



Qianxiao Li · Ting Lin · Zuowei Shen

Deep learning via dynamical systems: An approximation perspective

Received December 21, 2019

Abstract. We build on the dynamical systems approach to deep learning, where deep residual networks are idealized as continuous-time dynamical systems, from the approximation perspective. In particular, we establish general sufficient conditions for universal approximation using continuous-time deep residual networks, which can also be understood as approximation theories in L^p using flow maps of dynamical systems. In specific cases, rates of approximation in terms of the time horizon are also established. Overall, these results reveal that composition function approximation through flow maps presents a new paradigm in approximation theory and contributes to building a useful mathematical framework to investigate deep learning.

Keywords. Deep learning, approximation theory, controllability

1. Introduction and problem formulation

Despite the empirical success of deep learning, one outstanding challenge is to develop a useful theoretical framework to understand its effectiveness by capturing the effect of sequential function composition in deep neural networks. In some sense, this is a distinguishing feature of deep learning that separates it from traditional machine learning methodologies.

One candidate for such a framework is the *dynamical systems approach* [13, 21, 26], which regards deep neural networks as a discretization of an ordinary differential equation. Consequently, the latter can be regarded as the object of analysis in place of the former. An advantage of this idealization is that a host of mathematical tools from dynamical systems, optimal control and differential equations can then be brought to bear on various issues

Qianxiao Li: Department of Mathematics, National University of Singapore, Singapore 119076;
qianxiao@nus.edu.sg

Ting Lin: School of Mathematical Sciences, Peking University, Beijing, P.R. China;
lintingsms@pku.edu.cn

Zuowei Shen: Department of Mathematics, National University of Singapore, Singapore 119076;
matzuows@nus.edu.sg

Mathematics Subject Classification (2020): Primary 41A99; Secondary 68T07

faced in deep learning, and more importantly, shed light on the role of composition in function approximation and learning.

Since its introduction, the dynamical systems approach led to much progress in terms of novel algorithms [38,55], architectures [5,21,32,40,46,48,52] and emerging applications [7,17,31,57,59]. On the contrary, the present work is focused on the theoretical underpinnings of this approach. From the optimization perspective, it has been established that learning in this framework can be recast as a mean-field optimal control problem [14,26], and local and global characterizations can be derived based on generalizations of Pontryagin's classical maximum principle and the Hamilton–Jacobi–Bellman equation. Other theoretical developments include continuum limits and connections to optimal transport [44,45,49]. Nevertheless, other fundamental questions in this approach remain largely unexplored, especially when it comes to the function approximation properties of these continuous-time idealizations of deep neural networks. In this paper, we establish some basic results in this direction.

1.1. The supervised learning setting

We first describe the setting of the standard supervised learning problem we study in this paper. We consider a set of inputs $\mathcal{X} \subset \mathbb{R}^n$ and outputs $\mathcal{Y} \subset \mathbb{R}^m$ that are subsets of Euclidean spaces. In supervised learning, we seek to approximate some ground truth or target function, which is a mapping $F : \mathcal{X} \rightarrow \mathcal{Y}$. For example, in a typical image classification problem, each x specifies the pixel values of a $d \times d$ image ($n = d^2$), and y is its corresponding class label, which is a one-hot encoding corresponding to m different classes of images, e.g. for $m = 3$, $y = (0, 1, 0)$ corresponds to a label belonging to the second class. The ground truth F defines the label $y = F(x)$ associated with each image x , and it is the goal of supervised learning to approximate F from data. Concretely, one proceeds in two steps.

First, we specify a hypothesis space

$$\mathcal{H} = \{F_\theta : \mathcal{X} \rightarrow \mathcal{Y} : \theta \in \Theta\}, \quad (1.1)$$

which is a family of approximating functions parameterized by $\theta \in \Theta$. The parameter set Θ is some subset of a Euclidean space. For example, in classical linear basis regression models ($\mathcal{Y} = \mathbb{R}$), we may consider a set of orthonormal basis functions $\{\phi_i \in L^2(\mathcal{X}) : i = 1, 2, \dots\}$ forming a hypothesis space by linear combinations, i.e. $\mathcal{H} = \{\sum_i a_i \phi_i : a_i \in \mathbb{R}, \sum_i a_i^2 < \infty\}$. Of course, there are other hypothesis spaces one can consider, such as deep neural networks that we will discuss later.

Next, we find an approximant $\hat{F} \in \mathcal{H}$ of F by solving an optimization problem typically of the form

$$\inf_{G \in \mathcal{H}} \int_{\mathcal{X}} \ell(F(x), G(x)) d\mu(x), \quad (1.2)$$

Here, μ is a probability measure on \mathcal{X} modelling the input distribution and $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a loss function that is minimized when its arguments are equal. A common choice for

regression problems is the square loss, $\ell(y, y') = \|y - y'\|_2^2$, in which case the solution of (1.2) is a projection of F onto \mathcal{H} in $L^2(\mathcal{X}, \mu)$. For classification problems, typically one uses a surrogate loss function in place of the classification accuracy, e.g. cross entropy loss. Due to non-square loss functions and complex model architectures, in practice problem (1.2) is only solved approximately to give an \hat{F} as an approximation to F . Moreover, one typically does not have an explicit form for μ , but we have data samples from it: $x_i, y_i = F(x_i)$ for $i = 1, \dots, N$. In this case, we can set μ to be the empirical measure $\mu = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$, yielding the so-called empirical risk minimization problem

$$\inf_{G \in \mathcal{H}} \frac{1}{N} \sum_{i=1}^N \ell(F(x_i), G(x_i)). \tag{1.3}$$

This objective function is also called the training loss, since it is the loss function evaluated on the model predictions versus the true labels, averaged over the training samples.

1.2. Deep residual neural networks

In deep learning, the hypothesis space \mathcal{H} consists of functions in the forms of neural networks of varying architectures. In this paper, we will focus on a very successful class of deep network architectures known as residual networks (ResNet) [23]. These neural networks build the hypothesis space by iterating the difference equation

$$z_{s+1} = z_s + \tau f_{\theta_s}(z_s), \quad z_0 = x, \quad s = 0, \dots, S - 1. \tag{1.4}$$

The number S is the total number of layers of the network, and s indexes the layers. The number $\tau > 0$ is a step size, which is taken to be 1 in the original ResNet formulation, but we consider the slightly more general case here. The function f_{θ_s} specifies the architecture of the network, which depends on the trainable parameters θ_s at each layer s . For example, in the simplest case of a fully connected network, we have

$$\begin{aligned} f_{\theta_s}(z) &= V_s \sigma(W_s z + b_s), \\ \theta_s &= (W_s, V_s, b_s), \quad W_s \in \mathbb{R}^{q \times n}, \quad V_s \in \mathbb{R}^{n \times q}, \quad b_s \in \mathbb{R}^q. \end{aligned} \tag{1.5}$$

Here, $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is called the *activation function*, and is applied elementwise to a vector in \mathbb{R}^n . Popular examples include the rectified linear unit (ReLU) $\text{ReLU}(z) = \max(0, z)$, the sigmoid $\text{Sig}(z) = 1/(1 + e^{-z})$ and $\tanh(z)$. Depending on the application, more complex f_θ are employed, such as those involving blocks of fully connected layers or convolution layers. In this paper, we do not make explicit assumptions on the form of f_{θ_s} and consider the general difference equation (1.4) defining the class of residual network architectures. We remark that it is possible to have different parameter dimensions for each layer as is often the case in practice, but for simplicity of analysis we shall take them to have the same dimension and belong to a common parameter set Θ , by possibly embedding in higher-dimensional Euclidean spaces.

Now, let us denote by $\varphi_{S,\tau}(\cdot; \theta)$ the mapping $x \mapsto z_S$ via (1.4). This is the flow map of the difference equation, which depends on the parameters $\theta = (\theta_0, \dots, \theta_{S-1})$. To match

output dimension, we typically introduce another mapping g taken from a family \mathcal{G} of functions from \mathbb{R}^n to $\mathcal{Y} \subset \mathbb{R}^m$ at the end of the network (classification or regression layer, as it is typically called). We will hereafter call it the *terminal family*, and it is usually simple, e.g. some collection of affine functions. Together, they form the S -layer residual network hypothesis space

$$\begin{aligned} \mathcal{H}_{\text{resnet}}(S, \tau) &= \{g \circ \varphi_{S, \tau}(\cdot; \theta) : g \in \mathcal{G}, \theta \in \Theta^S\}, \\ \mathcal{H}_{\text{resnet}} &= \bigcup_{S \geq 0, \tau > 0} \mathcal{H}_{\text{resnet}}(S, \tau). \end{aligned} \tag{1.6}$$

One can see that this hypothesis space is essentially compositional in nature. First, the functions in \mathcal{H} involve the composition of a flow map $\varphi_{S, \tau}$ and the last layer g . Moreover, the flow map $\varphi_{S, \tau}$ itself is a composition of maps, each of which is a step in (1.4). One challenge in the development of a mathematical theory of deep learning is to understand the effect of compositions on approximation and learning, due to the lack of mathematical tools to handle function compositions.

1.3. The dynamical systems viewpoint

The results in this paper concerns a recent approach introduced in part to simplify the complexity arising from the compositional aspects of the residual network hypothesis space. This is the dynamical systems approach, where deep residual networks are idealized as continuous-time dynamical systems [13, 14, 26, 40]. Instead of (1.4), we consider its continuous-time idealization

$$\dot{z}(t) = f_{\theta(t)}(z(t)), \quad \theta(t) \in \Theta, \quad t \in [0, T], \quad z(0) = x. \tag{1.7}$$

That is, we replace the discrete layer numbers s by a continuous variable t , which results in a new continuous-time dynamics described by an ordinary differential equation (ODE). The limit as $\tau \rightarrow 0$ of (1.4) with $T = S\tau$ held constant gives (1.7) with the identification $t \approx \tau s$. Empirical work shows that this is justified since for trained deep residual networks, $z_{s+1} - z_s$ tend to be small [24, 51]. Consequently, the trainable variable θ is now an indexed by a continuous variable t . We will assume that each $f_{\theta(t)}$ is a Lipschitz continuous function on \mathbb{R}^n , so that (1.7) admits a unique solution.

As in discrete time, for a terminal time $T > 0$, $z(T)$ can be seen as a function of its initial condition x , and we denote it by $\varphi_T(\cdot, \theta) : \mathbb{R}^n \rightarrow \mathbb{R}^n$. The map φ_T is known as the *flow map* of the dynamical system (1.7). It depends on the parameter $\theta = \{\theta(t) \in \Theta : t \in [0, T]\}$, which is now a function of time. We impose a weak regularity condition on θ with respect to t by restricting θ to be essentially bounded, i.e. $\theta \in L^\infty([0, T], \Theta)$. As a result, we can replace the hypothesis space (1.6) by

$$\mathcal{H}_{\text{ode}}(T) = \{g \circ \varphi_T(\cdot, \theta) : g \in \mathcal{G}, \theta \in L^\infty([0, T], \Theta)\} \tag{1.8}$$

with the terminal time T playing the role of depth. In words, this hypothesis space contains

functions which are regression/classification layers g composed with flow maps of a dynamical system in the form of an ODE. It is also convenient to consider the hypothesis space of arbitrarily deep continuous-time networks as the union

$$\mathcal{H}_{\text{ode}} = \bigcup_{T>0} \mathcal{H}_{\text{ode}}(T) = \bigcup_{T>0} \{g \circ \varphi_T(\cdot, \theta) : g \in \mathcal{G}, \theta \in L^\infty([0, T], \Theta)\}. \quad (1.9)$$

The key advantage of this viewpoint is that a variety of tools from continuous time analysis can now be applied. This was pursued, for example, in [26, 27] for learning algorithms and in [5, 6, 40] on network stability. In this paper, we are concerned with the problem of approximation, which is one of the most basic mathematical questions we can ask given a hypothesis space. Let us outline the problem below.

1.4. The problem of approximation

The problem of approximation essentially asks how big \mathcal{H}_{ode} is. In other words, what kind of functions can we approximate using functions in \mathcal{H}_{ode} ? Before we present our results, let us first distinguish the concept of approximation and that of representation.

- We say that a function F can be *represented* by \mathcal{H}_{ode} if $F \in \mathcal{H}_{\text{ode}}$.
- In contrast, we say that F can be *approximated* by \mathcal{H}_{ode} if for any $\varepsilon > 0$, there exists an $\hat{F} \in \mathcal{H}_{\text{ode}}$ that is close to F up to error ε . Equivalently, F lies in the closure of \mathcal{H}_{ode} under some topology.

Therefore, representation and approximation are mathematically distinct notions. The fact that some class of mappings cannot be represented by \mathcal{H}_{ode} does not prevent it from being approximated by \mathcal{H}_{ode} to arbitrary accuracy. For example, it is well-known that flow maps must be orientation-preserving (OP), which is a very small set of functions in the Baire category sense [37]. At the same time, it is also known that OP diffeomorphisms are dense in L^p in dimensions larger than 1 [4]. However, what we need here is more than density of flow maps in general: the approximation set should have good structure to facilitate computation. In this paper, we investigate the density of flow maps with structural constraints.

We will work mostly in continuous time. Nevertheless, it makes sense to ask what the results in continuous time imply for discrete dynamics. After all, the latter is what we can actually implement in practice as machine learning algorithms. Observe that in the reverse direction, $z_{s+1} = z_s + \tau f_{\theta_s}(z_s)$ can be seen as a forward Euler discretization of (1.7). It is well-known that for a finite time horizon T and a fixed compact domain, Euler discretization has global truncation error of $\mathcal{O}(T/S)$ in supremum norm (see e.g. [25, Ch. 5]). In other words, any function in \mathcal{H}_{ode} can be uniformly approximated by those in $\mathcal{H}_{\text{resnet}}$. Consequently, density of \mathcal{H}_{ode} implies the density of flow maps of the corresponding discrete deep residual neural network. In this sense, we can see that approximation results in continuous time have immediate consequences for their discrete counterparts.

2. Main results and implications

In this section, we summarize our main results on the approximation properties of \mathcal{H}_{ode} and discuss their significance with respect to related results in the literature in the direction of approximation theory through the viewpoint of function composition, approximation properties of deep neural networks, and controllability problems in dynamical systems. We begin by fixing some notation.

2.1. Notation

Throughout this paper, we adopt the following notation:

- Let K be a Lebesgue measurable subset of \mathbb{R}^n . We denote by $C(K)$ the space of real-valued continuous functions on K , with norm $\|f\|_{C(K)} = \sup_{x \in K} |f(x)|$. Similarly, for $p \in [1, \infty)$, $L^p(K)$ denotes the space of p -integrable measurable functions on K , with norm $\|f\|_{L^p(K)} = (\int_K |f(x)|^p dx)^{1/p}$. Vector-valued functions are denoted similarly.
- A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called *Lipschitz* if $|f(x) - f(x')| \leq L|x - x'|$ for all $x, x' \in \mathbb{R}^n$. The smallest constant L for which this is true is denoted by $\text{Lip}(f)$. Vector-valued Lipschitz functions are defined similarly.
- Given a uniformly continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we denote by ω_f its *modulus of continuity*, i.e. $\omega_f(r) := \sup_{|x-x'| \leq r} |f(x) - f(x')|$.
- For any collection \mathcal{F} of functions on \mathbb{R}^n , we denote by $\overline{\mathcal{F}}$ its closure under the topology of compact convergence. In other words, $f \in \overline{\mathcal{F}}$ if for any compact $K \subset \mathbb{R}^n$ and any $\varepsilon > 0$, there exists $\hat{f} \in \mathcal{F}$ such that $\|f - \hat{f}\|_{C(K)} \leq \varepsilon$. We will refer to this as *approximation closure*.
- For any collection \mathcal{F} of functions on \mathbb{R}^n , we denote by $\text{CH}(\mathcal{F})$ its convex hull and by $\overline{\text{CH}}(\mathcal{F})$ the approximation closure of its convex hull, i.e. closure under the topology of compact convergence.
- We call a function $f : \mathbb{R} \rightarrow \mathbb{R}$ (resp. *strictly*) *increasing* if $x > y$ implies $f(x) \geq f(y)$ (resp. $f(x) > f(y)$).

