intervals, with an arbitrary transformation between points; finding an unconstrained ResNet for  $x \rightarrow -x$  is easy. While ODE-Nets used out-of-the-box have limited modeling capability, some evidence exists that this limitation can be overcome by changing the way way ODE-Nets are applied. Yet, the question whether they can be turned into universal approximators remains open.

## 1.2. Our Contribution

We analyze the approximation capabilities of ODE-Nets and i-ResNets. The results most closely related to ours have been recently provided by the authors of ANODE (Dupont et al., 2019), who focus on a  $p$ -ODE-Net followed by a linear layer. They provide counterexamples showing that such an architecture is not a universal approximator of  $\mathbb{R}^p \rightarrow \mathbb{R}$  functions. However, they show empirical evidence indicating that expanding the dimensionality and using  $q$ -ODE-Net for  $q > p$  instead of a  $p$ -ODE-Net has positive impact on training of the model and on its generalization capabilities. The authors of i-ResNet (Behrmann et al., 2019) also use expanded dimensionality in their experiments, observing that it leads to a modest increase in model’s accuracy.

Here, we prove that setting  $q = p + 1$  is enough to turn Neural ODE followed by a linear layer into a universal approximator for  $\mathbb{R}^p \rightarrow \mathbb{R}$ . We show similar result for i-ResNet. Our main focus is on modeling invertible functions – homeomorphisms – by exploring pure ODE-Nets and i-ResNets, not capped by a linear layer. We show a class of  $\mathcal{X} \rightarrow \mathcal{X}$  invertible mappings that cannot be expressed by these modeling approaches when they operate within  $\mathcal{X}$ . We then prove that any homeomorphism  $\mathcal{X} \rightarrow \mathcal{X}$ , for  $\mathcal{X} \subset \mathbb{R}^p$ , can be modeled by a Neural ODE / i-ResNet operating on an Euclidean space of dimensionality  $2p$  that embeds  $\mathcal{X}$  as a linear subspace. In our proofs, we relied on using universal approximation property of feed-forward networks, thus, the invertible models may need to be large.

## 2. Background on ODEs, Flows, and Embeddings

This section provides background on invertible mappings and ODEs; we recapitulate standard material, for details see (Utz, 1981; Lee, 2001; Brin & Stuck, 2002; Younes, 2010).

### 2.1. Flows

A mapping  $h : \mathcal{X} \rightarrow \mathcal{X}$  is a *homeomorphism* if  $h$  is a one-to-one mapping of  $\mathcal{X}$  onto itself, and both  $h$  and its inverse  $h^{-1}$  are continuous. Here, we will assume that  $\mathcal{X} \subset \mathbb{R}^p$  for some  $p$ , and we will use the term  $p$ -homeomorphism where dimensionality matters.

A *topological transformation group* or a *flow* (Utz, 1981) is an ordered triple  $(\mathcal{X}, \mathbb{G}, \Phi)$  involving an additive group  $\mathbb{G}$  with neutral element 0, and a mapping  $\Phi : \mathcal{X} \times \mathbb{G} \rightarrow \mathcal{X}$  such that  $\Phi(x, 0) = x$  and  $\Phi(\Phi(x, s), t) = \Phi(x, s+t)$  for all  $x \in \mathcal{X}$ , all  $s, t \in \mathbb{G}$ . Further, mapping  $\Phi(x, t)$  is assumed to be continuous with respect to the first argument. The mapping  $\Phi$  gives rise to a parametric family of homeomorphisms  $\phi_t : \mathcal{X} \rightarrow \mathcal{X}$  defined as  $\phi_t(x) = \Phi(x, t)$ , with the inverse being  $\phi_t^{-1} = \phi_{-t}$ . We will use the term  $p$ -flow to indicate that  $\mathcal{X} \subset \mathbb{R}^p$ .

Given a flow, an *orbit* or a *trajectory* associated with  $x \in \mathcal{X}$  is a subspace  $G(x) = \{\Phi(x, t) : t \in \mathbb{G}\}$ . Given  $x, y \in \mathcal{X}$ , either  $G(x) = G(y)$  or  $G(x) \cap G(y) = \emptyset$ ; two orbits are either identical or disjoint, they never intersect. A point  $x \in \mathcal{X}$  is a *fixed point* if  $G(x) = \{x\}$ . A *path* is a part of the trajectory defined by a specific starting and end points. A path is a subset of  $\mathcal{X}$ ; we will also consider a *space-time path* composed of points  $(x_t, t)$  if we need to make the time evolution explicit.

A *discrete flow* is defined by setting  $\mathbb{G} = \mathbb{Z}$ . For arbitrary homeomorphism  $h$  of  $\mathcal{X}$  onto itself, we easily get a corresponding discrete flow, an iterated discrete dynamical system,  $\phi_0(x) = x$ ,  $\phi_{t+1} = h(\phi_t(x))$ ,  $\phi_{t-1}(x) = h^{-1}(\phi_t(x))$ . Setting  $f(x) = h(x) - x$  gives us a ResNet  $x_{t+1} = x_t + f(x_t)$  corresponding to  $h$ , though not necessarily an i-ResNet, since there is no  $\text{Lip}(f) < 1$  constraint. For Neural ODEs, the type of flow that is relevant is a *continuous flow*, defined by setting  $\mathbb{G} = \mathbb{R}$ , and adding an assumption that the family of homeomorphisms, the function  $\Phi : \mathcal{X} \times \mathbb{R} \rightarrow \mathcal{X}$ , is differentiable with respect to its second argument,  $t$ , with continuous  $d\Phi/dt$ . The difference compared to a discrete flow is that the *flow at time  $t$* ,  $\phi_t(x)$ , is now defined for arbitrary  $t \in \mathbb{R}$ , not just for integers.

Informally, in a continuous flow the orbits are continuous, and the property that orbits never intersect has consequences for what homeomorphisms  $\phi_t$  can result from a flow. Unlike in the discrete case, for a given homeomorphism  $h$  there may not be a continuous flow such that  $\phi_T = h$  for some  $T$ .

We cannot just set  $\phi_T = h$ , what is required is a continuous family of homeomorphisms  $\phi_t$  such that  $\phi_T = h$  and  $\phi_0$  is identity – such family may not exist for some  $h$ . In such case, a Neural ODE would not be able to model  $h$ . While i-ResNets are discrete, the way they are constructed may also limit the space of mappings they can model to a subset of all homeomorphisms, even if each residual mapping is made arbitrarily complex within the Lipschitz constraint.

### 2.2. Continuous Flows and ODEs

Given a continuous flow  $(\mathcal{X}, \mathbb{R}, \Phi)$  one can define a corresponding ODE on  $\mathcal{X}$  by defining a vector  $V(x) \in \mathbb{R}^p$  for every  $x \in \mathcal{X} \subset \mathbb{R}^p$  such that  $V(x) = d\Phi(x, t)/dt|_{t=0}$ . Then, the ODE  $dx/dt = V(x)$  corresponds to continuous flow  $(\mathcal{X}, \mathbb{R}, \Phi)$ . Indeed,  $\Phi(x_0, T) = x_0 + \int_0^T V(x_t) dt$ ,  $\phi_0$  is identity, and  $\phi_{(s+T)}(x_0) = \phi_T(\phi_s(x_0))$  for time-independent  $V$ . Thus, for any homeomorphism family  $\Phi$  defining a continuous flow, there is a corresponding ODE that, integrated for time  $T$ , models the flow at time  $T$ ,  $\phi_T(x)$ .

The vectors of derivatives  $V(x) \in \mathbb{R}^p$  for all  $x \in \mathcal{X}$  are continuous over  $\mathcal{X}$  and are constant in time, and define a *continuous vector field* over  $\mathbb{R}^p$ . The ODEs evolving according to such a time-invariant vector field, where the right-hand side of eq. 2 depends on  $x_t$  but not directly on time  $t$ , are called *autonomous ODEs*, and take the form of  $dx/dt = f_\Theta(x_t)$ .

Any *time-dependent ODE* (eq. 2) can be transformed into an autonomous ODE by removing time  $t$  from being a separate argument of  $f_\Theta(x_t, t)$ , and adding it as part of the vector  $x_t$ . Specifically, we add an additional dimension<sup>1</sup>  $x[\tau]$  to vector  $x$ , with  $\tau = p + 1$ . We equate it with time,  $x[\tau] = t$ , by including  $dx[\tau]/dt = 1$  in the definition of how  $f_\Theta$  acts on  $x_t$ , and including  $x_0[\tau] = 0$  in the initial value  $x_0$ . In defining  $f_\Theta$ , explicit use of  $t$  as a variable is being replaced by using the component  $x[\tau]$  of vector  $x_t$ . The result is an autonomous ODE.

Given time  $T$  and an ODE defined by  $f_\Theta$ ,  $\phi_T$ , the flow at time  $T$ , may not be well defined, for example if  $f_\Theta$  diverges to infinity along the way. However, if  $f_\Theta$  is well behaved, the flow will exist at least locally around the initial value. Specifically, Picard-Lindelöf theorem states that if an ODE is defined by a Lipschitz-continuous function  $f_\Theta(x_t)$ , then there exists  $\varepsilon > 0$  such that the flow at time  $T$ ,  $\phi_T$ , is well-defined and unique for  $-\varepsilon < T < \varepsilon$ , with the range  $\varepsilon$  in which the flow is unique and within certain bounds depending the inverse of the Lipschitz constant. If exists,  $\phi_T$  is a homeomorphism, since the inverse exists and is continuous; simply,  $\phi_{-T}$  is the inverse of  $\phi_T$ .

<sup>1</sup>To avoid confusion with  $x_t$  indicating time, we use  $x[i]$  to denote  $i$ -th component of vector  $x$ .

### 2.3. Flow Embedding Problem for Homeomorphisms

Given a  $p$ -flow, we can always find a corresponding ODE. Given an ODE, under mild conditions, we can find a corresponding flow at time  $T$ ,  $\phi_T$ , and it necessarily is a homeomorphism. Is the class of  $p$ -flows equivalent to the class of  $p$ -homeomorphisms, or only to its subset? That is, given a homeomorphism  $h$ , does a  $p$ -flow such that  $\phi_T = h$  exist? This question is referred to as the problem of embedding the homeomorphism into a flow. In the embedding problem, equality  $\phi_T = h$  is need, which contrasts it with results on approximating  $h$  by  $\phi_T$  up to arbitrary error  $\varepsilon$  in  $L_p$  sense over a bounded domain, as studied in (Brenier & Gangbo, 2003) and in the context of deep models in (Li et al., 2019).

For a homeomorphism  $h : \mathcal{X} \rightarrow \mathcal{X}$ , its *restricted embedding into a flow* is a flow  $(\mathcal{X}, \mathbb{R}, \Phi)$  such that  $h(x) = \Phi(x, T)$  for some  $T$ ; the flow is restricted to be on the same domain as the homeomorphism. Studies of homeomorphisms on simple domains such as a 1D segment (Fort, 1955) or a 2D plane (Andrea, 1965) showed that a restricted embedding does not always exist.

An *unrestricted embedding into a flow* (Utz, 1981) is a flow  $(\mathcal{Y}, \mathbb{R}, \Phi)$  on some space  $\mathcal{Y}$  of dimensionality higher than  $p$ . It involves a homeomorphism  $g : \mathcal{X} \rightarrow \mathcal{Z}$  that maps  $\mathcal{X}$  into some subset  $\mathcal{Z} \subset \mathcal{Y}$ , such that the flow on  $\mathcal{Y}$  results in mappings on  $\mathcal{Z}$  that are equivalent to  $h$  on  $\mathcal{X}$  for some  $T$ , that is,  $g(h(x)) = \Phi(g(x), T)$ . While a solution to the unrestricted embedding problem always exists, it involves a smooth, non-Euclidean manifold  $\mathcal{Y}$ . For a homeomorphism  $h : \mathcal{X} \rightarrow \mathcal{X}$ , the manifold  $\mathcal{Y}$ , variously referred to as the twisted cylinder (Utz, 1981), or a suspension under a ceiling function (Brin & Stuck, 2002), or a mapping torus (Browder, 1966), is a quotient space  $\mathcal{Y} = \mathcal{X} \times [0, 1] / \sim$  defined through the equivalence relation  $(x, 1) \sim (h(x), 0)$ . The flow that maps  $x$  at  $t = 0$  to  $h(x)$  at  $t = 1$  and  $h(h(x))$  at  $t = 2$  involves trajectories in  $\mathcal{X} \times [0, 1] / \sim$  in the following way: for  $t$  going from 0 to 1, the trajectory tracks in a straight line from  $(x, 0)$  to  $(x, 1)$ ; in the quotient space  $(x, 1)$  is equivalent to  $(h(x), 0)$ . Then, for  $t$  going from 1 to 2, the trajectory proceeds from  $(h(x), 0)$  to  $(h(x), 1) \sim (h(h(x)), 0)$ .