2.2. Approximation results

Let us begin by slightly simplifying the form of the continuous-time hypothesis space. Let us denote by \mathcal{F} the set of functions that constitute the right hand side of (1.7):

$$\mathcal{F} = \{f_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^n : \theta \in \Theta\}. \tag{2.1}$$

Consequently, we can denote the family of flow maps generated by \mathcal{F} as

$$\Phi(\mathcal{F}, T) := \{x \mapsto z(T) : \dot{z}(t) = f_t(z(t)), f_t \in \mathcal{F}, z(0) = x, t \in [0, T]\}. \tag{2.2}$$

This allows us to write \mathcal{H}_{ode} compactly without explicit reference to parameterization:

$$\mathcal{H}_{\text{ode}} = \mathcal{H}_{\text{ode}}(\mathcal{F}, \mathcal{G}) = \bigcup_{T>0} \{g \circ \varphi : g \in \mathcal{G}, \varphi \in \Phi(\mathcal{F}, T)\}. \tag{2.3}$$

We will hereafter call \mathcal{F} a *control family*, since its elements control the dynamics induced by the differential equation (1.7). Unless specified otherwise, we assume \mathcal{F} contains only Lipschitz functions, which ensures existence and uniqueness of solutions to the corresponding ODEs (see Proposition 3.1). Recall that \mathcal{G} is the terminal family.

The central results in this paper establishes general sufficient conditions on \mathcal{F} and \mathcal{G} that induce an universal approximation property for \mathcal{H}_{ode} . To state the results we will need some definitions concerning properties of the control family. The first is the concept of well functions, which plays a fundamental role in constructing approximation dynamics.

Definition 2.1 (Well function). We say a Lipschitz function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is a *well function* if there exists a bounded open convex set $\Omega \subset \mathbb{R}^n$ such that

$$\{x \in \mathbb{R}^n : h(x) = 0\} = \overline{\Omega}. \tag{2.4}$$

Here the $\overline{\Omega}$ is the closure of Ω in the usual topology on \mathbb{R}^n .

Moreover, we say that a vector-valued function $h : \mathbb{R}^n \rightarrow \mathbb{R}^{n'}$ is a well function if each of its components $h_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is a well function in the sense above.

The name “well function” highlights the rough shape of this type of function: the zero set of a well function is like the bottom of a well. Of course, the “walls” of such a well need not point upwards and we only require that these functions are never zero outside of $\overline{\Omega}$.

We also define the notion of restricted affine invariance, which is weaker than the usual affine invariance.

Definition 2.2 (Restricted affine invariance). Let \mathcal{F} be a set of functions from \mathbb{R}^n to \mathbb{R}^n . We say that \mathcal{F} is *restricted affine invariant* if $f \in \mathcal{F}$ implies $Df(A \cdot + b) \in \mathcal{F}$, where $b \in \mathbb{R}^n$ is any vector, and D, A are any $n \times n$ diagonal matrices such that the entries of D are ± 1 or 0, and the entries of A are smaller than or equal to 1 in absolute value.

Now, let us state our main result on universal approximation of functions by flow maps of dynamical systems in dimension $n \geq 2$.

Theorem 2.3 (Sufficient condition for universal approximation). *Let $n \geq 2$ and let $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be continuous. Suppose that the control family \mathcal{F} and the terminal family \mathcal{G} satisfy the following conditions:*

- (1) *For any compact $K \subset \mathbb{R}^n$, there exists a Lipschitz $g \in \mathcal{G}$ such that $F(K) \subset g(\mathbb{R}^n)$.*
- (2) *$\overline{\text{CH}}(\mathcal{F})$ contains a well function (Definition 2.1).*
- (3) *\mathcal{F} is restricted affine invariant (Definition 2.2).*

Then, for any $p \in [1, \infty)$, compact $K \subset \mathbb{R}^n$ and $\varepsilon > 0$, there exists $\widehat{F} \in \mathcal{H}_{\text{ode}}$ such that

$$\|F - \widehat{F}\|_{L^p(K)} \leq \varepsilon. \tag{2.5}$$

In the language of approximation theory for neural networks, Theorem 2.3 is known as a universal approximation theorem, and \mathcal{H}_{ode} satisfying the conditions laid out is said to

have the *universal approximation property*. Note that this implies immediately, via Euler discretization, that the same universal approximation property holds for $\mathcal{H}_{\text{resnet}}$.

Here, the covering condition $F(K) \subset g(\mathbb{R}^n)$ is in some sense necessary. If the range of g does not cover $F(K)$, say it misses an open subset $U \subset F(K)$, then no flow maps composed with it can approximate F . Note that this condition is easy to satisfy, since any Lipschitz and surjective g suffices. For example, for $m = 1$ (regression problems), we may take $g(x) = w^\top x$ for any $w \neq 0$.

The requirement $n \geq 2$ is also necessary. In one dimension, the result is actually false, due to the topological constraint induced by flow maps of dynamical systems. More precisely, consider one-dimensional regression problems with $n = m = 1$. There is no need to match output dimensions, so we may take $\mathcal{G} = \{\text{id}\}$. Then one can show that each $\hat{F} \in \mathcal{H}_{\text{ode}}$ must be continuous and strictly increasing, and furthermore its closure also contains only increasing functions. Hence, there is no hope to approximate any function that is strictly decreasing on an open interval. However, we can prove the next best thing in this case: any continuous and increasing function can be approximated by a dynamical system driven by the control family \mathcal{F} .

Theorem 2.4 (Sufficient condition for universal approximation in 1D). *Let $n = m = 1$ and take $\mathcal{G} = \{\text{id}\}$. Then Theorem 2.3 holds under the additional assumption that F is increasing.*

Remark 2.5. In one dimension, Theorem 2.4 still holds if one replaces the $L^p(K)$ norm by the $C(K)$ norm, and furthermore one can relax the restricted affine invariance property to invariance with respect to only $D = \pm 1$ and $A = 1$ in Definition 2.2, i.e. we only require symmetry and translation invariance.

Let us now give some examples of control families satisfying the requirements of Theorems 2.3 and 2.4 in order to highlight the general applicability of these results.

Example 2.6 (ReLU Networks). Recall that the fully connected architecture (1.5) with ReLU activation corresponds to the control family

$$\mathcal{F}_{\text{ReLU}} = \{z \mapsto V \text{ReLU}(Wz + b) : W \in \mathbb{R}^{q \times n}, V \in \mathbb{R}^{n \times q}, b \in \mathbb{R}^q\} \quad (2.6)$$

where $\text{ReLU}(z)_i = \max(z_i, 0)$ and $q \geq n$. It is clear that restricted affine invariance holds. Moreover, one can easily construct a well function: Let $\Omega = (-1, 1)^n \subset \mathbb{R}^n$ be the open cube. Consider the function $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$ whose components are all equal and are given by

$$[h(z)]_i = \frac{1}{2n} \sum_{j=1}^n [\text{ReLU}(-1 - z_j) + \text{ReLU}(z_j - 1)], \quad i = 1, \dots, n. \quad (2.7)$$

Clearly, $h \in \text{CH}(\mathcal{F}) \subset \overline{\text{CH}}(\mathcal{F})$ and h is a well function with respect to Ω . Therefore, fully connected residual networks with ReLU activations possesses the universal approximation property as a consequence of our results. Note that this result can be proved using other methods with explicit architectural assumptions, e.g. [29, 34, 41–43].

We now discuss examples of architectural variations that can be handled with our approach. As far as we are aware, such results have not been established in the literature using other means.

Example 2.7 (Other activations). Our results also apply to other network architectures, e.g. with different choice of the activation function. As a demonstration, we consider another commonly used activation function known as the sigmoid activation

$$\text{Sig}(z) = \frac{1}{1 + e^{-z}} \tag{2.8}$$

in place of the ReLU activation in (2.6). We call this family \mathcal{F}_{Sig} . In this case, restricted affine invariance is again obvious. To build a well function, define the scalar soft-threshold function $s : \mathbb{R} \rightarrow \mathbb{R}$ by

$$s(z) = \frac{1}{2} \min(\max(|z| - 1, 0), 1) \tag{2.9}$$

and for positive integers M, N define the scalar function

$$s_{M,N}(z) = \frac{1}{2N} \sum_{k=1}^N [\text{Sig}(M(-q_k - z)) + \text{Sig}(M(z - q_k))]. \tag{2.10}$$

where $q_k = 1 + k/N$. We first show that $s_{M,N}$ can approximate s on any compact subset of \mathbb{R} if M, N are large enough, and it is sufficient to consider the subset to be an interval of the form $[-K, K]$. We now estimate $|s(z) - s_{M,N}(z)|$ directly, and by symmetry we only need to check $0 \leq z \leq K$. There are three cases:

Case 1: $0 \leq z < 1$. Here, $s(z) = 0$. Then $z - q_k, -z - q_k \geq -1/N$ for each k and hence $|s_{M,N}(z)| \leq \frac{1}{1 + \exp(M/N)}$.

Case 2: $q_l \leq z < q_{l+1}$ for some $l \geq 0$. Then $\text{Sig}(M(-q_k - z)) < \frac{1}{1 + \exp(M/N)}$ for all k . Also, $|1 - \text{Sig}(M(z - q_k))| < \frac{1}{1 + \exp(M/N)}$ for $k < l$, since $z - q_k > 1/N$, and $|\text{Sig}(M(z - q_k))| < \frac{1}{1 + \exp(M/N)}$ for $k > l + 1$, since $z - q_k < -1/N$. Combining these estimates we have

$$\left| \frac{l}{2N} - s_{M,N}(z) \right| < \frac{1}{2N} + \frac{1}{1 + \exp(M/N)}, \tag{2.11}$$

and since $|s(z) - \frac{l}{2N}| < \frac{1}{2N}$, we have $|s(z) - s_{M,N}(z)| < \frac{1}{N} + \frac{1}{1 + \exp(M/N)}$.

Case 3: $z \geq 2$. Here, $s(z) = 1$, and the estimates for Case 2 above also hold true. Thus, $|1 - s_{M,N}(z)| < \frac{1}{N} + \frac{1}{1 + \exp(M/N)}$.

Combining the cases above, for any $K > 0$ and $|z| \leq K$ we have

$$|s(z) - s_{M,N}(z)| < \frac{1}{N} + \frac{1}{1 + \exp(M/N)}. \tag{2.12}$$

By letting $M = N^2 \rightarrow \infty$, we can make the right hand side arbitrarily small, as required.

Now, define the function $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$ by

$$[h(z)]_i = \frac{1}{n} \sum_{j=1}^n s(z_j), \quad i = 1, \dots, n. \tag{2.13}$$

It is clear that h is a well function with respect to the cube $(-1, 1)^n$. Moreover, by estimate (2.12), h can be uniformly approximated on any compact subset by $[h_{M,N}]_i = \frac{1}{n} \sum_{j=1}^n s_{M,N}(z_j)$, which belongs to $\text{CH}(\mathcal{F}_{\text{Sig}})$. Thus, $h \in \overline{\text{CH}}(\mathcal{F}_{\text{Sig}})$ and we conclude using our results that continuous fully connected residual networks with sigmoid activations also possess the universal approximation property. Other activations such as tanh can be handled similarly. Importantly, we can see that in our framework, relatively little effort is required to handle such variations in architecture, but for existing approaches whose proofs rely on explicit architectural choices such as ReLU, this may be much more involved.

Example 2.8 (Residual blocks). As a further demonstration of the flexibility of our results, we can consider another type of variation of the basic residual network, which considers a “residual block” with more than one fully connected layer. For example, each block can be of the form

$$\begin{aligned} u_s &= \sigma(W_s^{(1)}z_s + b_s^{(1)}), \\ z_{s+1} &= z_s + V_s\sigma(W_s^{(2)}u_s + b_s^{(2)}), \end{aligned} \tag{2.14}$$

where σ is some non-linear activation function applied elementwise. In fact, the original formulation of residual networks has such a block structure, albeit with convolutional layers and ReLU activations [23]. Now, the corresponding control family for the continuous-time dynamics is

$$\begin{aligned} \mathcal{F}_{\text{block}} &= \{z \mapsto V\sigma(W^{(2)}\sigma(W^{(1)}z + b^{(1)}) + b^{(2)}) : \\ &V \in \mathbb{R}^{n \times q_2}, W^{(1)} \in \mathbb{R}^{q_1 \times n}, b^{(1)} \in \mathbb{R}^{q_1}, W^{(2)} \in \mathbb{R}^{q_2 \times q_1}, b^{(2)} \in \mathbb{R}^{q_2}\}, \end{aligned} \tag{2.15}$$

where $q_2, q_1 \geq n$. We now show that we can deduce universal approximation of this family from previous results. As before, restricted affine invariance holds trivially. Thus, it only remains to show that $\overline{\text{CH}}(\mathcal{F}_{\text{block}})$ contains a well function.

To proceed, we may set $q_1 = q_2 = n$ without loss of generality, since otherwise we can just pad the corresponding matrices/vectors with zeros. Now, let us assume that the “one-layer” control family

$$\mathcal{F}_\sigma = \{x \mapsto V\sigma(Wz + b) : W \in \mathbb{R}^{n \times n}, b \in \mathbb{R}^n\} \tag{2.16}$$

is such that $\overline{\text{CH}}(\mathcal{F}_\sigma)$ contains a well function that is non-negative. From the previous examples, we know that this is true for $\sigma = \text{ReLU}$ and for $\sigma = \text{Sig}$. In addition, we assume that the activation σ satisfies a non-degeneracy condition: there exists a closed interval $I \subset \mathbb{R}$ such that its preimage $\sigma^{-1}(I)$ is also a closed interval. Note that most activations we use in practice satisfy this condition.

Let us define a control family \mathcal{F} , which is a subset of $\mathcal{F}_{\text{block}}$, in which we set $W^{(1)} = I$ and $b^{(1)} = 0$. We also reparameterize the remaining variables as $W^{(2)} = \tilde{W}^{(2)}\tilde{W}^{(1)}$, $b^{(2)} = \tilde{W}^{(2)}\tilde{b}^{(1)} + \tilde{b}^{(2)}$ to obtain the smaller control family

$$\mathcal{F} = \{z \mapsto V\sigma(\tilde{W}^{(2)}[\tilde{W}^{(1)}\sigma(z) + \tilde{b}^{(1)}] + \tilde{b}^{(2)}) : V \in \mathbb{R}^{n \times n}, \tilde{W}^{(1)} \in \mathbb{R}^{n \times n}, \tilde{b}^{(1)} \in \mathbb{R}^n, \tilde{W}^{(2)} \in \mathbb{R}^{n \times n}, \tilde{b}^{(2)} \in \mathbb{R}^n\}, \quad (2.17)$$

We now show that $\overline{\text{CH}}(\mathcal{F})$ contains a well function. Since the activation function is applied elementwise, we may first consider the 1D case, as in the first two examples. Suppose that s is a scalar well function such that $z \mapsto (s(z_1), \dots, s(z_n)) \in \overline{\text{CH}}(\mathcal{F}_\sigma)$ and that s is non-negative (see Examples 2.6 and 2.7 for construction). Then we know that $z \mapsto (s(a\sigma(z_1) + b), \dots, s(a\sigma(z_n) + b)) \in \overline{\text{CH}}(\mathcal{F}) \subset \overline{\text{CH}}(\mathcal{F}_{\text{block}})$ for all $a, b \in \mathbb{R}$. It suffices to verify that $s(a\sigma(\cdot) + b)$ is a scalar well function for some a and b satisfying suitable conditions. Take a closed interval $I \subset \mathbb{R}$ such that $\sigma^{-1}(I)$ is also a closed interval. By rescaling and translating, we can take the zero set of s to be the interval $[-1, 1]$. Then, we choose a, b such that $z \mapsto az + b$ maps I to $[-1, 1]$, from which we can deduce that $s(a\sigma(\cdot) + b)$ is also a well function. We may now construct a well function in n dimensions analogously to the previous examples:

$$[h(z)]_i = \frac{1}{n} \sum_{j=1}^n s(z_j), \quad i = 1, \dots, n, \quad (2.18)$$

and by construction h is a well function in $\overline{\text{CH}}(\mathcal{F}) \subset \overline{\text{CH}}(\mathcal{F}_{\text{block}})$. By our results, this again induces the universal approximation property of its corresponding \mathcal{H}_{ode} .

The above examples serve to illustrate the flexibility of a general sufficient condition in deriving universal approximation results for many different architectures. It is possible that some or even all of these results can be derived using other means (such as reproducing other universal function classes), but such arguments are likely to be involved, and more importantly, they have to be handled on a case-by-case basis, without a systematic approach such as the one introduced in this paper.

We end this section with a remark on a negative example.

Remark 2.9. Observe that linear activations $\sigma(z) = z$ constitute a control family which does not contain a well function in $\overline{\text{CH}}(\mathcal{F})$. We can also immediately see that it cannot produce universal approximating flow maps, since the resulting flow maps are always linear functions.

Let us now discuss the implications of Theorems 2.3 and 2.4 in three broad directions: 1) approximation of functions by compositions; 2) approximation theory of deep neural networks and 3) control theory and dynamical systems.