The fact that the solution to the embedding problem involves a flow on a non-Euclidean manifold makes applying it in the context of gradient-trained ODE-Nets difficult.

## 3. Approximation of Homeomorphisms by Neural ODEs

In exploring the approximation capabilities of Neural ODEs for  $p$ -homeomorphisms, we will assume that the neural network  $f_{\Theta}(x_t)$  on the right hand side of the ODE is a universal approximator and, if needed, can be made large enough to closely approximate any desired function. Thus,

our concern is with what flows can be modeled by a  $q$ -ODE-Net assuming that  $f_{\Theta}(x_t)$  can have arbitrary internal architecture, including depth and dimensionality, as long as its input-output dimensionality remains fixed at  $q$ . We consider two scenarios,  $q = p$ , and  $q > p$ .

### 3.1. Restricting the Dimensionality Limits Capabilities of Neural ODEs

We show a class of functions that a Neural ODE cannot model, a class that generalizes the  $x \rightarrow -x$  one-dimensional example reported previously (Dupont et al., 2019).

**Theorem 1.** *Let  $\mathcal{X} = \mathbb{R}^p$ , and let  $\mathcal{Z} \subset \mathcal{X}$  be a set that partitions  $\mathcal{X}$  into two or more disjoint, connected subsets  $C_i$ , for  $i = [m]$ ; that is,  $\mathcal{X} = \mathcal{Z} \cup (\bigcup_i C_i)$ . Consider a mapping  $h : \mathcal{X} \rightarrow \mathcal{X}$  that*

- *is an identity transformation on  $\mathcal{Z}$ , that is,  $\forall z \in \mathcal{Z}, h(z) = z$ ,*
- *maps some  $x \in C_i$  into  $h(x) \in C_j$ , for  $i \neq j$ .*

*Then, no  $p$ -ODE-Net can model  $h$ .*

*Proof.* A  $p$ -ODE-Net can model  $h$  if a restricted flow embedding of  $h$  exists. Suppose that it does, a continuous flow  $(\mathcal{X}, \mathbb{R}, \Phi)$  can be found for  $h$  such that the trajectory (see Fig. 2, left) of  $\Phi(x, t)$  is continuous on  $t \in [0, T]$  with  $\Phi(x, 0) = x$ ,  $\Phi(x, T) = h(x)$  for some  $T \in \mathbb{R}$ , all  $x \in \mathcal{X}$ .

If  $h$  maps some  $x \in C_i$  into  $h(x) \in C_j$ , for  $i \neq j$ , the trajectory from  $\Phi(x, 0) = x \in C_i$  to  $\Phi(x, T) = h(x) \in C_j$  crosses  $\mathcal{Z}$  – there is  $z \in \mathcal{Z}$  such that  $\Phi(x, \tau) = z$  for some  $\tau \in (0, T)$ . From uniqueness and reversibility of ODE trajectories, we then have  $\Phi(z, -\tau) = x$ . From additive property of flows, we have  $\Phi(z, T - \tau) = h(x)$ .

Since  $h$  is identity over  $\mathcal{Z}$  and  $\mathcal{Z} \subset \mathcal{X}$ , thus  $h(z) = \Phi(z, T) = \Phi(z, 0) = z$ . That is, the trajectory over time  $T$  is a closed curve starting and ending at  $z$ , and  $\Phi(z, t) = \Phi(z, T + t)$  for any  $t \in \mathbb{R}$ . Specifically,  $\Phi(z, T - \tau) = \Phi(z, -\tau) = x$ . Thus,  $h(x) = x$ . We arrive at a contradiction with the assumption that  $x$  and  $h(x)$  are in two disjoint subsets of  $\mathbb{R}^p$  separated by  $\mathcal{Z}$ . Thus, no  $p$ -ODE-Net can model  $h$ . □

The result above shows that Neural ODEs applied in the most natural way, with  $q = p$ , are restricted in the way distinct regions of the input space can be rearranged in order to learn and generalize from the training set, and the restrictions go well beyond requiring invertibility and continuity. For example, a mirror reflection around a multidimensional plane meets the conditions in the theorem above, and thus cannot be modeled by any Neural ODE.



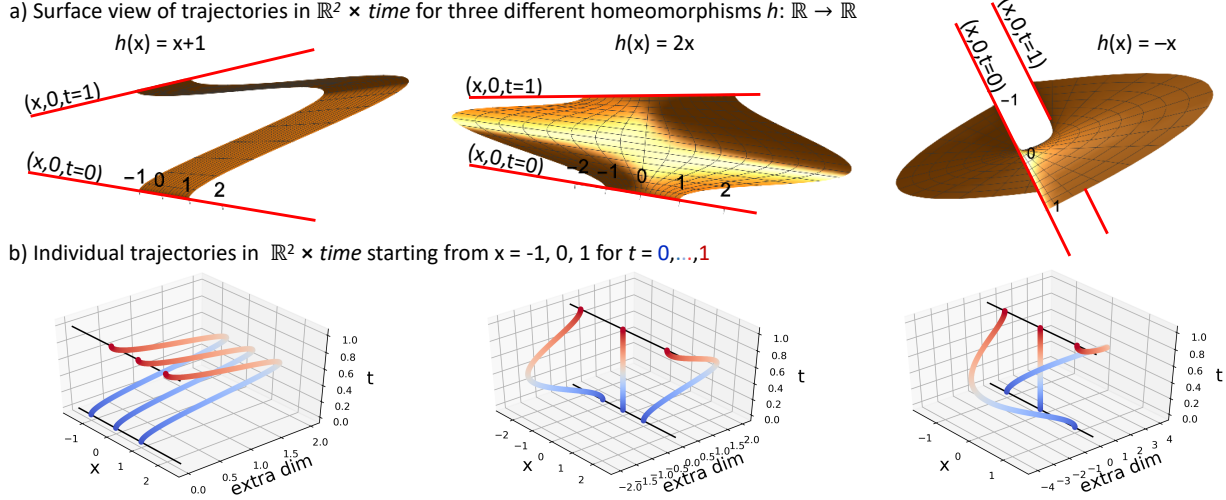


Figure 1. Trajectories in  $\mathbb{R}^{2p}$  that embed an  $\mathbb{R}^p \rightarrow \mathbb{R}^p$  homeomorphism, using  $f(\tau) = (1 - \cos \pi\tau)/2$  and  $g(\tau) = (1 - \cos 2\pi\tau)$ . Three examples for  $p = 1$  are shown, including the mapping  $h(x) = -x$  that cannot be modeled by Neural ODE on  $\mathbb{R}^p$ , but can in  $\mathbb{R}^{2p}$ . In a), shading is used to represent 3D shapes, in b) trajectory color changes with time from blue ( $t = 0$ ) to red ( $t = 1$ ).