2.3. Approximation of functions by composition

Let us first discuss our results in the context of classical approximation theory, but through the lens of compositional function approximation. In other words, we will recast classical

approximation methods in the form of a compositional hypothesis space (cf. (1.8))

$$\mathcal{H} = \{g \circ \varphi : g \in \mathcal{G}, \varphi \in \Phi\}, \tag{2.19}$$

which then allows us to compare the results with the setting considered in this paper. As before, we call \mathcal{G} the terminal family, and for convenience we will refer to Φ as a *transformation family*, to highlight the fact that it contains functions whose purpose is to transform the domain of a function g in order to resemble a target function F .

We start with the simplest setting in classical approximation theory, namely linear N -term approximation. Here, we consider a fixed, countable collection \mathcal{D} of functions from \mathbb{R}^n to \mathbb{R} , which we call a *dictionary*. The dictionary is assumed to have some structure so that its elements are simple to represent or compute. Common examples include polynomials, and simple periodic functions such as sines and cosines. Linear N -term approximation takes the first N elements ϕ_1, \dots, ϕ_N of \mathcal{D} and forms an approximant \hat{F} of F via their linear combinations

$$\hat{F}(x) = \sum_{i=1}^N w_i \phi_i(x), \quad w_i \in \mathbb{R}^m, \quad i = 1, \dots, N. \tag{2.20}$$

From the viewpoint of compositional function approximation, we may express the above hypothesis space by considering the linear terminal family

$$\mathcal{G}(N) = \left\{x \mapsto \sum_{i=1}^N w_i x_i : w_i \in \mathbb{R}^m, i = 1, \dots, N\right\} \tag{2.21}$$

Then, we obtain the compositional representation of the hypothesis space

$$\mathcal{H}(N) = \{g \circ \varphi : g \in \mathcal{G}(N), \varphi = (\phi_1, \dots, \phi_N)\}. \tag{2.22}$$

In other words, the basic N -term linear approximation can be recast as a compositional hypothesis space consisting of a linear terminal family and a transformation family containing of just one N -dimensional vector-valued function, whose coordinates are the first N elements of the dictionary \mathcal{D} . This is called *linear approximation*, because for two target functions F_1 and F_2 , whose best approximations (in terms of lowest approximation error, see Section 1) are \hat{F}_1 and \hat{F}_2 respectively, the best approximation of $\lambda_1 F_1 + \lambda_2 F_2$ is $\lambda_1 \hat{F}_1 + \lambda_2 \hat{F}_2$ for any $\lambda_1, \lambda_2 \in \mathbb{R}$.

Non-linear N -term approximation [11] takes this approach a step further, by lifting the restriction that we only use the first N elements in \mathcal{D} . Instead, we are allowed to choose, given a target F , which N functions to pick from \mathcal{D} . For this reason, the dictionaries \mathcal{D} used in non-linear approximation can be an uncountable family of functions. In this case, the compositional family is

$$\mathcal{H}(N) = \{g \circ \varphi : g \in \mathcal{G}(N), \varphi \in \mathcal{D}^N\}. \tag{2.23}$$

The term “non-linear” highlights the fact that the best approximations for functions do not remain invariant under linear combinations.

In these classical scenarios, the compositional formulations (2.22) and (2.23) share certain similarities. Most importantly, their transformation families have simplistic structures, and the terminal family is linear, hence also simple. Consequently, one must rely on having a large N in order to form a good approximation. A function can be efficiently approximated via linear approximation (i.e. requiring a small N) if its expansion coefficients in \mathcal{D} decay rapidly. On the other hand, efficient approximation through non-linear approximation relies on the sparsity of this expansion. Nevertheless, in both cases the complexity of their respective hypothesis spaces arises from a large number of linear combinations of simple functions.

Let us now contrast our results on \mathcal{H}_{ode} , which is also in this compositional form (see (1.8)). First, our results holds for more general terminal families that are not restricted to linear ones. Second, by looking at the form of \mathcal{H}_{ode} and comparing with (2.22) and (2.23), we observe that the complexity of \mathcal{H}_{ode} arises not due to linear combinations of functions, since universal approximation holds despite φ having fixed output dimension. Instead, the complexity of \mathcal{H}_{ode} arises from compositions, a point which we shall now expand on.

Observe that besides the overall compositional structure of a terminal family and a transformation family, a second aspect of composition is also involved in \mathcal{H}_{ode} : the transformation family $\Phi(\mathcal{F}, T)$ is itself generated by compositions of simple functions. To see this, observe that any flow map φ_T of an ODE up to time T can be written as

$$\varphi_T = \varphi_{\tau_M} \circ \varphi_{\tau_{M-1}} \circ \dots \circ \varphi_{\tau_2} \circ \varphi_{\tau_1}, \tag{2.24}$$

where $\tau_1 + \dots + \tau_M = T$, and each φ_{τ_i} represents the portion of the flow map from $t = \sum_{s \leq i-1} \tau_s$ to $t = \sum_{s \leq i} \tau_s$ (we set $\tau_0 = 0$). By increasing the number of such partitions, each φ_{τ_i} becomes closer to the identity mapping. More generally, the family of flow maps forms a continuous group under composition, with the identity element recovered when the time horizon of the flow map goes to 0. Therefore, each member of the transformation family $\Phi(\mathcal{F}, T)$ can be decomposed into a sequence of compositional mappings, each of which can be made arbitrarily close to the identity, as long as at the same time one increases the number of such compositions. While this decomposition holds true for any flow map of an ODE, the main results in this paper go a step further and show that the universal approximation property holds even when the flow map is restricted to one that is generated by some control family \mathcal{F} satisfying the assumptions in Theorem 2.3.

We remark that in classical non-linear approximation, the dictionary \mathcal{D} could also involve compositions. A prime example of this is wavelets [9, 36, 39] where one starts with some template function ψ and generates a dictionary by composing it with translations and dilations. For example, in one dimension the wavelet dictionary has the following compositional representation:

$$\mathcal{D} = \{\psi \circ T : T(x) = (x - x_0)/\lambda, x_0 \in \mathbb{R}, \lambda > 0\}. \tag{2.25}$$

The main contrasting aspect in \mathcal{H}_{ode} is that compositional transformations are much more complex. Instead of simple translations and dilations, the transformation family in \mathcal{H}_{ode} involves complex rearrangement dynamics in the form of an ODE flow that may be adapted to the specific target function F at hand.

In summary, in contrast to classical approximation schemes, \mathcal{H}_{ode} is built from compositions of functions from a simple terminal family \mathcal{G} and a complex transformation family $\Phi(\mathcal{F}, T)$, whose members can further be decomposed into a sequence of compositions of functions that are simple in two aspects: they are close to the identity map and they are generated by a potentially simple control family \mathcal{F} . Consequently, the complexity of \mathcal{H}_{ode} arises almost purely from the process of composition of these simple mappings. In other words, we trade complexity in T (compositions) for N (linear combinations), and can achieve universal approximation even when the transformation family has fixed output dimensions. From this viewpoint, the results in this paper highlight the power of composition for approximating functions.

2.4. Approximation theory of deep neural networks

As discussed previously, the transformation family in \mathcal{H}_{ode} consisting of flow maps is highly complex due to repeated compositions. At the same time, however, just like dictionaries in linear and non-linear approximation, it possesses structure that allows us to carry out approximation in practice. Concretely, recall that each flow map can be decomposed as in (2.24), where each component φ_{τ_i} is not only close to the identity, but is close in such a way that the perturbation from the identity is constrained by the control family \mathcal{F} . Thus, one just needs to parameterize each φ_{τ_i} by selecting appropriate functions from \mathcal{F} and then compose them together to form an approximating flow map. From this viewpoint, the family of deep residual network architectures is a realization of this procedure, by using a one-step forward Euler discretization to approximate each φ_{τ_i} . Concretely,

$$\varphi_{\tau_i}(z) = z + \int_0^{\tau_i} f(z(t)) dt \approx z + \tau_i f(z) \quad (2.26)$$

for τ_i small, which corresponds to the family of deep residual architectures motivated in Section 1. The standard convergence result for Euler discretization [25] allows one to carry approximation results in continuous time to the discrete case. In view of this, we now discuss our results in the context of approximation results in deep learning.

We start with the continuous-time case. Most existing theoretical work on the continuous-time dynamical systems approach to deep learning focus on optimization aspects in the form of mean-field optimal control [14, 30], or connections between the continuous-time idealization to discrete time [44, 45, 49]. The present paper focuses on the approximation aspects of continuous-time deep learning, which are less studied. One exception is the recent work of Zhang et al. [56], who derived some results in the direction of approximation. However, an important assumption there was that the driving force on the right hand side of ODEs (here the control family \mathcal{F}) are themselves universal approximators. Consequently, such results do not elucidate the power of composition and flows, since each “layer” is already complex enough to approximate an arbitrary function, and there is no need for the flow to perform any additional approximation.

In contrast, the approximation results here do not require \mathcal{F} , or even $\overline{\text{CH}}(\mathcal{F})$, to be universal approximators. In fact, \mathcal{F} can be a very small set of functions, and the

approximation power of these dynamical systems is by construction due to the dynamics of the flow. For example, the assumption that $\overline{\text{CH}}(\mathcal{F})$ contains a well function does not imply that \mathcal{F} that drives the dynamical system is complex, since the former can be much larger than the latter. In the 1D ReLU control family that induces the fully connected network with ReLU activations (see (1.4)), one can easily construct a well function with respect to the interval $\Omega = (q_1, q_2)$ by averaging two ReLU functions: $\frac{1}{2}[\text{ReLU}(q_1 - x) + \text{ReLU}(x - q_2)]$, but the control family $\mathcal{F} = \{v\text{ReLU}(w \cdot + b)\}$ is not complex enough to approximate arbitrary functions without further linear combinations. We have already illustrated this in Examples 2.6, 2.7 and 2.8.

We also note that in contrast to [56], the results here for $n \geq 2$ do not require embedding the dynamical system in higher dimensions to achieve universal approximation. The negative results given in [56] (and also [12]), which motivated embedding in higher dimensions, are basically on limitation of representation: flow maps of ODEs are orientation-preserving (OP) homeomorphisms (see Definition 3.2) and thus can only represent such mappings. However, these are not counter-examples for approximation. For instance, it is known that OP diffeomorphisms (and hence OP homeomorphisms) can approximate any L^∞ functions on open bounded domains in dimensions greater than or equal to 2 [4].

Although the present paper focuses on the continuous-time idealization, we should also discuss the results here in relation to the relevant work on the approximation theory of discrete deep neural networks. In this case, one line of work to establish universal approximation is to show that deep networks can approximate some other family of functions known to be universal approximators themselves, such as wavelets [35]. Others focus on certain specific architectures, such as in [3, 10, 15, 20, 29, 33, 54, 60], which sometimes allows for explicit asymptotic approximation rates to be derived for appropriate target function classes. Furthermore, non-asymptotic approximation rates for deep ReLU networks are obtained in [34, 41–43]. They are based on explicit constructions using composition, and hence are similar in flavour to the results here if we take an explicit control family and discretize in time.

With respect to all these works, the main difference of the results presented here is that we study sufficient conditions for approximation. In other words, we do not start with an *a priori* specific architecture (e.g. the form of the function f_θ , or the type of activation σ in (1.4)). In particular, none of the approximation results we present here depend on reproducing some other basis functions that are known to have the universal approximation property. Instead, we derive conditions on the control and terminal families \mathcal{F}, \mathcal{G} that imply the universal approximation property in \mathcal{H}_{ode} . One advantage of this viewpoint is that we can isolate the approximation power that arises from the act of composition from that which arises from the specific architectural choices themselves. As an example, the approximation result in [29] relies on approximating piecewise constant functions with finitely many discontinuities, hence its proof depends heavily on the ReLU activation. Furthermore, the high-dimensional results there require constructing the proximal grid indicator function, which is not straightforward with activations other than ReLU. We note that for the deep non-residual case, more precise approximation results including non-asymptotic rates for the ReLU architecture can be derived [34, 41–43]. In contrast, the

main results in this paper proceed in a more general way without assuming any precise architectures. Therefore, these results have greater applicability to diverse architectures (see Examples 2.6–2.8), and perhaps even novel ones that may arise in future deep learning applications.

2.5. Control theory and dynamical systems

Lastly, the results here are also of relevance to mathematical control theory and the theory of dynamical systems. In fact, the problem of approximating functions by flow maps is closely related to the problem of controllability in control theory [47]. However, there is one key difference: in the usual controllability problem on Euclidean spaces, our task is to steer one particular input x_0 to a desired output value $\varphi(x_0)$. However, here we want to steer the entire set of input values $x \in K$ to $\varphi(x) \in \varphi(K)$ by the *same* controls. This can be thought of as an infinite-dimensional function space version of controllability, which is a much less explored area, and present controllability results in infinite dimensions mostly focus on control of partial differential equations [2, 8].

In the theory of dynamical systems, it is well-known that functions represented by flow maps are subject to some restrictions. For example, [37] gives a negative result that the diffeomorphisms generated by C^1 vector fields are few in the Baire category sense. Some works also give explicit criteria for mappings to be represented by flows, such as [18] in \mathbb{R}^2 , [50] in \mathbb{R}^n , and more recently, [58] generalizes some results to the Banach space setting. However, these results are on exact representation, not approximation, and hence do not contradict the positive results presented in this paper. The results on approximation properties are fewer. A relevant one is [4], showing that every L^p mapping can be approximated by orientation-preserving diffeomorphisms constructed using polar factorization and measure-preserving flows. The results of the current paper give an alternative construction of a dynamical system whose flow also has such an approximation property. Moreover, Theorem 2.3 gives some weak sufficient conditions for any controlled dynamical system to have this property. In this sense, the results here further contribute to the understanding of the density of flow maps in L^p .

3. Preliminaries

In this section, we state and prove where necessary some preliminary results that are used to deduce our main results in the next section.

3.1. Background results on ordinary differential equations

Throughout this paper, we use some elementary properties and techniques in classical analysis of ODEs. For completeness, we collect these results in this section. The proofs of well-known results that are slightly involved are omitted. Readers familiar with basic ODE theory may skip this section, and the unfamiliar readers are referred to [1] for a comprehensive introduction.

The differential equation that generates the transformation family in \mathcal{H}_{ode} is of the form

$$\dot{z} = f_t(z), \quad f_t \in \mathcal{F}, \quad 0 < t \leq T, \quad z(0) = z_0, \tag{3.1}$$

where $z_0, z(t) \in \mathbb{R}^n$. Such equations are called time-inhomogeneous, since the right hand side changes in time. However, in subsequent proofs of approximation results, we usually consider $t \mapsto f_t$ that are piecewise constant, i.e. $f_t = f_i$ for all $t \in [T_i, T_{i+1}]$. In this case, for each interval on which f_t does not change, it is enough to consider the time-homogeneous equation

$$\dot{z} = f(z), \quad 0 < t \leq T, \quad z(0) = z_0, \tag{3.2}$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is fixed in time. An equivalent form of the ODE is the following integral form:

$$z(t) = z_0 + \int_0^t f(z(s)) ds. \tag{3.3}$$

The following classical result can be proved using fixed point arguments (see e.g. [1, Ch. 4]).

Proposition 3.1 (Existence, uniqueness and dependence on initial condition). *Let f be Lipschitz. Then the solution to (3.2) exists and is unique. Moreover, for each t , $z(t)$ is a continuous function of z_0 . If in addition, f is r -times continuously differentiable for $r \geq 1$, then for each t , $z(t)$ is an $(r - 1)$ -times continuously differentiable function of z_0 .*

Recall that the flow map $\varphi_T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ of (3.2) is the mapping $z_0 \mapsto z(T)$ where $\{z(t)\}$ satisfies (3.2). This is well-defined owing to Proposition 3.1. Thus, we hereafter assume f is Lipschitz. Let us now discuss an important constraint that such flow maps satisfy.

Definition 3.2 (Orientation-preserving for diffeomorphism case). We call a diffeomorphism $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ *orientation-preserving* (OP) if $\det J_\varphi(x) > 0$ for all $x \in \mathbb{R}^n$, where J_φ is the Jacobian of φ , i.e. $[J_\varphi(x)]_{ij} = \frac{\partial \varphi_i}{\partial x_j}(x)$.

Proposition 3.3. *Suppose f is twice continuously differentiable. Then the flow map φ_T of (3.2) is an OP diffeomorphism.*

The above follows from the observation that the flow map of $\dot{z} = -f(z)$ is the inverse of φ_T , and that

$$J_{\varphi_T}(x) = \exp\left(\int_0^T J_f(\varphi_t(x)) dt\right) \tag{3.4}$$

is positive definite.

Definition 3.2 requires differentiability. If the mapping φ is only bi-Lipschitz, we can define the Jacobian almost everywhere, and so a bi-Lipschitz mapping is OP if $\det J_\varphi > 0$ almost everywhere. If φ is merely continuous, a proper definition is subtle and can be given by homological techniques [22].

However, if we restrict our interest to low-dimensional spaces, such as on the real line ($n = 1$) or the plane ($n = 2$), the definition of OP can be easily given without any differentiability requirements. In this paper, we only need the $n = 1$ case, where a homeomorphism is OP if and only if it is strictly increasing. The definition is a natural extension of the diffeomorphism case. Below, we prove in the one-dimensional case that flow maps that are not necessarily differentiable must still be orientation-preserving homeomorphisms. This is sufficient to prove our subsequent results. Note that this result is well-known (see [1, Ch. 1]) but we include its proof for completeness.

Proposition 3.4. *Let $n = 1$ and φ_T be the flow map of (3.2). Then φ_T is increasing.*

Proof. It is enough to show that if z_1 and z_2 are solutions of (3.2), but with different initial values $x_1 < x_2$, then $z_1(t) < z_2(t)$ for all $t \geq 0$. Suppose not: $z_1(t_0) = z_2(t_0)$ for some t_0 . Consider the following ODE:

$$\dot{w} = -f(w), \quad w(0) = z_1(t_0). \tag{3.5}$$

Then both $z_1(t_0 - \cdot)$ and $z_2(t_0 - \cdot)$ are solutions to the above. By uniqueness we have $z_1(t_0 - t) = z_2(t_0 - t)$ for all t , which implies $x_1 = z_1(0) = z_2(0) = x_2$, a contradiction. Since both z_1 and z_2 are continuous in t , we have $z_1(t) < z_2(t)$ for all t . ■

Next we state a version of the well-known Grönwall inequality [19].

Proposition 3.5 (Grönwall’s inequality). *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a scalar function such that $f(t) \geq 0$ and $f(t) \leq A + B \int_0^t f(\tau) d\tau$. Then $f(t) \leq Ae^{Bt}$.*

Finally, we prove some practical results, which follow easily from classical results but are used in some proofs of the main body.

Proposition 3.6. *Let $n = 1$ and $z(\cdot; x)$ be the solution of the ODE (3.2) with initial value x . If x is in some compact set $K \subset \mathbb{R}$, then the finite time modulus of continuity*

$$\omega_{z(\cdot; x), [0, T]}(r) = \sup_{\substack{0 \leq t_1 \leq t_2 \leq T \\ |t_2 - t_1| \leq r}} |z(t_1; x) - z(t_2; x)| \tag{3.6}$$

converges to 0 as $r \rightarrow 0$ uniformly in $x \in K$.

Proof. We denote $a = \min K$ and $b = \max K$, By Proposition 3.4, we know that

$$z(t; a) \leq y(t; x) \leq y(t; b),$$

thus $H = \{z(t; x) : x \in K, t \in [0, T]\} \subset [\min_t z(t; a), \max_t z(t; b)]$ is compact, and so is $f(H) = \{f(h) : h \in H\}$. Let $M = \max_{h \in H} |f(h)|$. We have

$$\sup_{\substack{0 \leq t_1 \leq t_2 \leq T \\ |t_2 - t_1| \leq r}} |z(t_1; x) - z(t_2; x)| \leq rM, \tag{3.7}$$

implying the result. ■

The following proposition shows that in one dimension, if we have a well function, we can transport one point into another if they are on the same side of the well function’s zero interval. Note that by definition, the well function cannot change sign outside of this interval.

Proposition 3.7. *Let $n = 1$. Suppose $f(x) < 0$ for all $x \geq x_0$. Then for $x_0 < x_1 < x_2$, consider the ODE*

$$\dot{z} = f(z), \quad z(0) = x_2. \tag{3.8}$$

Then ultimately the ODE system will reach x_1 , i.e., for some T , $z(T) = x_1$.

Before proving it, we give a simple example to illustrate this proposition. Suppose that $f(z) = -\text{ReLU}(z - x_0)$. Then direct computation shows $z(T) = (x_2 - x_0) \exp(-T) + x_0$. In this case, $T = \ln((x_2 - x_0)/(x_1 - x_0))$ suffices. Intuitively, this proposition shows that under the stated conditions, x_0 is an attractor for the unbounded interval (x_0, ∞) .

Proof of Proposition 3.7. Notice that f is assumed to be continuous, and it suffices to give an estimate on $z(T)$. We only need to prove that for some T , $z(T) < x_1$, and this can be easily derived from the continuity of $T \mapsto z(T)$ and from $z(0) = x_2$. Choose an arbitrary $\tilde{x}_1 \in (x_0, x_1)$ and define $m = -\sup_{x \in [\tilde{x}_1, x_2]} f(x)$. We have

$$z(t) = x_2 + \int_0^t f(z(s)) ds. \tag{3.9}$$

Set $t = (x_2 - x_1)/m$. If $z(t) \leq \tilde{x}$ (which is smaller than x_1) then we are done by continuity. Otherwise, $z(t) \leq -m$ for all $t' \in [0, t]$, yielding

$$z(t) \leq x_2 + \int_0^t (-m) ds = x_2 - m \frac{x_2 - x_1}{m} = x_1, \tag{3.10}$$

which again implies our result by continuity. ■

With these results on ODEs in mind, we now present the proofs of our main results.

3.2. From approximation of functions to approximation of domain transformations

Now, we show that under mild conditions, as long as we can approximate any continuous domain transformation $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ using flow maps, we can show that \mathcal{H}_{ode} is a universal approximator. Consequently, we can pass to the problem of approximating an arbitrary φ by flow maps and establish our main results.

Proposition 3.8. *Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be continuous and $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be continuous. Let $K \subset \mathbb{R}^n$ be compact and suppose $g(\mathbb{R}^n) \supset F(K)$. Then, for any $\varepsilon > 0$ and $p \in [1, \infty)$, there exists a continuous function $\varphi : K \rightarrow \mathbb{R}^n$ such that*

$$\|F - g \circ \varphi\|_{L^p(K)} \leq \varepsilon. \tag{3.11}$$

Proof. This follows from a general result on function composition proved in [28]. We prove this in the special case here for completeness.

The set $F(K)$ is compact, so for any $\delta > 0$ we can form a partition $F(K) = \bigcup_{i=1}^N B_i$ with $\text{diam}(B_i) \leq \delta$. By assumption, $g^{-1}(B_i)$ is non-empty for each i , so let us pick $z_i \in g^{-1}(B_i)$. For each i we define $A_i = F^{-1}(B_i) \cap K$ so that $\{A_i\}$ forms a partition of K . By inner regularity of the Lebesgue measure, for any $\delta' > 0$ and for each i we can find a compact $K_i \subset A_i$ such that $\lambda(A_i \setminus K_i) \leq \delta'$ (λ is the Lebesgue measure) and the K_i 's are disjoint. By Urysohn's lemma, for each i there exists a continuous function $\varphi_i : K \rightarrow [0, 1]$ such that $\varphi_i = 1$ on K_i and $\varphi_i = 0$ on $\bigcup_{j \neq i} K_j$. Now, we form the continuous function

$$\varphi(x) = \sum_{i=1}^N z_i \varphi_i(x). \tag{3.12}$$

Define $K' = \{\sum_{i=1}^N \alpha_i z_i : \alpha_i \in [0, 1]\}$, which is compact and $\varphi(K) \subset K'$. Then

$$\begin{aligned} \|F - g \circ \varphi\|_{L^p(K)} &\leq \sum_{i=1}^N \|F - g \circ \varphi\|_{L^p(K_i)} + \sum_{i=1}^N \|F - g \circ \varphi\|_{L^p(A_i \setminus K_i)} \\ &\leq \sum_{i=1}^N \|F - g \circ \varphi_i\|_{L^p(K_i)} + [\|F\|_{C(K)} + \|g\|_{C(K')}] N \delta'. \end{aligned} \tag{3.13}$$

We take δ' small enough so that the last term is bounded by δ . Then

$$\|F - g \circ \varphi\|_{L^p(K)} \leq \sum_{i=1}^N \delta \lambda(K_i) + \delta \leq (1 + \lambda(K)) \delta. \tag{3.14}$$

Taking $\delta = \varepsilon / (1 + \lambda(K))$ yields the result. ■

We shall hereafter assume that $g(\mathbb{R}^n) \supset F(K)$, which as discussed earlier is easily satisfied. Hence we have the following immediate corollary.

Corollary 3.9. *Assume the conditions in Proposition 3.8 hold and g is Lipschitz. Let Φ be a collection of continuous functions from \mathbb{R}^n to \mathbb{R}^n such that for any $\delta > 0$, compact $K \subset \mathbb{R}^n$ and continuous $\varphi_1 : \mathbb{R}^n \rightarrow \mathbb{R}^n$, there exists $\varphi_2 \in \Phi$ with $\|\varphi_1 - \varphi_2\|_{L^p(K)} \leq \delta$. Then there exists $\varphi \in \Phi$ such that $\|F - g \circ \varphi\|_{L^p(K)} \leq \varepsilon$.*

Proof. By Proposition 3.8, there is a φ_1 such that $\|F - g \circ \varphi_1\|_{L^p(K)} \leq \varepsilon/2$. Now take $\varphi \in \Phi$ such that $\|\varphi_1 - \varphi\|_{L^p(K)} \leq \varepsilon / (2 \text{Lip}(g))$. Then

$$\begin{aligned} \|F - g \circ \varphi\|_{L^p(K)} &\leq \|F - g \circ \varphi_1\|_{L^p(K)} + \|g \circ \varphi_1 - g \circ \varphi\|_{L^p(K)} \\ &\leq \frac{\varepsilon}{2} + \text{Lip}(g) \left(\int_K \|\varphi_1(x) - \varphi(x)\|^p dx \right)^{1/p} \leq \varepsilon. \end{aligned} \tag{3.15}$$

■

3.3. Properties of attainable sets and approximation closures

Owing to Corollary 3.9, for the rest of the paper we will focus on proving universal approximation of continuous transformation functions φ from \mathbb{R}^n to \mathbb{R}^n by flow maps of the dynamical system

$$\dot{z}(t) = f_t(z(t)), \quad f_t \in \mathcal{F}, \quad z(0) = x, \tag{3.16}$$

after which we can deduce universal approximation properties of \mathcal{H}_{ode} via Corollary 3.9.

We now establish some basic properties of flow maps and closure properties. In principle, in our hypothesis space (2.3) we allow $t \mapsto f_t(z)$ to be any essentially bounded measurable mapping for any $z \in \mathbb{R}^n$. However, it turns out that to establish approximation results, it is enough to consider the smaller family of mappings piecewise constant in time, i.e. $f_t = f_j \in \mathcal{F}$ for $t \in [t_{j-1}, t_j)$. For a fixed f in the control family \mathcal{F} , to emphasize dependence on f we denote by φ_τ^f the flow map of the following ODE at time τ :

$$\dot{z}(t) = f(z(t)), \quad z(0) = x. \tag{3.17}$$

That is,

$$\varphi_\tau^f = z(\tau) \quad \text{where} \quad \dot{z}(t) = f(z(t)), \quad z(0) = x, \quad t \in [0, \tau]. \tag{3.18}$$

The attainable set of a finite time horizon T due to piecewise constant controls, denoted as $\mathcal{A}_{\mathcal{F}}(T)$, is defined as

$$\mathcal{A}_{\mathcal{F}}(T) = \{\varphi_{\tau_k}^{f_k} \circ \varphi_{\tau_{k-1}}^{f_{k-1}} \circ \dots \circ \varphi_{\tau_1}^{f_1} : \tau_1 + \dots + \tau_k = T, f_1, \dots, f_k \in \mathcal{F}, k \geq 1\}, \tag{3.19}$$

In other words, $\mathcal{A}_{\mathcal{F}}(T)$ contains the flow map of an ODE whose right hand side is f_i for $t \in [t_{i-1}, t_i)$, $j = 1, \dots, k$, with $\tau_i = t_i - t_{i-1}$ and $t_0 = 0$. It contains all the domain transformations that can be attained by an ODE by selecting a piecewise constant driving force from \mathcal{F} up to a terminal time T . The union of attainable sets over all possible terminal times, $\mathcal{A}_{\mathcal{F}} = \bigcup_{T>0} \mathcal{A}_{\mathcal{F}}(T)$, is the overall attainable set. In view of Corollary 3.9, to establish the approximation property of \mathcal{H}_{ode} it is sufficient to prove that any continuous transformation φ can be approximated by mappings in $\mathcal{A}_{\mathcal{F}}$. Note that $\mathcal{A}_{\mathcal{F}}(T)$ is a subset of $\Phi(\mathcal{F}, T)$ defined in (2.2).

Now, let us prove some properties of the approximation closure (i.e. closure with respect to the topology of compact convergence, see Section 2.1) of attainable sets. The first result below is immediate.

Lemma 3.10. *Let $n = 1$. If A is a family of continuous and increasing functions from \mathbb{R} to \mathbb{R} , then \overline{A} contains only increasing functions.*

Next, we prove an important property of approximation closures of control families: \mathcal{F} shares the same approximation ability as $\overline{\text{CH}}(\mathcal{F})$ when used to drive dynamical systems. A technicality is that the convex hull of a Lipschitz function family might not be a Lipschitz function family in general. Hence we adopt a slightly different description.

Proposition 3.11. *Let \mathcal{F} be a Lipschitz control family. Then, for any Lipschitz control family $\tilde{\mathcal{F}}$ such that $\mathcal{F} \subset \tilde{\mathcal{F}} \subset \overline{\text{CH}}(\mathcal{F})$, we have*

$$\overline{\mathcal{A}_{\mathcal{F}}} = \overline{\mathcal{A}_{\tilde{\mathcal{F}}}}. \tag{3.20}$$

Proposition 3.11 is an important result concerning the effect of continuous evolution, which can be regarded as a continuous family of compositions: any function family driving a dynamical system is as good as its convex hull in driving the system, which can be a significantly larger family of functions. Similar properties of flows have been observed in the context of variational problems [53]. This is a first hint at the power of composition in function approximation.

To prove Proposition 3.11 we need the following lemmas.

Lemma 3.12. *If $\mathcal{A}_{\mathcal{F}}$, $\mathcal{A}_{\tilde{\mathcal{F}}}$ are attainable sets of \mathcal{F} , $\tilde{\mathcal{F}}$ and $\mathcal{F} \subset \tilde{\mathcal{F}} \subset \overline{\mathcal{F}}$, then*

$$\overline{\mathcal{A}_{\mathcal{F}}} = \overline{\mathcal{A}_{\tilde{\mathcal{F}}}}. \tag{3.21}$$

Proof. It suffices to show that $\mathcal{A}_{\tilde{\mathcal{F}}} \subset \overline{\mathcal{A}_{\mathcal{F}}}$, which implies $\overline{\mathcal{A}_{\mathcal{F}}} \subset \overline{\mathcal{A}_{\tilde{\mathcal{F}}}} \subset \overline{\mathcal{A}_{\mathcal{F}}}$. Note that any $\tilde{\varphi} \in \mathcal{A}_{\tilde{\mathcal{F}}}$ is of the form

$$\tilde{\varphi} = \varphi_{t_k}^{\tilde{f}_k} \circ \varphi_{t_{k-1}}^{\tilde{f}_{k-1}} \circ \dots \circ \varphi_{t_1}^{\tilde{f}_1} \tag{3.22}$$

where each \tilde{f}_i is in $\tilde{\mathcal{F}}$. To prove the lemma, we have to show that for any compact $K \subset \mathbb{R}^n$ and any $\varepsilon > 0$, we can construct $\varphi \in \mathcal{A}_{\mathcal{F}}$ such that $\|\tilde{\varphi} - \varphi\|_{C(K)} \leq \varepsilon$. We prove this by induction on $k \geq 0$. First, the case $k = 0$ is obvious since it is just the identity mapping. Suppose now that the statement holds for $k - 1$ and fix any compact K . Write $\tilde{\varphi} = \varphi_{t_k}^{\tilde{f}_k} \circ \tilde{\psi}$, where $\tilde{\psi}$ is a composition of $k - 1$ flow maps driven by $\tilde{\mathcal{F}}$. By the inductive hypothesis, for any ε_1 (to be chosen later) there exists $\psi \in \mathcal{A}_{\mathcal{F}}$ such that $\|\tilde{\psi} - \psi\|_{C(K)} \leq \varepsilon_1$. Moreover, by the assumption $\tilde{\mathcal{F}} \subset \overline{\mathcal{F}}$, for any ε_2 and compact K' there is a function $f_k \in \mathcal{F}$ such that $\|f_k - \tilde{f}_k\|_{C(K')} \leq \varepsilon_2$. Here, we choose $K' = \{x : \inf_{y \in \psi(K)} \|x - y\| \leq 2(\varepsilon_1 + t_k)e^{t_k \text{Lip}(\tilde{f}_k)}\}$ and $\varepsilon_2 < 1$.