### 3.2. Neural ODEs with Extra Dimensions are Universal Approximators for Homeomorphisms

If we allow the Neural ODE to operate on Euclidean space of dimensionality  $q > p$ , we can approximate arbitrary  $p$ -homeomorphism  $\mathcal{X} \rightarrow \mathcal{X}$ , as long as  $q$  is high enough<sup>2</sup>. Here, we show that it suffices to take  $q = 2p$ . We construct a mapping from the original problem space,  $\mathcal{X} \in \mathbb{R}^p$  into  $\mathbb{R}^{2p}$  that<sup>3</sup>

- preserves  $\mathcal{X}$  as a  $p$ -dimensional linear subspace consisting of vectors  $[x, 0^{(p)}]$ ,
- leads to an ODE that maps  $[x, 0^{(p)}] \rightarrow [h(x), 0^{(p)}]$ .

Thus, we provide a solution with a structure that is convenient for out-of-the-box training and inference using Neural ODEs – it is sufficient to add  $p$  zeros to input vectors.

**Theorem 2.** *For any homeomorphism  $h: \mathcal{X} \rightarrow \mathcal{X}$ ,  $\mathcal{X} \subset \mathbb{R}^p$ , if a feed-forward network for mapping  $\delta(x) = h(x) - x$  can be constructed, then there exists a  $2p$ -ODE-Net  $\phi_T: \mathbb{R}^{2p} \rightarrow \mathbb{R}^{2p}$  for  $T = 1$  such that  $\phi_T([x, 0^{(p)}]) = [h(x), 0^{(p)}]$  for any  $x \in \mathcal{X}$ .*

*Proof.* We prove the existence in a constructive way, by showing a vector field in  $\mathbb{R}^{2p}$ , and thus an ODE, with the desired properties.

<sup>2</sup>The quotient space  $\mathcal{Y}$  from Section 2.3, the twisted cylinder resulting from a  $p$ -homeomorphism, is a smooth manifold of dimensionality  $p + 1$ . Strong Whitney embedding theorem (Whitney, 1944) ensures existence of an embedding of  $\mathcal{Y}$  in  $(2p + 2)$ -dimensional Euclidean space as its submanifold.

<sup>3</sup>We use upper subscript  $x^{(p)}$  to denote dimensionality of vectors; that is,  $0^{(p)} \in \mathbb{R}^p$ .

We start with the extended space  $(x, \tau)$  with a variable  $\tau$  corresponding to time added as the last dimension, as in the construction of an autonomous ODE from time-dependent ODE. We then define a mapping  $y(x, \tau): \mathbb{R}^p \times \mathbb{R} \rightarrow \mathbb{R}^{2p}$  that will represent paths starting from  $x$  at time  $\tau = 0$ . For  $\tau \in [0, 1]$ , the mapping (see Fig. 1) is defined through

$$y(x, \tau) = [x + f(\tau)\delta_x, \delta_x g(\tau)]. \quad (4)$$

For each  $x$ , let  $\delta_x \in \mathbb{R}^p$  be defined as  $\delta_x = h(x) - x$ . The functions  $f, g: \mathbb{R} \rightarrow \mathbb{R}$  are required to have continuous first derivative, and have  $f(0) = 0, f(1) = 1, g(\tau) = 0$  iff  $\tau \in \mathbb{Z}$ , and the derivatives  $df/d\tau$  and  $dg/d\tau$  are null at  $\tau \in \mathbb{Z}$  and only there. The mapping indeed just adds  $p$  dimensions of 0 to  $x$  at time  $\tau = 0$ , and at time  $\tau = 1$  it gives the result of the homeomorphism applied to  $x$ , again with  $p$  dimensions of 0

$$y(x, 0) = [x, 0^{(p)}],$$

$$y(x, 1) = [x + \delta_x, 0^{(p)}] = [h(x), 0^{(p)}] = y(h(x), 0).$$

For the purpose of constructing an ODE-Net with universal approximation capabilities,  $\tau \in [0, 1]$  suffices. However, more generally we can define the mapping for  $\tau \notin [0, 1]$ , by setting  $y(x, \tau) = y(h^{\lfloor \tau \rfloor}, \tau - \lfloor \tau \rfloor)$ ; for example,  $y(x, -1.75) = y(h^{-1}(h^{-1}(x)), 0.25)$ . Intuitively, the mapping  $y(x, \tau)$  will provide the position in  $\mathbb{R}^{2p}$  of the time evolution for duration  $\tau$  of an ODE on  $\mathbb{R}^{2p}$  starting from a position corresponding to  $x$ .

For two distinct  $x, x' \in \mathbb{R}^p$ , the paths in  $\mathbb{R}^{2p}$  given by eq. 4 do not intersect at the same position at the same point in time. First, consider the case where  $\delta_x$  is not parallel to  $\delta_{x'}$ . Then, the second set of  $p$  variables is equal

only if  $g(\tau) = 0$ , that is, only at integer  $\tau$ . At those time points, the first set of  $p$  variables takes iterates of  $h$ , that is,  $\dots, h^{-1}, x, h(x), h(h(x)), \dots$ , which are different for different  $x$  because  $h$  is a homeomorphism. Second, consider  $x, x'$  such that  $\delta_{x'} = c\delta_x$  for some  $c$ . Then, either  $c \neq 1$  and thus  $g(\tau) \neq cg(\tau)$  for all non-integer  $\tau$ , that is, the second set of  $p$  variables are always different except at  $\tau$  corresponding to iterates of  $h$ , which are distinct; or  $c = 1$ , and the second set of  $p$  variables are always the same. In the latter case, also  $f(\tau)\delta_x = f(\tau)\delta_{x'}$ , hence the first set of variables is  $x + f(\tau)\delta_x$  is only equal to  $x' + f(\tau)\delta_{x'}$  if  $x = x'$ . Thus, in  $\mathbb{R}^{2p}$ , paths starting from two distinct points do not intersect at the same point in time. Intuitively, we have added enough dimensions to the original space so that we can reroute all trajectories without intersections.