Now, consider two ODEs:

$$\tilde{z}(t) = \tilde{\psi}(x) + \int_0^t \tilde{f}_k(\tilde{z}(\tau)) d\tau \tag{3.23}$$

and

$$z(t) = \psi(x) + \int_0^t f_k(z(\tau)) d\tau. \tag{3.24}$$

We have $\tilde{\varphi}(x) = \tilde{z}(t_k)$ and $x \mapsto z(t_k)$ belongs to $\mathcal{A}_{\mathcal{F}}$, thus it remains to show that the solutions of these ODEs can be made arbitrarily close. In fact, we show that

$$\sup_{0 \leq t \leq t_k} |z(t) - \tilde{z}(t)| < 2(\varepsilon_1 + t_k \varepsilon_2)e^{t_k \text{Lip}(\tilde{f}_k)}. \tag{3.25}$$

Suppose not, then set

$$t = \inf \{s \geq 0 : |z(s) - \tilde{z}(s)| \geq 2(\varepsilon_1 + t_k \varepsilon_2)e^{t_k \text{Lip}(\tilde{f}_k)}\}. \tag{3.26}$$

By continuity we have $|z(t) - \tilde{z}(t)| \geq 2(\varepsilon_1 + t_k \varepsilon_2) e^{t_k \text{Lip}(\tilde{f}_k)}$. By assumption we also have $|\tilde{f}_k(z(\tau)) - f_k(z(\tau))| \leq \varepsilon_2$ for all $\tau < t$. By subtracting, we have

$$\begin{aligned} |\tilde{z}(t) - z(t)| &\leq \varepsilon_1 + \left| \int_0^t \tilde{f}_k(\tilde{z}(\tau)) d\tau - \int_0^t f_k(z(\tau)) d\tau \right| \\ &\leq \varepsilon_1 + \int_0^t |\tilde{f}_k(z(\tau)) - f_k(z(\tau))| d\tau + \int_0^t |\tilde{f}_k(\tilde{z}(\tau)) - \tilde{f}_k(z(\tau))| d\tau \\ &\leq \varepsilon_1 + \varepsilon_2 t + \text{Lip}(f_k) \int_0^t |\tilde{z}(\tau) - z(\tau)| d\tau. \end{aligned} \tag{3.27}$$

By Grönwall’s inequality, we have

$$\sup_{0 \leq t \leq t_k} |\tilde{z}(t) - z(t)| \leq (\varepsilon_1 + t_k \varepsilon_2) e^{t_k \text{Lip}(f_k)}, \tag{3.28}$$

which contradicts the choice of t . Hence (3.25) holds, and we can choose ε_1 and ε_2 arbitrarily small, which concludes the proof. ■

Lemma 3.13. *Suppose $f, g \in \mathcal{F}$ and $t > 0$. Then $\varphi_t^{(f+g)/2} \in \overline{\mathcal{A}_{\mathcal{F}}}$.*

Proof. We will show that $\varphi_{t/2N}^f \circ \varphi_{t/2N}^g \circ \dots \circ \varphi_{t/2N}^f \circ \varphi_{t/2N}^g$ can approximate $\varphi_t^{(f+g)/2} \in \overline{\mathcal{A}_{\mathcal{F}}}$ arbitrarily well by increasing N . The mapping $\varphi_t^{(f+g)/2}$ is the solution of

$$\begin{aligned} z(t) &= x + \int_0^t \left(\frac{f+g}{2} \right) (z(\tau)) d\tau \\ &= x + \left(\int_0^{t/2N} + \int_{2t/2N}^{3t/2N} + \dots + \int_{(2N-2)t/2N}^{(2N-1)t/2N} \right) (f+g)(z(\tau)) d\tau \\ &\quad + \left(\int_{t/2N}^{2t/2N} + \dots + \int_{(2N-1)t/2N}^t \right) \left[\left(\frac{f+g}{2} \right) (z(\tau)) - \left(\frac{f+g}{2} \right) \left(z\left(\tau - \frac{t}{2N}\right) \right) \right] d\tau \\ &= x + \left(\int_0^{t/2N} + \int_{2t/2N}^{3t/2N} + \dots + \int_{(2N-2)t/2N}^{(2N-1)t/2N} \right) f(z(\tau)) d\tau \\ &\quad + \left(\int_{t/2N}^{2t/2N} + \dots + \int_{(2N-1)t/2N}^t \right) g(z(\tau)) d\tau \\ &\quad + \left(\int_{t/2N}^{2t/2N} + \dots + \int_{(2N-1)t/2N}^t \right) \left[\left(\frac{f+g}{2} \right) (z(\tau)) - \left(\frac{f+g}{2} \right) \left(z\left(\tau - \frac{t}{2N}\right) \right) \right] \\ &\quad \quad \quad + \left[g\left(z\left(\tau - \frac{t}{2N}\right) \right) - g(z(\tau)) \right] d\tau. \end{aligned} \tag{3.29}$$

Thus if

$$\begin{aligned} w(t) &= x + \left(\int_0^{t/2N} + \int_{2t/2N}^{3t/2N} + \dots + \int_{(2N-2)t/2N}^{(2N-1)t/2N} \right) f(w(\tau)) d\tau \\ &\quad + \left(\int_{t/2N}^{2t/2N} + \dots + \int_{(2N-1)t/2N}^t \right) g(w(\tau)) d\tau, \end{aligned} \tag{3.30}$$

then

$$|z(t) - w(t)| \leq \int_0^t \max(\text{Lip}(f), \text{Lip}(g)) |z(\tau) - w(\tau)| d\tau + \frac{t}{2} \omega_{z, [0, t]} \left(\frac{t}{2N} \right) \left[\text{Lip} \left(\frac{f + g}{2} \right) + \text{Lip}(g) \right]. \tag{3.31}$$

Recall that ω is the modulus of continuity defined in Proposition 3.6. Again, by Grönwall’s inequality we have

$$|z(t) - w(t)| \leq \frac{t}{2} \omega_{z, [0, t]} \left(\frac{t}{2N} \right) \left[\text{Lip} \left(\frac{f + g}{2} \right) + \text{Lip}(g) \right] e^{\max(\text{Lip}(f), \text{Lip}(g))}. \tag{3.32}$$

For any selected compact set K , $\omega_{z, [0, t]} \left(\frac{t}{2n} \right) \rightarrow 0$ by Proposition 3.6, thus we obtain $\varphi_t^{(f+g)/2} \in \overline{\mathcal{A}_{\mathcal{F}}}$. ■

Now, we are ready to prove Proposition 3.11.

Proof of Proposition 3.11. Using the same technique as in the proof of Lemma 3.13, we can show that for $f_1, \dots, f_m \in \mathcal{F}$, we have $\varphi_t^h \in \overline{\mathcal{A}_{\mathcal{F}}}$, where $h = \sum_i q_i f_i$ for some rational numbers q_i . Let $\mathcal{A}_{\mathcal{F}'}$ be the attainable set with control family $\mathcal{F}' = \{\sum_{i=1}^m q_i f_i : q_i \in \mathbb{Q}, \sum_i q_i = 1, f_i \in \mathcal{F}, m \in \mathbb{N}\}$. Then $\overline{\mathcal{A}_{\mathcal{F}'}} = \overline{\mathcal{A}_{\mathcal{F}}}$. Since $\overline{\mathcal{F}'} = \overline{\mathcal{F}}$, we arrive at the desired result using Lemma 3.12. ■

4. Proof of main results

In this section, we prove the main results (Theorems 2.3 and 2.4). We start with the one-dimensional case to gain some insights on how a result can be established in general, and in particular, elucidate the role of well functions (Definition 2.1) in constructing rearrangement dynamics. This serves to motivate the extension of the results to higher dimensions.

4.1. Approximation results in one dimension and the proof of Theorem 2.4

We take $n = 1$ in this subsection. Proposition 3.4, together with the fact that compositions of continuous increasing functions are again continuous and increasing, implies that any function from $\mathcal{A}_{\mathcal{F}}$ must be continuous and increasing. This poses a restriction on the approximation power of $\overline{\mathcal{A}_{\mathcal{F}}}$ as the following result shows:

Proposition 4.1. *Let $n = 1$ and \mathcal{F} be a Lipschitz control family, whose attainable set is $\mathcal{A}_{\mathcal{F}}$. Then $\overline{\mathcal{A}_{\mathcal{F}}}$ contains only increasing functions.*

Proof. Proposition 3.4 implies any function in $\mathcal{A}_{\mathcal{F}}$ is continuous and increasing, since both properties are closed under composition. The proposition then follows from Lemma 3.10. ■

It follows from Proposition 4.1 that any continuous function φ that is strictly decreasing over an interval $[c, d]$ cannot be approximated by $\overline{\mathcal{A}_{\mathcal{F}}}$. Nevertheless, it makes sense to ask for the next best property: Can $\overline{\mathcal{A}_{\mathcal{F}}}$ approximate any continuous increasing function?

To investigate this problem, we first select an appropriate control family, which corresponds to deep neural networks with ReLU activations, and see if it can indeed approximate any such function. We will remove this explicit architectural assumption later. The ReLU control family in $n = 1$ is given by

$$\mathcal{F} = \{v\text{ReLU}(w \cdot +b) : v, w, b \in \mathbb{R}\}. \tag{4.1}$$

Notice that the ReLU control family (4.1) satisfies the restricted affine invariance condition of Definition 2.2.

We now show that in one dimension, flow maps of ODEs driven by the ReLU control family can in fact approximate any continuous increasing function.

Proposition 4.2. *Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be continuous and increasing and \mathcal{F} be the 1D ReLU control family (4.1). Then, for any $\varepsilon > 0$ and compact $K \subset \mathbb{R}$, there exists $\widehat{\varphi} \in \mathcal{A}_{\mathcal{F}}$ such that $\|\varphi - \widehat{\varphi}\|_{C(K)} \leq \varepsilon$. In other words, $\varphi \in \overline{\mathcal{A}_{\mathcal{F}}}$.*

Proof. We need the following lemma, from which we can deduce the desired result.

Lemma 4.3. *Let $M \geq 1$. Given $x_1 < \dots < x_M$ and $y_1 < \dots < y_M$, there exists a function $\psi \in \mathcal{A}_{\mathcal{F}}$ such that $\psi(x_i) = y_i$.*

We postpone the proof of Lemma 4.3 and first show how to deduce Proposition 4.2 from it. By replacing K with a larger set, we can always assume that K is a closed interval. Consider a partition Δ of K with nodes $x_1 < \dots < x_M$. By Lemma 4.3, we can find $\psi \in \mathcal{A}_{\mathcal{F}}$ such that $\psi(x_i) = \varphi(x_i)$ for all $i = 1, \dots, M$. Therefore

$$\psi(x) - \varphi(x) \leq \psi(x_{i+1}) - \varphi(x_i) \leq \varphi(x_{i+1}) - \varphi(x_i) \leq \omega_{\varphi}(|\Delta|) \tag{4.2}$$

whenever $x \in [x_i, x_{i+1}]$. Here $|\Delta| := \max_{1 \leq i \leq M} |x_i - x_{i-1}|$. We deduce that $\psi(x) - \varphi(x) \geq -\omega_{\varphi}(|\Delta|)$ for the same reason. Hence $\|\varphi - \psi\|_{C(K)} \leq \omega_{\varphi}(|\Delta|)$. Since φ is continuous, letting $|\Delta| \rightarrow 0$ and using Proposition 3.6 gives the desired result. ■

Now it remains to prove Lemma 4.3 constructively. To do this, first observe that the definition of well function (Definition 2.1) when specialized to one dimension is a function h_Q such that $h_Q(x) = 0$ if and only if $x \in Q = [q_1, q_2]$ for some $q_2 > q_1$. This can be constructed by the ReLU family (cf. Example 2.6) by

$$h_Q = \frac{1}{2}[\text{ReLU}(q_1 - x) + \text{ReLU}(x - q_2)]. \tag{4.3}$$

Obviously, $h_Q \in \text{CH}(\mathcal{F}) \subset \overline{\text{CH}(\mathcal{F})}$, so that the condition that the latter contains a well function is trivially satisfied.

Proof of Lemma 4.3. By Proposition 3.11, we denote $\widetilde{\mathcal{F}} = \mathcal{F} \cup \{h_Q : Q \subset K\}$. We will show that $\widetilde{\mathcal{F}}$ can produce the desired approximation property. We construct by induction a

mapping φ_k which maps x_i to y_i for $i = 1, \dots, k$. First we handle the base case $k = 1$. Take h_Q to be the well function with respect to $Q = [q_1, q_2]$. Since \mathcal{F} is translation invariant, we can suppose that both x_1 and y_1 are greater than q_2 . Since h_Q does not change sign in $[q_2, \infty)$, by Proposition 3.7 we know that either $\varphi_t^{h_Q}$ or $\varphi_t^{-h_Q}$ can map x_1 to y_1 for some t . Thus the base case is proved since \mathcal{F} is symmetric.

Suppose we have φ_k ; now we will construct φ_{k+1} based on φ_k . Applying φ_k , we may assume that $x_i = y_i, i = 1, \dots, k$. Again we assume h_Q is a well function with zero interval $Q = [q_1, q_2]$ ($q_2 < \min(x_1, y_1)$) and $h_{Q'}$ is a well function with $Q' = [q_0, q_1]$. We further assume that $h_{Q'}(x) < 0$ on $[q_1, \infty)$, $h_Q(x) < 0$ on $[q_2, \infty)$, otherwise we can use $-h_Q$ or $-h_{Q'}$ instead.

Let $t_1 = \inf \{t : \varphi_t^{h_{Q'}}(x_k) < q_2\}$ and $t_2 = \sup \{t : \varphi_t^{h_{Q'}}(x_{k+1}) > q_2, \varphi_t^{h_{Q'}}(y_{k+1}) > q_2\}$. Clearly $t_1 < t_2$. Choosing any $t' \in (t_1, t_2)$, and $\psi = \varphi_{t'}^{\pm h_Q}$ mapping $\varphi_{t'}^{h_{Q'}}(x_{k+1})$ to $\varphi_{t'}^{h_{Q'}}(y_{k+1})$, we construct

$$\varphi_{k+1} = \varphi_{t'}^{-h_{Q'}} \circ \psi \circ \varphi_{t'}^{h_{Q'}} \circ \varphi_k \tag{4.4}$$

as desired. By induction, we have completed the proof of Lemma 4.3. ■

Sufficient conditions for approximation of continuous and increasing functions and the proof of Theorem 2.4. We showed previously that all continuous increasing functions can be approximated by ReLU-driven dynamical systems. In this section, we shall do away with an explicit architecture, which leads to the proof of Theorem 2.4. The key observation from the proof of Lemma 4.3 is that all we really need is to have a well function contained in $\overline{\text{CH}}(\mathcal{F})$ that we can translate and change signs, which is achieved by the restricted affine invariance assumption. On the other hand, whether or not \mathcal{F} itself is a ReLU control family, or any other specific family, is irrelevant. This motivates us to pose the question of sufficiency: what assumptions on \mathcal{F} are enough to guarantee that it is a universal control family? Notice that instead of constructing an explicit well function in the form of the average of two ReLU functions, we can just use an arbitrary well function as defined in Definition 2.1 to drive the dynamics. The following result makes this precise.

Proposition 4.4. *Assume the control family \mathcal{F} is symmetric and translation invariant, which is equivalent to restricted affine invariance with $D = \pm 1$ and $A = 1$ in Definition 2.2, and that $\overline{\text{CH}}(\mathcal{F})$ contains a well function. Then the conclusion in Proposition 4.2 holds.*

Proof. The proof is almost identical to that of Proposition 4.2 with the well function constructed by averaging two ReLU functions replaced by a general well function contained in $\overline{\text{CH}}(\mathcal{F})$. Notice that since a well function does not change sign outside of I , by choosing a proper sign one can always shrink a finite point arbitrarily close to the interval. This follows from Proposition 3.7. ■

Proposition 4.4 combined with Corollary 3.9 implies Theorem 2.4. Clearly, Proposition 4.4 generalizes Proposition 4.2. It also follows that if $\overline{\text{CH}}(\mathcal{F})$ contains all continuous

functions, it must in particular contain a well function and so $\mathcal{A}_{\mathcal{F}}$ has the desired approximation property. However, this is not necessary for universal approximation to hold.

Remark 4.5. In one dimension, the ability for a dynamical system to approximate any continuous increasing function has the immediate consequence that if we were to embed the dynamical system in two dimensions, then we can approximate any continuous function φ of bounded variation, as long as we are allowed a linear transformation in the end, e.g. if g in Proposition 3.8 is linear. This is because a continuous function of bounded variation can always be written as a difference of two continuous increasing functions. However, this does require embedding in high dimensions. We will show later that for $n \geq 2$, embedding is not necessary to achieve universal approximation.