We have  $\tau$  correspond directly to time, that is,  $d\tau/dt = 1$ . The mapping  $y$  has continuous derivative with respect to  $t$ , defining a vector field over the image of  $y$ , a subset of  $\mathbb{R}^{2p}$

$$\frac{dy}{dt} = \left[ \frac{df}{dt}\delta_x, \frac{dg}{dt}\delta_x \right].$$

From the conditions on  $f, g$ , we can verify that this time-dependent vector field defined through derivatives of  $y(x, \tau)$  with respect to time has the same values for  $\tau = 0$  and  $\tau = 1$  for any  $x$

$$\frac{dy}{dt}(x, 0) = [0^{(p)}, 0^{(p)}] = \frac{dy}{dt}(x, 1) = \frac{dy}{dt}(h(x), 0)$$

Thus, the vector field is well-behaved at  $y(x, 1) = y(h(x), 0)$ , it is continuous over the whole image of  $y$ . The vector field above is defined over a closed subset  $y(x, \tau)$  of  $\mathbb{R}^{2p}$ , and can be (see (Lee, 2001), Lemma 8.6) extended to the whole  $\mathbb{R}^{2p}$ . A  $(2p)$ -ODE-Net with a universal approximator network  $f_\Theta$  on the right hand side can be designed to approximate the vector field arbitrarily well. The resulting ODE-Net approximates  $[x, 0^p]$  to  $[h(x), 0^p]$ .  $\square$

Based on the above result, we now have a simple method for training a Neural ODE to approximate a given continuous, invertible mapping  $h$  and, for free, obtain also its continuous inverse  $h^{-1}$ . On input, each sample  $x$  is augmented with  $p$  zeros. For a given  $x$ , the output of the ODE-Net is split into two parts. The first  $p$  dimensions are connected to a loss function that penalizes deviation from  $h(x)$ . The remaining  $p$  dimensions are connected to a loss function that penalizes for any deviation from 0. Once the network is trained, we can get  $h^{-1}$  by using an ODE-Net with  $-f_\Theta$  instead of  $f_\Theta$  used in the trained ODE-Net.

For any well-behaving invertible function  $h$ , given a training set of pairs  $(x, h(x))$ , a large-enough feed-forward network can approximate it on that training set. But even if  $h$  is a homeomorphism, training the network is unlikely to produce

an invertible mapping, it may have  $h(x') = h(x'')$  for some  $x' \neq x''$ , especially outside of the training set. Our results rely on univocal approximation capabilities of feed-forward networks. We required existence of a feed-forward network that can be used to defined the trajectories leading to a flow equal to  $h$ . In our constructions, we relied on a mapping  $\delta(x) = h(x) - x$  and on simple functions  $f$  and  $g$ , thus, the feed-forward network needed to learn the vector field should be comparable in size to a network needed to learn  $h$  without any invertibility guarantees.

## 4. Approximation of Homeomorphisms by i-ResNets

### 4.1. Restricting the Dimensionality Limits Capabilities of i-ResNets

We show that similarly to Neural ODEs, i-ResNets cannot model a simple  $f(x) \rightarrow -x$  homeomorphism  $\mathbb{R} \rightarrow \mathbb{R}$ , indicating that their approximation capabilities are limited.

**Theorem 3.** *Let  $F_n(x) = (I + f_n) \circ (I + f_{n-1}) \circ \dots \circ (I + f_1)(x)$  be an  $n$ -layer i-ResNet, and let  $x_0 = x$  and  $x_n = F_n(x_0)$ . If  $\text{Lip}(f_i) < 1$  for all  $i = 1, \dots, n$ , then there are is no number  $n \geq 1$  and no functions  $f_i$  for all  $i = 1, \dots, n$  such that  $x_n = -x_0$ .*

*Proof.* Consider  $a_0 \in \mathbb{R}$  and  $b_0 = a_0 + \delta_0$ . Then,  $a_1 = a_0 + f_1(a_0)$  and  $b_1 = a_0 + \delta_0 + f_1(a_0 + \delta_0)$ . From  $\text{Lip}(f_1) < 1$  we have that  $|f_1(a_0 + \delta_0) - f_1(a_0)| < |\delta_0|$  (see Fig. 2, center). Let  $\delta_1 = b_1 - a_1$ . Then, we have

$$\begin{aligned} \delta_1 &= a_0 + \delta_0 + f_1(a_0 + \delta_0) - a_0 - f_1(a_0) \\ &= \delta_0 + f_1(a_0 + \delta_0) - f_1(a_0), \\ \delta_1 &> \delta_0 - |\delta_0|, \\ \delta_1 &< \delta_0 + |\delta_0|. \end{aligned}$$

That is,  $\delta_1$  has the same sign as  $\delta_0$ . Thus, applying the reasoning to arbitrary  $i, i + 1$  instead of 0, 1, if  $a_i < b_i$ , then  $a_{i+1} < b_{i+1}$ , and if  $a_i > b_i$ , then  $a_{i+1} > b_{i+1}$ , for all  $i = 0, \dots, n - 1$ . Assume we can construct an i-ResNet  $F_n$  such that  $F_n(0) = 0$ ; then  $F_n(x) > 0$  for any  $x > 0$ , and  $F_n(x)$  cannot map  $x$  into  $-x$ .  $\square$

The result above leads to a more general observation about paths in spaces of dimensionality higher than one. As with ODE-Nets, we will use  $p$ -i-ResNet to denote an i-ResNet operating on  $\mathbb{R}^p$ .

**Corollary 4.** *Let the straight line connecting  $x_t \in \mathbb{R}^p$  to  $x_{t+1} = x_t + f(x_t) \in \mathbb{R}^p$  be called an extended path  $x_t \rightarrow x_{t+1}$  of a time-discrete topological transformation group on  $\mathcal{X} \in \mathbb{R}^p$ . In  $p$ -i-ResNet, for  $x_t \neq x'_t$ , extended paths  $x_t \rightarrow x_{t+1}$  and  $x'_t \rightarrow x'_{t+1} = x'_t$  do not intersect.*




Figure 2. Illustration of constructions used in proofs. **Left:** Theorem 1. **Center:** Theorem 3. **Right:** Corollary 4.

*Proof.* For two extended paths to intersect, vectors  $x_t, x'_t, x_{t+1}, x'_{t+1}$  have to be co-planar. If we restrict attention to dynamics starting from  $x_t, x'_t$ , we can view it as a one-dimensional system, with the space axis parallel to  $x_t - x'_t$ , and time axis orthogonal to it. If  $x_{t+1} - x'_{t+1}$  is parallel to  $x_t - x'_t$ , Theorem 3 shows that if  $x'_t$  is above  $x_t$ , then  $x'_{t+1}$  is above  $x_{t+1}$ ; extended paths do not intersect.

If  $x_{t+1} - x'_{t+1}$  is not parallel to  $x_t - x'_t$  (see Fig. 2, right), construct a new i-ResNet  $x \rightarrow x + g(x)$  with  $g(x_t) = cf(x_t)$  and  $g(x'_t) = c'f(x'_t)$ , with  $0 < c, c' \leq 1$  to preserve Lipschitz condition<sup>4</sup>. Assuming extended paths from  $x_t, x'_t$  intersect in the original i-ResNet, we can pick  $c, c'$  such that  $x_t + g(x_t) = x'_t + g(x'_t)$ . This preserves the intersection in the new i-ResNet, and makes  $x_{t+1}, x'_{t+1}$  equally far from the line  $x_t - x'_t$ . Now, Theorem 3 can be applied as above, leading to contradiction; the intersection cannot exist in the original i-ResNet.  $\square$

The result allows us to show that i-ResNets faces a similar constraint in its capabilities as Neural ODEs

**Theorem 5.** Let  $\mathcal{X} = \mathbb{R}^p$ , and let  $\mathcal{Z} \subset \mathcal{X}$  and  $h : \mathcal{X} \rightarrow \mathcal{X}$  be the same as in Theorem 1. No  $p$ -i-ResNet can model  $h$ .

*Proof.* Consider a  $T$ -layered i-ResNet on  $\mathcal{X}$ , giving rise to extended space-time paths in  $\mathcal{X} \times [0, T]$ , with integer  $t \in [0, T]$  corresponding to activations in subsequent layers. For any  $x \in \mathcal{Z}$ , the extended path in  $\mathcal{X} \times [0, T]$  starts at  $(x, 0)$  and ends at  $(x, T)$ . Since i-ResNet layers are continuous transformations, the union of all extended paths arising from  $\mathcal{Z}$  is a simply connected subset of  $\mathcal{X} \times [0, T]$ ; it has no holes and partitions  $\mathcal{X} \times [0, T]$  into separate regions. Since extended paths cannot intersect,  $(x, T)$  remains in the same region as  $(x, 0)$ , which is in contradiction with mapping  $h$ .  $\square$

<sup>4</sup>It preserves it for the  $x_t, x'_t$  pair. Then, function  $g$  that is  $\text{Lip}(g) < 1$  over whole  $\mathbb{R}^p$  can always be constructed around  $g(x_t) = cf(x_t)$  and  $g(x'_t) = c'f(x'_t)$  for  $0 < c, c' \leq 1$ ,  $\text{Lip}(f) < 1$ , e.g. by interpolating  $g$  linearly between  $g(x_t)$  and  $g(x'_t)$  over the segment  $[x_t, x'_t]$ , replicating the segment endpoints beyond its ends on the line passing through it, and replicating  $g$  from the line into the rest of  $\mathbb{R}^p$ .

The proof shows that the limitation in capabilities of the two architectures for invertible mappings analyzed here arises from the fact that paths in invertible mappings constructed through NeuralODEs and i-ResNets are not allowed to intersect and from continuity in  $\mathcal{X}$ .

#### 4.2. i-ResNets with Extra Dimensions are Universal Approximators for Homeomorphisms

Similarly to Neural ODEs, expanding the dimensionality of the i-ResNet from  $p$  to  $2p$  by adding zeros on input guarantees that any  $p$ -homeomorphism can be approximated, as long as its Lipschitz constant is finite and an upper bound on it is known during i-ResNet architecture construction.

**Theorem 6.** For any homeomorphism  $h : \mathcal{X} \rightarrow \mathcal{X}$ ,  $\mathcal{X} \subset \mathbb{R}^p$  with  $\text{Lip}(h) \leq k$ , if a feed-forward network for mapping  $\delta(x) = h(x) - x$  can be constructed, then there exists a  $2p$ -i-ResNet  $\phi : \mathbb{R}^{2p} \rightarrow \mathbb{R}^{2p}$  with  $\lfloor k + 4 \rfloor$  residual layers such that  $\phi([x, 0^{(p)}]) = [h(x), 0^{(p)}]$  for any  $x \in \mathcal{X}$ .

*Proof.* For a given invertible i-ResNet approximating  $h$ , define a possibly non-invertible mapping  $\delta(x) = (h(x) - x)/T$ , where  $T = \lfloor k + 1 \rfloor$ ; we have  $\text{Lip}(\delta(x)) < 1$ . An i-ResNet that models  $h$  using  $T + 3$  layers  $\phi_i$  for  $i = 0, \dots, T + 2$  can be constructed in the following way:

$$\begin{aligned} \phi_0([x, 0]) &\rightarrow [x, 0] + [0, \delta(x)], \\ \phi_i([z, y]) &\rightarrow [z, y] + [yT/(T + 1), 0] \\ &\quad \text{for } i = 0, \dots, T, \\ \phi_{T+2}([h(x), \delta(x)]) &\rightarrow [h(x), \delta(x)] + [0, -\delta(x)] \end{aligned}$$

The first layer maps  $x$  into  $\delta(x)$  and stores it in the second set of  $p$  activations. The subsequent  $T + 1$  layers progress in a straight line from  $[x, \delta(x)]$  to  $[h(x), \delta(x)]$  in  $T + 1$  constant-length steps, and the last layer restores null values in the second set of  $p$  activations.

All layers are continuous mappings. The residual part of the first layer has Lipschitz constant below one, since  $\text{Lip}(\delta(x)) < 1$ . The middle layers have residual part constant in  $z$  and contractive in  $y$ , with Lipschitz constant below one. The residual part of the last layer is a mapping

of the form  $[h, \delta] \rightarrow [0, -\delta]$ . For a pair  $x, x' \in \mathcal{X}$ , let  $h = h(x), h' = h(x'), \delta = \delta(x), \delta' = \delta(x')$ . We have  $\|[0, -\delta] - [0, -\delta']\| = \|\delta - \delta'\| \leq \|[h, \delta] - [h', \delta']\|$ , with equality only if  $h = h'$ . Since  $h(x)$  is a homeomorphism, from invertibility of  $h(x)$  we have that  $h = h'$  implies  $x = x'$  and thus  $\delta = \delta'$ ; hence, the residual part of the last layer also has Lipschitz constant below one.  $\square$

The construction above suggests that while on the order of  $k$  layers may be needed to approximate arbitrary homeomorphism  $h(x)$  with  $\text{Lip}(h) \leq k$ , only the first and last layers depend on  $h(x)$  and need to be trained; the middle layers are simple, fixed linear layers. The last layer for  $x \rightarrow h(x)$  is the same as the first layer of  $h(x) \rightarrow x$  would be, but the inverse of the first layer, but since i-ResNet construct invertible mappings  $x \rightarrow x + f(x)$  using possibly non-invertible  $f(x)$ , it has to be trained along with the first layer.