4.2. Approximation rates in one dimension

All results so far are on whether a given function can be approximated by a dynamical system with control families satisfying certain conditions. However, by the very definition of the attainable set we only considered dynamical systems of finite, but arbitrarily large time horizons. Just as in the development of traditional approximation theory, one may be interested in the following: given an approximation budget, how well can we approximate a given function? Perhaps a more pertinent question is this: what kind of functions can be efficiently approximated by dynamical systems? There are more than one way to define the notion of budget. Here, we will consider a natural one in continuous time: the time horizon T .

In this part, we give some results in this direction in the simplest case: the one-dimensional case ($n = 1$) and the fully connected ReLU control family. For convenience of exposition, we assume that our target function φ is defined on $[0, 1]$. We postpone results on general control families in higher dimensions to future work.

To properly quantify the approximation rate, we should eliminate the positive homogeneity of the ReLU control function, which masks the effect of the time horizon T due to the ability to arbitrarily rescale time. Therefore, we restrict $|v|, |w| \leq 1$ in $v\text{ReLU}(w \cdot + b)$ so that the time horizon becomes meaningful.

Remark 4.6. An alternative is using $\int |w| dt$ to measure the approximation cost in place of T . This notation is related to the Barron space analysis [16]. It can be checked that one can change T into $\int |w| dt$ in the following results. Lastly, it is also possible to measure the complexity of the variation of w, v, b in time. Here, we do not consider such cases.

First, we show in the following lemma that if φ is piecewise linear, then it can be represented by functions in $\mathcal{A}_{\mathcal{F}}(T)$ for some T large enough.

Before introducing the following lemma, we first define the total variation (TV) with a slight modification. Suppose u is a function defined on $[0, 1]$, we extend u to u_E such that $u_E = 0$ in $[-\varepsilon, 0) \cup (1, 1 + \varepsilon]$. We define

$$\|u\|_{\text{TV}} = \|u_E\|_{\text{TV}[-\varepsilon, 1+\varepsilon]},$$

the latter norm being defined as

$$\|f\|_{\text{TV}[-\varepsilon, 1+\varepsilon]} = \sup_{-\varepsilon=x_0 < \dots < x_M=1+\varepsilon} \sum_{i=1}^M |f(x_i) - f(x_{i-1})|. \tag{4.5}$$

Lemma 4.7. *If $\ln \varphi'$ is a piecewise constant function with N pieces, then φ can be written as*

$$\varphi = \varphi_{t_1}^{g_1} \circ \dots \circ \varphi_{t_{N-1}}^{g_{N-1}} + c. \tag{4.6}$$

Here, $g_i \in \mathcal{F}$, φ_i denotes the flow maps as defined in Section 3.2 and c is a constant. Moreover, $\varphi \in \mathcal{A}_{\mathcal{F}}(T)$ for $T \geq \|\ln \varphi'\|_{\text{TV}}$.

Proof. We first show an auxiliary result, suppose that

$$\varphi = \varphi_{t_1}^{g_1} \circ \dots \circ \varphi_{t_{N-1}}^{g_{N-1}} + c. \tag{4.7}$$

Then $\ln \varphi'$ can be written as the sum of $N - 1$ Heaviside functions. The proof is by induction. For the $N = 1$ case it can be checked by direct calculation, and suppose that $\varphi = \varphi_{t_1}^{g_1} \circ \varphi_2$ and we have

$$\varphi'(x) = (\varphi_{t_1}^{g_1})'(\varphi_2(x))\varphi_2'(x) \tag{4.8}$$

by the chain rule. Thus

$$\ln \varphi'(x) = \ln (\varphi_{t_1}^{g_1})'(\varphi_2(x)) + \ln \varphi_2'(x) \tag{4.9}$$

and the first term on the RHS is a Heaviside function, while the second term is a sum of $N - 2$ Heaviside functions by the induction hypothesis. Hence the result is proved by induction.

The proof of the original proposition is inductive by construction. Taking the derivative of $\varphi = \varphi_{t_1}^{g_1} \circ \dots \circ \varphi_{t_{N-1}}^{g_{N-1}} + c$, we obtain $\ln \varphi' = \ln (\varphi_{t_1}^{g_1})' + \dots + \ln (\varphi_{t_{N-1}}^{g_{N-1}})'$. Since $\ln \varphi'$ is a piecewise constant function, it can be written as

$$\ln \varphi' = H_1 + \dots + H_{N-1}, \tag{4.10}$$

where all H_i are Heaviside functions. Integrating $(\varphi_{t_{N-1}}^{g_{N-1}})'(x) = H_{N-1}(x)$ we obtain $\varphi_{t_{N-1}}^{g_{N-1}}$, and then integrating

$$(\varphi_{t_{N-2}}^{g_{N-2}})'(x) = H_{N-2}(\varphi_{t_{N-1}}^{g_{N-1}})^{-1}(x) \tag{4.11}$$

we obtain $\varphi_{t_{N-2}}^{g_{N-2}}$, and so on. One can easily verify that each φ_i is a flow map generated by some ReLU activation function. This yields the first part of the proposition.

For the second part, we notice that if $f = \text{ReLU}(wx + b)$, then $\ln (\varphi_t^f)'$ is a Heaviside function with a jump at $x = -b/w$. Since we can find a decomposition $\sum_i H_i$ such that $\sum_i |w_i| = \|\ln \varphi'\|_{\text{TV}}$, the second part of the proposition is proven. ■

From the proof of the previous lemma, we know that $T = \|\ln \varphi'\|_{\text{TV}}$ is optimal, since $\sum_i |w_i| \leq \|\ln \varphi'\|_{\text{TV}}$ ($\|\cdot\|_{\text{TV}}$ is a seminorm). In view of this lemma, we prove the following, which gives a quantitative approximation result.

Proposition 4.8. *Suppose $\varphi : [0, 1] \rightarrow \mathbb{R}$ is a piecewise smooth, increasing function with $T_0 := \|\ln \varphi'\|_{\text{TV}([0,1])} < \infty$. Then $\varphi \in \overline{\mathcal{A}_{\mathcal{F}}(T)}$ for $T \geq T_0$.*

Proof. The proof is separated into two parts. The first part is to show that the constant in Lemma 4.7 has negligible cost, by considering $\varphi_t^{\varepsilon \text{ReLU}(+M)} \varphi_t^{-\varepsilon \text{ReLU}(\cdot)}$. This provides a translation on $[0, 1]$: $x \mapsto x + (e^\varepsilon - 1)M$. By letting $M \rightarrow \infty$ we can construct any translation with negligible time cost.

The second part is that if $|\ln \varphi'(x) - \ln \psi'(x)| \leq \varepsilon$, and $\varphi(0) = \psi(0)$, then we have $|\varphi(x) - \psi(x)| \leq (e^\varepsilon - 1)\|\varphi'\|_{C[0,1]}$.

Now it suffices to prove a rather simple result: given a function $u = \ln \varphi'$, on each piece I' of u we can use a piecewise constant function $v|_{I'}$ to approximate $u|_{I'}$ (restriction to I') so that $\|v\|_{\text{TV}} \leq \|u\|_{\text{TV}}$ and $\|v - u\|_\infty \leq \varepsilon$. Thus we can find a function $\varphi \in \mathcal{A}_{\mathcal{F}}((1 - \varepsilon)T_0)$ such that $\ln \varphi' = (1 - \varepsilon)v$. By composing with a translation, there exists $\psi \in \mathcal{A}_{\mathcal{F}}(T_0)$ such that $\|\psi - \varphi\|_{C[0,1]} \leq \exp(\varepsilon\|\varphi\|_{C[0,1]} + \varepsilon)$. Thus we conclude that $\varphi \in \overline{\mathcal{A}_{\mathcal{F}}(T_0)}$. ■

The preceding results show that it is possible to constrain the approximation space to flow maps with time horizon up to some finite T_0 , provided that the target function φ is such that $\|\ln \varphi'\| \leq T$. In this sense, the total variation of the logarithm of φ is a measure of complexity under our compositional approximation procedure.

Let us now develop the quantitative results a little further for the case where T is not sufficiently large, i.e. $T < \|\ln \varphi'\|_{\text{TV}}$. This involves analyzing the error

$$E_T(\varphi) = \inf_{\psi \in \mathcal{A}_{\mathcal{F}}(T)} \|\varphi - \psi\|_{C[0,1]}, \tag{4.12}$$

which may be non-zero when $T < \|\ln \varphi'\|_{\text{TV}}$.

Proposition 4.9. *$E_T(\varphi)$ is given by the following optimization problem:*

$$\inf_{\psi} \{\|\varphi - \psi\|_{C[0,1]} : \|\ln \psi'\|_{\text{TV}} \leq T\}. \tag{4.13}$$

Notice that the existence of $\ln \psi'$ forces ψ to be a continuous increasing function, so we may consider the above optimization problem only for the case where ψ is continuous and increasing.

It is generally hard to work with optimization problems such as (4.13), since it involves total variations of logarithms of functions. Below, we formulate its relaxed version.

Proposition 4.10. *Consider the relaxed optimization problem*

$$\gamma(u, T) = \inf_v \{\|u - v\|_{C[0,1]} : \|v\|_{\text{TV}} \leq T\}. \tag{4.14}$$

Then $E_T(\varphi) \leq [\exp(\gamma(\ln \varphi', T)) - 1]\|\varphi'\|_{C[0,1]}$.

Proof. Choose v such that $\|v\|_{\text{TV}} \leq T$ and

$$\|\ln \varphi' - v\|_{C[0,1]} \leq \gamma(\ln \varphi', T) + \varepsilon. \tag{4.15}$$

Choose ψ such that $\ln \psi' = v$ and $\psi(0) = \varphi(0)$. Then since $\|\ln \frac{\psi'}{\varphi'}\|_{\text{TV}} \leq \gamma(\ln \varphi', T) + \varepsilon$, we have

$$|\psi'/\varphi' - 1| \leq \exp(\gamma(\ln \varphi', T) + \varepsilon) - 1.$$

Hence

$$|\varphi(x) - \psi(x)| \leq \int_0^x |\varphi'(x)| \left| 1 - \frac{\psi'(x)}{\varphi'(x)} \right| \leq \|\varphi'\|_{C[0,1]} [\exp(\gamma(\ln \varphi', T)\varepsilon) - 1]. \tag{4.16}$$

Letting $\varepsilon \rightarrow 0$, we arrive at the result. ■

In general, both (4.13) and (4.14) are hard to solve. However, for some simple cases of u , the problem (4.14) has an explicit solution. For example, if u itself is an increasing function, then the solution of (4.14) is $\frac{1}{2}(\|u\|_{\text{TV}} - T)$. If u is increasing in $[0, s]$ and decreasing in $[s, 1]$, then the solution of (4.14) is $\frac{1}{4}(\|u\|_{\text{TV}} - T)$. This gives approximation rates for specific cases, but a general investigation of these approximation rates is postponed to future work.

4.3. Approximation results in higher dimensions and proof of Theorem 2.3

In this section, we will generalize the universal approximation results to higher dimensions. The interesting finding is that in higher dimensions, the fact that $\mathcal{A}_{\mathcal{F}}$ contains only OP homeomorphisms no longer poses a restriction on approximations in the L^p sense. Moreover, the sufficient condition for universal approximation in higher dimensions is closely related to that in one dimension, where the rearrangement dynamics are driven by well functions. We will prove the following result, which together with Corollary 3.9 implies Theorem 2.3.

Proposition 4.11. *Let $n \geq 2$. Suppose \mathcal{F} is restricted affine invariant and $\overline{\text{CH}}(\mathcal{F})$ contains a well function. Then for any compact set K , $p \in [1, \infty)$, $\varphi \in C(\mathbb{R}^n)$, $\varepsilon > 0$, there exists a mapping $\tilde{\varphi} \in \mathcal{A}_{\mathcal{F}}$ such that $\|\tilde{\varphi} - \varphi\|_{L^p(K)} \leq \varepsilon$.*

We notice that for the purpose of approximation, the fact that $\overline{\text{CH}}(\mathcal{F})$ contains a well function h allows us to assume without loss of generality that \mathcal{F} contains a well function. To see this, denote by $\tilde{\mathcal{F}}$ the smallest restricted affine invariant set containing $\mathcal{F} \cup \{h\}$. We have $\mathcal{F} \subset \tilde{\mathcal{F}} \subset \overline{\text{CH}}(\mathcal{F})$. Proposition 3.11 then says that $\overline{\mathcal{A}_{\mathcal{F}}} = \overline{\mathcal{A}_{\tilde{\mathcal{F}}}}$, hence we can prove approximation results using $\tilde{\mathcal{F}}$ in place of \mathcal{F} .

4.3.1. Preliminaries. In order to prove Proposition 4.11, we require a few preliminary results. The key approach is similar to the one-dimensional case: we show that we can transform a finite number of distinct source points into a finite number of target points, which are not necessarily distinct. More precisely, we show the following proposition, which generalizes Lemma 4.3 given in the one-dimensional case.

Lemma 4.12. *Suppose \mathcal{F} contains a well function. Let $\varepsilon > 0$ and $x^1, \dots, x^m, y^1, \dots, y^m \in \mathbb{R}^n$ be such that $\{x^k\}$ are distinct points. Then there exists $\psi \in \mathcal{A}_{\mathcal{F}}$ such that $|\psi(x^k) - y^k| \leq \varepsilon$ for all $k = 1, \dots, m$.*

Lemma 4.12 follows from the combination of the following two lemmas.

Lemma 4.13. *Suppose \mathcal{F} contains a well function and x^1, \dots, x^m are distinct points. Then given any $\varepsilon > 0$, there exists a flow map $\psi \in \mathcal{A}_{\mathcal{F}}$ such that $|\psi(x^k) - x^k| \leq \varepsilon$ and for each $i = 1, \dots, n$, $[\psi(x^k)]_i$ (the i -th coordinate of $\psi(x^k)$), $k = 1, \dots, m$, are m distinct real numbers.*

Lemma 4.14. *Suppose \mathcal{F} contains a well function, x^1, \dots, x^m are distinct points and satisfy the result of Lemma 4.13, that is, $\{x_i^k\}$ are m distinct real numbers for any i . Then given any $\varepsilon > 0$ for m target points y^1, \dots, y^m , we have a flow map $\psi \in \mathcal{A}_{\mathcal{F}}$ such that $|\psi(x^k) - y^k| \leq \varepsilon$.*

Now we prove these two lemmas.

Proof of Lemma 4.13. It is enough to show that if there are two points x^j and x^k such that $x_I^j = x_I^k$ for some I , then we can find a flow map $\eta \in \mathcal{A}_{\mathcal{F}}$ such that η can separate x_I^j and x_I^k and at the same time, does not cause other pairs of points without initially distinct coordinates to overlap. Without loss of generality, we assume $j = 1, k = 2, I = 1$, and we only need to show that if $x_1^1 = x_1^2$, then there exists an $\eta \in \mathcal{A}_{\mathcal{F}}$ such that

- (1) $|\eta(x^k) - x^k| \leq \varepsilon_1 := \frac{1}{nm^2} \varepsilon$;
- (2) $[\eta(x^1)]_1 \neq [\eta(x^2)]_1$;
- (3) if $x_1^k \neq x_1^l$, then $[\eta(x^k)]_1 \neq [\eta(x^l)]_1$.

We briefly explain these requirements. Consider

$$X_1 = \{(k, l) : 1 \leq k < l \leq m, x_1^k = x_1^l\} \quad (4.17)$$

and

$$\eta(X_1) = \{(k, l) : 1 \leq k < l \leq m, [\eta(x^k)]_1 = [\eta(x^l)]_1\}. \quad (4.18)$$

(2) and (3) imply that $|X_1| > |\eta(X_1)| \geq 0$, hence $|X_1|$ is strictly decreasing after η .

Denote $d = \min\{|x_1^k - x_1^l| : x_1^k \neq x_1^l\}$. Since x^1 and x^2 are distinct points, we can find a coordinate index I ($\neq 1$) such that $x_I^1 \neq x_I^2$. Here we assume $x_I^1 < x_I^2$. Suppose f in \mathcal{F} is a well function with zero set Ω_1 . Written in coordinate form, f is given by

$$f = (f_1, \dots, f_n), \quad (4.19)$$

where $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ for each i . Since \mathcal{F} is translation-invariant, we can assume Ω_1 contains 0 without loss of generality.

Consider the dynamics

$$\dot{z}_1 = f_1(x_I + b), \quad \dot{z}_i = 0, \quad i = 2, \dots, n. \quad (4.20)$$

Notice that the boundedness and convexity of Ω_1 guarantees that the reduced 1D dynamics satisfies our previous discussion (it contains a 1D well function, since the intersection of a bounded convex set with a line is an interval). In other words, we choose $D = \text{diag}(1, 0, \dots, 0)$, $A_{ij} = \delta_{iI}\delta_{jI}$, $b = (0, 0, \dots, b_I, 0)^T$ in the form $\tilde{f} = Df(A \cdot + b)$ where b_I is chosen such that $x_I^1 + b_I \in \Omega_1$ but $x_I^2 + b_I \notin \overline{\Omega_1}$. The existence of b_I is implied by the boundedness of Ω_1 . We denote by P_t the flow map of this dynamics. We next choose a proper t such that (1)–(3) are satisfied.