The construction for i-ResNets is similar to that for NeuralODEs, except one does not need to enforce differentiability in the time domain, hence we do not need smooth accumulation and removal of  $\delta(x)$  in the second set of  $p$  activations, and the movement from  $x$  to  $h(x)$  in the original  $p$  dimensions does not need to be smooth. In both cases, the transition from  $x$  to  $h(x)$  progresses along a straight line in the first  $p$  dimensions, with the direction of movement stored in the second set of  $p$  variables.

## 5. Invertible Networks capped by a Linear Layer are Universal Approximators

We show that a Neural ODE or an i-ResNet followed by a single linear layer can approximate functions, including non-invertible functions, equally well as any traditional feed-forward neural network. Since networks with shallow-but-wide fully-connected architecture (Cybenko, 1989; Hornik, 1991), or narrow-but-deep unconstrained ResNet-based architecture (Lin & Jegelka, 2018) are universal approximators, so are ODE-Nets and i-ResNets. Consider a function  $\mathbb{R}^p \rightarrow \mathbb{R}^r$ . For any  $(x, y)$  such that  $y = f(x)$ , the mapping  $[x, 0] \rightarrow [x, y]$  is a  $(p + r)$ -homeomorphism, and as we have shown, can be approximated by a  $2(p + r)$ -ODE-Net or  $2(p + r)$ -i-ResNet;  $y$  can be extracted from the result by a simple linear layer. Through a simple construction, we show that using just  $p + r$  dimensions is sufficient.

**Theorem 7.** *Consider a neural network  $F : \mathbb{R}^p \rightarrow \mathbb{R}^r$  that approximates function  $f : \mathcal{X} \rightarrow \mathbb{R}^r$  that is Lebesgue integrable for each of the  $r$  output dimensions, with  $\mathcal{X} \subset \mathbb{R}^p$  being a compact subset. For  $q = p + r$ , there exists a linear layer-capped  $q$ -ODE-Net that can perform the mapping  $F$ . If  $f$  is Lipschitz, there also is a linear layer-capped  $q$ -i-ResNet for  $F$ .*

*Proof.* Let  $G$  be a neural network that takes input vectors

$x^{(a)} = [x^{(p)}, x^{(r)}]$  and produces  $q$ -dimensional output vectors  $y^{(a)} = [y^{(p)}, y^{(r)}]$ , where  $y^{(r)} = F(x^{(p)})$  is the desired transformation.  $G$  is constructed as follows: use  $F$  to produce  $y^{(r)} = F(x^{(p)})$ , ignore  $x^{(r)}$ , and always output  $y^{(p)} = 0$ . Consider a  $q$ -ODE-Net defined through  $dx/dt = G(x_t) = [0^{(p)}, F(x_t^{(p)})]$ . Let the initial value be  $x_0 = [x^{(p)}, 0^{(r)}]$ . The ODE will not alter the first  $p$  dimensions throughout time, hence for any  $t$ ,  $F(x_t^{(p)}) = y^{(r)}$ . After time  $T = 1$ , we will have

$$\begin{aligned} x_T &= x_0 + \int_0^1 G(x_t) dt = [x^{(p)}, 0^{(r)}] + \int_0^1 [0^{(p)}, y^{(r)}] dt \\ &= [x^{(p)}, F(x^{(p)})]. \end{aligned}$$

Thus, for any  $x \in \mathbb{R}^p$ , the output  $F(x)$  can be recovered from the output of the ODE-Net by a simple, sparse linear layer that ignores all dimensions except the last  $r$ , which it returns. A similar construction can be used for defining layers of i-ResNet. We define  $k$  residual layers, each with residual mapping  $[x^{(p)}, \dots] \rightarrow [x^{(p)}, \dots] + [0^{(p)}, F(x^{(p)})/k]$ . If  $\text{Lip}(F) < k$ , then the residual mapping  $[x^{(p)}, \dots] \rightarrow [0^{(p)}, F(x^{(p)})/k]$  has Lipschitz constant below 1.  $\square$

The result shows that invertible models can serve as a building block for arbitrary learning problems.


## 6. Experimental Results

### 6.1. i-ResNets

We tested whether i-ResNet operating in one dimension can learn to perform the  $x \rightarrow -x$  mapping, and whether adding one more dimension has impact on the ability learn the mapping. To this end, we constructed a network with five residual blocks. In each block, the residual mapping is a single linear transformation, that is, the residual block is  $x_{t+1} = x_t + W x_t$ . We used the official i-ResNet PyTorch package (Behrmann et al., 2019) that relies on spectral normalization (Miyato et al., 2018) to limit the Lipschitz constant to less than unity. We trained the network on a set of 10,000 randomly generated values of  $x$  uniformly distributed in  $[-10, 10]$  for 100 epochs, and used an independent test set of 2,000 samples generated similarly.

For the one-dimensional  $x \rightarrow -x$  and the two-dimensional  $[x, 0] \rightarrow [-x, 0]$  target mapping, we used MSE as the loss. Adding one extra dimension results in successful learning of the mapping, confirming Theorem 6. The test MSE on each output is below  $10^{-10}$ ; the network learned to negate  $x$ , and to bring the additional dimension back to null, allowing for invertibility of the model. For the i-ResNet operating in the original, one-dimensional space, learning is not successful (MSE of 33.39), the network learned a mapping  $x \rightarrow cx$  for a small positive  $c$ , that is, the mapping closest to negation of





**Figure 3. Left and center:** test set cross-entropy loss, for increasing number  $d$  of null channels added to RGB images. For each  $d$ , the input images have dimensionality  $32 \times 32 \times (3 + d)$ . Left: ODE-Net with  $k=64$  convolutional filters; center:  $k=128$  filters. **Right:** Minimum of test set cross-entropy loss across all epochs as a function of  $d$ , the number of null channels added to input images, for ODE-Nets with different number of convolutional filters,  $k$ .

$x$  that can be achieved while keeping non-intersecting paths, confirming experimentally Corollary 4.

We also performed a similar test for a 2D mapping, where the homeomorphism to be learned is a reflection around a random line. We used random data sampled uniformly from the  $[-1, 1]^2$  cube, with training and test sets sizes as above. For the variant with augmented dimensionality, we used a loss with equal weight in front of the MSE for approximating  $h(x)$  on the first two dimensions and the  $L_2$  penalty on the added two dimensions. The test set MSE for the model that uses only two dimensions is 0.166. Once two more dimensions are added, test set MSE drops to  $2^{-11}$  for approximating  $h$  on the first two dimensions, and the other two dimensions returns to almost zero, their mean magnitude is  $9^{-12}$ . These results confirm Theorem 6 beyond one-dimensional settings.