Since $\tilde{f}_1(x^1) = 0$ and $\tilde{f}_1(x^2) \neq 0$, we deduce that $[P_t(x^1)]_1 \neq [P_t(x^2)]_1$ whenever $t \neq 0$. Hence (2) is satisfied with no additional condition. Notice that when $|P_t(x^k) - x^k| \leq \min(\varepsilon_1, d/3)$, then both (1) and (3) are satisfied. Since $\text{CH}(\{x^k\})$ is bounded, we have $\|P_t - \text{id}\|_{C(\text{CH}(\{x^k\}))} \rightarrow 0$ as $t \rightarrow 0$ by Proposition 3.6. Therefore there exists $t_0 > 0$ such that $\|P_t - \text{id}\|_{C(\text{CH}(\{x^k\}))} \leq \min(\varepsilon_1, d/3)$. Hence we conclude that $\eta = P_{t_0}$ satisfies (1)–(3). ■

Proof of Lemma 4.14. Without loss of generality, we can assume that for each i , y_i^k are m distinct real numbers, since if not, we can always add a small perturbation to the point set $\{y_i^k : i = 1, \dots, m\}$ directly and this will not affect approximation. We also assume Ω_1 contains the origin, as we did in the proof of Lemma 4.13.

The basic idea is similar to Lemma 4.13: by choosing a proper linear transformation we can freeze some point while transporting other points. Since we need to control more than two points, we can take multiple transformations and evolve them sequentially. We only need to prove that for any i (without loss of generality $i = 1$), we can find an $\eta \in \mathcal{A}_{\mathcal{F}}$ such that $[\eta(x^k)]_1 = y_1^k$.

Upon relabelling, we can assume that $x_2^1 < x_2^2 < \dots < x_2^m$. Consider the dynamics

$$\dot{z}_1 = f_1(az_2 + b), \quad \dot{z}_i = 0, \quad i = 2, \dots, n. \tag{4.21}$$

In other words, we choose $D = \text{diag}(1, 0, \dots, 0)$, $A_{ij} = a\delta_{i2}\delta_{j2}$, $b = (0, b_2, \dots, 0)^T$ in the form $\tilde{f} = Df(A \cdot + b)$, where a is chosen sufficiently small such that all Ax^k lies in Ω_1 . We denote the flow map by $P_t(b_2)$, where the dependence on b_2 is emphasized. To simplify our notation, we use $P_{-t}(b_2)$ to denote the flow map of $\dot{z} = -Df(Az + b)$.

Now we wish to choose r^1, \dots, r^m such that $f_1(Ax^i + r^j) = 0$ if and only if $i < j$. Let $\{0\} \times (u_l, u_r) \times \dots = \Omega_1 \cap (\{0\} \times \mathbb{R} \times \dots)$, where (u_l, u_r) is the restriction of Ω_1 on coordinate 2. Then a choice of r^k is $r^k = u_r - ax_2^k + \frac{a}{2} \min_j (x_2^j - x_2^{j-1})$.

Now $\{\eta^{(k)}\}$ are defined recursively:

$$\begin{aligned} \eta^{(0)} &= \text{id}, \\ \eta^{(k)} &= P_{t^k}(r^k) \circ \eta^{(k-1)}, \quad \text{where } t^k = (y_1^k - [\eta^{(k-1)}(x^k)]_1) / f_1(Ax^k + r^k). \end{aligned} \tag{4.22}$$

We now prove that $\eta = \eta^{(m)} \circ \dots \circ \eta^{(1)}$ satisfies our requirement. By induction (on k), we show that $[\eta^{(k)}(x^i)]_1 = y_1^i$ for $i \leq k$. The case $k = 0$ is vacuous. Supposing $[\eta^{(k-1)}(x^i)]_1 = y_1^i$ for $i \leq k - 1$, since $[\eta^{(k-1)}(x^i)]_2 = x_2^i$, we have

$$\eta^{(k)}(x^i) = P_{t^k}(r^k)(\eta^{(k-1)}(x^i)) = \eta^{(k-1)}(x^i) = y_i. \tag{4.23}$$

By definition we know that $\eta^{(k)}(x^k) = y^k$. Hence the induction step is proved. ■

4.3.2. Proof of Proposition 4.11

Since K is compact, by extension it suffices to consider the case that K is a hyper-cube. We can for simplicity take the unit hyper-cube $K = [0, 1]^n$, as the general case is similar. For $\varphi \in L^p(K)$, by standard approximation theory φ can be approximated by piecewise constant functions, i.e. there exists

$$\hat{\varphi} = \sum_{\mathbf{i}} \varphi_{\mathbf{i}} \chi_{\mathbf{i}} \quad (4.24)$$

such that $\|\hat{\varphi} - \varphi\|_{L^p(K)} \leq \varepsilon/2$. Here $\mathbf{i} = [i_1, \dots, i_n]$ is a multi-index, $\varphi_{\mathbf{i}} \in \mathbb{R}^n$ and $\chi_{\mathbf{i}}$ is the indicator of the cube

$$\square_{\mathbf{i}} = \left[\frac{i_1}{N}, \frac{i_1 + 1}{N} \right] \times \dots \times \left[\frac{i_n}{N}, \frac{i_n + 1}{N} \right]. \quad (4.25)$$

We also denote $p_{\mathbf{i}} = (i_1/N, \dots, i_n/N)$. We define a shrunken cube

$$\square_{\mathbf{i}}^{\alpha} = \left[\frac{i_1}{N}, \frac{i_1 + \alpha}{N} \right] \times \dots \times \left[\frac{i_n}{N}, \frac{i_n + \alpha}{N} \right], \quad (4.26)$$

where $0 < \alpha \leq 1$. We have $K = \bigcup_{\mathbf{i}} \square_{\mathbf{i}}$, and we define $K^{\alpha} = \bigcup_{\mathbf{i}} \square_{\mathbf{i}}^{\alpha}$. We also construct a shrinking function in one dimension, $h^{\alpha} : [0, 1] \rightarrow [0, 1]$, such that $h^{\alpha}(x) = i/N$ if $i/N \leq x \leq (i + \alpha)/N$, and continuously increasing in $[0, 1]$. Using this, we can form an n -dimensional shrinking map by tensor product:

$$H^{\alpha}(x) = (h^{\alpha}(x_1), \dots, h^{\alpha}(x_n)). \quad (4.27)$$

The idea of the proof of Proposition 4.11 is quite simple: we just contract each grid $\square_{\mathbf{i}}$ into a point $p_{\mathbf{i}}$ approximately, then use the lemma above to transform each $p_{\mathbf{i}}$ into $\varphi_{\mathbf{i}}$. The latter is discussed in a preliminary step; we construct an ‘‘almost’’ contraction mapping in $\mathcal{A}_{\mathcal{F}}$ that approximates H^{α} .

Claim. *For a given tolerance $\varepsilon_1 > 0$, there exists a flow map $\tilde{H} \in \mathcal{A}_{\mathcal{F}}$ such that $\|\tilde{H} - H^{\alpha}\|_{C(K)} \leq \varepsilon_1$.*

Proof of the Claim. Since h is increasing and continuous, we wish to utilize our result in one dimension. Concretely, we demonstrate how to restrict the n -dimensional control family to one dimension.

Suppose \mathcal{F} is an n -dimensional control family. Then for each $f = (f_1, \dots, f_n) \in \mathcal{F}$ we define the dynamics driven by its restriction to first coordinate by

$$\dot{z}_1 = f_1(x_1), \quad \dot{z}_i = 0 \text{ for } i \geq 2, \quad (4.28)$$

i.e., take $D = A = \text{diag}(1, 0, \dots, 0)$. The set of such control systems is denoted by $\mathcal{F}_{R,1}$ (R means restriction and 1 means first coordinate). Clearly $\mathcal{F}_{R,1}$ is closed under composition. Moreover, $\mathcal{A}_{\mathcal{F}_{R,1}}$ coincides with

$$\mathcal{A}_{\mathcal{R}} \times \{\text{id}\} \times \{\text{id}\} \times \dots \times \{\text{id}\}, \quad (4.29)$$

where \mathcal{R} is a one-dimensional control family

$$\mathcal{R} = \{g(x) : g(x) = f_1(x, 0, \dots, 0), f \in \mathcal{F}\}. \tag{4.30}$$

Since \mathcal{F} contains a well function, so does \mathcal{R} . By Proposition 4.4 we can find $\tilde{h} \in \mathcal{R}$ such that $\|\tilde{h} - h^\alpha\|_{C([0,1])} \leq \varepsilon/n$.

By composition we know that $\tilde{H} := (\tilde{h}, \dots, \tilde{h})$ is in $\mathcal{A}_{\mathcal{F}}$, and $\|\tilde{H} - H^\alpha\|_{C(K)} \leq \varepsilon$. ■

We use the above notations, ψ for transporting p_i to φ_i , and \tilde{H} for the approximate contraction mapping, satisfying the following estimates:

$$|\psi(p_i) - \varphi_i| \leq \varepsilon_1 < 1, \tag{4.31}$$

$$\|\tilde{H} - H^\alpha\|_{C(K)} \leq \varepsilon_2 < 1, \tag{4.32}$$

where ε_1 and ε_2 are to be determined later.

Now we estimate the error $\|\psi \circ \tilde{H} - \varphi\|_{L^p(K)}$. For any α we can write

$$\begin{aligned} \|\psi \circ \tilde{H} - \varphi\|_{L^p(K)} &\leq \|\psi \circ \tilde{H} - \hat{\varphi}\|_{L^p(K)} + \|\hat{\varphi} - \varphi\|_{L^p(K)} \\ &\leq \|\psi \circ \tilde{H} - \hat{\varphi}\|_{L^p(K^\alpha)} + \|\psi \circ \tilde{H} - \hat{\varphi}\|_{L^p(K \setminus K^\alpha)} + \varepsilon/2 \\ &\leq J_1 + J_2 + \varepsilon/2. \end{aligned} \tag{4.33}$$

Estimation of J_1 . We have

$$\begin{aligned} J_1 &= \|\psi \circ \tilde{H} - \psi \circ H^\alpha\|_{L^p(K^\alpha)} + \|\psi \circ H^\alpha - \hat{\varphi}\|_{L^p(K^\alpha)} \\ &\leq \omega_\psi(\|\tilde{H} - H^\alpha\|_{C(K)}) + \sum_i |\psi(p_i) - \varphi_i| \cdot |\square_i|^{1/p} \\ &\leq \omega_\psi(\|\tilde{H} - H^\alpha\|_{C(K)}) + N^{n-n/p} \varepsilon_1 \leq \omega_\psi(\varepsilon_2) + N^{n-n/p} \varepsilon_1. \end{aligned} \tag{4.34}$$

Estimation of J_2 . Denote by $\tilde{K} = [-1, 2]^n$ an enlarged cube. We have

$$J_2 \leq |K \setminus K^\alpha| \cdot (\text{diam}(\psi(\tilde{K})) + \|\hat{\varphi}\|_{L^\infty(K)}). \tag{4.35}$$

We choose ψ such that $\varepsilon_1 \leq \varepsilon/8$, α such that

$$|K \setminus K^\alpha| \leq (\text{diam}(\psi(\tilde{K})) + \|\hat{\varphi}\|_{L^\infty(K)})^{-1} \varepsilon/4, \tag{4.36}$$

and finally \tilde{H} such that $\omega_\psi(\varepsilon_2) \leq \varepsilon/8$. Then $\|\psi \circ \tilde{H} - \varphi\|_{L^p(K)} \leq \varepsilon$, and we take $\tilde{\varphi} = \psi \circ \tilde{H} \in \mathcal{A}_{\mathcal{F}}$, yielding the result.

As in the 1D case, Proposition 4.11 together with Corollary 3.9 implies Theorem 2.3.

4.4. Approximation results in tensor-product type dynamical systems

Sometimes, we are interested in control families generated by tensor products. Such control families have the advantage that they can be parameterized by scalar functions of one variable, hence may enjoy greater flexibility in analysis and practice. In this last section, we

give some results that apply specifically to tensor product control families. Let us denote

$$\Pi\mathcal{F} = \{f(x) = (g(x_1), g(x_2), \dots, g(x_n)) : g \in \mathcal{F}\}, \tag{4.37}$$

where \mathcal{F} is a one-dimensional control family.

As in higher dimensions, we may wish to consider the n -dimensional control family $\mathcal{F}(n)$, which is the smallest set containing $\Pi\mathcal{F}$ that is also invariant under the transformation $f(\cdot) \mapsto Df(A \cdot + b)$, where D, A are diagonal matrices. However, all functions ψ in $\mathcal{A}_{\mathcal{F}}$ are separable: $\psi = (\psi_1(x_1), \dots, \psi_n(x_n))$. Moreover, we can deduce from the 1D results that ψ_i is continuous and increasing. Clearly, this $\mathcal{F}(n)$ has limited approximation ability. Instead, we will relax the requirement that A is diagonal so that it can be any matrix, leading to a stronger version of the restricted invariance requirement in Theorem 2.3. This then leads to the following approximation result:

Proposition 4.15. *Suppose that*

- (1) $\mathcal{F}(n)$ contains $\Pi\mathcal{F}$;
- (2) $\mathcal{F}(n)$ is invariant under $f(\cdot) \mapsto Df(A \cdot + b)$, where D is any diagonal matrix, A is any matrix, and b is any \mathbb{R}^d vector;
- (3) $\overline{\mathcal{A}_{\mathcal{F}}}$ contains all continuous increasing functions from \mathbb{R} to \mathbb{R} .

Then for any $\varepsilon > 0$, $p \in (1, \infty]$, compact set K , and $\varphi \in C(\mathbb{R}^n)$, there is $\tilde{\varphi} \in \mathcal{A}_{\mathcal{F}(n)}$ such that $\|\varphi - \tilde{\varphi}\|_{L^p(K)} \leq \varepsilon$.

Remark 4.16. This result is not a corollary of Proposition 4.11, even if $\overline{\text{CH}}(\mathcal{F})$ contains a well function, since in this case the zero set of the tensor product of the well function may be unbounded.

Similar to estimation of Proposition 4.11, we omit the main body of this proof but only restate preparations about ψ and \tilde{H} .

Estimations on \tilde{H}

Lemma 4.17. *Suppose $\mathcal{F}(n)$ and \mathcal{F} satisfy the conditions in Proposition 4.15. For a given tolerance $\varepsilon_1 > 0$, there exists a flow map $\tilde{H} \in \mathcal{A}_{\mathcal{F}(n)}$ such that $\|\tilde{H} - H^\alpha\|_{C(K)} \leq \varepsilon_1$.*

Proof. This is straightforward from the definition of the tensor product control family and $\overline{\mathcal{A}_{\mathcal{F}}}$ contains all continuous increasing functions. ■

Estimations on ψ . Before the main estimate, we first show a useful lemma.

Lemma 4.18. *If $g \in \mathcal{A}_{\mathcal{F}}$, then $(x_1, x_2, \dots, x_n) \mapsto (x_1 + g(x_2), x_2, \dots, x_n)$ is in $\mathcal{A}_{\mathcal{F}(n)}$.*

Proof. We divide the construction into two parts. First, set $D = \text{diag}(1, 1, 0, \dots, 0)$ and $A_{22} = A_{12} = 1$ and $A_{ij} = 0$ otherwise. We first only look at the second coordinate, knowing that for all $g \in \mathcal{A}_{\mathcal{F}}$, there exists a finite number $S \geq 1$ of flow maps $\varphi_{f_s}^{t_s}$, $s = 1, \dots, S$, such that

$$g = \varphi_{f_S}^{t_S} \circ \dots \circ \varphi_{f_1}^{t_1}. \tag{4.38}$$

If D and A are chosen as above, we deduce that $x_1 - x_2$ is constant under all mappings of the form

$$z \mapsto Df(Az + b). \tag{4.39}$$

If we select f_1, \dots, f_S sending x_2 to $g(x_2)$, we know $x_1 \mapsto x_1 + g(x_2) - x_2$. Hence

$$(x_1, x_2, \dots, x_n) \mapsto (x_1 + g(x_2), x_2, \dots, x_n) \tag{4.40}$$

is in $\mathcal{A}_{\mathcal{F}(n)}$.