## 6.2. Neural ODEs

To validate Theorem 2, we used the same two-dimensional problem as for i-ResNets, learning to approximate a reflection around a random line. The test set MSE for the model that only uses two dimensions is 0.166, same as with i-ResNet. Once two more dimensions are added, test set MSE drops to  $3^{-7}$  for approximating  $h$  on the first two dimensions, and the other two dimensions returns to almost zero, they have magnitude of  $6^{-9}$ .

We also performed experiments to validate if the  $q \geq 2p$ -dimensions threshold beyond which any  $p$ -homeomorphism can be approximated by a  $q$ -ODE-Net can be observed empirically in a practical classification problem. We used the CIFAR10 dataset (Krizhevsky, 2009) that consists of  $32 \times 32$  RGB images, that is, each input image has dimensionality of  $p = 32 \times 32 \times 3$ . We constructed a series of  $q$ -ODE-Nets with dimensionality  $q \geq p$ , and for each measured the cross-entropy loss for the problem of classifying CIFAR10 images into one of ten classes. We used the default split of the dataset into 50,000 training and 10,000 test images.

In designing the architecture of the neural network underlying the ODE we followed ANODE (Dupont et al., 2019). Briefly, the network is composed of three 2D convolutional layers. The first two convolutional layers use  $k$  filters, and the last one uses the number of input channels as the number of filters, to ensure that the dimensionalities of the input and output of the network match. The convolution stack is followed by a ReLU activation function. A linear layer, with softmax activation and cross-entropy loss, operates on top the ODE block. We used torchdiffeq package (Chen et al., 2018) and trained on a single NVIDIA Tesla V100 GPU.

To extend the dimensionality of the space in which the ODE operates, we introduce additional null channels on input, that is, we use input images of the form  $32 \times 32 \times (3 + d)$ . Then, to achieve  $q = 2p$ , we need  $d = 3$ . We tested  $d \in \{0, \dots, 7\}$ . To analyze how the capacity of the network interplays with the increases in input dimensionality, we also experimented with varying the number of convolutional filters,  $k$ , in the layers inside the ODE block.

The results in Fig. 3 show that once the networks with small network capacity, below 64 filters, behaves differently than networks with 64 or more filters. Once the network capacity is high enough, at 64 filters or more, adding dimensions past beyond 3, that is, beyond  $2p$ , results in slower decrease in test set loss. To quantify if this slowdown is likely to arise by chance, we calculated the change in test set loss  $\ell_d$  for dimensionality  $d$  as the  $d$  increases by one,  $\delta_d = \ell_d - \ell_{d-1}$ , for  $d=1, \dots, 7$ . We pooled the results from experiments with 64 convolution filter or more. Two-tailed nonparametric Mann-Whitney U test between  $\delta_1, \dots, \delta_3$  and  $\delta_4, \dots, \delta_7$  shows the change of trend is significant ( $p=.002$ ).

## Acknowledgements

T.A. is supported by NSF grant IIS-1453658.

## References

- Andrea, S. A. On homeomorphisms of the plane, and their embedding in flows. *Bulletin of the American Mathematical Society*, 71(2):381–383, 1965.
- Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D., and Jacobsen, J.-H. Invertible residual networks. In *International Conference on Machine Learning*, pp. 573–582, 2019.
- Brenier, Y. and Gangbo, W.  $L^p$  approximation of maps by diffeomorphisms. *Calculus of Variations and Partial Differential Equations*, 16(2):147–164, 2003.
- Brin, M. and Stuck, G. *Introduction to dynamical systems*. Cambridge university press, 2002.
- Browder, W. Manifolds with  $\pi_1 = z$ . *Bulletin of the American Mathematical Society*, 72(2):238–244, 1966.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pp. 6571–6583, 2018.
- Chen, T. Q., Behrmann, J., Duvenaud, D. K., and Jacobsen, J.-H. Residual flows for invertible generative modeling. In *Advances in Neural Information Processing Systems*, pp. 9913–9923, 2019.
- Cybenko, G. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Deco, G. and Brauer, W. Nonlinear higher-order statistical decorrelation by volume-conserving neural architectures. *Neural Networks*, 8(4):525–535, 1995.
- Dupont, E., Doucet, A., and Teh, Y. W. Augmented neural ODEs. *arXiv preprint arXiv:1904.01681*, 2019.
- Fort, M. K. The embedding of homeomorphisms in flows. *Proceedings of the American Mathematical Society*, 6(6): 960–967, 1955.
- Gholami, A., Keutzer, K., and Biros, G. ANODE: unconditionally accurate memory-efficient gradients for neural ODEs. *arXiv preprint arXiv:1902.10298*, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Hornik, K. Approximation capabilities of multilayer feed-forward networks. *Neural networks*, 4(2):251–257, 1991.
- Krizhevsky, A. Learning multiple layers of features from tiny images. 2009.
- Lee, J. M. *Introduction to smooth manifolds*. Springer, 2001.
- Li, Q., Lin, T., and Shen, Z. Deep learning via dynamical systems: An approximation perspective. *arXiv preprint arXiv:1912.10382*, 2019.
- Lin, H. and Jegelka, S. ResNet with one-neuron hidden layers is a universal approximator. In *Advances in Neural Information Processing Systems*, pp. 6169–6178, 2018.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- Pontryagin, L. S., Mishchenko, E., Boltyanskii, V., and Gamkrelidze, R. The mathematical theory of optimal processes. 1962.
- Rackauckas, C., Innes, M., Ma, Y., Bettencourt, J., White, L., and Dixit, V. DiffEqFlux.jl—a Julia library for neural differential equations. *arXiv preprint arXiv:1902.02376*, 2019.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538, 2015.
- Utz, W. The embedding of homeomorphisms in continuous flows. In *Topology Proc*, volume 6, pp. 159–177, 1981.
- Whitney, H. The singularities of a smooth  $n$ -manifold in  $(2n-1)$ -space. *Ann. of Math*, 45(2):247–293, 1944.
- Younes, L. *Shapes and diffeomorphisms*, volume 171. Springer, 2010.
- Zhang, T., Yao, Z., Gholami, A., Keutzer, K., Gonzalez, J., Biros, G., and Mahoney, M. ANODEV2: A coupled neural ODE evolution framework. *arXiv preprint arXiv:1906.04596*, 2019.