Also, by setting $D = \text{diag}(1, 0, \dots, 0)$ and $A_{ij} = \delta_{i2}\delta_{j2}$, we know that

$$(x_1, x_2, \dots, x_n) \mapsto (x_1, g^{-1}(x_2), \dots, x_n) \tag{4.41}$$

is in $\mathcal{A}_{\mathcal{F}(n)}$. Combining the two parts yields the result. ■

Lemma 4.19 (Analogous to Lemma 4.13). *Suppose $\mathcal{F}(n)$ and \mathcal{F} satisfy the conditions in Proposition 4.15 and x^1, \dots, x^m are distinct points. Then, given any $\varepsilon > 0$, there exists a flow map $\psi \in \mathcal{A}_{\mathcal{F}(n)}$ such that $|\psi(x^k) - x^k| \leq \varepsilon$ and for each $i = 1, \dots, n$, $[\psi(x^k)]_i$ (the i -th coordinate of $\psi(x^k)$), $k = 1, \dots, m$, are m distinct real numbers.*

Proof. Similar to Lemma 4.13, we prove that if $x_1^1 \neq x_1^2$, then we can find an η that separates them. The three requirements are the same as in Lemma 4.13.

Suppose $x_I^1 \neq x_I^2$ for some I . We can find two continuous increasing functions $P(\cdot)$ and $Q(\cdot)$ such that

$$\begin{aligned} P(x_I^1) - Q(x_I^1) &= 1, \\ P(x_I^2) - Q(x_I^2) &= -1, \\ P(x_I^k) - Q(x_I^k) &= 0 \quad \text{for the other } k\text{'s}. \end{aligned} \tag{4.42}$$

By the assumptions on \mathcal{F} , we can find $\tilde{P}(\cdot)$ and $\tilde{Q}(\cdot)$ such that

$$\|P(\cdot) - \tilde{P}(\cdot)\|_{C([0,1])}, \|Q(\cdot) - \tilde{Q}(\cdot)\|_{C([0,1])} \leq \min(d, 1)/4. \tag{4.43}$$

By Lemma 4.17, we know that $(x_1, \dots, x_n) \mapsto (x_1 + \tilde{P}(x_I) - \tilde{Q}(x_I), \dots, x_n)$ is in $\mathcal{A}_{\mathcal{F}(n)}$. It can be checked that this is our desired η . ■

Lemma 4.20 (Analogous to Lemma 4.14). *Suppose $\mathcal{F}(n)$ and \mathcal{F} satisfy the conditions in Proposition 4.15 and x^1, \dots, x^m are distinct points. Moreover, assume x_k is as in the conclusion of Lemma 4.13, that is, $\{x_i^k : k = 1, \dots, m\}$ are m distinct real numbers for any i . Then, given any $\varepsilon > 0$, for m target points y^1, \dots, y^m , there exists a flow map $\psi \in \mathcal{A}_{\mathcal{F}(n)}$ such that $|\psi(x^k) - y^k| \leq \varepsilon$.*

Proof. Similar to the proof of Lemma 4.14, we use x_2 to translate x_1 (denoted by η). We find two one-dimensional $P(\cdot)$ and $Q(\cdot)$, both continuous and increasing, such that

$x_1^k + P(x_2^k) - Q(x_2^k) = y_1^k$. By the assumptions on \mathcal{F} we can find $\tilde{P}(\cdot)$ and $\tilde{Q}(\cdot)$, such that

$$\|P(\cdot) - \tilde{P}(\cdot)\|_{C([0,1])}, \|Q(\cdot) - \tilde{Q}(\cdot)\|_{C([0,1])} \leq \varepsilon/2. \quad (4.44)$$

Since

$$\eta = (x_1, \dots, x_n) \mapsto (x_1 + \tilde{P}(x_2) - \tilde{Q}(x_2), \dots, x_n) \quad (4.45)$$

is in $\mathcal{A}_{\mathcal{F}(n)}$, and $|\eta(x^k)|_1 - y_1^k| \leq \varepsilon$, we conclude that η satisfies our requirement. ■

Combining these two lemmas, we obtain the following result from which we can deduce Proposition 4.15.

Lemma 4.21 (Analogous to Lemma 4.12). *Suppose $\mathcal{F}(n)$ and \mathcal{F} satisfy the conditions in Proposition 4.15. Let $\varepsilon > 0$ and $x^1, \dots, x^m, y^1, \dots, y^m \in \mathbb{R}^n$ be such that $\{x^k\}$ are distinct points. Then there exists $\psi \in \mathcal{A}_{\mathcal{F}(n)}$ such that $|\psi(x^k) - y^k| \leq \varepsilon$ for all $k = 1, \dots, m$.*

Funding. QL is supported in part by the start-up grant under the NUS PYP programme and the National Research Foundation, Singapore, under the NRF fellowship (NRF-NRFF13-2021-0005). ZS is supported by the Centennial Professorship of NUS.

References

- [1] Arnol'd, V. I.: Ordinary Differential Equations. Springer, Berlin (1992) Zbl [0744.34001](#) MR [1162307](#)
- [2] Balachandran, K., Dauer, J. P.: Controllability of nonlinear systems in Banach spaces: a survey. J. Optim. Theory Appl. **115**, 7–28 (2002) Zbl [1023.93010](#) MR [1937343](#)
- [3] Bao, C., Li, Q., Tai, C., Wu, L., Xiang, X.: Approximation analysis of convolutional neural networks. Submitted (2019)
- [4] Brenier, Y., Gangbo, W.: L^p approximation of maps by diffeomorphisms. Calc. Var. Partial Differential Equations **16**, 147–164 (2003) Zbl [1055.26008](#) MR [1956851](#)
- [5] Chang, B., Meng, L., Haber, E., Ruthotto, L., Begert, D., Holtham, E.: Reversible architectures for arbitrarily deep residual neural networks. In: Thirty-Second AAAI Conference on Artificial Intelligence, 2811–2818 (2018)
- [6] Chang, B., Meng, L., Haber, E., Tung, F., Begert, D.: Multi-level residual networks from dynamical systems view. arXiv:[1710.10348](#) (2017)
- [7] Chen, T. Q., Rubanova, Y., Bettencourt, J., Duvenaud, D. K.: Neural ordinary differential equations. In: Advances in Neural Information Processing Systems, 6571–6583 (2018)
- [8] Chukwu, E. N., Lenhart, S. M.: Controllability questions for nonlinear systems in abstract spaces. J. Optim. Theory Appl. **68**, 437–462 (1991) Zbl [0697.49040](#) MR [1097312](#)
- [9] Daubechies, I.: Ten Lectures on Wavelets. CBMS-NSF Reg. Conf. Ser. Appl. Math. 61, SIAM, Philadelphia, PA (1992) Zbl [0776.42018](#) MR [1162107](#)
- [10] Daubechies, I., DeVore, R., Foucart, S., Hanin, B., Petrova, G.: Nonlinear approximation and (deep) ReLU networks. arXiv:[1905.02199](#) (2019)
- [11] DeVore, R. A.: Nonlinear approximation. In: Acta Numerica, 1998, Acta Numer. 7, Cambridge Univ. Press, Cambridge, 51–150 (1998) MR [1689432](#)
- [12] Dupont, E., Doucet, A., Teh, Y. W.: Augmented neural ODEs. arXiv:[1904.01681](#) (2019)
- [13] E, W.: A proposal on machine learning via dynamical systems. Comm. Math. Statist. **5**, 1–11 (2017) Zbl [1380.37154](#) MR [3627592](#)

- [14] E. W., Han, J., Li, Q.: A mean-field optimal control formulation of deep learning. *Res. Math. Sci.* **6**, art. 10, 41 pp. (2019) Zbl [1421.49021](#) MR [3891852](#)
- [15] E. W., Ma, C., Wu, L.: A priori estimates of the population risk for two-layer neural networks. *Comm. Math. Sci.* **17**, 1407–1425 (2019) Zbl [1427.68277](#) MR [4044196](#)
- [16] E. W., Ma, C., Wu, L.: The Barron space and the flow-induced function spaces for neural network models. arXiv:[1906.08039v2](#) (2021)
- [17] Effland, A., Kobler, E., Kunisch, K., Pock, T.: Variational networks: an optimal control approach to early stopping variational methods for image restoration. *J. Math. Imaging Vision* **62**, 396–416 (2020) Zbl [1434.68626](#) MR [4082382](#)
- [18] Fort, M. K., Jr.: The embedding of homeomorphisms in flows. *Proc. Amer. Math. Soc.* **6**, 960–967 (1955) Zbl [0066.41306](#) MR [80911](#)
- [19] Gronwall, T. H.: Note on the derivatives with respect to a parameter of the solutions of a system of differential equations. *Ann. of Math. (2)* **20**, 292–296 (1919) Zbl [47.0399.02](#) MR [1502565](#)
- [20] Gühring, I., Kutyniok, G., Petersen, P.: Error bounds for approximations with deep ReLU neural networks in $W^{s,p}$ norms. *Anal. Appl. (Singap.)* **18**, 803–859 (2020) Zbl [1452.41009](#) MR [4131039](#)
- [21] Haber, E., Ruthotto, L.: Stable architectures for deep neural networks. *Inverse Problems* **34**, art. 014004, 22 pp. (2018) Zbl [1426.68236](#) MR [3742361](#)
- [22] Hatcher, A.: *Algebraic Topology*. Cambridge Univ. Press, Cambridge (2002) Zbl [1044.55001](#) MR [1867354](#)
- [23] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. IEEE Computer Soc. Conf. on Computer Vision and Pattern Recognition*, 770–778 (2016)
- [24] Jastrzębski, S., Arpit, D., Ballas, N., Verma, V., Che, T., Bengio, Y.: Residual connections encourage iterative inference. arXiv:[1710.04773](#) (2017)
- [25] LeVeque, R. J.: *Finite Difference Methods for Ordinary and Partial Differential Equations*. SIAM, Philadelphia, PA (2007) Zbl [1127.65080](#) MR [2378550](#)
- [26] Li, Q., Chen, L., Tai, C., E, W.: Maximum principle based algorithms for deep learning. *J. Mach. Learn. Res.* **18**, art. 165, 29 pp. (2017) Zbl [1467.68156](#) MR [3813814](#)
- [27] Li, Q., Hao, S.: An optimal control approach to deep learning and applications to discrete-weight neural networks. In: *Proc. 35th Int. Conf. on Machine Learning, Proc. Machine Learning Res.* 80, PMLR, 2985–2994 (2018)
- [28] Li, Q., Shen, Z., Tai, C.: Deep approximation of functions by composition. In preparation (2019)
- [29] Lin, H., Jegelka, S.: Resnet with one-neuron hidden layers is a universal approximator. In: *Advances in Neural Information Processing Systems*, 6169–6178 (2018)
- [30] Liu, H., Markowich, P.: Selection dynamics for deep neural networks. *J. Differential Equations* **269**, 11540–11574 (2020) Zbl [1448.49029](#) MR [4152217](#)
- [31] Lu, Y., Li, Z., He, D., Sun, Z., Dong, B., Qin, T., Wang, L., Liu, T.-Y.: Understanding and improving Transformer from a multi-particle dynamic system point of view. arXiv:[1906.02762](#) (2019)
- [32] Lu, Y., Zhong, A., Li, Q., Dong, B.: Beyond finite layer neural networks: bridging deep architectures and numerical differential equations. In: *Proc. 35th Int. Conf. on Machine Learning, Proc. Machine Learning Res.* 80, PMLR, 3276–3285 (2018)
- [33] Lu, Z., Pu, H., Wang, F., Hu, Z., Wang, L.: The expressive power of neural networks: A view from the width. In: *Advances in Neural Information Processing Systems*, 6231–6239 (2017)
- [34] Lu, J., Shen, Z., Yang, H., Zhang, S.: Deep network approximation for smooth functions. *SIAM J. Math. Anal.* **53**, 5465–5506 (2021) Zbl [07407717](#) MR [4319100](#)
- [35] Mallat, S.: Understanding deep convolutional networks. *Philos. Trans. Roy. Soc. A Math. Phys. Engrg. Sci.* **374**, art. 20150203, 16 pp. (2016)

- [36] Mallat, S.: *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, CA (1998) Zbl [0998.94510](#) MR [1614527](#)
- [37] Palis, J.: Vector fields generate few diffeomorphisms. *Bull. Amer. Math. Soc.* **80**, 503–505 (1974) Zbl [0296.57008](#) MR [348795](#)
- [38] Parpas, P., Muir, C.: Predict globally, correct locally: Parallel-in-time optimal control of neural networks. arXiv:[1902.02542](#) (2019)
- [39] Ron, A., Shen, Z.: Affine systems in $L_2(\mathbf{R}^d)$: the analysis of the analysis operator. *J. Funct. Anal.* **148**, 408–447 (1997) Zbl [0891.42018](#) MR [1469348](#)
- [40] Ruthotto, L., Haber, E.: Deep neural networks motivated by partial differential equations. *J. Math. Imaging Vision* **62**, 352–364 (2020) Zbl [1434.68522](#) MR [4082375](#)
- [41] Shen, Z., Yang, H., Zhang, S.: Deep network approximation characterized by number of neurons. *Comm. Comput. Phys.* **28**, 1768–1811 (2020) Zbl [07419151](#) MR [4188521](#)
- [42] Shen, Z., Yang, H., Zhang, S.: Nonlinear approximation via compositions. *Neural Networks* **119**, 74–84 (2019) Zbl [07209015](#)
- [43] Shen, Z., Yang, H., Zhang, S.: Optimal approximation rate of ReLU networks in terms of width and depth. *J. Math. Pures Appl.* (9) **157**, 101–135 (2022) Zbl [07453385](#) MR [4351074](#)
- [44] Sonoda, S., Murata, N.: Double continuum limit of deep neural networks. In: *ICML Workshop Principled Approaches to Deep Learning* (2017)
- [45] Sonoda, S., Murata, N.: Transport analysis of infinitely deep neural network. *J. Mach. Learn. Res.* **20**, art. 2, 52 pp. (2019) Zbl [07049722](#) MR [3911409](#)
- [46] Sun, Q., Tao, Y., Du, Q.: Stochastic training of residual networks: a differential equation viewpoint. arXiv:[1812.00174](#) (2018)
- [47] Sussmann, H. (ed.): *Nonlinear Controllability and Optimal Control*. CRC Press (1990) Zbl [0699.00040](#) MR [1061380](#)
- [48] Tao, Y., Sun, Q., Du, Q., Liu, W.: Nonlocal neural networks, nonlocal diffusion and nonlocal modeling. In: *Advances in Neural Information Processing Systems*, 496–506 (2018)
- [49] Thorpe, M., van Gennip, Y.: Deep limits of residual neural networks. arXiv:[1810.11741](#) (2018)
- [50] Utz, W. R.: The embedding of homeomorphisms in continuous flows. *Topology Proc.* **6**, 159–177 (1981) Zbl [0491.54035](#) MR [650486](#)
- [51] Veit, A., Wilber, M. J., Belongie, S.: Residual networks behave like ensembles of relatively shallow networks. In: *Advances in Neural Information Processing Systems*, 550–558 (2016)
- [52] Wang, B., Yuan, B., Shi, Z., Osher, S. J.: EnResNet: ResNets ensemble via the Feynman–Kac formalism for adversarial defense and beyond. *SIAM J. Math. Data Sci.* **2**, 559–582 (2020) Zbl [07482305](#) MR [4121885](#)
- [53] Warga, J.: Relaxed variational problems. *J. Math. Anal. Appl.* **4**, 111–128 (1962) Zbl [0102.31801](#) MR [142020](#)
- [54] Yarotsky, D.: Error bounds for approximations with deep ReLU networks. *Neural Networks* **94**, 103–114 (2017) Zbl [1429.68260](#)
- [55] Zhang, D., Zhang, T., Lu, Y., Zhu, Z., Dong, B.: You only propagate once: Painless adversarial training using maximal principle. arXiv:[1905.00877v1](#) (2019)
- [56] Zhang, H., Gao, X., Unterman, J., Arodz, T.: Approximation capabilities of neural ordinary differential equations. arXiv:[1907.12998v1](#) (2019)
- [57] Zhang, L., E, W., Wang, L.: Monge–Ampère flow for generative modeling. arXiv:[1809.10188](#) (2018)
- [58] Zhang, X.: Embedding diffeomorphisms in flows in Banach spaces. *Ergodic Theory Dynam. Systems* **29**, 1349–1367 (2009) Zbl [1177.37026](#) MR [2529653](#)
- [59] Zhang, X., Lu, Y., Liu, J., Dong, B.: Dynamically unfolding recurrent restorer: A moving endpoint control method for image restoration. arXiv:[1805.07709](#) (2018)
- [60] Zhou, D.-X.: Deep distributed convolutional neural networks: universality. *Anal. Appl. (Singap.)* **16**, 895–919 (2018) Zbl [1442.68214](#) MR [3871718](#)